

Thomas Seidelmann

**A universal, multi-mechanism,
simulation model for species
coexistence based on neutral
theory**



FAKULTÄT FÜR
INFORMATIK

Intelligent Cooperative Systems
Computational Intelligence

A universal, multi-mechanism, simulation model for species coexistence based on neutral theory

Master Thesis

Thomas Seidelmann

March 7, 2019

Supervisor: Prof. Dr.-Ing. habil. Sanaz Mostaghim

Advisor: Heiner Zille

Advisor: Dr. Christoph Steup

Thomas Seidelmann: *A universal, multi-mechanism, simulation model for species coexistence based on neutral theory*
Otto-von-Guericke Universität
Intelligent Cooperative Systems
Computational Intelligence
Magdeburg, 2019.

Abstract

The mystery of species coexistence, to this day, remains as one of the largest topics of contemporary research in ecology. And with good reason, considering that large portions of the factors which lead to sustained coexistence of competitors are poorly understood. With this thesis we propose a new multi-mechanism model, allowing the study of combined effects from multiple coexistence mechanisms at once, which have previously only been considered in isolation. We focus on storage effect, intransitivity and resource partitioning as research subjects, but the model is extendable. We show that our implementation produces results in agreement with contemporary research and then conduct experiments with combined mechanisms. Our results indicate strong influences between mechanisms in all tested scenarios. Overall it was observed that the majority of combinations promoted species coexistence strongly, although some combinations were found to be counterproductive. We show that the interactions between mechanisms alone can account for variations in coexistence times across orders of magnitudes. Additionally our findings suggest a significant dependence on community structure and size. These results underline the limitations of highly specialized models and imply that a more integrative approach may be necessary to understand the biodiversity we observe in nature.

Contents

List of Figures	VIII
List of Tables	IX
List of Algorithms	XI
1 Introduction	1
1.1 Goals	2
1.2 Structure of the Thesis	3
2 Scientific Background	5
2.1 Terminology & Basic Principles	5
2.2 Niche Theory	8
2.3 Neutral Theory	10
2.4 Coexistence Mechanisms	11
2.4.1 Storage Effect	11
2.4.2 Intransitivity	12
2.4.3 Resource Partitioning	13
2.4.4 Spatial Effects	14
2.4.5 Predation	15
2.4.6 Mutualism	16
2.5 Review of Species Coexistence Models	17
2.5.1 The classic neutral model	17
2.5.2 Neutral Model with Integrated Storage Effect	19
2.5.3 Cellular Automaton Model with Intransitive Competition	21
2.6 Summary	22

3	Concept of the Multi-Mechanism Model	23
3.1	Model Design Goals	23
3.2	Key Challenges	25
3.3	Scope & Requirements	25
3.4	The Extended Neutral Base Model	27
3.4.1	Environmental Fluctuations	28
3.4.2	Spatially Explicit Environment	30
3.4.3	Explicit Modeling of Physical Time	31
3.4.4	Explicit Modeling of Resources	33
3.4.5	Explicit Modeling of Competition	34
3.5	Integration of the Coexistence Mechanisms	36
3.5.1	Storage Effect	37
3.5.2	Intransitivity	38
3.5.3	Resource Partitioning	39
3.5.4	Spatial Effects	40
3.5.5	Predation	40
3.5.6	Mutualism	41
3.6	Concept Summary	42
3.6.1	Neutral Base Model	42
3.6.2	Coexistence Mechanisms	45
4	Implementation	47
4.1	Program Flow	48
4.2	Implementation Details	52
4.2.1	Environmental Fluctuation Functions	52
4.2.2	Target Selection Function	54
4.2.3	Event Based Time Steps	55
4.2.4	Mortality Rate Function	56
4.2.5	Survival Check Function	57
4.2.6	Decision Function for Food Consumption	58
4.3	Avoiding Biased Results	59
4.4	Parameter Reduction & Meta Parameters	60
4.5	Performance Considerations	62

5	Evaluation	63
5.1	Experimental Setup	63
5.1.1	Structure & Presentation	65
5.1.2	Collected Data	67
5.2	Validation Experiments	68
5.2.1	Demographic Stochasticity	69
5.2.2	Environmental Stochasticity	72
5.2.3	Storage Effect	75
5.2.4	Intransitivity	78
5.2.5	Resource Partitioning	84
5.3	Parameter Range Experiments	86
5.3.1	Base Scenario	87
5.3.2	Storage Effect	89
5.3.3	Intransitivity	93
5.3.4	Resource Partitioning	95
5.4	Combined Coexistence Mechanisms Experiments	97
5.4.1	Storage Effect & Intransitivity	98
5.4.2	Storage Effect & Resource Partitioning	100
5.4.3	Intransitivity & Resource Partitioning	102
5.4.4	Storage Effect & Intransitivity & Resource Partitioning .	108
6	Conclusion	111
	Bibliography	115

List of Figures

2.1	Example for a practical application of niche theory.	9
2.2	Illustration: Storage Effect	11
2.3	Illustration: Intransitivity	12
2.4	Illustration: Resource Partitioning	14
2.5	Illustration: Spatial Effects	14
2.6	Illustration: Predation	15
2.7	Local community dynamics of the classic neutral model.	18
2.8	Overview of the classic neutral model.	19
3.1	Adding environmental fluctuations.	29
3.2	Adding a spatially explicit environment.	30
3.3	Adding a concept of physical time.	32
3.4	Adding an environmental community with resources.	34
3.5	Adding a more refined competition system.	35
4.1	Screenshot of our visualization for the multi-mechanism model.	48
4.2	Plot of the mortality rate function.	57
5.1	Examples of frequently used diagrams.	66
5.2	Demographic stochasticity validation experiment boxplots.	71
5.3	Histogram to elaborate on Figure 5.2.	72
5.4	Environmental stochasticity validation experiment boxplots.	74
5.5	Histogram to elaborate on Figure 5.4.	75
5.6	Overall results of the storage effect validation experiment.	77
5.7	Histogram to elaborate on Figure 5.6	78

5.8	Scenario definitions for the intransitivity validation experiments.	81
5.9	Reference results for the intransitivity validation experiment.	82
5.10	Results of the intransitivity validation experiment.	83
5.11	Resource partitioning validation results with linear y-axis.	86
5.12	Resource partitioning validation results with logarithmic y-axis.	86
5.13	Persistence times of the multi-mechanism base scenario.	89
5.14	Storage effect parameter range experiment results.	92
5.15	Intransitivity parameter range experiment results.	94
5.16	Case studies of the intransitivity parameter range experiment.	95
5.17	Resource partitioning parameter range experiment results.	97
5.18	Storage effect and intransitivity results.	99
5.19	Storage effect and intransitivity case studies	100
5.20	Storage effect and resource partitioning results.	102
5.21	Intransitivity and r. partitioning results with linear z-axis.	105
5.22	Intransitivity and r. partitioning results with logarithmic z-axis.	106
5.23	Intransitivity and resource partitioning case studies.	107
5.24	Combined results of all mechanisms with linear y-axis.	110
5.25	Combined results of all mechanisms with logarithmic y-axis.	110

List of Tables

5.1	Default model parameters.	64
5.2	Tracked metrics during the experiments.	68
5.3	Demographic stochasticity validation parameters.	70
5.4	Environmental stochasticity validation parameters.	73
5.5	Storage effect validation parameters.	76
5.6	Intransitivity validation parameters.	80
5.7	Resource partitioning validation parameters.	84
5.8	Base scenario range experiment parameters.	88
5.9	Storage effect range experiment parameters.	89
5.10	Intransitivity range experiment parameters.	93
5.11	Resource partitioning range experiment parameters.	96
5.12	Storage effect and intransitivity parameters.	98
5.13	Storage effect and resource partitioning parameters.	101
5.14	Intransitivity and resource partitioning parameters.	103
5.15	Parameters for all mechanisms combined.	108

List of Algorithms

1	A single update step of the model.	49
2	Planning step for one individual.	50
3	Acting step for one individual.	50
4	Resolving step for one individual.	51

1 Introduction

“Survival of the fittest”, a phrase which almost everyone encounters sooner or later in life. From a layman’s perspective it seems to imply that only the strongest, best, most adapted individuals pass on their genes and survive collectively as a species. The competitive exclusion principle [1], suggesting that no two species which compete for the same resources can coexist in the long term, then apparently drives this point home. Thinking these assumptions through, one expects nature to be thoroughly dominated by just a few species. Falling out of favor due to environmental changes or genetic drift implies timely extinction. Actual investigation of natural communities however suggests otherwise: the evidence in favor of a vast biodiversity is overwhelming. This divergence between low diversity predictions of theory and high diversity observations in nature has been recognized for over 50 years by now [2, 1]. Furthermore, not only do many different species exist in any given area, but a very high amount of them is actually exceedingly rare, so much so, that singletons are consistently found to be significant contributors to the total biodiversity of any given community [3]. This characteristic of few highly abundant species and many rare ones is shared by virtually all natural communities and is therefore considered to be a true universal law of ecology [4]. But how is this possible? Although studied for over a century by now [5], the maintenance of biodiversity in natural systems remains a topic of strong debate and our understanding in regard to the factors which lead to sustained species coexistence is still lacking. Major advancements have been made, the most well known of which is probably niche theory, as it has become the traditional answer to the question of species coexistence: every existing species occupies a unique niche in which it specializes and performs better than any other species. Therefore it is only through the fundamental differences between species, that coexistence is possible. While niche theory is well supported [6], it does have its drawbacks [6, 7, 8, 9] and doesn’t seem plausible in all scenarios, a problem known as the paradox of the plankton [1]. In contemporary research another theory has recently gained a lot of traction: the unified neutral theory of biodiversity

and biogeography [10], and with it the whole class of neutral theory. Unlike niche theory, neutral models are characterized by their lack of differentiation between modeled species. In these models coexistence is merely a stochastic process between functionally identical species, differentiated only through their labels. Furthermore, irrespective of any particular theory, several different mechanisms which promote species coexistence are known to exist [7, 4, 6, 9]. A sizable amount of recent work has focused on these mechanisms and how they may help to explain the gaps in our understanding which are still left. In particular some studies have integrated such mechanisms into neutral models to improve their predictions, while keeping their simplicity intact. However, these studies only focus on a single mechanism at any given time and are often limited to two-species scenarios [11], disregarding complex interactions of the studied systems and possible emergent effects. At this point we contribute, by taking the efforts which have already been done one step further, towards a more holistic and integrative approach in studying the mysteries of species coexistence.

1.1 Goals

This thesis aims to provide new insight as to how different coexistence mechanisms may work together in nature to support the vast biodiversity we can observe. Of particular interest are the relations and interactions of these different mechanisms which rarely seem to be considered. There are however a sizable amount of studies focusing on just a single mechanism in isolation and how this impacts the coexistence times for a given scenario. This apparent lack of research on multi-mechanism models appears to be an oversight, since species coexistence in nature most certainly does not rely on any individual mechanism, but rather on an inseparable mix of them all. Any feasible mechanism imaginable can reasonably be expected to occur regularly in natural systems, and almost always in conjunction with others. This leads to the main goal of this thesis:

The creation of an extensive multi-mechanism model, highlighting the relations and interactions of the various coexistence mechanisms and their combined effects on species coexistence.

Achieving this goal is no trivial matter and requires careful planning. But once such a model is available, it may help to expand our knowledge by providing means to study the circumstances which lead to species coexistence in a more holistic manner, previously unattained. Further research questions follow:

1. How do the different mechanisms compare against each other? Which one has the highest influence, which one the lowest?
2. How do the different coexistence mechanisms influence each other? Are there any emergent effects?
3. Which combinations of mechanisms are most effective?

1.2 Structure of the Thesis

The thesis is structured as follows: First the relevant scientific background is briefly reviewed with a focus on the specific theories and models which are most relevant to this work. Next we derive and explain the concept of our proposed multi-mechanism model. Although it differs considerably from established ones and is comparatively complex, it is strongly rooted in contemporary neutral theory and can be described as a series of adjustments to the classic neutral model. This is utilized to break the complexity down into multiple comparatively simple parts which build on each other to eventually yield the complete model. At this point it is conceptually well defined, including how it works and why certain choices were made. The implementation chapter then describes our concrete realization of the previously defined concept. This includes a brief architectural overview of the written code and specifications for chosen solutions whenever the concept does not mandate a certain approach. It follows the evaluation chapter in which we outline our experimental setup, define multiple experiments, present their results and discuss them. Finally, we draw a conclusion and suggest avenues for the future work on this topic.

2 Scientific Background

This chapter begins by outlining the terminology used from here on out, defining frequently used and central basic principles in the thesis. Afterwards, the two major foundational theories on species coexistence are explained: niche theory and neutral theory. The majority of this thesis is firmly rooted in neutral theory and its school of thought, therefore it is paid more attention overall. An overview of the most notable and recognized coexistence mechanisms, as they are found in nature, follows. These are of central importance to this work as a whole, seeing as the main goal is to bring them together into a single model. Finally, some of the already proposed models are reviewed, focusing on those which had the biggest influence on and relate closest to the model which we propose.

2.1 Terminology & Basic Principles

In the following we define several terms for use in the context of this thesis.

Community The term community refers to the sum of all individuals within a certain distinct system. One could think of it as the whole contained biomass, all species included. If multiple communities exist, then they serve functionally different purposes. An example of this would be our distinction between active community and resource community: The first contains all the individuals which move around and interact with their surroundings. It is this communities population dynamics we are interested in. The second community on the other hand contains only passive resource individuals, which represent food sources in the environment for the active individuals to consume.

Population Compared to the community, which refers to all individuals and species within a distinct system, the term population is generally used to refer to the individuals of one specific species only.

Generation A generation in this context is defined to last until an amount of deaths occurred within a community, which is equal to the community's total size as of the beginning of the generation. Note that it is not necessary for each of the individuals that existed at the beginning of a generation to actually die before the next one can begin.

Zero-Sum Refers to a common simplifying assumption in species coexistence models, which mandates that modeled communities are of constant size. This forcefully limits the dynamics which occur in such a community to be purely relative in nature: one species cannot become more abundant without another one becoming less abundant at the same time.

Fitness The fitness of an individual, or of a whole species, is a general indicator of its reproductive potential. Meaning that individuals with higher fitness are more likely to produce offspring than those with lower fitness.

Trophic Levels Refers to the positions that individuals occupy in a food chain. Therefore a fox is on a higher trophic level than a rabbit, which in turn is on a higher trophic level than carrots.

Competition Competition as it is used here refers to a contest between trophically similar individuals or species over a common resource. It is therefore distinct from predation, which necessarily involves members of different trophic levels. Furthermore, predation generally results in the immediate death of the prey, while competition does not necessarily have any adverse effects on the participating individuals, aside from missing out on the contested resources.

Intraspecific / Interspecific Competition Refer to competition between members of the same, or different, species respectively.

Dispersal Dispersal in the broadest sense describes the patterns of spatial distribution for individuals or species. The focus is mostly on movement across generations, e.g. lasting changes, as opposed to temporary displacements like in seasonal animal migrations. It is most commonly used when discussing plant communities in which different species have developed various means to disperse their seeds. High dispersal describes that the seeds travel a long distance before sprouting and vice versa.

Recruitment Rate When talking about recruitment rates, the overall ability of a species to produce offspring is meant. Recruitment in this sense usually doesn't refer to any particular reproductive event or individual, instead it is to be understood as a metric, used to describe a species.

Reproduction Usually refers to the process of a particular individual giving birth to offspring. However, when speaking about reproductive ability, or rate, it becomes largely synonymous to recruitment rate.

Mortality Rate Is mostly used to describes the chances of an individual to die over a given time period in a given scenario. May sometimes refer collectively to all members of a species.

Immigration Immigration refers to the concept of individuals crossing over from one community into another one. The term may also be used to just describe one such event.

Speciation Speciation describes the concept of introducing new species to a community as an implied result of evolution. Most often a speciation event is handled as a modified reproduction in which the offspring does not inherit the species of the parent, but funds its own.

Invasion An invasion is any situation in which a species of low abundance is introduced to an existing community. Most often this occurs through immigration, but may also be a result of speciation. Invasion scenarios are frequently studied, observing how successful the invaders are in their new environment and if they can establish themselves, or even take over the community.

Environmental Fluctuations Sometimes also referred to as environmental stochasticity. In this context they describe how the fitness of a species changes over time due to varying, not explicitly modeled, outside influences. No species is unaffected by its environment and each one strives at specific levels of humidity, temperature and so on.

Coexistence Mechanism A coexistence mechanism is any observable, distinct pattern in natural communities which aids the maintenance of biodiversity.

Persistence Time The persistence time, often also referred to as coexistence time, is a measurement for the amount of generations a community experienced until one or more of its contained species went extinct.

Stability of a system The more stable a system is, the higher the chances are for a species to recover from low abundance, instead of going extinct.

Delta Time A shorthand term, referring to the amount of time represented by a single time step, or iteration, of a model with an explicit concept of time.

2.2 Niche Theory

Between all the different theories and developments in the vast field of ecology, niche theory is among the most fundamental ones and has been discussed in science for over a century by now [5]. Its influence is consequently far reaching. The popularity of the niche concept is hardly surprising, considering there is an abundance of case studies which support it [6]. Furthermore, the notion of specialized species is a very intuitive and straightforward explanation, which makes it desirable. However, there are some drawbacks associated with niche theory, just like with any other theory.

As the name already implies, niche theory revolves around the concept of ecological niches. What exactly constitutes such a niche is still controversial [5], but in simple terms a niche may be imagined as any sufficiently unique set of requirements which a species may have at its environment in order to survive

in it. This definition can include the presence of specific resources, weather and climate, competing species, predators and so forth. Each and every single one of these requirements constitutes a “niche axis”. Combining all the niche axes yields a high dimensional space in which each species occupies a unique hypervolume, which represents the specific niche of that species. Because every hypervolume is unique to a specific species, it follows that each possible niche can only ever be occupied by a single species over the long term [5]; a claim that is formally known as the “competitive exclusion principle”. While there exist many very prominent, influential and easily measured niche axes, such as temperature, humidity and resource availability, the identification and measurement of more obscure niche axes is often very difficult [7, 9]. This problem also demonstrates a weaknesses of niche theory in general: it tends to be highly specific to particular problems and is usually very parameter rich [6]. As a result the developed models are often highly complex and may not be universally applicable [6, 7, 8].

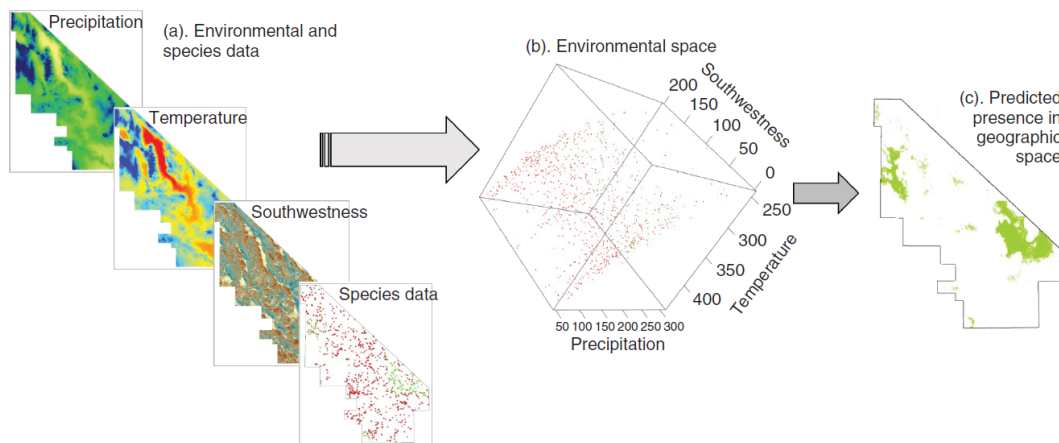


Figure 2.1: An example for a practical application of niche theory as depicted by Miller and Hallaway [5]. Empirical data on environmental conditions and species occurrence, shown on the left, are combined to identify the species niche. Then the potential habitats of the species can be calculated for the whole region, shown on the right.

Figure 2.1, initially published by Miller and Hallaway [5], illustrates a practical application of niche theory in the form of niche modeling. These types of models try to identify the niche of a species in order to make predictions about it, for example regarding distribution, abundance or invasion capabilities. The data for niche modeling is typically obtained through measurement of

selected environmental influences and sampling of species abundances within a given region. In this example three niche axes are considered: precipitation, temperature and southwestness. Together they form a three dimensional niche volume in which the species occupies a certain hypervolume. In order to calculate this hypervolume, and thereby the species niche, the environmental data is combined with a fourth dataset which indicates the geographical presence or absence of the species for a limited number of sampled spots. The combination of species occurrence and niche volume is depicted in the middle, it allows the identification of the species requirements at its environment. Once the species niche is then identified, predictions can be made. In this case all the possible regions in which the species is assumed to thrive in are calculated, effectively completing the previously fragmented dataset on species occurrence. Although one has to keep in mind that the result does not necessarily indicate the actual presence of the species, but rather the potential for it to persist in a given region. This insight may for example be used when trying to determine a good spot to return an animal to the wild, or to determine if a certain community is at risk of getting invaded by the species, in which case countermeasures could be employed in an informed manner.

2.3 Neutral Theory

Contrary to niche theory, which postulates profound species differences for long term coexistence, neutral theory thoroughly abstracts these differences away and explains species coexistence through a purely stochastic process. Although the concept of neutrality to explain ecological phenomena, such as the maintenance of biodiversity, is certainly not a new idea, the approach has only gained a sizable amount of attention in recent years [7, 4, 12]. More specifically since 2001, the year in which Hubbel published the unified neutral theory of biodiversity and biogeography [10]. The fundamental assumption of this theory is, that individuals of trophically similar species may exhibit different traits and behaviors, but these differences are ultimately irrelevant for their success and may therefore be disregarded. For this and other implicit assumptions the theory has received much criticism [7, 13, 14, 15, 16], however its usefulness is generally undoubted [7, 4, 6, 9]. Neutral models can be seen as null models, governed exclusively by stochastic fluctuations instead of intricate interspecific interactions. That said, neutral theory does not exclude complex

interactions between individuals or species, as long as those interactions are fully symmetrical. A closer look at the classic neutral model and variations of it follow in Section 2.5.

The main advantage of neutral models lies in their simplicity and clarity, which makes them well suited to be used as a basis for research specific modifications [8, 7, 17]. Simultaneously it offers an unbiased baseline to compare against [8, 7, 17]. For these reasons the model of this thesis was chosen to be based on neutral theory, though it does not fulfill the standard of neutrality to the fullest extent, since certain coexistence mechanisms require species differences to work.

2.4 Coexistence Mechanisms

It is well established that several different mechanisms interact in nature to support the coexistence of species [18, 19, 9, 12] and are therefore of central importance for the maintenance of biodiversity. On the following pages the different well established mechanisms are shortly presented.

2.4.1 Storage Effect

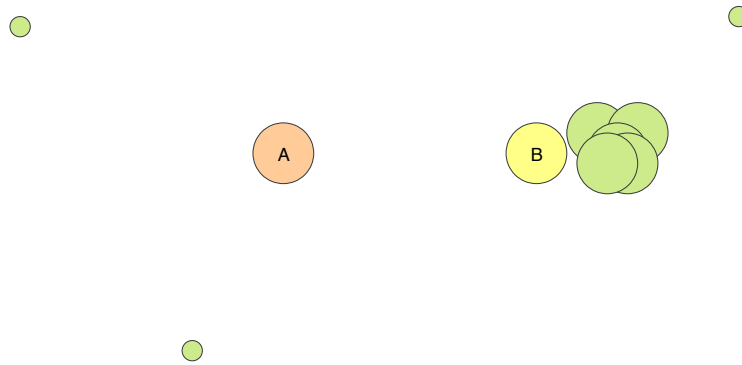


Figure 2.2: Illustration of the storage effect: Resources have become scarce, but individual *B* has stored food during good times to buffer against these conditions. *A*, on the other hand, is now in a bad situation.

The storage effect is probably one of the most studied mechanisms outlined here. It is a fluctuation dependent mechanism [20, 18], meaning that species

need to experience fitness altering changes over time which influence mortality, survival or recruitment rates. All species need to react differently to the changing environment. For example one type of flower may thrive in very humid environments while another one favors an arid climate. Then, when a species experiences a favorable environment it *stores* some of the gains as a buffer against bad times. When a shift of the environment happens which disfavors the species again, it can then survive the stress by avoiding competition and living out of the created buffer. This way the weakened species increases its chances of survival until another shift occurs, during which the population can recover the sustained losses [20, 18].

In natural systems storage in this sense is ubiquitous and there are many prominent examples of the effect: Beginning with hibernation of hedgehogs and bears, over food hoarding squirrels, up to soil seed banks which can ensure species survival over hundreds of years. But the storage effect is not only present on large scales, it also acts on small ones, for example when a bird fluffs its feathers to keep warm during cold days, or quite simply the pelt of a polar bear which has the same effect. Figure 2.2 illustrates the effect.

2.4.2 Intransitivity




Figure 2.3: Illustration of transitive and intransitive competition networks, respectively. On the left, *A* dominates all other species, but on the right, there is no universally best competitor and the system is self-stabilizing.

It is widely known that individuals in nature experience frequent competition with others over limiting resources, such as food and space. Certain species compete best for certain resources, but no species dominates in all respects.

The result is a complex network of competitive relations, which includes transitive, e.g. hierarchical, and intransitive components. Strongly hierarchical systems are found to be less stable than those with intransitive relations, because in strict hierarchical systems there exist no species which can restrict the top competitor [12, 21, 22]. Its population inevitably rises, eventually eliminating all inferior species. A system with strong intransitive loops on the other hand is self stabilizing as a population rise in one species indirectly aids its superior competitor which can then limit the species again [12, 21, 22]. In this sense the game rock-paper-scissors is a prime example of this: No dominant strategy exists, therefore all options are valid and “coexist”. Figure 2.3 illustrates this concept.

As outlined by Laird & Schamp [12] intransitivity is a particularly promising coexistence mechanism since it allows for strong competition among species, including differences in competitive potential, without sacrificing biodiversity on the whole. It is further suggested that competitive intransitivity is a promising candidate to help reconcile neutral theory with traditional ecology. Niche theory emphasizes species differences as the main factor for coexistence and observations of natural systems seem to support that. Laird & Schamp argue that these observations and neutral theory can both be accurate in the presence of strong intransitive competition networks: On small scales competition appears to follow transitive hierarchies, but on large scales intransitive dynamics take over and let the system appear neutral. While this is promising, the amount of intransitivity present in natural systems is controversial [23, 22, 21]. Empirical data on competition hierarchies are difficult to obtain, especially because small changes in the network can have significant consequences [22]. For these reasons most work on this coexistence mechanism is of theoretic nature.

2.4.3 Resource Partitioning

This mechanism is very well described by Griffin & Silliman [24]: Resource partitioning is one of the biggest drivers of niche theory and simultaneously one of the most trivially obvious coexistence mechanisms. It states that different species specialize in different resources or, if they are limited by the same resource, obtain it in different ways. This way species can avoid a large amount of interspecific competition, while simultaneously reinforcing their own




Figure 2.4: Illustration of resource partitioning. Although A and B both compete for the same resources, a significant amount of competition can be evaded due to their varying food dependencies.

niche. This concept is illustrated by Figure 2.4.

The most obvious examples of resource partitioning in nature is when species eat different food or inhabit non-overlapping areas. Perhaps the biggest common divide of resources is between aquatic and terrestrial life. However, just like with all the other mechanisms there are many more nuanced trade offs involved too. For example bumble bee species are found to divide resources among themselves based on the length of their proboscises and the altitude in which they are active, both enabling them to exploit different plants, [24].

2.4.4 Spatial Effects




Figure 2.5: Illustration of spatial effects in the form of density-dependence: Due to a high spatial clustering, B experiences stronger competition than A .

The presence of highly influential spatial effects in natural systems can hardly be denied. While animals are likely not as strongly constrained due to their mobile nature, plants depend heavily on their immediate local environment for all the resources they need. Spatial effects are often modeled in the form of density-dependent mortality (illustrated by Figure 2.5) or recruitment rates and dispersal limitations. Their frequent use can be explained by the profound effect they have on species coexistence [25], while simultaneously being comparatively easy to integrate into a model, because they don't require fluctuations to work [25]. In the case of density-dependence, the chances of death and birth are usually directly linked to the abundance of the respective species population. High abundance acts as a limiting factor by increasing mortality and slowing recruitment, while the opposite is true for low abundance, although variations of this are also used. Dispersal limitation on the other hand acts on the ability of a species to spread in space. For example plants do not move themselves, but spread through their seeds which may cover great distances before germinating in another place. Depending on the involved species and their dispersal limitations, coexistence can be prolonged [25].

2.4.5 Predation




Figure 2.6: Illustration of the stabilizing effect from generalist predators, represented by A, which naturally limit more abundant species stronger than less abundant ones.

It is controversial whether the workings of parasites and pathogens are similar enough to classic predator-prey interactions to include them in this category. But all of them can have strong stabilizing influences on a multi-species community [18]. This overview mainly focuses on classic predator-prey interactions, although most statements remain true for parasitic effects and pathogens as well. Predators are generally classified as being “generalists” or “specialists”, which indicates the diversity of species they eat. Generalists feed on many different species, and as illustrated by Figure 2.6, they consequently limit highly abundant species more than rare ones, simply because they are more readily encountered. In this form predation is usually a stabilizing, density-dependent effect [18]. Specialist predators on the other hand focus on just a selected few species and won’t switch to other prey, even if it is more abundant. As a result specialist predators can have stabilizing and destabilizing effects, depending on the specific circumstances [18]. Predation can seem like a simple explanation for stabilizing effects in natural communities, but has been recognized to be very intricate in practice [18].

2.4.6 Mutualism

Mutually beneficial interactions in natural systems are well established and documented [19]. It has also been shown that such positive interactions can occur even when species are competitors over a common limiting resource [19]. Very prominent examples of mutualism are the well known mycorrhizal networks of tropical forests which help the hosting plants by redistributing essential nutrients and minerals to those in need.

Mutualism is a coexistence mechanism which has rarely been considered, as most studies focus heavily on competitive species interactions. Nonetheless it has been shown that mutually beneficial relations can enable the coexistence of multiple species, even if they are limited by a single common resource and the net interspecific interactions remain negative [19]. That is to say the total effect of all interspecific interactions can still increase mortality rates of competitors more than they decrease them in such a scenario, but without endangering species coexistence.

2.5 Review of Species Coexistence Models

Many models exist which may be used to study coexistence times for specific communities and conditions, which is the main metric of interest here. Some of these model one or two specific coexistence mechanisms [9, 12, 22, 23, 26, 19, 25, 27], but none explicitly model and consider the interaction of more than two at once, nor are the combined effects focused. While not all of those models are said to be based on neutral theory, most of the described effects could easily be implemented on top of such a system. Contrary to the goal of this thesis, most of these models are driven phenomenologically, which is to say that they do not try to explain the underlying interactions in detail, but rather focus on the end result. A selected few of these models are detailed on the following pages, with a focus on those which had especially high influence on the design of our model, as we propose it with this thesis.

2.5.1 The classic neutral model

The classic neutral model can be seen as the foundation on which most contemporary neutral models are built, it is therefore of particular interest. This model, just like the others, is an individual based iterative simulation, although formulations as a purely mathematical model exist too [9]. As described by Rosindell, Hubbel and Etienne [6] this model can be characterized by consisting of two distinct communities: The local community and a very similar, but independent, metacommunity. Both may contain different individuals and species. The local community represents all the individuals within a restricted spatial area and is the focus of the model. The metacommunity on the other hand is analogous to the whole surroundings of that area and has its own inhabitants.

Both, the local community and the metacommunity, are initialized to contain a certain number of individuals, which do not have to be equal. Since this is a zero-sum model, the total amount of individuals in each community will remain constant throughout the process. A step in the simulation is then computed as follows: First one random individual of the local community is chosen to die, so that its “place” becomes unoccupied. The empty place is then




Figure 2.7: Graphical representation of the local community dynamics from the classic neutral model, taken from Rosindell et al. [6]. It shows one update iteration for the local community: A random individual dies and is then replaced either by a birth or immigration event. Note that this graphic implies a spatially explicit, grid based structure for the model, but this is only for visualization purposes.

promptly filled, either by a copy of another individual from the community, or by an immigrant from the metacommunity, which is to say a copy of a random individual out of the metacommunity. This updating process for the local community is also visualized in Figure 2.7. Although the mentioned figure seems to imply that the model itself is grid based and therefore spatially explicit, this is actually not the case. In the purest form of this model, species are solely defined through their label, the community to which they belong, and their abundance, nothing else. Once the local community has been updated through death and replacement, the metacommunity is updated in a very similar fashion. The only difference when updating the metacommunity compared to updating the local community is that the random dead individual cannot be replaced by an immigrant. Instead it can either be an offspring from a random member of the metacommunity, or a speciation event may occur, in which case an individual of a completely new species is born. The Figure 2.8 shows a simple representation of both, local and metacommunity, and includes flowcharts describing a single update process for each of them.

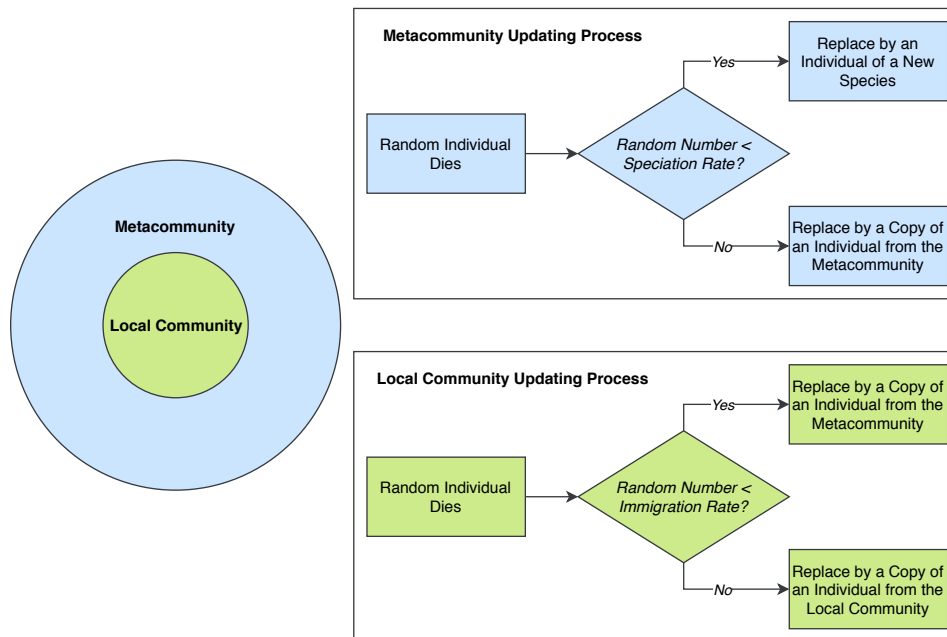


Figure 2.8: Overview of the classic neutral model, including a simple visualization of local community and its surrounding metacommunity on the left, and the corresponding updating processes for each respective community on the right.

2.5.2 Neutral Model with Integrated Storage Effect

Danino, Kessler & Shnerb published a general model for measuring species coexistence times based on the classic neutral model in 2016 [9]. They were able to show that the integration of pure environmental stochasticity, implemented as species specific fitness fluctuations, is a destabilizing factor. However, it opens up the possibility of integrating a storage effect mechanism which prolonged persistence times significantly. While neutrality is not preserved on any singular time step due to the species specific fitness fluctuations, the time-averaged fitness values are however identical for sufficiently large sample sizes. This particular model is restricted to two-species communities and omits the notion of a metacommunity. Eventual extinction of one species is guaranteed due to the stochastic nature of the model and lack of any newly generated species. The model focuses fully on the amount of generations until extinction occurs. The stability of the system can then be measured by observing the persistence times for different initial population ratios. This

is especially useful to measure the chances of invasion by a low-abundance species, which also translates to the stability of the system in this case.

The model as described in the paper by Danino et al. [9] can be divided into three distinct variations or stages, each building on top of the previous one. It is therefore best described by explaining these variants in order, finally arriving at the full model with integrated storage effect. These models are of great interest, because they initially served as a basis for the proposed multi-mechanism model and present a reference for our implementation of the storage effect coexistence mechanism to be compared against.

Model with Demographic Stochasticity: This version of the model serves as a basis for the two subsequent ones. It is equivalent to the classic neutral model (Section 2.7), but formulated specifically for two-species communities and without a metacommunity. This means that there are no immigration or speciation events, the process simplifies to only deaths and births of individuals within the local two-species community. In each time step two individuals of the community are selected randomly, the first of which dies, and is then replaced by a copy of the second one. This process continues until the eventual extinction of either species.

Model with Environmental Stochasticity: Highly similar to the aforementioned model, however the notion of fluctuating species-level fitness is now introduced. Each iteration consists mainly of a single competitive event: Two random individuals are selected, one of which will die while the other will reproduce, similar to the first variant. But the chance for either individual to survive and produce offspring is now proportional to the fitness relative to that of the other individual. The advantaged species as a whole is consequently more likely to reproduce and simultaneously less likely to die. The fitness values of each species fluctuate over time, so that in any time step one species is advantaged. It can be observed that this has a destabilizing effect on the system. Persistence times are consistently shorter than in the base model.

Model with Storage Effect: Once again this adaptation is very similar to the previously described variant. The notion of fluctuating species fitness is

kept, but the updating process is changed, so that fitness does no longer affect a species mortality rate, only its recruitment rate. The described competitive event is therefore replaced by a fitness-independent random death, followed by a birth event with proportionally higher chances for the advantaged species to increase its abundance. With this modification the model satisfies all formal requirements of the storage effect and this is reflected in the persistence times, which can be observed to be significantly greater than they were for either of the variants before.

2.5.3 Cellular Automaton Model with Intransitive Competition

A paper by Laird & Schamp from the year 2006 [12] investigates the potential for intransitivity of varying levels in the community competition structure to promote species coexistence. The full range, beginning with strict transitive hierarchies, up to maximally intransitive loops, is considered for multi-species communities. Consistent with their hypothesis, higher levels of intransitive competition were found to correlate with longer persistence times.

The model used by Laird & Schamp is an individual based, spatially explicit, cellular automaton with local competition. This model is also zero-sum based and does not feature speciation or immigration events, so that eventual all but one species are expected to go extinct, even in otherwise very stable systems. The competitive network is given in the form of a competition matrix, which defines pairwise competitive relations for all species combinations. Every relation is strictly unidirectional and deterministic, meaning that there is always a clear winner between species pairs. The initial state of the model is generated by placing individuals with random species, out of the defined species-pool, into the cells. Then the model is updated iteratively, with each iteration consisting of a single competition event: One random individual is chosen and its direct 3x3 neighborhood examined. If there are individuals in the neighborhood which are competitively advantaged, according to the competition matrix, a random one of them will replace the chosen individual. If multiple of the advantaged individuals belong to the same species, the chances of that particular species to reproduce are consequently higher. One

generation for this model is defined to last for as many competition events as there are cells present. We later use the results produced by this model to validate our implementation of the intransitivity coexistence mechanism.

The model was advanced for 500 generations in each respective experiment, examining the remaining species diversity and competitive network on each successive generation. More intransitive communities proved to be significantly more stable than highly transitive hierarchies.

2.6 Summary

In this chapter we first introduced the terminology and basic principles used in context of this thesis. Niche theory and neutral theory, the two major foundational theories in regards to species coexistence, were described afterwards. While niche theory assumes coexistence is a result of intricate species differences and niche specialization, neutral theory abstracts these differences away and explains population dynamics as a purely stochastic process. Irrespective of either theory, there are multiple well documented coexistence mechanisms present in natural systems, which promote biodiversity and are of central interest in this thesis. Intransitivity for example supports coexistence through non-hierarchical competition networks, analogous to rock-paper-scissors. We briefly introduced six of these mechanisms. Finally, three relevant contemporary models of species coexistence were reviewed: the classic neutral model, a neutral model with integrated storage effect and a cellular automaton with intransitivity. Our model is inspired by these and the last two also serve as a reference for comparison in some of our experiments. With this foundation laid down, the introduction of our proposed multi-mechanism model follows in the next chapter.

3 Concept of the Multi-Mechanism Model

This chapter is dedicated to the concept of the multi-mechanism model, which is one of the main contributions of this thesis. It highlights how the model works in principle, leaving implementation details intentionally open. The process used to obtain the final model is outlined, including the rationale behind the choices that were made. First we define the design goals for this model, a set of guiding principles. Afterwards the key challenges associated with the design process are briefly highlighted. Based on our goals and constraints the scope and functional requirements are then derived. Finally, the model itself is explained: First the basic neutral model in its adapted form, and then the ways with which the different coexistence mechanisms can be integrated into this base model.

3.1 Model Design Goals

Following we describe the main design goals for the multi-mechanism model. These goals are given in the form of idealistic characteristics, which a perfect realization should exhibit. This standard may not be achieved, but gives something to strive for and highlights where improvements can still be made.

Universal

The model should be able to represent all kinds of abstract animal and plant populations, including their basic interactions, simply by altering the already existing parameters to achieve the desired effect. Specific empirical data is not required to generate meaningful insight, but may aid in finding realistic sets of parameters.

Unified

All coexistence mechanisms realized in the model act simultaneously and may influence each other. The model is a single encompassing entity, not an accumulation of independent components which are mathematically stacked on top of each other.

Separable

While the different coexistence mechanisms operate in an interleaved fashion, their associated parameter interface should be as separate and well defined as possible. All the different mechanisms should have their own parameter set associated with them and be on an opt-in basis, meaning that the default behavior for each effect is to effectively be turned off, until specified otherwise.

Comparable

Each mechanism, taken individually, replicates its recognized and described key-characteristics in scientific literature. Furthermore, the model taken as a whole produces output on a similar scale to that of other models with related intent.

Extendable

The model should be designed in such a way, that coexistence mechanisms can be added, removed or altered in their implementation without having to come up with a new base model and without the need for major changes of it.

Natural

Abstraction is the whole purpose of building a model, but one should be careful in choosing how far to abstract when explanation of, and insight into, a natural system is desired. A model may generate sensible results, but for the wrong reasons. Therefore, in this thesis we try to keep the implementation fairly close to an intuitive abstraction of nature that manages to generate sensible output nonetheless.

3.2 Key Challenges

There were several key challenges associated with the design process of the model for it to eventually function as desired. Most notably the high inherent complexity due to modeling several coexistence mechanisms explicitly, including their own parameter sets, while trying to keep everything concise. With the design goal of extendability in mind, this problem was only accentuated more. These mechanisms are further required to naturally fit within this unified model at a comparatively fine resolution, so that the desired interaction among them can be achieved. Although inseparable with respect to their combined results, the mechanisms need to remain largely independent from each others parameters too. These challenges made a simulation based model essentially mandatory, to remain flexible enough to incorporate adjustments throughout the process. However, that choice brought a new concern to the table, namely that of performance and how to keep such a complex model from becoming too computationally intensive to be feasible. Interestingly this issue is closely related to another major one: how to avoid redundancy and condense the model and its mechanisms down to only the absolutely required complexity. Both of these problems needed to be considered during the models design, as well as its implementation.

3.3 Scope & Requirements

The created model should follow the design goals as closely as possible, but due to their idealistic nature, they cannot realistically all be achieved to the fullest extent. Furthermore, there were time constraints associated with this work, which made some compromises unavoidable. For this reason the goals themselves are not suited to act as the actual functional requirements. Instead in this section we translate these idealistic standards to a realistically achievable concept by taking the goals, research questions and constraints of reality into account. These requirements state what the model needs to achieve:

Species Scalability

The model has no arbitrary limit as to how many species can be considered, it is explicitly multi-species friendly.

Interacting Mechanisms

A minimum of three interacting coexistence mechanisms are possible. In this case the storage effect, intransitivity and resource partitioning were chosen, although the others still remain of interest.

Switch-ability

Each of the coexistence mechanisms is on an opt-in basis, meaning that their effects can be essentially deactivated if desired.

Independence

Each coexistence mechanism can work independently from all the others.

Mix-ability

All the different possible combinations of individual coexistence mechanisms are valid and supported.

Clear Interfaces

Each coexistence mechanism has a clearly defined interface with minimal interference from other parameters.

Internal Consistency

One of the models main concerns is internal consistency, meaning that the relations within the model itself should be close to their natural analog. On the other hand how exactly the model in its entirety relates to nature is undefined. It follows that there is not necessarily a way to adequately translate a real life scenario to the model or vice versa.

Nature Inspired

All of the mechanisms should function in a way that is reminiscent to the actual natural counterpart.

Spatially Explicit

The model is spatially explicit to ensure that a wide range of spatial effects can be modeled in a natural way.

Individual Based

The model is individual based to properly capture the interaction of multiple mechanisms. Additionally this allows adding of potentially important individual level variations [2].

Neutrality

The model adheres to neutral theory whenever feasible to reduce complexity and increase comparability. Although neutrality in and of itself is not a primary concern.

Neutral Base Model

There is a neutral base model into which the coexistence mechanisms can be integrated as desired. This ensures high modularity and establishes a baseline to compare the effects of the coexistence mechanisms against.

Possible Species Differences

Although firmly rooted in neutral theory the model allows for functionally different species to study the impacts of the various mechanisms in asymmetric scenarios. Certain mechanisms may also require some level of species variance.

Zero-Sum Communities

Zero-sum assumption for community sizes to remain comparable to most other models.

3.4 The Extended Neutral Base Model

It is arguable which parts exactly of the model belong to its neutral base and which do not. However, the distinction helps from a descriptive point of view. In this case we define the base as all those parts of the model, which can be used extensively without invoking any particular coexistence mechanism.

The basis of our model is derived from the classic neutral model, adapted to allow for the straightforward integration of several coexistence mechanisms. For this study the simulation of a metacommunity is of little value, since

immigration in and of itself cannot be considered a coexistence mechanism [18], just like speciation. As shown by numerous authors it is sufficient to simulate just the local community to measure a systems stability [12, 9, 22]. Therefore the model was chosen to solely focus on the local community and the interactions therein. The question then was, how to translate the classic neutral models dynamics to a spatially explicit simulation which can support several distinct and interacting coexistence mechanisms. In order to find a satisfying answer to this question, the coexistence mechanisms themselves and their respective requirements had to be kept in mind. By taking inspiration from nature and already existing models, consistent with the defined goals and reviewed literature, a solution was found. Compared to the classic neutral model there are essentially five major additions or changes:

1. Environmental fluctuations which affect species fitness
2. Spatially explicit simulation environment
3. Explicit modeling of physical time
4. Explicit modeling of resources
5. Explicit modeling of competition

These form the central elements of the model and are essential in integrating all the different coexistence mechanisms. The classic neutral model can be adapted step by step with these concepts, eventually forming the neutral base for everything else. Following each of the mentioned steps are described.

3.4.1 Environmental Fluctuations

As a first step the model is extended by incorporating environmental fluctuations, which modify species fitness. The resulting model is extremely similar to the one proposed by Danino et al. in their publication from 2016 [9], which was also covered in our literature review. It is the second variant of the neutral model with integrated storage effect, described in Section 2.5.2. In essence this model is equal to just the local community of the classic neutral model with added fitness fluctuations, which cause an individual of a species with higher fitness to be proportionally more likely to survive and reproduce than one of a disadvantaged species. Figure 3.1 illustrates as a simple example how such fitness fluctuations may look like when plotted over time for two

species.




Figure 3.1: Adding environmental fluctuations to the model causes species fitness to fluctuate over time. This example is not representative of our actually used fitness functions but illustrates the point. Different colors represent different species.

The environmental fluctuations themselves, in case of the two-species model by Danino et al. [9], are modeled by a very simple correlated dichotomous noise function. This means that each species can only attain one of two possible discrete fitness values: a high one or a low one. The difference between those fitness values is the “environmental fluctuation amplitude”. The function is correlated in the sense that if species one has the high fitness value, species two always has the low one. They are never equal unless the amplitude itself is zero in which case the environmental stochasticity is for all intents and purposes nonexistent. With a certain chance per time step the fitness values switch for both species so that the formerly disadvantaged species becomes favored and vice versa. The chance for this event is governed by the “environmental persistence time”, which is measured in units of generations. Because generations become longer with increasing community size, the environmental fluctuations are proportional to the size of the simulated community with a prefactor.

For our purposes the proposed dichotomous noise function by Danino et al. is rather unsuited, because that function is limited to two-species scenarios only. Furthermore, it is highly simplistic and more sophisticated alternatives might be desired. For these reasons other functions needed to be considered, although we still used this one for direct comparisons between the models. The function we ultimately used for other scenarios is very similar to the one outlined above and uses the same two parameters. However, the concept of our multi-mechanism model does not hinge on any particular solution here. Therefore the specific function to determine the environmental fluctuations is intentionally not defined in this conceptualization. The implementation chapter fills in these details.

3.4.2 Spatially Explicit Environment

The next step is to adapt the simulation to a spatially explicit format while keeping everything else unchanged. Taken by itself, this does not functionally impact the simulation in any way, but lays the necessary groundwork for what follows. Spatially explicit variations of the neutral model, or other very similar ones, already exist in the form of cellular automata. With our model we decided to go a step further by introducing a continuous spatial environment which offers even more resolution and conceptual freedom.

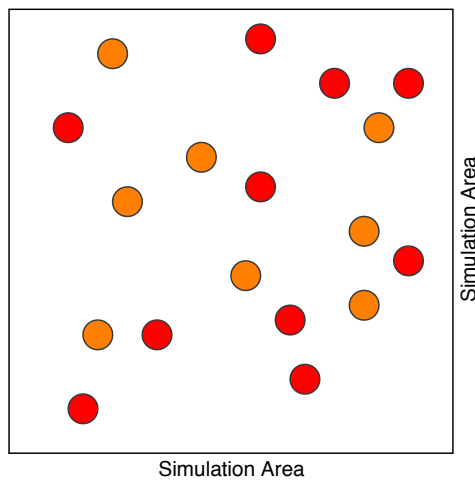


Figure 3.2: Illustration of the neutral base model with an explicit, continuous, spatial environment. Individuals are initially placed randomly into the area. Different colors correspond to different species.

To keep it simple the environment is two-dimensional and its outer boundaries are always defined by a square with side length equal to the parameter “simulation area”. We do not use periodic boundaries, so that boundary effects are to be expected. Since they affect all simulated individuals equally, these effects should not be a determining factor of species coexistence. The individuals themselves are modeled as points or particles within this square, initially placed randomly as illustrated by Figure 3.2. No physical properties such as mass or size are considered. Each individual can move freely within space with a certain speed. This value is the product of the individuals fitness and its “base speed”. Although until now there is no incentive for the individuals to actually move, there is nothing to gain by doing so. Note also that movement is necessarily tied to the time steps of the model, which

correlate directly to deaths of simulated individuals.

3.4.3 Explicit Modeling of Physical Time

The model is currently missing a concept of physical, continuous time. This is problematic due to two main reasons:

1. The resolution of movement becomes highly dependent on the simulated population sizes. For very large populations the movement may actually be fine enough to allow an individual to move in a somewhat realistic fashion, but for smaller population sizes this is certainly not the case: Regardless of population size, the smallest possible “time” unit is equal to the time it takes for the next individual to die. Looking at it from another angle, the amount of times an individual may move before it dies is directly proportional to the size of the simulated community. This is nonsensical and may significantly skew the models results.
2. Having movement tied to steps which represent an arbitrary and unknown amount of physical time is unnatural and makes speed values difficult to judge and compare.

In order to solve these problems a concept of physical time has to be modeled explicitly, the benefits of this are illustrated by Figure 3.3. This characteristic is atypical for these models and makes a few further changes necessary. At this point it should be emphasized however, that the modeled time is not intended to represent realistic time scales. It is purely meant to give the model an independent temporal structure, which may very well only be valid in context of the model itself. For convenience the time is deemed to be measured in “seconds”, but that is not intended to imply comparability to realistic time scales in natural systems.

Most importantly the death and birth of individuals can no longer sensibly be tied directly to time steps. Instead the death chance of each individual is now defined on a per-second basis and all individuals are independently checked for death at the end of each time step. For the time being the mortality rates are constant and fitness-independent, but this will be changed in the next steps. The zero-sum assumption is kept, but because multiple deaths per time step




Figure 3.3: Illustration of the fine and consistent movement possible when modeling an explicit physical concept of time.

are now a possibility, one has to be careful not to introduce a systematic bias due to the execution order of deaths and births. For example one such bias would be if the deaths of species *A* would always happen before the deaths of species *B*, which would advantage *B* because it has more opportunities to grow in abundance. Section 4.3 of the implementation chapter explains this problem in more detail.

Another complication lies in the definition of environmental fluctuations as described previously. The persistence time of the environment was measured in generations of the community, which was only possible because deaths occurred at a clearly defined rate of one per time step. With the aforementioned change to this, deaths are however no longer tied to time steps and occur unpredictably. Keeping the environmental persistence time defined on a per generation basis, when the time for a generation is unknown, obviously is problematic. Additionally all individuals are now simultaneously affected by the environmental fluctuations, since all are updated and checked for death in every time step, unlike before. This makes the implicit scaling of environmental stability with community size counterproductive. As a result the environmental persistence time is simply measured in seconds now instead of generations.

So far this step represents the biggest deviation from previously existing models. Up until now the local community dynamics of the classic neutral model could

be recovered to the fullest extent. With the inclusion of an explicit concept of time this is not quite the case anymore.

3.4.4 Explicit Modeling of Resources

So now that a concept of time exists and the individuals can move adequately through space, it is time to give them a reason to actually do so. This reason comes in the form of explicit resource modeling, meaning that individuals will need to consume food in order to survive, just like in natural systems.

To achieve this the simulation is divided into two distinct communities, similar to the classic neutral model which is divided into a local and a metacommunity. In this case the first community represents the “active” species and individuals, while the second one is the “environmental community” which represents the resources within the environment. The individuals which make up this environmental community solely exist to be consumed by those of the active community. Without exception all resource individuals belong to the same, generic food species for now. Both communities share the same simulation area, so that the active community has access to the resources. Members of the environmental community have no active component, so they are not updated in each time step like the active individuals are. This means that they cannot move, eat, or die through the passage of time. Both communities are independently zero-sum based. So if a resource individual is consumed (the only way in which it can die), it is registered for death and will also be replaced by another resource individual at the end of the time step. The zero-sum definition for the active community remains unchanged. This addition of resources is illustrated by Figure 3.4.

With the addition of explicit resources, the definition of mortality changes for the active individuals. Formerly the chance to die per second was a constant. With this change the chance to die at the end of a time step now increases continually with the advance of time. Mortality begins at a rate of 0% and then increases until eventually reaching 100%, at which point death is guaranteed. Consuming a resource individual however resets the mortality rate back to 0%. Afterwards it slowly begins rising again. This ensures that there is now a very strong incentive for individuals to move around, because in order to consume a resource individual, one has to move to its position. As a consequence all individuals should now be constantly on the move to acquire resources, which

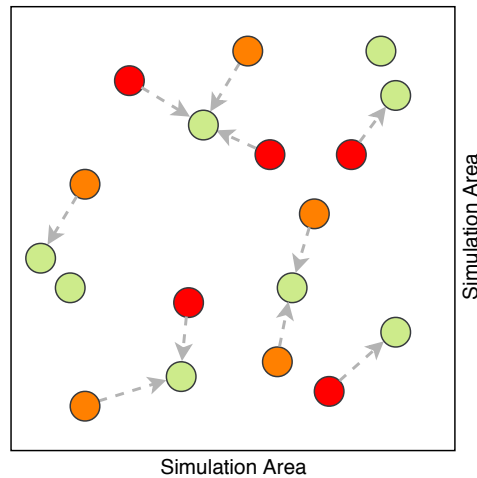


Figure 3.4: Illustration of the added resource community and its members, shown in green. The active individuals in orange and red are now constantly on the move in order to consume as many resources as possible, increasing their survival chances.

takes a certain amount of time. As defined previously environmental fluctuations impact the fitness of a species and the fitness is multiplied with the base speed of an individual to obtain its actual movement speed. Fitter individuals are therefore faster, which means that they will acquire more resources and will therefore be more likely to survive and reproduce.

3.4.5 Explicit Modeling of Competition

With the previous step there is arguably already a concept of competition present. Individuals compete for resources to consume and the faster an individual moves, the better its chances to claim the resource. However, this is a fairly weak model for competition in the sense that it does not allow for very good integration of further concepts which build upon competition. For example it is not possible to define competitive relations between species, or at least they would be very unintuitive. It also necessitates that all simulated, active individuals move around to survive. When simulating plant communities this may be very undesirable, considering their inability to move is important for their population dynamics [7, 25]. For these reasons the way competition is handled gets expanded significantly, shown by Figure 3.5 and explained in the following.

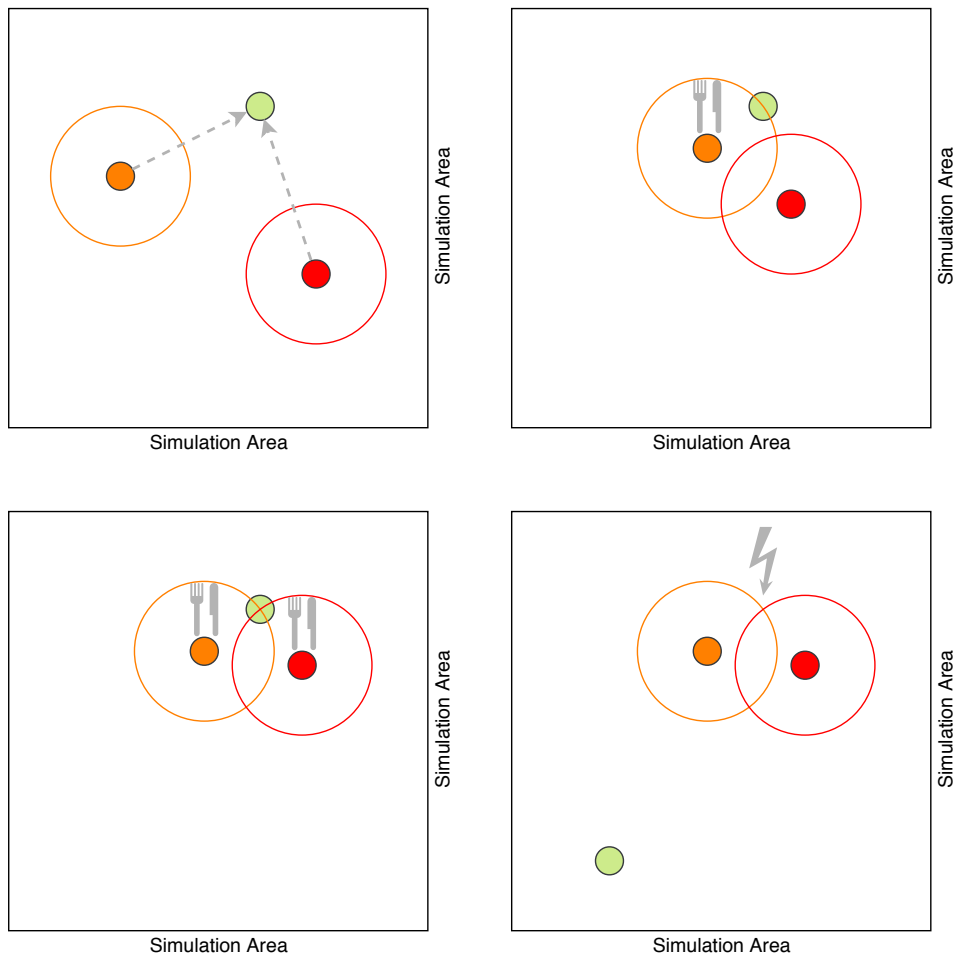


Figure 3.5: Illustration of the competition system. Individuals select resources to consume, moving towards them. As soon as an individual has the target in its interaction radius, the consumption countdown begins. If another individual begins the consumption process before the first finished, competition ensues. Only one competitor will get the food. After the consumption the resource dies and is replaced.

Previously an individual had to move directly onto a resource in order to consume it. With the introduction of an “interaction radius” this is no longer needed. Individuals can now consume resources within a certain radius, without having to get closer. This allows the simulation of slow moving, or even stationary individuals, which have to rely on required resources being in their vicinity. Until this point resource consumption is also still instantaneous, leaving little opportunity to include actual competitive interactions. Unless two individuals happen to target the same resource at the exact same time step, and

it is within both of their interaction radii, competition cannot occur. To solve this problem, and to prevent high interaction radii from giving trivial access to lots of resources, resource consumption is redefined to take a certain amount of time for completion during which the consumer cannot do anything else. Specifically the amount of time it takes to consume a resource, which is within the radius, is equal to the time it would take the individual to move towards the resource. There is an exception though for individuals with a base speed of zero. For them the time of resource consumption is calculated by dividing the distance by their fitness, instead of dividing it by their speed. This is done to accommodate the special case of not moving individuals if desired. Since resources in the interaction radius are not collected instantaneously anymore, there is now a good chance that multiple individuals attempt to collect the same resource. As soon as the first individual among them finishes its consumption process, the competition is resolved and the resource is given to one of the competitors, which is not necessarily the one who triggered the resolving. Victor is the last standing individual after a series of pairwise tournaments between randomly chosen competitors. The outcome of each of those pairwise tournaments in turn is decided deterministically by who has the higher fitness. Losers do not receive any further disadvantages besides having to give up on the resource they were competing for. With this adjustment the neutral base for the integration of various coexistence mechanisms is set.

3.5 Integration of the Coexistence Mechanisms

The neutral base model, which was described in the previous section, is the foundation for the following modifications. While the base model is required for each of these to work, there are no interdependencies between them. So unlike with the previously outlined changes, which were building upon each other to arrive at the base model, these changes can stand on their own. As a result it is not necessary, for example, to include the storage effect mechanism in order to add support for intransitive competition networks. This is very helpful, because it allows the model to remain highly modular with respect to how exactly a certain coexistence mechanism is realized, without having to change the fundamentals of the model as a whole. After all one

should keep in mind, that there are often multiple different ways to achieve them.

Although suggestions for most of the coexistence mechanisms are presented, the focus of this thesis is on the first three: storage effect, intransitivity and resource partitioning. As mentioned before this is necessary due to time constraints.

3.5.1 Storage Effect

As described by Chesson [20] the storage effect demands environmental fluctuations to which each species reacts with distinct fitness fluctuations. The fitter a species, the more competition it needs to experience. And finally the mortality rates of disadvantaged species need to be buffered against the negative impact. The first two requirements are already included with the base model: Environmental fluctuations modify species fitness over time, on a per species basis. And due to the increasing abundance and higher speed of fitter species, they are more often involved in competitive interactions. However, there is no buffering against low fitness values at all: Species at low fitness move slower and are unlikely to win any competitions, so their food intake is greatly hampered. Because the consumption of food resources is directly tied to the mortality rate of an individual, being the only way to reduce mortality, this has strong implications for disadvantaged species as a whole.

The required mortality buffering can however easily be added to allow the storage effect to work. The most intuitive way, which we propose here, is probably to literally let the individuals *store* their food, instead of being required to always consume it the moment it was picked up. This way an individual can build up a buffer during good times in order to survive during periods of low fitness by rationing the food it has still stored. Consequently each species now has an additional parameter associated with it: the maximum amount of food an individual of that species is allowed to store. The higher this value, the bigger the potential buffering effect becomes. Setting the parameter to zero recovers the neutral base model. Of course it is also important to consider the function used to ration the storage. But generally there is no right or wrong answer to that, therefore this conception leaves the question intentionally open, to instead be answered by the particular implementation of the model. Finally, it should be noted that the efficiency of this storage

mechanism is scenario specific to an extent. Most importantly the scenario has to allow for the accumulation of stored food in the first place. If for example food is so scarce, that no buffer can effectively be build up, because it always has to be consumed on the spot, then the storage effect outlined here won't do much in that scenario.

3.5.2 Intransitivity

Intransitivity requires interspecific competition with an impact on mortality. Then a network of pairwise competitive relations between species is needed to define which species is strong against which other species. If the network contains intransitive loops, then the coexistence mechanism can work. Although as shown by Gallien, Zimmermann, Levine and Adler [22], the possibility for varying networks of intransitivity rises quickly in scenarios involving a multitude of species and stable coexistence cannot trivially be predicted from the oversimplified assumption that any kind of intransitive composition will work equally. In the neutral base model interspecific competition is already present, but there is no network of pairwise species relations for competitive ability. The outcome of any competitive event is therefore purely governed by the fluctuating fitness values of the involved species. To change this, we now introduce pairwise competitive species relations in the form of a competition specific fitness multiplier, which is applied whenever two individuals of the given species compete. The multiplier is symmetric in the sense that defining the relation of species A to species B as 2, automatically also defines the relation in the opposite direction as 0.5. This fitness modifier is then only applied to one of the competing individuals, not both. As an example imagine A has currently a fitness of 0.9, while B has a fitness of 1.5. Now two individuals from those species are competing over a common food resource. The relation multiplier can be applied to either one of them, so let's take the individual from A : its fitness is now going to be 1.8, just for this moment, and it will win the competition against the B individual, even though it has a lower general fitness. When setting the competitive relations to 1 for all species, the neutral base model is recovered.

Note that due to the deterministic nature of pairwise competition events in our model, any amount of competitive advantage against another individual is sufficient to always win the competition, provided both individuals have the

same general fitness. This means that environmental fluctuations are necessary in our proposed model to achieve intransitive effects of intermediate strength. Without these fluctuations intransitivity either applies strictly, or not at all.

3.5.3 Resource Partitioning

The aspect of resource partitioning requires that species can evade interspecific competition by specializing on different resources, or through the way these resources are obtained [24]. Up until now however there is only a single common resource available for which all species need to compete equally. There are also no functional differences possible in how these resources may be obtained. An elegant way to solve this problem is to introduce the possibility of multiple different resource species being present in the environmental community, for which each active species can then have varying preferences defined.

The first task, to make multiple resource species possible, seems very straightforward and in many respects it is. But the zero-sum definition for the environmental community should now be reconsidered. Before there was only ever a single resource species present, so the resource abundance for each species was a constant regardless of the selection mechanism used to determine the offspring. This is no longer guaranteed now and using the birth selection mechanism of the active community for the environmental community would cause the resource availabilities to fluctuate and eventually certain resources would even vanish completely. In this case these fluctuations are not desired, as they would be highly influential while simultaneously being somewhat independent from the interactions of the active species. Therefore the zero-sum assumption is extended to cover each individual population within the environmental community, so that the resource availabilities remain constant for each species.

The second task then is to assign the active species preferences for the different resources, so that interspecific competition may be avoided. This can be realized by defining pairwise preferences, from active species to resource, which act as multipliers to the perceived distance of an individual to a resource. For example let's assume there are two resources of different species present: A

and B . An active individual with a preference of 1 for A and 0.5 for B must now decide which resource it is going to consume. The distance to A is 12, and to B is 8. B is therefore closer and more desirable. But the distances are now modified for the selection process based on the preferences as follows: $1/1 \cdot 12 = 12$ for A and $1/0.5 \cdot 8 = 16$ for B . So even though B is objectively closer, it is perceived to be further away by the active individual, due to its preferences, and therefore less desirable. Specifying the preferences of all active species to all resources as 1 recovers the neutral base model.

3.5.4 Spatial Effects

Spatial effects as coexistence mechanisms are generally density based, causing high intraspecific competition and mortality of abundant species due to spatial clustering of their individuals. This allows rare species to evade competition by avoiding the competitors cluster and limits the abundant species effectively. Spatial effects of this form could be added in many different ways to the model. One intuitive way is to change where newborn active individuals are initially placed. So far the birth positions were always chosen randomly from within the whole simulation area. If this position would instead be chosen randomly only from within the interaction radius of the selected parent, natural clustering of the species would occur. One drawback of this method is that the strength of the clustering becomes directly dependent on the relation between mortality rate and speed. The shorter individuals live and the less they can spread out through space within their lifetime, the stronger the clustering effect becomes. If the mortality rate is low enough in comparison to the speed, so that the initial positions of individuals become largely irrelevant, the neutral base model is approximately recovered. This mechanism is overall outside the scope of this particular study however.

3.5.5 Predation

Predation might be a challenging coexistence mechanism to adequately represent in the model, even though the act of predation itself is fairly straightforward to include: instead of only being able to consume resource species, other active species are now potential resources as well. Just like with resource partitioning before, this also mandates that each active species can be assigned a preference

for certain resources over others. Because the active species move constantly, it might be possible that a predator cannot catch up to its prey. For this reason an individual which was selected as target by a predator is forced to stop any movement until the predation attempt is over, either because the prey was consumed, the predator chose a different target, or the predator died. It is questionable whether the predators should be modeled alongside their prey in the same community however. It might be more appropriate to introduce a third community, similar to the environmental community, specific for the predators to differentiate between predation as a regulating coexistence mechanism from the populations of the active community itself. At this point it should also be noted that neutral theory is generally meant to model the dynamics of species within a similar trophic level [10], which would not be the case here if predators were mixed into the active community. Since predation does not belong to the core aspects of this thesis, the problem is not further explored from here on out, but it is important to acknowledge it for future work.

3.5.6 Mutualism

Mutualism as a coexistence mechanism requires that certain species influence the mortality or recruitment rates of other species positively. The model proposed here is strongly rooted in antagonistic relations and how this coexistence mechanism might be included is not obvious. One possibility would be to enable the sharing of a food resource between species that are marked as having positive interactions. So if two individuals compete over a food resource and those two individuals are from species with a positive relation, both could get the resource in the end, instead of only one. This would make it much easier for a low fitness species to still be able to accumulate food, and the more abundant the other species becomes, the higher the benefit would be. A disadvantage of this approach might be that it operates on a fairly coarse scale and cannot be regulated in its influence, it's either on or off. A solution to this could be to model food on a finer scale, so that the consumption of a resource individual does not necessarily yield 1 food unit, but an arbitrary amount which could be subdivided between the individuals through some rule. But once again this mechanism is not a main topic of this thesis and the implications of its design aren't explored further beyond this point.

3.6 Concept Summary

At this point we briefly recapture our proposed concept for the multi-mechanism model. The individual steps describing how to obtain the end result and the reasoning behind those steps were outlined just before, therefore this summary focuses only on the final result.

3.6.1 Neutral Base Model

This section reiterates all the changes which were made to the classic neutral model (as described in Section 2.5.1) in a condensed manner. The result is a neutral base model suited for integration of the various coexistence mechanisms.

Structural Overview

The base model consists of two communities: The active community, which is analogous to the local community in the classic neutral model, and the environmental community, which harbors the resources for consumption by the active individuals. Both communities share the same two-dimensional square space, which is defined by the length of its sides, the “simulation area”. Each individual is represented in that space as a dimensionless particle. The active community can be set up to contain an arbitrary amount of species, the environmental community on the other hand only consists of a single dedicated resource species. All species, without exception, can have an arbitrary amount of members. The only goal for the active individuals in the simulation is to consume as many resource individuals as possible. When multiple individuals want to consume the same resource, competition over it ensues. The environmental community, containing only resource individuals, is purely passive in nature and has no active component. The fitness values of the active species change over time due to environmental fluctuations. Both of the communities are independently zero-sum based, meaning that whenever one of the individuals in them dies, it is going to be replaced by a new one. The births and deaths are executed at the end of each time step: first, for every individual flagged as dead, a birth event is triggered. Then all the registered deaths are realized.

Updating Process

It is expedient to divide the updating process into several phases, not only for descriptive purposes, but also to avoid unwanted side effects, such as a bias towards certain species due to a fixed execution order (more on this in Section 4.3). Each phase should be completed for all individuals, or species respectively, before proceeding to the next. Following these phases are described in their execution order. Note that the updating process for the individuals only considers the active community, because the environmental community is of a purely reactive nature. It doesn't do anything by its own: the resource individuals can't move or eat and don't die through the passage of time. Unlike in many other models, every single individual of the active community is updated in every single time step.

Updating the Environmental Fluctuations

The first step in advancing the model is to update the environmental fluctuations for each species. How exactly this step proceeds depends on the underlying function which was chosen to represent the environmental fluctuations. Either way the function is evaluated and every species of the active community gets its fitness altered according to the result. A straightforward example would be a random walk of the fitness value over time, which is evaluated independently for each species.

Moving the Individuals

Next up is the movement of active individuals. If an individual has not already chosen a resource as target, it will need to choose one now. The function used to determine the target for each individual is an implementation detail, so it is left open here. Unless the individual happens to be close enough to the target, so that it is already in its interaction radius, it moves towards it. Finally, if the target is within the interaction radius at the end of this step, the individual initializes a consumption event for the target resource. Finalizing the consumption takes as much time, as it would take to move the rest of the way towards the resource. An exception from this are individuals with a speed of 0: in their case the fitness is used instead of their speed to determine the time needed for consumption. Until the consumption is complete the individual

cannot select another target and therefore cannot move or begin consuming other resources.

Resolving Competitions & Survival Checks

In this step any ongoing consumption events are completed, if their countdown has reached zero. When multiple individuals were consuming the same resource, this resolves the event simultaneously for all of them, regardless of their specific consumption timers. The winner of the ensuing competition is decided by a series of pairwise tournaments between randomly chosen competitors, until only one remains. The tournaments are decided in a deterministic manner by who has the higher fitness value. Losers of the competition gain nothing, while the winner gets its mortality rate reset to zero. The consumed resource individual is being registered for death.

The last thing to do in this step is to determine whether the current active individual should survive the time step, or die. The death chance is given per second in form of the mortality rate, a function which is left to be detailed by the implementation. If the individual was chosen to die, it is registered for death.

Resolving Birth & Death Events

As the last step in the updating process, any previously registered deaths are resolved. To keep the communities zero-sum based, every death must be countered by a functionally simultaneous birth. To ensure that the ordering of deaths plays no role, the reproductive events are handled first: For every death that will later be realized, a random individual from the respective community is selected with proportionally higher chances the higher its fitness is. The selected individual then reproduces, which means that a new individual of the chosen individuals species is being spawned at a random location in the simulation area. Once all birth events are done, the death events are handled: Every individual which is registered for death is being removed from the simulation. With this the updating process of the neutral base model is complete.

3.6.2 Coexistence Mechanisms

With the neutral base model in place, the various coexistence mechanisms can now be integrated to achieve the desired multi-mechanism model. We now briefly recapture one possibility per mechanism. Keeping in mind that these mechanisms can be added in different ways, the ones outlined here are not necessarily the best ones.

Storage Effect

The storage effect is integrated by allowing individuals to store the resources they collect, instead of having to consume them immediately. The individual can then ration the stored food to survive through periods of low fitness.

Intransitivity

Intransitivity is achieved by introducing pairwise competitive relations between active species. These relations modify the fitness values of competing individuals, such that certain species become advantaged against others.

Resource Partitioning

Resource Partitioning requires the presence of multiple different species in the environmental community. Food preferences for each resource type can then be defined for every active species. The defined preferences modify the perceived distances towards resources in the environment, so that individuals are inclined to select preferred resources, even if they are further away. All resource species populations are independently zero-sum based, therefore resource availabilities are constant throughout the simulation.

Spatial Effects

Spatial effects are introduced by modifying the initial placement of new individuals: Instead of placing offspring randomly within the simulation area, they are placed randomly only within the interaction radius of their parent.

Predation

Predation is modeled by adding a third community of predators. These predators feed on the members of the active community and have preferences for them defined analogous to those outlined for resource partitioning.

Mutualism

Mutualism works by allowing certain species pairs to share food when in competition with one another, instead of fighting over it to the end. The chances for resource sharing can be defined for all species pairs.

4 Implementation

This chapter is dedicated to our specific implementation of the previously outlined concept for a multi-mechanism model. It is this version which was used for all of the experiments we conducted. In our implementation the neutral base model with storage effect, intransitivity, resource partitioning, and spatial effects were realized according to the concept. The focus is however exclusively on the first three, as spatial effects are not further considered and won't effectively be used in the experiments either. It was necessary to omit the fourth effect due to time constraints. Having said that, we explain here how all the details, for which the concept does not postulate a certain solution, were realized. Including, for example, the exact function used to describe the mortality rate of an individual in relation to its food consumption. Additionally some hints in regard to the general code structure and performance optimizations are given when deemed important. The core of our model was written entirely in pure C++, but an interface to the Unreal Engine 4 [28] was created as a quick way to visualize and debug the simulation, which proved to be very helpful in the long run. There are no external dependencies of the core model on any libraries or codebases however, including the previously mentioned Unreal Engine.

Figure 4.1 shows a screenshot of how the model looks like when visualized with the unreal engine. Only the central viewport is really important here, the surroundings are engine tools which we don't use for visualization. On the very left of this viewport, above the black background, are several statistics displayed for the currently visualized model. In this screenshot all of them are 0, because the initial state of the model is shown. The model itself is displayed atop the checkerboard background, which is slightly smaller than the actual simulation area and is only there for visualization purposes. The different species are color coded, green dots are resource individuals, yellow and orange-red are active individuals. The circles correspond to the interaction

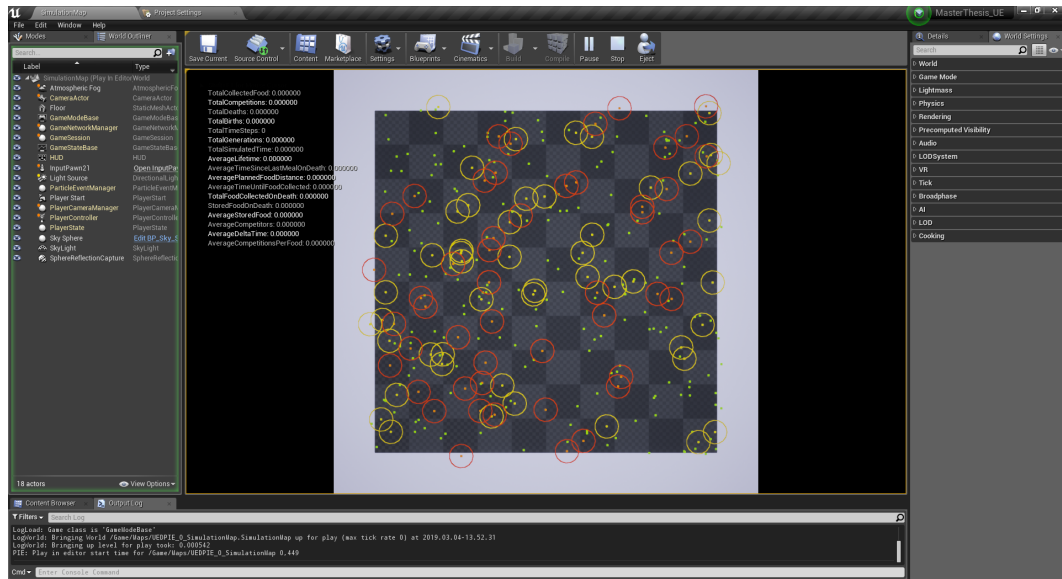


Figure 4.1: Screenshot of our unreal engine visualization for the multi-mechanism model.

radii of the active individuals. The visualized model can then be advanced in realtime, or step by step, to see how it unfolds.

4.1 Program Flow

In this section we give a high-level overview of the updating process for the model as we implemented it. The pseudo code presented here focuses explicitly on the general program flow and is not intended to be all-encompassing or representative of our actual code. Things like error handling, language specific constructs and special cases are omitted for clarity. Additionally we do not go down into the details of many individual functions, such as the target selection, to remain concise. Everything shown here was either already covered by the concept (Chapter 3), or will be explained in the next section, which is dedicated to our implementation specific details (Section 4.2).

Algorithm 1 shows one updating step of the model from the highest level of abstraction. All individuals of the active species go through three distinct phases: The first is to plan what they want to do in this time step and to report how long that would take. The model then chooses to advance up to the

Algorithm 1: A single update step of the model.

```

smallestDeltaTime = Infinite

foreach species ∈ activeSpecies do
  foreach individual ∈ species do
    requestedDeltaTime = individual.Plan()
    if requestedDeltaTime < smallestDeltaTime then
      smallestDeltaTime = requestedDeltaTime
    end
  end
end

foreach species ∈ activeSpecies do
  foreach individual ∈ species do
    individual.Act(smallestDeltaTime)
  end
end

foreach species ∈ activeSpecies do
  foreach individual ∈ species do
    individual.Resolve()
  end
end

BirthEvents()           // One for each registered death
DeathEvents()           // Removing those marked as dead
UpdateEnvironment()     // Fitness Fluctuations

```

closest reported event in the second phase, letting the individuals act out their plan for that amount of time. The third phase then ties up the time step by distributing food and marking individuals as dead, when applicable. Finally, for every registered death a birth event takes place. Then the deaths are actualized and the environmental fluctuations modify species fitness for the next time step.

Algorithm 2 outlines the planning step for an individual in more detail. It mainly consists of selecting a new target, in case the individual doesn't already have one and is roaming around. Afterwards the individual reports how long it would take to get that target into its interaction radius through movement. If,

Algorithm 2: Planning step for one individual.

```
// When currently on search for food
if state == roaming then
    if target == none then
        | target = TargetSelection()
    end
    return TimeTowards(target)
// When currently participating in a consumption event
else if state == eating then
    | return consumptionCountdown // Determined in the acting step
end
```

however, the individual is already locked into a consumption event, it simply reports how much time is left on the countdown until the event can be resolved.

Algorithm 3: Acting step for one individual.

```
if state == roaming then
    | MoveTowards(target, speed, deltaTime)
    if DistanceTowards(target) < interactionRadius then
        | state = eating
        | target.RegisterConsumer(this)
        | consumptionCountdown = DistanceTowards(target) / speed
    end
else if state == eating then
    | consumptionCountdown = consumptionCountdown - deltaTime
end
```

The acting phase follows, and as shown by 3, if the individual is roaming, it moves towards the target in accordance with its speed and how much time passed. If that brought the target resource into the individuals interaction radius, it will initialize the consumption process. If the individual was in a consumption process already, the countdown is simply reduced here.

Algorithm 4: Resolving step for one individual.

```
if state == eating then
  if consumptionCountdown <= 0 then
    /* Triggers competition in case of multiple registered
       consumers. The target is marked as dead and
       notifies consumers of the results. They all go back
       to roaming, target set to none, and the winner gets
       food to eat or store.                                     */
    target.ResolveConsumptionEvent()
  end
end

SurvivalCheck()      // Marks individual as dead if not passed
```

Algorithm 4 outlines the final phase, which consists of finalizing any ongoing consumptions, if their countdown reached 0, and a survival check for every individual.

4.2 Implementation Details

On the following pages we describe the details of our implementation for all those functions which are not already defined in their entirety by the concept.

4.2.1 Environmental Fluctuation Functions

The first implementation details to highlight are the environmental fluctuation functions which we used to change the active species fitness values over time. There are two functions to consider here: First of all the correlated dichotomous noise function as proposed by Danino et al. [9], used only when comparing our model to theirs. And second an independent multi-level noise function, which is an adaptation of the previous function and better suited to our needs. Both functions use the same two parameters: “environmental fluctuation amplitude” and “environmental persistence time”. The functions themselves are referred to as the “stochasticity type” of a scenario.

In our simulation each species has a base fitness of 1 and the environmental fluctuations act as a multiplier on that value. In each case the environmental fluctuation amplitude is then defined as the ratio between the highest attainable fitness multiplier and the lowest attainable fitness multiplier. An amplitude of 2 therefore means, that the highest possible fitness multiplier a species can obtain is twice as high as the lowest one. The maximum and minimum multipliers are therefore calculated as follows:

$$\begin{aligned} \text{MaxMultiplier} &= \sqrt{1 + \text{Amplitude}} \\ \text{MinMultiplier} &= \frac{1}{\text{MaxMultiplier}} \end{aligned}$$

This would yield a maximum multiplier of about 1.41 and a minimum multiplier of about 0.7 for the previous example with an environmental fluctuation amplitude of 2. The environmental persistence time on the other hand is given in seconds and determines the chance per time step that the fitness multiplier of a species changes. Since each time step can represent a varying amount of time, called delta time, this has to be taken into account as well:

$$\text{ChancePerSecond} = \max \left\{ 1, \frac{1}{\text{PersistenceTime}} \right\}$$

$$\text{ChanceOfTimeStep} = 1 - (1 - \text{ChancePerSecond})^{\text{DeltaTime}}$$

Note that this definition of the environmental persistence time does not effectively permit values below 1. This is acceptable in our case, because such low values are generally not considered, but it might be too restrictive in certain use cases. Both of the previously mentioned, and in the following presented, functions build upon these definitions.

Correlated Dichotomous Noise

This function only works for two-species scenarios and was already introduced in Section 3.4.1. It is characterized by the correlation between the two species fitness multipliers: When species *A* has the maximum fitness multiplier, species *B* has the minimum one and vice versa. This holds true for every time step of the simulation.

$$\text{Species A FitnessM.} = \begin{cases} \text{MaxMultiplier} & \text{if Species A Advantaged} \\ \text{MinMultiplier} & \text{if Species B Advantaged} \end{cases}$$

$$\text{Species B FitnessM.} = \begin{cases} \text{MaxMultiplier} & \text{if Species B Advantaged} \\ \text{MinMultiplier} & \text{if Species A Advantaged} \end{cases}$$

With a certain chance per time step, as detailed before, the advantaged species gets swapped, so that the fitness multipliers are then also reversed. Although this approach is surprisingly versatile [9], it has the drawback of not supporting scenarios with higher amounts of species. Additionally the resulting fitness values are coarse and predictable, which may not be desired. The function was kept mainly for comparison purposes between our model, and its storage effect implementation, and that of Danino et al. [9].

Independent Multilevel Noise

This function is an adaptation of the former. It is suited for multi-species scenarios and produces fitness fluctuations of varying strength, although still bound by the maximum and minimum values from before. The main change is that the function is no longer correlated between the species. Consequently it is evaluated independently for every species and in every time step. With

the same chance per time step, as defined before, the fitness multiplier of a species changes. If that is the case for a species, the new fitness multiplier is obtained as follows: first it is decided whether it will be greater or smaller than 1, with equal chances. Then the final value is a random one between 1 and the maximum or minimum multiplier respectively.

$$\textit{Species Fitness } M. = \begin{cases} U(1, \textit{MaxMultiplier}) & \text{if } U(0, 2) > 1 \\ U(\textit{MinMultiplier}, 1) & \text{else} \end{cases}$$

The possibility for intermediate fitness multipliers is important here, because the winner in a pairwise competition event is currently defined to always be the one with the higher fitness. Without allowing for intermediate degrees of fitness, the gap between advantaged and disadvantaged species is a constant. The intransitivity coexistence mechanism relies on fitness altering relations in interspecific competition. But with a constant fitness gap, the defined intransitive relations are either strong enough to always work, or never. As a result intermediate levels of intransitive strength become impossible in the presence of constant fitness gaps. If however the fitness gap is not a constant, but is fluctuating itself, even small amounts of intransitivity will suffice in some cases. This is how we did it in our implementation, but of course this problem could have also been alleviated by changing the competition tournaments to not be deterministic.

4.2.2 Target Selection Function

Here we describe how the individuals in our simulation select between the different food resources in the environment. It might seem tempting to simply always select the resource which is perceived to be the closest, but this has substantial side effects: Over time the individuals will invariably converge, forming dense clusters which are only broken up due to some of the individuals dying and getting placed elsewhere. This happens because there is nothing to discourage individuals from piling up. Once two individuals are close together they will repeatedly select the same resources as being closest. In order to avoid this, the individuals in our simulation only get to choose between preselected representatives from each resource species, excluding those for which they have no preference bigger than 0 defined. Among the representatives given, they can then choose the perceived closest one. The representatives are selected by

picking among 5 of the closest individuals out of each resource species, with proportionally increased chances to be selected the closer one is to the active individual. In our implementation the 5 is a constant, but other values could feasibly be chosen as well. It is suitable in our case, producing adequate results while allowing significant performance optimizations by using a spatial search structure that doesn't need to consider all resources for every query. Once an individual has selected a target, it will keep on to it, until it no longer exists due to being consumed. This, too, helps with performance by reducing the number of queries, but more importantly it also allows the simulation to use an event based time step system.

4.2.3 Event Based Time Steps

Each time step in the simulation represents a certain amount of time, the delta time, which determines how far the simulation advances in each iteration. This begs the question how to choose an appropriate value for the delta time. Unfortunately the answer to this question is not trivial, because the optimal delta time is not a constant and furthermore very implementation, scenario and situation specific. In our model each active individual can only perform a limited amount of actions per iteration. They can, for example, only move once and cannot finalize a consumption event in the same step it was started. Choosing a delta time which is too coarse then leads to artificially wasted time, which the individuals cannot utilize effectively, but which still influences the mortality rate. Time steps which are very fine on the other hand would cause little wasted potential, but incur an unnecessary amount of computational work, because the model advances very slowly. To eliminate all of these concerns we implemented an event based system, such that every time step is just as long as needed for the next event to happen. This is only possible, because each iteration can be planned through by each individual from the start, there are no interfering unknown factors involved: environmental fluctuations are updated before the planning phase, and only once per iteration. Resource individuals cannot move or act in any way by themselves. And the potential death of an active individual has no bearing on the plan of another active individual, or even itself. At the beginning of each time step every individual then reports how much time it would require to complete the next planned action, which is either a movement towards a resource and the begin of its consumption event, or the finalization of an ongoing consumption event. The model then

advances for the minimum reported time, so that no events are ever skipped, and it does so without wasting computation time on unnecessarily small time advancements.

4.2.4 Mortality Rate Function

This function establishes a connection between the time since an individual has last eaten, or more concise the time since last meal, and the mortality rate per second of said individual. It might also be interesting, and well justified, to not only consider the time since last meal, but also the age of the individual in this function. But for the sake of simplicity we did not choose to go that way. We use a bounded reciprocal function, which is characterized by two parameters: the “mortality span” and the “mortality index”. The x-axis corresponds to the time since last meal and the y-axis corresponds to the death chance per second for that time. The function is constructed in such a way that it always intersects the origin, coordinate $(0, 0)$, and $(mortality\ span, 1)$. Values below 0 on the x-axis are defined to always yield 0 and values above the mortality span are defined to always yield 1, though both of these cases usually do not occur. The mortality index defines how pronounced the curve of the function is. For very high values the function becomes approximately linear, for very low values it approximates a step function. The death chance per second for x-values between 0 and mortality span is therefore described by the following function:

$$DeathChance = \frac{Index}{Span + Index - x} - \frac{Index(Span - x)}{(Span + Index)Span}$$

Figure 4.2 shows a plot of the mortality rate function for a mortality span of 200 and a mortality index of 4. These values correspond to those we used in all of our experiments.

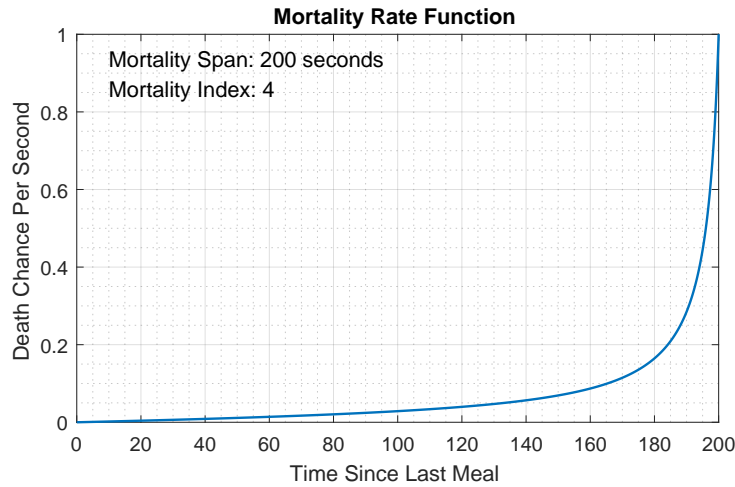


Figure 4.2: Plot of the mortality rate function with a mortality span of 200 and a mortality index of 4.

4.2.5 Survival Check Function

This function is used at the end of each time step, and for every individual, to decide if they will live until the next one or not. The mortality rate is not a constant value, but increases monotonically with increased time since last meal. Because every time step represents some amount of time, the death chance increases accordingly over its course, which should be taken into account. Another thing to keep in mind is that the individual has to be able to interfere with the increasing death chance by utilizing its food storage, if it has one, resetting the mortality rate. We decided to solve this problem by splitting the calculation down into a series of evaluations, each covering one second, with exception of the last, which might cover less if the total time is not evenly dividable by 1. The individual can then interfere before any of those evaluations become effective by utilizing its storage. In the end all of the independent evaluations are accumulated to form the net survival chance of the individual for this time step.

For example: Imagine an individual with a time since last meal of 10 seconds at the beginning of a time step which represents 2.5 seconds. In this case the process would be split into three evaluations: The first would cover the time

since last meal of 10 to 11. The death chance for this interval is approximated simply by evaluating the mortality rate function at 10. The survival chance is then $1 - \text{Death Chance}$, respectively. Since the mortality rate function is defined to yield death chances per second, it meshes well with this approach. With the calculated survival chance in mind, let's say it was 0.9, the individual must then decide to either take the risk or eat from its storage. Assuming it takes the risk here, the second evaluation would then cover the time since last meal of 11 to 12 by the same procedure as before. Assume now that the survival chance of this second became smaller than tolerable and the individual decided to eat from its stored food. With this decision the evaluation must be repeated with updated values, meaning for a time since last meal of 0 to 1, yielding a survival chance of 1, or 100% respectively. The final evaluation would then be for a time since last meal of 1 to 1.5, which is not a full second anymore. But just like before the mortality rate function is evaluated at the beginning of the interval, here at 1. Let's say it yielded 0.99. The decision to eat from storage or not still operates on this value directly, but for the final accumulation of all the independently calculated survival chances, this one has to be adjusted to fit the reduced interval. This is done by taking the survival chance to the power of the reduced time interval, e.g. $\text{Survival Chance}^{0.5}$ in this case. The accumulation itself is then done by simply taking the product of all the calculated survival chances: $0.9 \cdot 1 \cdot 0.99^{0.5} = 0.895$.

4.2.6 Decision Function for Food Consumption

How exactly does an individual determine when to eat from its storage? As is often the case this problem can be answered in a myriad of ways and solutions can become arbitrarily complex. We chose a very minimalist approach: Based on the mortality span of an individual a threshold for the lowest acceptable survival chance per second is set. If the threshold is surpassed, the individual eats from its storage, provided there actually exists some amount of stored food. We call this the “confidence threshold” and it is defined as follows:

$$\text{ConfidenceThreshold} = 1 - \frac{1}{\text{MortalitySpan}}$$

Assuming an individual with a mortality span of 200, this formula yields a confidence threshold of 0.995, or 99.5% respectively. It might seem that this limit is overly conservative, but since it is on a per second basis the chances

add up quickly. To put this into perspective: if this survival rate were a constant, the individual would only have a chance of approximately 50% to become older than 136 seconds. But of course it is not a constant, so well fed individuals are expected to survive longer than that, considering they won't be around the threshold levels of survival rates for long, while those starved of food will experience much harsher survival rates than those of the threshold, because they can't reset them.

The most prominent disadvantage of this method is the complete lack of awareness which individuals have in regards to their situation. As a direct consequence very fit individuals with an abundance of food hold their resources back much more than necessary, suffering worse survival rates than they probably should. The opposite is true as well and unfit individuals will not utilize their resources in the most promising way possible. Evidently this doesn't stop the storage effect from working as seen in the validation experiment (Section 5.2.3), but nonetheless this function might be worth reconsidering in future work.

4.3 Avoiding Biased Results

This topic has been mentioned before, such as in Section 3.6, and we now detail the issue. In general the problem can be summarized as the introduction of unintended advantages for certain species as a result of seemingly harmless implementation details. The cases we want to highlight in particular are those where multiple individuals of different species are involved and need to act one after another. This essentially includes the whole model updating process, but more specifically the handling of deaths and births, competition, movement and so on. The pitfall is characterized by choosing to handle the individuals in a predictable sequential order, even though the ordering does in fact have an influence on the outcome. This is especially true if the ordering derives directly from the index of their respective species. Imagine, for example, that the deaths and births would always be handled in such a manner. Further, imagine a scenario with two species, both having 100 members and equal fitness. Then, in this time step 10 individuals of each species were chosen to die. First the deaths of species *A* are handled and for every death a birth occurs. For the first 10 deaths species *A* can at best break even, although the

chance for that is very small. Likely it is going to lose several members to B , especially because the chances get increasingly skewed in favor of B with every abundance shift towards it. As B finally gets to experience the same 10 deaths, it is in a stronger starting position than A was and has higher chances to retain more members than A could. This gives B a tremendous and systematic advantage, strong enough to force A into extinction in virtually every repetition.

To avoid such systematic biases is paramount to achieve a usable model. Fortunately testing for them is very straightforward and can be done by running a few test cases involving multiple species and purely neutral dynamics. The results should show no bias in favor of any particular species. The best solutions to these cases of bias are to either randomize the ordering on an individual level across species, or to separate the execution order from the effect, which is what we usually chose to do.

4.4 Parameter Reduction & Meta Parameters

One drawback of our proposed multi-mechanism model lies in its inherent complexity and associated parameter richness. For practical reasons it is beneficial to reduce the amount of variables that need consideration to a minimum, focusing on the truly relevant ones. In our case this means that we are ultimately only interested in those variables that have a direct impact on species coexistence without redundancies. Some parameters with equal effects on all species can feasibly be set to a certain value and then be considered constant, especially if they are part of a redundant system. A good example of this would be the size of the simulation area, which has redundancies in form of the speed and interaction radii of the individuals. The redundancies here mean, that doubling the simulation area has largely the same effect as halving the speed and interaction radii of all species would have. As a consequence the simulation area as a parameter can be assumed constant without sacrificing information, although it does alter the absolute scaling of the modeled scenarios, which can however not be ascribed definitive meaning in any case.

Another related problem is, that we are sometimes interested in certain scenarios without knowing exactly which combination of parameters achieves that scenario. It is for example not trivial to determine how the simulation area in combination with species abundances, speed, interaction radii and environmental fluctuations should be chosen to achieve a system with high competitive pressure. This is especially problematic, because we do not actually care for some of those model parameters on an individual basis, but only their combined effect. To solve this problem we introduce two meta-parameters which can be used to drive other model parameters: the “competition ratio”, and the “time to food”. Both of these meta-parameters are to be understood as average values a certain scenario should achieve. The competition ratio describes the relation between food consumption events with and those without any competition. Hence a value of 0.5 would mean that every food consumption has a 50% chance to result in competition between two or more individuals. The time to food on the other hand is defined as the average time it takes an individual to consume a food resource.

The enforcement of these meta-parameters is handled by iterative processes, which operate on a given input scenario. The process then creates new scenarios with the relevant parameters adjusted, until eventually reaching desired values. Then a final optimized scenario can be constructed as output. In each optimization iteration the current scenario is run until enough samples, meaning completed food consumption events, were collected. Then based on the measured competition ratio, or time to food, the scenario is reconstructed with adjusted parameters. For the parameter optimization itself, an adjustment factor is always computed first, which is defined as the desired value divided by the measured value. When optimizing for a specific competition ratio, the abundance of all resource species is then divided by this factor and rounded up, while the interaction radius of all active species is multiplied by the factor. When optimizing for time to food on the other hand, the speed of all active individuals is simply multiplied by the factor. With these adjusted parameters the next optimization iteration is done and so on, until all iterations are finished. The described optimization functions are simple, but generally work very well and allow us to drive more important aspects of the scenarios directly, which could otherwise only be done indirectly through trial and error. However, it should be noted that these functions depend on a somewhat reasonable input scenario to work best and that they obviously cannot achieve logically

impossible scenarios. The enforcement of meta parameters is always the last step in defining a scenario, and the competition ratio is always enforced before the time to food.

4.5 Performance Considerations

Because our model is stochastic in nature and subject to a lot of randomness, scenarios need to be computed many times to obtain results with a reasonable amount of confidence. Combined with the fairly high inherent complexity of our model, this can lead to very long computation times. Code optimization is a topic beyond the scope of this thesis, but we want to briefly highlight the two optimizations we did, which had the biggest payoffs.

The first one was mentioned before in Section 4.2.2 and has to do with the target selection function. With our initial, naive, implementation of the model, individuals would update their targets in every time step by comparing the current distances towards all of the food resources. The function then was responsible for over 90% of the total computation time. We drastically reduced this by introducing a grid-based spatial search structure in which the n -closest food resources to any point in the simulation area can be found with little overhead and on a per species basis. For target selection we then only consider the 5 closest individuals from each resource species and keep selected targets over multiple time steps. With these modifications the function is now far from being the most computationally intensive one.

The second optimization is very straightforward, but nonetheless highly beneficial. As mentioned before each experiment generally consists of many independent simulation runs. These individual simulations are prime candidates for easy and efficient parallelization, taking full advantage of multi-core processors. Therefore all the simulations of our experiments were computed in parallel, reducing computation times significantly once more. But even with these optimizations in place the model remains computationally intensive, which we had to take into account for our experiments.

5 Evaluation

In this chapter we detail the experiments that were conducted, present their results and discuss them. First we give an outline of the general setup that was used for all of our experiments with the multi-mechanism model. Then we explain how the respective experiments and their results are presented, focusing on the formal aspects, and outline the data collection process. Having established the general framework, we can then concentrate on the experiments themselves, showing and discussing their respective setups and results.

5.1 Experimental Setup

In our experiments the simulations are generally run until the first active species goes extinct, unless otherwise mentioned. As outlined in the implementation chapter, certain model parameters can reasonably be assumed constant once a suitable value is found, which is especially true for those which are partly or fully redundant with respect to species coexistence. Other parameters have the same value in most experiments, but may change for specific ones, especially because few experiments actually make use of the full capabilities our model provides, such that parameters are often left in a functionally neutral value. As we want to concentrate only on the defining parameters for each experiment, so that not all have to be listed every time, we provide default values for many parameters in Table 5.1. Therefore if an experiment does not explicitly list the value for a certain parameter, then it is the one listed here. Further, unless otherwise stated all species from the same community are initialized with equal abundances. Each individual simulation we conducted was capped to compute a maximum of 10000 generations. This was necessary to keep simulation times feasible for certain scenarios with potentially enormous coexistence times, this limitation has to be kept in mind as it leads to an artificial ceiling for persistence times.

Table 5.1: Default model parameters, applicable unless we explicitly state otherwise for an experiment.

Parameter	Value
<i>Model parameters.</i>	
Base Fitness	1
Simulation Area	100
Mortality Span	200
Mortality Index	4
Stochasticity Type	Independent Multilevel
E. Fluctuation Amplitude	0
E. Persistence Time	0
Max. Stored Food	0
Resource Species	1
All Resource Preferences	1
All Competitive Relations	1
<i>Experiment parameters.</i>	
Max. Generations per Simulation	10000

We utilize the meta-parameters competition ratio and time to food (explained in Section 4.4) in all of our experiments to create the base scenario onto which the different starting conditions are then applied. The parameter values we list in the experiments always already include the enforced meta-parameters. We list them only for clarity and because they explain how we arrived at certain values. While the competition ratio can be set freely, as we see fit, the time to food value is necessarily fairly low in all of our experiments. This is important in our case, because it is tied to the spatial effect coexistence mechanism (Section 3.5.4) we included in our implementation, but don't actually want to test here. Setting the value low enough enables the active individuals to disperse evenly in space thanks to a long lifetime, making their starting positions largely irrelevant. Any initial clustering effects then become minimal and this coexistence mechanism is turned off for all intents and purposes. However, we cannot completely rule out that there might be some interference left.

5.1.1 Structure & Presentation

In each of our presented experiments we utilize the same general structure. Meaning that we first introduce the specific setup, including some exposition in regards to how and why we conducted this particular experiment. Most notably this includes a table which highlights all the parameters related to that experiment. In most cases we conduct multiple independent runs over which a single aspect is varied to set things into relation. These varying starting conditions for the runs are not specified in the parameter table, as that one only defines the general scenario, but become evident in the diagrams of the results, which are shown and discussed after the experiment is established. Often we use very similar diagrams to visualize the results of our simulations, which are read in very similar fashion. To avoid repetition down the line, we explain the common types of diagrams here by using some examples.

Figure 5.1a shows the most used diagram type. The y-axis indicates how many generations it took for one of the species to go extinct, while the x-axis represents the different runs. A boxplot is given for each of the starting conditions. The central red line corresponds to the median, while bottom and top edges of each box indicate the 25th and 75th percentiles, respectively. The whiskers extend upwards and downwards up to a maximum of 1.5 times the box height, with values beyond that considered outliers, marked as red crosses. These diagrams also feature a line in light blue, indicating the average values. Figures like 5.1b are used in certain cases to highlight one particular case out of the different starting conditions with more detail, effectively elaborating on the previously introduced boxplot diagram. The y-axis shows the frequency with which certain results occurred on a scale of 0 to 1, corresponding to percent, while the x-axis indicates the given persistence times. In these diagrams the average achieved coexistence time is indicated by a vertical red line. The surface plot, as shown by Figure 5.1c, on the other hand is utilized in bivariate experiments, when combining two coexistence mechanism. The changing starting conditions are indicated by the x and y-axis, while the resulting coexistence times are indicated by the height of the dots on the z-axis and their color. Finally, some experiments call for a closer analysis of individual simulations, in which case we plot the species abundances over time for them, Figure 5.1d is an example for this. The x-axis then indicates the generations, while the y-axis corresponds to population sizes. Each species abundance

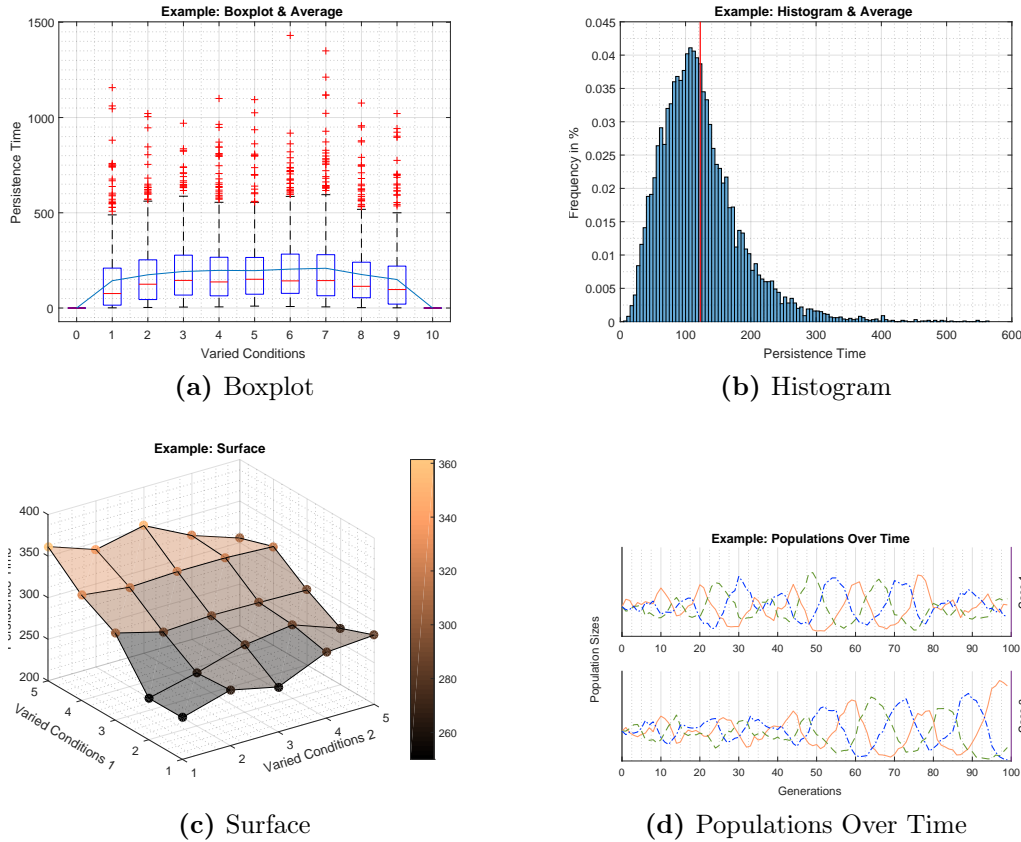


Figure 5.1: These figures show examples of the diagram types we used most frequently in our experiments.

is depicted as a colored line and different conditions of the simulations are annotated on the right.

Another point we want to address briefly here is our choice of using the average persistence times instead of the median whenever this choice had to be made. All of these simulations are characterized by a rather high variance, such that the same scenario might initially report a coexistence time of 10 generations, only to then go on for 500 in the next iteration, without those values even being considered outliers. Additionally the inherent variance and outliers are strongly skewed towards large positive values, consistently pushing the mean higher than the median. Considering that there is a natural lower bound at 0, but no natural ceiling for coexistence times, this isn't particularly surprising. Doesn't the median then paint a more accurate, less biased picture by ignoring those

rare cases of prolonged coexistence? We argue that the opposite is actually true here and the median should indeed be taken with a grain of salt: the main reason for this is, that an unusually long running simulation does not suddenly become less valid than any other, considering its contribution to the overall potential for coexistence of this scenario. Another way to think of this is to imagine a related, but slightly different kind of experiment: suppose we only simulate up to the 50th generation in each simulation. Then we count how many simulations could be finished before a total of 100 extinctions occurred. In this analogy the mean would correspond to counting each and every successful simulation into the final result, while the median would exclude long streaks of successful simulations, on the basis that they are rare and thus shouldn't count. In this analogy it becomes apparent that the latter approach is nonsensical. Why would the fact that streaks of successful simulations are statistically unlikely cause them to not count, considering they are identical to the others? The distinction seems arbitrary. Now admittedly this analogy is not perfect, as it presupposes that simulations could be cut down into chunks of smaller, capped, simulations with identical results, which is not quite the case. But it serves to illustrate the point and is not too far off to be valid.

5.1.2 Collected Data

The statistic we are generally interested in the most for all of the following experiments is the average time until extinction of one or multiple species for a given scenario, called “persistence time”. This metric serves as an indicator for the stability of the respective system, although we do not specifically conduct experiments with lopsided starting conditions, such as invasion scenarios, unless for comparison with other research. This simple metric already captures the core aspects of interest here, but it does not give any further insight as to how the value came to be. For this reason we capture many more statistics and write them out as csv files, usually with one data row per generation of each simulation. This is a somewhat coarse division, but it helps to keep the data manageable. Generally we keep track of all the metrics outlined in Table 5.2. Some of these metrics are not very important, and therefore aren't used much, but together they offer a rather comprehensive picture of what is happening within each simulation and help in validating that everything is working as expected.

Table 5.2: Tracked metrics during the experiments.

Metric	Notes
Average Lifetime	in seconds
Average Time Since Last Meal On Death	in seconds
Average Planned Food Distance	in meters
Average Time Until Food Collected	in seconds
Total Food Collected On Death	average across all individuals
Stored Food On Death	average across all individuals
Average Stored Food	average across all individuals
Average Competitors	per competition event
Average Competitions Per Food	-
Average Delta Time	in seconds
Population Size	of each active species
Total Collected Food	-
Total Competitions	-
Total Deaths	-
Total Births	-
Total Generations	that were simulated
Total Simulated Time	in seconds
Total Time Steps	-

5.2 Validation Experiments

The first array of experiments that was conducted serves to validate the neutral base model and the coexistence mechanisms we decided to focus on. This means in particular that we chose reference models out of the contemporary research, which represent certain aspects, such as the neutral base model with and without environmental fluctuations, or the intransitivity coexistence mechanism, with which we can compare the results of our simulations. In the design process we were aiming to remain comparable to already existing models, in terms of coexistence times, in an attempt to effectively unify multiple of those models. However, certain required design differences predictably prohibit us from completely recovering those models. The question therefore is, just how close we could get and where the differences are. The models we compare against were outlined in Section 2.5, which includes the neutral model with storage effect by Danino et al. [9], that was replicated by us to generate custom

reference scenarios, and the cellular automaton with intransitivity by Laird and Schamp [12], in which case we compare against their example case. Together they cover the neutral base model with and without environmental fluctuations, the storage effect, and intransitivity. The third coexistence mechanism, resource partitioning, on the other hand doesn't seem to have a compelling reference model in current research. Consequently we only compare it internally against the other two mechanisms and our expectations for validation.

5.2.1 Demographic Stochasticity

Here we compare the multi-mechanism model with the first variant of the neutral model with storage effect, as described in Section 2.5.2. This is the most basic scenario possible. There are no active coexistence mechanisms and no environmental fluctuations of any kind. It is a two-species simulation without any functional species differences. This means that in our extended model there is only a single resource species present, for which all species have the same preference of 1, and both species are competitively equivalent (in accordance with the default parameters listed in 5.1). The meta parameter competition ratio was set to an intermediate value, but should have little influence here in any case as all competitions are decided randomly in the absence of fitness modifiers. Time to food is set to a low value to avoid invoking the spatial effect coexistence mechanism in our implementation. Table 5.3 shows the relevant settings we used for this experiment. Any parameter not listed was already defined in Section 5.1.

The starting conditions used in this experiment are equivalent to those used in the paper by Danino et al. [9], and are defined by the abundance with which species 1 and 2 are initialized. A population ratio of 0.1 means that species 1 is initialized with an abundance equal to 10% of the total population size as specified in Table 5.3. The abundance of species 2 is then automatically defined as being equal to the remaining 90%. Varying the initial population sizes like this allows to quantify the stability of the system irrespective of the total achieved persistence times by observing how strongly the starting conditions influence the result. A certain scenario might achieve a very long coexistence time, but only if the initial populations are balanced precisely right, while less balanced populations descend into a monopoly quickly. Additionally the varying starting conditions help in confirming that the simulations work as

Table 5.3: Parameters used for demographic stochasticity validation.

Parameter	Reference	Our Model
<i>Final model parameters for the base scenario.</i>		
Active Species	2	2
Active Community Size	100	100
Resource Community Size	-	161
Base Speed	-	0.9
Interaction Radius	-	3.1
<i>Meta parameters that were used.</i>		
Competition Ratio	-	0.4
Time to Food	-	10
<i>Experiment parameters.</i>		
Simulated Scenarios	11	11
Iterations per Scenario	100000	1000

expected, because the results need to be symmetric around the 50% population ratio for a model with neutral dynamics.

Results

As can be seen in the figures 5.2 and 5.3 our model, for this particular scenario, comes very close to the reference. Even though the results are almost equivalent, there are a few interesting points to address though. First of all it should be noted that our model was simulated with much fewer iterations compared to the reference, which might account for some of the differences. The amount of data we collected should be enough for a reasonably safe approximation of the true persistence times, which is supported by the high symmetry exhibited by Figure 5.2, but when looking at the results in more detail, the relative lack of data points becomes apparent, which is demonstrated well by the histogram. This is a recurring concern with all of our experiments, caused by the relatively high computation times of our model, however the data gathered should suffice for our intents and purposes. Now unto the results themselves: the first unexpected observation is actually just how close the results really are. Not just when considering the average persistence times, but more so the inherent variance in these models. A somewhat broader distribution in the results was expected for the multi-mechanism model as it features an

additional, influential source of variance through the random food distribution within the simulated area. This could theoretically prolong simulations by providing weakened species better access to food, through sheer luck, or it could have the exact opposite effect. Consequently we expected to observe a consistently larger interquartile range, but this is hardly the case. Interesting is also the resulting curve of coexistence means, which has two peculiar aspects to it. Most noticeably there is a small, spiky, peak right around the middle. Phrased differently the curve in our model has a less steady curvature to it, instead the left and right halves seem to have some roughly linear segments to them. This is most visible when comparing the increases in persistence time from population ratio 0.3 to 0.5 for the reference model and ours, although the effect is small. This suggests that our model favors certain population ratios in terms of coexistence ever so slightly. The differences here are however truly minimal, which is also reflected by the histograms which show that even the more detailed distribution is almost equivalent. Overall this experiment was successful in showing that our model is capable of describing neutral dynamics driven by demographic stochasticity in a comparable manner to existing models.

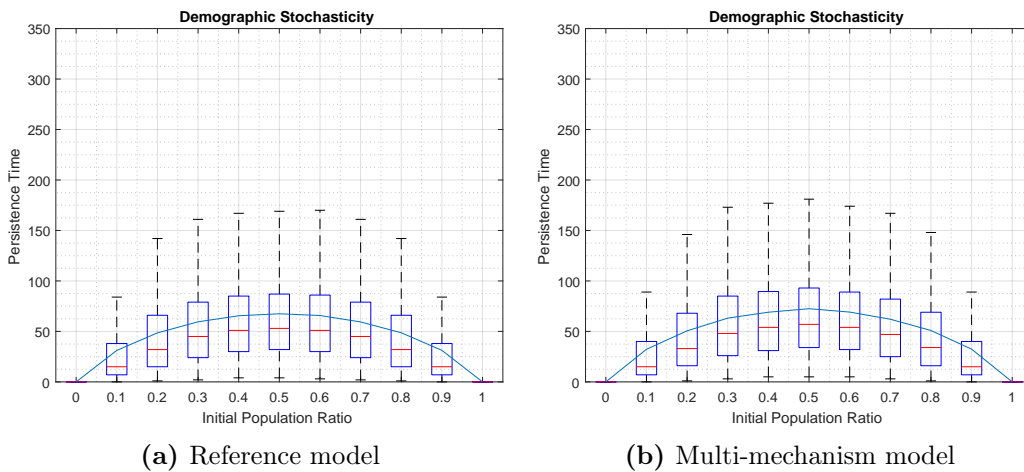


Figure 5.2: These figures show the overall results of the simulations driven by pure demographic stochasticity with respect to the achieved coexistence times. Outliers are omitted in these plots to avoid congestion. Scaling of the y-axis is fixed for easier comparison between these results and those of experiment 5.2.2 and 5.2.3.

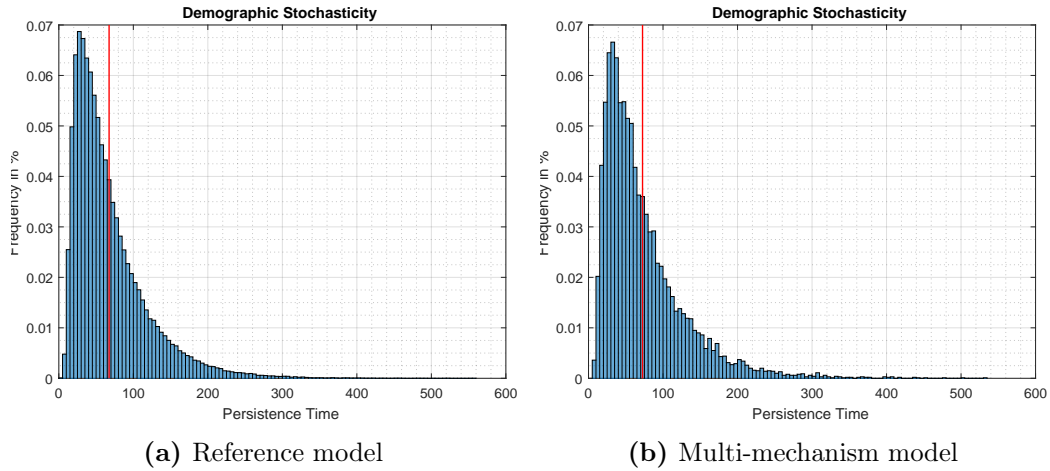


Figure 5.3: This histograms shows the distribution of persistence times which occurred for a population ratio of 0.5 in the simulation runs. It corresponds to the middle boxplot of Figure 5.2.

5.2.2 Environmental Stochasticity

This experiment is set up almost exactly like the one before (Section 5.2.1), but now environmental fluctuations are added to the mix. This means as reference model we now use the second variant of the neutral model with storage effect, as described in Section 2.5.2, and in our multi-mechanism model the environmental fluctuations are turned on. Everything else remains identical. Table 5.4 shows the relevant settings for this case. Note that we use the same environmental fluctuation amplitude as in the reference model, but a different environmental persistence time. This is necessary, because the reference model uses generations as a unit, while we use seconds (explained in Section 3.4.3).

Results

Once again the reference model could be approximated closely, as is evident in figures 5.4 and 5.5. However, there is an important thing to keep in mind here: the environmental fluctuations work fundamentally different between these models and no trivial conversion of the parameters is possible. This remark includes two main aspects. The first would be that environmental persistence time is measured differently, once in generations, once in seconds. The second would be that environmental fluctuations affect every individual simultaneously

Table 5.4: Parameters used for environmental stochasticity validation.

Parameter	Reference	Our Model
<i>Final model parameters for the base scenario.</i>		
Active Species	2	2
Active Community Size	100	100
Resource Community Size	-	154
Stochasticity Type	C. Dichotomous	C. Dichotomous
E. Fluctuation Amplitude	0.5	0.5
E. Persistence Time	0.2	20
Base Speed	-	0.875
Interaction Radius	-	3.24
<i>Meta parameters that were used.</i>		
Competition Ratio	-	0.4
Time to Food	-	10
<i>Experiment parameters.</i>		
Simulated Scenarios	11	11
Iterations per Scenario	100000	1000

in our simulations, while this is not the case for most models, including our reference. Imagine a scenario in which we have two species, each containing 1000 members. Suddenly the environment shifts greatly due to a flood, for an amount of time which usually equals half a generation in each model. One species cannot deal with this radical change and its fitness is essentially reduced to 0, while the other can adapt and remains with a fitness of 1. In our model the entire weakened species is subject to this high environmental stress at the same time. The individuals won't be able to acquire food and their mortality rate will raise dramatically. Long before the flood is over, the species will be extinct. Compare to the other model, in which individuals are only subjected to this stress one after another, like in a queue. If a weakened individual is selected for competition, it too will die. But the chances to be selected in the first place decrease quickly with decreasing abundance, regardless of the environment. This provides a buffering effect against the environmental conditions. Figuratively speaking the weaker individuals sit on top of the fitter ones, which can swim and therefore shield them from drowning in the water. Unless in every single time step one of the weak individuals is chosen for competition, which is exceedingly unlikely, our imaginative flood is not

enough to cause extinction. Due to this fundamental difference our model reacts more sensitively to the environmental conditions, which also becomes evident in a later experiment (Section 5.3.2). The environmental persistence time we chose here was actually the first we tried and it turned out to be a near perfect match, but this can be chalked up to coincidence. Auxiliary experiments have consistently shown that maintaining a fixed ratio of 1 to 100 between the models environmental persistence times is not a general solution. That said, there is fairly little to discuss for this experiment. Analyzing the results shown by 5.4, one can once again identify the less pronounced curvature with a distinct linearity over certain segments in our results, and a small, but noticeable, peak in the middle. This is evidence that the observed characteristics are likely a systematic effect. One other noticeable change is that our model seems to be more robust with respect to lopsided starting conditions. The relative difference between lowest measured persistence time and highest, excluding the outermost trivial cases, is approximately 50% for the reference model, while it is 40% for ours. Overall this experiment succeeds in showing that environmental stochasticity is adequately represented in our model and can produce results highly similar to those of the reference model, although the fundamental differences need to be taken into account.

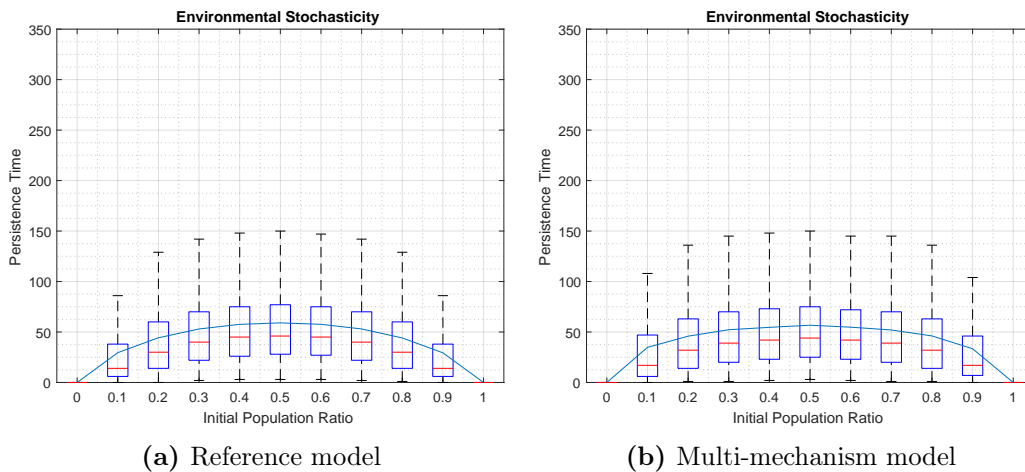


Figure 5.4: These figures show the overall results of the simulations including environmental stochasticity with respect to the achieved coexistence times. Just like in Figure 5.2 outliers are omitted again and the scaling of the y-axis remains fixed for comparison between these results and those of experiment 5.2.1 and 5.2.3.

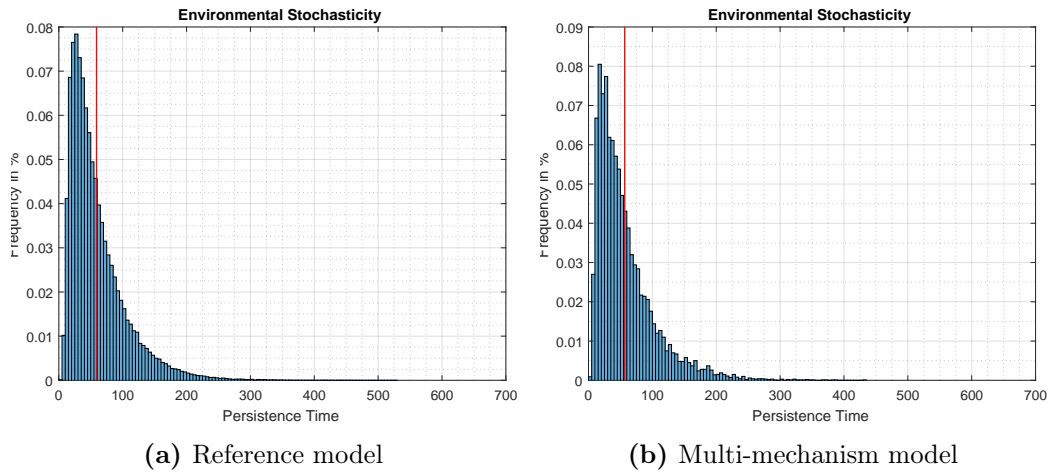


Figure 5.5: This histograms shows the distribution of persistence times which occurred for a population ratio of 0.5 in the simulation runs. It corresponds to the middle boxplot of Figure 5.4.

5.2.3 Storage Effect

Once more the set up for this experiment is highly similar to that before (Section 5.2.2). In this case the third variant of the neutral model described in Section 2.5.2 is used as the reference, now including the storage effect coexistence mechanism. Correspondingly we activated the storage effect in our model too. Aside from these changes the experiment is identical to the one before. Table 5.5 specifies the settings used here. Note that the storage effect is built into the reference model, not as a parameter, but through the model behavior and therefore cannot be adjusted in any way. The parameter listed here consequently only applies to our model.

Results

This experiment confirms all of our previous findings and matches the reference model quite well, although the differences between the models are unsurprisingly higher in this case compared to the previous experiments. That is because the storage effect, as we have implemented it in our model, cannot be adjusted arbitrarily in its strength. Taking this into account the results are actually closer than expected. Also unexpected is that the weakest form of storage which is representable by our model already matches the reference model.

Table 5.5: Parameters used for storage effect validation.

Parameter	Reference	Our Model
<i>Final model parameters for the base scenario.</i>		
Active Species	2	2
Active Community Size	100	100
Resource Community Size	-	161
Stochasticity Type	C. Dichotomous	C. Dichotomous
E. Fluctuation Amplitude	0.5	0.5
E. Persistence Time	0.2	20
Max. Stored Food	-	1
Base Speed	-	0.88
Interaction Radius	-	3.11
<i>Meta parameters that were used.</i>		
Competition Ratio	-	0.4
Time to Food	-	10
<i>Experiment parameters.</i>		
Simulated Scenarios	11	11
Iterations per Scenario	100000	1000

With higher amounts of maximum stored food, persistence times around 200 generations can easily be reached in an otherwise identical scenario. We were skeptical if the storage as we implemented it would really affect the simulation strongly, mainly because the effect is seemingly symmetric. Having the ability to store food and bridge periods of low fitness through that does not only help rare species, but all of them. And because this is a zero-sum model the relation between species is what makes the difference, not the absolute amount of time for which an individual can survive. Of course the fitter species gain a comparatively lower advantage, because they can readily acquire food anyway, such that the storage is not as needed, and due to their abundance and speed they experience more competition too. But it seemed these effects wouldn't be particularly impactful. However, as the figures 5.6 and 5.7 reveal our doubts were unfounded as this implementation is actually more effective than what we aimed for. Very interesting is, that the results once again exhibit the previously identified characteristics, having a visible peak around the middle, less pronounced curvature with almost linear segments around the population ratios 0.3 to 0.5 and 0.5 to 0.7 and a comparatively higher robustness against

lopsided starting conditions with a maximum persistence time difference of 40% in the reference and 34% in our model. More detailed experiments would be needed to identify the exact causes leading to these results. Overall this experiment was very much successful in proving that the storage effect as proposed and implemented works well and can reasonably be compared to the reference model.

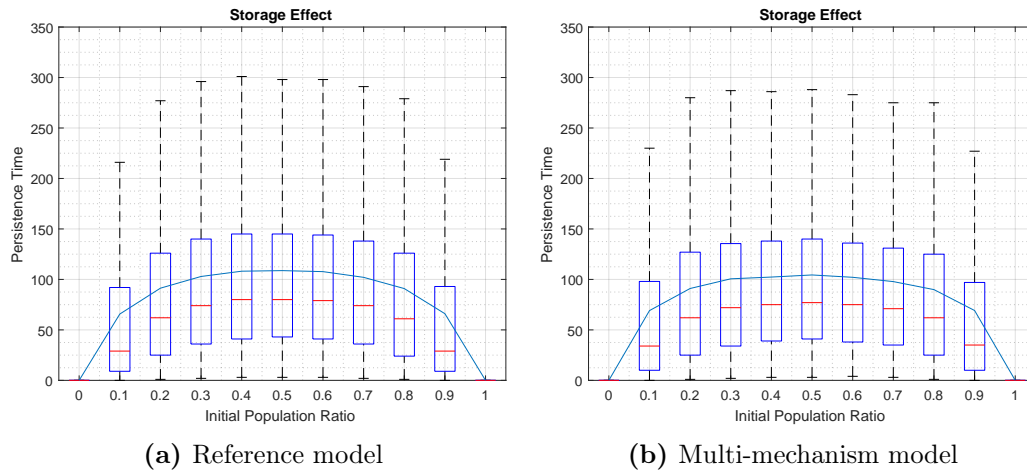


Figure 5.6: These figures show the overall results of the storage effect simulations with respect to the achieved coexistence times. Outliers are omitted again and the scaling of the y-axis remains fixed for comparison between these results and those of experiment 5.2.1 and 5.2.2.

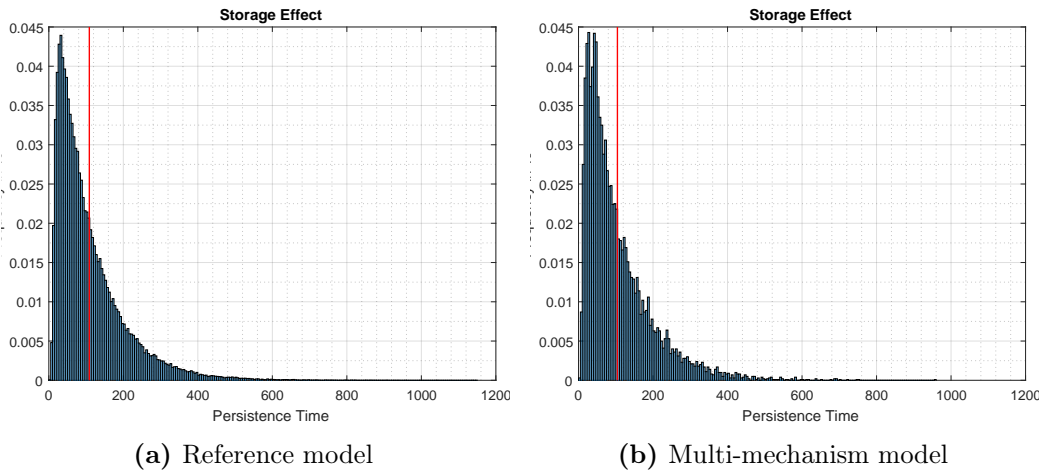


Figure 5.7: This histograms shows the distribution of persistence times which occurred for a population ratio of 0.5 in the simulation runs. It corresponds to the middle boxplot of Figure 5.6.

5.2.4 Intransitivity

In this experiment we aim to validate our implementation of the intransitivity coexistence mechanism. As reference we use a case study with varying levels of intransitivity published by Laird and Schamp [12]. The results comprising this case study were generated by their cellular automaton model, which was published in the same paper, and briefly reviewed by us in Section 2.5.3. Here we reenact the given scenarios with our model.

In this experiment 6 species are present, labeled from A to F, with strict pairwise competitive relations defined. In our model competitive outcomes are deterministic between species pairs if environmental fluctuations are absent and relations are set to any non-neutral value (not 1). This is the case here and any competitive relation greater than 1 from the superior to the inferior works equally well to produce strictly deterministic competitive results. Refer to Section 3.5.2 for more information on this. All species rely on the same resource species, of which there is only 1. Other coexistence mechanisms are turned off, although the time to food was chosen to be a little bit higher for this experiment to cut down on simulation time. A slight amount of spatial effects is always present in our model, but here it is consequently comparatively larger than in our other experiments. Interference with the simulation results

however remains minimal, as the outcomes are evidently strongly driven by the competitive relationships between the species. All of the active species are initialized with equal abundances in our case, while each cell is initially filled with an individual of a random species in the reference model, which leads us to a slightly different community size. Due to the very large populations in this experiment, we also used a higher simulation area than usual. The competition ratio is set rather aggressively in our model to better replicate the reference model, which features purely competitive dynamics, and to really capture the influence of the mechanism. Table 5.6 shows the relevant general model parameters, while Figure 5.8 details the five different scenarios with varying competitive relations which were simulated. In this figure arrows always point from the competitive superior species to the inferior. Correspondingly a 1 in the matrices means that the row species is strong against the column species and vice versa. The given t -values for each of the cases are a measurement for the amount of transitivity present, meaning that case 5 is maximally transitive, while case 1 is maximally intransitive. The indicated reversal of species relations between D and F in cases 4 and 5, serve to illustrate how intransitivity can be increased by swapping relations. In addition to defining the five scenarios for this experiment, the image also defines colors for each respective species, which we adopt for the presentation of our own results. Note that the top right reference to figure 2 is no longer valid in this context and should instead reference Figure 5.9.

Results

Figure 5.9 shows the results of the reference model, while Figure 5.10 shows the results of our model. To remain comparable species colors are equivalent in both cases, such that species A is always plotted in the same saturated blue, while B is consistently plotted in light blue and so on. Although the respective results are given in a different format they are directly comparable. We plotted species abundances for every generation and omitted spatial distributions, while they showed the distribution and abundance of species for a selected few generations. The main point of interest here, namely relative species abundances over time, are present in both. Going through the results from left to right and top to bottom respectively, the first two cases are trivial. Due to a strong systematic advantage of species A all others are completely displaced within less than 10 generations, which is confirmed by our model. Beginning with case 3 the

Table 5.6: Parameters used for intransitivity validation.

Parameter	Reference	Our Model
<i>Final model parameters for the base scenario.</i>		
Active Species	6	6
Active Community Size	10000	9996
Resource Community Size	-	2507
Simulation Area	-	5000
Base Speed	-	11
Interaction Radius	-	19.94
<i>Meta parameters that were used.</i>		
Competition Ratio	-	0.7
Time to Food	-	15
<i>Experiment parameters.</i>		
Simulated Scenarios	5	5
Iterations per Scenario	1	1
Generations per Iteration	500	500

effects of intransitivity begin to show clearly. As shown by Gallien et al. [22] intransitive loops with an even number of species invariably lead to the extinction of one species, only uneven intransitive competition networks are stable. This is confirmed by all of these scenarios, which feature six species each, but no amount of intransitivity in any of these cases manages to keep more than five of them around for any extended period of time, in the reference model, as well as ours. Also clearly visible is, that the models are in agreement concerning which species are involved in the stable loops and which are not. It is more difficult to determine whether the models also agree on the surviving species abundances, due to a lack of numerical data in the reference, but the results seem to be largely equivalent. Another interesting aspect which can be observed well through our plots is the influence of various different competitive networks. The case 1 scenario for example seems to be very stable, case 3 seems rather stable as well, but the case 4 scenario oscillates strongly and premature extinction of one species doesn't seem unlikely in this scenario. While in cases 1 and 2 all species have similar average abundance, there is a very distinctive gap in species abundances visible in case 3, in which A and F dominate. Taking the competitive relations into account this makes sense: A is strong against all other species, with exception of F, which is strong against A. So while A

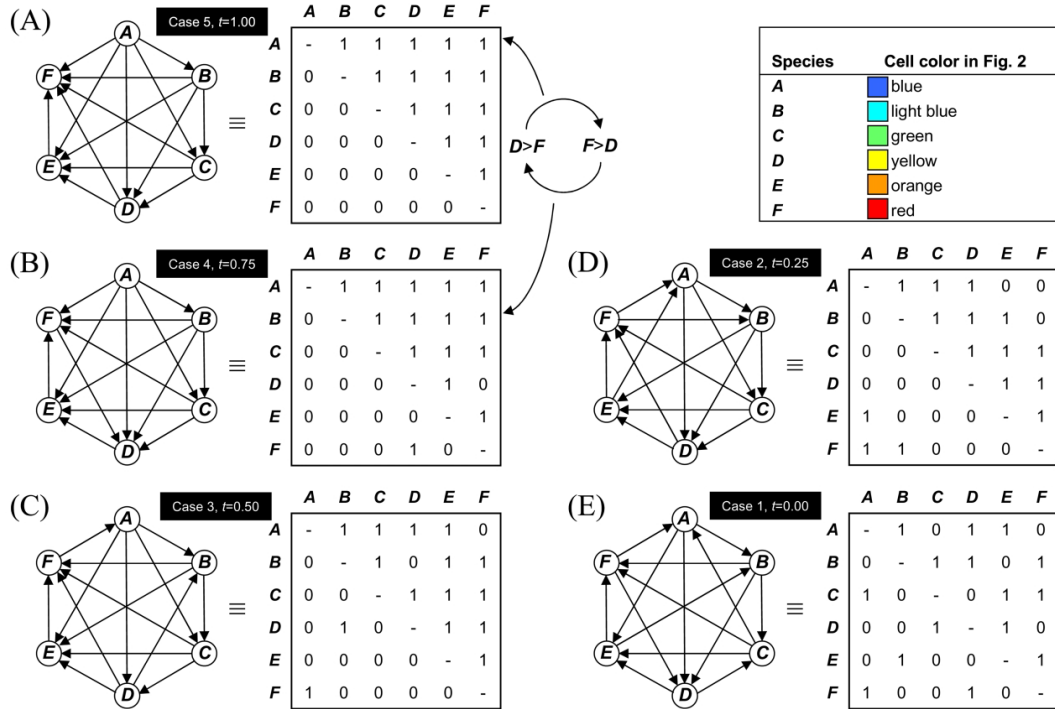


Figure 5.8: This image is taken from the paper of Laird and Schamp [12]. It details the five different competition matrices which were used to specify species relations in the reference scenarios and our own.

keeps the other species down, gaining abundance, it is kept in check by F which benefits from this. Even though the remaining species are strong against F, they cannot bring it down, because doing so automatically empowers A, which regulates them back down and simultaneously helps F again. The results of this experiment clearly show that our implementation of intransitivity is valid and produces expected results, which are in agreement with contemporary research.




Figure 5.9: This image represents the case study of Laird and Schamp, which we use as reference in our comparison. It was originally published in their paper from 2006 [12]. Shown is the state of their cellular automaton model for various generations and for each of the scenarios, or cases, defined by Figure 5.8. The important parts for comparison are the population sizes for each of the species and how they evolve over time.

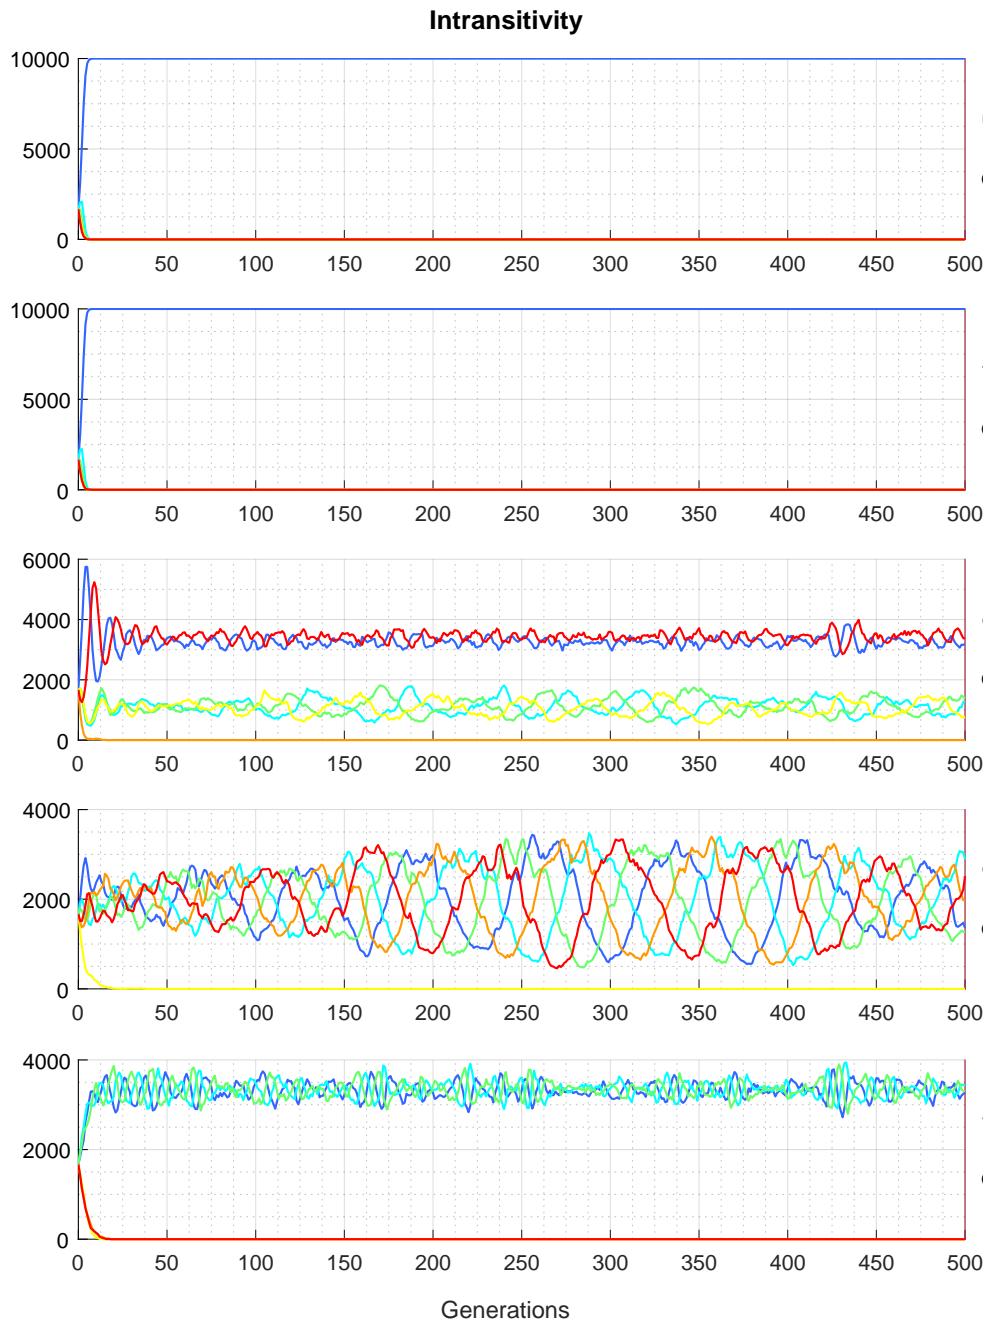


Figure 5.10: These diagrams show the results of our model for the scenarios outlined in Figure 5.8. They are displayed differently, but can be compared well to the results of the reference model, which are shown in Figure 5.9. The upper most diagram here corresponds to the most transitive test scenario (case 5), with following scenarios being increasingly intransitive. The x-axis represents the generations, while the y-axis indicates population sizes. Consequently for every generation the population sizes for all the species are plotted.

5.2.5 Resource Partitioning

The final coexistence mechanism to validate in our model is resource partitioning. In this case we could not find a well suited reference model to compare against. Therefore we compare solely against the results of the other coexistence mechanisms and our own expectations. In this experiment we use two active and two resource species. Then we vary the food preferences of the active individuals such that they increasingly focus on just one of the resources, ignoring the other. In particular we used food preferences of 1 as the maximum value, indicating the species focal resource. Then the preference for the other resource, which is focused by the other active species, was gradually decreased until it reached 0. This leads to reduced interspecific competition, while intraspecific competition increases, even though the resource abundances stay constant throughout the simulation. Table 5.7 lists the relevant parameters for this experiment.

Table 5.7: Parameters used for resource partitioning validation.

Parameter	Value
<i>Final model parameters for the base scenario.</i>	
Active Species	2
Active Community Size	100
Resource Species	2
Resource Community Size	182
Base Speed	0.85
Interaction Radius	2.75
<i>Meta parameters that were used.</i>	
Competition Ratio	0.4
Time to Food	10
<i>Experiment parameters.</i>	
Simulated Scenarios	11
Iterations per Scenario	100

Results

One thing immediately becomes obvious when looking at the results for this experiment, as depicted by the figures 5.11 and 5.12, namely that the achieved

coexistence times are huge. In fact even intermediate levels of resource partitioning are an order of magnitude more effective in regard to species coexistence than the storage effect in the scenario discussed before. And the base scenarios between these two experiments are actually very similar, so that a direct comparison is indeed possible. Not to mention that higher levels of resource partitioning are yet another order of magnitude higher and more. Unfortunately we very much ran into our computation limit of 10000 generations per simulation for the more effective values, so that all cases with a food preference overlap of 0.5 and smaller are artificially capped. Considering our expectations and the clearly visible trend, which seems to be exponential in nature, it can however safely be assumed that the real coexistence times for these cases reach well above 100000 generations. Not even the highest possible resource partitioning will ensure an indefinite coexistence in our model though. The zero-sum assumption, randomness of food distribution and randomness in mortality would eventually lead to extinction in every feasible scenario, if given enough time. This can also be seen in the diagrams, where even the cases with the highest amount of partitioning have produced outliers well below the expected results, but the median is unquestionably fixed at 10000 generations. The results obtained here conform well to our expectation, considering that very high amounts of resource partitioning are essentially a way of modeling perfect niches, which are generally accepted to be one of the most potent drivers for the maintenance of biodiversity. In further experiments we refrain from utilizing the full range of this effect, as it undermines the purpose of our research and models scenarios of trivial coexistence in which we are not interested here. This experiment demonstrates that the resource partitioning effect in our model produces expected results and is in line with the general understanding of resource partitioning in contemporary research.




Figure 5.11: This figure shows the results of the resource partitioning validation experiment with a linear y-axis.

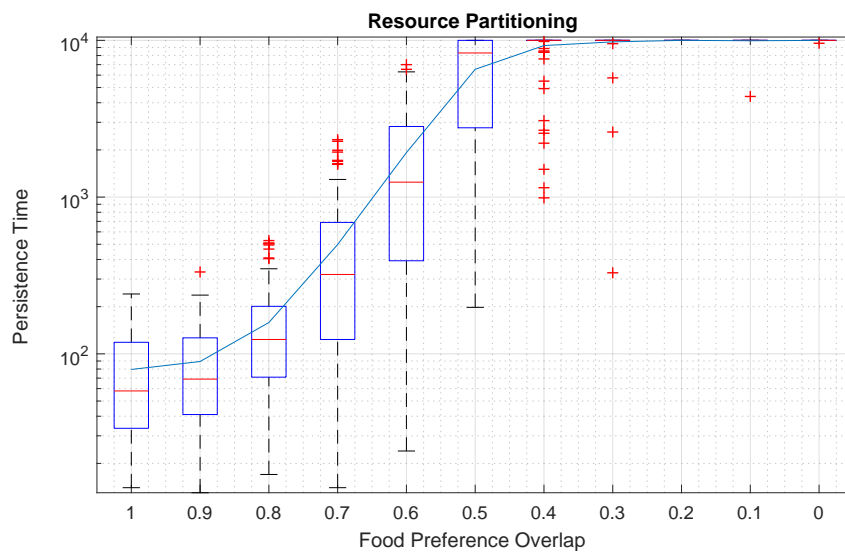


Figure 5.12: This figure shows the results of the resource partitioning validation experiment with a logarithmic y-axis, but otherwise represents the same data as Figure 5.11.

5.3 Parameter Range Experiments

Next up we conduct some parameter range experiments in preparation of the experiments with combined coexistence mechanisms. These serve to explore

the effects of the various coexistence mechanisms a little further, at least for one particular base scenario. An important aspect of these and the following experiments is, that all of them use the exact same base scenario definition on which various starting conditions are then applied. Specifically this holds true for the experiments in this section, as well as Section 5.4. Therefore the results of all following experiments can be compared directly to each other. Due to time constraints these experiments are conducted with fairly low iterations and few starting conditions, so that the results are rather coarse, however this does generally not pose a problem in identifying the trends we are looking for.

5.3.1 Base Scenario

This section defines the base scenario upon which all the following experiments build up. Consequently it was chosen in such a way, that everything for the three coexistence mechanisms we study is already in place and they just need to be turned on. In absence of the coexistence mechanisms the scenario is a simple neutral one, characterized by three active and three resource species, as well as the presence of environmental fluctuations. Although multiple resource species are present, every active species starts out to be indifferent towards them. To establish a baseline for species coexistence in the absence of any coexistence mechanisms we also conduct a brief experiment on the base scenario alone here. Once without the environmental fluctuations, once with them. In the later experiments they were always used.

Results

Although this scenario is a very simple one and the results shown by Figure 5.13 very straightforward, there are a few important points to address here. Firstly it should be noted, that these scenarios here are analogous to the validation experiments with demographic and environmental stochasticity (Sections 5.2.1 and 5.2.2) from before, only with another active species, higher community size and slightly different parameters. The increased population sizes lead to higher expected coexistence times [9], while the increased amount of active species decreases the time until the first extinction is expected to happen. Overall it can be observed that doubling the population sizes had a greater effect on prolonging coexistence times, than the addition of another species had on decreasing them, when environmental fluctuations are turned off. A persistence time of roughly

Table 5.8: Parameters used for the base scenario parameter range experiment.

Parameter	Value
<i>Final model parameters for the base scenario.</i>	
Active Species	3
Active Community Size	300
Resource Species	3
Resource Community Size	282
Simulation Area	200
Stochasticity Type	I. Multilevel
E. Fluctuation Amplitude	0.6
E. Persistence Time	20
Base Speed	1.33
Interaction Radius	5.3
<i>Meta parameters that were used.</i>	
Competition Ratio	0.6
Time to Food	12
<i>Experiment parameters.</i>	
Simulated Scenarios	1
Iterations per Scenario	100

95 generations is the result, compared to about 72 generations we observed in the validation experiment for demographic stochasticity. The next observation is that adding environmental fluctuations to this scenario decreased coexistence times much more strongly than was the case in the validation experiment for environmental stochasticity. We suspect that this is an interaction between having one more species which is suspect to randomness and said randomness having a stronger influence by virtue of a higher fluctuation amplitude and competition ratio. More experiments are needed to verify or disprove these suspicions. Other than that there is not a lot to discuss here.

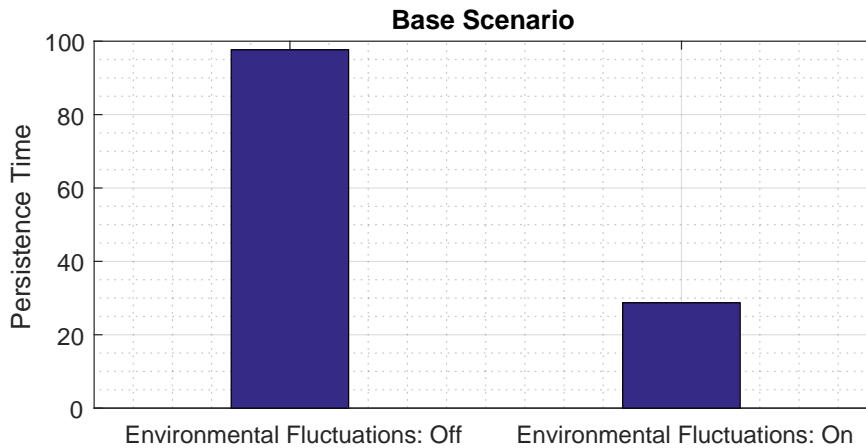


Figure 5.13: This figure shows the achieved persistence times of the base scenario with and without environmental fluctuations. The y-axis corresponds to the generations of coexistence, while the x-axis indicates the starting conditions.

5.3.2 Storage Effect

The first coexistence mechanism we took a closer look at was the storage effect. Model parameters directly linked to this mechanism are the environmental fluctuations and most significantly the max stored food parameter. In this parameter range experiment we varied the amount of food individuals can store, the environmental fluctuation amplitude and the environmental persistence time, one after another. Table 5.9 lists which values were used, when another one was varied. Aside from these parameters the base scenario definition from Section 5.3.1 applies.

Table 5.9: Parameters used for the storage effect parameter range experiment.

Parameter	Value
<i>Model parameters used when not varied.</i>	
E. Fluctuation Amplitude	0.6
E. Persistence Time	20
Max. Stored Food	4
<i>Experiment parameters.</i>	
Simulated Scenarios	13
Iterations per Scenario	100

Results

For this experiment we varied three parameters individually, keeping the other two constant. Therefore there are three main results for this experiment, all shown in Figure 5.14, and we will go over them one after another. First and most importantly for our model we tried different max stored food parameters. The value was increased in exponential steps, starting at 1, up to 16. This was done, because diminishing returns with respect to the achievable persistence times were expected. Evidently we were right with this assumption, as the increases in coexistence time per possible stored food are roughly linear over the exponential sequence. A closer inspection reveals that the slope of persistence times is positive and increasing from 1 to 4, but then starts to become smaller again. The resulting curve strongly resembles a sigmoid curve with a characteristic s-shape. This is in line with our expectations and seems reasonable. Surprising is however, that the storage effect seems to be comparatively weak in this scenario, when setting it into relation with the experiment of Section 5.2.3. Indeed a max stored food value of 4 is needed to even achieve coexistence times similar to the base scenario without any environmental fluctuations, while previously a value of 1 was enough to already surpass the variant with demographic stochasticity (Section 5.2.1). Even when setting the storage effect only into relation to the coexistence times which were achieved previously with environmental fluctuations, it performs only about half as well here for the case of one max stored food, as it did in the validation experiment. Most likely this is a result of this being a three species simulation, meaning that it is unrealistic to expect the storage effect to extend persistence times by the same relative amount regardless of how many species are simulated. Surely if we were to simulate a million active species, each starting with a population size of 100, no amount of storage would yield a persistence time greater than a single generation. Consequently this difference is significant and cannot be left out of the equation, but further research is needed nonetheless. The overall characteristic of the curve is however very consistent with our expectations.

Next we varied the environmental fluctuation amplitude. There is only a single clearly visible trend in this dataset: very high fluctuation amplitudes seem to support species coexistence, but this is not a gradual effect which holds true across the whole range of amplitudes, as no statistical significant differences

between the lower four values can be made out. The visible differences there might very well be caused by noise alone. Overall these findings are rather unspectacular. A much more visible and strict trend can be observed when varying the environmental persistence times: a faster changing environment greatly benefits species coexistence across the whole tested range. That large environmental persistence times have a destabilizing effect was very much expected, and the reasoning for that was outlined in the results of Section 5.2.2. But that the effect holds up strongly even among overall low environmental persistence times was somewhat unexpected. Our auxiliary experiments have consistently shown that these trends not only apply to our model, but also to the reference model we compared against in the first few experiments, which was proposed by Danino et al. [9]. Slower environmental changes also have a destabilizing effect there, while even among already fairly fast changing environments the faster ones remain consistently better in terms of achieved coexistence times. Albeit this effect is overall much weaker in their model than it is in ours, which is however unsurprising as we explained in aforementioned results section, because our model is more sensitive in regards to the specified environments. So on the one hand this result should have been expected by us and is in line with previous observations, on the other hand it seems to be an unrealistic characteristic, at least when it comes to very quickly changing environments. We suspect that the increase in coexistence times observed by the fast changing environments happens because it practically ensures that every individual gets its turn to shine, eliminating most bad luck from the simulation. It is highly debatable which values for the environmental fluctuations are most realistic and consequently it might be argued that the more extreme cases we looked at here are simply not, or rarely, occurring in nature. Either way further research into this might be beneficial.

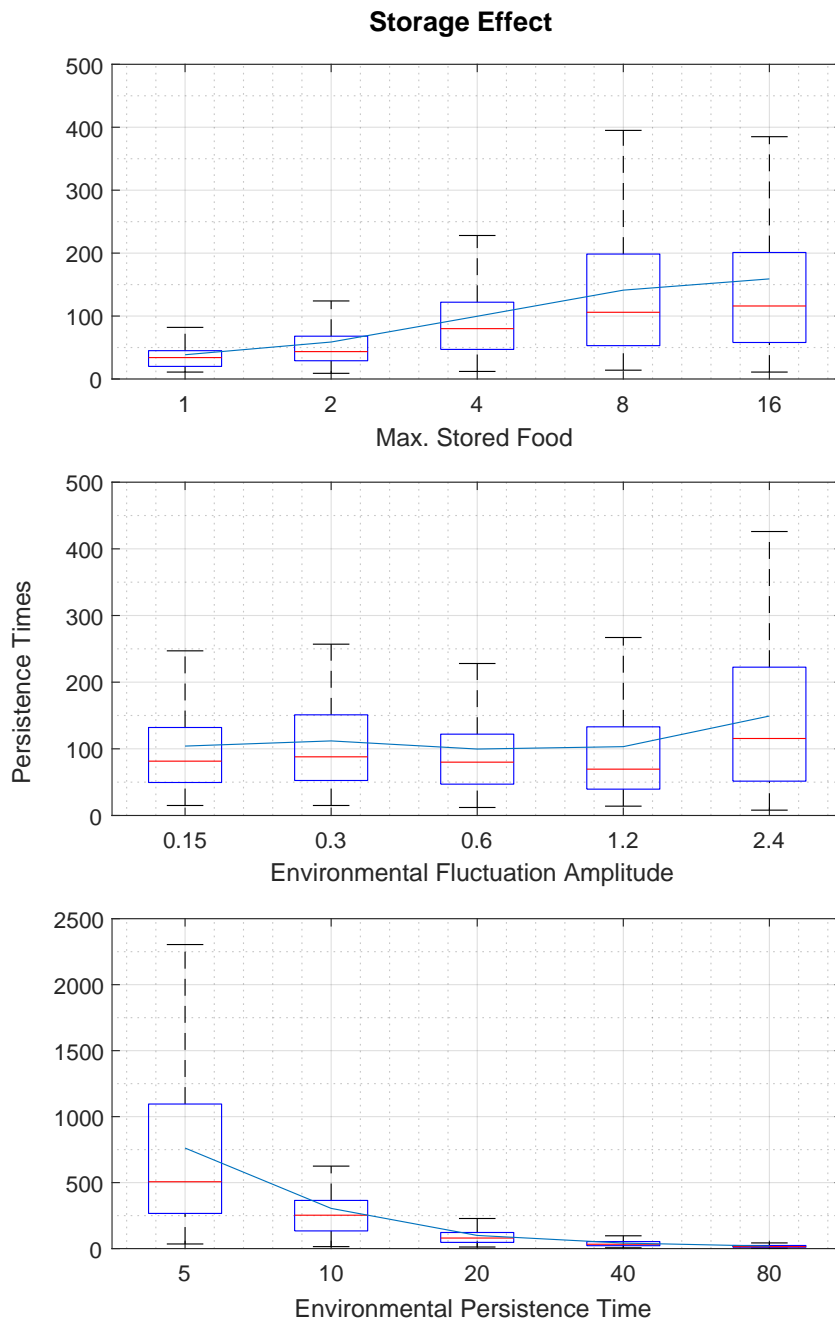


Figure 5.14: Results of the storage effect parameter range experiment. Only one parameter was varied at a time, and Table 5.9 lists the values used for the fixed parameters. Outliers are omitted for clarity.

5.3.3 Intransitivity

The next coexistence mechanism we explored more was intransitivity. Due to the environmental fluctuations present in the base scenario, the mechanism can be varied in its strength, which is otherwise not possible in our implementation as explained in Section 3.5.2. This functionality was also not tested by the validation experiment (Section 5.2.4), which did not include any environmental fluctuations and therefore only made use of this mechanism at its highest potency. We varied the strength of intransitive competitive relations, beginning with a value of 1.1 up to a value of 2. Since the environmental fluctuation amplitude of the scenario is set to 0.6, the highest possible fitness difference between two individuals in this scenario is 60%. This is important, because once the intransitive fitness multiplier becomes large enough to bridge the maximum difference, further increases of the competitive advantage become functionally ineffective. For this scenario the highest meaningful competitive relation is therefore 1.6, and values beyond that should produce equivalent results. Table 5.10 shows the parameters for this experiment. Aside from the varied intransitive strength, all parameters from the base scenario 5.3.1 apply.

Table 5.10: Parameters used for the intransitivity parameter range experiment.

Parameter	Value
<i>Experiment parameters.</i>	
Simulated Scenarios	10
Iterations per Scenario	100

Results

Similar to the storage effect parameter range experiment from before, the overall shape of the curve we can observe in Figure 5.15 matches very well with our expectations, but the scale of achieved persistence times with this coexistence mechanism does not. Even with the strongest possible intransitive scenarios, those with a competitive relation strength of 1.6 and above, the measured coexistence times are smaller than those in the simple base scenario without environmental fluctuations (Section 5.3.1), with just about 83 generations average, compared to about 97. It follows that the disturbance caused by the environmental fluctuations are greater than what the inclusion of intransitivity can do to stabilize the simulations for this particular scenario. In order to

gain a better understanding of what is happening here, we plotted a few randomly selected simulations out of those with a competitive relation of 1.6 in more detail, showing how the species abundances fluctuate over time. The corresponding Figure 5.16 confirms our intuition: the comparatively small community size we are simulating here doesn't seem sufficient to reliably buffer the fluctuations caused through intransitivity. The effect itself is very clearly visible and does behave as expected, however in relation to the small community size the fluctuations are so strong, that species are constantly on the brink of extinction and randomness does the rest. In support of this assumption we can look at the results of the intransitivity validation experiment, case 1, where the simulation quickly degraded into a three species community, but was then highly stable. The main difference between these simulations are the much higher species abundances with which the validation experiment could work. Another major difference of course is the presence of environmental fluctuations here, but it seems doubtful that they are the ultimate cause of the relatively larger population fluctuations, as the intransitive effect appears to dominate the population sizes strongly. Once again more dedicated research into this is required.

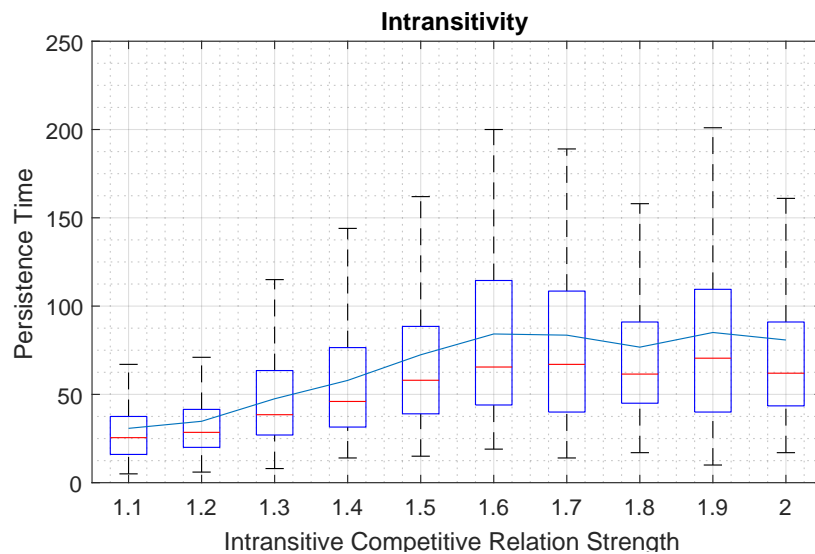


Figure 5.15: Results of the intransitivity parameter range experiment. The x-axis indicates varying levels of intransitive strength. Outliers are omitted for clarity.

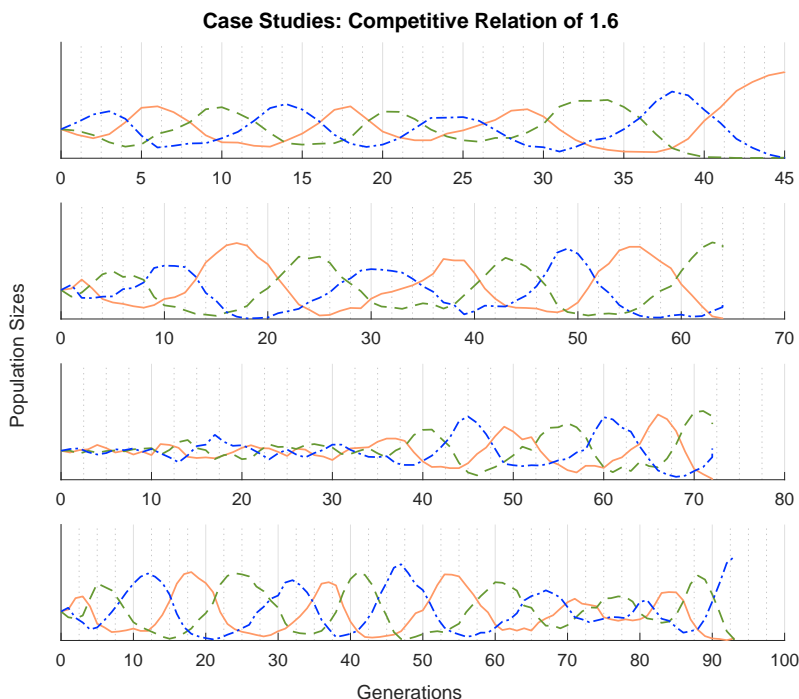


Figure 5.16: These graphs show a few individual simulations out of the performed ones with an intransitive competitive relation of 1.6 in detail. Simulations were chosen randomly.

5.3.4 Resource Partitioning

The last coexistence mechanism left to detail is resource partitioning. We concentrate on the less pronounced values of this mechanism here, as complete partitioning defeats the purpose of a competition oriented model and creates a scenario in which coexistence is trivially explained. With three resource species and three active species, there are two main kinds of resource partitioning possible, in the sense of a coexistence mechanism. In the first case every species has a single focal resource, which is not shared with any other species. The focal resource is the main source of food for the species, defined by a food preference value of 1, while the other two resources are equally less preferred with a value smaller than 1. In the second case every active species has two focal resources, and every resource is focused by two active species. Again, each focal resource is preferred equally with a value of 1 and the remaining resource with less than 1. In this case overlap is significantly greater than in the first, but even so no single abundant species is likely to deprive the others

of all their food, because one resource always remains comparatively untouched. However, the less abundant species both then have to share a single resource to evade competition with the abundant one, which remains less preferable than the first case. The parameters for this experiment are shown in Table 5.11, but aside from the varied preference overlaps, there are no changes to the base scenario defined in Section 5.3.1.

Table 5.11: Parameters used for the resource partitioning parameter range experiment.

Parameter	Value
<i>Experiment parameters.</i>	
Simulated Scenarios	20
Iterations per Scenario	100

Results

In this experiment we make observations very similar to those we made in the other parameter range experiments. The Figure 5.17 shows the results for this experiment and the obtained curves are shaped as expected. Resource Partitioning, even without utilizing food overlaps of less than 0.5, remains a very potent coexistence mechanism even for this scenario, but compared to the initial validation experiment from Section 5.2.5 it lags behind. When set in relation to the worse starting conditions of this scenario, the effect seems to be least impacted out of the three however. As expected the case with a single focal resource yielded coexistence times higher than that with two focal resources, although the differences only really begin to show substantially for higher values of partitioning. Interestingly the dataset we obtained from the simulations with two focal resources features characteristic bumps between certain value pairs. The increase in coexistence time from a food preference overlap value of 0.75 to 0.7 for example is much stronger than the increase from 0.8 to 0.75, which is almost nonexistent. Most likely this can be attributed to the few iterations per run we used, causing the data to be rather noisy, but curiously the effect seems rather strong here.

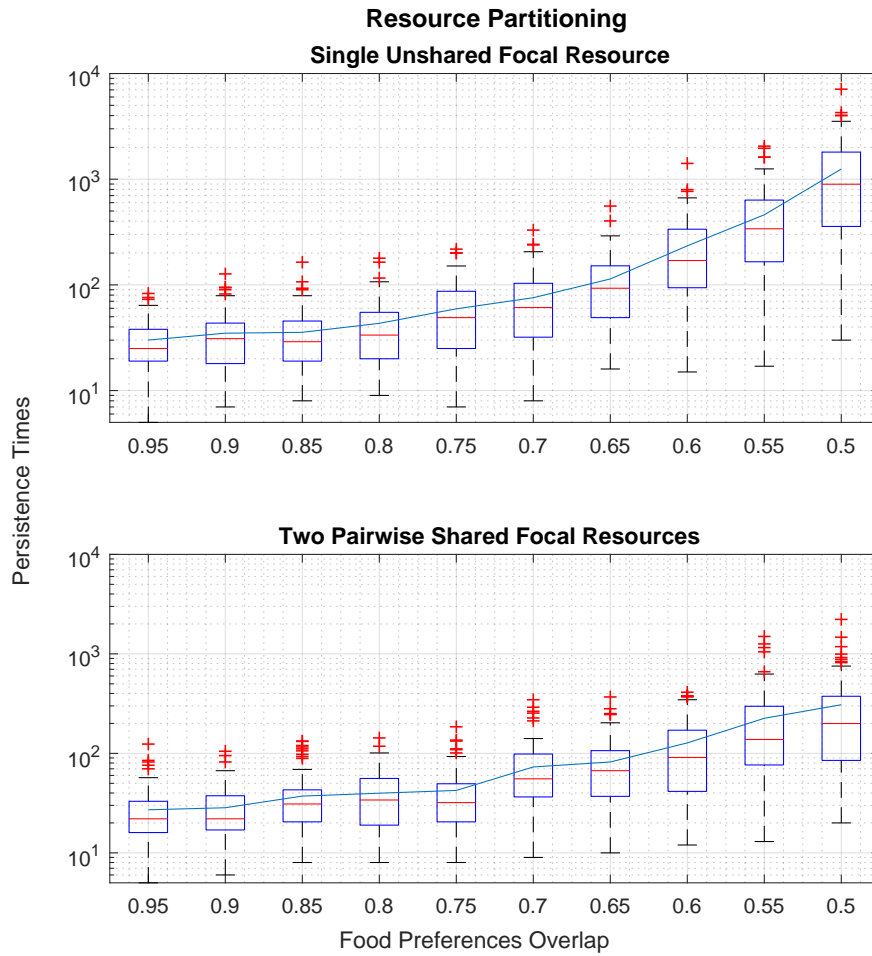


Figure 5.17: Results of the resource partitioning parameter range experiment. The x-axes here represent the varying amounts of preference overlaps for the non-focal resources of the active species. The top diagram represents the first, the bottom diagram the second case as outlined in Section 5.3.4.

5.4 Combined Coexistence Mechanisms Experiments

Having looked at each coexistence mechanism individually for a common base scenario, we now combine them to study their interactions. The base scenario used in the following experiments is identical to that of the parameter range experiments, as defined in Section 5.3.1. Additionally all of the parameter values used in the following combined experiments have known consequences

when isolated per coexistence mechanism due to the previous parameter range experiments. Therefore any deviations can be fully attributed to the interaction of the mechanisms alone. We conduct pairwise experiments between all three possible coexistence mechanisms to observe the influence they have on each other and one experiment with all three of them at once.

5.4.1 Storage Effect & Intransitivity

The first combination of mechanisms we consider is storage effect with intransitivity. For explanation of the relevant parameters and starting conditions used, refer to the previous parameter range experiments, sections 5.3.2 and 5.3.3, as this experiment makes use of the same parameters, only now in combination. The experiment related parameters are shown in Table 5.12. As stated the used base scenario is identical to that of Section 5.3.1 and the only differences lie in the respective parameter variations used.

Table 5.12: Parameters used for the experiment with storage effect and intransitivity.

Parameter	Value
<i>Experiment parameters.</i>	
Simulated Scenarios	25
Iterations per Scenario	100

Results

The combination of storage effect and intransitivity had surprising results because as shown by Figure 5.18, the addition of intransitive competitive relations actually had a negative impact on species coexistence in the presence of intermediate to high levels of storage. It was more advantageous in terms of coexistence to solely use the storage effect at higher levels, instead of combining it with intransitivity. For low storage values, this relation does not hold, however. We suspected that the increase in intransitivity might compromise the benefits of the storage effect by overshadowing its stabilizing effects with more pronounced abundance fluctuations. A similar effect was observed in the experiment of Section 5.3.3, which might promote premature species extinction. To investigate further we plotted species abundances over time for individual simulations, each featuring a max stored food parameter of 16 and a different

amount of intransitivity. We selected random simulations which ran for similar amounts of time for comparison. The plots are shown in Figure 5.19, but there is no obvious pattern visible that would explain why the increase in intransitivity hinders coexistence in these particular scenarios. It does seem like the abundance fluctuations become slightly more regular and distinctive with increased intransitivity, which would make sense. But the differences are small and overall these plots do not permit drawing conclusions about the ultimate cause of the observed relationship between storage and intransitivity. Further experiments are necessary. For simulations featuring larger community sizes, the relationship between storage and intransitivity we discovered may very well not apply, but at least in this scenario a combination of two mechanisms was not helpful in terms of coexistence.

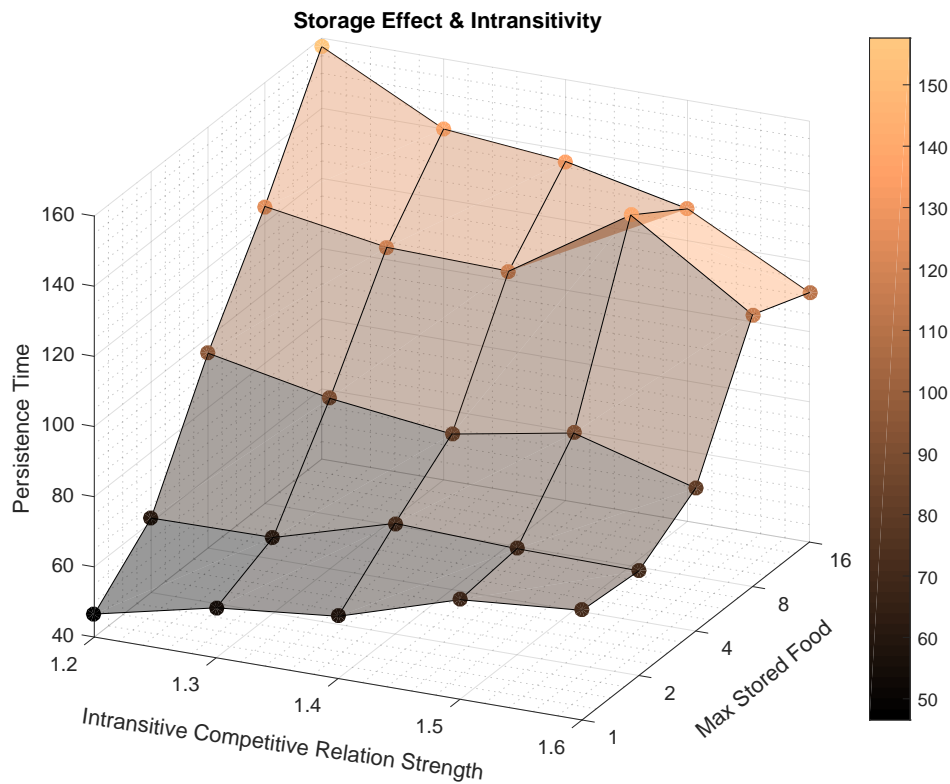


Figure 5.18: Results of the combined experiment with storage effect and intransitivity. The different levels of intransitive strength are indicated by the x-axis, while the maximum amount of stored food and therefore the strength of the storage effect is plotted on the y-axis.

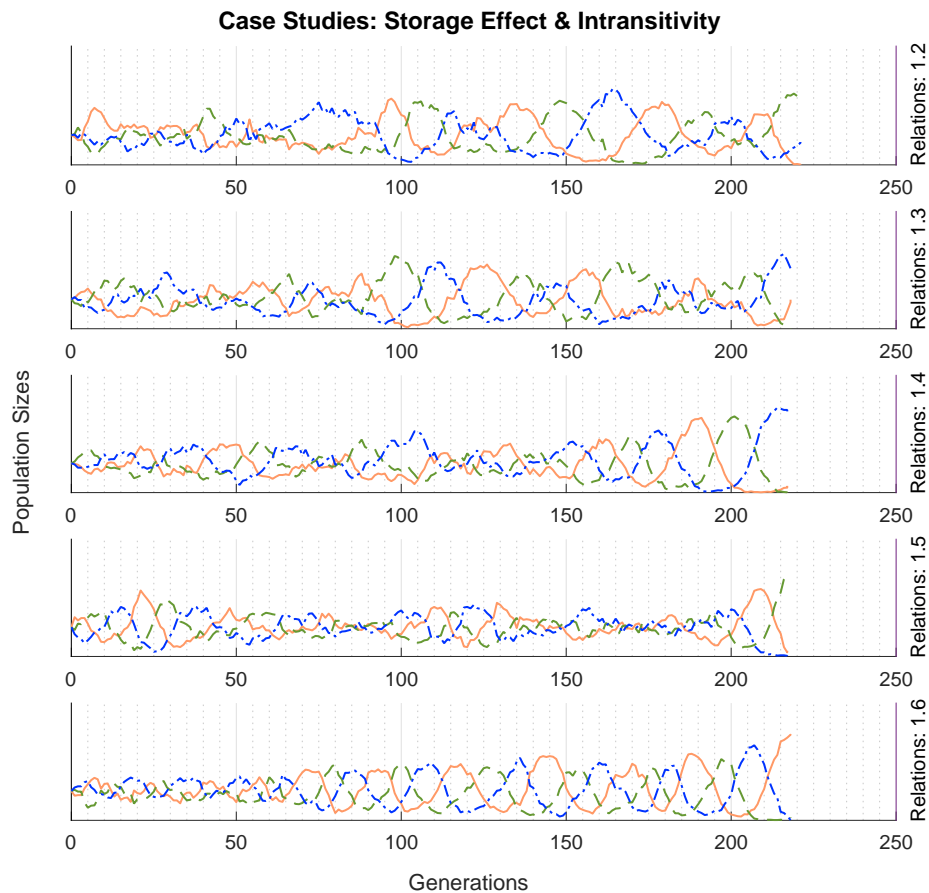


Figure 5.19: Exemplary comparison between individual simulations of similar length with a max stored food parameter of 16 and varying levels of intransitivity, starting from competitive relations of 1.2 at the top, up to competitive relations of 1.6 at the bottom.

5.4.2 Storage Effect & Resource Partitioning

The second combination is that of storage effect and resource partitioning. For explanation of the relevant parameters and starting conditions used, refer to the previous parameter range experiments, sections 5.3.2 and 5.3.4, as this experiment makes use of the same parameters, only now in combination. For resource partitioning we use the case with a single, unique, focal resource for each active species here. Otherwise the parameters for this experiment are shown in Table 5.13. As before, the used base scenario is identical to that of

Section 5.3.1 and the only differences lie in the respective parameter variations used.

Table 5.13: Parameters used for the experiment with storage effect and resource partitioning.

Parameter	Value
<i>Experiment parameters.</i>	
Simulated Scenarios	25
Iterations per Scenario	100

Results

The results of this experiment, depicted by Figure 5.20, are generally fairly close to what we expected. It can be observed that the combination of storage effect and resource partitioning leads to significantly greater coexistence times in all tested combinations than either coexistence mechanism could achieve on its own. The biggest benefit remains with resource partitioning, which is unsurprising considering the previous experiments. Still, the addition of even a minimal amount of storage almost tripled the coexistence time that was achieved for a food preference overlap of 0.6: Previously it was just about 234 generations without storage (Section 5.3.4), but is now about 586 with a max stored food of 1. This is an increase of over 350 generations, even though a max stored food of 1, when isolated, was only able to extend persistence for a mere 20 generations before (Section 5.3.2). Interestingly max stored food values greater than 1 generally do not lead to any further jumps in coexistence of this magnitude. Overall it can be observed that storage becomes more efficient in the presence of resource partitioning and vice versa, indicating that this combination scales stronger than the sum of its parts for most, if not all, parameter pairs. Very obviously there is a significant outlier among these results for a max stored food of 4 and minimal food preference overlap. The plot strongly suggests that this is an anomaly, caused by randomness and small sample sizes. But simultaneously the difference is big enough to be statistically significant, implying a true peak around there. Further experiments are required to identify the nature of this observation.

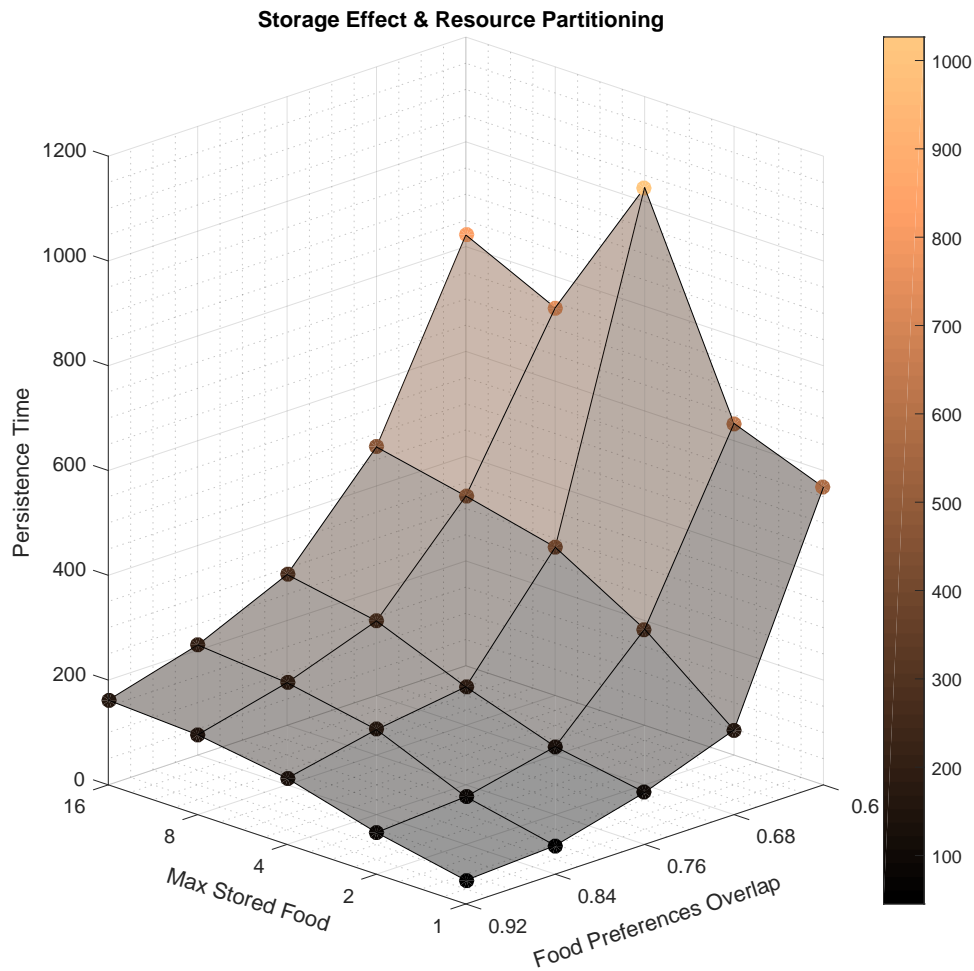


Figure 5.20: Results of the combined experiment with storage effect and resource partitioning. The amount of stored food, and therefore the strength of the storage effect, is plotted on the y-axis, while food preference overlap represents resource partitioning with a single, unique, focal resource with the values shown on the y-axis.

5.4.3 Intransitivity & Resource Partitioning

The third and last pairwise combination experiment is that of intransitivity and resource partitioning. For explanation of the relevant parameters and starting conditions used, refer to the previous parameter range experiments, sections 5.3.3 and 5.3.4, as this experiment makes use of the same parameters, only now in combination. For resource partitioning we use the case with a single, unique, focal resource for each active species here. Otherwise the parameters

for this experiment are shown in Table 5.14. As before the used base scenario is identical with that of Section 5.3.1 and the only differences lie in the respective parameter variations used.

Table 5.14: Parameters used for the experiment with intransitivity and resource partitioning.

Parameter	Value
<i>Experiment parameters.</i>	
Simulated Scenarios	25
Iterations per Scenario	100

Results

In this experiment we observe a very strong interaction between intransitivity and resource partitioning. As seen in Figure 5.21 this interaction demands higher levels of resource partitioning to work. The combined effect only really becomes apparent in the presence of food preference overlaps below 0.7, without which the added intransitivity remains comparatively ineffective. If the right conditions are met, the addition of intransitive competitive relations benefits coexistence times greatly in this combination: They were increased significantly beyond an entire order of magnitude. The full range of achieved coexistence is unfortunately not adequately represented by our data, because the imposed limit of 10000 generations per simulation was reached regularly for the three scenarios with the highest coexistence times. In fact the median of persistence times for a food preference overlap of 0.6 and intransitive competitive relations of 1.6 was 10000 generations, underlining the impact of this combination. Figure 5.22 allows a more detailed look at the less prominent parameter combinations, revealing that the combination of intransitivity and resource partitioning yields consistently higher coexistence times than either mechanism does in isolation. It can further be observed that the combined benefit of these mechanisms is greater than the individual benefits that they have in isolation. This holds true for each respective parameter combination, even for the lower ones, although this effect is not particularly strong there. Our intuition is, that the addition of resource partitioning to intransitivity is just strong enough, when using higher values, to ensure with some reliability that the species abundance fluctuations caused by intransitivity do not hit 0, thus achieving very high

coexistence times. We plotted species abundances over time for scenarios with the highest tested intransitivity and varying food preference overlaps. The simulations were randomly chosen out of those with at least 300 generations of coexistence, only plotting those for equally scaled diagrams. The plots are shown by Figure 5.23, but similar to the storage and intransitivity experiment from before (Section 5.4.1), there is no obvious difference in the behavior of the graphs from which we could derive the reasoning for the observed effects. It does seem like the cases involving more resource partitioning have overall tighter curves, but pronounced fluctuations remain and species still appear to be on the verge of extinction fairly regularly. This observation does not necessarily conflict with our intuition outlined above, but it underlines how subtle changes can make all the difference.

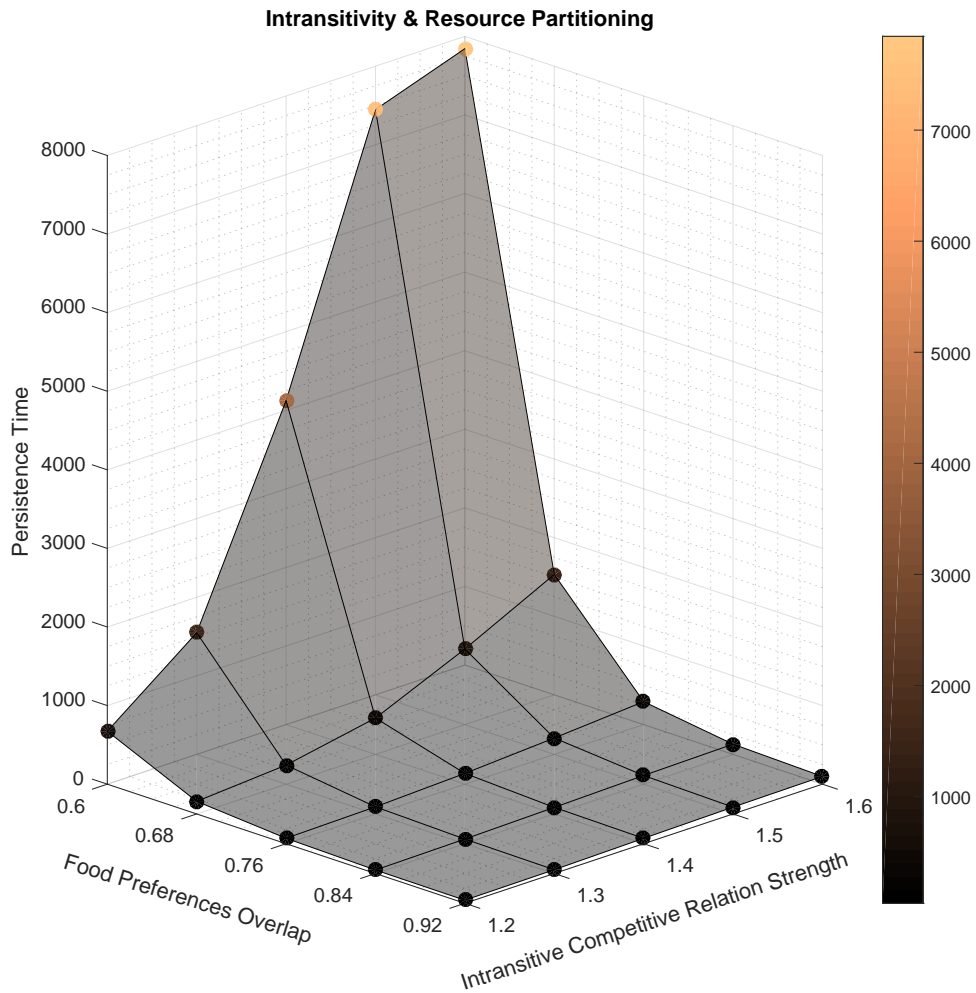


Figure 5.21: Results of the combined experiment with intransitivity and resource partitioning. Food preference overlap represents resource partitioning with a single, unique, focal resource with the values shown on the x-axis, while different levels of intransitive strength are indicated by the y-axis.




Figure 5.22: This diagram shows the same data as 5.22 but with a logarithmic z-axis for more detail.




Figure 5.23: Exemplary comparison between segments of individual simulations which were cut to the same length with intransitive competitive relations of 1.6 and varying levels of resource partitioning, starting from a food preference overlap of 0.92 at the top, up to an overlap of 0.6 at the bottom. Shown are the species abundances for each generation.

5.4.4 Storage Effect & Intransitivity & Resource Partitioning

After combining all of the mechanisms in pairwise fashion, we now put all three of them together in a single experiment. Unlike before we do not simulate every possible combination of parameters, which simplifies visualization and reduces the computational overhead. Instead we simulate five scenarios, steadily increasing the strength of all three coexistence mechanisms from one to the next. The exact same parameters used here were also used before in the three pairwise experiments (Sections 5.4.1, 5.4.2 and 5.4.3): The storage effect, modeled through the max stored food parameter, is increased exponentially from 1 to 16. Intransitivity, defined through the competitive relations between species, ranges linearly from 1.2 to 1.6. Resource Partitioning, controlled through the food preference overlap between species, is used with a single, unique, focal resource and varies from 0.92 to 0.6 in steps of 0.08. In the results these combinations are referred to as cases 1 to 5, in ascending order with increasing strength of the coexistence mechanisms. The parameters for this experiment are shown in Table 5.15. As before the used base scenario is identical with that of Section 5.3.1 and the only differences lie in the respective parameter variations used.

Table 5.15: Parameters used for the experiment with all three coexistence mechanisms combined.

Parameter	Value
<i>Experiment parameters.</i>	
Simulated Scenarios	5
Iterations per Scenario	100

Results

We have now combined mechanisms, which when considered pairwise either hindered themselves, stacked well, or had very strong positive interactions. The results of this experiment are shown by the figures 5.24 and 5.4.4. Based on the data we obtained, two main observations can be made by setting this experiment in relation to those before. The first observation is that having intermediate levels of all coexistence mechanisms active (cases 2, 3 and 4) yields significantly

higher persistence times than any pairwise combination of mechanisms could produce for these parameter ranges. The values we obtained are about 104, 377 and 2542 generations, respectively. Correspondingly the highest values we obtained in the equivalent pairwise simulations are 91, 187 and 962 generations. This difference is significant and the values obtained far outperform any single coexistence mechanism in isolation. It should be noted, however, that the very first and last cases did not follow this trend. Instead they performed worse than in some previously measured pairwise mechanism experiments, even though the achieved coexistence times are still overall large. For case 1 this may be attributed to noise, as the coexistence mechanisms are weak and most obtained values in these cases are close to each other. The results of case 5 on the other hand are statistically significant and show that the addition of the storage effect holds back the very potent combination of resource partitioning with intransitivity once their interaction reaches its highest point. Setting this in relation to our observations from Section 5.4.1, which already showed a negative interaction between storage and intransitivity, there is likely a connection here. But we currently lack enough data to pinpoint what exactly causes this negative interaction and, perhaps even more interestingly, what caused it to show up once again here. Our previous suspicion, that the storage effect is simply not stabilizing enough to counter the abundance fluctuations caused through intransitivity, and is thus undermined by its presence, most likely does not hold for this scenario, considering the overall high persistence times achieved. Especially not when resource partitioning with intransitivity together could already reach even greater persistence times reliably (Section 5.4.3). We would assume that the storage effect acts necessarily stabilizing, but based on our findings this assumption seems to break apart here. Our observation cannot be generalized, though, since the causes are unknown and may be entirely scenario dependent, or even an artifact of the zero-sum assumption or others.




Figure 5.24: Results of the combined experiment with all three coexistence mechanisms. The cases 1 to 5, as indicated by the x-axis, correspond to simultaneously increasing strengths of the used coexistence mechanisms. Storage effect goes from 1 to 16, intransitivity from 1.2 to 1.6 and resource partitioning from 0.92 to 0.6. For more details refer to the explanation in Section 5.4.4. Otherwise this diagram is read identical to that of Figure 5.2 with persistence times for the varying parameters plotted on the y-axis.




Figure 5.25: Results of the combined experiment with all three coexistence mechanisms. This diagram is identical to that of Figure 5.24, except that the y-axis is logarithmic for clarity.

6 Conclusion

In this thesis we proposed a new model for the study of species coexistence, focusing on the aggregate effects of several coexistence promoting mechanisms being active simultaneously. To arrive at this model, we first gave a brief overview of the related literature and scientific foundations for it. The focus was on coexistence mechanisms and models which already integrated these in a neutral context, or close to it. Based on our findings in the contemporary research, a multi-mechanism model was designed from scratch. Key challenges and requirements for such a model were identified and the scope defined accordingly to allow realization within the context of this thesis. In this work we restricted ourselves to focus on just three coexistence mechanisms, keeping as close to the related work as possible. Nonetheless, the multi-mechanism model was intentionally designed to be flexible enough in its assumptions and architecture, so that more mechanisms could be added with relative ease in the future. Due to the requirements and inherent complexity of this approach, we argued that an individual based simulation is best suited for the task. We then outlined the concept for our model by starting from the classic neutral model with just the local community and its dynamics, expanding upon it step by step. With this process a unifying solution for storage effect, intransitivity and resource partitioning on top of a neutral base was found. In the following chapter the previously defined concept was substantiated by presenting our respective implementation of it, going over all the details which were intentionally left open from the conceptual point of view. Furthermore, certain difficulties we faced, and which may be encountered when implementing such a model, were addressed. With a concrete realization of the multi-mechanism model now at hand, several experiments were conducted, which can be divided into three distinct groups: the first array of experiments was designed to test the feasibility and comparability of our approach, successfully validating that the multi-mechanism model is able to reproduce results of more specialized models from contemporary research with good accuracy. Scenarios which were compared against established references included those driven

by demographic stochasticity, environmental stochasticity, storage effect and intransitivity. There was no compelling reference model for resource partitioning available, so the mechanism was compared internally against the others and our expectations. Knowing that our approach can produce valid results, the next experiments were then designed to get a better feeling for how the different parameters influence our implementation and to establish a consistent baseline to reference later. Finally, pairwise experiments for all possible combinations of coexistence mechanisms were conducted and an experiment combining all three together. In these experiments we made several unexpected observations. We thought that each combination of mechanisms would roughly lead to a multiplicative increase in observed coexistence time, but this only hold true for the combination of storage effect and resource partitioning. Adding intransitive competition to a scenario with strong storage on the other hand proved to be detrimental to coexistence, reducing observed persistence times by about 30%. The combination of resource partitioning and intransitivity, however, was highly potent, producing coexistence times on another order of magnitude entirely. Interestingly the combination of all three mechanism produced scenarios significantly more stable than any other tested scenario but only if intermediate levels for all coexistence mechanisms were used. For very high levels of all mechanisms, the addition of storage to the otherwise highly potent combination of resource partitioning and intransitivity proved to be a hindrance to coexistence. Overall, our results strongly suggest that the combined effects of individual coexistence mechanisms are not trivial to predict and can have a significant impact on observed coexistence times, potentially outright dwarfing isolated predictions in magnitude. It must be assumed that the presence of more coexistence mechanisms in natural systems only complicates and stresses the intricate interplay of these mechanisms further. However, the results we produced are not thorough enough to allow for broad generalization of our particular findings and opened up as many questions as they answered, considering the causes behind our observations remain largely unknown. Results seem to be strongly scenario dependent and the full range of parameters remains to be tested. Consequently this work should be seen as a foundation on which future research can be build up, highlighting a new largely unexplored avenue for the study of species coexistence. Lastly, we want to suggest a few topics for future research and share some of the more promising ideas we had which may advance our understanding regarding the workings of nature:

-
- There is a need for more detailed analysis and experiments in regards to all of our findings. Time constraints made it impossible to dive deep enough into any single scenario to explore the underlying causes and a wealth of parameter combinations have been ignored. Why do intransitivity and storage effect have negative interactions? What causes coexistence times to explode for high levels of intransitive competition and resource partitioning? How come the combination of all three mechanisms works well only for intermediate levels of each mechanism?
 - Closely related to our first point, there seems to be a significant lack of research for many-species scenarios with large community sizes. Experiments often involve two-species scenarios only, with many models even being incapable of considering more than that, and very rarely more than ten. Similarly community sizes of more than a few thousand individuals are far and few between and our experiments here are no exception. There are good reasons for this, as it keeps the data manageable and results more tractable, but our findings suggest that the dynamics of species coexistence may be strongly influenced by the observed scenario size and complexity. Keeping them artificially simple may ultimately not be sufficient.
 - Here we implemented a model supporting several coexistence mechanisms at once, but our design is by no means complete and could be extended to support even more mechanisms. For some of the most well known ones, we have already given ideas on how they may be incorporated in Section 3.5, but that list is certainly not comprehensive of all possible mechanisms or their best representations.
 - In this thesis we aimed to remain as comparable to other models of contemporary research as possible, especially in adopting the zero-sum assumption for community sizes and neutrality. It might be very interesting to lift these restrictions, instead accepting asymmetric, dynamic and self regulating communities in accordance with the provided resources.
 - Similarly our model, as proposed here, focuses on the dynamics of a community with a fixed number of species. But this certainly does not have to be the case. Integration of the models dynamics into a larger framework, including speciation, migration, or even evolution, is possible and might offer a more complete picture of species coexistence.

- Finally, it might be useful to find a suitable conversion between our model and empirical observations. As it stands now, there is no such connection defined and model parameters, as well as statistics, such as the lifetime of an individual, are valid only in context of the model itself. It follows that our model, in its current version, is not well suited to represent specific natural communities, although possible in theory.

Bibliography

- [1] G. E. Hutchinson, “The paradox of the plankton,” *The American Naturalist*, vol. 95, no. 882, pp. 137–145, 1961. [Online]. Available: <http://www.jstor.org/stable/2458386>
- [2] J. S. Clark, “Individuals and the variation needed for high species diversity in forest trees,” *Science*, vol. 327, no. 5969, pp. 1129–1132, 2010.
- [3] V. Novotný and Y. Basset, “Rare species in communities of tropical insect herbivores: pondering the mystery of singletons,” *Oikos*, vol. 89, no. 3, pp. 564–572, 2000.
- [4] B. J. McGill, R. S. Etienne, J. S. Gray, D. Alonso, M. J. Anderson, H. K. Benecha, M. Dornelas, B. J. Enquist, J. L. Green, F. He, A. H. Hurlbert, A. E. Magurran, P. A. Marquet, B. A. Maurer, A. Ostling, C. U. Soykan, K. I. Ugland, and E. P. White, “Species abundance distributions: moving beyond single prediction theories to integration within an ecological framework,” *Ecology Letters*, vol. 10, no. 10, pp. 995–1015, oct 2007.
- [5] J. A. Miller and P. Holloway, *Niche Theory and Models*. American Cancer Society, 2017, pp. 1–10.
- [6] J. Rosindell, S. P. Hubbell, and R. S. Etienne, “The unified neutral theory of biodiversity and biogeography at age ten,” *Trends in Ecology & Evolution*, vol. 26, no. 7, pp. 340–348, jul 2011.
- [7] P. L. Wennekes, J. Rosindell, and R. S. Etienne, “The neutral—niche debate: A philosophical perspective,” *Acta Biotheoretica*, vol. 60, no. 3, pp. 257–271, feb 2012.
- [8] S. P. Hubbel, “Neutral theory in community ecology and the hypothesis of functional equivalence,” *Functional Ecology*, vol. 19, no. 1, pp. 166–172, feb 2005.

- [9] M. Danino, D. Kessler, and N. M. Shnerb, “Stability of two-species communities: Drift, environmental stochasticity, storage effect and selection,” *Theoretical population biology*, vol. 119, pp. 1–11, 2016.
- [10] S. P. Hubbel, *The Unified Neutral Theory of Biodiversity and Biogeography*. Princeton University Press, apr 2001.
- [11] S. Saavedra, R. P. Rohr, J. Bascompte, O. Godoy, N. J. B. Kraft, and J. M. Levine, “A structural approach for understanding multispecies coexistence,” *Ecological Monographs*, vol. 87, no. 3, pp. 470–486, 2017.
- [12] R. A. Laird and B. S. Schamp, “Competitive intransitivity promotes species coexistence,” *The American Naturalist*, vol. 168, no. 2, pp. 182–193, aug 2006.
- [13] R. S. Etienne, D. Alonso, and A. J. McKane, “The zero-sum assumption in neutral biodiversity theory,” *Journal of Theoretical Biology*, vol. 248, no. 3, pp. 522–536, oct 2007.
- [14] M. Dornelas, S. R. Connolly, and T. P. Hughes, “Coral reef diversity refutes the neutral theory of biodiversity,” *Nature*, vol. 440, no. 7080, pp. 80–82, mar 2006.
- [15] B. J. McGill, “A test of the unified neutral theory of biodiversity,” *Nature*, vol. 422, no. 6934, pp. 881–885, apr 2003.
- [16] F. A. S. dos Santos, K. Johst, and V. Grimm, “Neutral communities may lead to decreasing diversity-disturbance relationships: insights from a generic simulation model,” *Ecology Letters*, vol. 14, no. 7, pp. 653–660, may 2011.
- [17] I. Volkov, J. R. Banavar, F. He, S. P. Hubbell, and A. Maritan, “Density dependence explains tree species abundance and diversity in tropical forests,” *Nature*, vol. 438, no. 7068, pp. 658–661, dec 2005.
- [18] P. Chesson, “Mechanisms of maintenance of species diversity,” *Annual Review of Ecology and Systematics*, vol. 31, no. 1, pp. 343–366, 2000.
- [19] K. Gross, “Positive interactions among competitors can produce species-rich communities,” *Ecology Letters*, vol. 11, no. 9, pp. 929–936, sep 2008.
- [20] P. Chesson, “Multispecies competition in variable environments,” *Theoretical Population Biology*, vol. 45, no. 3, pp. 227–276, jun 1994.

- [21] S. Allesina and J. M. Levine, “A competitive network theory of species diversity,” *Proceedings of the National Academy of Sciences*, vol. 108, no. 14, pp. 5638–5642, mar 2011.
- [22] L. Gallien, N. E. Zimmermann, J. M. Levine, and P. B. Adler, “The effects of intransitive competition on coexistence,” *Ecology Letters*, vol. 20, no. 7, pp. 791–800, may 2017.
- [23] O. Godoy, D. B. Stouffer, N. J. B. Kraft, and J. M. Levine, “Intransitivity is infrequent and fails to promote annual plant coexistence without pairwise niche differences,” *Ecology*, vol. 98, no. 5, pp. 1193–1200, 2017.
- [24] J. N. Griffin and B. R. Silliman. (2011) Resource partitioning and why it matters. *Nature Education Knowledge* 3(10):49. [Online]. Available: <https://www.nature.com/scitable/knowledge/library/resource-partitioning-and-why-it-matters-17362658>
- [25] M. Detto and H. C. Muller-Landau, “Stabilization of species coexistence in spatial models through the aggregation–segregation effect generated by local dispersal and nonspecific local interactions,” *Theoretical Population Biology*, vol. 112, pp. 97–108, dec 2016.
- [26] M. Anazawa, “Interplay between habitat subdivision and minimum resource requirement in two-species competition,” *Theoretical Population Biology*, vol. 120, pp. 90–102, 2018.
- [27] J. Hidalgo, S. Suweis, and A. Maritan, “Species coexistence in a neutral dynamics with environmental noise,” *Journal of Theoretical Biology*, vol. 413, pp. 1–10, jan 2017.
- [28] Epic Games, Inc. Unreal engine 4 website. [Online]. Available: <https://www.unrealengine.com/en-US/what-is-unreal-engine-4>

Declaration of Authorship

I hereby declare that this thesis was created by me and me alone using only the stated sources and tools.

Thomas Seidelmann

Magdeburg, March 7, 2019