SUICIDE, SIGNALS, AND SYMBIONTS: EVOLVING COOPERATION IN AGENT-BASED SYSTEMS

By

Anya E. Vostinar

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Computer Science - Doctor of Philosophy;
Ecology, Evolutionary Biology and Behavior - Dual Major – Doctor of Philosophy

2017

ABSTRACT

SUICIDE, SIGNALS, AND SYMBIONTS: EVOLVING COOPERATION IN AGENT-BASED SYSTEMS

By

Anya E. Vostinar

Cooperation is ubiquitous in nature despite the constant pressure for organisms to cheat by receiving a benefit from cooperators, while not cooperating themselves. The continued evolution and persistence of countless forms of cooperation is a central topic in evolutionary theory. Extensive research has been done on the theoretical dynamics of cooperation through game theory and the natural examples of cooperation. However, it remains difficult to understand thoroughly the evolution of many cooperative systems, due in part to the ancient origins of these systems and the long time scales required to see cooperation evolve in any natural populations. I have systematically analyzed the evolution of three broad types of cooperation: programmed cell death, quorum sensing, and mutualisms (cooperation across species). I have provided evidence that programmed cell death can originate due to kin selection. I have also created two new systems to enable the extensive exploration of factors that affect the evolution of public goods cooperation and mutualism. Using these systems, I determine the effects of environmental factors on the evolution of public goods cooperation and mutualism. By uniting the expansive theoretical work on these forms of cooperation with a fully-controlled experimental system, I contributed to our understanding of how these forms of cooperation can emerge and be maintained in industrial and medical applications that rely on bacterial cooperation.

For my parents, Barbara and Keir Johnson.
Thank you for making sure I always believed I could do anything and showing me that women, and therefore I, belong in science and technology.

ACKNOWLEDGEMENTS

I am thankful to many, many people for helping me get to this point. First, I thank my family: Barbara and Keir Johnson, Grandma Dot, and Joe Vostinar for their support throughout my undergraduate and graduate experiences. I especially appreciate their time listening to me going on and on about some ant or bacterial species I just read about.

I have been very fortunate to have Dr. Charles Ofria as my Ph.D. advisor and am forever grateful for how much he taught me about being a scientist and a person. In my first year of my Ph.D., Dr. Heather Goldsby volunteered to temporarily advise me in Charles' absence and I cannot imagine any better start to graduate school. I greatly appreciate the many hours she spent teaching me how to write and conduct experiments, helping me to build a strong foundation for the rest of my career. Finally, my undergraduate research advisor, Dr. Sherri Goings, changed my life in several ways, but in particular I want to thank her for teaching an elective that sounded interesting and inspired me to pursue an academic career.

Finally, I had many fantastic conversations throughout my studies. For those, I wish to thank all the members (past and present) of the Digital Evolution Laboratory, the members of my Ph.D. committee – Dr. William Punch, Dr. Chris Waters, and Dr. Christoph Adami – and members of the BEACON community for inspirational seminars, workshops, and random discussion in the bullpen. In particular, I would like to thank Josh Nahum for finding a particular paper at a critical moment in my research, and Emily Dolson for her friendship and letting me stay in her spare bedroom.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Darwinian evolution is often portrayed as a "tooth and claw" [87] fight between individuals for survival [23]. However, interactions among organisms come in many forms, one major category of which is cooperation. Cooperation is any interaction that benefits a recipient at a potential cost to the actor, but that is favored by natural selection anyway due to either reciprocal cooperation or kin selection [94]. As Darwin suspected, researchers have found vast amounts of evidence that organisms of every size cooperate to improve their chances of survival [92, 27, 37].

1.1 Types of Cooperation

Cooperation in nature comes in many forms including among groups of closely-related individuals, among more distantly-related members of the same species, and even among organisms of different species [92].

Cooperation within a **group of related individuals** encompasses many of the best known types of cooperation, such as familial care among elephants traveling together, pack hunting in wolves or other large predators, and eusocial behavior in colonies of bees, ants, or other insects [24]. Even bacteria frequently engage in cooperation among groups of related individuals, including many cases that can affect human health, both positively (e.g., microbiomes) and negatively (e.g., diseases such as *E. coli*) [99].

Generally this form of cooperation works because the individual that is benefiting also has a copy of the cooperative gene, leading to the continued existence of that gene in the population even if the cooperative behavior kills (or otherwise reduces the fitness of) some of the individuals, as long as there is a net gain of fitness to kin that possess the gene [24]. Of course, in any cooperative system there is the possibility that a *cheater* will arise: a mutant that does not have the cooperative gene but still receives benefits from its relatives that do. The logic of natural selection seems to lead to that mutant doing disproportionately well, having many more offspring, and taking over

the population, leading to the collapse of cooperation. On the other hand, it's hard to imagine how that cooperative trait got going in the first place, since a single cooperator that has just been born surrounded by non-cooperators probably won't last long. There are many factors that could allow for this form of cooperation to evolve and flourish, but it is difficult to test all the possibilities in any single system. Therefore, in Section 2, I will be using a digital system to test some of the possibilities.

Cooperation is not limited to closely related individuals, however. Some notable forms of cooperation occur in **large groups within a single species**, but not necessarily closely related. Schools of fish, swarms of insects, and flocks of birds help the group avoid predators on average, no matter how distantly individuals are related. However, it is still better to be the fish on the inside of the school than the outside, leading to selective pressure to get a better position [24].

Many species of bacteria also engage in this form of cooperation. For example, *Vibrio fischeri* cooperate to produce the brilliant bioluminesence visible at night in many places of the ocean, and, less obviously, *Erwinia carotovora* – a plant pathogen – breaks down cell walls that are too difficult for an individual bacterium by waiting until they get to a high enough concentration in the plant before suddenly making the plant sick [75].

To coordinate these large groups of non-kin, there needs to be some kind of communication, information passed from one organism to another until the entire group knows. Often this communication is not the verbal kind that we are used to using; instead, bacteria use small chemicals that diffuse through the environment and detect how much of that chemical is around them [65]. This way, an individual can get a rough idea of whether the colony is big enough for some sort of mass action that is only beneficial in large quantities. Once there are enough individuals all producing the signal and listening, they can suddenly switch on and avoid wasting precious resources. This type of system is again susceptible to cheaters: individuals that don't switch on with the group but still get the benefit produced by the group. These individuals should then have more resources, produce more offspring, and again take over the population. Determining how that is prevented in natural systems and if it works in the same conditions where the cooperative communication first

evolved is vital to our being able to control the group action of many types of bacteria, both to allow us to increase and decrease the cooperation depending on the effect. To better analyze several possible factors, I created a new digital system to explore this form of cooperation in Section 3.

Finally, cooperation can occur among **individuals of different species**, and this type of cooperation is called a mutualism [99]. Virtually every known species is likely involved in a mutualism of some kind, from the well-known examples of cleaner fish and pollinators to the countless 'helpful' bacteria that we are discovering in ourselves, and every other species we examine, to mutualism between disease-causing bacteria and viruses that are difficult to fight due to rapid evolution [99]. When the benefit of cooperation is going to a member of a different species, it is clear that sharing genes cannot be the root cause as neatly as it can be within a species. This realization has made the evolution and maintenance of mutualism a topic that researchers have spent decades theorizing about, while also being a topic that is so complex that it has been difficult to test any of that theory [55]. By creating a digital system called *Symbulation*, I have been able to directly test and verify theoretical models as well as combine factors in ways that are simply impractical in a theoretical model. In Section 4 I present that system along with several experimental tests of long-debated and recent theoretical hypotheses.

1.2 Mathematical Frameworks

When systematically studying the evolution of cooperation, analyses must eventually boil down to measuring risks versus rewards and identifying the expected costs and benefits of the cooperative act. Therefore, we can use the mathematical framework of Game Theory to more formally categorize the three types of cooperation that I am considering in this work.

Game theoretic scenarios can often pose interesting puzzles for the players, where exact analysis is not always easy. A well-studied two-player game is the *Prisoner's Dilemma*. In a single round of this game, a player will always score better by defecting than by cooperating with their opponent, leading to an outcome where both players defect. The challenge comes in that players will do better if both cooperate rather than both defect, but far worse if they cooperate and their op-

		Player 2	
		C	D
Player 1	C	2, 2	0, 3
	D	3, 0	1, 1

Figure 1.1: **Example Prisoner's Dilemma payoffs.**

ponent defects, as shown by the payoffs in Fig 1.1. Indeed, game theoretic calculations [77] show that without information about the other players, defection is the only stable strategy in single iterations of these games. However, if there is information (often gained through repeated interactions between the same players), cooperation can become stable. A similar dynamic can also be seen in the Public Goods game, where multiple individuals have the opportunity to take advantage of a shared resource, often leading the to a tragedy of the commons [44]. Therefore, within either the Prisoner's Dilemma or the Public Goods game, the forms of cooperation can be classified by the type of information they rely on:

1. **Spatial information** - Perhaps the most basic form of information that an organism can rely on is the relative positions of themselves and the other players. If the environment is such that close genetic relatives tend to be near each other – either through some mechanism of self-selection or because offspring tend to stay near their parents – then cooperative organisms can assume they will disproportionately interact with other cooperators near them. As shown in Adami *et al.*, adding spatial dynamics to game theory makes analytical solutions difficult and therefore agent-based systems are more practical to determine the conditions under which cooperation can evolve and persist [2]. I analyze a form of cooperation via spatial information in Chapter 2.

2. **Communication information** - If organisms are able to use some sort of messaging, cooperation can be stable. Communication can allow organisms to identify other cooperators and

selectively interact with them. Therefore, organisms do not need to be spatially near other cooperators all the time for cooperation to succeed. Generally there is both a cost to communication and cooperation, leading to increasingly complicated dynamics. Game theory is able to analyze games where communication is enabled [3, 62], however spatial dynamics (in that biological organisms are almost always non-randomly assorted in space) again make agent-based simulations more practical to understand when cooperation will evolve and persist. I explore a form of communication called quorum sensing in Chapter 3, where organisms are able to communicate only their presence and detect the number of cooperators in their immediate area in a public goods game.

3. **Species information** - If cooperation is between different species, as is the case in biological mutualisms, the unique features of a species provide information to its partner. There is often some kind of communication between partners, which would make that instance a subset of the second category. However, because the organisms are from different species, they may not be in direct competition and as such there is the possibility of a net gain from cooperation due to resource specialization and differential need. For example, the acacia tree cannot physically attack herbivores that eat its leaves, but it can provide nutrition to an ant colony, and if that ant colony in return fights off herbivores, there is a net gain for organisms of both species. There are many theoretical models as to how mutualisms might evolve and be maintained, and there are many well-known examples of mutualisms in nature [55, 91, 59, 82]. The field is at a point where these two sources of information can be connected through experimental systems in ways that were previously impossible. In Chapter 4, I explore the factors that lead to the evolution and persistence of mutualism.

Of course, many other mechanisms are possible for gaining information about the other players in a game-theoretic scenario. For example, organisms may be able to play multiple games with each other to gain trust or observe each other interacting to create community-level reputations. As I am not focusing on any of these alternative mechanisms in this proposal, however, I am limiting

detailed descriptions to the three above.

1.3 Model Systems in Which to Study Cooperation

There are three main categories of model systems in which cooperation is frequently studied: mathematical models, biological model organisms, and agent-based computational models.

William D. Hamilton determined an elegant inequality, called Hamilton's Rule, that describes whether an organism should perform a costly behavior (with cost C) based on the relatedness of the other organism (r) and the benefit (B) to that other organism: $C < r * B$ [42]. As discussed previously, game theory is a field of mathematics devoted to the study of interacting agents, and it has often been used to study cooperation [45, 3, 2]. There have since been many studies that find conditions under which cooperation can persist, such as when individuals can remember what the other has done previously [62]. Evolutionary game theory has furthered these ideas by putting strategies into an ecological context (i.e. a population of individuals with one of several strategies) and determining what strategies dominate when a population of organisms is interacting [2]. By determining mathematically what benefit there should be in a system free of noise, researchers can form predictions of what may happen in a more realistic system. However, game theory assumes populations have infinite size and are under a strong-selection, weak-mutation regime (SSWM). SSWM means that it is assumed that when a mutation occurs, selection is strong enough that that mutation either sweeps the population or disappears before another mutation arises. However, as demonstrated in Adami *et al.* [2], even *E. coli* will, on average, have multiple mutations in a population at any given time, putting it, and likely most other biological organisms, in the weak-selection, strong-mutation regime (WSSM).

To truly understand cooperative traits in nature, we have to study biological organisms. There is a large amount of research into cooperation among natural organisms, from humans to insects, large carnivores to microscopic herbivores, and a myriad of other species in between [93, 76, 78]. Bacteria are especially useful for understanding the basic evolutionary dynamics of cooperation because there are many instances of cooperation between single-celled creatures, they have relatively

quick replication times, and they can be revived after being frozen making possible competitions between ancestral and derived strains [96]. However, even bacteria take hours to reproduce in a sustained experiment and therefore months or years to experiment with when studying long-term evolutionary dynamics. Further, while genetic testing has progressed at a staggering rate, it is still difficult and expensive to track every cooperative trait in which I may be interested in an evolving population of bacteria.

That ongoing challenge brings me finally to agent-based computational models of evolutionary dynamics, spanning from basic implementations of mathematical models, to complex and powerful artificial life platforms such as Avida [71]. By using computers to implement the necessary components of evolution – variation, inheritance and competition – we can 1) run evolutionary experiments for many more generations *per second*, 2) have perfect data recording and perfect control of the experiment, and 3) incorporate finite populations and a WSSM regime. Due to these benefits, I use Avida along side simpler artificial life systems, depending on the particular experimental needs of the project.

1.4 Methodology

As discussed previously, there are many possible systems in which to study the evolution and maintenance of cooperation, raising the question: which should be used to determine the factors leading to the evolution of cooperative traits? Observing and experimenting on the maintenance of cooperation is often done in a wet lab system, allowing researchers to catalog what cooperative systems actually exist in nature and determine how they will continue to evolve [92]. However, it is nearly impossible and certainly impractical to try to study the origin of cooperation in a wet lab system because either the system already has a cooperative trait (and thus you cannot study the origin) or it does not and you would have to hope that such a trait will evolve in your system if you wait long enough.

Some of the mathematical models that have been presented previously allow for studying the emergence of a cooperative trait when it is rare, thus giving an indication of how such a trait could

initially persist in a non-cooperative environment. However, the same issues with mathematical models discussed previously still apply: many interacting factors are likely involved in natural systems and mathematical models become unwieldy when stretched to include all known relevant factors [4]. To use a mathematical model, the researcher must make a number of assumptions about the system to simplify the dynamics. Due to the sometimes fragile nature of cooperation, this lack of noise and interacting dynamics can be misleading.

Therefore, here I use digital systems to bring together the strengths of mathematical and biological models. In a digital system, I am able to fully control the environment and organisms, similar to a mathematical model, but I am also able to combine all known relevant interacting factors, similar to a biological wet lab system. However, just as it is important to carefully choose a broad category of model, it is also important to choose which digital system to use. Some are large and as close to wet lab experiments as they can get [71], and in fact more complex than necessary for studying the evolution and maintenance of cooperation, even when factoring in many possible dynamics. Here, I use the complex digital system, Avida, when it is most appropriate, and I have also designed two new digital systems tailored to the needs of my research questions. Details about each will be presented in the chapter in which they are used.

1.5 Contributions

My central thesis in this work is that the origin and maintenance of cooperative traits are separate stages of evolution that have historically been conflated. The possible selective pressures for each stage are vital to our understanding of the system as a whole and therefore must be analyzed separately. The many possible factors affecting the evolution of both of these stages of cooperative evolution are far too complicated to tease apart in natural systems. Therefore, targeted models should be used to disentangle these possible factors, which can be used to inform future, more targeted studies of natural organisms.

By understanding the evolutionary pressures that lead to cooperation evolving and persisting in a population, I can improve our approaches to evolving communities to 1) better understand

natural populations and aid in conservation efforts, 2) identify methods of disrupting cooperation in pathogenic microbes to improve human health, 3) augment useful microbial populations to increase production of targeted biomolecules, and 4) use cooperation in evolutionary algorithms for solving computational and engineering problems.

Specifically, first I have contributed to our understanding of how altruism can come to exist in natural systems and what effects it has on the continued evolution of those systems. I did this work by determining the effects of factors leading to the evolution of cooperation with spatial information by creating two models of kin altruism in Avida. Specifically, I used a programmed cell death trait because by making the altruist unable to receive a later reciprocal benefit, I am able to clearly differentiate altruistic effects from reciprocal cooperation. Using both an direct-benefit and indirect-benefit example of programmed cell death, I have found that increased mutation rates and large population sizes increase the *de novo* evolution and persistence of programmed cell death. I have found that indirect-benefit programmed cell death evolves into a population more easily and is more persistent in the presence of cheating organisms. I also found that mutation rate strongly affects both forms of programmed cell death, whereas population size does not have a significant effect until very large population sizes are used [54, 53]. These results are presented in Section 2.

Second I have determined how cooperation evolves when basic communication information is available, by working directly with biologists to model a public goods behavior that is controlled by quorum sensing, based on *Vibrio harveyi*. I have created a digital system where organisms can produce a public good for a cost and receive a benefit if enough surrounding organisms are also producing that public good. To enable the organisms to collect information that will allow them to better decide whether producing the public good is worthwhile, they are also able to signal their presence and sense the presence of surrounding signaling organisms; that is, they can quorum sense. With this model, I have performed competitions (in parallel with my wet-lab collaborators) to compete defector, unconditional cooperator, and quorum-sensing cooperator strains to find what benefits enable the quorum-sensing cooperator to resist invasion by the defector. I have explored how spatial structure, unrestricted resources, and positive feedback mechanisms affect both the

origin and maintenance of this cooperative trait, as detailed in Section 3.

Finally, I have isolated some of the factors leading to the evolution of cooperation with species information (i.e. the evolution of mutualisms) by connecting the plethora of theoretical work in this area with a controllable experimental system. I have created a host-symbiont digital system to enable symbionts to inhabit hosts and cooperate by exchanging resources with the host, making them mutualists, or attempting to steal resources from a defensive host, making them parasites. With this system I have investigated how the likelihood of a mutualism evolving and persisting is affected by 1) the vertical transmission rate of the system – i.e. how likely it is that a symbiont will pass to a host offspring, 2) the starting conditions of the symbiont and host populations – for example if it is a parasite or commensalist (i.e. benefit without a cost), 3) the amount of division of labor between host and symbiont, and 4) the spatial structure of the environment. I have confirmed several recent theoretical findings with vastly relaxed assumptions from the mathematical models and found where those theoretical models break down, presented in Section 4.

To summarize, in the remainder of this dissertation I present the results of evolving two versions of programmed cell death under varying environmental conditions (Section 2), evolving cooperation controlled by quorum-sensing in a new digital system (Section 3), and the likelihood of mutualism evolving between two species in varying conditions in an additional new system (Section 4). Finally I conclude by observing general trends regarding the factors leading to the evolution of the origin or maintenance of cooperative traits and provide more details for how these results may promote further research beyond the scope of this dissertation.

CHAPTER 2

SUICIDAL SELECTION: PROGRAMMED CELL DEATH CAN EVOLVE IN UNICELLULAR ORGANISMS DUE SOLELY TO KIN SELECTION

2.1 Abstract

Unicellular organisms can engage in a process by which a cell purposefully destroys itself, termed programmed cell death (PCD). While it is clear that the death of specific cells within a *multicellular* organism could benefit the whole (i.e. during development), the origins of PCD in *unicellular* organisms is less obvious. Kin selection has been shown to help maintain instances of PCD in existing populations of unicellular organisms, however, competing hypotheses exist about whether additional factors are necessary for its origin. Those factors could include an environmental shift that causes latent PCD to be expressed, PCD hitch-hiking on a large beneficial mutation, and PCD being simply a common pathology. Here we present results using an artificial life model to demonstrate that kin selection can, in fact, be sufficient to give rise to PCD in unicellular organisms. Furthermore, when benefits to kin are direct – i.e. resources provided to nearby kin – PCD is more beneficial than when benefits are indirect – i.e. non-kin are injured, thus increasing the relative amount of resources for kin. Finally, the amount PCD is evolved to be used is influenced by the form (direct vs. indirect) of benefit.

2.2 Introduction

In programmed cell death, a cell destroys itself through internally controlled processes [67]. Initially, it may seem puzzling that such apparently undesirable behavior has not been eradicated by natural selection [67]. However, if the cell is part of a multicellular organism, programmed cell death has the potential to increase the fitness of the overall organism and thus be under positive selection [51]. For example, the purging of damaged or diseased cells to maintain health, or the restructuring of tissue to facilitate normal development [57]. However, many types of behavior

resembling programmed cell death have also been identified in unicellular organisms [60]. This situation is less clear given that programmed cell death in a unicellular organism kills the entire organism, eliminating any direct selective pressure for this trait. This observation raises the question: How can programmed cell death evolve at all in unicellular organisms?

One adaptive explanation for programmed cell death in unicellular organisms is that surviving kin receive some benefit from the death [37]. However, alternative non-adaptive hypotheses for the origin of programmed cell death have recently been proposed [67, 30]. These alternative hypotheses include: 1) programmed cell death is a side effect of behavior evolved in previous environments, 2) programmed cell death is caused by a hitchhiking and/or pleiotropic gene wherein a beneficial trait offsets this undesirable behavior, 3) programmed cell death is a pathological breakdown of the cell's functioning and is not under selection due to benefits provided to kin, or 4) the historically-accepted scenario of kin selection (where selection for programmed cell death is due to the benefits provided to kin instead of directly benefiting the organism itself [37]).

Determining the likelihood that these scenarios resulted in the evolution of programmed cell death is difficult for a number of reasons: scenarios one and two require knowledge of past evolutionary conditions to evaluate definitively; unfortunately, this detailed knowledge does not exist for many ancient organisms. Testing scenario three would require experimental evolution treatments that control whether a benefit is provided to kin or not. This is difficult to accomplish in an organic system due to both time constraints and current technological capabilities.

Because there is no simple way of evaluating these alternative hypotheses using historical data or within organic systems, we use self-replicating digital organisms, a type of computational modeling that evolves agents in a complex digital environment, to determine whether programmed cell death could evolve due to kin selection without any other possible scenarios. For this study, we use Avida, which is a digital evolution platform used to study evolutionary questions, including complexity [61], division of labor [41], and communication [9]. An Avida experiment consists of a virtual world where digital unicellular organisms, or 'unicells,' can evolve for thousands of generations over hours with perfect control and data collection. Using Avida, we implemented

programmed cell death behavior. We then observed under what conditions programmed cell death was able to evolve into a population *de novo*.

Our hypothesis is that programmed cell death can evolve as a response to specific types of kin selection. For this work, we categorize the potential benefits of programmed cell death into two major categories: direct benefit to kin and indirect benefit to kin through harm to competitors.

In directly beneficial programmed cell death, when an organism dies, its surrounding kin are provided a competitive advantage, for example *Dictyostelium discoideum* where individual organisms die to lift the spores of those remaining for better propagation [85].

In indirectly beneficial programmed cell death, when an organism dies the surrounding non-kin are set at a competitive disadvantage, such as colicin production in *E. coli*, which kills any organisms without the resistance gene after the focal organism has burst [22].

These two forms of programmed cell death occur throughout nature, but are difficult to directly compare due to the need to control for the many other differences between species that exhibit one form or the other [30, 11, 84].

Using the Avida software, we created two forms of programmed cell death that differed only in whether the effect directly increased the fitness of kin or decreased the fitness of competitors. We provide experimental evidence that a benefit to kin is sufficient to cause the evolution of programmed cell death. Furthermore, we confirmed that direct benefits made programmed cell death more beneficial than indirect benefits to kin. Finally, we found that indirect benefits led to a different pattern of response to the degree of kin discrimination than direct benefits. These findings are critical to our understanding of the evolution of programmed cell death as the form of benefit directly changes the conditions under which it will evolve.

2.3 Literature Reviewed

The phenomenon of programmed cell death has been of interest to researchers for the past 25 years, with many examples found, such as Table 1 in [67]. Nedelcu et al. explain how programmed cell death has been assumed to be adaptive due to providing benefits to surviving kin. They go on

to discuss how this assumption may be premature due to the difficulties in determining causality in the evolution of a trait after the fact.

This idea of a benefit to kin directly indicating an adaptive trait originated with Hamilton's concept of 'inclusive fitness' [43]. Traditionally an organism's fitness is measured as the number of offspring it has produced [100]. However, Hamilton realized that because related organisms share many of the same genes, an organism could help its genes be represented in the next generation by aiding in the survival and successful reproduction of its relatives. Dawkins expanded upon this idea by anthropomorphizing an individual gene to be 'selfish' to convey the idea that the gene is trying to get copies of itself into the next generation and an individual gene actually does not generally care what happens to the other genes it happens to share an organism with, except insofar as they are necessary for its own survival [24].

The general idea of an individual not necessarily being the unit of selection in natural selection was further explored by Lewontin, who proposed the idea that before there were organisms, individual molecules could have been under selection in the same way as genes or organisms: the molecule that will persist is the one that is able to replicate itself [63]. Lewontin also pushed this idea in the other extreme and suggested that groups of organisms could be units of selection if the group as a whole was able to propagate itself in the same way that multicellular organisms do. So-called 'group selection' is an alternative hypothesis to 'kin selection', though they are similar ideas and one can generally be translated into the other [69].

While there are many other forms of cooperation, reviewed in [93], programmed cell death must inherently be a form of altruism (if cooperative) because the acting organism is not able to gain a reward later, as is necessary for any form of reciprocal cooperation. Due to this clarity of the form of cooperation, programmed cell death is a useful form of extreme altruism that can shed light on less extreme forms of altruism.

Finally, while cooperation and altruism in particular can be found across all forms of life (from humans [13] to termites [12] to bacteria [95]), studying the evolution of altruism can be challenging due to the extensive evolutionary times required and lack of experimental control available.

However, the many forms of bacteria that perform altruistic behaviors has made them ideal systems for studying social behaviors [95]. Determining the factors contributing to the origin of these traits, however, even in bacteria, is challenging if not impossible. Therefore, we have turned to a digital system.

2.4 Avida Digital Evolution System

For this study, we use the Avida digital evolution system [71]. Avida is a flexible system that has been used for evolutionary studies with different biological analogies, including bacterial cells [9] and eusocial insects [41]. Here we consider each digital organism to be the digital equivalent of a unicellular organism, which we call 'unicells.' As shown in Figure 2.1, Avida consists of a virtual world where unicells compete for space as they reproduce with variation. Each grid position may contain a single unicell, which can replicate into the 8 neighboring grid positions.

Each unicell has a genome consisting of computational instructions from a Turing-complete programming language, a sample of which is shown in Figure 2.1. Each unicell also has a virtual CPU on which its genome is executed. This CPU includes storage space, which consists of registers and stacks.

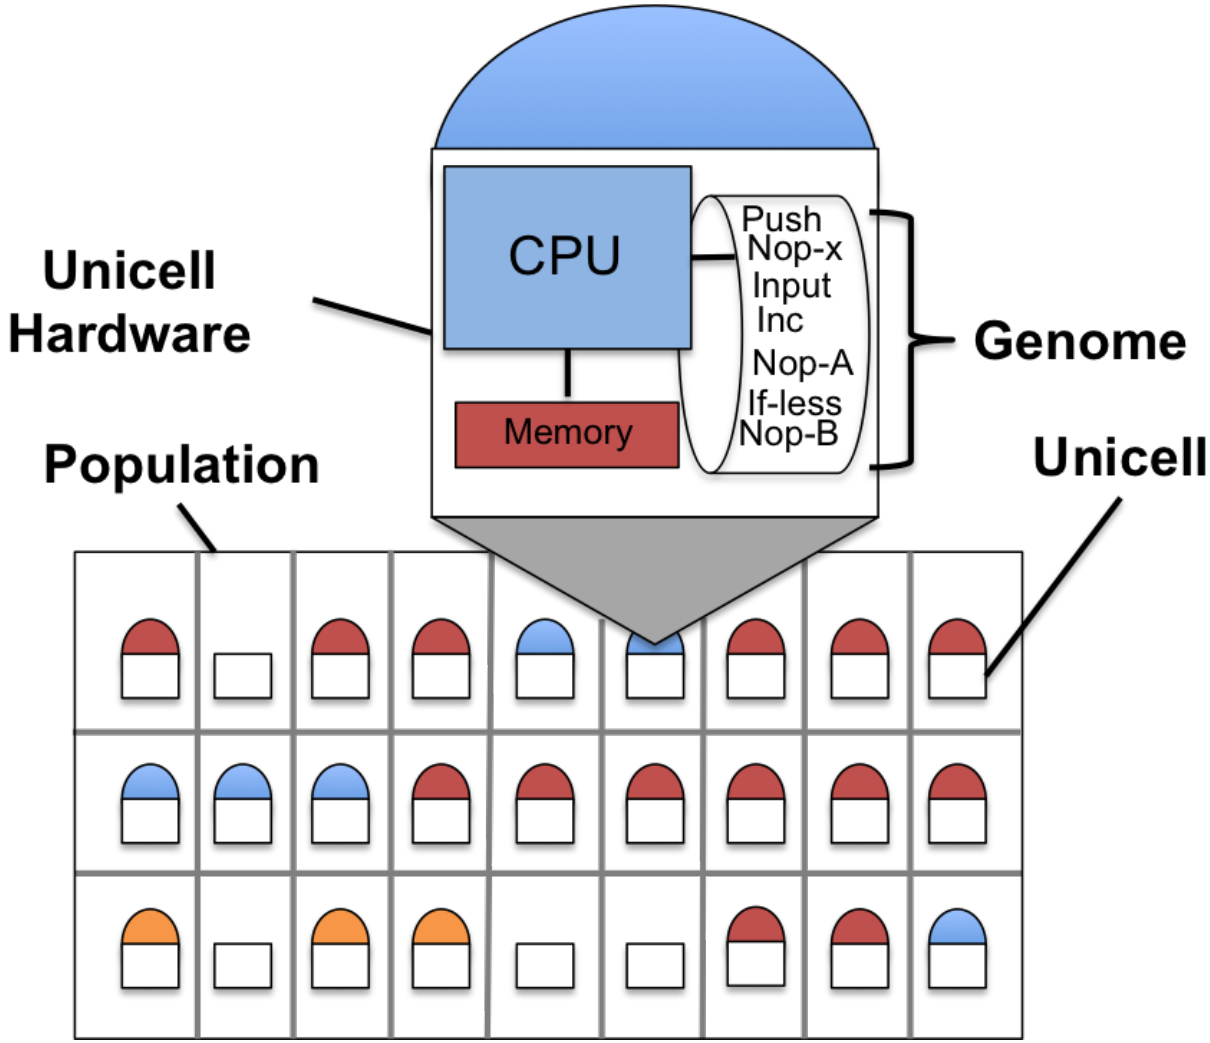

Figure 2.1: **A simplified example of an Avida world.** Half-circles are unicells where color indicates differences in genomes. One unicell's internal hardware is shown, including CPU, memory, and a genome of program instructions.

A unicell is able to reproduce by executing a series of instructions that fully copies its genome, followed by a "divide" instruction to place the resulting offspring into a neighboring space, killing any organism that was occupying that space. This copying process is imperfect, however, and mutations may be introduced. The rate at which these mutations occur is based on user-configured settings. These mutations produce the required inheritable variation for selection to act on, and when combined with competition for limited space, lead to evolution by natural selection.

Within Avida, the metabolic rate of a unicell determines how many CPU cycles it is given compared to the rest of the population. As such, metabolic rate is similar to the amount of energy a natural organism has to forage. A unicell may receive a bonus through the programmed cell death of another unicell. This bonus increases the metabolic rate (and thus number of CPU cycles) of any offspring the unicell has. These extra cycles will allow the offspring to execute its genome faster, and therefore potentially reproduce faster, allowing its genotype to spread more quickly through the population. Conversely, a unicell that is harmed by the programmed cell death of a non-kin unicell will have offspring with a lower metabolic rate (receiving fewer CPU cycles), which will therefore reproduce more slowly. We used the default settings of Avida (found in [70]) with the addition of the two novel programmed cell death behaviors described below.

2.4.1 Programmed Cell Death Instructions

We implemented two programmed cell death instructions to test the importance of direct vs. indirect kin selection. The *direct-kin-benefit programmed cell death instruction* (DIRECT-PCD), when executed successfully, kills the unicell, but increases the metabolic rate of a unicell's kin that are nearby. In contrast, the *indirect-kin-benefit programmed cell death instruction* (INDIRECT-PCD), when executed successfully, decreases the metabolic rate of any unrelated (non-kin) unicells that are nearby. To make sure that the effects of direct versus indirect kin benefits were comparable, we made the behavior of both types of instructions adjust metabolic rate: either directly increasing the metabolic rate of kin or directly decreasing the metabolic rate of non-kin. These two forms of programmed cell death are high-level abstractions of the many ways programmed cell death can

affect surrounding organisms in organic systems, such as producing a toxin, producing an external enzyme necessary for survival, or removing the threat of a bacteriophage from the colony [78].

When a unicell attempts to execute one of the programmed cell death instructions (either DIRECT-PCD or INDIRECT-PCD), there is a 5% chance the unicell will kill itself and cause benefit or harm to surrounding unicells in a 2-space radius on the grid. Clearly the effects of this behavior hinge upon which surrounding unicells are considered kin and which are not. There are two aspects to this determination of kin: 'kin inclusivity level' (i.e., how many genetic differences are tolerated in kin) and kin recognition (i.e. how accurate a unicell is at identifying related organisms).

We define 'kin inclusivity level' (KIL) to be a measure of how distantly related organisms may be before they are considered non-kin by the programmed cell death instructions. Specifically we use the number of genetic differences between two unicells to measure their kinship. For example, a kin inclusivity level of three – the default in our experiments – means that a nearby unicell whose genome has up to three genetic differences from the unicell executing the programmed cell death instruction will be considered kin (see Figure 2.2). In experiments with direct benefits (DIRECT-PCD-3), this kin unicell's metabolic rate will be multiplied by 5. Conversely, in experiments with indirect benefits (INDIRECT-PCD-3), any unicell in the radius with more than three genetic differences will have its metabolic rate divided by 5. The ancestor unicell's metabolic rate is initially the size of its genome, in this case 100. Therefore, unicells with a benefit will on average be able to copy their entire genome and produce an offspring before unicells without the benefit are able to copy one fifth of their genome. Conversely, unicells with the penalty will, on average, only be able to execute instructions at a rate of one fifth compared to non-penalized unicells. This effect is extreme with the goal of reflecting systems in which the focal organism releases colicins to kill nearby competitors or destroys a virus that would otherwise be deadly to surrounding kin [78].

Within our system, unicells' kin recognition is automated as part of the execution of the programmed cell death instructions. In particular, kin recognition within this system is perfect and uses the number of genetic differences between two unicells to compute kinship. The kin distance

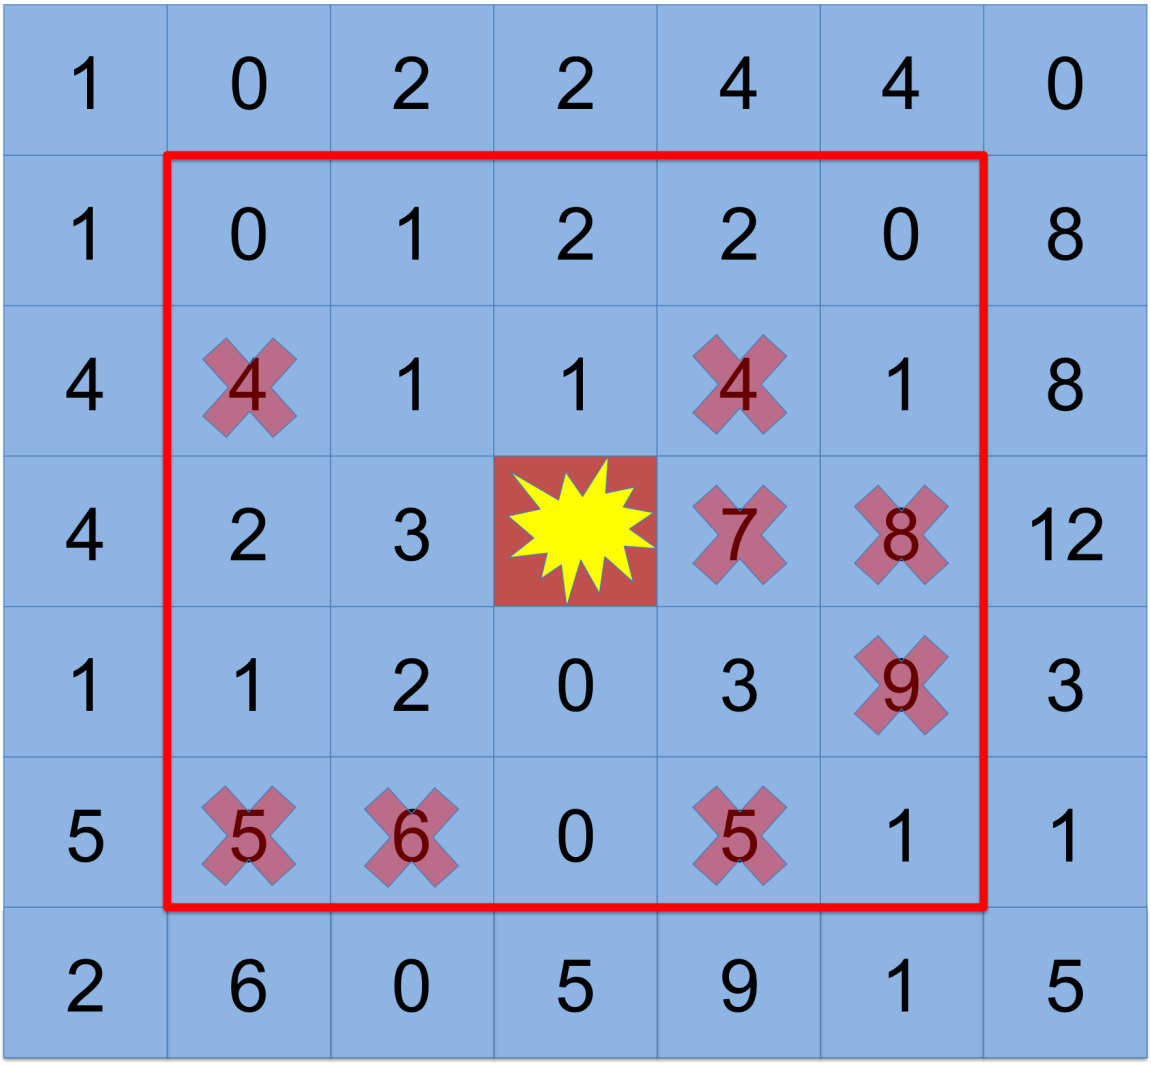

Figure 2.2: **An example of a unicell undergoing programmed cell death.** Each square represents a unicell. The number in each cell is a number of genetic differences between the unicell in that space and the focal unicell depicted as an explosion. Every unicell in a 2-space radius (red line) is evaluated as kin or non-kin. Non-kin are marked with red X's. If the programmed cell death has direct benefit, unicells within the radius without a red X will receive the benefit. If the behavior has indirect benefit, unicells with red X's will be harmed.

is the Hamming distance between the focal unicell and every other unicell in a 2-cell radius. Hamming distance is the measure of the number of differences between genomes when genomes are a fixed size. To decrease computational overhead, we required all unicells to have a fixed genome of length 100, allowing for Hamming distance to be used instead of Levenstein distance. This kin-recognition system is an idealized form of the many kin-recognition systems found in organic organisms [49] and can be thought of as focusing on, for example, only the genes that encode a particular external protein or the presence of a plasmid [22].

2.5 Methods

Here we outline the factors we considered in the *de novo* evolution of both forms of programmed cell death. We ran each experiment for 60,000 updates [1] or approximately 2,000 generations. For every treatment, we ran 30 replicates. All experiments were in structured populations. Avida is open-source at: https://github.com/devosoft/avida and our implementation is at commit 102dd2071558ccd981a2cdab7af47f99ff8d4299. All analysis scripts and data are at https://github.com/anyaevostinar/SuicidalAltruismDissertation.

2.5.1 Kin Inclusivity Level

Our previous work has suggested that the level of kin discrimination can greatly alter the benefits of programmed cell death [54]. If too many unicells are considered kin (i.e. kin inclusivity level is high and unicells with many genetic differences are considered kin), cheating unicells that do not posses the programmed cell death gene may gain the benefit, decreasing the relative inclusive fitness for that gene. (Inclusive fitness is a measure of fitness that takes into account all copies of that unicell's genetic variants in the population, including its siblings and their offspring, for example [86].) However, if too few unicells are considered kin (i.e. kin inclusivity is low and few genetic differences are tolerated), mutants that contain the programmed cell death gene or contain other beneficial traits are considered non-kin and thus fail to invade, decreasing the population's

[1] An update is the time unit in Avida during which unicells execute, on average, 30 instructions

evolvability.

To explore how kin inclusivity levels influence the evolution of programmed cell death, we tested a wide range of kin inclusivity levels (KIL): 0, 1, 3, 10, 20, 30, and 100. Because our unicells were required to have genomes of length 100, these treatments span the full possible range of kin inclusivity from only clones being considered kin (KIL 0) to every possible unicell being considered kin (KIL 100).

2.5.2 No-Effect Controls

To assess whether kin selection can be the sole driver of the evolution of programmed cell death, we ran controls for every replicate where the programmed cell death killed the unicell but had no other direct or indirect effect on kin. In the experimental treatments, the effect is 5: for direct benefits, metabolic rate of kin was multiplied by 5; for indirect benefits, metabolic rate of non-kin was divided by 5. In our controls, we simply set the effect to the identity 1, so that nothing changed in our configurations except the effect to metabolic rate. This control allows us to measure the exact effect of there being a benefit – direct or indirect – to kin on the evolution of programmed cell death.

2.6 Results and Discussion

The key question addressed by this study is: can programmed cell death evolve strictly due to the benefit conferred to surrounding kin? To begin to answer this question, we created an environment where it was not possible for the alternative hypotheses presented previously to have an effect – unicells could not die due to pathology, gain fitness through other tasks, or maintain adaptations to a previously different environment.

We then made the programmed cell death instruction (DIRECT-PCD) available via mutation. This programmed cell death behavior was stochastic in that it had only a 5% probability of successfully killing the unicell and conferring benefits. Successful use of this behavior increased the metabolic rate of surrounding kin unicells. We compared the number of times that the behavior

was attempted and benefited kin to how much it was attempted when no benefit was given to kin (a control where the active organisms merely died with no benefit). As shown in Figure 2.3, we found that 12.52% of unicells executed the programmed cell death instruction when there was a benefit to surrounding kin, compared to 0.09% when there was not a benefit to surrounding kin (Mann-Whitney U test $p < 2.2e-16$). This result provides evidence that, under these environmental conditions, kin selection is sufficient as the sole selective pressure for the evolution of programmed cell death.

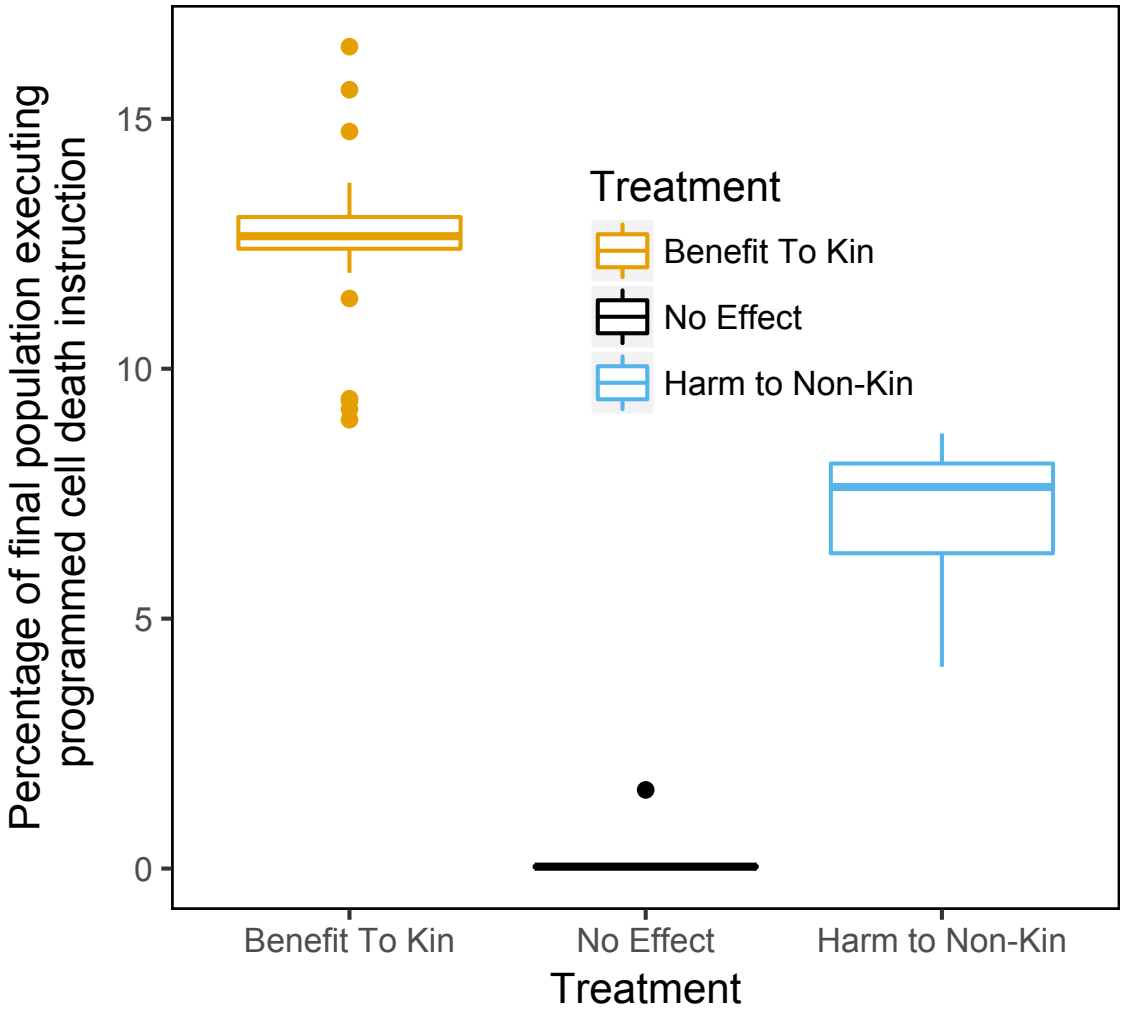

Figure 2.3: **When there is a direct benefit to kin, programmed cell death evolves to occur in 12.52% of the final population, on average.** However, when that benefit is removed, the behavior does not evolve into the population. When there is an indirect benefit to kin, programmed cell death occurs in 7.11% of the final population on average. However, when the indirect benefit is removed, the behavior does not evolve into the population.

2.6.1 How do indirect benefits affect programmed cell death?

There is another category of programmed cell death, however, that could exhibit different evolutionary dynamics: programmed cell death that damages non-kin instead of directly benefiting kin. Many forms of programmed cell death involve the focal organisms producing harmful substances that only impact non-kin, such as colicinogenic *E. coli* [22]. However, clearly such an indirect benefit only aids kin when non-kin are present within the local environment. The indirect nature of this benefit could decrease the viability of kin selection for programmed cell death.

To investigate this possibility, we provided organisms with the INDIRECT-PCD instruction, which they could incorporate into their genomes. When executed, this instruction had a 5% chance of triggering PCD, which would kill the host organism as well as decrease the execution speed (metabolic rate) of any non-kin in the surrounding area. We again compared our experimental results to a control treatment where there was no effect on surrounding non-kin or kin unicells. As shown in Figure 2.3, we found that when there was an indirect benefit to kin resulting from harm to non-kin; 7.11% of unicells executed the programmed cell death instruction compared to 0.09% when there was no harm done to non-kin (Mann-Whitney U test $p < 2.2e-16$). This result shows that an indirect benefit to kin is also sufficient for programmed cell death to evolve due to kin selection in conditions where it otherwise would not, in agreement with previous studies that suggested kin selection as a possible mechanism for this behavior [1, 72, 34].

There are, however, differences in the evolution of programmed cell death when it provides direct benefits, as compared to when it provides indirect benefits. As seen when comparing treatments in Figure 2.3, indirect benefits reduce how strongly selected programmed cell death is. When the benefit to kin is direct, 12.52% of unicells executed the programmed cell death instruction by the end of the experiment compared to 7.11% when the benefit was indirect (Mann-Whitney U test $p-value = 2.493e-13$). This nearly twofold difference in use demonstrates that the form of benefit significantly impacts the selection for programmed cell death and should be considered when analyzing organic systems. This result suggests that kin selection could favor a direct benefit, such as freeing resources [30], over an indirect benefit as found in colicinogenic *E. coli* [22].

2.6.2 How does kin inclusivity level affect programmed cell death?

Finally, the accuracy and degree of kin discrimination determines how many unicells are considered kin or non-kin and therefore is likely to have a large impact on the evolution of programmed cell death under kin selection. To explore this behavioral factor, we varied the kin inclusivity level (the number of genetic differences necessary to qualify as non-kin) used for both direct and indirect benefits. As discussed in the Methods section, our system uses perfect genetic information of length-100 genomes to determine kinship and should be considered a simulation of a kin-recognition system based on, for example, the form of an external protein [22].

As shown in Figure 2.4, we found that the kin inclusivity level has a significant effect on the evolution of programmed cell death for both direct and indirect benefits. When the benefit is direct, at a kin inclusivity level of zero, 7.58% of unicells execute the programmed cell death instruction on average compared to 0.05% when the kin inclusivity level is 100 (Mann-Whitney U test $p - value < 2.2e - 16$).

Similarly, at a kin inclusivity level of zero, 7.33% of unicells execute the programmed cell death instruction with indirect benefits, as compared to 0.09% when the kin inclusivity level was 100 (Mann-Whitney U test $p - value = 2.897e - 11$). Note that the value of performing programmed cell death at kin inclusivity level zero is very low, as a result of a lack of variation in the population. At that level, mutants of any kind are severely penalized and therefore little genetic variation is able to persist. This result shows that programmed cell death can evolve due only to kin selection if there is an accurate and discriminatory form of kin recognition, though the exact degree of discrimination for an optimal benefit from programmed cell death clearly will rely on factors such as average mutation rate, as we have found previously [54].

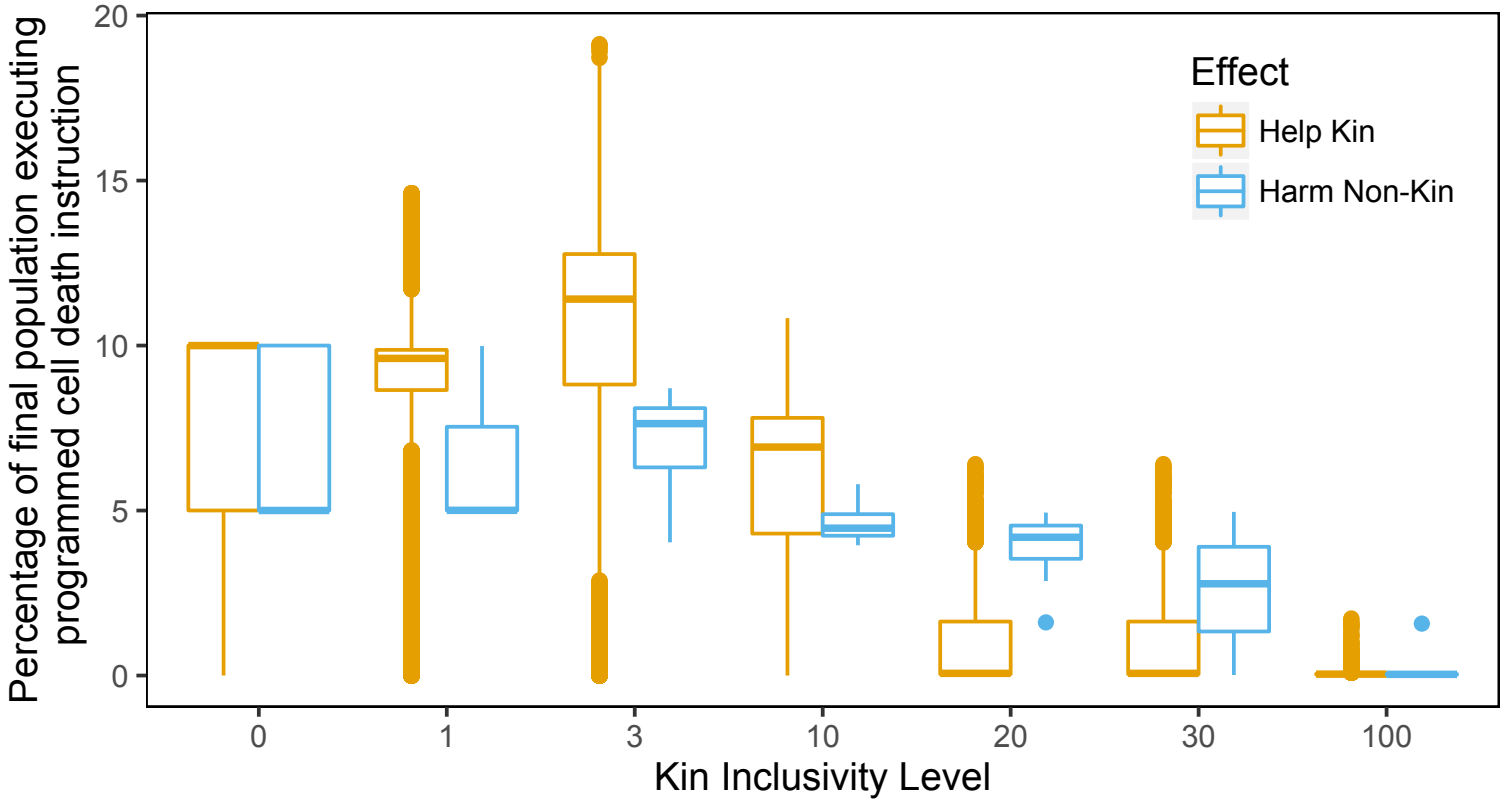

Figure 2.4: **Percentage of final population that performed programmed cell death when benefits to kin were direct or indirect with varying kin inclusivity levels.** At DIRECT-PCD-3, the programmed-cell-death behavior is used most frequently. For indirect benefits, a KIL 3 or lower leads to the most use of the behavior.

At intermediate kin inclusivity levels, the effect of kin selection on programmed cell death differs depending on whether benefits are direct or indirect (as shown in Figure 2.4). DIRECT-PCD-3 was executed significantly more (10.32% of the unicells) than
DIRECT-PCD-0 (7.58% of unicells, Mann-Whitney U test $p-value < 2.2e-16$). However, the number of times the INDIRECT-PCD instruction was executed was not significantly different at INDIRECT-PCD-3 (7.11% of unicells) compared to INDIRECT-PCD-0 (7.33% of unicells, Mann-Whitney U test $p-value = 0.4639$). Only at INDIRECT-PCD-10 or higher is the instruction executed significantly less frequently (4.67% unicells at INDIRECT-PCD-10, Mann-Whitney U test $p-value = 6.637e-09$ compared to INDIRECT-PCD-3). This result shows that the degree of kin discrimination differentially affects the likelihood of programmed cell death evolving via kin selection depending on whether the benefits to kin are direct or indirect.

2.6.3 What causes the different response to kin inclusivity level?

This difference between the use of programmed cell death in contexts where it provides direct or indirect benefits is due to the likelihood of the instructions having an effect on the neighboring unicells.

When the benefit is direct, there is a trade off between the effectiveness of programmed cell death events during the initial invasion and the effectiveness of the programmed cell death events to resist subsequent invasion of unicell cheaters without the PCD gene. Initially, programmed cell death is most beneficial when a sufficient number of unicells are considered kin, as shown in Figure 2.5a. At DIRECT-PCD-3, enough unicells are considered kin for a large benefit to be generated upon an individual's programmed cell death. If too few unicells are considered kin (such as at DIRECT-PCD-0), the average benefit generated is significantly less per PCD event compared to DIRECT-PCD-3 (at update 500 DIRECT-PCD-0 mean kin is 7.32 unicells, DIRECT-PCD-3 kin is 19.67 unicells, Mann-Whitney U test $p-value = 2.847e-07$). Later in evolution, if too many unicells are considered kin (such as at DIRECT-PCD-100), unicells without the programmed cell death gene can gain a benefit from programmed cell death events and invade, as shown in Figure 2.6a (at

update 2000 DIRECT-PCD-3 mean kin without PCD gene is 4.72 unicells, DIRECT-PCD-100 mean kin without PCD gene is 18.31 unicells, Mann-Whitney U test $p - value = 8.988e - 11$).

When the benefit is instead indirect, programmed cell death is only beneficial if a sufficient number of unicells are considered non-kin, as shown in Figure 2.5b. A kin inclusivity level of zero harms the maximum number of unicells with a programmed cell death (mean 14.23 unicells affected per PCD event at update 500). When the kin inclusivity level of an indirect benefit increases to INDIRECT-PCD-3, the number of unicells harmed significantly decreases (mean 0.27 unicells affected per PCD event at update 500, Mann-Whitney U test $p - value = 5.332e - 10$), thereby decreasing the benefit of the trait. This effect is even more extreme at INDIRECT-PCD-100 where no unicells are considered non-kin at update 500, making the programmed cell death worthless (yet still lethal to the cell engaging in it). As shown in Figure 2.6b INDIRECT-PCD-0 also prevents unicells with the PCD gene from invading throughout evolution.

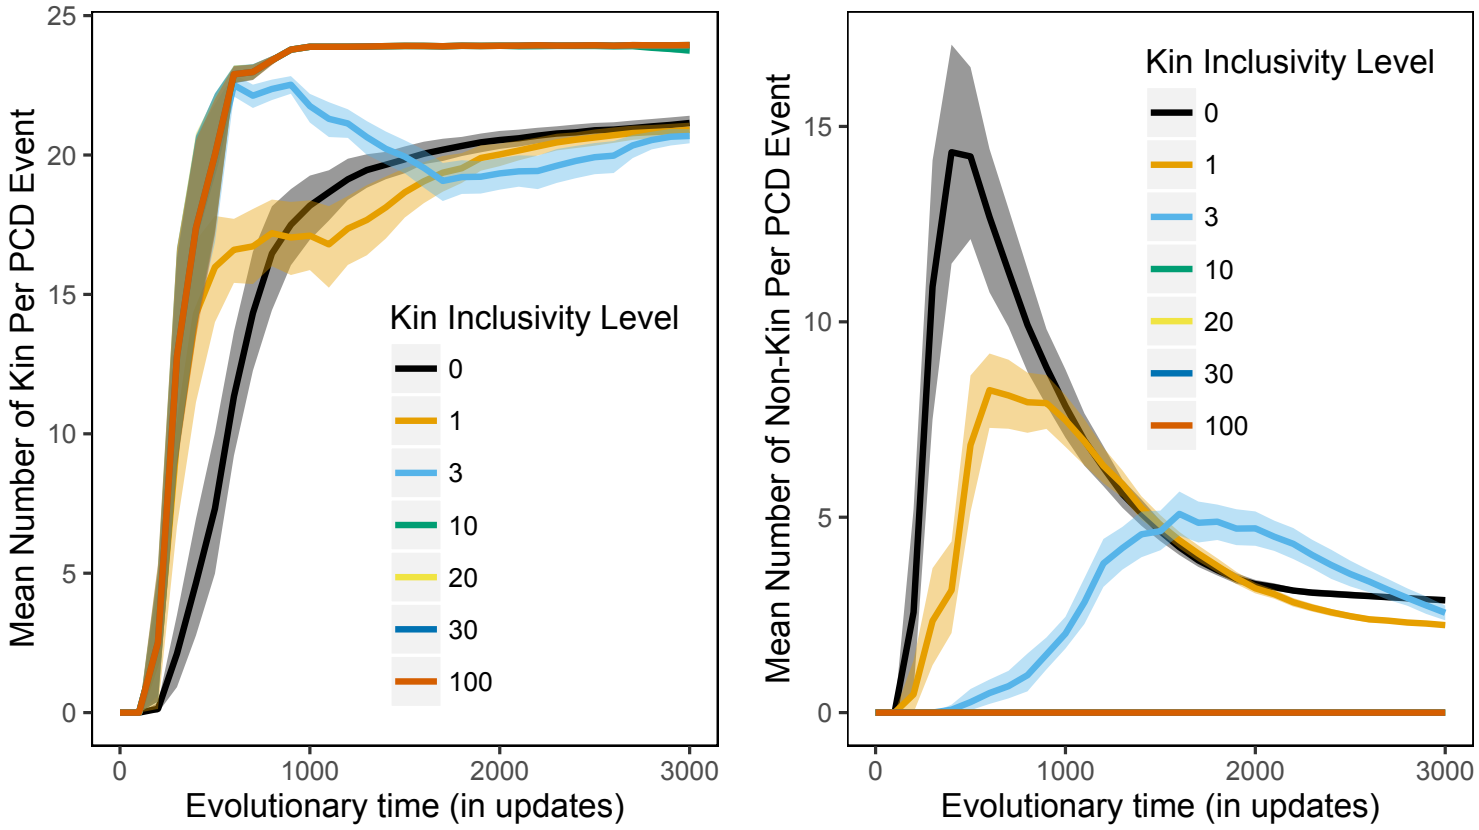

(a) Direct benefits to kin

(b) Indirect benefits to kin through harm to non-kin

Figure 2.5: **The average number of surrounding unicells directly affected per programmed cell death event across varying kin inclusivity levels when trait first emerges.** (a) When benefits are direct, DIRECT-PCD-100 have the highest amount of surrounding kin during initial evolution of the trait, leading to the most initial benefit. (b) When benefits are indirect, non-kin unicells are directly affected and at INDIRECT-PCD-0, the most surrounding unicells are considered non-kin and therefore each programmed cell death event has the largest effect.

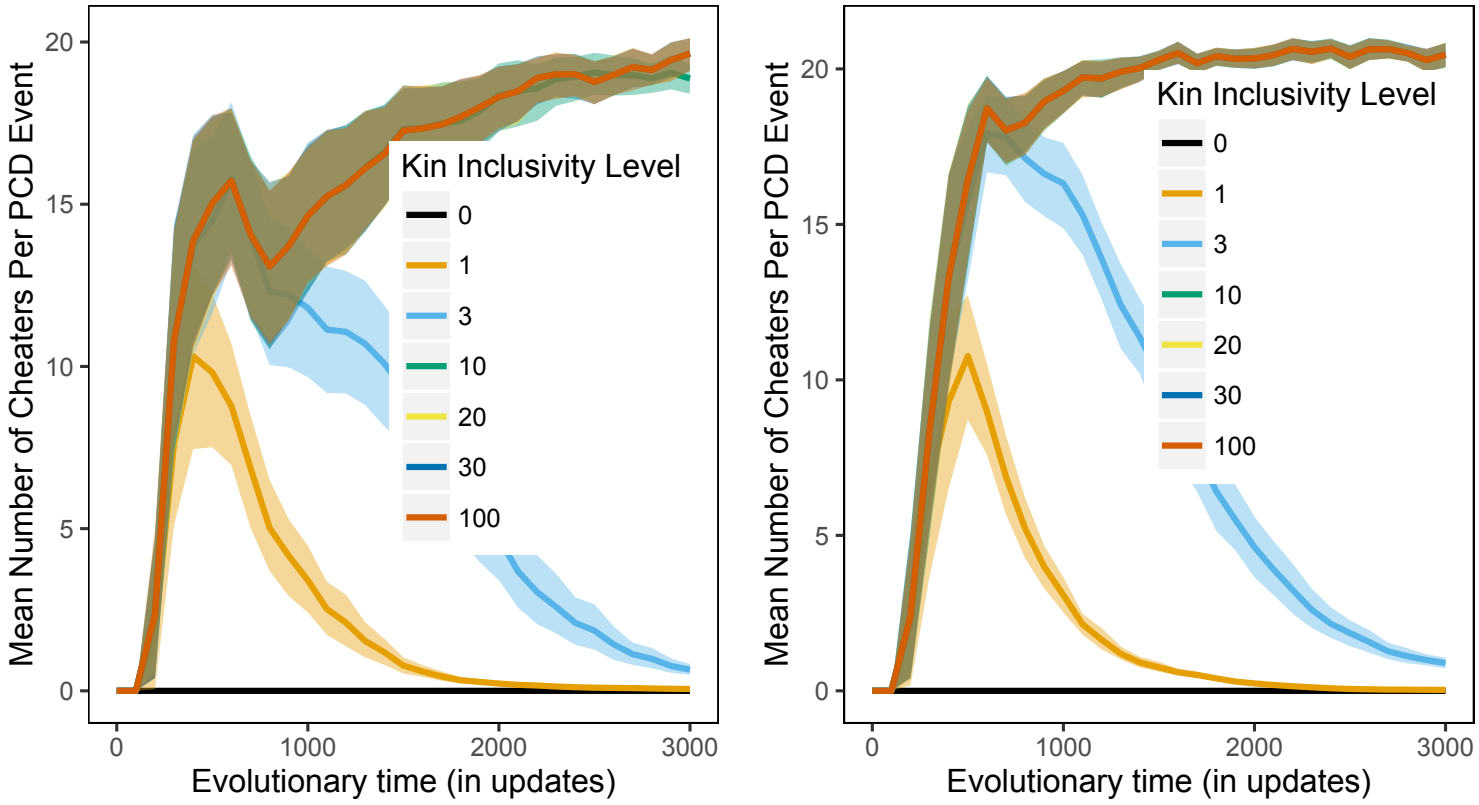

(a) Direct benefit to kin

(b) Indirect benefit through harm to non-kin

Figure 2.6: **Average number of cheating unicells, i.e. unicells that are considered kin but do not contain the PCD instruction and therefore will never engage in the behavior.** (a) While populations with DIRECT-PCD-30 to 100 have a high initial benefit (Fig 2.5a), unicells without the PCD gene are considered kin throughout the experiment, decreasing the relative benefit of the PCD trait for unicells with the PCD gene. DIRECT-PCD-3, however, enables unicells expressing the PCD trait to balance a high initial benefit with excluding unicells without the PCD gene after the initial emergence of the trait (cheaters). (b) At INDIRECT-PCD-0, unicells without the PCD trait (cheaters) are prevented from invading completely throughout evolution. Unicells have no way of detecting presence of the PCD gene, the KIL requires unicells to use overall genetic difference as a proxy for the likelihood of another unicell having the PCD gene or not.

As Dawkins discussed in *The Selfish Gene*, whether an altruistic trait is selected ultimately depends on whether it is able to increase the number of copies of itself in the population [24]. Relatedness in Hamilton's Rule is a proxy for the likelihood of an organism containing an altruistic genetic variant [43], but here we are able to directly measure the effect on unicells containing the altruistic genetic variant of interest. We have shown that the degree of kin discrimination affects kin selection differently depending on whether the programmed cell death behavior directly affects kin or non-kin. When the behavior is indirect, the lowest level of kin inclusivity is beneficial initially for a large effect and later in evolution to resist invasion. However, if the benefit is direct, there is a trade off between high initial benefit by having high kin inclusivity and the ability to resist subsequent invasion, which requires a lower kin inclusivity level. Due to the wide variety of kin discrimination mechanisms and direct or indirect benefits found in organic systems [11, 39, 33], these varying responses should be taken into consideration when analyzing programmed cell death behavior.

2.7 Conclusions

We have shown in this work that programmed cell death can evolve due to kin selection under conditions where it otherwise would not have evolved. This work provides a proof of concept that kin selection can be the driving force behind the evolution of programmed cell death in unicellular organisms. We further analyzed how the type of benefit associated with programmed cell death (direct or indirect) and the degree of kin discrimination affected evolutionary outcomes. We found that accurate kin discrimination was necessary for kin selection to evolve programmed cell death. However, programmed cell death that confers a direct benefit to kin evolved to significantly higher levels of use when many unicells are considered kin, whereas when the benefit was indirect the most restrictive kin discrimination led to the highest use of programmed cell death.

While it is difficult to fully control organic systems, several experimental systems have suggested that programmed cell death could have been under direct selection due to inclusive fitness including *E. coli* [78], *Dictyostelium discoideum* [64], and *Streptococcus mutans* [74]. Further-

more, while this work demonstrates that programmed cell death can arise as the results of benefits conferred to kin, it does not rule out other factors and conditions that may result in the evolution of programmed cell death as well. Indeed, further explorations using this system could test the alternative mechanisms that may lead to the evolution of programmed cell death.

This work has focused on unicellular organisms, but multicellular organisms also engage in whole-organism programmed death, such as bee workers that remove themselves from the hive when their health is compromised or termites that produce an explosive substance on their backs to attack predators [83]. Given the importance of these eusocial insects to our ecosystem and the current concern for the survival of many bee species [31], expanding this model to consider these more complex organism interactions could prove useful to understanding these behaviors in eusocial insects.

The presence or *de novo* evolution of programmed cell death in unicellular organisms is an exciting possible mechanism for improving human health, either by triggering programmed cell death in pathogenic bacteria or to reduce viral load in beneficial bacteria. However, to harness that power, we must understand what selective forces are acting on the behavior now and in the past. This work contributes to our understanding of how such behavior could have evolved and provides a system that can be used to understand how it will continue to change.

CHAPTER 3

SIGNALS IN THE DARK: WHAT FACTORS SELECT FOR THE EVOLUTION OF COOPERATION CONTROLLED BY QUORUM SENSING?

Quorum sensing is a primary mechanism for coordination in many species of organisms, including various bacteria [90, 18, 90]. It occurs when individual cells produce a small amount of a signal molecule – called an *autoinducer* – and detect the amount of that molecule in their surroundings. By sensing the concentration of autoinducer, cells can get a rough estimate of how many related organisms are likely to be around them. If the size of the colony is large enough (i.e. it has reached a *quorum*), the organisms will start performing an action that is only beneficial when there are a large number of cooperating cells. In this way, a colony is able to cooperate to do things like bioluminesce [17], digest resources too large for individuals [29], or wait to become virulent until they can overwhelm an immune system [8].

Like most forms of cooperation, quorum sensing is susceptible to cheating. Cheating organisms do not contribute to the costly cooperative behavior, but they gain from the rest of the cells' cooperation [93]. For example, cooperators may produce an enzyme that cheaters can gain from by taking in freed resources without paying the cost of producing the enzyme themselves. Some especially insidious cheaters may even produce the autoinducer signal that makes other organisms believe that a quorum has been reached and that they should start the costly behavior [18].

Vibrio harveyi is an example of a marine bacterium that uses quorum sensing to control the production of a public good [18]. The public good is an extracellular protease that can break down proteins outside of the cell, but must be excreted in large amounts to be useful. The restricted nutrients produced when the proteins are broken down are then generally available in the environment for the surrounding bacteria to ingest. There is a known defector strain that produces negligible amounts of autoinducer and protease, but still intakes the nutrients freed by surrounding cooperators.

Bruger *et al.* has found in their experiments that their wildtype *V. harveyi* is resistant to de-

fectors while an unconditional cooperator is not, indicating that the quorum-sensing control gives the cooperators the ability to repel invading defectors [19]. Therefore, the question becomes: what factors select for the origin and maintenance of cooperation controlled by quorum sensing, thereby preventing the known defector from invading the population? Further, are the factors that select for the origin of the trait the same as those that select for the maintenance of the trait?

We hypothesized that population structure, the availability of unrestricted resources, and a positive feedback loop in the quorum sensing mechanism are all possible factors to select for cooperation in a quorum-sensing system. We tested the effects of these factors in an agent-based computer system called Empirical and found that increased population structure and decreased unrestricted resources both stabilize pre-existing cooperation and select for cooperation to evolve *de novo*.

3.1 Literature Reviewed

Quorum sensing is found in many bacterial systems, a variety of which are reviewed in [90]. Waters et al. further explain how quorum sensing is a behavior that can enable unicellular organisms to act similarly to a multicellular organism by allowing coordination within a large area. Quorum sensing could even allow for communication between bacteria and multicellular organisms because the molecules that are used have substantial overlap across species.

As with altruism, quorum sensing, and specifically the group-generated benefits often produced by quorum sensing, can be vulnerable to cheaters [89]. Because quorum sensing controls other cooperative behaviors, cheaters could theoretically cheat by failing to produce the autoinducer-sensing molecule at all, or by participating in the quorum sensing, but not producing the costly public good when quorum is reached.

Diggle et al. showed experimentally that both types of cheaters can invade an established population of cooperators when there was no spatial structure [27]. However, when spatial structure is introduced, such as in a biofilm, cooperators are more likely to be near their offspring and therefore kin selection can stabilize the cooperation. Even when the population is not completely structured,

if there are sufficient subpopulations, cooperation may be able to persist [73]. Cheating can also be frequency-dependent, meaning that cheaters have a higher fitness when rare but if instead cooperators are rare, the cooperators have higher fitness [79]. This dynamic can occur when the population is structured such that groups of cooperators can grow faster than groups of cheaters.

Another mechanism found to decrease the likelihood of cheaters invading a population of cooperators is bottlenecking, discussed in [14]. A bottleneck is when the majority of the population (usually 99%) is killed and the survivors are allowed to grow to fill the population again. If cheaters grow slower than cooperators, they will represent less of the population and therefore be less likely to make it through the bottleneck. Typically the bottleneck leaves the surviving population spread out over the available space, so even if cheaters are common, they are unlikely to be near enough to a cooperator to gain a fitness advantage [14]. If these bottleneck events occur at the right frequency, they may be able to stabilize a cooperative trait.

Many models of cooperation, including some discussed above, use only two phenotypes: cooperator and cheater. However that may not be biologically realistic and Killingback et al. found that when intermediate cooperative phenotypes were considered, cooperation was able to stabilize as long as there was also spatial structure [58].

In many of the model systems used, researchers are still trying to understand the actual costs and benefits to quorum-sensing-based cooperation. Heilmann et al. provides a theoretical framework to test the many quorum-sensing-controlled public-goods strategies already known about in natural systems [47]. Because the optimal public goods strategy depends on the shape of the benefits curve, researchers can use this framework to determine what strategy various bacteria are likely to use and then manipulate the benefits they are receiving to our own gain. Using two mathematical models, Heilmann et al. explored how the benefit of the public good changes with the population size and found that the best strategy for production would be either to continuously increase production as population size increases ("continuous") or to not produce anything until a tipping point in the population size ("discontinuous"). Which of these categories holds the optimal strategy depends on how beneficial the public good is at small concentrations. If the benefit of the

public good is sigmoidal so that it initially decelerates as the concentration increases, then the cost of producing the public good at small concentrations (the cost is a linear function) outweighs the benefit. It is only when the concentration of public good hits a critical amount that there is a net benefit, leading to a discontinuous strategy of production being best. If, instead, the benefit of the public good accelerates initially, a continuous production strategy is best [47].

Finally, one of the more complex mathematical models produced for quorum sensing, collective action, is presented in [17]. Using this model, Brown et al. predicted that as the number of strains of quorum-sensing bacteria increased, the strength of signaling would also increase. This dynamic would be the result of individuals increasing their signal to try to trigger competitors to perform the costly action. This additional noise would then decrease the value of the signal, which could explain why there is such diversity in the types of signals and cooperative responses in bacteria [17]. They also pointed out that a costly signal could prevent this devaluation.

3.2 Methods

To investigate the factors that select for and stabilize quorum-sensing-controlled public goods, we created an agent-based system in the artificial life platform Empirical. While more complex artificial life systems, like Avida [71], have been used previously for similar work, the questions we are addressing are more focused and will benefit from initial investigations with a targeted simulation. Therefore, we implemented the following key features:

1. a system with the fundamental components necessary for evolution (mutation, inheritance, and differential fitness)

2. organisms that can sense their neighbors, identify quorum, and conditionally produce a restricted good

3. a population with varying spatial structures

4. rewards that can be set independently for unrestricted and restricted resources

5. an optional positive feedback mechanism that enables cooperators to produce more autoinducer once quorum is reached.

Specifically, our system creates a toroidal world of organisms that are able to engage in quorum sensing and public goods production that frees an otherwise restricted resource. Each organism has a simple genome consisting of one value: the probability of attempting to cooperate. The quorum threshold – the proportion of neighbors producing an autoinducer necessary to produce a public good – is user-configurable and set to 60% by default. This default value was determined empirically, by testing a range of possible values, as seen in Figure 3.1, determining that 60% would be a reasonable default setting.

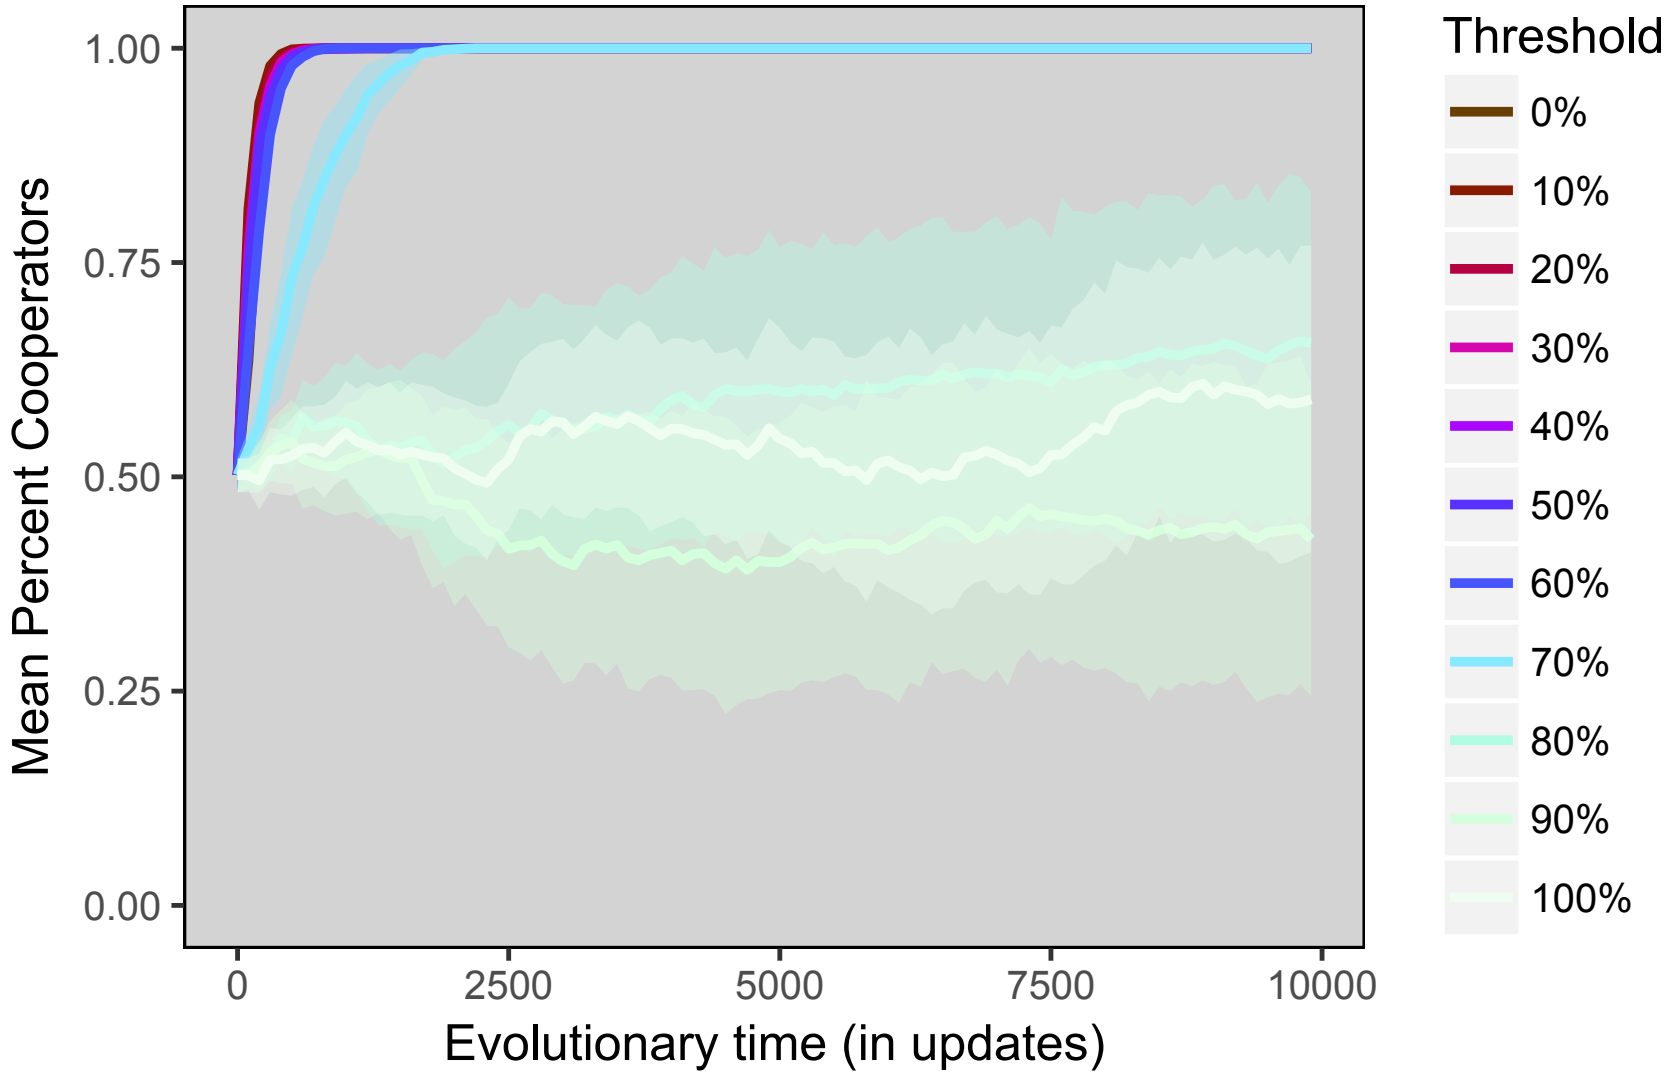

Figure 3.1: **Relative fitness of wildtype and defector across quorum threshold values.**

Whether an organism will produce the autoinducer is determined at birth by its probability of attempting to cooperate. As such, we make cooperation and autoinducer production pleiotropically linked and we remove the possibility of a defector that 'lies' by producing autoinducer. At every 'update' each organism is allowed to try to cooperate by testing its probability of cooperation. If an organism tries to cooperate, its neighbors within a specified radius (10 by default) are checked for autoinducer production. If enough neighbors are producing an autoinducer compared to the set quorum threshold, the organism pays the cost to produce the public good (by default 9 units of resource) and all organisms within a one-cell radius receive a portion of the restricted resource freed by the public good, including the producer, all cooperators, and any cheaters fortunate enough to be part of the group. By default there are 45 resource units split among up to nine organisms round-robin starting with the producer.

After the public goods benefits are determined, each organism is tested for reproduction. If an organism has reached 50 or more units of resource, it is allowed to reproduce. Both the parent and the offspring are mutated (essentially creating two offspring and replacing the parent). If mutations are enabled, the probability to cooperate value has a user-defined chance of mutating, 0.001 by default, and if it mutates, a new value is pulled from a uniform distribution between 0 and 1.

We started all experiments with an initial population of 1% of the spaces filled, randomly distributed across the world. An initial starting amount of unrestricted resource (i.e. a resource that does not need to be freed via public goods production) of 80,000 is distributed evenly throughout the world. This unrestricted resource is necessary to allow the wildtype to increase to sufficient density to trigger public goods production.

Most experiments included regular 'bottleneck' events, simulating a daily transfer in a wet lab. The population is randomly reduced back to 1% of the world and 80,000 more units of unrestricted resource are distributed evenly. These bottleneck events occur every 100 updates, which is approximately how long it takes for the population to fill the world in monocultures, in agreement with standard wetlab protocol.

Finally, we created three hand-coded organism types that can be injected at varying proportions into the starting population. The 'wildtype' simulates quorum-sensing-controlled public goods production by having a cooperation probability of 1.0 and a quorum threshold of 0.6 (i.e. 60% of neighbors producing autoinducer within 10 spaces). The 'unconditional' type simulates an unconditional or constitutive cooperator by having a cooperation probability of 1.0 and a quorum threshold of 0.0 (i.e. no neighbors need to produce autoinducer). Finally, the 'defector' type simulates a defecting organism by a cooperation probability of 0.0 and it would have the set default quorum threshold if it were to mutate. Spaces in the world can be empty, but only one organism can occupy a space at a time.

3.3 Results and Discussion

We determined how factors such as population structure, units of unrestricted resources, and a positive feedback mechanism affected the stability of pre-established populations of cooperators when challenged with defectors and the *de novo* evolution of the cooperation trait from a non-cooperating phenotype.

3.3.1 Verification of Simulation Accuracy

Before starting experiments with our simulation, we calibrated the parameters with the goal of qualitatively matching the *Vibrio harveyi* system that we aimed to simulate. In Figure 3.2 the *V. harveyi* growth of each phenotype in monoculture is shown on the right and the simulation is shown on the left.

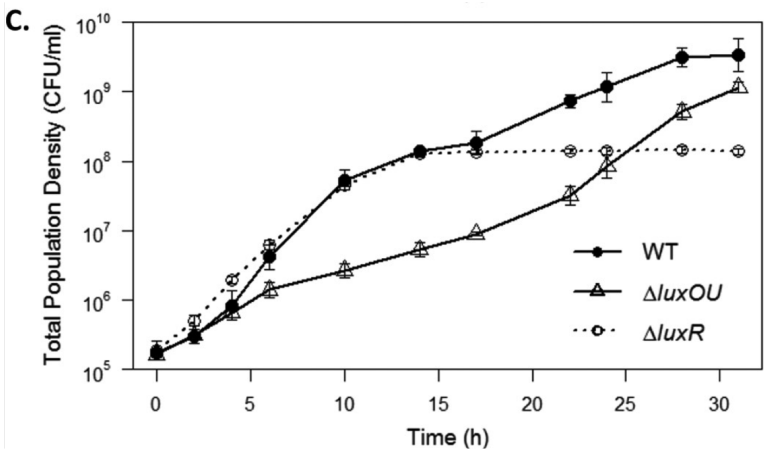

(a) Monocultures in *Vibrio harveyi*

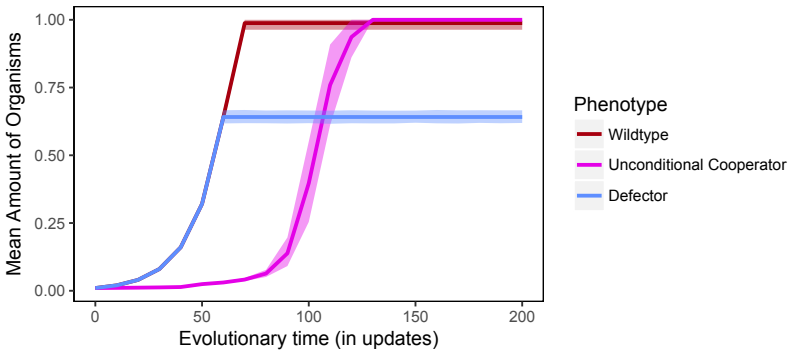

(b) Monocultures in simulation

Figure 3.2: **Monocultures of wildtype, unconditional cooperator, and defector in the *Vibrio harveyi* system (left) and our simulation (right).** The qualitative similarity between the two systems verifies that our simulation is correctly calibrated.

3.3.2 Does population structure increase selection for quorum-sensing-controlled public goods?

We first hypothesized that population structure plays a large role in the initial evolution and continued stability of public goods cooperation controlled by quorum sensing. The fitness gained from the public good is determined by how many surrounding organisms are producing the public good, therefore the more cooperators are able to isolate themselves from defectors, the more successful the cooperators will be, as we found in previous work [53] and as found in other agent-based cooperation models [2].

As predicted, population structure does enable the unconditional cooperator to outcompete the defector, whereas in the well-mixed environment, the unconditional cooperator quickly loses. Further, the wildtype is also able to outcompete the defector only when spatial structure is applied. Notably, the wildtype is never able to grow to a density to activate public goods production in a well-mixed environment (and therefore each replicate drifts to one dominate phenotype, averaged in the figure).

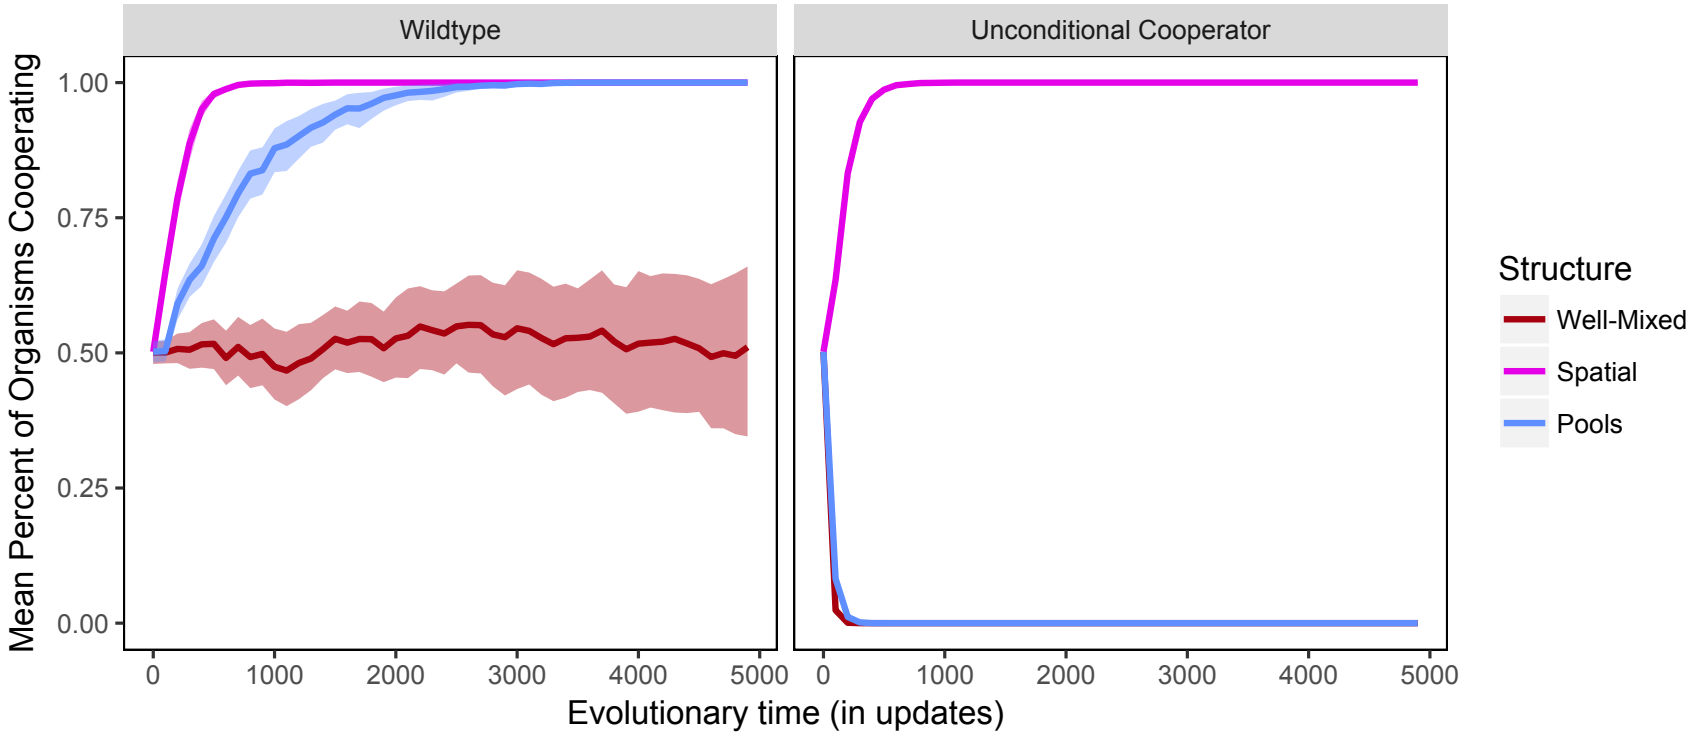

Figure 3.3: **Competition assays between defectors and two types of cooperator across population structures.** When the population is well-mixed, neither cooperator is able to consistently outcompete the defector. When a grid spatial structure is introduced, both are able to completely invade the defector. When the spatial structure is partway between those extremes, with well-mixed but isolated subpopulations, the wildtype cooperator is able to invade the defector, demonstrating Simpson's Paradox.

The final population structure we implemented is a set of subpopulations in the form of pools similar to a 96-well plate setup. Each subpopulation was a 10 x 10 well-mixed pool, and organisms from one subpopulation were only able to move to another subpopulation during bottleneck events when the surviving 1% of organisms were randomly distributed. It is in this structure that the effect of quorum-sensing can be seen most strongly. The unconditional cooperators are quickly driven to extinction. However, the wildtype is able to eventually outcompete the defector, despite these pools having well-mixed internal structures. The wildtype finding is the same result seen in 96-well plates with *V. harveyi* [19] and is a demonstration of Simpson's Paradox [73], where a dynamic is observed in subpopulations, but that dynamic disappears when those populations are combined. Bruger *et al.* found, however, that the unconditional cooperator was also able to to outcompete the defector in 96-well plates [19]. This difference is likely due to the timing of the bottleneck event and will be explored in future work.

After verifying that we were seeing the same ecological results as the wetlab experiments, we proceeded to evolutionary time scales that are more difficult to achieve in the wetlab. By enabling the cooperation probability to mutate, we were able to determine what effect population structure may have on evolving populations of *V. harveyi* in conditions such as the open ocean compared to a biofilm.

As seen in Figure 3.4, when the populations start with the cooperative trait at 0% probability, i.e. a defector population, only in a spatially-structured environment does the cooperative trait evolve to significantly higher than the control value. The control variable is a number that is mutated at the same rate as the cooperative trait but is not under selection and therefore shows how the cooperative trait would evolve if it were not under selection, positive or negative. When compared to the previous ecological results, it is noteworthy that while the wildtype could invade the defectors in a pools environment when there were no mutations, the cooperative trait is not similarly adaptive when mutations are enabled. This result suggests that cooperation of this type would need to first emerge in a spatially structured environment though it could then persist in well-mixed pools. As such, this result is an example of where observing the conditions under

which a trait is maintained is not necessarily sufficient for determining the origin of that trait.

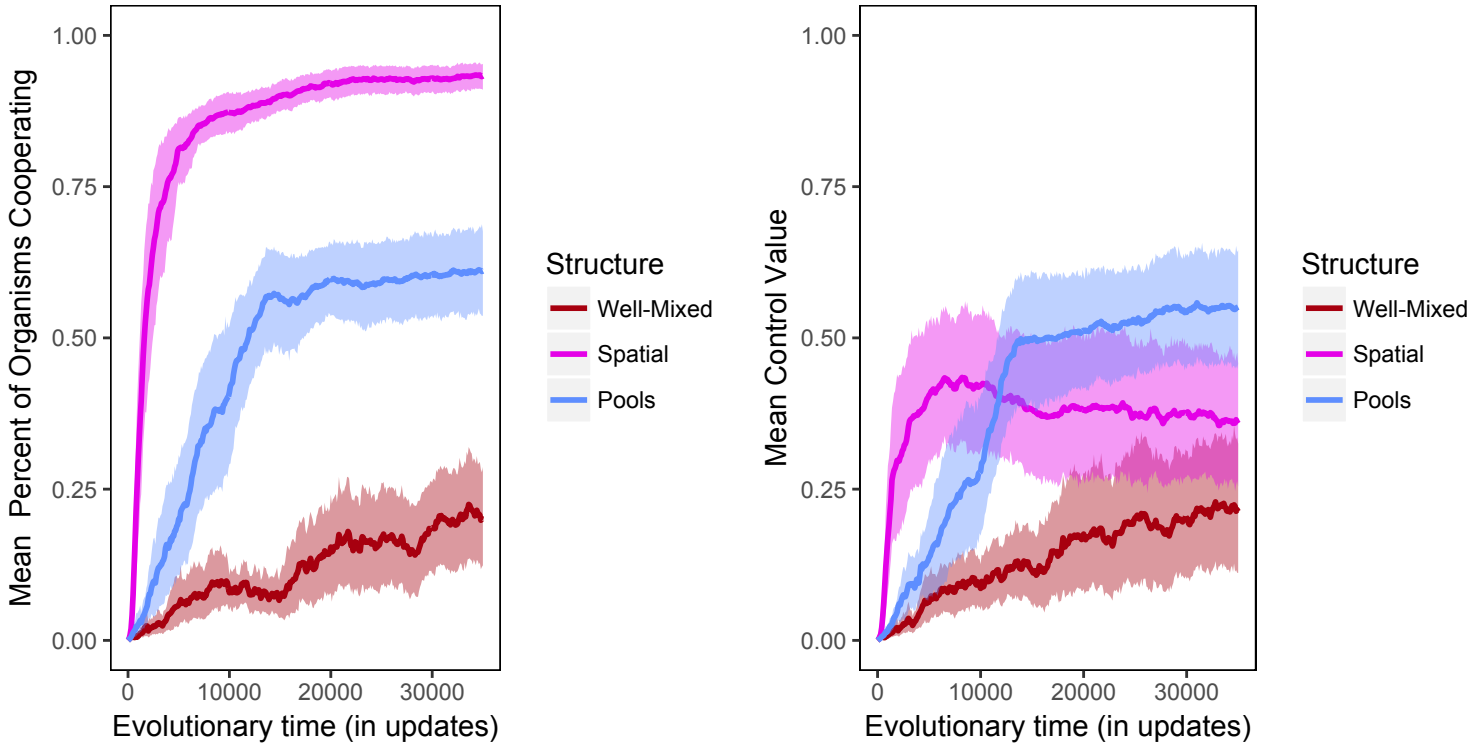

(a) Evolution of cooperation across population structures.

(b) Control Value

Figure 3.4: **When the populations start with no cooperative trait, only the spatially structured population evolves a cooperative trait value significantly different from the control value.** The variation in the control variable is due to the varying genetic drift resulting from different population structures.

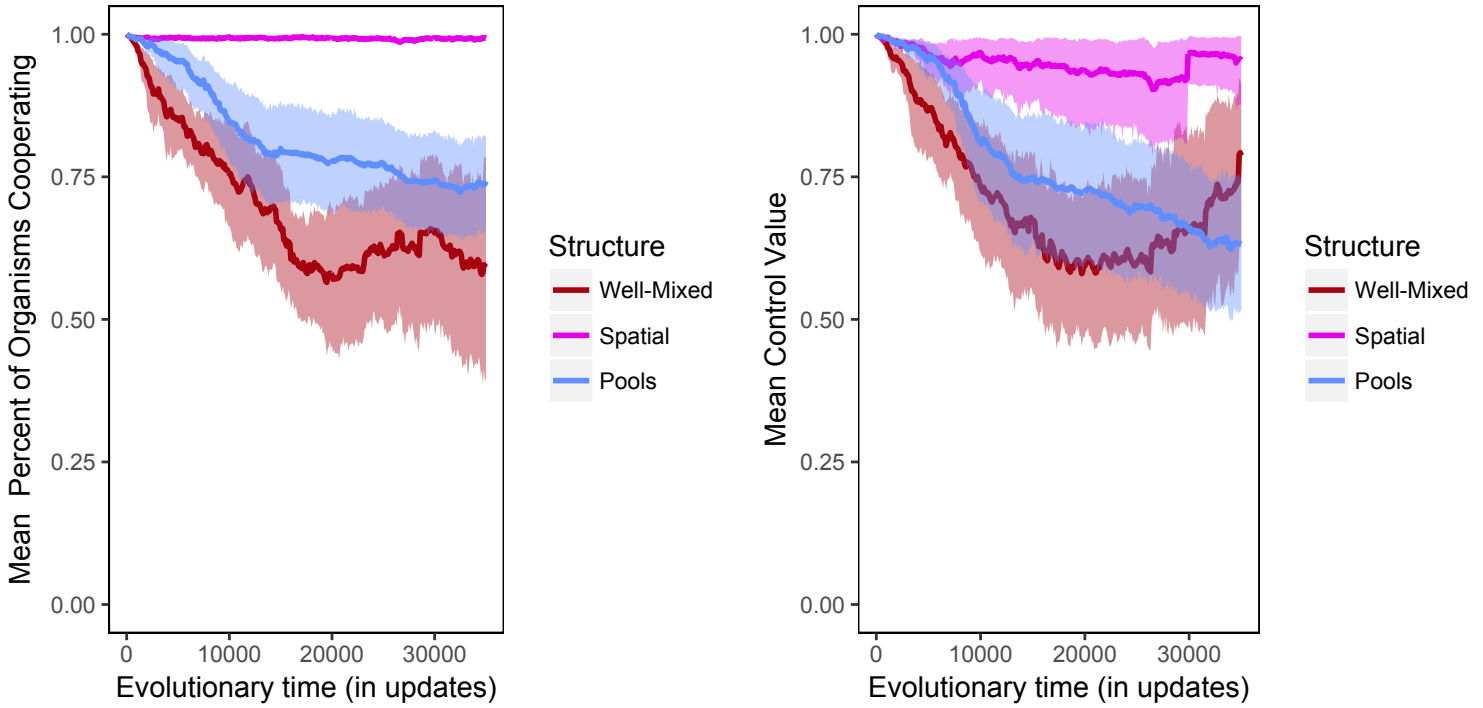

(a) Control variable.

(b) Maintaining cooperation over evolutionary time.

Figure 3.5: **Maintaining cooperation across population structures. In all population structures, the cooperative trait is under selective pressure and is significantly different than the control value.** However, in well-mixed populations, it is under negative selection and in pools and spatial, it is under positive selection.

The final stage of evolution is the long-term maintenance of a trait. Due to the ancient nature of the many natural systems with quorum-sensing-controlled public goods cooperation, it is clear that such a trait can be maintained for extended time scales. However, it is not guaranteed that the same factors that provide a short-term competitive advantage will also hold for long-term evolution. Therefore, we started populations with full wildtype cooperators (i.e. they had 100% chance to cooperate though still determined by quorum sensing) and enabled mutations to assess the stability of the cooperative trait over evolutionary time. As shown in Figure 3.5, spatial structure is still the most beneficial environment for the cooperative trait, maintaining cooperation at near 100% and significantly higher than the control value (Wilcoxon rank sum pairwise $p - value < 2.2e - 16$). Furthermore, when the environment is structured into well-mixed pools, the percent of cooperation stabilizes at around 75%, higher than the 61% when evolving from a non-cooperative population, though similar. The cooperation probability in the well-mixed pools is significantly higher than the control value, showing that the value of 75% is adaptive (Wilcoxon rank sum pairwise $p < 0.005$). Finally, the cooperative trait is not significantly different than the control in the well-mixed environment (Wilcoxon rank sum pairwise $p - value = 0.09512$). These results indicate that while a population may achieve high cooperation in a pools environment, they may not be able to persist at that high level of cooperation long term.

3.3.3 Does the availability of unrestricted resources decrease the benefit of public-goods cooperation?

Because our specific cooperative behavior is focused around restricted resources (i.e. freed by a public good), the relative distribution of unrestricted resources and restricted resources is likely to influence what strategy is most successful [40]. There is a small amount of unrestricted resource in all of our experiments because it is required for a population of purely wildtype to be able to start producing the public good and freeing the restricted resource. However, we hypothesized that increasing amounts of unrestricted resource would destabilize cooperation and allow defectors to more easily invade because defectors would be less reliant on being able to use restricted resources

freed by cooperators.

We created a new set of experimental treatments where unrestricted resources were introduced after each bottleneck event. The treatments differed in the amount of unrestricted resource provided and allowed us to explore the effects of varying levels of unrestricted resources on the evolution of public-goods cooperation. As shown in Figure 3.6, in well-mixed and spatially structured environments, the resource levels we tested – 80k, 100k, 120k, 140k, 160k, 180k and 200k – did not have a significant effect on the final proportion of cooperators. This result disputes our hypothesis that increased unrestricted resource would destabilize cooperation, indicating that spatial structure has a larger effect than the amount of unrestricted resource.

However, in the intermediately structured pools environment, increased unrestricted resources prevented cooperators from out-competing defectors. This result confirms that when there is enough unrestricted resource, defectors are able to avoid their lower yield weakness and prevent the cooperators from reaching quorum.

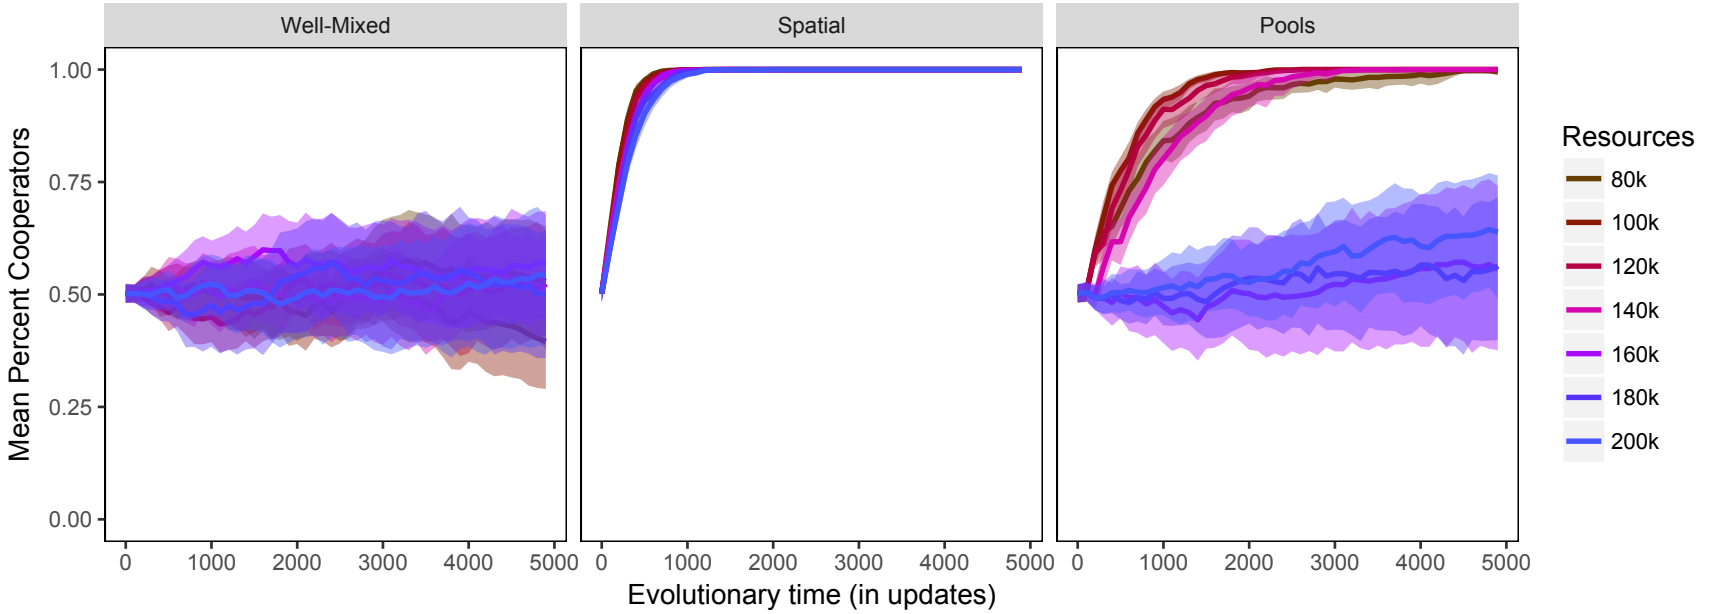

Figure 3.6: **The effect of increased unrestricted resource on ecological competition between wildtype and defectors across population structures.** Increased unrestricted resources had no significant effect in well-mixed and spatial environments, however it did reduce the competitive advantage of wildtype in pools.

For a trait to be successful in ecological competitions, a genome must first arise that possesses the trait and then increase in frequency in a population. To determine the likelihood of the cooperative trait evolving from a fully defector background across resource levels, we initiated experiments where the population started with the defector phenotype of 0% chance of cooperating and allowed for mutations upon reproduction. As shown in Figure 3.7, we again compared the value of the cooperation probability to a control variable to determine whether the trait was under selection. We found that in well-mixed environments, cooperation is not under positive selection. However, in spatially-structured environments, the cooperative trait quickly invades the population of defectors at all resource levels. This result indicates that spatial structure is sufficient to enable cooperators to overcome defectors even when the defectors do not face an unrestricted resource limitation. Finally, at resource levels 80-140k in the pools environment, cooperation is able to evolve *de novo* at the same resource levels in which cooperators were able to outcompete defectors ecologically. This result indicates that within the resource levels I tested, the factors that select for the ecological maintenance of the cooperative trait also select for the *de novo* evolution of it.

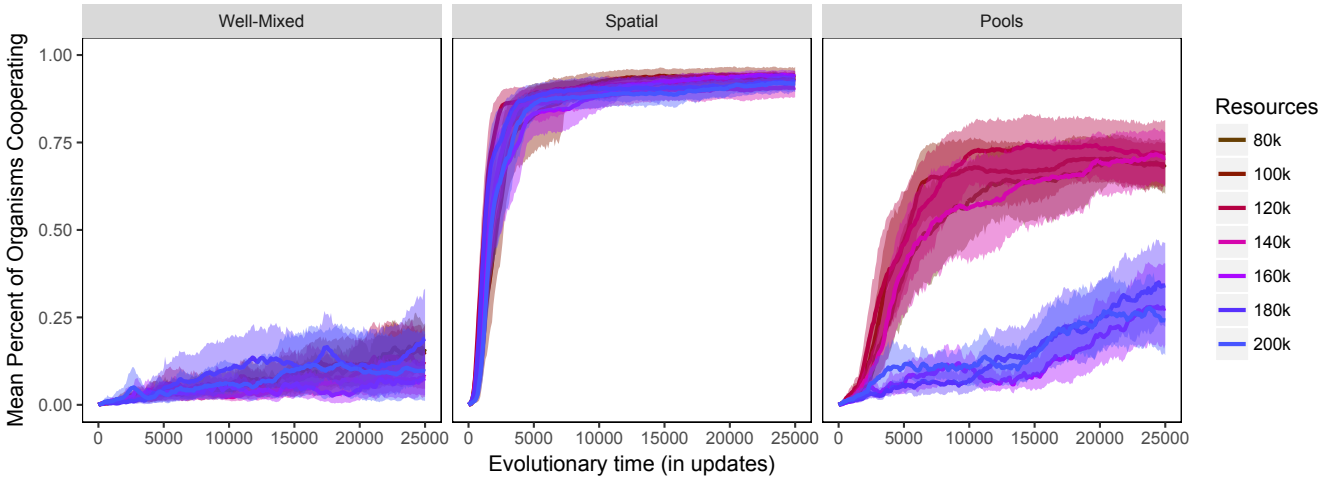

(a) Cooperation evolving across structures and resource values.

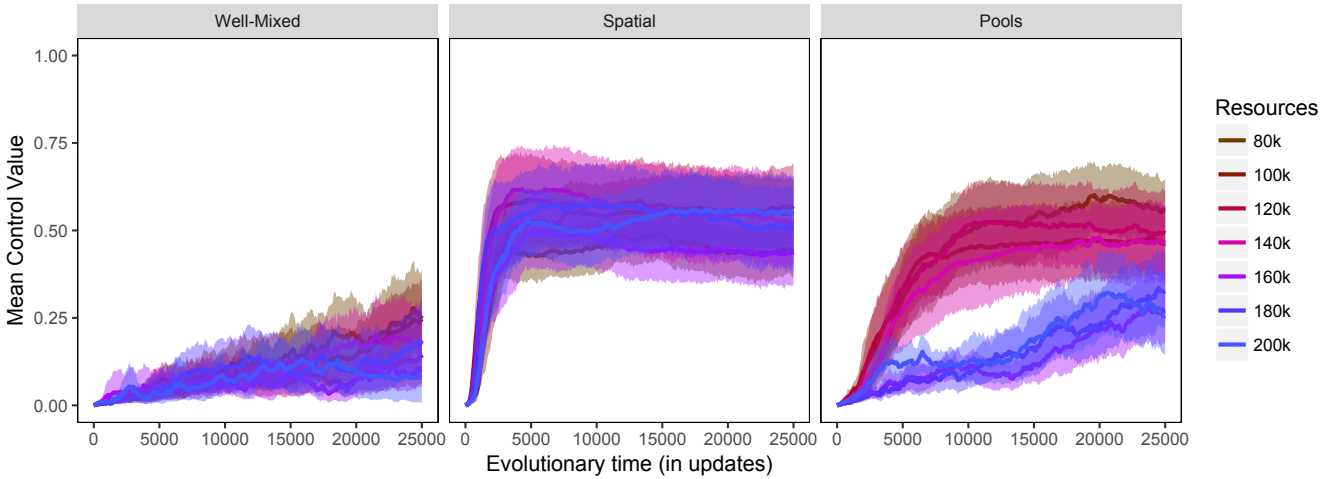

(b) Control value across structures and resource values.

Figure 3.7: **Cooperation evolving *de novo* across resource values and spatial structure.** In spatial environments, the resource levels do not affect the evolution of cooperation. In the pools environment, lower levels of resource select for cooperation whereas higher levels do not.

Finally, I again tested the long-term stability of the cooperative trait when mutations were enabled. As seen in Figure 3.8, in well-mixed and spatial environments, the resource values do not affect the stability of the cooperative trait; as seen previously, well-mixed environments select against cooperation and spatial environments select for cooperation. However as seen in the *de novo* evolution of the cooperative trait over resource values, in the pools environment, when 80,000 unrestricted resource units are introduced regularly, cooperation is maintained significantly higher than the control (Wilcoxon rank sum pairwise $p-value < 2.2e-16$), whereas at the higher resource values, they are not significantly different than the control. This result indicates that the resource values tested here generally do not affect the long-term stability of cooperation, except in the case of the pools environment.

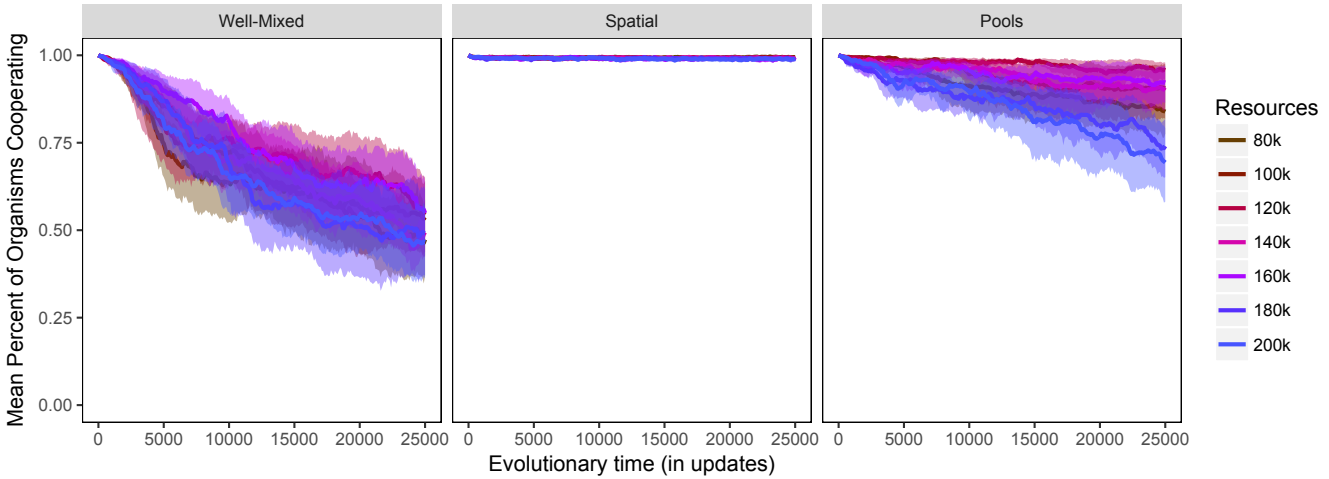

(a) Cooperation stability across structures and resource values.

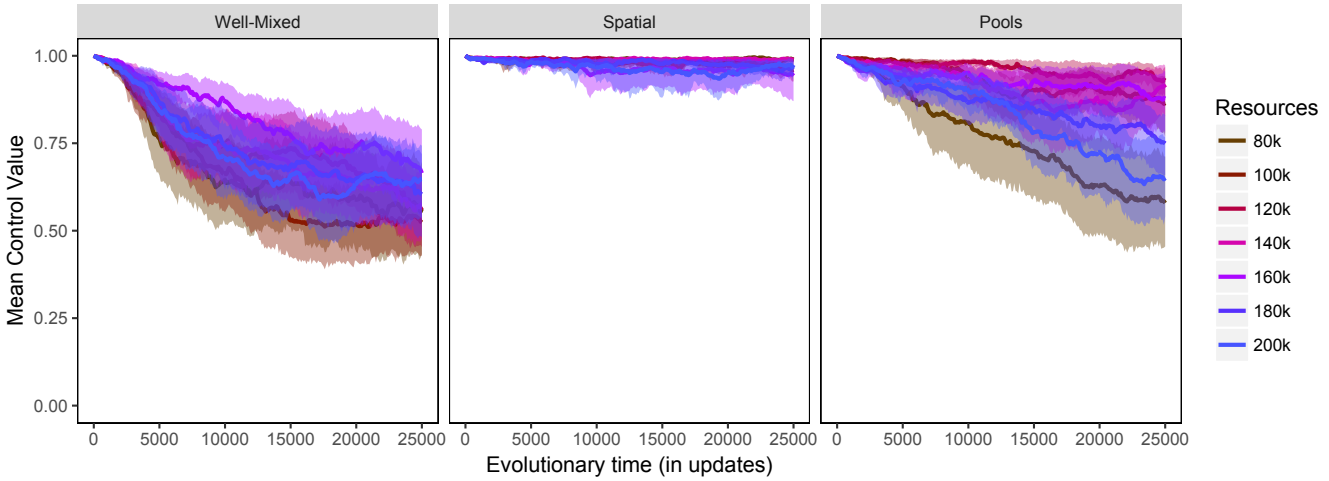

(b) Control value stability across structures and resource values.

Figure 3.8: **Resource values do not significantly affect the long-term stability of the cooperative trait, with the exception of 80k unrestricted resource units in the pools environment, where cooperation is maintained at a value significantly above the control.**

Taken together, the previous results show that the resource values did not strongly affect any stages of the evolution of this cooperative trait, with the exception of the pools environment. In that environment, lower resource values selected for cooperation in ecological competition and *de novo* evolution. However, only the lowest resource amount significantly increased the long-term stability of the cooperative trait. This result shows that ecological competitions do not necessarily reflect the long-term evolutionary trends of the trait.

3.3.4 Does a positive feedback loop on quorum sensing increase stability of cooperation?

Many natural quorum-sensing systems have a positive feedback mechanism for the production of autoinducers [18]. With a positive feedback mechanism, once an individual detects enough autoinducer for a quorum, it increases the amount of autoinducer it produces to make more of the colony switch states to produce the public good. We hypothesize that this mechanism could increase the initial evolution and stability of quorum sensing, as long as defectors are somehow policed. If defectors are not policed, they would gain a huge fitness increase by deliberately pushing surrounding cooperators into a high-production state even though there are not enough cooperators for it to be beneficial. In our system, we do not allow defectors to produce an autoinducer, simulating one of various mechanisms that could prevent 'lying' defectors [85].

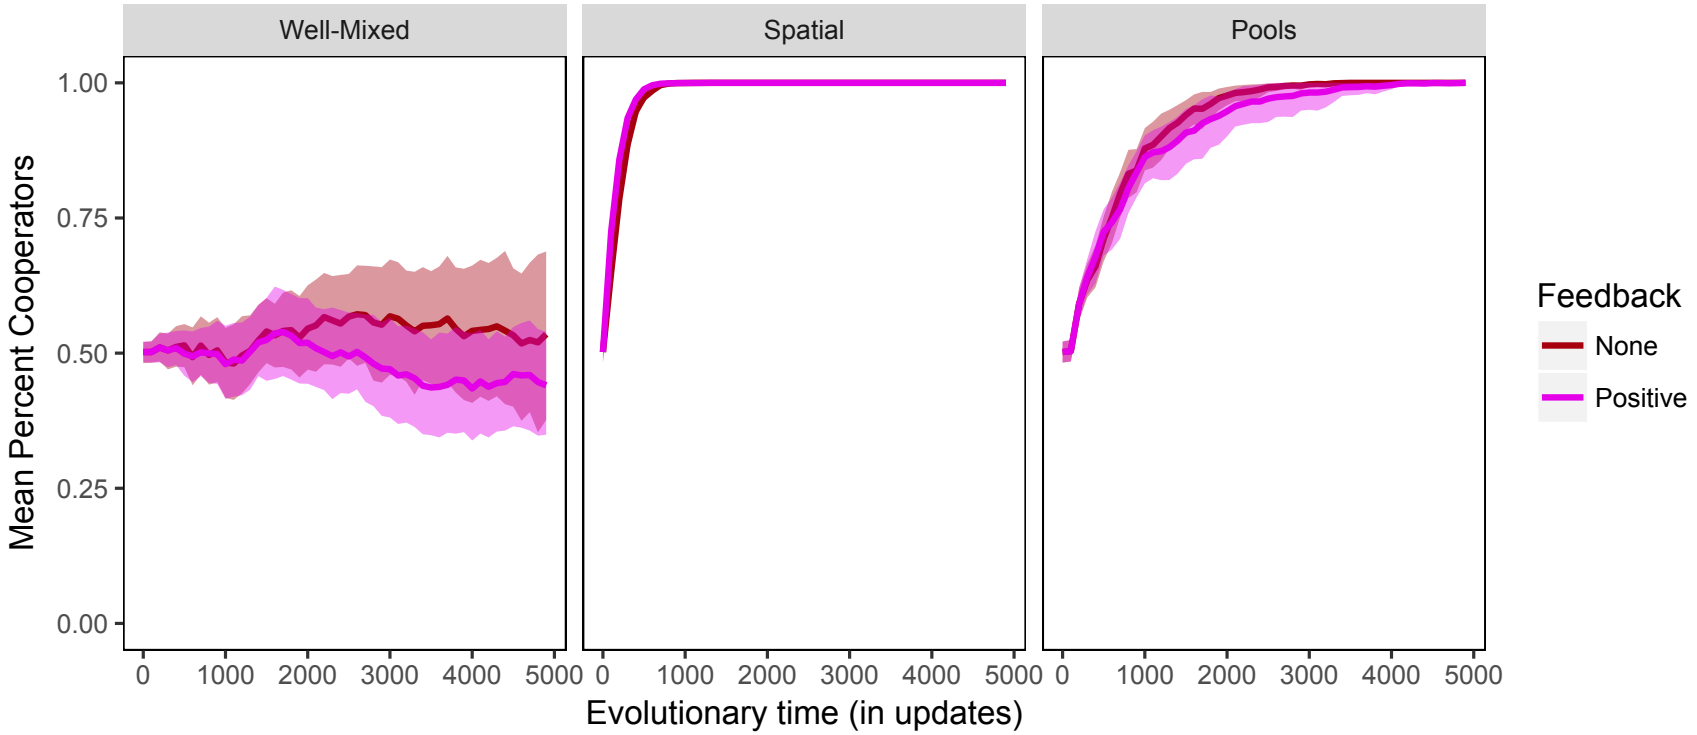

Figure 3.9: **The addition of a positive feedback mechanism does not significantly change the evolution of cooperation in spatial or pools environments.** In the well-mixed environment, the positive feedback reduces the fitness advantage of the cooperator significantly (Mann-Whitney U test $p = 1.911e-06$).

I found that in the well-mixed environment, this positive feedback loop caused cooperators to be selected against, as seen in Figure 3.9 (Mann-Whitney U test $p - value = 1.911e - 06$). In the spatial and pools environments, however, the availability of the positive feedback mechanism did not significantly affect the evolution of the cooperative trait. Because the pools consist of smaller well-mixed environments, this result indicates that the Simpson's Paradox dynamic – whereby cooperative subpopulations grow faster than defector subpopulations – has a stronger effect than a positive feedback mechanism. Further, these findings indicate that the positive feedback mechanism implemented here is neither necessary nor sufficient for cooperation to evolve.

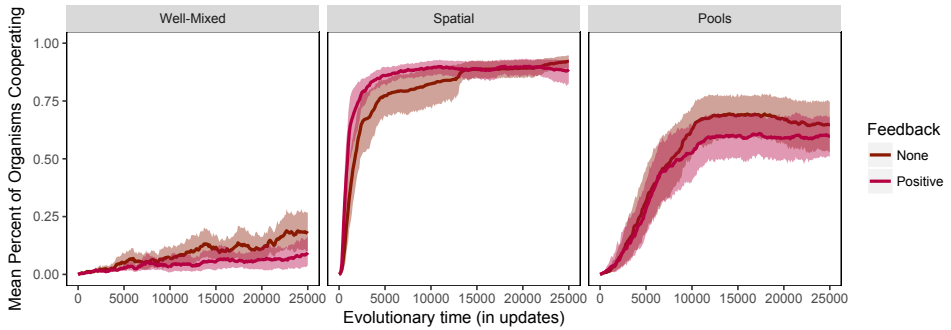

(a) Cooperation evolving across structures and feedback.

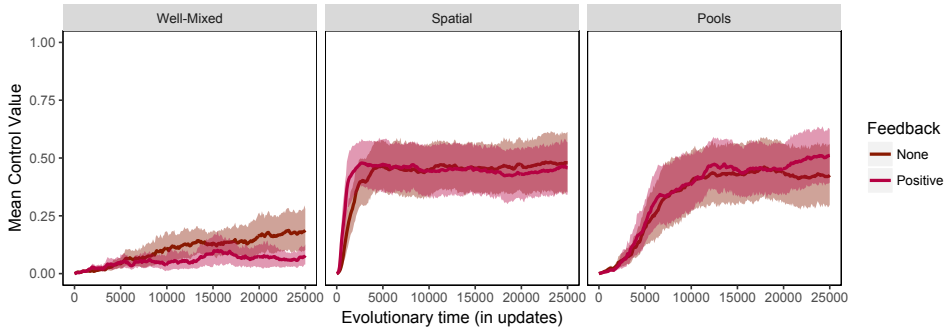

(b) Control value evolving across structures and feedback.

Figure 3.10: **Positive feedback does not significantly affect the evolution of the cooperative trait over most spatial structures. Only in the pools environment does the positive feedback mechanism significantly decrease the final amount of cooperation in the populations.**

I again tested how the *de novo* evolution of the cooperative trait may differ from the ecological competitions by enabling mutations in a population that started with the defector phenotype. As seen in 3.10, I found that generally the positive feedback mechanism did not have a significant effect, except in the pools structure. When evolving *de novo* in the pools structure, the positive feedback mechanism significantly reduced the level of cooperation evolved (Mann-Whitney U test $p - value < 2.2e - 16$). This result indicates that when evolving *de novo*, the selective pressure against the positive feedback cooperator in well-mixed structures does affect the pools structure, though cooperation is still able to evolve.

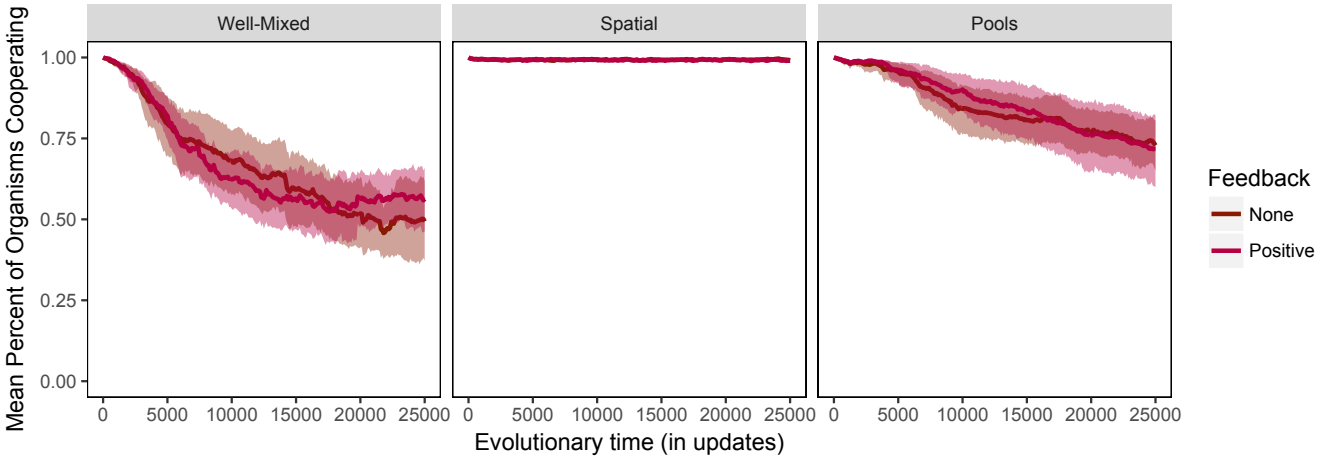

(a) Cooperation evolving across structures and feedback.

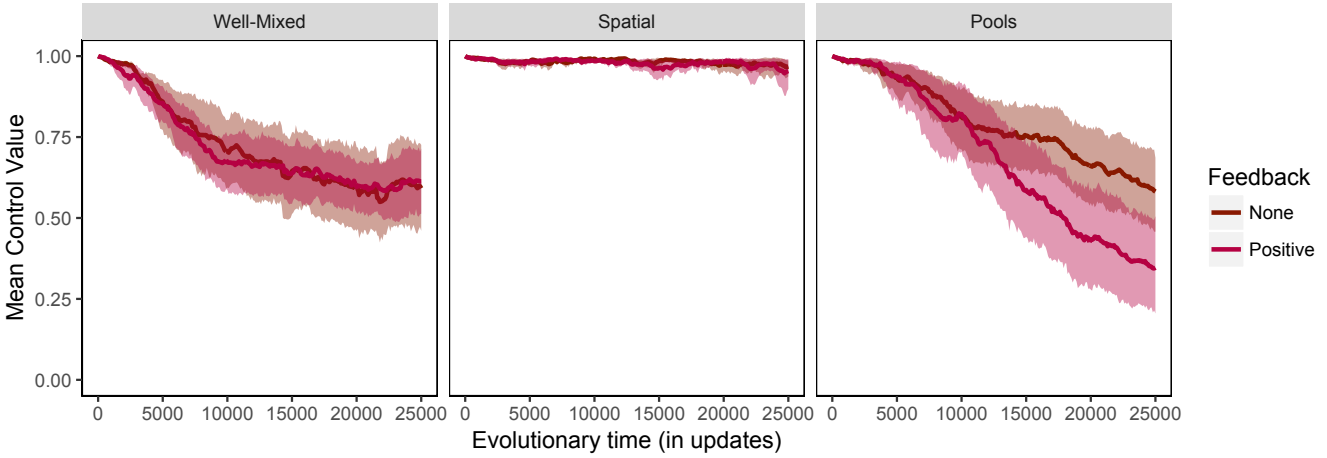

(b) Control value evolving across structures and feedback.

Figure 3.11: **Positive feedback does not significantly affect the long term maintenance of the cooperative trait across environmental structures.**

Finally I examined the long-term evolution of the cooperative trait with and without the positive feedback mechanism. As seen in Figure 3.11, I found that the positive feedback mechanism did not have a significant affect on the stability of the cooperative trait in any of the environmental structures.

These results indicate that the positive feedback mechanism I implemented does not, in general, significantly affect the evolution of the cooperative trait, with the exception of some effect in well-mixed environments and sub-populations. In ecological competitions, the positive feedback decreased cooperation in well-mixed environments, whereas in *de novo* evolution that effect was only apparent in the pools (which consists of 100 well-mixed sub-populations). Finally, neither structure showed an effect of positive feedback on the long-term stability of the cooperative trait.

3.3.5 Conclusion

In this chapter, I have presented a digital system for studying the factors contributing to the three stages of evolution of a quorum sensing trait. I found that population structure was the dominant factor, with unrestricted resource amounts and the presence of a positive feedback mechanism having only a small, if any, effect on the cooperative trait. Further, I found that whether there is an effect and the direction of selection can vary depending on whether the trait is evolving *de novo*, maintaining short term, or maintaining long term. The final cooperation values of the population are not guaranteed to stabilize at the same value regardless of the stage of evolution.

These findings show that in order to fully understand the evolutionary dynamics of a cooperative trait, we must study its origin, short-term stability, and long-term stability; we must not make assumptions that the factors selecting for the origin or long-term stability will be the same as those found to select for short-term stability.

As we begin to explore using evolving populations in medical and agricultural interventions, we must identify when the interventions rely on a cooperative trait and study how that cooperative trait may evolve over the long term. Digital systems such as the one presented here are ideal for testing such scenarios and making more informed decisions about matters relating to human health

and safety.

CHAPTER 4

SYMBIOTIC RELATIONSHIPS: WHAT FACTORS SELECT FOR THE EVOLUTION OF MUTUALISM?

Researchers have been fascinated by the existence of mutualisms – cooperation between members of different species – since Darwin [23], and therefore there are many well-documented example systems, such as pollinators with flowers, cleaner fish with other fish, and soil bacteria with plants [20, 6, 25]. Mutualisms between bacterial microbiomes and host species have been found to be especially common, including in many parts of the human body [7], making them relevant to human health. Further, more intricate mutualisms continue to be discovered in complex communities at an astounding rate [91, 66, 80, 56, 81]. Because these systems range across abundant sizes and species, it has been difficult to identify general principles of the evolution of mutualism. Further, most of these mutualisms are ancient, and therefore difficult to recreate or determine the factors that led to the original evolution of the behavior.

Like all cooperative systems, mutualisms are at risk of one partner cheating the other and destroying the mutually-beneficial arrangement [55]. This conflict raises the question of how mutualisms are able to arise in evolving populations and be maintained over many generations of organisms.

There are many hypotheses about what factors could influence the evolution and maintenance of mutualisms [55], several of which highlight the influence of the natural history of the partners [56], the amount of vertical and/or horizontal transmission of the symbiont [82], and the nature of the benefit conveyed by the symbiont [16]. While examining extant mutualisms can provide insight into how these factors affect their stability, it is difficult to experimentally test the influence of these factors on the *de novo* evolution and long-term maintenance of a mutualism.

The evolution of mutualisms is difficult to test in biological systems due to the long generation times of partners (even bacterial species take months to evolve [81]) and our imperfect control of the experimental system. To circumvent these challenges, I created a simulation to evolve a

mutualism under varying conditions, called Symbulation. Using this system, I was able to vary specific factors while guaranteeing conditions were otherwise identical, enabling me to thoroughly test the effects of natural history, vertical transmission, and benefits from the symbiont on the *de novo* evolution of a mutualism.

4.1 Literature Reviewed

Mutualism – i.e. cooperation between members of different species – has been a subject of interest to field researchers since Darwin [23]. Due to this interest, there are countless examples of mutualisms in the natural world, a number of which are reviewed in depth by Herre et al. [48]. Because two organisms of different species are highly unlikely to share a cooperative gene, kin selection and the idea of selfish genes do not easily explain this type of cooperation, making it all the more intriguing, and therefore leading to an abundance of hypotheses regarding their evolution and maintenance.

Historically it was a common assumption that because mutualism was better overall for both species, parasites would be pressured to evolve into commensalists, which would then have pressure to evolve to mutualists [32]. However Ewald et al. introduced the idea that, while parasitism to mutualism is a spectrum, we must remember that selection often pressures individuals of a species to cheat a cooperative trait because in a cooperative background, cheaters will have a fitness advantage [32].

Frank et al. proposed that mutualisms are still selected due to benefits to kin, just indirectly. If the original donor is likely to have kin nearby that would benefit from reciprocation from the original recipient, it could lead to the cooperation being under more consistent positive selection [38]. Van Cleve et al. [88] modeled a system with a high benefit relative to the cost of the mutualism, as well as high within-species relatedness and high between-species partner fidelity – the idea that the maintenance of evolution depends on partners remaining with each other for extended periods of time, thus aligning their fitnesses [21]. They found that with those assumptions, feedback mechanisms between partners were more important than the genetic correlations of individuals within

each of the species involved [88].

Other researchers have turned to inspiration from mate choice, drawing the parallel that a female choosing a male (or vice versa) is in many ways similar to a mutualist choosing a partner among the individuals of the other species [68]. Using game theoretic models of rational traders in human markets, assuming that partner choice is the dominant force in the system, Noe et al. found that the cost of sampling partners determined whether the host was choosy or not. If it is cheap to figure out the quality of a partner, it makes sense to try many and keep the best one. Conversely, if it expensive to try out a partner, hosts cannot switch symbionts whenever they please and so will be more tolerant of less-than-ideal partners [68]. This idea of partner-choice being a determining factor in the evolution of mutualism first arose in [21] in opposition to the concept of partner fidelity. Conversely, partner-choice is the idea discussed previously where a host has some way of sampling potential partners and picking the highest quality, thus possibly selecting for symbionts to be of higher quality [21].

Partner fidelity can also take place across generations with vertical transmission. Vertical transmission is when offspring of a host are infected by offspring of the host's symbiont [35]. This idea of vertical transmission was first discussed in terms of parasites passed from parent to offspring, which Fine et al. modeled. However the same idea can and should be applied to potentially cooperative symbionts. In Yamamura et al., researchers presented a differential equation model of a mutualism system where vertical transmission can evolve [97]. They found with their model that when the vertical transmission rate reaches a tipping point, the system will be pushed toward perfect vertical transmission. That tipping point's exact value can vary based on other factors, however. Bronstein et al. also found that ecological conditions must be taken into account when trying to understand the costs and benefits of mutualisms because they can cause a large shift along the parasitism-to-mutualism spectrum [15].

Much of the research on cooperation is focused on how cheaters – or low-quality partners when considering mutualisms – are prevented from invading a system [46]. However Heath et al. pointed out that, when natural systems have been closely studied, a large amount of natural

variation is generally present, which many models do not take into account [46]. Therefore, when determining how mutualisms evolved and are maintained in many natural systems, models must incorporate large amounts of standing variation in partners [46]. Foster et al. found that factoring in standing variation in partner quality may actually have a stabilizing effect on mutualisms [36]. When all partners are of high quality, there is no reason to be choosy. When choosiness is no longer under strong selection, low-quality individuals can start to take advantage of non-choosy partners [36]. Therefore, if a variety of low-quality partners are introduced every generation, selection for choosiness would be maintained and therefore cooperation as well.

Yamamura et al. explored how spatial structure and dispersal could affect the evolution of a mutualism in [98]. They created a mathematical model and simulation with a mutualist and non-mutualist phenotype with a high benefit-to-cost ratio and a low reproductive rate. They found that unlimited dispersal prevented the mutualist phenotype from being able to invade non-mutualists, and conversely with neighborhood reproduction only, the mutualists were able to invade non-mutualists [98]. Doebeli et al. also found that spatial structure was vital to the evolution of mutualism in their model [28]. They also required that increased investment in the mutualism led directly to increased yield and with those assumptions, mutualisms were actually relatively easy to evolve. This finding led them to suggest that, in some situations, it may not be an issue for mutualism to evolve once a symbiont is engaged with a host; it may instead be more an issue of the host's initial defenses [28].

Akcay reviews the mechanisms discussed previously and identifies selective pathways that can lead to stability in a mutualism [4]. Beyond these mechanisms, however, Akcay comes to the conclusion that most of the existing models that he found incorporated only one of the many mechanisms listed. This finding highlights a critical flaw, as Van Cleve et al. finds that when reciprocity, genetic relatedness, and synergy (the additional fitness advantage created by the mutualism) are all included in a mathematical model, they interact in non-trivial ways and cannot be assumed to be simply additive [88]. This issue is further discussed in Hoeksema et al., which reviewed the theoretical findings of the field of mutualism evolution and noted that while many of the models

work well for a specific aspect, none attempted to bring together the many factors [50]. Of course, they then point out that doing so would be difficult if not impossible in an analytical model, which is compounded by the issue that naturalists have had difficulty in determining the actual costs and benefit values in any specific system [50]. A bridge needs to be created between the theoretical and empirical researchers, and I have endeavored to begin the creation of that bridge in this work.

4.2 Methods

I created an agent-based simulation, Symbulation, that models coevolution between host organisms and symbionts, providing a simple environment to examine the conditions under which symbionts engage in strategies that are mutualistic, at one extreme, or parasitic, at the other.

4.2.1 Symbulation

In Symbulation, host and symbiont genomes each consist of a single number, the *resource behavior*. The resource behavior can be inclusively between -1 (antagonistic) and 1 (cooperative) and determines how the organism acts toward its partner, if it has one.

A resource behavior value below 0 means the organism is antagonistic toward its partner. A host with a negative resource behavior will invest that proportion of the new resources available to it into defense. A symbiont with a negative resource behavior attempts to steal resources from the host. It attacks the host based on the magnitude of its resource behavior and the amount that the symbiont's value is more negative than the host's value determines the proportion of resources the symbiont is able to steal from the host. For example, as shown in 4.1, if the host has a value of -0.3 and the symbiont a value of -0.8, the host first invests 30% of its resources into defense, meaning that those resources are no longer available for host or symbiont reproduction. If 20 resources came into the host cell, 6 have now been used up and 14 remain for reproduction. The symbiont then attacks with -0.8 and because that value is negative and of greater magnitude than the host's defense value, the symbiont is able to overcome the host defense. The symbiont is therefore able to steal 50% of the *remaining* resources. Therefore the symbiont is able to steal 7 resources and

apply them to its own reproductive progress. The host then has 7 remaining resources that it is able to use for its own reproduction.

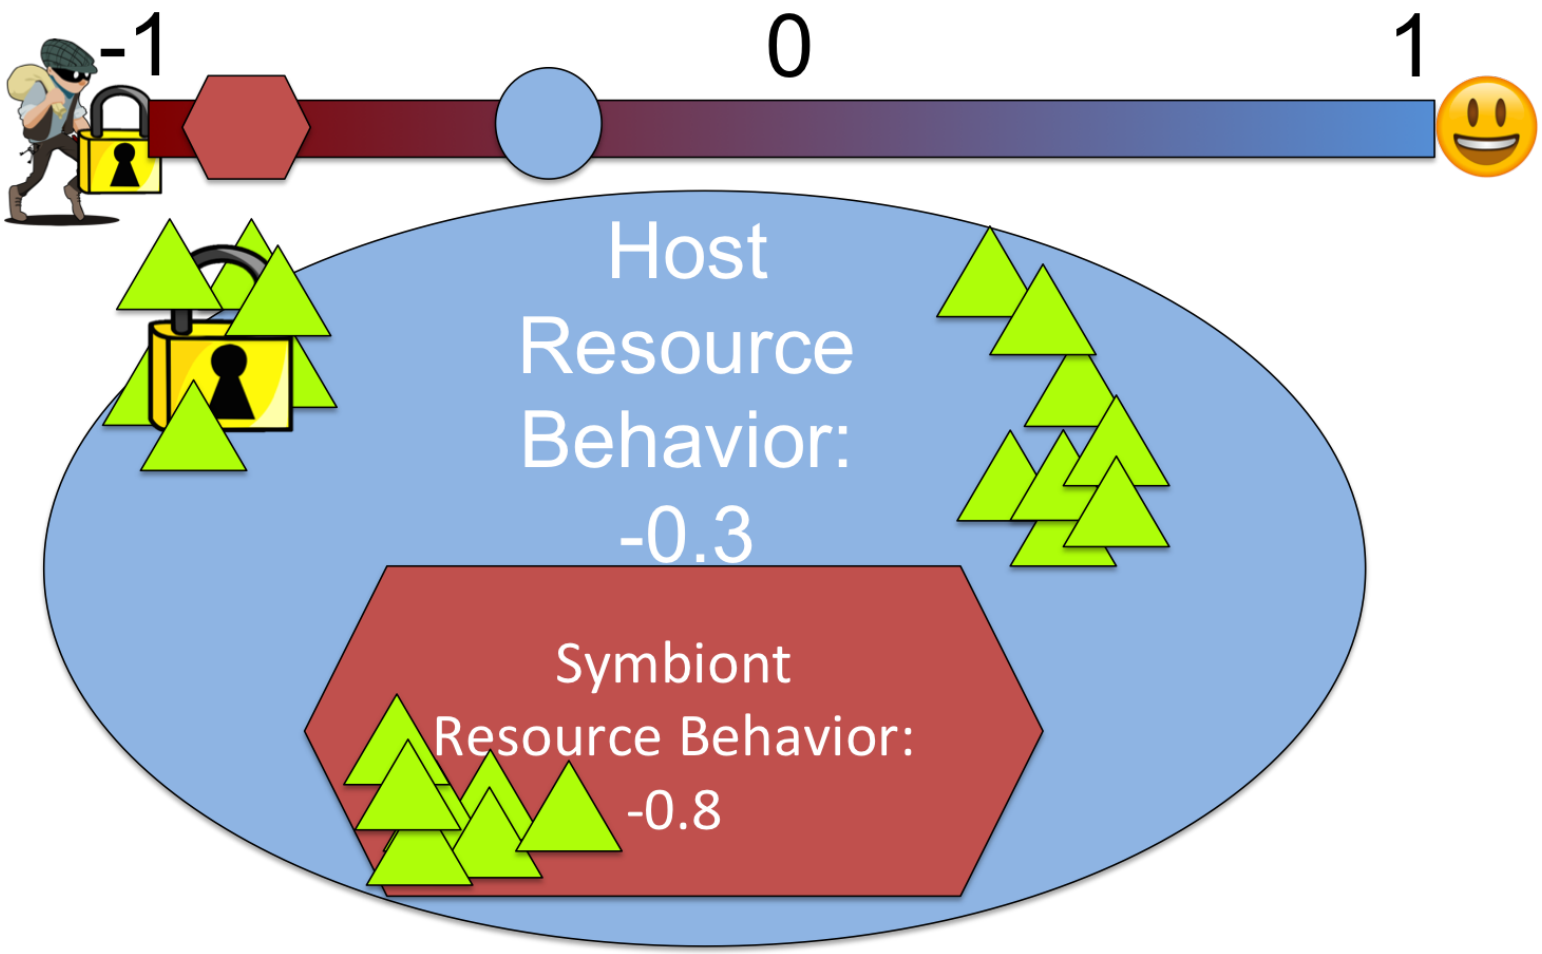

Figure 4.1: **Example host and symbiont with antagonistic behavior.** The host spends 30% of resources on defense (lock), the symbiont steals 50% of what remains, and the host retains what the symbiont does not steal.

Table 4.1: Interaction results for host and symbiont phenotypes

Host Phenotype	Symbiont Phenotype	Result
$X > 0$	$Y > 0$	Host donates proportion X to symbiont, symbiont donates proportion Y back, which is multiplied by 5
$X < 0$	$Y < 0$	Host invests proportion X in defense, symbiont steals proportion X - Y, host gets what remains
$X > 0$	$Y < 0$	Host donates proportion X to symbiont, symbionts steals additional proportion Y, host gets what remains
$X < 0$	$Y > 0$	Host invests X in defense, symbiont has no resources to donate, host gets remaining resources

Conversely, a resource behavior value above 0 means the organism is acting cooperatively toward its partner. The host donates that proportion of resources to its symbiont (if it has one), and the symbiont, in turn, donates a proportion back to the host – determined by the symbiont's resource behavior – with a synergy factor of multiplying by five applied. All combinations of positive and negative phenotypes are presented in Table 4.1.

When either a host or a symbiont reproduces, the offspring are placed in a random location in the world. A symbiont offspring must infect a host to survive and if the host targeted is already infected, the offspring fails and dies. Reproduction is imperfect, of course, and the parent and offspring both have the chance of mutation changing their resource behavior value and their chosen resources. The resource behavior is altered by a random number pulled from a Gaussian distribution with a mean of zero and a standard deviation of 0.002. The chosen resources set is altered by a random resource type being removed and a random possible resource type being added.

The symbionts have two ways of replicating themselves: horizontal and vertical transmission. In horizontal transmission, the symbiont accrues enough resources to make an offspring and send that offspring into the world to try to infect a random host. Alternatively, *vertical transmission* may occur when a symbiont's host reproduces. The *vertical transmission rate* determines the likelihood of the symbiont being copied into the offspring along with the host's genome. This transmission rate can be set by the user or evolved by the hosts.

Determining whether a mutualism has actually emerged in any system is difficult, both because it is difficult to measure the required values and because the literature is not agreed on what the measurements or values should actually be [55]. Because Symbulation allows complete control of data, I am able to avoid any indirect proxies and instead measure exactly the behavior each partner exhibits to the other, that is, the resource behavior value. However, there is still the question of how positive a value needs to be to be a mutualism and which partner needs to be positive. I am using the definition of a mutualism that the host is better off with its partner than it would be without (and conversely worse with a parasite than without). Therefore, a symbiont that has a positive resource value less than 1 (100% of resources back to host) could still be a mutualist if what it is contributing is more than the host would have without that partner, even if there are also better partners in the world. I consider a mutualism to be any relationship that appears stable over evolutionary time and both partners have positive resource behavior value. While this is not a direct test of whether the host is better off with the symbiont, selective pressure will select for hosts with a positive resource behavior value only if donating to the symbiont is of value to the host. Similarly, because defense is costly to the host, a host will only have a negative value if the symbiont is causing it more harm than not having a partner; parasitism will therefore be a stable relationship where both partners have negative resource behavior values.

4.3 One Infinite Resource

4.3.1 Methods

Each section of this chapter will have a methods subsection describing what additional features I made to examine the questions posed in that section. One infinite resource is the default environment and therefore has nothing unique to it. All future types of resource will have these settings unless otherwise specified in the respective methods subsection.

4.3.2 How does vertical transmission rate affect the maintenance of mutualism?

Vertical transmission is imperfect in natural systems, and the probability of vertical transmission succeeding is likely to affect the investment a symbiont places into a mutualism [55]. The more a symbiont can rely on spreading its offspring when its host reproduces, the more beneficial it would be to increase its host's fitness. However, if it is unlikely a symbiont will be able to spread when its host reproduces, it would make more sense for the symbiont to spread via horizontal transmission [82]. Therefore, I first focused on determining how the probability of vertical transmission would influence the evolution of a mutualism. I tested values of 0, 10, 20, 30, 40, 50, 60, 70, 80, 90, and 100% vertical transmission. These set probabilities are in addition to the other requirements for vertical and horizontal transmission discussed previously. An artificial synergy factor of five was applied.

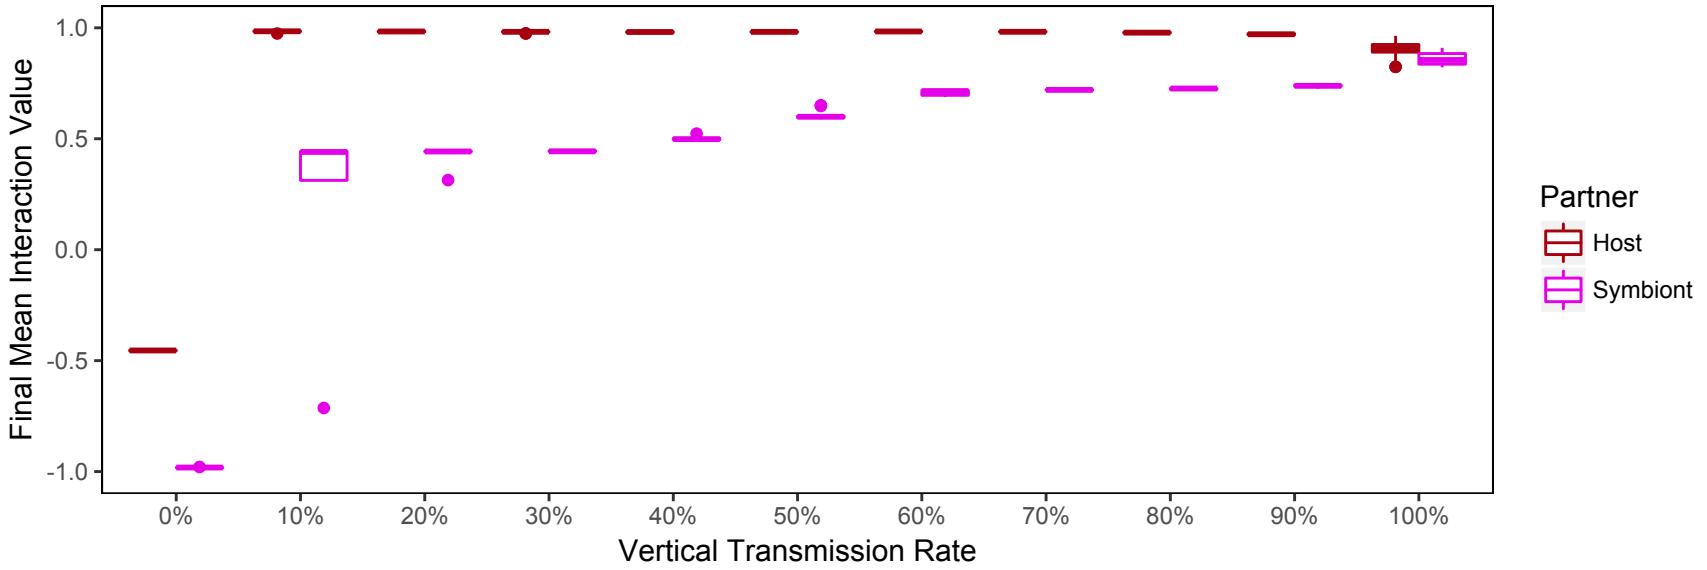

Figure 4.2: **Evolution of mutualisms across vertical transmission rates.** Final resource behavior values are shown after evolution across vertical transmission rates from 0% to 100%. Vertical transmission rate clearly strongly influences the final resource behavior value of each partner. At 0% vertical transmission rate, the symbiont becomes fully parasitic and the host defensive. At 100% vertical transmission rate, the symbiont and host are fully invested in a mutualism, to the point that the symbiont has lost the ability to reproduce horizontally.

As shown in Figure 4.2, I found that a vertical transmission rate of 0% led to a clear parasitism where the symbiont is on average attacking the host as much as it can and the host is investing half of its resource into defense. This confirms my hypothesis that when there is no chance of the symbiont being vertically transmitted, it will not cooperate with the host and instead steal as many resources as possible for horizontal transmission.

Conversely, at 100% vertical transmission rate, an extreme mutualism evolves where the host donates all of its resources to the symbiont and the symbiont donates all of its resources back to the host. This result means that the symbiont has lost the ability to horizontally transfer because it can no longer amass any resources to do so. Similarly, the host is completely dependent on its symbiont such that it would be unable to reproduce without it (assuming the resources it tried to donate were therefore wasted). This extreme result confirms the hypothesis that a guaranteed vertical transmission will align the symbiont's selective pressure with the host's and possibly lead to a major transition over evolutionary time.

At intermediate vertical transmission rates of 10-90%, the final relationship supports the idea of imperfect mutualisms [55]. In these cases, the symbiont is not a perfect partner, but it could contribute more resources back to the host instead of keeping some for itself. However, the host still has an resource behavior value near to 1. This result is an example of when the host is still selected to invest in the mutualism with an imperfect partner because having an imperfect partner is better than no partner at all. Because of the synergy effect (x5), the host only needs to get 10% of its contributed resources back to break even. Any higher resource behavior value above 10% in a symbiont is beneficial to the host.

The transition from most likely to evolve to a mutualism and most likely to evolve to a parasitism occurs between 10 and 0% vertical transmission rate. To better understand the dynamic of this transition, I ran additional experiments at every whole percentage vertical transmission rate between 0 and 20%. In Figure 4.3 it is clear that the transition between effects of vertical transmission rates are sharp between 7% and 11% with intermediate final population states at 8, 9, and 10% vertical transmission rates. This result suggests that there are few vertical transmission rates that

can lead to an intermediate population state without other confounding factors, but it is possible.

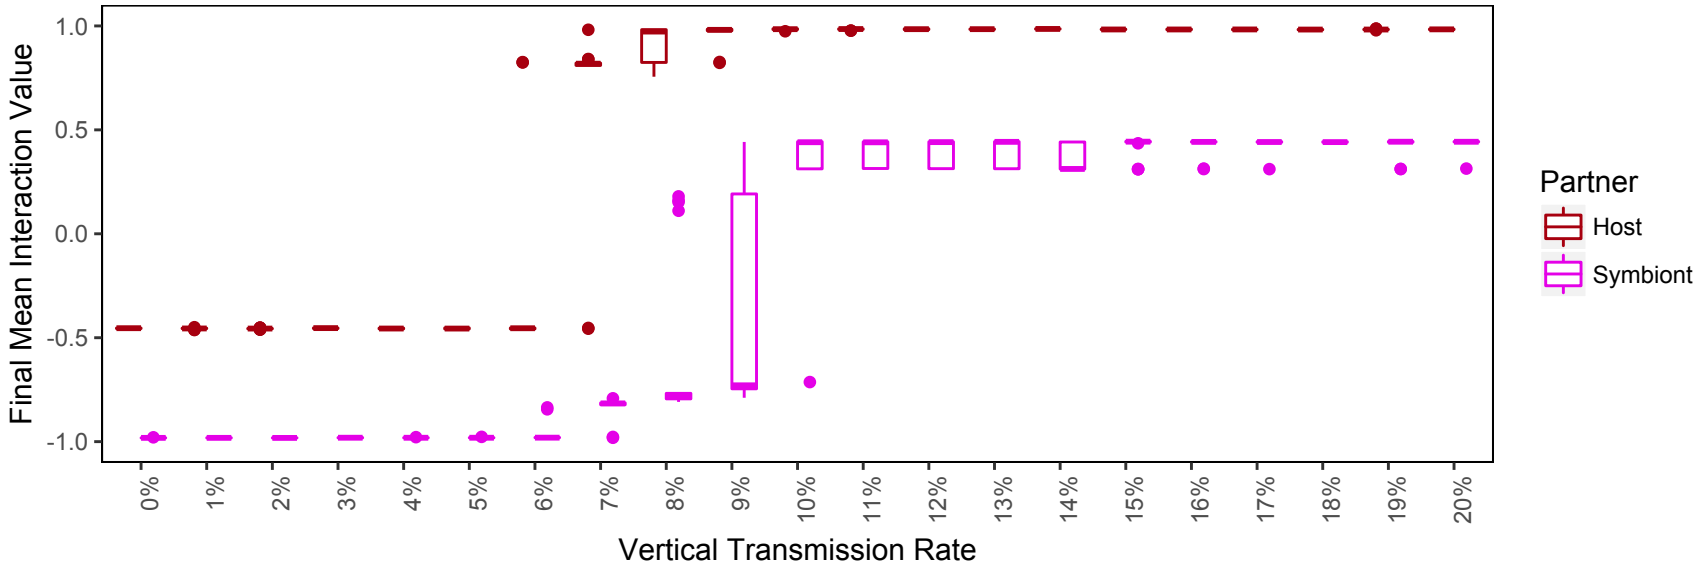

Figure 4.3: **Evolution of behavior between 0% and 20% vertical transmission rates.** Vertical transmission rates of 8, 9 and 10% show an intermediate final state. However, the other treatments make it clear that generally there is a sharp tipping point between a population ending in mutualism or parasitism.

The intermediate state of populations at 8 and 9% vertical transmission rate is particularly strange because the average is -0.598 and -0.286, respectively, suggesting that symbionts in these populations are at a behavior of below 0 yet the hosts are still investing in these useless symbionts. However, when the phenotypes of the symbionts are not reduced to a single average, as in Figure 4.4, it is clear that these are not populations of purely parasitic symbionts and mutualistic hosts. Instead, there appears to be a stable coexistence of parasitic and mutualistic symbionts.

However, when the phenotypes of individual populations are examined, as in Figure 4.5, the three possible scenarios become clear. In 17/20 replicates, a coexistence between extreme parasites and moderate to extreme mutualists was able to persist to 100,000 updates. Notably in the 3/20 replicates that did not have a stable coexistence, a mutualistic phenotype was dominant.

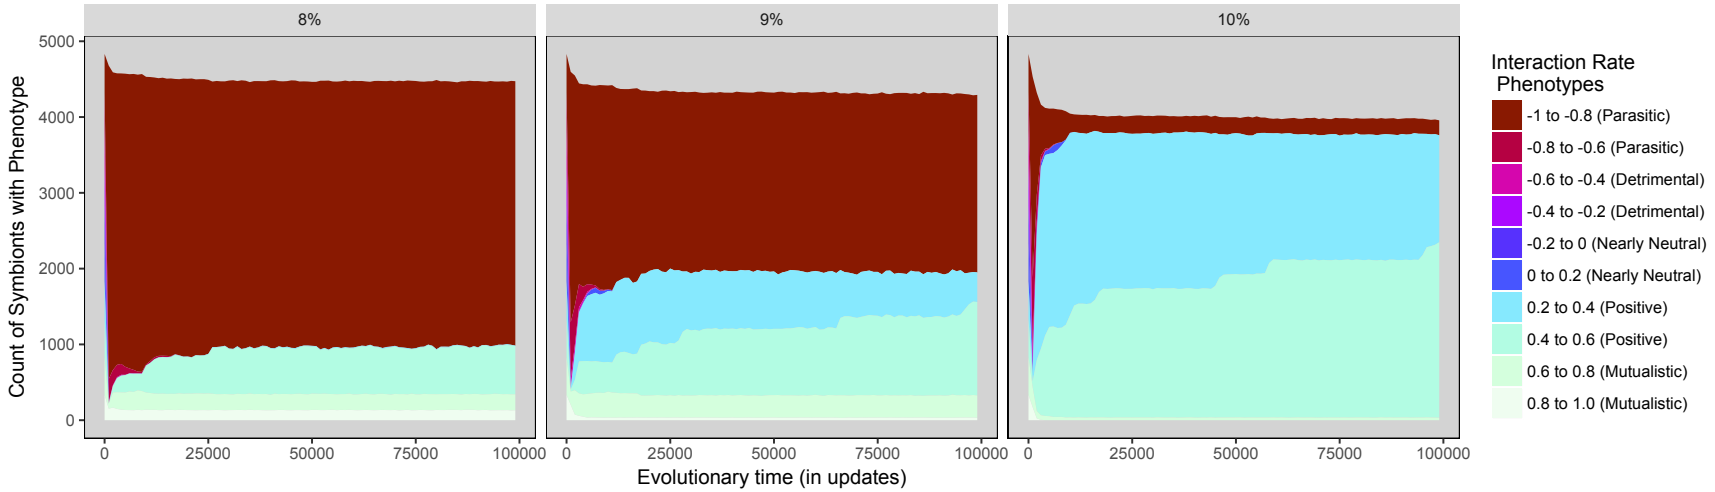

Figure 4.4: **Phenotypes of symbiont at vertical transmission rates 8, 9, and 10%.** At each intermediate vertical transmission rate a stable coexistence between parasitic and mutualistic symbionts persists through evolutionary time.

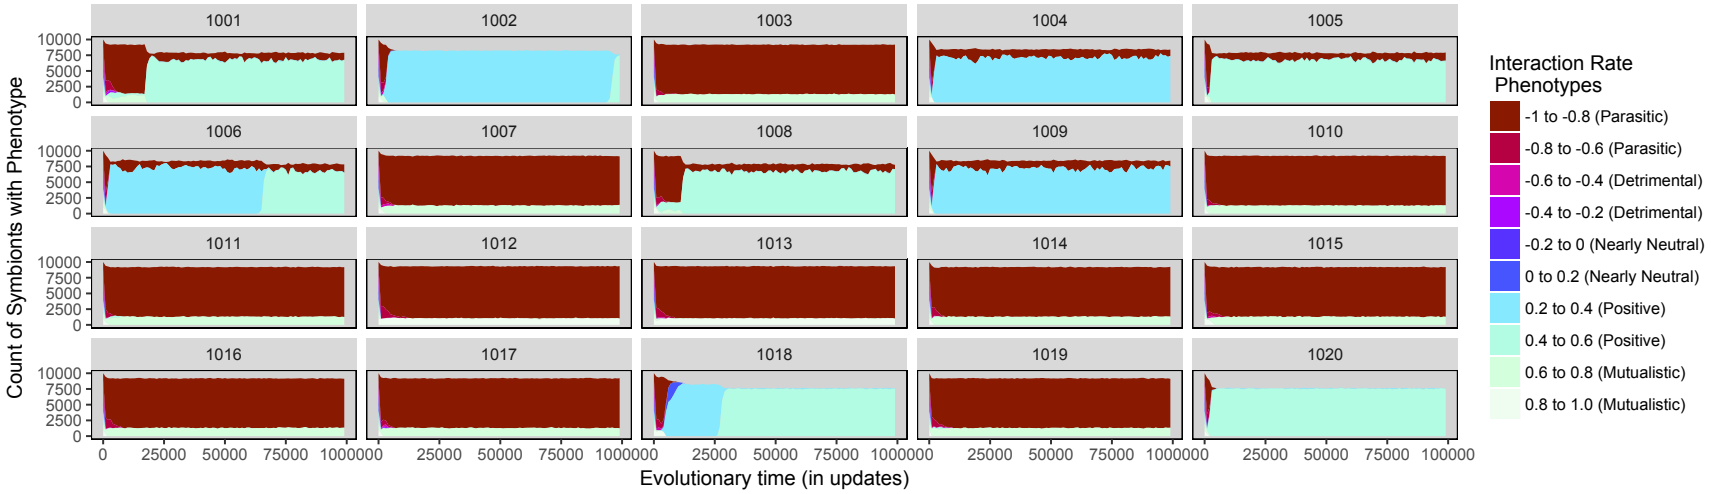

Figure 4.5: **Phenotypes of each replicate symbiont population at vertical transmission rate 9%.** Twenty replicate populations evolved at 9% vertical transmission rate are shown. Two stable types of co-existence exist with either imperfect mutualists or extreme parasites as dominant.

Further, there are two distinct types of coexisting populations: 1) majority moderate mutualists and minority parasites (6/20 replicates, ex: 1001 in 4.5) and 2) majority parasites and minority extreme mutualists (11/20 replicates, ex: 1003 in 4.5).

Together these results show that even in a simple agent-based system, complex dynamics can emerge and there are multiple stable states that populations can evolve to. I have confirmed that increased vertical transmission rate does more strongly select for mutualism, however at intermediate vertical transmission rates, stable coexistence of parasites and mutualists is possible.

4.3.3 How does vertical transmission affect the evolution of mutualism?

The previous section focused on whether mutualistic traits could invade an environment when starting as an established subset of the population. However, a population with high standing variation in a mutualistic trait will not necessarily have the same dynamics as a population starting with a neutral phenotype (neither mutualistic nor antagonistic) and therefore the *de novo* evolution of a mutualistic trait against a neutral background must be explored separately.

I created starting populations where all organisms started with a phenotype of 0, meaning they neither donated to their partners nor acted antagonistically. As seen in Figure 4.6, at all vertical transmission rates except for 100%, starting with a neutral background led to the extinction of the symbionts before they were able to evolve to parasitism or mutualism. At 100% vertical transmission rate, a small proportion of more mutualistic hosts starts to emerge at the end of the experiment, as seen in Figure 4.7, however these final populations are clearly less mutualistic than populations that start with standing diversity as discussed in the previous section.

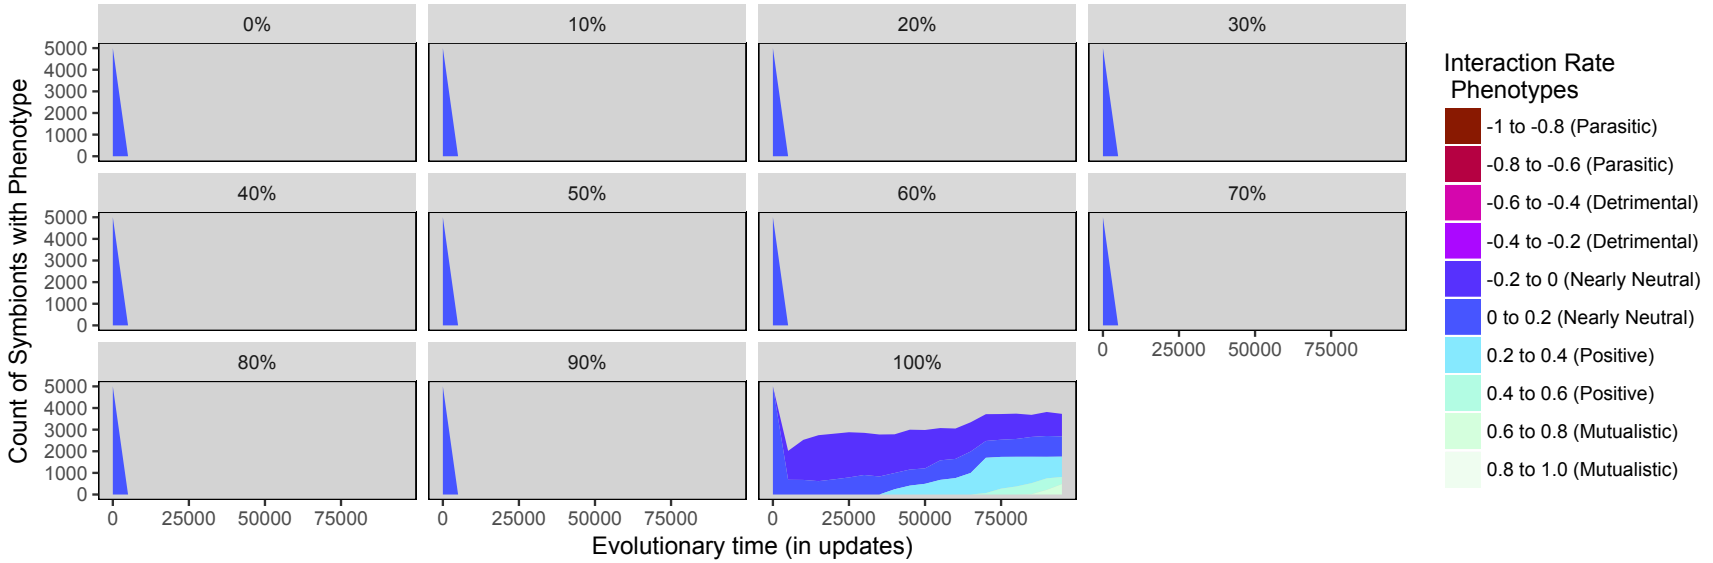

Figure 4.6: **Phenotypes of Symbionts across vertical transmission rates when hosts start at neutral phenotype.** Symbionts go extinct except when the vertical transmission rate is 100%.

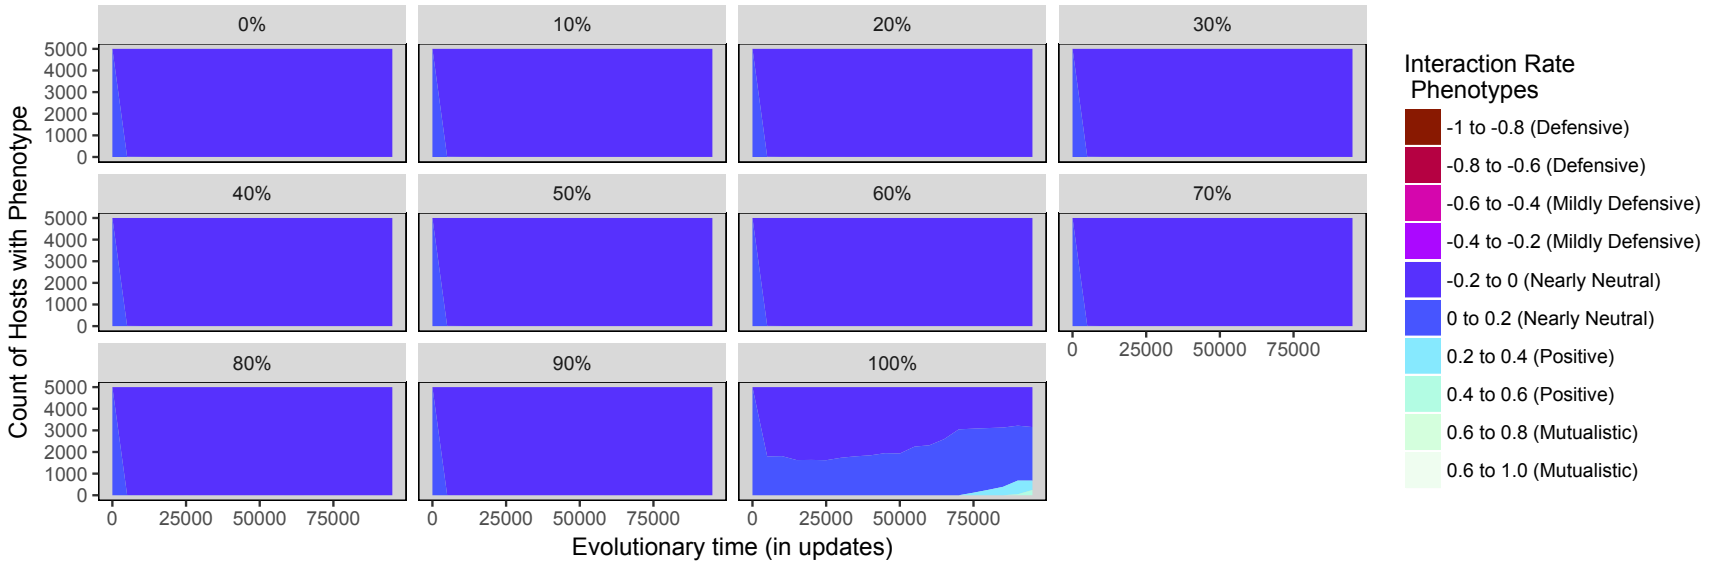

Figure 4.7: **Phenotypes of hosts across vertical transmission rates when ancestral hosts all start with the neutral phenotype.** Because symbionts go extinct (see 4.6), host resource behavior value is not under selective pressure and does not evolve except at 100% vertical transmission rate.

These results suggest that the conditions that allow mutualisms to become established and stable in a population do not necessarily also enable a mutualism to emerge *de novo* in a population of neutral phenotypes. This finding would imply that other mechanisms would need to be in place to allow a mutualism to grow to a sufficient proportion of the population or allow a population to already have a diverse set of organisms when selection on the mutualist trait becomes active.

4.3.4 How does the ancestral symbiont behavior affect the evolution and maintenance of mutualism?

The initial behaviors of symbionts is likely to influence the likelihood that mutualisms successfully evolve or can be maintained in the system. Symbionts frequently evolve to infect a new host species [55], thus beginning their co-evolution with that host as potentially already parasitic or mutualistic. If a symbiont at the beginning of an evolutionary experiment has more parasitic behaviors, I hypothesized that a stable mutualistic relationship would be less likely to arise in the longer term than if the symbiont has initially random or mutualistic behaviors. Therefore, I designed three treatments to test how the symbionts' behavior at the start of evolution changes the ending relationship between symbiont and host. My control treatment and the default is for a new symbiont's genome to be randomly assigned values at the start of evolution. I then created a parasitic population of symbionts to start the parasitic behavior treatment by forcing the symbionts' resource behavior values to be negative. Conversely, I created ancestor symbionts with more mutualistic behavior by changing their resource behavior values to be positive, thus exploring the maintenance of mutualisms. The values could still be between -1 and 0 or 0 and 1, respectively.

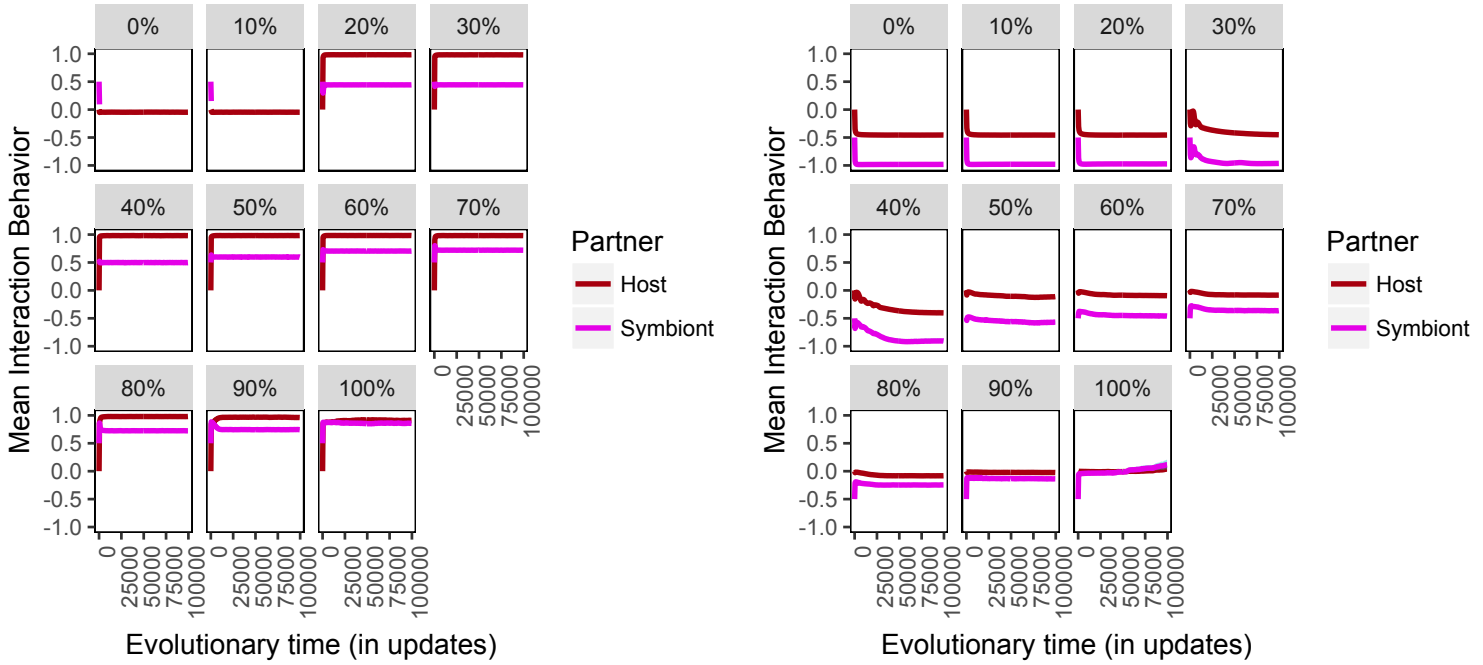

(a) **Symbiont starts more mutualistic.**

(b) **Symbiont starts more parasitic.**

Figure 4.8: **Evolution of mutualisms across vertical transmission rates when symbiont starts (a) more mutualistic than random or (b) more parasitic than random.** (a) When a symbiont starts as more mutualistic in behavior, mutualisms are able to evolve at vertical transmission rates 20% and higher, where tended to fail when the symbiont started with a random behavior. (b) When a symbiont starts as more parasitic in behavior, it is more difficult for a mutualism to evolve than when the symbionts start with random behaviors. However at vertical transmission rate 100% the mutualism is able to recover from the parasitic starting behavior.

As seen in 4.8 the general starting phenotype of the symbiont significantly affects whether a mutualism will evolve or be maintained at vertical transmission rates between 20% and 90%. This result indicates that the evolutionary history of the symbiont when it moves to a new host can completely alter the relationship with that host.

Once again the population averages in the intermediate vertical transmission rates of the parasitic starting treatments are near 0%, indicating that there is a coexistence or multiple stable phenotypes. In Figure 4.9 a coexistence between two parasitic symbionts occurs at the vertical transmission rates where a mutualism would evolve if the symbiont had started with random phenotypes. This result indicates that this type of coexistence between one extreme phenotype and one less extreme phenotype is not dependent on a mutualistic host.

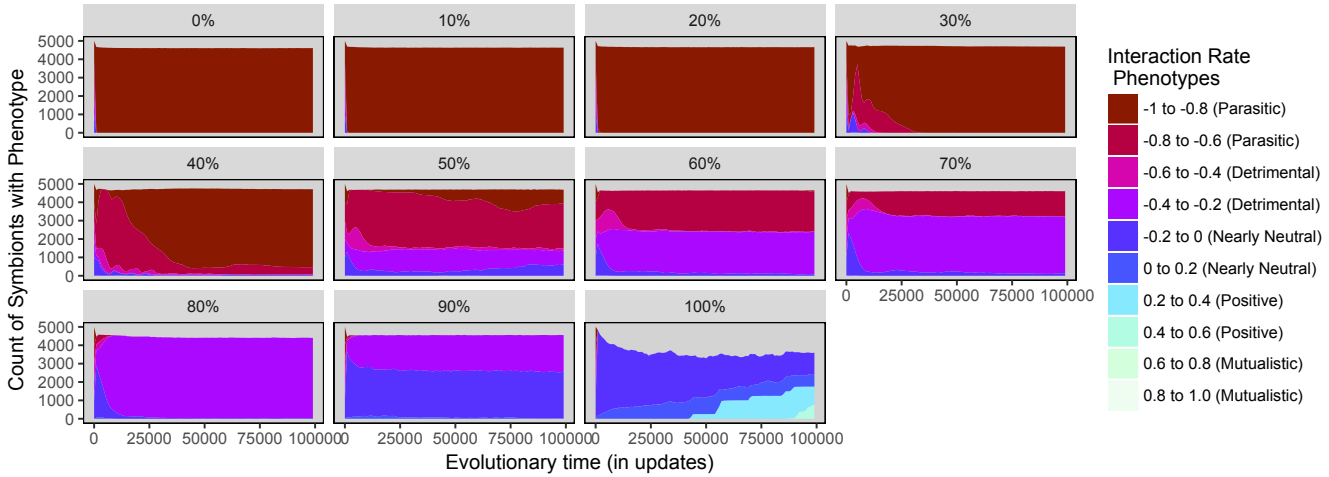

Figure 4.9: **Phenotypes of symbionts starting more parasitic than random across vertical transmission rates.** When starting at more parasitic phenotypes, stable coexistence of parasitic and mutualistic phenotypes can evolve only at 90% or 100% vertical transmission rate as opposed to all rates about 10% when starting random.

While the previous section shows that a mutualism cannot evolve de novo in this system from a neutral starting population, these results indicate that a new mutualism could evolve from a host range change of a symbiont. If the symbiont starts interacting with a new host and is more mutualistic than random, a mutualism can evolve under a broader range of vertical transmission rates.

4.4 Multiple Infinite Resources

Many natural mutualisms are successful because the endosymbiont performs a task, such as defense, that the host is incapable of doing itself (as in the classic example of the ants and acacia trees [52]). This division of labor leads to a synergy effect where the symbiont is able to increase the energy yield for the host when the host donates resources to it. In the previous experiments this effect was artificially set. Here I introduce resource types and allow the synergy effect to emerge due to true division of labor between host and symbiont. I then compare this to a control where the host is able to choose all resource types and therefore has no need of the symbiont.

4.4.1 Methods

If resource types are enabled, at least two types of resources are available in the environment in unlimited quantities, but only a set amount of each resource can be brought into the host in a given time step. If the symbiont chooses different resources than the host, its host is able to process that additional type of resource efficiently, increasing the overall amount of resources available to the pair at that time step. For example, if there are two possible resources, A and B, and the host and symbiont can choose only one resource type each, the host might choose resource A and the symbiont might choose resource B and enable the host to efficiently digest resource B. A host that overlaps with its symbiont is able to efficiently digest only a single type of resource and therefore fail to obtain half of the available resources.

The host first efficiently digests all of its available specialty resource. Its resource behavior value then determines the proportion of non-specialty resource types it gives to its symbiont. The

symbiont then processes whatever types it can and returns a proportion of those resources to the host. The host is able to efficiently make use of the symbiont's specialty resource after the symbiont processes it.

When the relationship is antagonistic and resource types are enabled, as shown in Figure 4.10, the chosen resource types determine the type of defense the host invests in and the type of attack the symbiont uses. By default 25 (example uses 10 for clarity) of each type of resource is again available to the host, but the symbiont is trying to steal some of them and the host is only able to defend one by default.

If the host and symbiont do not have matching chosen resources, the host uses a portion – determined by its resource behavior value – of its non-specialty resources, in the example resource B, to protect its specialty resource. If the symbiont has a more negative resource behavior value than the host, it is able to overcome the host's defenses and steal a portion of the host's specialty resource in addition to what remains of the host's non-specialty resources. If the symbiont has the same specialty resource as the host, the symbiont is only able to access that specialty resource, and the non-specialty resource is wasted. If the host is able to fend off the symbiont, each partner receives half of the B resource (host's non-specialty), and the host keeps all of its specialty resource.

To summarize, when the host and symbiont are mutualistic, jointly they are more fit when they have different chosen resource types. However, when the relationship is antagonistic, the hosts will be most fit if they choose the *same* resource type as the symbiont. Conversely, the symbionts will be most fit if they choose a *different* resource type from the host, leading to conflict between hosts and parasites.

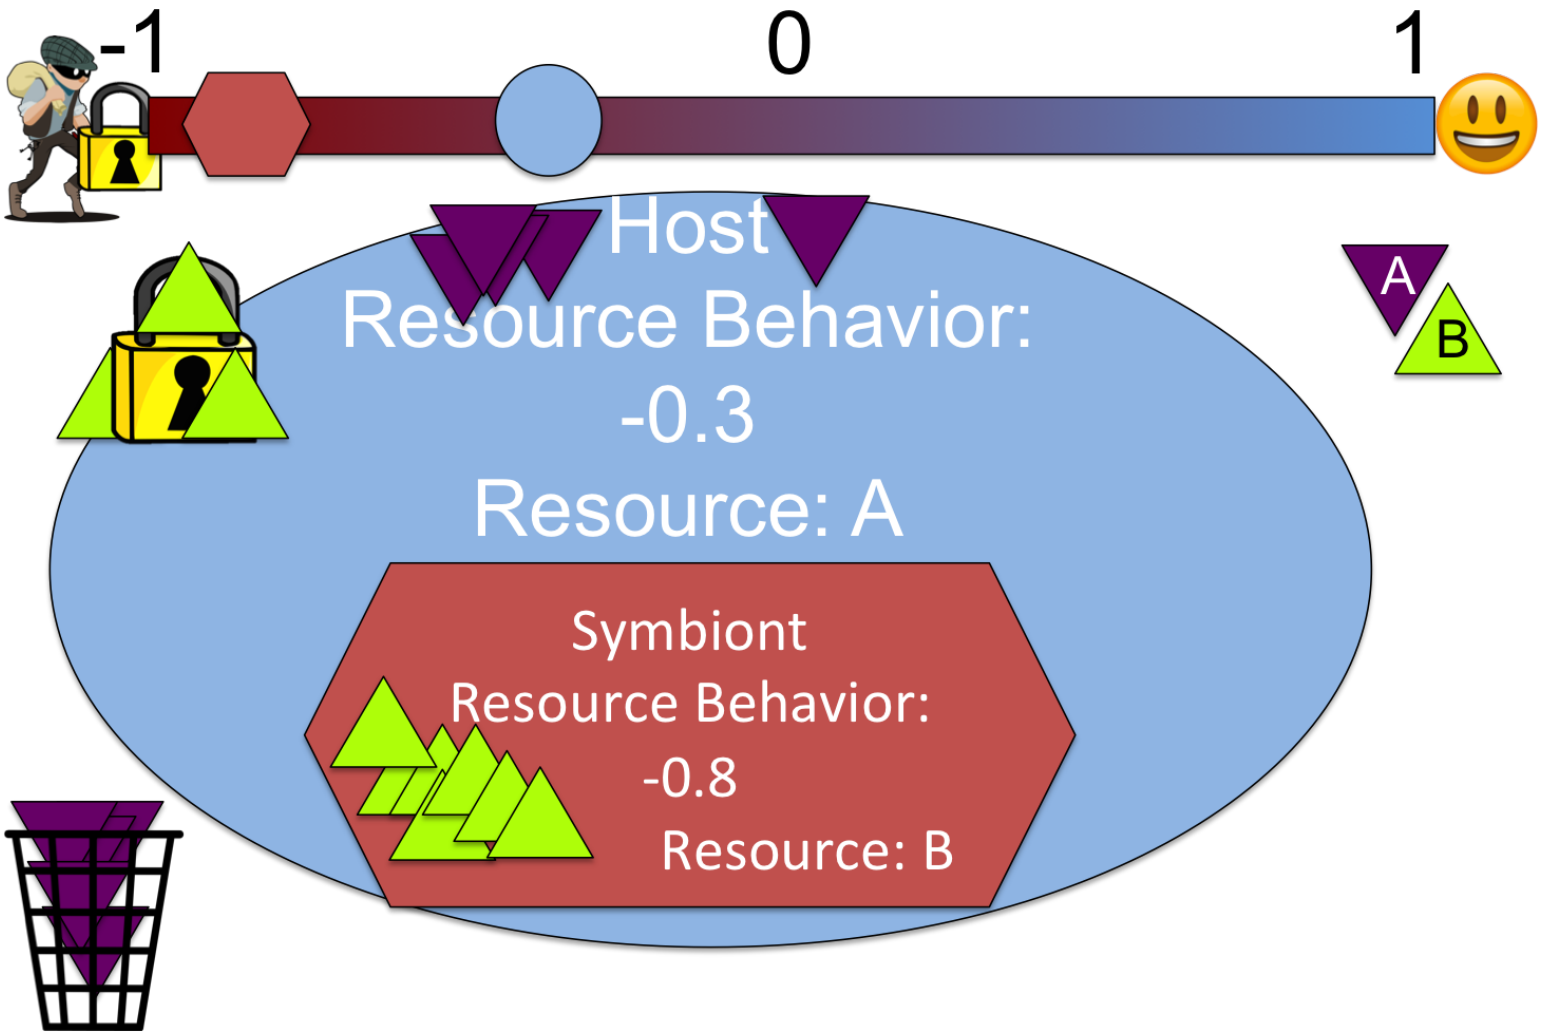

Figure 4.10: **Example host and symbiont antagonistic resource behavior with non-matching chosen resource types.** The host uses up 3 B in defense of resource A. The symbiont is able to destroy some of A and steal what is left of B.

Table 4.2: Interaction results for host and symbiont phenotypes with resource types

Host Phenotype	Symbiont Phenotype	Chosen Resource Types	Result
$X > 0$	$Y > 0$	Types match (ex: both A)	Host digests all of A, donates X of B to symbiont, inefficiently digests 1-X of B; symbiont cannot digest B and so X of B is wasted
$X > 0$	$Y > 0$	Types mismatch (ex: Host A, Sym B)	Host digests A and donates X of B to symbiont, host inefficiently digests 1-X of B; symbiont processes B and returns Y proportion processed B, symbiont gains 1-Y of processed B; host digests processed B
$X < 0$	$Y < 0$	Types match (ex: both A)	Host spends -X of B for general defense; symbiont tries to steal A, if $Y < X$, symbiont wastes Y-X of A and steals -Y of remaining A; if $Y >= X$ each get half of A; B is wasted
$X < 0$	$Y < 0$	Types mismatch (ex: Host A, Sym B)	Host spends X of B for general defense; symbionts tries to steal B, if $Y < X$, symbiont wastes Y-X of A, steals -Y of B; if $Y >= X$ each get half of B; host gets remaining A
$X > 0$	$Y < 0$	Types match (ex: both A)	Host donates X of A to symbiont, symbiont takes all remaining A, B is wasted

Table 4.2 – *(cont'd)*

Host Phenotype	Symbiont Phenotype	Chosen Resource Types	Result
$X > 0$	$Y < 0$	Types mismatch (ex: Host A, Sym B)	Host donates X of A to symbiont, symbiont takes all of B, host gains what remains of A
$X < 0$	$Y > 0$	Types match (ex: both A)	Host digests A; Host spends X of B on defense; Symbiont tries to digest A but none left; B is wasted
$X < 0$	$Y > 0$	Types mismatch (ex: Host A, Sym B)	Host digests A; Host spends X of B on defense, another 0.5 of B is lost due to symbiont having to fight to get to it; Symbiont donates Y of what B it gets to host as processed and keeps 1-Y of the B it got for itself

All combinations of resource behavior value and chosen resource type are presented in Table 4.2. These interaction dynamics were selected because they produce a fitness landscape that creates conflict between the partners. As seen in Figure 4.11, the region of the fitness landscape that is best for the symbiont is the opposite of that that is ideal for the host. This conflict is an abstraction of the many mutualism-parasitism systems where it is assumed that there is this conflict due to observed interactions.

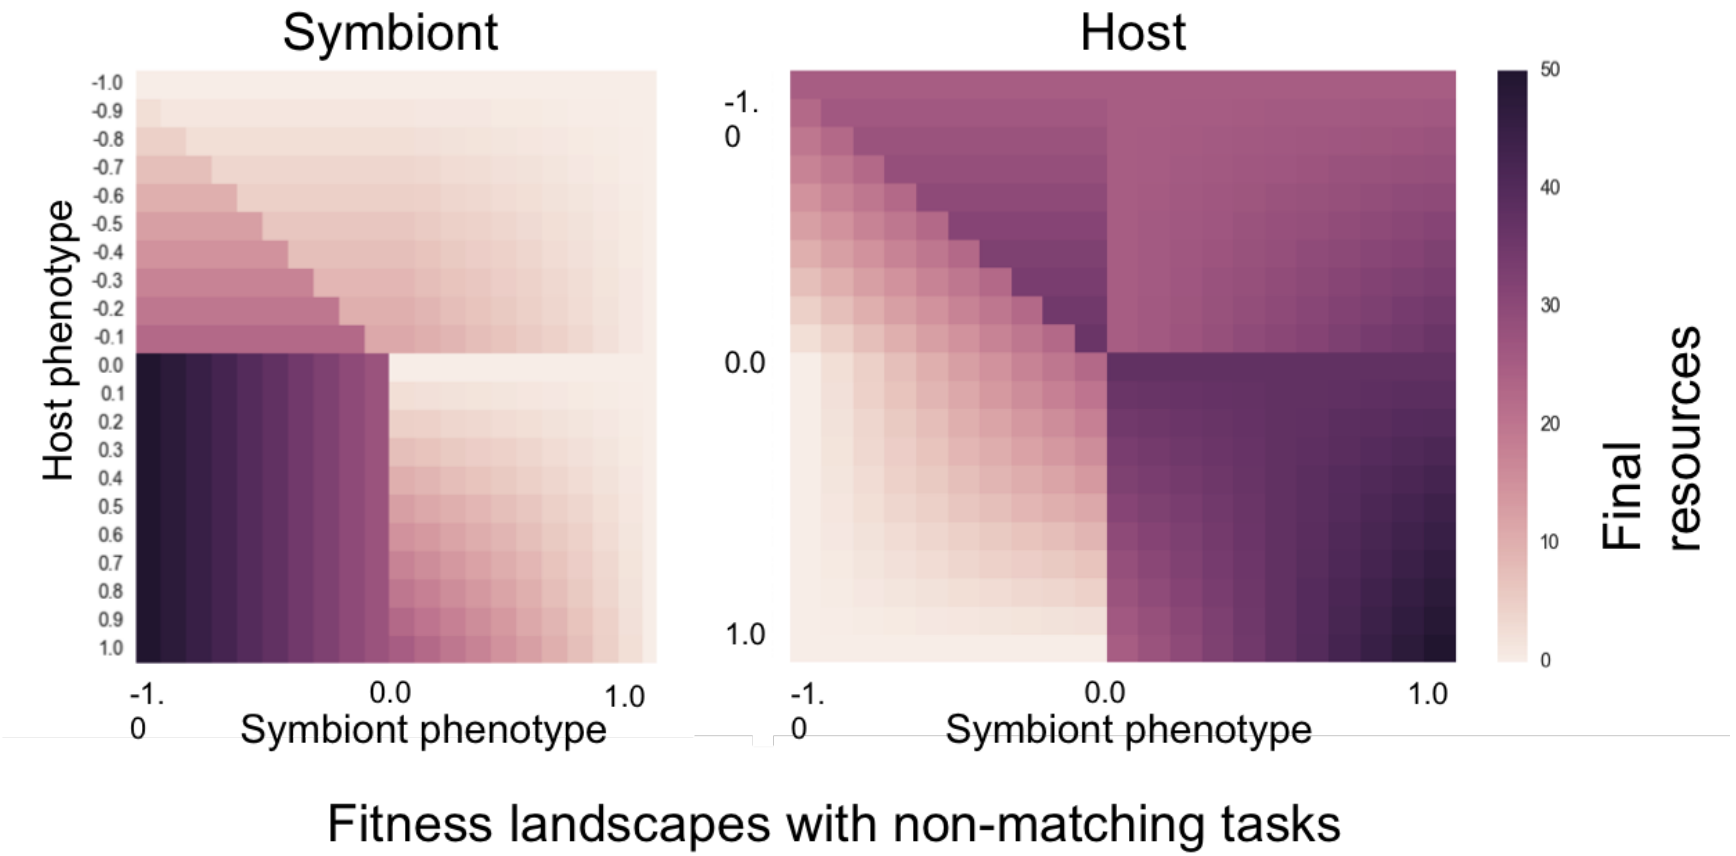

Figure 4.11: **Fitness landscapes for Symbiont and Host with non-matching resource types.** The symbiont and host fitness landscapes show that the two partners conflict on which phenotypes are best for each of them.

4.4.2 Is a mutualism able to evolve without a forced synergy effect?

By having the organisms choose resources to use, I am no longer enforcing an artificial synergy by multiplying the returned symbiont's resources by 5. Instead, the host and symbiont only have a synergy if they use different resources. I first verified that the full range of parasitism to mutualism was able to evolve across vertical transmission rates, as seen in 4.12

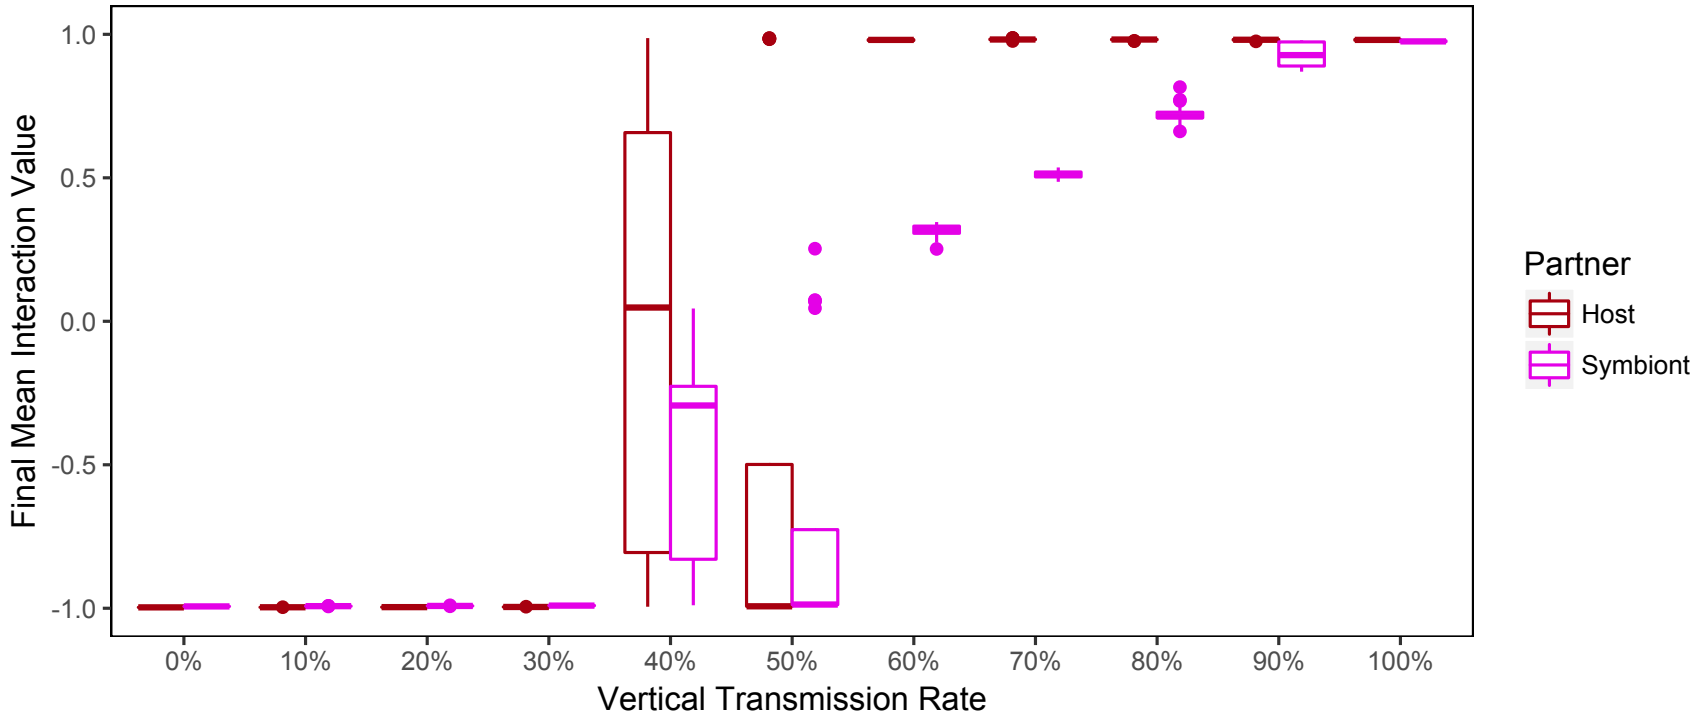

Figure 4.12: **Evolution of mutualism across vertical transmission rates without artificial synergy.** When an artificial synergy factor is replaced with a set of resource types that are divided based on the host and symbiont phenotypes, mutualisms are still able to evolve at higher vertical transmission rates.

I then determined that the coexistence of multiple symbiont phenotypes was not an artifact of the artificial synergy, as shown in 4.13.

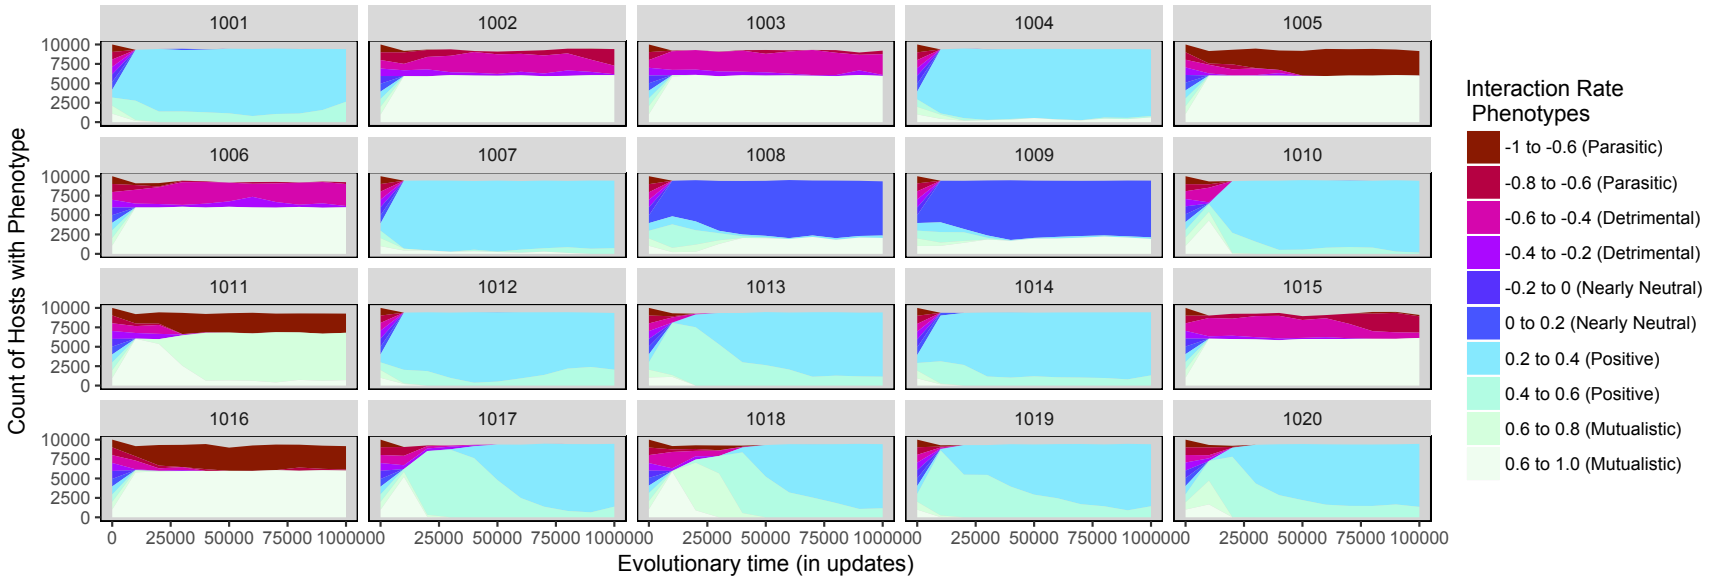

Figure 4.13: **Count of symbiont phenotypes across evolution of mutualism when vertical transmission is 60%** As with a forced synergy, multiple phenotypes of symbiont, some parasitic and some mutualistic, are able to stably coexist with a more natural resource type system.

Further, the host population at vertical transmission rate 60% has multiple coexisting phenotypes, some defensive and some mutualistic, as shown in 4.14.

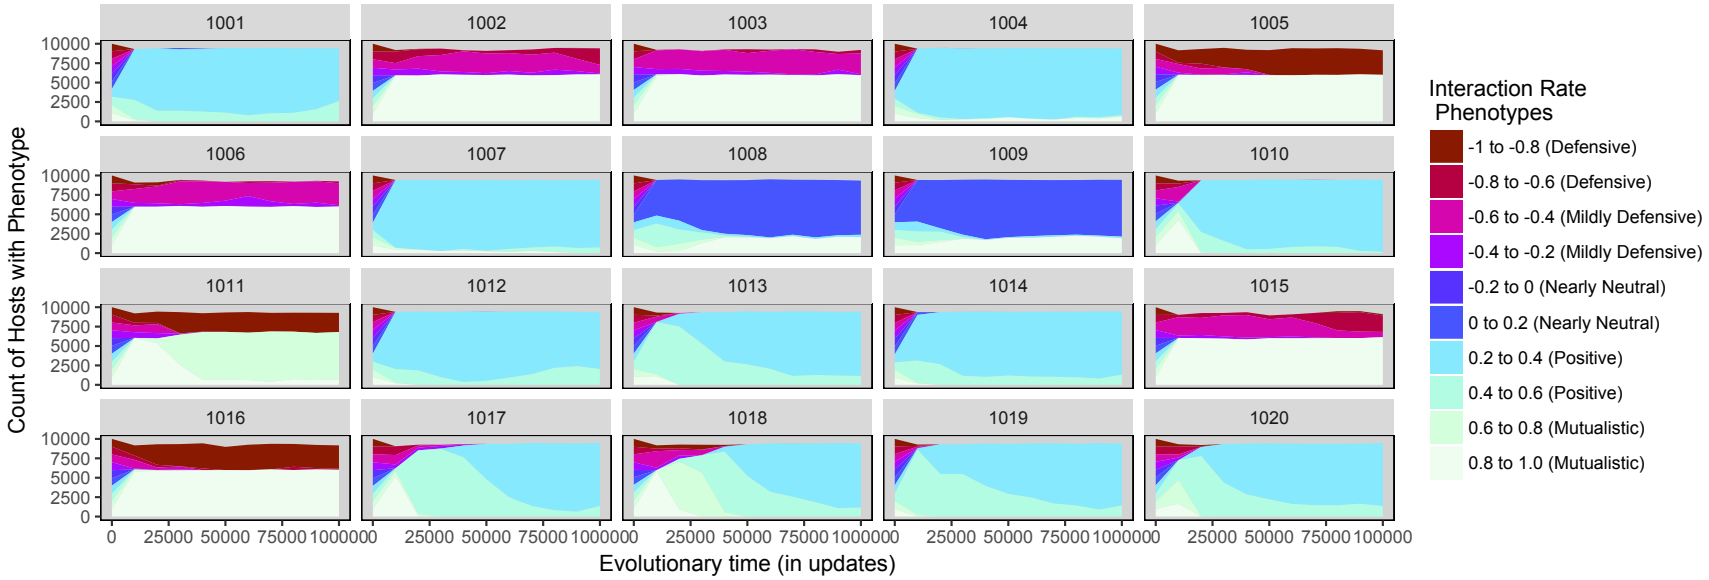

Figure 4.14: **Count of host phenotypes across evolution of mutualism when vertical transmission rate is 60%** Host populations evolved coexistence of phenotypes.

Finally, as seen in 4.15, the resource type system enables the host to attempt to defend itself from parasitic symbionts by changing resource types. This change would then pressure the symbiont population to evolve to again match the host resource type, then pushing the hosts back to the original resource type. This dynamic is found in many natural systems of parasites and hosts [26], and therefore the emergence of this dynamic in Symbulation indicates that the Red Queen dynamic can occur in the simplest of host-parasite interactions.

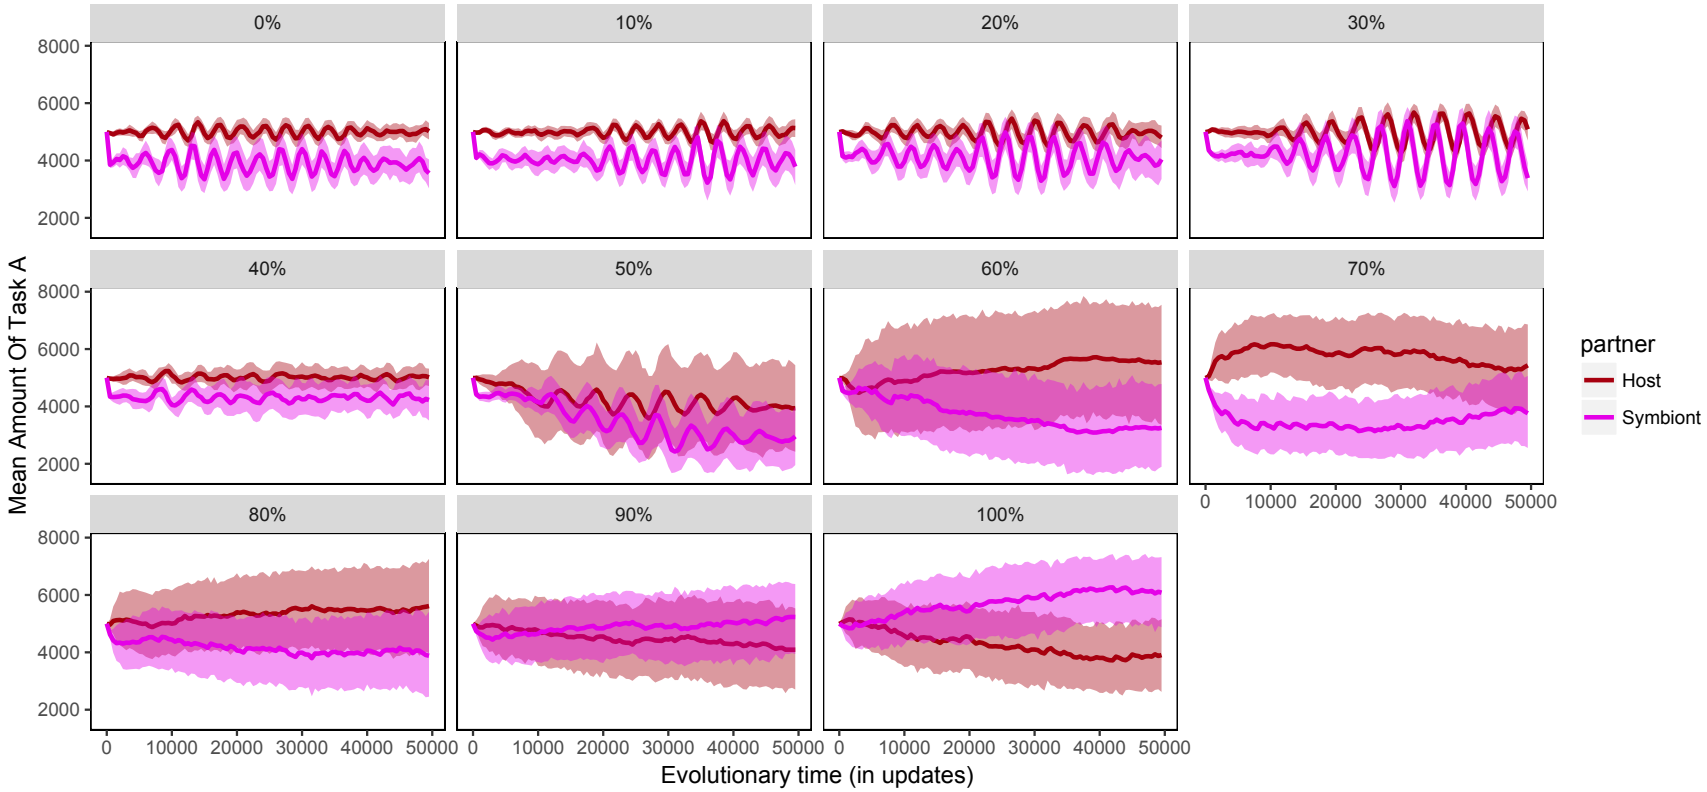

Figure 4.15: **The number of hosts and symbionts choosing resource A over time across vertical transmission rates.** When the vertical transmission rate is low and the majority of symbionts are parasitic, Red Queen oscillations emerge quickly in the populations.

4.4.3 Is division of labor between the partners necessary for the evolution of mutualism?

While I have shown that a mutualism can evolve and persist with a natural synergy created by a required division of labor, that does not mean that a division of labor is necessary for the evolution of a mutualism. To determine its necessity, I removed the limitation on how many resource types the host could choose, thus allowing the host to evolve to choose both resources A and B and no longer requiring a division of labor between host and symbiont.

When the host population is able to choose both resources A and B, it quickly evolves resistance to the symbiont, as seen in Figure 4.16. The symbiont population in fact goes extinct, as seen in Figure 4.17.

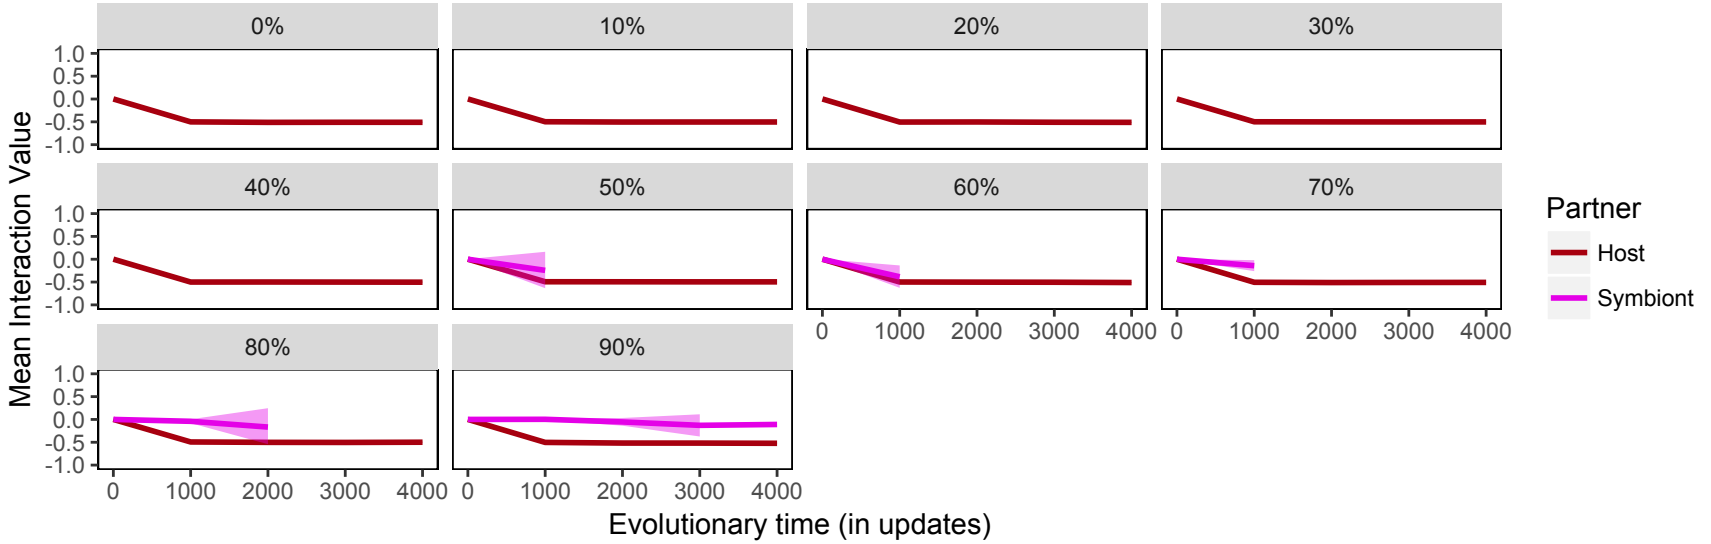

Figure 4.16: **The resource behavior values of hosts and symbionts over time when the host can choose both resource types.** The host quickly evolves to be extremely defensive and the symbiont population cannot evolve to parasitic before going extinct.

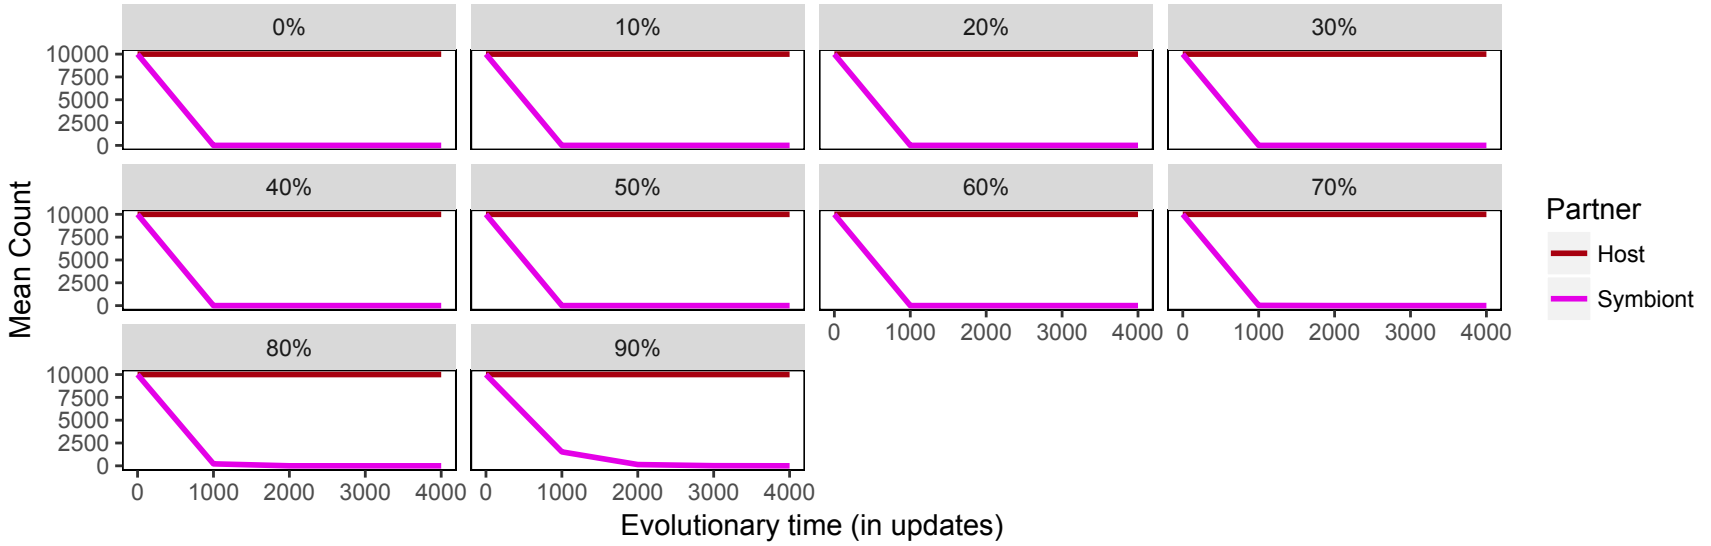

Figure 4.17: **The population counts of hosts and symbionts over time when the host can choose both resource types.** The symbiont quickly goes extinct at all vertical transmission rates except 100% (which was omitted because the host cannot force the symbiont extinct then).

This result indicates that some form of division of labor is necessary for the evolution of a mutualism, thus all of the following experiments have the required division-of-labor mechanism active.

4.5 Spatial Structure

As discussed in Chapter 3, spatial structure is often able to increase the fitness advantage of cooperators because it allows them to preferentially interact with other cooperators and exclude cheaters. Therefore, the introduction of spatial structure to a mutualism system has been hypothesized to increase the evolution of mutualism [98]. However, Akcay et al. presented a mathematical model where spatial structure in fact selected for less mutualism [5]. This result was due to the decreased diversity of symbionts in a given area, due to neighborhood dispersal of the symbionts. Akcay et al.'s system allowed hosts to be choosy regarding their symbiont partners and therefore a reduced diversity of symbionts to choose from meant that hosts in some areas were stuck with lower quality choices than the population as a whole. Conversely, in a well-mixed population, all hosts have access to the highest quality symbiont partners, potentially leading to more stable mutualism. Of course, if choosiness can evolve as well, as discussed in [36], consistently high quality partners can lead to reduced selection for choosiness. Mathematical models are limited in their ability to show how complex dynamics may combine in noisy and imperfect environments, making Symbulation an ideal system for exploring the effects of spatial structure on mutualism empirically.

4.5.1 Methods

For these experiments, we enabled neighborhood reproduction in Symbulation with two infinite resources (as discussed previously). When a host was able to reproduce, its offspring was randomly placed in one of the eight spaces surrounding the parent in a Von Neumann neighborhood. Any previous occupant of the chosen space is killed, just as before in the well-mixed reproductive model. Symbionts are able to vertically transmit with these offspring with a set probability just as

previously. If a symbiont is able to horizontally transmit, however, its offspring attempts to inject into one of the eight neighboring spaces. If there is not an uninfected host in the chosen space, the symbiont transmission fails, just as previously in the well-mixed model.

There is not an explicit mechanism for host choice in these experiments. Once a host has a symbiont, it cannot oust the symbiont. However, when a host is overwritten by another's offspring, that offspring may be born uninfected. If that offspring is fairly similar to the organism in that space previously, the dynamic is nearly indistinguishable from a host ousting its current symbiont. A host with a less cooperative symbiont is more likely to be overwritten by a reproducing neighbor, thus making it more likely that a parasite will be ousted from its current host and replaced with an uninfected host. Conversely, a host with a more cooperative symbiont is more likely to produce offspring that have some probability of inheriting the host parent's symbiont. Thus if a neighboring host is overwritten by a new host infected by a vertically transmitted symbiont, this dynamic approximates a host choosing a more mutualistic symbiont. I will refer to this dynamic as *implicit partner choice* as it closely parallels within-lifetime partner choice if host mutation rate is low enough.

4.5.2 How does local reproduction change the evolution of mutualism across vertical transmission rates?

Spatial structure or local reproduction is generally found to increase cooperation in single-species systems (see previous chapters). However, there is evidence that this might not be the case when a second species is included [5] due to the limited number of partner choices available in a spatially structured environment. Vertical transmission rate also affects the potential diversity of local organisms because the Symbulation system assumes that an established symbiont (whether passed vertically or horizontally) cannot be forced out of its host. Therefore, a higher vertical transmission rate would increase the number of hosts born already infected, potentially halting the spread of a horizontally transmitted strain and thus reducing local partner diversity. To explore how vertical transmission rate and spatial structure interact to influence the evolution of mutualism, I enabled

spatial structure and compared the resulting data to the previous section's.

As seen in 4.18, at vertical transmission rates of 60% and 70%, the final average resource behavior value is significantly higher with global reproduction than with local reproduction (60%: local mean -0.218, global mean 0.663, Mann-Whitney U test $p << 0.005$; 70%: local mean 0.385, global mean 0.759, Mann-Whitney U test $p << 0.005$). Conversely, when the vertical transmission rate is 30%, the average resource behavior value is significantly lower with global reproduction than with local reproduction (local mean: -0.693, global: -0.993, Mann-Whitney U test $p << 0.005$). This result indicates that the vertical transmission rate is a key determining factor in whether spatial structure selects for or against mutualism.

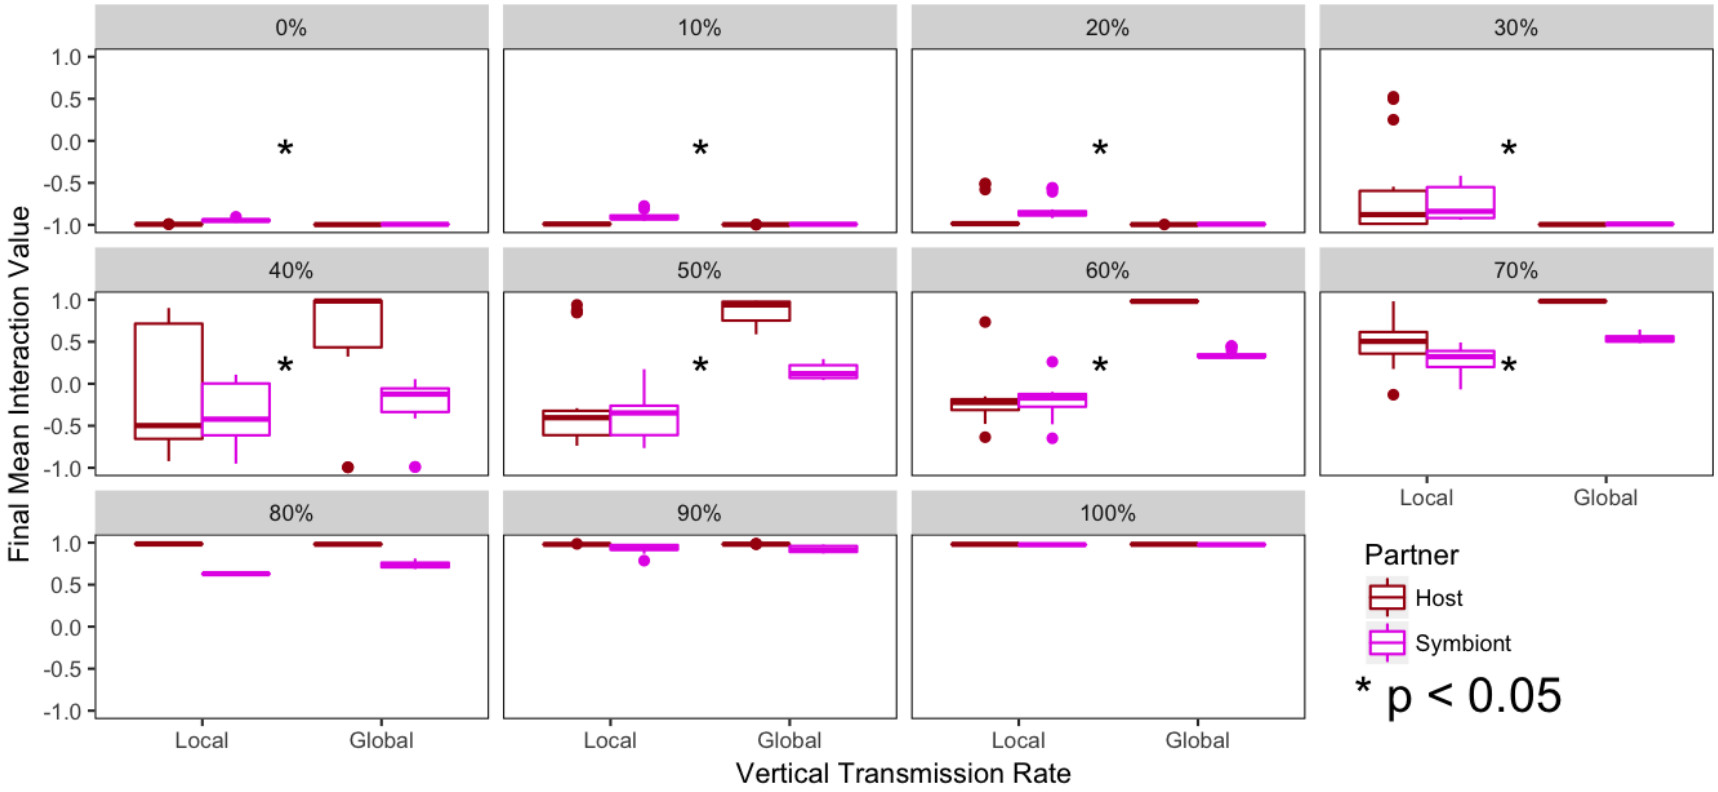

Figure 4.18: **Resource behavior values after 100,000 updates across vertical transmission rates with global and local reproduction of host and symbiont.** Considering the vertical transmission rates tested, at rates of 40% and lower, global reproduction resulted in significantly lower levels of mutualism than local reproduction. Conversely, at vertical transmission rates of 60% and 70%, global reproduction resulted in significantly less mutualism than local reproduction.

Akcay et al. predicted that spatial structure would reduce the fitness advantage for mutualisms because the hosts would not have access to the full population of symbionts when choosing partners [5]. While my experiments do not have an explicit partner choice mechanism – a host cannot oust a symbiont even if it is parasitic – this dynamic is approximated by population dynamics: a host with a low-quality partner is more likely to be overwritten before it has the chance to reproduce and possibly vertically transmit its partner. When that unfortunate host is replaced, it is likely to be by either an uninfected host or a host infected with a vertically transmitted mutualistic symbiont (since hosts with mutualists are more likely to be able to reproduce). Thus the same dynamic of low-quality symbionts being less likely to be maintained in a host through partner choice can be approximated by differential survival.

To determine whether my results are due to the same mechanism as outlined by Akcay et al., I calculated the average Shannon diversity of symbionts in a local neighborhood in both treatments. When reproduction is local, the Shannon diversity of symbionts will capture how many different phenotypes of symbiont a host lineage may 'choose'. When reproduction is global, the Shannon diversity of the immediate neighboring symbionts will be a proxy for the Shannon diversity of the full population of symbionts.

As seen in Figure 4.19, when the vertical transmission rate is 60%, global reproduction results in a significantly higher Shannon diversity of local symbionts than local reproduction. However at a vertical transmission rate of 30%, there is not a significant difference, as seen in Figure 4.20.

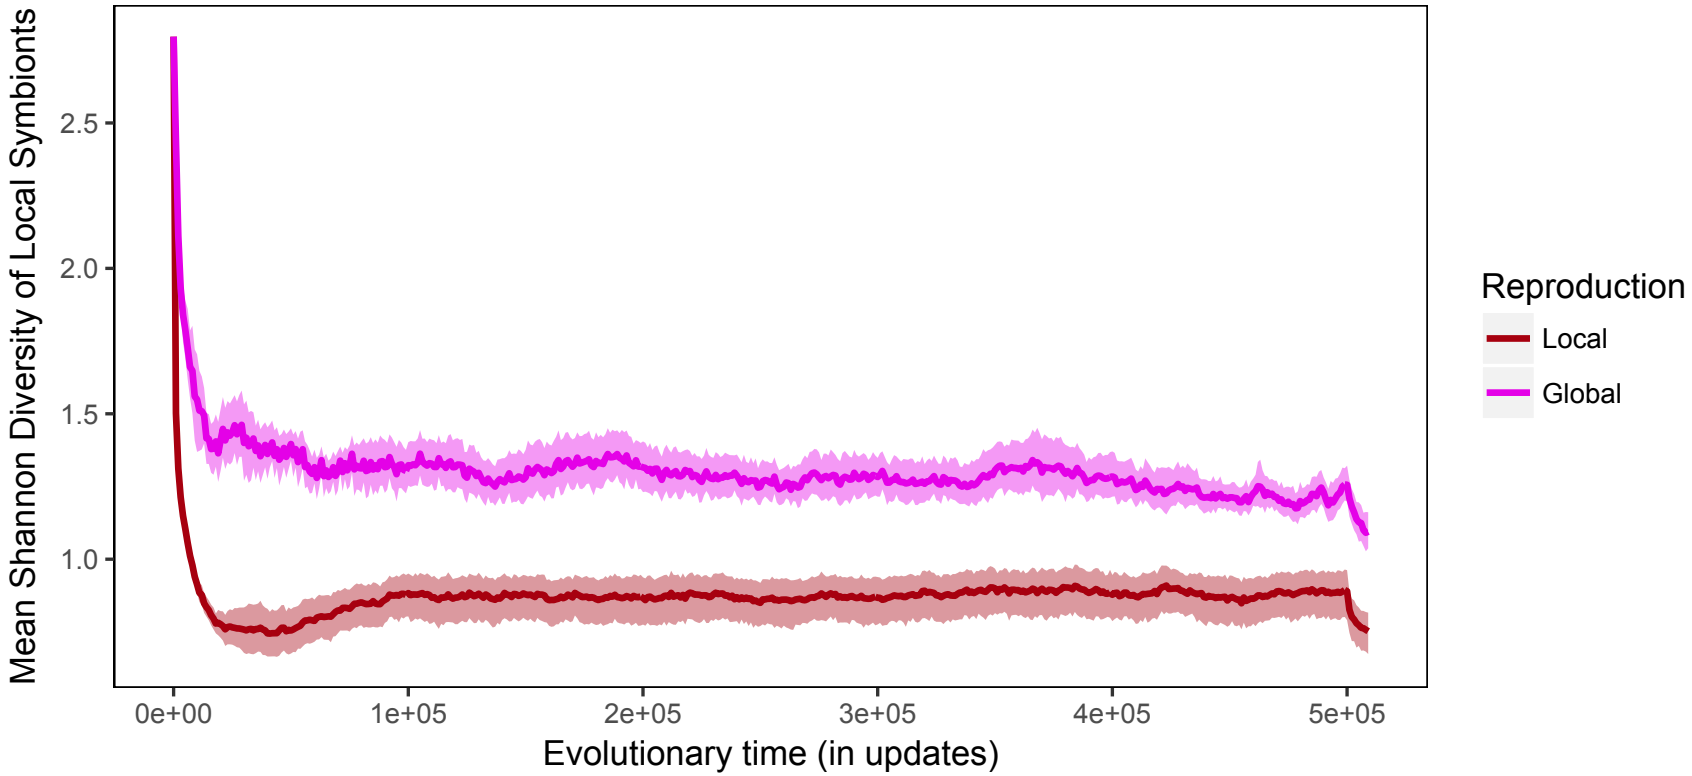

Figure 4.19: **The average Shannon diversity of local symbionts when reproduction is local and global when vertical transmission rate is 60%.** When reproduction is local, the Shannon diversity is lower than when reproduction is global, indicating the hosts have fewer choices of symbiont partner.

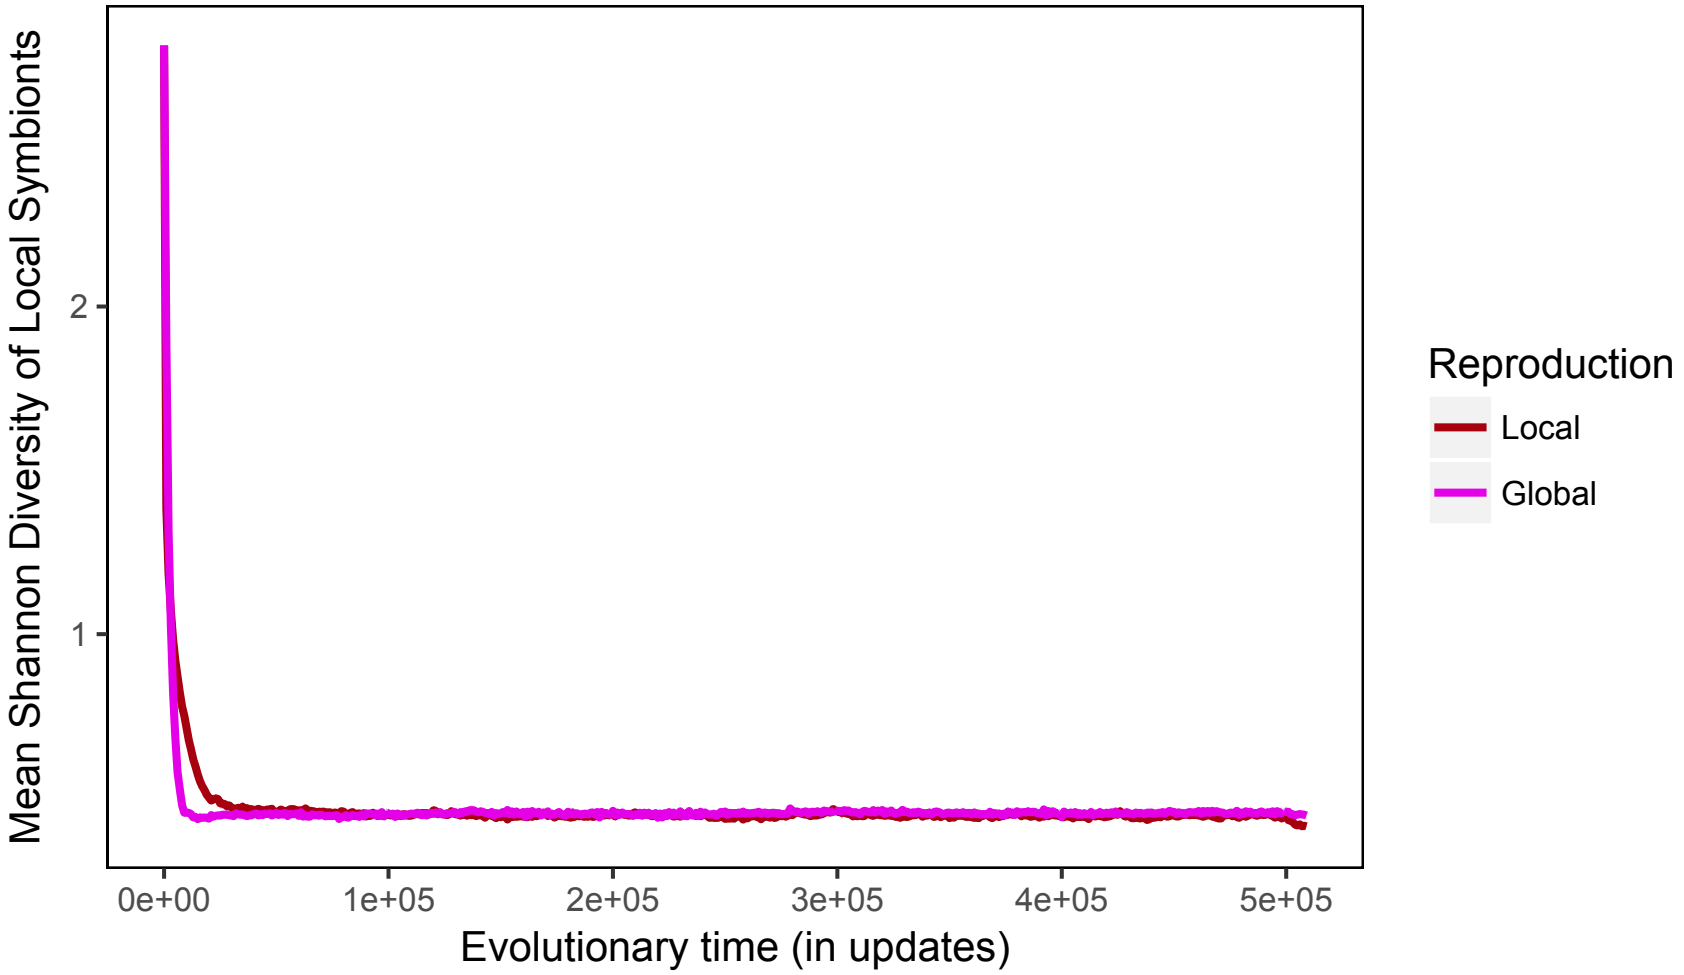

Figure 4.20: **The average Shannon diversity of local symbionts when reproduction is local and global when vertical transmission rate is 30%.** There is no significant difference between local and global reproduction.

This result confirms that spatial structure does reduce the partner choice options available to the host lineage at intermediate vertical transmission rates, leading to reduced evolution of mutualism. This result also shows that the theoretical findings hold for a system with intermediate phenotypes and additional noise.

4.5.3 Does spatial structure increase the likelihood of mutualism evolving *de novo?*

Finally, I again return to the question of how the origin of a trait may differ from the maintenance of that trait. In the spatially structured environment, a mutualism that already exists in the population is able to invade in some vertical transmission rates, even if it is less than in a well-mixed environment. However, previously I showed that mutualisms were not able to evolve from neutral hosts and symbionts when I used an artificial synergy effect. The symbionts quickly died out and the hosts evolve defenses at all vertical transmission rates except 100%. While spatial structure reduces host partner choice, it also causes a newly emerged cooperative phenotype to stay next to its relatives, making it more likely a pocket of cooperation could grow to large enough numbers to become established. To determine the effect of spatial structure on the *de novo* evolution of mutualism, I again created a population of hosts and symbionts starting with neutral phenotypes, neither aggressive nor cooperative, and allowed evolution to proceed. As shown in Figure 4.21, the symbionts do not go extinct in the spatially-structured population, but a strong mutualism also does not evolve. At vertical transmission rates of 60% and below, parasites and defensive hosts become dominant by the end of the experiment. At higher vertical transmission rates, neutral or slightly positive hosts and symbionts are able to persist, though they do not evolve to be strong mutualists. This result suggests that, while a high vertical transmission rate is not sufficient for the *de novo* evolution of mutualism as it is for the maintenance, spatial structure does allow the symbionts to remain extant. This extra time could allow another mechanism to select for a mutualism, such as partner sanctions.

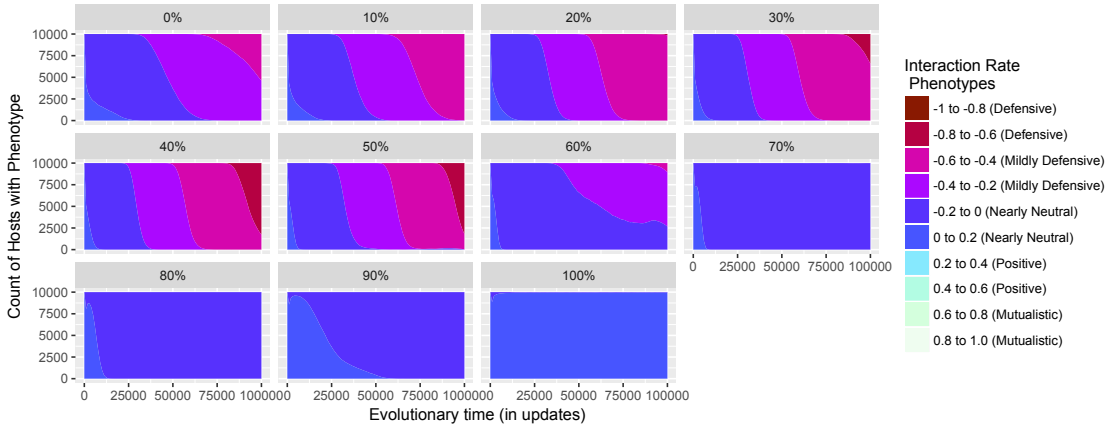

(a) Host phenotypes

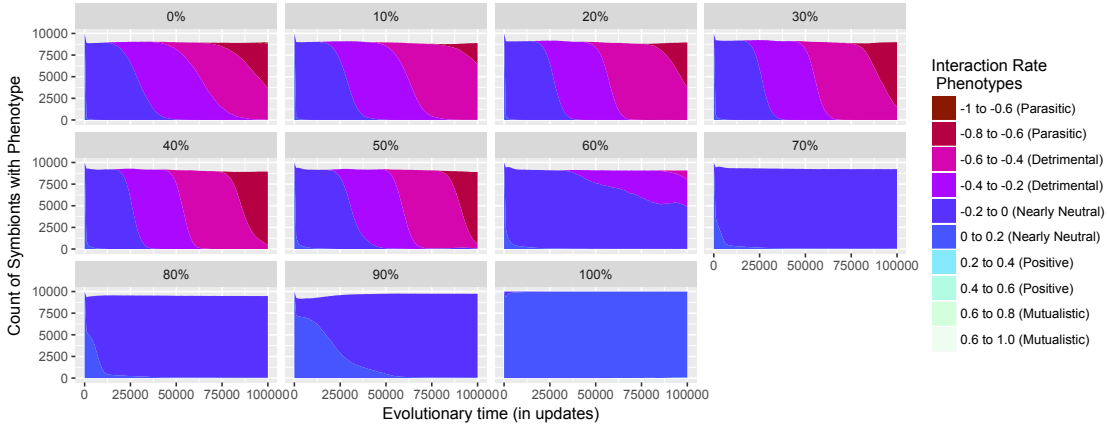

(b) Symbiont phenotypes

Figure 4.21: **Host and symbiont phenotypes across vertical transmission rates when evolving from a neutral starting phenotype.** At low vertical transmission rates, successive waves of more parasitic and defensive phenotypes evolve over time. At higher vertical transmission rates, neutral or slightly mutualistic phenotypes remain dominant over time.

4.6 Conclusion

This chapter has presented the system Symbulation to study the many theoretical hypotheses related to the evolution of mutualism. Due to the limited nature of both biological systems and mathematical models, many of these hypotheses have been difficult if not impossible to test. Here, I have shown that a mutualism can evolve with imperfect vertical transmission and imperfect partners. I have empirically tested the theoretical hypothesis that the evolution of mutualism is in fact hampered by spatial structure at some vertical transmission rates due to lower partner choice. Finally, I have shown that the conditions under which a mutualism is able to increase in a population are not sufficient for the *de novo* evolution of mutualism in well-mixed or spatially-structured populations.

Taken together, these results argue the value of Symbulation as a new tool in the study of the evolution of mutualism that can be used to experimentally confirm theoretical findings and determine conditions that are likely to lead to the evolution of a mutualism in a natural system. There are many expansions I plan to make to Symbulation to encompass more of the mutualism theory, including: within-lifetime partner choice, multi-infection, multi-births for symbionts, lysogeny, and partner sanctions. With these and other optional additional features, Symbulation will be a system that can be used to predict biological evolution of the human gut microbiome in response to changing conditions; pathogens of humans, crops, and livestock; and phage therapy in response to antibiotic resistant bacteria.

CHAPTER 5

CONCLUSIONS

In this work, I have contributed to the understanding of the evolution and maintenance of three forms of cooperation: programmed cell death via spatial information, public goods cooperation via communication, and mutualism via species information. I have found that some factors can select for both the origin and maintenance of a cooperative trait, such as population structure in the quorum sensing system. However, others enable the maintenance of a cooperative trait – such as intermediate vertical transmission rates in mutualisms – but are not sufficient to give rise to a new cooperative trait. Combined, these results make it clear that the relationship between factors selecting for the origin or maintenance of a cooperative trait are complicated and must be assessed individually at first, before trying to understand them in combination. This dissertation begins to address this need by isolating several factors for three types of cooperation and demonstrating the effect of those factors on the origin and maintenance of the cooperative trait over evolutionary time. Specifically, I have presented three systems in which to study the origin and maintenance of cooperative traits, demonstrated the need to study interactions among environmental factors, and established the ground work for further examination of increasingly complex cooperative systems in the future.

5.1 Systems

In this work, I have made use of several digital systems and different levels of complexity, depending on the type of question that I was asking, allowing me to explore a range of cooperative dynamics. When the cooperative trait of interest required more exploration of open-ended evolutionary process, I found an open-ended system such as Avida to be most appropriate. However, when the hypotheses that I was investigating involved a more mathematically rigorous hypothesis regarding the dynamics of a cooperative trait, I was able to make use of a more targeted numerical simulation that enabled a clear answer to the validity of those hypotheses.

Each of the systems that I used can be further expanded based on the work presented here. First, the programmed cell death system in Avida can be expanded to test the remaining alternative hypotheses, such as pleitropy and changing environments, to determine if they are also sufficient to lead to programmed cell death becoming established in a population. Second, the quorum sensing system has shown that phenotypically plastic phenotypes may be more fit than fixed phenotypes and therefore can be expanded to test the fitness advantage of that phenotypic plasticity. Finally, Symbulation has many possible expansions due to the plethora of theory regarding the evolution of mutualisms; these possibilities include multi-infection, an independent symbiont stage, extreme differing generational times between partners, partner choice, partner fidelity, partner sanctions, and a third trophic level interaction.

In this work I have established the utility of these three systems and developed the most salient features for each of them. This work therefore enables future explorations to finesse the finer details of these intricate cooperative systems.

5.2 Interactions

I have also demonstrated in this work that some factors interact in non-additive ways while others completely overpower seemingly important environmental factors. In biological systems, it is often impossible to isolate these factors and vary them relative to each other in controlled ways. This limitation makes it difficult to causally determine how the interaction of factors changes the evolution of the cooperative trait. As discussed in [4], mathematical models have the opposite problem: it is relatively simple to model a single factor's influence, but combining a virtually unlimited number of those factors is unwieldy if not impossible. The three digital systems I have used here, however, present no difficulty when used to examine both isolated factors and enable those factors to interact in an evolving cooperative system. This ability makes digital systems such as mine vital to future work in the evolution of cooperation because natural systems obviously interact in infinitely complex ways, and therefore we must be able to gain experimental insight into those interactions.

Beyond interactions of factors within a single cooperative system, these cooperative systems themselves routinely interact in the natural world. Quorum sensing bacteria such as *Vibrio fischeri* exist in a mutualism with squid and other bioluminescing organisms [59]; *E. coli* execute programmed cell death when the presence of a parasitic pathogen endangers the colony [10]. The additional complications in these natural systems of even more organisms to control and care for make studying the interaction of the cooperative traits prohibitive despite their clear importance. By creating digital systems that can be combined and controlled at every level, future work will be able to indicate fruitful directions for the challenging experiments on natural systems and suggest what assumptions can be made in new analytical models.

5.3 Complexity

Finally, each of the forms of cooperation I have examined in this dissertation can interact with the evolution of complexity both positively and negatively. Because cooperative traits generally confer a benefit to related organisms, new mutants can be selected against. For example, in a population with an established programmed cell death trait that is sensitive to the slightest phenotypic difference, mutants with otherwise beneficial and more complex mutations can be prevented from invading the population. However, for increasing complexity to evolve, cooperation is required, either between genes to form a successful organism or cells to form a multicellular individual or between individuals to form colonies. Therefore, each of these digital systems can be used to determine under what conditions cooperation selects against increased complexity and under what conditions cooperation is necessary for continued evolution of complexity.

In this work I have completed the first steps in building and analyzing each of these forms of cooperation. By further expanding each system and eventually combining them, I will be able to experimentally demonstrate how the evolution and maintenance of many forms of cooperation lead to the increasingly complex and varied natural world.

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Martin Ackermann, Bärbel Stecher, Nikki E Freed, Pascal Songhet, Wolf-Dietrich Hardt, and Michael Doebeli. Self-destructive cooperation mediated by phenotypic noise. *Nature*, 454(7207):987–990, August 2008.

[2] C Adami, J Schossau, and A Hintze. Evolutionary game theory using agent-based methods. *arXiv.org*, April 2014.

[3] Christoph Adami, Jory Schossau, and Arend Hintze. Evolution and stability of altruist strategies in microbial games. *Physical Review E*, 85(1):011914, January 2012.

[4] E Akçay. Evolutionary Models of Mutualism. *Mutualism*, pages 57–76, 2015.

[5] Erol Akçay. Population structure reduces the benefits from partner choice in mutualism. *bioRxiv*, page 068445, August 2016.

[6] M Anstett. Figs and fig pollinators: evolutionary conflicts in a coevoled mutualism. *Trends in Ecology & Evolution*, 12(3):94–99, March 1997.

[7] Fredrik Bäckhed, Ruth E Ley, Justin L Sonnenburg, Daniel A Peterson, and Jeffrey I Gordon. Host-Bacterial Mutualism in the Human Intestine. *Science*, 307(5717):1915–1920, March 2005.

[8] A M L Barnard, S D Bowden, T Burr, S J Coulthurst, R E Monson, and G P C Salmond. Quorum sensing, virulence and secondary metabolite production in plant soft-rotting bacteria. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1483):1165–1183, July 2007.

[9] Benjamin E Beckmann and Philip K McKinley. Evolving quorum sensing in digital organisms. In *GECCO '09: Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 97–104. ACM Request Permissions, July 2009.

[10] Thomas W Berngruber, Sébastien Lion, and Sylvain Gandon. Evolution of suicide as a defence strategy against pathogens in a spatially structured environment. *Ecology Letters*, 16(4):446–453, January 2013.

[11] Kay D Bidle. Programmed Cell Death in Unicellular Phytoplankton. *Current Biology*, 26(13):R594–R607, July 2016.

[12] C Bordereau, A Robert, V Van Tuyen, and A Peppuy. Suicidal defensive behaviour by frontal gland dehiscence in Globitermes sulphureus Haviland soldiers (Isoptera). *Insectes Sociaux*, 44(3):289–297, August 1997.

[13] S Bowles. Group Competition, Reproductive Leveling, and the Evolution of Human Altruism. *Science*, 314(5805):1569–1572, December 2006.

[14] Michael A Brockhurst. Population Bottlenecks Promote Cooperation in Bacterial Biofilms. *PLoS ONE*, 2(7):e634, July 2007.

[15] Judith L Bronstein. Conditional outcomes in mutualistic interactions. *Trends in Ecology & Evolution*, 9(6):214–217, June 1994.

[16] Judith L Bronstein. The Costs of Mutualism. *American Zoologist*, 41(4):825–839, 2001.

[17] S P Brown and R A Johnstone. Cooperation in the dark: signalling and collective action in quorum-sensing bacteria. *Proceedings of the Royal Society B: Biological Sciences*, 268(1470):961–965, May 2001.

[18] E Bruger and C Waters. Sharing the sandbox: Evolutionary mechanisms that maintain bacterial cooperation. *F1000Research*, 4, December 2015.

[19] Eric L. Bruger and Christopher M. Waters. Bacterial quorum sensing stabilizes cooperation by optimizing growth strategies. *In preparation*.

[20] Redouan Bshary and Alexandra S Grutter. Image scoring and cooperation in a cleaner fish mutualism. *Nature*, 441(7096):975–978, June 2006.

[21] J J Bull and W R Rice. Distinguishing mechanisms for the evolution of co-operation. *Journal of theoretical biology*, 149(1):63–74, March 1991.

[22] L Chao and B R Levin. Structured habitats and the evolution of anticompetitor toxins in bacteria. *Proceedings of the National Academy of Sciences*, 78(10):6324–6328, October 1981.

[23] Charles Darwin. On the origin of species. 1859.

[24] R Dawkins. The Selfish Gene:30th Anniversary edition: 30th Anniversary edition - Richard Dawkins - Google Books. 2006.

[25] Jennifer M Dean, Mark C Mescher, and Consuelo M De Moraes. Plant–rhizobia mutualism influences aphid abundance on soybean. *Plant and Soil*, 323(1-2):187–196, February 2009.

[26] Ellen Decaestecker, Sabrina Gaba, Joost A M Raeymaekers, Robby Stoks, Liesbeth Van Kerckhoven, Dieter Ebert, and Luc De Meester. Host–parasite 'Red Queen' dynamics archived in pond sediment. *Nature*, 450(7171):870–873, November 2007.

[27] S P Diggle, A S GRIFFIN, G S Campbell, and S A WEST. Cooperation and conflict in quorum-sensing bacterial populations. *Nature*, 450(7168):411–414, 2007.

[28] Michael Doebeli and Nancy Knowlton. The evolution of interspecific mutualisms. *Proceedings of the National Academy of Sciences*, 95(15):8676–8680, July 1998.

[29] K Drescher, C D Nadell, H A Stone, and N S Wingreen. Solutions to the Public Goods Dilemma in Bacterial Biofilms. *Current Biology*, 2014.

[30] Pierre M Durand, Stuart Sym, and Richard E Michod. Programmed Cell Death and Complexity in Microbial Systems. *Current Biology*, 26(13):R587–R593, July 2016.

[31] Jay D Evans, Claude Saegerman, Chris Mullin, Eric Haubruge, Bach Kim Nguyen, Maryann Frazier, Jim Frazier, Diana Cox-Foster, Yanping Chen, Robyn Underwood, et al. Colony collapse disorder: a descriptive study. *PloS one*, 4(8):e6481, 2009.

[32] PAUL W EWALD. Transmission Modes and Evolution of the Parasitism-Mutualism Continuum. *Annals of the New York Academy of Sciences*, 503(1):295–306, July 1987.

[33] F Fiegna and G J Velicer. Competitive fates of bacterial social parasites: persistence and self-induced extinction of Myxococcus xanthus cheaters. *Proceedings of the Royal Society B: Biological Sciences*, 270(1523):1527–1534, July 2003.

[34] Francesca Fiegna and Gregory J Velicer. Exploitative and Hierarchical Antagonism in a Cooperative Bacterium. *PLOS Biology*, 3(11):e370, November 2005.

[35] Paul E M Fine. VECTORS AND VERTICAL TRANSMISSION: AN EPIDEMIOLOGIC PERSPECTIVE. *Annals of the New York Academy of Sciences*, 266(1):173–194, November 1975.

[36] Kevin R Foster and Hanna Kokko. Cheating can stabilize cooperation in mutualisms. *Proceedings of the Royal Society of London B: Biological Sciences*, 273(1598):2233–2239, September 2006.

[37] Kevin R Foster, Tom Wenseleers, and Francis L W Ratnieks. Kin selection is the key to altruism. *Trends in Ecology & Evolution*, 2006.

[38] Steven A Frank. Genetics of Mutualism: The Evolution of Altruism between Species. *Journal of theoretical biology*, 170(4):393–400, October 1994.

[39] Masaki Fukuyo, Akira Sasaki, and Ichizo Kobayashi. Success of a suicidal defense strategy against infection in a structured habitat. *Scientific Reports*, 2, January 2012.

[40] A Gardner and S A West. Greenbeards. *Evolution*, 64(1):25–38, January 2010.

[41] Heather J Goldsby, Anna Dornhaus, Benjamin Kerr, and Charles Ofria. Task-switching costs promote the evolution of division of labor and shifts in individuality. *Proceedings of the National Academy of Sciences*, 109(34):13686–13691, August 2012.

[42] W D Hamilton. The Evolution of Altruistic Behavior. *The American Naturalist*, 1963.

[43] W D Hamilton. ScienceDirect.com - Journal of Theoretical Biology - The genetical evolution of social behaviour. I. *Journal of theoretical biology*, 1964.

[44] Garrett Hardin. The tragedy of the commons. *Journal of Natural Resources Policy Research*, 1(3):243–253, 2009.

[45] Christoph Hauert, Miranda Holmes, and Michael Doebeli. Evolutionary games and population dynamics: maintenance of cooperation in public goods games. *Proceedings of the Royal Society B: Biological Sciences*, 273(1600):2565–2571, October 2006.

[46] Katy D Heath and John R Stinchcombe. EXPLAINING MUTUALISM VARIATION: A NEW EVOLUTIONARY PARADOX? *Evolution*, 68(2):309–317, February 2014.

[47] Silja Heilmann, Sandeep Krishna, and Benjamin Kerr. Why do bacteria regulate public goods by quorum sensing? How the shapes of cost and benefit functions determine the form of optimal regulation. *Frontiers in Microbiology*, 6, July 2015.

[48] E HERRE, N KNOWLTON, U MUELLER, and S REHNER. The evolution of mutualisms: exploring the paths between conflict and cooperation. *Trends in Ecology & Evolution*, 14(2):49–53, February 1999.

[49] H I Ho, S Hirose, A Kuspa, and G Shaulsky. Kin Recognition Protects Cooperators against Cheaters. *Current Biology*, 2013.

[50] Jason D Hoeksema and emilio m bruna. Pursuing the big questions about interspecific mutualism: a review of theoretical approaches. *Oecologia*, 125(3):321–330, November 2000.

[51] Michael D Jacobson, Miguel Weil, and Martin C Raff. Programmed Cell Death in Animal Development. *Cell*, 88(3):347–354, February 1997.

[52] Daniel H Janzen. Coevolution of Mutualism Between Ants and Acacias in Central America. *Evolution*, 20(3):249, September 1966.

[53] A E Johnson, E Strauss, R Pickett, C Adami, I Dworkin, and H J Goldsby. More Bang For Your Buck: Quorum-Sensing Capabilities Improve the Efficacy of Suicidal Altruism. *arXiv.org*, June 2014.

[54] Anya E Johnson, Heather J Goldsby, Sherri Goings, and Charles Ofria. The evolution of kin inclusivity levels. In *the 2014 conference*, pages 177–184, New York, New York, USA, 2014. ACM Press.

[55] Emily I Jones, Michelle E Afkhami, Erol Akçay, Judith L Bronstein, Redouan Bshary, Megan E Frederickson, Katy D Heath, Jason D Hoeksema, Joshua H Ness, M Sabrina Pankey, Stephanie S Porter, Joel L Sachs, Klara Scharnagl, and Maren L Friesen. Cheaters must prosper: reconciling theoretical and empirical perspectives on cheating in mutualism. *Ecology Letters*, 18(11):1270–1284, November 2015.

[56] Martin Kaltenpoth, Kerstin Roeser-Mueller, Sabrina Koehler, Ashley Peterson, Taras Y Nechitaylo, J William Stubblefield, Gudrun Herzner, Jon Seger, and Erhard Strohm. Partner choice and fidelity stabilize coevolution in a Cretaceous-age defensive symbiosis. *Proceedings of the National Academy of Sciences*, 111(17):6359–6364, April 2014.

[57] J F R Kerr, A H Wyllie, and A R Currie. Apoptosis: A Basic Biological Phenomenon with Wide-ranging Implications in Tissue Kinetics. *British Journal of Cancer*, 26(4):239, August 1972.

[58] T Killingback and M Doebeli. Variable investment, the continuous prisoner's dilemma, and the origin of cooperation. *Proceedings of the Royal Society of London B: Biological Sciences*, 1999.

[59] Eric J Koch, Tim Miyashiro, Margaret J McFall Ngai, and Edward G Ruby. Features governing symbiont persistence in the squid–vibrio association. *Molecular Ecology*, 23(6):1624–1634, March 2014.

[60] E V Koonin and L Aravind. Origin and evolution of eukaryotic apoptosis: the bacterial connection. *Cell death and differentiation*, 2002.

[61] R E Lenski, C Ofria, R T Pennock, and C Adami. The evolutionary origin of complex features. *Nature*, 423:139–144, May 2003.

[62] Hannah M Lewis and Alex J Dumbrell. Evolutionary games of cooperation: Insights through integration of theory and data. *Ecological Complexity*, 16:20–30, December 2013.

[63] R C Lewontin. The units of selection. *Annual Review of Ecology and Systematics*, 1970.

[64] Anagha K Matapurkar and Milind G Watve. Altruist Cheater Dynamics in Dictyostelium: Aggregated Distribution Gives Stable Oscillations. *dx.doi.org*, 1997.

[65] M B Miller and B L Bassler. Quorum sensing in bacteria. *Annual Reviews in Microbiology*, 55:165–199, 2001.

[66] Nancy A Moran, John P McCutcheon, and Atsushi Nakabachi. Genomics and Evolution of Heritable Bacterial Symbionts. *Annual Review of Genetics*, 42(1):165–190, November 2008.

[67] Aurora M Nedelcu, William W Driscoll, Pierre M Durand, Matthew D Herron, and Armin Rashidi. On The Paradigm Of Altruistic Suicide In The Unicellular World. *Evolution*, 65(1):3–20, January 2011.

[68] Ronald Noë and Peter Hammerstein. Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism and mating. *Behavioral Ecology and Sociobiology*, 35(1):1–11, 1994.

[69] Martin A Nowak, Corina E Tarnita, and Edward O Wilson. The evolution of eusociality. *Nature*, 466(7310):1057–1062, August 2010.

[70] C Ofria and C O Wilke. Avida: A Software Platform for Research in Computational Evolutionary Biology. *Artificial Life*, 10(2):191–229, March 2004.

[71] Charles Ofria, David M Bryson, and Claus O Wilke. Avida. In *Artificial Life Models in Software*, pages 3–35. Springer London, 2009.

[72] E A Ostrowski, M Katoh, G Shaulsky, and D C Queller. PLOS Biology: Kin Discrimination Increases with Genetic Distance in a Social Amoeba. *PLOS Biology*, 2008.

[73] Alexandra S Penn, Tim C R Conibear, Richard A Watson, Alex R Kraaijeveld, and Jeremy S Webb. Can Simpson's paradox explain co-operation in Pseudomonas aeruginosa biofilms? *FEMS Immunology & Medical Microbiology*, 65(2):226–235, July 2012.

[74] Julie A Perry, Dennis G Cvitkovitch, and Céline M Lévesque. Cell death in Streptococcus mutans biofilms: a link between CSP and extracellular DNA. *FEMS Microbiology Letters*, 299(2):261–266, October 2009.

[75] M Pirhonen, D Flego, R Heikinheimo, and E T Palva. A small diffusible signal molecule is responsible for the global control of virulence and exoenzyme production in the plant pathogen Erwinia carotovora. *The EMBO Journal*, 12(6):2467, June 1993.

[76] David G Rand and Martin A Nowak. Human cooperation. *Trends in Cognitive Sciences*, 17(8):413–425, August 2013.

[77] Anatol Rapoport and Albert M Chammah. *Prisoner's dilemma: A study in conflict and cooperation*, volume 165. University of Michigan press, 1965.

[78] Dominik Refardt, Tobias Bergmiller, and Rolf Kümmerli. Altruism can evolve when relatedness is low: evidence from bacteria committing suicide upon phage infection. *Proceedings of the Royal Society of London B: Biological Sciences*, 280(1759):20123035–16, May 2013.

[79] Adin Ross Gillespie, Andy Gardner, Stuart A West, and Ashleigh S Griffin. Frequency Dependence and Cooperation: Theory and a Test with Bacteria. *The American Naturalist*, 170(3):331–342, September 2007.

[80] Joel L Sachs, Ryan G Skophammer, Nidhanjali Bansal, and Jason E Stajich. Evolutionary origins and diversification of proteobacterial mutualists. *Proceedings of the Royal Society of London B: Biological Sciences*, 281(1775):20132146, January 2014.

[81] Joel L Sachs, Ryan G Skophammer, and John U Regus. Evolutionary transitions in bacterial symbiosis. *Proceedings of the National Academy of Sciences*, 108(Supplement 2):10800–10807, June 2011.

[82] J W Shapiro and P E Turner. The impact of transmission mode on the evolution of benefits provided by microbial symbionts. *Ecology and evolution*, 2014.

[83] JR Shorter and Olav Rueppell. A review on self-destructive defense behaviors in social insects. *Insectes Sociaux*, 59(1):1–10, 2012.

[84] Deborah R Smith. Behavioral Interactions between Parasites and Hosts: Host Suicide and the Evolution of Complex Life Cycles. July 1980.

[85] J E Strassmann and D C Queller. Colloquium Paper: Evolution of cooperation and control of cheating in a social microbe. *Proceedings of the National Academy of Sciences*, 108(Supplement 2):10855–10862, June 2011.

[86] P D Taylor. Altruism in viscous populations — an inclusive fitness model. *Evolutionary Ecology*, 6(4):352–356, July 1992.

[87] Alfred Tennyson Baron Tennyson. *In Memoriam AHH*. Bankside Press, 1900.

[88] Jeremy Van Cleve and Erol Akçay. Pathways to social evolution: reciprocity, relatedness, and synergy. *arXiv.org*, November 2012.

[89] G J Velicer. Social strife in the microbial world. *Trends in Microbiology*, 2003.

[90] C M Waters and B L Bassler. Quorum sensing: cell-to-cell communication in bacteria. *Annual Review of Cell and Developmental Biology*, 21:19–46, June 2005.

[91] Gijsbert D A Werner, William K Cornwell, Johannes H C Cornelissen, and E Toby Kiers. Evolutionary signals of symbiotic persistence in the legume-rhizobia mutualism. *Proceedings of the National Academy of Sciences of the United States of America*, 112(33):10262–10269, August 2015.

[92] S A WEST. Cooperation and Competition Between Relatives. *Science*, 296(5565):72–75, April 2002.

[93] S A WEST, A S GRIFFIN, and A GARDNER. Evolutionary Explanations for Cooperation. *Current Biology*, 17(16):R661–R672, 2007.

[94] S A WEST, A S GRIFFIN, and A GARDNER. Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *Journal of Evolutionary Biology*, 20(2):415–432, March 2007.

[95] Stuart A West, Stephen P Diggle, Angus Buckling, Andy Gardner, and Ashleigh S Griffin. The Social Lives of Microbes. *Annu. Rev. Ecol. Evol. Syst.*, December 2007.

[96] Michael J Wiser, Noah Ribeck, and Richard E Lenski. Long-term dynamics of adaptation in asexual populations. *Science*, 342(6164):1364–1367, 2013.

[97] Norio Yamamura. Evolution of mutualistic symbiosis: A differential equation model. *Researches on Population Ecology*, 38(2):211–218, 1996.

[98] Norio Yamamura, Masahiko Higashi, Narayan Behera, and Joe Yuichiro Wakano. Evolution of mutualism through spatial effects. *Journal of theoretical biology*, 226(4):421–428, February 2004.

[99] Ed Yong. *I Contain Multitudes: The Microbes Within Us and a Grander View of Life*. Random House, 2016.

[100] Carl Zimmer and Douglas John Emlen. *Evolution: Making sense of life*. Roberts Greenwood Village, CO, 2013.