Environmental Stress Can Favor Cooperation by Triggering an Adaptive Race Between Cooperators and Cheaters

Adam James Waite

A dissertation

submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

University of Washington

2013

Reading Committee:

Wenying Shou, Chair

Benjamin Kerr

Harmit Malik

Program Authorized to Offer Degree:

Molecular and Cellular Biology

University of Washington

Abstract

Environmental Stress Favors Cooperation by Triggering an Adaptive Race Between Cooperators and Cheaters

Adam James Waite

Chair of the Supervisory Committee:
Assistant Member Wenying Shou
Fred Hutchinson Cancer Research Center

Resolving the apparent paradox of the almost universal success of cooperation despite its vulnerability to cheating has been the focus of decades of theoretical and experimental research. Cooperators pay a cost to produce common goods that can be used by any individual, including cheaters. Cheaters consume but do not contribute fairly to the production of common goods, and are therefore more fit than cooperators. Cheaters should therefore displace cooperators, resulting in the eventual demise of cooperation known as “the tragedy of the commons.” How can cooperation persist despite the threat of cheaters? Previous research has shown that mechanisms such as spatial clustering of cooperators or cheater-sanctioning can help cooperation survive cheating. I have discovered an entirely new mechanism—adaptation to environmental stress—that favors cooperation over cheating.

I started by testing whether the known pro-cooperation mechanisms are the

only mechanisms that can stabilize cooperation. I employed a synthetic yeast cooperator-cheater system based on the release or exploitation of essential metabolites to systematically remove all known cheater control mechanisms, which would have been impossible if using natural cooperative systems. Surprisingly, I found that despite cooperators being less fit than cheaters, some cocultures rapidly purged cheaters, while in others cheaters dominated and drove cocultures and themselves extinct. Using deep sequencing and phenotypic analyses, I demonstrated the molecular mechanism responsible for this unanticipated outcome. Essentially, the nutrient-limited cooperative environment acted as a stress on both cooperators and cheaters. I found that this stress created a selective pressure which triggered an "adaptive race" between cooperators and cheaters to sample from the same pool of fitness-enhancing mutations. If cooperators sampled the best mutation, cooperators would win the race and continue to proliferate. If cheaters won the race, rapid self-extinction ensued due to the tragedy of the commons. I found evidence both in my system and in a natural cooperative bacteria system that the adaptive race can be triggered by previously-encountered environmental stresses.

Importantly, I showed stress can synergize with a spatially-structured environment to promote cooperation. I mathematically modeled the biologically realistic scenario of separate populations containing both cooperators and cheaters which interact through occasional migration. I found that cooperators are less likely to be destroyed by cheaters if environmental stresses are present. In collaboration with Ben Kerr lab at UW, we tested my model using the cooperative bacterium

Pseudomonas aeruginosa, which produces proteases to digest extracellular proteins. Cheater cells do not produce proteases and are fitter than cooperators. We showed that cooperation can be stabilized even when spatial structure is random and transient, but only when an environmental stress–antibiotics in this case–elicits an adaptive race between cooperators and cheaters. We are testing whether this phenomenon is due to a process called "niche-hiking": Isolated cooperators, by creating an environment suitable for growth, may sample more beneficial mutations in the stressful environment than isolated cheaters, which cannot grow. The deleterious cooperative phenotype can thus "hitchhike" on the beneficial mutations sampled as a result of cooperator-associated niche construction.

Finally, I found considerable non-genetic heterogeneity in response to the environmental stress of nutrient limitation, and tracked it to its molecular source: differential ability of individual cells to express the high-affinity lysine permease Lyp1p, at the cell surface. This heterogeneity could potentially provide enough variation to fuel the adaptive race when populations are too small to contain any high-fitness mutants.

In sum, previous theoretical and experimental work on cooperation and cheating has always assumed a constant environment in which adaptation is precluded. I used mathematical modeling as well as engineered and natural systems to demonstrate, for the first time, the importance of adaptation to stressful environmental changes—a biological reality as ubiquitous as spatial structure—in facilitating cooperation.

Table of Contents

Abstract i
List of Figures vi
List of Tables vii
Chapter 1. Introduction 1
1.1. The problem of cooperation 1
1.2. Ways to avert the tragedy 2
1.3. Using an engineered yeast system to understand cheater tolerance in an incipient cooperative system 5
1.4. Modeling the tragedy in a metapopulation 7
1.5. Causes and consequences of non-genetic variability 12
Chapter 2. Environmental stress selects for mutants that allow cooperators to purge cheaters stochastically 16
2.1. Cheaters are fitter than cooperators 16
2.2. Cheaters are purged from co-cultures stochastically 17
2.3. Extremely fit mutations drive stochastic cooperator dominance 18
2.4. The adaptive race is fueled by mutations in a small set of genes involved in nutrient transport 22
2.5. Improved cooperators can stochastically defeat newly-arising cheaters 25
2.6. Materials and methods 26
2.7. Figures 33
Chapter 3. The adaptive race can synergize with sources of abiotic variation to promote cooperation in a metapopulation 46
3.1. Experimental observations 48
3.2. Model assumptions 50
3.3. Initial conditions 53
3.4. High rates of migration and death lead to cooperator domination 54
3.5. The adaptive race and empty space help cooperators beat cheaters 56
3.6. Testing the model with an experimental metapopulation 58

3.7. Materials and Methods.....60
3.8. Figures.....63
Chapter 4. Environmental stress induces non-heritable variation in otherwise clonal populations.....73
4.1. Growth on limiting lysine reveals non-genetic heterogeneity.....73
4.2. Population growth rate is dependent upon nutrient concentration.....74
4.3. Limiting lysine increases colony size heterogeneity.....77
4.4. Heritable decrease in growth rate observed in respiratory-proficient and respiratory-deficient cells.....77
4.5. Heritable and non-heritable decrease in death rate observed after prolonged lysine starvation.....78
4.6. Individual cells behave differently in limiting lysine.....81
4.7. Heterogeneous expression of high-affinity lysine permease is responsible for population heterogeneity.....83
4.8. Diploid prototrophs show growth heterogeneity during glucose limitation.....87
4.9. Materials and Methods.....92
4.10. Figures.....94
Chapter 5. Conclusions and future directions.....103
5.1. Adaptation is a variance-generating mechanism that allows stochastic cooperator dominance 103
5.2. The variance generated by the adaptive race can be maintained by migration 108
5.3. Can genetic “niche-hiking” explain the success of cooperation in an experimental metapopulation?.....111
5.4. The molecular origins of non-genetic heterogeneity.....113
5.5. Final thoughts: Variation and the success of cooperation.....118
Appendix A. The Monod “constant” is not a constant.....123
Appendix B. Location of unpublished data used to generate figures.....126
References 127

List of Figures

Figure 2.1: A yeast model of cooperation and cheating. 34
Figure 2.2: Stochastic cheater outcome in initially identical cooperator-cheater cocultures.....36
Figure 2.3: Rapid divergence in the growth of cooperators and cheaters in cooperator-dominated and cheater-dominated co-cultures. 38
Figure 2.4: Evolved cooperators and cheaters are superior to ancestors.....40
Figure 2.5: Evolved cooperators and cheaters show a tradeoff between the ability to grow at low lysine concentrations and the maximum growth rate achieved at high lysine concentrations.....42
Figure 2.6: Stochastic cheater purging observed in co-cultures composed of cooperators and cheaters previously exposed to the cooperative environment. 43
Figure 3.1: The fitness of cooperators and cheaters decreases with increasing cheater frequency.....63
Figure 3.2: Evolved strains die more slowly than their corresponding ancestors.....64
Figure 3.3: Summary of fitness relationships used in the model.....65
Figure 3.4: The adaptive race increases cooperator-dominated parameter space.....67
Figure 3.5: Empty space synergizes with the adaptive race to help cooperation: global migration.....69
Figure 3.6: Empty space synergizes with the adaptive race to help cooperation: local migration.....71
Figure 3.7: Environmental stress helps cooperators defeat cheaters in an experimental metapopulation.....72
Figure 4.1: Growth on limiting lysine is not a heritable phenotype.....94
Figure 4.2: Growth rate is nutrient dependent.....95
Figure 4.3: Size heterogeneity is not solely due to increased petite frequency.....96
Figure 4.4: Growth rates of previously starved colonies.....97
Figure 4.5: Death due to lysine starvation is density-independent and multiphasic.....98

Figure 4.6: Differing fates of cells in identical environments. 99
Figure 4.7: The ability to grow in limiting lysine is dependent upon Lyp1p surface expression......101
Figure 4.8: Heterogeneity observed in wild diploid prototrophic strains when limited for glucose.....102

List of Tables

Table 2.1: Constructed strains used in this study......29
Table 2.2: Mutations potentially sufficient for the large observed fitness advantage in the nutrient-limited cooperative environment......44
Table 2.3: Strains with chromosomal duplications.....45
Table 3.1: Summary of variables used in describing the model.....51
Table 3.2: Summary of values used in the model......54
Table 3.3: Possible outcomes when single cooperators and cheaters migrate to new locations that are either empty or contain one individual of the indicated type. 56
Table 4.1: Estimates of parameters of the Monod Equation......76
Table 4.2: Some published values of the Km for Lyp1p......76
Table 4.3: Summary of individual cell responses to environments with different lysine concentrations.....83
Cells of the "residual growth" class are not included......83
Table B.1: Folders containing unpublished data used to generate figures for this dissertation.....126

Acknowledgments

First and foremost, I would like to thank my advisor, Wenying Shou, for her guidance and support. Her dedication, curiosity, willingness to learn new things, and scientific rigor are just four qualities among many that I hope I have absorbed to some degree. Wenying does not consider herself a patient person. However, this quality was required in abundance while mentoring me to think quantitatively and mathematically about biological problems.

I would also like to acknowledge the other members of the Shou Lab, especially Justin Burton, Babak Momeni, Chi-Chun Chen, and David Skelding, for their willingness to discuss, at length, any and all aspects of my various projects. Their intelligence and critical insight helped all of my projects immensely.

My committee members, Dan Gottschling, Ben Kerr, Harmit Malik, and Chip Asbury, for their guidance and support. I would like to specifically mention Dan's generosity in “adopting” me as an unofficial member of the lab, and including me in his group meetings; Harmit's help in guiding me through the difficult (but exciting) process of identifying and deciding upon a post-doctoral position; and Ben's limitless energy, enthusiasm and extreme generosity with his very limited time.

The Hutch has an amazing shared resources facility. The members of the Flow Cytometry Department, helped me immensely with my experiments and taught me a lot about how flow cytometry works. Andy Marty and Ryan Basom were critical in enabling the high-throughput sequencing portion of my dissertation. Sarah Holte was an amazing statistics mentor and helped me through some difficult and non-standard ways of analyzing my data. I thank the Scientific Computing Division for their patience in helping me learn how to set up and run simulations on our cluster...even though I had a tendency to use it in ways they probably did not intend. I would also like to give special thanks Luna Yu, whose general computing expertise was invaluable, both in helping me set up my computers, but also in teaching me about them. Her positive attitude and love of her job were, and will continue to be, an inspiration.

The administrative support in both the Molecular and Cellular Biology Program and at the Hutch was fantastic, and was critical for helping me navigate the non-scientific portion of obtaining my PhD. I would especially like to thank MaryEllin Robinson at UW and Michele Karantsavelos at the Hutch for their help, which consistently was well above and beyond my expectations.

On a less personal note, I would like to thank the open source community for their tireless dedication to the production and maintenance of free and open software. Due to the efforts of this community of enthusiasts, I was able to do all of my dissertation research, analysis, and reporting using open source software. I will continue to do so in the future.

On a more personal note, I would like to express my gratitude towards the amazing group of friends I have met during my time in Seattle. In addition to

scientists, I have met a diverse array of musicians, artists, and outdoor enthusiasts who have encouraged my participation in non-scientific endeavors which I consider crucial for maintaining a creative, well-balanced life.

Finally, I would like to thank my family, whose love and support made this all possible. While they have not always understood exactly what it is I am doing—or why I am doing it—they have always provided a warm, loving environment that has served as a foundation from which I can draw strength and inspiration. I especially thank my sister, Alexis, for always being available to talk me through the occasional tough time, and her visits which were always a blast. I could not ask for better parents, who have always understood and never questioned my desire to pursue a career in basic research, despite its inherent difficulties. With all they have done for me, I hope that this dissertation can serve as a token of my appreciation; a symbol of the work I was able to achieve because of them.

To Mom, Dad, and Alexis

Chapter 1. Introduction

1.1. The problem of cooperation

The cooperative act of paying a cost to produce a publicly-available good is a common biological phenomenon. Human volunteers contribute their time to build Wikipedia, which can be used by anyone with Internet access. Throughout the animal kingdom, alarm calls are produced by individuals to warn others of danger, even though producing the call makes the caller more conspicuous (Hollén and Radford, 2009). Microbes excrete a plethora of costly compounds which can be used by the producers and their neighboring cells to acquire hard-to-obtain nutrients, access favorable environments, or improve antibiotic resistance (Lee et al., 2010; West et al., 2007). In biological systems, publicly-available goods are generally "common goods," because consumption by one individual reduces their availability to others. "Cheaters" utilize the common good without paying a cost to produce it. Thus, because the common good is equally accessible to all members of a population, cheaters, introduced through migration or mutation, will be more fit than cooperators, increase in frequency, and eventually exhaust the common good, leading to the "tragedy of the commons" (Hardin, 1968). For instance, while cooperative viruses produce diffusible shared proteins required for viral reproduction, selfish viruses synthesize less but sequester more of these proteins and thereby displace cooperative viruses, lowering overall infectivity (Turner and Chao, 1999). Cancers, a leading cause of death globally (World Health

Organization, 2008), cheat by exploiting the common good produced by normal cells that cooperate to form a functional human body.

Despite exploitation of common goods by naturally-arising cheaters, cooperation persists (Dobata and Tsuji, 2009; Naumov et al., 1996; Schaber et al., 2004; Strassmann et al., 2000; Vos and Velicer, 2009). How does cooperation survive cheating? The simultaneous ubiquity and vulnerability of cooperation has puzzled evolutionary biologists since Darwin first recognized it as a challenge to the validity of his theory of evolution by natural selection (Darwin, 1859), and a large amount of effort has been directed to understanding this apparent paradox. This research has been fruitful, and in the next section I discuss the major mechanisms allowing the persistence of cooperation.

1.2. Ways to avert the tragedy

Frequently, through unexpected genetic or physical processes, what initially appear to be "cheaters" or "common goods" are not, and thus the tragedy of the commons does not apply. For instance, a gene required for cooperation can have pleiotropic effects, such that a cell defective in paying the cost of cooperation is also incapable of enjoying the cooperative benefit. In this case, "cheaters" will end up suffering a net fitness cost. This situation has been found to occur in the social amoeba *Dichtyostelium discodium*. *D. discodium* responds to starvation by aggregating and forming a fruiting body. During fruiting body formation, some cells form a self-sacrificing, non-reproductive stalk which lifts other cells that differentiate into

reproductive spores. The gene encoding the receptor necessary for differentiation into stalk cells is also necessary for proper spore formation, and thus cheaters trying to avoid the stalk fate cannot become spores (Foster et al., 2004). Another possibility is that the apparently "common good" is partially privatized by its producer. For instance, the budding yeast *Saccharomyces cerevisiae* secrets invertase to hydrolyze the disaccharide sucrose, which it cannot import, into glucose and fructose, which can be imported and metabolized. These monosaccharides were initially thought to be strictly common goods (Greig and Travisano, 2004), although it was later found that ~1% are retained by the producing cell (Gore et al., 2009). Even such a seemingly insignificant level of privatization can allow cooperators to invade a population of cheaters (Doebeli and Hauert, 2005; Gore et al., 2009). This explains the initially puzzling coexistence of invertase-producing cooperative cells with non-producing cheating cells in wild populations (Naumov et al., 1996). However, the benefits of privatization can only be realized when all cooperators produce and consume the same common good (homotypic cooperation). In mutualism, privatization is pointless because each partner requires a common good produced only by the other partner.

Even if cheaters have a net fitness advantage over cooperators who produce true common goods, several mechanisms can avert the tragedy of the commons. First, when individuals interact through the production and/or consumption of inexpensive common goods in randomly-formed groups that assemble and disassemble cyclically, as long as an increase in the availability of the common good leads to a less-than-proportional

increase in the fitness of its consumers, a stable equilibrium between cooperators and cheaters is expected (Archetti and Scheuring, 2011; Foster, 2004). This "diminishing return" of the common good (Foster, 2004; Gore et al., 2009) can account for the surprising observation that, in the yeast invertase system, maximum group size is attained with a mixture of cheaters and cooperators: cooperators produce more invertase than they can use, and cheaters convert this excess benefit into additional biomass (MacLean et al. 2010). The cost-to-benefit ratio can be kept low if the common good is produced facultatively (i.e., only when needed), which is the case for most organisms, or if the durability of common goods is high, as found in siderophore production in *Pseudomonas aeruginosa* (Kümmerli and Brown, 2010; Kümmerli et al., 2009). Second, for cooperation based on scarce common goods, mechanisms of "positive assortment" that increase the frequency of interactions between cooperators (Fletcher and Doebeli, 2009) can facilitate the persistence of cooperation. Positive assortment can involve specifically directing benefits to other cooperators and excluding or punishing cheaters based on recognition or previous experience (Axelrod and Hamilton, 1981; Hamilton, 1964; Trivers, 1971). This can occur even in organisms lacking nervous systems, as microbes can achieve "recognition" through cell adhesion and chemical communication (Strassmann et al., 2011) and legumes "reward" and "punish" beneficial and cheating rhizobia, respectively (Heath and Tiffin, 2009; Kiers et al., 2003), Another mechanism of positive assortment is "population viscosity," brought about by limited dispersal in spatially-structured environments, which keeps cooperators clustered with

their relatives in homotypic cooperation (Chao and Levin, 1981; Hamilton, 1964; Maynard Smith, 1964; Nowak and May, 1992), or with their partners in heterotypic cooperation (Harcombe, 2010). Thus, natural cooperative systems, whether homotypic or heterotypic, utilize many different mechanisms to mitigate the tragedy of the commons, allowing cooperation via common goods to be a successful evolutionary strategy.

1.3. Using an engineered yeast system to understand cheater tolerance in an incipient cooperative system

Existing cooperative systems may have evolved for millions of years, and can now deploy mechanisms discussed in the previous section—such as pleiotropy, facultative production of common goods, and cheater recognition—to prevent destruction by cheaters. But what about cheater control at the origins of these systems? A spatially-structured environment can stabilize cooperation against cheating, but under certain circumstances it can hinder cooperation (Hansen et al., 2007; Hauert and Doebeli, 2004; Verbruggen et al., 2012) through enhancing competition among cooperators (Griffin et al., 2004; Queller, 1992; West et al., 2002). And what about motile species, which may behave as if they were locally well-mixed even if they live in a spatially-structured environment. To search for novel cheater-control mechanisms that might stabilize nascent cooperative systems, I extended a synthetically-engineered cooperative system (Shou et al., 2007) to include cheating (Figure 2.1A). In this system, which is based on the yeast *Saccharomyces cerevisiae*, cooperation occurs between the

red-fluorescent $R_{\rightarrow A}^{\leftarrow L}$ strain, which requires lysine ($\leftarrow L$) and provides adenine ($\rightarrow A$), and the yellow-fluorescent $Y_{\rightarrow L}^{\leftarrow A}$ strain, which requires adenine ($\leftarrow A$) and provides lysine ($\rightarrow L$). The presence of both $R_{\rightarrow A}^{\leftarrow L}$ and $Y_{\rightarrow L}^{\leftarrow A}$ are necessary for growth in minimal media lacking adenine and lysine (SD) (Shou et al., 2007). The cheater is a cyan-fluorescent $C^{\leftarrow L}$ strain which requires lysine but does not provide any nutrients.

This system allowed us to systematically remove the mechanisms currently thought to avert the tragedy of the commons. Engineering the cooperative interactions through mutations that constitutively overproduce metabolites eliminates the possibility of pleiotropy or facultative production. Using a two-partner heterotypic cooperative system renders privatization of common goods futile, as neither strain can directly utilize the common good it produces. After an initial period of asymmetric nutrient release and cell death (Shou et al., 2007), our cooperative co-cultures approach a stable doubling time (~12 hrs) that is much slower than those of the corresponding monocultures supplemented with the necessary nutrient (~2 hrs). Thus, common goods are very limited. Additionally, since cooperation is based on metabolic manipulations not found in the ancestor strain, it is unlikely that the cooperating partners could recognize one another or exclude the cheater. As all strains are derived from S288C, which does not flocculate (Liu et al., 1996), growth in well-mixed liquid culture prevents spatial clustering.

In Chapter 2 I describe how I used this simple, well-defined biological model

(Momeni et al., 2011) free of known mechanisms of positive assortment to discover a new mechanism of cheater control. Essentially, the stress imposed by a novel environment—in this case, the nutrient-limited environment requiring cooperation—selects for cooperator and cheater mutants that have much higher fitness than their corresponding ancestor types. Because this fitness increase is much greater than the cost of cooperation, the cooperative phenotype can "hitchhike" its way to high frequency, giving cooperators a chance, equivalent to their initial frequency in the population, to purge cheaters. Of course, a cheater can obtain the highest fitness mutation and destroy cooperation but, since cooperation is essential for survival, these high-fitness cheaters ensure their own self-destruction. In a sense, cooperators and cheaters "race" to adapt to this novel environment, and so I refer to this mechanism as the "adaptive race."

1.4. Modeling the tragedy in a metapopulation

Despite the initially random nature of the adaptive race, the fact that successful cooperators continue to grow while successful cheaters ensure their own demise appears to favor cooperation over cheating. This view inherently assumes that successful cooperators and cheaters are isolated from other populations long enough to reach high frequency. It also assumes that the most successful populations would be able to expand into previously unoccupied territory. Both of these assumptions implicitly rely on the concept of a metapopulation, or a "population of populations." In a metapopulation, individual populations can interact through occasional migration, expand into empty territory, and be removed due to extinction.

Like mutation and selection, migration changes gene frequencies of populations, and therefore contributes to their evolution. In fact, if only one population is considered, the effect of the arrival of migrants is completely analogous to mutants arising in this population (Wright, 1931). Unlike mutation alone, which can only alter gene frequencies within a population, migration can alter gene frequencies between populations. If the number of migrants produced by a population is proportional to its size, which seems like a biologically realistic assumption to make, then larger populations will generate more migrants. These migrants will then alter the gene frequencies of neighboring populations and, given enough time, the gene frequencies of neighboring populations will resemble the gene frequencies of the most successful population.

Situations where selection acts in the same direction on both individuals within a population and between populations are easy to understand. More interesting are situations where the selective forces on individuals within populations are opposed to the selective forces acting between populations. This occurs with common goods cooperation since individual cheaters are more fit than individual cooperators, but a pure population of cooperators is more fit (e.g., more productive or more persistent) than a pure population of cheaters. These dynamics are typically modeled as an *N*-person prisoner's dilemma (NPD), in which cooperators pay a cost to produce a benefit which is shared among all members of the group. Cheaters benefit from the cooperative act without paying a cost. Because payoffs are proportional to fitness, cheaters increase in

frequency in groups that contain cooperators.

As with the problem of cooperation at the individual level, solutions in metapopulations involve generating positive assortment, so that interactions between cooperators are maximized. Both theoretical (Hauert et al., 2006; Killingback et al., 2006) and experimental (Kerr et al., 2006) work has shown that cooperators are doomed to extinction when migration rates are high and migrants can visit any population in the metapopulation ("global" migration). On the other hand, if migration rates are relatively low and migration is restricted to nearest neighbors ("local" migration), it is possible for cooperation to persist (Kerr et al., 2006). It is also possible for cooperation to survive if the migration rate is itself a trait that is subject to selection (Parvinen, 2013; Suzuki and Kimura, 2011). Depending on the model assumptions, either cooperators or cheaters can evolve the higher migration rate. However, it should be noted that an NPD represents an extreme case where the payoff is a linear function of the number of cooperators in a group. If the payoff function is non-linear, such that the payoff saturates after enough members have cooperated, and cooperation occurs with a probability, then cooperation is stable when groups are transiently and randomly formed (Archetti and Scheuring, 2011; Pacheco et al., 2009).

In models where the migration rate is a function of the quality of the local environment, such that individuals in cheater-dominated populations have a higher probability of migrating (Ichinose and Arita, 2008; Pepper, 2007; Pepper and Smuts, 2002), cooperation can survive when migration is global. The ability of this

“eco-evolutionary feedback” to generate positive assortment and promote cooperation recently received experimental validation in a yeast metapopulation model (Datta et al., 2013). In this model, which uses invertase as the common good, cooperators have an equilibrium frequency of 15% when mixed with cheaters because of partial privatization of the public good. As populations grow, they can expand into neighboring wells by migration. Cooperators were able to achieve a higher frequency (~45%) in this “expanding front” than in equilibrium mixtures with cheaters (~15%). Intriguingly, this was only achieved when minimum population densities were relatively high.

Population size also plays a crucial role in the overall success of cooperation. Conditions that lead to small population sizes increase the variance in cooperator frequency, and can lead to positive assortment purely by chance. This can be achieved by limited dispersal, even if dispersal is global (Killingback et al., 2006). High variance can also be achieved through extinction of populations. Extinction provides empty space into which small numbers of migrants from the most successful populations (which were more likely to be cooperators) can expand. Extinction can occur with cheater-dominated populations, which do not produce enough common good to support the group. In some models, population extinction occurs at random, regardless of cheater frequency (Parvinen, 2011). Even in intrademic models, in which groups form randomly and transiently to reproduce before dispersal (Wade, 1978), cooperation can survive due to the presence of empty space. For example, in the model of Hauert et al. (2006), birth rates of individuals were proportional to the amount of empty space available and death

was explicitly modeled. Cheater-dominated groups died rapidly, reducing population density and the average group size until some groups were mostly composed of cooperators. While much of the parameter space led to oscillations in cooperator frequency, cooperators could fix if the benefit to cost ratio and the death rate were high enough.

All of these mechanisms result in Simpson's paradox, a statistical effect related to second-order interactions in contingency tables (Blyth, 1972; Simpson, 1951). In the context of cooperation and cheating in metapopulations, Simpson's paradox allows cooperators to increase in frequency globally even though cheater frequency increases within each population. This is due to the high covariance between cooperator frequency and population growth rate (Chuang et al., 2009; Hauert et al., 2002). In other words, groups dominated by cooperators expand in size much faster than groups dominated by cheaters.

As will become clear in Chapter 2, the adaptive race creates extreme inter-population variance: either cooperators or cheaters reach high frequency in each population. Is the additional variance due to migration and population extinction enough to maintain this variance and allow cooperators to increase in frequency globally? In Chapter 3, I present a mathematical model of the adaptive race in a metapopulation, and show that the adaptive race promotes the global increase in cooperator frequency over a wide range of migration rates and cheater-induced death rates. I then describe some experiments performed by members of the Kerr lab at the University of Washington that

test the adaptive race in an experimental metapopulation using the cooperative bacterium *Pseudomonas aeruginosa*. They were able to show that, even in an intrademic model of population formation, the variance induced by the adaptive race was sufficient to allow cooperators to achieve and maintain a high frequency.

1.5. Causes and consequences of non-genetic variability

In Chapter 2, I demonstrate how heritable variation due to mutation allows the adaptive race to occur. Theoretically, however, selection can act on *any* variation in a population—genetic or otherwise. And, in microbes, this "otherwise" certainly exists. Ingenious experimenters observed differences between individual bacterial growth rate (Hughes, 1955), enzyme production (Benzer, 1953; Novick and Weiner, 1957), motility behavior (Spudich and Koshland, 1976), and viral burst size (Delbrück, 1945), long before the modern techniques of microfluidcs and flow cytometry made analysis of single-cell behavior a relatively simple task. These experiments established that individual bacteria within a clonal population are not perfect reproductions of one another. However, from an evolutionary point of view, variation is not enough. Beneficial traits must be stably heritable to appreciably improve the mean fitness of a population over time and lead to sustained phenotypic change, and so less permanent traits are expected to have less impact on the evolution of a population than more permanent ones. Understandably, then, most research in evolutionary biology has focused on genetic mutation as the sole source of variation in a population.

On the other hand, recent work on microbes has revealed that the amount of non-genetic variation in a given phenotype can itself be a genetically heritable trait which may be tuned by evolution (Bishop et al., 2007; Colman-Lerner et al., 2005; Davidson and Surette, 2008; Gordon et al., 2009; Kaufmann et al., 2007; Levy et al., 2012; Raser and O'Shea, 2004). In unchanging environments, or environments that experience temporary fluctuations, fitness is maximized by minimizing phenotypic variation. At a cellular level, this is achieved regulatory feedback loops (Bennett et al., 2008; Yu et al., 2008). However, in environments that change drastically or for extended periods of time, non-genetic variation may be useful, especially when this change occurs randomly but with a probability greater than zero. When "expecting the unexpected," a population can "hedge its bet" that the environment will remain constant by setting aside a portion of its population that is less fit now, but will survive a catastrophic environmental change later (Balaban et al., 2004; Beaumont et al., 2009; Ni et al., 2012). For this to occur, the probability of environmental change must be sufficiently high, and the size of the "intentionally" less fit portion of the population must be sufficiently small. Otherwise, mutations that reduce variation would rapidly sweep through the population.

While it may be impossible to determine how frequently environments meeting the criteria to select for increased non-genetic heterogeneity are met in nature, experimental evidence shows that it is at least possible. For example, naturally-occurring populations of *E. coli* undergo phenotypic switching between a

fast-growing, antibiotically sensitive type and a slow-growing "persister" type that can survive antibiotic exposure (Balaban et al., 2004). Beaumont et al. (2009) were able to experimentally evolve a phenotypic switch in *Pseudomona fluorescence* by imposing an alternating selection scheme. If fluctuating environmental pressures are consistent enough, the trait might become permanent through "genetic assimilation." First proposed by Waddington (1942), he later demonstrated it experimentally in the laboratory using *Drosophila*. He noticed that exposing flies to environmental stress (heat shock or ether) while they were pupae increased the frequency of appearance of different phenotypic abnormalities. In the absence of these stresses, even abnormal flies would have normal offspring. After several generations of selectively breeding these abnormal flies, he was able to evolve lines that stably produced the same abnormalities in the absence of the environmental stress (Waddington, 1953, 1956, 1957).

One result of genetic assimilation would be the reproducible appearance of a seemingly predictive response to environmental cues. Evidence that this has evolved outside the laboratory has recently been found in both bacteria and yeast. Upon entering a mammalian host, a series of reproducible environmental changes take place. The increased temperature (to 37 °C) encountered in the mouth precedes the anoxic conditions of the digestive tract. Once in the digestive tract, lactose is encountered before maltose. Correspondingly, when exposed to 37 °C, *E. coli* up-regulates genes required for growth in low oxygen conditions—even if maintained in a fully oxic environment (Tagkopoulos et al., 2008). Similarly, exposure to lactose induce genes

required for maltose metabolism, but growth on maltose has no effect on genes related to lactose metabolism (Mitchell et al., 2009). Importantly, in both cases the investigators could break these genetic correlation structures thorough experimental evolution, demonstrating that these correlations were adaptive and not the result of some kind of biochemical necessity. In yeast, heat stress due to fermentation precedes ethanol stress, which in turn precedes oxidative stress encountered during respiration. It was found that pre-exposure to heat or ethanol increased survival upon subsequent oxidative stress, but oxidative stress offered no such protection (Mitchell et al., 2009). Furthermore, pre-exposure to other, unrelated stresses had no effect on surviving oxidative stress.

It is clear that there is a link between environmental stress and the induction of non-genetic heterogeneity in a population (Badyaev, 2005; Rutherford and Lindquist, 1998; Zhuravel et al., 2010). As mentioned earlier, the cooperative environment is extremely nutrient-limited, which should impose a considerable amount of stress on cooperators and cheaters alike. In Chapter 4, I demonstrate considerable heterogeneity in terms of lag time, growth rate, and death rate in otherwise clonal populations of yeast using a combination of population and single-cell based techniques. Importantly, this phenotypic heterogeneity could provide the variation necessary to trigger an adaptive race when populations are too small to generate high-fitness genetic mutants.

Chapter 2. Environmental stress selects for mutants that allow cooperators to purge cheaters stochastically

2.1. Cheaters are fitter than cooperators

Many biosynthetic pathways are regulated by end-product feedback-inhibition (Armitt and Woods, 1970; Blickling and Knäblein, 1997), which suggests that metabolite overproduction carries an evolutionarily significant cost. To quantify the cost of adenine overproduction in our cooperative yeast system, I competed $R^{\leftarrow L}_{\rightarrow A}$ with $C^{\leftarrow L}$ in SD supplemented with excess (164 μM) lysine, and found that the latter carried a fitness advantage of 1.8% (95% CI: 1.6% – 1.9%; Figure 2.1B). In the absence of lysine, which occurs during the prolonged delay in lysine release by $Y^{\leftarrow A}_{\rightarrow L}$ (Shou et al., 2007), $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ died at similar rates (95% CI of the difference in death rate: -0.003–0.0003 hr^{-1}; Figure 2.1C). In both conditions, the fitness difference between expressing CFP and dsRed is small. Given the overall fitness advantage of $C^{\leftarrow L}$ over $R^{\leftarrow L}_{\rightarrow A}$, the paucity and futility of privatization of common goods, and the lack of genetic mechanisms for cheater control, I predicted that in a well-mixed environment, $C^{\leftarrow L}$ would increase in frequency and eventually destroy the cooperative system.

2.2. Cheaters are purged from co-cultures stochastically

I experimentally tested the prediction of deterministic cheater dominance by mixing $R^{\leftarrow L}_{\rightarrow A}$, $Y^{\leftarrow A}_{\rightarrow L}$, and $C^{\leftarrow L}$ at a ratio of 1:1:1 in SD (Figure 2.2). This master mix was split into replicate co-cultures which were monitored for growth and diluted to ensure that nutrients other than adenine and lysine were never limiting (Text S1). Surprisingly, within 50 generations I observed a bifurcation in growth rates (Figure 2.2B): some co-cultures were growing slowly or not at all (Figure 2.2A, gray triangles), while others continued to grow (Figure 2.2A, orange triangles) at rates very similar to cheater-free co-cultures (Figure 2.2A, black circles). At ~450 hours, I quantified the frequency of each population in each co-culture using flow cytometry (Figure 2.2C). The slow-growing co-cultures (gray) contained mostly dead or dying cells that had lost membrane integrity and therefore reacted with the nucleic acid dye TO-PRO-3. The remaining live cells were predominantly $C^{\leftarrow L}$ ("cheater-dominated"). Furthermore, cheater takeover took much less time than predicted: in less than 40 generations the average $C^{\leftarrow L}$: $R^{\leftarrow L}_{\rightarrow A}$ ratio was 280:1, as opposed to the 2:1 ratio predicted by the 1.8% fitness advantage of $C^{\leftarrow L}$ (Figure 2.1B). Even more surprising were the fast-growing co-cultures (orange). These co-cultures contained mostly live $R^{\leftarrow L}_{\rightarrow A}$ and $Y^{\leftarrow A}_{\rightarrow L}$ cells ("cooperator-dominated"), with an average $R^{\leftarrow L}_{\rightarrow A}$: $C^{\leftarrow L}$ ratio of 16:1. I confirmed in subsequent experiments with a $C^{\leftarrow L}$ strain differentially marked with drug resistance that $C^{\leftarrow L}$ was driven extinct in cooperator-dominated co-cultures.

2.3. Extremely fit mutations drive stochastic cooperator dominance

To understand what had caused rapid divergence in population growth and stochastic cheater outcomes, I investigated the detailed population dynamics of these co-cultures by frequently sampling replicate co-cultures using flow cytometry. All $R_{\rightarrow A}^{\leftarrow L}$ and $C^{\leftarrow L}$ populations, regardless of whether they eventually became cooperator- or cheater-dominated (Figure 2.3A, solid or dashed lines respectively), were nearly identical for the first ~60 hours of growth, suggesting that the vast majority of cells were behaving identically and, most likely, ancestrally. Afterward, $R_{\rightarrow A}^{\leftarrow L}$ and $C^{\leftarrow L}$ diverged rapidly (Figure 2.3A). On the other hand, $Y_{\rightarrow L}^{\leftarrow A}$ from cooperator-dominated and cheater-dominated co-cultures behaved identically until after the divergence between $R_{\rightarrow A}^{\leftarrow L}$ and $C^{\leftarrow L}$ (>80 hours, Figure 2.3A), suggesting that whatever was occurring in the $R_{\rightarrow A}^{\leftarrow L}$ and $C^{\leftarrow L}$ populations preceded any changes in the $Y_{\rightarrow L}^{\leftarrow A}$ population.

One possible explanation for the rapid population divergence was the presence of variants with large fitness advantages relative to their ancestor. Consider the cooperator-dominated co-culture (Figure 2.3A, solid lines). If the $R_{\rightarrow A}^{\leftarrow L}$ population obtained the most fit variant ($R!$), then by the time of population divergence (~60 hr), $R!$ must have proliferated sufficiently to visibly influence the growth kinetics of the red-fluorescent population. To estimate the fitness advantage of $R!$ over non- $R!$, I would need to know their relative abundance over time. I could not distinguish the two

sub-populations directly from flow cytometry. However, because the $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ populations initially behaved similarly, and I expected both to be influenced by the presence of $R!$ in a similar way, I assumed that the observed behavior of $C^{\leftarrow L}$ was similar to that of the non- $R!$ portion of the red-fluorescent population. Thus, the population size of $R!$ could be approximated by the difference between the red-fluorescent and cyan-fluorescent populations. Assuming that all $R!$ cells were descendants of a single cell present at the beginning of the experiment, this cell must have doubled on average every ~3.4 hours to achieve its estimated abundance at 80 hours, a large improvement from the ancestral average doubling time of ~18.5 hours (see Methods). The same argument can be applied to the cheater-dominated co-culture, which suggests that cooperation was destroyed by similar, highly adaptive variants that arose in the $C^{\leftarrow L}$ population. These variants must not have been rare (>1 in 10^6) because otherwise in many co-cultures, $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ would have failed to sample any variant and population divergence would have been slow.

$R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ thus appeared to be engaged in an "adaptive race" to obtain the most fit variant. If the same pool of variants was sampled by both populations, then the final ratio of cooperator-dominated : cheater-dominated co-cultures should be determined by the initial ratio of $R^{\leftarrow L}_{\rightarrow A}$: $C^{\leftarrow L}$. Furthermore, in our system, these two ratios should be approximately equal because the fitness advantage of $C^{\leftarrow L}$ over $R^{\leftarrow L}_{\rightarrow A}$ is insignificant compared to the fitness gain of a successful variant. I tested this idea by

setting up co-cultures at different initial ratios of $R^{\leftarrow L}_{\rightarrow A}$: $C^{\leftarrow L}$ (Figure 2.3B; $R^{\leftarrow L}_{\rightarrow A}$: $Y^{\leftarrow A}_{\rightarrow L}$ was always 1:1). The slope relating the two ratios was 1.0 (95% CI: 0.7 – 1.3). Thus, consistent with our hypothesis that $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ were sampling from the same pool of variants, the frequency of being cooperator-dominated was determined by the initial frequency of $R^{\leftarrow L}_{\rightarrow A}$.

I next examined whether the improved fitness of these variants in the context of cooperation was a heritable phenotype. I isolated a $Y^{\leftarrow A}_{\rightarrow L}$, $R^{\leftarrow L}_{\rightarrow A}$, and $C^{\leftarrow L}$ clone from a cooperator-dominated co-culture, grew them in monoculture in SD supplemented with adenine or lysine, washed them free of supplements, and mixed them 1:1:1. Like the parental co-culture, all six co-cultures became cooperator-dominated (Figure 2.4A, ii), which was significantly more than what was observed for 1:1:1 all-ancestor co-cultures (Figure 2.4A, i; 6/6 vs 31/64 cooperator-dominated, Fisher's Exact Test P<0.03). The isolated $Y^{\leftarrow A}_{\rightarrow L}$ clone did not contribute to cooperator-dominance (Figure 2.4A iii; 3/6 vs 31/64 cooperator-dominated, Fisher's Exact Test P>0.9). This is consistent with the idea that, short of partner-specific recognition, any changes in $Y^{\leftarrow A}_{\rightarrow L}$ would affect $C^{\leftarrow L}$ and $R^{\leftarrow L}_{\rightarrow A}$ equally. On the other hand, I tested one $R^{\leftarrow L}_{\rightarrow A}$ isolate from each of seven independent cooperator-dominated co-cultures, and found that they dominated ancestral $C^{\leftarrow L}$ in a nearly deterministic fashion (Figure 2.4A iv; 41/42 vs 31/64 cooperator-dominated, Fisher's Exact Test $P < 3\text{x}10^{-8}$). I then tested five isolates of

$C^{\leftarrow L}$ from three independent cheater-dominated co-cultures and found that they deterministically dominated ancestral cooperators (Figure 2.4A vii; 0/48 vs 31/64 cooperator-dominated, Fisher's Exact Test $P < 10^{-9}$). Thus, changes occurring in $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ were heritable and sufficient for dominating ancestral $C^{\leftarrow L}$ and $R^{\leftarrow L}_{\rightarrow A}$, respectively.

If $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ sampled from the same pool of mutations during the adaptive race, then even the "losing" types—i.e., $R^{\leftarrow L}_{\rightarrow A}$ from cheater-dominated and $C^{\leftarrow L}$ from cooperator-dominated co-cultures—may have nevertheless improved relative to their respective ancestors. I tested two "losing" $C^{\leftarrow L}$ isolates from cooperator-dominated co-cultures, and both were significantly better than ancestral $C^{\leftarrow L}$ at dominating ancestral $R^{\leftarrow L}_{\rightarrow A}$ (Figure 2.4A v; 0/12 vs 31/64 cooperator-dominated, Fisher's Exact Test $P<0.002$). I tested one "losing" $R^{\leftarrow L}_{\rightarrow A}$ isolate, and it was better than ancestral $R^{\leftarrow L}_{\rightarrow A}$ at dominating ancestral $C^{\leftarrow L}$ (Figure 2.4A vi; 6/6 vs 31/64 cooperator-dominated, Fisher's Exact Test $P<0.03$).

In addition to dominating ancestral $C^{\leftarrow L}$ and the evolved $C^{\leftarrow L}$ that they had raced against, evolved $R^{\leftarrow L}_{\rightarrow A}$ also improved in the sense that they lowered the minimal cell density required to initiate a viable cooperative co-culture (Shou et al., 2007). I tested a panel of six $R^{\leftarrow L}_{\rightarrow A}$ strains isolated from cooperator-dominated (Figure 2.4B, orange) or cheater-dominated (Figure 2.4B, gray) co-cultures. I mixed each isolate 1:1

with ancestral $Y_{\rightarrow L}^{\leftarrow A}$ and serially diluted each co-culture into SD. Compared to ancestral co-cultures, co-cultures initiated with evolved $R_{\rightarrow A}^{\leftarrow L}$ required less than one third the initial cell density to achieve growth in 50% of replicate co-cultures (Figure 2.4B). Co-cultures initiated with evolved $Y_{\rightarrow L}^{\leftarrow A}$ showed a similar degree of improvement (Figure 2.4C).

2.4. The adaptive race is fueled by mutations in a small set of genes involved in nutrient transport

I used whole-genome re-sequencing to identify the mutations underlying the rapid evolution in cooperators and cheaters. Some strains contained only one unambiguous, high-quality, non-synonymous SNP. These mutations defined a set of genes, exactly one of which was mutated in every sequenced strain that contained any non-synonymous mutations. I therefore reasoned that mutations in these genes were responsible for improvements in cooperation and cheating relative to the ancestor strains. All alleles of these genes and the strains that contain them are reported in Table 2.2. A complete list of sequenced strains and their SNPs can be found in Dataset S1 (available online at http://www.pnas.org/content/109/47/19079/suppl/DCSupplemental. Of the 17 unique alleles found, *ECM21* and *DOA4* had 6 each, accounting for 70% of the total. A simple estimate based on the Poisson distribution suggests that I found ~97% of the genes responsible for the initial fitness increase. The small number of genes could still allow for rapid divergence between populations if each of the many

alleles conferred a different fitness gain, since it would be unlikely for two populations to sample the same allele of the same gene.

Most of the identified mutations suggested that evolved $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ enhanced import of the limiting lysine provided by $Y^{\leftarrow A}_{\rightarrow L}$ by decreasing transporter turnover. Ecm21p is an arrestin-like adaptor protein that allows the E3 ubiquitin ligase Rsp5p to ubiquitinate and degrade the high-affinity lysine permease Lyp1p in response to stress (Lin et al., 2008). All but one of the *ECM21* mutations introduced a premature stop codon before the PY-motif necessary for interaction with Rsp5p (Lin et al., 2008). Doa4p is a deubiquitination protein required to maintain free ubiquitin pools (Swaminathan et al., 1999) and is specifically implicated in deubiquitination of plasma membrane proteins (Dupré and Haguenauer-Tsapis, 2001). Significantly, Doa4p represses expression of the general amino acid permease GAP1 (Springael et al., 1999), which can import lysine (Grenson et al., 1970a; Regenberg et al., 1999a). Less frequently observed were mutations in *RSP5* itself, in *BRO1* which is necessary for the function of Doa4p by recruiting it to endosomes (Luhtala and Odorizzi, 2004), and in the high-affinity lysine permease, *LYP1* (Grenson, 1966; Regenberg et al., 1999a).

To test whether nutrient transport was enhanced in evolved strains, I compared the ability of evolved and ancestral cells to grow into microcolonies (defined as >5 cells) under limiting concentrations of lysine. While ancestral $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ could grow on 4 μM but not 1 μM lysine (Figure 2.5A; 'Anc.'), every evolved $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ isolate

could grow as well on 1 μM lysine as they did on 4 μM lysine (Figure 2.5A; 'Evo.'). This improvement came at a cost, as almost all of the evolved cooperators and cheaters grew slower in non-limiting lysine than their ancestors (Figure 2.5B), which is consistent with their low frequency at the start of our experiments. Thus, adaptation to the nutrient-limited cooperative environment involved a trade-off in fitness in non-limiting nutrient.

To confirm that the observed mutations were sufficient to improve growth on limiting lysine, I replaced full-length *ECM21* in ancestral $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ with the truncated version found in the evolved $R^{\leftarrow L}_{\rightarrow A}$ strain CT22 (Table 1). Indeed, nearly 100% of these cells were able to grow on media containing 1 μM lysine (Figure 2.5A; '*ecm21*').

In addition to these SNPs, many strains contained duplications of one or more chromosomes. Of particular interest is CT10, which has a SNP in LYP1. It also has a duplication of chrXIV, which is where LYP1 resides. I checked IGV, and this strain is homogenous for this SNP, suggesting that both copies of chrXIV have this SNP. I believe the most parsimonious explanation is that the SNP preceeded the duplication, which is opposite of the order I think we are more comfortable with (i.e., duplication is a "quick fix" that precedes acquisition of SNPs).

2.5. Improved cooperators can stochastically defeat newly-arising cheaters

If during adaptation to the nutrient-limited cooperative environment the maximum fitness gain of cooperators exceeds that of cheaters by at least the cost of cooperation, cooperators will win the race (Figure 2.2A), and eventually purge the inferior cheaters. However, a mutation could turn an evolved cooperator into a cheater, which would now be as well-adapted to the nutrient-limited cooperative environment as the cooperators and would therefore rise in frequency. Since successive fitness gains during adaptation to the same environment are expected to decrease over time (Lenski and Travisano, 1994), I tested whether evolved cooperators could stochastically dominate otherwise isogenic cheaters during a subsequent round of the adaptive race. I chose two improved cooperator strains (CT22 and CT75, Table 2.2) carrying different truncation alleles of *ECM21* and derived, via allele replacement, matched cooperators and cheaters marked by different drug resistances (two independent pairs from CT22 and one pair from CT75; Table 2.2 and Table 2.1). As a control, similarly matched cooperators and cheaters were derived from ancestral $R^{\leftarrow L}_{\rightarrow A}$ (Table 2.1).

As before, $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ pairs were mixed with ancestral $Y^{\leftarrow A}_{\rightarrow L}$ at a ratio of 1:1:1. The frequencies of each cell type were periodically measured by plating onto media containing the appropriate antibiotic. As expected, the derived ancestral co-cultures showed stochastic cooperator dominance (2/12 cooperator-dominated). Of the co-cultures derived from the two *ecm21* alleles (Figure 2.6; $*ecm21^{\leftarrow L}_{\rightarrow A}$ and

$*ecm21^{\leftarrow L}$), 11 were cooperator-dominated, 4 were cheater-dominated, and 2 were indeterminate. Thus, even after cooperators had obtained major fitness improvements in the cooperative environment, a fraction of them were able to fend off cheaters that were just as well-adapted.

2.6. Materials and methods

Strains and cell culture

All strains were derived from the BY designer deletion strains of the S288C background (Brachmann et al., 1998), and created essentially as described in (Shou et al., 2007). See Table S1 for a list of constructed strains used in this study. To derive cooperating and cheating strains from a parent strain, DNA fragments PCR amplified off a cassette containing either wild-type *ADE4* and resistance to Hygromycin B (*hph*; WSB150), or an *ADE4* overproduction allele (*PUR6*) and resistance to ClonNAT (*nat*; WSB151), were transformed into the parent strain. The primers contained regions homologous to the site of integration. Transformants were selected on media containing the appropriate antibiotic. Presence of *ADE4* or *PUR6* in the lysine-auxotrophic parent strain was determined by patching a small amount of transformant onto SD+lys plates pre-spread with *ade8* diploid cells and cut into individual squares. Overproduction of adenine resulted in adenine release and visible growth of *ade8* cells into satellite colonies, while strains with *ADE4* did not. *ecm21*$^{\Delta 143}$ strains were created by replacing *ECM21* after the 142nd amino acid with DNA fragments PCR amplified off a *hph*--containing cassette (WSB117).

Cells were grown at 30 °C in the rich medium YPD (1% Bacto-yeast extract, 2% Bacto-peptone, 2% Dextrose, 2% Bacto-agar) or the minimal medium SD (0.67% Difco Yeast Nitrogen Base without amino acids, 2% Dextrose) with or without supplemented amino acids, as necessary (Guthrie and Fink, 1991). To prepare for use, strains were

struck from the freezer and grown on YPD plates at 30 °C for 2–3 days. An individual colony was picked and grown at 30 °C in liquid YPD until saturated (1–3 days). Unless otherwise indicated, strains were diluted >1:1000 into SD + lysine or adenine and grown for ~16 hrs to exponential phase (~$2x10^6$ – $4x10^6$ cells/ml) before being assayed.

Strain name	Derived from	Genotype
WY954 ($Y^{\leftarrow A}_{\rightarrow L}$)		*MAT*a *ste3*::*kanMX4 ade8*Δ0 *lys21*::*LYS21*(fbr) *trp1-289*::pRS404(TRP)-ADHp-venus-YFP
WY950 ($R^{\leftarrow L}_{\rightarrow A}$)		*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::*PUR6* *trp1-289*::pRS404(TRP)-ADHp-DsRed.T4
WY962 ($C^{\leftarrow L}$)		*MAT*a *ste3*::*kanMX4 lys2Δ0* *trp1-289*::pRS404(TRP)-ADHp-CFP
WY1356 WY1357	WY962	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::TEFp-*hph*-TEFt-*ADE4* *trp*1-289::pRS404(*TRP*)-ADHp-DsRed.T4
WY1358 WY1359	WY950	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::TEFp-*nat*-TEFt-*PUR6* *trp*1-289::pRS404(*TRP*)-ADHp-DsRed.T4
WY1390	WY950	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::*PUR6* *trp1-289*::pRS404(TRP)-ADHp-DsRed.T4 *ecm21*::*ecm21*$^{\Delta 143}$-loxP$_2$-*hph*
WY1391	WY950	*MAT*a *ste3*::*kanMX4 lys2Δ0* *trp1-289*::pRS404(TRP)-ADHp-CFP *ecm21*::*ecm21*$^{\Delta 143}$-loxP-*hph*-loxP
WY1403 ($*ecm21^{\leftarrow L}$) WY1404 ($*ecm21^{\leftarrow L}$)	CT22	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::TEFp-*hph*-TEFt-*ADE4* *trp*1-289::pRS404(*TRP*)-ADHp-DsRed.T4 *ecm21*^{Q143*}
WY1405 ($*ecm21^{\leftarrow L}_{\rightarrow A}$) WY1406 ($*ecm21^{\leftarrow L}_{\rightarrow A}$)	CT22	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::TEFp-*nat*-TEFt-*PUR6* *trp*1-289::pRS404(*TRP*)-ADHp-DsRed.T4 *ecm21*^{Q143*}
WY1407† ($*ecm21^{\leftarrow L}$)	CT75	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::TEFp-*hph*-TEFt-*ADE4* *trp*1-289::pRS404(*TRP*)-ADHp-DsRed.T4 *ecm21*^{Q544*}
WY1408† ($*ecm21^{\leftarrow L}_{\rightarrow A}$)	CT75	*MAT*a *ste3*::*kanMX4 lys2*Δ0 *ade4*::TEFp-*nat*-TEFt-*PUR6* *trp*1-289::pRS404(*TRP*)-ADHp-DsRed.T4 *ecm21*^{Q544*}

Table 2.1: Constructed strains used in this study.
'†' indicates the presence of additional high-quality SNPs in this strain.

Calculation of average doubling times

The average doubling times of $R^{\leftarrow L}_{\rightarrow A}$ and $R!$ were calculated using the formula $\frac{t}{\log_2\left(\frac{N(t)}{N(0)}\right)}$, where t is time in hours, $N(0)$ is the population size at time 0, and $N(t)$ is the population size at time t.

Competition assays

To measure the fitness difference between cooperators and cheaters (Fig 1B), exponentially growing cells were washed and resuspended in SD, diluted to $1.7\text{x}10^5$ cells/ml into SD (for death competition) or SD + 164 μM lysine (for growth competition) in triplicate, and incubated at 30 °C on a rotator. Growing co-cultures were periodically diluted to maintain exponential growth, while keeping population size above $3\text{x}10^4$ cells/strain. Ratios of the competing strains were determined by flow cytometry every 24 hours. The relationship between the $C^{\leftarrow L}$: $R^{\leftarrow L}_{\rightarrow A}$ ratio and time in generations of $R^{\leftarrow L}_{\rightarrow A}$ (*t*) was fit to the form $\frac{C^{\leftarrow L}}{R^{\leftarrow L}_{\rightarrow A}} = A\exp(rt)$ using the method of weighted non-linear least-squares, where *r* is the difference in Malthusian parameters between $C^{\leftarrow L}$ and $R^{\leftarrow L}_{\rightarrow A}$. Because the growth rate of any population ($R^{\leftarrow L}_{\rightarrow A}$ here) is $\ln(2)$ /generation by definition, percent fitness advantage with respect to $R^{\leftarrow L}_{\rightarrow A}$ was calculated as $r/\ln(2)$. To determine fitness differences between $R^{\leftarrow L}_{\rightarrow A}$ and its evolved descendants

(Fig 5B), which were all red-fluorescent, an otherwise isogenic yellow-fluorescent $Y^{\leftarrow L}_{\rightarrow A}$ strain was competed against each strain. The difference in fitness between $Y^{\leftarrow L}_{\rightarrow A}$ and $R^{\leftarrow L}_{\rightarrow A}$ (0.1%) was an order of magnitude less than the difference in fitness between $R^{\leftarrow L}_{\rightarrow A}$ and the fastest-growing evolved strain (1.7%, CT22).

Flow cytometry

Population compositions were measured by flow cytometry using a FACSCanto II (BD Biosciences) or DxP10 (Cytek) with three lasers (405nm, 488nm, and 633nm). If necessary, fluorescent beads of known concentration were added to determine cell densities. Data analyses were performed using FloJo (TreeStar) or custom software written in R (R Development Core Team, 2012) using the Bioconductor (Gentleman et al., 2004) packages flowCore (Hahne et al.), flowMeans (Aghaeepour, 2010), and flowViz (Sarkar et al., 2008).

Co-culture growth and assay of population compositions

Exponentially growing cells were washed in SD and counted using a CoulterCounter Z2 (BeckmanCoulter), or by flow cytometry. After counting, cells were diluted into SD and mixed. Then, 3 ml aliquots were put in 13 mm tubes and cultured on a rotator at 30 °C. OD_{600} was monitored at least once a day using a Gensys 20 (Thermo Spectronic), and population compositions were determined by flow cytometry. To estimate densities and ratios for the **ecm21* strains (Fig 6), co-cultures were sampled

into a microtiter plate and serially diluted in ten-fold steps from 1- to 10^{-4}-fold using a multichannel pipette, spotted onto agar plates containing YPD, YPD+clonNAT (100 µg/ml), and YPD+ Hygromycin B (200 µg/ml), and incubated at 30 °C for up to 7 days. Colony counts were used to calculate population sizes of the entire co-culture (YPD), $*ecm21^{\leftarrow L}_{\rightarrow A}$ (YPD+clonNAT), and $*ecm21^{\leftarrow L}$ (YPD+Hygromycin B).

Viability assay

Exponentially growing cells were washed free of supplements. Ancestral $Y^{\leftarrow L}_{\rightarrow A}$ and ancestral and evolved $R^{\leftarrow L}_{\rightarrow A}$ strains were mixed 1:1 and serially diluted in a deep multi-well trough (~10 ml/well) (Fig 4B). Six or twelve 200 µl aliquots of each prepared density were transferred into wells of a microtiter plate using a multi-channel pipette, which was then sealed using parafilm and placed in a 30 °C incubator in moisturized tupperware bins. Visible growth by eye was considered successful initiation of cooperation. Data were analyzed using a generalized linear model (GLM) with a binomial link function. Strain and the logarithm of cell density were fixed effects.

Microcolony growth assay

Exponentially growing cells were washed and resuspended in SD, then diluted to ensure that they would remain at low density after residual growth. After starvation in SD for 16–20 hours, cell density was estimated by OD_{600}, diluted to ~10^4 cells/ml, and plated onto SD plus 4 µM or 1 µM lysine (Fig 5A). Plates were observed using light

microscopy under a 10x objective after 24 (4 μM lysine) or 48 (1 μM lysine) hours. Clusters of 5 or more cells were considered microcolonies. Percentages and 95% confidence intervals for the ancestor and evolved data in Figure 2.5A were estimated by a Generalized Linear Mixed-Effect Model (GLMM) with a binomial link function using the package lme4 (Bates et al., 2011). Cooperator/cheater ("type"), ancestor/evolved ("state"), lysine concentration, and lysine concentration-state interaction were fixed effects, while strain and strain-experiment interaction were random effects.

Sequencing and analysis

Genomic DNA was extracted using the Genomic-tip 20/G kit (QIAGEN). Libraries were prepared using the tagmentation reaction through Nextera DNA Sample Preparation Kit, Nextera Index Kit (96-indices), and the TruSeq Dual Index Sequencing Primers (Illumina). All cleanup steps were performed with DNA clean & concentrator-5 (Zymo research). Libraries from 30 strains (including the ancestral $R_{\rightarrow A}^{\leftarrow L}$ and $C^{\leftarrow L}$) were multiplexed and run on a HiSeq2000 (Illumina) using 50-cycles paired-end reading. Sequence data was analyzed using a custom perl script that incorporated the bwa aligner (Li and Durbin, 2009) and samtools (Li et al., 2009).

2.7. Figures

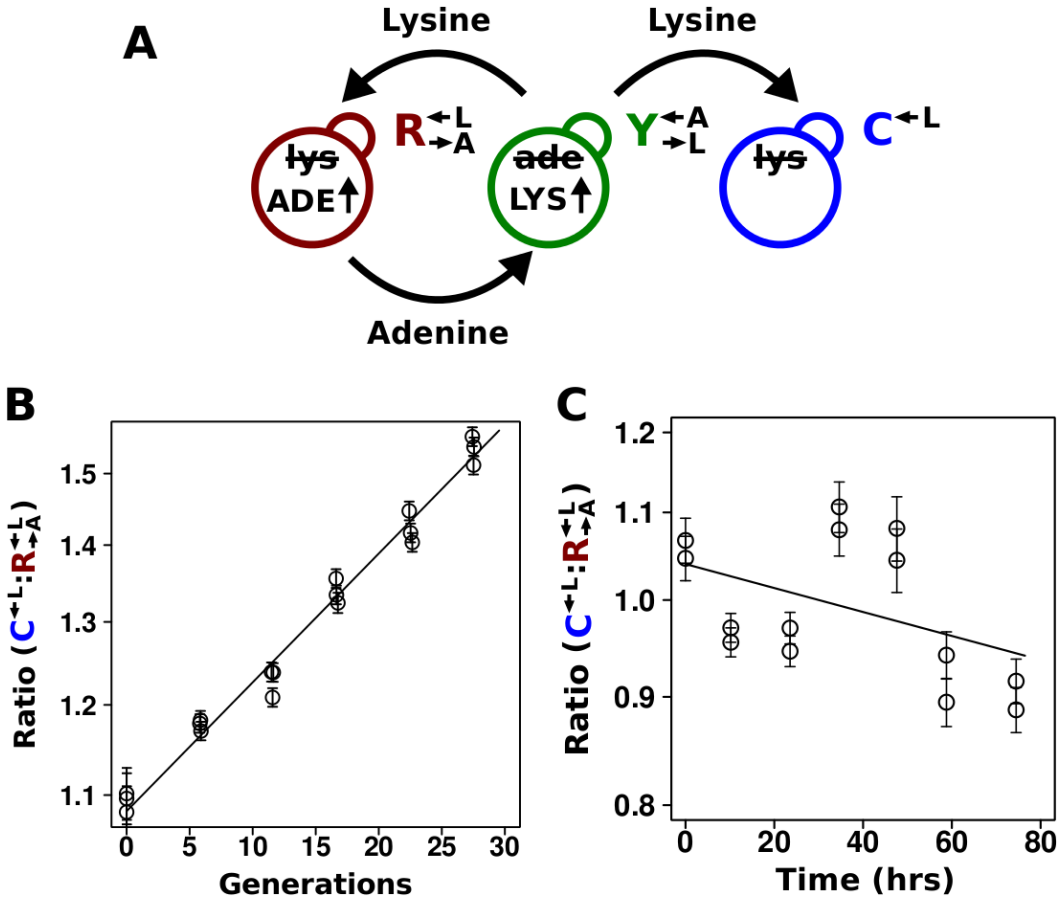
A
Lysine
Lysine
R L A
Y A L
C L
lys
ADE
ade
LYS
lys
Adenine
B
Ratio (C L:R L A)
1.1
1.2
1.3
1.4
1.5
0
5
10
15
20
25
30
Generations
C
Ratio (C L:R L A)
0.8
0.9
1.0
1.1
1.2
0
20
40
60
80
Time (hrs)

Figure 2.1: A yeast model of cooperation and cheating.
A) Our cooperative system is composed of two non-mating yeast strains. The red-fluorescent $R^{\leftarrow L}_{\rightarrow A}$ strain requires lysine and over-produces adenine, while the yellow-fluorescent $Y^{\leftarrow A}_{\rightarrow L}$ strain requires adenine and over-produces lysine. The cheater is a cyan-fluorescent $C^{\leftarrow L}$ strain that requires lysine and doesn't over-produce adenine. **B**) *Cheaters have a fitness advantage over cooperators*. $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ strains were mixed and competed in SD supplemented with non-limiting (164 μM) lysine. Co-cultures were periodically diluted into fresh supplemented medium to maintain exponential growth. The ratio of $C^{\leftarrow L}$ to $R^{\leftarrow L}_{\rightarrow A}$ was determined using flow cytometry. The ratio was fit using weighted non-linear least squares regression to the form $A\exp(rt)$, where *A* is the initial ratio, *r* is the cheater growth rate minus the cooperator growth rate, and *t* is time in generations of $R^{\leftarrow L}_{\rightarrow A}$. Note the logarithmic y-axis. $C^{\leftarrow L}$ has a 1.8% (95% CI: 1.6% – 1.9%) advantage over $R^{\leftarrow L}_{\rightarrow A}$. **C**) *Cheaters and cooperators die at similar rates*. $R^{\leftarrow L}_{\rightarrow A}$ and $C^{\leftarrow L}$ strains were mixed in SD. The ratio of $C^{\leftarrow L}$ to $R^{\leftarrow L}_{\rightarrow A}$ was determined using flow cytometry. The ratio was fit using weighted non-linear least squares regression to the form $A\exp(rt)$, where *A* is the initial ratio, *r* is the death rate difference of cheater over cooperator, and *t* is time in hours. Note the logarithmic y-axis is. No significant fitness difference was observed (95% CI of difference in death rate: -0.003–0.0003 hr^{-1}).

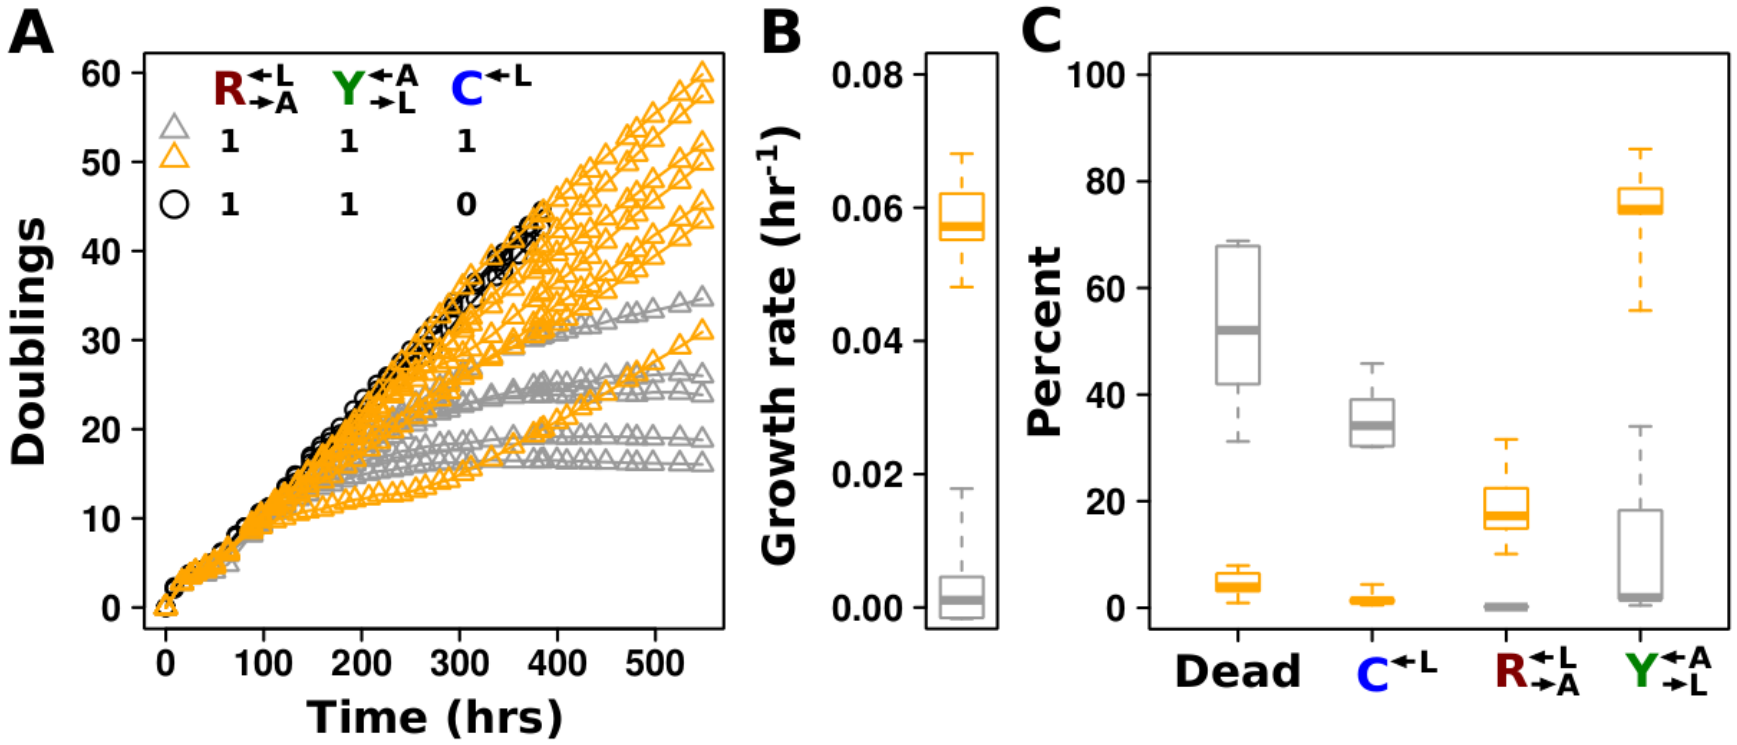
A
Doublings
Time (hrs)
0
10
20
30
40
50
60
0
100
200
300
400
500
R ← L
→ A
Y ← A
→ L
C ← L
1
1
1
1
1
0
B
Growth rate (hr-1)
0.00
0.02
0.04
0.06
0.08
C
Percent
0
20
40
60
80
100
Dead
C ← L
R ← L
→ A
Y ← A
→ L

Figure 2.2: Stochastic cheater outcome in initially identical cooperator-cheater cocultures.

A, B) *Distinct growth behavior in replicate cooperator-cheater co-cultures*. Exponential cultures of $R^{\leftarrow L}_{\rightarrow A}$, $Y^{\leftarrow A}_{\rightarrow L}$, and $C^{\leftarrow L}$ were washed to remove residual supplements and mixed as indicated to a final density of 4.2×10^5 cells/ml per strain. Twelve plus-cheater (triangles) and three minus-cheater (black circles) co-cultures were propagated, with total cell density maintained in the sub-saturation range of ~10^6 cells/ml to 1.5×10^7 cells/ml by dilution into fresh SD when necessary. **B**) After ~400 hrs, the growth rates of plus-cheater co-cultures were bimodally distributed into fast- (orange) and slow-growing (gray) groups. **C**) *Fast-growing co-cultures are dominated by cooperators*. At ~450 hrs, the frequencies of dead cells reacting to the nucleic acid dye TOPRO-3 (Dead), $C^{\leftarrow L}$, $R^{\leftarrow L}_{\rightarrow A}$, and $Y^{\leftarrow A}_{\rightarrow L}$ cells in fast- (orange) and slow-growing (gray) co-cultures were quantified using flow cytometry. Boxes extend from the first to the third quartile of the data, "whiskers" extend to the most extreme observations, and the thick bar inside each box represents the median.

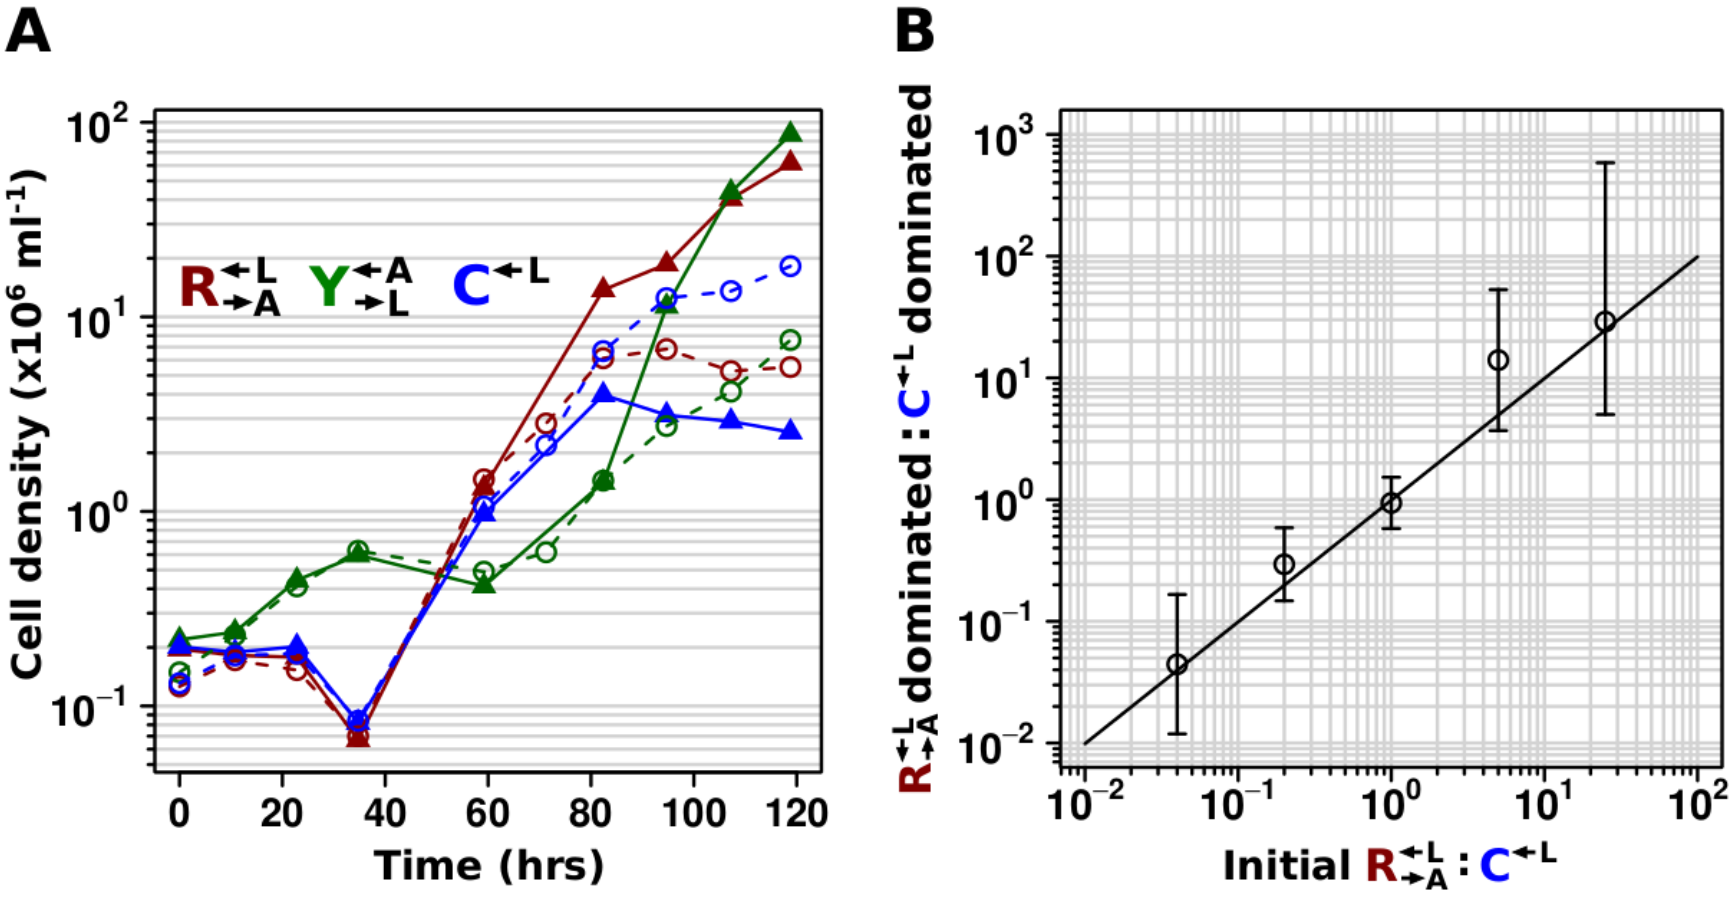
A
B
Cell density (x10^6 ml^-1)
Time (hrs)
R←L →A
Y←A →L
C←L
R←L →A dominated : C←L dominated
Initial R←L →A : C←L

Figure 2.3: Rapid divergence in the growth of cooperators and cheaters in cooperator-dominated and cheater-dominated co-cultures.
A) Co-cultures were initiated at a density of 1.7×10^5 cells/ml per strain. Cell densities of the three sub-populations were measured by flow cytometry by adding fluorescent beads of a known concentration to each sample. After initially following similar trajectories, rapid divergence resulted in dominance of $R^{\leftarrow L}_{\rightarrow A}$ cells (triangles connected by solid lines), or $C^{\leftarrow L}$ cells (circles connected by dashed lines). A representative example of each class is shown. **B**) *The frequency of cooperator-dominated co-cultures is determined by the initial frequency of cooperators*. Co-cultures were prepared as in **A**, except the initial density of $C^{\leftarrow L}$ or $R^{\leftarrow L}_{\rightarrow A}$ was varied from 8.3×10^5 to 4.2×10^6 cells/ml per strain. Total initial population size influenced how quickly co-cultures diverged, but not the eventual frequency of cooperator-dominated co-cultures. Error bars represent 95% confidence intervals assuming that the number of viable co-cultures followed a binomial distribution with parameters *n* and *p*, where *n* is the total number of co-cultures and *p* is the probability of being cooperator-dominated. The slope of the logistic regression fitting line $\log\left(\frac{p}{1-p}\right) = \beta \log\left(\frac{c}{1-c}\right)$, where *c* is the initial cooperator frequency, is 1.0 (95% CI: 0.7 – 1.3). This suggests that the odds of being a cooperator-dominated co-culture is nearly equal to the initial proportion of $R^{\leftarrow L}_{\rightarrow A}$. This occurs because, in our system, the fitness cost of cooperation is small.

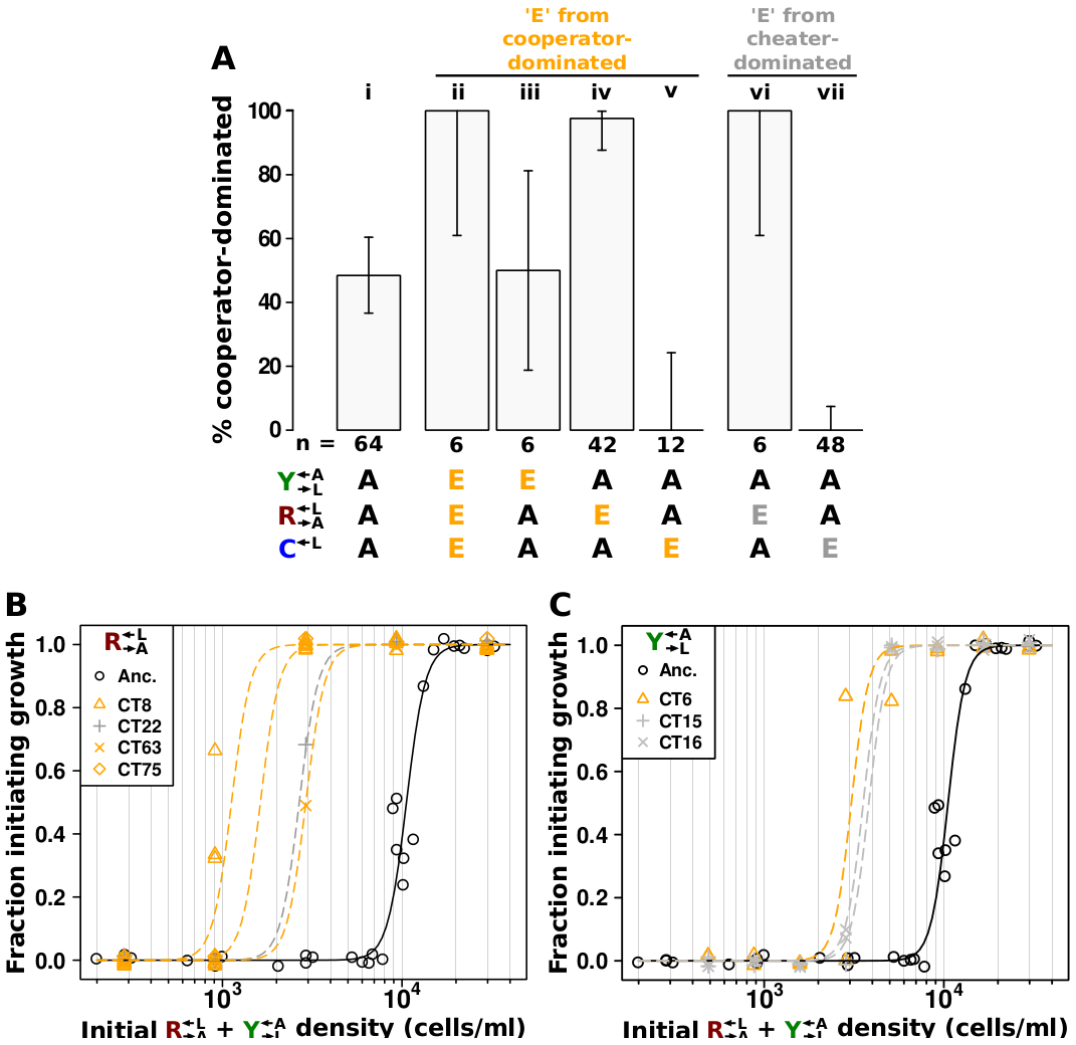
A
'E' from cooperator-dominated
'E' from cheater-dominated
i
ii
iii
iv
v
vi
vii
% cooperator-dominated
100
80
60
40
20
0
n = 64 6 6 42 12 6 48
Y A A E E A A A A
R A A E A E A E A
C A A E A A E A E
B
Fraction initiating growth
1.0
0.8
0.6
0.4
0.2
0.0
R
Anc.
CT8
CT22
CT63
CT75
10^3
10^4
Initial R + Y density (cells/ml)
C
Fraction initiating growth
Y
Anc.
CT6
CT15
CT16
10^3
10^4
Initial R + Y density (cells/ml)

Figure 2.4: Evolved cooperators and cheaters are superior to ancestors.
A) *Improved cooperation and cheating are heritable and occur in both "winning" and "losing" strains*. Co-cultures consisting of ancestral strains (**A**) and strains isolated from cooperator-dominated (orange) or cheater-dominated (gray) co-cultures (**E**) were formed as indicated at equal initial densities. Percentages of cooperator-dominated co-cultures were measured from replicate co-cultures whose numbers ('n') are indicated. Error bars represent the 95% confidence interval calculated as in Figure 2.3. **B**) *Improved* $R^{\leftarrow L}_{\rightarrow A}$ *can initiate cooperation at lower densities than their ancestor*. Ancestral $R^{\leftarrow L}_{\rightarrow A}$ (black) and $R^{\leftarrow L}_{\rightarrow A}$ isolated from cooperator-dominated (orange) and cheater-dominated (grey) co-cultures were mixed 1:1 with ancestral $Y^{\leftarrow A}_{\rightarrow L}$ and serially diluted into a microtiter plate. After 1 month at 30 °C, wells showing visible growth were scored. Data are jittered slightly in the vertical direction to aid visualization. **C**) Improved $Y^{\leftarrow A}_{\rightarrow L}$ cooperators can also initiate cooperation at lower densities than their ancestor. Ancestral $Y^{\leftarrow A}_{\rightarrow L}$ (black) and $Y^{\leftarrow A}_{\rightarrow L}$ isolated from cooperator-dominated (orange) and cheater-dominated (grey) co-cultures were mixed 1:1 with ancestral $R^{\leftarrow L}_{\rightarrow A}$ and serially diluted into a microtiter plate. After 1 month at 30 °C, wells showing visible growth were scored. Data are jittered slightly in the vertical direction to aid visualization.

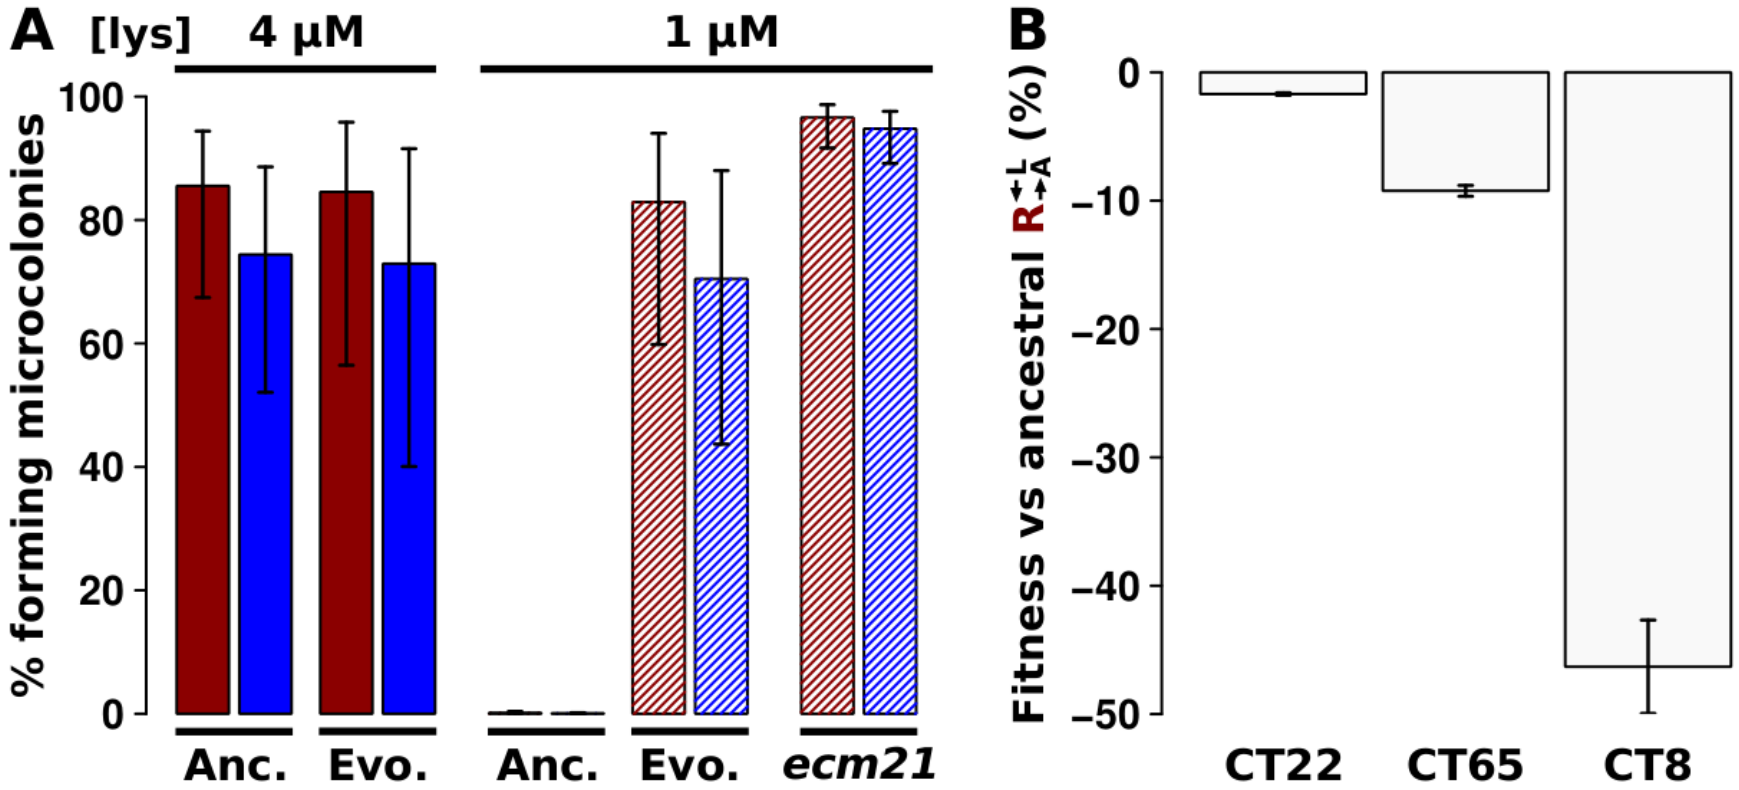
A
[lys]
4 µM
1 µM
% forming microcolonies
100
80
60
40
20
0
Anc.
Evo.
Anc.
Evo.
ecm21
B
Fitness vs ancestral R↔L A (%)
0
−10
−20
−30
−40
−50
CT22
CT65
CT8

Figure 2.5: Evolved cooperators and cheaters show a tradeoff between the ability to grow at low lysine concentrations and the maximum growth rate achieved at high lysine concentrations.
A) *Evolved strains grow better at low lysine concentrations than their ancestors*. Exponentially growing ancestral (Anc.) or evolved (Evo.) cooperator (red) and cheater (blue) strains were washed free of lysine and starved for 16-20 hours to deplete vacuolar storage of lysine. Cells were plated at ~200 cells/cm^2 onto minimal medium agar containing 4 μM (solid bar) or 1 μM (shaded bar) lysine. After 1–3 days, clusters of >5 cells were counted as microcolonies. Percentages and standard errors were estimated using a Generalized Linear Mixed-Effect Model (Text S1). The *ecm21*$^{\Delta 143}$ allele from evolved cooperator CT22 was cloned into the ancestral cooperator and cheater strains and their growth on low lysine plates was measured (*ecm21*). The percentage is based on scoring >100 cells. Error bars indicate the 95% confidence interval assuming a binomial distribution with parameters *n* and *p*, as in Figure 2.3B. **B**) *Evolved cooperators have lower growth rates than their ancestor in non-limiting lysine*. Relative fitness compared to ancestral $R_{\rightarrow A}^{\leftarrow L}$ was measured for CT22 and CT65 through direct competition, and calculated for CT8 using monoculture growth rates.

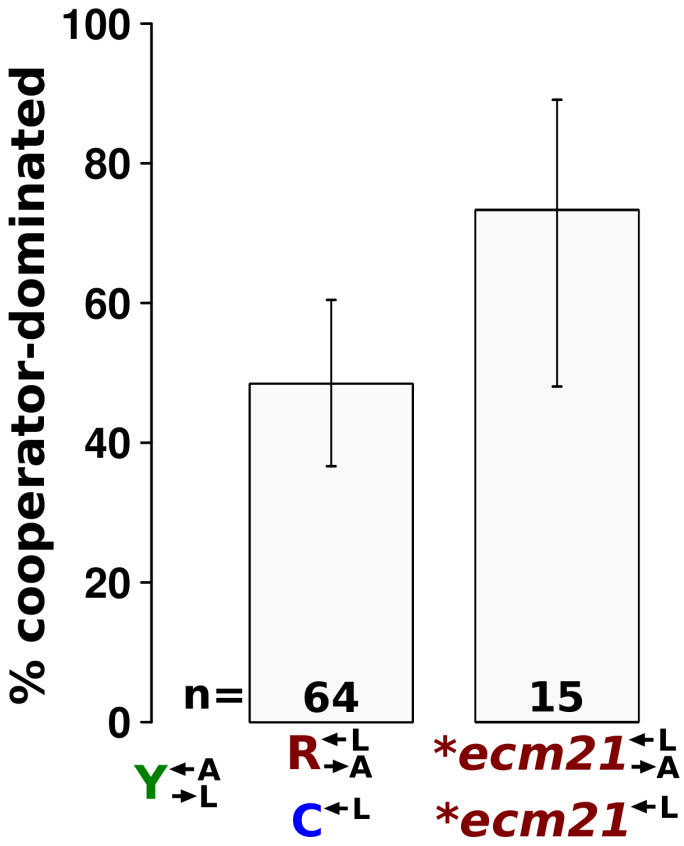

Figure 2.6: Stochastic cheater purging observed in co-cultures composed of cooperators and cheaters previously exposed to the cooperative environment. We replaced the adenine overproduction allele of *ADE4* (*PUR6*) in two evolved $R^{\leftarrow L}_{\rightarrow A}$ strains containing truncation alleles of *ECM21* (CT22 and CT75; Table 1) with *PUR6* or wild-type *ADE4* marked with different antibiotic resistances. Co-cultures were initiated at 1:1:1 . Frequencies of cell types were periodically determined by plating onto media containing the appropriate antibiotic. A co-culture was considered "cooperator-dominated" if the ratio of $R^{\leftarrow L}_{\rightarrow A}$ to $C^{\leftarrow L}$ was >100, or if the frequency of $C^{\leftarrow L}$ fell below 1% of its initial frequency. It was considered "cheater-dominated" if the ratio of $C^{\leftarrow L}$ to $R^{\leftarrow L}_{\rightarrow A}$ was >100, or if the frequency of $R^{\leftarrow L}_{\rightarrow A}$ fell below 1% of its initial frequency, and "indeterminate" if none of these conditions were satisfied after ~2000 hours of propagation. 'n' is the sum of cooperator-dominated and cheater-dominated co-cultures. Since results across pairs were not significantly different from one another, the data were combined and are presented as $*ecm21^{\leftarrow L}_{\rightarrow A}$ and $*ecm21^{\leftarrow L}$. Error bars are the 95% confidence interval calculated as in Figure 2.3B.

Gene	Allele	Strain	Cell type	Co-culture dominated by	Mutant freq.
ECM21	Q143*	CT22	$R^{\leftarrow L}_{\rightarrow A}$	cheat	0.35
	E268*	CT11†	$C^{\leftarrow L}$	cheat	
	S300*	CT78	$R^{\leftarrow L}_{\rightarrow A}$	coop	
	L475R	CT82	$C^{\leftarrow L}$	cheat	
	Q544*	CT75†	$R^{\leftarrow L}_{\rightarrow A}$	coop	
	L780*	CT63†	$R^{\leftarrow L}_{\rightarrow A}$	coop	
DOA4	D213G	CT60	$R^{\leftarrow L}_{\rightarrow A}$	cheat	0.35
	P235L	CT69	$C^{\leftarrow L}$	cheat	
	A643P	CT12	$C^{\leftarrow L}$	cheat	
	Q660K	CT7†	$R^{\leftarrow L}_{\rightarrow A}$	coop	
	Ins 'A' → C735*	CT17	$C^{\leftarrow L}$	coop	
	Y916*	CT61†	$R^{\leftarrow L}_{\rightarrow A}$	coop	
RSP5	L131W	CT8	$R^{\leftarrow L}_{\rightarrow A}$	coop	0.12
	N746K	CT110†	$C^{\leftarrow L}$	cheat	
BRO1	Q434*	CT83†	$C^{\leftarrow L}$	cheat	0.12
	S673*	CT57†	$R^{\leftarrow L}_{\rightarrow A}$	cheat	
LYP1	D583N	CT10	$C^{\leftarrow L}$	cheat	0.06

Table 2.2: Mutations potentially sufficient for the large observed fitness advantage in the nutrient-limited cooperative environment.
'' indicates a STOP codon. '†' indicates the presence of additional high-quality SNPs in this strain. When multiple isolates from the same co-culture were analyzed, some would harbor a single high-quality SNP, while others contained this SNP plus additional high-quality SNPs. Only the former type is included in the table.*

Strain	Duplicated chromosome
CT10	XIV
CT11	XIV
CT81	XII
CT59	IX
CT60	IX
CT64	XI, XIV
CT61	XVI
CT62	XII, XIV
CT65	XII, XIV
CT77	XIV
CT78	I, V, XII, XIV, XVI
CT82	XI, XII
CT83	XIV
CT106	XI
CT110	XIV

Table 2.3: Strains with chromosomal duplications.

Chapter 3. The adaptive race can synergize with sources of abiotic variation to promote cooperation in a metapopulation

Migration is necessary for the success of cooperation. Without migration, cooperators, even if they defeated cheaters in an adaptive race, cannot migrate to and expand in previously unoccupied locations or locations emptied by cheater-induced extinction. Since cooperators rarely rise from cheaters (however, see (Fiegna et al., 2006)), but cheaters can easily rise from cooperators through reducing or eliminating the production of costly common goods, cooperators are always threatened by cheaters. The deteriorating environment caused by the increasing cheater load, or any other environmental stress, could trigger a new round of adaptive race in which cooperators could lose. In addition, the waiting-time between mutations that confer fitness benefits greater than the intrinsic cost of cooperation is expected to lengthen (Lenski and Travisano, 1994). Thus, it may take many generations to sample a mutation that significantly enhances fitness (Blount et al., 2008) or produces a suitable cheater tolerance mechanism. During this time, cheaters might have already destroyed cooperation.

On the other hand, too much migration will also destroy cooperation. Essentially, a high migration rate transforms a metapopulation into a single well-mixed population with nowhere else to migrate to, which, as I have argued above, cannot

support cooperation.

What are “ideal” migration rates? Are they biologically plausible? Successful experimental studies have used robots to manipulate metapopulations in which subpopulations are defined by wells of microtiter plates (Nahum et al. 2011; Benjamin Kerr et al. 2006). However, these studies utilized fast-growing bacteria and viruses. Because the yeast cooperator-cheater system has a doubling time of ~12 hours, the experimental duration would be unrealistically long.

I developed an agent-based mathematical model to simulate metapopulations. A mathematical model has two advantages. First, it allows variations in assumptions to be easily tested. For example, I have tried a large number of migration rates and migration patterns (e.g., local versus global migration, discussed in more detail below) to get a better sense for how these parameters effect the overall cooperator abundance. Second, since metapopulation dynamics depends upon stochastic processes, a large number of populations must be followed for a sufficiently long time to obtain a distribution of “steady-state” results. By using experimentally-realistic parameters and a small set of assumptions, I attempted to generate predictions that could be tested experimentally in the future.

The adaptive race in my experiments was between a cooperator-type and a cheater-type that consumed the same resource. Thus, one cooperator type (*C*) and one cheater type (*D*) were explicitly modeled. In the terminology of Chapter 2, *C*

corresponds to $R_{\rightarrow A}^{\leftarrow L}$ and *D*to $C^{\leftarrow L}$. Even though experimentally a variety of mutations were able to confer a high-fitness phenotype during nutrient limitation, for simplicity the model includes one evolved cooperator type (*C**) and one evolved cheater type (*D**). In section 3.1, I describe the experiments I performed to guide my thinking in developing the model assumptions. The resulting mathematical model is presented in section 3.2.

3.1. Experimental observations

Fitness declines as cheater frequency increases

An important component of the model is the way that the fitnesses of each type change as cheater frequency increases. Experimentally, I observed a reduction in coculture growth rate with increased cheater frequency (Figure 2.2A). This is due to the reduction in the fitness of each sub-population as nutrient limitation becomes more severe. In general, the specific form of this fitness reduction will vary between cocultures and within coculutures as a function of time due to the sequential appearance of a variety of possible mutants.

To get a better idea of how the fitness of each sub-population changes as the frequency of cheaters increases, I used data from an experiment where I sampled cocultures twice a day for ~ 250 hours. Three replicate cocultures were initiated without cheaters, while an additional six cocultures were initiated with an initial $R_{\rightarrow A}^{\leftarrow L}$ to $C^{\leftarrow L}$ ratio of 1:1 (Figure 3.1A). By chance, $C^{\leftarrow L}$ took over all six replicate cocultures, yielding information about the variety of ways cocultures respond to an increasing

frequency of cheaters. I used data after the 100 hour time point, because this was when the cheater-free coculture achieved a relatively steady doubling time. At this point, $C^{\leftarrow L}$ had already increased from their initial frequency of 0.5 in every coculture. Because cheater takeover was so rapid, most of the data points were for cheater frequencies > 0.8. Two cocultures ("A" and "F") provided data for cheater frequencies < 0.6. As expected, the growth rates of $R_{\rightarrow A}^{\leftarrow L}$ (Figure 3.1B) and $C^{\leftarrow L}$ (Figure 3.1C) decrease as the frequency of $C^{\leftarrow L}$ increases. While the individual trajectories are different from one another, they all converge to a similar growth rate of ~ 0. In addition, in a given coculture cheater fitness decreases more slowly than cooperator fitness.

Evolved strains die more slowly during starvation

Evolved strains grow better than their ancestors in the nutrient-limited environment of cooperation (Figure 2.4A,B and Figure 2.5A), but how are their death rates affected? To compare the death rates of ancestor and evolved strains, I developed a fluorescence microscopy assay to measure loss of fluorescence in minimal media (see Section 3.7). This assay showed that evolved strains also die slower than the ancestor (Figure 3.2). However, this apparent discrepancy between cooperator and cheater death rates is due to differences in the fluorophore used to mark each strain (data not shown). Whereas multimeric dsRed is likely to be retained in cells with compromised membrane integrity, monomeric GFP easily leaks out such cells. This causes a difference in the decline rate of total fluorescence intensity, which is used to measure rate of cell death.

3.2. Model assumptions

To incorporate these experimental observations, I modeled fitness as a decreasing linear function of cheater frequency (*x*). Let $r_z(x)$ be the growth rate of sub-population *z* when the frequency of cheaters is *x*. Without loss of generality, I defined a maximum death rate (d_{max}) for all sub-populations in terms of the maximum growth rate of *C*, such that

$$d_{max} = -\frac{r_C(1)}{r_C(0)}. \tag{3.1}$$

When $x=0$, *C** had a growth rate advantage δ over C, and each cheater type (*D*, *D**) had a growth rate advantage α over its corresponding cooperator type (*C*, *C**). Let the genotype be denoted by (*i*, *j*), where *i*=0 for cooperators, *i*=1 for cheats, *j*=0 for ancestral, and *j*=1 for adapted mutants. Fitness at cheater frequency *x* was calculated as

$$r_{(i,j)}(x) = \alpha^i \delta^j r_C(0) + (r_{min} - \alpha^i \delta^j r_C(0))x . \tag{3.2}$$

These assumptions are summarized graphically in Figure 3.3. The social dilemma was satisfied because $r_D(x) > r_C(x)$ for all $0 \leqslant x < 1$, and $r_z(0) > r_z(1)$.

This fitness model has two important differences from otherwise similar models (Hauert et al., 2006; Killingback et al., 2006; Parvinen, 2011, 2013; Suzuki and Kimura, 2011). First, because this is a model that implicitly includes the cooperative partner, I did not assume that there needed to be more than one *C* for any payoff to be generated.

Second, all sub-populations had different maximum growth rates but the same d_{max}. (Note that this is a simplification, as Figure 3.2 clearly shows that evolved types die more slowly than ancestral types.) Thus, the difference in fitness between any two types decreased as the environment became less favorable (i.e., more cheater-dominated), and the full difference between cooperators and cheaters was only achievable in the most favorable conditions for growth (i.e., cooperator-dominated). This is more bioloogically realistic, at least in the case of the yeast model.

Symbol	Description
x	Cheater frequency.
$r_z(x)$	Growth rate of sub-population *z*.
d_{max}	Maximum death rate
α	Fitness advantage of *D*over *C*.
δ	Fitness advantage of *C** over *C*.

Table 3.1: Summary of variables used in describing the model.

Dilution and migration

During biological experiments, cocultures were periodically diluted with fresh minimal media to ensure that they were only limited for the nutrients provided by cooperation. To model this carrying capacity, and to prevent population sizes from becoming too large, which may have resulted in computational overflow errors, I included dilution in the model. I chose a carrying capacity of 10^7 cells. Any population in a metapopulation achieving this size would trigger a global two-fold dilution of all populations in that metapopulation. This differed from my experiments, where dilutions

were performed on a tube-by-tube basis. I chose global dilution to prevent the "penalty" imposed by a limited population size on the most successful populations. In addition, global dilutions would speed up the dynamics of the simulation, by driving populations that were growing very slowly extinct. I also tried periodic dilution of all populations, which may be more realistic, and found the results to be qualitatively similar.

Migrants were sampled from each population and traveled as a group (or "propagule" (Wade, 1978)). I modeled both global and local migration on a square lattice with periodic boundary conditions. Global migration allowed migrants to end up in any location except the one from which they originated. Local migration allowed migrants to travel to one of the 8 nearest neighbors around their original location.

For both dilution and migration, transitions from one cell to zero were handled probabilistically. If one cell remained in a population before dilution/migration, the dilution/migration amount was interpreted as the probability that the remaining cell would be diluted to nothing or migrated to a new location. For example, if one cell remained in a population, at the next two-fold dilution, this cell would be diluted to nothing with a probability of 0.5; otherwise, one cell would remain after dilution.

Algorithm

A single time step of the simulation proceeded as follows. First, the size of all populations would be checked to see if a dilution would occur at this time step. For each location containing cells, migrants were collected and new locations for them were

assigned. Dilution would occur if necessary, and then the population would grow. The new population size was calculated as

$$N_t = N_{t-1}\exp(r_z(x)), \quad (3.3)$$

where N_{t-1} is the population size of that type in the previous time step. Populations containing less than one cell were set to zero and considered extinct. After all locations had been visited, migrants would be distributed to their new locations.

3.3. Initial conditions

Initially, I ran simulations in a small world containing 10 rows and columns. I reasoned that a small world would be less favorable for cooperation, as it would be easier for cheaters to find cooperator-dominated populations. Thus, if there were conditions favorable for cooperation in a small world, they would probably be favorable in large worlds, as well.

I varied the fraction of occupied locations from 1 to 0.25 in steps of 0.25. More empty locations should favor cooperation, since they represent chances for single cooperator migrants to escape cheaters. Each population initially contained 10^5 cells, with a 1:1 *C*:*D* ratio. The number of *C** and *D** added to each population was drawn from a Poisson distribution with a mean 1, 10, or 100 mutants expected in each population. Populations lacking mutants were run as controls. I covered a range of migration rates (10^{-1} to 10^{-12} hr^{-1}, in 10-fold steps) and d_{max} values (1 to 10^{-5}, in 10-fold steps). I set the baseline growth rate of *C*, ($r_C(0)$, to $\ln(2)/2 \approx 0.69$ hr^{-1} (i.e., a two

hour doubling time). Each simulation was run with one minute time steps for 100,000 hours, and the state of the metapopulation was saved every 200 hours. In cases where neither cooperators or cheaters fixed, the final 5 cooperator frequencies (spanning 1000 hours) were averaged. Seventeen independent simulations were run for each condition, and the average of their final cooperator frequencies, including extinction and fixation events, is presented.

Parameter	Values explored
α	1.2
δ	2.5
d_{max}	10^{-5} to 1, 10-fold steps
Initial mutant frequency	0, 10^{-5}, 10^{-4}, 10^{-3}
Migration rate	10^{-12} to 10^{-1} hr^{-1}, 10-fold steps
Fraction initially occupied	0.25, 0.50, 0.75, 1

Table 3.2: Summary of values used in the model.

3.4. High rates of migration and death lead to cooperator domination

I set $\alpha = 1.2$, ten times greater than the experimentally-determined cheater advantage (Figure 2.1). When all locations were initially occupied by 1:1 C:D mixtures and no mutants were present, cooperation was disfavored over a majority of the parameter space (Figure 3.4A). For all but the highest values of d_{max} ($\leq 10^{-2}$), regardless of the migration rate, cooperator frequency asymptotically approached zero. When d_{max} was 0.1 the max growth rate of *C* or greater, and at all but the highest

migration rates ($\leq 10^{-3}$ hr^{-1}), the entire metapopulation went extinct.

Surprisingly, at high d_{max} (≥ 0.1) and high migration rates ($\geq 10^{-2}$ hr^{-1}) there was a possibility that cooperators could fix. At the highest d_{max} (1), and migration rate (10^{-1} hr^{-1}), cooperators fixed in every run. This initially seemed counterintuitive, since high migration rates should act to reduce variance between populations and favor cheaters. Careful examination of individual trajectories revealed that this occurs because of the increased variability introduced by a high death rate. Initially, cheaters rapidly rise in frequency, leading to near total extinction of the metapopulation. In a few locations, however, both cooperator and cheater populations are < 2 and a single cheater migrates, while the cooperator does not. This makes the cheater population < 1 and the cooperator population > 1. Thus, the cheater population is set to 0, and the cooperator population keeps growing. If the single cheater migrant finds an empty location, it drives itself extinct. If the single cheater migrant finds a location occupied by cooperators, the cycle happens again. On the other hand, if a single cooperator migrates, its population in the current location is set to zero, and the remaining cheater quickly goes extinct. If the single cooperator migrant finds an empty location, it grows until a cheater finds it. If it lands in a location with a cheater population (which is likely to be low, since most positions are low density or empty), then the cycle happens again.

Table 3.3 summarizes the possible outcomes when a single cooperator or cheater migrates into a location that is either empty, contains a cooperator, or contains a cheater. Two out of three types of location guarantees that cooperators will increase in number,

while a cheater can only land in a location with a cooperator to increase its population size.

Migrant	Occupant of new location		
	Empty	**Cooperator**	**Cheater**
Cooperator	Growth	Growth	Extinction or migration
Cheater	Extinction	Extinction or migration	Extinction

Table 3.3: Possible outcomes when single cooperators and cheaters migrate to new locations that are either empty or contain one individual of the indicated type. 'Growth' indicates that the population size of that location will increase. 'Extinction' indicates that the population size of that location will go to zero. 'Migration' indicates that the indicated type may migrate to a new location.

This inherent asymmetry allows cooperators to gain a foothold at the brink of population extinction and rapidly increase in number. If cheaters do not go extinct, the process happens again, but each cycle favors cooperators because of their ability to grow in empty patches.

3.5. The adaptive race and empty space help cooperators beat cheaters

When the frequency of stress-adaptive mutants was on average 1 mutant per population, cooperators dominated at migration rates of 10^{-6} hr^{-1} or below at all but the highest value of d_{max} (Figure 3.4B). At this frequency, many populations had no mutants, while many contained either a *C** or a *D**. Variance between populations was large (Coefficient of variation of each population in the metapopulation = CV = 141%).

At low migration rates, populations containing a *D** will rapidly go extinct, while populations initialized with a *C** will rapidly rise in population size and dilute out populations that do not contain any mutants. This opens up empty locations, which, as described above, tend to favor cooperation.

An initial mutant frequency of 10^{-4} (Figure 3.4C) represents the estimated frequency of mutants in the experiments performed in Chapter 2. On average, 10 mutants were expected per population of 10^5 cells, so the variance between initial *C** and *D** population sizes was still relatively large (CV = 45.7%), but it was very unlikely to have a population initialized without at least one *D** (*P* = 0.007). Now at lower migration rates a population was more likely to be driven extinct by *D** before a lone cooperator had a chance to migrate into a recently emptied location. However, in some populations the initial ratio of *C** to *D** would be high by chance. The number of *C** in a population could be substantially higher than the number of *D**. This would cause extinction of *D** before *C** in a population, and would occasionally lead to the overall domination of *C**, even at low to moderate migration rates ($\leq 10^{-6}$ hr^{-1}).

When the mutant frequency was 10^{-3} (Figure 3.4D), so that 100 mutants were expected per population, the odds of not having *D** were impossibly small (< 1 in $2x10^{22}$). The expected variance of a sub-population was also much smaller (CV = 14%). At this mutant frequency, outcomes became deterministic and the cooperators were worse off than in the no mutant case.

Interestingly, the presence of initially empty locations was enough to allow

cooperators to dominate a wide swath of parameter space—even in the absence of mutants (Figure 3.5, left column). At 75% occupancy, cooperators always dominated at moderate migration rates ($\leq 10^{-7}$ hr^{-1} and $\geq 10^{-9}$ hr^{-1}) and all but the most extreme death rates. Cooperators dominated at the highest death rate when the migration rate was also high ($\geq 10^{-6}$ hr^{-1}). Decreasing occupancy to 25% produced qualitatively similar results.

The effect of initially empty locations on cooperator dominance was dependent upon the frequency of stress-adaptive mutants. At the lowest mutant frequency (Figure 3.5, second column from left; 1 mutant expected per population), the additional variance added by empty space had little effect on the overall outcome. When 10 mutants were expected per population (third column), the addition of initially empty locations had a large effect, with the range of deterministically cooperator-dominated parameters increasing with decreasing initial occupancy. Finally, the case of 100 mutants expected per population (right column) looked like the no mutant case, except cooperators dominated over a narrower range of parameters.

The results from simulations run with local migration (Figure 3.6) were qualitatively similar to the results for global migration, but, as expected, cooperators dominated a larger parameter range.

3.6. Testing the model with an experimental metapopulation

My modeling results suggested that the presence of rare mutants that have a fitness advantage upon exposure to environmental stress, coupled with low to moderate

rates of migration would allow cooperators to defeat cheaters in a metapopulation. To test this idea, our lab is collaborating with the Kerr Lab at the University of Washington. Our system uses the bacterium *P. aeruginosa*, which is an opportunisitic pathogen associated with chronic infections in immunocompromised humans, such as burn victims and cystic fibrosis (CF) patients (Jimenez et al., 2012). *P. aeruginosa* cooperates through the quorum-sensing (QS) mediated production of common goods such as siderophores for iron uptake and proteases for digestion of extracellular proteins (Cornelis, 2010; Darch et al., 2012; Diggle et al., 2007; Griffin et al., 2004; Jimenez et al., 2012; Kümmerli and Brown, 2010; Kümmerli et al., 2009). In this system, there are two types of cheaters. The "specific" type can respond to the QS signal, but lacks the genes responsible for production of a particular common good, such as parts of the biosynthetic pathway responsible for siderophore or protease production. The "general" type lacks one or more components of the QS pathway itself, in which case it cannot produce any costly public good. Both types have been found in the lungs of CF patients (Jimenez et al., 2012; Schaber et al., 2004).

We are studying the "general" type of cheater, which lacks *lasR* and *rhlR*, the two major transcription factors necessary for QS (Jimenez et al., 2012). This cheater has a ~17% fitness advantage over the QS-competent cooperator (Figure 3.7A). Cheaters and cooperators were mixed at a ratio of 3:1 and distributed to the wells of a microtiter plate. In the control condition they grew for 48 hrs in media that contains casein–a milk protein that requires elastase to be used as a nutrient source–as the dominant nutrient source. In

the experimental condition, the antibiotic streptomycin was added to the media at a concentration sufficient to kill all but a few cells. After growth, the wells of each microtiter plate were mixed. At this point, the ratio of cooperator to cheater was determined by plating on skim milk agar plates. Each mixture was then diluted into wells of new microtiter plates. This completed one "transfer." Four such transfers were performed. Because of the mixing step, any positive assortment that may have developed during the growth phase was erased during the transfer, a situation that is highly unfavorable to the maintenance of cooperation.

As expected, without antibiotic stress, cheaters deterministically took over (Figure 3.7B, black circles). However, in the presence of antibiotic, cooperators maintained a high frequency through all five transfers (Figure 3.7B, orange plus-signs). These preliminary results demonstrate that adaptation to environmental stress allows cooperators to achieve and maintain a high frequency in a metapopulation that was only periodically structured.

3.7. Materials and Methods

Determining trajectory of growth rate as a function of cheater frequency

At each time point, the density (cells/ml) of each subpopulation ($R_{\rightarrow A}^{\leftarrow L}$, $Y_{\rightarrow L}^{\leftarrow A}$, or $C^{\leftarrow L}$) was determined by spiking a known amount of beads into each sample, as previously described (see Section 2.6). These data were then converted to cumulative

densities by multiplying by the cumulative coculture dilution amount. For all time points but the first, coculture growth rate was calculated as $\ln\left(N_t / N_{t-1}\right)/\left(\Delta h\right)$, where N_t, N_{t-1} were the cumulative live population sizes at time points t and $t-1$, respectively, and Δh the hours elapsed between these two time points. Starting at the third time point, sub-population growth rate for time point t was estimated by averaging the growth rates calculated at $t-1$, t, and $t+1$.

Data from each cheater-containing coculture was concatenated to cheater-free data and fit using linear regression. Cocultures "B", "C", "D", and "F" were tested as a group and found to not be significantly different from one another. These data were combined and tested against cocultures "A" and "E." The interaction between cheater frequency and coculture was significantly positive (slope significantly less) for coculture "A" (P<0.007) and significantly negative (slope significantly more) for "E" (P<10^{-4}), than the rest of the cocultures.

Metapopulation simulations

The simulation was implemented in the Java programming language. Simulations were run on the high-performance scientific computing cluster at the Fred Hutchinson Cancer Research Center. Data was analyzed using the R programming language (R Development Core Team, 2012).

Experimental metapopulation

All experiments were performed by Sarah Hammarlund and Katie Dickenson in the laboratory of Dr. Ben Kerr at the University of Washington. Wild-type or Δ*lasR* Δ*rhlR* cells were scraped from the freezer into 0.4% casamino acid minimal media (CAA) and placed in a shaking incubator at 37 °C for 24 hrs. Each tube was diluted 1:100 into fresh CAA and grown for an additional 24 hrs. These tubes were diluted 1:10 into the wells of a microtiter plate (one for each strain) containing CAA and grown for an additional 24 hrs. The wells of the microtiter plate for each strain were pooled, mixed together, and diluted 1:10 into microtiter plates containing CAA with or without 200 μg/ml streptomycin. Each microtiter plate was set up in triplicate. A portion of each reservoir was serially diluted and plated on skim milk agar to check the frequency of mutants, which cannot form a "halo" of digestion around their colonies. This point is "Transfer 0" in Figure 3.7.

After 48 hrs of growth, the contents of each microtiter plate were pooled in separate reservoirs. Each reservoir was mixed and diluted 1:10 into new microtiter plates containing the appropriate media. To track cooperator frequency, a portion of the contents of each reservoir were plated as described above.

3.8. Figures

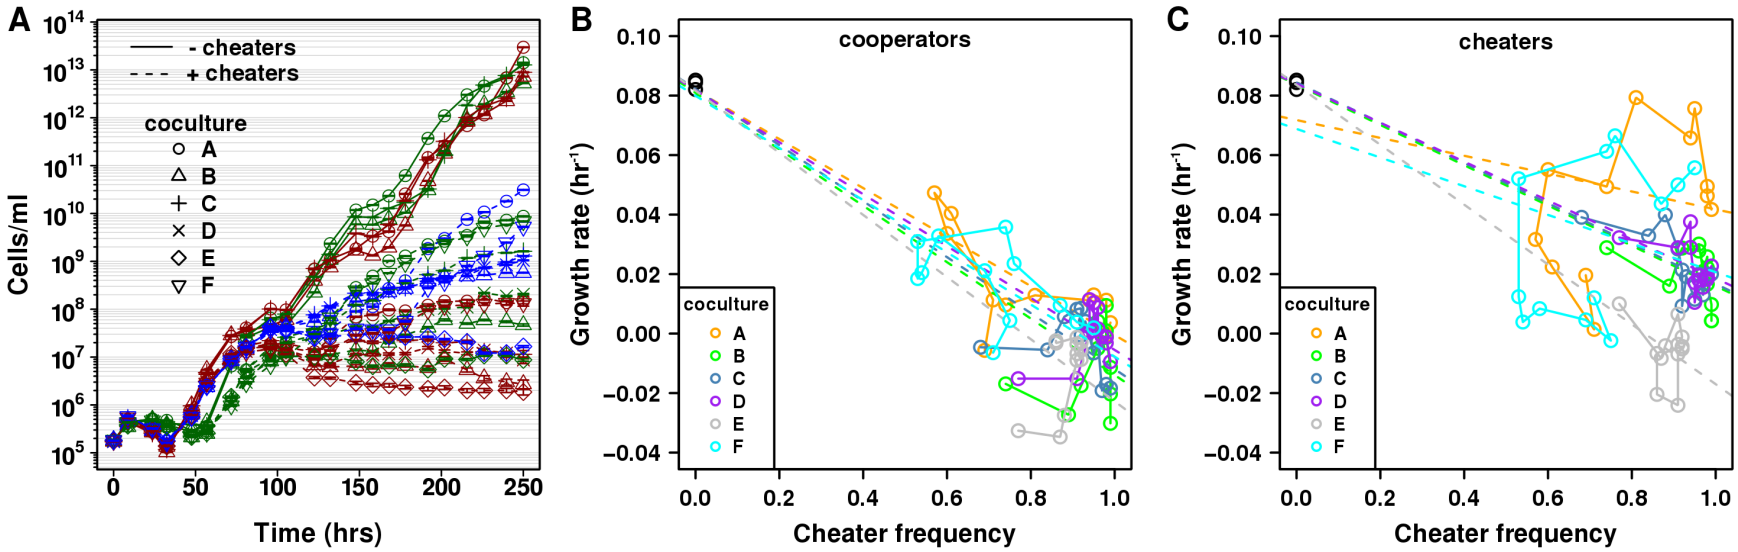

Figure 3.1: The fitness of cooperators and cheaters decreases with increasing cheater frequency. $R^{\leftarrow L}_{\rightarrow A}$ (red), $Y^{\leftarrow A}_{\rightarrow L}$ (green), and $C^{\leftarrow L}$ (blue) were washed in SD, mixed at a ratio of 1:1:0 (solid lines) or 1:1:1 (dashed lines) and diluted to a final concentration of 1.7×10^5 cells/ml/strain into SD. **A**) Two cheater-free and six cheater-containing co-cultures were monitored using flow cytometry (see Section 3.7) twice a day. For each co-culture in **A**, plots of cooperator (**B**) or cheater (**C**) growth rates versus cheater frequency were generated as described in Section 3.7.

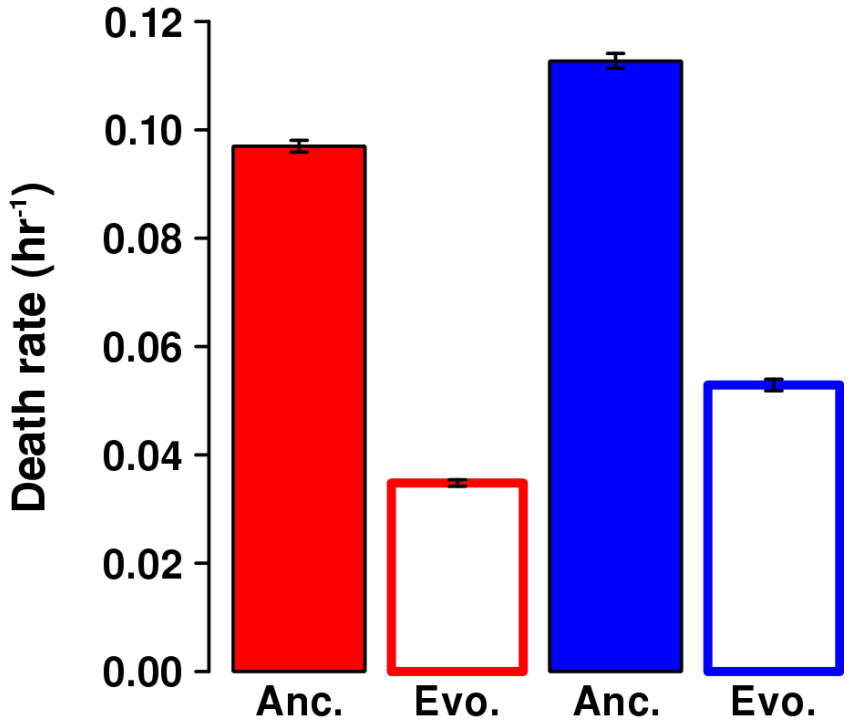

Figure 3.2: Evolved strains die more slowly than their corresponding ancestors. Representative strains of ancestor (WY950; solid red) and evolved (CT8; empty red) cooperators and ancestor (WY962; solid blue) and evolved (CT12; empty blue) cheaters were struck from the freezer onto YPD. Single colonies were picked and grown to saturation in YPD. These overnights were diluted into SD + 164 μM lysine and grown to exponential phase (~16 hrs), at which point they were washed 3x in SD and resuspended in SD. Washed cells were diluted into the wells of a microtiter plate containing SD to a final concentration of $5.1x10^5$ cells/ml. The fluorescence intensity at four locations per well was recorded every 30 min for ~ 100 hrs, and analyzed using a custom ImageJ macro. Death rates were multiphasic; the initial fast death rate is reported.

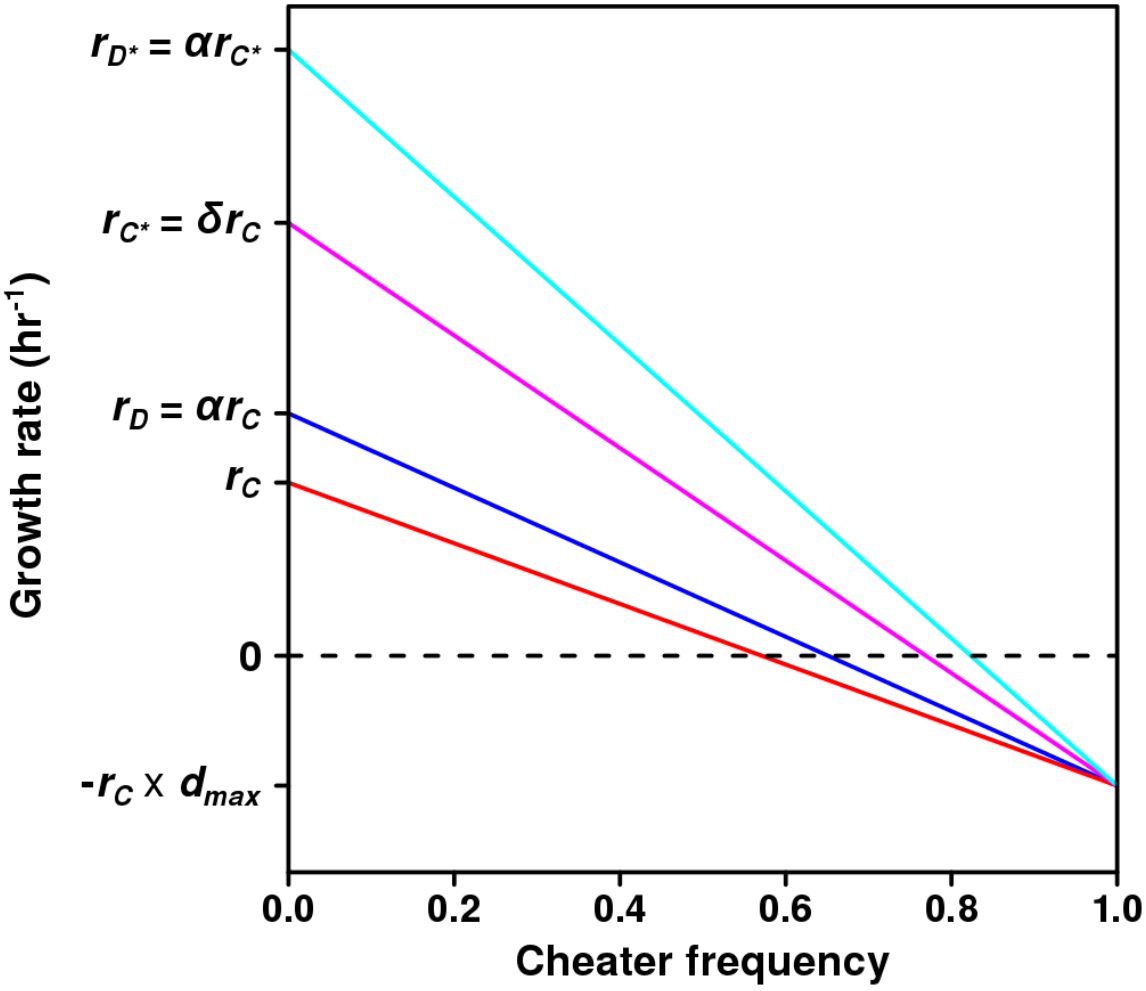

Figure 3.3: Summary of fitness relationships used in the model.
Growth rate (fitness) is a linear decreasing function of cheater frequency. Cheaters (*D*, *D**) have an advantage α over their corresponding cooperators (*C*, *C**). The evolved cooperator *C** has an advantage δ over the ancestor cooperator *C*. The maximum death rate is identical for all populations, and is a negative fraction (d_{max}) of the maximum growth rate of *C*.

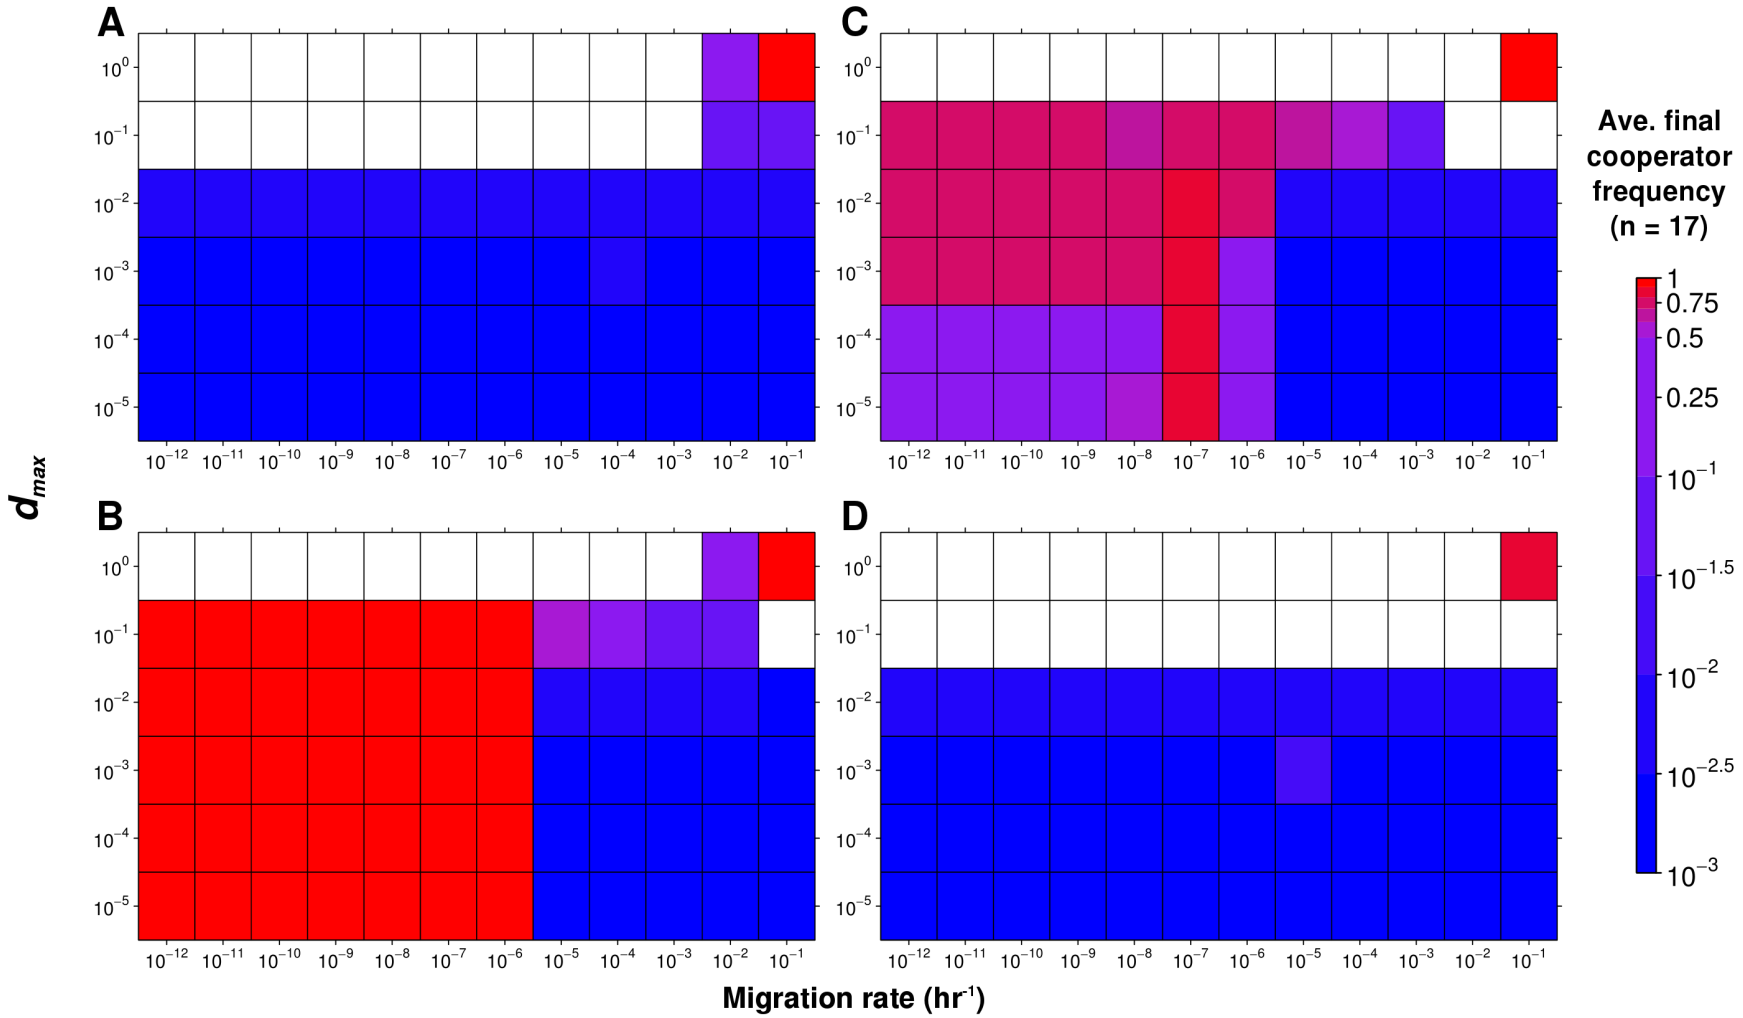
A
B
C
D
d_{max}
10^{0}
10^{-1}
10^{-2}
10^{-3}
10^{-4}
10^{-5}
10^{-12} 10^{-11} 10^{-10} 10^{-9} 10^{-8} 10^{-7} 10^{-6} 10^{-5} 10^{-4} 10^{-3} 10^{-2} 10^{-1}
Migration rate (hr^{-1})
Ave. final cooperator frequency (n = 17)
1
0.75
0.5
0.25
10^{-1}
$10^{-1.5}$
10^{-2}
$10^{-2.5}$
10^{-3}

Figure 3.4: The adaptive race increases cooperator-dominated parameter space. Simulation results for initially fully-occupied metapopulations with global migration. Each box represents a metapopulation run with one combination of d_{max} and migration rate parameters. Empty boxes indicate every run resulted in population extinction. Each combination was run 17 independent times (i.e., different seeds for the random number generator). Mutant frequencies of 0 (**A**), 10^{-5} (**B**), 10^{-4} (**C**), or 10^{-3} (**D**) corresponding to 0, 1, 10, or 100 mutants on average in the initial populations, were tried.

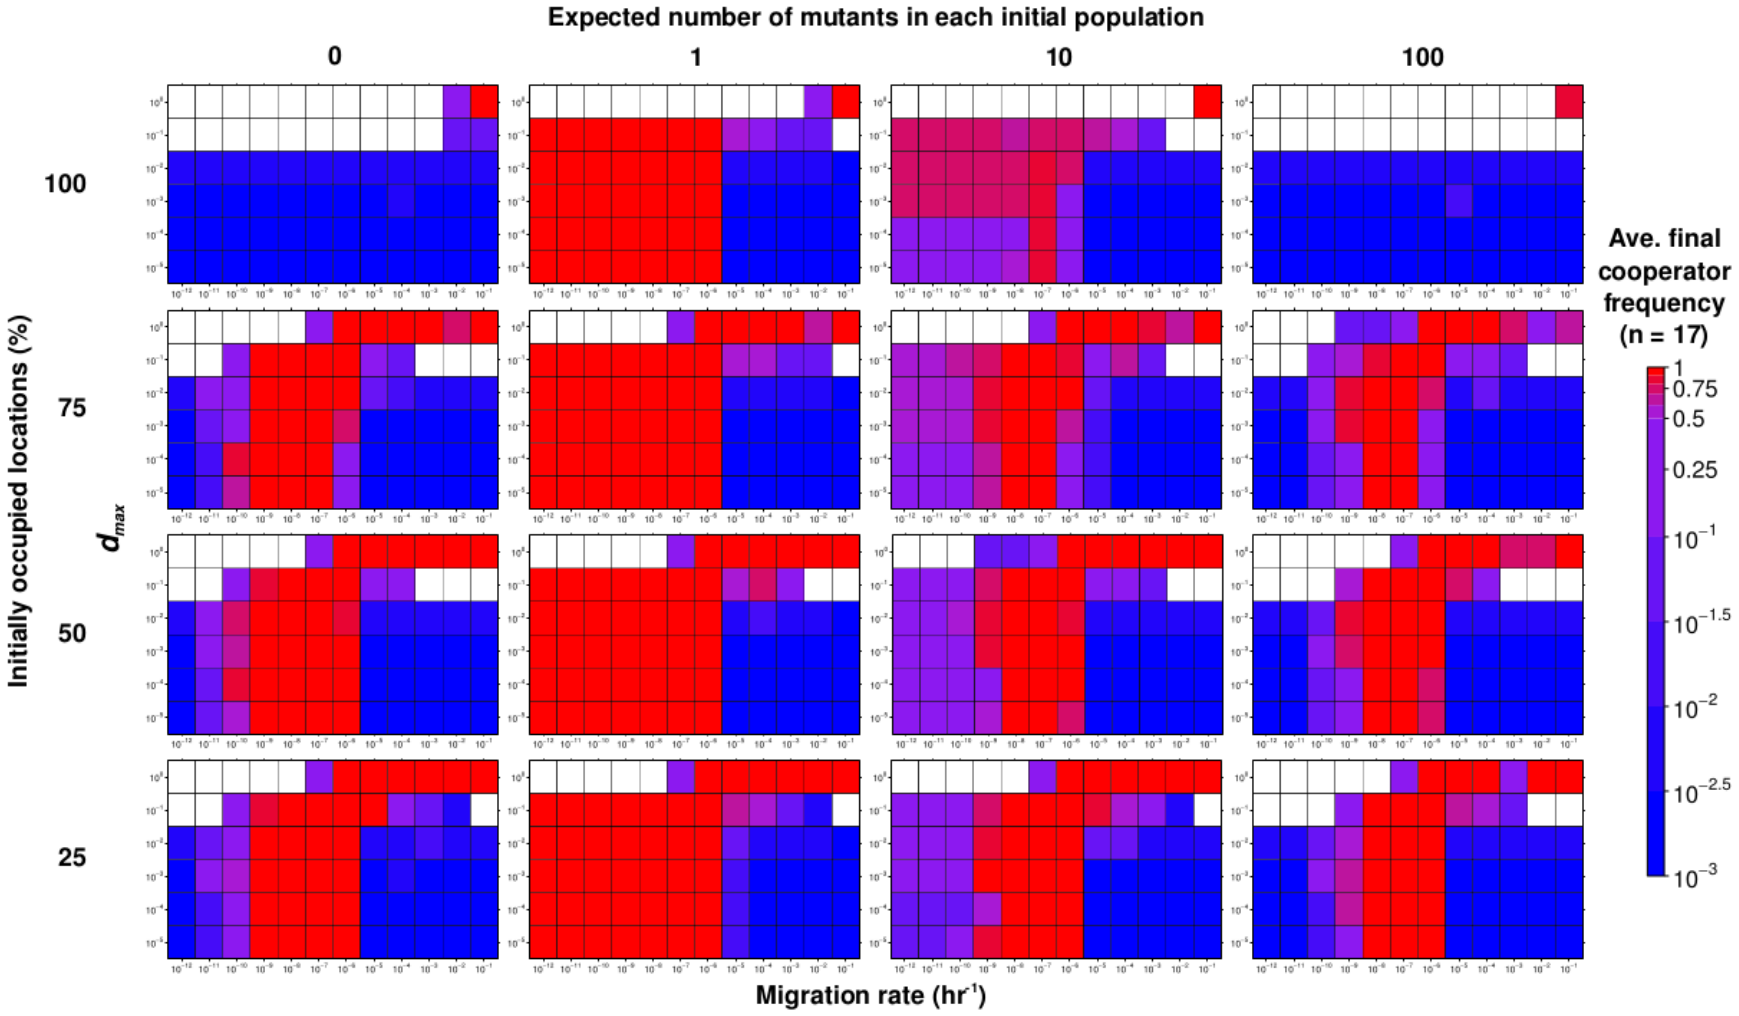
Expected number of mutants in each initial population
0
1
10
100
Initially occupied locations (%)
100
75
50
25
d_{max}
Migration rate (hr^{-1})
Ave. final cooperator frequency (n = 17)
1
0.75
0.5
0.25
10^{-1}
$10^{-1.5}$
10^{-2}
$10^{-2.5}$
10^{-3}

Figure 3.5: Empty space synergizes with the adaptive race to help cooperation: global migration.
Additional simulations were run with metapopulations that varied in their initial occupancies. Axes are the same as in Figure 3.4.

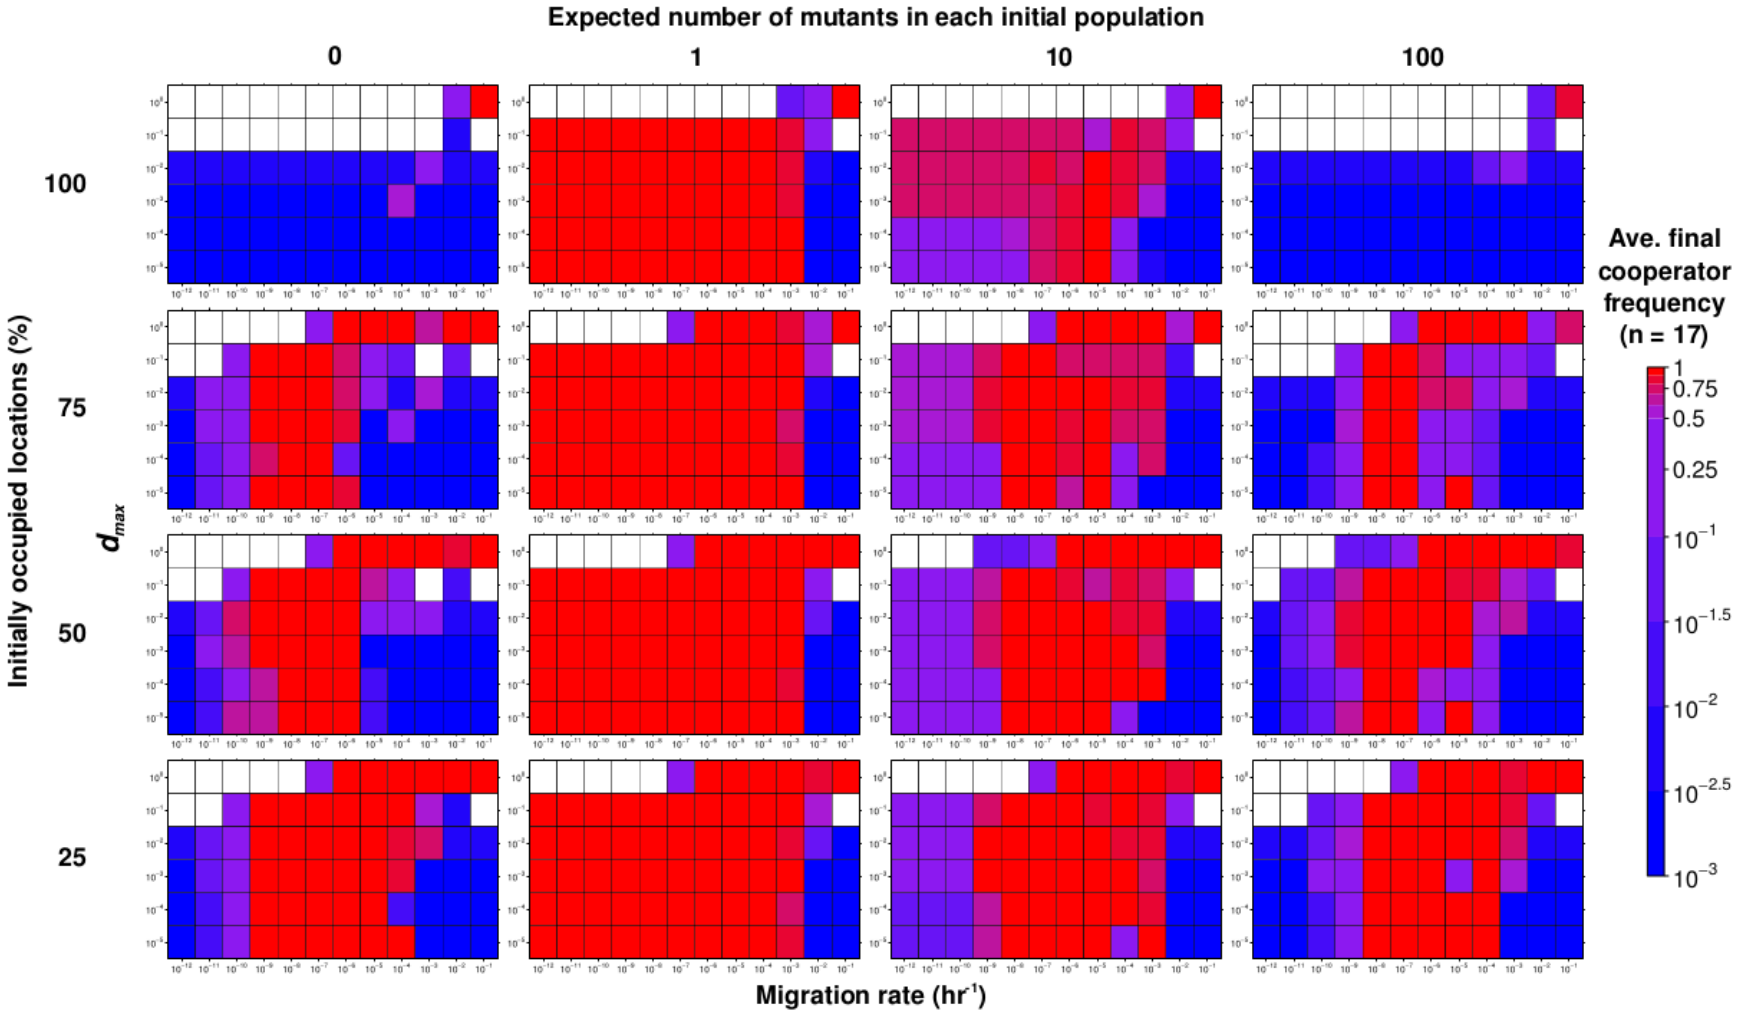
Expected number of mutants in each initial population
0
1
10
100
Initially occupied locations (%)
100
75
50
25
d_{max}
Migration rate (hr^{-1})
Ave. final cooperator frequency (n = 17)
1
0.75
0.5
0.25
10^{-1}
$10^{-1.5}$
10^{-2}
$10^{-2.5}$
10^{-3}

Figure 3.6: Empty space synergizes with the adaptive race to help cooperation: local migration.
Simulations otherwise identical to those in Figure 3.5 were run with 8-neighbor local migration. Axes are the same as in Figure 3.4.

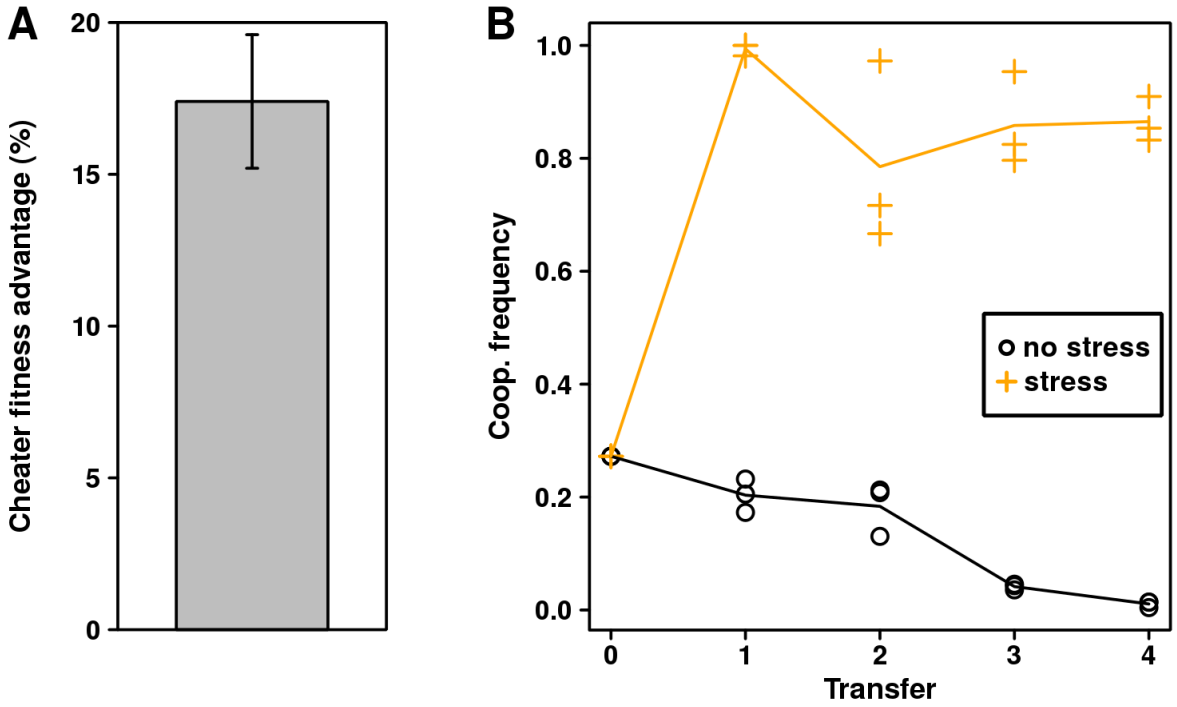

Figure 3.7: Environmental stress helps cooperators defeat cheaters in an experimental metapopulation.
A) Quorum sensing deficient (Δ*lasR* Δ*rhlR*) cheaters have a fitness advantage over wild-type cooperators. Data is mean of three replicates, error bar is twice the standard deviation. **B**) Cheaters and cooperators were inoculated into wells of a microtiter plate at an initial ratio of 3:1. Three replicates were initiated. The media either did not (black circles) or did (orange plus) contain the antibiotic streptomycin. At the end of each 48 hour growth period, all wells were pooled, mixed, and diluted into the wells of a new microtiter plate. During each transfer, a portion of the pooled sample was diluted and plated on skim milk petri dishes to check the frequency of cheaters. Data for all three replicate microtiter plates is shown. Lines connect replicate means.

Chapter 4. Environmental stress induces non-heritable variation in otherwise clonal populations

In the previous chapters, I showed how variation generated by mutation, migration, and extinction can help cooperators outcompete cheaters despite the fact that cheaters have a fitness advantage over cooperators. Interestingly, the plating assay I used to quantify the ability of evolved strains to grow at low levels of lysine revealed another form of variation that expressed itself during nutrient limitation. Essentially, each strain had a characteristic fraction of cells that could grow at limiting amounts of lysine. This heterogeneity had the potential to provide an additional, non-genetic source of variation that could fuel the adaptive race. In addition, since nutrient limitation exerts a strong influence on the evolution of organisms, regardless of whether cooperation or cheating is involved, I thought it worthwhile to study this phenomenon outside of the context of cooperation and cheating *per se*.

4.1. Growth on limiting lysine reveals non-genetic heterogeneity

I began by performing more plating experiments on a lysine-auxotrophic strain that had no changes besides a deletion of *LYS2* (and *HO* which is common to all laboratory strains). This strain could grow on 1 μM lysine ~ 2% of the time (95% CI = 1.4 – 3.1%). Because the fraction of cells that could form microcolonies was much

higher than what one might expect from mutations, I reasoned that growth on 1 μM lysine was not a genetically stable phenotype. To test this, I replica plated one of these plates onto YPD to allow any microcolonies to grow to a larger size. I compared this plate to the original to ensure that I was only picking colonies that had original formed microcolonies on the 1 μM plate. I then tested two of these colonies for growth on glycerol, to ensure that they were not petites. These two isolates were then pre-starved and re-plated on 1 μM lysine plates. If this phenotype were heritable, I would expect nearly 100% of the colonies to grow on 1 μM lysine. However, the frequency of microcolony growth was not statistically different from the ancestor strain (2.5%, 95% CI = 0.8 – 5.7%, Figure 4.1), confirming that this was an instance of non-genetic heterogeneity.

4.2. Population growth rate is dependent upon nutrient concentration

Having established growth on limiting lysine was a non-genetic phenotype, I wanted to see if it had any implications for the growth of the population as a whole. To do this, I followed the methodology of Novick and Szilard (1950) who asked a similar question about the response of populations of tryptophan-requiring bacteria to limiting tryptophan. Cells grown in different concentrations of lysine were diluted and plated onto rich media at defined intervals. These populations displayed a lag time, followed by exponential growth or death (Figure 4.2A). From counting the number of colonies on each plate, I could estimate the growth rate. Plotting growth rate versus lysine

concentration reveals saturating kinetics (Figure 4.2B). This relationship has historically been modeled using the "Monod Equation":

$$r = \frac{V_{max} \times S}{K_m + S}, \tag{4.1}$$

where r is the growth rate at concentration S of the limiting nutrient, V_{max} is the maximum growth rate achievable by the organism, and K_m is the concentration of limiting nutrient required to achieve half-maximal growth (Ferenci, 1999; Monod, 1949).

Between 0 μM and 1 μM lysine, populations die at a rate that increases with decreasing amounts of lysine. This implies that cells require a non-zero "maintenance concentration" (S_m) of lysine to survive. Incorporating this into Equation (4.1) gives

$$r = \frac{V_{max} \times (S - S_m)}{K_m + (S - S_m)}. \tag{4.2}$$

The results of fitting the data in Figure 4.2B to Equation 4.2 are presented in Table 4.1, in the row labeled "Plating." Under these conditions, one might naïvely expect th "Monod" constant, K_m , to equal the "Michaelis-Menten" constant of LYP1p ($K_{m,LYP1}$), the high-affinity lysine permease (Regenberg et al., 1999b; Sychrova and Chevallier, 1993). This is because the growth of these lysine-requiring cells is limited by lysine, which seems to suggest that the concentration of lysine satisfying half-maximal lysine uptake will also cause half-maximal growth rate. Published values

of the $K_{m,LYP1}$ vary by almost an order of magnitude, even when the same strain is used (Table 4.2). Even the lowest value is an order of magnitude greater than the upper bound of 95% confidence interval for our estimate of K_m. One way to account for this discrepancy is to consider K_m a parameter that is a function of time (Ferenci, 1999). This would be expected for transporters whose expression is inducible, as it is with LYP1p (Barnett, 2008), as lysine availability would increase with increasing LYP1p surface expression.

Experiment	V_{max} (hr^{-1})	K_m (μM)	S_m (μM)
Plating	0.40 (0.26 – 0.56)	0.39 (0.11 – 1.4)	0.30 (0.11 – 0.77)
Microtiter (all data)	0.52	1.4	0.39
Microtiter (without 0 μM data)	0.41	0.33	0.38
Microtiter (without 0.25 μM data)	0.42	0.69	0.23

Table 4.1: Estimates of parameters of the Monod Equation.
Data is from the population-level plating assay (Figure 4.2B) or from rates of fluorescence intensity increase in the microtiter assay (Figure 4.3D). Because the data collected at 0 and 0.25 μM lysine have large errors, parameter estimates without this data is included, as well. Values in parentheses indicate the 95% confidence interval.

Strain	K_m (μM)	Reference
ATC9896	12	(Morrison and Lichstein, 1976)
Σ1278b	25	(Grenson, 1966)
Σ1278b	78	(García and Kotyk, 1988)
Σ1278b	19	(Sychrova and Chevallier, 1993)

Table 4.2: Some published values of the K_m for Lyp1p.

4.3. Limiting lysine increases colony size heterogeneity

Most small colonies are respiratory-deficient

During the course of the plating assay, I noticed an increasing heterogeneity of colony size. This heterogeneity was more pronounced and increased more rapidly in lower lysine conditions (Figure 4.3A).

Starvation for amino acids, vitamins (Nagai, 1969), nucleotides (Barclay and Little, 1978; Nagai, 1969), and fatty acids (Marzuki et al., 1974) increases the proportion of respiratory-deficient "petite" cells whose mitochondrial DNA is defective (ρ^-) or lacking completely (ρ^0). These cells grow more slowly and therefore create smaller colonies. In addition, S288C is known to produce a relatively large number of petites compared to other *S. cerevisiae* strains (Dimitrov et al., 2009).

To test whether the small colonies were petite, I transferred individual colonies to plates containing the non-fermentable carbon source glycerol (YPG) using a toothpick. The same toothpick was used to to transfer cells onto YPD. I considered any colony that grew on YPD and not YPG petite. This experiment showed that most (78%, $n = 50$) but not all small colonies were petite (Figure 4.3B), suggesting that dysfunctional mitochondria could not completely explain the observed heterogeneity.

4.4. Heritable decrease in growth rate observed in respiratory-proficient and respiratory-deficient cells

The observed colony size heterogeneity could be due to differences in lag time or

growth rate between the cells that initiated the colonies. In addition, these differences could be heritable or non-heritable.

To test for heritable changes, I picked individual colonies of various sizes (roughly characterized as “small”, “medium”', or “large”) that had grown up on YPD plates after being starved of lysine for 97.5 hours, inoculated them into fresh YPD media, and allowed them to grow to saturation. Their ability to utilize glycerol was tested. All of the large and medium colonies and none of the small colonies could respire glycerol.

These cultures were used to test growth rates. Each culture was used to inoculate eight wells of a 96-well plate, and their growth rates in SD supplemented with non-limiting lysine were determined (Figure 4.4). As expected, the large colonies grew at wild type rates, while the small colonies grew at rates typical of petites. Interestingly, cultures inoculated from a medium sized colony grew at an intermediate rate. Thus, this clone represents a mutant with a growth deficiency not related to respiratory deficiency.

4.5. Heritable and non-heritable decrease in death rate observed after prolonged lysine starvation

To investigate possible heterogeneity in cell death rate, I performed more detailed experiments in SD lacking any lysine (Figure 4.5).

Death rate is multiphasic

It is possible that death rate is density-dependent. For example, cellular

components from recently lysed cells could be used as nutrient for other cells, and the available amount of nutrient would increase with cell density. I observed that death is independent of density at or below 10^6 cells/ml (Figure 4.5A).

If cells taken from an exponentially growing culture with lysine are washed and resuspended in SD without lysine, a characteristic sequence of events takes place. Initially, a 3-4–fold increase in population is observed while cells use up vacuolar reserves of lysine. After ~20 hours, viability of the culture drops dramatically. If each cell responded identically to starvation–i.e., had a fixed probability of death–then the drop in CFUs should be characterized by monoexponential decay. However, our results are not compatible with this hypothesis (Figure 4.5B). After an initial rapid drop in CFUs, a transition to a slower death rate is observed at ~50 hours. Surprisingly, viable colonies were observed out to 250 hours (Figure 4.5B), suggesting a third, even slower rate of death.

There are at least two ways to interpret these results. The probability of death could be a function of time, so that the population size (N) changes as

$$N(t) = e^{k_1(t_1-t_0)} e^{k_2(t_2-t_1)} e^{k_3(t_3-t_2)}, \tag{4.3}$$

where k_1, k_2, and k_3 are death rates that begin at times t_1, t_2, and t_3, respectively. Alternatively, subpopulations with different death rates could be present in the initial population, so that

$$N(t) = Ae^{at} + Be^{bt} + Ce^{ct}, \tag{4.4}$$

where *A*, *B*, *C* are the initial sizes of the hypothetical subpopulations, which have death rates *a*, *b*, and *c*, respectively.

Fitting the data to the subpopulation model (Equation 4.4) roughly 96% of the population die at a rate of 0.3 hr^{-1}, 4% die at a rate of 0.05 hr^{-1}, and 0.1% die at a rate very close to zero.

The "death-resistant" phenotype is heritable for petites

Both grande and petite cells (as verified by inability to grow on YPG, data not shown) were able to survive 250 hours of lysine starvation. To determine whether or not this "death-resistant" phenotype is heritable, colonies from the plates seeded at 250 hours after inoculation in SD were picked and grown to saturation in YPD. These overnight cultures were then used to seed wells of a microtiter plate containing SD.

I compared the death rates of five petites and one grande to the strain used to begin the experiment (WY928). Five positions in each well were monitored by automatic timelapse fluorescence microscopy over a period of 70 hours. I defined death as loss of fluorescence, as this corresponds to cell lysis using the Venus fluorophore (data not shown). Importantly, this definition of death is fundamentally different from the "ability to form a colony", which was used in the plating assay. For example, it is possible for a metabolically active cell with an intact membrane to fail to grow into a colony when plated on rich media. This cell would be considered "alive" in the fluorescence-based assay and "dead" in the plating assay. Therefore, I expected a lower

death rate than what was determined from the plating experiment.

The results indicate radically different death rates for the petites and grandes (Figure 4.5C). The grande strain displays similar kinetics to WY928. In fact, the initial death rate of this culture is slightly faster than that of WY928 (0.11 hr^{-1} 95% CI 0.10 – 0.12 hr^{-1} vs. 0.098 95% CI 0.086 – 0.11 hr^{-1}). Petites, on the other hand, initially die at rate greater than either grande (0.2 hr^{-1}; 95% confidence interval 0.17 – 0.23 hr^{-1}) but, starting around 30 hours, shift to a much slower death rate (0.005 hr^{-1}; 95% CI 0.001 – 0.008 hr^{-1}). This change is specific to petites from later plates, as petites starving for six hours have death kinetics indistinguishable from their grande counterparts (data not shown).

4.6. Individual cells behave differently in limiting lysine

Results from the plating assays indicated that individual cells were responding differently to limiting lysine conditions. To investigate how individual members of a population respond to limiting lysine, I set up a microtiter plate with wells containing various concentrations of lysine and added pre-starved cells. As in the plating assay, I picked a concentration of cells such that several rounds of doubling could occur before the concentration of lysine was significantly reduced. I used automated time lapse microscopy to record images every three hours in five positions of each well used in this study.

I found that counting cell number increase is not possible in a microtiter plate, as

cells quickly grow into three-dimensional clumps. To obtain growth rate information, I used a strain expressing the yellow fluorescent protein 'Venus' (Nagai et al., 2002) under the control of the Alcohol Dehydrogenase 1 promoter (ADH1p), which is expressed throughout fermentative growth (de Smidt et al., 2008), and used total fluorescence intensity as a measure of changing biomass (Figure 4.6).

Four mutually exclusive fates were observed for each cell in the microtiter plate. A cell could grow and divide, and its daughters could grow and divide, forming a colony. This is considered “growth.” A cell could lyse, which is a distinct morphological event in bright field and coincides with a dramatic loss loss of fluorescence signal. This is considered “death.” Some cells were observed to divide several times, while none of the newly-born daughters themselves divided. I consider this a form of “residual growth.” Finally, some cells did not divide or lyse over the course of the experiment (69 hours). I call this fate “static.” The fates of individual cells are summarized in Table 4.3. While most cells were able to grow in 4 μM and 2 μM lysine, only a few were able to grow in 1 μM lysine, and none were observed to grow in 0.5 μM or 0.25 μM lysine.

[Lysine] (μM)	Grew	Died	Static	Total
0.25	0	8	5	13
0.5	0	6	0	6
1.0	3	7	2	12
2.0	10	2	0	12
4.0	11	2	1	14

Table 4.3: Summary of individual cell responses to environments with different lysine concentrations.
Cells of the "residual growth" class are not included.

Even within a single field of view (670 μm x 896 μm), striking variability in cell fates was observed (Figure 4.6A). Somewhat surprisingly, the variability of growth rates for cells able to grow in 1 μM lysine (SD = 0.03, N = 5) was equivalent to that observed for 4 μM lysine (SD = 0.03, N = 5), despite the small sample size. Differences in cell-to-cell starvation tolerance is apparent in 0.5 μM and 0.25 μM lysine. Fast-dying, slow-dying, and death-resistant populations are evident in 0.25 μM lysine (Figure 4.6C).

As with the population study, growth rates of cells able to grow were dependent upon the lysine concentration (Figure 4.6D). The mean growth rate in 4 μM lysine (0.33 ± 0.001 hr^{-1}, n = 5) agreed well with the growth rates observed at 5 μM and 164 μM. In this case, however, the fit is poor unless either the 0 μM or 0.25 μM data is excluded.

4.7. Heterogeneous expression of high-affinity lysine permease is responsible for population heterogeneity

Making a functional Lyp1p-FP fusion

There are three possible routes of lysine uptake in *S. cerevisiae*: The general

amino acid permease Gap1p (K_m for lysine = 3.1 μM) (Grenson et al., 1970b), the arginine permease Can1p (K_m for lysine = 200 μM) (Grenson, 1966), and the lysine-specific Lyp1p (K_m 25 μM) (Grenson, 1966). Gap1p expression depends on the type of nitrogen supplied and varies from strain to strain (Magasanik and Kaiser, 2002).

To investigate whether *LYP1* alone is necessary for lysine uptake in S288C, I mated a *lys2*Δ *can1*Δ strain with a *lyp1*Δ strain. After sporulation, 10 tetrads were dissected onto YPD or media with proline as the only nitrogen source (YNB-Pro) to repress *GAP1* expression. In both cases, 10 out of 40 haploids were inviable. Genotyping the haploids revealed that the double mutant *lys2*Δ *lyp1*Δ and the triple mutant *can1*Δ *lys2*Δ *lyp1*Δ were absent. Thus, without a functional copy of *LYP1*, *lys2*Δ haploids are not able to germinate in rich media. This suggests that Lyp1p is the dominant mode of lysine uptake under both nitrogen-rich and nitrogen-poor conditions.

I made a C-terminal *LYP1*-Venus fusion in a *lys2*Δ background, using a method that removes the selectable marker (Reid et al., 2002). This was done to avoid premature mRNA degradation due to separating the *LYP1* gene from its 3'-UTR (Muhlrad and Parker, 1999). To better characterize any fitness cost associated with a fused Lyp1p, I compared the growth rates of the tagged strain to its untagged predecessor. In SD supplemented with lysine and uracil, the tagged strain had a growth rate that was 90% of the untagged strain (0.408 ± 0.002 hr^{-1}, n = 14 for untagged versus 0.365 ± 0.001 hr^{-1}, n = 14 for tagged), while no significant difference was observed in YPD (0.560 ± 0.002 hr^{-1},

n = 14 for untagged versus 0.566 ± 0.002 hr^{-1}, n = 14 for tagged).

Microfluidics reveals heterogeneous growth in limiting lysine

Results from population (Figure 4.2 and microtiter plate growth (Figure 4.6) showed that lysine does not limit growth rate down to at least 4 μM, and becomes limiting between 2 μM and 4 μM. Furthermore, results from the microtiter plate experiments revealed that individual cells differ in their responses to growth-limiting lysine concentrations. To better understand the cause of this heterogeneity, I observed expression and localization patterns of Lyp1p using a microfluidics device. Microfluidics provides a constant nutrient environment, and the design I use limits cell growth to a monolayer, facilitating accurate quantification of cellular parameters (Bennett and Hasty, 2009).

I observed growth of pre-starved cells in the microfluidics device at 164 μM ("non-limiting") and 1 μM ("limiting") lysine (Figure 4.7). In 164 μM lysine, all observed cells were able to grow and divide, while in 1 μM lysine, only 19% of cells were able to bud at all. I measured the time it took the initial cells to start budding after being exposed to nutrient ("lag time"), as well as the time between all budding events ("budding time"). I used the coefficient of variation (standard deviation / mean) to quantify variation in each population (Lande, 1977). Compared to cells exposed to 164 μM lysine, cells exposed to 1 μM lysine were more variable in both lag time (36% vs 21%; $P < 5x10^{-13}$ Fligner-Killeen test, Figure 4.7A) and budding time (47% vs 39%; $P <$

0.005, Figure 4.7B).

Of the cells that could bud in 1 μM lysine, only one was able to bud more than once. Comparison of this cell to its immediate neighbors (Figure 4.7C) revealed how growth was related to expression of Lyp1p. All cells initially started with bright vacuolar expression (Figure 4.7C, top row). After ~ 4 hours of exposure to 1 μM lysine, all cells lost this expression (Figure 4.7C, second row). Because the apparently constant pool of vacuolar Lyp1p reflects a steady-state between Lyp1p production, trafficking to the vacuole, and degradation, it's sudden disappearance is most likely due to cessation of new Lyp1p production. After ~5 hours, surface expression of Lyp1p could be seen on one of the four cells (Figure 4.7C, third row), and this cell proceeded to bud a daughter cell that also had bright surface expression of Lyp1p (Figure 4.7C, bottom row).

To get a better idea about the relationship between Lyp1p surface expression and growth in 1 μM lysine, I made a pedigree of all the cells in this region of the microfluidics device (Figure 4.7D). The left half of Figure 4.7D shows the pedigree for the cell that formed a microcolony in 1 μM lysine, while the right half shows the fates of all the other cells in the same region of the microfluidics device. The cell with bright Lyp1p expression produced two more daughters before lysing, and these daughters went on to bud themselves, eventually forming a microcolony. In general, the ability to express Lyp1p on the plasma membrane decreased with the number of generations, as did the ability to bud, suggesting that this expression was only a partially-heritable phenotype and was not genetic in origin. Lyp1p surface expression was highly

correlated with the ability to have daughters. However, one cell that was able to have daughters was uniformly dim Figure 4.7D, left panel). This might indicate that surface expression is not necessary for cells that were born in the nutrient-limited environment. On the other hand, in this and subsequent experiments (data not shown) I have never observed a cell that was initially loaded into the device divide if it did not express Lyp1p. I did observe two cells that were able to express Lyp1p on the plasma membrane but did not have any daughters. Taken together, it appears that Lyp1p surface expression is necessary but not sufficient for pre-starved cells to divide in 1 μM lysine.

4.8. Diploid prototrophs show growth heterogeneity during glucose limitation

Wild yeast are prototrophic. Thus, removing the ability to biosynthesize lysine is expected to be extremely disruptive to the cell, as has been observed with other auxotrophies (Boer et al., 2008). This is not only due to the absence of lysine when the cell requires it. As with many other biosynthetic pathways, lysine biosynthesis is controlled by feedback inhibition (Blickling and Knäblein, 1997). Since a lysine-starved cell never gets this feedback, it will continuously try to biosynthesize more lysine. This causes many potentially toxic biosynthetic intermediates to accumulate, which may also disrupt regulation of other metabolic pathways, such as general nitrogen metabolism and the TCA cycle (Ljungdahl and Daignan-Fornier, 2012). I was ultimately interested in how this heterogeneity could influence the evolution of a population facing an unfamiliar environment. However, because cells lack any regulatory capacity to deal

with lysine auxotrophy, it is possible that the heterogeneity I observed in lysine auxotrophs was due to the massive regulatory failure caused by lysine auxotrophy, and that such heterogeneity would not be observed during starvation for more "natural" nutrients such as glucose, nitrogen, and phosphate.

In addition, while haploids are used in laboratory research for their genetic simplicity, wild *S. cerevisiae* spends most of its life-cycle in the diploid state. After sporulation and germination, haploids quickly mate with one-another, with the majority of matings occurring within the same ascus (Tsai et al., 2008). Since ploidy influences the expression of many genes (Galitski et al., 1999), I wanted to ensure that I used the more natural diploid state for these experiments.

To address these issues, I used a wild vineyard strain of *S. cerevisiae* known as RM11 (Mortimer et al., 1994), the genome of which is readily available. I constructed a diploid prototroph of this strain (see Section 4.9 for details), as only haploid auxotrophic versions of this strain were available.

Unlike lysine, glucose is a nutrient as well as a growth factor, so its removal signals growth arrest for cells that have not passed the START checkpoint in the cell cycle (Gray et al., 2004). Cells that have passed START finish budding, and then arrest. To determine how long this took at the single-cell level, I observed cells in minimal media lacking glucose in the wells of microtiter plates using time-lapse microscopy. As expected, cells ceased dividing almost immediately. Partially-budded cells completed their cell cycle, while unbudded cells did not have any more daughters. After growth

arrest, cells can survive for many hours without glucose, utilizing glycogen and trehelose reserves (Gray et al., 2004). For the experiments described below, cells were starved in glucose-free media for 2 days. This was done to minimize the effect of different cells entering this growth-arrested state at different times.

In an attempt to find a concentration of glucose that resulted in a heterogeneous growth rate, I inoculated RM11 into the wells of a microtiter plate containing different amounts of glucose. Using time-lapse microscopy, I could get a rough estimate of the growth rate of each cell. My initial experiment revealed that the budding times of cells were similar to one another (within 0.3 hour) down to 0.30 mM glucose. At the next lowest concentration (0.16 mM), I observed a wide range of budding times (between 2.9 and 4.7 hours, data not shown). At 0.08 mM glucose or below, cells were not able to divide more than once. Based on these observations, I decided to use 27.5 mM (0.5%) glucose as the non-limiting condition, and 0.15 mM (0.003%) glucose as the nutrient-limited condition. Similar heterogeneity was observed at these concentration when the cells were starved for 1 day.

As before, I recorded lag (Figure 4.8A) and budding times (Figure 4.8B) for pre-starved cells exposed to non-limiting and limiting nutrient. I checked whether distance from the nutrient source influenced lag or budding time during nutrient limitation. While there was no evidence for this with lag times, there was evidence of a increase in budding time with distance when cells were > 17 μm from the nutrient source (data not shown). Thus, for the nutrient-limited conditions, data for lag times includes

all cells, while budding time data only includes cells ≤ 17 μm from the nutrient source.

Contrary to the case of limiting lysine, cells exposed to limiting glucose had a significantly *lower* CV in lag time (empty blue bars) compared to cells exposed to non-limiting (translucent orange bars) glucose (28%, n = 300 vs 37%, n = 136; $P < 3x10^{-16}$ Fligner-Killeen test), although the mean lag time was significantly higher for the nutrient-limited cells (10.7 vs 2.6 hrs, Wilcoxson rank sum test $P < 3x10^{-16}$).

The response of budding time to glucose limitation was similar to what I observed for lysine limitation. Overall, budding times were significantly longer in limiting glucose (mean of 3.4 hours vs 1.5 hours; Wilcoxson rank sum $P < 3x10^{-16}$). Variation also significantly increased ($P < 3x10^{-16}$, Fligner-Killeen test). The CV of budding time was 25% in limiting (n = 148; empty blue bars) versus 13% for non-limiting (n = 136; translucent orange bars) glucose. To see the effect of starvation alone on the variability of the cells, I also recorded budding times of unstarved cells exposed to non-limiting glucose (Figure 4.8B, solid black bars), and found that this population had a mean budding time of 1.3 hours and a CV of 7.6%. (I could not record lag times for this group, as they were not pre-starved and therefore never stopped budding.) This experiment showed that pre-starvation alone nearly doubles the CV in budding time of a population and significantly increases its mean budding time (Wilcoxson rank sum test, $P < 3x10^{-16}$). In addition, exposure to limiting glucose further doubles the CV in budding time.

The high variation in budding time of cells exposed to limiting glucose may have

been due to an amplification of initial variation in the population. For example, the initial mother cells might be at different stages of cell cycle when exposed to limiting glucose, which may cause them to vary in any number of factors that could influence the rate of growth upon exposure to limiting nutrient. Under this hypothesis, these factors would become less heterogeneous in successive generations of daughters cells exposed to an identical environment. To test this idea, I grouped the budding time data by generation (Figure 4.8C). I considered the cells initially present in the microfluidics device to be “mothers”; the daughters of these cells were considered 'daughter 1'; their grand-daughters 'daughter 2', etc.

Unstarved cells exposed to non-limiting amounts of glucose maintain a CV between 5-10% (Figure 4.8C, black). Pre-starved cells exposed to a non-limiting concentration of glucose (orange) initially display a higher variation in budding time of just under 15%, which is maintained for the first two daughters. Consistent with the hypothesis above, this additional variation disappears by the time the third daughter is born. On the other hand, pre-starved cells exposed to limiting glucose initially display an even higher CV of 23%, and maintain this variation for at least four generations. This shows that the variation is not some kind of “memory” effect, and is likely an inherent property of growth in limiting glucose.

In summary (Figure 4.8D), I observed an increase in variation in both auxotrophic and prototrophic strains when they were exposed to growth-limiting concentrations of a necessary nutrient. While variation in lag time only increased in the

auxotrophic cells, variation in budding time increased in the prototrophic cells to an even greater extent than it did in auxotrophic cells. Furthermore, in the prototrophic strain this variation was sustained for at least four generations, and cannot be fully explained by the pre-starvation treatment.

4.9. Materials and Methods

Starvation Treatment

I used a strain of the budding yeast *Saccharomyces cerevisiae* derived from S288C that is auxotrophic for lysine as a model system. *S. cerevisiae* can store large amounts of vacuolar lysine (Kim et al., 2003; Kitamoto et al., 1988), enough for several rounds of doubling after removal of exogenous lysine (data not shown). To remove stored lysine, exponentially growing cells were washed with sterile water, inoculated into SD, and incubated on a rotator between 17 and 20 hours before exposure to experimental conditions.

Determination of maximum population in limited lysine

A growing culture will continuously use up the limiting nutrient. Assuming each cell requires 10 fmol of lysine to make a new cell (literature value is 5.4 fmol/cell (Shou et al., 2007)), media containing 0.02 μM lysine can support up to 100 cells/ml without changing the concentration of lysine by more than 5%. To monitor several rounds of doubling, cultures were initially inoculated at a concentration 20 cells/ml, and subsequent concentrations were determined by concentration/dilution followed by

plating on YPD.

Growth rate determination

The cultures grown at each concentration showed a lag phase followed by either exponential growth or death. This can be described by the following piecewise equation:

$$N_0(t) = \begin{cases} N_0, & \text{if } t < t_{lag} \\ N_0 \exp r(t - t_{lag}), & \text{if } t \geq t_{lag} \end{cases}, \quad (4.5)$$

where N is the number of cells/ml, N_0 is the initial number of cells/ml, r is the growth rate, and t_{lag} is the lag time. Using non-linear regression, the best-fit values of r and t_{lag} were obtained.

Construction of an RM11 diploid prototroph

I obtained a *MAT*α *lys2 ura3* haploid from the Gottschling Lab, and a *MAT*a *leu2 ura3* haploid which harbored the S288C allele of the gene *AMN1* (*AMN1*-BY) from the Bedelov Lab. The RM11 allele of *AMN1* prevents efficient separation of daughter cells from their mothers, which leads to cell clustering (Yvert et al., 2003). I inserted a functional URA3 gene amplified from S288C genomic DNA into the *MAT*α *lys2 ura3* haploid to create a *MAT*α *lys2* strain. I crossed this strain with the *MAT*a *leu2 ura3* strain and sporulated it to generate *MAT*a and *MAT*α prototrophs that carried *AMN1*-BY. These strains were mated to obtain a diploid prototrophic version of RM11.

4.10. Figures

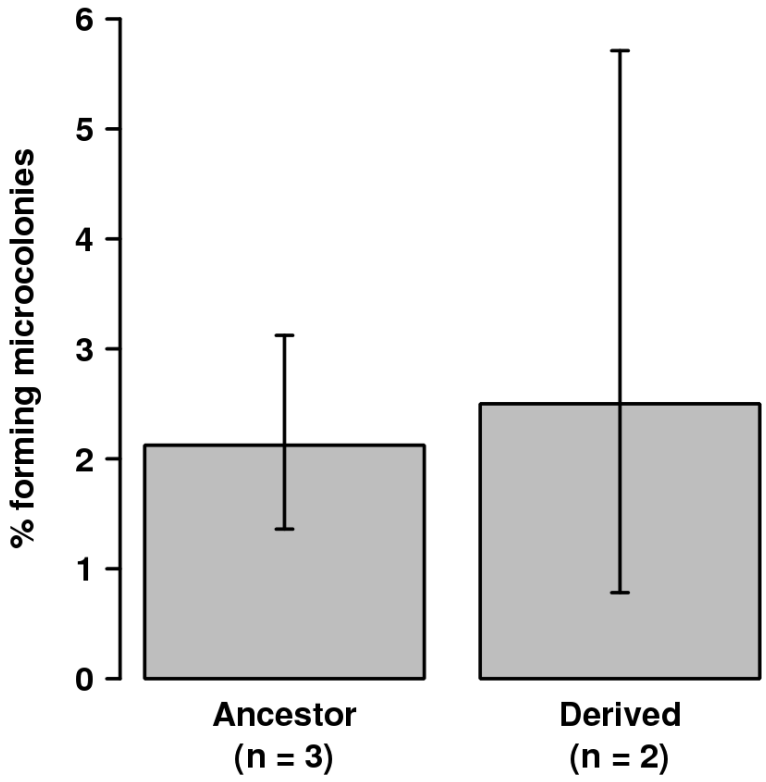

Figure 4.1: Growth on limiting lysine is not a heritable phenotype.
The plating assay was carried out as described (see Section 2.6) using a lysine auxotroph ("Ancestor"). After ~ 1 week of growth, pinpoint-sized colonies could be seen. At this point, one plate was replicated onto glycerol-containing media, followed by glucose-containing rich media (YPD). One large and one small colony that were present on all three plates were picked off of the YPD plate into SD + 164 μM lysine and the plating assay was repeated. Since they did not significantly differ, data from the large and small colony were combined and are presented as "Derived."

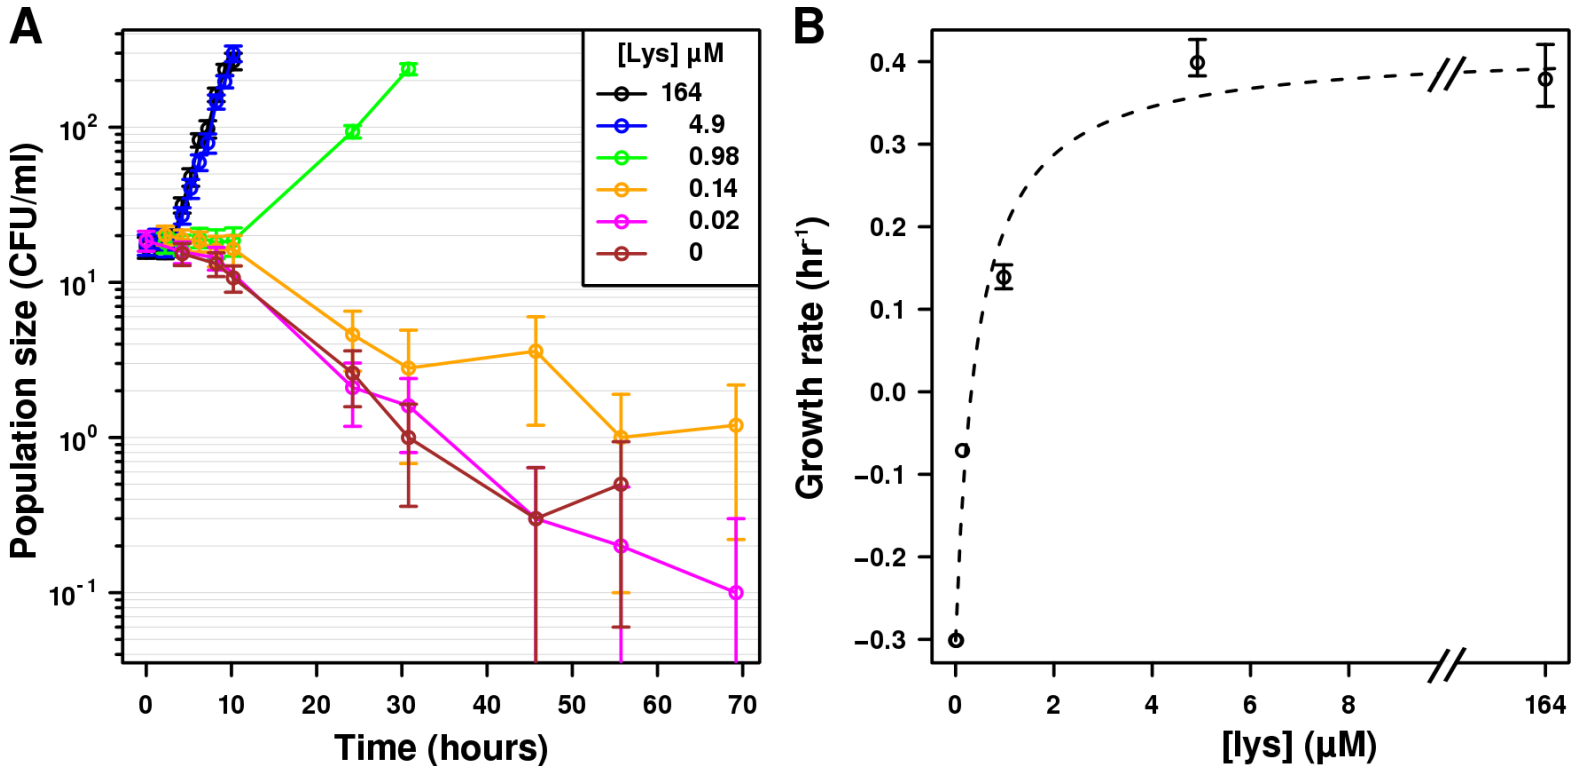

Figure 4.2: Growth rate is nutrient dependent.
A) Exponentially growing *lys2*Δ cells were transferred to minimal media (SD) to deplete vacuolar lysine storage. After 24 hours, cells were inoculated at low density into SD containing various amounts of lysine (0 hours on graph). Colonies were counted at various times after plating. Error bars represent the 95% confidence interval assuming a Poisson sampling process during plating. **B**) Growth rates from **A** were calculated using a model incorporating the lag time (see section 4.9 for details). Non-linear regression was used to fit the Monod Equation to these data (dashed line).

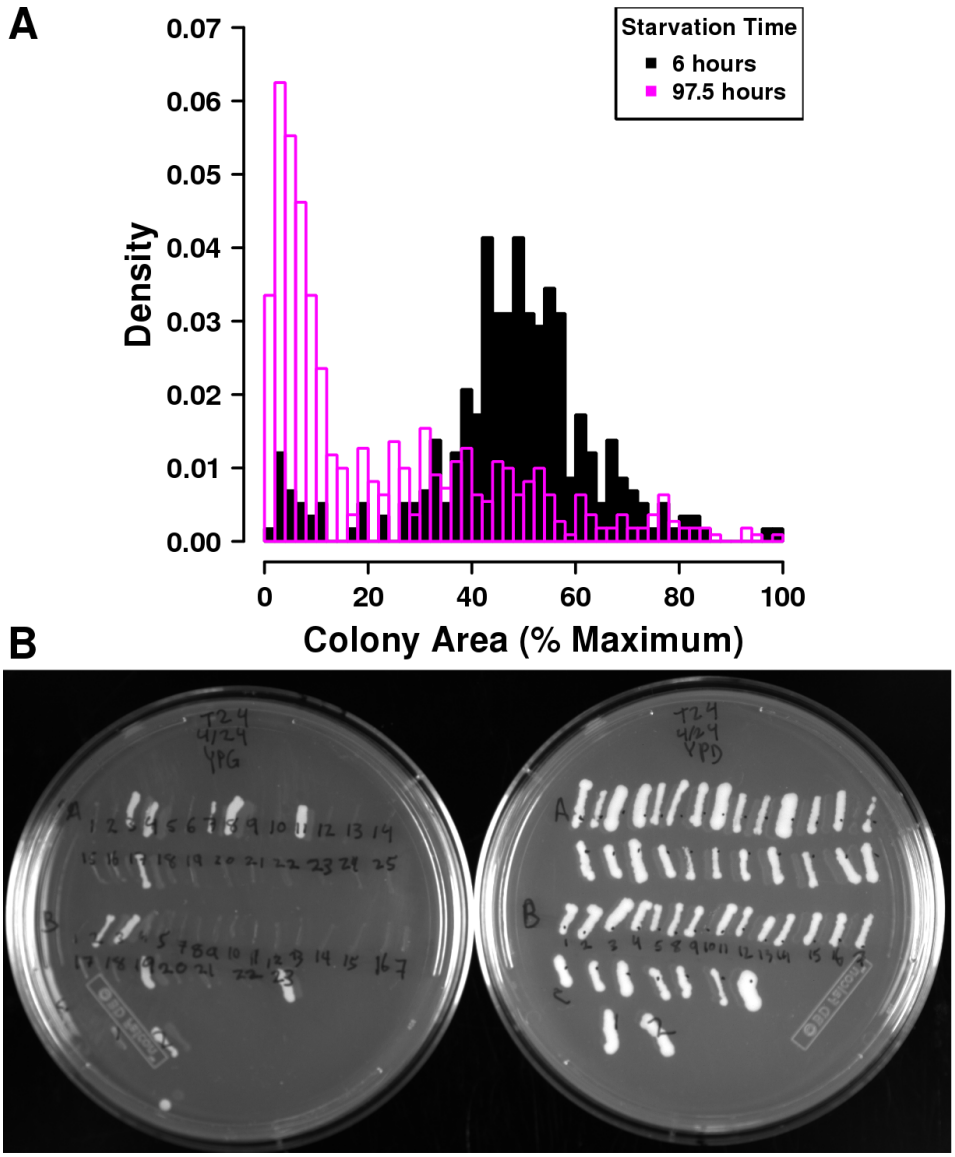

Figure 4.3: Size heterogeneity is not solely due to increased petite frequency. **A**) Lysine auxotrophic cells were starved in a lysine-free medium for various amounts of time (corresponding to Figure 4.2) before being plated on rich medium Plates were scanned and analyzed with image processing software (see Chapter 6.3) to compare colony sizes. Colony area is presented as percent of maximum to facilitate comparison between colonies that had been growing for different amounts of time. **B**) Mitochondrial deficiency does not account for all observed heterogeneity. Cells in media lacking lysine were plated after 24 hours. Small colonies were picked and struck onto YPG (left plate) followed by YPD (right plate) with the same toothpick. Failure to grow on YPG indicates mitochondrial deficiency.

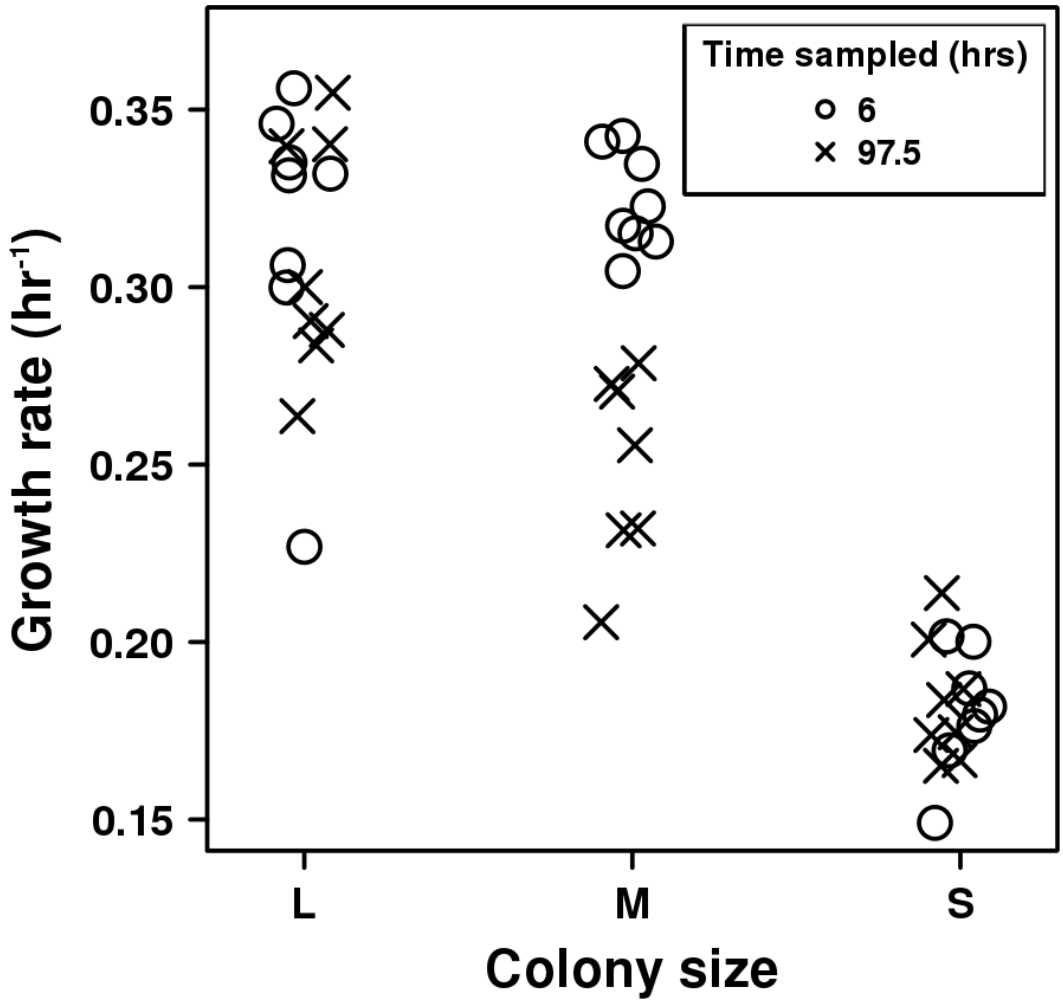

Figure 4.4: Growth rates of previously starved colonies.
Cultures starved for 6 (circles) or 97.5 (crosses) hours in SD were plated onto YPD. A large ("L"), medium ("M"), and small ("S") colony was picked from each plate and grown to saturation in YPD. These YPD cultures were diluted into the wells of a microtiter plate containing SD + 164 μM lysine in 8x replicates and incubated at 30 °C with shaking. ODs were read every 10 minutes for ~2 days. Growth rates were obtained for each curve in a semi-automated fashion by manually picking the range of ODs giving exponential growth and fitting it to an exponential function of the form $A\exp(rt)$.

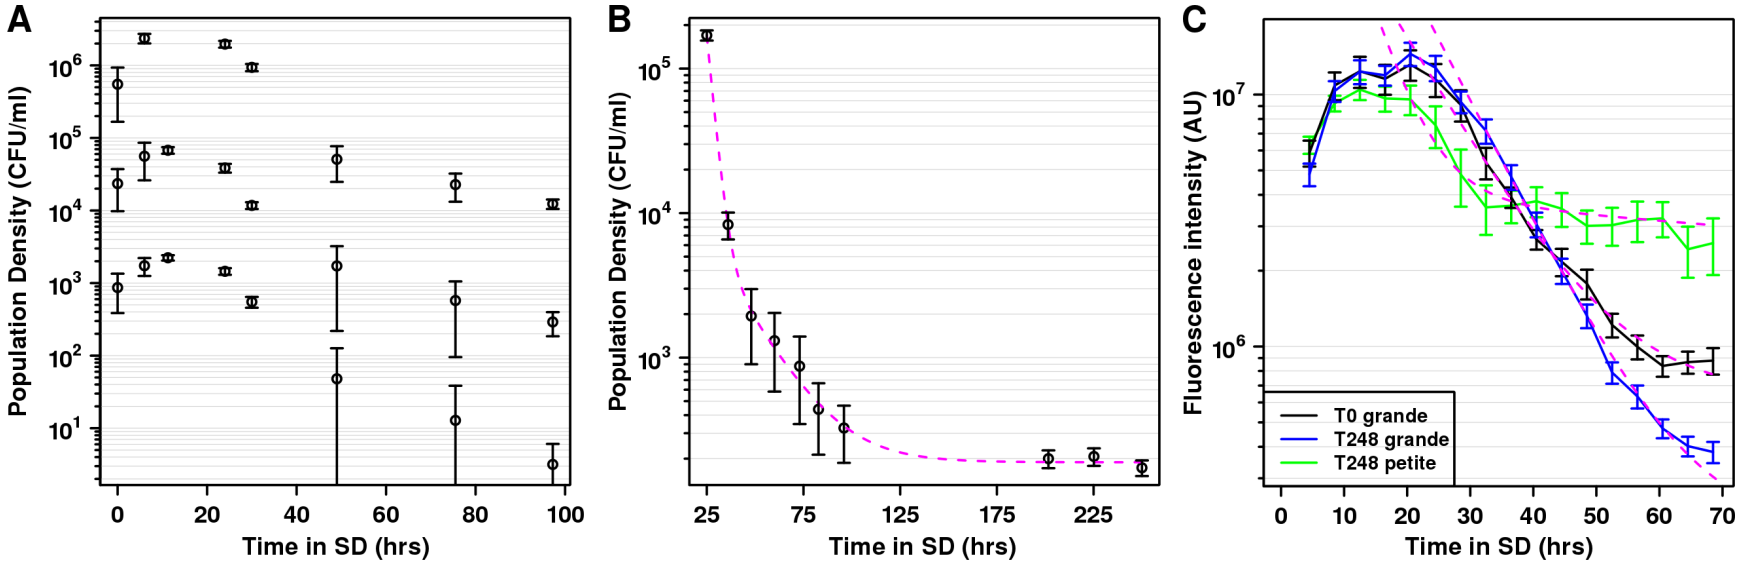

Figure 4.5: Death due to lysine starvation is density-independent and multiphasic.
Exponentially growing cells in SD supplemented with 164 μM lysine were washed three times in SD to remove residual lysine, and diluted into SD (Time 0). At the indicated times, cells were diluted into sterile water, plated onto YPD, and grown for up to 5 days at 25–30 °C. Error bars indicate the 95% confidence interval assuming a Poisson sampling process. **A**) Cells from identical exponentially growing culture were diluted to three different initial densities. **B**) In a separate experiment, a culture was tracked out to 250 hours (black points). The drop in CFU was fit to a biexponential plus a constant (magenta dashed line; see text for parameter values). **C**) Slower death is heritable for petites but not grandes. A grande colony (blue) and five petite colonies (green) from cells plated after 250 hours of lysine starvation were grown separately to saturation in YPD. These overnight cultures, plus the strain used to initiate the experiment (black) were used to inoculate wells of a microtiter plate containing minimal media. Fluorescence intensity for five positions of each well were monitored. Non-linear least squares regression was used to estimate death rates (magenta dashed line; see text for parameter values).

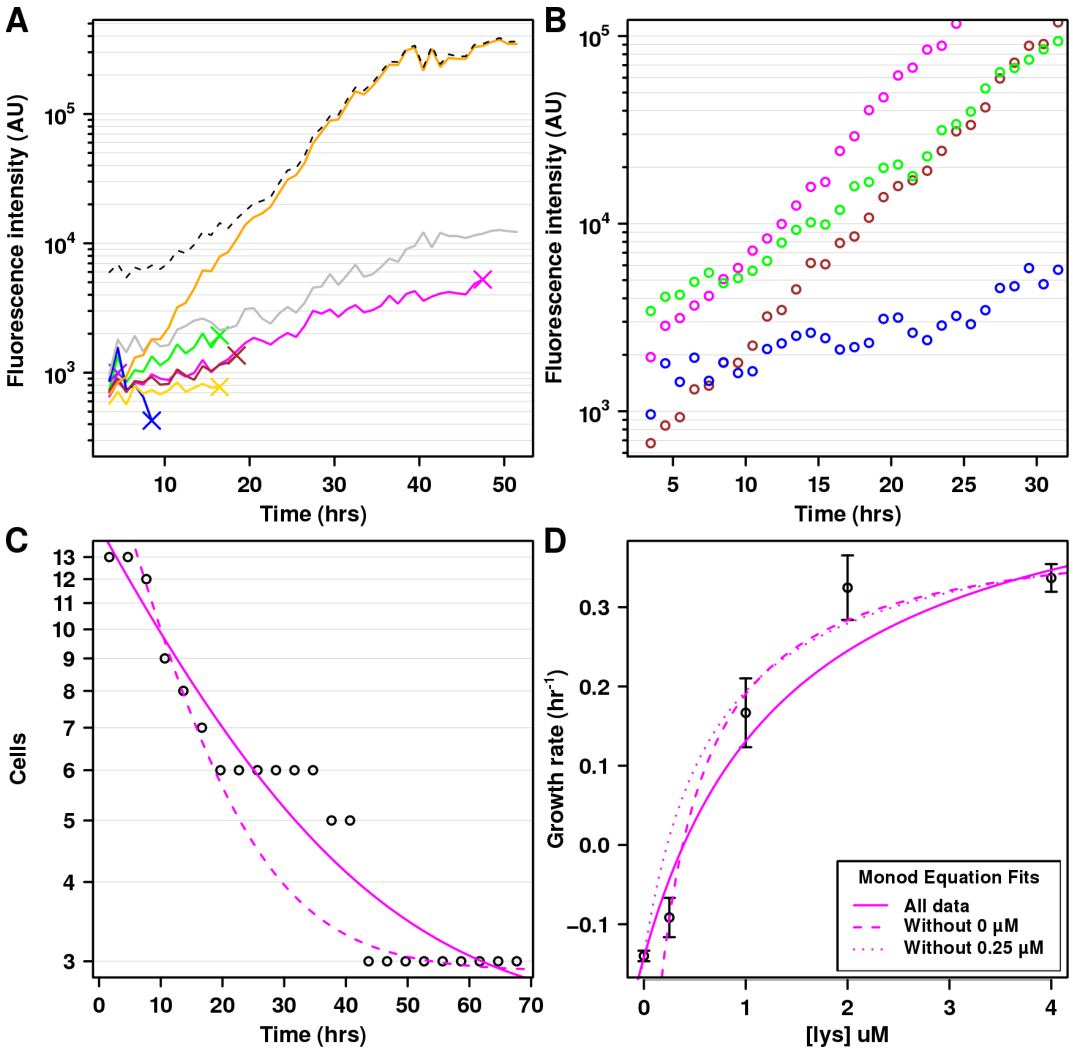

Figure 4.6: Differing fates of cells in identical environments.
A) The fates of all cells (n = 8) from a single position in a microtiter well containing 1 µM lysine were tracked individually by cropping the image to the region around each cell (solid lines). A dramatic decrease in fluorescence indicates cell death. A "lag time" is apparent if total fluorescence for the entire image is used (dashed black line). **B**) Growth rates vary between cells able to grow in 1 µM lysine. **C**) Cells die at different rates in 0.25 µM lysine. Because some cells did not die according to a simple exponential growth model (i.e., no decrease in cell number from ~20 to ~40 hours), data is fit is to a single exponential plus a constant with (solid magenta line) or without (dashed magenta line) these cells. Calculated rate is 0.09 (0.08–0.10) hr^{-1}. **D**) Growth and death rates based on single-cell data are plotted against lysine concentration. Estimates of growth and death rates in various concentrations of lysine are poorly fit by the Monod Equation, even if different data points are removed.

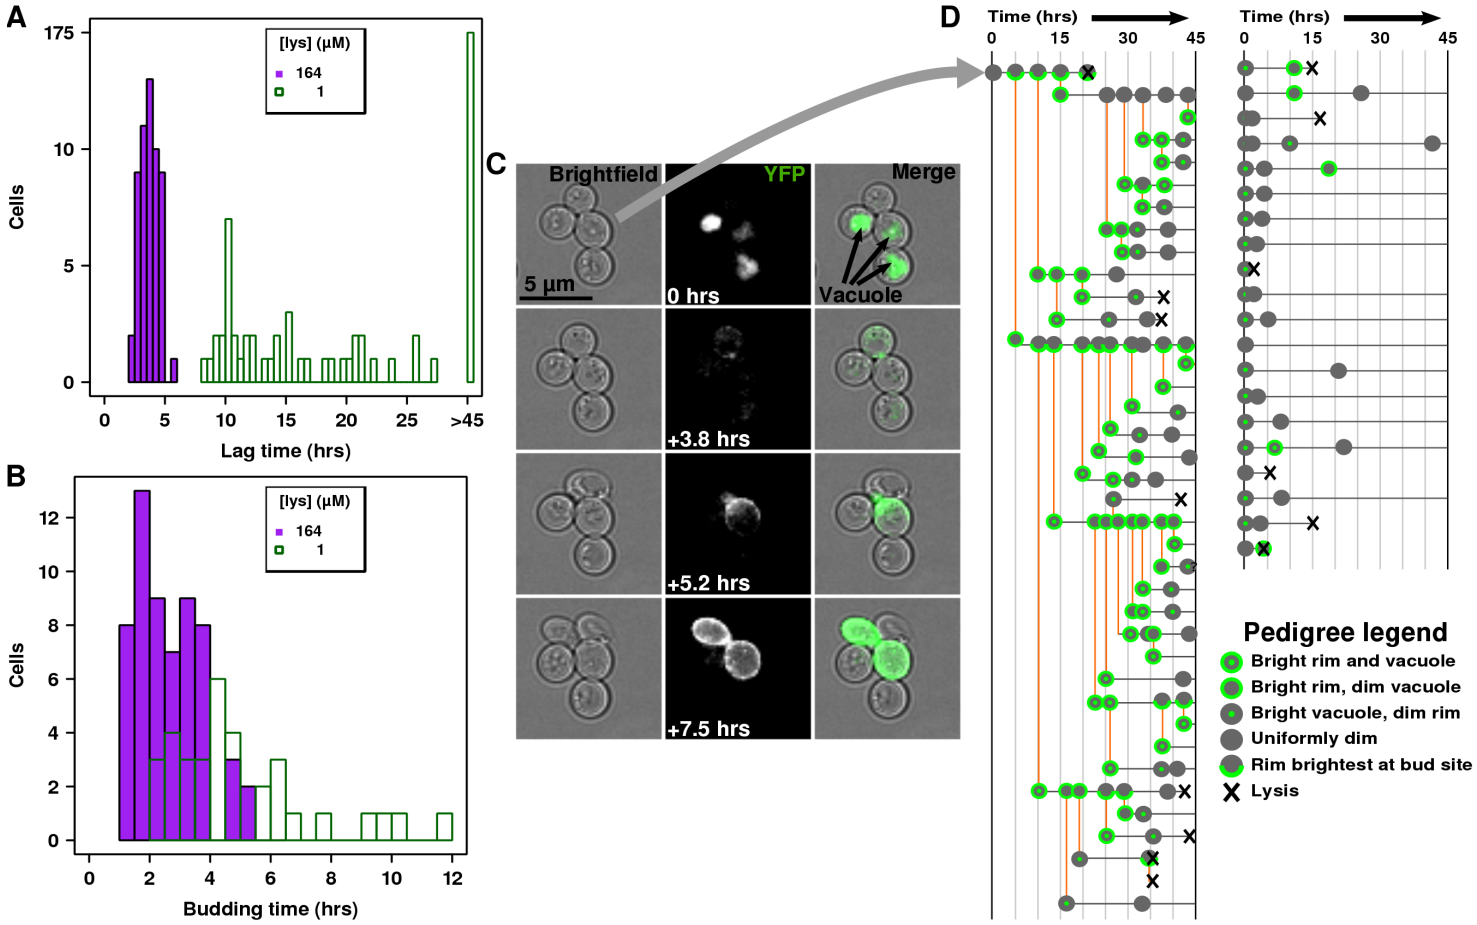
A
[lys] (μM)
164
1
Cells
Lag time (hrs)
B
[lys] (μM)
164
1
Cells
Budding time (hrs)
C
Brightfield
YFP
Merge
5 μm
0 hrs
Vacuole
+3.8 hrs
+5.2 hrs
+7.5 hrs
D
Time (hrs)
Time (hrs)
Pedigree legend
Bright rim and vacuole
Bright rim, dim vacuole
Bright vacuole, dim rim
Uniformly dim
Rim brightest at bud site
Lysis

Figure 4.7: The ability to grow in limiting lysine is dependent upon Lyp1p surface expression.
A) Pre-starved cells were loaded into the observation chamber of a microfluidics device and exposed to 164 μM (purple solid bars) or 1 μM (green empty bars) lysine. The lag time (i.e., the amount of time it took a cell to bud upon exposure to nutrient) was recorded. While 100% of cells in 164 μM were able to bud, only 19% were observed to bud in 1 μM lysine. **B**) An increase in the variability of budding time (i.e., for a given cell, the time between the appearance of bud *n* to the appearance of bud *n*+1) was also observed in 1 μM lysine. **C**) Timelapse images of cells exposed to 1 μM lysine. Initially bright vacuolar localization disappears. All cells able to bud increased expression and localization of Lyp1p to cell surface, with the greatest expression at the future budding site. Time in 1 μM lysine is indicated in middle column. **D**) Pedigree analysis revealed that surface expression of Lyp1p is necessary, but not sufficient, for growth in 1 μM lysine. Horizontal lines indicate the time progression of individual cells, with changes in phenotype indicated by different symbols along this line. Budding events are indicated by orange vertical lines and connect mothers to daughters. The cell that was able to form a microcolony in **C** is indicated by the gray arrow.

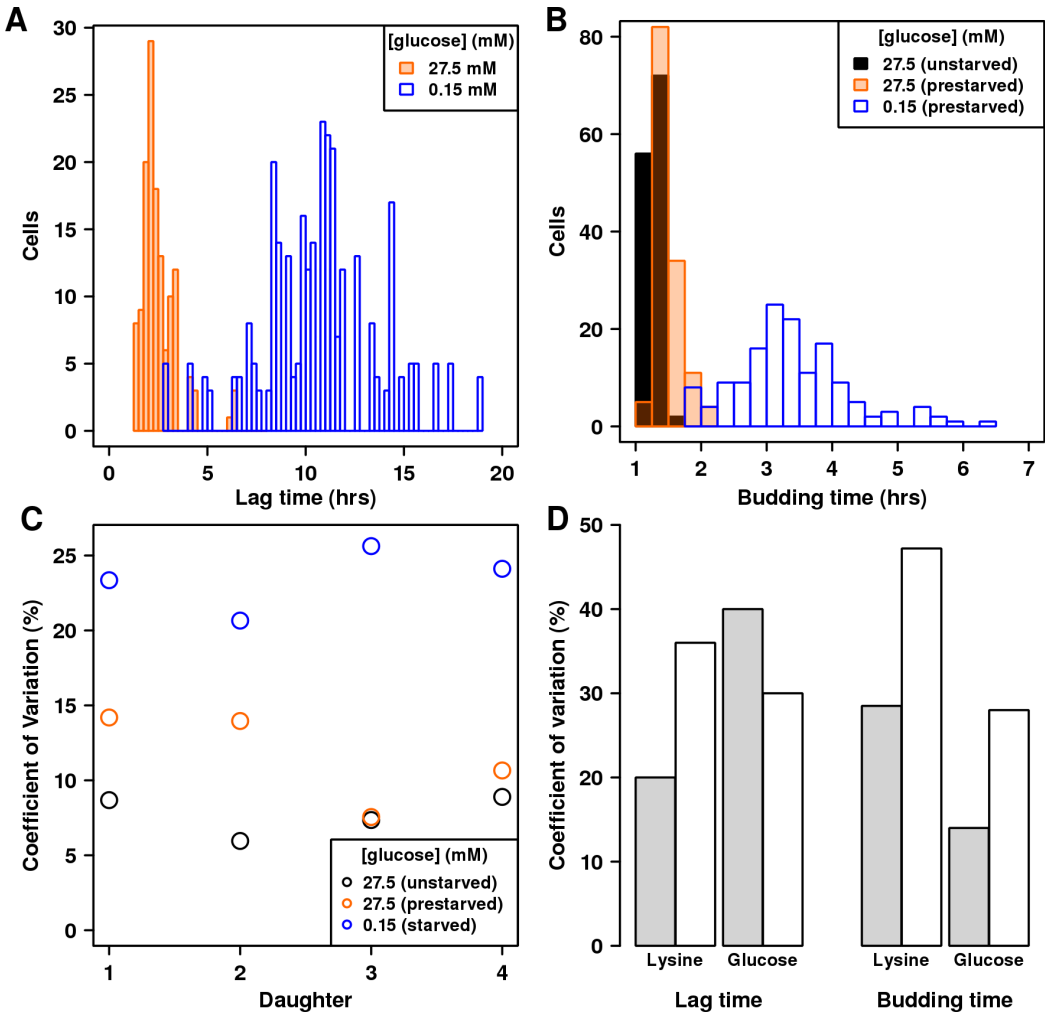

Figure 4.8: Heterogeneity observed in wild diploid prototrophic strains when limited for glucose.

A and **B**) Pre-starved cells were loaded into a microfluics device and exposed to concentrations of glucose that were either non-limiting (27.5 mM; translucent orange) or limiting (0.15 mM; blue border) for growth rate. Lag time (**A**) and budding time (**B**) were scored as in Figure 4.7. To control for the effects of starvation alone, unstarved cells (solid black lines in **B**) were scored for budding time. **C**) All observed cells were grouped by which daughter they were, and the coefficient of variation (standard deviation / mean, "CV") of budding time across the population was calculated. Prestarved cells exposed to non-limiting glucose (orange) regain the CV of unstarved cells (black) after budding four times, whereas cells exposed to limiting glucose (blue) maintain a high CV even after budding five times. **D**) Comparison of heterogeneity in non-limiting (filled) and limiting (open) nutrient. Lysine data is from Figure 4.7.

Chapter 5. Conclusions and future directions

5.1. Adaptation is a variance-generating mechanism that allows stochastic cooperator dominance

To uncover novel mechanisms that allow cooperation based on the production of common goods to survive cheating, I examined an engineered cooperative system that bypasses known mechanisms of cheater control. After adding cheaters, I expected slow but deterministic cheater takeover. Instead, identically initiated co-cultures showed rapid population divergence: whereas cheaters rapidly destroyed cooperation in some co-cultures, cooperators rapidly displaced cheaters in other co-cultures. I determined that this process was an "adaptive race" between cheaters and cooperators that was driven by strong selection for improved growth in the novel, nutrient-limited cooperative environment.

The race between cooperators and cheaters to adapt to a new environment can eventually favor cooperation, despite the initial fitness advantage of cheaters. As has been shown theoretically (Morgan et al., 2012a; Santos and Szathmáry, 2008), if the frequency of recombination is negligible, and if the fitness gain of the most adaptive mutation in cooperators exceeds that of cheaters by at least the fitness cost of cooperation, the cooperative trait can "hitchhike" on environmental adaptation and rise to high frequency (Maynard Smith and Haigh, 1974; Morgan et al., 2012a; Santos and Szathmáry, 2008). Otherwise, cheaters dominate. Thus, the probability of either type

eventually dominating a co-culture is related to its initial abundance in the population (Figure 2.3B), as a larger population is more likely to sample better mutations. However, this initial symmetry is quickly broken: if cheaters dominate, the cooperative system will collapse (Figure 2.2, gray) because the common good necessary for survival is no longer produced in sufficient quantities. In fact, the better the mutation obtained by cheaters, the sooner self-extinction will occur. On the other hand, if cooperators dominate, they continue to produce the common good and proliferate (Figure 2.2, orange). These cooperators have improved their ability to initiate cooperation (Figure 2.4B), and are able to defeat evolved cheaters that are superior to the ancestral cheater (Figure 2.4A). While proliferating, cooperators can potentially acquire further beneficial mutations (Schwilk and Kerr, 2002).

Two factors can synergize with the adaptive race to stabilize cooperation. First, different types of environmental change can provide different opportunities for adaptation. Improved cooperators that have survived one round of the adaptive race will eventually generate cheating mutants. Initially, the small population size of these cheating mutants makes them vulnerable to extinction by genetic drift and less likely to acquire the best mutation during subsequent rounds of the adaptive race. Eventually, however, cheaters will crash the cooperative system if the environment remains constant long enough. This is because the magnitude of possible fitness gains from subsequent adaptations will diminish over time (Lenski and Travisano, 1994). However, other environmental changes in, for example, temperature or osmolarity, will select for distinct

mutations that confer large fitness gains, triggering another adaptive race. Second, if individuals can migrate between spatially separated populations, “Simpson's paradox” allows cooperators to dominate globally even though cheaters grow faster in each population. This can occur even when the population structure is periodically destroyed by global mixing (Chuang et al., 2009). The only requirement is sufficiently variable cheater frequency between groups, which can be achieved through dilution (Chuang et al., 2009), or though mutations obtained during an adaptive race. Thus, in conjunction with spatially-structured environments that change in different ways—two unavoidable realities of nature—the adaptive race “buys time” for a cooperative system to acquire additional mutations which may result in more reliable mechanisms of cheater tolerance, such as partner recognition.

Our cooperative system was engineered to cooperate in a manner completely foreign to its ancestor, so it is pertinent to consider whether the adaptive race occurs in more natural systems. I consider two examples in natural microbial cooperative communities where the adaptive race may have operated. The first example involves the social bacterium *Myxococcus xanthus*, members of which, when starved, exchange developmental signals, aggregate, and differentiate to form a fruiting body. During this process, a majority of cells die while the remaining cells become tough spores (Velicer and Vos, 2009). This provides an opportunity for a strain to cheat by disproportionately forming spores when mixed with another strain (Fiegna and Velicer, 2003). One cheater strain, which could not form spores on its own, repeatedly rose to a high frequency when

mixed with the wild type strain and drove the entire system to very low population sizes (Fiegna and Velicer, 2003). Reminiscent of our system, some of the replicate co-cultures suffered complete collapse, while the remaining co-cultures emerged cheater-free. If not solely due to drift, I suspect that this stochastic cheater outcome may be the result of an adaptive race similar to what I report here. This could occur if the cheater created a novel environment unfamiliar to the wild-type cooperating ancestor. If the emerging cooperators demonstrate an improved ability to defeat ancestral cheaters, then sequencing evolved cooperators might reveal the type of novel environment imposed by cheaters, as well as the mechanisms deployed to adapt to it.

A second example of a potential adaptive race was recently shown in the siderophore-producing bacterium *Pseudomonas fluorescens* (Morgan et al., 2012b). During iron-limitation, *P. fluorescens* produces and releases iron-chelating siderophores which can be taken up by any nearby individual. When non-producing cheaters, which have a large fitness advantage over siderophore-producing cooperators, were mixed with cooperators in an iron-limited environment, cheaters introduced at low frequency failed to invade (Morgan et al., 2012c). The authors propose that this occurred because the numerically-dominant cooperators had a greater chance of obtaining a high-fitness mutation adaptive to iron limitation that would sweep through the population. However, the actual mechanisms were not determined. This implies that, although *P. fluorescens* must have experienced similar types of iron limitation throughout its evolutionary history—as evidenced by the very existence of siderophore production—the

experimentally-imposed iron limitation was still sufficiently “novel” to elicit an adaptive race. This supports our experimental finding that the environment does not need to be completely novel for the adaptive race to occur (Figure 2.6).

Similar genetic hitchhiking in a not-so-novel environment was recently observed outside the context of cooperation. In an experiment initially designed to track deterministic takeover of sterile yeast mutants (which have greater fitness than non-sterile cells), the researchers instead observed complex patterns consistent with the rise of multiple non-sterile mutants with greater fitness than the sterile mutants (Lang et al., 2011). Importantly, this occurred in rich media thought to impose minimal selective pressure. These examples illustrate that the conditions necessary to initiate the adaptive race may be met more frequently than is currently assumed.

In summary, by using an engineered cooperating-cheating yeast system to bypass all known cheater-control mechanisms, I discovered that the adaptive race allowed cooperators to stochastically defeat cheaters. This simple mechanism requires low rates of recombination and can be triggered by any environment where adaptation offers fitness gains greater than the cost of cooperation. Experimental evidence from naturally occurring cooperative systems suggests that these conditions may be met even when the environment is not completely novel (Fiegna and Velicer, 2003; Morgan et al., 2012c). Because adaptation to a changing environment is the norm in biology, I propose that the adaptive race between cooperators and cheaters, especially in conjunction with spatially structured populations, is a fundamental mechanism for cooperative systems to survive

the perpetual onslaught of cheaters.

5.2. The variance generated by the adaptive race can be maintained by migration

The experiments discussed in Chapter 2 considered single populations—represented by individual tubes—over a relatively small number of generations. Although the contents of each tube were well-mixed, the tubes themselves imposed an implicit spatial structure, because no migration was allowed between tubes. Thus, cooperators and cheaters raced to obtain the best *local* mutation within their tube, but these mutations could not compete on an inter-tube level. Occasional migration between populations is a biologically realistic assumption, so I was interested in how migration could interact with the adaptive race to potentially stabilize cooperation against cheating.

To investigate how migration would influence the outcome of the adaptive race, I developed a mathematical model that allowed me to simulate this process over a large parameter space. I found that, in this model, the presence of a small number of high-fitness mutants was critical for the ability of cooperation to succeed in metapopulations that were initially fully occupied by 1:1 mixtures of cheaters and cooperators. The presence of initially empty locations allowed cooperators to defeat cheaters under some combinations of migration rates and maximum death rates even in the absence of mutants. In general, more initially empty space was better for the domination of cooperation. In these cases, the presence of a small number of mutants still improved the parameter range that favored cooperation over cheating.

My model differs in important ways from otherwise similar ones found in the literature. Classically, continuous versions of the public goods dilemma assume that the fitness enhancing cooperative benefit increases with the number of cooperators, and cooperators pay a fixed cost of cooperation (Hauert et al., 2006; Ichinose and Arita, 2008; Killingback et al., 2006; Parvinen, 2011; Suzuki and Kimura, 2011). Death is either not explicitly modeled (Ichinose and Arita, 2008; Killingback et al., 2006), or is modeled as a distinct process that either occurs with a constant probability (Hauert et al., 2006; Suzuki and Kimura, 2011), or is an increasing function of how much cooperative benefit is provided by the individual (Parvinen, 2011). I took an alternative approach and modeled growth rates instead of costs and benefits. I assumed that all types (i.e., ancestor and evolved cooperators and cheaters) shared the same maximum death rate, which meant the advantage of cheaters over cooperators decreased with increasing cheater frequency. In general, this assumption should make my model more favorable to cooperation.

Explicitly modeling fitness as a growth (or death) rate allowed me to observe that the combination of two conditions that intuitively should be detrimental to the survival of cooperation—namely, high migration rates and high cheater-induced death rates—resulted in deterministic fixation of cooperation (Figure 3.4).

Another unique aspect of my model was the dilution scheme, under which *all* populations were diluted by a fixed amount any time *one* population reached some maximum population size. Similar models enforce a maximum population size by

explicitly modeling density dependent effects, such that birth rates slow down as a function of population size (Parvinen, 2011) or available space (Hauert et al., 2006). However, this punishes populations that do well by making them slow down or stop. My dilution scheme kept populations in the evolutionarily relevant exponential growth phase without punishing the most successful populations. Future experiments will explore other dilution schemes. A simple modification would be to only dilute the population that reaches maximum size. Another option would be periodic dilution of all populations in a metapopulation. This will simulate the more ecologically realistic case of periodic nutrient availability. I expect this regime to be less favorable to cooperation, since the most successful populations will not grow or migrate nearly as much as they do in the current model. However, I do not expect this to negate the observation that a small number of mutants present in the initial population is beneficial for cooperation.

The high-fitness mutants present at the beginning of the simulations represented preexisting mutants that are adaptive to the nutrient-limited cooperative environment. Cheaters reduce growth rates even further by taking away a portion of this already limited resource. Because cheaters act to increase the selective pressure of nutrient limitation, their very presence could allow the adaptive race to occur in populations that have already adapted to the level of nutrient provided by a cheater-free cooperative environment. Thus, by degrading the environment, cheaters could select for the low-frequency, high-fitness mutants ultimately responsible for the overall success of cooperation. I plan to test this idea in a future version of the model.

5.3. Can genetic "niche-hiking" explain the success of cooperation in an experimental metapopulation?

I wanted to test the predictions of my simulations. Since our lab was not set up to do full metapopulation experiments, we started a collaboration with the Kerr lab at the University of Washington. These experiments used the bacterium *P. aeruginosa*, and the environmental stress was antibiotics instead of nutrient limitation.

The metapopulation experiments were designed to be unfavorable to cooperation. Growth in the structured habitat of microtiter wells was disrupted at every transfer by completely mixing the contents of the plate before dilution and transfer to a new microtiter plate, which corresponded to global migration. Nevertheless, a high frequency of cooperation was achieved, and maintained, over the course of four transfers (Figure 3.7).

The increase in cooperator frequency in the first transfer could be explained by the fact that, on average, less than one cell per well was able to survive. Thus, the bottleneck of antibiotic treatment isolated individual surviving cooperator and cheater cells into separate wells. Cooperators could grow unimpeded, while cheaters could not grow and perhaps died due to starvation. After this growth period, however, the entire plate was mixed and replated. Now, because of their low frequency, surviving cheater cells were likely to end up in a well with a cooperator. At this point, the situation once again represents a tragedy of the commons scenario, and cheaters should deterministically increase in frequency. Instead, after an initial drop, the cooperator

frequency appears to remain stable.

What could cause the apparent stability of cooperation in this nearly well-mixed metapopulation? One possibility is that this is an example of "niche-hiking" (Schwilk and Kerr, 2002). Normal genetic hitchhiking occurs when a deleterious trait (such as cooperation) is brought to high frequency by being linked to a strongly advantageous one (Maynard Smith and Haigh, 1974). Niche-hiking occurs when a cooperative trait generates a niche favorable for growth, which allows cooperators to sample a relatively greater number of beneficial mutations than non-cooperators. Importantly, unlike hitchhiking, niche-hiking does not require genetic linkage between the cooperative trait and the beneficial mutations.

According to the niche-hiking hypothesis, because the isolated cooperators were able to grow in the antibiotic environment they sampled more mutations than the isolated cheaters. They were thus more likely to obtain mutations that compensated for the fitness reduction that typically accompanies antibiotic resistance. After the first transfer, then, the tragedy was averted because cooperators were actually more fit than cheaters. We plan to test this hypothesis by sequencing the genomes of emerging cooperators and cheaters to look for mutations that could compensate for streptomycin resistance. If the only mutations found are not known to compensate for streptomycin resistance, they can be tested by placing them, one at a time, into a streptomycin resistant strain and testing for improved growth rate.

5.4. The molecular origins of non-genetic heterogeneity

I realized that the assay I was using to test the ability of strains to grow on limiting lysine was revealing another source of variation in each population. This variation was not heritable (Figure 4.1), only expressed itself during nutrient limitation (Figure 4.7A and B), and, similar to the observations of Novick and Szilard (1950), resulted in a reduction in population growth rate (Figure 4.2). Using a combination of fluorescence microscopy and microfluidics, I was able to track the source of this variation to differential expression of the high-affinity lysine permease, Lyp1p (Figure 4.7).

Of course, this answer only pushes the question one level back. Why are some members of the population able to express Lyp1p under these conditions, while others are not?

One likely scenario is that expression of Lyp1p is similar to the control of the *lac* operon in *E. coli* (Görke and Stülke, 2008). In this system, expression of the operon, which encodes β-galactosidase (LacZ) and the lactose permease (LacY), is imperfectly repressed (Choi et al., 2008) by the inhibitor LacI, which is constitutively active. The fact that repression is imperfect means that each cell will have a low number of LacY molecules, and low numbers means high variance due to sampling error. As first observed by Novick (Novick and Weiner, 1957), cells with less than the required number of permeases (now known to be ~400 (Choi et al., 2008)) cannot respond to low amounts of inducer, while those above the threshold rapidly achieve full expression of the operon.

Upon division, the daughters of induced cells inherit this “molecular success”, which allows growth in low inducer environments to be heritable and not genetically mediated.

In *S. cerevisiae*, several high-affinity amino acid permeases have been shown to be constituatively expressed, but are sent to the vacuole for degradation when their substrate concentration is too high (Beck et al., 1999; Umebayashi and Nakano, 2003), or when the cell is stressed (Lin et al., 2008). The turnover process is controlled by the arrestin-like family of proteins (*ECM21*, in the case of Lyp1p), which allow ubiquitination of plasma membrane proteins by the E3 ubiquitin ligase *RSP5* (Lin et al., 2008). Cells degrade permeases to recycle amino acids for cellular protein synthesis. Because high-affinity permease is always produced, this type of regulation allows for a “*lac*-like” situation to occur. Any imperfection in the process of ubiquitination and turnover could allow a small fraction of permease to be mis-sorted to the cell surface. This could allow some cells to accumulate enough permease at the plasma membrane to trigger a positive feedback loop, facilitating uptake of external amino acids present at concentrations that are normally too low for normal cells to stabilize Lyp1p on the cell surface.

One way this cell-to-cell variability could occur is if Ecm21p and Rsp5p are present at low abundance in the cell. Purely random fluctuations in the number of these proteins means that some cells would not have enough of them to prevent some Lyp1p from reaching the cell surface. Ecm21p is estimated to be present at ~ 800 molecules/cell under normal growth conditions (Ghaemmaghami et al., 2003). While no

estimate exists for the abundance of Rsp5p, only dim expression is seen with our own C-terminal tags of this protein (data not shown). Unlike the *lac* operon case, however, the expected amount of "inducer" (Rsp5p or Ecm21p) in the daughter cell is the *normal* amount, which would effectively prevent more surface expression of Lyp1p. Dilution due to division and normal turnover would rapidly reduce the amount of surface Lyp1p until it falls back below the threshold. Thus, growth in limiting lysine would only appear to be weakly heritable, which is what I observed using microfluidics (Figure 4.7D).

A testable prediction generated by this hypothesis is that other permeases regulated by Ecm21p/Rsp5p should be present at the plasma membrane of cells that are able to grow in 1 μM lysine, even though these permeases are unrelated to lysine uptake. Research in our lab performed by Chi-Chun Chen has shown that the purine/cytosine permease Fcy2p is also regulated by Ecm21p and Rsp5p. Thus, cells that can grow in low lysine should also show surface expression of Fcy2p. My preliminary experiments indicate that Lyp1p and Fcy2p are both present in 90% of growing cells, while in non-growing cells both proteins are in the vacuole (data not shown). However, in the remaining 10% of cells, some show surface expression of Lyp1p and not Fcy2p, while none express Fcy2p at the cell surface in the absence of Lyp1p surface expression. These observations allow the possibility that surface expression of Lyp1p – and *not* low numbers of Ecm21p or Rsp5p – is the first step in this process. Followup experiments to test this idea, which include time-lapse microscopy to capture expression dynamics,

are currently in progress.

The hypothesis that daughter cells inherit their mothers' Lyp1p is also testable. The RITE system allows rapid and irreversible conversion of a fluorescence tag from green to red upon addition of estradiol (Verzijlbergen et al., 2010). I could add estradiol to the microfluidics observation chamber upon addition of limiting lysine. This would convert the Lyp1p in the initial cells from green to red. If the daughters of cells that grow contain red fluorescent Lyp1p, then inheritance has occurred. If only green fluorescent Lyp1p is present, then the daughters have made Lyp1p *de novo*. It is also possible that some Lyp1p is inherited, while the rest is made *de novo*. This could be quantified by calculating the ratio of red to green fluorescence in growing daughter cells.

My experiments also revealed heterogeneity in cell death. If all cells died with a constant probability, which is expected in a completely homogenous population, a single exponential would fit the death curve data. However, my data shows that the best fit contains at least two exponentials, which suggests at least two sub-populations are present. In addition, this heterogeneity is also non-genetic, as grande cells that survived ~250 hours of starvation were not mutants, and died just as quickly (if not more so) than the “ancestor” population (Figure 4.5C). On the other hand, petites isolated from this time point showed heritable death resistance, probably due to mutations arising from their inherently unstable genomes (Veatch et al., 2009).

The source of heterogeneity in grande cells remains unknown, but is a potentially interesting area for future study. One obvious source of heterogeneity is cell cycle

position, as it is possible that cells in different parts of the cell cycle display differential sensitivity to auxotrophic starvation. This could be easily tested using a synchronized population. Another set of experiments would involve finding differences in whole-genome expression profiles of the "fast dying" and "slow dying" populations using RNA-seq (Wang et al., 2009), and then develop a model to describe how these differences occur. This data would also produce candidate genes whose deletion would produce homogenous fast- or slow-dying populations. Previous research has identified *PPM1*, *TOR1*, and *SD9* as genes contributing to fast death during auxotrophic starvation (Boer et al., 2008). Interestingly, the death profiles for these auxotrophs (leucine and uracil) were consistent with a homogenous population, suggesting that multi-phasic death is a property of lysine auxotrophy. Thus, the way cells die could reveal important aspects of cell metabolism.

Nutrient limitation and starvation for auxotrophic nutrients have revealed interesting insights into cellular protein expression and metabolism. However, because these cells are auxotrophic, the observed variation is not likely to be evolutionary relevant. To gain a better understanding of the evolutionary relevance of heterogeneity during nutrient limitation, I switched to a wild isolate that was diploid and prototrophic, and observed its growth under limiting concentrations of glucose. The importance of glucose as a nutrient can be summarized by the fact that no less than 20 transporters must be removed before cells can no longer utilize it (Wieczorke et al., 1999). While variation in lag time (as a fraction of the mean lag time) decreased, I observed a large

increase in budding time variation. Since budding time is equivalent to fitness, the observed heterogeneity could be highly relevant to the evolution of *S. cerevisiae*.

In this case, I was unable to establish whether this heterogeneity was driven by heterogeneity in high-affinity permeases, as it is in the case of lysine limitation. *S. cerevisiae* has two major high-affinity permeases, *HXT6* and *HXT7*, which are almost identical in sequence (Boles and Hollenberg, 1997). I have made strains containing fluorescently-labelled *HXT6* and *HXT7*, but preliminary microfluidics experiments have not revealed an obvious connection between cells that can grow and the expression of these proteins. More thorough followup experiments and quantitative analysis may discern more subtle differences that are nevertheless important for growth in these conditions.

In summary, nutrient limitation increases heterogeneity in an otherwise clonal population of cells. This was observed in both auxotrophic strains limited for lysine, as well as prototrophic strains starved for glucose. The heterogeneous response to lysine limitation was driven by differential surface expression of the high-affinity lysine transporter, Lyp1p. Increased heterogeneity in the glucose-limited populations suggest that this variation is potentially relevant to the evolution of *S. cerevisiae* as a species.

5.5. Final thoughts: Variation and the success of cooperation

Lower relatedness within populations can be beneficial to cooperation

Population homogeneity is essential for the survival of cooperation, since it

increases the odds that cooperators will interact with other cooperators (Fletcher and Doebeli, 2009). In the absence of this "positive assortment", the fitness advantage of cheaters predicts the deterministic destruction of cooperation. How is this "positive assortment" achieved? Mechanisms include the ability to alter strategies based upon recognition and/or past experience (Axelrod and Hamilton, 1981; Trivers, 1971), living in a spatially-structured environment (Nowak et al., 2010; Queller, 1992), and privatization of the common good (Gore et al., 2009).

In Chapter 2 I looked for other mechanisms by removing these known one from an engineered yeast system. I discovered that the environmental stress imposed by the nutrient-limited cooperative environment selected for mutations conferring fitness advantages much larger than the cost of cooperation, allowing cooperators to purge cheaters stochastically. In other words, initial variation within the population at sites unrelated to cooperation allowed natural selection to remove variation at sites related to cooperation. Overall, population homogeneity was increased as a result of this process, which ultimately allowed the success of cooperation. However, without some initial variation within the population, this process could not have taken place. Thus, far from being universally detrimental to cooperation, variation within a population at loci other than those responsible for cooperation provides another mechanism by which cooperation can survive cheating. In other words, less than perfect relatedness among members of a population can be more beneficial to cooperation than clonality.

To illustrate this point, consider two populations that are locally well-mixed but

do not interact with each other. Each population is composed of cooperators, all of which which pay the same cost *c* to produce a common good, which provides benefit *b*. In one population (*X*), all individuals are identical to one another. The other population (*Y*) is identical to *X* but also contains mutants which are mildly deleterious which are maintained at low frequency through mutation-selection equilibrium. If relatedness (*r*) is measured by identity at the cooperative locus, $r_Y = r_X$. If relatedness is measured genome-wide, $r_Y < r_X$. Thus, no matter how *r* is measured, $r_Y \leq r_X$. Now assume one cooperator from each population mutates into a cheater which does not pay *c*. In *Y* it occurs in one of the numerically dominant wild-type individuals. These two cheaters will increase in frequency identically in their respective populations, and will eventually cause their populations to collapse. However, if the environment changes, natural selection can act on *Y*, since previously deleterious mutations can be highly advantageous in an altered environment (just like the *ecm21* and *rsp5* mutants in my study). Cooperators now have a chance proportional to their abundance to rise to high frequency and purge cheaters. On the other hand, since all members of *X* are identical, they will all respond in the same way to the new environment, and the cheaters will deterministically take over the population. Even though $r_Y \leq r_X$, *Y* is better off than *X*. In general, this will be true whenever natural selection is acting on loci other than those responsible for cooperation.

The proceeding example is especially relevant to studies that do not directly measure *r* at the cooperative locus, because, in the terminology of the example,

cooperation has a better chance of surviving in *Y* than *X* even though genome-wide $r_Y < r_X$. For example, in natural populations, the full set of genes responsible for the cooperative trait of interest are frequently unknown. Because of this, genome-wide *r* is used to determine relatedness (Boomsma, 2007; Cornwallis et al., 2010; Gilbert et al., 2007; Lukas and Clutton-Brock, 2012). Surprisingly, some studies using well-characterized, genetically tractable microbial systems make claims about relatedness without measuring *r* in any sense. Instead, relatedness is inferred from population structure (Brockhurst, 2007; Griffin et al., 2004; Refardt et al., 2013).

In Chapter 3 I extended the ideas of the adaptive race to a metapopulation to address the biologically realistic scenario of multiple populations that interact through migration. Migration reduces positive assortment by bringing cheaters to groups of cooperators, and cooperators to groups of cheaters. On the other hand, migration can allow successful cooperators to expand into new territory. Could the adaptive race synergize with migration to allow global domination of cooperation? In a metapopulation, because populations composed of mostly cooperators grow better than cheater-dominated ones, the global frequency of cooperation can increase even though cheaters increase in frequency within any mixed population (a phenomenon known as "Simpson's Paradox" (Chuang et al., 2009; Hauert et al., 2002)). When inter-population variance is maximized and intra-population variance is minimized, each population is composed of only cooperators (which grow well) or only cheaters (which die quickly). Thus, mechanisms that increase inter-population variance can support the overall success

of cooperation. In addition to migration, metapopulations introduce two other important abiotic sources of variation—extinction and unoccupied locations—which, depending on how frequently they occur, can either increase or decrease inter-population variance. I found that the adaptive race can synergize with these abiotic sources of variation to ensure the long-term survival of cooperation.

The adaptive race hinges upon the assumption that a few high-fitness mutations are available in a population. What would happen if populations were too small to have any of these mutations? This question is addressed in Chapter 4, where I observed variation in response to the nutrient-limited cooperative environment at the single-cell level. Somewhat surprisingly, when I observed small populations exposed to limiting nutrient, only a fraction were able to grow. Although the fraction that could grow was reproducible, the ability to grow was not heritable. I was able to trace the source of this variation to an "all-or-none" expression pattern of the high-affinity lysine permease, Lyp1p. Is the inherent variation resulting from bimodal expression of a high-affinity permease enough to act as an adaptive race at the single-cell level? Conceptually it seems possible, but the ultimate answer awaits future experimentation.

Appendix A. The Monod "constant" is not a constant

The relationship between the growth rate of a cell and the concentration of limiting nutrient has traditionally been modeled by the "Monod Equation":

$$r = \frac{R_{max} \times S}{K_r + S}, \tag{A.1}$$

where r is the growth rate, R_{max} is the maximum growth rate achievable by the organism, S is the concentration of the limiting nutrient, and K_r is the concentration of limiting nutrient required to achieve half-maximal growth. This equation essentially black-boxes all steps that occur between nutrient uptake and its conversion to the production of a new cell. The parameter R_{max}, being an absolute limit to growth, is easily measured under ideal growth conditions. S is easily controlled by the experimenter. However, K_r has been a notoriously difficult number to pin down for any given organism, with differences of an order of magnitude or more (Ferenci, 1999).

At the cell surface, transporters facilitate the uptake of nutrients. They function to lower the "activation energy" of moving a molecule across a cell membrane in a repeatable manner, and thus behave as catalysts. The relationship between an enzyme and its substrate is described by the Michaelis-Menten Equation

$$v = n\frac{V_{max} \times S}{K_m + S}, \tag{A.2}$$

where v is the instantaneous rate of enzymatic turnover of substrate to product, V_{max} is

the maximum rate of the enzyme, S is the concentration of substrate (which is the limiting nutrient), and K_m is the concentration of substrate at which $v = 1/2 V_{max}$. The number of transporters at the cell surface is captured by n.

If each cell requires x molecules of nutrient to produce a new cell, then Eq. A.1 and Eq. A.2 can be combined to get the division time

$$\frac{\ln 2}{r} = \frac{x}{v} = x \frac{K_m + S}{nV_{max} \times S} . \tag{A.3}$$

Suppose there are two populations of cells with identical R_{max}. Population 1 has n transporters per cell and population 2 has kn transporters per cell. The time it takes population 1 to reach 50% R_{max} at $S = K_{s1}$ is

$$2 \ln 2 R_{max} = x \frac{K_m + K_{s1}}{nV_{max} \times K_{s1}} . \tag{A.4}$$

Similarly, the time it takes population 2 to reach 50% R_{max} at $S = K_{s2}$ is

$$2 \ln 2 R_{max} = x \frac{K_m + K_{s2}}{knV_{max} \times K_{s2}} . \tag{A.5}$$

Equating Eq. A.4 and Eq. A.5 gives

$$\frac{K_m + K_{s1}}{nV_{max} \times K_{s1}} = \frac{K_m + K_{s2}}{knV_{max} \times K_{s2}} . \tag{A.6}$$

Multiplying both sides by nV_{max} and rearranging gives

$$K_{s2} = \frac{K_{s1} K_m}{k K_m + k K_{s1} - K_{s1}} \,. \qquad \text{(A.7)}$$

Thus, the Monod "constant" for growth, K_{s2}, is a function of the number of transporters, k.

Appendix B. Location of unpublished data used to generate figures

Figure	Location
3.1	Experiments/CoSMO/cheater_tolerance/cocultures/LC9-influence_of_cheaters
3.2	Experiments/CoSMO/cheater_tolerance/monoculture/death/ 101230-death_characterization_III
3.3	Experiments/Modeling/metapop
3.4, 3.5, 3.6	Code/Java/metapop/experiments/fig/fitness_before_load/static_env/fast_doubling/global_dilution/high_cheat_adv/reps/10x10
4.1	Experiments/transporter_localization/experiments/plating/data/microcolony_counts
4.2A, B	Experiments/BatchCulture/LimitingLysine
4.3A, 4.5A	Experiments/BatchCulture/Death
4.4	Experiments/BatchCulture/Death/powerwave/090318_sml_growth_rates
4.5B	Experiments/BatchCulture/Death/glucose_titration/090420
4.5C	Experiments/Microtiter/090706-96well_starved
4.6A, B	Experiments/Microtiter/090826-mf_effluent
4.6C, D	Experiments/Microtiter/090807-lys_titr
4.7	Experiments/Microfluidics
4.8	Experiments/RM11_nutrient_limitation/glucose/microfluidics

Table B.1: Folders containing unpublished data used to generate figures for this dissertation.
All folders are under /home/ajwaite/Documents.

References

Aghaeepour, N. (2010). flowMeans: Non-parametric Flow Cytometry Data Gating.

Archetti, M., and Scheuring, I. (2011). Coexistence of Cooperation and Defection in Public Goods Games. Evolution *65*, 1140–1148.

Armitt, S., and Woods, R.A. (1970). Purine-excreting mutants of Saccharomyces cerevisiae. I. Isolation and genetic analysis. Genet. Res *15*, 7–17.

Axelrod, R., and Hamilton, W.D. (1981). The evolution of cooperation. Science *211*, 1390–1396.

Badyaev, A.V. (2005). Stress-induced variation in evolution: from behavioural plasticity to genetic assimilation. Proc. R. Soc. B *272*, 877–886.

Balaban, N.Q., Merrin, J., Chait, R., Kowalik, L., and Leibler, S. (2004). Bacterial persistence as a phenotypic switch. Science *305*, 1622–1625.

Barclay, B.J., and Little, J.G. (1978). Genetic damage during thymidylate starvation in Saccharomyces cerevisiae. Mol. Gen. Genet *160*, 33–40.

Barnett, J.A. (2008). A history of research on yeasts 13. Active transport and the uptake of various metabolites. Yeast *25*, 689–731.

Bates, D., Maechler, M., and Bolker, B. (2011). lme4: Linear mixed-effects models using S4 classes.

Beaumont, H.J.E., Gallie, J., Kost, C., Ferguson, G.C., and Rainey, P.B. (2009). Experimental evolution of bet hedging. Nature *462*, 90–93.

Beck, T., Schmidt, A., and Hall, M.N. (1999). Starvation Induces Vacuolar Targeting and Degradation of the Tryptophan Permease in Yeast. J Cell Biol *146*, 1227–1238.

Bennett, M.R., and Hasty, J. (2009). Microfluidic devices for measuring gene network dynamics in single cells. Nat. Rev. Genet *10*, 628–638.

Bennett, M.R., Pang, W.L., Ostroff, N.A., Baumgartner, B.L., Nayak, S., Tsimring, L.S., and Hasty, J. (2008). Metabolic gene regulation in a dynamically changing environment. Nature *454*, 1119–1122.

Benzer, S. (1953). Induced synthesis of enzymes in bacteria analyzed at the cellular

level. Biochimica et Biophysica Acta *11*, 383–395.

Bishop, A.L., Rab, F.A., Sumner, E.R., and Avery, S.V. (2007). Phenotypic heterogeneity can enhance rare-cell survival in “stress-sensitive” yeast populations. Molecular Microbiology *63*, 507–520.

Blickling, S., and Knäblein, J. (1997). Feedback inhibition of dihydrodipicolinate synthase enzymes by L-lysine. Biol. Chem. *378*, 207–210.

Blount, Z.D., Borland, C.Z., and Lenski, R.E. (2008). Historical contingency and the evolution of a key innovation in an experimental population of Escherichia coli. Proceedings of the National Academy of Sciences *105*, 7899–7906.

Blyth, C.R. (1972). On Simpson’s Paradox and the Sure-Thing Principle. Journal of the American Statistical Association *67*, 364–366.

Boer, V.M., Amini, S., and Botstein, D. (2008). Influence of genotype and nutrition on survival and metabolism of starving yeast. Proc. Natl. Acad. Sci. U.S.A *105*, 6930–6935.

Boles, E., and Hollenberg, C.P. (1997). The molecular genetics of hexose transport in yeasts. FEMS Microbiology Reviews *21*, 85–111.

Boomsma, J.J. (2007). Kin Selection versus Sexual Selection: Why the Ends Do Not Meet. Current Biology *17*, R673–R683.

Brachmann, C.B., Davies, A., Cost, G.J., Caputo, E., Li, J., Hieter, P., and Boeke, J.D. (1998). Designer deletion strains derived from <I>Saccharomyces cerevisiae</I> S288C: A useful set of strains and plasmids for PCR-mediated gene disruption and other applications. Yeast *14*, 115–132.

Brockhurst, M.A. (2007). Population bottlenecks promote cooperation in bacterial biofilms. PLoS ONE *2*, e634.

Chao, L., and Levin, B.R. (1981). Structured habitats and the evolution of anticompetitor toxins in bacteria. Proc Natl Acad Sci U S A *78*, 6324–6328.

Choi, P.J., Cai, L., Frieda, K., and Xie, X.S. (2008). A Stochastic Single-Molecule Event Triggers Phenotype Switching of a Bacterial Cell. Science *322*, 442–446.

Chuang, J.S., Rivoire, O., and Leibler, S. (2009). Simpson’s Paradox in a Synthetic Microbial System. Science *323*, 272–275.

Colman-Lerner, A., Gordon, A., Serra, E., Chin, T., Resnekov, O., Endy, D., Pesce, C.G., and Brent, R. (2005). Regulated cell-to-cell variation in a cell-fate decision system.

Nature *437*, 699–706.

Cornelis, P. (2010). Iron uptake and metabolism in pseudomonads. Applied Microbiology and Biotechnology *86*, 1637–1645.

Cornwallis, C.K., West, S.A., Davis, K.E., and Griffin, A.S. (2010). Promiscuity and the evolutionary transition to complex societies. Nature *466*, 969–972.

Darch, S.E., West, S.A., Winzer, K., and Diggle, S.P. (2012). Density-Dependent Fitness Benefits in Quorum-Sensing Bacterial Populations. PNAS *109*, 8259–8263.

Darwin, C. (1859). On the origin of species by means of natural selection, or, the preservation of favoured races in the struggle for life (London: J. Murray).

Datta, M.S., Korolev, K.S., Cvijovic, I., Dudley, C., and Gore, J. (2013). Range expansion promotes cooperation in an experimental microbial metapopulation. PNAS.

Davidson, C.J., and Surette, M.G. (2008). Individuality in Bacteria. Annual Review of Genetics *42*, 253–268.

Delbrück, M. (1945). The burst size distribution in the growth of bacterial viruses (bacteriophages). Journal of Bacteriology *50*, 131.

Diggle, S.P., Griffin, A.S., Campbell, G.S., and West, S.A. (2007). Cooperation and conflict in quorum-sensing bacterial populations. Nature *450*, 411–414.

Dimitrov, L.N., Brem, R.B., Kruglyak, L., and Gottschling, D.E. (2009). Polymorphisms in Multiple Genes Contribute to the Spontaneous Mitochondrial Genome Instability of Saccharomyces cerevisiae S288C Strains. Genetics *183*, 365–383.

Dobata, S., and Tsuji, K. (2009). A cheater lineage in a social insect. Commun Integr Biol *2*, 67–70.

Doebeli, M., and Hauert, C. (2005). Models of cooperation based on the Prisoner's Dilemma and the Snowdrift game. Ecology Letters *8*, 748–766.

Dupré, S., and Haguenauer-Tsapis, R. (2001). Deubiquitination Step in the Endocytic Pathway of Yeast Plasma Membrane Proteins: Crucial Role of Doa4p Ubiquitin Isopeptidase. Mol. Cell. Biol. *21*, 4482–4494.

Ferenci, T. (1999). Growth of bacterial cultures' 50 years on: towards an uncertainty principle instead of constants in bacterial growth kinetics. Research in Microbiology *150*, 431–438.

Fiegna, F., and Velicer, G.J. (2003). Competitive fates of bacterial social parasites: persistence and self-induced extinction of Myxococcus xanthus cheaters. Proc. Biol. Sci *270*, 1527–1534.

Fiegna, F., Yu, Y.-T.N., Kadam, S.V., and Velicer, G.J. (2006). Evolution of an obligate social cheater to a superior cooperator. Nature *441*, 310.

Fletcher, J.A., and Doebeli, M. (2009). A simple and general explanation for the evolution of altruism. Proceedings of the Royal Society B: Biological Sciences *276*, 13 –19.

Foster, K.R. (2004). Diminishing returns in social evolution: the not-so-tragic commons. Journal of Evolutionary Biology *17*, 1058–1072.

Foster, K.R., Shaulsky, G., Strassmann, J.E., Queller, D.C., and Thompson, C.R.L. (2004). Pleiotropy as a mechanism to stabilize cooperation. Nature *431*, 693–696.

Galitski, T., Saldanha, A.J., Styles, C.A., Lander, E.S., and Fink, G.R. (1999). Ploidy Regulation of Gene Expression. Science *285*, 251–254.

García, J.C., and Kotyk, A. (1988). Uptake of L-lysine by a double mutant of Saccharomyces cerevisiae. Folia Microbiol. (Praha) *33*, 285–291.

Gentleman, R.C., Carey, V.J., Bates, D.M., Bolstad, B., Dettling, M., Dudoit, S., Ellis, B., Gautier, L., Ge, Y., Gentry, J., et al. (2004). Bioconductor: open software development for computational biology and bioinformatics. Genome Biol *5*, R80.

Ghaemmaghami, S., Huh, W.-K., Bower, K., Howson, R.W., Belle, A., Dephoure, N., O'Shea, E.K., and Weissman, J.S. (2003). Global analysis of protein expression in yeast. Nature *425*, 737–741.

Gilbert, O.M., Foster, K.R., Mehdiabadi, N.J., Strassmann, J.E., and Queller, D.C. (2007). High relatedness maintains multicellular cooperation in a social amoeba by controlling cheater mutants. Proceedings of the National Academy of Sciences *104*, 8913–8917.

Gordon, A.J.E., Halliday, J.A., Blankschien, M.D., Burns, P.A., Yatagai, F., and Herman, C. (2009). Transcriptional infidelity promotes heritable phenotypic change in a bistable gene network. PLoS Biol *7*, e44.

Gore, J., Youk, H., and van Oudenaarden, A. (2009). Snowdrift game dynamics and facultative cheating in yeast. Nature *459*, 253–256.

Görke, B., and Stülke, J. (2008). Carbon catabolite repression in bacteria: many ways to make the most out of nutrients. Nat Rev Micro *6*, 613–624.

Gray, J.V., Petsko, G.A., Johnston, G.C., Ringe, D., Singer, R.A., and Werner-Washburne, M. (2004). “Sleeping beauty”: quiescence in Saccharomyces cerevisiae. Microbiol. Mol. Biol. Rev *68*, 187–206.

Greig, D., and Travisano, M. (2004). The Prisoner’s Dilemma and polymorphism in yeast SUC genes. Proceedings of the Royal Society B: Biological Sciences *271*, S25–S26.

Grenson, M. (1966). Multiplicity of the amino acid permeases in Saccharomyces cerevisiae. II. Evidence for a specific lysine-transporting system. Biochim. Biophys. Acta *127*, 339–346.

Grenson, M., Hou, C., and Crabeel, M. (1970a). Multiplicity of the amino acid permeases in Saccharomyces cerevisiae. IV. Evidence for a general amino acid permease. J. Bacteriol *103*, 770–777.

Grenson, M., Hou, C., and Crabeel, M. (1970b). Multiplicity of the Amino Acid Permeases in Saccharomyces Cerevisiae IV. Evidence for a General Amino Acid Permease. J. Bacteriol. *103*, 770–777.

Griffin, A.S., West, S.A., and Buckling, A. (2004). Cooperation and competition in pathogenic bacteria. Nature *430*, 1024–1027.

Guthrie, C., and Fink, G.R. (1991). Guide to yeast genetics and molecular biology (Academic Press, Inc.).

Hahne, F., LeMeur, N., Brinkman, R.R., Ellis, B., Haaland, P., Sarkar, D., Spidlen, J., Strain, E., and Gentleman, R. flowCore: a Bioconductor package for high throughput flow cytometry. BMC Bioinformatics *10*, 106–106.

Hamilton, W.D. (1964). The genetical evolution of social behaviour I and II. Journal of Theoretical Biology *7*, 1–52.

Hansen, S.K., Rainey, P.B., Haagensen, J.A.J., and Molin, S. (2007). Evolution of species interactions in a biofilm community. Nature *445*, 533–536.

Harcombe, W. (2010). Novel cooperation experimentally evolved between species. Evolution *64*, 2166–2172.

Hardin, G. (1968). The Tragedy of the Commons. Science *162*, 1243–1248.

Hauert, C., and Doebeli, M. (2004). Spatial structure often inhibits the evolution of cooperation in the snowdrift game. Nature *428*, 643–646.

Hauert, C., Monte, S.D., Hofbauer, J., and Sigmund, K. (2002). Volunteering as Red Queen Mechanism for Cooperation in Public Goods Games. Science *296*, 1129–1132.

Hauert, C., Holmes, M., and Doebeli, M. (2006). Evolutionary games and population dynamics: maintenance of cooperation in public goods games. Proc. R. Soc. B *273*, 2565–2571.

Heath, K.D., and Tiffin, P. (2009). Stabilizing Mechanisms in a Legume-Rhizobium Mutualism. Evolution *63*, 652–662.

Hollén, L.I., and Radford, A.N. (2009). The development of alarm call behaviour in mammals and birds. Animal Behaviour *78*, 791–800.

Hughes, W.H. (1955). The Inheritance of Differences in Growth Rate in Escherichia coli. J Gen Microbiol *12*, 265–268.

Ichinose, G., and Arita, T. (2008). The role of migration and founder effect for the evolution of cooperation in a multilevel selection context. Ecological Modelling *210*, 221–230.

Jimenez, P.N., Koch, G., Thompson, J.A., Xavier, K.B., Cool, R.H., and Quax, W.J. (2012). The Multiple Signaling Systems Regulating Virulence in Pseudomonas aeruginosa. Microbiol. Mol. Biol. Rev. *76*, 46–65.

Kaufmann, B.B., Yang, Q., Mettetal, J.T., and van Oudenaarden, A. (2007). Heritable stochastic switching revealed by single-cell genealogy. PLoS Biol *5*, e239.

Kerr, B., Neuhauser, C., Bohannan, B.J.M., and Dean, A.M. (2006). Local migration promotes competitive restraint in a host–pathogen “tragedy of the commons.”Nature *442*, 75–78.

Kiers, E.T., Rousseau, R.A., West, S.A., and Denison, R.F. (2003). Host sanctions and the legume–rhizobium mutualism. Nature *425*, 78–81.

Killingback, T., Bieri, J., and Flatt, T. (2006). Evolution in Group-Structured Populations Can Resolve the Tragedy of the Commons. Proceedings: Biological Sciences *273*, 1477–1481.

Kim, Y., Ramirez-Montealegre, D., and Pearce, D.A. (2003). A role in vacuolar arginine transport for yeast Btn1p and for human CLN3, the protein defective in Batten disease.

Proceedings of the National Academy of Sciences of the United States of America *100*, 15458–15462.

Kitamoto, K., Yoshizawa, K., Ohsumi, Y., and Anraku, Y. (1988). Dynamic aspects of vacuolar and cytosolic amino acid pools of Saccharomyces cerevisiae. J. Bacteriol *170*, 2683–2686.

Kümmerli, R., and Brown, S.P. (2010). Molecular and regulatory properties of a public good shape the evolution of cooperation. Proceedings of the National Academy of Sciences *107*, 18921 –18926.

Kümmerli, R., Jiricny, N., Clarke, L.S., West, S.A., and Griffin, A.S. (2009). Phenotypic plasticity of a cooperative behaviour in bacteria. Journal of Evolutionary Biology *22*, 589–598.

Lande, R. (1977). On Comparing Coefficients of Variation. Systematic Zoology *26*, 214.

Lang, G.I., Botstein, D., and Desai, M.M. (2011). Genetic Variation and the Fate of Beneficial Mutations in Asexual Populations. Genetics *188*, 647–661.

Lee, H.H., Molla, M.N., Cantor, C.R., and Collins, J.J. (2010). Bacterial charity work leads to population-wide resistance. Nature *467*, 82–85.

Lenski, R.E., and Travisano, M. (1994). Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations. Proceedings of the National Academy of Sciences *91*, 6808 –6814.

Levy, S.F., Ziv, N., and Siegal, M.L. (2012). Bet Hedging in Yeast by Heterogeneous, Age-Correlated Expression of a Stress Protectant. PLoS Biol *10*, e1001325.

Li, H., and Durbin, R. (2009). Fast and accurate short read alignment with Burrows–Wheeler transform. Bioinformatics *25*, 1754 –1760.

Li, H., Handsaker, B., Wysoker, A., Fennell, T., Ruan, J., Homer, N., Marth, G., Abecasis, G., Durbin, R., and 1000 Genome Project Data Processing Subgroup (2009). The Sequence Alignment/Map format and SAMtools. Bioinformatics *25*, 2078 –2079.

Lin, C.H., MacGurn, J.A., Chu, T., Stefan, C.J., and Emr, S.D. (2008). Arrestin-Related Ubiquitin-Ligase Adaptors Regulate Endocytosis and Protein Turnover at the Cell Surface. Cell *135*, 714–725.

Liu, H., Styles, C.A., and Fink, G.R. (1996). Saccharomyces cerevisiae S288C has a mutation in FLO8, a gene required for filamentous growth. Genetics *144*, 967.

Ljungdahl, P.O., and Daignan-Fornier, B. (2012). Regulation of Amino Acid, Nucleotide, and Phosphate Metabolism in Saccharomyces Cerevisiae. Genetics *190*, 885–929.

Luhtala, N., and Odorizzi, G. (2004). Bro1 Coordinates Deubiquitination in the Multivesicular Body Pathway by Recruiting Doa4 to Endosomes. J Cell Biol *166*, 717–729.

Lukas, D., and Clutton-Brock, T. (2012). Cooperative breeding and monogamy in mammalian societies. Proc. R. Soc. B *279*, 2151–2156.

MacLean, R.C., Fuentes-Hernandez, A., Greig, D., Hurst, L.D., and Gudelj, I. (2010). A Mixture of “Cheats” and “Co-Operators” Can Enable Maximal Group Benefit. PLoS Biol *8*, e1000486.

Magasanik, B., and Kaiser, C.A. (2002). Nitrogen regulation in Saccharomyces cerevisiae. Gene *290*, 1–18.

Marzuki, S., Hall, R.M., and Linnane, A.W. (1974). Induction of respiratory incompetent mutants by unsaturated fatty acid depletion in Saccharomyces cerevisiae. Biochem. Biophys. Res. Commun *57*, 372–378.

Maynard Smith, J. (1964). Group Selection and Kin Selection. Nature *201*, 1145–1147.

Maynard Smith, J., and Haigh, J. (1974). The Hitch-Hiking Effect of a Favourable Gene. Genetics Research *23*, 23–35.

Mitchell, A., Romano, G.H., Groisman, B., Yona, A., Dekel, E., Kupiec, M., Dahan, O., and Pilpel, Y. (2009). Adaptive prediction of environmental changes by microorganisms. Nature *460*, 220–224.

Momeni, B., Chen, C.-C., Hillesland, K., Waite, A., and Shou, W. (2011). Using artificial systems to explore the ecology and evolution of symbioses. Cellular and Molecular Life Sciences *68*, 1353–1368.

Monod, J. (1949). The growth of bacterial cultures. Annual Reviews in Microbiology *3*, 371–394.

Morgan, A.D., Quigley, B.J.Z., Brown, S.P., and Buckling, A. (2012a). Selection on non-social traits limits the invasion of social cheats. Ecology Letters n/a–n/a.

Morgan, A.D., Quigley, B.J.Z., Brown, S.P., and Buckling, A. (2012b). Selection on non-social traits limits the invasion of social cheats. Ecology Letters *15*, 841–846.

Morgan, A.D., Quigley, B., Brown, S., and Buckling, A. (2012c). Selection on non-social

traits limits the invasion of social cheats. Ecology Letters.

Morrison, C.E., and Lichstein, H.C. (1976). Regulation of lysine transport by feedback inhibition in Saccharomyces cerevisiae. J. Bacteriol *125*, 864–871.

Mortimer, R.K., Romano, P., Suzzi, G., and Polsinelli, M. (1994). Genome renewal: a new phenomenon revealed from a genetic study of 43 strains of Saccharomyces cerevisiae derived from natural fermentation of grape musts. Yeast *10*, 1543–1552.

Muhlrad, D., and Parker, R. (1999). Aberrant mRNAs with extended 3' UTRs are substrates for rapid degradation by mRNA surveillance. RNA *5*, 1299–1307.

Nagai, S. (1969). High-frequency production of respiratory mutants in yeast under nutritional deficiencies. Mutat. Res *8*, 557–564.

Nagai, T., Ibata, K., Park, E.S., Kubota, M., Mikoshiba, K., and Miyawaki, A. (2002). A variant of yellow fluorescent protein with fast and efficient maturation for cell-biological applications. Nat Biotech *20*, 87–90.

Naumov, G.I., Naumova, E.S., Sancho, E.D., and Korhbla, M.P. (1996). Polymeric SUC genes in natural populations of Saccharomyces cerevisiae. FEMS Microbiology Letters *135*, 31–35.

Ni, M., Decrulle, A.L., Fontaine, F., Demarez, A., Taddei, F., and Lindner, A.B. (2012). Pre-Disposition and Epigenetics Govern Variation in Bacterial Survival upon Stress. PLoS Genet *8*, e1003148.

Novick, A., and Szilard, L. (1950). Experiments with the Chemostat on spontaneous mutations of bacteria. Proc Natl Acad Sci U S A *36*, 708–719.

Novick, A., and Weiner, M. (1957). Enzyme induction as an all-or-none phenomenon. Proc Natl Acad Sci U S A *43*, 553–566.

Nowak, M.A., and May, R.M. (1992). Evolutionary games and spatial chaos. Nature *359*, 826–829.

Nowak, M.A., Tarnita, C.E., and Antal, T. (2010). Evolutionary dynamics in structured populations. Phil. Trans. R. Soc. B *365*, 19–30.

Pacheco, J.M., Santos, F.C., Souza, M.O., and Skyrms, B. (2009). Evolutionary dynamics of collective action in N-person stag hunt dilemmas. Proceedings of the Royal Society B: Biological Sciences *276*, 315 –321.

Parvinen, K. (2011). Adaptive Dynamics of Altruistic Cooperation in a Metapopulation:

Evolutionary Emergence of Cooperators and Defectors or Evolutionary Suicide? Bull Math Biol *73*, 2605–2626.

Parvinen, K. (2013). Joint evolution of altruistic cooperation and dispersal in a metapopulation of small local populations. Theoretical Population Biology *85*, 12–19.

Pepper, J.W. (2007). Simple Models of Assortment through Environmental Feedback. Artificial Life *13*, 1–9.

Pepper, J.W., and Smuts, B.B. (2002). A Mechanism for the Evolution of Altruism among Nonkin: Positive Assortment through Environmental Feedback. The American Naturalist *160*, 205–213.

Queller, D.C. (1992). Does population viscosity promote kin selection? Trends in Ecology & Evolution *7*, 322–324.

Raser, J.M., and O'Shea, E.K. (2004). Control of stochasticity in eukaryotic gene expression. Science *304*, 1811–1814.

R Development Core Team (2012). R: A Language and Environment for Statistical Computing (Vienna, Austria).

Refardt, D., Bergmiller, T., and Kümmerli, R. (2013). Altruism can evolve when relatedness is low: evidence from bacteria committing suicide upon phage infection. Proc. R. Soc. B *280*.

Regenberg, B., Düring-Olsen, L., Kielland-Brandt, M.C., and Holmberg, S. (1999a). Substrate specificity and gene expression of the amino-acid permeases in Saccharomyces cerevisiae. Curr. Genet. *36*, 317–328.

Regenberg, B., Düring-Olsen, L., Kielland-Brandt, M.C., and Holmberg, S. (1999b). Substrate specificity and gene expression of the amino-acid permeases in <i>Saccharomyces cerevisiae</i> Current Genetics *36*, 317–328.

Reid, R.J.D., Lisby, M., and Rothstein, R. (2002). Cloning-free genome alterations in Saccharomyces cerevisiae using adaptamer-mediated PCR. Meth. Enzymol *350*, 258–277.

Rutherford, S.L., and Lindquist, S. (1998). Hsp90 as a capacitor for morphological evolution. Nature *396*, 336–342.

Santos, M., and Szathmáry, E. (2008). Genetic hitchhiking can promote the initial spread of strong altruism. BMC Evolutionary Biology *8*, 281.

Sarkar, D., Meur, N.L., and Gentleman, R. (2008). Using flowViz to visualize flow cytometry data. Bioinformatics *24*, 878–879.

Schaber, J.A., Carty, N.L., McDonald, N.A., Graham, E.D., Cheluvappa, R., Griswold, J.A., and Hamood, A.N. (2004). Analysis of Quorum Sensing-Deficient Clinical Isolates of Pseudomonas Aeruginosa. J Med Microbiol *53*, 841–853.

Schwilk, D.W., and Kerr, B. (2002). Genetic niche-hiking: an alternative explanation for the evolution of flammability. Oikos *99*, 431–442.

Shou, W., Ram, S., and Vilar, J.M.G. (2007). Synthetic cooperation in engineered yeast populations. Proceedings of the National Academy of Sciences *104*, 1877–1882.

Simpson, E.H. (1951). The Interpretation of Interaction in Contingency Tables. Journal of the Royal Statistical Society. Series B (Methodological) *13*, 238–241.

De Smidt, O., du Preez, J.C., and Albertyn, J. (2008). The alcohol dehydrogenases of Saccharomyces cerevisiae: a comprehensive review. FEMS Yeast Res *8*, 967–978.

Springael, J.Y., Galan, J.M., Haguenauer-Tsapis, R., and André, B. (1999). NH4+-induced down-regulation of the Saccharomyces cerevisiae Gap1p permease involves its ubiquitination with lysine-63-linked chains. J. Cell. Sci. *112 (Pt 9)*, 1375–1383.

Spudich, J.L., and Koshland, D.E. (1976). Non-genetic individuality: chance in the single cell. Nature *262*, 467–471.

Strassmann, J.E., Zhu, Y., and Queller, D.C. (2000). Altruism and social cheating in the social amoeba Dictyostelium discoideum. Nature *408*, 965–967.

Strassmann, J.E., Gilbert, O.M., and Queller, D.C. (2011). Kin Discrimination and Cooperation in Microbes. Annual Review of Microbiology *65*, 349–367.

Suzuki, S., and Kimura, H. (2011). Oscillatory dynamics in the coevolution of cooperation and mobility. Journal of Theoretical Biology *287*, 42–47.

Swaminathan, S., Amerik, A.Y., and Hochstrasser, M. (1999). The Doa4 deubiquitinating enzyme is required for ubiquitin homeostasis in yeast. Mol. Biol. Cell *10*, 2583–2594.

Sychrova, H., and Chevallier, M.R. (1993). Cloning and sequencing of the Saccharomyces cerevisiae gene LYP1 coding for a lysine-specific permease. Yeast *9*, 771–782.

Tagkopoulos, I., Liu, Y.-C., and Tavazoie, S. (2008). Predictive Behavior Within

Microbial Genetic Networks. Science *320*, 1313–1317.

Trivers, R.L. (1971). The Evolution of Reciprocal Altruism. The Quarterly Review of Biology *46*, 35–57.

Tsai, I.J., Bensasson, D., Burt, A., and Koufopanou, V. (2008). Population genomics of the wild yeast Saccharomyces paradoxus: Quantifying the life cycle. PNAS *105*, 4957–4962.

Turner, P.E., and Chao, L. (1999). Prisoner's dilemma in an RNA virus. Nature *398*, 441–443.

Umebayashi, K., and Nakano, A. (2003). Ergosterol is required for targeting of tryptophan permease to the yeast plasma membrane. J Cell Biol *161*, 1117–1131.

Veatch, J.R., McMurray, M.A., Nelson, Z.W., and Gottschling, D.E. (2009). Mitochondrial Dysfunction Leads to Nuclear Genome Instability via an Iron-Sulfur Cluster Defect. Cell *137*, 1247–1258.

Velicer, G.J., and Vos, M. (2009). Sociobiology of the Myxobacteria. Annu. Rev. Microbiol. *63*, 599–623.

Verbruggen, E., Mouden, C.E., Jansa, J., Akkermans, G., Bücking, H., West, S.A., and E. Toby Kiers (2012). Spatial Structure and Interspecific Cooperation: Theory and an Empirical Test Using the Mycorrhizal Mutualism. The American Naturalist *179*, E133–E146.

Verzijlbergen, K.F., Menendez-Benito, V., Welsem, T. van, Deventer, S.J. van, Lindstrom, D.L., Ovaa, H., Neefjes, J., Gottschling, D.E., and Leeuwen, F. van (2010). Recombination-induced tag exchange to track old and new proteins. PNAS *107*, 64–68.

Vos, M., and Velicer, G.J. (2009). Social Conflict in Centimeter-and Global-Scale Populations of the Bacterium Myxococcus xanthus. Current Biology *19*, 1763–1767.

Waddington, C.H. (1942). Canalization of Development and the Inheritance of Acquired Characters. Nature Nature *150*, 563–565.

Waddington, C.H. (1953). Genetic Assimilation of an Acquired Character. Evolution *7*, 118–126.

Waddington, C.H. (1956). Genetic Assimilation of the Bithorax Phenotype. Evolution *10*, 1–13.

Waddington, C.H. (1957). The genetic basis of the "Assimilated Bithorax" stock. Journal

of Genetics *55*, 241–245.

Wade, M.J. (1978). A Critical Review of the Models of Group Selection. The Quarterly Review of Biology *53*, 101–114.

Wang, Z., Gerstein, M., and Snyder, M. (2009). RNA-Seq: a revolutionary tool for transcriptomics. Nat Rev Genet *10*, 57–63.

West, S.A., Pen, I., and Griffin, A.S. (2002). Cooperation and Competition Between Relatives. Science *296*, 72–75.

West, S.A., Diggle, S.P., Buckling, A., Gardner, A., and Griffin, A.S. (2007). The Social Lives of Microbes. Annual Review of Ecology, Evolution, and Systematics *38*, 53–77.

Wieczorke, R., Krampe, S., Weierstall, T., Freidel, K., Hollenberg, C.P., and Boles, E. (1999). Concurrent knock-out of at least 20 transporter genes is required to block uptake of hexoses in Saccharomyces cerevisiae. FEBS Letters *464*, 123–128.

World Health Organization (2008). The global burden of disease: 2004 update (Geneva: World Health Organization).

Wright, S. (1931). Evolution in Mendelian Populations. Genetics *16*, 97–159.

Yu, R.C., Pesce, C.G., Colman-Lerner, A., Lok, L., Pincus, D., Serra, E., Holl, M., Benjamin, K., Gordon, A., and Brent, R. (2008). Negative feedback that improves information transmission in yeast signalling. Nature *456*, 755–761.

Yvert, G., Brem, R.B., Whittle, J., Akey, J.M., Foss, E., Smith, E.N., Mackelprang, R., and Kruglyak, L. (2003). Trans-acting regulatory variation in Saccharomyces cerevisiae and the role of transcription factors. Nature Genetics *35*, 57–64.

Zhuravel, D., Fraser, D., St-Pierre, S., Tepliakova, L., Pang, W.L., Hasty, J., and Kærn, M. (2010). Phenotypic impact of regulatory noise in cellular stress-response pathways. Syst Synth Biol *4*, 105–116.