Phenotypic heterogeneity and evolutionary games in microbial populations

by

David W. Healey
B.S. Biochemistry
Brigham Young University (2009)

Submitted to the Department of Biology
In Partial Fulfillment of the Requirements for the Degree of

DOCTOR OF PHILOSOPY

At the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Signature of Author……
Department of Biology
May 16, 2015

Certified by……
Jeff Gore
Assistant Professor of Physics
Thesis Supervisor

Accepted by……
Michael Hemann
Associate Professor of Biology
Co-Chair, Graduate Committee

Phenotypic heterogeneity and evolutionary games in microbial populations

by

David W. Healey

B.S. Biochemistry
Brigham Young University, 2009

Submitted to the Department of Biology on May 16, 2015 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy in Biology

ABSTRACT

One of the most interesting discoveries of the last decade is the surprising degree of phenotypic variability between individual cells in clonal microbial populations, even in identical environments. While some variation is an inevitable consequence of low numbers of regulatory molecules in cells, the magnitude of the variability is nevertheless an evolvable trait whose quantitative parameters can be "tuned" by the biochemical characteristics and architecture of the underlying gene network. This raises the question of what adaptive advantage might be conferred to cells that implement high variation in their decision-making. Currently, the predominant answer in the field is that stochastic gene expression allows cells to "hedge their bets" against unpredictable and potentially catastrophic environmental shifts. We proposed and experimentally demonstrated an alternative solution: that heterogeneity implements the evolutionarily stable mixed strategy (or mixed ESS), from the field of evolutionary game theory. In a mixed ESS, phenotypic heterogeneity is a result of competitive interactions between cells in the population rather than a response to uncertain environments, so unlike with bet-hedging, in a mixed ESS the evolutionary fitness of different phenotypes is frequency dependent. Each phenotype can invade the other when rare, and the resulting equilibrium—the stable mix of the two—is not necessarily the one that maximizes the population's fitness. We demonstrated these and other predictions of the mixed ESS using engineered "pure strategist" strains of the yeast GAL network. We demonstrated also that the wild type mixed strategist can invade both pure strategists and is uninvasible by either. Taken together, our results provide experimental evidence that evolutionary hawk-dove games between identical cells can explain the phenotypic heterogeneity found in clonal microbial populations.

Thesis supervisor: Jeff Gore
Title: Assistant Professor of Physics

ACKNOWLEDGEMENTS

There are many people without whom this work would not have been possible. Firstly I must thank my advisor Jeff Gore. Jeff has been an outstanding research mentor; I was truly lucky to have been invited to join his lab. He has spent an enormous amount of time and effort meeting with me and guiding the research along; much more than is typical for graduate advisors. He has championed tirelessly the cause of the evolutionarily stable mixed strategies; if frequency dependence gains recognition as a possible explanation for microbial phenotypic heterogeneity it will be because of Jeff. I am also very grateful for his guidance in my professional career, which has often taken a somewhat unorthodox track. He has always been supportive and generous in letting me explore options outside of academia; his motto has always been *the important thing is to find what you want to do.* It's safe to say I would not have gotten to this point if it had not been for his encouragement to keep at it. It's common knowledge in graduate school that your advisor makes or breaks your experience, and I could not have found a better one.

I am grateful to Mike Laub and Matt Vander Heiden for serving on my thesis committee these years, especially for their input and for their encouragement in my career plans. I thank also Mike Springer for helpful conversations about the yeast GAL network, for serving on my thesis defense committee, and for doing so much of the recent work on the yeast GAL network that has helped shed light on the biochemistry of galactose consumption. Another faculty mentor I would like to thank is Frank Solomon, who has given me (and so many others) some of the best practical career advice I have ever had.

All of my Gore group labmates deserve special recognition as well. Kevin Axelrod, of course, who worked closely with me on the experimental aspect of this project in its infancy, and now again at the very end, and whose gazelle-like dodgeball dodging skills inspired a generation. Alvaro Sanchez, now of the Rowland Institute, for talking over some of the more slippery theoretical points early on, and Andrea Velenich for letting me talk his ear off in the office about whatever topic was distracting me that day. Eugene Yurtsev and Yonatan Friedman deserve also special recognition for their help with computational work, especially their developing of the phenomenal FlowCytometryTools Python library. May it make them famous one day. I am tempted to name everyone else in the lab by name as well, but I will move on.

I should also recognize the support of my parents, LeAnn Gowans and Troy Healey, for raising me to be curious, for educating, supporting and encouraging me, and for teaching me how to work and think critically. Ryan Davis, Thad Gillespie, and Bryan Haslam have spent many hours listening to me either rant or rave about my project, depending on how the experiments were going that day. Their input was invaluable to understanding and communication about my research.

Finally, and most importantly, I must acknowledge the support of my wife, Anne. We met and married during my PhD, and I am grateful to her for her endless encouragement and tolerance of late-night trips to the lab. Marrying her has given my life meaning and peace and joy beyond anything I ever dreamed, and our little family, which now includes a three-month old son, Maxwell, is my reason for everything.

TABLE OF CONTENTS

CHAPTER ONE: PHENOTYPIC HETEROGENEITY AND SURVIVAL STRATEGIES

PHENOTYPIC HETEROGENEITY IN CLONAL CELLULAR POPULATIONS

Historical perspective

For decades, it was taken for granted in the field of biology that the phenotype of a cell was basically a deterministic function of the cell's genetic material and its environment. The cell's DNA sequence and epigenetics would contain information on what proteins to make and when to make them, while the environment would modulate the levels of gene expression through protein signaling cascades. In this model, two cells with identical genetic material and in identical environments will adopt identical phenotypes. In other words, genetic clones would also be phenotypic clones.

This model of gene expression is somewhat simplistic, however. It was observed as early as the 1940s and 50s that even apparently clonal microbial populations exhibited meaningful phenotypic differences, such as variable phage burst sizes (*1*), cell division times (*2*), flagellar phases (*3*), and β-galactosidase concentrations (*4*). Some variation arises inevitably from Poisson distributions of small numbers of important regulator molecules(*5*). If a bacterium, for example, has exactly three copies of a certain transcription factor when it divides, at best, one daughter will have half the concentration of that factor than her sister, while there would be a relatively high probability that one daughter will have all three, and the other none at all. Still, for decades, non-genetic individuality was considered a statistical side-

effect, and the exception rather than the rule. Perhaps because microbial gene expression was studied from the average behavior of cells in a population, few would have guessed the magnitude of phenotypic heterogeneity that became apparent decades later.

In the late 1990s and early 2000s, technological advances in single-celled measurements allowed the first in-depth look at cell-to-cell variation in gene expression . Clonal populations, rather than being phenotypically similar, showed a remarkable amount of phenotypic heterogeneity even in homogenous environments (*6, 7*). Furthermore, the magnitude of variation was not uniform from gene to gene; controlling for cell-wide differences ("extrinsic" variation) such as size, volume, and stage of the cell cycle, there was still a high degree of individual variation (*6, 8, 9*). This observation raised two important questions: First, *how* do cells introduce variable degrees of stochasticity into their cellular decision-making, and secondly, *why*?

Phenotypic heterogeneity is an evolvable trait

The question of how clonal microbial populations managed such high phenotypic variation drove a whole new field of biochemical inquiry into the mechanisms underlying stochastic gene expression. Several mechanisms have since been described. We will consider some of them here:

Firstly, as mentioned earlier, low numbers of molecules in cells necessarily produce a high variance from cell to cell, and variation of a gene can thus be increased by lowering the average numbers of regulatory molecules (or mRNA) of that gene within a single cell (*5, 10*). All other things equal, if a cell has fewer genetic copies of a gene, or very low fidelity transcription, it will produce few mRNAs on average with high variance from cell to cell. Inefficient transcription combined with efficient translation is alone sufficient to create a high amount of variation between cells (*11*). Moreover, rather than occurring at a roughly uniform rate, transcription can also take place via "bursting," wherein a background of low transcription is punctuated by sudden furious bouts of transcription that occur at intervals (*12-14*) . Most intriguingly, perhaps, the topology of the gene network itself can promote randomness—positive feedback loops can create thresholds and tipping points in gene expression. In some extreme cases, genes are expressed bimodally in a clonal population: each cell expresses either at high levels or not at all (*15*), as in the case of the yeast GAL network (*16*) or the lambda phage lysis and lysogeny genes (*17-19*).

What is interesting about these mechanisms is that they are all evolutionarily tunable. Gene copy numbers, binding affinities, and network structures can be altered by genetic mutation and maintained across generations (*20*). Consequently, not only the phenotypes themselves can be selected for, but also the level of stochasticity in the gene expression levels. (*21-24*). Variation is an evolvable trait.

Which brings us to the second question about phenotypic heterogeneity: *why variability?* In biology, randomness is typically associated with imprecision and noise, to be controlled and limited if possible (*25, 26*). Given some environment, different phenotypes will confer different fitness on the individuals expressing them, so the winning evolutionary strategy would seem to be: "adopt the most fit phenotype for every environment as precisely as possible." However, given that heterogeneity can be evolutionarily tuned, and given that some genetic systems exhibit a great deal of heterogeneity, what might be the adaptive benefit of a high degree of randomness in gene expression?

We will first consider the predominant answer in the literature: namely that heterogeneity allows populations to evolutionarily "hedge their bets" in unpredictable environments (*27*). **Bet-hedging** is by far the most commonly-cited evolutionary explanation for phenotypic heterogeneity; indeed, the term has nearly become synonymous with the phenomenon of phenotypic heterogeneity itself (*28*). We will next consider a frequently-overlooked alternative explanation: the **evolutionarily stable mixed strategy (or mixed ESS)** from evolutionary game theory, which, although theoretically well-established, is rarely considered in the context of microbial phenotypic heterogeneity, and has yet to be demonstrated experimentally in that context. Although both methods promote phenotypic heterogeneity, they are distinct in several ways that make experimentally testable predictions. The main difference is that a mixed ESS is the stable result of interactions within the population, and can happen even in the absence of

environmental uncertainty, while bet-hedging is independent of interactions and only occurs in uncertain environments. We will briefly contrast the mixed ESS with a similar but also distinct third form of heterogeneity, the altruistic **division of labor**. Before we attempt to demonstrate that any particular microbial heterogeneity is more consistent with one theory than another, we will need to explore the three theories in more detail.

BET-HEDGING

One reason why an organism may not benefit from completely deterministic decision-making is that the environment may be somewhat uncertain, and organisms may not have the time or ability to respond to sudden (and potentially catastrophic) shifts. They may therefore benefit from "hedging their bets" by adopting a phenotype that is suboptimal for their current environment, but which may increase their overall chances of survival in some possible future environment (*27, 29-31*). The old investment adage about not putting "all your eggs in one basket" may just as well apply to microbial survival strategies as it does to finance.

Many examples of bet-hedging abound in the literature. Cohen (*32, 33*) initially developed the theory in the 1960s in the context of annual plants, which can either germinate, yielding more seeds in numbers that depend on environmental conditions, or remain dormant in the soil to germinate a different year. Dormant seeds decrease the plant's yield during favorable years, but in unfavorable years, germinating plants may die while dormant seeds increase the chances for survival.

The optimal strategy may therefore be for each seed to germinate sometimes and remain dormant sometimes (*34-37*). Ecologists have subsequently described bet-hedging strategies in a diverse range of plants and animals, in everything from body mass in wild boars (*38*) to mating choice in Salmon (*39*).

Some of the early examples of phenotypic heterogeneity in microbial populations resemble the plant germination problem and were candidates for being modeled as bet-hedging. For example, in multiple experimental bacterial systems it was observed that identical cells stochastically switch between fast-growing phenotypes and slow-growing "persister" cells or spores (*40, 41*), Persister cells and spores have much lower fitness in plentiful resources, but also have much higher survival of extreme environmental stress such as in the presence of antibiotics. Since then, bet hedging has been used to explain scores of observed microbial phenotypic heterogeneity (*42-48*). It has become so popular as an explanation, however, that it is common for researchers to explain phenotypic heterogeneity as a response to environmental uncertainty with little empirical support and without considering alternative explanations (for review, see (*28*)).

Definition of bet-hedging

The defining concept of bet-hedging is that maximizing the geometric mean fitness of an individual over time often comes at the cost of lower arithmetic mean fitness (*27, 34, 35*). The geometric mean fitness is always less than the corresponding arithmetic mean by a function of the variance. Bet-hedging strategies maximize

geometric mean fitness by reducing the fitness variation across time. For example, consider two different cases of crop yields across two consecutive years:

Table 1. Crop yields across consecutive yields

	Case 1	Case 2
Year 1 yield	10	4
Year 2 yield	2	6
Mean yield	**6**	5
Combined yield (y1 x y2)	20	**24**
Geometric mean yield	4.47	**4.89**

In Case 2 (which might be thought of as a bet-hedging strategy relative to Case 1), even though the mean yield is lower, the combined yield over the two years is higher. Given that the geometric mean is the square root of the combined yield, this example illustrates the tradeoff between geometric and arithmetic mean that is at the center of bet-hedging.

A simple model of bet-hedging considers an organism with two possible phenotypes: *X* and *Y*. The organism exists in an environment that for each period of time can be in one of two states, which we will call Good (*G*) and Bad (*B*). The environment fluctuates between Good and Bad randomly, with the probability over time of Good environments denoted as P_G, and bad as $P_B = (1\text{-}P_G)$. The evolutionary yield, $Y_{i,j}$, is defined as the number of progeny produced by in individual of

phenotype *i* in in environment *j*). An individual's fitness in environment *j, Wj*, is an average of the fitness of each phenotype in that environment, weighted by f_i, the probability of the individual's stochastically adopting phenotype *i*:

$$W_j = Y_{x,j} f_x + Y_{y,j}(1 - f_x)$$

As mentioned above, bet-hedging strategies maximize the geometric mean fitness of the individual across time, which is given as:

$$\overline{W} = [Y_{x,G} f_x + Y_{y,G}(1 - f_x)]^{P_G} [\, Y_{x,B} f_x + Y_{y,B}(1 - f_x)]^{(1-P_G)}$$

Solving for f_x^*, the probability of adopting phenotype x that maximizes geometric mean fitness, yields:

$$f_x^* = \frac{Y_{x,B} Y_{y,G} P_G - Y_{x,B} Y_{y,G} - Y_{x,G} Y_{y,B} P_G + Y_{y,G} Y_{y,B}}{\left(Y_{x,G} - Y_{y,G}\right)\left(Y_{x,B} - Y_{y,B}\right)}$$

In a common example of bet-hedging (shown in figure 1), the phenotypes are to either grow (X) or remain dormant (Y). In a good environment, a growing phenotype produces another individual ($Y_{x,g} = 2$), while in a bad environment, the growing phenotype has a 90% chance of dying ($Y_{x,b} = 0.1$). Remaining in a dormant state results merely in the survival of the individual regardless of environment ($Y_{y,g}$

$= Y_{y,b} = 1$). If good environments occur with a probability 0.8, then the optimal probability of adopting the growth phenotype would be about 0.69.

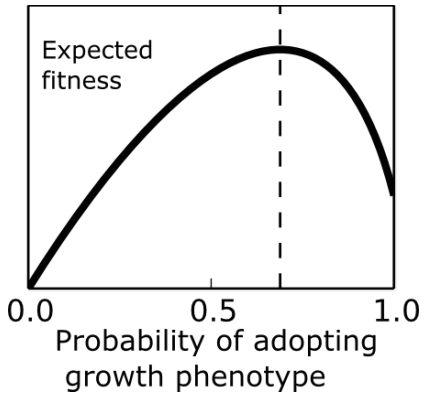

Figure 1. A simple example of bet-hedging. Geometric mean fitness for an individual seed is shown as a function of the probability that the seed will germinate vs remain dormant.

We should note here that the technical definition of bet-hedging encompasses more than just a diversified phenotypic strategy. In fact there are two categories of bet-hedging: diversified bet hedging, which we have been discussing, and conservative bet-hedging (*49*). In conservative bet-hedging, there is no stochastic "coin-flipping" between alternative strategies. Each individual makes a deterministic phenotypic choice, only the phenotype it chooses is one with a lower temporal variance and lower average fitness than the phenotype that produces the highest average fitness. While they may appear somewhat different, both strategies involve sacrificing average yield while minimizing temporal variation, so are therefore forms of bet-hedging. Following the terminology in the field of microbial heterogeneity, and since

we are primarily interested in phenotypic heterogeneity, for our purposes the term "bet-hedging" will refer specifically to diversified bet-hedging.

The benefits of bet-hedging are best intuited from a population level, but we should be careful not to assume it is primarily a population-level survival strategy the way an altruistic division of labor (discussed below) is. True, a population of bacteria which hedges its bets against the possibility of encountering antibiotics will contain some percentage of slow-growing "persister" cells at any given time, and those individuals may survive a catastrophic shock and replenish the population as a whole. But although bet-hedging, does preserve (and indeed optimizes) the long-term fitness of the population, bet-hedging solutions to uncertainty exist independently of population-level dynamics. Adopting a bet-hedging strategy is a way for cells to maximize their individual expected fitness; it is an *individual level* survival strategy. A bet-hedging clonal population is merely a collection of individuals that is each implementing an identical stochastic survival strategy for itself. Put differently, the optimal probability distribution over the available phenotypes is exactly identical for a single cell existing alone as it is for a cell in a population of a million clonal sisters. Furthermore, for an individual existing in a population of a million, the optimal probability is the same regardless of which phenotype the sister cells adopt. Note that in the above equation, the optimal probability for the cell to adopt a specific phenotype is a function only of the probabilities of "good" and "bad" environments and the fitness of each phenotype in

each environment. It is not a function of the numbers or phenotypes of any other individuals.

Of course, in the real world, it is likely that the fitness of each phenotype does, in fact, depend somewhat on the phenotypes the other members of the population. Resources can be scarce, and crowding of resources reduces payoffs for those consuming them. Accordingly, some bet-hedging models incorporate frequency or density-dependence. However, it is vital to recognize that *environmental uncertainty alone* is sufficient to drive the evolution of phenotypic heterogeneity, in the absence of any interaction with other individuals. And likewise, as we will discuss in the following two sections, frequency dependent interactions with other individuals is also alone sufficient to drive the evolution of phenotypic heterogeneity, even in the absence of any environmental uncertainty. Therefore, although uncertain environments and population interactions can coexist, they are distinct drivers of heterogeneity.

EVOLUTIONARILY STABLE MIXED STRATEGIES

Game theory and the Nash equilibrium

The second kind of heterogeneity that we will discuss is driven by frequency-dependent interactions, and is described by the evolutionarily stable mixed strategy, or mixed ESS, from the field of evolutionary game theory (EGT). Evolutionary game theory is a subset of traditional economic game theory, which concerns itself with

the question of which strategy is optimal when the payoffs of strategies depend on the strategies adopted by others. It therefore deals with inherently social situations. Such situations are called "games," and the individuals are called "players." In game theory, a player's "strategy" is a specification of what an individual will do in any situation it may find itself.

There is often confusion when discussing strategies and phenotypes in a game theory context, since both terms broadly refer to the characteristics of the cell. In order to avoid confusion, then, we will consider the immediate characteristics of the cells to be *phenotypes*, and *strategies* will be complete specifications of what individuals will do. For example, in the bet-hedging example above, the phenotypes would be "grow" and "remain dormant," while the strategies may be things like "always grow", "always remain dormant" or "remain dormant with probability p." Contrary to some common usage, then, we will consider bet-hedging to be a *strategy*, rather than a *phenotype*. This distinction will become especially important when discussing the differences between bet-hedging and evolutionarily stable mixed strategies.

The Nash Equilibrium

For simple games, the solution concept is called a Nash Equilibrium. A Nash Equilibrium occurs when all players are "content" with their chosen strategy, in the sense that they have no incentive to change their strategy, given the strategies of others(*50*). Perhaps the most canonical game in game theory is the Prisoner's

Dilemma, which illustrates the concept of a Nash Equilibrium well. In the Prisoner's Dilemma, two individuals agree to cooperate with each other, but each has a choice to actually cooperate (C) or defect (D), and they have no knowledge of the other's actions. If they both cooperate, they both receive high payoffs. If they both defect, they receive low payoffs. However, if one player cooperates, then the other can get an even higher payoff by defecting, and the cooperator receives the lowest payoff of all. The following payoff matrix—typical notation in game theory—illustrates the payoffs faced in the Prisoner's Dilemma.

Table 2. The Prisoner's Dilemma

		Player 2	
		Cooperate	Defect
Player 1	Cooperate	(3,3)	(0, 5)
	Defect	(5,0)	(1,1)

It is easy to see that, for the prisoner's dilemma, the unique Nash Equilibrium is that both players defect, because regardless of whether an opponent cooperates or defects, the best option is to defect, and when both players defect, neither has an incentive to cooperate. This illustrates an important principle: that Nash equilibria are not necessarily the "optimal" outcomes for either player. Unlike the solution to

the in the bet-hedging example described previously, the solution concept in game theory—the Nash equilibrium—is about stability rather than optimality. In the Prisoner's Dilemma, both players could get a higher payoff if they cooperated; however, cooperation is not stable.

Pure vs. mixed strategies

With the Prisoner's Dilemma, we have only considered deterministic, or "pure", strategies ("cooperate" and "defect"). There are other strategies known as "mixed" strategies, where the strategy is a probabilistic distribution over the pure strategies (for instance, in the Prisoner's Dilemma, a player's strategy could be to just flip a coin: cooperate with probability ½ and defect with probability ½.) Mixed strategies are central concepts in game theory, partly because there are some games whose most important (or only) Nash equilibria require players implement mixed strategies. For example, in the popular game of rock-paper-scissors, the only situation in which neither player has an incentive to change strategies is if both players adopt a strategy of randomizing between rock, paper, and scissors with probability 1/3 each. The link between phenotypic heterogeneity and mixed strategy equilibria will become more clear in the following sections, when we will discuss the hawk-dove game.

Evolutionary game theory and evolutionarily stable strategies

Evolutionary game theory (EGT) considers the same principles as traditional game theory, but in an evolutionary context with biological players. Evolutionary game theory differs from traditional game theory in a few respects. First, payoffs are measures of evolutionary fitness (numbers of progeny). Secondly, all strategies are determined genetically and passed on to subsequent generations, with more "rational" strategies coming to dominate the populations in the course of natural selection. Therefore, while game theory generally assumes rational players will remain at Nash equilibria through rational analysis of the relevant strategies and payoffs, in EGT the players are individuals or genes that are not able to "think" through the payoffs, but must stumble upon new strategies through random mutation and are kept at stable equilibria not because they are rationally satisfied, but because deviations are competed away through natural selection. In evolutionary game theory, then, natural selection takes the place of rationality.

Evolutionary game theory largely concerns itself with finding evolutionarily stable strategies (ESS) rather than specific Nash equilibria (*51, 52*). An evolutionarily stable strategy is one that, "if all members of the population adopt it, then no mutant strategy could invade the population under the influence of natural selection." (*52*) Evolutionarily stable strategies are considered subsets of Nash equilibria because when a population implements an evolutionarily stable strategy, no single individual can improve its fitness by switching to any other strategy, pure or mixed.[1]

[1] The statement that evolutionarily stable strategies are a subset of Nash equilibria is useful, but somewhat misleading. An ESS is a type of strategy, while a Nash

However, though every ESS implements a Nash equilibrium, not all Nash equilibria involve evolutionary stability. For example, in the hawk-dove game, discussed below, one player playing "hawk" and the other playing "dove" is a Nash equilibrium because neither player has an incentive to deviate. However, as we will see, neither of these strategies are evolutionarily stable.

Just as there exist some games from classical game theory whose only Nash equilibrium is a mixed strategy Nash equilibrium (ie rock, paper, scissors), there also exist a class of biological games whose only evolutionarily stable strategy is a mixed strategy. The simple two player versions of these games are commonly called *anti-coordination* games, also known as *snowdrift* or *chicken* games. The defining characteristic of these games is that the optimal strategy is the opposite of the opponent's strategy. If interactions are randomized (ie the population is well-mixed), then rare strategies are favored over common ones.

The hawk-dove game

The canonical example of an evolutionary anti-coordination game is the hawk-dove game, first described by Maynard Smith and Price (*53*) to explain—from a selfish gene perspective rather than a "good of the species" perspective—why some animals back down from violent conflict over resources. In this game, two animals are competing over a resource of value *V*. The animals will display aggressive

---

equilibrium is a type of outcome. An ESS implements a Nash equilibrium when every individual in the population adopts it.

behavior towards each other, but ultimately each animal has two pure strategies available: to escalate the contest to a fight (Hawk) or back down (Dove). Note that Hawks and Doves are not intended to be separate species, but rather behavioral variants of the same species. There are three possible pairings in a symmetric two-player Hawk-Dove game:

1. If Hawk meets Dove *(H,D),* Dove flees and Hawk gets the resource (*V*). Dove retreats and receives a payoff of zero.
2. If Hawk meets Hawk *(H,H)* they fight. Each has a ½ probability of getting the resource, but only after each incurs an injury cost (*C*) greater than half the value of the resource.[2] Their expected payoff is ½ (*V*-*C*). The key point is that the expected payoff of Hawk against Hawk, *E(H, H),* is less than the payoff from retreating against Hawk (H,D).
3. If Dove meets Dove (*D,D*) they share the resource. Their expected payoff is ½ (*V*)s

In matrix notation, the payoffs are as follows:

[2] If, instead, $C < \frac{1}{2}V$, then the game becomes a Prisoner's Dilemma with Hawk being the only stable strategy.

Table 3. The two-player hawk-dove game

		Player 2: Hawk	Player 2: Dove
Player 1	Hawk	$(\frac{1}{2}V - C, \frac{1}{2}V - C)$	(V, 0)
	Dove	(0 , V)	(½ V , ½ V)

Since the expected payoff of playing Hawk against Hawk, *E(H,H)* is less than the expected payoff of playing Dove against Hawk, it is apparent that the Hawk strategy cannot be evolutionarily stable. In a population of all Hawks, a single Dove will receive a higher payoff than the rest, and invade the population. Likewise Dove is also not evolutionarily stable, since against a Dove opponent, hawks can win the entire resource. *E(H,D)* > *E(D,D)*; a single Hawk will invade a population of all Doves. Thus, the pure strategies are mutually invasible. If we assume that all individuals play pure strategies (they must either be Hawks of Doves), then the stable result will be a coexistence of the two kinds of individuals. Importantly, as with the Prisoner's Dilemma earlier, the coexistence ratio will not necessarily be *optimal* for the population as a whole; indeed, in this case, it is easy to see that a population of 100% Doves gives the highest average payoff, of ½ V (see Figure 1).

Suppose now that individuals can implement a mixed strategy rather than being only Hawk or only Dove. We will define a mixed strategy as "play Hawk with

probability *p*, and Dove with probability *(1-p)"*. The ESS of our hawk-dove game, P_{ESS} is found, not by optimizing the expected as in our bet-hedging example, but by finding the fraction of hawks and doves in a population that would result in *an equal expected payoff* for playing hawk and playing dove. For an individual in a population consisting of both Hawks and Doves, the expected payoff of adopting a strategy depends on the fraction of Hawks and Doves in a population. If Hawks exist with frequency *f* and Doves with frequency *(1-f),* then for any strategy *I*, the expected payoff of adopting *I, E(I)*, is a weighted average, given as:

$$E(I) = f * E(I,H) + (1-f) * E(I,D)$$

The evolutionarily stable mixed strategy will be to choose *p* such that *E(H) = E(D).* Since for a clonal population playing a mixed strategy, $p = f$, solving the above equation generally for E(D) = E(H) yields an evolutionary stable probability:

$$P^{ESS} = \frac{E(D,D) - E(H,D)}{E(H,H) - E(D,H) - E(H,D) + E(D,D)}$$

Substituting the payoffs in table 3, we have:

$$P^{ESS} = \frac{V}{2C}$$

The key characteristic to note in this example is that evolution will favor equal payoffs rather than optimal payoffs, because only when the payoffs are equal for the phenotypes is there no incentive for any individual to switch to the other

phenotype. For $V = \frac{2}{3}C$, for example, the optimal scenario—the one which provides the highest expected fitness—is for all individuals to be Doves all the time, while the mixed ESS is to play Hawk with frequency $1/3$ (see Fig 1).

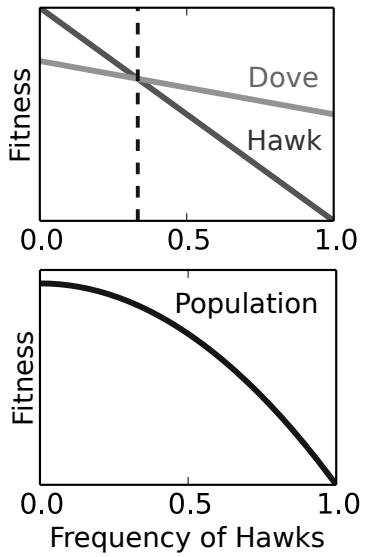

Figure 1. A simple hawk-dove game with $C = 3V$. The ESS occurs when Hawks are $1/3$ of the population, so by playing Hawk with frequency $1/3$, a clonal population can be evolutionarily stable. Note also that the payoff for the population as a whole (below) is maximized when there are no Hawks, but that such a population is prone to invasion by Hawks. Thus, the growth-optimal frequency is not necessarily evolutionarily stable.

The key characteristic of a mixed ESS, like that in the hawk-dove game, is that the stable point is not the one that maximizes the mean payout. A non-game-theoretic way to consider hawk-dove games is in the context of frequency-dependent selection. Any negative-frequency dependent scenario can be thought of as a sort of hawk-dove game, as long as the frequency dependence is strong enough to produce mutual invasibility. Given that negative frequency dependence often results in stable polymorphisms (coexistence of different species) an evolutionarily stable

mixed strategy can be thought of as a way for a single genetic entity to mimic a stable polymorphism through phenotypic heterogeneity.

MIXED ESS vs BET-HEDGING

To this point, we have introduced two unique survival strategies that both rely on phenotypic heterogeneity: bet-hedging and the evolutionarily stable mixed strategy. The question now becomes: when faced with phenotypic heterogeneity, how to tell which of the two survival strategies it might be (or might it be both?) Is the clonal population hedging its bets in response to environmental uncertainty, or has it been driven to a stable heterogeneity by a hawk-dove game between the individual phenotypes?[3]. Hitherto, explanations of microbial heterogeneity have centered almost exclusively around bet-hedging. Though recognized by some as a theoretical possibility (*28, 54-56*) evolutionarily stable mixed strategies have been largely overlooked in the literature to this point; indeed, there are no empirical demonstrations of mixed ESSs in the context of microbial phenotypic heterogeneity). It is possible, however, to experimentally probe which survival strategy (or neither, or both) is being implemented, because the models make different predictions about the fitness dynamics between the phenotypes in question. We will list the three main differences here, and go into more detail.

1. In a hawk-dove game, pure phenotypes are mutually invasible, while in bet-hedging there is no frequency dependent fitness.

[3] Given that there are still other survival strategies that produce heterogeneity

2. At the mixed ESS, all individuals in the population have the same fitness, while in a bet-hedge, each pure phenotype has a fitness advantage in different environments.
3. A bet-hedging "optimum" maximizes expected yield over time, while an evolutionarily stable strategy does not necessarily maximize yield.

1. Mutual invasibility

Evolutionarily stable mixed strategies arise from situations in which neither pure strategy is evolutionarily stable, because a population composed almost entirely of one phenotype can be invaded by the other, and each phenotype's fitness is negatively correlated with its prevalence in the population. In contrast, phenotypic fitness in bet-hedging is not a function of the population composition. By probing the phenotypes for negative frequency dependence, one can establish that a hawk-dove game is likely (at least partially) responsible for the observed heterogeneity.

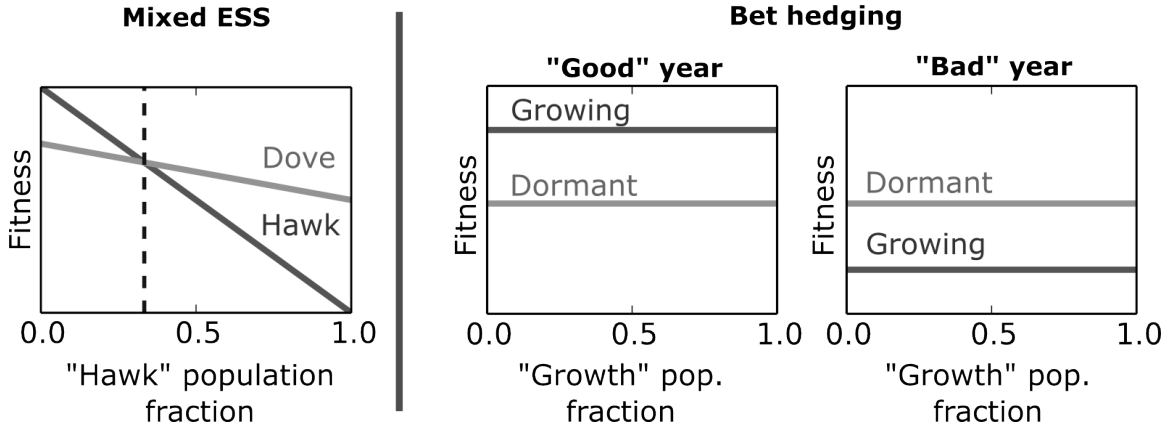

Figure 3. A side-by-side comparison of a mixed ESS (left) and a bet-hedge (right), showing fitness as a function of population frequency of each

phenotype. The parameters are those previously discussed for the hawk-dove game and the seed germination examples, respectively. The cross-shaped fitness curve in the mixed ESS show the negative frequency dependence and mutual invasibility of the pure phenotypes, while the pure phenotypes in the bet-hedge are frequency-independent in both environment types.

2. Equal fitness of any strategy at the stable point

When a clonal population implements an evolutionarily stable mixed strategy, the phenotypic fitness of all the relevant pure strategies is equal, as is every possible mixed strategy combination of them. This is best intuited by considering the counterfactual in the context of the hawk-dove game. Suppose that there existed some other mixed strategy such that when that strategy was implemented, hawks received a higher expected payoff than doves in the population. Then it would be true that a single individual would receive a higher fitness from the "hawk" phenotype than the "dove," and a single mutant individual which used the strategy "only play hawk" would have a higher expected fitness than the mixed strategy, and would increase in the population over time. This trend would increase the fraction of hawks in the overall population, which would in turn decrease the fitness of the hawks relative to the doves until it they were again equal. Thus, natural selection will generally drive the mixed strategy towards the point where hawks and doves have equal fitness, and, indeed where either those strategies or any mixed strategy of hawk and dove all have equal fitness. (Exceptions to this rule are the altruistic divisions of labor discussed below). In contrast, consider the bet-hedging example from earlier. In any given environment, one of the phenotypes is more fit than the

other—adopting the lower-fitness phenotype is merely a "bet" that the environment will change.

That is all well and good for a single environment, but we might be tempted to argue that because it is inherently a multi-environment strategy, it is somewhat unfair to limit bet-hedging to a single environment. We may wish to make the claim that with bet hedging as with mixed strategies, all strategies are still equal *over a long period of time*, taking into account all the environmental shifts. We might reason, then, that if, *over time,* one phenotype had an advantage over the other, there would exist an incentive to alter the stochastic bet-hedging strategy to include a higher probability of the more fit phenotype, and so the optimal bet-hedging phenotype must take.

However, this reasoning ignores an important difference between bet-hedging and mixed ESS: while it's true that evolution will drive a bet-hedging population towards its optimal mixing frequency, this does not mean that *every strategy* has equal fitness, as it does in the mixed ESS. Quite the reverse: in a bet hedge, over time the optimal frequency has higher fitness than any other strategy. By way of illustration, let us contrast the previous example of a mutant pure strategist hawk arising in a mixed ESS population with its counterpart in the bet-hedging example of growing and dormant phenotypes. Recall that if a mutant "pure strategist hawk" arises in a population implementing a mixed ESS, the pure strategist mutant will have *identical fitness* with the rest of the population, because in a mixed ESS the population fractions of each phenotype are such that hawks and doves have equal fitness.

However, if a population is in an optimal bet-hedge and a mutant arises with the strategy "always remain dormant", it is easy to see that, because it will never grow, this mutant will have lower fitness than the bet-hedging population. Likewise "always grow" will have lower fitness because it will perish in bad years. Indeed, any other strategy—randomized or otherwise--will have a lower geometric average fitness than the bet-hedging optimum. This is an essential (and experimentally verifiable) difference between mixed strategies and bet-hedges: in bet-hedging, the optimal mix of phenotypes is better than any other over time, while a population playing the mixed ESS will render a single newcomer or mutant totally indifferent to any of the possible strategies it could adopt, whether they be pure or mixed.

3. Not necessarily growth optimal

As we have mentioned a few times before, stable mixed strategies are unlike both bet-hedging and altruistic divisions of labor in that they are not necessarily growth-optimal for the population. This is because the optimal fitness strategies for a population often involve some sort of cooperative behavior. As an illustration, consider the *optimal* outcomes of our two games, the Prisoner's Dilemma and the Hawk-Dove game. In both the Prisoner's Dilemma and the Hawk-dove game, the "optimal" payoff is reached through *unstable pure strategies*: "all players cooperate" and "all players choose dove", respectively. (They are unstable because a mutant defector or hawk can invade the population.) If the optimal average payout is instead reached through an *unstable mixed strategy*—as in the case of colicin production discussed below— instead of a pure strategy, the unstable optimum can

be considered an altruistic division of labor[4]. Since altruistic divisions of labor are, like their pure strategy counterparts, inherently unstable, they must likewise be maintained through inclusive fitness effects like kin selection or group selection. The uninvasible strategy does not necessarily correspond to the strategy that causes the population to grow the fastest.

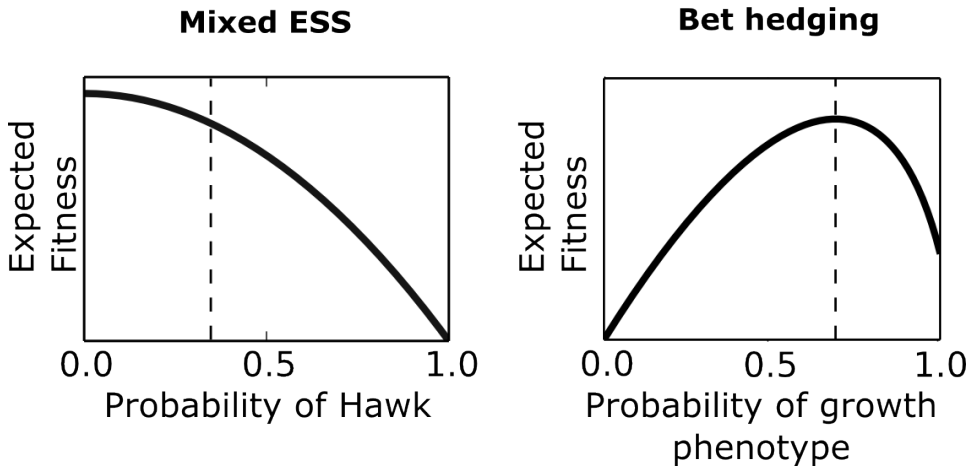

Figure 4. For both our mixed ESS (left) and bet hedging (right) examples, the expected fitness (of an individual or a population) is given as a function of the probability of its adopting one of the phenotypes. The parameters for the models were given previously, and fitness is measured in geometric mean fitness over time. The shape of these curves can very (the mixed ESS curve may have mixed strategy optimum, for example). The curves serve to highlight that the bet hedging optimum is always the fitness optimum, while the mixed ESS may be sub-optimal in fitness.

[4] If the optimal payout is achieved through a *stable* mixed strategy wherein all phenotypes have equal fitness, then the optimal division of labor is not altruistic and coincides with the *selfish* stable point, the mixed ESS.

MIXED ESS VS ALTRUISTIC DIVISIONS OF LABOR

A third distinct survival strategy bears mentioning here: the altruistic division of labor (DoL)[5]. Altruistic DoLs are related to stable mixed strategies in that they rely on interactions within the population rather than on uncertain environments. The primary difference is that evolutionarily stable mixed strategies are *selfish* and *stable*, while altruistic divisions of labor are *cooperative* and *unstable (o*r at least can only be stabilized through kin selection).

Altruistic divisions of labor occur when, through phenotypic heterogeneity, individuals with some phenotype "sacrifice" some of their immediate fitness in order to increase the immediate fitness of other phenotypes. Often called "public goods cooperation" or merely "altruism," examples from the animal world abound—including, most canonically, eusocial insects where a large subset of the population is sterile and serves only to increase the reproduction of a separate subset (*57, 58*). Some altruistic divisions of labor have been described in clonal microbial populations as well, most famously in the production of some colicin toxins, wherein the toxin is only released upon lysis of the cell (*59*). Only a subset of the clonal

[5] Although we use the term "altruistic" here, we recognize that in evolutionary theory, nothing is considered truly altruistic. Every survival strategy must either increase or maintain the associated alleles in the population or be competed away. We use the term rather to distinguish situations wherein apparent altruism creates and incentive for an individual to "cheat" (ie an individual adopts a lower fitness phenotype that increases the fitness of other individuals). This behavior may ultimately increase the genes in the population through kin or group selection.

population releases the colicins, and the colicin-producing phenotype is self sacrificing (*60*). A similar example is the self-sacrificing virulent phenotype of *S. typhimurium,* which causes host inflammation at a cost to its fitness, while the nonvirulent phenotype takes advantage of the inflammation but does not need to produce the virulence factors *(61, 62).* Such divisions of labor are unstable because at any point, a single individual has a clear incentive to "cheat" by never adopting the self-sacrificing phenotype, and such cheaters will then increase in the population, driving it towards the evolutionary stable state and decreasing the fitness of the population as a whole. Because of this constant incentive to cheat, divisions of labor cannot be maintained by classical natural selection, and must be maintained by kin selection, group selection, or some other inclusive fitness effect.

Divisions of labor can be thought of as a phenotypically heterogeneous equivalent to the "optimal" strategies mentioned in the context of the Prisoner's Dilemma and Hawk-Dove game ("always cooperate" and "always play Dove", respectively). In those strategies, the population as a whole is better off than in the evolutionarily stable point, but they are unstable because there is an incentive to "cheat" or "play hawk," respectively. The only difference between these optimal states and the altruistic division of labor, is that in a division of labor, the optimal state is a mixed, rather than a pure strategy.

Divisions of labor are also similar to bet-hedging, in that both kinds of strategies feature asymmetric fitness between phenotypes: some phenotypes have higher

fitness than others. The difference between them is that bet-hedging is *individually selfish* in that 1) the lower-fitness phenotypes exist as a bet on the survival of the individual and 2) the presence of the lower-fitness phenotype lowers, rather than raises, the immediate fitness of the population as a whole. In contrast, the presence of self-sacrificial "cooperators" in altruistic divisions of labor serve to increase the fitness of the other phenotypes in the population.

SUMMARY

Phenotypic heterogeneity can arise from multiple distinct survival strategies. Here we have discussed three: bet-hedging, evolutionarily stable mixed strategies, and altruistic divisions of labor. Though all three strategies involve similar-looking phenotypic heterogeneity, they are all driven by unique evolutionary forces: negative frequency dependence in the case of mixed ESS, uncertain environment in the case of bet-hedging, and kin or group selection in the case of altruistic divisions of labor. Furthermore, each has a unique profile of observable phenotypic fitness characteristics that allows observers to fairly simply classify specific instances of heterogeneity.

Table 1 shows an example of separate fitness profiles that characterize the three strategies. In each case, the relevant evolutionary driver (negative frequency dependent interactions, uncertain environment, and kin or group selection) is sufficient to drive evolution of phenotypic heterogeneity independently of the others. They are unique strategies in response to different environmental

conditions. However, the evolutionary drivers are not mutually exclusive; an environment may contain two or even all three evolutionary drivers. In cases of overlap, the resulting phenotypic heterogeneity can be considered a mixture of multiple strategies.

For example animal bet-hedging models sometimes include frequency dependence that modifies the bet-hedging optimum (*63*). While not yet described in a clonal microbial system, in the context of our "growth vs dormancy" bet-hedging example, overlap with the mixed ESS could for instance occur in the (rather likely) event that limited resources caused negative frequency dependence whereby higher numbers of growing individuals result in a lower payoff for the growing phenotype. One might expect then that the stable mix of strategies would be some combination of the bet-hedging optimum and the mixed ESS. Furthermore, at least one microbial system, *S. typhimurium*, has been described as an altruistic division of labor that is simultaneously bet-hedging as well. (*64*) The virulent phenotype sacrifices its fitness to enhance the fitness of the nonvirulent phenotype, but also has higher survival in the presence of antibiotics, and so may be thought of as a combination of bet-hedging and altruistic division of labor. Such combinations of strategies muddy the experimental water a bit, but the three strategies should still be considered separate survival strategies because they each drive the evolution of phenotypic heterogeneity independently of the others and through different means.

Table 1. Survival strategies implementing phenotypic heterogeneity

	Mixed ESS	**Bet-hedging**	**Altruistic division of labor**
Evolutionary driver	Negative frequency dependence	Uncertain environment	Kin or group selection
Relevant game theory model	Hawk-dove/ Snowdrift/ Chicken	**N/A** (No interactions)	Prisoner's dilemma
Individual fitness depends on phenotypic composition of population	✔	✗	✔
Phenotypes are mutually invasible	✔	✗	✗
At equilibrium, fitness is the same for all phenotypes	✔	✗	✗
"Optimal" mix of phenotypes maximizes population growth	✗	✔	✔
Presence of a low-fitness phenotype increases the fitness of other phenotypes	**N/A** (No "low fitness" phenotype)	✗	✔

Note: Green checks indicate that the observation in the row is predicted by the survival strategy (column). Red Xs indicate the strategy does not predict the observation.

CHAPTER TWO: Phenotypic heterogeneity implements a game theoretic mixed strategy in a clonal yeast population

David Healey[1], Jeff Gore[2]

[1]Department of Biology and [2]Department of Physics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA.

D.H. and J.G. designed the study and performed analysis. D.H. performed the experiments.

Abstract
Genetically identical cells in microbial populations often exhibit a remarkable degree of phenotypic heterogeneity even in homogenous environments. While such heterogeneity is often thought to be a bet-hedging strategy against unpredictable environments, evolutionary game theory also predicts phenotypic heterogeneity as a stable response to frequency-dependent cellular interactions, in which rare strategies are favored over common ones. Here we provide experimental evidence for this game theoretic explanation in the context of the well-studied yeast GAL network. In an environment containing the two sugars glucose and galactose, the GAL network displays stochastic bimodal activation. We show that the two relevant phenotypes (GAL-ON and GAL-OFF) can each invade the opposite strategy when rare, indicating frequency dependence between the two. Consistent with the Nash equilibrium of an evolutionary "hawk-dove" game, the stable mix of pure strategists does not necessarily maximize the growth of the overall population, as it would in a bet-hedging or labor-dividing population. Yeast with the wild type GAL network can invade populations of both pure strategists while remaining uninvasible by either. We show that, beginning with populations of pure strategists, a mixed sugar condition drives both the evolution of a clonal population of mixed strategists similar to the wild type GAL network, and the evolution of a stable coexistence of opposite pure strategists. Taken together, our results provide experimental evidence that frequency-dependent interactions between identical cells can underlie the phenotypic heterogeneity found in clonal microbial populations.

Introduction

Stochastic gene expression is ubiquitous in biological systems (*6, 9, 24, 61*). While some noise in gene expression is inevitable, phenotypic heterogeneity is an evolvable trait whose quantitative parameters can be tuned by the architecture and properties of the underlying gene network (*20-23*). This raises the question of what adaptive advantage might be conferred to cells that implement stochastic decision-making (*54, 65*). Microbial phenotypic heterogeneity is most often thought to be a response to environmental uncertainty; populations that "hedge their bets" by stochastically adopting a range of phenotypes can gain a fitness advantage if the environment shifts unexpectedly (*22, 23, 41-43, 65-68*). For example, bacteria may, at some frequency, stochastically adopt a dormant or slow-growing "persister" state, which has reduced fitness when times are good, but which is more likely to survive in the event of catastrophic environmental stress (*36, 41*).

Other evolutionary drivers of heterogeneity, such as altruistic divisions of labor and evolutionary "hawk-dove" (or snowdrift) games, are distinct from bet-hedging in that they result from interactions between individuals within populations, and can manifest even in deterministic environments (*28, 56, 69, 70*). Altruistic divisions of labor occur when one phenotype sacrifices its fitness to increase the fitness of the remaining population. Canonical examples include the self-sacrificing virulent phenotype of *S. typhimurium*(*61, 62*) and colicin production, in which toxin is only released upon cell lysis(*60, 61*). Because there exists the potential for an individual to gain a fitness advantage by never adopting the low-fitness phenotype, such

altruistic divisions of labor must generally be maintained by inclusive fitness effects such as kin or group selection(*67, 71*). In contrast, phenotypes in hawk-dove games are mutually invasible. Such games tend toward an equilibrium wherein all phenotypes have equal fitness: either a stable coexistence of pure strategies or a single evolutionarily stable mixed strategy (mixed ESS). Identifying which of these three evolutionary drivers (hawk-dove games, uncertain environment, altruistic division of labor) is at work in a given phenotypically heterogeneous population is complicated by the possibility that multiple of these phenomena can coexist in a given system (*70, 72*). While bet-hedging and altruistic division of labor have been observed in microbial populations (*41-43, 62, 72*), the relevance of hawk-dove games to microbial phenotypic heterogeneity remains largely unexplored.

Evolutionary game theory concerns itself with situations in which the fitness of a phenotype is a function not only of the individual's own phenotype, but of the phenotypes adopted by other individuals. In the hawk-dove game of animal conflict, neither pure strategy (ie "play hawk" or "play dove") is evolutionarily stable, since populations of each can be invaded by a minority population of the other. Because of this mutual invasibility, if allowed to evolve, the system reaches a stable equilibrium that contains a mix of phenotypes. In microbial populations, such negative frequency dependent interactions have been shown to stabilize the coexistence of different genes (*73-75*).. However, a genetically identical population can theoretically achieve the same stable mix of phenotypes, provided each individual randomizes between strategies with the appropriate probabilities. Such stochastic

choices between strategies are called *mixed strategies,* and the specific probabilistic strategy that implements a stable equilibrium is the *evolutionarily stable mixed strategy* (or mixed ESS).

Mixed ESS in hawk-dove games have several experimentally observable characteristics that distinguish them from bet-hedging strategies and altruistic divisions of labor (Chapter 1 Table 1). The primary defining characteristic of the hawk-dove game is that the pure strategies, or phenotypes, are mutually invasible: each pure strategy can invade the other when rare. Secondly, if the mixed strategy is evolutionarily stable, all individuals in the population—and potential invaders implementing any other strategy, pure or mixed—will receive an equal payoff (*52*). Thirdly, while other strategies of phenotypic variation maximize some measure of population growth (see discussion), evolutionarily stable mixed strategies are not necessarily growth optimal for a population (*76*).

Results

To study stable mixed strategies in the laboratory, we investigated the decision of the budding yeast *S. cerevisiae* regarding which carbon source to consume. Yeast prefers the sugar glucose, but when glucose is limited yeast can consume other carbon sources (*77*). The well-studied yeast GAL network contains the suite of genes required to metabolize the sugar galactose. The GAL network activates as a phenotypic switch: each cell is either GAL-ON or GAL-OFF (*16, 78*). The network is under catabolite repression by glucose (*77*); however, yeast can still activate the

GAL genes in the presence of modest glucose concentrations provided there is galactose in the media as well. The presence of a negative feedback loop involving *GAL80* serves to reduce the heterogeneity of GAL expression under activation and deactivation (*79*). However, due to a positive feedback loop from *GAL3* (Figure 1), in a wide range of glucose and galactose environments the GAL network is neither uniformly activated or deactivated across the population, but is expressed bimodally (*78, 80*) (Figure 3). In mixed sugar conditions, a tradeoff exists between activation and repression of the GAL network: activation of the GAL network in the presence of glucose may provide some benefits to the cell in consuming galactose, but expression of the GAL genes also imposes a significant metabolic cost (*80, 81*) (Figure 6). Similar tradeoffs in catabolite-repressed networks have been characterized previously (*23, 44, 82, 83*).

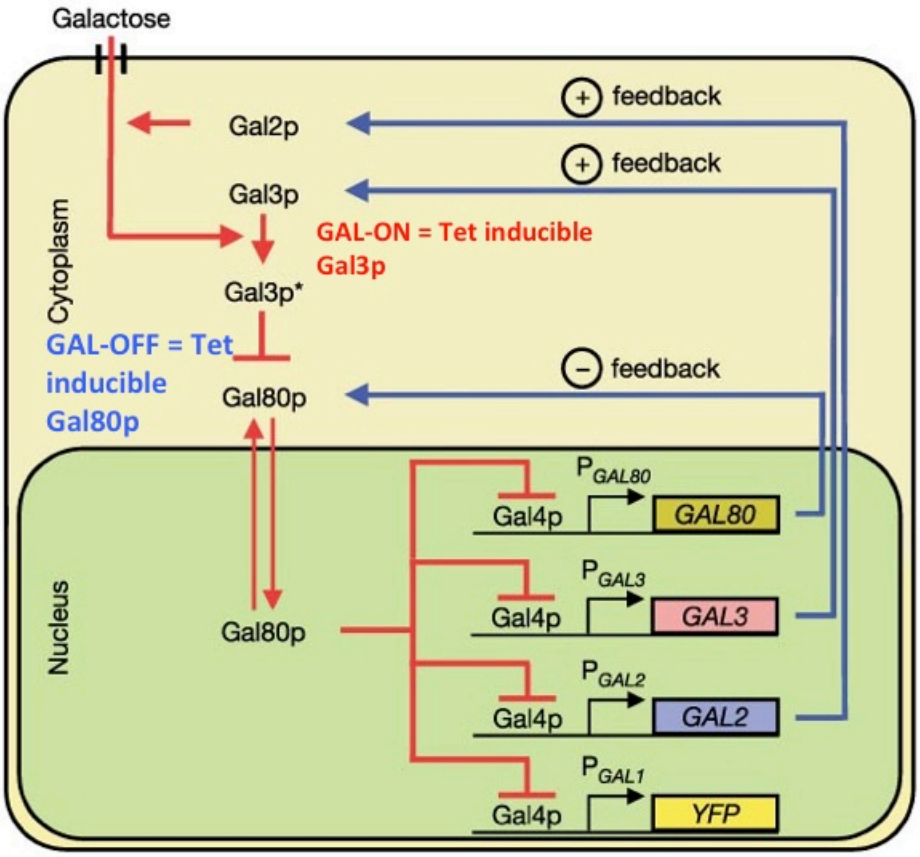

Figure 1. Feedback loops in the GAL network modulate phenotypic heterogeneity
The GAL network is composed of three feedback loops through GAL4p, the universal activator of the GAL network. The sole negative feedback loop is responsible for inhibiting phenotypic heterogeneity of the ON and OFF states (*79*), while the positive feedback loops through *GAL2* and *GAL3* create a bistable "switch," which ultimately increases phenotypic heterogeneity significantly. The combination of the three feedback loops produces the characteristic bimodal behavior of the GAL network in mixed glucose and galactose. Figure modified from Acar et al. 2005.

GAL bimodality in mixed glucose and galactose suggests a similarity to the following hawk-dove-like foraging game. In this game, an isogenic population is confronted with a phenotypic decision to "specialize" in consuming one or the other of two limited food sources, A and B (Figure 2a-b). The more individuals who adopt the

pure strategy “specialize in A,” the more quickly A will be consumed, reducing the payout to individuals who chose that strategy. Hence, if all individuals choose “specialize in A,” an incentive may exist for an individual to choose instead “specialize in B,” and vice versa. The resulting equilibrium consists of a stable mix of the two pure strategies; therefore, an isogenic population that can adopt that stable mix via phenotypic heterogeneity would be uninvasible.

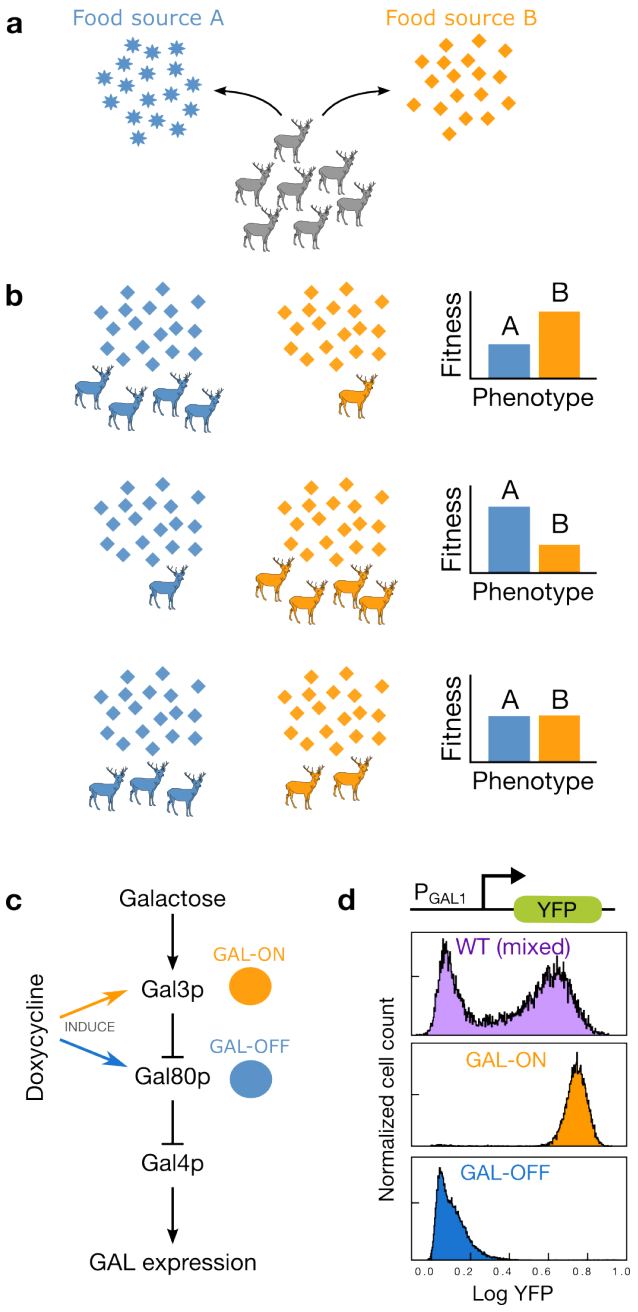

Figure 2. A simple foraging game with multiple food sources can favor phenotypic heterogeneity

a, A simple foraging game with a mixed Nash equilibrium. Each member of a group of foragers is confronted with a binary decision about whether to specialize in consuming food source A or B. We assume that individuals choose simultaneously and without knowledge of the actions of others. Resource limitation makes it a game; each individual's payoff is a function of the actions of other individuals. **b,** If all other members of the population

adopt some pure strategy (e.g. "specialize in food A"), an individual opting for the opposite pure strategy (e.g. "specialize in food B") gains a fitness advantage (top and middle panels). The Nash equilibrium of the simple foraging game is reached when the population divides between the two sources such that both phenotypes receive the same fitness and there is no fitness incentive for any single individual to change strategies (lower panel). In such a game, if each individual adopts the mixed strategy that stochastically chooses between pure strategies with the equilibrium probabilities, then that mixed strategy is evolutionarily stable. Though this game is not necessarily representative of real life foraging scenarios, it serves to illustrate why we might expect environments with multiple food sources to favor the evolution of mixed strategies. **c,** Gene expression in the yeast GAL network is regulated in part by *GAL4, GAL80,* and *GAL3* (full network not shown). A GAL-OFF pure strategist is engineered by inducing the expression of *GAL80*, whose protein product inhibits GAL expression. Likewise a GAL-ON pure strategist can be engineered by inducing expression of *GAL3*, which inhibits GAL80 in the presence of galactose. **d,** In a mixed sugar environment (0.03% glucose, 0.05% galactose), "GAL-ON" and "GAL-OFF" pure strategists remain unimodally activated and inactivated, respectively, while the wild type GAL network exhibits bimodal gene expression. Cultures in Figure 1d were initially grown overnight in 0.01% (w/v) glucose and 1ug/mL doxycycline to saturation, then diluted to an OD of 0.002 and grown 8 hours in mixed glucose and galactose before measuring GAL activation via flow cytometry.

Given the bimodal expression of the yeast GAL network in some conditions, we sought to probe experimentally whether this phenotypic heterogeneity might be the implementation of an evolutionarily stable mixed strategy in response to a foraging game. Since mutual invasibility of phenotypes is the defining characteristic of a hawk-dove game, we began by competing mutant "pure strategists" at many initial population frequencies. As a GAL-OFF pure strategist, we used a yeast strain whose native *GAL80* (a repressor of the GAL network, Figure 2c) was replaced with a mutant version containing a tet-inducible promoter (*78*). As a GAL-ON pure strategist, we used a mutant whose *GAL3* (a repressor of *GAL80*, Figure 2c) was similarly tet-inducible. We confirmed that, in the range of glucose and galactose

concentrations that induce bimodality in the wild type yeast, our doxycycline-induced GAL-OFF and GAL-ON pure strategists are unimodally inactivated and activated, respectively, for GAL gene expression (Figure 2d, Figure 3b-c).

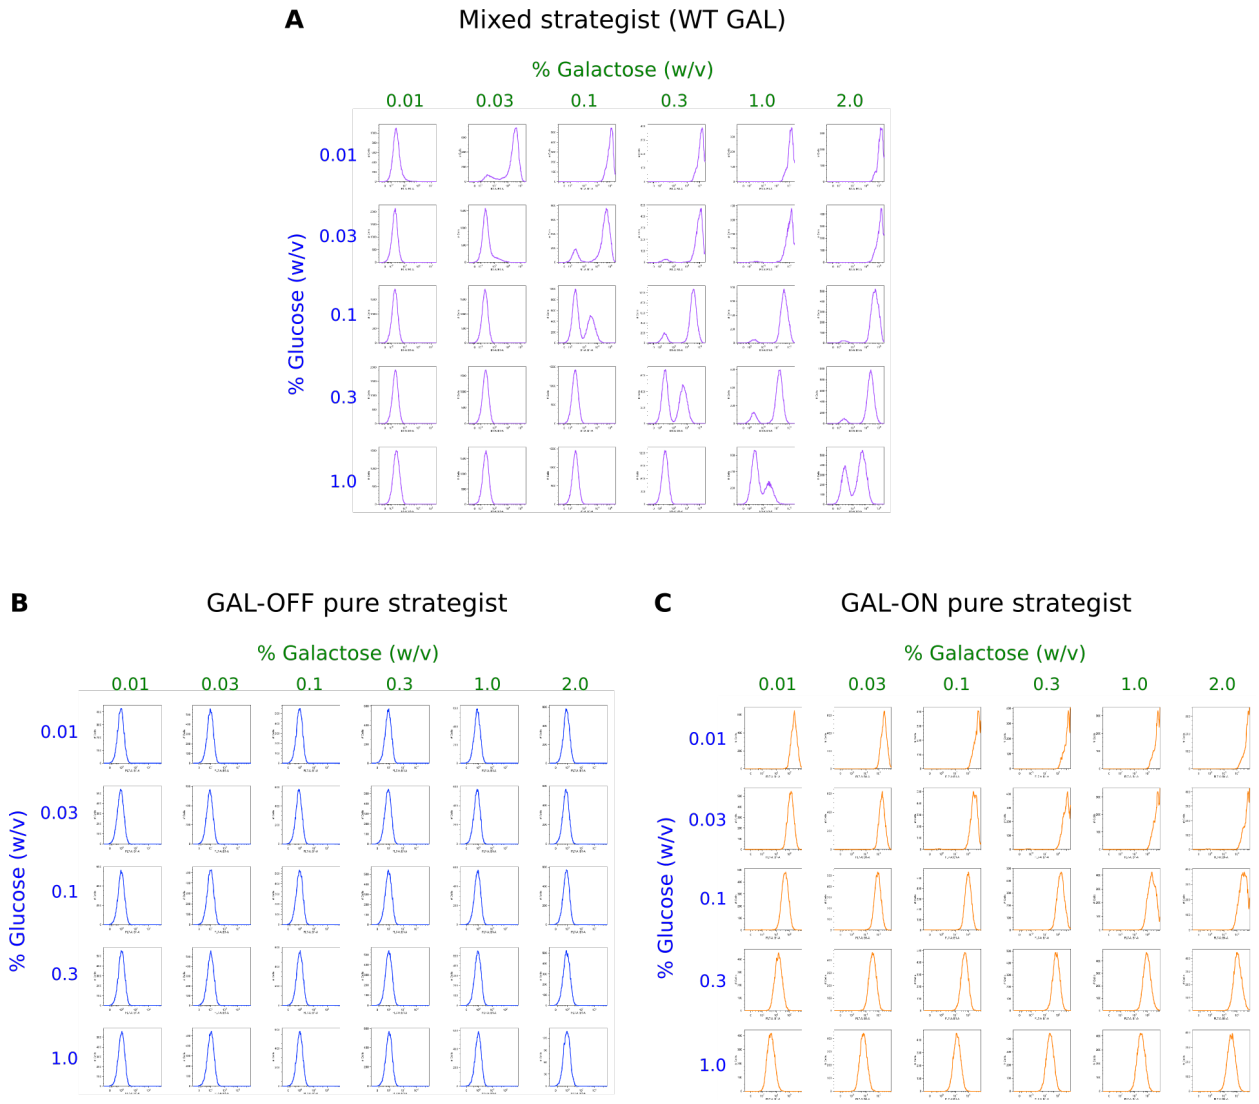

Figure 3. Response of the GAL network to mixed glucose and galactose. Wild type mixed strategist (**a**), GAL-OFF pure strategist (**b**), and GAL-ON pure strategist (**c**) strains were induced in 1 µg/mL doxycycline and 0.01% glucose for 8 generations, then diluted and incubated in 1 µg /mL doxycycline and various concentrations of glucose and galactose as shown. GAL activation states were measured after 8 hours via flow cytometry, and normalized histograms plotted against log YFP with a smoothing of 3. Broadly speaking, the wild type exhibits bimodal GAL expression over a wide

range of roughly equal ratios of glucose and galactose concentration. At higher ratios of glucose to galactose, it remains GAL-OFF, while at lower ratios it is unimodally GAL-ON. In all of the conditions wherein the wild type adopts a mixed strategy, the GAL-OFF pure strategist is inactivated, and the GAL-ON pure strategist is activated. However, in the cases of both the wild type and the GAL-ON pure strategist, increasing ratios of glucose to galactose result in lower activation levels of the GAL network.

To test for negative frequency dependence between the pure strategists, we mixed six biological replicate pairs of RFP-labeled GAL-ON and CFP-labeled GAL-OFF strains at a total of 60 different initial frequencies, and incubated them for 20 hours in a mixed glucose and galactose environment. To calculate precise fitness values for both strains, we measured population frequencies before and after incubation via flow cytometry (Figure 4). We found that small populations of each pure strategist were indeed able to invade majority populations of the other (Figure 5b). Our experimental yeast populations therefore display mutual invasibility between the two pure strategists. Furthermore, there was a unique stable equilibrium frequency of GAL-ON cells that resulted in the same fitness for both pure strategies. Importantly, we find that the frequency of GAL-ON cells that is evolutionary stable is not the frequency that maximizes population growth. Populations with much higher fractions of GAL-ON cells than the equilibrium population grow to saturating density more quickly than the evolutionarily stable population (Figure 5c).

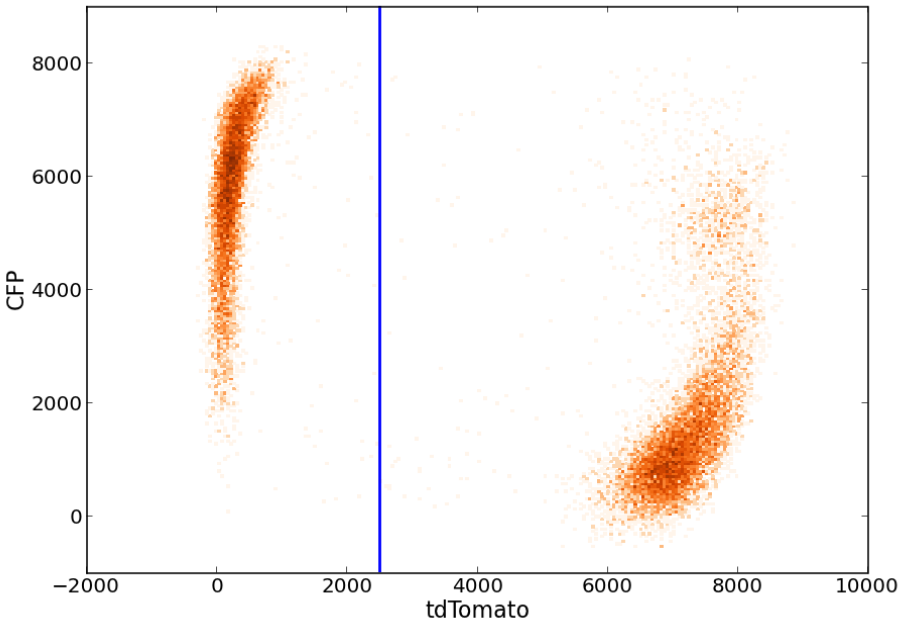

Figure 4. Absolute and relative fitness determinations

a, In order to calculate relative fitness, we first labeled each strain with either RFP (tdTomato) or CFP driven on a TEF1 constitutive promoter We determined population frequencies before and after competition (f_i and f_f, respectively)$f_i$andf_f) for each strain via flow cytometry using a Miltenyi MACSquant flow cytometer (20,000+ cells per well). Separation of the two strains was very clean. Absolute fitness for each strain was calculated as the number of doublings:

$$W_{abs} = \log_2\left[\frac{OD_f * f_f}{OD_i * f_i}\right]$$

Relative Fitness for each strain is calculated as follows:

$$W_{rel} = \ln\left[\frac{OD_f * f_f}{OD_i * f_i}\right] / \ln\left[\frac{OD_f * (1-f_f)}{OD_i * (1-f_i)}\right]$$

where f_i and f_f are the initial and final fractions of the strain in the population, and OD_iOD$_i$ and $OD_f$are the initial and final population densities as measured by absorbance at 600 nm in a microplate spectrophotometer.

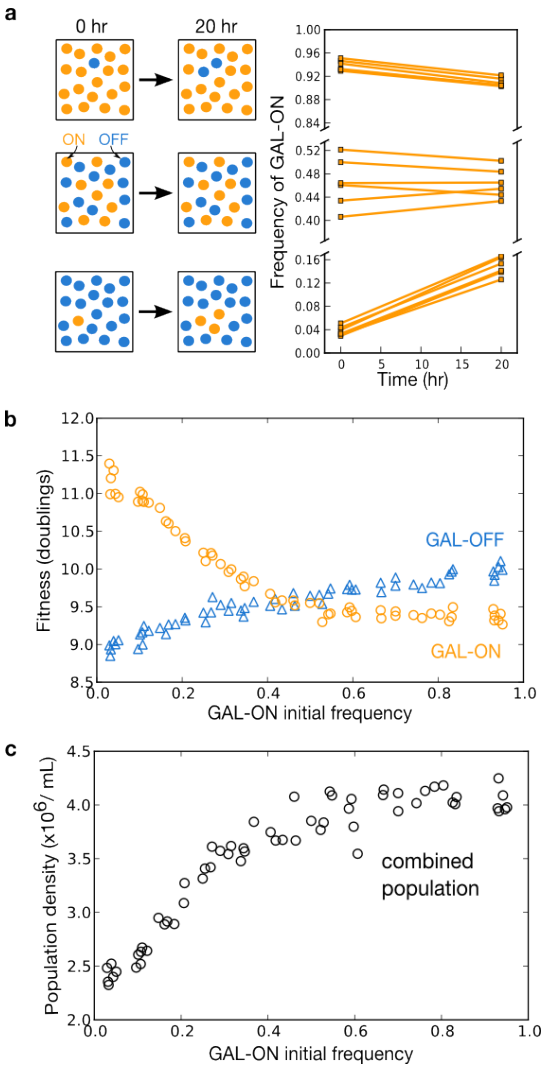

Figure 5. Characterization of the game played between GAL-OFF and GAL-ON pure strategists

a, Pure strategists are mutually invasible. Population frequency of the *GAL3*-induced GAL-ON pure strategist (orange circles) relative to the *GAL80*-induced GAL-OFF pure strategist (blue triangles) is plotted at the beginning and end of a 20-hour competition. Six independent cultures of each pure strategist were mixed at high (top panel), intermediate (middle panel), and low (bottom panel) initial frequency of the GAL-ON strain. Each pure strategist invades the other when rare. **b,** Game payoffs (in number of doublings) for both pure strategists are plotted for 60 initial starting frequencies of the GAL-ON strain. The crossing point corresponds to the Nash equilibrium of the pure strategists for the experimental foraging game. **c,** The evolutionarily stable equilibrium is not necessarily growth optimal. Population densities of the mixed populations are shown at 16 hours, before all cultures have reached saturation. Mixed cultures with high initial

frequency of the GAL-ON strain grew faster than cultures near the evolutionarily stable mix.

A more in-depth investigation of the dynamics between the pure strategists indicates that the negative frequency dependence is related to the depletion of glucose in the media. Both pure strategists adopt a diauxic growth model; they consume primarily glucose until it is depleted (Figure 6e). Indeed, recent evidence argues that the advantage provided to GAL-ON cells may be primarily due to their ability to consume galactose quickly when the glucose is exhausted (*80*). The GAL-ON pure strategist suffers a fitness disadvantage while glucose is still relatively abundant, but outcompetes the GAL-OFF pure strategist when the glucose becomes low and galactose remains. The GAL-OFF and GAL-ON strategies can therefore be thought of as "specialists" in glucose and galactose, respectively. The hawk-dove game arises because the galactose "payoff" goes to the GAL-ON cells, but the more cells that activate their GAL networks in a population, the slower the glucose gets depleted, and the higher the resulting payoff to glucose specialists (Figure 6).

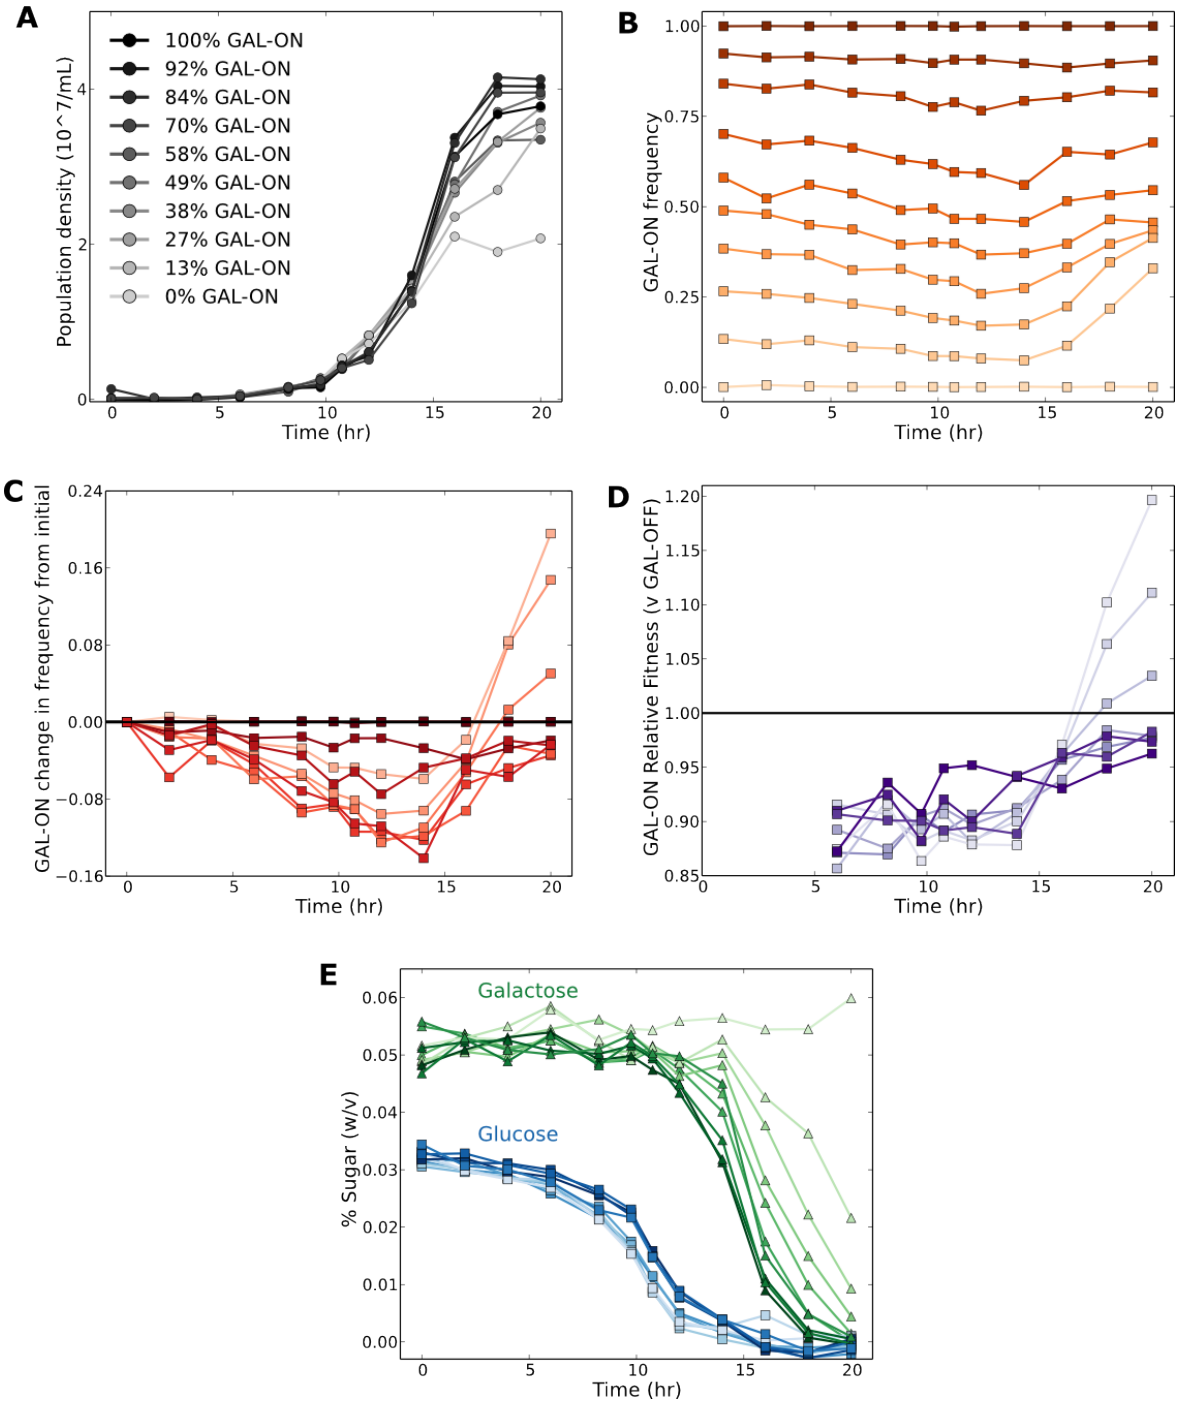

Figure 6. Dynamics of growth, fitness, and sugar consumption for the 20-hr competition between GAL-ON and GAL-OFF pure strategists. Competitions took place in the presence of 1 μg/mL doxycycline, 0.03% (w/v) glucose and 0.05% (w/v) galactose. In all plots, color gradients indicate the increasing initial frequencies of GAL-ON as described in panel A legend.

Game theory predicts that varying the payoff structure of a hawk-dove game correspondingly alters the Nash equilibrium fractions (Box 1). In the context of the simple foraging game, this simply means that if food source A increases, then the stable equilibrium should shift towards a larger fraction of the population specializing in food source A. To test for this phenomenon in the GAL network, we replicated the initial competition of our two pure strategists in eight different concentrations of glucose and galactose. More galactose yields a higher equilibrium fraction of GAL-ON cells, while more glucose yields a lower equilibrium fraction of GAL-ON cells (Figure 7). The stable equilibrium between our pure strategists therefore shifts as predicted by a negative frequency dependent game. However, while this well-behaved shifting of the equilibrium is robust within a range of relatively low sugars, the pattern breaks down in environments of high total sugar concentrations (>0.1%), where carbon may not be limiting in the same way.

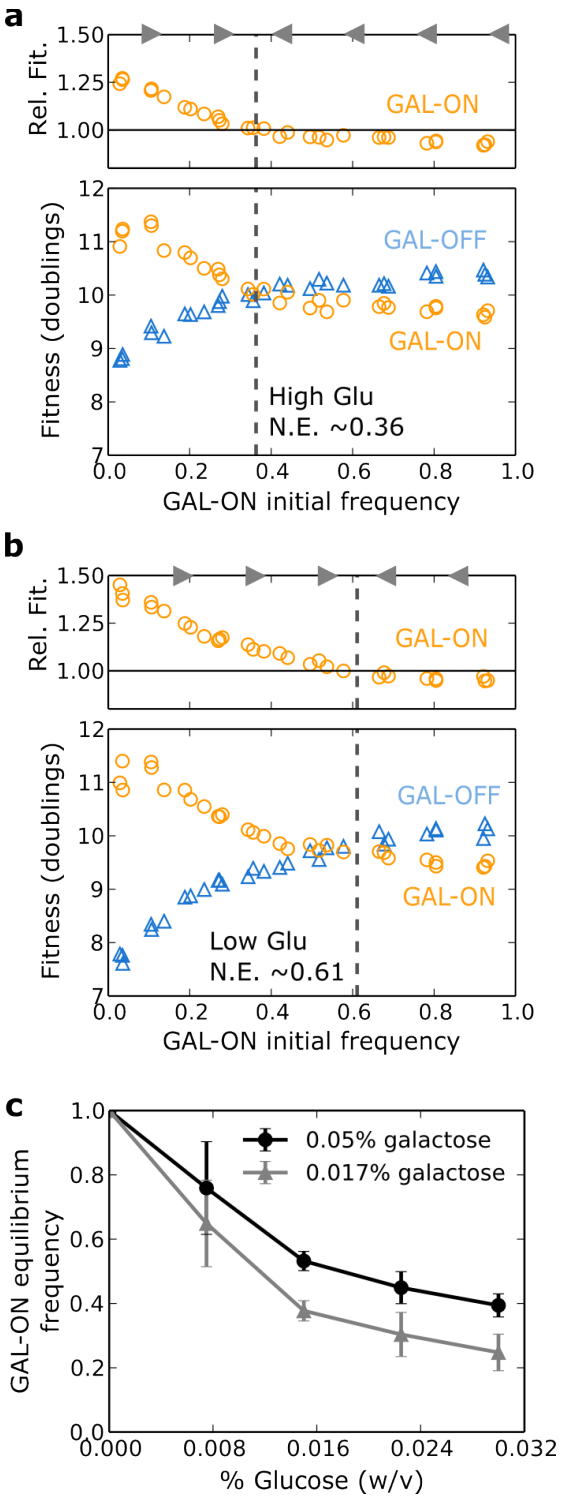

Figure 7. Altering sugar concentration adjusts game payoffs and equilibrium fractions accordingly.
a-b, Relative fitness of the GAL-ON pure strategist, and absolute fitness (in number of doublings) of both pure strategists is shown for 30 different populations at varying initial frequency of GAL-ON. Data is shown for .05%

galactose and two conditions: high glucose (.03%, **a**), and low glucose (.017%, **b**). The payoff for the GAL-ON pure specialists remains roughly the same between the two conditions, while the GAL-OFF pure strategists receive a higher payoff in higher glucose. Lower glucose results in a higher equilibrium frequency of GAL-ON cells, as expected in a hawk-dove like game. **c,** Equilibrium GAL-ON pure strategist frequencies as a function of increasing glucose concentrations. Data is shown for high (0.05%, circles) and low (0.017%, triangles) galactose. All Nash equilibria were calculated by polynomial spline fitting of relative fitness curves (error bars are 95% confidence intervals; n = 3.)

Because the Nash equilibrium mix of pure strategies is a function of sugar concentrations, we next tested whether the wild type mixed cells naturally alter the frequency of mixing based on the concentrations of glucose and galactose. Just as the Nash equilibrium of the pure strategists shifts with varying sugar concentrations, we observed that the mixing frequency of the wild type yeast also shifts: in a higher concentration of galactose, yeast adopt a higher GAL-ON frequency. This type of responsiveness is one of the hallmark predictions of the mixed strategy model. Yeast is able to sense even small differences in the ratio of glucose to galactose (*84*) and adopts a pure OFF, pure ON, or appropriate mixed strategy accordingly (Figure 8).

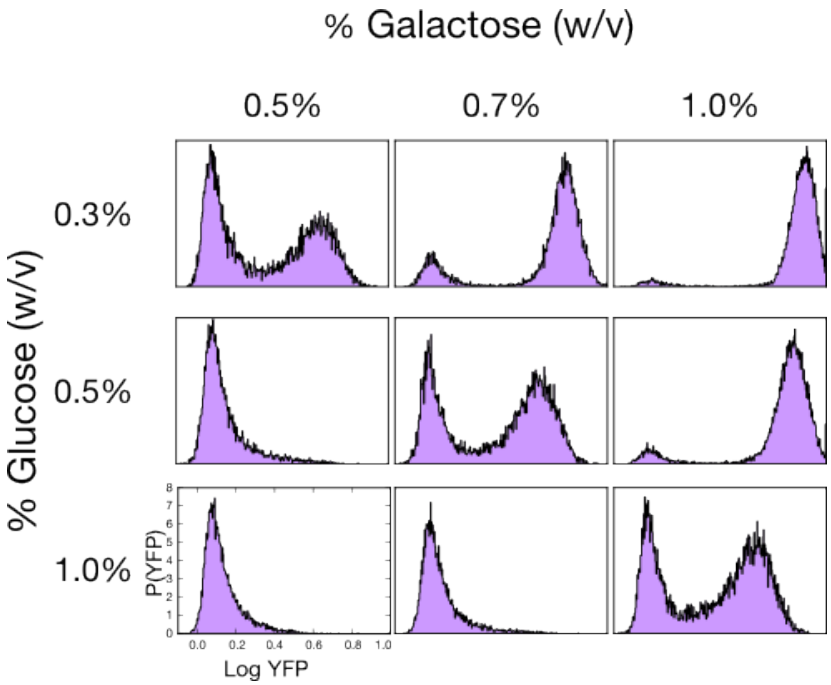

Figure 8. Mixed strategist (wild type GAL network) senses payoff ratio and alters strategy accordingly. The mixing frequency of the wild type mixed strategist is highly responsive to sugar concentrations. GAL network activation level is shown for nine different mixtures of glucose and galactose. From a bimodal expression state, more galactose in the media results in a higher frequency of cells with GAL activation, while more glucose in the media has the opposite effect. This trend is essential to the implementation of an evolutionary stable mixed strategy.

Another prediction of the hawk-dove game is that a strain adopting a mixed ESS cannot be invaded by either pure strategist. However, as the population frequency of the mixed strategist approaches one, it becomes only neutrally uninvasible. In other words, in the limit of a population consisting entirely of mixed strategists, any single invading cell adopting any strategy (pure or mixed) will receive a payoff equal to the mixed strategist. By competing the pure strategist strains (GAL-ON/OFF) with a strain containing the wild type GAL network (mixed strategist), we determined that the mixed strategist is indeed uninvasible by either pure strategist. Additionally, a competition between pure GAL-OFF and the mixed strategist displays the neutral uninvasibility predicted from the game theoretic model (Figure 9). The wild type mixed strategy can spread in a population of GAL-OFF cells, but as the wild type strategy increases in frequency, its advantage disappears. Moreover, the wild type mixed strategist cells are uninvasible by the GAL-ON pure strategist at all frequencies (Figure 9b), though the interaction does not display strong frequency dependence. This lack of strong frequency dependence between this pair suggests that the dynamics of yeast in mixed sugar environments have some subtle deviations from a simple foraging game.

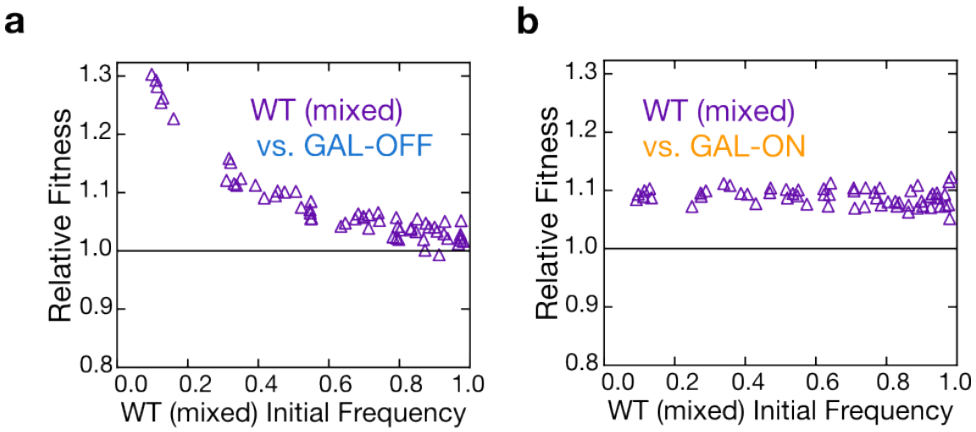

Figure 9. Wild type mixed strategist invades both pure strategists and is uninvasible by either.
Relative fitness of the wild type mixed strategist over the GAL-OFF pure strategist (**a**) and GAL-ON pure strategist (**b**) is shown. Low frequencies of the mixed strategist invade strongly in populations dominated by either pure strategist. As expected of an evolutionarily stable mixed strategy, the relative fitness of the mixed strategist to the GAL-OFF pure strategist approaches one in populations dominated by the mixed strategist. However, the mixed strategist does not display frequency dependence against the GAL-ON pure strategist.

To more directly investigate whether a mixed-sugar environment drives the evolution of phenotypic heterogeneity, we evolved eight replicates of each pure strategist strain over 250 generations in the presence of doxycycline and three resource conditions: pure glucose, pure galactose, and a mixture of glucose and galactose. Cultures were diluted daily 1000x into fresh media, and GAL activation levels over time were determined via flow cytometry (Figure 10). Starting from the GAL-OFF glucose specialist strain (Figure 10a), all eight populations in mixed sugars evolved a mix of GAL-OFF and GAL-ON. In contrast, a pure galactose condition drove the rapid evolution of GAL-ON pure strategist strains in six of the eight

populations while the remaining two replicates were driven to extinction. The eight GAL-OFF populations in pure glucose remained GAL-OFF throughout.

Starting from the GAL-ON galactose specialist strain (Figure 10b), all eight populations in mixed glucose and galactose similarly evolved a phenotypic mix of GAL-OFF and GAL-ON, while the strain in pure galactose remained essentially GAL-ON throughout. (In pure glucose, the GAL-ON strain adopted a low-level wide distribution of GAL activation, which remained unchanged throughout the experiment.)

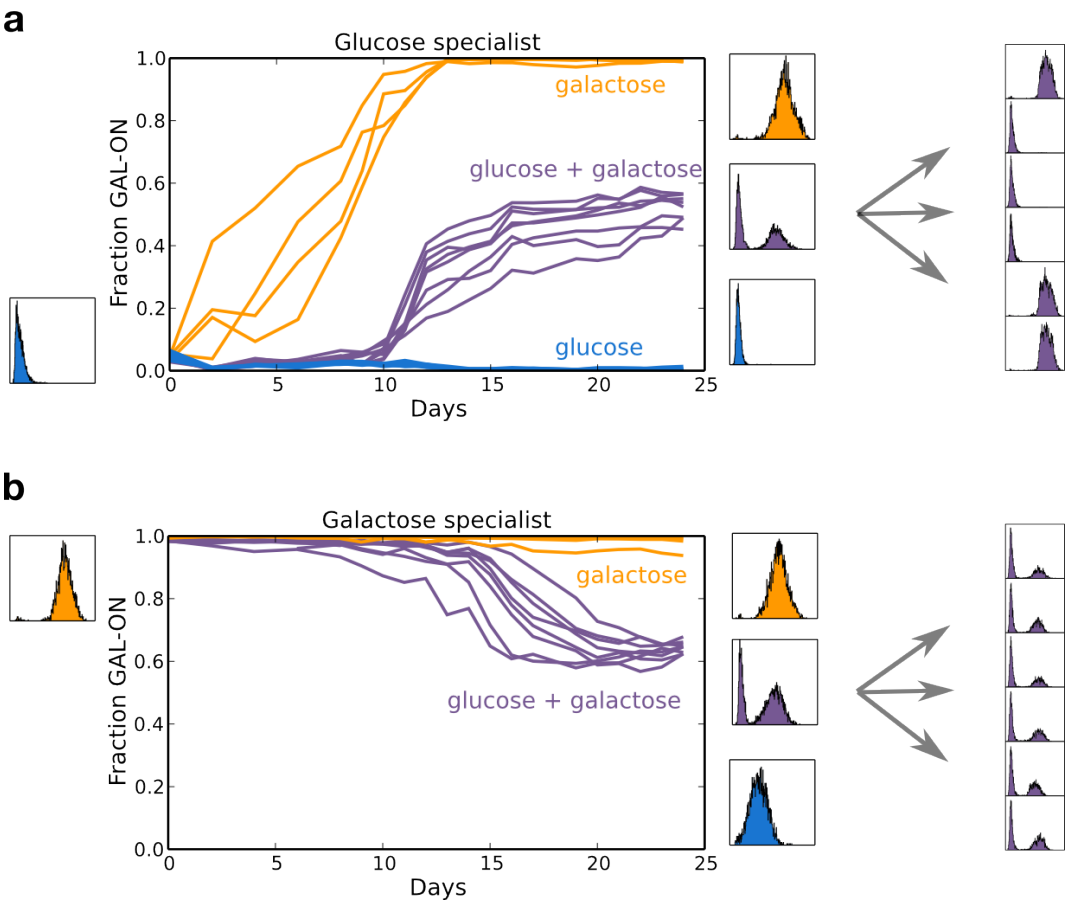

Figure 10. Frequency dependence from a mixed resource environment drives the evolution of both phenotypic and genetic heterogeneity. Eight replicates of each of the two specialist strains (GAL-OFF, **a**, and GAL-ON, **b**) were incubated in the presence of doxycycline and three separate sugar conditions: 0.1% glucose (blue), 0.1% galactose (orange), and a mixture of 0.03% glucose and 0.05% galactose (purple.) Cultures were diluted 1000x daily into fresh media after reaching saturation. To determine the composition of the final population, cultures were plated on agar and individual colonies were grown separately in appropriate sugars. **a)** Starting from a glucose specialist strain and in the presence of galactose, mutant pure strategist GAL-ON strain arose spontaneously. In pure galactose, the strain eventually took over the population (orange), while in the mixed resource condition, they evolved towards a stable equilibrium with the GAL-OFF strain (right). **b)** Starting from a galactose specialist strain in the presence of mixed sugars, the population similarly evolved to a stable mix of GAL-ON and GAL-OFF, but colony purification revealed that it had evolved to a clonal population of mixed strategists rather than a coexistence of pure strategists. Importantly, while all three drivers of heterogeneity (uncertain environment, hawk-dove game, and king/group selection) promote evolution of clonal phenotypic heterogeneity, the hawk-dove game is the only one that also promotes the stable coexistence of pure strategists.

To determine whether the heterogeneous populations that evolved in the mixed resource conditions represented mixed strategists or the coexistence of pure strategists, we isolated individuals through colony purification on agar plates, and incubated them separately in the presence of doxycycline and in mixed glucose and galactose. We found that the GAL-OFF history population had evolved into a stable coexistence of GAL-ON and GAL-OFF pure strategists, while the GAL-ON population had evolved into a clonal population of mixed strategists similar to the wild type strain. This result strongly indicates that frequency dependent interactions drove the evolution of phenotypic heterogeneity in the mixed sugar conditions, because while all three evolutionary drivers of heterogeneity (uncertain environments, negative frequency dependent games, and public goods cooperation) drive the

evolution of mixed strategies in clonal populations, only negative frequency dependence also drives the stable coexistence of pure strategies. This experiment directly demonstrates that phenotypic heterogeneity in clonal populations can arise from frequency-dependent foraging interactions.

Discussion

When observing phenotypic heterogeneity in microbial populations, it is important to consider the underlying evolutionary reasons for heterogeneity and distinguish between the different explanations where possible. While it is very difficult (if not impossible) to prove claims about historical reasons for the evolution of stochasticity in specific systems like the yeast GAL network, different survival strategies do make unique and experimentally verifiable predictions about the fitness dynamics between the associated phenotypes (see Chapter 1 Table 1 and deJong and Kupers' review(*28*)). In this work we have demonstrated a simple way of probing whether observed phenotypic heterogeneity might be implementing an evolutionarily stable mixed strategy. By isolating the pure strategies and probing them for mutual invasibility, we have determined that a hawk-dove style foraging game is being played between the GAL-ON and GAL-OFF strategy in the presence of mixed glucose and galactose. This frequency dependent mutual invasibility distinguishes a mixed ESS from bet-hedging (which is not frequency dependent) and altruistic division of labor (in which the altruistic phenotype is always less fit). We have also verified the theoretical prediction that the evolutionarily stable mixed strategy is not necessarily optimal for growth, and confirmed that a strain

implementing a mixed strategy invades populations of both pure strategists, and is uninvasible by either. Additionally, we have shown directly that mixed resource environments can drive the evolution of phenotypic heterogeneity by evolving a mixed strategist strain from a pure strategist over 250 generations in mixed sugars, a condition which separately drove the evolution of stable coexistence of opposite pure strategists.

In the mixed sugar conditions we have shown, the wild type mixing frequency is roughly the same as the Nash equilibrium between the mutant pure strategists (compare Figure 2 d and Figure 4b). However, we do not expect the quantitative agreement to be general, since budding yeast did not evolve its mixing frequency in laboratory cultures of mixed glucose and galactose. There are also slight differences between the mutants and the wild-type phenotypes. For example, the GAL-repressed subpopulation of the mixed strategist adopts a diauxic growth phenotype: it activates its GAL network upon glucose depletion (Figure 6), but because of the doxycycline induction of *GAL80*, the GAL-OFF pure strategist does not transition to GAL-ON within the time frame of the competition (Figure 11). Consequently, in a mixed sugar scenario, the GAL-suppressed fraction of the mixed strategist is likely to be more fit than the GAL-OFF pure strategist. Additionally, the GAL-ON pure strategist's induction activates the GAL network to a greater degree than the induction in the wild type (Figure 2d), resulting in slightly different costs for expressing the GAL network. Also, while the majority of the wild type yeast's stochastic decision to be GAL-OFF or GAL-ON is determined early, there is a small

amount of stochastic switching between the states (Figure 12), which does not occur in the mutant pure strategists.

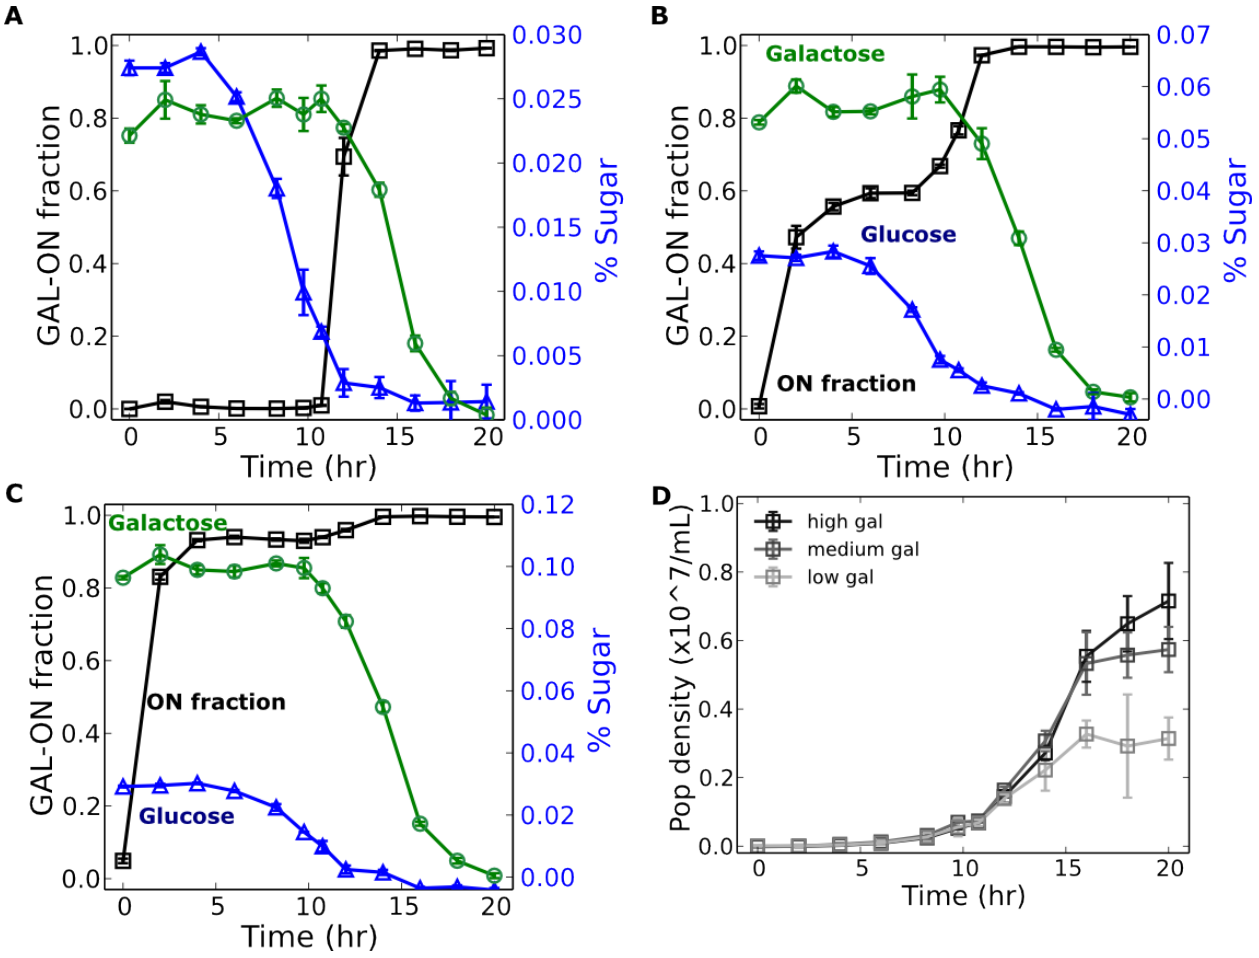

Figure 11. Dynamics of wild type mixed strategist yeast in mixed glucose and galactose. Yeast containing a wild-type GAL network were incubated in 0.03% (w/v) glucose and varying levels of galactose. **a,** In low (0.025%) galactose, the wild type remains unimodally off (black squares) until the glucose (blue triangles) is mostly consumed, then rapidly turns on its GAL network and consumes galactose (green circles). **b,** At intermediate (.05%) galactose, the wild type exhibits bimodal expression of the GAL network (see Figure 4c) until the glucose is depleted, then rapidly switches to all ON. **c,** At high galactose, the GAL network is rapidly activated in all cells. The 0-10 hr data in panels **b** and **c** suggest that even when the GAL network is activated in either half **(b)** or all **(c)** of the population, the yeast still consumes primarily glucose while it is still available.

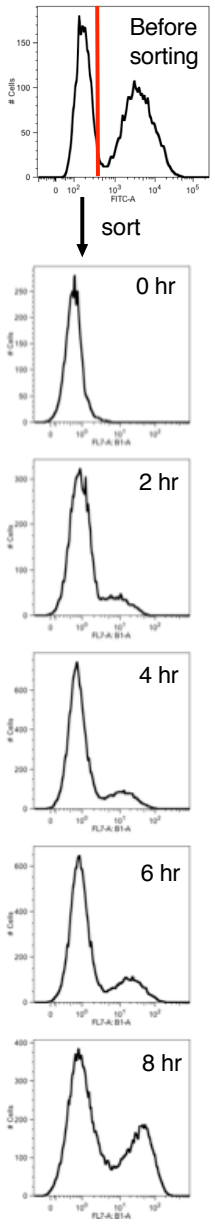

Figure 12. Stochastic switching of the wild type GAL network in mixed glucose and galactose. To measure the stochastic switching rate between the GAL-ON and GAL-OFF states, we grew the mixed strategist RFP-labeled strain for 8 hours in 0.03% glucose and 0.05% galactose from a density of about 2×10^4 cells/mL. After 8 hours, cells were sorted by FACS into GAL-OFF and GAL-ON fractions, and diluted back into the same concentration of mixed sugars. Activation states were monitored by flow cytometry. We observe a low rate of stochastic switching from GAL-OFF to GAL-ON. The reverse switching rate was much lower or in some cases not observable. We replicated the experiment at 0.1% glucose and 0.2% galactose and observed no stochastic switching between the two states. Taken with Supplementary Fig. S3b, this data suggests that the bulk of distribution of GAL activation occurs soon after introduction into the mixed sugar environment.

Negative frequency dependent interactions are often invoked as reasons for stable coexistence, and evolutionary stable mixed strategies (in the context of hawk-dove games) are central to evolutionary game theory. Yet this broad class of interactions has received almost no attention as an evolutionary reason for phenotypic heterogeneity in clonal populations. To our knowledge this work constitutes the first experimental evidence that phenotypic diversity in an isogenic microbial population implements—at least in part—a game theoretic mixed strategy in response to a frequency-dependent foraging game. Given the abundance of multi-resource environments, we suspect that such foraging games may be fairly common; however it remains to be seen to what degree such games underlie the widespread phenotypic heterogeneity found in clonal microbial populations.

Materials and Methods

Strains: The three strains of *Saccharomyces cerevisiae* (wild type mixed strategist, GAL-OFF specialist and GAL-ON specialist), are modified from those used in Acar et al. (2005), which were derived from the diploid W303 strain of *S. cerevisiae.* All strains have a ADE2-P_{GAL1}-YFP reporter construct inserted at one *ade2* site for monitoring activation of the GAL network. Since one *ura3* locus was already occupied by inducible forms of *GAL80* or *GAL3*, yeast was first sporulated to isolate the remaining *ura3* locus. Identity of the haploids was confirmed by replica plating.

Haploids containing *ura3* were then transformed with the yeast integrating vector pRS306 containing URA3 and either RFP(tdTomato) or CFP cloned downstream of a TEF1 promoter. Constitutive fluorescence was confirmed by microscopy and flow cytometry. Fluorescent cells were then mated with the appropriate haploid to produce the desired strain. All strains were maintained on synthetic media his- and ura- agar dropout plates supplemented with 2% glucose.

The Gal80-inducible (GAL-OFF pure strategist) strain has a double *GAL80* deletion. P_{TETO2}*-GAL80* is inserted at one *ura3* locus, while P_{MYO2}*-rtTA* is inserted at an *ade2* locus. The GAL3-inducible (GAL-ON pure strategist) strain has a double *GAL3* deletion with P_{TETO2}-GAL3 inserted at a *ura3* locus and PMYO2-rtTA inserted at an *ade2* locus. Complete genotypes for the strains are found in the Supplementary Information.

Competitions: To initiate doxycycline induction in pure strategists, strains were initially mixed at desired initial frequencies from plated colonies, then incubated in 1.0 μg/mL doxycycline and 0.01% (w/v) glucose for 24 hours from a starting density of ~$3x10^4$ cells/mL to a saturating density of ~$6x10^6$ cells/mL, then diluted to ~$3x10^4$ cells/mL in synthetic media supplemented with glucose and galactose as indicated. Fractions were measured before and after incubation using a Miltenyi MACSquant flow cytometer (20,000+ cells per well), and population density was measured as absorbance at 600 nm in a microplate spectrophotometer (conversions assume 3×10^7 cells/mL at A_{600} = 1.0). Competitions involving the wild type GAL

network were necessarily limited to a single day because the wild type GAL-OFF fraction switches to GAL-ON when all the galactose is consumed. Subsequent days, beginning from the GAL-ON history, do not adopt a bimodal state in mixed sugars. Rather, the population adopts a very wide intermediate distribution of GAL activation which remains essentially unchanged over the course of multiple rounds of growth in mixed sugars.

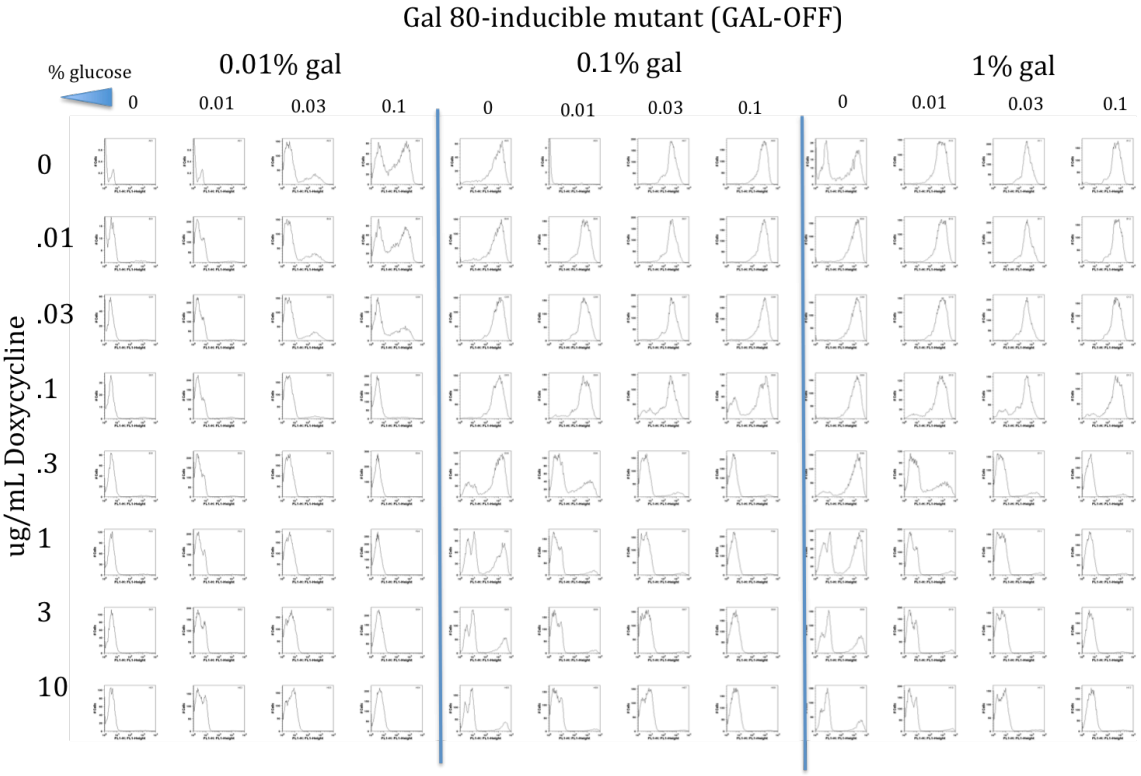

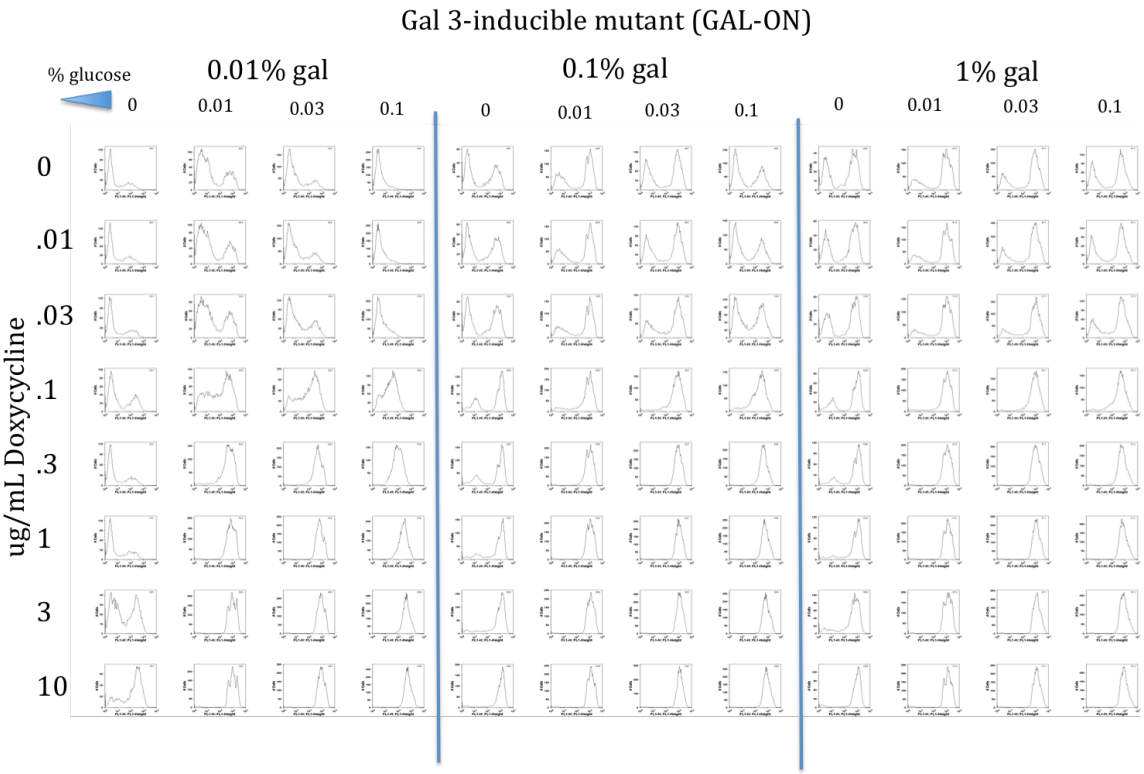

Figure 13. Doxycycline-dependent GAL induction states of GAL-OFF and GAL-ON mutants.
Doxycycline induction of Gal80 (top) and Gal3 (bottom) produces various GAL-induction states, dependent on Glucose and Galactose environmental composition as well. GAL activation in the Gal80 inducible mutant at 12 hours was decreasing in doxycycline concentration and glucose concentration, and increasing in galactose concentration. GAL activation in the Gal3 inducible mutant at four hours was increasing doxycycline concentration and galactose, and decreasing in glucose concentration. For competition conditions we settled on 1ug/mL doxycycline, as the lowest concentration that produced the desired pure strategies across the relevant ranges of glucose and galactose for the study.

Evolution of mixed strategists from pure strategists: Pure strategist Gal80-inducible (Gal80i or GAL-OFF) and Gal3-inducible (Gal3i or GAL-ON) strains were initially doxycycline-induced in media containing YNB, CSM, 1ug/mL doxycycline, and 0.01% glucose for 24 hours to saturation (~OD 0.25). Then they were diluted 100-

fold into 0.1% glucose, 0.1% galactose, and a mixture of 0.03% glu/0.05% gal. After consuming all the sugars and reaching saturation, they were diluted 1000x into fresh media and allowed to resume growing. This process was repeated on a daily basis for 26 days, however, due to two early rounds of 200x and 500x dilution, respectively, the total number of generations of growth over the 26 days was ~250. GAL-activation states were analyzed via flow-cytometry. To explore the composition of the evolved populations 30uL of evolved saturated culture was streaked on agar plates containing his- and ura- dropout YNB/CSM and 2% glucose. After 2 days of growth, individual colonies were suspended in separate wells of a 96-well plate and grown for 2 additional cycles in the presence of doxycycline and the same sugar mix they evolved in, and the resulting populations' GAL-activation states were measured via flow cytometry. All evolution experiments were performed in flat-bottomed 96-well culture plates incubated at 30°C.

Table 1. Yeast strains used in this study

Strain	Genotype
WT-R (mixed strategist)	*MAT***a**/α, *ura*3/*URA*3-P*TEF1-tdTomato, his3*::*HIS*3/*his*3, *ade*2/*ade*2::*ADE*2-P*GAL1-YFP*
WT-C (mixed strategist)	*MAT***a**/α, ura3/URA3-P*TEF1*his3::*HIS*3/*his*3, *ade*2/*ade*2::*ADE*2-P*GAL1-YFP*
Gal80i-R (GAL-OFF)	*MAT***a**/α, *URA3*-P*TEF1-tdTomato*/*ura*3::*URA*3-P*TETO2-GAL80*, *his*3::*HIS*3/*his*3, *ade*2::*ADE*2- P*MYO2-rtTA/ade*2::*ADE*2-P*GAL1-YFP, gal80*Δ::*KanMX*/*gal80*Δ::*KanMX*
Gal80i-C (GAL-OFF)	*MAT***a**/α, *URA3*-P*TEF1-CFP* /*ura*3::*URA*3-P*TETO2-GAL80*, *his*3::*HIS*3/*his*3, *ade*2::*ADE*2- P*MYO2-rtTA/ade*2::*ADE*2-P*GAL1-YFP, gal80*Δ::*KanMX*/*gal80*Δ::*KanMX*

Gal3i-C (GAL-ON)	*MAT***a**/α, *URA3*-P*TEF1-CFP*/*ura*3::*URA*3-P*TETO2-GAL3*, *his*3::*HIS*3/*his*3, *ade*2::*ADE*2- P*MYO2-rtTA/ade*2::*ADE*2-P*GAL1-YFP, gal3*Δ::*KanMX*/*gal3*Δ::*KanMX*
Gal3i-R (GAL-ON)	*MAT***a**/α, *URA3*-P*TEF1-RFP*/*ura*3::*URA*3-P*TETO2-GAL3*, *his*3::*HIS*3/*his*3, *ade*2::*ADE*2- P*MYO2-rtTA/ade*2::*ADE*2-P*GAL1-YFP, gal3*Δ::*KanMX*/*gal3*Δ::*KanMX*

Acknowledgements

The authors would like to thank A. Sanchez, A. Velenich, T. Gillespie, M. Springer, M. Laub and M. Vander Heiden, for helpful discussions, K. Axelrod for experimental help, and the members of the Gore Laboratory for helpful comments on the manuscript. M. Acar provided the strains from which the current strains were modified

Appendix: Foraging game models and simulations

Introduction

In Chapter 1, we used a simple two-player hawk-dove game to model basic predictions of scenarios with mutually invasible phenotypes. In chapter 2, we outline a similar scenario which we call the "simple foraging game," which we hypothesize should also exhibit mutual invasibility of phenotypes with a stable coexistence. However, in the context of microbial foraging, even the simple foraging game has significant conceptual deviations from the canonical two-player hawk-dove game.

In order to better understand the dynamics of foraging in environments with multiple resources, we constructed mathematical models and ran simulations. These models simulated three different and increasingly complex versions of a foraging game: 1) the simple foraging game outlined in figure 1 of chapter two, wherein two phenotypes are available and each can only consume one of the two resources, 2) the same game, but allowing for phenotypic switching with diauxic lag upon depletion of one of the resources, and 3) a game more closely approximating that played by our mutant GAL-ON and GAL-OFF pure strategists. This last game resembles the simple version, but allows the GAL-ON mutant to consume glucose even when galactose is still present.

Insights from these three simulated games include:

1. While the stable mixed strategy is the one in which both phenotypes have equal fitness, the "growth optimal" mixed strategy (defined as the mix that produces the highest average growth rate for the population) can be intuitively thought of in terms of resource consumption. The growth optimal mix is generally the phenotypic mix whose resulting population finishes both of the resources at the same time.
2. A growth rate "cost" for consuming one of the two resources is sufficient to distinguish the growth-optimal mixed strategy from the evolutionarily stable mixed strategy.
3. In several parameters—including the magnitude of the difference between the growth rates of the two phenotypes—the stable mix and the optimal mix diverge.
4. The canonical diauxic growth model, in which the resources are consumed sequentially in order of decreasing growth rate, is the unique solution (both stable and optimal) to the case in which the individuals are allowed to switch to the opposite resource with a relatively short diauxic lag time.
5. "Ratio sensing"—dividing between GAL-ON and GAL-OFF using the ratios of the two resources rather than their absolute concentrations (*84*)—will generally either allow the cell to closely adopt the optimal or the stable mix, but not both. Which solution concept can be adopted faithfully by ratio sensing depends on the parameters of the model.

Simple Foraging game

In the basic version of the foraging game, individuals are confronted with two finite resources, *A* and *B*, and can adopt one of two corresponding mutually exclusive phenotypes, *N* and *M*. Phenotype *N* is necessary to consume resource *A*, while phenotype *M* is necessary to consume resource *B*. Growth is exponential; per capita growth is constant while corresponding resources are abundant, and zero when the resource is depleted. Table 1 lists relevant parameters. For comparability with the later simulations of the yeast glucose/galactose foraging game, simulation values for the simple general foraging game are approximated from experimental data in Chapter 2 2S2 if *A* is glucose and *B* is galactose, and populations grow in a volume of $200uL$.

Table 1. Parameters for simulation of simple foraging game

Parameter	Description	Default simulation value
Y_0	*Initial population density*	*0.001 OD*
f_n	*Starting fraction of phenotype n*	*0.5*
f_m	*Starting fraction of phenotype m*	*0.5*
N_0	*Initial density of phenotype n*	$f_n Y_0$
M_0	*Initial density of phenotype m*	$f_m Y_0$
A_0	*Initial density of resource A*	*0.03% (w/v)*
B_0	*Initial density of resource B*	*0.05% (w/v)*

r_n	*Growth rate of phenotype n on resource A*	$0.5\ hr^{-1}$
r_m	*Growth rate of phenotype m on resource B*	$0.425\ hr^{-1}$
$\boldsymbol{\lambda}_A$	*Rate of consumption of resource A by phenotype n*	$0.0003\ OD^{-1}hr^{-1}$
$\boldsymbol{\lambda}_B$	*Rate of consumption of resource B by phenotype m*	$0.0003\ OD^{-1}hr^{-1}$

Growth conditions are as follows:

$$\frac{\boldsymbol{dN}}{\boldsymbol{dt}} = \begin{cases} r_n N, & \boldsymbol{A > 0} \\ \mathbf{0}, & \boldsymbol{A = 0} \end{cases}$$

$$\frac{\boldsymbol{dM}}{\boldsymbol{dt}} = \begin{cases} r_n M, & \boldsymbol{B > 0} \\ \mathbf{0}, & \boldsymbol{B = 0} \end{cases}$$

$$\frac{\boldsymbol{dA}}{\boldsymbol{dt}} = \begin{cases} -\boldsymbol{\lambda}_A N, & \boldsymbol{A > 0} \\ \mathbf{0}, & \boldsymbol{A = 0} \end{cases}$$

$$\frac{\boldsymbol{dB}}{\boldsymbol{dt}} = \begin{cases} -\boldsymbol{\lambda}_B M, & \boldsymbol{B > 0} \\ \mathbf{0}, & \boldsymbol{B = 0} \end{cases}$$

Simulations were run in 1-minute increments until all the resources were consumed. Figure 1 shows an example of a single simulated run with equal starting fractions of phenotypes *n* and *m*, and with a growth cost, *c=0.15,* for consuming sugar B such that:

$$r_m = r_n(1 - c) \qquad 0 \leq c \leq 1$$

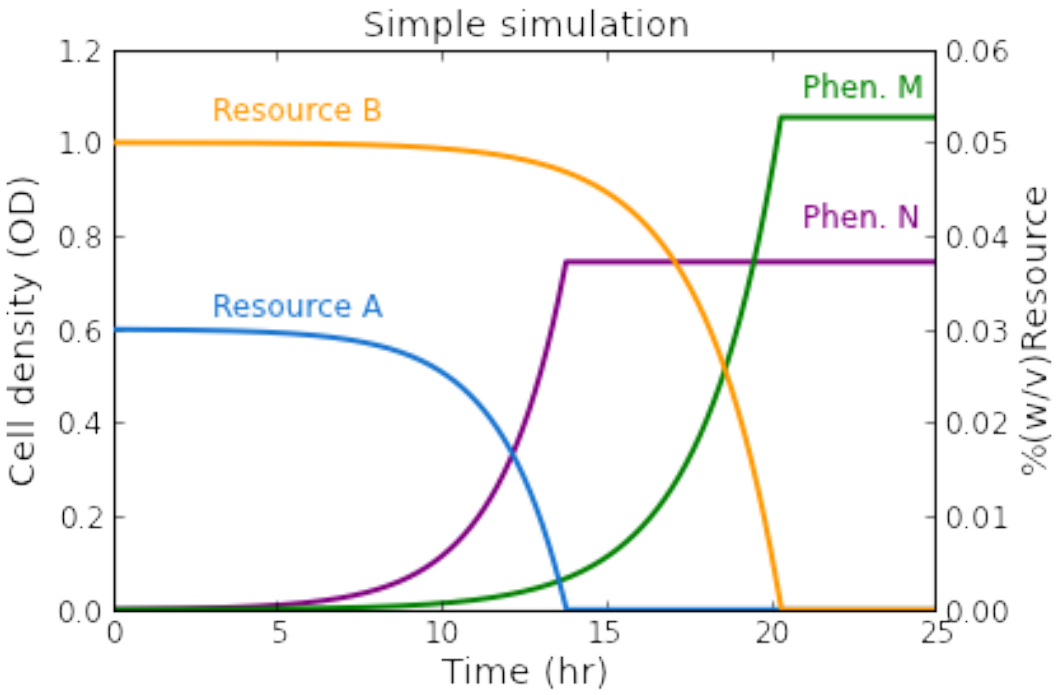

Figure1. Simple foraging game with two phenotypes and two resources
In this simulation of a simple foraging game, phenotype N and M can consume resource A and B, respectively, but not the opposite resource. Resource B is more abundant than resource A, but consuming resource B incurs a 15% growth cost. Phenotypes N and M start out at equal abundance in the population.

It is important to note that in the simulations, phenotype *m* consumes resources at the same rate as phenotype n, but it has a smaller growth rate. This simulates the presumably common scenario wherein a proportion of the energy extracted from a food source goes towards paying the metabolic cost of producing the proteins and enzymes necessary to consume the resource.

We next investigated the frequency-dependent game played between individuals of the two metabolic phenotypes. By running the simulation with a range of different

starting fractions of phenotype *m,* f_m, we show that even in this very simplified foraging game, our experimental results are borne out in simulation. There is strong negative frequency dependence and mutual invasibility between the phenotypes (Figure 2 a-b), with an intermediate stable equilibrium at about 60% phenotype *m.* We also calculated an average growth rate for the population defined as the final population density (which is basically the same regardless of population composition) divided by the time it took the population to reach saturation. We found, as we did with our experimental results, that the "optimal mix"—the composition that maximizes the population's growth—contains much more phenotype *m* than the evolutionarily stable mix (Figure 2d, grey dashed line). Since the final population densities are equal (Figure 2c), average growth rate can be thought of as mostly a function of the time it takes a population to consume all of the sugars. The optimal mix, with higher proportion of phenotype *m,* finishes the sugars the fastest, and is also the one that finishes the sugars at the same time. Indeed, in all cases where phenotype *m* has lower growth rate than phenotype *n,* we find the optimal mix contains a higher proportion of phenotype *m* than does the stable mix.

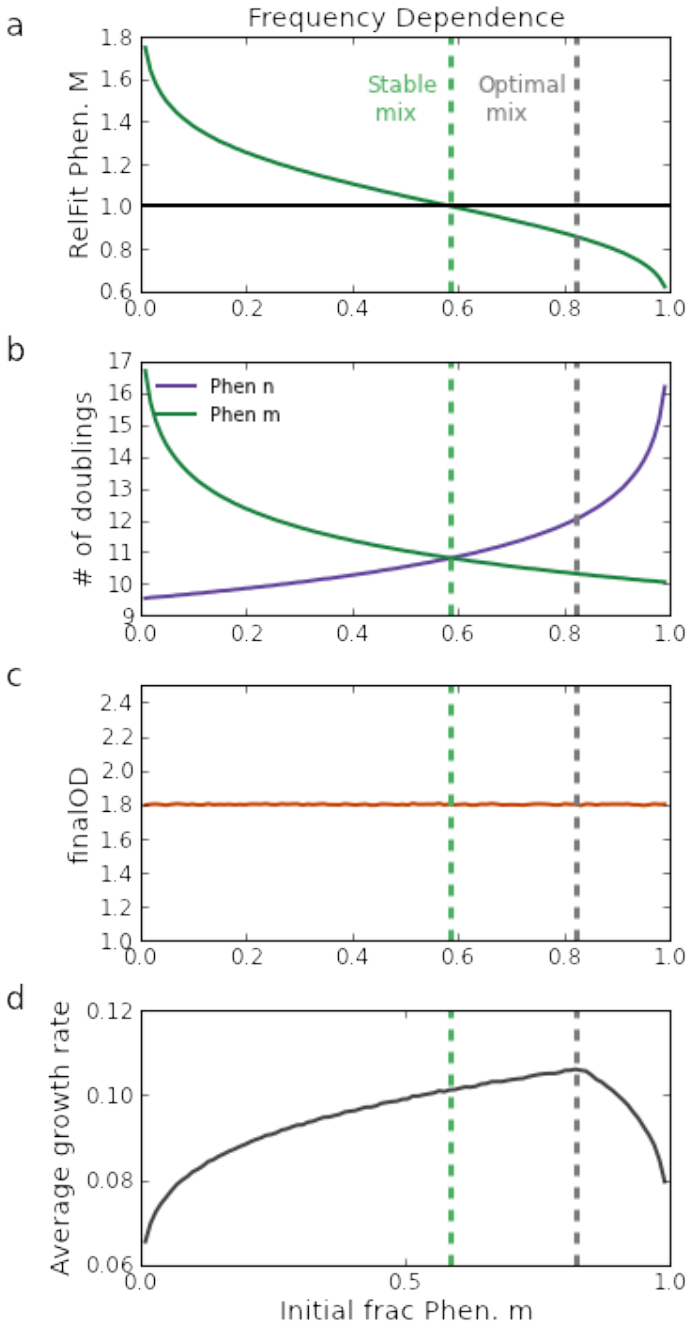

Figure 2. Frequency dependence of simple foraging game
As a function of initial population frequency of *m* phenotype: ***a)*** Relative fitness of the *m* phenotype and **b)** absolute fitness of both *n* (purple) and *m* (green) phenotypes indicate negative frequency dependence and mutual invasibility between phenotypes. Each phenotype is more fit than the other when rare. **c)** the final total population density remains essentially unchanged by the phenotypic composition of the population. However, since the time to reach the final OD does change depending on the composition of the population, then **(d)** the average growth rate of the population is

maximized at an intermediate frequencies. However, the optimal mix for the population (grey dotted line) is not the same as the evolutionarily stable mix (green dotted line).

Divergence of stable and optimal mixes

We next investigated the stable and optimal mixes as a function of phenotypic growth rates. Unsurprisingly, both the stable and the optimal mixes changed depending on the growth disparity between the phenotypes, but remarkably, they changed in opposite directions. As phenotype *m* gets less and less fit relative to phenotype *n,* the evolutionarily stable strategy is to adopt phenotype *m* less frequently, while the growth optimal strategy is to adopt phenotype *m* more frequently. This difference can be understood intuitively by recalling the informal definitions of the stable and the optimal mix.

The stable mix is the one in which all individuals have the same fitness regardless of phenotype. Hence, with negative frequency dependent fitness, if phenotype m becomes less fit, more individuals must adopt phenotype *n* until their fitness is correspondingly lower. On the other hand, the growth optimal mix is the one in which the two resources are consumed at the same time (so that no phenotype sits idle). A lower fitness for phenotype *m* means that resource *B* will take longer to consume than resource *A*, so to maximize the growth of the population, individuals should adopt phenotype *m* to consume resource *B* more quickly. Figure 3 illustrates the divergence of the stable and the optimal mixes as a function of the fitness cost of phenotype *m.* Intriguingly, although the "optimal mix" population grows much

faster than the “stable mix” population, an evolutionary competition between the two mixed strategies in the absence of opportunities for group or kin selection would favor would favor the slower-growing stable mix. A spatially-structured environment, however, may favor the optimal mix through group or kin selection(*67*). The simulations demonstrate that the two solution concepts diverge even with a relatively simple two-resource scenario.

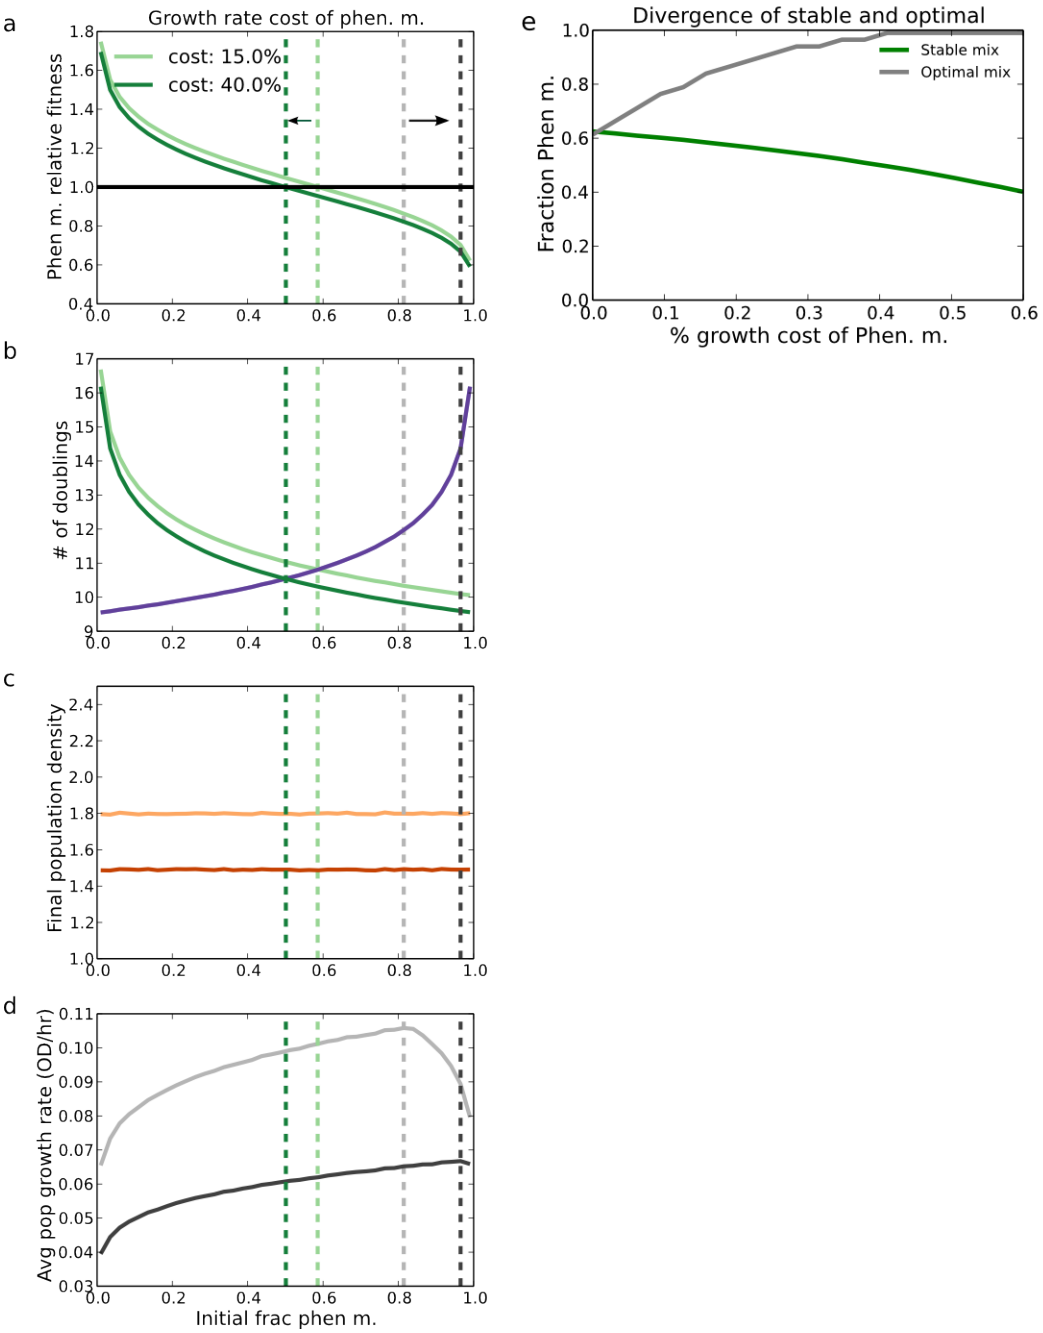

Figure 3 Stable and optimal mixes diverge with increasing metabolic cost of phenotype *m.*

(a-d) Frequency dependent fitness, similar to Figure 2, is plotted for two different growth costs of adopting phenotype *m*: *c = 0.15%* (light) and *c=0.4%* (dark). In other words, in the light case, phenotype *m* grows 15% slower and in the dark case it grows 40% slower than phenotype *n.* Higher phenotype *m* growth cost lowers the total number of divisions possible on resource B, shifting the stable equilibrium left, while the growth optimal mix (**d,** grey dotted lines) shifts right. **e)** The stable mix (green) and optimal mix

(gray) are plotted as a function of the growth cost of adopting phenotype *m.* The fact that the two solution concepts not only are different but also diverge serves to highlight that mixed ESS and optimal divisions of labor are very different solution concepts.

Foraging game with facultative phenotypic switching

The simple foraging game outlined above assumes that each phenotype can only consume one of the resources exclusively. However, this assumption is bears little resemblance to most microbial environments with multiple carbon sources. In reality, individuals can generally sense the presence of a resource and adopt the appropriate phenotype for consuming it, switching from one to the other if necessary. Indeed, the canonical response to multiple sugars is thought to be a diauxic response: consume the most advantageous carbon source first, then switch to the next most advantageous, and so on. A reasonable follow-up to the simple game, therefore, is to relax the consumption constraints and allow individuals to switch their phenotype when the resource runs out. However, given that many metabolic regimes such as galactose or lactose consumption involve the production of large numbers of specialized metabolic enzymes, a diauxic lag phase is typically observed during the switch. Given that little growth happens during this time, such a lag phase represents a cost to switching. Accordingly, for our simulations, we altered the simple foraging simulation to allow such phenotypic switching, and we introduced a new parameter, θ, the diauxic lag time:

Parameter	Description	Default simulation value
θ	*Diauxic shift lag time*	*2.0 hr*

Figure 4 shows simulations with the same default growth parameters as in Figure 1, but with ability to switch phenotypes after sugars are consumed and with a lag growth time of 2 hrs.

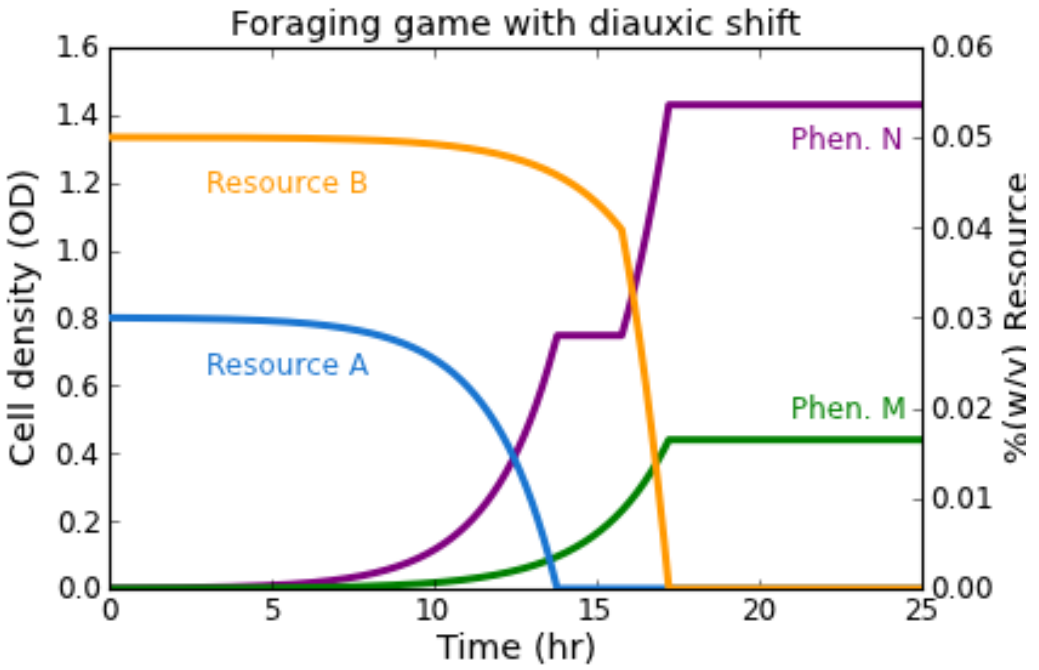

Figure 4. A simulated foraging game allowing for diauxic growth with a diauxic lag time of 2 hr. Individuals of phenotypes *n* and *m* begin at 80% and 20% of the population, respectively, and consume resource A and resource B, respectively. At around 13hr, resource A is depleted and individuals of phenotype N switch to consuming resource B after a diauxic lag of 2 hours.

Diauxic lag, stability, and optimality

To explore the effects of the switching costs on the stable mixed and optimal mixed strategies, we ran simulations with many different switching costs. As expected, at the limit of high diauxic lag time—or very costly switching—frequency dependence is identical to that in the simple foraging game (Figure 5a, darkest curve). At the limit of very low diauxic lag time, frequency dependence vanishes entirely (Figure 5a, lightest curve), and phenotype *n* is more fit than phenotype *m* regardless of population composition. Below a threshold lag time, the optimal and stable solution to the foraging game is the classic diauxic growth strategy: all individuals consume the best resource first.

With intermediate diauxic lag times, the situation is somewhat more complicated. The relative fitness of phenotype *m* as a function of its population frequency can be essentially divided into three sections: low, intermediate, and high fm. At intermediate f_m—close to the optimal mix—the fitness of the two phenotypes are identical to those in the simple foraging game (Figure 5a) because the two resources are consumed close enough together that there is not time enough for the faster phenotype to switch before the other resource is gone also. The higher the diauxic lag time, the wider this intermediate fitness zone is. One corollary of this observation is that, since the optimal mix is the one in which both resources are consumed at the same time, the optimal mix is not really responsive to differences in the diauxic lag time (Figure 5d). And in the parameters we have outlined, the stable mix follows a similar pattern, though there is a remarkable phenomenon where in low-to-intermediate lag times (1.5-2.0 hr in Figure 5d) the population mix

that grows the fastest consists almost entirely of phenotype *m*, while the population mix that is evolutionarily stable is 100% phenotype *n.*

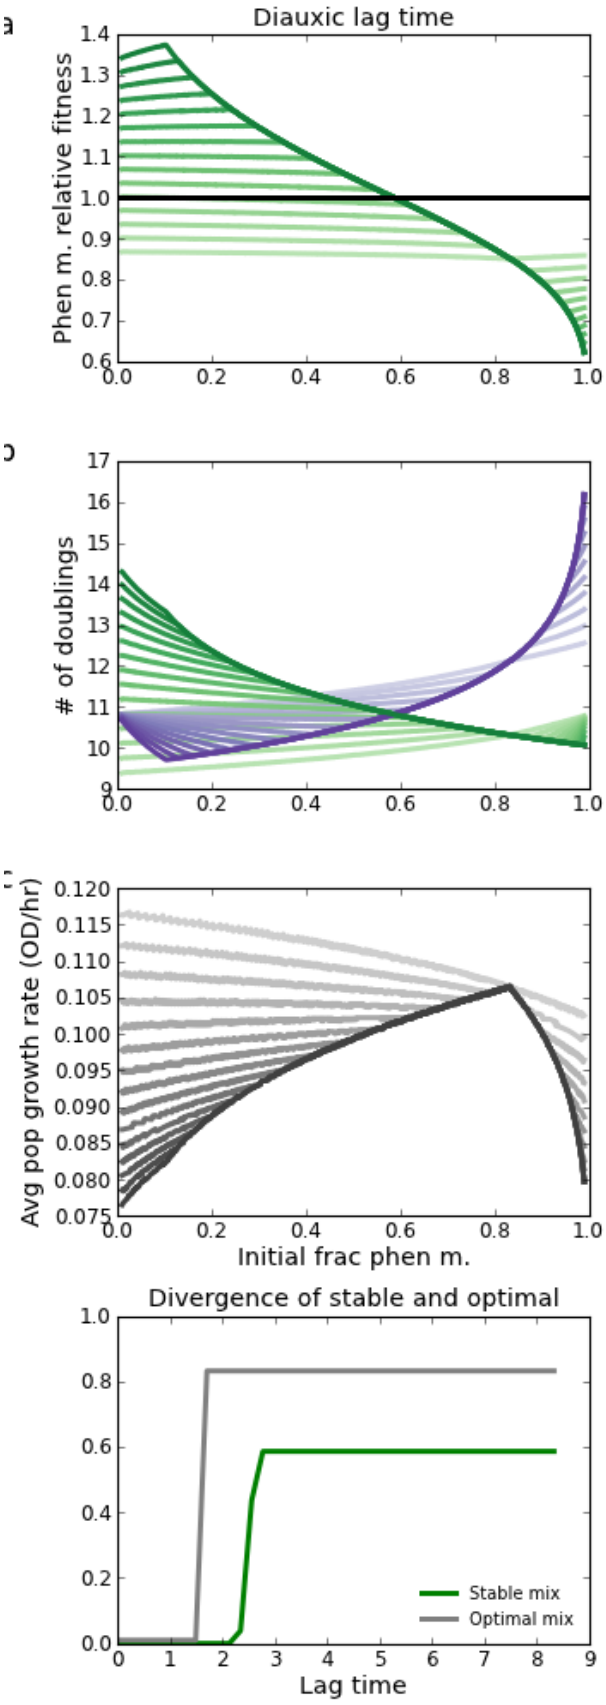
a
Diauxic lag time
Phen m. relative fitness
1.4
1.3
1.2
1.1
1.0
0.9
0.8
0.7
0.6
0.0
0.2
0.4
0.6
0.8
1.0
b
# of doublings
17
16
15
14
13
12
11
10
9
0.0
0.2
0.4
0.6
0.8
1.0
Avg pop growth rate (OD/hr)
0.120
0.115
0.110
0.105
0.100
0.095
0.090
0.085
0.080
0.075
0.0
0.2
0.4
0.6
0.8
1.0
Initial frac phen m.
Divergence of stable and optimal
1.0
0.8
0.6
0.4
0.2
0.0
Stable mix
Optimal mix
0
1
2
3
4
5
6
7
8
9
Lag time

Figure 5 The effect of diauxic lag times on the stable and optimal solutions to the foraging game
Frequency dependent **a)** relative fitness, **b)** number of divisions, and **c)** average population growth are shown for a range of diauxic lag times from 0 to 8 hours. At high diauxic lag times, the frequency dependence is identical to the simple model simulation, and at the limit of low diauxic lag times, frequency dependence disappears altogether and phenotype *n* is more fit than *m* at all frequencies. **d)** The stable (green) and optimal (gray) mixed strategies are shown as a function of diauxic lag time. Both solutions display transition between no phenotype *m* at low lag times to intermediate frequencies of phenotype *m* at high lag times.

Simulation of GAL-OFF and GAL-ON pure strategists

To try to better understand the foraging game being played between our GAL-ON and GAL-OFF pure strategist strains, we undertook a simulation that roughly approximates our understanding of that game. Because the GAL-OFF pure strategists do not switch to GAL-ON during the course of the experiment, the game resembles the simple foraging game, with the exception that in this case, phenotype *m* (the GAL-ON phenotype) can consume resource *A* (glucose) even when resource *B* is abundant. This necessitates us replacing the r_m and parameter with sugar-specific parameters.

Parameter	Description	Default simulation value
r_m^A	*GAL-ON growth rate on glucose*	$0.425\ hr^{-1}$
r_m^B	*GAL-ON growth rate on galactose*	$0.3825\ hr^{-1}$

Because in the mixed sugar conditions the GAL-ON cells are still GAL-ON even when consuming glucose, their growth rate on glucose reflects the 15% growth cost we imputed to them earlier. GAL-ON's growth rate on galactose was also estimated from experimental values. The most salient question that arose during these simulations was the amount of galactose that GAL-ON pure strategists consume. Some researchers have claimed that, in the presence of substantial glucose, even GAL-activated yeast do not consume substantial amounts of galactose (*80*), while others have observed some galactose consumption (*81, 84*). Our own data is somewhat ambiguous (see chapter 2 Figure 6). It appears that before the glucose is entirely consumed, the galactose has been somewhat consumed, though not to the same degree. Under growth parameters estimated from our experimental data, in order to see stable mixed strategies of intermediate frequencies we must either make the GAL network more costly to run than we anticipated (15-30% growth cost) or allow the GAL-ON mutants to consume significant galactose even in the presence of glucose (in that case, we assume that galactose is displacing rather than adding to the glucose in the metabolic pathways). Figure 6 illustrates the frequency-dependent fitness as a function of the proportion of GAL-ON's sugar consumption being galactose.

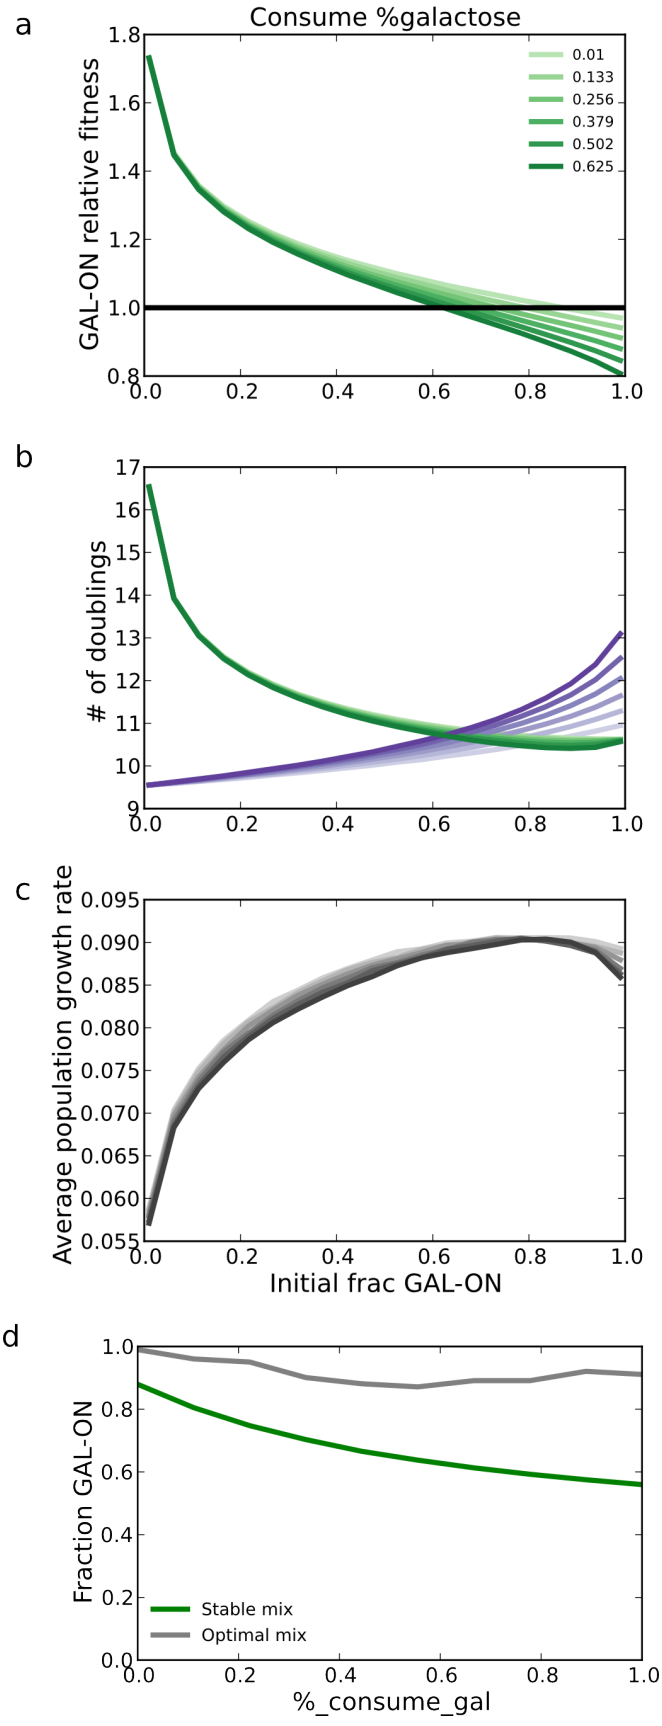
a
Consume %galactose
GAL-ON relative fitness
0.01
0.133
0.256
0.379
0.502
0.625
b
# of doublings
c
Average population growth rate
Initial frac GAL-ON
d
Fraction GAL-ON
Stable mix
Optimal mix
%_consume_gal

Figure 6. GAL-ON and GAL-OFF frequency dependence in .03/.05% glu/gal as a function of the distribution of sugar uptake for GAL-ON in mixed sugars.
Frequency dependent **a)** relative fitness, **b)** number of divisions, and **c)** average population growth are shown for a range different galactose-consumption conditions, ranging from GAL-ON only consuming 1% galactose while glucose is still present (lightest curve) to GAL-ON's consuming galactose in proportion to its initial fraction of total sugars (62.5%, darkest curve). The amount of galactose that GAL-ON consumes in the presence of glucose does not have a large effect on the optimal mix **(c,d),** but does have a large effect on the stable mix **(a,d) d)** The stable (red) and optimal (green) mixed strategies are shown as a function of the proportion of galactose consumed in the presence of both glucose and galactose.

References

1. M. Delbruck, The Burst Size Distribution in the Growth of Bacterial Viruses (Bacteriophages). *Journal of bacteriology* **50**, 131 (Aug, 1945).
2. E. O. Powell, An outline of the pattern of bacterial generation times. *Journal of general microbiology* **18**, 382 (Apr, 1958).
3. J. Lederberg, T. Iino, Phase Variation in Salmonella. *Genetics* **41**, 743 (1956).
4. A. Novick, M. Weiner, Enzyme Induction as an All-or-None Phenomenon. *P Natl Acad Sci USA* **43**, 553 (1957).
5. J. L. Spudich, D. E. Koshland, Non-Genetic Individuality - Chance in Single Cell. *Nature* **262**, 467 (1976).
6. M. B. Elowitz, A. J. Levine, E. D. Siggia, P. S. Swain, Stochastic gene expression in a single cell. *Science* **297**, 1183 (Aug 16, 2002).
7. A. Lobner-Olesen, Distribution of minichromosomes in individual Escherichia coli cells: implications for replication control. *Embo J* **18**, 1712 (Mar 15, 1999).
8. P. S. Swain, M. B. Elowitz, E. D. Siggia, Intrinsic and extrinsic contributions to stochasticity in gene expression. *Proc Natl Acad Sci U S A* **99**, 12795 (Oct 1, 2002).
9. M. Thattai, A. van Oudenaarden, Intrinsic noise in gene regulatory networks. *Proc Natl Acad Sci U S A* **98**, 8614 (Jul 17, 2001).
10. J. Paulsson, Summing up the noise in gene networks. *Nature* **427**, 415 (Jan 29, 2004).
11. E. M. Ozbudak, M. Thattai, I. Kurtser, A. D. Grossman, A. van Oudenaarden, Regulation of noise in the expression of a single gene. *Nature genetics* **31**, 69 (May, 2002).
12. A. Raj, C. S. Peskin, D. Tranchina, D. Y. Vargas, S. Tyagi, Stochastic mRNA synthesis in mammalian cells. *PLoS biology* **4**, 1707 (Oct, 2006).
13. H. H. McAdams, A. Arkin, Stochastic mechanisms in gene expression. *P Natl Acad Sci USA* **94**, 814 (Feb 4, 1997).
14. D. Hebenstreit, Are gene loops the cause of transcriptional noise? *Trends Genet* **29**, 333 (Jun, 2013).
15. D. Dubnau, R. Losick, Bistability in bacteria. *Mol Microbiol* **61**, 564 (Aug, 2006).
16. S. R. Biggar, G. R. Crabtree, Cell signaling can direct either binary or graded transcriptional responses. *Embo J* **20**, 3167 (Jun 15, 2001).
17. A. Arkin, J. Ross, H. H. McAdams, Stochastic kinetic analysis of developmental pathway bifurcation in phage lambda-infected Escherichia coli cells. *Genetics* **149**, 1633 (Aug, 1998).
18. G. Balazsi, A. van Oudenaarden, J. J. Collins, Cellular Decision Making and Biological Noise: From Microbes to Mammals. *Cell* **144**, 910 (Mar 18, 2011).

19. J. Hasty, J. Pradines, M. Dolnik, J. J. Collins, Noise-based switches and amplifiers for gene expression. *Proc Natl Acad Sci U S A* **97**, 2075 (Feb 29, 2000).
20. T. Cagatay, M. Turcotte, M. B. Elowitz, J. Garcia-Ojalvo, G. M. Suel, Architecture-dependent noise discriminates functionally analogous differentiation circuits. *Cell* **139**, 512 (Oct 30, 2009).
21. J. M. Raser, E. K. O'Shea, Control of stochasticity in eukaryotic gene expression. *Science* **304**, 1811 (Jun 18, 2004).
22. H. J. E. Beaumont, J. Gallie, C. Kost, G. C. Ferguson, P. B. Rainey, Experimental evolution of bet hedging. *Nature* **462**, 90 (Nov 5, 2009).
23. M. Acar, J. T. Mettetal, A. van Oudenaarden, Stochastic switching as a survival strategy in fluctuating environments. *Nature genetics* **40**, 471 (Apr, 2008).
24. H. Maamar, A. Raj, D. Dubnau, Noise in gene expression determines cell fate in Bacillus subtilis. *Science* **317**, 526 (Jul 27, 2007).
25. B. Lehner, Selection to minimise noise in living systems and its implications for the evolution of gene expression. *Mol Syst Biol* **4**, 170 (2008).
26. I. Lestas, G. Vinnicombe, J. Paulsson, Fundamental limits on the suppression of molecular fluctuations. *Nature* **467**, 174 (Sep 9, 2010).
27. M. Slatkin, Hedging Ones Evolutionary Bets. *Nature* **250**, 704 (1974).
28. I. G. de Jong, P. Haccou, O. P. Kuipers, Bet hedging or not? A guide to proper classification of microbial survival strategies. *Bioessays* **33**, 215 (Mar, 2011).
29. J. W. Veening, W. K. Smits, O. P. Kuipers, Bistability, epigenetics, and bet-hedging in bacteria. *Annu Rev Microbiol* **62**, 193 (2008).
30. J. Muller, B. A. Hense, T. M. Fuchs, M. Utz, C. Potzsche, Bet-hedging in stochastically switching environments. *J Theor Biol* **336**, 144 (Nov 7, 2013).
31. J. M. McNamara, Optimal life histories for structured populations in fluctuating environments. *Theoretical Population Biology* **51**, 94 (Apr, 1997).
32. D. Cohen, Optimizing Reproduction in a Randomly Varying Environment. *J Theor Biol* **12**, 119 (1966).
33. D. Cohen, Optimizing Reproduction in a Randomly Varying Environment When a Correlation May Exist between Conditions at Time a Choice Has to Be Made and Subsequent Outcome. *J Theor Biol* **16**, 1 (1967).
34. J. R. Gremer, D. L. Venable, Bet hedging in desert winter annual plants: optimal germination strategies in a variable environment. *Ecol Lett*, (Jan 7, 2014).
35. T. Philippi, Bet-hedging germination of desert annuals: beyond the first year. *The American naturalist* **142**, 474 (Sep, 1993).
36. J. T. Lennon, S. E. Jones, Microbial seed banks: the ecological and evolutionary implications of dormancy. *Nature Reviews Microbiology* **9**, 119 (Feb, 2011).
37. M. J. Clauss, D. L. Venable, Seed Germination in Desert Annuals: An Empirical Test of Adaptive Bet Hedging. *The American naturalist* **155**, 168 (Feb, 2000).
38. M. Gamelon *et al.*, The relationship between phenotypic variation among offspring and mother body mass in wild boar: evidence of coin-flipping? *Journal of Animal Ecology* **82**, 937 (Sep, 2013).
39. M. L. Evans, M. Dionne, K. M. Miller, L. Bernatchez, Mate choice for major histocompatibility complex genetic divergence as a bet-hedging strategy in

the Atlantic salmon (Salmo salar). *P Roy Soc B-Biol Sci* **279**, 379 (Jan 22, 2012).
40. N. Q. Balaban, J. Merrin, R. Chait, L. Kowalik, S. Leibler, Bacterial persistence as a phenotypic switch. *Science* **305**, 1622 (Sep 10, 2004).
41. E. Kussell, R. Kishony, N. Q. Balaban, S. Leibler, Bacterial persistence: a model of survival in changing environments. *Genetics* **169**, 1807 (Apr, 2005).
42. S. F. Levy, N. Ziv, M. L. Siegal, Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant. *PLoS biology* **10**, e1001325 (2012).
43. J. W. Veening *et al.*, Bet-hedging and epigenetic inheritance in bacterial cell development. *P Natl Acad Sci USA* **105**, 4393 (Mar 18, 2008).
44. A. Solopova *et al.*, Bet-hedging during bacterial diauxic shift. *Proc Natl Acad Sci U S A*, (May 5, 2014).
45. S. E. Jones, J. T. Lennon, Dormancy contributes to the maintenance of microbial diversity. *P Natl Acad Sci USA* **107**, 5881 (Mar 30, 2010).
46. P. Piggot, Epigenetic Switching: Bacteria Hedge Bets about Staying or Moving. *Current Biology* **20**, R480 (Jun 8, 2010).
47. M. P. Stumpf, Z. Laidlaw, V. A. Jansen, Herpes viruses hedge their bets. *Proc Natl Acad Sci U S A* **99**, 15234 (Nov 12, 2002).
48. R. Meadows, Yeast survive by hedging their bets. *PLoS biology* **10**, e1001327 (2012).
49. D. Z. Childs, C. J. Metcalf, M. Rees, Evolutionary bet-hedging in the real world: empirical evidence and challenges revealed by plants. *Proc Biol Sci* **277**, 3055 (Oct 22, 2010).
50. J. F. Nash, Equilibrium Points in N-Person Games. *P Natl Acad Sci USA* **36**, 48 (1950).
51. P. Y. Quenette, J. F. Gerard, Does frequency-dependent selection optimize fitness? *J Theor Biol* **159**, 381 (Dec 21, 1992).
52. J. Maynard Smith, *Evolution and the theory of games*. (Cambridge University Press, Cambridge ; New York, 1982), pp. viii, 224 p.
53. J. M. Smith, G. R. Price, Logic of Animal Conflict. *Nature* **246**, 15 (1973).
54. D. M. Wolf, V. V. Vazirani, A. P. Arkin, Diversity in times of adversity: probabilistic strategies in microbial survival games. *J Theor Biol* **234**, 227 (May 21, 2005).
55. D. M. Wolf, V. V. Vazirani, A. P. Arkin, A microbial modified prisoner's dilemma game: how frequency-dependent selection can lead to random phase variation. *J Theor Biol* **234**, 255 (May 21, 2005).
56. C. Adami, J. Schossau, A. Hintze, Evolution and stability of altruist strategies in microbial games. *Physical Review E* **85**, (Jan 20, 2012).
57. M. A. Nowak, C. E. Tarnita, E. O. Wilson, The evolution of eusociality. *Nature* **466**, 1057 (Aug 26, 2010).
58. W. D. Hamilton, Genetical Evolution of Social Behaviour I. *J Theor Biol* **7**, 1 (1964).
59. M. A. Riley, D. M. Gordon, The ecological role of bacteriocins in bacterial competition. *Trends Microbiol* **7**, 129 (Mar, 1999).

60. B. Kerr, M. A. Riley, M. W. Feldman, B. J. Bohannan, Local dispersal promotes biodiversity in a real-life game of rock-paper-scissors. *Nature* **418**, 171 (Jul 11, 2002).
61. M. Ackermann *et al.*, Self-destructive cooperation mediated by phenotypic noise. *Nature* **454**, 987 (Aug 21, 2008).
62. M. Diard *et al.*, Stabilization of cooperative virulence by the expression of an avirulent phenotype. *Nature* **494**, 353 (Feb 21, 2013).
63. J. Ripa, H. Olofsson, N. Jonzen, What is bet-hedging, really? *Proc Biol Sci* **277**, 1153 (Apr 22, 2010).
64. M. Arnoldini *et al.*, Bistable expression of virulence genes in salmonella leads to the formation of an antibiotic-tolerant subpopulation. *PLoS biology* **12**, e1001928 (Aug, 2014).
65. M. Viney, S. E. Reece, Adaptive noise. *P Roy Soc B-Biol Sci* **280**, (Sep 22, 2013).
66. B. B. Pradhan, S. Chatterjee, Reversible nongenetic phenotypic heterogeneity in bacterial quorum sensing. *Mol Microbiol*, (Mar 7, 2014).
67. M. A. Nowak, Five rules for the evolution of cooperation. *Science* **314**, 1560 (Dec 8, 2006).
68. A. M. Simons, Modes of response to environmental change and the elusive empirical evidence for bet hedging. *Proc Biol Sci* **278**, 1601 (Jun 7, 2011).
69. C. S. Wylie, A. D. Trout, D. A. Kessler, H. Levine, Optimal strategy for competence differentiation in bacteria. *PLoS Genet* **6**, e1001108 (Sep, 2010).
70. C. D. Nadell, J. B. Xavier, K. R. Foster, The sociobiology of biofilms. *FEMS Microbiol Rev* **33**, 206 (Jan, 2009).
71. J. A. Shapiro, Thinking about bacterial populations as multicellular organisms. *Annu Rev Microbiol* **52**, 81 (1998).
72. M. Arnoldini *et al.*, Bistable Expression of Virulence Genes in Salmonella Leads to the Formation of an Antibiotic-Tolerant Subpopulation. *PLoS biology* **12**, (Aug, 2014).
73. O. X. Cordero, M. F. Polz, Explaining microbial genomic diversity in light of evolutionary ecology. *Nature Reviews Microbiology* **12**, 263 (Apr, 2014).
74. J. Gore, H. Youk, A. van Oudenaarden, Snowdrift game dynamics and facultative cheating in yeast. *Nature* **459**, 253 (May 14, 2009).
75. J. B. Xavier, K. R. Foster, Cooperation and conflict in microbial biofilms. *Proc Natl Acad Sci U S A* **104**, 876 (Jan 16, 2007).
76. G. A. Parker, J. M. Smith, Optimality Theory in Evolutionary Biology. *Nature* **348**, 27 (Nov 1, 1990).
77. J. M. Gancedo, Yeast carbon catabolite repression. *Microbiol Mol Biol R* **62**, 334 (Jun, 1998).
78. M. Acar, A. Becskei, A. van Oudenaarden, Enhancement of cellular memory by reducing stochastic transitions. *Nature* **435**, 228 (May 12, 2005).
79. S. A. Ramsey *et al.*, Dual feedback loops in the GAL regulon suppress cellular heterogeneity in yeast. *Nature genetics* **38**, 1082 (Sep, 2006).
80. O. S. Venturelli, I. Zuleta, R. M. Murray, H. El-Samad, Population diversification in a yeast metabolic program promotes anticipation of environmental shifts. *bioRxiv*, (2014).

81. J. Wang *et al.*, Natural variation in preparation for nutrient depletion reveals a cost-benefit tradeoff. *PLoS biology* **13**, e1002041 (Jan, 2015).
82. A. M. New *et al.*, Different levels of catabolite repression optimize growth in stable and variable environments. *PLoS biology* **12**, e1001764 (Jan, 2014).
83. M. L. Friesen, G. Saxer, M. Travisano, M. Doebeli, Experimental evidence for sympatric ecological diversification due to frequency-dependent competition in Escherichia coli. *Evolution* **58**, 245 (Feb, 2004).
84. R. Escalante-Chong *et al.*, Galactose metabolic genes in yeast respond to a ratio of galactose and glucose. *Proc Natl Acad Sci U S A* **112**, 1636 (Feb 3, 2015).