Edinburgh Research Explorer

Phenotypic Switching Can Speed up Microbial Evolution

Phenotypic switching can speed up biological evolution of microbes

Andrew C. Tadrowski[1], Martin R. Evans[1], and Bartłomiej Waclaw[1,2]

[1] *SUPA, School of Physics and Astronomy, The University of Edinburgh,*
Peter Guthrie Tait Road, Edinburgh EH9 3FD, United Kingdom
[2] *Centre for Synthetic and Systems Biology, The University of Edinburgh*

Stochastic phenotype switching has been suggested to play a beneficial role in microbial populations by leading to the division of labour among cells, or ensuring that at least some of the population survives an unexpected change in environmental conditions. Here we use a computational model to investigate an alternative possible function of stochastic phenotype switching - as a way to adapt more quickly even in a static environment. We show that when a genetic mutation causes a population to become less fit, switching to an alternative phenotype with higher fitness (growth rate) may give the population enough time to develop compensatory mutations that increase the fitness again. The possibility of switching phenotypes can reduce the time to adaptation by orders of magnitude if the "fitness valley" caused by the deleterious mutation is deep enough. Our work has important implications for the emergence of antibiotic-resistant bacteria. In line with recent experimental findings we hypothesise that switching to a slower growing but less sensitive phenotype helps bacteria to develop resistance by exploring a larger set of beneficial mutations while avoiding deleterious ones.

Keywords: evolution | population genetics | stochastic phenotype switching

INTRODUCTION

Biological evolution relies on two mechanisms which are instrumental in natural selection: preferential survival of better adapted individuals (selection) and variations among individuals (phenotypic variation). One of the sources of phenotypic variability is genetic alteration due to mutations and recombination. However, even genetically identical organisms will often behave differently because the same genotype may lead to many different phenotypes - observable traits of an organism - as a result of environmental factors and the organism's history. Although ubiquitous and easily observed in animals and plants, phenotypic diversity can already be amply demonstrated in microorganisms. Examples range from different cell sizes depending on growth medium [49], through bistability in utilization of different food sources [35], to diversification between motile/non-motile cells [20]. Microorganisms are often able to switch between these phenotypes in response to a change in external conditions such as the arrival of a new food source or depletion of the currently used one. A typical example is diauxic shift - a switch to another food source, for example from glucose to cellobiose in *L. lactis* when glucose becomes depleted [41], which involves altering gene expression levels without changing the genetic code.

Some microorganisms switch seemingly randomly between two or more phenotypes even in the absence of external stimuli. This causes the population to become phenotypically heterogeneous. Several explanations have been proposed as to why *stochastic phenotype switching* has evolved [1]. One of them is the division of labour [2] in which different microbial cells perform different functions and this maximizes the benefit to the population. Another theory, called bet hedging [6, 26], proposes that

in a fluctuating and unpredictable environment it pays to have a fraction of the population in a different state, which is perhaps maladapted to the present environment but better suited to possible future environments. Since only a small fraction of the population expresses the maladapted phenotype at any one time, this strategy conserves resources while allowing the population to stay prepared for an unexpected change. Examples include bacterial persisters [4, 25], flu(Ag43)/fim switch [3, 4, 19, 44] and competence to non-competence switching in the bacterium *B. subtilis* [29]. Animal cells also exhibit this behaviour. Epithelial-to-mesenhymal transition – a landmark in cancer progression – is thought to be a phenotypic, epigenetic change [33]. There is also some evidence that phenotypic switching may be involved in resistance to chemotherapy [21], although this remains controversial [50].

Here we investigate an alternative possible role for the evolution of phenotype switching in microbes: the facilitation of genetic evolution in a static environment. Previous theoretical work [28] has demonstrated that switching to a persistent phenotype can provide a larger pool of bacteria to evolve genetic resistance to antibiotics. That work only addressed the situation in which resistance is caused by a single, beneficial mutation. However, a significant change of fitness, defined here as a measure of a microorganism's growth rate, often requires multiple mutations. Although many fitness landscapes have at least one accessible pathway [14, 46], along which the fitness increases monotonically, in some cases a fitness valley - a genotype with a lower fitness than all its neighbours - is unavoidable. For example, developing resistance to the antibiotic streptomycin involves a fitness cost which must be counteracted by compensatory mutations, i.e. subsequent mutations that increase the microorganism's
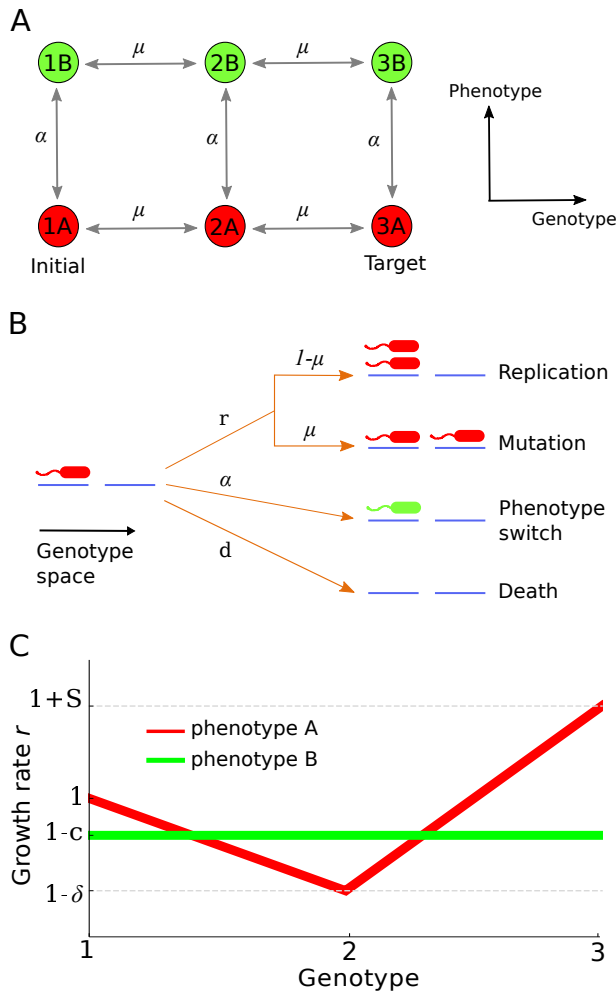
Figure 1: The model. (A) Diagram showing the six possible states of a cell and the available transitions between them. The genotypes are labelled 1, 2 and 3, the phenotypic states are labelled A and B. Transitions between genotypes/phenotypes occur at rates μ and α, respectively. All cells are initially in state 1A. Evolution continues until a single cell reaches the target state 3A. (B) Each cell can replicate, switch phenotype, or die with rates b, α and d respectively. Upon replication a cell has the probability μ of producing a mutant of each neighbouring genotype. (C) The fitness landscapes for both phenotypes. Phenotype A has a fitness valley at 2A while phenotype B has uniform fitness across all genotypes.

fitness to at least that of the ancestral strain [30, 38, 39].

Here we study evolution in the presence of such a fitness valley, in which a population must acquire two mutations to reach the best-adapted genotype. The first mutation is deleterious and corresponds to a valley in the fitness landscape that is crossed by a second, compensatory mutation. Theory and computer modelling suggests that fitness valleys significantly affect biological evolution [9, 42, 47, 48], often, but not always, slowing it down. However, in our model cells can also switch to an alternative phenotype in which the fitness landscape is

flat. We show that this switching enables the population to avoid the costly valley, even if the alternative phenotype is less fit than the initial state. We also demonstrate the existence of an optimal range for the switching rate. In this range the time to evolve the best-adapted genotype can be reduced by many orders of magnitude compared to the case without switching. Finally, we show that if the switching rate is allowed to evolve it will converge to values within this optimal range.

MODEL

We consider a population of haploid cells. Each cell can be in one of two phenotypic states, A and B, and can assume one of three genotypes ("genetic states"), as shown in Fig. 1A,B. Cells replicate stochastically with rate $r_i(1 - N/K)$, where i labels one of the six possible states, N is the total population size and K is the carrying capacity of the environment. The logistic-like factor $(1 - N/K)$ causes growth to cease when N becomes as large as K, which limits the total population size. During replication a cell will produce a mutated offspring of an adjacent genotype with probability μ. Mutation does not change the phenotypic state. All cells switch randomly between the states A and B at the symmetric switching rate α. Cells are randomly removed from the population at a constant death rate d. If the carrying capacity K is small ($K = 10 \cdots 10^4$) the model is appropriate to describe a small microbial population growing in a microfluidic chemostat with constant dilution rate [5]. For larger K ($K = 10^4 \cdots 10^9$) the model is relevant to populations cultured in mesoscopic (cm-size) chemostats.

The population initially consists of all individuals of type 1A, i.e. of genotype 1 in phenotypic state A, which has the growth rate $r_{1A} = 1$ (arbitrary units). The initial population size, unless otherwise stated, is equal to the equilibrium size $(1 - d)K$. State 3A is the global maximum of both fitness landscapes with the growth rate $r_{3A} = 1 + S$, where $S > 0$ is its selective advantage over the 'wild-type' 1A (Fig. 1C). Such a beneficial mutant has the probability of fixation, from a single individual in a population of cells in state 1A, approximately given by $\frac{1-e^{-S}}{1-e^{-KS}} \approx S$ [34] when $1/K \ll S \ll 1$. We are interested in the *adaptation time T* that it takes for the population to evolve the first individual in state 3A, conditioned on the population not going extinct (otherwise the time would be infinite).

We shall begin by considering the case in which the growth rate $r_{2A} = 1 - \delta$, while the growth rates $r_{1B} = r_{2B} = r_{3B} = 1$, which leads to a "fitness valley" in phenotype A and a flat fitness landscape in phenotype B (Fig. 1C). In the absence of phenotype switching the only way a population in state 1A can evolve an individual in state 3A is to go through the low-fitness state 2A. However, phenotype B allows the population to traverse an

alternative route that avoids state 2A.

RESULTS

Optimal range of the switching rate exists in which evolution is fastest

Figure 2A shows simulation results examining the effect of phenotype switching on the adaptation time T for a range of the parameters of the model. In all cases the presence of low frequency switching, i.e. small α, decreases the adaptation time by orders of magnitude. Holding all other parameters fixed and changing only the switching rate reveals an optimal range for α (Fig. 2B) in which T is minimized provided μ is sufficiently small. For the smallest mutation probability ($\mu = 10^{-5}$) considered in Fig. 2B the minimal adaptation time is approximately two orders of magnitude lower than in the absence of switching ($\alpha = 0$), or when switching is frequent (α very large). Only if the mutation rate is unrealistically large does switching not speed up evolution (SI and Fig. S1).

In the absence of transitions between states 2A and 2B (i.e. removing this step from the model in Fig. 1A) the adaptation time decreases monotonically with α (SI and Fig. S2).

Fastest evolutionary trajectories avoid the valley

To understand the evolutionary trajectory selected in the optimal range of switching frequency we examined the histories of successful cells, i.e. the states from Fig. 1A visited during evolution from state 1A to the final state 3A. We then represented each cell's trajectory in the state space as a sequence of links connecting the visited states (Fig. 3A). These were then classified as one of 21 classes by grouping sequences with multiple (back-and-forth) transitions between the states together with those without (Fig. 3A). Fig. 3B shows the most probable trajectory as a function of μ and α. As expected, simulated evolution favours different trajectories depending on the values of μ, α. The (μ, α)-space can be approximately separated into three regions, each corresponding to one or more trajectory classes. Region 1 corresponds to trajectories that go through the fitness valley at 2A. In this region mutations are frequent enough to offset the loss of fitness incurred when passing through the valley. In region 2 the dominant trajectory avoids the valley by switching to phenotype B. This region corresponds to the fastest adaptation times from Fig. 2B. Region 3 is characterised by a mixture of trajectory types using both phenotype B and visiting state 2A. This region corresponds to large α which not only enables transitions to phenotype B but also from state 2B to the deleterious state 2A. The latter is responsible for the increase in the adapta-
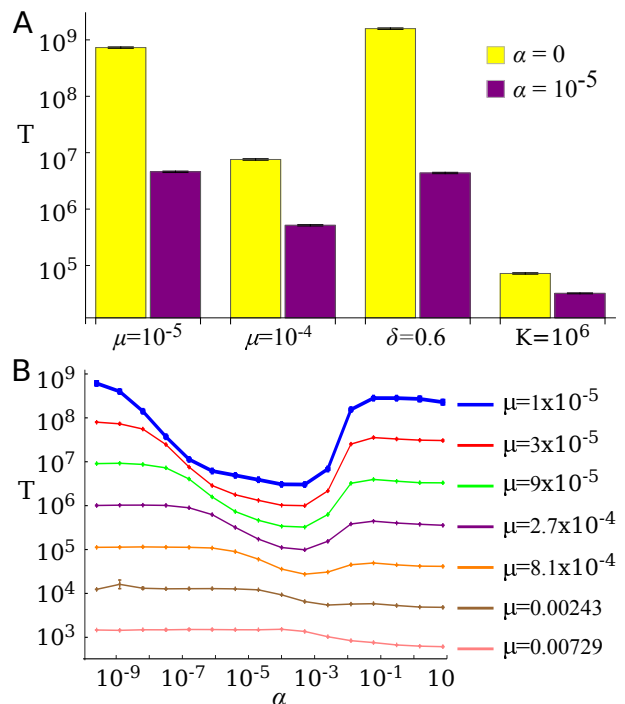

Figure 2: Adaptation time T for different scenarios – note logarithmic scale. (A) A bar chart comparing pairs of T values with and without switching phenotypes ($\alpha = 10^{-5}$ and $\alpha = 0$ respectively) for different parameter values. Left-most pair of bars: $K = 100, \mu = 10^{-5}, \delta = 0.4$ and $d = 0.1$. A label underneath each pair of bars indicates which variable has been changed compared to the left-most pair. (B) T as a function of the switching rate α for a range of mutation probabilities μ. Parameters as in (A). For small enough μ an optimal (minimizing T) switching rate can be seen.

tion time for large α in Fig. 2B. A similar decomposition into three regions can be obtained from the probability of individual transitions occurring between the states of the system (Fig. S3).

We next examined which trajectories would be preferred if we varied K instead of μ, as the carrying capacity is more likely to vary significantly in real microbial populations and is much easier to control experimentally than the mutation rate. Since we ascertained earlier (cf. the previous paragraph) that the time to adaptation is most significantly affected by whether the trajectory visits state 2A or not, we further reduced the number of trajectory classes to just two (Fig. 3C): the "u" (upper) trajectory type that avoids state 2A or the "v" (valley) trajectory type that visits it. We found a large region in the space (K, α) that favours trajectories that avoid the state 2A (Fig. 3D). This region extends to large values of $K \sim 10^9$. The adaptation time depends non-monotonically on the population size K (Fig. S4); similar non-monotonic behaviour has been seen in models without phenotype switching [13, 32].
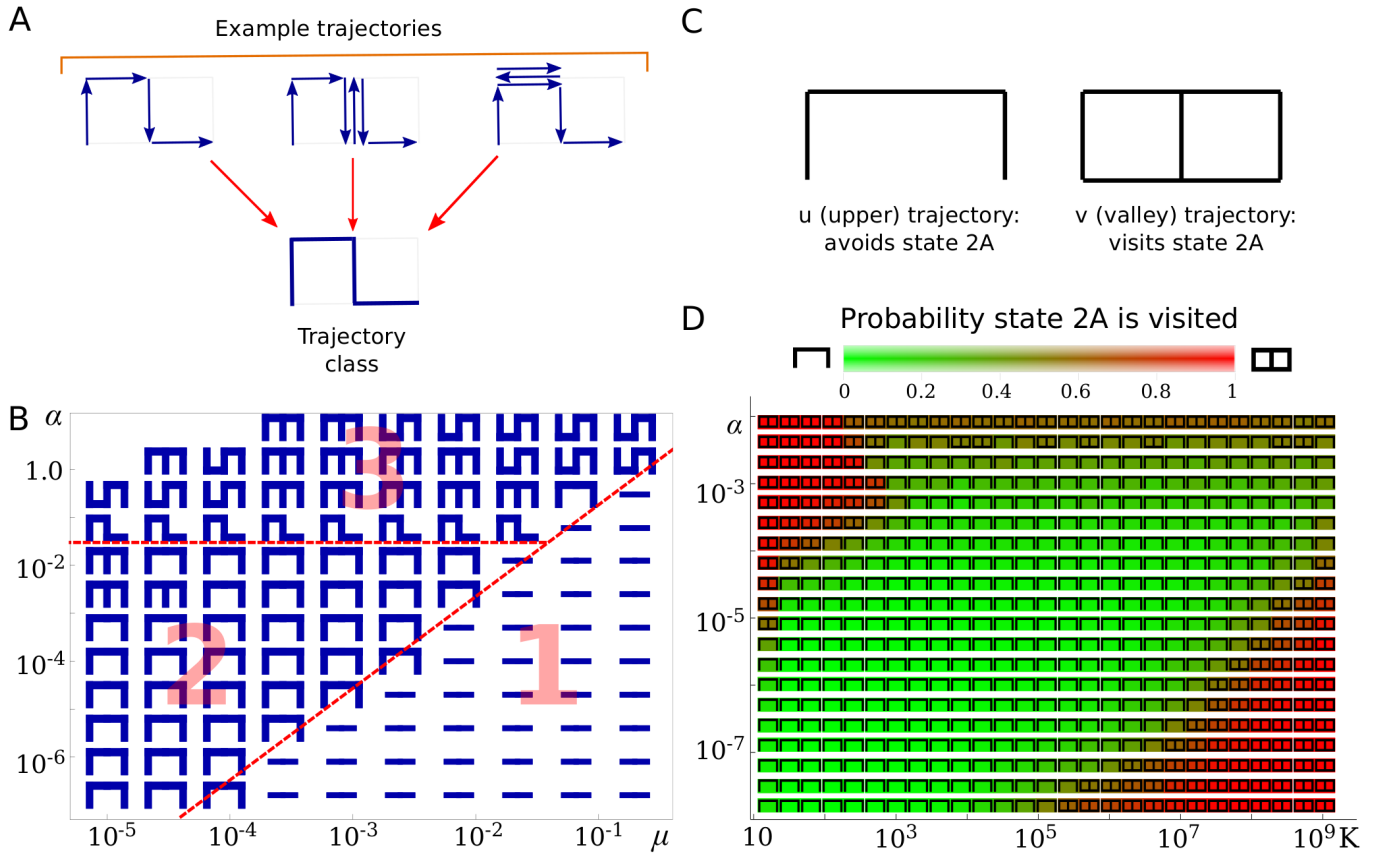
Figure 3: Trajectories of successful cells in genotype/phenotype space. (A) Diagram showing how different trajectories that go through the same set of states are grouped together into the same trajectory class. This class is represented as a symbol in which blue lines correspond to transitions made by the successful cell. (B) The most probable trajectory class as a function of μ and α, for $K = 100$. Three regions labelled $1, 2, 3$ can be distinguished. (C) Trajectories can be further grouped regarding whether they avoid (trajectory "u") or visit (trajectory "v") the fitness valley at state 2A. (D) The most common trajectory group (symbols as in panel C) and the probability of a trajectory visiting state 2A (colours, see the colour bar) as a function of switching rate α and carrying capacity K, for $\mu = 10^{-6}$. In all simulations $\delta = 0.4$ and $d = 0.1$.

Results are robust to small fitness costs of phenotype switching

So far we have assumed that the alternative phenotype B has the same fitness as the wild-type state 1A. We shall now consider the case in which phenotype B is less adapted to the environment than the state 1A. This is a common scenario; for example, phenotypes more resistant to antibiotics often have a lower growth rate than susceptible phenotypes do in the absence of the drug [4]. To model this, we assume that the growth rates $r_{1B} = r_{2B} = r_{3B} = 1 - c$, where $c > 0$ is the fitness cost of switching to phenotype B (see Fig. 1C). A cell that switches from 1A to 1B grows more slowly but switching can still be advantageous if the alternative route by phenotype B is faster than crossing the fitness valley at state 2A. Fig. 4 shows the adaptation time T as a function of switching rate α for different fitness costs. For sufficiently small c the same qualitative behaviour as for $c = 0$ is observed.

The same optimal switching rate range as in Fig. 2 can be seen in Fig. 4 for costs $c \leq 0.05$ of state 1A's fitness, although it reduces in significance as c increases. This is due to the increased time taken by trajectories that avoid state 2A as a result of the reduced growth rates of the phenotype B states. For large c, such as $c = 0.1$ in Fig. 4, there is no optimal switching range but the adaptation time T decreases monotonically with α. Stochastic phenotype switching thus speeds up evolution even in the presence of (moderate) fitness costs.

Switching rate evolves until it reaches the optimal range

So far we have shown that there often exists an optimal range for the switching rate that results in the shortest adaptation time. This range coincides with successful cells avoiding the deleterious state 2A by switching to phenotype B. To see whether a variable switching rate
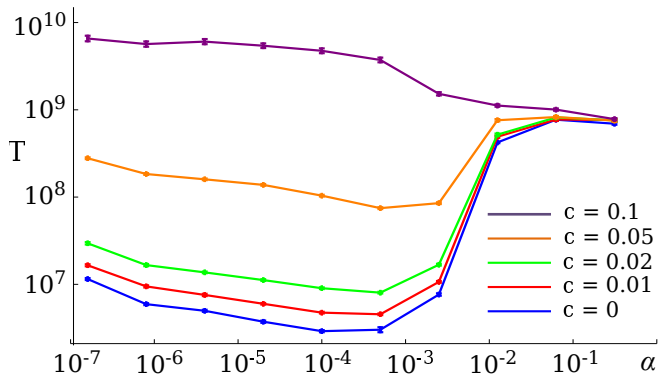
Figure 4: Adaptation time in the presence of fitness costs. The adaptation time T as a function of the switching rate α for different fitness costs $c = \{0, 0.01, 0.02, 0.05, 0.1\}$ of phenotype B and for parameters $K = 100, \mu = 10^{-5}, \delta = 0.9$ and $d = 0.1$.

space at α values below the optimal α range. However, the plot of T in Fig. 5C suggests this will be a slow process. The latter observation can be explained through these cells first evolving a large switching rate, allowing them to switch quickly to phenotype B, before their switching rate decreases again. This picture is corroborated by observing that at low α there is little phenotype switching (purple links in Fig. 5B) and yet the likelihood of mutations in phenotype B across genotype space is large (green links in Fig. 5B). This provides an alternative efficient trajectory for a population to cross the fitness valley. However, such trajectories need at least one extra mutation compared to those that evolve α directly to within the optimal range. Therefore evolution is dominated by the latter trajectories, in particular for low mutation probabilities μ.

DISCUSSION

Phenotypic plasticity has been recently shown to affect Darwinian evolution of animals [15]. Our theoretical research suggests that stochastic phenotypic switching is also able to significantly speed up biological evolution of microbes, perhaps by many orders of magnitude. Switching phenotypes enables cells to evade a fitness valley which may be difficult to cross otherwise, particularly for small populations. Although our model uses only three genotypes, we believe the results presented here are much more general and apply to realistic, multi-dimensional fitness landscapes [36, 43, 46]. Epistasis causes such landscapes to contain many local fitness maxima separated by genotypes with lower growth rate [24]. In a small population, in the absence of phenotype switching, evolution can spend much time in a local maximum before "tunnelling" through one of the fitness valleys. Phenotypic switching opens up a second fitness landscape in which minima and maxima can be located at different genotypes compared to the first landscape. Phenotype switching provides alternative routes between otherwise isolated fitness maxima in both landscapes, and thus enables evolution to proceed faster. This mechanism, similar to the proposed interplay between genetic and epigenetic mutations [23], adds another dimension to the complexity of evolutionary pathways [36, 37, 46] which is relevant for the predictability of evolution [10].

We now discuss how plausible our computational results are for real microbial populations. Although our theory is quite general, we focus on antibiotic resistance evolution. We believe our model represents best a situation in which a population of microbes is exposed to sub-lethal concentrations of an antibiotic. In this scenario, state 1A corresponds to a 'wild-type' genotype whose growth rate is slightly lowered by the presence of the antibiotic, but cells are still able to reproduce. The fitness landscape (Fig. 1C) for phenotype A then

would evolve to be in this optimal range, thus optimising the evolutionary process, we simulated a population of cells that started at 1A with $\alpha \approx 0$ initially. The switching rate was subsequently allowed to evolve: during replication, offspring were assigned a new α value with probability μ, the same as for mutations between genotypes $1 \leftrightarrow 2 \leftrightarrow 3$. The new α was randomly and uniformly selected from a fixed set of possible values ($\alpha = 2.56 \times 10^{-10} \times 5^i$ for integer $i \in [0, 15]$); similar results were obtained when α evolved incrementally (Fig. S5). In this extended model the evolutionary trajectory of successful cells spans three dimensions: two for the genotype, one of which evolves the switching rate, and one for the two phenotypes A and B (Fig. 5A). The end point of the trajectory remains a cell produced at state 3A regardless of its switching rate. We collected trajectories of successful cells in the expanded state space, including states with different α, and calculated transition probabilities between any two states. Fig. 5B shows these probabilities as links of different thickness between the states in state space. We observe that successful trajectories are unlikely to feature genotype mutations in phenotype A (very few red links). Instead, most successful cells first switch to state 1B, mutate to 3B, and switch back to state 3A. The evolved switching rate often falls within the optimal range found in Fig. 2B for the same set of parameters K, μ and d. This is further illustrated in Fig. 5C where we overlay the adaptation time from Fig. 2B with a bar chart that shows the probability that a particular α was selected by evolution.

Fig. 5C reveals two more interesting features. First, we see that cells that have evolved α above the optimal range ($\alpha \geq 0.02$) are unlikely to be successful due to fast transitions between states 2B and 2A (fitness valley). Such transitions effectively lower the fitness of state 2B and create a new fitness valley at this state. Second, we see that a large number of trajectories cross genotype
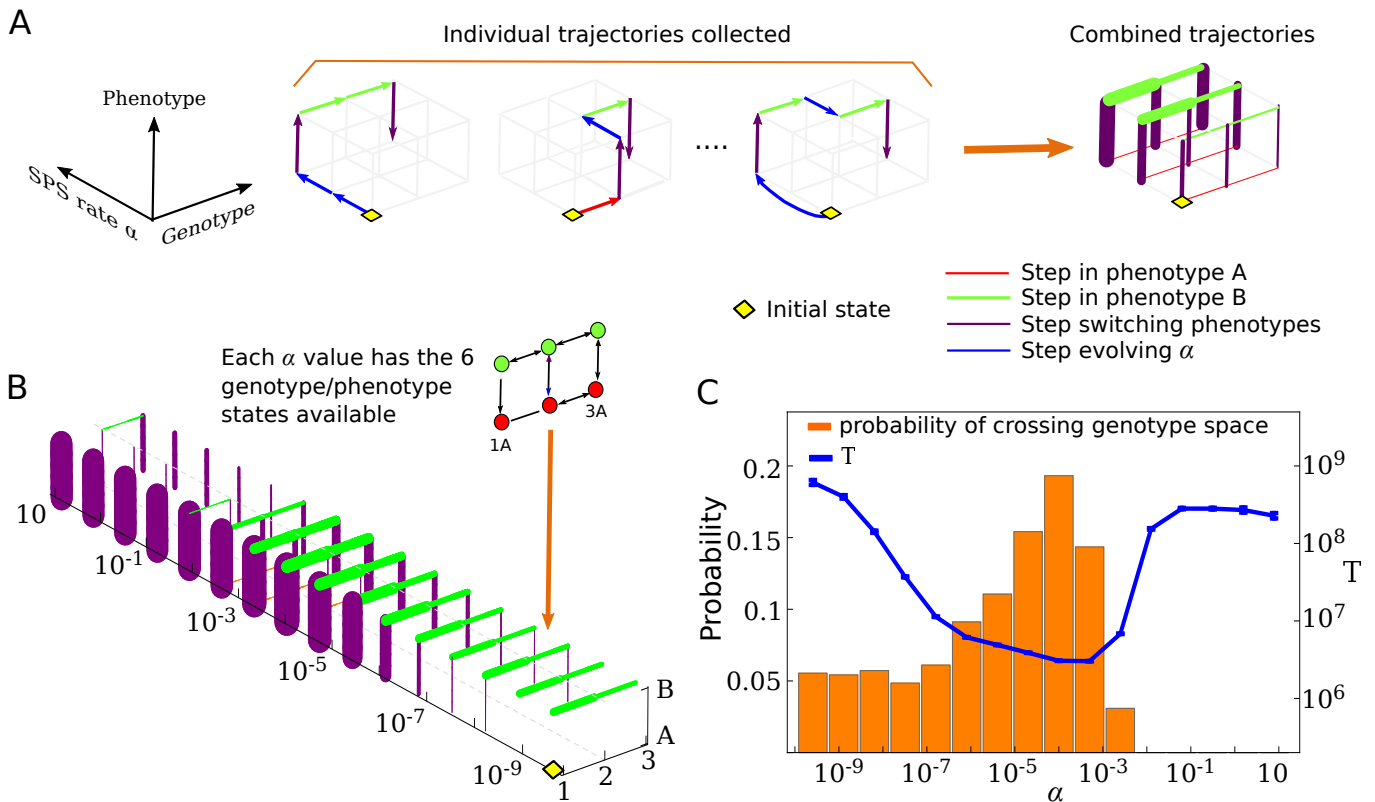
Figure 5: Evolution selects switching rates within the optimal range. (A) Left: examples of evolutionary trajectories in the extended model in which the stochastic phenotype switching (SPS) rate α also evolves. (A) Right: trajectories are used to calculate transition probabilities between the states of the system, which are then represented by the thickness of links connecting the states. Red links correspond to mutations in phenotype A, green links to mutations phenotype B and purple links to switching between phenotypes. (B) Graph of transition probabilities where line thicknesses are proportional to the probability that a successful trajectory will take that step. The parameters are $K = 100, \mu = 10^{-5}, \delta = 0.4$ and $d = 0.1$, the same that were used in Fig. 2B (blue line). The population begins at the wild-type 1A with $\alpha = 2.56 \times 10^{-10}$ and evolves until a cell in state 3A is produced. (C) The probability that genotype space is crossed at a given α in either phenotype A or B. The probability has a maximum where the adaptation time T for fixed α (plot from Fig. 2B superimposed on the same graph) has its minimum.

corresponds to a situation in which an initial mutation further lowers the growth rate by e.g. making the target enzyme less susceptible to the drug but also less efficient [31, 38, 39], while a second mutation compensates for the loss of efficiency and increases the growth rate beyond that of the ancestral strain. In the same example the flat landscape for phenotype B can result from switching to a state with higher gene expression of multidrug efflux pumps or antibiotic-degrading enzymes. This has been observed for a biofilm-forming bacterium *P. aeruginosa* which can switch to a resistant phenotype in the presence of antibiotics [8, 12]. Alternatively, phenotype B can be a persister - although the conventional view is that persistent cells do not replicate, this has been recently challenged [45]. Slower (but non-zero) replication rates can enable persisters to mutate and explore a larger set of genotypes than the non-persister (phenotype A) would be able to do in the absence of switching.

In microbes, the probability of single nucleotide mu-

tation is $10^{-10} - 10^{-9}$ per replication [11, 27] but can be higher (10^{-7}) in mutator strains [27, 40] or in the presence of antibiotics [18]. Assuming $\mu = 10^{-6}$ which would correspond to ~ 10 point mutations in a mutator strain or any loss-of-function mutation in a non-mutator strain (assuming a typical gene length 1kbp [22]), Fig. 3D indicates that phenotype switching would be the preferred mode of evolution for $\alpha < 10^{-4}$ in a population of $K = 10^3$ cells, for $10^{-7} < \alpha < 5 \times 10^{-3}$ for intermediate size $K = 10^6$, and for $\alpha \approx 10^{-3}$ for $K = 10^9$. For lower mutation probabilities μ we expect this behaviour to extend to lower switching rates, consistent with the increasing optimal ranges seen in Fig. 2B, while maintaining approximately the same upper limit of $\alpha \approx 10^{-3}, \ldots, 10^{-2}$. Switching rates encountered in nature can be as large as this [3, 44] and thus we believe that stochastic phenotypic switching can be a significant feature affecting the biological evolution of microbes.

The model we have proposed is a mathematical ideal-

isation. Although it may be possible to test the quantitative predictions of our model in a carefully designed experiment, many of the assumptions we made would have to be relaxed to model naturally occurring populations of bacteria. However, recent work suggests that mechanisms similar to what we described here may be at play in real microbial populations. In particular, Refs. [7, 51] show that the bacterium *E. coli* switches to a filamentous phenotype ("phenotype B") as a result of exposure to sub-inhibitory concentration of the antibiotic ciprofloxacin, and filamentous cells occasionally produce normal-length cells ("phenotype A") resistant to ciprofloxacin. It has been hypothesized that phenotype B provides a "safe niche" giving bacteria enough time to evolve resistance. Although the switching between the phenotypes in this experiment is not entirely random but a response to stress, the observed behaviour is similar to that predicted by our model.

METHODS

Custom written Java programs were used to simulate the model. For smaller K ($K \leq 100$) each cell was assigned a current type (one of the 6 possible states from Fig. 1A) and a sequence of past states beginning at state 1A. A variant of the Gillespie kinetic Monte Carlo algorithm[17] was used to evolve the system. In each time step, a random cell and action (replication, death, or phenotype switch) was selected with a probability proportional to the rate with which that action occurs. A second random number drawn from an exponential distribution, with mean equal to the inverse of the total rate of all possible actions in the system, was used as a measure of time elapsed during that step. For larger $K > 100$ (or when $c > 0$) the evolution of the system was approximately modelled using a type of $\tau-$leaping algorithm [16].

Statistical analysis was performed in Mathematica using a custom-written code. Mean values presented in all plots were obtained by averaging over 10^4 runs, except for Fig. 3D where points were averaged over $10^2 - 10^3$ runs. Error bars are s.e.m.

Acknowledgments

[1] Martin Ackermann. A functional perspective on phenotypic heterogeneity in microorganisms. *Nature Reviews Microbiology*, 13(8):497–508, August 2015.

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

[3] Aileen M Adiciptaningrum, Ian C Blomfield, and Sander J Tans. Direct observation of type 1 fimbrial switching. *EMBO reports*, 10(5):527–532, May 2009.

[4] Nathalie Q. Balaban, Jack Merrin, Remy Chait, Lukasz Kowalik, and Stanislas Leibler. Bacterial Persistence as a Phenotypic Switch. *Science*, 305(5690):1622–1625, September 2004.

[5] Frederick K. Balagaddé, Lingchong You, Carl L. Hansen, Frances H. Arnold, and Stephen R. Quake. Long-Term Monitoring of Bacteria Undergoing Programmed Population Control in a Microchemostat. *Science*, 309(5731):137–140, July 2005.

[6] Hubertus J. E. Beaumont, Jenna Gallie, Christian Kost, Gayle C. Ferguson, and Paul B. Rainey. Experimental evolution of bet hedging. *Nature*, 462(7269):90–93, November 2009.

[7] Julia Bos, Qiucen Zhang, Saurabh Vyawahare, Elizabeth Rogers, Susan M. Rosenberg, and Robert H. Austin. Emergence of antibiotic resistance from multinucleated bacterial filaments. *Proceedings of the National Academy of Sciences*, 112(1):178–183, January 2015.

[8] Elena B. M. Breidenstein, César de la Fuente-Núñez, and Robert E. W. Hancock. Pseudomonas aeruginosa: all roads lead to resistance. *Trends in Microbiology*, 19(8):419–426, August 2011.

[9] Arthur W. Covert, Richard E. Lenski, Claus O. Wilke, and Charles Ofria. Experiments on the role of deleterious mutations as stepping stones in adaptive evolution. *Proceedings of the National Academy of Sciences*, 110(34):E3171–E3178, August 2013.

[10] J. Arjan G. M. de Visser and Joachim Krug. Empirical fitness landscapes and the predictability of evolution. *Nature Reviews Genetics*, 15(7):480–490, July 2014.

[11] J. W. Drake. A constant rate of spontaneous mutation in DNA-based microbes. *Proceedings of the National Academy of Sciences*, 88(16):7160–7164, August 1991.

[12] Eliana Drenkard. Antimicrobial resistance of Pseudomonas aeruginosa biofilms. *Microbes and Infection*, 5(13):1213–1219, November 2003.

[13] S. F Elena, C. O Wilke, C. Ofria, and R. E Lenski. Effects of population size and mutation rate on the evolution of mutational robustness. *Evolution*, 61(3):666–674, 2007.

[14] J. Franke, A. Klözer, J. A.G.M de Visser, and J. Krug. Evolutionary accessibility of mutational pathways. *PLoS computational biology*, 7(8):e1002134, 2011.

[15] Cameron K. Ghalambor, Kim L. Hoke, Emily W. Ruell, Eva K. Fischer, David N. Reznick, and Kimberly a. Hughes. Non-adaptive plasticity potentiates rapid adaptive evolution of gene expression in nature. *Nature*, September 2015.

[16] D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *The Journal of Chemical Physics*, 115(4):1716–1733, 2001.

[17] Daniel T. Gillespie. Exact stochastic simulation of cou-

pled chemical reactions. *The journal of physical chemistry*, 81(25):2340–2361, 1977.

[18] Stephen H. Gillespie, Shreya Basu, Anne L. Dickens, Denise M. O'Sullivan, and Timothy D. McHugh. Effect of subinhibitory concentrations of ciprofloxacin on Mycobacterium fortuitum mutation rates. *Journal of Antimicrobial Chemotherapy*, 56(2):344–348, August 2005.

[19] H. Hasman, M.A. Schembri, and P. Klemm. Antigen 43 and type 1 fimbriae determine colony morphology of Escherichia coli K-12. *Journal of bacteriology*, 182(4):1089, 2000.

[20] Daniel B. Kearns and Richard Losick. Cell population heterogeneity during growth of Bacillus subtilis. *Genes & Development*, 19(24):3083–3094, December 2005.

[21] Kristel Kemper, Pauline L. de Goeje, Daniel S. Peeper, and Renée van Amerongen. Phenotype switching: tumor cell plasticity as a resistance mechanism and target for therapy. *Cancer Research*, 74(21):5937–5941, November 2014.

[22] Michael J. Kerner, Dean J. Naylor, Yasushi Ishihama, Tobias Maier, Hung-Chun Chang, Anna P. Stines, Costa Georgopoulos, Dmitrij Frishman, Manajit Hayer-Hartl, Matthias Mann, and F. Ulrich Hartl. Proteome-wide analysis of chaperonin-dependent protein folding in Escherichia coli. *Cell*, 122(2):209–220, July 2005.

[23] Filippos D. Klironomos, Johannes Berg, and Sinéad Collins. How epigenetic mutations can affect genetic evolution: Model and mechanism. *BioEssays*, 35(6):571–578, 2013.

[24] R. Korona, C.H. Nakatsu, L.J. Forney, and R.E. Lenski. Evidence for multiple adaptive peaks from populations of bacteria evolving in a structured habitat. *Proceedings of the National Academy of Sciences of the United States of America*, 91(19):9037–41, September 1994.

[25] Edo Kussell, Roy Kishony, Nathalie Q. Balaban, and Stanislas Leibler. Bacterial Persistence A Model of Survival in Changing Environments. *Genetics*, 169(4):1807–1814, April 2005.

[26] Edo Kussell and Stanislas Leibler. Phenotypic Diversity, Population Growth, and Information in Fluctuating Environments. *Science*, 309(5743):2075–2078, September 2005.

[27] Heewook Lee, Ellen Popodi, Haixu Tang, and Patricia L. Foster. Rate and molecular spectrum of spontaneous mutations in the bacterium Escherichia coli as determined by whole-genome sequencing. *Proceedings of the National Academy of Sciences of the United States of America*, 109(41):E2774–E2783, October 2012.

[28] Bruce R. Levin and Daniel E. Rozen. Non-inherited antibiotic resistance. *Nature Reviews Microbiology*, 4(7):556–562, 2006.

[29] Hédia Maamar and David Dubnau. Bistability in the Bacillus subtilis K-state (competence) system requires a positive feedback loop: Bistability in B. subtilis competence. *Molecular Microbiology*, 56(3):615–624, March 2005.

[30] Sophie Maisnier-Patin, Otto G. Berg, Lars Liljas, and Dan I. Andersson. Compensatory adaptation to the deleterious effect of antibiotic resistance in Salmonella typhimurium. *Molecular Microbiology*, 46(2):355–366, October 2002.

[31] Linda L Marcusson, Niels Frimodt-Moller, and Diarmaid Hughes. Interplay in the selection of fluoroquinolone resistance and bacterial fitness. *PLoS pathogens*, 5(8):e1000541, August 2009.

[32] O. C Martin, L. Peliti, et al. Population size effects in evolutionary dynamics on neutral networks and toy landscapes. *Journal of Statistical Mechanics: Theory and Experiment*, 2007:P05011, 2007.

[33] Oliver G. McDonald, Hao Wu, Winston Timp, Akiko Doi, and Andrew P. Feinberg. Genome-scale epigenetic reprogramming during epithelial-to-mesenchymal transition. *Nature Structural & Molecular Biology*, 18(8):867–874, August 2011.

[34] Martin A. Nowak. *Evolutionary Dynamics*. Harvard University Press, September 2006.

[35] Ertugrul M. Ozbudak, Mukund Thattai, Han N. Lim, Boris I. Shraiman, and Alexander van Oudenaarden. Multistability in the lactose utilization network of Escherichia coli. *Nature*, 427(6976):737–740, February 2004.

[36] Adam C. Palmer, Erdal Toprak, Michael Baym, Seungsoo Kim, Adrian Veres, Shimon Bershtein, and Roy Kishony. Delayed commitment to evolutionary fate in antibiotic resistance fitness landscapes. *Nature Communications*, 6, June 2015.

[37] Frank J. Poelwijk, Daniel J. Kiviet, and Sander J. Tans. Evolutionary Potential of a Duplicated Repressor-Operator Pair: Simulating Pathways Using Mutation Data. *PLoS Computational Biology*, 2(5):e58, 2006.

[38] S J Schrag and V Perrot. Reducing antibiotic resistance. *Nature*, 381(6578):120–121, May 1996.

[39] Stephanie J Schrag, Véronique Perrot, and Bruce R Levin. Adaptation to the fitness costs of antibiotic resistance in Escherichia coli. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 264(1386):1287–1291, September 1997.

[40] P D Sniegowski, P J Gerrish, and R E Lenski. Evolution of high mutation rates in experimental populations of E. coli. *Nature*, 387(6634):703–705, June 1997.

[41] Ana Solopova, Jordi van Gestel, Franz J. Weissing, Herwig Bachmann, Bas Teusink, Jan Kok, and Oscar P. Kuipers. Bet-hedging during bacterial diauxic shift. *Proceedings of the National Academy of Sciences*, 111(20):7427–7432, 2014.

[42] Ivan G. Szendro, Jasper Franke, J. Arjan G. M. de Visser, and Joachim Krug. Predictability of evolution depends nonmonotonically on population size. *Proceedings of the National Academy of Sciences*, 110(2):571–576, January 2013.

[43] Ivan G. Szendro, Martijn F. Schenk, Jasper Franke, Joachim Krug, and J. Arjan G. M. de Visser. Quantitative analyses of empirical fitness landscapes. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):P01005, January 2013.

[44] Marjan W van der Woude and Ian R Henderson. Regulation and function of Ag43 (flu). *Annual Review of Microbiology*, 62:153–169, 2008.

[45] Y. Wakamoto, N. Dhar, R. Chait, K. Schneider, F. Signorino-Gelo, S. Leibler, and J. D. McKinney. Dynamic Persistence of Antibiotic-Stressed Mycobacteria. *Science*, 339(6115):91–95, January 2013.

[46] D. M. Weinreich. Darwinian Evolution Can Follow Only Very Few Mutational Paths to Fitter Proteins. *Science*, 312(5770):111–114, April 2006.

[47] D.M. Weinreich and L. Chao. Rapid evolutionary escape by large populations from local fitness peaks is likely in nature. *Evolution*, 59(6):1175–1182, 2005.

[48] D.B. Weissman, M.M. Desai, D.S. Fisher, and M.W.

Feldman. The rate at which asexual populations cross fitness valleys. *Theoretical population biology*, 75(4):286–300, 2009.

[49] Zhizhong Yao, Rebecca M. Davis, Roy Kishony, Daniel Kahne, and Natividad Ruiz. Regulation of cell size in response to nutrient availability by fatty acid biosynthesis in Escherichia coli. *Proceedings of the National Academy of Sciences*, 109(38):E2561–E2568, September 2012.

[50] Stefano Zapperi and Caterina A. M. La Porta. Do cancer cells undergo phenotypic switching? The case for imperfect cancer stem cell markers. *Scientific Reports*, 2, June 2012.

[51] Q. Zhang, G. Lambert, D. Liao, H. Kim, K. Robin, C.-k. Tung, N. Pourmand, and R. H. Austin. Acceleration of Emergence of Bacterial Antibiotic Resistance in Connected Microenvironments. *Science*, 333(6050):1764–1767, September 2011.

SUPPLEMENTARY INFORMATION

Phenotypic switching will not always speed up evolution

In Fig. 2A we compared the adaptation time T with and without phenotypic switching for various parameters. In all cases switching decreased T. However, if the mutation probability μ is large enough, the pathway that goes directly through the valley genotype 2A may take less time than any pathway involving the alternative phenotype B. Figure S1 shows T as a function of mutation probability μ for the case with and without switching. For large enough μ, switching makes no observable difference to the adaptation time as successful cells do not use an alternative pathway. This is the same behaviour observed in region 1 of Fig. 3B.
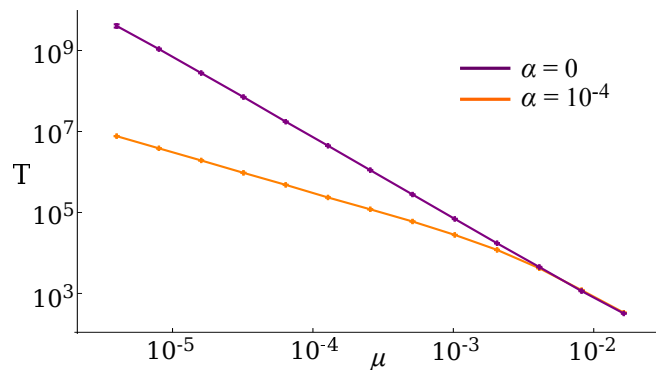

Figure S1: Phenotypic switching does not speed up evolution for large mutation probabilities μ. Shown is the adaptation time T from computer simulations as a function of μ with and without phenotype switching ($\alpha = 10^{-4}$ and $\alpha = 0$ respectively), for $K = 100, \delta = 0.4$ and $d = 0.1$.

Adaptation time decreases monotonously with α in the absence of switching between states 2A and 2B

We identified that region 3 in Fig. 3B appears due to the existence of the step between states 2A and 2B. For large α values a trajectory via phenotype B is likely to involve switching to the deleterious state 2A. This reduces the average fitness of the population and causes the adaptation time to increase with α (Fig. 2B). To demonstrate this consider the model without switching between states 2A and 2B (Fig. S2A). The adaptation time T as a function of switching rate α decreases monotonically in this modified model (Fig. S2B).
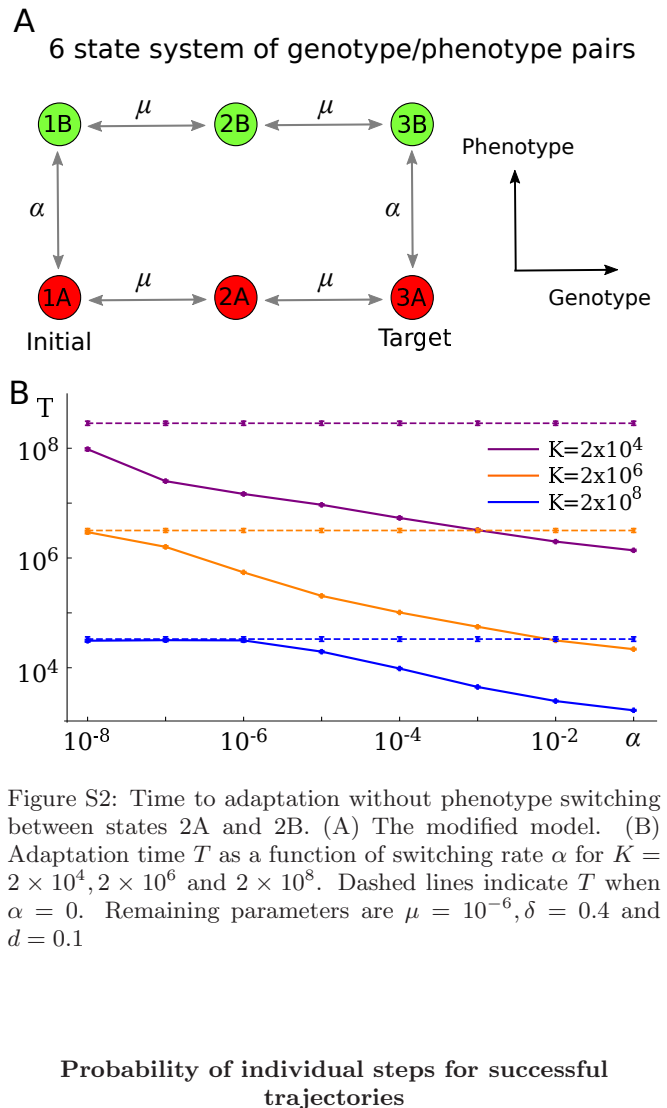

Figure S2: Time to adaptation without phenotype switching between states 2A and 2B. (A) The modified model. (B) Adaptation time T as a function of switching rate α for $K = 2 \times 10^4, 2 \times 10^6$ and 2×10^8. Dashed lines indicate T when $\alpha = 0$. Remaining parameters are $\mu = 10^{-6}, \delta = 0.4$ and $d = 0.1$

Probability of individual steps for successful trajectories

In Fig. 3B we showed the most common trajectory classes for successful cells as a function of (μ, α). A more detailed picture can be obtained by looking at the probability that each step occurs in a successful cell's trajectory. For each pair of (μ, α) we create a diagram of transitions between states (i.e. 1A,1B,...,3A,3B) in which
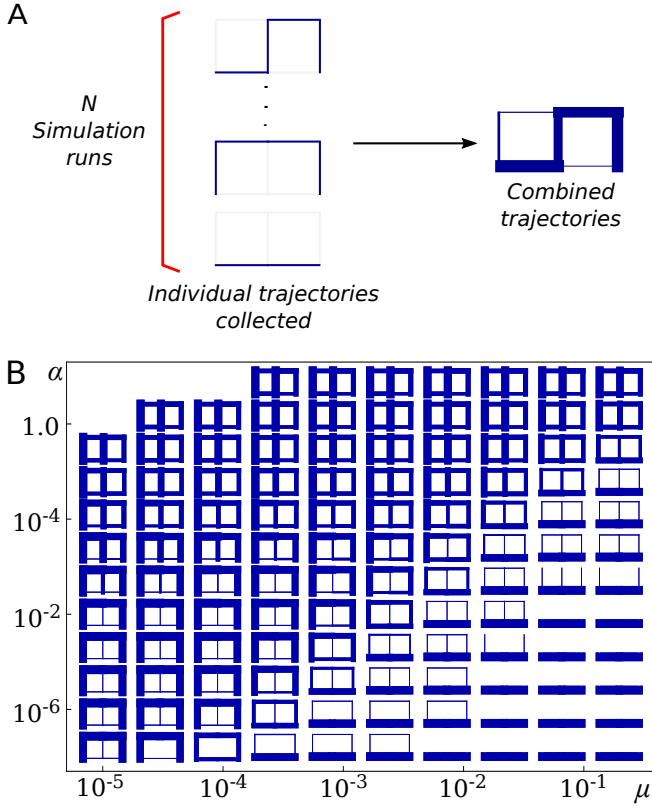
A

B

Figure S3: The probability of each step being taken by successful cells. (A) Constructing images that display the probability that each step is taken. By collecting many successful trajectory classes we can compile them into a single image that consists of links of different thicknesses, where the thickness of each link is proportional to the probability that that step is taken by a successful cell. (B) The probability each step is taken as a function of μ and α. Remaining parameters are $K = 100, \delta = 0.4$ and $d = 0.1$.

the thickness of each link corresponds to the probability that that step is taken (Fig. S3A). Figure S3B shows a graph in which diagrams for different pairs of (μ, α) have been put together. This graph is very similar to Fig. 3B. This confirms that the most-common trajectories plotted in Fig. 3B are good representatives of the "typical" trajectories of successful cells.

Adaptation time as function of carrying capacity

Using the same simulation data as in Fig. 3D but plotting it differently we can explore how the adaptation time T depends on the carrying capacity K. Figure S4 shows T for different values of the switching rate α. A

non-monotonic behaviour can be seen: T is maximal for intermediate K and falls off for both small and large carrying capacities.
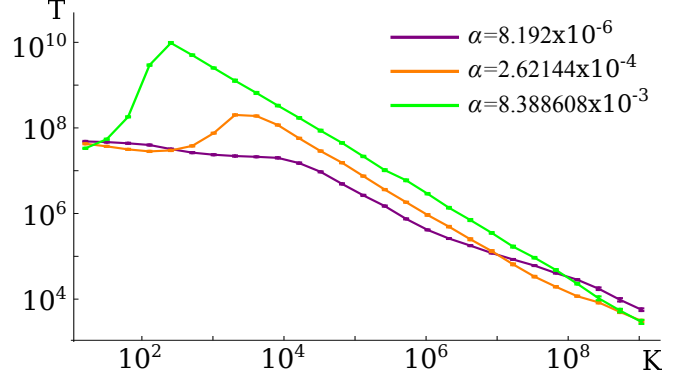

Figure S4: Adaptation time as a function of carrying capacity K for different values of α. Remaining parameters are $\mu = 10^{-6}, \delta = 0.4$ and $d = 0.1$

Evolving the switching rate in fixed increments

Fig. 5 showed simulation results in which newly created cells could mutate their switching rate α by drawing a new value from a discrete set of allowed switching rates. We observed successful cell trajectories to consist largely of those that evolved switching rates to within the optimal range identified in Fig. 2B. Here we consider the same set of allowed α values but, upon mutation of α, the new value is chosen as one of the two closest values to the switching rate of the ancestor cell. This corresponds to α evolving in fixed multiplicative steps, with increase or decrease of α being equally likely. The rest of the algorithm is the same as before.

We again collected trajectories of successful cells in the expanded state space, including states with different α, and calculated transition probabilities between any two states. Fig. S5B shows these probabilities as links of different thickness between the states in state space. We can clearly see some new behaviour here. First, successful trajectories are more likely to feature genotype mutations in phenotype A (more red lines) than in Fig. 5B. Second, the probability of crossing the genotype space at large α (in either phenotype) is much smaller due to the increased number of steps it takes a cell to evolve such large α values. Otherwise, Fig. S5 is similar to Fig. 5. This shows that the details of how mutations of α are implemented do not have a significant effect upon the results of our model.
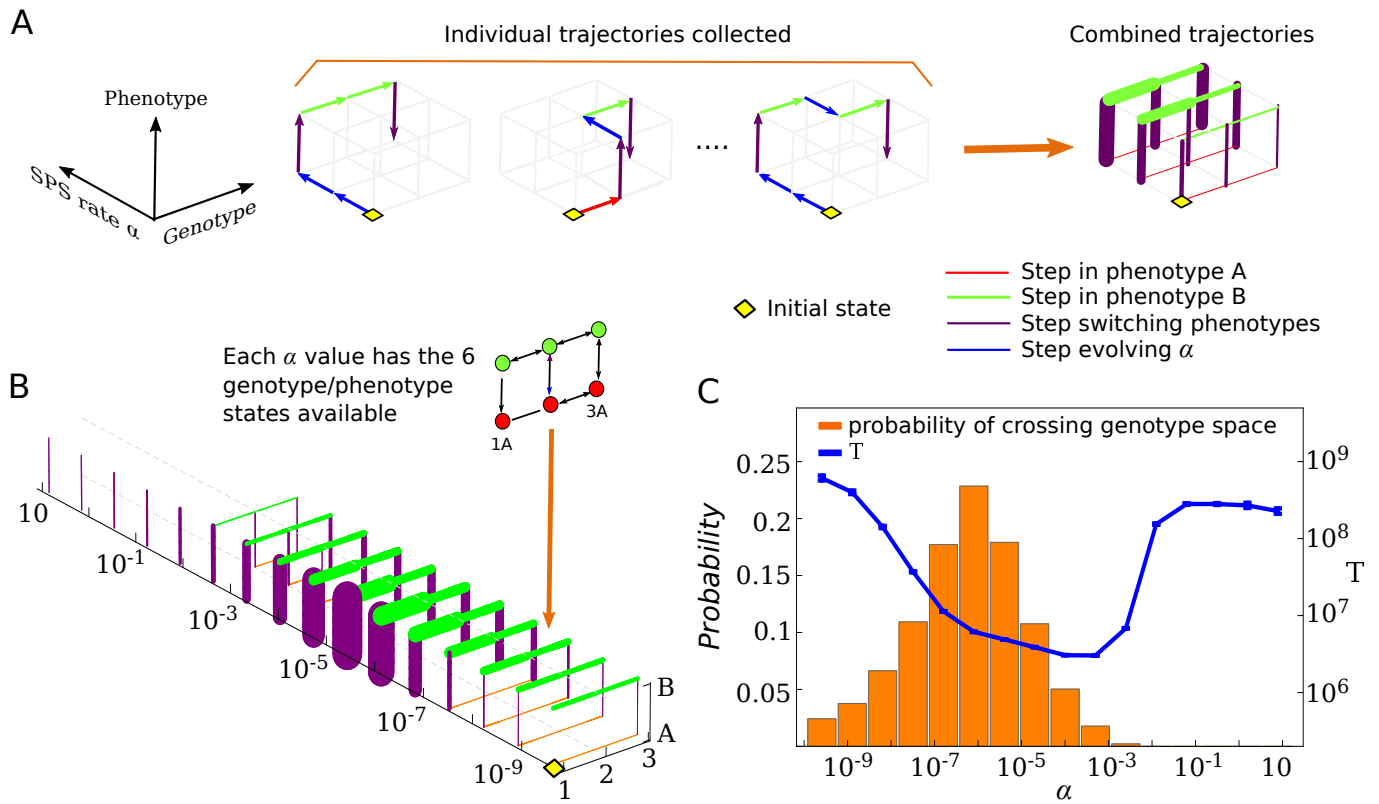
Figure S5: Evolutionary trajectories when α evolves in fixed multiplicative steps. (A) Left: examples of evolutionary trajectories. Upon mutation of α a neighbouring value is selected (with increase or decrease equally likely in the set of allowed values $\alpha = 2.56 \times 10^{-10} \times 5^i$ for $i \in [0, 15]$). (A) Right: trajectories are used to calculate transition probabilities between the states of the system, which are then represented by the thickness of links connecting the states. Red, green and purple links identify steps in phenotype A, phenotype B and between phenotypes respectively. (B) Graph of transition probabilities where line thicknesses are proportional to the probability that a successful trajectory involves that step. The parameters are $K = 100, \mu = 10^{-5}, \delta = 0.4$ and $d = 0.1$, the same that were used in Fig. 2B (blue line). The population begins at the wild-type 1A with $\alpha = 2.56 \times 10^{-10}$ and evolves until a cell in state 3A is produced. (C) The probability that genotype space is crossed at a given α in either phenotype A or B. Superimposed on this plot is the adaptation time T as a function of α obtained in a simulation with fixed α (Fig. 2B).