

# Exploiting evolutionary non-commutativity to prevent the emergence of bacterial antibiotic resistance

Daniel Nichol<sup>1,5</sup>, Peter Jeavons<sup>1</sup>, Alexander G. Fletcher<sup>2</sup>, Robert A. Bonomo<sup>3</sup>, Philip K. Maini<sup>2</sup>, Jerome L. Paul<sup>4</sup>, Robert A. Gatenby<sup>5</sup>, Alexander R.A. Anderson<sup>5</sup> & Jacob G. Scott<sup>2,5</sup>

<sup>1</sup> Department of Computer Science, University of Oxford, Oxford, UK

<sup>2</sup> Wolfson Centre for Mathematical Biology, Mathematical Institute, University of Oxford, Oxford, UK

<sup>3</sup> Department of Medicine, Louis Stokes Department of Veterans Affairs Hospital, Cleveland, OH, USA

<sup>4</sup> School of Computing Sciences and Informatics, University of Cincinnati, Cincinnati, OH, USA

<sup>5</sup> Department of Integrated Mathematical Oncology, H. Lee Moffitt Cancer Center and Research Institute, Tampa, FL, USA

Correspondence:

Jacob G Scott, e-mail: [jacob.g.scott@gmail.com](mailto:jacob.g.scott@gmail.com) and

Daniel Nichol, e-mail: [dan.nichol@cs.ox.ac.uk](mailto:dan.nichol@cs.ox.ac.uk)

Keywords: theoretical biology, antibiotic resistance, bacteria, mathematical model, evolution, fitness landscape, strong selection weak mutation

# Abstract

The increasing rate of antibiotic resistance and slowing discovery of novel antibiotic treatments presents a growing threat to public health. Here, we develop a Markov Chain model of evolution in asexually reproducing populations which is an encoding of the Strong Selection Weak Mutation model of evolution on fitness landscapes. This model associates the global properties of the fitness landscape with the algebraic properties of the Markov Chain transition matrix and allows us to derive general results on the non-commutativity and irreversibility of natural selection as well as antibiotic cycling strategies. Utilizing this formalism, we analyse 15 empirical fitness landscapes of *E. coli* under selection by different beta-lactam antibiotics and demonstrate that the emergence of resistance to a given antibiotic can be both hindered and promoted by different sequences of drug application. Further, we derive optimal drug application sequences with which we can probabilistically ‘steer’ the population through genotype space to avoid the emergence of resistance. This suggests a new strategy in the war against antibiotic–therapy–resistant organisms: drug sequencing to shepherd evolution through genotype space to states from which resistance cannot emerge and by which to maximize the chance of successful therapy.

**Background:** The increasing rate of antibiotic resistance and slowing discovery of novel antibiotic treatments presents a growing threat to public health. Previous studies of bacterial evolutionary dynamics have shown that populations evolving on fitness landscapes follow predictable paths. In this article, we develop a general mathematical model of evolution and hypothesise that it can be used to understand, and avoid, the emergence of antibiotic resistance.

**Methods and Findings:** We develop a Markov Chain model of evolution in asexually reproducing populations which is an encoding of the Strong Selection Weak Mutation model of evolution on fitness landscapes. This model associates the global properties of the fitness landscape with the algebraic properties of the Markov Chain transition matrix and allows us to derive general results on the non-commutativity and irreversibility of natural selection as well as antibiotic cycling strategies. Utilizing this formalism, we analyse 15 empirical fitness landscapes of *E. coli* under selection by different  $\beta$ -lactam antibiotics and demonstrate that the emergence of resistance to a given antibiotic can be both hindered and promoted by different sequences of drug application. We show that resistance to a given antibiotic is promoted in 61.4%, 68.6% and 70.3% of possible orderings of single, pair or triple prior drug administrations, respectively. Further, we derive optimal drug application sequences with which we can probabilistically ‘steer’ the population through genotype space to avoid the emergence of resistance.

**Conclusions:** Our model provides generalisable results of interest to theorists studying evolution as well as providing specific, testable predictions for experimentalists to validate our methods. Further, these results suggest a new strategy in the war against antibiotic-therapy-resistant organisms: drug sequencing to shepherd an evolving population through genotype space to states from which resistance cannot emerge and from which we can maximize the likelihood of successful therapy using existing drugs. While our results are derived using a specific class of antibiotics, the method is generalisable to other situations, including the emergence of resistance to targeted therapy in cancer and how species could change secondary to changing climate or geographical movement.

# 1 Introduction

2 Resistance to antibiotic treatments within bacterial pathogens poses an increasing threat to pub-  
 3 lic health, which coupled with the slowing discovery of novel antibiotics, could soon reach crisis  
 4 point [Spellberg et al., 2013, French, 2010]. Novel classes of antibiotics discovered since 1987 are  
 5 few in number [Silver, 2011] . Thus, it is becoming ever clearer that if we are to combat highly  
 6 resistant bacterial infections, then we must find new ways to prevent resistance and new applica-  
 7 tions of existing antibiotics to these pathogens. Indeed, public health efforts have attempted to  
 8 stem the emergence of resistance by reducing unnecessary prescription of antibiotics [Shlaes et al.,  
 9 1997, Bartlett, 2011, Leuthner and Doern, 2013] and stopping the addition of sub-therapeutic an-  
 10 tibiotics in livestock feed [Mathew et al., 2007]. However, such policies require global adoption to  
 11 be truly effective [Moody et al., 2012], which they have not yet achieved, and is likely infeasible.  
 12 Recently, there have been efforts to explore how existing antibiotics can be used in new ways to pro-  
 13 vide effective treatments for resistant pathogens, for example through combination therapy [Chait  
 14 et al., 2007] or through cycling strategies [Goulart et al., 2013, Pena-Miller and Beardmore, 2010,  
 15 Bergstrom et al., 2004] – long term hospital-scale treatment protocols that cycle which antibiotics  
 16 are prescribed over timescales of weeks, months or years. However, our understanding of the mecha-  
 17 nisms underlying these strategies remains limited. Indeed, from a mathematical model Pena-Miller  
 18 and Beardmore [2010] predict that cycling strategies have the potential to perform both better and  
 19 worse than mixed strategies and hence tools to find optimal strategies are needed.

20 In order to understand how to minimize the emergence of resistant pathogens, and to decide  
 21 how best to treat them, we must first understand how their evolution is driven by the selective  
 22 pressures of different antibiotic drugs — a fundamental problem of biology [de Visser and Krug,  
 23 2014]. In particular, if we understand which traits are likely to be selected for by which treatments,  
 24 then we may be able to avoid selecting for those traits which confer resistance. Recent insights  
 25 into the evolutionary process have yielded some actionable information. Specifically, Weinreich  
 26 et al. [2005, 2006] showed that if the genome of a pathogen exhibits sign epistasis, where a given  
 27 mutation is beneficial on some genetic backgrounds and deleterious on others, then there can exist  
 28 inaccessible evolutionary trajectories. Further, Tan et al. [2011] studied the evolutionary trajec-  
 29 tories of *Escherichia coli* under different antibiotics and found that adaptive mutations gained under  
 30 one antibiotic are often irreversible when a second is applied. These irreversible trajectories can  
 31 occur when resistance conferring mutations for one environment carry a cost in another which can  
 32 be mitigated by other compensatory mutations [zur Wiesch et al., 2010, Tanaka and Valckenborgh,  
 33 2011]. These findings lead us to hypothesize that one antibiotic could be used to irreversibly steer  
 34 the evolution of a population of pathogens to a genotype (or combination of genotypes) which

is sensitive to a second antibiotic but also from which it is unlikely to acquire resistance to that antibiotic. This hypothesis was partly verified by the work of Imamovic and Sommer [2013] who demonstrated that evolving *E. coli* to become resistant to certain antibiotics can increase sensitivity to others. However, this work is limited in two ways. Firstly, in that the experiments do not exhaustively consider all evolutionary trajectories but instead highlight only those that arose in a small number of replicates and secondly, they do not consider how evolution will proceed once the second drug is applied and whether resistance can then emerge.

In this paper we present a Markov Chain model of evolution which we use to derive general results on the non-commutativity of evolution and requirements for cycling strategies. Then, using previously measured landscapes for 15  $\beta$ -lactam anti-biotics, we illustrate that selective pressures are non-commutative, and that the emergence of resistance can be both hindered and promoted by different orderings of these pressures. These findings suggest new treatment strategies which use rational orderings of collections of drugs to shepherd evolution through genotype space to a configuration which is sensitive to treatment, as in the work of Imamovic and Sommer [2013], but also from which resistance cannot emerge.

## Evolution on Fitness Landscapes

We begin with the concept of a fitness landscape introduced by Wright [1932] and used by Weinreich et al. [2005] and Tan et al. [2011] to study evolutionary trajectories in asexually reproducing populations. We represent the genotype of an organism by a bit string of length  $N$  and model mutation as the process of flipping a single bit within this string. This model of mutation only accounts for point mutations and ignores the possibility of other biologically relevant mutations such as gene insertions, gene deletions and large structural changes to the genotype. This gives a set of  $2^N$  possible genotypes in which individuals of a given genotype, say  $g$ , can give rise to mutated offspring which take genotypes given by one of the  $N$  mutational neighbors of  $g$  — precisely those genotypes  $g'$  for which the Hamming distance [Hamming, 1950],  $\text{Ham}(g, g')$ , from  $g$  is 1. As such, our genotype space can be represented by an undirected  $N$ -dimensional (hyper-)cube graph with vertices in  $\{0, 1\}^N$  representing genotypes and edges connecting mutational neighbors (Figure 1a).

We define a selective pressure on our graph that drives evolution, for example through an environmental change or drug application, as a fitness function

$$f : \{0, 1\}^N \rightarrow \mathbb{R}^{\geq 0}. \quad (1)$$

This fitness function represents a genotype-phenotype map in the simplest sense — assigning to each

genotype a single real-valued fitness. Gillespie [1983, 1984] showed that if the mutation rate  $u$  and population size  $M$  of a population satisfy  $Mu \gg 1 \gg Mu^2$ , and if we assume that each mutation is either beneficial or deleterious, then each beneficial mutation in the population will either reach fixation or become extinct before a new mutation occurs. Further, selection will be sufficiently strong that any deleterious mutation will become extinct with high probability and hence we may assume that this always occurs. In the case that  $Mu^2 \approx 1$  stochastic tunneling [Iwasa et al., 2004] through double mutations can occur and we cannot ignore deleterious mutations. Assuming  $Mu \gg 1 \gg Mu^2$ , then after each mutation the population will stabilize to consist entirely of individuals with the same genotype and this genotype will be eventually replaced by a fitter neighboring genotype whenever one exists. This observation gives rise to the Strong Selection Weak Mutation (SSWM) model, which models a population as occupying a single vertex on a directed graph on the set of  $2^N$  possible genotypes,  $\{0, 1\}^N$ , in which there exists an edge from vertex  $a$  to a neighboring vertex  $b$  if, and only if,  $f(b) > f(a)$  (see Figure 1b and 1c). This population undergoes a stochastic walk in which the population moves to an adjacent fitter genotype with some probability. Several ‘move rules’ have been proposed which can be used to select an adjacent fitter neighbor during this stochastic walk [Orr, 2005] and which of these move rules is most accurate depends on the population size [de Visser and Krug, 2014]. Common move rules include selecting the fittest neighbor [Kauffman and Levin, 1987, Fontana et al., 1993], selecting amongst fitter neighbors at random [Macken and Perelson, 1989, Macken et al., 1991, Flyvbjerg and Lautrup, 1992] or selecting fitter neighbors with probability proportional to the fitness increase conferred [Gillespie, 1983, 1984, 1991]. We encapsulate each of these variants of the SSWM model within our model.

## A Markov Model of Evolution

The SSWM model of evolution reduces the evolutionary process to a random walk on a directed graph and hence can be modeled by a Markov Chain [Grinstead and Snell, 1998]. For a fitness function  $f : \{0, 1\}^N \rightarrow \mathbb{R}^{\geq 0}$  we can define a transition matrix  $P = [\mathbb{P}(i \rightarrow j)]_{i,j \in \{0,1\}^N}$  for a time-homogeneous absorbing Markov Chain by setting, for  $i \neq j$ ,

$$\mathbb{P}(i \rightarrow j) = \begin{cases} \frac{(f(j)-f(i))^r}{\sum_{\substack{g \in \{0,1\}^N, \text{Ham}(i,g)=1 \\ f(g)-f(i)>0}} (f(g)-f(i))^r} & \text{if } f(j) > f(i) \text{ and Ham}(i, j) = 1 \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$




Figure 1: (a) The space of genotypes comprising bit strings of length  $N = 3$ . The vertices represent genotypes, and edges connect those genotypes which are mutational neighbors. (b) An example fitness landscape. (c) The directed evolutionary graph according to the landscape in (b) where the vertices still represent genotypes but are labeled by the associated fitness for clarity. The directed graph edges are determined by the fitness function and represent those mutations which can fix in a population (those which confer a fitness increase). (d) The Markov Chain constructed for the same landscape according to equations (2) and (3) with  $r = 0$ .

91 and

$$\mathbb{P}(i \rightarrow i) = \begin{cases} 1 & \text{if } i \text{ has no fitter one-step mutational neighbors} \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

92 for each  $i$  (see Figure 1d). Here the parameter  $r \geq 0$  determines the extent to which the fitness  
 93 increase of a mutation affects its likelihood of determining the next population genotype. In the case  
 94  $r = 0$ , we have the random move SSWM model (as in Macken and Perelson [1989], Macken et al.  
 95 [1991], Flyvbjerg and Lautrup [1992]), in the limit  $r \rightarrow \infty$  we have the steepest gradient ascent  
 96 SSWM model (as in Kauffman and Levin [1987], Fontana et al. [1993]), and for  $r = 1$  we have

probability proportional to fitness increase (as in Gillespie [1983, 1984, 1991]). This model differs from the Markov model used by Sella and Hirsh [2005] to study the neutral theory of evolution as we do not allow deleterious mutations to fix in the population.

Using this Markov Chain we can explore the possible evolutionary trajectories of a population on a given fitness landscape  $f$ . We next define a collection of population row vectors  $\mu^{(t)}$  for each  $t \in \mathbb{N}$ , where  $\mu^{(t)}$  has length  $2^N$  and  $k^{\text{th}}$  component which gives the probability that the population has the  $k^{\text{th}}$  genotype at time  $t$  (where the genotypes are ordered numerically according to their binary value). These time steps  $t$  are an abstraction which discretely measure events of beneficial mutations occurring and fixing in the population. As such, the actual time between steps  $t$  and  $t+1$  is not constant but may be considered drawn from a distribution parameterized by the mutation rate, reproductive rate and the number of beneficial mutations that can occur. This distribution could, for example, be determined by a Moran [Moran et al., 1962] or Wright–Fisher [Wright, 1932, Fisher, 1958] process depending on how we choose to interpret the fitness values given by  $f$ . If the population has a genotype corresponding to a local optimum of the fitness landscape at time  $t$  then there are no beneficial mutations that can occur and this definition of a time step is not well defined. In this case there can be no more changes to the population under the SSWM assumptions and for mathematical convenience we define the probability of a local optimum population genotype remaining unchanged as one in equation 3 to ensure our model is a Markov Chain. In this case the step  $t$  to  $t+1$  can be chosen to take some fixed arbitrary time.

The distribution of a population at time  $t$  is related to its initial distribution,  $\mu^{(0)}$ , by

$$\mu^{(t)} = \mu^{(0)} P^t. \quad (4)$$

Since the Markov Chain is absorbing we know that there exists some  $k$  such that  $P^k P = P^k$  [Grinstead and Snell, 1998]. Consequently, we know that the matrix

$$P^* = \lim_{t \rightarrow \infty} P^t \quad (5)$$

exists and in fact this limit is reached after only finitely many matrix multiplications. Thus a given initial population distribution  $\mu^{(0)}$  will converge to a stationary distribution  $\mu^*$  after a finite number of steps in our model. Furthermore, if  $P^*$  is known then we compute the stationary distribution  $\mu^*$  as

$$\mu^* = \mu^{(0)} P^*. \quad (6)$$

In particular, provided we assume a drug is applied for sufficiently long to ensure that the disease



population reaches evolutionary equilibrium, we can explore the effects of applying multiple drugs sequentially by considering the matrices  $P^*$  for the associated fitness landscapes. In the following discussion we make this assumption.

By encoding the evolutionary dynamics in a Markov Chain we can investigate the evolutionary process from an algebraic perspective. In particular, as the transition matrix  $P$  encodes all of the evolutionary dynamics of the associated fitness landscape  $f$ , we can explore global properties of  $f$  by considering the algebraic properties of  $P$ . In the following section we present two simple, yet powerful, consequences of this observation.

## Non-Commutativity and Cycling of Natural Selection

We use the Markov Chain model to formally prove that for a large class of fitness landscape pairs, there is non-commutativity in the evolutionary process as described by the SSWM assumptions. More precisely, consider two drugs,  $X$  and  $Y$ , with corresponding fitness landscapes  $x$  and  $y$ . We wish to determine what, if any, difference there is between applying  $X$  followed by  $Y$  to a population as opposed to applying  $Y$  followed by  $X$  to that population. If we construct the transition matrices  $P_x$  and  $P_y$  corresponding to  $x$  and  $y$ , respectively, and take the limits  $P_x^*$  and  $P_y^*$ , then our model predicts that the ordering makes no difference to the final population distribution on an initial population taking genotype  $i$  if, and only if,

$$\mu_i P_x^* P_y^* = \mu_i P_y^* P_x^*, \quad (7)$$

where  $\mu_i$  is row vector of length  $2^N$  whose  $i^{\text{th}}$  component is one and all of whose other components are zero.

If we do not know the starting population genotype we can only guarantee that the order of application is irrelevant when the outcome is the same regardless of the starting genotype. We require for all possible length  $2^N$  unit vectors  $\mu_i$  that  $\mu_i P_x^* P_y^* = \mu_i P_y^* P_x^*$ . Since these unit vectors form a basis of  $\mathbb{R}^N$  this occurs precisely when

$$P_x^* P_y^* = P_y^* P_x^*. \quad (8)$$

Hence drug application will only commute when the corresponding limit matrices commute. In practice we may be able to narrow down which genotypes are likely to constitute the population through bacterial genotyping or by observing that certain strains are not viable in the wild due to the high fitness cost of certain mutations. To determine how common commutativity is we first

tested each pair of fitness landscapes from the landscapes empirically determined by Mira et al. [2014] for *E. coli* in the presence of  $N = 4$  possible resistance conferring mutations under 15 different antibiotics (listed in Table 1). Of these 15 antibiotics we found no commuting pairs. We then tested  $10^6$  pairs of random fitness landscapes with varying ruggedness generated according to Kauffman’s NK model for generating “tunably rugged” fitness landscapes [Kauffman and Levin, 1987, Kauffman and Weinberger, 1989] using a random neighborhood Boolean function for determining the fitness contributions of each locus. We fixed  $N = 5$  and generated each landscape by first drawing  $K$  uniformly from  $\{0, \dots, N - 1\}$  and then using Kauffman’s model to build a landscape. We found that 0.132% of the landscape pairs generated had limit matrices which commuted, suggesting that commutativity is rare.

We now turn our attention to finding antibiotic cycling strategies as in Goulart et al. [2013]. Unless  $x$  is a flat landscape (taking equal values for all genotypes) there must exist a (not necessarily unique) genotype  $j$  whose fitness is a minimum and which has a fitter neighbor. Such a genotype satisfies  $\mathbb{P}[(i \rightarrow j)] = 0$  for all genotypes  $i$ . Hence if  $x$  is not flat, the limit matrix  $P_x^*$  has at least one column of all zeros and is singular, so there cannot exist a second landscape  $y$  for which  $P_x^* P_y^* = I$ . Hence there exists a unit row vector  $\mu_i$  for which  $\mu_i P_x^* P_y^* \neq \mu_i$ . As the  $\mu$  vectors encode probability distributions this means that natural selection in our model is irreversible in the sense that for a given (non-flat) landscape we cannot find another which is guaranteed to reverse its effects. This result precludes the existence of a general cycling strategy through which we can utilize a sequence of drugs to drive evolution through a cycle in genotype space such that the disease population returns to its original genotype, regardless of that starting genotype. If we do in fact know the starting genotype, as we might if the disease is contracted in the wild where resistance conferring mutations often carry a cost [Andersson and Hughes, 2010] making the wild-type genotype most likely, then cycling strategies can be found efficiently by our model. In particular, for a given starting genotype  $i$  the initial population distribution will be  $\mu_i$  and a sequence of drugs  $X_1, \dots, X_k$  with fitness landscapes  $x_1, \dots, x_k$  will constitute a cycling strategy precisely when  $\mu_i P_{x_1}^* \dots P_{x_k}^* = \mu_i$ . This criterion will be satisfied when  $\mu_i$  is a left 1-eigenvector of  $P_{x_1}^* \dots P_{x_k}^*$ . As such, we can find cycling strategies using matrix algebra and avoid the graph-search technique used by Goulart et al. [2013].

## Evolutionary Steering can both Prevent and Promote Resistance

Prescriptions of sequences of drugs occur frequently in the clinic, and often without any guidelines as to which orderings are preferable. Common examples of this include, but are not limited to,




Figure 2: The evolutionary graphs for the fitness landscapes of *E. coli* with the antibiotics (a) Ampicillin (Amp), (b) Ampicillin + Sulbactam (Sam) and (c) Cefprozil (Cpr) for 4 possible substitutions found in *bla*<sub>TEM-50</sub>. Arrows represent fitness conferring mutations which can fix under the SSWM assumptions, the absence of an arrow in either direction corresponds to a neutral mutation which cannot fix under our assumptions. Squares denote local optima within the fitness landscape.

treatment of *H. pylori* [Gisbert et al., 2010], Hepatitis B [Hanazaki, 2004] and the ubiquitous change from broad to narrow spectrum antibiotics [Heenen et al., 2012]. The ordering of the sequence is therefore often determined arbitrarily, by the individual clinician’s personal, or historical experience or from laboratory data. Ideally, we would like to be able to identify drug orderings that lower the probability of a highly resistant disease population emerging during the treatment. To consider optimal drug orderings in the context of our model we first need to know the fitness landscapes (or proxies of the fitness landscapes) of a number of antibiotics used to treat a given bacterial infection. Experimentally determining these landscapes requires one to consider all possible  $2^N$  combinations of genotypes in a set of  $N$  genes, a task which is prohibitively difficult for all but small values of  $N$ . De Visser and Krug [2014] found that there have been less than 20 systematic empirical studies of fitness landscapes and that landscapes have been calculated for a number of model organisms including *E. coli* [Khan et al., 2011, Tan et al., 2011, Schenk et al., 2013, Goulart et al., 2013], *Saccharomyces cerevisiae* [Hall et al., 2010], *Plasmodium falciparum* [Lozovsky et al., 2009] and type 1 Human Immunodeficiency Virus [da Silva et al., 2010]. Recent work by Hinkley et al. [2011] which utilizes regression methods to approximate large fitness landscapes from samples of the space could help ameliorate the complexity of experimentally determining fitness landscapes.

Mira et al. [2014] investigated the fitness landscapes of *E. coli* under 15 different  $\beta$ -lactam antibiotics using the mean minimum inhibitory concentration (MIC) of drug as a proxy for fitness for a total of  $N = 4$  resistance conferring mutations. Figure 2 shows the evolutionary graphs




Figure 3: The probability distributions for accessibility of the peaks of the Amp landscape for different steering regimes. The starting distribution is  $\mu = [1/2^N, \dots, 1/2^N]$ . When Amp is given first any of the three peaks of the landscape are accessible with most resistant genotype 1111 being most likely. If Sam is given first to steer the population to its sole peak 1111, then resistance to Amp will be guaranteed when it is applied. Alternatively, if Sam is given followed by Cpr then the population evolves to the local optima genotype 0110 of the Cpr landscape. If Amp is applied to this primed population the global optimum is inaccessible.

of the fitness landscapes for three of these antibiotics, Ampicillin (Amp), Ampicillin+Sulbactam (Sam) and Cefprozil (Cpr). We will use these three fitness landscapes to demonstrate the steering hypothesis explicitly. In the case of a single peaked landscape, such as that for Ampicillin + Sulbactam, we cannot reduce the likelihood of resistance as all evolutionary trajectories lead to the global fitness optimum. It is only when a drug has a multi-peaked landscape that we may be able to avoid resistance through careful choice of preceding drugs. Of the 15 landscapes determined empirically by Mira et al. [2014] only the landscape for Ampicillin+Sulbactam is single peaked. In their review of empirical fitness landscapes de Visser and Krug [2014] find that biological landscapes show a variable but substantial level of ruggedness suggesting that multi-peaked landscapes could be quite common. Poelwijk et al. [2007, 2011] showed that reciprocal sign epistasis within a landscape is necessary for the landscape to be multi-peaked.

In the following we take  $r = 0$  in equation (2) and note that changing the value of  $r$  will not change the accessibility of an evolutionary trajectory, hence by taking a different value of  $r \geq 0$  we will only change the result quantitatively (the probabilities may change) but not qualitatively. We begin by supposing that we do not know the initial population genotype. We can model this situation by taking as our prior population distribution  $\mu = [1/2^N, \dots, 1/2^N]$  which specifies that

each genotype is equally likely to constitute the starting population.

If we apply the drug Amp to this population distribution we find that in the expected distribution  $\mu^* = \mu P_{Amp}^*$  (shown in the first diamond in the top row of Figure 3) each of the fitness peaks could be found. In particular, the most highly resistant genotype 1111 can arise in the population with probability 0.62. Suppose instead we apply Sam first. In this case as the landscape is single peaked the population will converge to the global optimum genotype 1111. This genotype is also the global optimum of the Amp landscape and hence if we apply Amp after Sam we will encounter high resistance. We have steered the population with one drug to a configuration which increases the probability of resistance to a second. Next suppose instead we apply Cpr after Sam, in this case the population is guaranteed to evolve to a local optimum 0110 of the Cpr landscape. 0110 is the least fit local optimum of the Amp landscape. Thus if we apply Amp to the population primed by Sam and Cpr then evolution to the global optimum 1111 is not possible. This example demonstrates the steering hypothesis, that evolution can be shepherded through careful orderings of multiple drugs to increase or decrease the likelihood of resistance emerging.

To test our steering hypothesis we performed for each of the 15 drugs with landscapes derived by Mira et al. [2014] an *in silico* test of steering using combinations of one, two or three preceding drugs. Table 1 shows for each of the 15 antibiotics the combinations which reduce the probability of evolution to a peak fitness genotype to the lowest possible when applied in order followed by the final drug. We found that for 3 of the 15 drugs there exists another which steers an initial population  $\mu$  to a configuration which prevents evolution to the global fitness optimum of the landscape entirely. This number rose to 6 when pairs of drugs applied sequentially are used to steer the population and to 7 when triples applied in sequence were considered.

We then performed a second *in silico* experiment to find combinations of steering antibiotics that maximize the probability that evolution proceeds to the least fit of the local optima of a final antibiotic. Table 3 show the results of this experiment. We found that, excluding the single peaked landscape for Ampicillin with Sulbactam, there exist 0 drugs for which a single other drug is able to steer the population to a configuration from which evolution to only the least fit optimum is possible. If pairs of drugs are used to steer there are 3 such drugs (including the example presented in the above demonstration) and if triples of steering drugs are considered there remains only 3. These findings suggest that through careful choice of steering drugs we may be able to prevent the emergence of resistance. During these experiments we found that 14/15 of the antibiotics in our experiment (Cefpodoxime (CPD) being excluded) appeared in an optimal steering combination of some length.

Whilst careful selection of drugs for steering can prevent the emergence of resistance, arbitrary

drug orderings can also promote it. We performed an exhaustive *in silico* search of all singles, pairs of, and triples of steering drugs applied sequentially to prime the initial population  $\mu$  for a final application of each of the 15 antibiotics. We found that steering with a single drug increased the likelihood of the most resistant genotype emerging in 57.3% of cases and decreased the likelihood in 29.8% of cases. Steering with pairs of drugs increased the likelihood in 64.1% of cases and decreased it in 28.4% of cases and steering with triples increased the likelihood in 65.6% of cases and decreased it in 27.5%. For each of the antibiotics except Cefaclor, Cefprozil and Ampicillin+Sulbactam (which is single peaked making steering irrelevant) we found that a random steering combination of length one, two or three is more likely to increase the chances of resistance than to reduce it. Indeed, for Piperacillin+Tazobactam and Ceftizoxime we found that a random steering combination will increase the probability of the most highly resistance genotype occurring in more than 80% of cases, suggesting that sequential multidrug treatments which use these antibiotics should proceed with caution. These findings suggest that the present system of determining sequential drug orderings without quantitative optimisation based guidelines could in fact be promoting drug resistance and that to avoid resistance we must carefully consider the order in which drugs are applied.

## Discussion

We have encoded Strong Selection Weak Mutation evolutionary dynamics on fitness landscapes as a Markov Chain. Through this encoding we can explore the dynamics of evolution by considering the algebraic properties of the associated Markov transition matrix. In particular, we have demonstrated that evolution on fitness landscapes is non-commutative through parallels with the non-commutativity of matrix multiplication and that antibiotic cycling strategies can be determined through matrix multiplication. We argue then that the ordering in which a collection of drugs is applied can significantly impact the population that exists after the application is complete.

We have shown that we can find sequences of drugs that can be applied to both avoid and promote the emergence of resistance in the population. In light of the slow pace of novel antibiotic discovery and the rapid emergence of resistance to the presently most utilized antibiotics, this finding suggests a new treatment strategy — one in which we use a sequence of drugs (or even treatment breaks which themselves impose a selective pressure [Andersson and Hughes, 2010]) to steer, in an evolutionary sense, the disease population to avoid resistance from developing. Further, the drugs used to prime the disease population for treatment by an effective antibiotic do not themselves need to be the most effective drugs available. This means that there could be a large pool of potential steering drugs in the form of antibiotics which have gone unused for many years

Final drug	# of Peaks	Fittest peak genotype	Probability of ending at highest fitness genotype (no steering)	Best single steering drug	Probability of ending at highest fitness genotype (with best single steering drug)	Best ordered pair of steering drugs	Probability of ending at highest fitness genotype (with best steering pair)	Best ordered triple of steering drugs	Probability of ending at highest fitness genotype (with best steering triple)
Ampicillin (AMP)	3	1111	0.62	CPR	0.25	AMP→CPR	0.0	-	-
Amoxicillin (AMX)	2	1101	0.75	CTX	0.62	SAM→CPR	0.51	CRO→SAM→CPR	0.51
Cefaclor (CEC)	3	0011	0.19	SAM	0.0	-	-	-	-
Cefotaxime (CTX)	4	1111	0.17	CPR	0.04	AMP→CPR	0.0	-	-
Ceftizoxime (ZOX)	2	0111	0.86	AMX	0.81	SAM→AMX	0.75	CEC→CPR→AMC	0.75
Cefuroxime (CXM)	2	0111	0.56	AMC	0.33	CXM→AMC	0.19	AMC→CXM→AMC	0.11
Ceftriaxone (CRO)	4	1111	0.28	CEC	0.20	TZP→CEC	0.05	AMC→TZP→CEC	0.0
Amoxicillin +Clav (AMC)	2	1101	0.67	CXM	0.39	AMC→CXM	0.23	CXM→AMC→CXM	0.13
Ceftazidime (CAZ)	3	0110	0.39	TZP	0.08	AMC→TZP	0.0	-	-
Cefotetan (CTT)	5	0111	0.18	AMC	0.0	-	-	-	-
Ampicillin +Sulbactam (SAM)	1	1111	1.0	*	1.0	*	1.0	*	1.0
Cefprozil (CPR)	3	0101	0.25	AMP	0.0	-	-	-	-
Cefpodoxime (CPD)	2	1111	0.81	CRO	0.74	CTX→FEP	0.57	SAM→CEC→CRO	0.50
Piperacillin +Tazobactam (TZP)	2	0101	0.80	CTT	0.72	CTT→CTT	0.72	CTX→FEP→CTT	0.68
Cefepime (FEP)	4	1111	0.48	CTX	0.39	CAZ→CEC	0.36	CAZ→CEC→CTX	0.33

Table 1: For each of the 15 antibiotics we have derived the (ordered) sets of one, two and three steering drugs which reduce the probability of evolution to the maximal fitness genotype to the lowest possible. In the case that an ordered set of steering drugs reduced the probability to 0 we have not considered ordered sets of greater length (marked as - in the table). \* – as the landscape for SAM is single peaked there can be no combination of steering drugs which reduce the probability of finding the global optimum. In all experiments the initial population distribution is taken as  $\mu = [1/2^N, \dots, 1/2^N]$  and  $r = 0$



Final drug	# of single steering drugs better (* /15)	# of single steering drugs worse (* /15)	# of steering pairs better(* /225)	# of steering pairs worse (* /225)	# of steering triples better (* /3375)	# of steering pairs worse (* /3375)
Ampicillin (AMP)	7 (46.7%)	7 (46.7%)	97 (43.1%)	126 (56.0%)	1373 (40.7%)	1998 (59.2%)
Amoxicillin (AMX)	5 (33.3%)	9 (60.0%)	60 (26.7%)	164 (72.9%)	787 (23.3%)	2578 (76.4%)
Cefaclor (CEC)	9 (60.0%)	5 (33.3%)	132 (58.7%)	88 (39.1%)	2087 (61.8%)	1264 (37.5%)
Cefotaxime (CTX)	4 (26.7%)	10 (66.7%)	67 (29.8%)	155 (68.9%)	989 (29.3%)	2335 (69.2%)
Ceftizoxime (ZOX)	2 (13.3%)	12 (80.0%)	31 (13.8%)	193 (85.8%)	444 (13.2%)	2930 (86.8%)
Cefuroxime (CXM)	5 (33.3%)	9 (60.0%)	95 (42.2%)	128 (56.9%)	1486 (43.5%)	1885 (55.9%)
Ceftriaxone (CRO)	5 (33.3%)	9 (60.0%)	61 (27.1%)	163 (72.4%)	810 (24.0%)	2564 (76.0%)
Amoxicillin +Clav (AMC)	7 (46.7%)	7 (46.7%)	99 (44.0%)	124 (55.1%)	1428 (42.3%)	1943 (57.6%)
Ceftazidime (CAZ)	4 (26.7%)	10 (66.7%)	76 (33.8%)	147 (65.3%)	1218 (36.1%)	2153 (63.8%)
Cefotetan (CTT)	4 (26.7%)	10 (66.7%)	58 (25.8%)	166 (73.8%)	843 (25.0%)	2531 (75.0%)
Ampicillin +Sulbactam (SAM)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)	0 (0.0%)
Cefprozil (CPR)	7 (46.7%)	7 (46.7%)	113 (50.2%)	110 (48.9%)	1703 (50.5%)	1668 (49.4%)
Cefpodoxime (CPD)	3 (20.0%)	11 (73.3%)	26 (11.6%)	195 (86.7%)	282 (8.4%)	3077 (91.2%)
Piperacillin +Tazobactam (TZP)	2 (13.3%)	12 (80.0%)	11 (5.9%)	213 (94.7%)	81 (2.4%)	3293 (97.6%)
Cefepime (FEP)	3 (20.0%)	11 (73.3%)	34 (15.1%)	190 (84.4%)	402 (11.9%)	2972 (88.1%)
Overall	67 (29.8%)	129 (57.3%)	960 (28.4%)	2162 (64.1%)	13933 (27.5%)	33200 (65.6%)
Overall (Without steering for SAM)	67 (31.9%)	129 (61.4%)	960 (30.5%)	2162 (68.6%)	13933 (29.5%)	33200 (70.3%)

Table 2: For each of the 15 antibiotics we determined how many of the possible single, ordered double and ordered triple combinations of steering drugs (allowing drugs to appear multiple times in the combination and also allowing the target drug to appear as a steering drug) improved or worsened the probability of the most resistant genotype being found after the steering drugs are applied in order followed by the final target drug. In each case the initial population was given by  $\mu = [1/2^N, \dots, 1/2^N]$  and  $r = 0$ . As the SAM landscape is single peaked no combination of steering drugs will improve or worsen the outcome. As such, we have computed the overall numbers both with and without the contribution of the SAM row.



due to their inefficacy.

However, a major difficulty in using sequential drug treatments to steer disease populations is that in order to predict the outcomes we must know the fitness landscapes of the drugs involved. De Visser and Krug [2014] state that there exist less than 20 systematic studies of fitness landscapes and that these studies consider between 3 and 9 possible mutations. For steering to be fully effective we must account for all likely fitness conferring mutations and their effects on fitness under many drugs. Thus, many of the studies reviewed by de Visser and Krug are insufficient for determining clinically actionable steering strategies for certain diseases. Fortunately, for a number of highly resistant infectious diseases [Woodford and Ellington, 2007, Jensen and Lyon, 2009] and cancers [Lord and Ashworth, 2013, Lito et al., 2013] there are only a small number of mutations which seem to contribute to resistance. Further, recent work by Hinkley et al. [2011] in HIV has introduced a method to approximate large fitness landscapes from relatively fewer data points using a regression method. It follows that determining the landscapes is not an entirely intractable problem. A further complication in determining steering strategies is that fitness landscapes can be dependent on the disease microenvironment and have the potential to change from patient to patient or throughout the course of treatment. The consequences of such effects on fitness landscapes have not yet been experimentally determined.

Even in the absence of actual fitness landscapes our findings should be taken as a cautionary warning for multi-drug treatments, particularly those used to treat complex diseases such as multiple infections or cancers. In the same way that the drug ordering can be used to steer away from resistance we have shown it can also be used to make resistance more likely. Our results show that we may be inadvertently selecting for highly resistant disease populations through arbitrary drug ordering in the same way that highly resistant disease can emerge through irresponsible drug dosing. This result corroborates the findings of Pena-Miller and Beardmore [2010] that antibiotic cycling strategies can vary greatly in their efficacy and be both worse and better than mixing strategies. If we are to avoid resistance to our most effective drugs we must carefully consider how they are used together, both in combination and in sequence, with other drugs and take appropriate steps to reduce the risk.

Two major assumptions within our modeling are that drugs are administered for sufficiently long that evolution can converge to a local fitness optimum and that this convergence is guaranteed to occur. This assumption poses two potential problems in converting our model predictions to predictions of real-world bacterial evolution. The first is that if selection is strong and mutations are rare then there is a possibility of the population being driven to extinction before an adaptive mutation occurs. We have chosen to ignore this possibility within our modeling as in the context

of treating bacterial infections this would constitute a success. The second is that the time to convergence could be prohibitively long for steering to constitute a realistic treatment strategy. We believe that the assumption of reasonable convergence times could be valid as adaptive walks in rugged landscapes are often short [Orr, 2005]. However, it has been shown that for certain landscapes there can exist adaptive walks of length exponential in the number of loci [Kaznatcheev, 2013], but since we get to choose those drugs with which to steer we can avoid landscapes for which the convergence time is prohibitively long. Further, our model is not necessarily restricted to the dynamics within a single patient. Goulart et al. [2013] used fitness landscapes to explore whole hospital scale antibiotic treatment strategies and our model, as an encoding of evolution on fitness landscapes, is capable of making predictions at this scale also. As such, even if evolutionary convergence is experimentally determined to be prohibitively slow for steering to be effective as a treatment for bacterial infection within a single patient, our results will still hold in scenarios which admit longer timescales. Specifically, evolutionary steering could provide an effective means to avoid the emergence of drug resistance at the hospital scale or in life-long diseases such as HIV.

The Strong Selection Weak Mutation model we have used here is a highly simplified, yet well studied model of evolution. The model ignores much of the complexity of the evolutionary process, specifically simplifying the genotype-phenotype map, ignoring the disease microenvironment and making the assumption of a monomorphic disease population in which deleterious and neutral mutations cannot fix. Under certain regimes of population size and mutation rates these simplifying assumptions break down. For example, if the population is sufficiently large then stochastic tunneling [Iwasa et al., 2004] – the situation where double mutations can occur allowing the crossing of fitness valleys – can arise causing a breakdown of the Strong Selection assumption. Similarly, if the mutation rate is sufficiently high then the population ceases to be monomorphic and forms a quasispecies [Nowak, 1992, Bull et al., 2005]. Conversely, if the population is sufficiently small then it becomes possible for deleterious mutations to fix [Moran et al., 1962, Wright, 1932, Fisher, 1958]. Finally, we have ignored the possibility of neutral spaces in the fitness landscape which have been shown to have significant impact on whether non-optimal genotypes can fix in the population as well as the time taken for evolution to find a locally optimal genotype [Schaper and Louis, 2014, Schaper et al., 2012]. We believe that each of these breakdowns of the SSWM model will have important consequences for the possibility and efficacy of steering and hence a proper treatment of their implications is beyond the scope of this paper. In our future work we aim to undertake a comprehensive study of the implications of population size, mutation rate, neutral drift and evolutionary convergence times on the steering hypothesis.

The SSWM model is a simplification of the evolutionary process and given that non-commutativity

is present in this highly simplified model it is unlikely that commutativity will emerge as more complexity is introduced. It follows that the cautionary message regarding sequential drug application which results from our simplified model merits serious consideration. Whether or not measuring fitness landscapes provides sufficient information to correctly identify optimal drug orderings *in vivo* is a question that cannot be answered through mathematical modeling alone. It is only by verifying the predictions of steering strategies given by our model through biological experiment that we can determine whether they are viable. Supposing our model predictions are indeed viable then knowledge of some approximation to the fitness landscapes of the presently most used antibiotics could, in combination with our model, provide at least a good heuristic for how to proceed with multi-drug treatments, future antibiotic stewardship programs and clinical trial design.

## Acknowledgment and Supplementary Information

We would like to offer the code underlying this model to any interested parties openly. DN would like to thank the Engineering and Physical Sciences Research Council (EPSRC) for generous funding for his doctoral studies. JGS would like to thank the NIH Loan Repayment Program for support. AGF is funded by the EPSRC and Microsoft Research, Cambridge through grant EP/I017909/1. ARAA and JGS gratefully acknowledge funding from the NCI Integrative Cancer Biology Program (ICBP) grant U54 CA113007 and they and PKM and RAG also thank the NCI Physical Sciences in Oncology Centers U54 CA143970 grant. We would also like to thank the Veterans Affairs Merit Review Program (RAB), the National Institutes of Health (Grant AI072219-05 and AI063517-07 to RAB), the Geriatric Research Education and Clinical Center VISN 10 (RAB), and an unrestricted research grant from the STERIS Corporation.

Final drug	# of Peaks	Least fit peak genotype	Probability of ending at lowest fitness genotype (no steering)	Best single steering drug	Probability of ending at lowest fitness genotype (with best single steering drug)	Best ordered pair of steering drugs	Probability of ending at lowest fitness genotype (with best steering pair)	Best ordered triple of steering drugs	Probability of ending at lowest fitness genotype (with best steering triple)
Ampicillin (AMP)	3	0110	0.19	CPR	0.51	SAM→CPR	1.0	-	-
Amoxicillin (AMX)	2	0010	0.25	CTX	0.38	SAM→CPR	0.49	AMP→SAM→CPR	0.49
Cefaclor (CEC)	3	0100	0.41	TZP	0.70	CXM→AMC	0.78	AMC→CXM→AMC	0.87
Cefotaxime (CTX)	4	1010	0.21	FEP	0.27	CTX→FEP	0.43	SAM→CEC→CRO	0.50
Ceftizoxime (ZOX)	2	1001	0.14	AMX	0.19	SAM→AMX	0.25	AMP→SAM→AMX	0.25
Cefuroxime (CXM)	2	0100	0.44	AMC	0.67	CXM→AMC	0.81	AMC→CXM→AMC	0.89
Ceftriaxone (CRO)	4	0100	0.30	CXM	0.59	AMC→CXM	0.76	CXM→AMC→CXM	0.86
Amoxicillin +Clav (AMC)	2	0100	0.33	CXM	0.61	AMC→CXM	0.77	CXM→AMC→CXM	0.87
Ceftazidime (CAZ)	3	0011	0.26	FEP	0.28	CAZ→CEC	0.40	CAZ→AMX→FEP	0.45
Cefotetan (CTT)	5	1101	0.19	AMX	0.75	SAM→AMX	1.0	-	-
Ampicillin +Sulbactam (SAM)	1	1111	1.0	*	1.0	*	1.0	*	1.0
Cefprozil (CPR)	3	0011	0.25	ZOX	0.26	CAZ→CEC	0.34	CAZ→CEC→FEP	0.43
Cefpodoxime (CPD)	2	1010	0.19	CRO	0.26	CTX→FEP	0.43	SAM→CEC→CRO	0.50
Piperacillin +Tazobactam (TZP)	2	1000	0.20	CTT	0.28	CTT→CTT	0.28	CTX→FEP→CTT	0.32
Cefepime (FEP)	4	0000	0.14	FEP	0.14	TZP→CEC	0.17	CXM→AMC→CTX	0.20

Table 3: For each of the 15 antibiotics we have derived the (ordered) sets of one, two and three steering drugs which increase the probability of evolution to the least fit optimum genotype to the highest possible. In the case that an ordered set of steering drugs increases the probability to 1 we have not considered ordered sets of greater length (marked as - in the table). \* – as the landscape for SAM is single peaked there can be no combination of steering drugs which change the probability of finding the global optimum. In all experiments the initial population distribution is taken as  $\mu = [1/2^N, \dots, 1/2^N]$  and  $r = 0$

# Bibliography

- Brad Spellberg, John G Bartlett, and David N Gilbert. The future of antibiotics and resistance. *New England Journal of Medicine*, 368(4):299–302, 2013.
- Gary L French. The continuing crisis in antibiotic resistance. *International Journal of Antimicrobial Agents*, 36:S3–S7, 2010.
- Lynn L Silver. Challenges of antibacterial discovery. *Clinical Microbiology Reviews*, 24(1):71–109, 2011.
- David M Shlaes, Dale N Gerding, Joseph F John, William A Craig, Donald L Bornstein, Robert A Duncan, Mark R Eckman, William E Farrer, William H Greene, Victor Lorian, et al. Society for Healthcare Epidemiology of America and Infectious Diseases Society of America Joint Committee on the Prevention of Antimicrobial Resistance: Guidelines for the Prevention of Antimicrobial Resistance in Hospitals. *Clinical Infectious Diseases*, 25(3):584–599, 1997.
- John G Bartlett. A call to arms: the imperative for antimicrobial stewardship. *Clinical Infectious Diseases*, 53(suppl 1):S4–S7, 2011.
- Kimberly D Leuthner and Gary V Doern. Antimicrobial stewardship programs. *Journal of Clinical Microbiology*, 51(12):3916–3920, 2013.
- Alan G Mathew, Robin Cissell, and S Liamthong. Antibiotic resistance in bacteria associated with food animals: a united states perspective of livestock production. *Foodborne Pathogens and Disease*, 4(2):115–133, 2007.
- Julia Moody, Sara E Cosgrove, Russell Olmsted, Edward Septimus, Kathy Aureden, Shannon Oriola, Gita Wasan Patel, and Kavita K Trivedi. Antimicrobial stewardship: a collaborative partnership between infection preventionists and healthcare epidemiologists. *Infection Control and Hospital Epidemiology*, 33(4):328–330, 2012.

396 Remy Chait, Allison Craney, and Roy Kishony. Antibiotic interactions that select against resistance.  
397 *Nature*, 446(7136):668–671, 2007.

398 Christiane P Goulart, Mentar Mahmudi, Kristina A Crona, Stephen D Jacobs, Marcelo Kallmann,  
399 Barry G Hall, Devin C Greene, and Miriam Barlow. Designing antibiotic cycling strategies by  
400 determining and understanding local adaptive landscapes. *PloS one*, 8(2):e56040, 2013.

401 R Pena-Miller and Robert Beardmore. Antibiotic cycling versus mixing: the difficulty of using  
402 mathematical models to definitively quantify their relative merits. 2010.

403 Carl T Bergstrom, Monique Lo, and Marc Lipsitch. Ecological theory suggests that antimicrobial  
404 cycling will not reduce antimicrobial resistance in hospitals. *Proceedings of the National Academy*  
405 *of Sciences of the United States of America*, 101(36):13285–13290, 2004.

406 J Arjan GM de Visser and Joachim Krug. Empirical fitness landscapes and the predictability of  
407 evolution. *Nature Reviews Genetics*, 2014.

408 Daniel M Weinreich, Richard A Watson, and Lin Chao. Perspective: sign epistasis and genetic  
409 constraint on evolutionary trajectories. *Evolution*, 59(6):1165–1174, 2005.

410 Daniel M Weinreich, Nigel F Delaney, Mark A DePristo, and Daniel L Hartl. Darwinian evolution  
411 can follow only very few mutational paths to fitter proteins. *Science*, 312(5770):111–114, 2006.

412 Longzhi Tan, Stephen Serene, Hui Xiao Chao, and Jeff Gore. Hidden randomness between fitness  
413 landscapes limits reverse evolution. *Physical Review Letters*, 106(19):198102, 2011.

414 Pia Schulz zur Wiesch, Jan Engelstädter, and Sebastian Bonhoeffer. Compensation of fitness costs  
415 and reversibility of antibiotic resistance mutations. *Antimicrobial agents and chemotherapy*, 54  
416 (5):2085–2095, 2010.

417 Mark M Tanaka and Frank Valckenborgh. Escaping an evolutionary lobster trap: drug resistance  
418 and compensatory mutation in a fluctuating environment. *Evolution*, 65(5):1376–1387, 2011.

419 Lejla Imamovic and Morten OA Sommer. Use of collateral sensitivity networks to design drug  
420 cycling protocols that avoid resistance development. *Science translational medicine*, 5(204):  
421 204ra132–204ra132, 2013.

422 Sewall Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In  
423 *Proceedings of the Sixth International Congress on Genetics*, volume 1, pages 356–366, 1932.

Richard W Hamming. Error detecting and error correcting codes. *Bell System Technical Journal*, 29(2):147–160, 1950.

John H Gillespie. A simple stochastic gene substitution model. *Theoretical Population Biology*, 23(2):202–215, 1983.

John H Gillespie. Molecular evolution over the mutational landscape. *Evolution*, pages 1116–1129, 1984.

Yoh Iwasa, Franziska Michor, and Martin A Nowak. Stochastic tunnels in evolutionary dynamics. *Genetics*, 166(3):1571–9, Mar 2004.

H Allen Orr. The genetic theory of adaptation: a brief history. *Nature Reviews Genetics*, 6(2):119–127, 2005.

Stuart Kauffman and Simon Levin. Towards a general theory of adaptive walks on rugged landscapes. *Journal of Theoretical Biology*, 128(1):11–45, 1987.

Walter Fontana, Peter F Stadler, Erich G Bornberg-Bauer, Thomas Griesmacher, Ivo L Hofacker, Manfred Tacker, Pedro Tarazona, Edward D Weinberger, and Peter Schuster. RNA folding and combinatorial landscapes. *Physical Review E*, 47(3):2083, 1993.

Catherine A Macken and Alan S Perelson. Protein evolution on rugged landscapes. *Proceedings of the National Academy of Sciences of the United States of America*, 86(16):6191–6195, 1989.

Catherine A Macken, Patrick S Hagan, and Alan S Perelson. Evolutionary walks on rugged landscapes. *SIAM Journal on Applied Mathematics*, 51(3):799–827, 1991.

Henrik Flyvbjerg and Benny Lautrup. Evolution in a rugged fitness landscape. *Physical Review A*, 46(10):6714, 1992.

John H Gillespie. *The causes of molecular evolution*. Oxford University Press, 1991.

Charles Miller Grinstead and James Laurie Snell. *Introduction to Probability*. American Mathematical Soc., 1998.

Guy Sella and Aaron E Hirsh. The application of statistical physics to evolutionary biology. *Proceedings of the National Academy of Sciences of the United States of America*, 102(27):9541–9546, 2005.

Pat A P Moran et al. The statistical processes of evolutionary theory. 1962.

452 Ronald A Fisher. *The genetical theory of natural selection.* , 1958.

453 Portia M Mira, Kristina Crona, Devin Greene, Juan C Meza, Bernd Sturmfels, and Miriam Barlow.  
454 Rational design of antibiotic treatment plans. *arXiv preprint arXiv:1406.1564*, 2014.

455 Stuart A Kauffman and Edward D Weinberger. The NK model of rugged fitness landscapes and  
456 its application to maturation of the immune response. *Journal of Theoretical Biology*, 141(2):  
457 211–245, 1989.

458 Dan I Andersson and Diarmaid Hughes. Antibiotic resistance and its cost: is it possible to reverse  
459 resistance? *Nature Reviews Microbiology*, 8(4):260–271, 2010.

460 Javier P Gisbert, Xavier Calvet, Anthony O’Connor, Francis Mégraud, and Colm A O’Morain.  
461 Sequential therapy for helicobacter pylori eradication: a critical review. 44(5):313–325, 2010.

462 Kazuhiro Hanazaki. Antiviral therapy for chronic hepatitis b: a review. *Current Drug Targets-  
463 Inflammation & Allergy*, 3(1):63–70, 2004.

464 Sarah Heenen, Frédérique Jacobs, and Jean-Louis Vincent. Antibiotic strategies in severe nosoco-  
465 mial sepsis: Why do we not de-escalate more often? *Critical Care Medicine*, 40(5):1404–1409,  
466 2012.

467 Aisha I Khan, Duy M Dinh, Dominique Schneider, Richard E Lenski, and Tim F Cooper. Negative  
468 epistasis between beneficial mutations in an evolving bacterial population. *Science*, 332(6034):  
469 1193–1196, 2011.

470 Martijn F Schenk, Ivan G Szendro, Merijn LM Salverda, Joachim Krug, and J Arjan GM de Visser.  
471 Patterns of epistasis between beneficial mutations in an antibiotic resistance gene. *Molecular  
472 biology and evolution*, 30(8):1779–1787, 2013.

473 David W Hall, Matthew Agan, and Sara C Pope. Fitness epistasis among 6 biosynthetic loci in the  
474 budding yeast *saccharomyces cerevisiae*. *Journal of Heredity*, 101(suppl 1):S75–S84, 2010.

475 Elena R Lozovsky, Thanat Chookajorn, Kyle M Brown, Mallika Imwong, Philip J Shaw, Sumalee  
476 Kamchonwongpaisan, Daniel E Neafsey, Daniel M Weinreich, and Daniel L Hartl. Stepwise  
477 acquisition of pyrimethamine resistance in the malaria parasite. *Proceedings of the National  
478 Academy of Sciences of the United States of America*, 106(29):12025–12030, 2009.

479 Jack da Silva, Mia Coetzer, Rebecca Nedellec, Cristina Pastore, and Donald E Mosier. Fitness  
480 epistasis and constraints on adaptation in a human immunodeficiency virus type 1 protein region.  
481 *Genetics*, 185(1):293–303, 2010.



- 482 Trevor Hinkley, João Martins, Colombe Chappey, Mojgan Haddad, Eric Stawiski, Jeannette M  
483 Whitcomb, Christos J Petropoulos, and Sebastian Bonhoeffer. A systems analysis of mutational  
484 effects in hiv-1 protease and reverse transcriptase. *Nature genetics*, 43(5):487–489, 2011.
- 485 Frank J Poelwijk, Daniel J Kiviet, Daniel M Weinreich, and Sander J Tans. Empirical fitness  
486 landscapes reveal accessible evolutionary paths. *Nature*, 445(7126):383–386, 2007.
- 487 Frank J Poelwijk, Sorin Tănase-Nicola, Daniel J Kiviet, and Sander J Tans. Reciprocal sign epistasis  
488 is a necessary condition for multi-peaked fitness landscapes. *Journal of Theoretical Biology*, 272  
489 (1):141–144, 2011.
- 490 Niel Woodford and Matthew J Ellington. The emergence of antibiotic resistance by mutation.  
491 *Clinical Microbiology and Infection*, 13(1):5–18, 2007.
- 492 Slade O Jensen and Bruce R Lyon. Genetics of antimicrobial resistance in staphylococcus aureus.  
493 *Future Microbiology*, 4(5):565–582, 2009.
- 494 Christopher J Lord and Alan Ashworth. Mechanisms of resistance to therapies targeting brca-  
495 mutant cancers. *Nature Medicine*, 19(11):1381–1388, 2013.
- 496 Piro Lito, Neal Rosen, and David B Solit. Tumor adaptation and resistance to raf inhibitors. *Nature*  
497 *Medicine*, 19(11):1401–1409, 2013.
- 498 Artem Kaznatcheev. Complexity of evolutionary equilibria in static fitness landscapes. *arXiv*  
499 *preprint arXiv:1308.5094*, 2013.
- 500 Martin A Nowak. What is a quasispecies? *Trends in Ecology & Evolution*, 7(4):118–121, 1992.
- 501 James J Bull, Lauren Ancel Meyers, and Michael Lachmann. Quasispecies made simple. *PLoS*  
502 *computational biology*, 1(6):e61, 2005.
- 503 Steffen Schaper and Ard A Louis. The arrival of the frequent: How bias in genotype-phenotype  
504 maps can steer populations to local optima. *PLOS ONE*, 9(2):e86635, 2014.
- 505 Steffen Schaper, Iain G Johnston, and Ard A Louis. Epistasis can lead to fragmented neutral spaces  
506 and contingency in evolution. *Proc Biol Sci*, 279(1734):1777–83, 2012.