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

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

[1] Department of Computer Science, University of Oxford, Oxford, UK

[2] Wolfson Centre for Mathematical Biology, Mathematical Institute, University of Oxford, Oxford, UK

[3] Department of Medicine, Louis Stokes Department of Veterans Affairs Hospital, Cleveland, OH, USA

[4] School of Computing Sciences and Informatics, University of Cincinnati, Cincinnati, OH, USA

[5] 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 and
Daniel Nichol, e-mail: dan.nichol@cs.ox.ac.uk

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

Statement of impact

In a time when we receive almost daily warnings of a 'post-antibiotic era' from the CDC and other groups, and drug development is stalling, we find ourselves in desperate need for novel strategies in the fight against bacterial evolution. Herein, we abstract the process of evolution on a fitness landscape to a Markov chain and, using this abstraction, demonstrate how different orderings of commonly used antibiotic therapies can steer bacterial evolution to genotypes from which highly resistant states are inaccessible. These results suggest a strategy by which, using drugs which may have been been considered less efficacious, we can prevent the emergence of resistance before it arises.

Abstract

The increasing rate of antibiotic resistance and slowing discovery of novel antibiotic treatments presents a growing threat to public health. In the present study we develop a Markov Chain model of evolution in asexually reproducing populations which we use to illustrate that different selection pressures do not commute. We demonstrate that the emergence of resistant individuals can be both hindered and promoted by careful orderings of drug application. This suggests a new strategy in the war against antibiotic therapy resistant organisms: rational drug ordering to shepherd evolution through genotype space to states corresponding to greater sensitivity to antibiotic treatment. The model we present is an encoding of the 'Strong Selection Weak Mutation' model of evolution on fitness landscapes within a Markov Chain, which associates the global properties of the fitness landscape with the algebraic properties of the Markov Chain transition matrix. Through this association we derive results on the non-commutativity and irreversibility of natural selection.

Introduction

Resistance to antibiotic treatments within bacterial pathogens poses an increasing threat to public health, which coupled with the slowing discovery of novel antibiotics, could soon reach crisis point [Spellberg et al., 2013, French, 2010]. Novel classes of antibiotics discovered since 1987 are few in number [Silver, 2011] . Thus, it is becoming ever clearer that if we are to combat highly resistant bacterial infections, then we must find new ways to prevent resistance and new applications of existing antibiotics to these pathogens. Indeed, public health efforts have attempted to stem the emergence of resistance by reducing unnecessary prescription of antibiotics [Shlaes et al., 1997, Bartlett, 2011, Leuthner and Doern, 2013] and stopping the addition of sub-therapeutic antibiotics in livestock feed [Mathew et al., 2007]. However, such policies require global adoption to be truly effective [Moody et al., 2012], which they have not yet achieved, and which is likely infeasible. Recently, there have been efforts to explore how existing antibiotics can be used in new ways to provide effective treatments for resistant pathogens, for example through combination therapy [Chait et al., 2007], though our understanding of the mechanisms of these processes remains limited.

In order to understand how to minimize the emergence of resistant pathogens, and to decide how best to treat them, we must first understand how their evolution is driven by the selective pressures of different antibiotic drugs — a fundamental problem of biology. In particular, if we understand which traits are likely to be selected for by which treatments, then we may be able to avoid selecting for those traits which confer resistance. Recent insights into the evolutionary process have yielded some actionable information. Specifically, Weinreich et al. [2005, 2006] showed that if the genome of a pathogen exhibits sign epistasis, where a given mutation is beneficial on some genetic backgrounds and deleterious on others, then there can exist inaccessible evolutionary trajectories. Further, Tan et al. [2011] studied the evolutionary trajectories of *Escherichia coli* under different antibiotics and found that adaptive mutations gained under one antibiotic are often irreversible when a second is applied. These findings lead us to hypothesize that one antibiotic could be used to irreversibly steer the evolution of a population of pathogens to a genotype (or combination of genotypes) from which it is much more difficult to acquire resistance to a second antibiotic. In this paper we present a Markov Chain model of evolution which we use to illustrate that selective pressures are non-commutative, and that the emergence of resistant genotypes can be both hindered and promoted by different orderings of selective pressures. These findings suggest a new strategy in the war against therapy resistant organisms: rational drug ordering to shepherd evolution through genotype space to more sensitive states.

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. Further, we assume that the mutation rate is sufficiently low that stochastic tunneling [Iwasa et al., 2004] through double mutations does not occur. 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], $\mathrm{Ham}(g, g')$, from g is 1. As such, our genotype space can be represented by an undirected N-dimensional 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

$$f : \{0,1\}^N \to \mathbb{R}^{\geq 0}. \tag{1}$$

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 $M \log M = \mathrm{o}(u)$, then each beneficial mutation in the population will either reach fixation or become extinct before a new mutation occurs and that a deleterious mutation will become extinct with sufficiently high probability that we may assume that this always occurs. It follows that after each mutation the population will stabilize to consist entirely of individuals with the same genotype, and that 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 uphill 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], including 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.

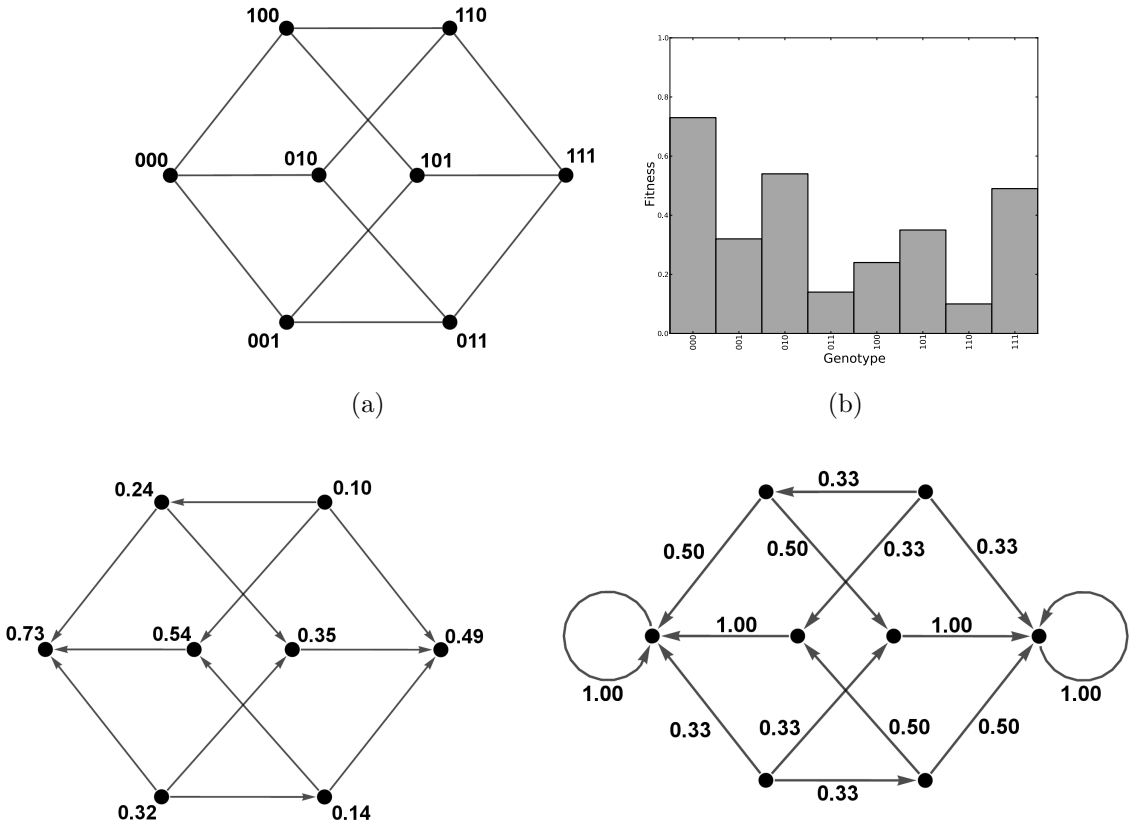

Figure 1: (a) The space of genotypes comprising bit strings of length $N = 3$. The vertices represent genotypes, and edges connect those genotypes which are mutational neighbors. (b) An example fitness landscape. (c) The directed evolutionary graph according to the landscape in (b) where the vertices still represent genotypes but are labelled 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 $r = 0$.

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 modelled by a Markov Chain [Grinstead and Snell, 1998]. For a fitness

function $f : \{0,1\}^N \to \mathbb{R}^{\geq 0}$ we can define a transition matrix $P = [\mathbb{P}(i \to j)]_{i,j \in \{0,1\}^N}$ for a time–homogeneous absorbing Markov Chain by setting, for $i \neq j$,

$$\mathbb{P}(i \to j) = \begin{cases} \dfrac{\big(f(j)-f(i)\big)^r}{\displaystyle\sum_{\substack{g \in \{0,1\}^N,\ \mathrm{Ham}(i,g)=1 \\ f(g)-f(i)>0}} \big(f(g)-f(i)\big)^r} & \text{if } f(j) > f(i) \text{ and } \mathrm{Ham}(i,j) = 1 \\ 0 & \text{otherwise} \end{cases}, \tag{2}$$

and

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

for each i (see Figure 1d). Here the parameter $r \geq 0$ determines the extent to which the fitness increase of a mutation affects its likelihood. In the case $r = 0$, we have the random move SSWM model (as in Macken and Perelson [1989], Macken et al. [1991], Flyvbjerg and Lautrup [1992]), in the limit $r \to \infty$ we have the steepest gradient ascent 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^{th} component which gives the probability that the population has the k^{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 occuring 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

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

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

$$P^* = \lim_{t \to \infty} P^t \tag{5}$$

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

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

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 a simple, yet powerful, consequences of this observation.

Non-Commutativity and Irreversibility 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,

$$\mu_i P_x^* P_y^* = \mu_i P_y^* P_x^*, \tag{7}$$

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

In practice we are unlikely to know the starting population genotype and as such 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 μ_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

$$P_x^* P_y^* = P_y^* P_x^*. \tag{8}$$

Hence drug application will only commute when the corresponding limit matrices commute. To test how common commutativity is we tested 10^7 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]. We fixed $N = 5$ and generated each landscape by first drawing K uniformly from $\{0, \ldots, N-1\}$ and then using Kauffman's model. We found that 3.22% of the landscape pairs generated had limit matrices which commuted, suggesting that commutativity is rare. In fact, this figure is an over-estimate as when $K = 0$ in the NK model the landscape will be single peaked with either $[0, 0, 0, 0, 0]$ or $[1, 1, 1, 1, 1]$ being the peak. As such, a pair of single peaked landscapes in this experiment will commute in 50% of cases as opposed to 3.125% as for general single peaked landscapes.

Further, unless x is a flat landscape (taking equal values for all genotypes) there must exist a genotype j whose fitness is strictly less than all others and which has a fitter neighbor. Such a genotype satisfies $\mathbb{P}[(i \to j)] = 0$ 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 μ_i for which $\mu_i P_x^* P_y^* \neq \mu_i$. As the μ 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.

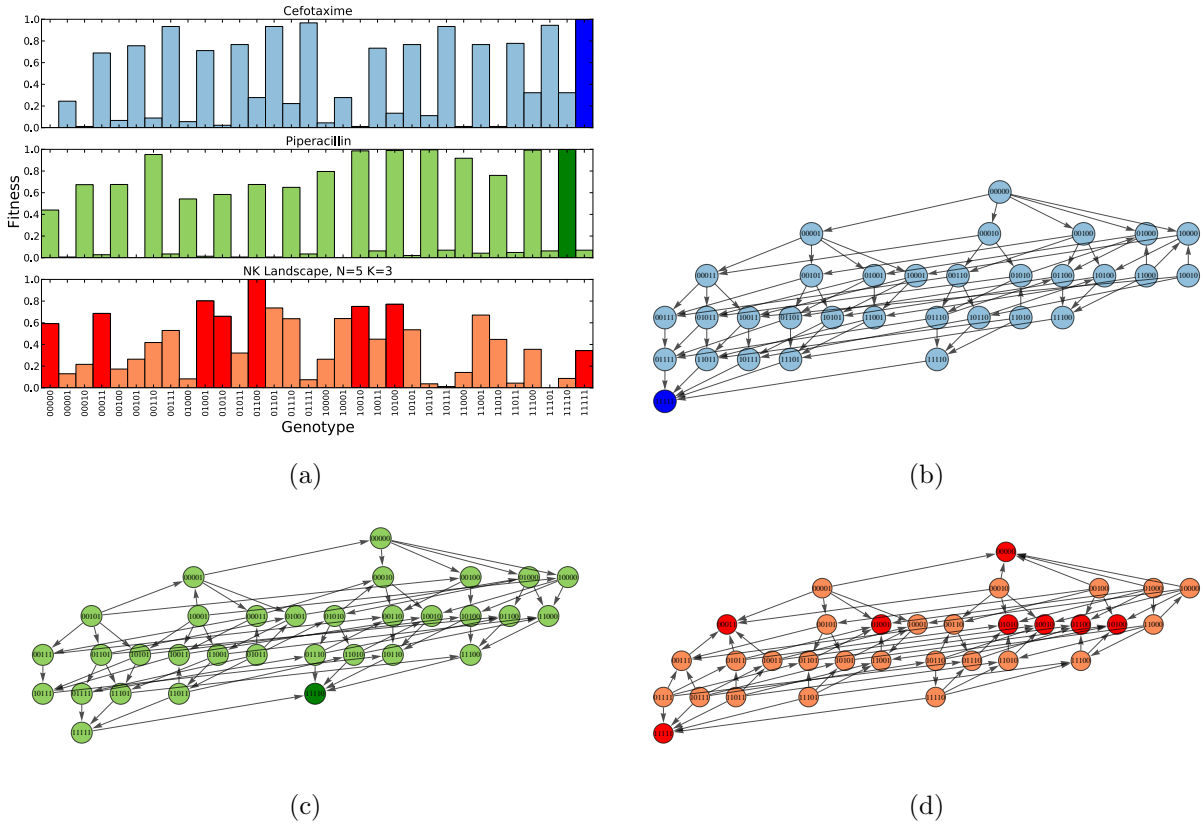

Figure 2: (a) Three fitness landscapes normalised to the range $[0, 1]$. Blue is the fitness landscape for Cef taken from Tan et al. [2011], green is the fitness landscape for Pip/I from Tan et al. [2011] and red is a landscape generated by the NK algorithm for $N = 5$, $K = 3$. (b) The evolutionary graph for the Cef landscape. (c) The evolutionary graph for the Pip/I landscape. (d) The evolutionary graph for the NK landscape. The local optima of each landscape are highlighted in bold. Note that as we have taken $r = 0$ in equation (2) only the ordering of genotype fitness values is relevant, hence the normalization of the fitness values to the range $[0, 1]$ does not change our results.

Treatment Optimization

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, 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. Unfortunately, experimentally determining these landscapes requires us 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. In spite of these difficulties, empirical fitness landscapes have been calculated for a number of model organisms including *E. coli* [Khan et al., 2011, Tan et al., 2011], *Saccharomyces cerevisiae* [Hall et al., 2010], *Plasmodium falciparum* [Lozovsky et al., 2009] and type 1 Human Immunodeficiency Virus [da Silva et al., 2010]. However, these landscapes are studied for a single environment and so are unsuitable for investigating the effects of a major change in the environment (e.g. a different drug application) on the evolutionary dynamics of the population.

Tan et al. [2011] investigated the fitness landscapes of *E. coli* under two β-lactam antibiotics, Cefotaxime (Cef) and Piperacillin with an inhibitor, clavulanic acid (Pip/I), using the mean minimum inhibitory concentration (MIC) of drug as a proxy for fitness. Figure 2 shows the fitness landscapes derived by Tan et al. [2011] normalised to the range $[0, 1]$ along with layered graph plots of the associated directed evolutionary graphs in which an edge indicates a mutation which increases fitness. We see that both of these landscapes are single peaked (peaks are highlighted) and that the peak genotypes correspond to global basins of attraction in the evolutionary graph. Hence, for any sequential application of fitness landscapes for which Cef or Pip/I are the final landscapes it does not matter which landscapes precede them, as the evolutionary trajectory will always proceed to the single global fitness optima.

We may be able to avoid the emergence of resistance (or equivalently avoid reaching some high fitness peak in the landscape) through careful choice of preceding drug landscapes only when the final landscape in the sequence is multi-peaked. Lozovsky et al. [2009] observed a multi-peaked landscape in four mutations in a gene responsible for resistance to an antimalarial, demonstrating that multi-peaked landscapes can occur naturally. A necessary condition for such landscapes to exist is that they exhibit reciprocal sign epistasis [Poelwijk et al., 2007, 2011]. In the following demonstration of evolutionary steering we use a multi–peaked landscape, denoted F (shown as red in Figure 2), generated according to Kauffmann's NK model, an algorithm which generates "tunably rugged" fitness landscapes, with $N = 5$, $K = 3$ [Kauffman and Levin, 1987, Kauffman and Weinberger, 1989]. This landscape has several fitness peaks with the genotype $[0, 1, 1, 0, 0]$ being the global optimum. For this demonstration we assume that this genotype corresponds to a highly resistant phenotype which we wish to avoid.

In the following demonstration we take $r = 0$ in equation (2) as this allows us to normalize the fitness landscapes to the range $[0, 1]$ without changing the final population distributions. Changing the value of r will not change the accessability of an evolutionary trajectory. Hence, by taking a

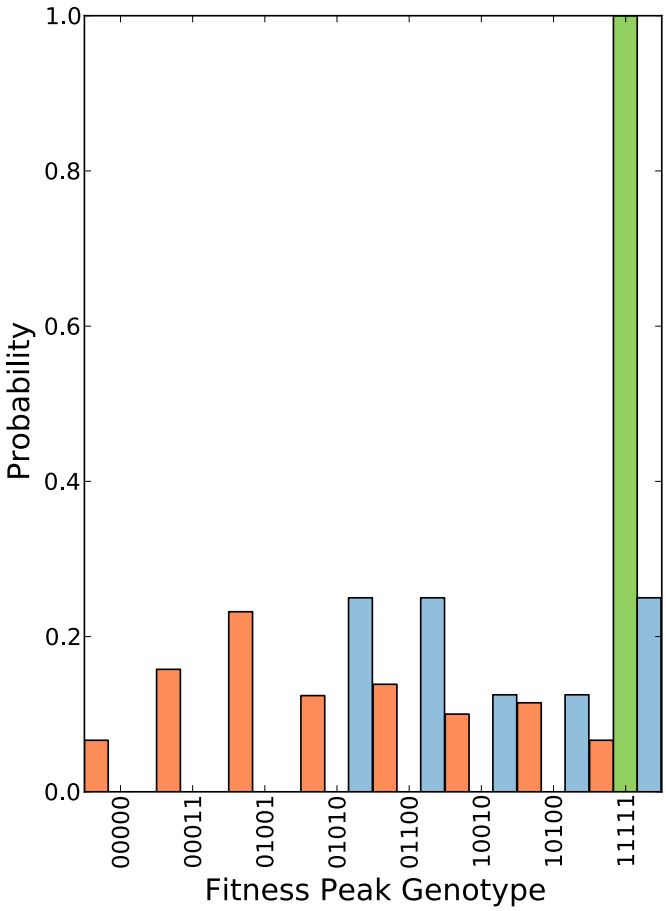

Figure 3: The probabilitiy distributions for which peaks of the NK landscape F, shown as red in Figure 2, are accessible with different steering landscapes applied first. Only genotypes corresponding to the peaks of F are shown as the others will always have probability zero of being the final population genotype. Red is the distribution for F applied to the initial population distribution $\mu = [1/2^N \ldots, 1/2^N]$. As expected all of the peaks highlighted in Figure 2 are accessible. Blue is the distribution if the Cef landscape is applied first to the initial distribution μ and F is applied afterwards. Green is the distribution if the Pip/I landscape is applied first.

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 have no information about which mutants are in the population. We can model this situation by taking as our prior population distribution

$$\mu = [1/2^N, \ldots, 1/2^N],$$

which specifies that each genotype is equally likely to constitute the starting population. If we

apply the drug F to this population we find the expected distribution is

$$\mu^* = \mu P_F^*,$$

which is shown in red in Figure 3. As expected, if we do not know the initial population then any of the local optima of F have a chance of arising through evolution. In particular, we have no guarantee that the highly resistant genotype $[0,1,1,0,0]$ will not arise in the population. Next, suppose instead we apply Pip/I first. In this case, after the application we will have the expected population distribution

$$\mu^* = \mu P_{pip}^* = [0,0,\ldots,1,0],$$

in which the population has been steered through natural selection to a single peak, $[1,1,1,1,0]$, in the Pip/I landscape. If we now apply the drug F after Pip/I we find that the distribution is

$$\mu^{**} = \mu^* P_F^* = (\mu P_{pip}^*) P_F^* = [0,0,\ldots,1]$$

which is shown as green in Figure 3. The population is guaranteed to evolve under F to genotype $[1,1,1,1,1]$, which is the least fit of the local optima in the landscape F (Figure 2(a)). Hence by applying Pip/I first we can prime the population such that the drug F can be applied without the possibility of the highly resistant genotype arising. Conversely, if we apply Cef first and F second then we have the final population distribution

$$\mu^{**} = (\mu P_{cef}^*) P_F^*,$$

shown as blue in Figure 3. In this case we have again primed the population such that some local optima of the landscape F are inaccessible when it is applied. However, in doing so we have made the highly resistant genotype $[0,1,1,0,0]$ more likely to emerge. It follows that through the selection of a sequence of drugs we can make the emergence of resistance more or less likely.

Discussion

We have encoded Strong Selection Weak Mutation evolutionary dynamics on fitness landscapes as a Markov Chain model. Through this encoding we can explore the dynamics of evolution by considering the algebraic properties of the associated transition matrix. In particular, we have demonstrated that evolution on fitness landscapes is non-commutative through parallels with the non-commutativity of 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 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 steering drugs in the form of antibiotics which no longer are considered useful as they have little efficacy on their own and have gone unused for many years.

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. At present, only very few of these landscapes are known for model organisms and determining fitness landscapes is difficult (we must check the fitness of all 2^N possible mutational variants with respect to each of the drugs). Furthermore, the landscapes can be dependent on the disease microenvironment and could change from patient to patient or throughout the course of treatment. 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 and hence determining the landscapes is not an entirely intractable problem.

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. Hence, we may be inadvertently selecting for highly resistant disease populations through irresponsible drug ordering in the same way that highly resistant disease can emerge through irresponsible drug dosing. 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.

The Strong Selection Weak Mutation model we have used here is a highly simplified, but well studied model of evolution. The model ignores much of the complexity of the evolutionary process, specifically simplifying the genotype–phenotype map and ignoring the disease microenvironment and the possibility of heterogeneous disease populations. 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]. Further, we have ignored several mechanisms of (epi-)genomic change such as deletion, aneuploidy and double mutations. However, each of these omissions 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 the fitness landscapes provides sufficient information to correctly identify optimal drug orderings *in vivo* is a question that cannot be answered through mathematical modelling alone and in the absence of experimentally determined landscapes cannot at present be answered at all. However, if at least some approximation to the fitness landscapes of the most used antibiotics were known then they would provide a good heuristic for how to proceed with multi-drug treatments and future modelling efforts to support 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.

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.

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

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

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

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

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

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

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.

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 combinatory 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.

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

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

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

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

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

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

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

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

Frank J Poelwijk, Daniel J Kiviet, Daniel M Weinreich, and Sander J Tans. Empirical fitness landscapes reveal accessible evolutionary paths. *Nature*, 445(7126):383–386, 2007.

Frank J Poelwijk, Sorin Tănase-Nicola, Daniel J Kiviet, and Sander J Tans. Reciprocal sign epistasis is a necessary condition for multi-peaked fitness landscapes. *Journal of Theoretical Biology*, 272 (1):141–144, 2011.

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

Niel Woodford and Matthew J Ellington. The emergence of antibiotic resistance by mutation. *Clinical Microbiology and Infection*, 13(1):5–18, 2007.

Slade O Jensen and Bruce R Lyon. Genetics of antimicrobial resistance in staphylococcus aureus. *Future Microbiology*, 4(5):565–582, 2009.

Christopher J Lord and Alan Ashworth. Mechanisms of resistance to therapies targeting brca-mutant cancers. *Nature Medicine*, 19(11):1381–1388, 2013.

Piro Lito, Neal Rosen, and David B Solit. Tumor adaptation and resistance to raf inhibitors. *Nature Medicine*, 19(11):1401–1409, 2013.

Steffen Schaper, Iain G Johnston, and Ard A Louis. Epistasis can lead to fragmented neutral spaces and contingency in evolution. *Proc Biol Sci*, 279(1734):1777–83, 2012.

Steffen Schaper and Ard A Louis. The arrival of the frequent: How bias in genotype-phenotype maps can steer populations to local optima. *PLOS ONE*, 9(2):e86635, 2014.