

Stochasticity in the Genotype-Phenotype Map: Implications for the Robustness and Persistence of Bet-Hedging

Daniel Nichol,^{*,†,1} Mark Robertson-Tessi,[†] Peter Jeavons,^{*} and Alexander R.A. Anderson^{†,1}

^{*}Department of Computer Science, University of Oxford, OX1 3QD, United Kingdom, and [†]Department of Integrated Mathematical Oncology, H. Lee Moffitt Cancer Center and Research Institute, Tampa, Florida 33612

ABSTRACT Nongenetic variation in phenotypes, or bet-hedging, has been observed as a driver of drug resistance in both bacterial infections and cancers. Here, we study how bet-hedging emerges in genotype–phenotype (GP) mapping through a simple interaction model: a molecular switch. We use simple chemical reaction networks to implement stochastic switches that map gene products to phenotypes, and investigate the impact of structurally distinct mappings on the evolution of phenotypic heterogeneity. Bet-hedging naturally emerges within this model, and is robust to evolutionary loss through mutations to both the expression of individual genes, and to the network itself. This robustness explains an apparent paradox of bet-hedging—why does it persist in environments where natural selection necessarily acts to remove it? The structure of the underlying molecular mechanism, itself subject to selection, can slow the evolutionary loss of bet-hedging to ensure a survival mechanism against environmental catastrophes even when they are rare. Critically, these properties, taken together, have profound implications for the use of treatment-holidays to combat bet-hedging-driven resistant disease, as the efficacy of breaks from treatment will ultimately be determined by the structure of the GP mapping.

KEYWORDS evolution; bet-hedging; genotype–phenotype map; bacterial persistence; drug resistance

TREATMENT resistance in many diseases is driven by the pre-existence of resistant phenotypes within the population. Why such phenotypes coexist (with sensitive phenotypes), and persist in environments never exposed to drug treatment, remains a significant unanswered question. Phenotypic heterogeneity has been observed within isogenic populations of a number of organisms, and at many scales (Balázs *et al.* 2011), from the unicellular—bacteria (Veening *et al.* 2008), fungi (Levy *et al.* 2012), or cancer cells (Gupta *et al.* 2011)—through insects (Danforth 1999; Hopper

1999), plants (Childs *et al.* 2010), and even aspects of human development (Tonegawa 1983). Importantly, this intercellular variation has been observed even in homogeneous and constant environments, suggesting that aspects of organismal phenotype may be stochastically determined.

In environments that fluctuate unpredictably, this phenomenon can serve as a survival mechanism by increasing the likelihood that at least some offspring are well-adapted to future environments. Thus, nongenetic, nonenvironmentally-driven variation in phenotypes has been termed *bet-hedging*, as a species diversifies the phenotypes within the population in order to “hedge its bets” against environmental change [see Seger (1987) for justification of this naming, and de Jong *et al.* (2011) for a discussion of what evolutionary phenomena can be considered bet-hedging]. Oscillatory environments are common in a range of ecological settings, including fluctuating climates, immune–pathogen interactions, or cyclic hypoxia within tumors, and the range of phenotypic traits that are thought to display stochastic determination is just as broad.

Bet-hedging can offer a survival mechanism in the event of rare catastrophic environmental change. An important clinical

Copyright © 2016 Nichol *et al.*

doi: 10.1534/genetics.116.193474

Manuscript received July 6, 2016; accepted for publication October 6, 2016; published Early Online October 20, 2016.

Available freely online through the author-supported open access option.

This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

Supplemental material is available online at <http://www.genetics.org/cgi/content/full/genetics.116.193474/DC1>.

¹Corresponding authors: Department of Computer Science, University of Oxford, Wolfson Building, Parks Rd, Oxford OX1 3QD, United Kingdom. E-mail: daniel.nichol@cs.ox.ac.uk; and Department of Integrated Mathematical Oncology, H. Lee Moffitt Cancer Center and Research Institute, SRB-4, 12902 Magnolia Drive, Tampa, FL 33612. E-mail: Alexander.Anderson@Moffitt.org

example is that of persister cells that arise stochastically within isogenic populations of infectious bacteria such as *Escherichia coli* (Balaban *et al.* 2004; Lewis 2006; Veening *et al.* 2008). These cells, which constitute a small fraction of the population [$< 1\%$, Lewis (2006)], have reduced metabolism, and shut down nonessential cellular functions. In this dormant state, the persister cells are tolerant to the cytotoxic effects of a number of antibiotic agents. Although dormant, these cells can retain the ability to proliferate (although at a drastically reduced rate), and, when antibiotic treatment ceases, persisters will proliferate, producing non-persisters and driving the re-emergence of the bacterial population. Hence, bet-hedging, by creating a small subpopulation impervious to those therapies that act on proliferating cells, proves to be an effective survival mechanism against antibiotic treatment. Indeed, bacterial persisters are thought to be a contributing factor to multidrug resistance in a number of diseases (Keren *et al.* 2004; Lewis 2006; Nikaido 2009), and are implicated in the dormancy of chronic diseases, such as tuberculosis, which can be suppressed but not eradicated (Zhang *et al.* 2012). Novel treatment strategies capable of effectively killing persister cells are desperately needed, and this need will continue to grow with the increasing incidence of resistance to our presently most effective antibiotics.

In cancer, bet-hedging has been minimally studied; however, a number of aspects of disease course suggest that bet-hedging mechanisms may be important for understanding how tumors evade therapy. Significant regression of tumors post-therapy leads to a period of remission, followed by the regrowth of aggressive, therapy-resistant lesions. These dynamics can be explained by the clonal model of cancer (Greaves and Maley 2012), wherein recurring drug-resistant cells are those that have stochastically acquired resistance mechanisms through genetic mutation. However, the high frequency of tumor recurrence in many cancers suggests that therapeutic escape cannot be based solely on mutational “luck.” Experimental results have shown evidence of transitory resistance (Kurata *et al.* 2004; Yano *et al.* 2005) indicative of the existence of a small drug-resistant subpopulation that re-establishes a drug-sensitive cancer cell population. Recent experiments have identified the existence of such populations of “cancer persister cells” in a cell line of EGFR+ nonsmall cell lung cancer (Sharma *et al.* 2010), indicating that bet-hedging may play a role in the emergence of cancer drug resistance (Ramirez *et al.* 2016). Thus, an understanding of bet-hedging in normal and abnormal (e.g., cancer) cell function may help us understand why certain types of therapies fail while others succeed.

Müller *et al.* (2013), as well as others (Thattai and Van Oudenaarden 2004; Kussell and Leibler 2005; Wolf *et al.* 2005), have demonstrated mathematically the selective advantage of bet-hedging strategies in stochastically fluctuating environments. Showing that fitness is maximized when the probability of individuals taking certain phenotypes

matches the likelihood of the environment selecting for that phenotype, provided that fluctuations are not sufficiently slow that adaptation through genetic mutation can occur, or so fast that no individuals of any phenotype can survive and reproduce. Further theoretical work by Botero *et al.* (2015) considers when bet-hedging can offer a greater fitness advantage than *phenotypic plasticity*, where phenotypes are modulated via the environmental variation (Via and Lande 1985). This previous work derives constraints on the cost of sensing, predictability of environmental fluctuations, and the fitness effects of environmental change to determine when bet-hedging, plasticity, or determinism offers a selective advantage.

It has been suggested that drug-insensitive cells that arise stochastically in an isogenic population can facilitate the emergence of genetically driven resistance by providing a window of opportunity in which resistance conferring mutations can arise Brock *et al.* (2009). However, this window is not indefinite as drug-insensitive cells will revert to a sensitive state, and likely die in the presence of a drug. Charlebois *et al.* (2011) explored this phenomenon through a mathematical model that incorporates switching from a drug-insensitive to a drug-sensitive phenotype as the stochastic relaxation from a state of high to low gene expression. This latter study demonstrated that the timescale of relaxation necessary to facilitate a high likelihood of genetic resistance is comparable to timescales measured for certain genes implicated in human cancers. In further work, Charlebois *et al.* (2014) introduced a model of a feed-forward transcriptional regulatory network to demonstrate that the network architecture can extend the time that drug-insensitive cells maintain their phenotype, and, thus, can increase the likelihood of therapeutic escape occurring through genetically driven mechanisms. This work highlights that, to understand bet-hedging-driven drug resistance, it may be necessary to look beyond the genetic scale, and toward the gene–gene interactions that comprise genotype–phenotype (GP) mapping.

The mathematical results of Müller *et al.* (2013) and others (Thattai and Van Oudenaarden 2004; Kussell and Leibler 2005; Wolf *et al.* 2005) suggest a paradox when compared to clinical observations of bet-hedging, for example, bacterial persistence, as a survival mechanism against rare catastrophic events. Specifically, in hospitable environments, bacterial persisters reproduce more slowly than cells with a proliferative phenotype, reducing population fitness. It follows that natural selection will act to remove or minimize the number of persisters in the population. Where catastrophic events are rare, we should expect bet-hedging strategies to be lost before the event occurs, resulting in extinction of the population when it eventually does. In this work, we suggest, following a similar argument to that provided by Charlebois *et al.* (2014), that the architecture of the molecular interactions networks may slow the evolutionary loss of bet-hedging to preserve it as a survival mechanism.

The causes of bet-hedging

A number of causes of bet-hedging have been identified across different species, but in many cases the cause remains an open question. The difficulty in identifying the precise drivers lies in distinguishing between the variability (or *noise*) introduced at different biological scales. For example, gene promoter, transcription, and translation dynamics are driven by inherently stochastic molecular interactions that result in the expression of gene products that vary both temporally and between isogenic individuals (Elowitz *et al.* 2002; Kaern *et al.* 2005; Blake *et al.* 2006; Raj and van Oudenaarden 2008). These gene products interact in nonlinear, molecular networks, often forming feedback loops that have the potential to suppress (Becskei and Serrano 2000), or amplify, noise (Hasty *et al.* 2000), or induce oscillations in the concentrations of molecules (Cardelli and Csikász-Nagy 2012; Cardelli 2014). This intracellular system is further modulated by variability in environmental factors and intercellular signaling.

Current biological thought is that noise in the levels of specific intracellular proteins may drive nongenetic phenotype differentiation. Indeed, under certain regimes of promoter switching, transcription, and translation, protein abundance can follow a bimodal distribution (Kaern *et al.* 2005), inducing two distinct phenotypes in a population. Further, theoretical modeling, coupled with experimental validation, has highlighted how bistable autoregulatory genetic motifs can induce bimodal protein distributions (Hasty *et al.* 2000; Becskei *et al.* 2001; Isaacs *et al.* 2003).

The evolutionary origin of bet-hedging is unclear. It is not known whether bet-hedging emerged as an adaptation to unpredictable environments, or as a spandrel (Gould and Lewontin 1979), arising from the inherently stochastic nature of the biochemical reactions governing cellular behavior, and later co-opted as a survival mechanism. What is clear is that bet-hedging strategies, manifested as phenotypically heterogeneous populations, are subject to natural selection. Beaumont *et al.* (2009) demonstrated the *de novo* evolution of bet-hedging in the phenotypic trait of colony morphology of the bacterium *Pseudomonas fluorescens* by imposing stochastically fluctuating environments through replating. The genetic driver underpinning this switching behavior was partially elucidated by Gallie *et al.* (2015), who identified a single nucleotide change in the gene *carB* as responsible for the emergence of phenotype switching; however, identifying the precise molecular pathways through which this mutation acts to produce bet-hedging remains an open problem. Following the recent development of persister isolation techniques, a number of genetic drivers thought to contribute to bacterial persistence in *E. coli* have been identified (Lewis 2006). However, while overexpression, or deletion, of these genes were shown to impact the *proportion* of bacterial persisters within a population, none was found to completely inhibit the persister phenotype, suggesting that bet-hedging may arise from the interactions of multiple gene products.

To address the difficulty in identifying genetic drivers of bet-hedging, we introduce a model GP mapping, wherein phenotypes emerge through the stochastic interactions of proteins in intracellular molecular pathways. Specifically, we simulate minimal interaction networks encoding bistable switches among two or three chemical species. The dynamics of these switches are implemented as chemical reaction networks (CRNs) simulated stochastically. Through this model we explore bet-hedging, which arises naturally, from the perspective of *network intrinsic noise*, as opposed to the *gene intrinsic* perspective of previous modeling (Figure 1). Under this model, stable configurations of a reaction network are analogous to local minima in an epigenetic landscape [for example, as used by Huang (2009, 2012) to study phenotypic heterogeneity in cancer].

There exist a number of regulatory motifs that induce bistable switching; however, the evolutionary implications of structural differences between these motifs have remained unexplored. Specifically, what properties of the molecular networks are likely to be selected for, and whether such properties can be exploited to identify potentially novel therapies, remain open questions. By considering simple, minimalistic models of bistable switches built from direct and indirect feedback mechanisms, we demonstrate that the structure of the networks governing phenotype differentiation can result in bet-hedging that is robust to major alterations. We argue that this robustness offers a potential explanation for the difficulty in identifying single genetic drivers of bet-hedging. Further, we demonstrate that network structure can alter the rate of evolutionary convergence to fitness optima, reducing evolvability and preventing the loss of bet-hedging in fixed environments. This result suggests a solution to the apparent paradox of bet-hedging: how can it persist for long periods of time in environments where natural selection acts to remove it? Finally, we discuss the implications of this result for the design of treatments for diseases which display nongenetic phenotypic heterogeneity.

Materials and Methods

In this work we implement model GP mappings that stochastically determine one of two phenotypes, *A* or *B*, from a single genotype, *g*, encoding gene expression. The GP mapping can be considered to determine a genotype-dependent probability, $p(g)$, of an individual taking phenotype, *A*. We aim to investigate the rate of evolutionary loss of bet-hedging by calculating the probability that an individual mutant invades an otherwise isogenic population. This probability is dependent on the distribution of phenotypes in the resident population. As such, this section is structured as follows. First, we derive the distribution of phenotypes, and average growth rate of a population in which each individual has probability, p , of having phenotype *A* at birth (and $1 - p$ of having phenotype *B*). Second, we derive the probability of a single mutant individual with probability p' (corresponding to genotype g') invading this resident population. Third, we outline the model for the GP map,

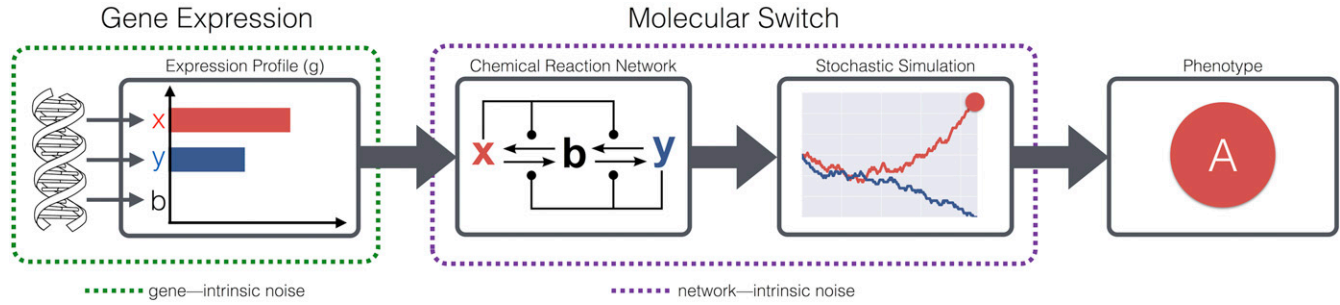


Figure 1 Schematic representation of the CRN model for determining phenotypes from genotypes. The gene expression profiles (g) are assumed to be fixed for each genotype, and the dynamics of expression from the biological genotype (green dashed leftmost box) are ignored. These modeling assumptions allow us to explore the implications of network-intrinsic noise (purple dashed centered box) independently of gene-intrinsic noise.

wherein $p(g)$ is determined from g by a bistable CRN. Finally, we describe simulations of long-term evolution and drug therapy that form the basis of our results.

Population dynamics

Assume a fixed GP map that determines a probability, p , of an individual with genotype g having phenotype A at birth. This phenotype is fixed throughout the life of the individual. As we wish to study evolutionary loss of bet-hedging, we may assume $p \in (0, 1)$, as when $p = 1$ we will end our simulations. We also assume a fixed environment. A discrete time model is used to simulate the population dynamics as follows. Denote by $\mathbf{x}(t) = [x_A(t), x_B(t)]^T$ the number of individuals of phenotype A and B at discrete timestep t . We assume Wright-Fisher sampling, wherein each individual in the population at time t can contribute any number of individuals to the population at time $t + 1$. Denote by w_A (respectively, w_B) the expected number of individuals in the population at time $t + 1$ that are descended from a single individual of phenotype A (respectively, B) present in the population at t . We assume $w_A, w_B > 0$. Each new offspring takes phenotype A with probability p (and B with probability $1 - p$), where p is determined by the genotype, g , and the GP mapping.

For a fixed population size, the dynamics of the population can be modeled using mathematics from the theory of quasispecies that describe a population in mutation–selection balance (Wilke 2005). Specifically, if the average number of individuals at time $t + 1$ produced by a single individual at time t is given by $\langle w \rangle$, the population dynamics are governed by the projection matrix,

$$\mathbf{P} = \frac{1}{\langle w \rangle} \begin{bmatrix} w_A p & w_B p \\ w_A(1-p) & w_B(1-p) \end{bmatrix}. \quad (1)$$

The population distribution after one discrete time step is given by $\mathbf{x}(t+1) = \mathbf{P}\mathbf{x}(t)$. As $p \in (0, 1)$, the matrix \mathbf{P} is positive, and the Perron-Frobenius theorem (Cushing 1998; Li and Schneider 2002) tells us that the normalized eigenvector corresponding to the dominant eigenvalue of \mathbf{P} gives the long-term stationary distribution of the two phenotypes A and B . As we have only two phenotypes, we can easily determine this dominant eigenvalue as

$$\begin{aligned} 0 = |\mathbf{P} - I\lambda| &= \begin{vmatrix} \frac{w_A p}{\langle w \rangle} - \lambda & \frac{w_B p}{\langle w \rangle} \\ \frac{w_A(1-p)}{\langle w \rangle} & \frac{w_B(1-p)}{\langle w \rangle} - \lambda \end{vmatrix} \\ &= \left(\frac{w_A p}{\langle w \rangle} - \lambda \right) \left(\frac{w_B(1-p)}{\langle w \rangle} - \lambda \right) - \frac{w_A w_B p(1-p)}{\langle w \rangle^2} \\ &= \lambda \left(\lambda - \frac{w_A p}{\langle w \rangle} - \frac{w_B(1-p)}{\langle w \rangle} \right). \end{aligned}$$

Thus, $\lambda = \frac{1}{\langle w \rangle} [p w_A + (1-p) w_B]$. It is easy to verify then that $\mathbf{x}^* = (p, 1-p)$ is the normalized eigenvector corresponding to this eigenvalue as

$$\mathbf{P} \begin{pmatrix} p \\ 1-p \end{pmatrix} = \frac{1}{\langle w \rangle} \begin{bmatrix} p^2 w_A + p(1-p) w_B \\ p(1-p) w_A + (1-p)^2 w_B \end{bmatrix} \quad (2)$$

$$= \frac{1}{\langle w \rangle} [p w_A + (1-p) w_B] \begin{pmatrix} p \\ 1-p \end{pmatrix} \quad (3)$$

$$= \lambda \begin{pmatrix} p \\ 1-p \end{pmatrix} \quad (4)$$

Note that this is also the distribution of phenotypes in the case where the population grows without bound (when the factor of $1/\langle w \rangle$ is omitted). From this stationary distribution, we can derive $\langle w \rangle = p w_A + (1-p) w_B$. This phenotype equilibrium will be used in the next section to determine the invasion probability of a different (*i.e.*, mutant) bet-hedging population into an existing one.

In reality, the values w_A and w_B are dependent on a number of stochastic processes, most importantly reproduction and death. Later, we will consider the effect of an increased death rate associated with drug treatment, so it is informative to consider the relationship between parameters governing the explicit processes of birth and death, and the offspring numbers w_A and w_B . Suppose individuals of type A (respectively, B) die with probability d_A (respectively, d_B) over each time step. Further, suppose that individuals of phenotype A (respectively, B) that survive reproduce with probability f_A (respectively, f_B) over the timestep. In this case

$$w_i = (1 - d_i)(1 + f_i) \text{ for } i \in \{A, B\}. \quad (5)$$

Note that multiple pairs of values for f_i and d_i can yield the same w_i . The population dynamics we present here, and the invasion dynamics presented below, are identical for all such pairs.

Invasion dynamics

Our aim is to determine the long-term evolutionary dynamics of bet-hedging populations endowed with differing GP-maps. It is intractable to determine these trajectories through explicit simulation alone. Instead we derive an analytic solution for the probability of a mutant genotype invading an existing isogenic population. Consider a large fixed-size population, and assume that mutation is sufficiently rare (explicitly that the mutation rate μ and population size N satisfy $N\mu\log N \ll 1$) that we may consider strong-selection weak mutation (SSWM) evolutionary dynamics (Gillespie 1983, 1984). Under these assumptions, we can assume that the population is isogenic, and that each time a new mutant appears in the population it either fixes as the new population genotype or becomes extinct before another can arise.

Suppose a single mutant of genotype g' arises in an isogenic population of genotype g and denote by $\pi(g')$, the probability that this mutant reaches fixation as the population genotype. This probability is dependent on the phenotype of this initial mutant, and is given by

$$\pi(g') = p(g')\pi(g'|A) + [1 - p(g')]\pi(g'|B). \quad (6)$$

Denote $\boldsymbol{\pi} = [\pi(g'|A), \pi(g'|B)]$, and suppose that the population size, N , is sufficiently large that we may approximate it by the limit $N \rightarrow \infty$. As we assume Wright-Fisher sampling for reproduction, the value of $\boldsymbol{\pi}$ can be determined from the theory of branching processes. In particular, $\boldsymbol{\pi}$ can be calculated numerically as the solution to the equation

$$1 - \boldsymbol{\pi} = e^{-\mathbf{P}\boldsymbol{\pi}} \quad (7)$$

where \mathbf{P} is the matrix governing the population dynamics defined above.

A proof of this identity, modified from the theory of viral quasispecies (Wilke 2003), is presented in Supplemental Material, [File S1](#). As the value of $\langle w \rangle$ can be calculated from the population dynamics described above, this equation can be solved numerically. Figure 2 shows an example heatmap of invasion probability for an invader with genotype corresponding to probability p_2 of phenotype A into a resident population with probability p_1 . The parameters are $w_A = 2.0$ and $w_B = 1.01$.

CRNs as a model GP map

To study the evolution of bet-hedging, we consider the genetic drivers of changes to the probability, p , of an individual having phenotype A. We implement a model GP mapping in which phenotypes emerge with proportions that are determined from the interactions of expressed gene products. In this model, the genotype, g , is represented in an abstracted way,




Figure 2 Invasion probabilities for a single mutant with genotype corresponding to a probability p_2 of producing phenotype A into a resident population with probability p_1 . The probabilities are calculated from Equations 6 and 7 with parameters $w_A = 2.0$, $w_B = 1.01$. Note that deleterious and neutral mutations cannot fix under our model of invasion dynamics; hence, invasion in the case $p_2 \leq p_1$ (above the antidiagonal of the plot) is impossible.

as the numbers of chemical species, labeled x and y , that are present in the cell at birth. Thus, $g = (x_0, y_0) \in \mathbb{N}^2$.

The model relies on the stochastic resolution of a CRN through the Gillespie algorithm (Gillespie 1977, 2001) to determine a phenotype from the genotype g . CRNs are defined by a collection of labeled chemical species and a list of reactions, with associated rates, between these species. The Gillespie algorithm determines a stochastic progression of reactions within a CRN, and returns the sequence of reactions that occur, along with the times at which they occur. We consider the class of CRNs that encode bistable switches (Cardelli and Csikász-Nagy 2012), wherein the sequence of reactions will almost surely (in the probabilistic sense) terminate in one of two stable configurations (see Figure 3). These different final configurations can be considered different states of a stochastic switch, and to represent the configurations that ultimately result in different phenotypes, A and B. The probability that the CRN progresses to a specific switch state is dependent on the initial conditions for the network. Thus, we can define a model GP mapping according to

$$g \xrightarrow{\text{GP}} \begin{cases} A, & \text{if the simulation progresses to stable state 1} \\ B, & \text{if the simulation progresses to stable state 2.} \end{cases} \quad (8)$$

A schematic representation of this model is shown in Figure 1. In this model, an individual has a phenotype that is determined soon after birth and fixed until that individual reproduces. This model of the GP map can be considered to be a stochastic, irreversible developmental program in the case of differentiating cancer cells or higher organisms. In the

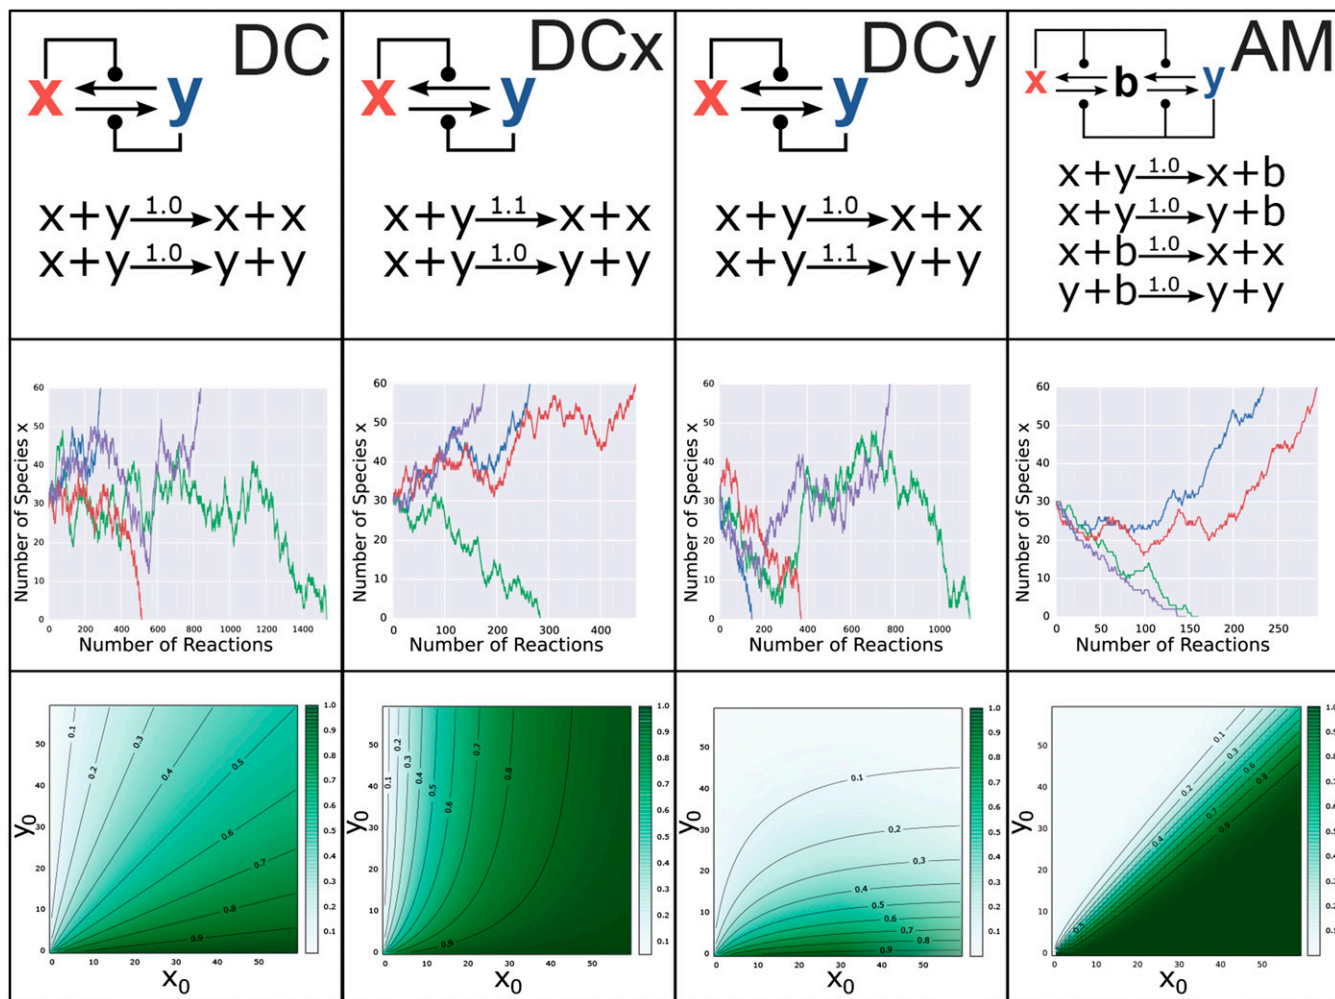


Figure 3 Example molecular switches as GP maps. Each column shows the characteristics of one of the four switches (DC, DCx, DCy, and AM) introduced in the main text. The first row shows the name, CRN structure, and precise definition of each switch. The second row shows stochastic trajectories of the number of molecule x in the system for four different simulations of each switch. The starting condition in all simulations is $x = y = 30$, ($b = 0$ for the AM network). Note that all of the switches are able to resolve to either of the stable conditions, $x = 60$ or $x = 0$, which correspond to the phenotypes A and B , respectively. Row three shows contour plots displaying the probability of switching to phenotype A for each possible initial condition with $0 < x_0, y_0 \leq 60$ and $x_0 + y_0 \neq 0$ ($b = 0$ for the AM switch). Contour lines show subspaces of genotype space of equal hedging probability for hedges equal to 0.1, 0.2, ..., 0.9.

context of bacteria, the developmental perspective is less appropriate, and the model can be justified by considering gene expression bursts that are cell-cycle dependent (specifically, bursts that occur at the start of the G1 phase).

Note that the genotype g corresponding to a stable gene expression profile is the sole heritable determinant of phenotype in this model. The state of the molecular switch, *i.e.*, the phenotype, of a parent individual has no influence on the phenotype of the offspring. This assumption can be justified as follows. The bistable switches we discuss may represent only a small subnetwork of the complex and dynamic interaction network governing the GP mapping. Thus, we can expect the chemical species comprising the molecular switch, x and y , to be further transformed, or consumed, in additional unmodelled reactions that determine phenotype. Further, even when this is not the case, we can expect the x and y molecules to

decay over time. Thus, the omission of epigenetic inheritance of switch state can be considered an assumption that the time scale of molecular decay is much faster than that of cellular division. Weakening this assumption, and permitting epigenetic inheritance, represents a potential extension of the model that is briefly explored in our *Discussion*.

Finally, we assume that the series of chemical reactions that result in a stable configuration for the network all occur within a sufficiently short time period (in comparison to the cell cycle) that we may take them to have all occurred instantaneously. This assumption permits us to ignore the timing information provided by the Gillespie algorithm. As such, the abscissas of all figures showing stochastic simulations of CRNs presented in this work measure time discretely, in terms of the number of reactions that have occurred, instead of continuously. The time between successive reactions in the

Gillespie algorithm is dependent on a propensity function that accounts for reaction rates and the volume of the container (*i.e.*, cell cytoplasm). As we are only interested in the probability of finding each stable configuration, and not the precise time taken to reach this configuration, we may set the volume to be an arbitrary constant, say 1.0 mm^3 , and also normalize the reaction rates such that at least one reaction has rate 1.0 sec^{-1} .

Four model bistable switches that can serve as model GP maps are shown in Figure 3, along with examples of stochastic realizations of the switches (second row), and heatmaps showing $p[(x_0, y_0)] = \mathbb{P}[g \rightarrow \text{GPA}]$ for $x_0, y_0 \in \{0, \dots, 60\}$ (third row). The value of $p[(x_0, y_0)]$ is required to determine the population and invasion dynamics as described above. Estimating this value numerically through multiple samples of the Gillespie algorithm is prohibitively slow. Instead, the value can be determined analytically for the DC, DCx, and DCy switches as they correspond to the classical “drunkard’s walk” of probability theory. For the AM switch, no such analytic solution is possible. To determine $p[(x_0, y_0)]$ in this case, we construct the Markov chain on the space of possible configurations of (x, y, b) and numerically solve for $p[(x_0, y_0)]$. The details are provided in [File S1](#).

Simulating evolutionary loss of bet-hedging

To investigate the impact of GP mapping on the evolutionary loss of bet-hedging, we implement a stochastic simulation of mutation and selection. We consider mutations to a genotype $g = (x_0, y_0)$ as changes to the initial abundances x_0 and y_0 . The possible mutations are thus modeled by

$$\mu[(x_0, y_0)] = \{(x_0 \pm 1, y_0 \pm 1) \mid \text{provided } x_0 + y_0 = g_{\max}\}. \quad (9)$$

Note that we have restricted mutations such that total expression is conserved, and the genotype is determined entirely by the value $x_0 (y_0 = g_{\max} - x_0)$. For the remainder of this work, we omit reference to y_0 , and equate g with x_0 . Owing to the computational complexity of our simulations, we take $g_{\max} = 60$. Changes to the value of g_{\max} do not change the qualitative results, but will change them quantitatively as the time of evolutionary convergence to a nonhedging strategy increases as g_{\max} is increased (see [File S1](#)).

The mutations defined by Equation 9 differ from previous network models of the GP mapping; for example, the models of Huang (2009) or Gerlee and Anderson (2009), as they modify the initial conditions of a network-defined dynamical system, as opposed to the system itself. This choice of mutation is appropriate to the level at which genotypes are modeled. For example, the phenotypes of the gene regulatory network model studied by Huang (2009) are stable gene expression profiles. By contrast, in our model, expression profiles are taken as the genotypes, the initial expression levels of x and y . This notion of genotype is chosen to allow us to investigate the impact of network-intrinsic noise on the evolution of bet-hedging (see Figure 1).

As discussed above, we assume a large asexually reproducing bet-hedging population exists under SSWM dynamics (Gillespie 1983). Our evolutionary simulation proceeds by repeatedly generating a mutant of the population genotype g according to Equation 9, computing the probability π that this mutant fixes as the new population genotype according to Equation 7, and then stochastically deciding whether the mutation fixes [by sampling $q \sim \text{Unif}(0, 1)$ and comparing the value to π]. The simulation terminates when the genotype satisfies $x_0 = g_{\max}$, and the total number of mutations that are sampled in the simulation (including those that do not fix) is returned as a proxy for the time until evolutionary loss of the bet-hedging. The initial population genotype, g_0 , is chosen such that $p(g_0) \approx 0.5$.

Simulating treatment holidays

To explore the clinical impact of treatment holidays on treating disease with bet-hedging driven resistance, we implemented a simulation comprising two parts. In the first, an evolutionary simulation similar to the one described above is performed to determine the population genotype following a treatment holiday. From an initial genotype g_0 with $p(g_0) \approx 0.5$, the expected population genotype g following a treatment holiday of length T (measured in mutational events) is computed. The expected time, in mutational events, taken for the genotype $g + 1$ to arise by mutation and fix in the population is $1/[0.5\pi(g + 1)]$. Here, the factor of 0.5 arises as only half of mutations are beneficial ($g \rightarrow g - 1$ is not). Using this fact, an expected postholiday genotype can be easily determined by repeatedly incrementing g , while keeping a sum of the expected number of mutational events required for that new g to fix in the population. This process is terminated when this number of events exceeds T .

Using this postholiday genotype, and the associated GP-map-dependent hedging probability $p(g)$, a stochastic death–birth process without mutation was performed to determine an approximate time to extinction. A population of size 10^{10} was initialized with a proportion $p(g)$ of phenotype A and $[1 - p(g)]$ of phenotype B. The simulation proceeds as depicted in Figure 4, and is terminated when the population is extinct, or 20,000 simulated hr has elapsed. Each time step corresponds to 1 hr. Unlike the simulation of evolutionary timescales, the treatment simulation is dependent on the parameters of per-hour probability of death, d_i^{drug} , and reproduction f_i^{drug} , for individuals of each phenotype in the drug-treated environment. This environment is different to the hospitable environment taken for evolutionary simulation.

This simulation was performed 2000 times for each treatment holiday length $T \in (0, 3000, 5000, 50,000, \text{ and } 100,000)$, and each molecular switch from Figure 3. The extinction times were collated to form the histograms shown in Figure 8.

Data availability

All models described in this work were implemented in Python. All scripts are available from D.N. upon request.

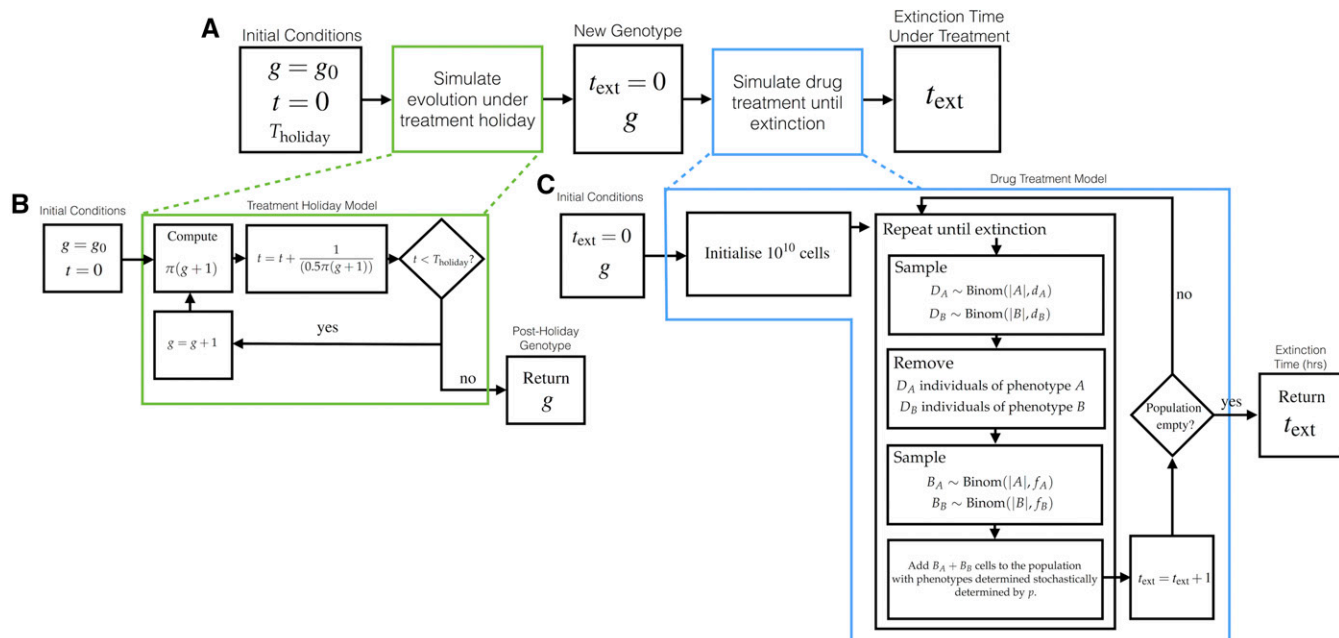


Figure 4 The model for simulating treatment holidays. (A) The overall model consists of first determining a post-treatment holiday genotype, g , and then simulating drug treatment on a population with bet-hedging determined by that genotype. (B) The expected postholiday genotype is determined from the invasion probabilities $\pi(g)$, which are, in turn, dependent on the molecular switch. (C) A stochastic death–birth process is used to determine an extinction time, t_{ext} .

Results

We have implemented a model GP mapping in which phenotypes emerge with proportions that are determined from a genotype and molecular switch. It is infeasible to explicitly model the full array of chemical interactions comprising the translation from genes to phenotypes. However, investigation of smaller network motifs can provide insight into the properties of the full molecular network. A similar approach was taken by Cardelli (2014), who studied emulation between CRNs, the phenomenon where one network is capable of reproducing the exact mass-action kinetics of another. Identifying emulations provides a method to extend results gained from studying of simple CRN motifs to larger molecular pathways.

We consider four different bistable switches that can serve as model GP mappings. These switches are constructed from minimal interaction encoding direct and indirect feedback among two or three chemical species, and represent the simplest possible implementations of bistable networks. The switches, along with examples of their dynamics, are presented in Figure 3. The *Direct Competition* (DC) switch (along with DCx, DCy) consists of a pair of autocatalytic reactions. The *Approximate Majority* switch, studied by Angluin *et al.* (2008) and later Cardelli and Csikász-Nagy (2012), consists of two catalytic and two autocatalytic reactions. A biological implementation of the AM switch is presented by Dodd *et al.* (2007) as a potential mechanism for epigenetic cell memory. By picking appropriate genotypes (*i.e.*, initial conditions for the molecular network), any switching probability, and,

equivalently, any ratio of phenotypes A and B, can be closely approximated using any of the switches (Figure 3, Row 3). It follows that bet-hedging can arise solely from network-intrinsic noise introduced by the stochastic interactions among as few as two chemical species. Here, we explore how the topology of simple stochastic networks influences the evolutionary fate of bet-hedging.

Robustness and redundancy in molecular switches

By introducing redundancy, we demonstrate how bet-hedging can be robust to major perturbations to the underlying network. Figure 5A shows a version of the DC switch from Figure 3 in which the species x and y are duplicated. In this network, which we call DCdup, the set of stable configurations are determined by $x + x' = 0$ or $y + y' = 0$. If we associate the phenotypes A and B with these two configurations, respectively, then the switching probability on initial conditions (x_0, x'_0, y_0, y'_0) is identical to the switching probability of DC with initial conditions $(x_0 + x'_0, y_0 + y'_0)$ (a simple mathematical argument to establish this proceeds by symmetry and relabeling the species). The potential benefit of DCdup is that it maintains its switching properties, even if chemical species are removed. Figure 5A shows numerical solutions for the CRN switching probability when the species x is deleted (middle network), and then when both x and y are deleted (right hand network). These induced CRNs maintain switching behavior similar to the original network DCdup. The network induced by deleting x (or by symmetry y) behaves precisely as DCdup with initial condition $(x_0, 0, y_0, y'_0)$ [by

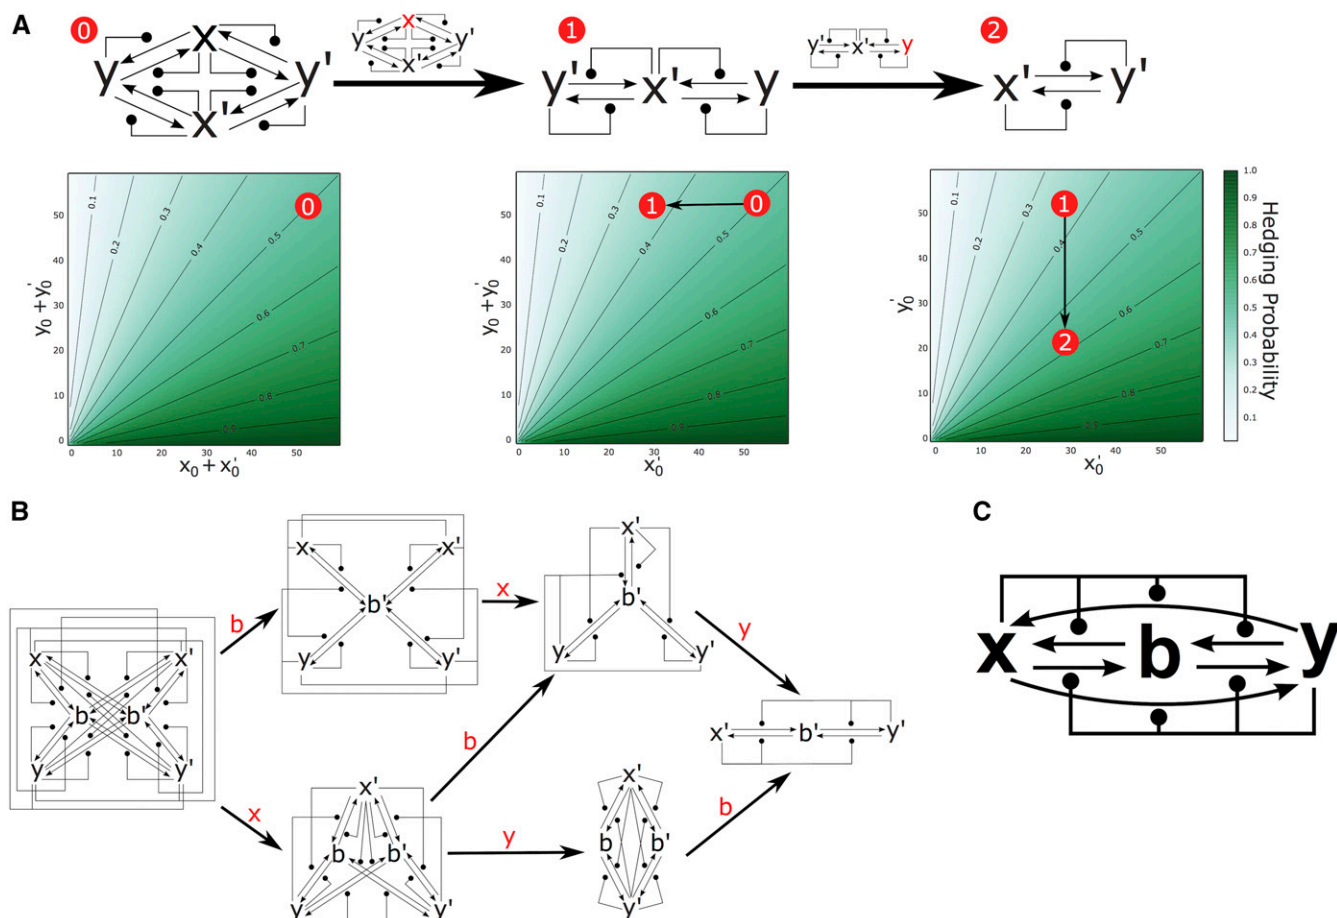


Figure 5 Redundancy results in bet-hedging that is robust to mutation. (A) Redundancy in the CRN implementing the DC switch maintains molecular switching when chemical species are deleted. Marked in red is the switching probability for initial conditions (20, 30, 30, 20) before deletion (0), after the deletion of x (1), and after the deletion of x and y (2). Contour lines show initial conditions of equal switching behavior. (B) Redundancy in the CRN implementing the AM molecular switch. Switching is maintained if the species x, y and b are removed in any order. We omit the case where y is removed before x due to symmetry. (C) A molecular switch that can reduce to either AM or DC when specific reactions are inhibited.

symmetry $(x_0, x'_0, 0, y'_0)$. Further, removing both x and y from DCdup creates a version of the DC switch in the species x' and y' that behaves precisely as the DCdup switch on initial conditions $(0, x'_0, 0, y'_0)$. It follows that deletion of chemical species will change the likelihood of switching (and thus the proportion of phenotypes A and B) if initial numbers of all other proteins remain fixed. An example is shown by the red circles in Figure 5A, where deletion of chemical species shifts the switching probability, but does not inhibit bet-hedging entirely. Hence, the DCdup switch is robust to the removal of chemical species—a mutational event that, in our model, can be interpreted as deletion or downregulation of a gene. Reversing this argument, the switching behavior of the CRNs in Figure 5A demonstrate how bet-hedging is robust to gene duplications or upregulating mutations.

A similar redundant implementation of the AM molecular switch is shown in Figure 5B, and is robust to the removal of species, x, y , and b in any order. As with the DCdup network, the removal of any chemical species will change the switching probability, and shift the proportions of phenotypes in the population, but will not inhibit one phenotype entirely. Fur-

ther, we note that the switching is not only robust to mutations that remove chemical species, but also to alterations in the rates of reaction between them. For example, each of the switches presented in Figure 3 can be derived from the larger network shown in Figure 5C by inhibition of specific reactions. Removal of the autocatalytic reactions $x + y \rightarrow x + x$ and $x + y \rightarrow y + y$ yields the AM switch. Alternatively, removing the four reactions involving the chemical species b yields the DC switch.

Evolutionary loss of bet-hedging

For fixed gene expression, the specific switching mechanism responsible for producing multiple phenotypes is irrelevant, as the proportion of phenotypes remains fixed. However, over longer timescales, the structure of the molecular switch has a significant impact on the evolution of bet-hedging. Throughout the remainder of this work, we take the two phenotypes A and B to correspond to a high fitness, proliferative phenotype, and a low fitness, slow proliferating phenotype, respectively, mirroring the phenomenon of bacterial persistence.

The invasion probability is computed independently of the (assumed to be large) population size, and is dependent only on the (stable) distribution of phenotypes in the population at equilibrium. This distribution can be computed from the values w_A and w_B . We parametrize the model determined as follows. The discrete timesteps are taken to be $t = 60$ mins, mirroring an expected division time of *E. coli*. The time from birth until reproduction for a phenotype *A* individual is distributed exponentially, with rate parameter $\lambda_A = 1.0$ per 60 min. The number of reproductions of a single individual is then Poisson distributed with parameter $\lambda_A = 1.0$. Hence, on average, each phenotype *A* individual present at time t reproduces once in the 60 min, and produces two individuals at time $t + 1$ (itself and its offspring). It follows, $w_A = 2.0$. The reproductive rate of persister-type cells is unknown. To match their behavior qualitatively, we take our persister-like phenotype *B* cells to have an expected reproduction time 100 times slower than phenotype *A* individuals, at 6000 min. Thus, reproduction time for phenotype *B* cells is distributed exponentially with rate parameter $\lambda_b = 0.01$ per 60 min. The expected number of reproductions over the 60 min timestep is then 0.01 and $w_B = 1.01$. The environment remains fixed (w_A, w_B are unchanged) throughout the evolutionary simulation. Although this parametrization is only an approximation, it is sufficient as an illustrative model demonstrating the importance of the structural properties of GP mapping. The effects of changing w_A and w_B are discussed in [File S1](#).

Figure 6A shows how changes in the population genotype manifest themselves as changes in the average population fitness. The expected population fitness increase associated with a mutation from x_0 to $x_0 + 1$ is not constant, and instead is dependent on the underlying molecular switch. As a result, invasion probabilities are dependent on both the population genotype, invading genotype, and the form of the molecular network. Figure 6C shows the probability of a single mutant genotype, x'_0 , invading a resident population of genotype x_0 . Note that only mutations that increase the proportion of phenotype *A* are beneficial, and, hence, as our invasion probabilities are determined from the theory of branching processes, are the only mutations that can fix. Figure 6B shows the probability of successive beneficial mutations, $x_0 + 1$, invading a resident population of genotype x_0 . In this figure, we see the impact of the GP map on the evolutionary dynamics. For the DC, DCx, and AM switches, the probability of the next beneficial mutation fixing reduces for each successive mutation. The magnitude of this decrease is dependent on the switch, and, in the case of DCx and AM, approaches zero. Conversely, for the DCy switch, each successive mutation is more likely to fix.

Consider the evolutionary trajectories of phenotype proportions determined by each of the molecular switches from an initial gene expression profile corresponding closely to a population consisting of $\sim 50\%$ of each phenotype. As the DC and AM switches are symmetric, the genotype corresponding to a 0.5 hedging probability is $x_0 = 30$. For DCx, the closest genotype to a 50% hedge is $x_0 = 7$, which corresponds to a prob-

ability of 0.49. For DCy, the closest genotype to a 0.5 hedging probability is $x_0 = 53$, and corresponds to a probability of 0.51. As deleterious and neutral mutations cannot fix under our model of population dynamics, the population genotype will be periodically incremented until $x_0 = 60$, at which point the bet-hedge is lost. Figure 7 shows the evolutionary trajectories toward this loss of bet-hedging, highlighting considerably different convergence dynamics. For the DC, DCx, and DCy switches, the expected convergence times can be determined as the expectation of a sum of nonidentical independent geometric distributions. We find that the expected number of mutational events required for a complete loss of bet-hedging, as highlighted in Figure 7, are given by

$$\begin{aligned}\mathbb{E}[\text{time to loss of bet-hedge} | \text{DCy switch}, x_0 = 53] &= 195.8 \\ \mathbb{E}[\text{time to loss of bet-hedge} | \text{DC switch}, x_0 = 30] &= 3382.8 \\ \mathbb{E}[\text{time to loss of bet-hedge} | \text{DCx switch}, x_0 = 7] &= 68,000.1 \\ \mathbb{E}[\text{time to loss of bet-hedge} | \text{AM switch}, x_0 = 30] &= \infty\end{aligned}$$

In the case of the AM network, each subsequent mutation provides a diminishing increase in fitness until mutations are approximately neutral. The probability of neutral mutations fixing within our model of invasion dynamics, which models the population size as tending to infinity, is zero. In reality, the actual convergence times in the AM will depend on the population size. For large populations, as is our assumption, the timescales will be sufficiently long that we take it as equivalent to the evolutionary trajectory never converging. (This assumption can be justified by the observation that over these time scales either unmodelled mutations [such as mutations to the GP map itself, to other genes governing the phenotypes *A* and *B*, or to other aspects of the phenotype], or unmodelled changes in the environment or ecosystem, will occur, rendering our model unsuited to the situation.)

Simulation of therapeutic intervention

To demonstrate the importance of the underlying molecular switch on the efficacy of treatment holidays for diseases with bet-hedging-driven resistance, we performed a two-part simulation. First, a treatment holiday was simulated through an evolutionary simulation in a hospitable environment to determine a postholiday population genotype. A nonspatial, individual-based model was then used to simulate drug treatment. The parameter values used in the evolutionary simulation of treatment holidays were chosen to coincide with the values $\mathbf{w} = [w_A^{\text{hosp}}, w_B^{\text{hosp}}] = [2.0, 1.01]$ used in our simulation of evolutionary loss of bet-hedging. For the simulation of treatment, explicit values of f_i^{drug} and d_i^{drug} (for $i \in \{A, B\}$) are required. First, we determined the values for the nondrug environment that correspond to those used in the evolutionary simulation. For the hospitable environment, we assumed a fixed death probability for both *As* and *Bs* of $d_A^{\text{hosp}} = d_B^{\text{hosp}} = 0.005$. The per-hour likelihood of reproduction for the surviving individuals was taken as $f_A^{\text{hosp}} = 1.0$ and $f_B^{\text{hosp}} = 0.015$. These parameters correspond to overall expected offspring for each phenotype

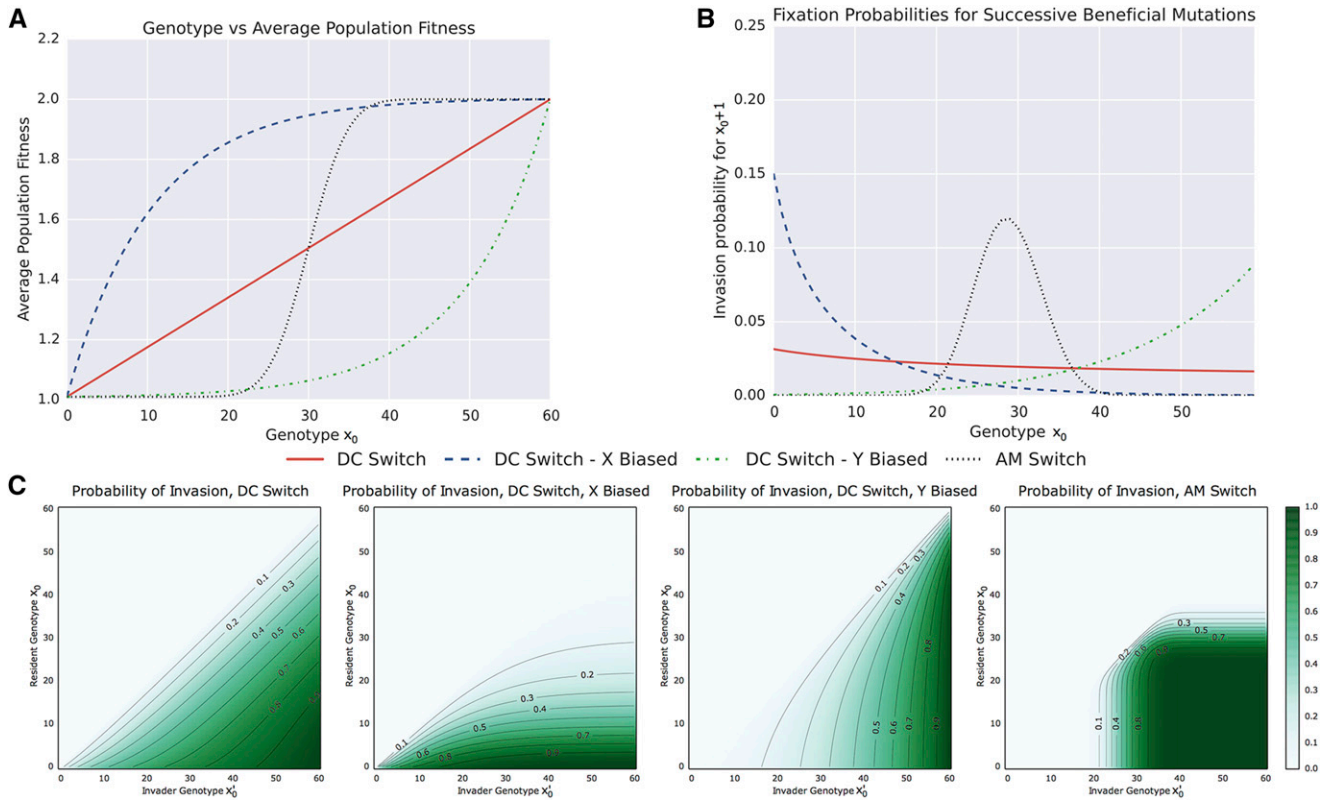


Figure 6 GP mapping determines the dynamics of invasion. (A) The relationship between population genotype, x_0 , and average population fitness for each molecular switch. (B) Invasion probabilities for a new genotype $x_0 + 1$ into a resident population of genotype x_0 . (C) Invasion probabilities for resident and invader genotypes.

$$\begin{aligned} [w_A^{\text{hosp}}, w_B^{\text{hosp}}] &= \left[(1 - d_A^{\text{hosp}}) (1 + f_A^{\text{hosp}}), (1 - d_B^{\text{hosp}}) \right. \\ &\quad \left. \times (1 + f_B^{\text{hosp}}) \right] \approx [2.0, 1.01]. \end{aligned}$$

We consider three regimes of drug–environment parameter sets corresponding to a purely cytostatic drug, affecting only f_A , a purely cytotoxic drug, affecting only d_A , and a drug that is a mixture of the two, affecting both f_A and d_A intermediately. In the following discussion the “drug” superscript is omitted for readability.

For each regime we assume that the parameters for phenotype B , mirroring a drug-impermeable persister-like phenotype, are unchanged. Consider a fixed drug parameter set. For each switching network, the expected postholiday population genotype, g , was calculated for each treatment holiday length $T \in (0, 3000, 5000, 50,000, 10,000)$ from initial genotype g_0 chosen, as before, to correspond closely to an equal proportion of each phenotype. From each g , we then performed 2000 simulations of drug treatment to determine a distribution of extinction times. Where simulation time exceeded 20,000 hr, the simulation was halted and extinction was determined to not occur. A standard course of antibiotic treatment lasts between 7 and 10 days. Using this period as a guide, we defined extinction times of <240 hr as corresponding to a potentially successful course of therapy, and thus, a viable treatment holiday strategy. The extinction

time histograms lying within this period are shaded green and marked by a * in the histograms. Mutations were not modeled during the simulations of therapy.

The cytostatic drug parameter set was taken to correspond to a substantial reduction in the proliferative rate of phenotype A individuals. We note that extinction will almost surely never occur if

$$\begin{aligned} \langle w^{\text{drug}} \rangle &= p(g)w_A^{\text{drug}} + [1 - p(g)]w_B^{\text{drug}} = p(g)(1 - d_A) \\ &\quad \times (1 + f_A) + [1 - p(g)](1 - d_B)(1 + f_B) \geq 1. \end{aligned}$$

As $(1 - d_B)(1 + f_B) > 1.0$, it follows that a cytostatic drug can only drive extinction in this case if $(1 - d_A)(1 + f_A) < 1.0$. As $d_A = 0.005$, this occurs only if $f_A < 0.995^{-1} - 1 \approx 0.005$. For our simulation of cytostatic drug treatment, we assume that the drug is entirely effective at inhibiting reproduction, $f_A = 0$. For this parameter, we find that extinction never occurs within the 240 hr time period, indicating that, regardless of the underlying switching mechanism, no length of treatment holiday will result in complete cure with a cytostatic drug. The histograms of extinction times, where extinction does occur, are presented in Figure 3 of File S1.

For the cytotoxic drug parameter set, we consider an increased death probability of $d_A = 0.8$, holding all others the same. The histograms of extinction times are presented in Figure 8. Here, we see that the length of treatment holiday

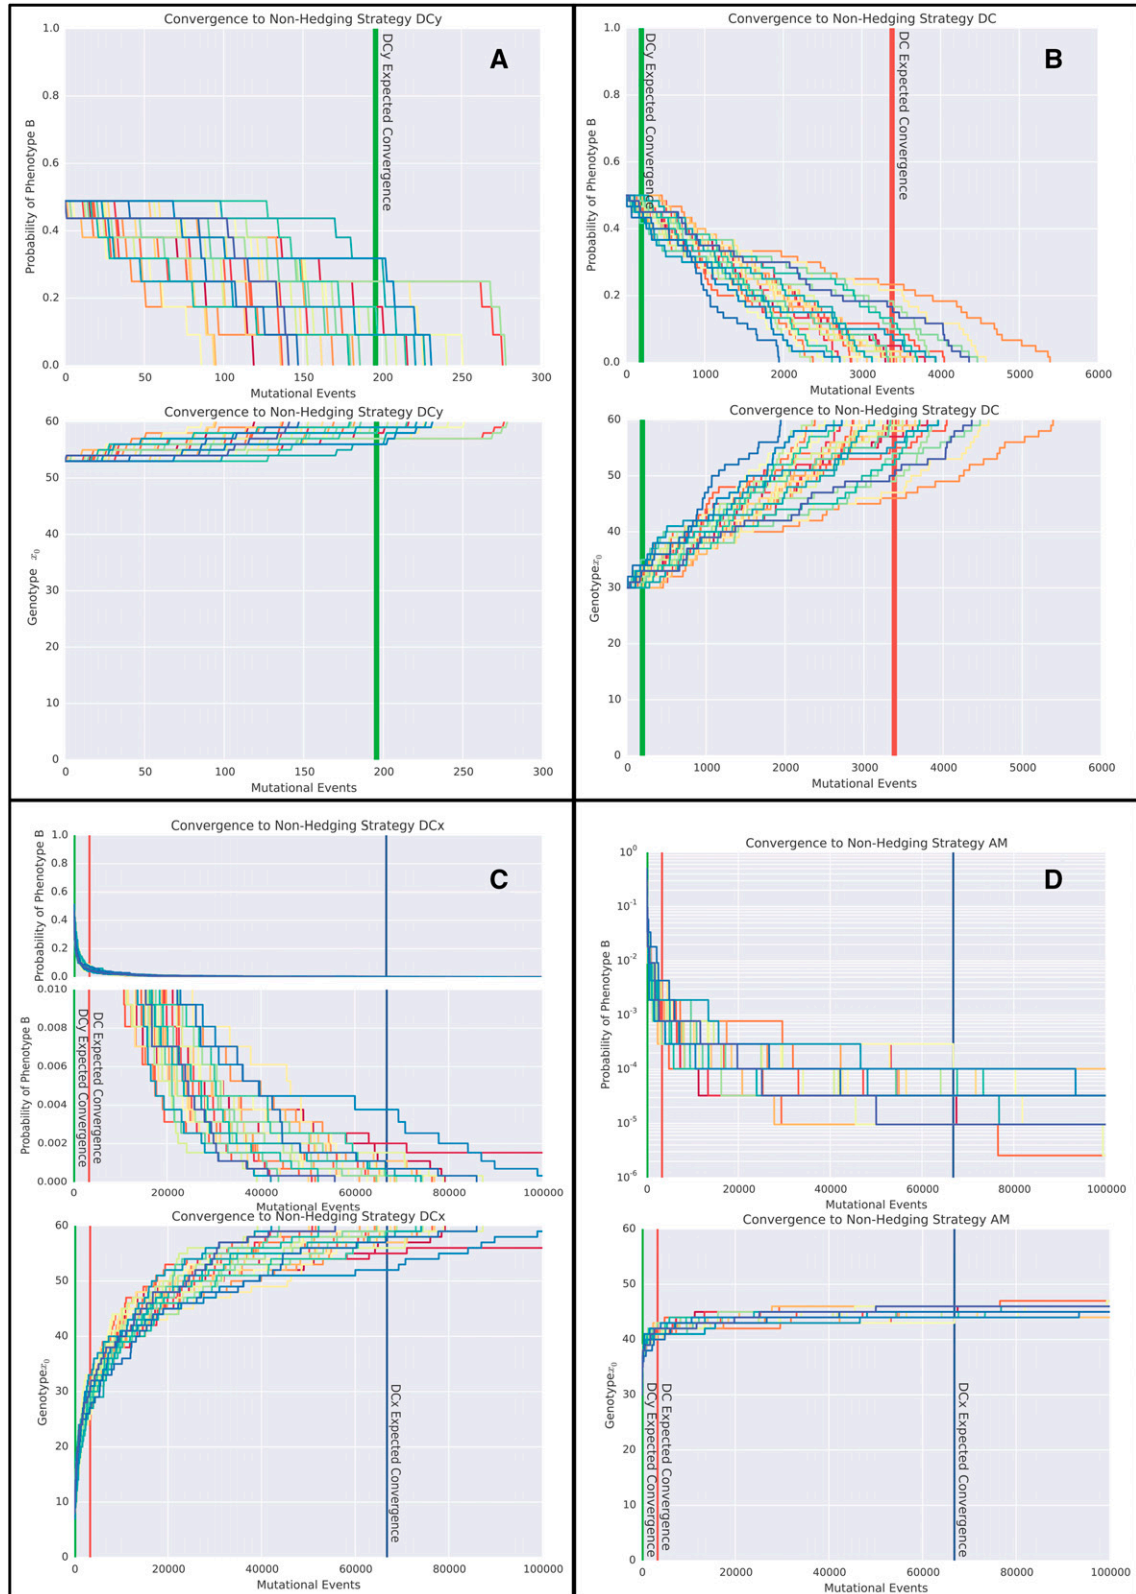


Figure 7 Convergence dynamics through genotype and probability space for the GP maps defined by DC, DCx, DCy, and AM; 30 stochastic realizations of the evolutionary simulation through both genotype and probability space are shown for (A) the DCy switch, (B) the DC switch, (C) the DCx switch, and (D) the AM switch. Due to the rapid initial change in hedging probability for the DCx switch, the convergence dynamics are also shown on a restricted scale. As the probability of phenotype B rapidly approaches zero in the AM switch simulation but never converges, the dynamics are shown on a logarithmic scale. The expected convergence time for the DCy switch is marked in green, for the DC switch is marked in red, and for the DCx switch is marked in blue.

necessary for a successful (<240 hr) follow up course of therapy varies by orders of magnitude depending on the underlying switch. It follows that the efficacy of treatment holidays as a potential therapeutic intervention for a disease with bet-hedging-driven resistance is dependent on the underlying driver of bet-hedging.

To test the robustness of this result, we considered an intermediate parameter set of $d_A = 0.4$ and $f_A = 0.5$ as a trade-off between cytostasis and cytotoxicity. The histograms of extinction times in this simulation are presented in Figure 4 of File S1. We find that, although the specific extinction times change, the qualitative results remain the same. Specifically, those combinations of treatment holiday, T , and switching mechanism that resulted in an effective follow-up treatment with a cytotoxic drug (Green, Figure 8) also permit effective treatment under this mixed parameter set. The converse is also true, with the blue (extinction but not within 240 hr) and gray (no extinction with 20,000 hr) histograms from Figure 8 being preserved.

Discussion

We have introduced a model for the GP map that uses minimal networks of stochastic interactions to determine phenotypes. Other models of nongenetic phenotypic heterogeneity have utilized deterministic interaction networks, for example, the models of Gerlee and Anderson (2007) or Huang (2009, 2012), or stochastic simulation of empirically derived molecular pathways, for example Charlebois *et al.* (2014). The model presented here differs as we investigate minimal stochastic instances of switching behavior. These networks, like those investigated by Gerlee and Anderson or Huang, are, at present, hypothetical. However, this work could be readily extended to empirically observed reaction networks in future work.

Remarkably simple network motifs have been demonstrated to implement switching behavior that can produce populations in which different proportions of cells, determined by the initial conditions of the network, take on different phenotypes. While abstract in its representation of reactions within a network, the reactions in our model are closely related to the physical mechanisms that govern intracellular regulatory networks, providing valuable insight into the impact of network architecture on the stochastic determination of phenotypes.

We have demonstrated how redundancy, a common feature of many biological systems, can result in bet-hedging that is robust to the addition or removal of chemical species. This redundancy, which can arise initially through neutral or nearly-neutral mutations of network structure, can ensure that bet-hedging is not lost through gene deletions or duplications. Critically, this observation may explain the failure to identify genes responsible for bacterial persistence. For example, Lewis (2006) highlights mutations to the genes *hipA*, *rmf*, *sulA*, and toxin-antitoxin (TA) loci *relBE*, *dinJ* and *mazEF* as possible drivers of bacterial persistence. However,

deletion of *rmf*, *relBE* or *mazEF* has been demonstrated to have no effect on the phenomenon of persistence, owing possibly to redundancy in TA modules, while deletion or over expression of *hipA* can change the proportion of bacterial persisters but not eradicate them. This is consistent with our results that indicate deletion of a single species in the CRN will not inhibit phenotypic heterogeneity but may alter phenotype proportions (Figure 5). The conclusion to be taken from the results reviewed by Lewis (2006) need not be that the factors identified are not the ones driving bacterial persistence, but instead it may be that the search for a single genetic factor responsible for bet-hedging is doomed to fail. It may be that bet-hedging emerges from the interactions of a collection of genetic factors in the sense of the epigenetic landscape studied by Brock *et al.* (2009) and Huang (2009). If this is the case, then to identify the biological mechanisms responsible for bet-hedging, we need to move beyond a gene-centric perspective, and to identify those networks of interactions governing the determination of phenotypes.

Mutations in cancer have often been associated with their direct effect on phenotypes—the concept of a driver mutation being that it induces a novel adaptive phenotype leading to clonal expansion (McFarland *et al.* 2014). However, our results suggest another phenomenon, wherein mutations do not induce novel phenotypes, but rather alter frequencies of pre-existing phenotypes within the population. This change in phenotypic ratio can have implications for cancer progression—a phenomenon previously explored by Charlebois *et al.* (2011). Consider the phenomenon of tumorigenic cells, where it is thought that only cells of a certain phenotype can form a growing tumor mass (Pardal *et al.* 2003; Ricci-Vitiani *et al.* 2007; Meacham and Morrison 2013). Genetic heterogeneity can explain the existence of a tumorigenic subpopulation if certain driver mutations are responsible for the tumorigenic phenotype. However, an alternative mechanism is that stem-like tumor cells give rise to a population of heterogeneous phenotypes. In the traditional stem cell model, a hierarchy exists where the stem cells produce a range of tumor cell phenotypes (Reya *et al.* 2001). Cancer stem cells divide to produce either more cancer stem cells (self-renewal) or cells with nonstem phenotypes down the hierarchy. This cellular decision is often taken to be stochastic (an example of bet-hedging), and our results highlight the potential for mutations to alter the probabilities of self-renewal or differentiation that have been shown to have significant impact on many aspects of tumor progression (Enderling *et al.* 2013).

An alternative bet-hedging mechanism for tumorigenicity is that the tumorigenic phenotype is transient and stochastically determined (potentially with microenvironmental influence). Evidence for this phenomenon was highlighted by Quintana *et al.* (2010), who demonstrated that the tumorigenic potential of individual melanoma cells is similar, despite high heterogeneity of many markers in the initializing cell. No driver or stem population was found, and, indeed,

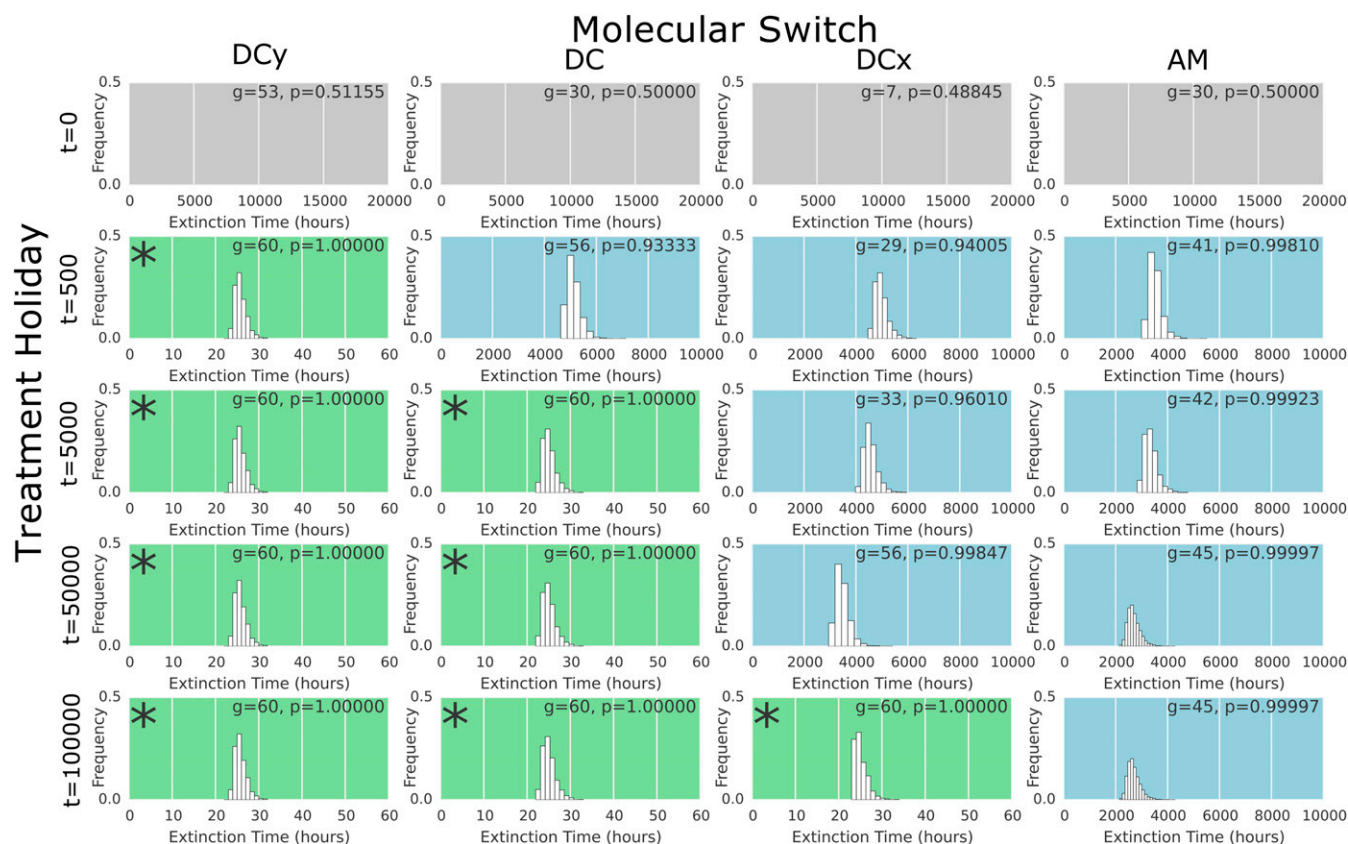


Figure 8 Treatment dynamics for populations endowed with the different switching networks after differing timescales of treatment holidays for a cytotoxic drug regime. Each histogram shows the distribution of extinction times over 2000 simulations of treatment in an individual-based model. The switch used as the GP map is shown as the column heading. The genotype and associated probability of phenotype A (shown inset to each subfigure) are determined by an evolutionary simulation of a treatment holiday for a timescale, measured in mutational events, determined by the row. The gray background (top row) indicates that no extinction occurred within the simulated 20,000 hr of treatment. The blue background indicates extinction times longer than a timeframe viable for an antibiotic treatment (240 hr), a green background (or inset star) indicates extinction times within this time frame.

the heterogeneity of marker expression was recapitulated by most tumorigenic cells, regardless of the starting pattern of expression. Such a mechanism would have different implications from the stem cell model presented above, as potentially sensitive cells would be more difficult to define. However, our predictions remain the same: that genetic mutations could shift the frequency of tumorigenic phenotypes and profoundly impact cancer progression.

Our results further demonstrate that the structure of the network governing phenotypic differentiation also has important implications for the evolutionary loss of bet-hedging. By considering mutations to the expression levels of genes, we find that the time taken for a two-phenotype bet-hedge to be lost, in an environment favorable to only one of the phenotypes, can vary by orders of magnitude depending on the network structure. If bet-hedging serves as a survival mechanism in the event of rare catastrophic environmental change (e.g., drug treatment), then the GP mapping can prevent loss of this survival mechanism over the long timescales in which catastrophe does not occur. For example, if bet-hedging is driven by a network such as the AM network (Figure 3), then each successive mutation toward a one-phenotype strategy induces a diminishing increase in the probability of generat-

ing that phenotype, and, thus, a diminishing increase in expected fitness. Eventually mutations become (essentially) neutral, and unable to fix in a large population. It is the structure of the molecular switch that substantially slows evolutionary convergence. This result provides a possible solution to the apparent paradox of bet-hedging, *i.e.*, how can bet-hedging persist in environments where natural selection acts to remove it? The structure of the molecular mechanism, itself subject to natural selection, can slow the loss of bet-hedging strategies to ensure a survival mechanism, even where environmental catastrophes are very rare.

These results have important implications for therapeutic strategies to treat diseases displaying bet-hedging-driven drug resistance. Theoretical strategies to combat bet-hedging-induced drug resistance have focused on identifying novel agents capable of killing persister cells, or identifying genetic mechanisms that can be targeted to prevent the persister phenotype from emerging. This latter strategy bears a striking resemblance to the targeted therapy revolution in the treatment of many cancers. The identification of molecular targets whose inhibition induces death (or inhibits growth) has led to the discovery of a number of novel therapies for melanoma, nonsmall cell lung cancer, and colorectal cancers. These drugs

are, in the short term, remarkably effective; however, the effects are rarely durable. Mutations that abrogate the effects of targeted therapies quickly emerge during treatment, driving resistance, and, ultimately, mortality. The results of our chemical reaction model shed light on this Darwinian adaptation, and suggest that targeted therapies to prevent bet-hedging may either be impossible, or, where they do exist, prone to fail due to the re-emergence of bet-hedging through evolution. More precisely, the discovery of a single “magic bullet” genetic factor (Strebhardt and Ullrich 2008), which, when targeted, can switch off multi-drug resistant dormant phenotypes, is unlikely, owing to redundancy in the network architecture. However, we should not rule out the potential of targeted therapies entirely. It may be possible to identify multiple targets for which simultaneous inhibition prevents bet-hedging. Alternatively, targets may be identified that shift the proportion of resistant or dormant individuals within a population to a manageable level, either permitting treatment with other cytotoxic agents, or driving the disease into a dormant state.

A second theoretical treatment strategy suggested to combat resistance in cancers and resistant infections (Bigger 1944) is the introduction of treatment holidays. The traditional doctrine for therapy is one of *maximal dose-density*, i.e., that we should treat diseases using the most potent drug with the highest tolerable dose for the longest possible time until the disease is cured, the therapy ceases to be effective, or the drugs become too toxic. Mathematical models of disease progression assuming genetically driven resistance indicate that this approach could drive the emergence of resistance through an ecological principle called *competitive release* (Alto *et al.* 2013; Adkins and Shabbir 2014). Before treatment, cells compete with one another for limited resources within a spatially constrained population. In a nontreated environment, pre-existing resistant cells are often less fit than sensitive ones, and, thus, owing to competition, do not grow to large numbers within the population. Selective pressures, for example within a growing tumor, are often not sufficiently strong for clonal sweeps (the fixation of a single genetic clone) to occur (Robertson-Tessi and Anderson 2015; Sottoriva *et al.* 2015) and the population contains a heterogeneous mix of phenotypes. When this population is exposed to a maximum dose-density therapy, the sensitive cells are killed, allowing the rapid outgrowth of the previously small and resistant population. This population then drives the recurrence of drug-resistant disease. Treatment holidays have been suggested as a potential therapeutic strategy to avoid drug resistance driven by competitive release (Enriquez-Navas *et al.* 2016).

Here, we implemented an individual-based model of the dynamics of a bet-hedging population under treatment to explore the efficacy of treatment holidays in combating bet-hedging-driven resistance. Coupled with a long-term evolutionary simulation, we explored the impact of the mechanism driving bet-hedging on the efficacy of treatment holidays. Our model suggests that it is the GP map, and, in particular, how it

hinders or promotes the rate of evolution, that determines the efficacy of treatment breaks in reversing drug resistance. If the mechanism driving bet-hedging permits evolutionary loss in a short-to-medium timescale, then treatment holidays may drive the loss of bet-hedging and re-establish drug sensitivity. However, if instead, the driving mechanism slows the evolutionary loss of bet-hedging, then a treatment break is unlikely to re-establish drug sensitivity in a time frame relevant to disease progression. We note that interfering with the mechanism driving bet-hedging through targeted therapy could alter the switching to allow the fast evolutionary loss of bet-hedging where previously this would not occur.

In this paper, we have taken the initial conditions of our reaction networks to be genetically determined and stable, allowing us to explore the implications of the structure of the network on the evolution of bet-hedging. The model we present could also be used to study additional aspects of nongenetic heterogeneity other than those presented here. For example, the model could be extended to directly account for the dynamics of gene expression [for example using the model presented by Kaern *et al.* (2005) or the GRN model presented by Huang (2009)]. We have assumed that the chemical species comprising the reaction networks are entirely depleted before reproduction and determination of the next phenotype. Omitting this assumption to permit the epigenetic inheritance of molecular switch state represents an extension of our model to account for phenotypic memory or carry-over. Alternatively, the initial conditions could be taken to be environmentally determined, and the network itself to be genetically determined. In this case the model would give a GP map similar to the neural network model used by Gerlee and Anderson (2007) and Gerlee *et al.* (2015) to study phenotypic plasticity. However, our model would differ in that the determination of phenotypes would be stochastic, permitting the study of environment-dependent bet-hedging strategies.

Finally, it is worth noting that the model of Gerlee and Anderson (2007) and Gerlee *et al.* (2015) is an extension of the classical concept of the reaction norm (Via and Lande 1985) to nonlinear and higher dimensional functions. The model presented here offers a further natural extension, breaking down the assumption of functionality within the reaction norm by introducing stochasticity, bringing the reaction norm concept more closely in line with biological reality. This theoretical extension of GP mapping to a stochastic nonfunctional process could be further extended to account for environmental factors, as discussed above. Such an extension will bring the model more closely in line with the maxim of developmental biology, that both environment and genotype are equally important in determining the phenotype (“ $g + e = p$ ”). Pigliucci (2010) suggests, in a review of theoretical models of the GP mapping, that we must build models that take both environment and genotype as equal partners in determining the phenotype, and attempt to bridge the divide between developmental biology and the modern evolutionary synthesis. The model presented here

represents a potential step toward this goal, but importantly offers something more than previous models that have set out along this path: an attempt to account for the role of chance.

Acknowledgments

D.N. thanks the Engineering and Physical Sciences Research Council (EPSRC) for generous funding for doctoral studies (ref: OUCL/DN/2013). M.R.T. and A.R.A.A. acknowledge the National Cancer Institute funded Physical Science Oncology Center grant, U54CA193489.

Literature Cited

- Adkins, S., and A. Shabbir, 2014 Biology, ecology and management of the invasive parthenium weed (*Parthenium hysterophorus* L.). *Pest Manag. Sci.* 70: 1023–1029.
- Alto, B. W., R. L. Lampman, B. Kesavaraju, and E. J. Muturi, 2013 Pesticide-induced release from competition among competing *Aedes aegypti* and *Aedes albopictus* (Diptera: Culicidae). *J. Med. Entomol.* 50: 1240–1249.
- Angluin, D., J. Aspnes, and D. Eisenstat, 2008 A simple population protocol for fast robust approximate majority. *Distrib. Comput.* 21: 87–102.
- Balaban, N. Q., J. Merrin, R. Chait, L. Kowalik, and S. Leibler, 2004 Bacterial persistence as a phenotypic switch. *Science* 305: 1622–1625.
- Balázs, G., A. van Oudenaarden, and J. J. Collins, 2011 Cellular decision making and biological noise: from microbes to mammals. *Cell* 144: 910–925.
- Beaumont, H. J., J. Gallie, C. Kost, G. C. Ferguson, and P. B. Rainey, 2009 Experimental evolution of bet hedging. *Nature* 462: 90–93.
- Becskei, A., and L. Serrano, 2000 Engineering stability in gene networks by autoregulation. *Nature* 405: 590–593.
- Becskei, A., B. Séraphin, and L. Serrano, 2001 Positive feedback in eukaryotic gene networks: cell differentiation by graded to binary response conversion. *EMBO J.* 20: 2528–2535.
- Bigger, J., 1944 Treatment of staphylococcal infections with penicillin by intermittent sterilisation. *Lancet* 244: 497–500.
- Blake, W. J., G. Balázs, M. A. Kohanski, F. J. Isaacs, K. F. Murphy *et al.*, 2006 Phenotypic consequences of promoter-mediated transcriptional noise. *Mol. Cell* 24: 853–865.
- Botero, C. A., F. J. Weissing, J. Wright, and D. R. Rubenstein, 2015 Evolutionary tipping points in the capacity to adapt to environmental change. *Proc. Natl. Acad. Sci. USA* 112: 184–189.
- Brock, A., H. Chang, and S. Huang, 2009 Non-genetic heterogeneity—a mutation-independent driving force for the somatic evolution of tumours. *Nat. Rev. Genet.* 10: 336–342.
- Cardelli, L., 2014 Morphisms of reaction networks that couple structure to function. *BMC Syst. Biol.* 8: 84.
- Cardelli, L., and A. Csikász-Nagy, 2012 The cell cycle switch computes approximate majority. *Sci. Rep.* 2: 656.
- Charlebois, D. A., N. Abdennur, and M. Kaern, 2011 Gene expression noise facilitates adaptation and drug resistance independently of mutation. *Phys. Rev. Lett.* 107: 218101.
- Charlebois, D. A., G. Balázs, and M. Kaern, 2014 Coherent feed-forward transcriptional regulatory motifs enhance drug resistance. *Phys. Rev. E Stat. Nonlin. Soft Matter Phys.* 89: 052708.
- Childs, D. Z., C. Metcalf, and M. Rees, 2010 Evolutionary bet-hedging in the real world: empirical evidence and challenges revealed by plants. *Proc. Biol. Sci.* 277: 3055–3064.
- Cushing, J. M., 1998 *An Introduction to Structured Population Dynamics*, Vol. 71, SIAM, Philadelphia, PA.
- Danforth, B. N., 1999 Emergence dynamics and bet hedging in a desert bee, *Perdita portalis*. *Proc. R. Soc. Lond. B Biol. Sci.* 266: 1985–1994.
- de Jong, I. G., P. Haccou, and O. P. Kuipers, 2011 Bet hedging or not? A guide to proper classification of microbial survival strategies. *BioEssays* 33: 215–223.
- Dodd, I. B., M. A. Micheelsen, K. Sneppen, and G. Thon, 2007 Theoretical analysis of epigenetic cell memory by nucleosome modification. *Cell* 129: 813–822.
- Elowitz, M. B., A. J. Levine, E. D. Siggia, and P. S. Swain, 2002 Stochastic gene expression in a single cell. *Science* 297: 1183–1186.
- Enderling, H., L. Hlatky, and P. Hahnfeldt, 2013 Cancer stem cells: a minor cancer subpopulation that redefines global cancer features. *Breast* 11: 200.
- Enriquez-Navas, P. M., Y. Kam, T. Das, S. Hassan, A. Silva *et al.*, 2016 Exploiting evolutionary principles to prolong tumor control in preclinical models of breast cancer. *Sci. Transl. Med.* 8: 327ra24.
- Gallie, J., E. Libby, F. Bertels, P. Remigi, C. B. Jendresen *et al.*, 2015 Bistability in a metabolic network underpins the de novo evolution of colony switching in *Pseudomonas fluorescens*. *PLoS Biol.* 13: e1002109.
- Gerlee, P., and A. R. Anderson, 2007 An evolutionary hybrid cellular automaton model of solid tumour growth. *J. Theor. Biol.* 246: 583–603.
- Gerlee, P., and A. R. Anderson, 2009 Modelling evolutionary cell behaviour using neural networks: application to tumour growth. *Biosystems* 95: 166–174.
- Gerlee, P., E. Kim, and A. R. Anderson, 2015 Bridging scales in cancer progression: mapping genotype to phenotype using neural networks. *Semin. Cancer Biol.* 30: 30–41.
- Gillespie, D. T., 1977 Exact stochastic simulation of coupled chemical reactions. *J. Phys. Chem.* 81: 2340–2361.
- Gillespie, D. T., 2001 Approximate accelerated stochastic simulation of chemically reacting systems. *J. Chem. Phys.* 115: 1716–1733.
- Gillespie, J. H., 1983 A simple stochastic gene substitution model. *Theor. Popul. Biol.* 23: 202–215.
- Gillespie, J. H., 1984 Molecular evolution over the mutational landscape. *Evolution* 38: 1116–1129.
- Gould, S. J., and R. C. Lewontin, 1979 The spandrels of San Marco and the Panglossian paradigm: a critique of the adaptationist programme. *Proc. R. Soc. Lond. B Biol. Sci.* 205: 581–598.
- Greaves, M., and C. C. Maley, 2012 Clonal evolution in cancer. *Nature* 481: 306–313.
- Gupta, P. B., C. M. Fillmore, G. Jiang, S. D. Shapira, K. Tao *et al.*, 2011 Stochastic state transitions give rise to phenotypic equilibrium in populations of cancer cells. *Cell* 146: 633–644.
- Hasty, J., J. Pradines, M. Dolnik, and J. J. Collins, 2000 Noise-based switches and amplifiers for gene expression. *Proc. Natl. Acad. Sci. USA* 97: 2075–2080.
- Hopper, K. R., 1999 Risk-spreading and bet-hedging in insect population biology 1. *Annu. Rev. Entomol.* 44: 535–560.
- Huang, S., 2009 Non-genetic heterogeneity of cells in development: more than just noise. *Development* 136: 3853–3862.
- Huang, S., 2012 The molecular and mathematical basis of Waddington's epigenetic landscape: a framework for post-Darwinian biology? *BioEssays* 34: 149–157.
- Isaacs, F. J., J. Hasty, C. R. Cantor, and J. J. Collins, 2003 Prediction and measurement of an autoregulatory genetic module. *Proc. Natl. Acad. Sci. USA* 100: 7714–7719.
- Kaern, M., T. C. Elston, W. J. Blake, and J. J. Collins, 2005 Stochasticity in gene expression: from theories to phenotypes. *Nat. Rev. Genet.* 6: 451–464.

- Keren, I., D. Shah, A. Spoering, N. Kaldalu, and K. Lewis, 2004 Specialized persister cells and the mechanism of multi-drug tolerance in *Escherichia coli*. *J. Bacteriol.* 186: 8172–8180.
- Kurata, T., K. Tamura, H. Kaneda, T. Nogami, H. Uejima *et al.*, 2004 Effect of re-treatment with gefitinib ('Iressa', ZD1839) after acquisition of resistance. *Ann. Oncol.* 15: 173–174.
- Kussell, E., and S. Leibler, 2005 Phenotypic diversity, population growth, and information in fluctuating environments. *Science* 309: 2075–2078.
- Levy, S. F., N. Ziv, and M. L. Siegal, 2012 Bet hedging in yeast by heterogeneous, age-correlated expression of a stress protectant. *PLoS Biol.* 10: 952.
- Lewis, K., 2006 Persister cells, dormancy and infectious disease. *Nat. Rev. Microbiol.* 5: 48–56.
- Li, C.-K., and H. Schneider, 2002 Applications of Perron–Frobenius theory to population dynamics. *J. Math. Biol.* 44: 450–462.
- McFarland, C. D., L. A. Mirny, and K. S. Korolev, 2014 Tug-of-war between driver and passenger mutations in cancer and other adaptive processes. *Proc. Natl. Acad. Sci. USA* 111: 15138–15143.
- Meacham, C. E., and S. J. Morrison, 2013 Tumour heterogeneity and cancer cell plasticity. *Nature* 501: 328–337.
- Müller, J., B. Hense, T. Fuchs, M. Utz, and C. Pötzsche, 2013 Bet-hedging in stochastically switching environments. *J. Theor. Biol.* 336: 144–157.
- Nikaido, H., 2009 Multidrug resistance in bacteria. *Annu. Rev. Biochem.* 78: 119.
- Pardal, R., M. F. Clarke, and S. J. Morrison, 2003 Applying the principles of stem-cell biology to cancer. *Nat. Rev. Cancer* 3: 895–902.
- Pigliucci, M., 2010 Genotype–phenotype mapping and the end of the 'genes as blueprint' metaphor. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 365: 557–566.
- Quintana, E., M. Shackleton, H. R. Foster, D. R. Fullen, M. S. Sabel *et al.*, 2010 Phenotypic heterogeneity among tumorigenic melanoma cells from patients that is reversible and not hierarchically organized. *Cancer Cell* 18: 510–523.
- Raj, A., and A. van Oudenaarden, 2008 Nature, nurture, or chance: stochastic gene expression and its consequences. *Cell* 135: 216–226.
- Ramirez, M., S. Rajaram, R. J. Steininger, D. Osipchuk, M. A. Roth *et al.*, 2016 Diverse drug-resistance mechanisms can emerge from drug-tolerant cancer persister cells. *Nat. Commun.* 7: 10690.
- Reya, T., S. J. Morrison, M. F. Clarke, and I. L. Weissman, 2001 Stem cells, cancer, and cancer stem cells. *Nature* 414: 105–111.
- Ricci-Vitiani, L., D. G. Lombardi, E. Pilozzi, M. Biffoni, M. Todaro *et al.*, 2007 Identification and expansion of human colon-cancer-initiating cells. *Nature* 445: 111–115.
- Robertson-Tessi, M., and A. Anderson, 2015 Big bang and context-driven collapse. *Nat. Genet.* 47: 196–197.
- Seger, J., 1987 What is bet-hedging? *Oxford Surveys in Evolutionary Biology* 4: 182–211.
- Sharma, S. V., D. Y. Lee, B. Li, M. P. Quinlan, F. Takahashi *et al.*, 2010 A chromatin-mediated reversible drug-tolerant state in cancer cell subpopulations. *Cell* 141: 69–80.
- Sottoriva, A., H. Kang, Z. Ma, T. A. Graham, M. P. Salomon *et al.*, 2015 A big bang model of human colorectal tumor growth. *Nat. Genet.* 47: 209–216.
- Strebhardt, K., and A. Ullrich, 2008 Paul Ehrlich's magic bullet concept: 100 years of progress. *Nat. Rev. Cancer* 8: 473–480.
- Thattai, M., and A. Van Oudenaarden, 2004 Stochastic gene expression in fluctuating environments. *Genetics* 167: 523–530.
- Tonegawa, S., 1983 Somatic generation of antibody diversity. *Nature* 302: 575–581.
- Veening, J.-W., W. K. Smits, and O. P. Kuipers, 2008 Bistability, epigenetics, and bet-hedging in bacteria. *Annu. Rev. Microbiol.* 62: 193–210.
- Via, S., and R. Lande, 1985 Genotype–environment interaction and the evolution of phenotypic plasticity. *Evolution* 39: 505–522.
- Wilke, C. O., 2003 Probability of fixation of an advantageous mutant in a viral quasispecies. *Genetics* 163: 467–474.
- Wilke, C. O., 2005 Quasispecies theory in the context of population genetics. *BMC Evol. Biol.* 5: 1.
- Wolf, D. M., V. V. Vazirani, and A. P. Arkin, 2005 Diversity in times of adversity: probabilistic strategies in microbial survival games. *J. Theor. Biol.* 234: 227–253.
- Yano, S., E. Nakataki, S. Ohtsuka, M. Inayama, H. Tomimoto *et al.*, 2005 Retreatment of lung adenocarcinoma patients with gefitinib who had experienced favorable results from their initial treatment with this selective epidermal growth factor receptor inhibitor: a report of three cases. *Oncol. Res.* 15: 107–111.
- Zhang, Y., W. W. Yew, and M. R. Barer, 2012 Targeting persisters for tuberculosis control. *Antimicrob. Agents Chemother.* 56: 2223–2230.

Communicating editor: L. M. Wahl

Supplementary Materials

Proof of the Identity for π

Here we provide a brief proof of the identity

$$1 - \pi = e^{-\mathbf{P}\pi} \quad (1)$$

where

$$\mathbf{P} = \frac{1}{\langle w \rangle} \begin{bmatrix} w_A p & w_B p \\ w_A(1-p) & w_B(1-p) \end{bmatrix}. \quad (2)$$

The proof is modified from the theory of viral quasispecies and was originally presented by Wilke [2003].

Under Wright–Fisher sampling the probability that a randomly chosen individual in the next generation is the offspring of a given individual of phenotype $i \in \{A, B\}$ in the current generation is $\frac{w_i}{\langle w \rangle N}$. Thus, the probability that an individual is the offspring of a particular parent of phenotype i and genotype g and also has phenotype r is

$$\xi_r = \frac{w_i \mathbb{P}(g \xrightarrow{\text{GP}} r)}{\langle w \rangle N} \quad (3)$$

as not all offspring will have phenotype r . It follows that the probability that a given individual of phenotype i has precisely k_r offspring of phenotype r in the next generation is given by

$$p(k_r|i) = \binom{N}{k_r} \xi_r^{k_r} (1 - \xi_r)^{N-k_r} \quad (4)$$

We can extend this argument to both phenotypes. The probability that an individual of phenotype i has k_A offspring of phenotype A and k_B offspring of phenotype B is

$$p(k_A, k_B|i) = \frac{N!}{k_A! k_B! (N - k_A - k_B)!} \xi_A^{k_A} \xi_B^{k_B} (1 - \xi_A - \xi_B)^{N-k_A-k_B} \quad (5)$$

$$= \frac{N!}{k_A! k_B! (N - k_A - k_B)!} \frac{\mathbf{P}_{i1}^{k_A}}{N} \frac{\mathbf{P}_{i2}^{k_B}}{N} \times \left(1 - \frac{\mathbf{P}_{i1}}{N} - \frac{\mathbf{P}_{i2}}{N} \right)^{N-k_A-k_B} \quad (6)$$

where the second equality follows from the definition of \mathbf{P} . Assume now that the population size, N , is sufficiently large that we may approximate $p(k_A, k_B|i)$ by taking the limit as $N \rightarrow \infty$. This limit is a bivariate Poisson distribution

$$p(k_A, k_B|i) = \frac{\mathbf{P}_{i1}^{k_A}}{k_A!} \times \frac{\mathbf{P}_{i2}^{k_B}}{k_B!} \times e^{-\mathbf{P}_{i1} - \mathbf{P}_{i2}} \quad (7)$$

In the following derivation we will use the theory of branching processes (Harris [2002]). This method, along with the assumption of infinite population size, restrict the resulting theory to fixation events of mutations which increase the average fitness of the population. Denote by χ_i the probability that the lineage of a single mutant individual with phenotype i in a population of average fitness $\langle w \rangle$ becomes extinct after finitely many generations. From the theory of branching processes we know that the vector of extinction probabilities $\chi = (\chi_A, \chi_B)$ satisfies $\chi = \mathbf{f}(\chi)$ where

$\mathbf{f}(\mathbf{z}) = (f_A(\mathbf{z}), f_B(\mathbf{z}))$ is the probability generating function of the offspring probabilities $p(k_A, k_B|i)$ given by,

$$f_i(\mathbf{z}) = \sum_{k_A} \sum_{k_B} p(k_A, k_B|i) z_1^{k_A} z_2^{k_B}. \quad (8)$$

Substituting equation 7 into equation 8 we have

$$f_i(\chi) = \sum_{k_A} \sum_{k_B} \left(\frac{\mathbf{P}_{i1}^{k_A}}{k_A!} \times \frac{\mathbf{P}_{i2}^{k_B}}{k_B!} \times e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \right) \chi_1^{k_A} \chi_2^{k_B} \quad (9)$$

$$= e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \sum_{k_A} \sum_{k_B} \left(\frac{\mathbf{P}_{i1}^{k_A}}{k_A!} \chi_1^{k_A} \right) \left(\frac{\mathbf{P}_{i2}^{k_B}}{k_B!} \chi_2^{k_B} \right) \quad (10)$$

$$= e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \sum_{k_A} \left(\frac{(\mathbf{P}_{i1}\chi_1)^{k_A}}{k_A!} \right) \sum_{k_B} \left(\frac{(\mathbf{P}_{i2}\chi_2)^{k_B}}{k_B!} \right) \quad (11)$$

$$= e^{-\mathbf{P}_{i1}-\mathbf{P}_{i2}} \times e^{\mathbf{P}_{i1}\chi_1} \times e^{\mathbf{P}_{i2}\chi_2} \quad (12)$$

$$= e^{\mathbf{P}_{i1}(\chi_1-1)+\mathbf{P}_{i2}(\chi_2-1)}. \quad (13)$$

Taking the convention that $e^{\mathbf{x}} = (e^{x_1}, \dots, e^{x_n})$ this gives

$$\mathbf{f}(\chi) = e^{\mathbf{P}(\chi-1)}. \quad (14)$$

Now note that the probability of fixation of the mutant individual with phenotype i is precisely the probability that its lineage does not go extinct in finitely many generations, $\pi = \mathbf{1} - \chi$. Hence

$$1 - \pi = \mathbf{f}(1 - \pi) = e^{-\mathbf{P}\pi}. \quad (15)$$

Robustness to Parameter Variation

The parameters g_{\max} , w_A and w_B use in this paper take specific values chosen for reasons of computational efficiency (for g_{\max}) or to mirror qualitative behaviour (for w_A and w_B). Here, we demonstrate the robustness of our results to variations in these parameters. Figure 1 shows the effect of increasing g_{\max} on convergence time. For each type of switch, the convergence time increases monotonically with g_{\max} and the difference in convergence time between switching networks, on which our results are predicated, holds as g_{\max} increases. Figure 2 shows the effect of increasing w_A on convergence time when $w_B = 1.01$ is fixed. The differing dynamics of evolutionary convergence are preserved for increasing w_A . However, convergence time is monotonically decreasing as w_A increases. The AM switch is omitted as there is no convergence for $w_A \leq 10$.

Analytic and Numerical Calculation of $p(g)$

Computing $p(g)$ by sampling many stochastic runs of the Gillespie algorithm for each molecular switch is computationally expensive. Instead, we analytically determine $p(g)$ for the DC, DCx and DCy switches and numerically approximate it for the AM switch.

For the DC, DCx and DCy switch we note that the stochastic process corresponds to the classical ‘drunkards walk’ in probability theory. Thus, we know




Figure 1: **Comparison of Convergence Times for Increasing g_{max} .**

$$\mathbb{P}(\text{Phenotype } A \mid \text{DC switch}, g = (x_0, y_0)) = \frac{x_0}{x_0 + y_0} \quad (16)$$

$$\mathbb{P}(\text{Phenotype } A \mid \text{DCx switch}, g = (x_0, y_0)) = \frac{1 - \left(\frac{10}{11}\right)^{x_0}}{1 - \left(\frac{10}{11}\right)^{x_0 + y_0}} \quad (17)$$

$$\mathbb{P}(\text{Phenotype } A \mid \text{DCy switch}, g = (x_0, y_0)) = \frac{1 - \left(\frac{11}{10}\right)^{x_0}}{1 - \left(\frac{11}{10}\right)^{x_0 + y_0}}. \quad (18)$$

The switching probabilities in the AM network can be approximated numerically by constructing a Markov chain on the state space of possible configurations of x , y and b molecules.

Assume that there are $n \in \mathbb{N}$ molecules in the system. Denote a configuration of the system by (r, s) for $r, s \in \mathbb{N}$ with $0 < r + s \leq n$ where r and s denote the number of molecule x and y respectively. Call a configuration stable if no more reactions can occur. As each reaction in the AM network conserves the total number of molecules we know that the number of molecules b is equal to $n - r - s$. Denote

$$p(r, s) = \mathbb{P}[\text{The AM-network with initial config.} \quad (19)$$

$$(r, s) \text{ ends in the stable config. } (n, 0)]. \quad (20)$$

In the absence of y molecules, either the system has reached the stable configuration $(n, 0)$ or the only reaction that can occur is $x + b \rightarrow x + x$ and the system must reach $(n, 0)$ eventually. Hence

$$p(r, 0) = 1 \text{ for all } r \in [1, n]. \quad (21)$$




Figure 2: **Comparison of Convergence Times for Increasing w_A , $w_B = 1.0$. $g_{\max} = 60$ for each curve. The AM network dynamics are omitted as convergence does not occur.**

Further, by the symmetry of the reaction network, we must have that

$$p(r, s) = 1 - p(s, r) \text{ for all } s, r \in [0, n] \text{ with } 0 < r + s \leq n \quad (22)$$

and in particular we have

$$p(r, s) = \frac{1}{2} \text{ for all } r = s \text{ with } 0 < r + s \leq n. \quad (23)$$

By this symmetry, we need only determine $p(r, s)$ on the set

$$S = \{(r, s) : r, s \in [0, n], 0 < r + s \leq n \text{ and } r \geq s\} \quad (24)$$

Denote the subset of S for which we know $p(r, s)$ by S^* . Within $S \setminus S^*$ we have a recurrence relation for $p(r, s)$.

$$\begin{aligned} p(r, s) = & \mathbb{P}[x + y \longrightarrow x + b : r, s]p(r, s - 1) \\ & + \mathbb{P}[y + x \longrightarrow y + b : r, s]p(r - 1, s) \\ & + \mathbb{P}[x + b \longrightarrow x + x : r, s]p(r + 1, s) \\ & + \mathbb{P}[y + b \longrightarrow y + y : r, s]p(r, s + 1) \end{aligned} \quad (25)$$

Where the probabilities of each reaction are given by the propensity functions,

$$\begin{aligned}
\mathbb{P}[x + y \longrightarrow x + b : r, s] &= \frac{rs}{2rs + r(n - r - s) + s(n - r - s)} \\
\mathbb{P}[y + x \longrightarrow y + b : r, s] &= \frac{rs}{2rs + r(n - r - s) + s(n - r - s)} \\
\mathbb{P}[x + b \longrightarrow x + x : r, s] &= \frac{r(n - r - s)}{2rs + r(n - r - s) + s(n - r - s)} \\
\mathbb{P}[y + b \longrightarrow y + y : r, s] &= \frac{s(n - r - s)}{2rs + r(n - r - s) + s(n - r - s)}.
\end{aligned} \tag{26}$$

Now, as we know the values of $p(r, s)$ along two boundaries we can solve this recurrence numerically with a Markov Chain. Note that we can view the chemical reaction simulation as a random walk inside the configuration space S with transition probabilities given by the likelihoods of each reaction. We define the $|S| \times |S|$ transition matrix for the Markov chain as

$$M((r, s), (r', s')) = \begin{cases} 1 & \text{if } (r, s) \in S^* \text{ and } r = r', s = s' \\ 0 & \text{if } (r, s) \in S^* \text{ and } r \neq r' \text{ or } s \neq s' \\ \mathbb{P}[R : r, s] & \text{if } (r, s) \notin S^* \text{ and there exists a reaction} \\ & R \text{ taking } (r, s) \text{ to } (r', s') \\ 0 & \text{otherwise} \end{cases} \tag{27}$$

in which we have taken those values (r, s) for which we know $p(r, s)$ to be absorbing states. If we define the limit matrix

$$M^* = \lim_{k \rightarrow \infty} M^k \tag{28}$$

then $M^*((r, s), (r', s'))$ is only non-zero when (r', s') is a state for which we already know $p(r', s')$. In this case $M^*((r, s), (r', s'))$ denotes the probability that the random walk in S which simulates a series of chemical reactions will reach (r', s') before any other absorbing state. As such we can determine $p(r, s)$ by summing over all of the absorbing states of the Markov Chain

$$p(r, s) = \sum_{\substack{r', s' \in [0, n], 0 < r' + s' \leq n \\ s' = 0 \text{ or } r' = s'}} p(r', s') M^*((r, s), (r', s')). \tag{29}$$

This relation is used to determine the switching probabilities for the AM switch, for example in row 3 of Figure 3 in the main text.

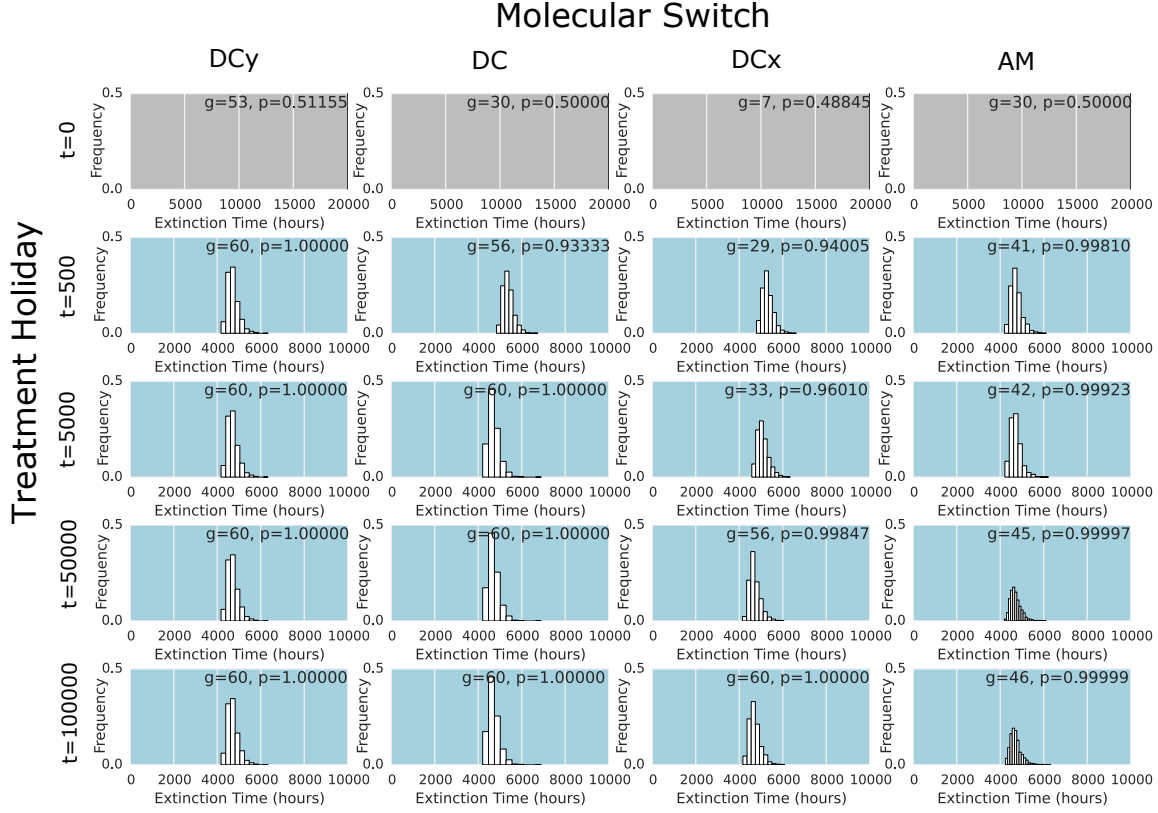


Figure 3: **Treatment dynamics for populations endowed with the different switching networks after differing timescales of treatment holidays for the cytostatic drug regime.** Each histogram shows the distribution of extinction times over 2000 simulations of treatment in an individual-based model. The switch used as the GP-map is shown as the column heading. The genotype and associated probability of phenotype A (shown inset to each subfigure) are determined by an evolutionary simulation of a treatment holiday for a time scale, measured in mutational events, determined by the row. The grey background (top row) indicates that no extinction occurred within the simulated 20000 hours of treatment. The blue background indicates extinction times longer than a timeframe viable for an antibiotic treatment (240 hrs), a green background (or inset star) indicates extinction times within this time frame.

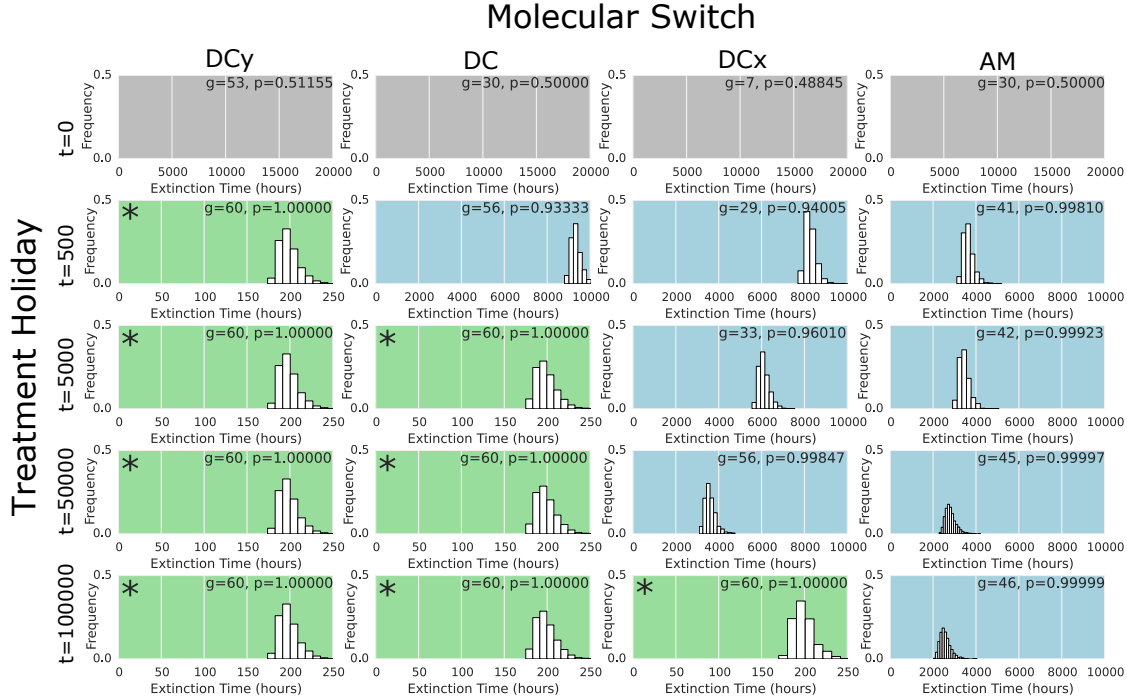


Figure 4: **Treatment dynamics for populations endowed with the different switching networks after differing timescales of treatment holidays for the mixed drug regime.** Each histogram shows the distribution of extinctintimes over 2000 simulations of treatment in an individual-based model. The switch used as the GP-map is shown as the column heading. The genotype and associated probability of phenotype A (shown inset to each subfigure) are determined by an evolutionary simulation of a treatment holiday for a time scale, measured in mutational events, determined by the row. The grey background (top row) indicates that no extinction occurred within the simulated 20000 hours of treatment. The blue background indicates extinction times longer than a timeframe viable for an antibiotic treatment (240 hrs), a green background (or inset star) indicates extinction times within this time frame.

Bibliography

Claus O Wilke. Probability of fixation of an advantageous mutant in a viral quasispecies. *Genetics*, 163(2):467–474, 2003.

Theodore E Harris. *The theory of branching processes*. Courier Corporation, 2002.