PROCEEDINGS OF THE ROYAL SOCIETY **B** | **BIOLOGICAL SCIENCES**

Exclusion rules, bottlenecks and the evolution of stochastic phenotype switching

Eric Libby and Paul B. Rainey

Supplementary data	"Data Supplement" http://rspb.royalsocietypublishing.org/content/suppl/2011/04/12/rspb.2011.0146.DC1.html
References	**This article cites 42 articles, 10 of which can be accessed free** http://rspb.royalsocietypublishing.org/content/early/2011/05/03/rspb.2011.0146.full.html#ref-list-1
P<P	Published online 13 April 2011 in advance of the print journal.
Subject collections	Articles on similar topics can be found in the following collections systems biology (211 articles) ecology (2347 articles) evolution (2638 articles)
Email alerting service	Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

Advance online articles have been peer reviewed and accepted for publication but have not yet appeared in the paper journal (edited, typeset versions may be posted when available prior to final publication). Advance online articles are citable and establish publication priority; they are indexed by PubMed from initial publication. Citations to Advance online articles must include the digital object identifier (DOIs) and date of initial publication.

To subscribe to *Proc. R. Soc. B* go to: **http://rspb.royalsocietypublishing.org/subscriptions**

Exclusion rules, bottlenecks and the evolution of stochastic phenotype switching

Eric Libby[1],* and Paul B. Rainey[1],[2]

[1]*New Zealand Institute for Advanced Study and Allan Wilson Centre for Molecular Ecology and Evolution, Massey University, Private Bag 102904, North Shore Mail Centre, North Shore City 0745, Auckland, New Zealand*
[2]*Max Planck Institute for Evolutionary Biology, 24306 Plön, Germany*

Stochastic phenotype switching—often considered a bet hedging or risk-reducing strategy—can enhance the probability of survival in fluctuating environments. A recent experiment provided direct evidence for an adaptive origin by showing the de novo evolution of switching in bacterial populations propagated under a selective regime that captured essential features of the host immune response. The regime involved strong frequency-dependent selection realized via dual imposition of an exclusion rule and population bottleneck. Applied at the point of transfer between environments, the phenotype common in the current environment was assigned a fitness of zero and was thus excluded from participating in the next round (the exclusion rule). In addition, also at the point of transfer, and so as to found the next bout of selection, a single phenotypically distinct type was selected at random from among the survivors (the bottleneck). Motivated by this experiment, we develop a mathematical model to explore the broader significance of key features of the selective regime. Through a combination of analytical and numerical results, we show that exclusion rules and population bottlenecks act in tandem as potent selective agents for stochastic phenotype switching, such that even when initially rare, and when switching engenders a cost in Malthusian fitness, organisms with the capacity to switch can invade non-switching populations and replace non-switching types. Simulations demonstrate the robustness of our findings to alterations in switching rate, fidelity of exclusion, bottleneck size, duration of environmental state and growth rate. We also demonstrate the relevance of our model to a range of biological scenarios such as bacterial persistence and the evolution of sex.

Keywords: stochastic switching; origins; immune response; exclusion rule; bottleneck

1. INTRODUCTION

Organisms must keep pace with environmental change to avoid extinction. A common mechanism is the generation of phenotypic variation. Recognized strategies range from phenotypic plasticity [1], whereby individual organisms modulate their behaviour or phenotype to match prevailing conditions, to bet hedging [2–5], where phenotypes arise stochastically without prior knowledge of the environment. While there is general recognition that plasticity is favoured in environments with reliable informational cues, and bet hedging in temporally fluctuating, or course-grained (unpredictable) environments [5–7], defining the precise set of conditions for the evolution of any one strategy poses numerous challenges [8].

The capacity to switch stochastically between heritable phenotypic states is common in bacteria. Observed initially as variation in the morphology of colonies arising from single clones of certain bacterial pathogens [9] (and underpinned by specific mutational mechanisms [10]), recent work shows stochastic switching to be a near universal feature of living systems [11,12], arising from little other than molecular noise [12–16].

There are at least three instances among bacteria where the case for stochastic phenotype switching as adaption has been argued. In the case of bacterial persistence, cells switch stochastically between growing and non-growing (persister) states [17]. This can be adaptive in the face of periodic encounters with antibiotics despite the cost associated with non-growing cells [18,19]. A similar argument explains the competence to non-competence switch for natural DNA transformation in the soil bacterium *Bacillus subtilis* [14]. Like the persister state, competence is associated with periods of non-growth in an otherwise growing population and can be beneficial, despite the cost, provided the population periodically encounters conditions that kill growing cells [20].

The third example comes from obligate commensals—and sometime pathogens—of humans, such as *Haemophilus influenzae* (reviewed in [10,21]). Survival of *H. influenzae* depends on the avoidance of recognition by the host immune response. Given that moment-by-moment fluctuations in the state of the immune response cannot be predicted [22], *H. influenzae* would appear to survive by hedging its evolutionary bets. This it achieves via mutational mechanisms which cause genes involved in critical interactions with the host to switch stochastically between expression states. In conjunction with population growth, the capacity to switch means that highly polymorphic populations emerge rapidly from limiting and initially

uniform inocula. The net effect is to ensure that the risk of immune detection is spread among variable offspring, each of which has some chance of avoiding recognition. What ecological circumstances might promote the evolution of such a strategy?

Consider once again *H. influenzae*: during the course of colonizing a new host, the bacterium faces fluctuating and unpredictable conditions as the specific effects wrought by the immune response are numerous. For example, *H. influenzae* experiences environmental fluctuations with varying dynamics and degrees of uncertainty; whether or not bet hedging evolves depends on many factors [19,23–29], including the existence and reliability of environmental cues [5–7], the capacity of the population to respond via mutation and selection [24,30], the nature of the fitness landscape [27,31] and the cost–benefit balance of different strategies [19,23,27,30].

But fluctuating selection is likely to exert additional population effects. For example, as *H. influenzae* populations increase in size, types not detected by the immune response stand a chance of becoming common: however, common types are likely to be detected and eliminated. At the moment of detection the population experiences strong frequency-dependent selection: types that were common are eliminated and concomitantly the population collapses. Re-establishment of the population is via rare types that avoided immune detection.

In a recent experiment, Beaumont *et al.* [32] allowed populations of bacteria to evolve in the face of a selective regime that mimicked the dynamic fluctuations described above. Specifically, populations of *Pseudomonas fluorescens*—a bacterium that does not undergo visible phenotypic switching—were subjected to strong frequency-dependent selection wrought by repeated imposition of an *exclusion rule* and *bottleneck*. Applied at the point of transfer between environments, the phenotype common in the current environment was assigned a fitness of zero and was thus excluded from participating in the next round (the exclusion rule). In addition, also at the point of transfer, and so as to found the next bout of selection, a single phenotypically distinct type was selected at random from among the survivors (the bottleneck). In two of 12 replicate lines, stochastic switching types evolved after eight successive rounds of fluctuating selection—each punctuated by concomitant imposition of the exclusion rule and bottleneck. The authors suggested that the exclusion rule selected for phenotypic innovation, whereas the bottleneck negated the cost of bet hedging by eliminating competition with conspecifics.

Here, motivated by the experiment of Beaumont *et al.* [32], we use a simple mathematical model to explore the competitive benefits of switching in populations subjected to repeated bouts of frequency-dependent selection imposed via exclusion rules and bottlenecks. We do so in order to assess the robustness and generality of the ecological conditions defined by the experiment of Beaumont *et al.* [32] for the evolution of stochastic switching. Using mathematical and computer simulation models, we show that even when initially rare and when switching engenders a cost in Malthusian fitness, organisms with this capacity can invade non-switching populations and replace non-switching phenotypes.

2. MODEL AND METHODS

To examine effects due to exclusion rules and bottlenecks—two separable facets of frequency-dependent selection—we propose a simple model of four competing phenotypes: X_1, X_2, S_A and S_B (figure 1). Each phenotype reproduces asexually giving rise to a copy of itself, or, with a smaller probability, any of the other three phenotypes. Although the system is fully connected, the S phenotypes are designated 'switchers' and the X phenotypes 'non-switchers'. The difference in nomenclature reflects the nature and magnitude of the transition between phenotypes. We assume, although it is not critical, that the X phenotypes are different genotypes and a mutation is required for the transition between X_1 and X_2. In contrast, S_A and S_B are different epigenetic states of the same genotype and, thus, transitions between S_A and S_B occur more often.

While in biological systems we expect different phenotypes to possess different growth rates or transition probabilities, here we assume that these parameters are identical for each of the switching and non-switching phenotypes. This assumption allows a simplification of the analytical calculations in order to focus on the effects of transition rates. Later in the simulations, these assumptions are relaxed for the switching phenotypes in order to investigate the case where switchers grow at different rates.

In practice, many different types of non-switching phenotypes can arise by mutation. We consider X_2 to encompass the complete repertoire of non-switching phenotypes that can arise from X_1. We expect such events to be more common than transitions between switching and non-switching types. Thus, the transition between non-switching phenotypes (m_x) is greater than the mutation rate between non-switching and switching types ($m_x > m_{t2}, m_{t1}$).

The system is closed, mitigating effects attributable to migration, but the model can account for such effects. The probability that a new type arises depends on the mutation supply rate, thus in exponential growth mutation supply will be high. If the probability of an organism invading is of the same order of magnitude as the mutation supply rate (or less), then the model can accommodate new phenotypes by invasion rather than mutation.

We describe the model depicted in figure 1 with a system of ordinary differential equations (equations (2.1)) similar to those used elsewhere [23,25,27,30]. Here the growth rates (k_1 and k_2) are several orders of magnitude greater than the transitions between phenotypes. The system at ($t = 0$) is founded by X_1, so any types present at future times must originate from X_1. Types grow exponentially for short periods and without any frequency-dependent effects.

$$
\left.
\begin{aligned}
\frac{\mathrm{d}X_1}{\mathrm{d}t} &= (k_1 - m_x - 2m_{t1})X_1 + m_xX_2 + m_{t2}(S_A + S_B),\\
\frac{\mathrm{d}X_2}{\mathrm{d}t} &= (k_1 - m_x - 2m_{t1})X_2 + m_xX_1 + m_{t2}(S_A + S_B),\\
\frac{\mathrm{d}S_A}{\mathrm{d}t} &= (k_2 - m_s - 2m_{t2})S_A + m_sS_B + m_{t1}(X_1 + X_2)\\
\frac{\mathrm{d}S_B}{\mathrm{d}t} &= (k_2 - m_s - 2m_{t2})S_B + m_sS_A + m_{t1}(X_1 + X_2).
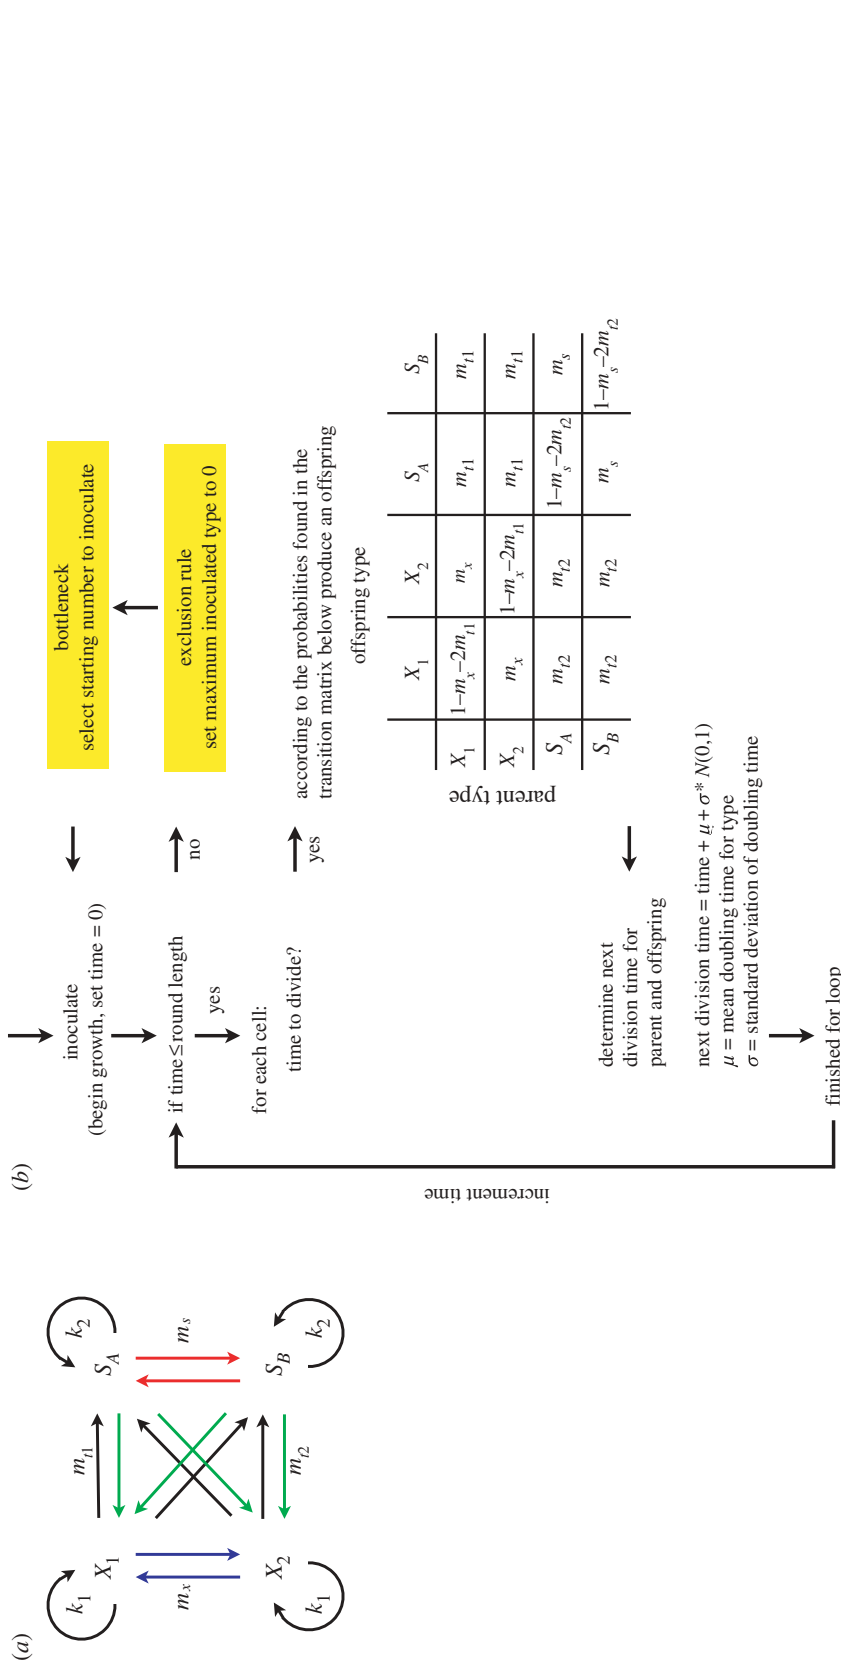\end{aligned}
\right\}
$$

(2.1)

Figure 1. Model of switcher evolution. (*a*) X_1 and X_2 are non-switching phenotypes that grow according to rate k_1. X_1 can generate X_2 and vice versa by mutation (m_x). S_A and S_B phenotypes grow at a rate k_2 and are the same genotype: they stochastically switch at a rate greater than spontaneous mutation (m_s). The routes between X and S phenotypes, m_{t1} and m_{t2}, occur via mutation. The model depicts the system during one round of growth in a constant environment. When populations experience multiple rounds of growth they do so in the face of a regime that involves frequency-dependent selection realized via dual imposition of an exclusion rule and population bottleneck (see text). Applied at the end of each round, the phenotype common during the current round is assigned a fitness of zero and is thus excluded from participating in the next round (the exclusion rule). Concomitantly, the population passes through a bottleneck and is re-established from one (or a small number) of phenotypically distinct types selected at random from among the survivors (blue, m_x; red, m_s; black, m_{t1}; green, m_{t2}). (*b*) A flowchart depicting the algorithm governing the stochastic simulations during one round of growth and selection. Inoculum added to a pristine environment begins a round of growth which continues for a set time (the round length), after which an exclusion rule and bottleneck (highlighted) are applied. The bottleneck determines the inoculum for the next round of growth and selection.

To determine the effect of different transition rates, growth rates for all phenotypes are set to the same value k (i.e. $k_1 = k_2 = k$). In subsequent simulations, we handicap switchers by reducing their growth rates. At this stage, there is no cost to switching. We solve the system (equations (2.1)) for the time-varying concentrations and obtain a sum of exponential terms for each phenotype (equations (2.2)).

$$
\begin{aligned}
X_1(t) &= \frac{m_{t1}}{2m_{t1} + 2m_{t2}} e^{(k - 2m_{t1} - 2m_{t2})t} + \frac{m_{t2}}{2m_{t1} + 2m_{t2}} e^{kt} \\
&\quad + \frac{1}{2} e^{(k - 2m_x - 2m_{t1})t} \\
X_2(t) &= \frac{m_{t1}}{2m_{t1} + 2m_{t2}} e^{(k - 2m_{t1} - 2m_{t2})t} + \frac{m_{t2}}{2m_{t1} + 2m_{t2}} e^{kt} \\
&\quad - \frac{1}{2} e^{(k - 2m_x - 2m_{t1})t}, \\
S_A(t) &= \frac{-m_{t1}}{2m_{t1} + 2m_{t2}} e^{(k - 2m_{t1} - 2m_{t2})t} + \frac{m_{t1}}{2m_{t1} + 2m_{t2}} e^{kt} \\
\text{and} \quad S_B(t) &= \frac{-m_{t1}}{2m_{t1} + 2m_{t2}} e^{(k - 2m_{t1} - 2m_{t2})t} + \frac{m_{t1}}{2m_{t1} + 2m_{t2}} e^{kt}.
\end{aligned}
$$
$$(2.2)$$

The equations for S_A and S_B are identical as they start with the same concentration and their dynamics are governed by the same equation. The equations for X_1 and X_2 differ only in the third term which is the sole appearance of the term m_x. In contrast, the switching rate between S_A and S_B, m_s does not appear in the solution. Because the initial population began with an X type, the growth of both S_A and S_B depends on m_{t1} not on the switching rate. Furthermore, since the m_s term is not biased, it has no net effect and does not factor into their growth. If the initial concentrations were different there would be an $e^{(k - 2m_s - 2m_{t2})t}$ term (the fourth eigenvalue of the system), but here its coefficient is zero.

Since the system is fully connected, the proportion of each phenotype tends towards a fixed value as $t \to \infty$. The transition rates are small enough in magnitude ($\ll 10^{-2}$, and as low as approx. 10^{-7}) that for the time intervals considered ($t < 10$) the proportions of each type are far from equilibrium.

If after a short amount of growth ($tm_s, tm_x \ll 10^{-1}$) a type is randomly chosen without excluding any from the lottery, the probability of selecting a non-switching phenotype is solely dependent on the transition rates m_{t1} and m_{t2} (equation (2.3)), where $S(t) = S_A(t) + S_B(t)$). The m_s term does not appear because of the initial conditions; the m_x term does not appear because there is no advantage to switching when the initial type (X_1) is not excluded. Thus far the conditions are symmetric so had the initial phenotype been a switcher, S_A or S_B, equation (2.3) would still apply, however, the terms m_{t1} and m_{t2} would be swapped. Consequently, in the absence of an exclusion rule, there is no advantage to switching

$$
\frac{X_1(t) + X_2(t)}{S(t) + X_1(t) + X_2(t)} = \frac{m_{t2}}{m_{t1} + m_{t2}} + \frac{m_{t1}}{m_{t1} + m_{t2}} e^{-2t(m_{t1} + m_{t2})}.
$$
$$(2.3)$$

If, however, a type is randomly selected while excluding the initial phenotype X_1, the probability of selecting

a non-switching phenotype will depend on the rate m_x as well as on the transition rate m_{t1} (equation (2.4), where $S(t) = S_A(t) + S_B(t)$). Owing to the symmetrical nature of the system, if switcher S_A were the founding type then the probability of selecting switcher S_B would follow equation (2.4), with m_s replacing m_x and m_{t2} replacing m_{t1}. Here, we see a direct advantage for switching. Indeed, with the exclusion rule in place selection rewards types that switch at high rates. For example, the probability of losing a non-switcher after it is inoculated (choose S instead of X_2 if X_1 were the initial phenotype) can be as high as 10^4 times more likely than losing a switcher [$m_s = 10^{-2}$, m_{t2}, $m_{t1} = 10^{-7}$, $m_x = 10^{-6}$]

$$
\frac{X_2(t)}{S(t) + X_2(t)} = \frac{m_x}{2m_{t1} + m_x}.
$$
$$(2.4)$$

3. SIMULATIONS AND RESULTS

The above calculations show the benefit of stochastic phenotype switching via the proportion of switching and non-switching phenotypes after just one period of growth. Yet, over the course of evolution many rounds of growth and selection occur. After each period of growth the number of switchers selected for future rounds follows a hyper-geometric distribution and depends on the actual number of each species in the population. In addition, our approach holds for large numbers of bacteria yet stringent bottlenecks can reduce the population to a single entity. Stochastic events early in a round of growth such as a chance mutation will have downstream effects when selection acts. To capture the stochastic dynamics of exponential growth expanding from an initially small number of bacteria as well as the long-term effects of repeated rounds of growth and selection, we used the Monte Carlo scheme outlined in figure 1b.

In the simulation, every time either S_A or S_B divides there is a probability of 10^{-3} that it will generate the other switcher phenotype. This value is in the same order of magnitude as experimental models of switchers [26,33,34]. While we expect biological mutation probabilities between 10^{-6} and 10^{-9} [33,35,36], we chose higher values to make a mutation likely to occur during each round of the simulation. Thus, there is a probability of 10^{-4} that a non-switching type (X_1 or X_2) will produce a mutant of the other non-switching type, i.e. an X_1 will produce an X_2 or vice versa. This conservative estimate means that the non-switching type can be viewed as a 'slower switcher', only 10 times slower than S. We set the probability at which a switching type will mutate to a non-switching type, or vice versa, to 10^{-5}.

Each time an organism reproduces, we randomly sample from a uniform distribution between 0 and 1 to determine whether it breeds true or produces a different phenotype. For example, when an X_1 divides it will produce another X_1 with a probability $1 - 2m_{t1} - m_x$, an X_2 with probability m_x, an S_A with probability m_{t1}, or an S_B with probability m_{t1}. Although, we assume that an entity cannot change state during a round, relaxing this assumption does not fundamentally alter the results. Permitting such state changes is equivalent to increasing the probabilities of transitions.

The division times for cells are governed by a Gaussian distribution [37] with a mean division time determined by the phenotype and a constant coefficient of variation
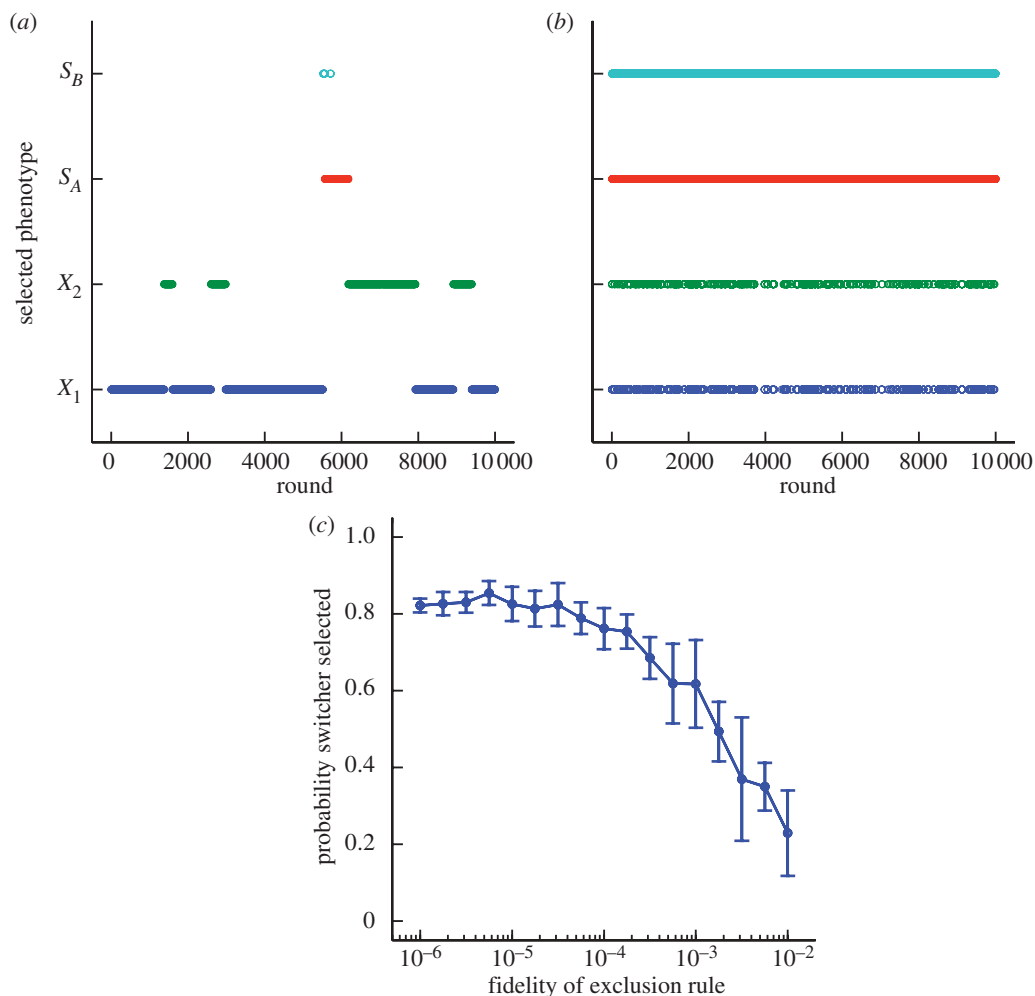
Figure 2. The exclusion rule. (*a*) The model in figure 1 simulated over 10 000 rounds of growth and selection without application of an exclusion rule. Each round corresponds to 3.5 time units of growth (doubling times are 0.20 and 0.25). At the end of a round of growth, one organism is selected randomly to inoculate the next round. The phenotype of the selected organism is indicated at different heights on the vertical axis. X_1 was selected 6514 times, $X_2 = 2825$, $S_A = 602$ and $S_B = 59$. (*b*) The same process as in (*a*) is repeated except the exclusion rule is applied at the end of each round, preventing the same phenotype from winning consecutive bouts. The number of bouts in which the switcher (S_A and S_B) invades from rare is enhanced by implementation of an exclusion rule. X_1 was selected 863 times, $X_2 = 857$, $S_A = 4136$ and $S_B = 4144$. (*c*) The proportion of the original inoculum is varied: the probability an immediate ancestor survives selection is shown on the horizontal axis (complete exclusion is 10^{-6} and no exclusion is 1). Each point represents the mean of five samples of 10 simulations and the error bars show the standard deviation. The simulations were run for 125 rounds of growth and selection with the first 25 discarded to avoid bias due to the initial inoculation.

(0.1). Unless specified, the mean division time for non-switching phenotypes is 0.20 time units and is 0.25 time units for switching phenotypes. Using exponentially distributed division times instead of Gaussian distributed division times did not alter the main findings. We continue these divisions for 3.5 time units which is of the same order as employed by Beaumont *et al.* [32] and the time the human immune system takes to identify and respond to a pathogen [38]. The computer code for the simulations is provided in the electronic supplementary material.

The mathematical analysis shows that the exclusion rule favours switching phenotypes when the growth rates for non-switching and switching types are identical. In biological systems, however, we expect switching to impose a burden [25,27,31]. Because a switcher stochastically generates different phenotypes, it is likely that at least one type will be maladapted to the prevailing conditions. The cost of switching, therefore, is diminished

ecological performance and consequently a reduced growth rate for at least one of the phenotypes. In the simulation, the cost is encoded as a 25 per cent longer doubling time for the switcher (both S_A and S_B): 0.25 time units as opposed to 0.20 time units for non-switching phenotypes. Unlike previous work, we apply a cost to *both* switcher types so that it is less fit in *both* environments. The switcher, therefore, is equivalent to a slow-growing mutant whose only advantage is being able to randomly generate two phenotypes.

Figure 2 shows the results of the simulation for 10 000 rounds of growth and selection. At each round the population passes through a bottleneck such that each new round starts from a single randomly chosen entity (at $t =$ round length, $X_1(t) + X_2(t) + S_A(t) + S_B(t) =$ bottleneck size). The population size before the bottleneck depends on the inoculated type and the stochasticity in the time to division. Typically, if inoculated with one entity, the population reaches 10^4 to 10^5. In

the absence of the exclusion rule, the combination of the switcher's longer doubling time and the conservative mutation rate between non-switching phenotypes (m_x) means that the switcher is rarely selected (figure 2a). As predicted from the mathematical analysis, lack of an exclusion rule favours the inoculated phenotype. Occasionally a transition occurs because of a chance mutation that happens early. If, on the other hand, the exclusion rule is applied (at $t =$ round length, inoculated type, say $X_1(t), = 0$) then we see that the switcher does significantly better—even outperforming the non-switching phenotypes (figure 2b). Since the exclusion rule prevents the same phenotype from being selected twice in a row, the growth advantage of the non-switching phenotypes is offset by its reliance on mutation to generate new phenotypes. In environments governed by an exclusion rule, therefore, it is more advantageous to evolve switching phenotypes than faster growth rates.

To assess how stringent the exclusion rule must be to favour switchers, we varied the proportion of the immediate ancestor excluded. The switcher is selected more than half of the time until the exclusion rule is relaxed to the point where the amount of the immediate ancestor available to seed the next round is of the same order of magnitude as the novel phenotypes (figure 2c). The exclusion rule, therefore, can be relaxed without jeopardizing the success of the switcher.

With the exclusion rule in place, switchers overcome the cost of growing 25 per cent more slowly. To determine how robust switching is to a reduction in growth rate, we simulated the model for a range of doubling times from parity to a twofold increase (figure 3a). The switcher was selected over 60 per cent of the rounds until the doubling time of the switcher was twice that of the non-switching phenotypes. We repeated the simulations extending the round length from 3.5 to 4.0 time units: the switcher was still selected approximately 50 per cent of the rounds at a doubling time of 0.4. Extending doubling time to 0.8 does not significantly change the probability of selecting a switcher. This suggests that if round lengths are short enough (i.e. selection occurs quickly) then the evolutionary advantage to switching can more than compensate for the cost associated with a reduced growth rate.

The exclusion rule hampers non-switching phenotypes because they must rely on reduced mutation rates to produce a new phenotype. So far we have held the mutation probability between X_1 and X_2 to be 10^{-4}. We vary this probability from 10^{-5} to 10^{-3}, the lowest and highest probabilities used in the simulations, and calculate the number of times a switcher is selected (figure 3b). At a mutation probability of 10^{-5} the switcher wins over 90 per cent of the rounds, but as the mutation probability increases past 3.0×10^{-4} the switcher loses more often than the non-switching phenotypes.

The previous simulations used rounds lasting 3.5 time units, mimicking the speed of immune responses [38] and the short round lengths used by Beaumont *et al.* [32]. Yet, the round length affects the number of bacteria and, consequently, both the likelihood of a mutation and the impact of the growth disparity. We varied the round length between 3.0 and 6.0 time units and calculated how often a switcher was selected (figure 3c). For very short round lengths ($t < 3.0$), few cell divisions occur, offering few opportunities for new types to arise by

mutation. Thus, when an experiment is founded with a non-switching phenotype there is insufficient time for a switching phenotype to appear and vice versa. As the round length increases (above 3.0 time units) opportunity for new types to arise by mutation increases: should a switcher be selected as the next inoculum its capacity to generate a novel type at high frequency means that it will readily outcompete non-switching types. With further increases in round length, faster growing non-switching phenotypes outpace the slower dividing switcher.

Throughout these simulations, we have maintained a strict bottleneck in which only one bacterium is chosen to inoculate a future round. Previous research has investigated the role of bottlenecks in selecting for fitness-improving mutations and subsequent effects on population fitness [35,39–41]. Here, our interest is on the effect of bottleneck events on the evolution of switching types with reduced growth rates. The major effect of such events is to eliminate competition with conspecifics. We repeated the simulations using different bottleneck sizes (from 1 to 501 in figure 3d). In each case, the bottleneck size also determined the initial number of bacteria for the next round. We applied the exclusion rule by preventing the most frequent phenotype inoculated from passing on to the next round. Because the switcher does not grow as fast as the non-switching types, the switcher benefits from smaller bottleneck sizes. As the bottleneck size increases from 1 to 10, there is an almost 20 per cent decrease in the probability of selecting a switcher. Larger bottlenecks permit faster growing, non-switching types to pass through to the next round where they can outgrow switching types. The benefit accruing to switchers arises from the fact that competition with more fit, non-switching types is eliminated [42]. The curve in figure 3d reaches an asymptote because for large bottlenecks, i.e. large inoculation sizes, there is a high probability that all possible phenotypes will arise in the first round of cell division.

The defining characteristic of a switcher is its rate of generating a different phenotype. Both experimental and theoretical studies have shown that if one phenotype grows poorly in an environment then the switching rate should be tuned to the rate at which the environments fluctuate [19,23,25–27,30,31]. In these studies there was a fitness trade-off such that the each type had a preferred environment. In our model, the switcher shows no such specialization and the exclusion rule no such specificity. We calculate the switching rate that ensures a switcher the best chance of being selected over two rounds of growth and selection (figure 4a). Because our switching phenotypes do not specialize to environments the optimal switching rate is symmetric so that S_A produces S_B at the same rate as S_B produces S_A. Additionally, the switcher is rewarded for infidelity: the best strategy is to immediately produce the opposite phenotype, ensuring the switcher a chance to be selected for the next round. The contours demonstrate that a switcher can tolerate a biased switching rate and still be selected; as long as the switching rate is at least an order of magnitude above the non-switching phenotype's mutation rate, the switcher dominates. If the switcher is biased such that one phenotype grows poorly (figure 4b), the optimal switching rate is also biased. Like previous theoretical work [25,27,31,38], the switching rate favours the faster growing phenotype. Nonetheless, the switcher can be
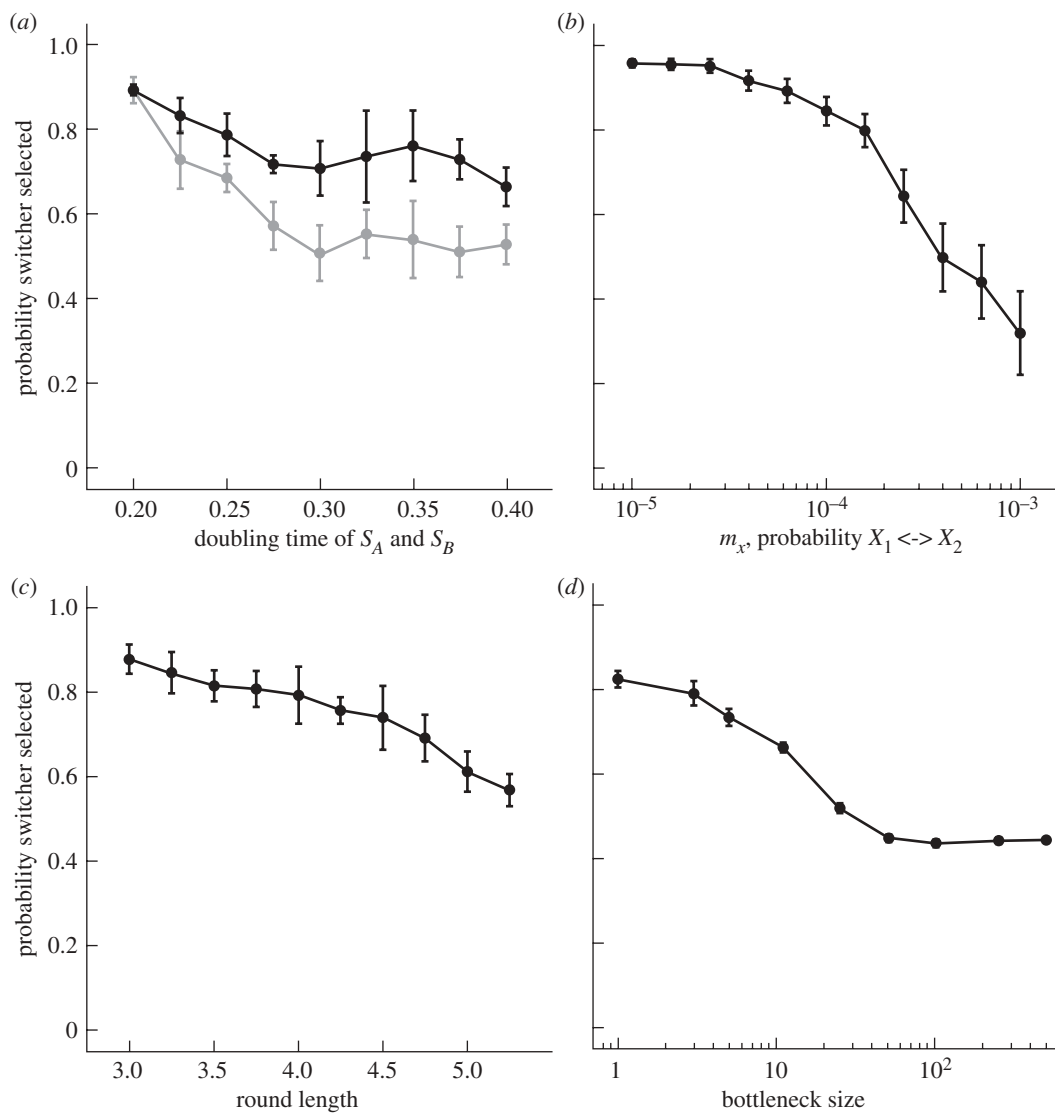
Figure 3. Resilience of a switcher. (a) To examine the effect of growth rate on the invasion of a switcher from rare, we varied the doubling time for S_A and S_B from 0.2 to 0.4 and calculated the number of times a switcher was selected for rounds of length 3.5 (black) and 4.0 (grey) time units. As in figure 2c (and the rest of figure 3), each point is a mean of five samples of 10 simulations run for 125 rounds with the first 25 discarded to avoid biasing the results in favour of the initial inoculum (the X phenotype). The error bars show the standard deviation. The switcher can invade from rare despite longer doubling times. (b) We repeat the process as in (a) except the mutation probability is varied between non-switching phenotypes from 10^{-5} to 10^{-3}. The switcher wins more than 50% of the time until a mutation probability above 0.0003, when the switching rate is only 3.3 times greater. (c) The round length ranges from 3.0 to 6.0 (a round length of 3.5 is used through the rest of the study). Increasing the round length gives more time for the faster growing non-switchers to outcompete the switcher. Even at a round length of 6.0 when the switcher divides up to six times less than the non-switching phenotypes, it still is selected in almost 50% of the rounds. (d) The bottleneck size is varied and the exclusion rule is implemented by preventing the organism that represents the majority of an inoculum from continuing on to the next round. The switcher wins more often than non-switching phenotypes so long as the bottleneck is below 10, otherwise faster growing phenotypes outcompete the switcher because they are no longer hampered by their diminished mutation probability.

selected over two rounds with greater than 50 per cent probability if the switching rate is biased against the faster grower.

4. DISCUSSION

Our mathematical model and ensuing simulations show that the de novo evolution of stochastic phenotype switching observed in the experiment of Beaumont *et al.* [32] was unlikely to have been a result of chance. Together, exclusion rules and population bottlenecks—two complementary faces of frequency dependent selection—define a set of ecological conditions that select for stochastic phenotype switching, such that even when initially rare, and when switching engenders a cost in Malthusian fitness, organisms with the capacity to switch can invade non-switching populations and replace non-switching types. Simulations demonstrate the robustness of our findings to alterations in switching rate, fidelity of exclusion, bottleneck size, duration of environmental state and growth rate.

In addition to a general recipe for the evolution of stochastic switching, the theoretical work makes a number of experimentally testable predictions. Firstly, it suggests that evolution of switcher genotypes in the experiment of Beaumont *et al.* [32], while requiring the exclusion rule and bottleneck, was not dependent upon alternating
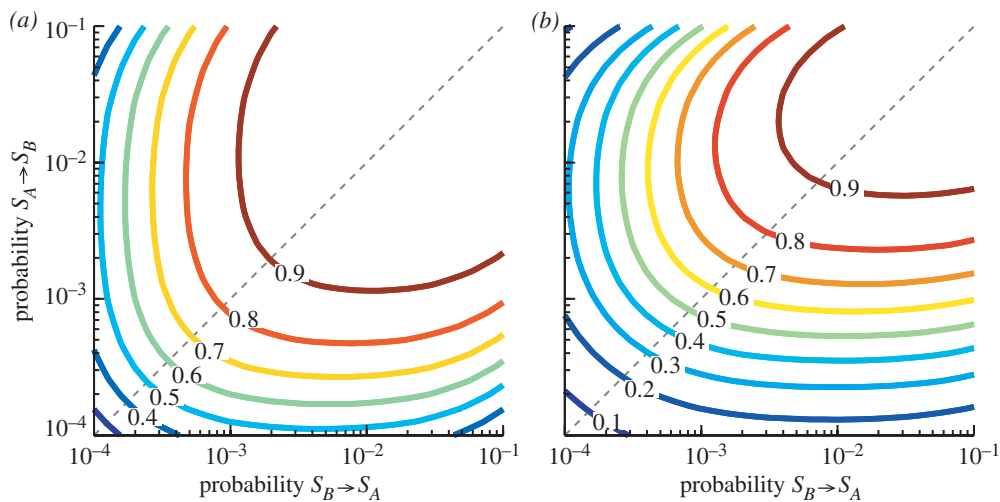
Figure 4. Optimal switching rate. (*a*) The probability a switcher is selected over two rounds plotted for different switching probabilities. The switching phenotypes are identical so the contours are symmetric through the identity line (dashed). The surface is highest when each phenotype is more likely to produce the opposite phenotype. This strategy ensures that there is at least one organism of each phenotype guaranteeing a high probability of passing through to the next round. (*b*) Contour plot as in (*a*) except the S_B phenotype grows with a doubling time of 0.5 rather than 0.25. The switching rate provides an advantage to the faster growing phenotype, but switchers have higher probabilities of surviving two rounds of selection, despite a suboptimal switching probability.

periods of selection in contrasting environments. Frequent imposition of the exclusion rule and bottleneck to populations propagated in either static or shaken microcosms would have alone been sufficient. However, from a practical point of view, transferring populations between static and shaking environments is likely to have increased the chances of identifying phenotypically distinct types because of opportunity to experimentally impose the exclusion rule. Secondly, the model predicts that fluctuating selection is a necessary, but not sufficient condition for the evolution of stochastic switching. Fluctuating selection need not generate strong frequency-dependent selection: in the absence of phenotypic exclusion switchers are unlikely to evolve. Thirdly, altering the experimental design to increase round length or decrease potency of the bottleneck is predicted to make selection of switchers less likely. Fourthly, the model predicts that switchers will persist over the long term, provided exclusion rules and bottlenecks remain features of the selective environment.

The model also predicts that the phenotypic states achieved by the stochastic switching types are of secondary significance. A selective regime involving strong frequency-dependent selection selects entities that generate phenotypic novelty: our model shows that these entities can be adaptive despite poor ecological performance of each phenotypic state. In this regard there is a close parallel with the interaction between pathogenic bacteria and the host immune response: antigen recognition ensures that bacteria experience frequent and repeated population bottlenecks and phenotypic exclusion [10,25]. Survival in the face of such a challenge stems from avoidance of recognition (being different), rather than generation of types fit to different manifestations of the external environment. Both the experiment and theory developed here, emphasizes bet hedging as an adaptive response, not just to changes in the environment, but to change itself [2,4,5].

In contrast to previous theoretical work [19,25,27, 30,31], we did not require a fit between phenotypic and environmental states: switcher types grew more slowly than their non-switching counterparts in *all* environments. This removed the constraint that a switcher must not only evolve a switch but also produce two phenotypes adapted to different environmental states. By focusing solely on the emergence of the capacity to switch—and showing that switching can invade from rare despite significant costs—our work sheds light on the evolutionary origins of bet hedging and has further relevance for the more general problem of the evolution of genetic switches. Clearly, the capacity to stochastically switch can be in-and-of-itself adaptive, even when the ecological performance of variant types is poor: the key evolutionary innovation would thus appear to be the capacity to switch, with the fit between phenotypic and environmental states likely to be the product of further evolutionary refinement. A further distinction between our model and other recent studies [19,25,27,30,31] is incorporation of a stringent bottleneck. The bottleneck means that arithmetic gains in population growth during one round of selection are lost upon transfer to the next round. This reduces the barrier a switcher faces in order to go to fixation.

While we have not stated a specific mechanism of stochastic switching our model is likely to be applicable to a range of possible scenarios. The fully connected model (figure 1) applies to switchers that move between phenotypic states by an epigenetic mechanism where X_1 generates by a single mutation types S_A or S_B (or X_2). In many organisms switching between states involves DNA sequence change such that X_1 generates S_A (or X_2), with S_B realizable following an additional mutation. Omitting direct connections between X_1 and S_B (and also S_A and X_2) should not significantly change the results as it would amount to reducing the m_{t1} and m_{t2} rates by 0.5. This does not affect either m_x or m_s and essentially makes transitions between non-switching phenotypes and switching phenotypes occur less often. Thus, a switcher may take longer to rise to prominence but once it does it will dominate for longer periods.

In order to keep the model both simple and general, several assumptions were made that mark a departure from strict biological reality. Firstly, we assume a large pool of mutations generating novel phenotypes. Secondly, we make no distinction among mutational effects. Thirdly, perfect bottlenecks and exclusion rules were imposed in the absence of biological noise. In real systems, it is unlikely that environments will fully exclude the immediate phenotype nor restrict population size to the same number during each switch of the environment. While the effect of noise, mutations of different fitness effects and other evolutionary processes are yet to be determined, our simulations explored a number of relaxations of the strict exclusion rule and bottleneck and despite highly conservative choices of mutation and fitness parameters, we found entities capable of stochastic phenotype switching to invade from rare over a broad range of conditions.

An additional issue arises from the fact that we place our work in the context of stochastic switching in bacteria and yet ignore the diversity of hosts and migration of bacteria between hosts. Although, we do not explicitly consider this broader scale of heterogeneity our model readily accommodates this dimension. For example, a rapidly evolving (but non-switching) and readily transmissible pathogen can reasonably be considered a switcher. Imagine a collection of hosts each carrying a non-switching pathogen that mutates to generate a new phenotype. Assuming opportunity for transmission then different non-switching pathogens migrate to different hosts and continue evolving. This is equivalent to an increase in the rate at which a non-switcher type generates a novel phenotype, i.e. $X_1 \rightarrow X_2$ happens with higher probability and is shown in figure 3.

Finally, the simple and general nature of our model means it can be readily extended to a range of biological settings characterized by fluctuating selection. For example, exclusion rules, strong bottlenecks and short round lengths are hallmarks of an effective immune response an environment where there is a high prevalence of stochastic phenotype switching [9,10,25]. Similar ecological conditions are also features of antagonistic coevolution between parasites and their hosts. Bet hedging in this instance takes the form of sexual reproduction [43] (see the electronic supplementary material for applications to the evolution of sex). The selective conditions captured by our model also tie together models of bacterial persistence [23,30] with work that investigates the effects of bottlenecks on fitness-improving mutations [35,39–41]. Exclusion rules and bottlenecks acting in tandem are analogous to a death rate or catastrophic event. In this regard our model may prove useful for exploring ecological conditions, such as those generated by antibiotic therapy, likely to promote the evolution of bacterial persistence [23,30]. The accompanying electronic supplementary material also shows how our model might be applied to the problem of bacterial persistence.

REFERENCES

1. Via, S. & Lande, R. 1985 Genotype–environment interaction and the evolution of phenotypic plasticity. *Evolution* **39**, 505–522. (doi:10.2307/2408649)
2. Cohen, D. 1966 Optimizing reproduction in a randomly varying environment. *J. Theoret. Biol.* **12**, 119–129. (doi:10.1016/0022-5193(66)90188-3)
3. Slatkin, M. 1974 Hedging one's evolutionary bets. *Nature* **250**, 704–705. (doi:10.1038/250704b0)
4. Seger, J. B. & Brockmann, H. 1987 What is bet hedging? In *Oxford surveys in evolutionary biology*, vol. 4 (eds P. Harvey & L. Partridge), pp. 182–211. Oxford, UK: Oxford University Press.
5. Bull, J. J. 1987 Evolution of phenotypic variance. *Evolution* **41**, 303–315. (doi:10.2307/2409140)
6. Levins, R. 1962 Theory of fitness in a heterogeneous environment. I. The fitness set and the adaptive function. *Am. Nat.* **96**, 361–373. (doi:10.1086/282245)
7. Donaldson-Matasci, M. C., Lachmann, M. & Bergstrom, C. T. 2008 Phenotypic diversity as an adaptation to environmental uncertainty. *Evol. Ecol. Res.* **10**, 493–515.
8. Meyers, L. A. & Bull, J. J. 2002 Fighting change with change: adaptive variation in an uncertain world. *Trends Ecol. Evol.* **17**, 551–557. (doi:10.1016/S0169-5347(02)02633-2)
9. Andrewes, F. W. 1922 Studies in group-agglutination. I. The *Salmonella* group and its antigenic structure. *J. Path. Bacteriol.* **25**, 505–521. (doi:10.1002/path.1700250411)
10. Moxon, E. R., Rainey, P. B., Nowak, M. A. & Lenski, R. E. 1994 Adaptive evolution of highly mutable loci in pathogenic bacteria. *Curr. Biol.* **4**, 24–33. (doi:10.1016/S0960-9822(00)00005-1)
11. Kaern, M., Elston, T. C., Blake, W. J. & Collins, J. J. 2005 Stochasticity in gene expression: from theories to phenotypes. *Nat. Rev. Genet.* **6**, 451–464. (doi:10.1038/nrg1615)
12. Smits, W. K., Kuipers, O. P. & Veening, J.-W. 2006 Phenotypic variation in bacteria: the role of feedback regulation. *Nat. Rev. Microbiol.* **4**, 259–271. (doi:10.1038/nrmicro1381)
13. Elowitz, M. B., Levine, A. J., Siggia, E. D. & Swain, P. S. 2002 Stochastic gene expression in a single cell. *Science* **297**, 1183–1186. (doi:10.1126/science.1070919)
14. Maamar, H., Raj, A. & Dubnau, D. 2007 Noise in gene expression determines cell fate in *Bacillus subtilis*. *Science* **317**, 526–529. (doi:10.1126/science.1140818)
15. Lim, H. N. & van Oudenaarden, A. 2007 A multistep epigenetic switch enables the stable inheritance of DNA methylation states. *Nat. Genet.* **39**, 269–275. (doi:10.1038/ng1956)
16. Freed, N. E., Silander, O. K., Stecher, B., Böhm, A., Hardt, W.-D. & Ackermann, M. 2008 A simple screen to identify promoters conferring high levels of phenotypic noise. *PLoS Genet.* **4**, e1000307. (doi:10.1371/journal.pgen.1000307)
17. Keren, I., Kaldalu, N., Spoering, A., Wang, Y. & Lewis, K. 2004 Persister cells and tolerance to antimicrobials. *FEMS Microbiol. Lett.* **230**, 13–18. (doi:10.1016/S0378-1097(03)00856-5)
18. Balaban, N. Q., Merrin, J., Chait, R., Kowalik, L. & Leibler, S. 2004 Bacterial persistence as a phenotypic switch. *Science* **305**, 1622–1625. (doi:10.1126/science.1099390)

19 Kussell, E. & Leibler, S. 2005 Phenotypic diversity, population growth, and information in fluctuating environments. *Science* **309**, 2075–2078. (doi:10.1126/science.1114383)

20 Johnsen, P. J., Dubnau, D. & Levin, B. R. 2009 Episodic selection and the maintenance of competence and natural transformation in *Bacillus subtilis*. *Genetics* **181**, 1521–1533. (doi:10.1534/genetics.108.099523)

21 Moxon, R., Bayliss, C. & Hood, D. 2006 Bacterial contingency loci: the role of simple sequence DNA repeats in bacterial adaptation. *Annu. Rev. Genet.* **40**, 307–333. (doi:10.1146/annurev.genet.40.110405.090442)

22 Maizels, N. 2005 Immunoglobulin gene diversification. *Annu. Rev. Genet.* **39**, 23–46. (doi:10.1146/annurev.genet.39.073003.110544)

23 Kussell, E., Kishony, R., Balaban, N. Q. & Leibler, S. 2005 Bacterial persistence: a model of survival in changing environments. *Genetics* **169**, 1807–1814. (doi:10.1534/genetics.104.035352)

24 King, O. D. & Masel, J. 2007 The evolution of bet-hedging adaptations to rare scenarios. *Theoret. Popul. Biol.* **72**, 560–575. (doi:10.1016/j.tpb.2007.08.006)

25 Thattai, M. & van Oudenaarden, A. 2004 Stochastic gene expression in fluctuating environments. *Genetics* **167**, 523–530. (doi:10.1534/genetics.167.1.523)

26 Acar, M., Mettetal, J. T. & van Oudenaarden, A. 2008 Stochastic switching as a survival strategy in fluctuating environments. *Nat. Genet.* **40**, 471–475. (doi:10.1038/ng.110)

27 Gaal, B., Pitchford, J. W. & Wood, A. J. 2010 Exact results for the evolution of stochastic switching in variable asymmetric environments. *Genetics* **184**, 1113–1119. (doi:10.1534/genetics.109.113431)

28 Wolf, D. M., Vazirani, V. V. & Arkin, A. P. 2005 A microbial modified prisoner's dilemma game: how frequency-dependent selection can lead to random phase variation. *J. Theoret. Biol.* **234**, 255–262. (doi:10.1016/j.jtbi.2004.11.021)

29 Donaldson-Matasci, M. C., Bergstrom, C. T. & Lachmann, M. 2010 The fitness value of information. *Oikos* **119**, 219–230. (doi:10.1111/j.1600-0706.2009.17781.x)

30 Visco, P., Alled, R. J., Majumdar, S. N. & Evans, M. R. 2010 Switching and growth for microbial populations in catastrophic responsive environments. *Biophys. J.* **98**, 1099–1108. (doi:10.1016/j.bpj.2009.11.049)

31 Salathe, M., Van Cleve, J. & Feldman, M. W. 2009 Evolution of stochastic switching rates in asymmetric fitness landscapes. *Genetics* **182**, 1159–1164. (doi:10.1534/genetics.109.103333)

32 Beaumont, H. J. E., Gallie, J., Kost, C., Ferguson, G. & Rainey, P. B. 2009 Experimental evolution of bet-hedging. *Nature* **462**, 90–93. (doi:10.1038/nature08504)

33 Gordon, A. J. E., Halliday, J. A., Blankschien, M. D., Burns, P. A., Yatagai, F. & Herman, C. 2009 Transcriptional infidelity promotes heritable phenotypic change in a bistable gene network. *PLoS Biol.* **7**, e1000044. (doi:10.1371/journal.pbio.1000044)

34 Boothroyd, C. E., Dreesen, O., Leonova, T., Ly, K. I., Figueiredo, L. M., Cross, G. A. M. & Papavasiliou, F. N. 2009 A yeast-endonuclease-generated DNA break induces antigenic switching in *Trypanosoma brucei*. *Nature* **459**, 278–281. (doi:10.1038/nature07982)

35 Levin, B. R., Perrot, V. & Walker, N. 2000 Compensatory mutations, antibiotic resistance and the population genetics of adaptive evolution in bacteria. *Genetics* **154**, 985–997.

36 Pal, C., Macia, M. D., Oliver, A., Schachar, I. & Buckling, A. 2007 Coevolution with viruses drives the evolution of bacterial mutation rates. *Nature* **450**, 1079–1081. (doi:10.1038/nature06350)

37 Irwin, P. L., Nguyen, L.-H. T, Paoli, G. C. & Chen, C.-Y. 2010 Evidence for a bimodal distribution of *Escherichia coli* doubling times below a threshold initial cell concentration. *BMC Microbiol.* **10**, 207. (doi:10.1186/1471-2180-10-207)

38 Wolf, D. M., Vazirani, V. V. & Arkin, A. P. 2005 Diversity in times of adversity: probabilistic strategies in microbial survival games. *J. Theoret. Biol.* **234**, 227–253. (doi:10.1016/j.jtbi.2004.11.020)

39 Wahl, L. M., Gerrish, P. J. & Saika-Voivod, I. 2002 Evaluating the impact of population bottlenecks in experimental evolution. *Genetics* **162**, 961–971.

40 Bergstrom, C. T., McElhany, P. & Real, L. A. 1999 Transmission bottlenecks as determinants of virulence in rapidly evolving pathogens. *Proc. Natl Acad. Sci. USA* **96**, 5095–5100. (doi:10.1073/pnas.96.9.5095)

41 Hubbarde, J. E. & Wahl, L. M. 2008 Estimating the optimal bottleneck ratio for experimental evolution: the burst-death model. *Math. Biosci.* **213**, 113–118. (doi:10.1016/j.mbs.2008.03.006)

42 Willi, Y., Van Buskirk, J. & Hoffmann, A. A. 2006 Limits to the adaptive potential of small populations. *Annu. Rev. Ecol. Evol. Syst.* **37**, 433–458. (doi:10.1146/annurev.ecolsys.37.091305.110145)

43 Williams, G. 1975 *Sex and evolution*. Princeton, NJ: Princeton University Press.