

Epistasis and the dynamics of reversion in molecular evolution

David M. McCandlish^{*1}, Premal Shah², and Joshua B. Plotkin¹

¹Department of Biology, University of Pennsylvania, Philadelphia, PA

²Department of Genetics, Rutgers University, Piscataway, NJ

Abstract

Recent studies of protein evolution contend that the longer an amino acid substitution is present at a site, the less likely it is to revert to the amino acid previously occupying that site. Here we study this phenomenon of decreasing reversion rates rigorously, and in a much more general context. We show that, under weak mutation and for arbitrary fitness landscapes, reversion rates decrease with time for any site that is involved in at least one epistatic interaction. Specifically, we prove that, at stationarity, the hazard function of the distribution of waiting times until reversion is strictly decreasing for any such site. Thus, in the presence of epistasis, the longer a particular character has been absent from a site, the less likely the site will revert to its prior state. We also explore several examples of this general result, which share a common pattern whereby the probability of having reverted increases rapidly at short times to some substantial value before becoming almost flat after a few substitutions at other sites. This pattern indicates a characteristic tendency for reversion to occur either almost immediately after the initial substitution or only after a very long time.

Keywords: weak mutation, fitness landscape, entrenchment, reversible Markov chain

^{*}To whom correspondence should be addressed. E-mail: davidmc@sas.upenn.edu

Introduction

In the context of evolutionary theory, reversion describes a population that returns to an ancestral character state (Porter and Crandall 2003). While many early (Dollo 1893; Muller 1939; Simpson 1953; Gould 1970) and more recent (Teotónio and Rose 2001; Collin and Miglietta 2008; Bridgham et al. 2009; Tan et al. 2011) discussions of reversion consider an environmental change that confers a selective advantage to an ancestral phenotype, reversion may also occur at the level of nucleic acids or protein sequences, with evolution proceeding under longterm purifying selection (Kimura 1983). Such reversions occur both because of the strictly limited number of character states (four possible nucleotides or twenty possible amino acids, Jukes and Cantor 1969) and because selection on molecular function may constrain a given position to only a subset of these possible character states (Rokas and Carroll 2008; Breen et al. 2012).

It has long been hypothesized that epistatic interactions should lower the rate of reversion, rendering evolution effectively irreversible (Muller 1918, 1939). This issue has been especially important recently, due to ongoing debate in the field of protein evolution about how position-specific preferences for amino acids may change over time (Pollock et al. 2012; Naumenko et al. 2012; Ashenberg et al. 2013; Pollock and Goldstein 2014; Risso et al. 2015; Shah et al. 2015; Usmanova et al. 2015; Goldstein et al. 2015; Bazykin 2015). In particular, several groups have suggested that once an amino acid substitution occurs at a particular position, epistatic interactions with subsequent substitutions at other positions should tend to increase the selective preference for the derived amino acid relative to the ancestral state (Pollock et al. 2012; Naumenko et al. 2012; Shah et al. 2015, c.f. Fisher 1930, pg. 95), a phenomenon known as entrenchment (Shah et al. 2015). This means that a mutation that was nearly neutral when it originally went to fixation may become increasingly deleterious to revert, which would cause a decreasing propensity to revert as time elapses.

However, the above verbal argument is not entirely convincing. While it is easy to imagine some forms of epistasis that would cause reversion rates to decrease over time, evolutionary dynamics on high-dimensional fitness landscapes can have many counter-intuitive properties (Conrad 1990; Gavrillets 1997; Carneiro and Hartl 2010; Kondrashov and Kondrashov 2015; McCandlish et al. 2013, 2015b). Here we undertake a rigorous mathematical investigation into the relationship between the rate of reversion and the presence of epistasis, for arbitrary fitness landscapes. We study this problem under the assumption that mutation is weak relative to drift, so that the evolution of the population can be modeled as a Markov chain on the set of genotypes (for a review see McCandlish and Stoltzfus 2014). Our main conclusion is that, for any site involved in at least one epistatic interaction, the rate of reversion for a substitution is a strictly decreasing function of the time since the initial substitution.

Our first task is to provide a rigorous definition for the rate of reversion. For a population evolving under weak mutation, we can consider the population as a single particle that jumps

from one genotype on the fitness landscape to another at each substitution event. We consider some focal set of genotypes that includes the starting state of the population. If we observe the population for long enough, the population will eventually leave this focal set and trace a path through the space of genotypes. At each point along this path, it has some propensity to fix a genotype in the focal set, that is, to revert. Eventually, this propensity is realized and the population returns to the focal set. If we continue to watch the population for long enough, this process will repeat itself many times, and we can ask the question: given that the population left the focal set t time units ago and has not yet returned to it, what is the instantaneous expected propensity for that population to return to the focal set, i.e. what is the rate of reversion as a function of time?

To study reversion it is helpful to note the following relationship between the rate of reversion and the distribution of waiting times until a reversion event occurs. As we observe the population evolve on the fitness landscape, every time the population leaves the focal subset, we can record the waiting time until it first returns. And again, if we observe the population for long enough, these waiting times will converge to a particular distribution. Importantly, the reversion rate described above is equal to the hazard function of this probability distribution of reversion times, that is, the probability density of this distribution at time t , conditioned on drawing a value t or greater. Thus, we can study how the rate of reversion changes the longer the population has been absent from the focal subset by studying the hazard function of this distribution of reversion times.

The case of a bi-allelic non-epistatic (i.e. additive) fitness landscape provides an instructive, introductory example. In this case, it is easy to show that the distribution of return times for a particular allele at a particular site is always exponentially distributed, corresponding to a constant hazard function. That is, for a bi-allelic site on a non-epistatic fitness landscape, the reversion rate does not change as a function of time. We want to understand how this simple situation changes in the presence of epistasis.

Here we show that if the focal site interacts epistatically with at least one other site, then the hazard function of the distribution of reversion times, and therefore the rate of reversion, is strictly decreasing in time. This implies that the longer a population has been away from the focal set of genotypes, the longer the expected waiting time until it returns to that set. Moreover, this decreasing reversion rate is due to two factors, each of which would individually result in a decreasing reversion rate.

The first factor is co-evolution between sites as suggested by, e.g., Pollock et al. (2012). As long as the derived allele is resident at the focal site, it forms part of the genetic background for other substitutions, and this causes the population to tend to spend more time at genotypes where the derived allele is selectively favored.

To isolate the effect of this first factor, we consider a modified process where we do not allow the focal site to return to its original state after the initial substitution. Thus, the dynamics after

the initial substitution capture the acclimatization of the rest of the genome to the derived state at the focal site. While reversion events cannot occur under this modified model, we nonetheless keep track of the reversion rate that would occur if we were to suddenly allow reversions. We show that the reversion rate for this modified process is decreasing for sites involved in at least one epistatic interaction, which demonstrates that co-evolution between sites leads to reversion rates that decrease in time.

The second factor that produces decreasing rates of reversion is statistical in nature. This second factor arises because in order to focus on the first time a population returns to the ancestral state we must condition on that return not having yet occurred when we calculate the rate of reversion. If the population is at a genotype with a high reversion rate, it tends to actually revert, so that the high reversion rate no longer contributes to the expectation. This alone results in reversion rates that decrease in time. To put this in a more biological light, populations that have been gone for a long time from the focal subset tend to have low propensity to return to the focal subset, because if they had a high propensity they would have returned already.

To isolate the effects of this second factor, we can consider a different, modified model in which the population never moves to another genotype outside the focal set once it leaves the focal set. Thus, each time the population leaves the focal set, its propensity to return to the focal set is constant. Nonetheless, the reversion rate under this model will be strictly decreasing in time if there is any variation among these genotypes in the propensity to return to the focal set.

In actuality, these two factors operate simultaneously, and their effects on the time-evolution of the rate of reversion are coupled. Nonetheless, when both factors operate, we will show that the reversion rate is still decreasing with the time since the substitution of the derived allele at the focal site.

We will also consider what occurs when more than two alleles are available at a site and the more general case of reversion to subsets of genotypes, e.g. reversion to the set of codons corresponding to a particular amino acid. The key observation in this context is that the two factors above operate when there is any genotype-to-genotype variation in the propensity to return to the focal set. While for models with bi-allelic sites epistasis is the only way of producing variation in these propensities, for more general models with more than two alleles per site the rate of reversion may be decreasing even in the absence of epistasis.

In addition to our main results, which concern populations that have already been evolving on the same fitness landscape for a long time, we will briefly explore how changes to the fitness landscape effect the dynamics of reversion. Finally, we will explore several simple examples to gain intuition for the magnitude and evolutionary importance of reversion rates that decrease in time.

Materials and methods

Population-genetic model

We consider a population evolving in continuous time under weak mutation on an arbitrary finite-state fitness landscape (e.g. Iwasa 1988; Sella and Hirsh 2005; McCandlish et al. 2015b). In this regime, we can model the population as a single particle that moves from genotype to genotype at each fixation event (see McCandlish and Stoltzfus 2014, for a review). More formally, we model evolution as a continuous time Markov chain with a rate matrix \mathbf{Q}_{full} :

$$\mathbf{Q}_{\text{full}}(i, j) = \begin{cases} \frac{\mathbf{F}(j) - \mathbf{F}(i)}{1 - e^{-(\mathbf{F}(j) - \mathbf{F}(i))}} \mathbf{M}_{\text{full}}(i, j) & \text{for } i \neq j \\ -\sum_{k \neq i} \mathbf{Q}_{\text{full}}(i, k) & \text{for } i = j \end{cases} \quad (1)$$

where $\mathbf{F}(i)$ is the fitness of genotype i and $\mathbf{M}_{\text{full}}(i, j)$ is the mutation rate from i to j . We further assume that the fitness landscapes is connected, so that there exists a mutational path between any two genotypes i and j and that the Markov chain defined by \mathbf{Q}_{full} is reversible, so that there exists a stationary probability distribution $\boldsymbol{\pi}_{\text{full}}$ of the chain defined by \mathbf{Q}_{full} such that $\boldsymbol{\pi}_{\text{full}}(i) \mathbf{Q}_{\text{full}}(i, j) = \boldsymbol{\pi}_{\text{full}}(j) \mathbf{Q}_{\text{full}}(j, i)$ for all i, j . This latter condition will be satisfied whenever the neutral mutational dynamics produce a reversible Markov chain (see e.g., Sella and Hirsh 2005; McCandlish et al. 2015b); a simple sufficient condition is that the mutation rates are pairwise symmetric, $\mathbf{M}(i, j) = \mathbf{M}(j, i)$ for all i, j .

We are interested in the situation where a population has just left some subset of states A and want to study the waiting time for the population to return to that subset A . Without loss of generality, we can order the states so that all states in A come after the states in the complement of A , so that we can write \mathbf{Q}_{full} in block matrix form as:

$$\mathbf{Q}_{\text{full}} = \left(\begin{array}{c|c} \mathbf{Q}_{A^c} & \mathbf{E}_{A^c, A} \\ \hline \mathbf{E}_{A, A^c} & \mathbf{Q}_A \end{array} \right) \quad (2)$$

where A^c is the complement of A and $\mathbf{E}_{A^c, A}$ gives the transition rates from A^c to A and \mathbf{E}_{A, A^c} gives the transition rates from A to A^c .

Our main object of study will be the absorbing Markov chain with rate matrix \mathbf{Q}_{A^c} , where absorption corresponds to a return to the subset A . Because we will assume that a population starts at time 0 having just left the subset of states A , this means that an absorption event is also a reversion event, so that we can study the dynamics of reversion by studying the waiting time until absorption for the Markov chain with rate matrix \mathbf{Q}_{A^c} . For brevity, we will simply call this matrix \mathbf{Q} .

The row sums of $-\mathbf{Q}$ (or equivalently, the row sums of $\mathbf{E}_{A^c, A}$) give the propensity for a population currently fixed at genotype i to return to A , and we write the rate at which such an

event occurs for a population fixed at genotype i as $\mathbf{g}(i)$. Let $\mathbf{x}_t(i)$ be the probability that the population is fixed for genotype i and has not yet reverted at time t . Then the time evolution of \mathbf{x}_t is given by

$$\mathbf{x}_t^T = \mathbf{x}_0^T e^{\mathbf{Q}t}, \quad (3)$$

where $\mathbf{x}_0(i)$ gives the probability that the population initially left the subset A by becoming fixed for genotype i .

The hazard function of reversion times

Consider a population that first leaves subset A by fixing genotype i . The probability that the population reverts, that is first becomes fixed for a genotype in the subset A , during the time interval $[t, t + dt)$ is given by $f_i(t) dt$ where

$$f_i(t) \equiv \mathbf{x}_t^T \mathbf{g} \quad (4)$$

and $\mathbf{x}_0(j)$ is 1 for $i = j$ and 0 otherwise. Thus, $f_i(t)$ is the probability density function of the distribution of reversion times for a population that initially leaves subset A by fixing genotype i ; we note that $f_i(t)$ is indeed a proper probability density since \mathbf{Q}_{full} defines an ergodic Markov chain and so populations return to the subset A with probability 1.

Now, a population that has already been evolving on a fitness landscape for a long time is much more likely to leave the subset A by fixing some genotypes rather than others. To capture this effect, we can choose the initial distribution \mathbf{x}_0 by considering a population whose genotype is described by the stationary distribution $\boldsymbol{\pi}_{\text{full}}$ and then condition on leaving the subset A in the interval $[0, dt)$. We thus specify the distribution \mathbf{x}_0 as

$$\mathbf{x}_0(i) \propto \sum_{j \in A} \boldsymbol{\pi}_{\text{full}}(j) \mathbf{Q}_{\text{full}}(j, i) \quad (5)$$

$$= \sum_{j \in A} \boldsymbol{\pi}_{\text{full}}(j) \mathbf{Q}_{\text{full}}(j, i) \quad (6)$$

$$= \boldsymbol{\pi}_{\text{full}}(i) \mathbf{g}(i). \quad (7)$$

With this choice of initial distribution, the probability density function for the distribution of reversion times is given by

$$f(t) \equiv \sum_{i \in A^c} \mathbf{x}_0(i) f_i(t). \quad (8)$$

Note that this is the same distribution that we would get if we watched a single population evolve for an infinite amount of time and recorded, each time the population left A , the waiting time to return to A .

We now turn to formally defining the rate of reversion. What we want to understand is how the rate at which a population first returns to some set of states changes the longer a population has been outside that set. This suggests that we should define the reversion rate as the probability density of a population returning to set A for the first time in the time interval $[t, t + dt)$ given that the population has not already returned to set A before time t . In the more general context of non-negative probability distributions, this quantity is known as the hazard function (or failure rate, or force of mortality). Thus, we define the reversion rate to be hazard function of the probability distribution of reversion times.

For instance, consider the distribution of reversion times at stationarity, with density given by $f(t)$, cumulative distribution function $F(t) \equiv \int_0^t f(t)$, and complementary cumulative distribution function $\bar{F}(t) \equiv 1 - F(t)$. Then the reversion rate at time t is given by the hazard function $h(t)$ where

$$h(t) \equiv \frac{f(t)}{\bar{F}(t)}. \quad (9)$$

If $h(t)$ is increasing in t it means that populations that have been away from the set A for a long time on average have instantaneous rates of return to A larger than the average rate of return to A of populations that have only been away for a short time. Likewise, if $h(t)$ is decreasing in t it means that populations that have been away from the set A for a long time on average have instantaneous rates of return to A smaller than the average rate of return to A of populations that have only been away for a short time. Another quantity of interest is the expected remaining time until reversion, given that the population has not yet reverted at time t . This expected waiting time can be expressed in terms of the hazard function as

$$m(t) = \int_0^\infty e^{-\int_t^{t+\tau_2} h(\tau_1) d\tau_1} d\tau_2. \quad (10)$$

Furthermore, from the above equation it is easy to see that if the reversion rate is increasing in time, then the expected remaining waiting time is decreasing in time, whereas if the reversion rate is decreasing in time, then the expected remaining waiting time is increasing.

Results

General theory of reversions

Our first main result is that for a population that has already been evolving on the same fitness landscape for a long time, the rate of reversion to A , $h(t)$, is a non-increasing function of t , where t is the time since the population left A . Moreover, the reversion rate $h(t)$ is strictly decreasing unless all genotypes in A^c have the same reversion rate (i.e. $g(i)$ is constant) in which

case $h(t)$ is also constant. This result shows that populations that have spent a longer time away from A cannot possibly have higher rates of return to the ancestral character state or shorter expected remaining waiting times until return.

This result is a simple consequence of the fact, well-known in the mathematical literature (Kielson 1979; Aldous and Fill 2002), that the distribution of return times to a subset for a stationary, finite-state, reversible, continuous-time Markov chains take the form of a mixture of exponential distributions (in the mathematical literature, such a distribution is known as a “completely monotone” distribution). It is easy to show that the hazard function for such a distribution is strictly decreasing except in the case of a pure exponential distribution. In Appendix 1, we provide elementary proofs of these facts as well as the fact that that the distribution of return times is a pure exponential only in the case when $g(i)$ is constant.

Connection to epistasis

An immediate implication of this result is that the reversion rate for any site in a fitness landscape that is involved in an epistatic interaction must be strictly decreasing in time. This is because epistatic interactions result in variation in the genotype-specific rates of return to A , $g(i)$, since different back-mutations have different fitness consequences.

More formally, consider a bi-allelic fitness landscape with L sites and with forward and backward mutation rates of μ_l and ν_l , respectively, at site l . Thus, $M_{\text{full}}(i, j) = \mu_l$ if j differs from i by a forward mutation at site l , $M_{\text{full}}(i, j) = \nu_l$ if j differs from i by a back mutation at site l and $M_{\text{full}}(i, j) = 0$ otherwise. Furthermore, let us pick a focal site l^* and order the states such that all genotypes that have the allele produced by the forward mutation are in set A and indexed first. We are interested in the distribution of reversion times at site l^* .

We say a site is not involved in any epistatic interactions if its fitness effect is constant, in the sense that if genotype i differs from genotype j by a forward substitution then the scaled selection coefficient $\mathbf{F}(i) - \mathbf{F}(j)$ is equal to some constant S_{l^*} . In this case $g(i)$ is constant and equal to $\nu_{l^*}(i)S_{l^*}/(1 - e^{-S_{l^*}})$, so that $h(t)$ is constant. On the other hand, suppose l^* is involved in an epistatic interaction. Then there exist $i, i' \in A^c$ and $j, j' \in A$ such that both i differs from j and i' differs from j' by a forward mutation at site l^* but $\mathbf{F}(i) - \mathbf{F}(j) \neq \mathbf{F}(i') - \mathbf{F}(j')$. But then $g(i) \neq g(i')$ because $Q_{\text{full}}(i, j) = \nu_{l^*}(\mathbf{F}(j) - \mathbf{F}(i))/(1 - e^{-(\mathbf{F}(j) - \mathbf{F}(i))}) \neq \nu_{l^*}(\mathbf{F}(j') - \mathbf{F}(i'))/(1 - e^{-(\mathbf{F}(j') - \mathbf{F}(i'))}) = Q_{\text{full}}(i', j')$. Thus, in this case $h(t)$ is strictly decreasing.

The above argument shows that $h(t)$ is strictly decreasing for a biallelic fitness landscape if and only if the focal site is involved in at least one epistatic interaction. For fitness landscapes that include more than two alleles at a site, by contrast, the notation and argument become somewhat more involved (Appendix 2), but the end result is weakened to the statement that

epistasis is sufficient for $h(t)$ to be strictly decreasing. Indeed, even for non-epistatic multi-allelic fitness landscapes we generically expect $h(t)$ to be strictly decreasing. This is because even if fitness is additive between sites, the different alleles within a site will have different fitness differences from the focal allele and therefore different probabilities of fixation that result in a non-constant $g(i)$.

Do reversions also become more deleterious?

Besides the relationship between decreasing reversion rates and epistasis, it is interesting to ask about the relationship between decreasing reversion rates and the mean selection coefficient of reversions. This relationship is not necessarily simple. First, this is because different genotypes might have different mutation rates to the subset A , so that the decrease in $h(t)$ might be realized by the non-reverting subset of the population becoming concentrated at genotypes with low mutation rates to the subset A rather than by having more negative selection coefficients for such mutations. However, even if all genotypes in A^c produce mutations to subset A at the same rate the decrease in $h(t)$ might not correspond to reversion mutations becoming more deleterious. This is due to the non-linearity of the function relating the scaled selection coefficient S to the rate of evolution; in particular if S is the scaled selection coefficients, this function is $K(S) = S/(1 - e^{-S})$. However, using the fact that $K(S)$ is concave (McCandlish et al. 2015a), if each genotype in A^c produces mutations at rate μ to a single corresponding genotype in set A , then Jensen's inequality tells us that $h(t)$ is greater than μ times $K(S_t)$, where S_t is the average scaled selection coefficient of the reversion at time t and this average is taken with respect $\mathbf{x}_t/F(t)$. Thus, if for any t we have $h(t) < \mu K(S_0)$, then the mean selection coefficient of a reversion among those populations that have not yet reverted is less than or equal to the mean selection coefficient of a reversion immediately after the initial substitution away from A . This provides a sufficient condition for the mean selection coefficient of reversions to have decreased in time.

Expected reversion times

If the reversion rate is non-increasing over time, then the expected waiting time until reversion must be greater than what would be predicted based on the initial reversion rate alone, i.e. we must have $m(0) \geq 1/h(0)$. In fact, there is a simple formula for the expected time until reversion that holds even if the Markov chain describing the weak mutation dynamics is not reversible. In particular, consider the expected rate of returns to A at stationarity when we condition on a population being in the subset A^c , i.e. $\sum_{i \in A^c} \pi_{\text{full}}(i) g(i) / \sum_{i \in A^c} \pi_{\text{full}}(i)$. Then $m(0)$ is simply the reciprocal of this rate (see Appendix 3).

However, because the reversion rate is non-increasing in time, this mean value does not necessarily provide a very informative summary of the dynamics of reversion. For instance, consider the case where $h(t)$ is initially high, but drops to a very low value. Under these circumstances, it can be the case that a large proportion of populations revert rapidly, but the subset of populations that do not revert rapidly may have extremely long expected reversion times, so that the mean reversion time $m(0)$ is really an average over two very different subsets of populations. Intuitively, this situation will often arise when the set A is the basin of attraction of a fitness peak. Most populations that cross into the fitness valley around A will return to A rather than cross the valley, but a small subset will cross the fitness valley and end up at another fitness peak. The expected return time for this subset of populations might then be extremely long.

Why is the reversion rate decreasing?

In order to gain an intuitive understanding for why the reversion rate is decreasing in time, it is helpful to distinguish between two different phenomena that each contribute to this decrease.

The first phenomenon is the acclimatization of the rest of the genotype to being in the subset A^c . For instance, in the context of protein evolution, once a mutation has fixed at a focal site it forms part of the genetic background that determines the fitness effects of mutations at other sites. For sites that interact epistatically with the focal site, this will tend to favor substitutions whose effects are more positive when the derived allele is present at the focal site.

Such acclimatization of the genome tends to result in a decreasing reversion rate at the focal site. To make this idea precise, let us consider a modified process where we do not allow populations to revert to A after they enter the subset A^c . Following this initial entrance, the dynamics are thus governed by a rate matrix \mathbf{Q}^* where $\mathbf{Q}^* \equiv \mathbf{Q} + \mathbf{D}_g$ and \mathbf{D}_g is the diagonal matrix whose i, i -th entry is $g(i)$. For ease of exposition, let us also assume that the subset A^c is connected (i.e. one can always evolve from any state in A^c to any other state in A^c without returning to A ; the more general case is handled in Appendix 1).

Even though we do not allow reversions to occur back to A under this modified process, we can still keep track of the reversion rate that would occur were we to allow reversions. When the population initially leaves the subset A , it is distributed as $\mathbf{x}_0(i) \propto \pi_{\text{full}}(i)g(i)$. However, as time elapses under the modified process, the probability that the population is fixed for genotype $i \in A^c$ tends to a distribution that is $\propto \pi_{\text{full}}(i)$. Now, the initial reversion rate under the modified process is equal to the expected value of $g(i)$ with respect to the first of these distributions, while the asymptotic reversion rate under the modified process is equal to the expected value of $g(i)$ with respect to the second. Clearly, the reversion rate is higher under the first distribution, since this distribution differs from the second only in that it is more concentrated at values with high $g(i)$. Indeed in Appendix 1 we show the stronger result that the reversion rate is strictly

decreasing for the modified process unless $g(i)$ is constant, in which case the reversion rate is also constant.

The second phenomenon is a statistical phenomenon having to do with the fact that we are interested in the first return to A . Even if there was no acclimatization to the initial substitution (and therefore no co-evolution) the reversion rate would tend to be decreasing in time due to genotype-specific variation in the rates of return to A , i.e. the $g(i)$. This effect occurs because populations that have not reverted even after a long time have likely spent most of this time at genotypes where returns to A are unlikely. To put this another way, if the population had spent a great deal of time at genotypes with high return rates, then it would have returned already. This phenomenon is well-known in the reliability and demography literature, where unaccounted for heterogeneity can result in failure or mortality rates that decrease in time even when individual failure or mortality rates are constant (Vaupel and Yashin 1985).

To isolate this second phenomenon, we consider a second modified process in which a population that leaves A by fixing genotype i experiences no additional substitutions until it returns to A , i.e. the return rate for such a population is always $g(i)$. The rate matrix for this process is thus $-\mathbf{D}_g$. The derivative of the reversion rate of this process is equal to -1 times the variance in g conditional on not having yet reverted. Since variances are non-negative, the derivative of the reversion rate is non-positive, so that the reversion rate is non-increasing.

While we have constructed these two modified processes to separate the effects of acclimatization of the genome and the effects of conditioning on not having yet reverted, these two effects interact to determine the dynamics of the original process. The fact that both phenomena tend to lead to decreasing reversion rates when there is variation in g (or more precisely for the second phenomenon, variation in its non-zero elements), helps clarify why the reversion rate is non-increasing for the original process.

Reversion rates under non-stationary evolution

So far, we have concentrated on the case of a population that has already been evolving on a fixed fitness landscape for a long time, so that each possible way of leaving the subset A is weighted by its stationary probability. However, we can also gain some strong intuitions for the behavior of the reversion rate in the more general case, where we allow the initial probability vector \mathbf{x}_0 to be arbitrary.

In particular, in this more general setting we can show that the time-evolution of the (sign-reversed) reversion rate is isomorphic to the time-evolution of the mean fitness of an infinite population on a suitable fitness landscape. The basic idea is that the time-evolution of the probability distribution describing the genotype of a population conditional on not having yet reverted can be viewed as the time-evolution of the frequencies of genotypes in an infinite

population whose mutational dynamics are specified by the off-diagonal entries of the matrix \mathbf{Q} and where the Malthusian fitness of genotype i is given by $-\mathbf{g}(i)$. That is, $\mathbf{g}(i)$ plays the role of a genotype-specific death rate.

More formally, let $\mathbf{p}_t(i)$ be the probability that a population is fixed for genotype i at time t given that it has not returned to the subset A by time t . Then $\mathbf{p}_t(i) = \mathbf{x}_t(i) / \sum_j \mathbf{x}_t(j)$. Writing \mathbf{I} for the identity matrix $\mathbb{1}$ for the vector of all 1s, and $\mathbf{D}_{-\mathbf{g}}$ for the diagonal matrix with $-\mathbf{g}$ down its main diagonal, we have

$$\frac{d}{dt} \mathbf{p}_t^\top = \frac{(\mathbf{x}_t^\top \mathbb{1}) \mathbf{x}_t^\top \mathbf{Q} - \mathbf{x}_t^\top (-\mathbf{x}_t^\top \mathbf{g})}{(\mathbf{x}_t^\top \mathbb{1})^2} \quad (11)$$

$$= \mathbf{p}_t^\top \mathbf{Q} + \mathbf{p}_t^\top (\mathbf{p}_t^\top \mathbf{g}) \quad (12)$$

$$= \mathbf{p}_t^\top \mathbf{Q}^* + \mathbf{p}_t^\top (\mathbf{D}_{-\mathbf{g}} - (\mathbf{p}_t^\top (-\mathbf{g})) \mathbf{I}). \quad (13)$$

This is simply the “parallel” version of the replicator equation (i.e. where mutation occurs independently from reproduction), with \mathbf{Q}^* as the rate matrix of the mutational process and $-\mathbf{g}$ as the vector of fitnesses. Furthermore, the reversion rate at time t is then given by $\mathbf{p}_t^\top \mathbf{g}$, which is simply the mean fitness $\mathbf{p}_t^\top (-\mathbf{g})$ with its sign reversed. The derivate of the reversion rate is then:

$$\frac{d}{dt} (\mathbf{p}_t^\top \mathbf{g}) = \mathbf{p}_t^\top \mathbf{Q}^* \mathbf{g} + \mathbf{p}_t^\top (\mathbf{D}_{-\mathbf{g}} - (\mathbf{p}_t^\top (-\mathbf{g})) \mathbf{I}) \mathbf{g} \quad (14)$$

$$= \mathbf{p}_t^\top \mathbf{Q}^* \mathbf{g} - \text{Var}_{\mathbf{p}_t} \mathbf{g}, \quad (15)$$

where $\text{Var}_{\mathbf{p}_t} \mathbf{g}$ is the variance in the return rate with respect to the probability distribution \mathbf{p}_t . Here, the term $\mathbf{p}_t^\top \mathbf{Q}^* \mathbf{g}$ gives the effect on the reversion rate due to acclimatization of the genotype to being in the subset A^c , while $-\text{Var}_{\mathbf{p}_t} \mathbf{g}$ captures the effects of conditioning on not having yet reverted. Unlike under stationarity, the effects of acclimatization for the non-stationary case can either increase or decrease the reversion rate. However, conditioning on not having yet reverted still always produces a bias towards a decreasing reversion rate, since $-\text{Var}_{\mathbf{p}_t} \mathbf{g}$ is non-positive.

Examples

The preceding results describe a general tendency for the reversion rate—quantified as the hazard function of the distribution of return times—to be decreasing in the time. Here we present some simple examples to explore the magnitude and evolutionary consequences of this effect.

Crossing fitness valleys

Consider the case of a biallelic fitness landscape with a symmetric fitness valley. In particular, consider the case where genotypes ab and AB have equal fitnesses and scaled selection coefficient S with respect to the valley genotypes Ab and aB and mutations occur independently at each locus at rate 1. Let us now examine the rate of reversion at the first site, so that the set $A = \{ab, aB\}$ and leaving A corresponds to an $a \rightarrow A$ substitution. Given that the population has already been evolving on this fitness landscape for a long time and that an $a \rightarrow A$ substitution has just occurred, we want to understand the distribution of times until an $A \rightarrow a$ substitution occurs as a function of the depth of the fitness valley S .

This distribution of times is easy to understand intuitively. For S less than ≈ 2 the dynamics are approximately neutral and the waiting time for reversion is approximately exponential with mean 1. For larger S , the dynamics are substantially affected by the fitness valley. The best way to understand these dynamics is by considering where the population is immediately after an $a \rightarrow A$ substitution. At stationarity, we know that the frequency of substitutions into the fitness valley must be equal to the frequency of substitutions out of the valley. This means that at stationarity, half the time the population will have just fixed the valley genotype Ab and half the time it will have fixed the peak genotype AB . Populations that fix the peak genotype AB are unlikely to revert in the short-term, whereas populations that fix the valley genotype Ab are likely to move to one or the other peak in the short-term, with equal probability. Thus, roughly speaking, after a short time the population will either have reverted (with probability $1/4$, since half the time it fixes the valley genotype and then immediately reverts half the time) or be fixed at the fitness peak AB (with probability $3/4$). This means that reversions happen either very shortly after the initial substitution—and this occurs $1/4$ of the time—or only after the long waiting time needed for deleterious fixations to occur.

Figure 1 shows these dynamics as a function of S . The top row of panels shows the probability that an $A \rightarrow a$ reversion has occurred as a function of time and the bottom row of panels shows the reversion rate, $h(t)$. The left-most column show the neutral case. The reversion rate is constant and the probability that a reversion has occurred by time t is given by $1 - e^{-t}$. The middle and right columns show the case where $S = 2.5$ and $S = 5$ respectively. In both cases, the reversion rate is high initially, starting at approximately $S/2$ (probability $1/2$ of starting at the valley genotype in which case returns to A occur at roughly rate S). The reversion rate drops rapidly as populations leave the valley genotype, approaching an asymptotic rate of approximately $3/2$ the substitution rate of a deleterious fixation $-S/(1 - e^S)$ (populations that have not yet reverted are likely at the AB fitness peak; reversion occurs either via a direct substitution of a at rate $-S/(1 - e^S)$ or via fixation of the valley genotype Ab , which occurs at rate $-S/(1 - e^S)$, followed by reversion with probability $1/2$). Thus, the reversion rate is initially high but rapidly approaches a much lower rate, which results in a characteristic “knee” in the probability that a population has reverted by time t , where a population has a substantial prob-




Figure 1: Reversion rates for a biallelic fitness landscape with a symmetric fitness valley. Genotypes ab and AB have equal fitness and scaled selective advantage S relative to genotypes Ab and aB , and mutations occur at each site at rate 1 in each direction. The dynamics are shown for returns to the focal subset $A = \{ab, aB\}$ assuming that the population has just left the subset A at stationarity. Each column corresponds to a different value of S . The first row shows the probability that the population has reverted as a function of time, and the second row shows the reversion rate, $h(t)$.

ability of having reverted at short times (corresponding to a steep initial increase in the fraction reverted as a function of time), but populations that do not revert during this initial period have much lower reversion rates (resulting in a slow increase in the fraction reverted after this initial period), see Figure 1 first row, middle and right columns.

Mutational robustness and the rate of reversion

As a second example, we consider a biallelic fitness landscape with one focal site and L other sites, where each genotype is viable with probability p and all viable genotypes are neutral relative to each other. All sites experience forward and reverse mutations at rate 1. This example is motivated by the model used by Kondrashov and colleagues to study long-term purifying selection (Breen et al. 2012; Usmanova et al. 2015). We want to understand the dynamics of reversion at the focal site.

We proceed with a heuristic treatment based on the assumption that $Lp \gg 1$. In this regime, we treat each genotype as having a constant number Lp of neutral neighbors accessible by mutations at non-focal sites. The key idea is that immediately after a substitution at the focal site, the back-substitution is viable and so reversions initially occur at rate 1. However, once an additional substitution at a non-focal site accrues, the probability that the reversion mutation at the focal site is viable is only p . Thus, the reversion rate drops very rapidly from 1 to approximately p .

In particular, after a substitution at the focal site, the expected substitution rate is $1 + Lp$ and the probability that the first substitution is a reversion is $1/(1 + Lp)$. Thus, we can approximate the probability distribution of waiting times until reversion as a mixture of two exponential distributions. The first distribution has rate 1 and is weighted by the probability that revert before any other substitution occurs, $1/(1 + Lp)$. The second exponential distribution has rate p and is weighted by the probability that a population experiences a substitution at another site before the reversion event occurs, $Lp/(1 + Lp)$. This suggests that the probability of having reverted by time t can be approximated by:

$$\bar{F}(t) \approx 1 - \frac{1}{1 + Lp} e^{-(1+Lp)t} - \frac{Lp}{1 + Lp} e^{-t}. \quad (16)$$

Based on a similar argument, we can approximate the reversion rate as

$$h(t) \approx e^{-(1+Lp)t} + p(1 - e^{-(1+Lp)t}), \quad (17)$$

where the first term corresponds to the initial reversion rate of 1, which applies with probability approximately $e^{-(1+Lp)t}$ (the probability of not having yet having experienced the first substitution) and the second term corresponds to the expected reversion rate given that one substitution has occurred, p , weighted by the probability that a substitution has already occurred, $1 - e^{-(1+Lp)t}$.

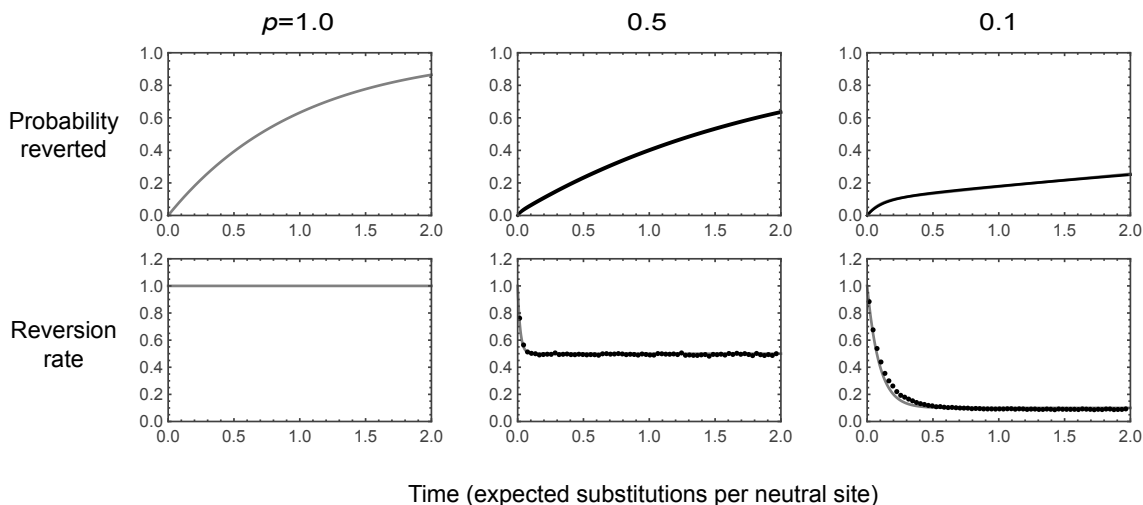


Figure 2: Dynamics of reversion for a biallelic fitness landscape with 100 sites other than the focal site, where each genotype is viable with probability p , all viable genotypes are neutral relative to each other, and each site experiences forward and backward mutations at rate 1. The first row shows the probability that a population has reverted as a function of time, and the second row shows the reversion rate. Theoretical results are in gray, and simulations are in black. Dots in bottom row calculate the reversion rate in terms of temporal bins of width .05, where the reversion rate in a bin is calculated as the number of simulations that revert during that bin divided by the average of the number of simulations that have not yet reverted at the beginning of that bin and the number of simulations that have not yet reverted at the end of that bin. Simulation results are based on 1,000,000 runs for each value of p .

Although an exact treatment of any particular realization of the model is possible using the associated rate matrix \mathbf{Q}_{full} , the construction of this matrix is not computationally feasible for large L . We thus proceeded by simulation. A population that has just experienced a substitution at the focal site must be at a viable genotype and have just come from a viable genotype that differs only at the focal site. We simulated the subsequent evolution of such a population under weak mutation, keeping track of the genotypes produced by mutation and whether they are viable (with probability p) or inviable (with probability $1 - p$) until the first reversion event at the focal site. For each choice of parameters, we repeated this procedure 1,000,000 times.

The resulting reversion times are shown in Figure 2 for $L = 100$. The first row shows the probability of having reverted by time t , and the second row shows the reversion rate as a function of time. The first column shows the exact theoretical results for the case where all genotypes are viable (gray lines). The second and third columns show the results for $p = .5$ and $p = .1$, respectively, with the approximate theoretical results given by Equations 16 and 17 in gray and the simulation results in black. For the case $p = .5$ the reversion rate drops rapidly

from 1 to approximately p . Indeed, this drop is so rapid that relatively few populations actually revert at the higher initial rate, which is as we expect, since the probability that a population reverts before experiencing another substitution is only $\approx 1/(1 + 100 \times .5) \approx .02$. On the other hand, for the case $p = .1$ we see a modest knee in the time evolution of the fraction of populations that have reverted. This is because a larger fraction $\approx 1/(1 + 100 \times .1) \approx .09$ of populations revert before experiencing another substitution, and the subsequent substitution rate is lower, approximately .1 instead of .5.

Intuitively, the knee shape illustrated in Figure 2 arises for small p and large Lp arises because for small p there are fewer other loci at which a substitution can occur, and so populations stay at the genotype they initially arrived at—with its elevated reversion rate—longer, which leads to a substantial probability of having reverted at short times. However, eventually such populations leave the initial genotype, at which time they experience lower reversion rates, $\approx p$. This produces a pattern where the probability of having reverted becomes substantial at short times but thereafter grows very slowly. Although the theory is complicated and so we do not discuss it in detail here, it is worth noting when p is small enough that Lp becomes small the waiting time until reversion again becomes short. Intuitively, this occurs because for very small Lp the only viable mutation after a substitution at the focal site is the back mutation at the focal site, and thus essentially all reversions occur at the initial elevated rate.

More generally, this simple model suggests that the dynamics of reversion should depend on the degree of mutational robustness even when fitness is allowed to take more than two values. The key insight is that just after a substitution at a focal site, the population is not at a random genotype, but rather one where the back substitution rate is likely to be unusually high. Now, if there is a high degree of mutational robustness, then the neutral substitution rate is high, and the population will leave this special genetic background quickly (c.f. McCandlish 2013), so that when most reversions occur, the substitution rate for the back mutation is much lower than it was initially. When the degree of mutational robustness is very low, by contrast, the population will often stay at the initial genotype until the reversion occurs by a direct back substitution, which occurs at the initial rate. When the degree of mutational robustness is intermediate we see a combination of these two dynamics: a substantial fraction of populations revert via the direct back substitution from the initial genotype, and another substantial fraction revert at the lower rates that are typical of other genetic backgrounds. This intermediate degree of robustness thus produces a characteristic knee in the time-evolution of the probability that the population has reverted.

Amino acid reversion in codon models

The two examples we have considered so far describe the dynamics of reversion at a single site in a bi-allelic fitness landscape. For our final example we consider the dynamics to a single

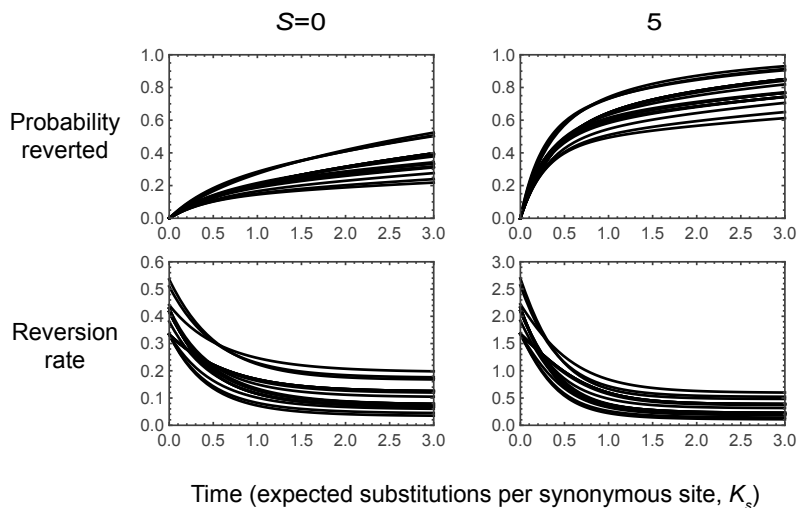


Figure 3: Dynamics of reversion for a codon model where the focal amino acid has scaled selection coefficient S over all other amino acids and mutations occur at rate 1 at each site with an equal probability of producing each of the alternative nucleotides. The top row shows the probability of having reverted by time t and the bottom row shows the reversion rate, $h(t)$, where each of the 20 curves corresponds to a different choice of one of the 20 amino acids as the focal amino acid. Under moderate selection for the focal amino acid ($S = 5$), the time-evolution of the fraction of populations that have already reverted shows a pronounced knee.

amino acid under a codon model. Stop codons are treated as inviable, and mutations occur at each site with equal probability to each of the three alternative nucleotides (Jukes and Cantor 1969) with a total rate of 1 so that time is measured in the expected number of substitutions per synonymous site (K_s). We implement selection by assuming that genotypes with the focal amino acid are neutral relative to each other and have a scaled selective advantage S over all other amino acids.

Figure 3 shows the dynamics of reversion under this model for the neutral case, $S = 0$, and also for the case where the focal amino acid has a moderate selective advantage, $S = 5$ (each amino acid corresponds to one line; the fraction reverted is shown in the first row and $h(t)$ is shown in the second row). The key insight for understanding these dynamics is that, just after a population leaves the set of codons that code for the focal amino acid, it can still mutate back to the focal amino acid. However, after additional substitutions accrue, the population is likely to be at a genotype that is no longer mutationally adjacent to the focal amino acid, leading to a decreasing reversion rate over time. In particular, for both $S = 0$ and $S = 5$, the reversion rate drops from its initial value to nearly its asymptotic value by $K_s = 1$. In the case of $S = 5$, this produces a pronounced knee in the shape of the curve describing the time-evolution of the probability that a population has reverted: more than a third of populations are expected

to revert almost immediately, whereas the remaining populations take a long time before they eventually revert.

Discussion

We have analyzed reversions in a population that has just left some focal set of genotypes A . Our main interest has been the rate at which the population first returns to the set A —that is, the rate at which the population experiences a reversion—and how this rate changes over time. For a population that has already been evolving on a constant fitness landscape for a very long time, we have shown that the rate of reversion to A is non-increasing in time. Furthermore, when the set A consists of all genotypes with a particular allele at a focal site in a bi-allelic fitness landscape, then we showed the reversion rate is strictly decreasing if and only if epistasis is present at the focal site. We have also explored some simple examples of these reversion dynamics, to gain intuition about their magnitude and evolutionary impact.

We have also analyzed the case of a population that has only been evolving on the fitness landscape for a short time, and whose initial state is therefore not drawn from the stationary distribution of the evolutionary dynamics. Here we have shown that the time-evolution of the reversion rate is mathematically identical to the time-evolution of the mean Malthusian fitness for an infinite population evolving on an altered fitness landscape where the substitution rates of the original fitness landscape play the role of mutation rates and the rates of returns to A play the role of genotype-specific death rates.

One consequence of reversion rates that decrease in time is a distinctive pattern, whereby a population either reverts very quickly after the initial substitution with some moderate probability or else takes a very long time to revert. If we consider probability of having reverted as function of time, such a pattern can be seen as a “knee”, where this curve rises rapidly at short times until the probability of having reverted is substantial, and then shifts to rising much more slowly. This pattern occurs because the expected rate of substitutions back to A at stationarity is particularly high if we condition on just having left A . As time passes and other substitutions accrue, the features of the genetic background that resulted in the unusually high substitution rate to A are lost. One consequence of this knee-like phenomenon is that the mean waiting time until reversion may not be very informative about the actual dynamics of reversion.

Our results show that at stationarity, a strictly decreasing reversion rate occurs whenever there is genotype-to-genotype variation in the rate of substitutions to the set A . When the set A corresponds to the set of genotypes with a particular allele at a particular site, the presence of an epistatic interaction between the focal site and at least one other site is sufficient to produce such genotype-specific variation, and hence sufficient for the reversion rate to be strictly decreasing.

However, other factors besides epistasis can also be responsible for genotype-to-genotype variation in the rate of substitutions to A . For instance, in our codon example, which has more than two alleles per site, the structure of the genetic code itself results in variation in the substitution rate to any particular amino acid, because the focal amino acid will not be mutationally accessible from all other codons. More generally, the genotype-phenotype map will tend to produce genotype-to-genotype variation in the rate at which the focal phenotype becomes fixed in the population, so that the reversion rate for phenotypes should typically be decreasing. Finally, it is worth mentioning that while we define a site at the genotypic level by the feature that its mutational dynamics are independent of the states of the other sites, context-dependent mutation could also produce genotype-to-genotype variation in rates of substitution to A , which would again be sufficient for decreasing reversion rates (provided the resulting Markov chain is still reversible).

It is also interesting to ask the converse question: what conditions could possibly produce reversion rates that increase in time? To do this, it is helpful to consider a way of restating our main result: if evolution can be modeled as a reversible Markov chain on the set of genotypes, then at stationarity the reversion rate to any subset of genotypes is non-increasing with the time since the population left that subset. A reversible Markov chain is a Markov chain such that the probability of moving through any cycle of states in one direction is equal to the probability of moving around that cycle in the reverse direction (this is known as Kolmogorov's criterion, see e.g. Kelly 1979, Section 1.5). Thus, a reversible Markov chain is simply a Markov chain with no cyclic biases. This means we can restate our main result as saying that reversion rates can increase in time, starting from stationarity, only if some factor induces a cyclic bias in the Markov chain describing evolution among genotypes.

It is known that, under weak mutation, adding frequency and time-independent natural selection does not add any cyclic bias if no such bias is present in the mutational dynamics (i.e. the weak mutation Markov chain under constant frequency independent selection is reversible if the mutational dynamics are reversible, see e.g. Sella and Hirsh 2005). However, we may expect reversion rates to sometimes increase in time under conditions where fitness is non-transitive (e.g. Kerr et al. 2002) or when environments change in a cyclic manner (e.g. Leslie et al. 2004; Hensley et al. 2009; Bergland et al. 2014). Under these types of conditions, natural selection tends to push populations through a cycle of genotypic states in a periodic manner, which means that reversion becomes more likely as time passes.

Another condition that might produce increasing reversion rates is when the evolutionary dynamics are not stationary, for instance at the beginning of an adaptive transient. We have provided intuition for such circumstances by noting that, under time- and frequency-independent selection, the time evolution of the reversion rate can be recast as the time-evolution of the mean fitness of an infinite population on an alternative fitness landscape. This intuition suggests that the reversion rate should tend to go down just as the mean fitness tends to go up, unless the infinite population starts at an unusually high fitness or the reversion rate starts at an

unusually low value. In the case of an adaptive transient, early adaptive substitutions will tend to be strongly favored by natural selection meaning that the initial reversion rate is unusually low. In these circumstances we may expect the reversion rate to increase, particularly in the case of diminishing-returns epistasis, where mutations that have strong positive effects early in adaptation have much smaller effects towards the end of adaptation (Draghi and Plotkin 2013; Kryazhimskiy et al. 2014, c.f. Hartl et al. 1985; Hartl and Taubes 1996).

Our results help bring clarity to a recent controversy concerning site-specific amino acid preferences during protein evolution (Pollock et al. 2012; Naumenko et al. 2012; Ashenberg et al. 2013; Pollock and Goldstein 2014; Risso et al. 2015; Goldstein et al. 2015; Shah et al. 2015; Usmanova et al. 2015; Bazykin 2015). First, our results show that in the presence of epistasis, reversion rates will be decreasing in time and that the longer a population has left a set of genotypic states, the longer the expected time until reversion. However, our results show that when considering reversion to an amino-acid state, these results already hold even if there is no epistasis at the amino-acid level, because of the structure of the genetic code itself. Indeed, it is worth noting that the epistasis already present in the genetic code is sufficient to produce many of the dynamical signatures of epistatic evolution even if fitnesses are additive at the amino-acid level. For instance, it is easy to specify site-specific amino acid preferences that produce multi-peaked fitness landscapes at individual codons; such codons can have extremely long equilibration times, contrary to the analysis of Kondrashov et al. (2010) and Breen et al. (2013) who only consider the case where every amino acid can mutate to every other amino acid. Overall, decreasing reversion rates are a generic feature of evolution under long-term purifying selection and not a definitive signature of epistasis at the amino-acid level.

Second, it is helpful to distinguish between what is expected when we consider a population that is conditioned to remain outside a subset of states versus one that is simply restricted from entering that subset (c.f. Ashenberg et al. 2013; Pollock and Goldstein 2014). The difference is whether, after a substitution, one only considers populations that have not yet reverted or whether all populations are prevented from reverting (but where we nonetheless keep track of the average propensity to return to the focal subset if such substitutions were to be permitted). We have shown that these two processes have different, but related, mathematical characteristics. If there is any variation in the genotype-specific rates of return to the focal subset then the reversion rate will be strictly decreasing for both processes. While acclimatization of the rest of the genome to being in a new region of the fitness landscape contributes to decreasing reversion rate for both processes, for the conditioned process there is also a statistical effect because populations that spend more time at genotypes with rapid rate of return to the focal subset are likely to revert quickly. In addition, there is a simple mathematical relationship between these two processes: the asymptotic rate of transitions back to the focal subset when populations are restricted from entering it is the rate of an exponential distribution with mean equal to the expected reversion time, and this rate is bounded between the initial reversion rate and the asymptotic reversion rate when conditioning. Thus, while (Ashenberg et al. 2013) have criticized simulations where reversions are prevented from occurring as being unrealistic and

misleading, such simulations are in fact more relevant to understanding the evolutionary process than would appear at first glance.

Finally, because of non-linearity in the probability of fixation and genotype-to-genotype variation in mutation rates, the dynamics of the reversion rate and the dynamics of the expected selection coefficient of a reversion mutation may have qualitatively different shapes. Nonetheless, it is the substitution rate that is most closely related to the evolutionary dynamics. This is why we can derive strong results on the reversion rate, but not on the time-evolution of the average selection coefficient of reversions.

Acknowledgements

This work was funded by the Burroughs Wellcome Fund, the David and Lucile Packard Foundation, US Department of the Interior Grant D12AP00025, and Foundational Questions in Evolutionary Biology Fund Grant RFP-12-16, NIH training Grant 2T32AI055400-11 and US Army Research Office Grant W911NF-12-1-0552.

Literature Cited

- Aldous, D. and J. A. Fill, 2002. Reversible markov chains and random walks on graphs. Unfinished monograph, recompiled 2014, available at <http://www.stat.berkeley.edu/~aldous/RWG/book.html>.
- Ashenberg, O., L. I. Gong, and J. D. Bloom, 2013. Mutational effects on stability are largely conserved during protein evolution. *Proceedings of the National Academy of Sciences* 110:21071–21076.
- Bazykin, G. A., 2015. Changing preferences: deformation of single position amino acid fitness landscapes and evolution of proteins. *Biology letters* 11:20150315.
- Bergland, A. O., E. L. Behrman, K. R. O'Brien, P. S. Schmidt, and D. A. Petrov, 2014. Genomic evidence of rapid and stable adaptive oscillations over seasonal time scales in drosophila. *PLoS Genet* 10:e1004775. URL <http://dx.doi.org/10.1371/journal.pgen.1004775>.
- Breen, M. S., C. Kemena, P. K. Vlasov, C. Notredame, and F. A. Kondrashov, 2012. Epistasis as the primary factor in molecular evolution. *Nature* 490:535–538.
- , 2013. Breen et al. reply. *Nature* 497:E2–E3.
- Bridgham, J. T., E. A. Ortlund, and J. W. Thornton, 2009. An epistatic ratchet constrains the direction of glucocorticoid receptor evolution. *Nature* 461:515–519.
- Carneiro, M. and D. L. Hartl, 2010. Adaptive landscapes and protein evolution. *Proc. Natl. Acad. Sci. U. S. A.* 107:1747 – 1751. URL <http://www.pnas.org/content/107/suppl.1/1747.abstract>.
- Collin, R. and M. P. Miglietta, 2008. Reversing opinions on Dollo's law. *Trends in Ecology & Evolution* 23:602–609.
- Conrad, M., 1990. The geometry of evolution. *Biosystems* 24:61 – 81. URL <http://www.sciencedirect.com/science/article/pii/0303264790900305>.
- Dollo, L., 1893. The laws of evolution. *Bull. Soc. Bel. Geol. Paleontol* 7:164–166.
- Draghi, J. A. and J. B. Plotkin, 2013. Selection biases the prevalence and type of epistasis along adaptive trajectories. *Evolution* 67:3120–3131. URL <http://dx.doi.org/10.1111/evo.12192>.
- Fisher, R. A., 1930. *The Genetical Theory of Natural Selection*. Clarendon Press, London.
- Gavrilets, S., 1997. Evolution and speciation on holey adaptive landscapes. *Trends. Ecol. Evol.* 12:307 – 312.

- Goldstein, R. A., S. T. Pollard, S. D. Shah, and D. D. Pollock, 2015. Nonadaptive amino acid convergence rates decrease over time. *Molecular biology and evolution* 32:1373–1381.
- Gould, S. J., 1970. Dollo on Dollo's law: irreversibility and the status of evolutionary laws. *Journal of the History of Biology* 3:189–212.
- Grimmett, G. and D. Stirzaker, 2001. *Probability and random processes*. Oxford university press.
- Hartl, D. L., D. E. Dykhuizen, and A. M. Dean, 1985. Limits of adaptation: The evolution of selective neutrality. *Genetics* 111:655–674. URL <http://www.genetics.org/content/111/3/655.abstract>.
- Hartl, D. L. and C. H. Taubes, 1996. Compensatory nearly neutral mutations: selection without adaptation. *J. Theor. Biol.* 182:303–309. URL <http://dx.doi.org/10.1006/jtbi.1996.0168>.
- Hensley, S. E., S. R. Das, A. L. Bailey, L. M. Schmidt, H. D. Hickman, A. Jayaraman, K. Viswanathan, R. Raman, R. Sasisekharan, J. R. Bennink, et al., 2009. Hemagglutinin receptor binding avidity drives influenza a virus antigenic drift. *Science* 326:734–736.
- Iwasa, Y., 1988. Free fitness that always increases in evolution. *J. Theor. Biol.* 135:265 – 281. URL <http://www.sciencedirect.com/science/article/pii/S0022519388802431>.
- Jukes, T. H. and C. R. Cantor, 1969. Evolution of protein molecules. Pp. 21–132, in H. N. Munro, ed. *Mammalian Protein Metabolism*, vol. 3. New York.
- Kelly, F. P., 1979. *Reversibility and Stochastic Networks*. Series in Probability and Mathematical Statistics. Wiley, Chichester.
- Kerr, B., M. A. Riley, M. W. Feldman, and B. J. Bohannan, 2002. Local dispersal promotes biodiversity in a real-life game of rock–paper–scissors. *Nature* 418:171–174.
- Kielson, J., 1979. Markov Chain Models — Rarity and Exponentiality, *Applied Mathematical Sciences*, vol. 28. Springer-Verlag, New York.
- Kimura, M., 1983. *The neutral theory of molecular evolution*. Cambridge University Press, Cambridge.
- Kondrashov, A. S., I. S. Povolotskaya, D. N. Ivankov, and F. A. Kondrashov, 2010. Rate of sequence divergence under constant selection. *Biol. Direct* 5:5.
- Kondrashov, D. A. and F. A. Kondrashov, 2015. Topological features of rugged fitness landscapes in sequence space. *Trends in Genetics* 31:24–33.

- Kryazhimskiy, S., D. P. Rice, E. R. Jerison, and M. M. Desai, 2014. Global epistasis makes adaptation predictable despite sequence-level stochasticity. *Science* 344:1519–1522. URL <http://www.sciencemag.org/content/344/6191/1519.abstract>.
- Leslie, A., K. Pfafferott, P. Chetty, R. Draenert, M. Addo, M. Feeney, Y. Tang, E. Holmes, T. Allen, J. Prado, et al., 2004. HIV evolution: CTL escape mutation and reversion after transmission. *Nature medicine* 10:282–289.
- McCandlish, D. M., 2013. On the findability of genotypes. *Evolution* 67:2592 – 2603. URL <http://dx.doi.org/10.1111/evo.12128>.
- McCandlish, D. M., C. L. Epstein, and J. B. Plotkin, 2015a. Formal properties of the probability of fixation: identities, inequalities and approximations. *Theoretical population biology* 99:98–113.
- McCandlish, D. M., J. Otwinowski, and J. B. Plotkin, 2015b. Detecting epistasis from an ensemble of adapting populations. *Evolution* 69:2359–2370.
- McCandlish, D. M., E. Rajon, P. Shah, Y. Ding, and J. B. Plotkin, 2013. The role of epistasis in protein evolution. *Nature* 497:E1–E2.
- McCandlish, D. M. and A. Stoltzfus, 2014. Modeling evolution using the probability of fixation: History and implications. *The Quarterly Review of Biology* 89:pp. 225–252. URL <http://www.jstor.org/stable/10.1086/677571>.
- Muller, H. J., 1918. Genetic variability, twin hybrids and constant hybrids, in a case of balanced lethal factors. *Genetics* 3:422.
- , 1939. Reversibility in evolution considered from the standpoint of genetics. *Biological Reviews* 14:261–280.
- Naumenko, S. A., A. S. Kondrashov, and G. A. Bazykin, 2012. Fitness conferred by replaced amino acids declines with time. *Biology letters* 8:825–828.
- Pollock, D. D. and R. A. Goldstein, 2014. Strong evidence for protein epistasis, weak evidence against it. *Proceedings of the National Academy of Sciences* 111:E1450–E1450.
- Pollock, D. D., G. Thiltgen, and R. A. Goldstein, 2012. Amino acid coevolution induces an evolutionary Stokes shift. *Proc Natl Acad Sci USA* 109:E1352–9.
- Porter, M. L. and K. A. Crandall, 2003. Lost along the way: the significance of evolution in reverse. *Trends in Ecology & Evolution* 18:541–547.

- Risso, V. A., F. Manssour-Triedo, A. Delgado-Delgado, R. Arco, A. Barroso-delJesus, A. Ingles-Prieto, R. Godoy-Ruiz, J. A. Gavira, E. A. Gaucher, B. Ibarra-Molero, and J. M. Sanchez-Ruiz, 2015. Mutational studies on resurrected ancestral proteins reveal conservation of site-specific amino acid preferences throughout evolutionary history. *Molecular Biology and Evolution* 32:440–455. URL <http://mbe.oxfordjournals.org/content/32/2/440.abstract>.
- Rokas, A. and S. B. Carroll, 2008. Frequent and widespread parallel evolution of protein sequences. *Molecular biology and evolution* 25:1943–1953.
- Sella, G. and A. E. Hirsh, 2005. The application of statistical physics to evolutionary biology. *Proc. Natl. Acad. Sci. U. S. A.* 102:9541 – 9546. URL <http://www.pnas.org/content/102/27/9541.abstract>.
- Shah, P., D. M. McCandlish, and J. B. Plotkin, 2015. Contingency and entrenchment in protein evolution under purifying selection. *Proceedings of the National Academy of Sciences* 112:E3226–E3235. URL <http://www.pnas.org/content/112/25/E3226.abstract>.
- Simpson, G. G., 1953. *The Major Features Of Evolution*. Columbia University Press, New York.
- Tan, L., S. Serene, H. X. Chao, and J. Gore, 2011. Hidden randomness between fitness landscapes limits reverse evolution. *Physical review letters* 106:198102.
- Teotónio, H. and M. R. Rose, 2001. Perspective: reverse evolution. *Evolution* 55:653–660.
- Usmanova, D. R., L. Ferretti, I. S. Povolotskaya, P. K. Vlasov, and F. A. Kondrashov, 2015. A model of substitution trajectories in sequence space and long-term protein evolution. *Molecular biology and evolution* 32:542–554.
- Vaupel, J. W. and A. I. Yashin, 1985. Heterogeneity's ruses: some surprising effects of selection on population dynamics. *The American Statistician* 39:176–185.

Appendix 1 The distribution of reversion times at stationarity

Our analysis of the distribution of reversion times at stationarity relies on the fact that the matrix \mathbf{Q} , which gives the rates for the absorbing Markov chain describing the evolution dynamics until the time of reversion, admits an eigendecomposition. We will develop this eigendecomposition and its properties, and then use these results to show that the distribution of reversion times at stationarity can be expressed as a mixture of exponential distributions. We will then use a very similar analysis to understand the modified process where reversion events are not permitted to occur.

We derive the eigendecomposition of \mathbf{Q} by first noting some features of the larger matrix \mathbf{Q}_{full} . Because the Markov chain defined by \mathbf{Q}_{full} is reversible it satisfies the detailed balance relation $\pi_{\text{full}}(i)\mathbf{Q}_{\text{full}}(i, j) = \pi_{\text{full}}(j)\mathbf{Q}_{\text{full}}(j, i)$ for all i, j . As a consequence, the matrix $\mathbf{D}_{\pi_{\text{full}}}^{1/2}\mathbf{Q}_{\text{full}}\mathbf{D}_{\pi_{\text{full}}}^{-1/2}$ is symmetric, where $\mathbf{D}_{\mathbf{y}}$ denotes the diagonal matrix with the vector \mathbf{y} as its main diagonal.

Now, define the vector π of length $n = |A^c|$ such that $\pi(i) = \pi_{\text{full}}(i) / \sum_{j \in A^c} \pi_{\text{full}}(j)$ for $i = 1, \dots, n$. The matrix $\mathbf{D}_{\pi}^{1/2}\mathbf{Q}\mathbf{D}_{\pi}^{-1/2}$ is thus symmetric, since it is simply a constant times a diagonal block of the symmetric matrix $\mathbf{D}_{\pi_{\text{full}}}^{1/2}\mathbf{Q}_{\text{full}}\mathbf{D}_{\pi_{\text{full}}}^{-1/2}$. We can thus expand $\mathbf{D}_{\pi}^{1/2}\mathbf{Q}\mathbf{D}_{\pi}^{-1/2}$ in terms of its eigenvalues and eigenvectors as

$$-\mathbf{D}_{\pi}^{1/2}\mathbf{Q}\mathbf{D}_{\pi}^{-1/2} = \sum_{k=1}^n \lambda_k \mathbf{u}_k \mathbf{u}_k^{\text{T}}, \quad (\text{A1})$$

where $0 < \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \dots \leq \lambda_n$ are the eigenvalues of $-\mathbf{D}_{\pi}^{1/2}\mathbf{Q}\mathbf{D}_{\pi}^{-1/2}$ and the eigenvectors \mathbf{u}_k are orthonormal (the eigenvalues are real because $-\mathbf{D}_{\pi}^{1/2}\mathbf{Q}\mathbf{D}_{\pi}^{-1/2}$ is symmetric and negative because they are same as those of $-\mathbf{Q}$ where \mathbf{Q} is the generator of an absorbing Markov chain so that all of its eigenvalues have negative real parts).

We can now use this decomposition of $\mathbf{D}_{\pi}^{1/2}\mathbf{Q}\mathbf{D}_{\pi}^{-1/2}$ to likewise decompose \mathbf{Q} . In particular, multiplying Equation A1 by $\mathbf{D}_{\pi}^{-1/2}$ from the left and $\mathbf{D}_{\pi}^{1/2}$ from the right gives us:

$$-\mathbf{Q} = \sum_{k=1}^n \lambda_k \mathbf{r}_k \mathbf{l}_k^{\text{T}}, \quad (\text{A2})$$

where $\mathbf{l}_k = \mathbf{D}_{\pi}^{1/2}\mathbf{u}_k$ and $\mathbf{r}_k = \mathbf{D}_{\pi}^{-1/2}\mathbf{u}_k$ are the left and right eigenvectors of $-\mathbf{Q}$ associated with λ_k .

Using this decomposition, we can then write the probability density function of the distribution of reversion times at stationarity as

$$f(t) = \mathbf{x}_t^\top \mathbf{g} \quad (\text{A3})$$

$$= \mathbf{x}_0^\top e^{\mathbf{Q}t} \mathbf{g} \quad (\text{A4})$$

$$= \mathbf{x}_0^\top \left(\sum_{k=1}^n e^{-\lambda_k t} \mathbf{r}_k \mathbf{l}_k^\top \right) \mathbf{g} \quad (\text{A5})$$

$$= \frac{1}{\boldsymbol{\pi}^\top \mathbf{g}} \mathbf{g}^\top \mathbf{D} \boldsymbol{\pi} \left(\sum_{k=1}^n e^{-\lambda_k t} \mathbf{r}_k \mathbf{l}_k^\top \right) \mathbf{g} \quad (\text{A6})$$

$$= \frac{1}{\boldsymbol{\pi}^\top \mathbf{g}} \sum_{k=1}^n e^{-\lambda_k t} (\mathbf{l}_k^\top \mathbf{g})^2, \quad (\text{A7})$$

where we have used the fact that at stationarity $\mathbf{x}_0(i) \propto \boldsymbol{\pi}(i)\mathbf{g}(i)$ (Equation 7). Thus, $f(t)$ is a mixture of exponential densities with rates $\lambda_1, \dots, \lambda_n$ where the exponential density with rate λ_k has weight $(\mathbf{l}_k^\top \mathbf{g})^2 / (\lambda_k \boldsymbol{\pi}^\top \mathbf{g})$.

We just showed that the distribution of reversion times at stationarity is a finite mixture of exponential distributions. We will now show that the hazard function for any finite mixture of exponential distributions is non-increasing. In particular, if $f(t)$ is a finite mixture of exponential densities

$$f(t) = \sum_m a_m \theta_m e^{-\theta_m t} \quad (\text{A8})$$

with $a_m, \theta_m > 0$, $\sum_m a_m = 1$, then the hazard function is

$$\frac{f(t)}{\bar{F}(t)} = \frac{\sum_m a_m \theta_m e^{-\theta_m t}}{\sum_m a_m e^{-\theta_m t}}. \quad (\text{A9})$$

Differentiating the hazard function, we see that that the sign of the derivative depends only on the sign of

$$\bar{F}(t)f'(t) - \bar{F}'(t)f(t) = \bar{F}(t)f'(t) + (f(t))^2 \quad (\text{A10})$$

$$= - \sum_{i,j} a_i a_j \theta_j^2 e^{-(\theta_i + \theta_j)t} + \sum_{i,j} a_i a_j \theta_i \theta_j e^{-(\theta_i + \theta_j)t} \quad (\text{A11})$$

$$= - \left(\sum_{i,j \text{ with } i \neq j} a_i a_j e^{-(\theta_i + \theta_j)t} (\theta_j^2 - \theta_i \theta_j) \right) \quad (\text{A12})$$

$$= - \left(\sum_{i,j \text{ with } i > j} a_i a_j e^{-(\theta_i + \theta_j)t} (\theta_j^2 - 2\theta_i \theta_j + \theta_i^2) \right) \quad (\text{A13})$$

$$= - \left(\sum_{i,j \text{ with } i > j} a_i a_j e^{-(\theta_i + \theta_j)t} (\theta_i - \theta_j)^2 \right), \quad (\text{A14})$$

which is non-positive since each term in the sum is a product of non-negative quantities and hence non-negative. Thus, the hazard function is non-increasing.

We turn now to the analysis of the modified process with rate matrix $\mathbf{Q}^* = \mathbf{Q} + \mathbf{D}_g$, where \mathbf{D}_g is the diagonal matrix with \mathbf{g} on its main diagonal. Let us start by assuming that A^c is mutationally connected; we will return to the case of disconnected A^c momentarily. First note that \mathbf{Q}^* only differs from \mathbf{Q} on its diagonal entries. Thus, following our analysis for \mathbf{Q} , $\mathbf{D}_\pi^{1/2} \mathbf{Q}^* \mathbf{D}_\pi^{-1/2}$ is symmetric and can be expanded in terms of its eigenvalues and eigenvectors as

$$-\mathbf{D}_\pi^{1/2} \mathbf{Q}^* \mathbf{D}_\pi^{-1/2} = \sum_{k=1}^n \lambda_k^* \mathbf{u}_k^* \mathbf{u}_k^{*\top}, \quad (\text{A15})$$

where $0 = \lambda_1^* < \lambda_2^* \leq \lambda_3^* \leq \dots \leq \lambda_n^*$ are the eigenvalues of $-\mathbf{D}_\pi^{1/2} \mathbf{Q}^* \mathbf{D}_\pi^{-1/2}$ and the eigenvectors \mathbf{u}_k^* are orthonormal (we have $0 = \lambda_1^* < \lambda_2^*$ because A^c is connected and so the continuous time Markov chain generated by \mathbf{Q}^* is ergodic). With this decomposition in hand, we can write $\mathbf{l}_k^* = \mathbf{D}_\pi^{1/2} \mathbf{u}_k^*$ and $\mathbf{r}_k^* = \mathbf{D}_\pi^{-1/2} \mathbf{u}_k^*$ as the left and right eigenvectors of $-\mathbf{Q}^*$ associated with λ_k^* and write the reversion rate under the modified process as:

$$f^*(t) = \mathbf{x}_0^\top e^{\mathbf{Q}^* t} \mathbf{g} \quad (\text{A16})$$

$$= \mathbf{x}_0^\top \left(\sum_{k=1}^n e^{-\lambda_k^* t} \mathbf{r}_k^* \mathbf{l}_k^{*\top} \right) \mathbf{g} \quad (\text{A17})$$

$$= \frac{1}{\boldsymbol{\pi}^\top \mathbf{g}} \mathbf{g}^\top \mathbf{D}_\pi \left(\sum_{k=1}^n e^{-\lambda_k^* t} \mathbf{r}_k^* \mathbf{l}_k^{*\top} \right) \mathbf{g} \quad (\text{A18})$$

$$= \boldsymbol{\pi}^\top \mathbf{g} + \frac{1}{\boldsymbol{\pi}^\top \mathbf{g}} \sum_{k=2}^n e^{-\lambda_k^* t} (\mathbf{l}_k^{*\top} \mathbf{g})^2, \quad (\text{A19})$$

where we've used the fact that the row sums of \mathbf{Q}^* are all 0 so that \mathbf{r}_1^* is the vector of all 1s and thus $\mathbf{l}_1^* = \mathbf{D}_\pi \mathbf{r}_1^* = \boldsymbol{\pi}$. Because $\lambda_k^* > 0$ for $k \geq 2$, this expression is clearly strictly decreasing in t unless $n = 1$ (in which case $\mathbf{g}(i)$ is obviously constant) or $\mathbf{l}_k^{*\top} \mathbf{g} = 0$ for $k \geq 2$. In the latter case, we have $\mathbf{u}_k^{*\top} \mathbf{D}_\pi^{1/2} \mathbf{g}$ for all $k \geq 2$, so that $\mathbf{D}_\pi^{1/2} \mathbf{g}$ must be orthogonal to all the \mathbf{u}_k^* for $k \geq 2$. Since the \mathbf{u}_k^* are an orthonormal basis of \mathbb{R}^n , the only direction remaining for $\mathbf{D}_\pi^{1/2} \mathbf{g}$ is \mathbf{u}_1^* , so that \mathbf{g} must be in the direction of \mathbf{r}_1^* , which is the vector of all 1s. Thus, in that case \mathbf{g} is also constant.

So far, we've shown that provided that A^c is mutationally connected, at stationarity the reversion rate under \mathbf{Q}^* is constant if $\mathbf{g}(i)$ is constant and otherwise it is strictly decreasing. Now we consider the case where A^c has several disconnected components, A_1, \dots, A_l . Clearly the total reversion rate is a weighted average of the reversion rates for each disconnected component when analyzed separately as above. If $\mathbf{g}(i)$ is non-constant within any one of these components, then that term in the average is strictly decreasing in time and hence the reversion rate as a

whole is strictly decrease. Thus, for general A^c the reversion rate at stationarity under Q^* is constant if $g(i)$ is piece-wise constant in the sense of being constant for each of the disconnected components of A^c , otherwise it is strictly decreasing.

Appendix 2 Role of epistasis in multiallelic fitness landscapes

In the main text, we showed that for biallelic fitness landscapes, the reversion rate at stationarity for a site is a decreasing function of time if and only if that site is involved in an epistatic interaction with another site. Here we show that for multi-allelic fitness landscapes the presence of an epistatic interaction with another site is a sufficient condition under stationarity for reversion rates at the focal site to be decreasing.

First we must establish some notation for multi-allelic fitness landscapes. If we have a fitness landscape over L sites, we can label the $n_l \geq 2$ alleles at the l -th site, $\alpha_{l,1}, \dots, \alpha_{l,n_l}$ and write the mutation rate between $\alpha_{l,k}$ and $\alpha_{l,m}$ as $\mu_{\alpha_{l,k} \rightarrow \alpha_{l,m}}$. Thus, if genotypes i and j differ at more than one site we have $M_{\text{full}}(i, j) = 0$ otherwise they differ at exactly one site l in which case $M_{\text{full}}(i, j) = \mu_{\alpha_{l,k} \rightarrow \alpha_{l,m}}$ where genotype i has the $\alpha_{l,k}$ allele at site l and genotype j has the $\alpha_{l,m}$ allele at site l .

Without loss of generality we can define the focal set A to be the set of genotypes with allele $\alpha_{1,1}$ in site 1. We want to show that if site 1 has an epistatic interaction with at least one other site, then the $g(i)$ are not constant, since this is sufficient to show that the reversion rate is decreasing at stationarity.

Now, either $\mu_{\alpha_{1,k} \rightarrow \alpha_{1,1}} > 0$ for all k or not. If not, then $\mu_{\alpha_{1,k} \rightarrow \alpha_{1,1}} > 0$ for some k , since otherwise the Markov chain defined by Q_{full} would not be reversible (populations could never return from A^c to A). Now, if $\mu_{\alpha_{1,k} \rightarrow \alpha_{1,1}} > 0$ for some k but $\mu_{\alpha_{1,m} \rightarrow \alpha_{1,1}} = 0$ for some m , then g is positive for genotypes with allele $\alpha_{1,k}$ at site l , but 0 for genotypes with allele $\alpha_{1,m}$ at site l . Thus, g is non-constant.

We thus turn to the case where $\mu_{\alpha_{1,k} \rightarrow \alpha_{1,1}} > 0$ for all k . Our strategy is to show that if all alleles at a site can mutate to all the other alleles, then the presence of any epistatic interaction between one site and another is sufficient, by the transitivity of fitness differences, to guarantee the existence of a set of four genotypes to which we can apply the argument in the main text for biallelic sites. Now, if site 1 is epistatic with some other site, then without loss of generality we can let this other site be site 2. Because site 1 is epistatic with site 2 there exist, by the definition of epistasis, genotypes i, i', j, j' and alleles $\alpha_{1,k}, \alpha_{1,m}, \alpha_{2,k'}, \alpha_{2,m'}$ such that:

1. Genotype j and i' are both single mutants with respect to genotype i , and j' is the corresponding double mutant. In particular, genotype j is formed from genotype i by a

single $\alpha_{1,k} \rightarrow \alpha_{1,m}$ mutation, genotype i' is formed from i as a single $\alpha_{2,k'} \rightarrow \alpha_{2,m'}$, and the double mutant genotype j' is formed from genotype i by the combination of an $\alpha_{1,k} \rightarrow \alpha_{1,m}$ mutation and an $\alpha_{2,k'} \rightarrow \alpha_{2,m'}$ mutation.

$$2. \mathbf{f}(i) - \mathbf{f}(j) \neq \mathbf{f}(i') - \mathbf{f}(j').$$

If either $k = 1$ or $m = 1$ then \mathbf{g} is non-constant by the same argument as given for the two allele case in the main text. Therefore suppose $k \neq 1$ and $m \neq 1$. In that case we can find genotypes h and h' such that h is identical to i and j except that it has the $\alpha_{1,1}$ allele at site 1 and likewise h' is identical to i' and j' except that it has the $\alpha_{1,1}$ allele at site 1. Now, if it were the case that both $\mathbf{f}(h) - \mathbf{f}(j) = \mathbf{f}(h') - \mathbf{f}(j')$ and $\mathbf{f}(i) - \mathbf{f}(h) = \mathbf{f}(i') - \mathbf{f}(h')$, then we would have

$$\begin{aligned} \mathbf{f}(i) - \mathbf{f}(j) &= (\mathbf{f}(i) - \mathbf{f}(h)) + (\mathbf{f}(h) - \mathbf{f}(j)) \\ &= (\mathbf{f}(i') - \mathbf{f}(h')) + (\mathbf{f}(h') - \mathbf{f}(j')) \\ &= \mathbf{f}(i') - \mathbf{f}(j'). \end{aligned}$$

But $\mathbf{f}(i) - \mathbf{f}(j) \neq \mathbf{f}(i') - \mathbf{f}(j')$, so either $\mathbf{f}(h) - \mathbf{f}(j) \neq \mathbf{f}(h') - \mathbf{f}(j')$ or $\mathbf{f}(i) - \mathbf{f}(h) \neq \mathbf{f}(i') - \mathbf{f}(h')$. Thus, we could have chosen h and h' instead of either i and i' or j and j' above, in which case the non-constancy of \mathbf{g} again follows by the same argument given for the two allele case in the main text.

Appendix 3 Expected time until reversion

We want to derive the expected time until reversion to a set A at stationarity given that the population has just left A . To study this time, it is helpful to define a modified Markov chain on a state space $A^c \cup \{a\}$ where the new state a takes the place of the set A in the original chain.

In particular, consider the chain with rate matrix

$$\mathbf{Q}_{\text{collapsed}} = \left(\begin{array}{c|c} \mathbf{Q}_{A^c} & \mathbf{g} \\ \hline k \mathbf{x}_0^T & -k \end{array} \right), \quad (\text{A20})$$

where k is the expected substitution rate from A to A^c at stationarity for \mathbf{Q}_{full} conditional on the population being in A , and $\mathbf{x}_0(i)$ is the probability of the population being at genotype $i \in A^c$ under \mathbf{Q}_{full} at stationarity, conditional on a substitution from A to A^c having just occurred. By construction, this new chain with rate matrix $\mathbf{Q}_{\text{collapsed}}$ has the same distribution of reversion times to a at stationarity as did the original chain.

Furthermore, if A^c contains n elements, it is easy to verify that the stationary distribution of $\mathbf{Q}_{\text{collapsed}}$ is given by the vector $(\boldsymbol{\pi}_{\text{full}}(1), \dots, \boldsymbol{\pi}_{\text{full}}(n), \boldsymbol{\pi}_{\text{full}}(A))$, where $\boldsymbol{\pi}_{\text{full}}(A) = \sum_{i \in A} \boldsymbol{\pi}_{\text{full}}(i)$,

since the new chain conserves both the stationary rate of substitutions from A to i and from i to A for all states $i \in A^c$.

Analyzing this new chain, a straightforward application of the reward-renewal theorem for the renewal process defined by returns to a and reward equal to the waiting time before leaving a after each return gives us:

$$\pi_{\text{full}}(A) = \frac{1}{\tau_a k} \quad (\text{A21})$$

(Grimmett and Stirzaker 2001, Theorem 10.5.10), where τ_a is the expected waiting time for a population currently at a to leave a and then return to a for the first time.

Now, a population leaving a and then returning to a must first leave a and then revert to a . The expected time before the population leaves a is $1/k$ and then the expected waiting time to revert to a is $m(0)$, the quantity that we are trying to develop an expression for. Thus, we have

$$\tau_a = \frac{1}{k} + m(0) \quad (\text{A22})$$

Plugging Equation A22 into Equation A21 and solving for $m(0)$ gives us:

$$m(0) = \frac{1 - \pi_{\text{full}}(A)}{k \pi_{\text{full}}(A)}. \quad (\text{A23})$$

At stationarity, the total rate that populations arrive at a must be equal to the rate that they leave a , so we have $k \pi_{\text{full}}(A) = \sum_{i \in A^c} \pi_{\text{full}}(i) \mathbf{g}(i)$, and of course $1 - \pi_{\text{full}}(A) = \sum_{i \in A^c} \pi_{\text{full}}(i)$. Substituting these expressions into Equation A23 yields

$$m(0) = \left(\frac{\sum_{i \in A^c} \pi_{\text{full}}(i) \mathbf{g}(i)}{\sum_{i \in A^c} \pi_{\text{full}}(i)} \right)^{-1}, \quad (\text{A24})$$

i.e. the expected waiting time reversion is equal to the reciprocal of the expected substitution rate to A at stationarity conditional on the population being in A^c , as required.