Empirical Complexities in the Genetic Foundations of Lethal Mutagenesis

James J. Bull,*,†,‡,1 Paul Joyce,§, Eric Gladstone,* and Ian J. Molineux†,††**

*Institute for Cellular and Molecular Biology, †Center for Computational Biology and Bioinformatics, ‡Department of Integrative Biology, and ††Department of Molecular Biosciences, University of Texas, Austin, Texas 78712, and §Department of Mathematics and **Department of Statistics, University of Idaho, Moscow, Idaho 83844

ABSTRACT From population genetics theory, elevating the mutation rate of a large population should progressively reduce average fitness. If the fitness decline is large enough, the population will go extinct in a process known as lethal mutagenesis. Lethal mutagenesis has been endorsed in the virology literature as a promising approach to viral treatment, and several *in vitro* studies have forced viral extinction with high doses of mutagenic drugs. Yet only one empirical study has tested the genetic models underlying lethal mutagenesis, and the theory failed on even a qualitative level. Here we provide a new level of analysis of lethal mutagenesis by developing and evaluating models specifically tailored to empirical systems that may be used to test the theory. We first quantify a bias in the estimation of a critical parameter and consider whether that bias underlies the previously observed lack of concordance between theory and experiment. We then consider a seemingly ideal protocol that avoids this bias—mutagenesis of virions—but find that it is hampered by other problems. Finally, results that reveal difficulties in the mere interpretation of mutations assayed from double-strand genomes are derived. Our analyses expose unanticipated complexities in testing the theory. Nevertheless, the previous failure of the theory to predict experimental outcomes appears to reside in evolutionary mechanisms neglected by the theory (*e.g.*, beneficial mutations) rather than from a mismatch between the empirical setup and model assumptions. This interpretation raises the specter that naive attempts at lethal mutagenesis may augment adaptation rather than retard it.

LETHAL mutagenesis is extinction of a viral population by artificially elevating its mutation rate. The concept was originally inspired by Eigen and Schuster's "error catastrophe" theory that has commonly been construed to represent extinction at high mutation rate, and then by evidence that an antiviral drug known as ribavirin acts in part by elevating viral mutation rates (Eigen 1971; Eigen and Schuster 1977; Eigen *et al.* 1988; Crotty *et al.* 2001; Graci and Cameron 2002). In the past decade, a variety of studies have also shown that mutagenic drugs can extinguish viral populations in tissue culture (Crotty *et al.* 2001; Anderson *et al.* 2004; Chung *et al.* 2007; Martin *et al.* 2008).

It is obvious that a sufficiently high rate of lethal mutations will extinguish a population. Less obvious is the minimum mutation rate that will cause extinction over time and what deleterious fitness effects are needed for extinction. One model proposed for lethal mutagenesis of asexual organisms is based on the principle that a large population subjected only to deleterious mutations every generation will eventually decay to a fitness equilibrium of

$$\overline{W}_\infty = W_0\ \mu(0), \tag{1a}$$

where W_0 is the initial mean fitness of the population prior to mutagenesis and $\mu(0)$ is the probability that an offspring escapes deleterious mutation (Kimura and Maruyama 1966). Under a Poisson model of mutation with an average of U deleterious mutations per generation, this equation becomes

$$\overline{W}_\infty = W_0\ \mathrm{e}^{-U}. \tag{1b}$$

Extinction ensues if the baseline fecundity W_0 is not large enough to offset the fitness decline from mutation, but conversely, a sufficiently high fecundity can offset any negative impact of mutations and avoid extinction (Bull *et al.* 2007). The Poisson model lends itself to many useful extensions

doi: 10.1534/genetics.113.154195
Manuscript received June 10, 2013; accepted for publication July 29, 2013
Available freely online through the author-supported open access option.
[1]Corresponding author: Department of Integrative Biology, University of Texas, 1 University Station Stop C0990, Austin, TX 78712. E-mail: bull@utexas.edu

and will be assumed in much of what follows. However, a Poisson process need not operate, and (1a) gives the general result.

A limitation of both models is that beneficial mutations are omitted; beneficial mutations are no doubt inescapable in most systems and obviously raise fitness above that predicted by deleterious mutations alone. Recent work has begun to address mutation processes that include deleterious and beneficial mutations (Gerrish *et al.* 2013). A second and subtle point is that the equilibria in (1a) and (1b) hold only up to certain types of error catastrophe (Eigen and Schuster 1977; Eigen *et al.* 1988; Wiehe 1997; Bull *et al.* 2007). The standard model of an error catastrophe is loss of the mutation-free genotype in favor of a mutationally robust genotype or network, a robustness that buffers the population against extinction at higher mutation rates (Wilke *et al.* 2001; Wilke and Adami 2003). However, the full set of processes considered as error catastrophes is diverse and not easily connected to biology (Wiehe 1997), so the results in (1) cannot be generalized to account for error catastrophes.

These equilibria have four noteworthy properties. First, the equilibria apply to large populations; stochastic effects do not enter. Second, effect size and epistasis of the mutations do not matter to the equilibrium. However, equilibrium applies in the long term; effect sizes and epistasis do influence the rate of approach to equilibrium. Third, and true only for (1b), the mutation rate enters as an exponent; thus small changes in the mutation rate have large effects on fitness. Last, genome size does not enter directly, only the deleterious mutation rate per genome.

The single empirical test of this model, applied to the virus T7, failed to observe any decline in fitness after 200 generations of continuous mutagenic growth in large experimental populations, even though the predicted decline based on this equilibrium was large (Springman *et al.* 2010). To the extent that those results reflect an underlying misunderstanding of the fitness consequences of mutation, this failure suggests that attempts to induce lethal mutagenesis may be misguided and have the unintended consequence of providing raw material for adaptation. The failure also inspires attempts to understand the basis of the discrepancy between predicted and observed fitness.

A central question motivated by that empirical test is whether the model failure resides at a fundamental level, such as failing to capture fundamental genetic and evolutionary principles or instead resides at a technical level whereby the specific implementation of the test is a poor match to the assumptions of the model. The purpose here is to develop the theory of lethal mutagenesis as it applies to the types of empirical systems and assays that may be easily employed with viruses, with the specific goal of resolving the question.

The T7 Mutagenic Protocol

In prior work, T7 was grown in cultures of hosts infused with the mutagen *N*-methyl-*N*′-nitro-N-nitrosoguanidine (Springman *et al.* 2010). Hosts were replaced every few generations, but phage populations were propagated sequentially. The major direct effect of this mutagen on the phage is manifested intracellularly, during viral replication, as an increase in mutation rate. The measured number of mutations per genome per infection was four to six, depending on protocol, the lower value being the approximate number of viable mutations. However, mutations accumulate continually, and their cumulative effect is expected to ultimately decrease fitness toward the equilibrium, hence the need for serial propagation across hundreds of generations.

Both the bacteria and phage are adversely affected by mutagen, however, and the bacteria are greatly impaired in their ability to support phage growth. Thus there is a strong physiological effect of the mutagen on phage fitness that is not manifested through mutations, and this physiological effect must be partitioned into W_0 to parameterize the model. Phage fitness was measured as growth rate in the mutagenic environment, a $\log_2$ transform of the per-capita increase in phage numbers per unit time.

Structure of the Article

The question motivating this study is the previously observed failure of fitness decline under mutagenesis in phage T7. This article considers possible reasons for this failure and an alternative empirical approach to testing the model, using the following organization:

1. The first part of the article addresses the possibility that the methodology for estimating baseline fitness is flawed; new theoretical results are proposed, the original T7 study is reconsidered, and a new set of empirical results is offered.
2. Second, the article considers an alternative empirical system that seems ideal for testing the theory: mutagenic treatment of free virus particles. Prior empirical work on this type of system is reviewed, and new theory is developed to address an apparent empirical anomaly. Data using the mutagen hydroxyamine on T7 are presented and used to infer limitations of that system.
3. Third, a complexity is noted in the conversion between measured mutation rate and the fitness impact of those mutations.

Results

Estimating fitness in the mutagenic environment: Deleterious mutations confound the baseline estimate

The fitnesses in Equations 1a and 1b apply to the environment in which the virus is grown while being subjected to mutation. Thus, it is necessary to measure both the baseline fitness, W_0, and the final fitness, $\overline{W}_\infty$, in that environment. With T7 and virtually all viruses subjected to mutagenic drugs clinically, that environment is one in which the cells are exposed to mutagen. The problem this presents when

parameterizing the model is that measuring W_0 in the mutagenic environment introduces mutations, during the assay, that confound a precise estimate of W_0.

The fitness consequences of viral growth in the mutagenic environment may be partitioned as follows:

1. Mutagen impairs bacterial physiology and growth, reducing the number of phage progeny produced by the cell; although fewer progeny are produced, the progeny produced are not affected by this mechanism.
2. Progeny viruses from hosts grown in mutagen potentially have structural defects in their virions; these effects are not heritable but they do affect the ability of those progeny to initiate infections.
3. Replication of viruses in hosts grown in mutagen introduces nonlethal mutations in progeny viral genomes; these effects are heritable, affecting fitness across future generations.
4. Viral replication in hosts grown in mutagen introduces lethal mutations in some progeny; progeny with lethal mutations leave no descendants and cannot be counted.

The difficulty is that the nonheritable effects are confounded with heritable effects when measuring (initial) fitness, but the model requires that only the nonheritable effects be included. Suppose the burst size (fitness in this example) of the wild-type genome in the absence of mutagen is 100. The fitness assay conducted in the presence of mutagen reveals a (viable) burst of 30 for the wild type, giving the impression that W_0 is 30. However, suppose the host has reduced ability to support phage growth—effects 1 and 2 in the list above—and that the true W_0 is 50; lethal mutations kill all but 60% of those progeny. The value of W_0 is thus 50, not 100 nor 30. Yet there is no obvious way to disentangle the nonheritable effects of mutagen on the host from the effect of mutagen in killing 40% of the actual progeny. Ignoring the problem and treating the observed 30 as W_0 leads to predicting too low of a final fitness; hence the model would be rejected even if it was correct. In the case of T7, the initial fitness W_0 was based on an assay spanning six generations, so the effect is possibly substantial. How serious, then, is the expected impact of the early mutations?

The decline in fitness with continual mutagenesis

The fitness impact of mutagenesis is easily calculated under the Poisson model. Johnson (1999) provided a set of recursions by which the impact of progressive mutation may be calculated iteratively, although the formula we need remains to be derived. Let all mutations be deleterious. Within the set of deleterious mutations there are n classes according to their effect size: those with deleterious effect s_i arise at genomic rate U_i. The combined deleterious rate across all mutations is $\sum_i U_i = U$. We further assume the convenient property that a genome's fitness declines multiplicatively with its set of mutations. Mean fitness after k generations of mutation and selection is

$$\overline{W}_k = W_0\, \mathrm{e}^{-U+\sum_{i=1}^{n}\left[U_i(1-s_i)^k\right]} \qquad (2)$$

(*Appendix*). Mean fitness in the first two generations of exposure and selection follows $\overline{W}_1 = \mathrm{e}^{-\bar{s}U}$ and $\overline{W}_2 = e^{-U+\sum_{i=1}^{n}U_i(1-s_i)^2} = e^{U(-2\bar{s}+\bar{s}^2+\mathrm{Var}(s))}$ (for $\bar{s} = \sum_i^n U_i s_i/U$). $\overline{W}_{k+1}$ is always smaller than $\overline{W}_k$ except in the extreme case that all mutations are either neutral or lethal. With appropriate limits on the mean and higher-order moments of s, mean fitness declines approximately as the power k early in the process.

The mutation load experienced in the kth generation of mutagenesis is $W_0 - \overline{W}_k$. The residual load, *i.e.*, that experienced beyond generation k, is thus $\overline{W}_k - \overline{W}_\infty$. This residual load is what is measurable in protocols that assay initial fitness during exposure to mutagen, for assays lasting k generations. As $\bar{s} \to 1$, the residual load vanishes even at low k because the full effect of mutation is experienced after few generations of mutagenesis.

From (2), there are important consequences of a mutagenesis protocol using fitness measured in the mutagenic environment. Overall, the observed decline in fitness between the initial and final estimates will be less than the true negative impact of mutations on fitness, as given by (1b). The observational bias is worse still if the "final" population has not attained equilibrium. This bias arises because the observed initial fitness is confounded by deleterious mutations accumulating during the assay, and the bias increases the larger the mutational effects and the longer the assay continues. Failure to account for this bias—by assuming that the initial fitness estimate represents W_0 instead of $\overline{W}_k$—will overpredict the expected fitness decline from a high mutation rate. The theory would then be tested inappropriately.

A few comments about model (2) are warranted. First, the model assumes an infinite number of sites in each class i. Thus a genome can theoretically experience and accumulate infinitely many mutations in each class. Second, the model admits beneficial mutations ($s_j < 0$), but with even a single class of beneficial mutations, fitness increases without bound as k increases to infinity ($U_j(1 - s_j)^k \to \infty$). The unbounded fitness is due to the infinite sites assumption, which allows genomes to accumulate ever-increasing numbers of beneficial mutations (as favored by selection). Consequently, use of (2) with beneficial mutations should be limited to the initial generations of mutagenesis. The long-term equilibrium may be calculated for a finite number of beneficial mutations by assuming that initial genotype has already acquired the beneficial mutations (with W_0 adjusted accordingly).

The average effect of a random deleterious mutation (including lethals) has been estimated as $\sim$ 0.68 in VSV (vesicular stomatitis virus, calculated from Table 1 in Sanjuán *et al*. 2004) and as 0.40 3 and 0.45 in two phages (Domingo-Calap *et al*. 2009). With these values, approximately half the long-term deleterious fitness effect of a high mutation rate is

experienced in the first generation and will not be detectable by many methods that assay fitness in the mutagenic environment. The average fitness effect of all mutations is −0.3 and −0.36 for the two phages, and even more negative for VSV. By this measure as well, there should be a substantial fitness decline in the first few generations.

T7 revisited

In the study of bacteriophage T7 mutagenesis, initial growth rate of the wild-type virus in the mutagenic environment was 18.3 doublings/hr (Springman *et al.* 2010). From a directly measured 4.3 viable mutations per genome per generation, the viable deleterious mutation rate was converted to 2.6 by assuming a deleterious fraction of 0.6 (based on direct estimates from other viruses). A model similar to that in Equation 1b but tailored to viral growth rate gave a predicted equilibrium of 8.6 doublings/hr. The observed fitness after 200 generations of propagation was 21.9, higher than fitness of the initial virus and a profound failure of the theory. Per 20 min viral generation, the difference between 21.9 and 8.6 is equivalent to a 20-fold difference in number of offspring.

Several possibilities were addressed in attempting to account for those results (Springman *et al.* 2010):

1. The phage may have evolved a lower mutation rate during the adaptation. T7 encodes its own DNA polymerase (DNAP) for genome replication, and resistance to mutagen could easily have evolved in that gene. However, a replicate evolution in which the phage DNAP was prevented from evolving exhibited a similar lack of fitness decline.
2. Deleterious fitness effects may be too small to expect a fitness drop in 200 generations. This possibility is inconsistent with the fitness effects of mutations reported in other viruses.
3. With 200 generations of evolution, beneficial mutations may have been able to ascend and offset the decline. The trajectory of fitness evolution supported this possibility, but with only a small magnitude of effect: fitness at 15 generations was not significantly below that of the estimated W_0, but fitness at 80 generations showed a significant but small decline of about one doubling per hour below that of W_0. Fitness at generation 200 thus implied a rise from generation 80 of about four doublings per hour. The question was why fitness exhibited such a modest decline at the earlier time points.

The question raised by the preceding analysis is whether the predicted impact of mutations was affected by a biased estimate of W_0. It is of course not possible to explain a fitness *increase* from this bias, but an upward correction of the predicted final fitness certainly helps reduce the discrepancy. The solution suggested by (2) is to correct for the early accumulation of deleterious mutations. However, if the distribution of deleterious fitness effects is not known, the correction is not straightforward.

As one approach, Springman *et al.* (2010) developed a model that partitioned lethals separately from other deleterious mutations and then parameterized the model so that the predicted fitness drop was restricted to the nonlethal component. This correction mirrors the one suggested above—the fitness impact of lethals on mean fitness is the same from the first generation onward so is independent of the duration of the fitness assay. In contrast, the fitness impact of nonlethals accumulates gradually and is thus possibly detectable when comparing assays of different durations.

Partitioning the deleterious mutation rate U into a lethal rate (U_L) and a nonlethal rate (U_{nL}), Equation 1b may be rewritten as $\overline{W}_\infty = (W_0\, e^{-U_L}) \cdot e^{-U_{nL}}$. The product in parentheses is the initial viable fitness. This value corresponds to the 30 in our example three sections above and is easily measurable. Using this adjusted fitness as initial fitness, the predicted decline from other deleterious mutations is now the component due to nonlethals, $e^{-U_{nL}}$. This decline should accumulate gradually and thus be detectable with the methods used in the study, to the extent that the average effect of nonlethal, deleterious mutations is small. Of course, this drop does not represent the full negative impact of deleterious mutations, but understanding its magnitude could facilitate resolving the enigma of T7 evolving higher fitness under mutagenesis.

As the standard fitness assay spans several generations and thus might have partitioned some of the nonlethal component into $W_0 e^{-U_L}$, the T7 study estimated initial fitness in two ways. One method directly measured fitness (growth rate) from viral titers compared between three and six generations, so that estimate would indeed have been depressed below its true value by the early accumulation of nonlethal mutations. The other estimate, which avoids this problem, parameterized a fitness function using fitness components obtained during the first generation of mutagen exposure. This latter method should avoid the effects of deleterious mutations, save for lethal mutations (which were already accounted for by using the viable burst size).

The two fitness estimation methods gave broadly similar values, albeit variance in the fitness function estimate was not quantified. Consequently, there was no obvious bias attributable to nonlethals accumulated during the first six generations of mutagenic growth. For both fitness estimates, the long-term predicted fitness decline was the component based on the nonlethal mutation rate, and a directly measured nonlethal mutation rate was used to parameterize that component. The predicted decline was substantial and profoundly at odds with the fitness increase observed at the end of the experiment. Thus it does not seem that the failure of the theory is due to a biased initial fitness estimate.

A surrogate fitness measure in T7

To test the model appropriately, the assays used to estimate fitness must represent fitness as it applies during the

mutagenic evolution; thus the most appropriate fitness assay environment would use mutagenic growth. For short-term evolution, however, fitness measured in a nonmutagenic environment may correspond to fitness in the mutagenic environment so that alternative fitness assays are acceptable. Specifically, if most mutations deleterious in the mutagenic environment are also deleterious in a surrogate environment, the fitness decay over a few generations of mutagenesis may parallel each other in the two environments. A fitness decline may then be more easily documented in a nonmutagenic environment and be free of the estimation biases noted above. In the long term, adaptive evolution specific to the mutagenic environment may destroy any correspondence in fitness between the two environments, but this problem will not arise in the very short term.

We thus attempted an alternative fitness assay of T7 exposed to mutagen to shed light on these anomalies of initial fitness measures. As in prior work, T7 was grown on cells exposed to nitrosoguanidine to introduce mutations, but here fitnesses of those stocks were assayed in the absence of mutagen. We obtained samples of phage exposed for zero generations (mutagen-free control), exposed for one generation, or exposed for three generations. These phage were then added to separate, mutagen-free cultures of cells, with titers determined at 2 min and 13 min; the sample at 2 min represents the number of viable parent phage and the sample at 13 min represents the number of progeny after a single infection cycle. For the mutation-free control, burst occurs at just under 11 min (Heineman and Bull 2007), so the titer at 13 min should be an effective burst size. For mutated genomes, the effective burst may be reduced by three mechanisms—reduced adsorption rate, reduction of progeny in the cell, and delay of burst (an infected cell not yet burst plates as a single progeny phage). If there is a rapid accumulation of deleterious, nonlethal mutations during the first few generations of mutagenesis, these effective burst sizes should decline progressively with the number of generations of exposure.

Figure 1 reveals that the highest effective burst size is with no exposure (red squares), but no further decline is seen between 1 and 3 generations of exposure (solid blue and open black circles). There is thus an apparent contradiction here: from Equation 2, a decline in the first generation should be followed by additional declines in the next few generations unless $\bar{s}$ is near 1. Herein lies a methodological problem, however. The observed decline in the first generation need not be due to classic mutations; it could be due to mutagen effects on the nongenomic part of the virion or to epigenetic effects on the genome (see the following section for an example of the former). Thus these data fail to provide clear evidence of a short-term fitness decline from nonlethal mutations. These results are consistent with those of Springman *et al.* (2010), whose assay of initial fitness based on six generations of mutagenic growth matched the estimate based on one generation. The failure of T7 fitness to decline on long-term mutagenic propagation appears to reside at a fundamental level of evolution, such as neglecting beneficial mutations or overestimating the deleterious fraction of mutations, not to a confounded estimate of W_0.

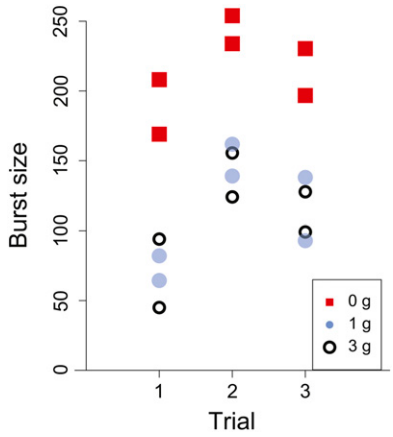

Figure 1 Effective burst sizes after exposure to nitrosoguanidine for zero, one, and three generations. Across three trials, there is a clear pattern of decline from zero- to one-generation exposure and no evident decline from one- to three-generations exposure. For each trial (conducted with a different, freshly mutagenized phage stock) the three pairs of values exhibit statistically significant heterogeneity. Methods: Populations of $T7_{Hi}$ from Heineman and Bull (2007) were grown in 10 μg/ml of the mutagen nitrosoguanidine for a single infection cycle (one generation, or 1 g), or for 70 min (three generations) or grown without mutagen (zero generations) and collected over chloroform to kill cells. These phage stocks were then added to mutagen-free cells that had been grown for 1 hr to a density of 1×10^8 cells/ml and plated at 2 min and 13 min after infection without killing cells. Burst size was calculated as the phage titer at 13 min divided by the titer at 2 min. Points shown are the raw data. Broth was LB and cells *Escherichea coli* IJ1133; methods, bacterial strain genotype and recipes otherwise follow Springman *et al.* (2010).

Avoiding problems with in vivo mutagenesis: Virions exposed to mutagen

If viral growth in the presence of mutagen confounds mutagenic effects on the virus with impairment of host, one solution is to mutate the genome separately from the host, by treating free virus particles (virions) with mutagen. The fitness of both the wild-type and mutated viruses can be measured in the mutagen-free environment, fully avoiding the estimation biases highlighted in the preceding section. If sequencing is employed to measure the average mutation rate, and a Poisson distribution is assumed for the incidence of mutations per genome, the surviving fraction of virions can be used to calculate the lethal mutation rate. An average burst size among the survivors may even be used to estimate an average fitness effect of nonlethal mutations.

Virion survival upon exposure to mutagens was assayed extensively nearly half a century ago (*e.g.*, Benzer 1961; Tessman *et al.* 1964, 1965; Tessman 1968; Botstein and Shortle 1985). For chemicals supposedly creating base changes (hydroxylamine, nitrous acid, EMS, MMS), the typical decay in viability was log-linear with time of exposure.

Log-linear decay is expected if there is a constant rate of lethal mutations per unit time, even if the fraction of mutations that are lethal differs across genes. For example, if the mutation rate per unit time is U_i for gene i (U is the total genomic rate and $U_i = p_iU$) and the lethal fraction within that gene is L_i, the viable fraction of genomes viable (V) under a Poisson model is

$$\begin{aligned} V &= \prod_i e^{-U_iL_i} \text{ or} \\ \log(V) &= -\sum_i U_iL_i \\ &= -U\sum_i p_iL_i = -U\bar{L}, \end{aligned} \tag{3}$$

where $\sum_i p_i = 1$. $\bar{L}$ is independent of the total mutation rate, so if the mutation rate U is linear with time of exposure, the plot of $\log(V)$ will decline linearly with time, as commonly reported. In phage T4, for example, the log linearity spans the entire viability plot, nearly 6 orders of magnitude survival (Tessman 1968).

At face value, the virion survival methods developed in this early work are ideally suited to test the foundations of lethal mutagenesis theory, because they seem to provide genome survival rates that could be aligned with base substitution rates. However, there are reasons to question whether those mutagens act solely or even primarily through classical mutations. Consider the mutagen hydroxylamine, reported to induce largely C → T transitions. Using phage T4, Tessman (1968) observed log-linear declines in viability across 5–6 orders of magnitude with time of exposure. Mutations in the survivors were detected phenotypically in a diagnostic phage gene, indicating that at least some of the mutagen's effect was an induction of standard mutations. However, base substitutions may not be the only cause of virion death. Hydroxylamine is also known to cause breaks in DNA and to cleave specific protein bonds, both of which will cause virion death (Rhaese and Freese 1968; Tessman 1968; Taylor *et al.* 1970; Rauko *et al.* 1993; Robins *et al.* 2013). Thus, virion decay may be due to a combination of effects, only some of which are due to classic mutations.

Is log linearity expected from point mutations?

Another reason to question whether log-linear survival under chemical mutagenesis is due chiefly to base changes comes from recent work measuring retention of protein function for genes subjected to various base substitution rates. In these studies, mutations were generated by error-prone PCR and average base substitution rates from libraries of clones were evaluated directly by sequencing. Relationships between base substitution rate and protein survival did not obey log linearity but instead exhibited a shoulder effect, with an accelerating downward slope at higher mutation rates—as if robustness to mutation deteriorated once a few mutations were acquired in the protein (Bloom *et al.* 2005, 2006, 2007). This curvature implies negative epistasis on a multiplicative fitness scale, the deleterious effects of mutations together being stronger than when alone. The absence of log-linear survival for individual proteins seems incompatible with log linearity of genome survival. However, it is not intuitive how much deviation from log linearity to expect at the genome level when log linearity is violated at the gene level. We thus offer a model bridging the two processes. The model is simplified for analytical tractability.

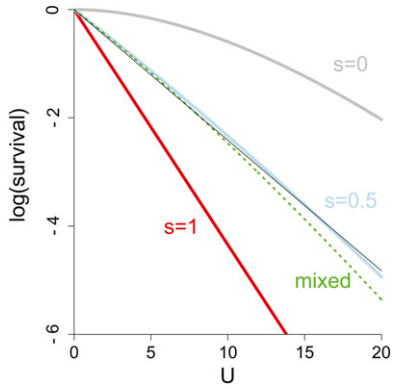

Figure 2 Log survival per mutation rate for different values of s for a genome with 30 genes (n = 30) from Equation 4; all genes experience the same value of s except where indicated. The gray curve represents maximal epistasis within genes (s = 0), the red line is no epistasis (s = 1), and the blue curve represents intermediate epistasis s = 0.5 (epistatic if fitness is multiplicative within genes). The curve for s = 1 is strictly linear. The thin black line overlaid on the blue curve is also strictly linear, so the blue curve is seen to bend only slightly. The dashed green curve represents a genome with an equal mix of genes with s = 1 and s = 0 and is seen to have slightly greater curvature than the curve for $s = \frac{1}{2}$. However, both the blue and green curves deviate from linearity only slightly over four logs—suggesting that a modest level of epistasis within genes might not be detectable in these plots.

Suppose that there are n genes in the genome. Genome survival is the product of gene survival rates across all genes. Let survival of gene i be 1 in the absence of mutation, $1 - s$ for a single mutation, and 0 for two or more mutations in that gene; the latter assumption implies negative epistasis within genes on the multiplicative scale except for $s = 1$. (Lethality of two or more mutations is assumed because it greatly facilitates the calculations.) Let the mutation rate of gene i be U_i ($\sum_i U_i = U$), and for further mathematical necessity, all genes have the same mutation rate,: $U_i = U/n$, and all mutations have the same fitness effect, s.

The surviving fractions for genomes with different numbers of mutations are calculated as follows. Let Y_i be the number of mutations arising in gene i at rate U/n. The probability of k single mutations in any set of k genes is $\binom{n}{k}$ times the probability that the singles fall in the first k genes (since all possible combinations have the same probability) and is

$$\begin{aligned} &\binom{n}{k} P(Y_1 = 1, Y_2 = 1, \cdots, Y_k = 1, Y_{k+1} = 0, \cdots, Y_n = 0) \\ &= \binom{n}{k} \left[\frac{U}{n} e^{-U/n}\right]^k \left[e^{-U/n}\right]^{n-k} = \binom{n}{k} \left(\frac{U}{n}\right)^k e^{-U}. \end{aligned}$$

Table 1 Effect of hydroxylamine treatment on T7

Exposure time (hr)	Survival[a]	Muts/genome[b]	C → T/genome
0	1.0	2.3	0.7
21	0.05, 0.05, 0.025	6.9	4.9
45	0.0005, 0.001, 0.0002	11.1	8.1

Methods. Virions of $T7_{Hi}$ from Heineman and Bull (2007) were suspended in 0.1 M sodium phosphate buffer (pH 6.0) with 1 mM EDTA and 3.5 μM phenol red, incubated with or without hydroxylamine (HA) at 0.3M for the time indicated (Tessman 1968). At the end of exposure, HA-treated cultures were neutralized with acetone at 1.5% and NaOH at 0.1 M, the latter to maintain the pH near 7.0. DNA was extracted from all exposed virions regardless of viability and subjected to 454 pyrosequencing; reads were mapped with breseq (Barrick *et al.* 2009) and mutations enumerated on a per-read basis; bases were counted only if the quality score was at least 25. Our inference of the *in vivo* mutation spectrum assumes that the *in vitro* sequencing method accurately mimics the *in vivo* conversion of lesions into base changes.

[a] survival range across three separate trials, each standardized to the HA-free sample.

[b] Mutations observed in the control may be due to partly or entirely to sequencing error.

A genome with k single mutations has survival $(1 - s)^k$, so viability over all possible numbers of single mutations per genes becomes

$$V = \sum_{k=0}^{n} \binom{n}{k} \left(\frac{U(1-s)}{n}\right)^k e^{-U} = \sum_{k=0}^{n} \binom{n}{k} \left(\frac{U(1-s)}{n}\right)^k (1)^{n-k} e^{-U} = \left(1 + \frac{U(1-s)}{n}\right)^n e^{-U}. \quad (4)$$

When $s = 1$, we recover the result that survival is merely e^{-U}. If $s = 0$, and then as n increases without bound, the survival probability approaches $e^U e^{-U} = 1$. This result is easily understood: with infinitely many genes and a finite number of mutations, two mutations never fall in the same gene.

As seen in Figure 2, the deviation from log-linear viability is pronounced for strong epistasis ($s = 0$, gray curve) but barely detectable for moderate epistasis ($s = \frac{1}{2}$, blue line). The slopes are greatly affected by s, but it is the curvature that is of interest when interpreting the published data, as we have no basis for interpreting that slope. The observation of log linearity for whole genomes is thus visually compatible with moderate epistasis at the level of individual genes. For comparison to the $s = \frac{1}{2}$ case, the figure also shows a curve for a mixed genome in which half the genes experience $s = 1$ and the other half experience $s = 0$ (viability follows $(1 + (U/n))^{n/2} e^{-U}$). The latter shows a slightly greater curvature than does $s = \frac{1}{2}$ (dashed green line) but is still not pronounced; curvatures for s between $\frac{1}{2}$ and 1 are likewise slight (not shown). While the generality of these results remains to be demonstrated with more realistic models, the suggestion here is that log linearity is only slightly violated by moderate levels of epistasis within genes.

Hydroxylamine and T7

Using T7, we revisited the effect of hydroxylamine on virion survival and also measured mutation rates of DNA extracted from exposed virions over a three- to four-log drop in survival (Table 1). The viability varies fivefold across the three trials for the longest treatment time (45 hr) and twofold for the 21-hr treatments. When subtracting baseline error/mutation rates, nearly the total increase in average mutations per genome is due to C → T transitions, as expected for this mutagen. To attribute the loss in viability to the measured increase in point mutations, consider the 21-hr data. The most generous, high survival rate is 0.05, and the excess mutations per genome is 4.6. The lethal fraction (L) of those mutations under a Poisson model is found as $e^{-4.6L} = 0.05$, or $L = 0.65$. Considering that approximately one-third of the genome is nonessential, that the lethal rate observed in a variety of small genome viruses ranges from 0.2 to 0.4 (Sanjuán *et al.* 2004; Domingo-Calap *et al.* 2009), and that only ~1/5 of residues of an essential T7 protein are lethal when mutated (Robins *et al.* 2013), it appears that the loss in viability from exposure to hydroxylamine is too high to be explained by the observed rate of point mutations (under a Poisson model).

A possible concern with virion mutagenesis is an unequal distribution of mutations across the genome. DNA density in phage heads is 500 mg/ml, approaching crystalline density (Earnshaw and Casjens 1980). If only a small fraction of the genome is exposed to mutagen, as may be plausible, the observed relationship between average mutation rate and virion survival may poorly reflect the genome-wide tolerance of mutations. A plot of C → T mutation frequency across hydroxylamine-treated genomes suggests only modest variation in exposure across different regions (Figure 3). As the sliding window method obscures site–site variation on a local scale, we also counted the incidence of sites with no mutations, singles, and so on (17586, 837, 134, 40, 11, 7, 1, 1), observing one site each with six and seven mutations. Ignoring the site–site variation in number of reads, these numbers show a strong violation of Poisson due to an excess of multiply hit sites, consistent with the much earlier results of Benzer (1961). Thus, there appears to be quantitative, nonrandom variation in exposure to the mutagen, but not strong enough that large portions of the genome completely avoid exposure.

Measuring the relevant mutation rate in double-strand genomes

Although it is easy to model an arbitrary mutation rate, fitting the relevant empirical rate may be challenging. In a double-stranded genome, mutagenesis typically converts a base on one side of the duplex, not also its complement. If

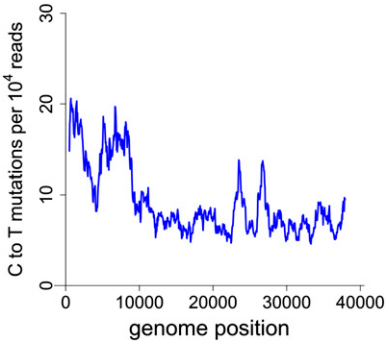

Figure 3 A sliding window analysis of the number of C to T mutations in a $T7_{Hi}$ genome treated with hydroxylamine 45 hr exhibits a modest deviation from uniformity, suggestive of unequal exposure across the genome. The counts are 10,000 times the total number of C to T mutations observed in all reads within the window, divided by the number of reads at which a C was observed as the template base within the window. A window span of 1001 nucleotides was applied stepwise across the genome and plotted every 50 bases. Methods are as in Table 1, except that a minimum quality score of 20 was applied here.

no subsequent replication occurs before the genome infects, virions will then carry different mutations on each strand. There are several ramifications of this heteroduplex asymmetry that depend on the viral life cycle, the fitness effect of the mutations, and on recombination. For tractability in everything that follows, we consider a single episode of mutation—progeny are not exposed to further mutagenesis.

Consider first that all mutations have no fitness effect (are neutral): all infections are equally viable regardless of mutation content on either strand. Upon infection by a heteroduplex genome, and in the absence of DNA repair, semiconservative replication of the parent genome will distribute mutations on each strand to half the genomes destined for progeny, but now as homoduplexes. In T7 at least, recombination among the homoduplexes will create progeny genomes with a mix of mutations from the two parental strands; for convenience the mix may be assumed to obey linkage equilibrium, each mutation at a frequency of 0.5. If the mutation rate per parental strand is U, progeny carry an average total of $2U$ mutations in their double-strand genomes. With free recombination, individual progeny inherit half the total, an average of U. The distribution (Poisson) will be the same as in single strands of the parent genomes. The same result would apply if there was no recombination during the process.

The situation changes when mutations have deleterious effects. Consider the extreme case that all mutations are lethal and furthermore that they are lethal regardless of strand. If the lethal rate per strand is U_L, parent genomes are now subject to a lethal rate of $2U_L$, twice the per-strand rate. In the absence of other types of mutations, sequencing would observe a mutation rate U_L, whereas the survival rate would be observed to decline at $2U_L$, an apparent impossibility under the standard Poisson model. As no real system experiences only lethal mutations, the ramifications of this process would merely be an overestimate of the fraction of mutations that are lethal. Recombination makes no difference in this case.

A less extreme case is one that applies to T7: all genes are transcribed from the same strand, and mutations on that strand will largely determine viability of the infection (Molineux 2005). Again assume that only lethal mutations occur, albeit a lethal mutation on the non-transcribed strand (NTS) does not manifest an effect until its complement is generated by replication using products made from the transcribed strand of the infecting genome. Lethals arise at rate U_L on each genomic strand, independently of the opposing strand. The transcribed strand escapes lethals with probability e^{-U_L}, which is thus the fraction of virions that produce viable progeny. Of interest, then is the fraction of progeny in those viable infections that inherit lethal mutations from the NTS.

Within viable infections and with complete recombination, the surviving progeny for different numbers of mutations on the NTS are enumerated in Table 2. The fraction of all progeny in viable infections escaping inheritance of lethals from the NTS sums as

$$e^{-U_L}\left[1+\frac{U_L}{2}+\left(\frac{U_L}{2}\right)^2\frac{1}{2!}+\cdots\right]=e^{-U_L+U_L/2} \qquad (5)$$

and is thus $e^{-U_L/2}$. A plot of surviving infectious particles (parent genomes producing at least some viable progeny) would be log-linear with slope $-U$, but the number of viable progeny from mutagenized parents would have slope $-1.5U$. As before, therefore, mutations arising independently on each strand will have a greater lethal effect than expected from the mutation rate observed using sequence methods (which is per strand).

Abolishing recombination among progeny genomes raises fitness in this case. In the absence of recombination, minimally half the progeny of a viable infection are themselves viable because their genome is descended from the mutation-free transcribed strand of the parent; progeny survival follows $e^{-U}(1 + e^{-U})/2$. Progeny survival without recombination is invariably larger than with recombination, markedly so at high mutation rates. Without recombination, survival is approximately $\frac{1}{2}e^{-U}$ and would not necessarily be empirically distinguishable from e^{-U}.

Discussion

The genetic principle behind lethal mutagenesis is that a high mutation rate will drive the population to a mutation–selection balance in which fitness is depressed well below the starting value. A proposed theoretical framework holds that the equilibrium genetic load depends only on the genomic deleterious mutation rate (Kimura and Maruyama 1966)—a profound generality if true. The alternative outcomes of population survival *vs.* extinction also depend on whether the intrinsic fecundity is high enough to offset the drop in genetic fitness (Bull *et al.* 2007), but the focus here has been on the genetic impact of mutations.

Table 2 Progeny surviving lethal mutations from the NTS

NTS mutations	Probability	Recombinant progeny survival
0	e^{-U_L}	1
1	$U_L e^{-U_L}$	$\frac{1}{2}$
2	$\frac{U_L^2 e^{-U_L}}{2}$	$\frac{1}{4}$
3	$\frac{U_L^3 e^{-U_L}}{3!}$	$\frac{1}{8}$
⋮	⋮	⋮

The underlying genetic foundation of lethal mutagenesis has been formally tested in one case. T7 was subjected to a genomic viable mutation rate of over four viable mutations per genome per generation; despite a predicted fitness decline of 60% on a log scale, a slight fitness increase was observed at 200 generations (Springman *et al.* 2010). A second example can also be interpreted in this framework. In Hayden *et al.* (2011), a ribozyme of 197 bases was subjected to a genomic rate of one mutation per generation. Across 10 generations of mutagenesis, fitness (measured as the ribozyme activity selected) declined to about two-thirds of the initial value by generation 4 but recovered to the initial value by the 10th. The model here predicts a fitness equilibrium of $e^{-1} = 0.37$ times that of the initial value if all mutations are deleterious. The molecular basis of the ribozyme fitness recovery was not investigated, although the modal number of mutations per genome had increased by ~2 at the 10th generation.

The work here has attempted to reconcile the basic model with the specifics of T7 *in vivo* mutagenesis; a secondary goal was to consider the feasibility of more tractable empirical systems. This effort has supported the original interpretation of the T7 work and finds no suggestion that failure of the theory resides at a mismatch between methods and interpretation of results. Instead, it appears that the sustained moderately high mutation rate simply fails to cause a progressive fitness decline over the course of 200 generations. Thus, a mutation rate much higher than 4–6 is required to cause a major fitness decline.

In light of the present results, we can identify three possible causes of the failure of T7 to match the predicted fitness decline:

1. Adaptive evolution. Equations 1 do not admit beneficial mutations. Beneficial mutations have the effect of raising initial fitness after the fact, as if the fitness of the initial virus is higher than its measured value. The predicted equilibrium fitness under mutagenesis is scaled relative to the initial fitness, so neglecting the ultimate contribution of beneficial mutations that would have effectively boosted initial fitness will predict an equilibrium fitness too low. In the T7 study, there was direct evidence of beneficial mutations ascending during the period of evolution but no unambiguous way to determine how much fitness benefit they contributed.
2. Recombination with negative epistasis. The Kimura–Maruyama model underlying Equations 1 assumes an absence of recombination; it makes no assumption about epistasis because epistasis is irrelevant in the absence of recombination. The protocol used for T7 allowed recombination. If mutations interact so that the deleterious effects attributed to mutations are worse in combination with other deleterious mutations than alone (negative epistasis), recombination can reduce the mutation load, profoundly so in extreme cases (Kimura and Maruyama 1966). Thus, if the mutations arising in T7 tend to have negative epistasis, recombination in T7 may account for a substantially higher equilibrium fitness than predicted by e^{-U}. The average levels of epistasis observed in studies of mutation interaction in microbes have typically not supported the negative model (*e.g.*, Jasnos *et al.* 2008; Lali and Elena 2012), but those studies have been limited to interactions among few mutations, not necessarily shedding light on interaction terms for large numbers of mutations in the T7 populations.
3. Overestimated deleterious mutation rate. The genomic viable mutation rate was estimated directly for T7 as 4.3 per generation in the mutagenic environment Springman *et al.* (2010). However, the fraction of those mutations that were deleterious was taken at 0.6 based on estimates of other viruses, all viruses with smaller genomes than T7 (40 kb for T7 *vs.* 11, 5.4, and 4.2 kb for the others). In T7, that fraction may be <0.6. Predicted fitness declines approximately linearly with the deleterious mutation rate (from 21.8 with $U = 0$ to 8.6 with $U = 2.6$), so even a halving of U would still result in a profound fitness drop. Thus, this violation of the model is, by itself, not likely to account for a qualitative change in the outcome.

The study of Hayden *et al.* (2011) admits the first and last of these explanations as well, although it is ambiguous whether recombination could have existed in the amplification step as well.

The T7 study remains enigmatic, but two avenues of further exploration seem worthy of attention. First and as noted above, our model accommodates beneficial mutations somewhat restrictively, yet beneficial mutations are known to have evolved. Other work has also begun to incorporate beneficial mutations into models of evolution at high mutation rate (Gerrish *et al.* 2013). A particularly attractive result from their study that provides conditions sufficient for population decline is

$$-U\delta > \sigma^2, \tag{6}$$

where U is the total genomic mutation rate, δ is the average effect of a mutation (combining beneficial and deleterious mutations), and σ^2 is the variance in fitness. Although δ and σ^2 both change during evolution, the latter will be near 0 for a clonal population (as true of the wild-type T7 at the start of mutagenesis). For the three viruses with measures of the distribution of fitness effects, the average effect is consistently and strongly negative (*e.g.*, on the order of 0.3; Sanjuán *et al.* 2004; Domingo-Calap *et al.* 2009). If those

numbers apply to T7, this model also predicts a rapid fitness decline in T7 over the first few generations of mutagenesis (as does our model), although a long-term increase is certainly compatible with this model. More generally, Gerrish *et al.* (2013) point the way toward a more realistic understanding of the practicality of lethal mutagenesis.

The distribution of deleterious deleterious fitness effects in T7 is an additional focus for work attempting to resolve the empirical enigmas. Fitness effect sizes in T7, with a genome of 40 kb, need not reflect those of the small genome viruses for which fitness effects have been measured directly (Sanjuán *et al.* 2004; Domingo-Calap *et al.* 2009). Indeed, deleterious effect sizes in some essential bacterial genes have proved remarkably small (Lind *et al.* 2010). Small average effect sizes could explain why fitness is slow to decline upon onset of mutagenesis.

Despite the unsolved high fitness in two studies of adapted genomes exposed to a measured level of mutagenesis, the collective quantitative work on lethal mutagenesis suggests a practical approach to extinction. When the deleterious effects of mutation from a single generation of mutagenesis (as given by Equation 2 for $k = 1$) are severe enough by themselves to enact population decline, the lethal mutagenesis threshold may be considered to satisfy a *strong* criterion—population size will begin declining immediately. Alternatively, when the deleterious effects of mutation from a single generation of mutagenesis are not severe enough to prevent population expansion but the long-term effects (as given by Equation 1a) do ensure population decline, the threshold may be considered to satisfy a *weak* criterion. When only the weak criterion is satisfied, the population may expand hundreds of generations before the deleterious load is expected to halt growth (Bull and Wilke 2008). Thus other evolutionary processes outside the model, especially beneficial mutations, may intervene in time to avoid extinction.

When lethal mutagenesis succeeds, the mechanisms contributing to population decline may encompass more than base changes. Attributing fitness decline to just the observed base mutation rate may often be a mistake. Nonetheless, immediate population decline should be a robust indicator of successful lethal mutagenesis by whatever mechanisms, except when extreme forms of resistance to mutagen can arise with single mutations.

A type of "automatic" lethal mutagenesis has been proposed for asexual populations evolving toward an optimum: ever-greater mutation rates are selected because mutator alleles disproportionately generate the favored mutants and thus hitchhike to fixation. The genetic load increases much more slowly than does the benefit but eventually drives the population extinct (Gerrish *et al.* 2007; Sniegowski and Gerrish 2010). Experimental evidence with mutator strains supports the process (Gentile *et al.* 2011). Although this form of lethal mutagenesis is distinct from one in which the high mutation rate is imposed, it does rely on a slow buildup of the genetic load—which may underlie the observed failure in T7. The genetics of the bacterial case may thus provide a useful comparison for viral work.

Acknowledgments

We thank I. Wang and S. Abedon for strains of non-T7 phages used in our (unsuccessful) attempts to estimate multiple kill curve parameters. L. Wahl identified an initial mistake in our (A1) and worked with us to confirm legitimacy of text Equation 2; we thank her, Phil Gerrish, and Paul Sniegowski for comments that improved the manuscript. This work was funded by National Institutes of Health grant GM 57756 (J.J.B) and GM 76040 (P.J.). J.J.B. receives support from the Miescher Regents Professorship at the University of Texas.

Literature Cited

Anderson, J. P., R. Daifuku, and L. A. Loeb, 2004 Viral error catastrophe by mutagenic nucleosides. Annu. Rev. Microbiol. 58: 183–205.

Barrick, J. E., D. S. Yu, S. H. Yoon, H. Jeong, T. K. Oh *et al.*, 2009 Genome evolution and adaptation in a long-term experiment with *Escherichia coli*. Nature 461(7268): 1243–1247.

Benzer, S., 1961 On the topography of the genetic fine structure. Proc. Natl. Acad. Sci. USA 47(3): 403–415.

Bloom, J. D., J. J. Silberg, C. O. Wilke, D. A. Drummond, C. Adami *et al.*, 2005 Thermodynamic prediction of protein neutrality. Proc. Natl. Acad. Sci. USA 102(3): 606–611.

Bloom, J. D., S. T. Labthavikul, C. R. Otey, and F. H. Arnold, 2006 Protein stability promotes evolvability. Proc. Natl. Acad. Sci. USA 103(15): 5869–5874.

Bloom, J. D., F. H. Arnold, and C. O. Wilke, 2007 Breaking proteins with mutations: threads and thresholds in evolution. Mol. Syst. Biol. 3: 1744–4292.

Botstein, D., and D. Shortle, 1985 Strategies and applications of in vitro mutagenesis. Science 229(4719): 1193–1201.

Bull, J. J., and C. O. Wilke, 2008 Lethal mutagenesis of bacteria. Genetics 180: 1061–1070.

Bull, J. J., R. Sanjuan, and C. O. Wilke, 2007 Theory of lethal mutagenesis for viruses. J. Virol. 81(6): 2930–2939.

Chung, D. H., Y. Sun, W. B. Parker, J. B. Arterburn, A. Bartolucci *et al.*, 2007 Ribavirin reveals a lethal threshold of allowable mutation frequency for Hantaan virus. J. Virol. 81(21): 11722–11729.

Crotty, S., C. E. Cameron, and R. Andino, 2001 RNA virus error catastrophe: direct molecular test by using ribavirin. Proc. Natl. Acad. Sci. USA 98: 6895–6900.

Domingo-Calap, P., J. M. Cuevas, and R. Sanjuan, 2009 The fitness effects of random mutations in single-stranded RNA and DNA bacteriophages. PLoS Genet. 5(11): e1000742.

Earnshaw, W. C., and S. R. Casjens, 1980 DNA packaging by the double-stranded DNA bacteriophages. Cell 21(2): 319–331.

Eigen, M., 1971 Selforganization of matter and the evolution of biological macromolecules. Naturwissenschaften 58: 456–523.

Eigen, M., and P. Schuster, 1977 The hypercycle, a principle of natural self-organization. A. Emergence of the hypercycle. Naturwissenschaften 64: 541–565.

Eigen, M., J. McCaskill, and P. Schuster, 1988 Molecular quasispecies. J. Phys. Chem. 92: 6881–6891.

Gentile, C. F., S.-C. Yu, S. A. Serrano, P. J. Gerrish, and P. D. Sniegowski, 2011 Competition between high- and higher-mutating strains of *Escherichia coli*. Biol. Lett. 7(3): 422–424.

Gerrish, P. J., A. Colato, A. S. Perelson, and P. D. Sniegowski, 2007 Complete genetic linkage can subvert natural selection. Proc. Natl. Acad. Sci. USA 104(15): 6266–6271.

Gerrish, P. J., A. Colato, and P. D. Sniegowski, 2013 Genomic mutation rates that neutralize adaptive evolution and natural selection. J. R. Soc. Interface 10(85): 20130329.
Graci, J. D., and C. E. Cameron, 2002 Quasispecies, error catastrophe, and the antiviral activity of ribavirin. Virology 298(2): 175–180.
Hayden, E. J., E. Ferrada, and A. Wagner, 2011 Cryptic genetic variation promotes rapid evolutionary adaptation in an RNA enzyme. Nature 474(7349): 92–95.
Heineman, R. H., and J. J. Bull, 2007 Testing optimality with experimental evolution: lysis time in a bacteriophage. Evolution 61: 1695–1709.
Jasnos, L., K. Tomala, D. Paczesniak, and R. Korona, 2008 Interactions between stressful environment and gene deletions alleviate the expected average loss of fitness in yeast. Genetics 178: 2105–2111.
Johnson, T., 1999 The approach to mutation-selection balance in an infinite asexual population, and the evolution of mutation rates. Proc. Biol. Sci. 266(1436): 2389–2397.
Kimura, M., and T. Maruyama, 1966 The mutational load with epistatic gene interactions in fitness. Genetics 54: 1337–1351.
Lali, J., and S. F. Elena, 2012 Magnitude and sign epistasis among deleterious mutations in a positive-sense plant RNA virus. Heredity 109(2): 71–77.
Lind, P. A., O. G. Berg, and D. I. Andersson, 2010 Mutational robustness of ribosomal protein genes. Science 330(6005): 825–827.
Martin, V., A. Grande-Perez, and E. Domingo, 2008 No evidence of selection for mutational robustness during lethal mutagenesis of lymphocytic choriomeningitis virus. Virology 378(1): 185–192.
Molineux, I., 2005 The T7 group, pp. 277–301 in *The Bacteriophages*, edited by R. Calendar. Oxford University Press, Oxford.
Rauko, P., L. Novotn, and E. Balzov, 1993 DNA breakage and inactivation resulting from hydroxylamine and/or cis-diamminedichloroplatinum(II) interactions with plasmid DNA. Int. J. Biochem. 25(10): 1475–1481.
Rhaese, H. J., and E. Freese, 1968 Chemical analysis of DNA alterations. I. Base liberation and backbone breakage of DNA and oligodeoxyadenylic acid induced by hydrogen peroxide and hydroxylamine. Biochim. Biophys. Acta 155(2): 476–490.
Robins, W. P., S. M. Faruque, and J. J. Mekalanos, 2013 Coupling mutagenesis and parallel deep sequencing to probe essential residues in a genome or gene. Proc. Natl. Acad. Sci. USA 110 (9): E848–E857.
Sanjuán, R., A. Moya, and S. F. Elena, 2004 The distribution of fitness effects caused by single-nucleotide substitutions in an RNA virus. Proc. Natl. Acad. Sci. USA 101: 8396–8401.
Sniegowski, P. D., and P. J. Gerrish, 2010 Beneficial mutations and the dynamics of adaptation in asexual populations. Philos. Trans. R. Soc. Lond. B Biol. Sci. 365(1544): 1255–1263.
Springman, R., T. Keller, I. J. Molineux, and J. J. Bull, 2010 Evolution at a high imposed mutation rate: adaptation obscures the load in phage T7. Genetics 184: 221–232.
Taylor, A. T., S. B. Crist, and O. W. Jones, 1970 Alteration of native DNA transcription by the mutagen hydroxylamine. Cancer Res. 30(1): 95–99.
Tessman, I., 1968 Mutagenic treatment of double- and single-stranded DNA phages T4 ans S13 with hydroxylamine. Virology 35(2): 330–333.
Tessman, I., R. K. Poddar, and S. Kumar, 1964 Identification of the altered bases in mutated single-stranded DNA. I. Hydroxylamine, ethyl methanesulfonate and nitrous acid. J. Mol. Biol. 9: 352–363.
Tessman, I., H. Ishiwa, and S. Kumar, 1965 Mutagenic effects of hydroxylamine in vivo. Science 148(3669): 507–508.
Wiehe, T., 1997 Model dependency of error thresholds: the role of fitness functions and contrasts between the finite and infinite sites models. Genet. Res. 69(1): 127–136.
Wilke, C. O., and C. Adami, 2003 Evolution of mutational robustness. Mutat. Res. 522(1–2): 3–11.
Wilke, C. O., J. L. Wang, C. Ofria, R. E. Lenski, and C. Adami, 2001 Evolution of digital organisms at high mutation rates leads to survival of the flattest. Nature 412(6844): 331–333.

Communicating editor: C. H. Langley

Appendix: Mean Fitness per Generation Under Mutagenic Exposure

We consider the fitness of a genotype G in generation $k + 1$ based on its mutational and selective history through the previous generations. Let $P(M_j)$ represent the Poisson probability of the set of M mutations that were acquired by genotype G in generation j: $m_{j,1}$ mutations acquired in class 1 at rate U_1, $m_{j,2}$ in class 2 at rate U_2, and so on for $j = 1$ though k. Mean fitness in generation j is $\overline{W}_j$. The fitness of a genotype G in generation $k + 1$ is

$$\begin{aligned} W_{G_{k+1}} = {} & P(M_1)\,\frac{(1-s_1)^{m_{1,1}}\ldots(1-s_n)^{m_{1,n}}}{\overline{W}_0} \\ & \times P(M_2)\,\frac{(1-s_1)^{m_{1,1}+m_{2,1}}\ldots(1-s_n)^{m_{1,n}+m_{2,n}}}{\overline{W}_1} \\ & \vdots \\ & \times P(M_{k+1})\,\frac{(1-s_1)^{m_{1,1}+\ldots m_{k+1,1}}\ldots(1-s_n)^{m_{1,n}+\ldots m_{k+1,n}}}{\overline{W}_k}, \end{aligned} \tag{A.1}$$

where $\overline{W}_0 \equiv 1$.

Summing over the $m_{j,i}$ for all genotypes,

$$\overline{W}_{k+1} = E\left(\frac{\prod_{i=1}^{n}(1-s_i)^{\sum_{j=1}^{k+1} m_{j,i}(k-j+2)}}{\prod_{j=0}^{k}\overline{W}_j}\right). \tag{A.2}$$

Solving for $\overline{W}_1$ and $\overline{W}_2$ directly, it becomes clear that several terms cancel. As one proceeds to other $\overline{W}_l$ we hypothesize

$$\overline{W}_l = E\left(\prod_{i=1}^{n}(1-s_i)^{lm_{1,i}}\right). \tag{A.3}$$

It follows directly from (A.2) that (A.3) is true when $l = 1$. By induction we assume that (A.3) is true for all $l \leq k$ and then we will show that (A.3) is true for $l = k + 1$. We begin by separating the $j = 1$ term in (A.2) to get

$$\overline{W}_{k+1} = \frac{E\left(\prod_{i=1}^{n}(1-s_i)^{\sum_{j=2}^{k+1}m_{j,i}(k-j+2)}\right)}{\prod_{j=0}^{k}\overline{W}_j} E\left(\prod_{i=1}^{n}(1-s_i)^{(k+1)m_{1,i}}\right). \tag{A.4}$$

We now show that the numerator and denominator in the first term in (A.4) cancels. To show this we start by changing the index by letting $l = k - j + 2$ and use the fact that all of the Poison random variables are independent to get

$$\overline{W}_{k+1} = \frac{\prod_{i=1}^{n}\prod_{l=1}^{k}E(1-s_i)^{lm_{k-l+2,i}}}{\prod_{j=0}^{k}\overline{W}_j} E\left(\prod_{i=1}^{n}(1-s_i)^{(k+1)m_{1,i}}\right).$$

Note that $m_{1,i}$ has the same distribution as $m_{k-l+2,i}$ so we can rewrite the above as

$$\begin{aligned}
\overline{W}_{k+1} &= \frac{\prod_{i=1}^{n}\prod_{l=1}^{k}E(1-s_i)^{lm_{k-l+2,i}}}{\prod_{j=0}^{k}\overline{W}_j} E\left(\prod_{i=1}^{n}(1-s_i)^{(k+1)m_{1,i}}\right) \\
&= \frac{\prod_{i=1}^{n}\prod_{l=1}^{k}E(1-s_i)^{lm_{1,i}}}{\prod_{j=0}^{k}\overline{W}_j} E\left(\prod_{i=1}^{n}(1-s_i)^{(k+1)m_{1,i}}\right) \\
&= \frac{\prod_{l=1}^{k}E\left(\prod_{i=1}^{n}(1-s_i)^{lm_{1,i}}\right)}{\prod_{j=0}^{k}\overline{W}_j} E\left(\prod_{i=1}^{n}(1-s_i)^{(k+1)m_{1,i}}\right) \\
&= E\left(\prod_{i=1}^{n}(1-s_i)^{(k+1)m_{1,i}}\right).
\end{aligned} \tag{A.5}$$

The last line follows by the induction hypothesis.

For independent random variables X_i,

$$E\left(\prod_{i=1}^{n} r_i^{X_i}\right) = \prod_{i=1}^{n} G_{X_i}(r_i)$$

where $G_X(r)$ is the probability generating function of the variable X. For a Poison random variable X with rate U it is well known that the probability generating function is given by $G_X(r) = e^{(r-1)U}$. We can thus rewrite (A.3) as

$$\overline{W}_{k+1} = \prod_{i=1}^{n} G_{m_{1,i}}\left[(1-s_i)^{(k+1)}\right] = \mathrm{e}^{\sum_{i=1}^{n}\left((1-s_i)^{(k+1)}-1\right)U_i} \tag{A.6}$$