Shifts in mutation spectra enhance access to beneficial mutations

Mrudula Sane[1], Gaurav D Diwan[1,2,3], Bhoomika A Bhat[4], Lindi M Wahl[5] and Deepa Agashe[1]*

[1] National Centre for Biological Sciences, GKVK Campus, Bellary Road, Bengaluru, India

[2] CellNetworks, Bioquant, University of Heidelberg, 69120 Heidelberg, Germany

[3] Biochemie Zentrum Heidelberg (BZH), University of Heidelberg, 69120 Heidelberg, Germany

[4] Undergraduate Programme, Indian Institute of Science, Bengaluru, India

[5] Western University, London, Ontario, Canada

*Correspondence:

dagashe@ncbs.res.in

SUMMARY

Mutations are critical for evolutionary change. Although mutation spectra (mutational biases) vary
greatly across organisms, we lack direct evidence of their evolutionary consequences. Here, we show
experimentally that a biased spectrum alters the distribution of fitness effects in a bacterium, and should
facilitate adaptation by increasing beneficial mutations and reducing genetic load. Adaptive walk
simulations show that this advantage arises via better sampling of mutational classes that were poorly
explored by the ancestor. Hence, reversing the spectrum of a biased ancestor is generally
advantageous, while introducing bias in an unbiased ancestor is selectively neutral. Indeed, across
bacterial lineages, evolutionary transitions in DNA repair enzymes – which shape spectra – typically
reverse ancestral bias. These broad consequences of mutation spectra imply a major role in shaping
evolutionary dynamics.

Keywords

DFE, pleiotropy, mutation bias, DNA repair, adaptive walk, ancestral reconstruction

Mutations are a critical source of genetic variation, and their frequency, spectrum and fitness impacts are thus important to predict evolutionary change. The mutation spectrum refers to the frequency of various classes of sampled mutations (e.g. transversions vs. transitions (Tv/Ts), or base pair vs. copy number changes). It is often biased towards specific mutational classes (mutational bias) [1, 2], driven by genetic or plastic changes in enzymes that repair replication- or damage-induced errors in DNA. Such changes in DNA repair function typically alter both mutation rate and spectra, as observed in "mutator" genotypes with loss-of-function mutations [3], and during stress-induced mutagenesis [4, 5]. Since the mutation spectrum determines the pool of available genetic variation on which selection can act, it has the potential to dramatically alter evolutionary outcomes. However, in contrast to our understanding of mutation rate, the evolutionary impact of the mutation spectrum has been poorly explored [6].

Recent studies strongly suggest important evolutionary consequences of the spectrum, although the generality of these effects remains unclear. For instance, biased mutation spectra can shape key genome features such as nucleotide [7, 8], codon, and amino acid composition [9]. Simulations show that a bias favouring beneficial classes of mutations can alter adaptation in specific fitness landscapes under diverse conditions [6, 10], and influence the evolutionary success of mutator genotypes [11, 12]. By altering the pool of mutations available for selection, mutation spectra may determine the genetic basis of adaptation [10, 13, 14] and drive convergent evolution [15, 16]. Finally, from a practical standpoint, mutational biases can shape the evolution of resistance to antibiotics [11, 17, 18] and anti-cancer drugs [19]. However, we lack direct experimental evidence demonstrating the fitness consequences of mutation spectra. It is also unclear whether (and under what conditions) shifts in the mutation spectrum could be selectively favoured; and if such shifts occur frequently across evolutionary time.

We addressed these outstanding questions using experiments, simulations, and phylogenetic analyses. To measure the immediate evolutionary impact of the mutation spectrum, we estimated the genome-wide distribution of fitness effects (DFE) of new mutations in *E. coli* (Fig. S1). We manipulated the mutation spectrum by deleting a DNA repair gene (Δ*mutY*, "mutator") from our wild type ("WT") strain, altering both the mutation spectrum and rate [3]. We evolved independent lineages of WT and mutator under mutation accumulation (MA), and sequenced whole genomes of evolved isolates to identify strains carrying a single new mutation each. We thus minimized the impact of selection on the mutation spectrum and the DFE, and decoupled the effects of the mutator's spectrum from its high mutation rate. Measuring the fitness effect of each mutation in twelve different environments, we show that a shift in the mutation spectrum alters the DFE, reducing the genetic load and increasing the supply of beneficial mutations. Next, we demonstrate the generality and longer-term evolutionary consequences of these results by simulating adaptive walks across NK fitness landscapes [20]. We find that a reversal in the mutation spectrum is often beneficial – especially for well-adapted genotypes (such as our WT) – because it allows populations to sample previously unexplored mutational space. Finally, by inferring past evolutionary transitions in DNA repair enzymes, we show that most bacterial lineages have

experienced multiple reversals in mutation spectra. Together, our work suggests that shifts in mutation

spectra can facilitate rapid evolution in new environments, fuelling adaptation via previously unsampled

evolutionary paths.

The mutator has a distinct DFE with more beneficial mutations, reduced load, and altered

pleiotropic effects

From our MA experiments, we obtained 80 evolved WT (*21*) and 79 mutator strains, each carrying a

single distinct mutation with respect to its ancestor (Tables S1, S2). We measured the fitness effect

(relative maximal growth rate) of each mutation in 12 environments with different carbon sources,

constructing global and environment-specific empirical DFEs (Fig. 1, Fig. S2). The mutator and WT had

distinct global DFEs (Fig. 1A; KS test: D=0.23, p=2.2×10^{-16}); but the environment alone did not

significantly impact fitness effects (ANOVA: relative fitness ~ strain x environment, p_{strain}=1.25×10^{-18},

p_{env}=0.69, $p_{strainxenv}$=1.8×10^{-5}). New mutations had a median fitness advantage of ~3% in the mutator,

and an equivalent fitness disadvantage in the WT (Fig. 1A). This pattern was consistent for

environment-specific DFEs in 9 of 12 environments (Fig. 1B, Table S3), indicating that the mutator has

a consistently beneficial-shifted DFE compared to the WT. The probability of finding these differences

by chance was 0.0006 (Fig. S3A), and our sample size of ~80 mutations per strain adequately captured

the presumed "true" global DFE (Fig. S3B, Table S4). The DFE differences were also robust to removing

four conditionally lethal mutations found in evolved WT (Fig. S4, Table S3). Our single-mutation

genome-wide DFEs have more beneficial mutations than previously reported DFEs (*22*), which were

typically estimated using populations evolved under selection, carrying many or unknown numbers of

mutations (i.e. confounded by epistasis and selection). Since there is no precedence for single-mutation

genome-wide DFEs, we re-analysed a single-mutation single-gene DFE for antibiotic resistance (*23*)

and found ~31% beneficial mutations. Thus, our estimated DFEs likely represent the genome-wide DFE

of single mutations, suggesting that beneficial mutations may not be as rare as is usually believed.

To quantify the impact of the mutator's distinct DFE, we coarsely classified each mutation as neutral

(0.95 > relative fitness > 1.05, conservatively accounting for ~5% fitness measurement error), beneficial

(relative fitness >1.05), or deleterious (relative fitness <0.95). On average, the mutator had a greater

fraction of beneficial mutations than the WT, but fewer deleterious mutations (Fig 2A, Table S5).

Accounting for its order-of-magnitude higher mutation rate (Table S6), we estimated (following *24*) that

the mutator should have an ~14-fold greater genome-wide supply of beneficial mutations, and an ~0.45-

fold lower genetic load than the WT (Tables S7, S8). In contrast, if we ignored the mutator's distinct

DFE, the supply of beneficial mutations would only be ~1.5 fold greater than WT, and the genetic load

would be ~4 fold higher (Tables S7, S8). Thus, depending on their mutation rate and DFE, mutators

may have a substantially higher supply of beneficial mutations, and effectively no genetic load.

The pleiotropic effect of beneficial mutations (Fig. 2B) is important in new environments, because it can

facilitate adaptation (via synergistic pleiotropic benefits) and shape fitness tradeoffs (via antagonistic

pleiotropy). The mutator had a distinct distribution of pleiotropic effects in 11 of 12 environments (Fig.

S5, Table S9). These differences were driven by a higher incidence of beneficial synergistic pleiotropy and lower deleterious synergistic pleiotropy in the mutator; but equally low antagonistic pleiotropy in both strains (Fig. 2C–D, Fig. S5). Beneficial mutations in the mutator were also beneficial across many more environments (KS test, D=0.36, p=6.7x10^{-5}); whereas deleterious mutations were deleterious in fewer environments than WT (KS test, D=0.51, p=1.9x10^{-9}) (Fig. 2G, Fig. S6). Thus, new beneficial mutations in the mutator are more likely to facilitate adaptation across many environments, but are no more likely to generate tradeoffs. We speculated that these differences may arise because the high proportion of beneficial mutations in the mutator might increase the likelihood of beneficial synergistic pleiotropy. Indeed, simulating an increase in the median fitness effect of the WT DFE without changing its shape ("WT-beneficial shift") mimicked pleiotropic effects observed in the mutator in all but one environment (compare Fig. 2C-D, Fig. S5, Table S9). Conversely, simply reducing the median fitness effect in the WT DFE ("WT-deleterious shift") lowered beneficial synergistic pleiotropy in 7 of 12 environments (compare Fig. 2C and 2F, Fig. S5, Table S9). Ignoring the mutation spectrum can thus cause overestimation of a mutator's genetic load, and underestimation of both the beneficial supply of mutations as well as their pleiotropic effects during adaptation.

The mutator's distinctive DFE arises from its distinct mutation spectrum
We had hypothesized that the biased spectrum of the mutator could generate a distinct DFE. However, other global effects of the initial *mutY* gene deletion could also lead to a DFE shift. For instance, if deleting *mutY* reduced the mutator's fitness, a larger fraction of new mutations might be beneficial (*25*). However, the mutator ancestor had lower fitness than WT in only two environments (Fig 3A, Table S10). Alternatively, epistatic interactions with *mutY* could increase the deleterious effect of new mutations in the WT. However, paired strains carrying the same mutation in either the WT or mutator background had similar fitness in all environments (Fig. 3B, Table S11). Thus, the observed shift in the mutator DFE cannot be explained by global effects of the original *mutY* deletion.

To directly test the effect of the spectrum, we asked whether the mutator over-samples specific classes of mutations that happen to be more beneficial (or less deleterious). The mutator strongly favours transversions, GC→AT mutations, base pair substitutions (BPS) and coding mutations (Fig. 3C, Table S6). The resulting skew in sample size of mutation classes limited our ability to determine the association between mutation bias and fitness effects across strains. However, in the WT, transversions and GC→AT mutations were indeed less deleterious than transitions and AT→GC mutations (Fig. 3D–E; fitness effects were pooled across environments, Table S12). We could not test the effect of BPS/indel bias, and the coding/non-coding bias had no effect (Fig. S7, Table S12). Other properties of the mutations (e.g. their location in the genome or specific genes, or the resulting amino acid changes) did not differ across strains (Figs. S8, S9; Tables S13, S14). Thus, the mutation spectrum – specifically the strong Tv and GC→AT bias – shapes the distinct DFE of the mutator.

Simulations demonstrate a general benefit of reversing mutation spectra

To test the generality and impact of our empirical results, we simulated adaptive walks (following (26)),
modelling sequences of length N such that mutations could be classified as Ts vs. Tv or GC→AT vs.
AT→GC. Starting from a genetically uniform population (i.e. a strong-selection, weak-mutation regime)
with a randomly chosen ancestor sequence and mutation spectrum (absolute Tv bias = fraction of
transversions), we allowed populations to explore successive mutational steps at randomly chosen loci
(bases). At various points during the walk, we generated a DFE by simulating 300 possible substitutions,
and computing their fitness given the underlying fitness landscape (affected by K other randomly chosen
loci, to incorporate epistasis). From the DFE, we calculated the fraction of beneficial (f_b) and deleterious
(f_d) mutations, and their effect size. We initially set N=200 and K=1 (i.e. mild epistasis, with each locus
epistatically coupled to 1/199=0.5% loci), and present the average outcomes for 500 adaptive walks on
500 randomly-generated fitness landscapes (Fig. S10 shows variation across walks).

As expected, the mean population fitness increased during the adaptive walk, with a concomitant
reduction in f_b and beneficial effect size, an increase in f_d, and a relatively constant deleterious effect
size (Fig. 4A). Setting ancestral Tv bias to 0.45 (mimicking our WT), we compared the DFE generated
by the ancestor at various time points, to the DFE that would be created if the bias were changed
(mimicking a mutator). The mutator thus started with the same fitness and sequence as the ancestor,
differing only in spectrum. As the population evolves (i.e. ancestor f_b decreases), mutators with a
stronger bias (i.e. higher Tv) sample proportionally more beneficial mutations than the ancestor (Fig.
4B). More generally, exploring well-adapted populations (f_b = 0.04) but varying the ancestral mutation
spectrum, we found that mutators that reverse the ancestral Tv bias have the greatest advantage. If the
ancestor is biased toward Ts, mutators with higher Tv show a greater increase in f_b values compared
to the ancestor (Fig. 4C), with larger effect sizes (Fig. 4D). Simultaneously, f_d also decreases (Fig. 4E),
with a small increase in the effect size (Fig. 4F). These results were consistent for more rugged fitness
landscapes with higher epistasis (up to 16% or 43% interacting loci, Fig. S11).

The symmetric nature of these results implies that a specific mutation class or change in spectrum is
not always beneficial. In the example presented here, a bias towards more transversions was beneficial
only because the ancestor underwent a prior adaptive walk with a spectrum that favoured transitions.
Analogous results hold for GC→AT bias: reversing the WT bias yields more beneficial mutations, and
the effect is stronger with better-adapted WT populations (Fig. 4G, Fig. S12). Thus, these results can
be generalized for any axis of the mutation spectrum. Interestingly, for an unbiased ancestor (Tv bias
= 0.66 since 2/3 of possible mutations are transversions), introducing any bias is selectively neutral,
regardless of how long it has evolved in the landscape (Fig. 4H–I). However, if the ancestor has evolved
with even a slightly biased spectrum (e.g. our WT), a shift in the opposite direction (as observed in our
mutator) allows the population to explore mutational space that was not previously well-sampled,
increasing the chances of finding new beneficial mutations. After a period of evolving with the new
spectrum, a change in the fitness landscape (e.g. due to epistasis or environmental change) may again
render a spectrum shift advantageous. Together, these results lead to two predictions: (a) Since an
unbiased state is not selectively favoured, biased spectra can evolve through drift (b) Once a biased

spectrum evolves, a bias reversal is beneficial (including a return to the unbiased state), and may thus
occur frequently over long evolutionary periods.
190
Evolutionary transitions in DNA repair enzymes indicate mutation spectrum reversals in most
bacterial lineages
Next, we inferred the long-term dynamics of mutation spectra by mapping evolutionary transitions in 11
bacterial DNA repair enzymes whose loss changes the Tv or GC→AT bias in *E. coli* (Table S15; limited
data indicate consistent effects across species, Table S16). Identifying gene orthologues in 1093 extant
bacterial genomes, we used ancestral reconstruction to infer enzyme gains and losses across the
bacterial phylogeny (following (*27*); Fig 5A). Broadly consistent with a previous analysis (*28*), we found
frequent evolutionary transitions in most enzymes. Using the predicted set of enzymes found in
hypothetical ancestors (nodes), we estimated their Tv and GC→AT bias relative to WT *E. coli* and
traced the change in bias over time (Fig. 5B, Fig. S14). Over >80% lineages experienced a reversal of
bias during their evolutionary history, with ~50% experiencing one reversal in Tv bias or two reversals
in GC→AT bias (Fig. 5C). In contrast, very few lineages reinforced the ancestral bias through
successive enzyme gain or loss events (Fig. 5C, reinforced vs. reversed: Tv, χ^2=45.4, p=1.5x10^{-11};
GC→AT, χ^2=181.2, p=2.2x10^{-16}). These patterns likely arise from the simultaneous impact of DNA
repair enzymes on mutation rate and spectrum. While it is difficult to separate these effects, note that
enzymes with a weak effect on the mutation rate but a strong impact on the spectrum are also highly
dynamic (e.g. *ung* and *mutM*, whose impact on mutation spectrum we report here for the first time;
Table S15). The relative biases estimated here should not be taken literally, since we cannot account
for factors such as other repair enzymes whose effects on the mutation spectrum are unknown.
Nevertheless, 33 additional genes from multiple DNA repair modules also show frequent evolutionary
transitions (Fig. S15). Thus, sustained changes in mutation spectra likely occur frequently across
evolutionary timescales, with long-term impacts on the set of genetic variants available for selection.
213
Conclusions
Together with prior studies showing that mutation biases are pervasive and may contribute to the
genetic basis of adaptation (see Introduction), our results demonstrate that mutation spectra may be
key to driving adaptation and innovation in natural populations under myriad scenarios. We also uncover
why pervasive mutation biases might evolve, complementing emerging understanding of conditions
when such biases can alter adaptation (*29*, *30*). Our results have clear implications for the evolutionary
dynamics of mutators, which are governed by their genetic load and supply of beneficial mutations (*24*,
31, *32*). In a new environment (away from fitness optima), these parameters are primarily driven by
mutators' high mutation rate (*33*) – particularly under strong selection (*34*) – allowing them to hitchhike
with beneficial mutations (*35*). We show that the beneficial supply rate is strongly influenced by the
mutation spectrum even in well-adapted populations under weak selection, perhaps explaining why
mutators with relatively small increases in mutation rate are abundant in natural bacterial populations
(*36*). Thus, the mutation rate and spectrum may jointly govern the rise, persistence, and fall of mutators
under varying degrees of selection (*24*, *37*). These effects deserve further attention (*12*).

Our work also offers new insights into the evolution of genome-wide mutational biases. GC-rich bacterial

genomes typically show more GC→AT than AT→GC substitutions (*38*). Although the excess is

attributed to selection, its mechanistic basis has remained a puzzle (*7*). Similarly, the pervasive

transition bias observed across taxa does not appear to be broadly advantageous (*39*). Instead, our

simulations suggest that these patterns may arise because sampling from mutational classes that were

previously under-sampled (given the prevalent bias) are more likely to yield new beneficial mutations,

and are hence more likely to fix. Conversely, the population would have already sampled most beneficial

mutations within the class favoured by the existing bias; hence, such mutations will be under-

represented among fixed substitutions. For instance, the protein-disruptive effects of GC→AT mutations

are predicted to be less deleterious in GC-rich genomes, and vice versa (*12*). The symmetry of these

results supports our suggestion that the genomic context and prior evolutionary history may be sufficient

to explain widespread yet varied genomic biases, without invoking selection or mechanisms favouring

specific mutational classes. The rarity of species with unbiased mutation spectra (*2*) is therefore

counterintuitive but not surprising, given that transitions from unbiased to biased spectra are not

disadvantageous and may occur easily and frequently via changes in DNA repair function. Hence, a

biased spectrum may be the norm rather than exception.

Finally, our work leads to the exciting prediction that populations could hop along similar or different

axes of the mutation spectrum and sustain long-term adaptation, even when relatively well-adapted. At

a point where the beneficial mutation supply is limited by the existing spectrum, an antiparallel (opposite

direction on the same axis) or orthogonal change (on a different axis of the spectrum) could allow further

sampling of new beneficial mutations, facilitating rapid adaptation. Our phylogenetic analysis likely

underestimated such evolutionary shifts in bacterial mutation spectra, which may occur frequently on

shorter timescales via horizontal transfer and recombination (*40*) and drive polymorphism in spectra

across natural bacterial isolates (*36*, *41*). We predict multiple cascading effects of shifts in mutation

spectra, including a reduction in the waiting time for beneficial mutations, decreased likelihood of

mutational meltdown in populations evolving under drift, and distinct genetic pathways of adaptation.

We hope that future work will test these predictions.

MATERIALS AND METHODS

Bacterial strains

We obtained the wild-type (WT) strain of *E. coli* K-12 MG1655 from the Coli Genetic Stock Centre

(CGSC, Yale University), streaked it on LB (Luria Bertani) with 2% agar, and chose one colony at

random as the WT ancestor for subsequent experiments. We similarly obtained the mutator strain of *E.*

coli (Δ*mutY*) from the Keio collection (BW25113 strain background; (*42*)) of gene knockouts at the same

stock centre. These gene knockouts were made by replacing open reading frames by a Kanamycin

resistance cassette, such that removing the cassette generates an in-frame deletion of the gene. The

design of gene deletion primers ensured that downstream genes were not disrupted due to polar effects

(*42*). We moved the knockout locus from the BW25113 background into the MG1655 (WT) background
using P1-phage transduction (*43*). After transduction, we removed the kanamycin resistance marker by
transforming kanamycin-resistant transductants with pCP20, a plasmid carrying the flippase
recombination gene and Amp^R resistance marker. We grew ampicillin resistant transformants at 42 ºC
in LB broth overnight to cure pCP20. We streaked out 10 µL of these cultures on LB plates. After 24
hours, we replica-plated several of these colonies on both kanamycin-LB agar plates and ampicillin-LB
agar plates, to screen for the loss of both kanamycin and ampicillin resistance. We PCR-sequenced the
knockout locus from several colonies to confirm removal of the kanamycin cassette. We used the same
protocol to generate Δ*ung* and Δ*mutM* strains, to determine their mutation spectra.

To test whether genetic background alters the fitness effects of single mutations observed during
mutation accumulation, we deleted the *mutY* locus from 19 evolved WT MA lines (each carrying a single
mutation with respect to the WT ancestor). We first raised P1 phage on the Δ*mutY*::kanR strain in the
MG1655 genetic background, as described above. We used the lysate to infect evolved WT MA lines
carrying single mutations. After transduction, we selected kanamycin resistant colonies, removed the
kanamycin resistance cassette, and cured the pCP20 plasmid as described above. We confirmed the
mutY deletion by amplifying and sequencing the *mutY* locus. Thus, for each of the 19 mutations, we
obtained paired strains with either a WT or Δ*mutY* genetic background.

Experimental evolution under mutation accumulation (MA)
We founded 38 WT MA lines from 38 different colonies (two lines per Petri plate) incubated at 37ºC, as
described earlier (*21*). For each line, every 24 hours we streaked out a random colony (closest to a pre-
marked spot) on a fresh agar plate. Every 4-5 days, we inoculated a part of the transferred colony in LB
broth at 37ºC for 2-3 hours and froze 1 mL of the growing culture with 8% DMSO at -80 ºC. For the
current study, we used stocks frozen on days 39, 104, 140, 200, 250 and 300 (Fig S1A). For the mutator,
we similarly founded 300 MA lines, and evolved them for 12 days (~330 generations; Fig S1A). This
time was expected to be sufficient for the accumulation of one mutation per line on average, based on
previous estimates of mutation rate (*3*). We made stocks of all lines as described for WT. To determine
the impact of two previously untested DNA repair enzymes (*ung* and *mutM*) on the mutation spectrum,
we evolved 20 independent MA lines of each of deletion strain for approximately 4400 and 5775
generations respectively.

Whole-genome sequencing to identify single mutations and determine mutation spectra
We sequenced individual colonies from MA experiments to identify clones carrying only a single
mutation relative to the WT or Δ*mutY* ancestor. For WT, we sequenced MA lines at days 39, 104, 140,
200, 250, 300 as described earlier (*21*), to obtain 80 single mutational steps (Fig S1B) that we used for
the analyses in the present study. For Δ*mutY*, we sequenced whole genomes of the 300 colonies from
the MA experiment frozen on day 12, and the Δ*mutY* MA ancestor (Fig S1B). To do this, we inoculated
2 µL of the frozen stock of each evolved MA line (or the ancestor) in 2 mL LB, and allowed the cells to
grow overnight at 37ºC with shaking at 200 rpm. We extracted genomic DNA (GenElute Bacterial

Genomic DNA kit, Sigma-Aldrich) and quantified it (Qubit HS dsDNA assay, Invitrogen). We prepared

paired-end libraries from each line using the Illumina Nextera XT DNA library preparation kit using the

manufacturer's instructions. We sequenced the Δ*mutY* MA ancestor on the Illumina MiSeq platform

using the 2x250bp paired-end reaction chemistry, and obtained a total of 1.2M reads with quality > Q30,

corresponding to an average per base coverage of ~120x. We sequenced libraries from the evolved

MA lines on the Illumina HiSeq 2500 platform using the 2x100bp paired-end reaction chemistry. We

obtained sufficient reads for 299 out of the 300 samples (median 3.3 million reads per sample; range:

1.15 – 7.13 million reads per sample), which were retained for further analysis. We discarded reads

with quality scores less than Q30, retaining >=90 % of the reads in all 299 samples. With these reads,

we obtained a median per base coverage of ~ 92x (range: 32x – 198x). For each sample, we aligned

quality-filtered reads to the NCBI reference *E. coli* K-12 MG1655 genome (RefSeq accession ID

GCA_000005845.2) using the Burrows-Wheeler short-read alignment tool, BWA (*44*). We generated

pileup files using SAMtools (*45*) and used VARSCAN to extract a list of base-pair substitutions and

short indels (< 10bp) (*46*). We only retained mutations with >80% frequency that were represented by

at least 5 reads on both strands for further analysis.

After removing ancestral mutations from the evolved lines, we identified 79 Δ*mutY* MA evolved lines

having a single new mutation compared to the ancestor. We performed all further experiments and

analyses with these 80 WT and 79 Δ*mutY* isolates. To calculate mutation rates and estimate overall

mutation spectra, we used mutations called from all sequenced isolates at day 300 for WT and day 12

for Δ*mutY*, irrespective of whether they had single or multiple mutations. Thus, we estimated mutation

spectra for both strains using a large pool of mutations (total 237 mutations in WT and 669 in mutator).

Similarly, we determined the mutation spectra of Δ*ung* and Δ*mutM* strains by sequencing 20 isolates

per strain (total 157 and 138 mutations respectively).

We next calculated mutational biases from the fraction of different types of mutations observed in our

evolved strains (Table S6). For example, we calculated the WT Tv bias as the fraction of

transversions/(fraction of transitions + fraction of transversions), i.e. 0.447/(0.553 + 0.447) = 0.447.

Therefore, we used WT Tv bias = 0.45 for our adaptive walk simulations (Fig. 4). Similarly, we calculated

the WT GC→AT bias as fraction of GC→AT mutations/(fraction of AT→GC mutations + fraction of

GC→AT mutations) = 0.52/(0.36 + 0.52) = 0.59 (ignoring 26 AT→TA and GC→CG mutations that did

not contribute to the bias; see Table S6). For convenience, for the adaptive walk simulations, we

assumed the fraction of GC→AT mutations = 0.55 and the fraction of AT→GC mutations = 0.35, thus

obtaining WT GC→AT bias = 0.61 (Fig. 4).

Measuring growth rate as a fitness proxy

We measured growth rates of evolved isolates with single mutations and their respective ancestors in

liquid culture media: LB broth (Miller, Difco), or M9 minimal salts (Difco) + 5 mM carbon source (glucose,

trehalose, fructose, lactose, maltose, galactose, succinate, pyruvate, melibiose, fumarate or malate;

Sigma-Aldrich). We inoculated each isolate from its freezer stock into M9 minimal salts medium with

0.4% glucose, and allowed it to grow at 37 ºC with shaking at 200 rpm for 16 hrs. We inoculated 6 µL

of this culture into 594 µL growth media (LB broth or M9 minimal salts + 5 mM carbon source) in 48-

well plates (Costar) and incubated in a shaking tower (Liconic) at 37 ºC. Plates were read by an

automated growth measurement system (Tecan, Austria) every 45 minutes for 18 hrs. We measured

the growth rate of three technical replicates per isolate per growth medium; meaningful biological

replicates could not be obtained since we had a single colony at the end of each MA line. In each 48-

well plate, we included the WT ancestor to estimate variation in growth rates across plates and to

calculate relative fitness of evolved isolates. We used the average growth rate across technical

replicates as the estimated ancestral growth rate, in calculations of relative fitness of the evolved strain.

We estimated maximum growth rate using the Curve Fitter software (*47*). We used the average growth

rate of three technical replicates of each mutant to calculate relative fitness as: Growth rate of evolved

strain/Growth rate of ancestral strain. Relative fitness < 1 indicated a deleterious mutation, while relative

fitness > 1 indicated a beneficial mutation. The error in measurement of growth rates across technical

replicates (run on the same day) was less than 5%. Hence, we considered mutants with relative fitness

> 0.95 and relative fitness < 1.05 as showing no change in fitness, or neutral.

Generating and subsampling from the pooled DFE for each strain

We wanted to understand how often could the DFEs differ significantly across our two strains just by

chance, in any particular environment. A powerful method to ask this question is to pool fitness values

across strains to remove their association with the genotype, and then randomly subsample and

compare two distributions drawn from the pool. Therefore, we pooled fitness effects across WT and

Δ*mutY* (i.e. total 159 fitness values), for each of the 12 environments. We randomly picked 2 sets of 80

fitness values from this pool for each environment, and compared the distributions of fitness values in

the two sets using Kolmogorov-Smirnov (KS) tests. We performed 1000 iterations of this process for

each environment, and calculated the probability that randomly sampled DFEs differ significantly from

each other.

Next, we wanted to understand whether our sample size of 80 fitness values was sufficient to accurately

inform us of the underlying "global" or "real" DFE. We reasoned that the global DFE for a given genotype

may be better estimated by pooling the fitness values of single mutations for each genotype across

multiple environments. To determine the probability that a sample of a given size accurately represents

the larger distribution, we can then randomly subsample smaller distributions of varying sizes from this

"global" distribution, and compare them to the "global" distribution. Therefore, for each strain we pooled

fitness values across all environments, obtaining a pooled set of 912 fitness values for WT (after

excluding lethal mutations), and 948 fitness values for Δ*mutY*. For each strain, we randomly

subsampled sets of fitness values of size 50, 60, 70, 80, 90, 100, 120, 150, 200, 300, 400 and 500.

Note that our actual sample size (~80) is included in this range. For each subsample size, we performed

500 iterations, generating 500 DFEs per sample size for each strain. We then compared each of these

500 DFEs to the large pooled set using KS tests. We expected that as sample size increases, the

proportion of iterations in which the subsampled distribution is significantly different from the pooled

distribution should saturate. If our actual sample size (~80) was too low to capture the global DFE, we

expected to find that the proportion of iterations in which the subsampled distribution is significantly

different from the pooled distribution should be large. On the other hand, if our actual sample size was

sufficient to capture the underlying global DFE, we expected the proportion of iterations in which the

subsampled distribution is significantly different from the pooled distribution should be small.

Re-analysing published single-mutation single-gene datasets

We analysed raw data from (*23*) to obtain the median proportion of beneficial mutations, for their dataset

containing fitness effects of single mutations in the TEM-1 β-lactamase gene in *E. coli*. They had

measured the fitness effects of randomly generated point mutations and codon mutations in this gene,

by deep sequencing pooled libraries of cells harbouring the TEM-1 alleles, grown at 13 different

Ampicillin concentrations. They then used the sequencing read counts for each allele at the different

antibiotic concentrations to calculate fitness, and normalized all fitness values by the fitness of the WT,

and reported these in their study. Thus, normalized fitness values > 1 indicated beneficial mutations

and normalized fitness values < 1 indicated deleterious mutations. We used these normalized data,

provided in their Supplementary Data Table S1, for our analysis. We used the subset of fitness values

corresponding to single base changes (i.e. point mutations; Data Table S1 – Base Changes column =

1) to obtain a set of fitness values for 2583 point mutations. From this set, we found that 806 mutations

(31% of total) had fitness > 1.

Estimating expected genetic load and supply of beneficial mutations

We calculated the deleterious load and supply of beneficial mutations for each strain as described

earlier (*24*). Briefly, using the fraction of beneficial and deleterious mutations from the empirical DFEs

of WT and mutator, we calculated genetic load due to deleterious mutations as:

$$L_d = f_d \, N \, \mu$$

where L_d = deleterious genetic load; f_d = fraction of deleterious mutations; N = number of sites (here, *E.*

coli genome size, 4.64×10^6 bp); μ = average genome-wide mutation rate of each strain (bp^{-1} generation^{-1}

^{1}). Similarly, we calculated the supply of beneficial mutations as:

$$S_b = f_b \, N \, \mu$$

where S_b = beneficial supply and f_b = fraction of beneficial mutations. Using the L_d and S_b values for

each strain, we calculated the fold change in beneficial supply and genetic load of the mutator relative

to WT. Next, we calculated the mutator L_d and S_b assuming that the WT and mutator DFEs were

identical, and allowing for variation only in μ, as done in previous studies:

$$S_{b(WT\,DFE)} = f_{b(WT)} \, N \, \mu_{\Delta mutY}$$

$$L_{d(WT\,DFE)} = f_{d(WT)} \, N \, \mu_{\Delta mutY}$$

To measure the impact of accounting for the mutator's distinct DFE, we compared the fold change in

the mutator's S_b and L_d values compared to WT, with or without assuming distinct DFEs.

Measuring the incidence of pleiotropy among new single mutations

We estimated the incidence of pleiotropy among new single mutations in the mutator (Δ*mutY*)
background, as reported earlier for single mutations in WT *E. coli* (*21*). Briefly, for each possible pair of
environments (total 110 pairs, given 11 minimal medium environments), we calculated the proportion
of new mutations that increased fitness in one environment while decreasing fitness in another
(antagonistic pleiotropy), simultaneously increased or decreased fitness in both environments
(synergistic pleiotropic increases or decreases), changed fitness only in one environment, or did not
change fitness in either environment. For each focal environment, we then calculated the mean
proportion of each type of pleiotropic effect (considering all pairs that included that resource). Next, we
tested whether the differences in proportions of pleiotropy between WT and mutator could be
recapitulated solely by shifting the WT DFE towards more beneficial mutations ("WT-beneficial shift").
We calculated the difference in median fitness effects of mutator and WT in each resource, and added
this value to each of the WT fitness effects in that resource. The resulting DFE had a median equivalent
to the mutator DFE, but retained the shape of the WT DFE. Similarly, we generated a WT DFE shifted
towards more deleterious mutations ("WT-deleterious shift") by subtracting the same value from the
relative fitness values of WT mutations. We used these "shifted" DFEs to calculate the proportion of
fitness effects showing each type of (pairwise) pleiotropy. Finally, to measure pleiotropic effects across
multiple environments, we counted the number of environments in which each mutation had a
consistent effect (either beneficial or deleterious); and tested whether the proportion of mutations with
consistent effects differed significantly across WT and mutator.

Estimating the effect of mutation spectrum changes via simulations of adaptive walks
We used an established evolutionary model to test whether and to what extent changes in mutation
spectra are generally expected to affect a DFE. Following Stolzfus (*26*), we simulated adaptive walks
on the tunably-rugged NK fitness landscape (*20*), modelling sequences composed of strings of four
bases (A, C, G, T), allowing mutations to be classified as transitions or transversions, GC→AT or
AT→GC, and so on. We used neighbourhood size K=1, which corresponds to a relatively smooth fitness
landscape, but includes a degree of epistasis (ruggedness). Each locus in the sequence was randomly
assigned K=1 neighbors from among all loci in the model. The fitness contribution of locus i (ω_i) depends
on the state of that locus and the state of its neighbour. To create the fitness landscape, for each locus,
4^2 = 16 possible values of ω_i were drawn from a uniform distribution on (0,1); these values define ω_i
if the state of the locus and its neighbour are AA, AC, AG... through TT. The fitness of the sequence (W)
is then given by the sum of contributions from each locus, $W = \sum_{i=1}^{N} \omega_i$. When N=200, each landscape
contained 4^{200} = 2.58 x 10^{120} unique sequences and fitness values. We simulated adaptive walks by
assuming a genetically uniform population (strong-selection-weak-mutation regime). When simulating
Tv/Ts bias, starting at a randomly chosen resident (ancestor) sequence, we chose a single locus at
random for mutation. The mutation is a transition with probability Ts and a transversion with probability
Tv = 1 − Ts. Note that within transversions, each of the two possible outcomes (e.g. A to C or A to T) is
chosen uniformly at random. The selective effect of the mutation was computed as s=W_{mu}/W − 1, where
W_{mu} is the fitness of the mutant sequence, and W is the fitness of the current population sequence. The
mutant sequence replaces the population sequence (a "step" in the walk occurs) with probability 2s if

s > 0. The mutational spectrum (value of Ts) does not change during the adaptive walk. To investigate

higher degrees of epistasis (more rugged landscapes), we varied both N and K. Adaptive walks and

DFEs were simulated for N=8, 20 and 200, with K = 1, 2, or 3. Thus each locus is epistatically coupled

with between 1/199 = 0.5% and 3/7 = 42.8% of other loci in the genome.

We also implemented an alternative mutation scheme to test if the results varied for GC/AT biased

mutation spectra, where after evolved at a very high bias, the available sites constrain what mutations

can occur. In these simulations, mutations were constrained such that 10% were AT→TA or GC→CG

(mimicking our WT), and of the remaining mutations, fraction *x* were AT→GC, while (0.9-*x*) were

GC→AT. For this mutation scheme, GC bias is calculated as the fraction of AT→GC mutations among

all AT→GC or GC→AT mutations.

Inferring past evolutionary changes in mutation spectra

From previously published and our own mutation accumulation studies, we tabulated the predicted

effects of the loss of 11 genes encoding DNA repair enzymes on the mutation spectrum, focusing on

the Tv/Ts and GC/AT bias. To determine the evolutionary history of these genes, we detected

orthologues and mapped predicted gain and loss events for each gene across the bacterial phylogeny,

as described previously (*27*). Briefly, we downloaded the reviewed sequences for each gene from

UniProt, built a Hidden Markov Model (*48*), and searched for significant orthologues in the genomes of

1093 fully sequenced bacteria. Next, we determined the presence/absence of each gene in each extant

genome and carried out ancestral reconstruction of gene presence/absence on a previously described

phylogenetic tree (*49*), to infer the enzyme state at all nodes in the phylogeny. We used stochastic

character mapping (*50*) from the R package 'phytools' (*51*) with the following parameters: transition rate

matrix determination – Bayesian MCMC, prior distribution on the root node of the tree – estimated,

number of simulations – 500. We used a posterior probability cut-off of 0.7 to score evolutionary

transitions in character states. We used the predicted evolutionary transitions to count the total number

of gains and losses of each enzyme in the phylogeny. Next, we used the inferred state of all enzymes

at every node in the tree (including the eubacterial ancestor) to predict the overall Tv/Ts and GC/AT

bias of the hypothetical ancestors at all nodes. Where possible, we used experimental data on double

or triple knockouts of DNA repair mutants (Table S15); otherwise, we assumed additive effects of the

missing enzymes. For each of the 1093 lineages in the tree, we traced the change in bias over time, up

to the extant taxon (Fig. S13). We used these data to calculate the number and direction of changes in

bias in each lineage. Finally, we inferred past evolutionary transitions across a broader set of DNA

repair enzymes (without associated experimental data on their effects on the mutation spectrum), as

described above (Fig. S14).

ACKNOWLEDGEMENTS

We thank Shyamsunder Buddh, Joachim Krug, Christopher Marx, Saurabh Mahajan and Mukund

Thattai for their suggestions; Joshua Miranda for assistance with MA experiments; and the NGS facility

at the Bangalore Life Sciences Cluster. We acknowledge funding and support from the National Centre

for Biological Sciences (NCBS-TIFR), the Council for Scientific and Industrial Research India (Senior

Research Fellowship to MS), the University Grants Commission of India (Senior Research Fellowship

to GDD), the Department of Science and Technology India (KVPY Fellowship to BAB), the Natural

Sciences and Engineering Research Council of Canada (LMW), and the Wellcome Trust (GDD is

supported by grant 210585/B/18/Z to Robert B. Russell; and DBT/Wellcome Trust India Alliance grant

IA/I/17/1/503091 to DA). We also thank the International Centre for Theoretical Sciences (ICTS) for

facilitating our collaboration, initiated during the Bangalore School on Population Genetics and

Evolution (code: ICTS/popgen2020/01).

DATA ACCESSIBILITY

All data used for experimental and phylogenetic analysis are available as Supplementary Information

files. Genome sequences will be deposited in GenBank, and simulation and analysis code will be

available on Github.

AUTHOR CONTRIBUTIONS

MS designed and conducted experiments, analysed data, and drafted the manuscript. GDD designed

and conducted phylogenetic analyses. BAB conducted experiments and analysed data. LMW designed

and conducted simulation analyses. DA conceived and designed experiments and phylogenetic

analyses, analysed data, obtained funding, and wrote the manuscript.

FIGURES

Figure 1. Mutator DFE is shifted towards more beneficial mutations. (A) Global DFE (distribution of fitness effects) of single mutations in wild type (WT) and mutator (Δ*mutY*). The number of mutations per strain is given in parentheses in the key (n; each DFE was constructed with n x 12 data points, pooling across 12 environments). Black vertical line: relative fitness of 1 (neutral); arrows: median fitness; asterisk: significant difference. (B) Environment-specific DFEs (except LB, M9 minimal medium + 5mM indicated carbon source). Black dots: medians; horizontal grey shaded area: 5% measurement error margin (0.95 < relative fitness < 1.05); horizontal black line: relative fitness = 1; asterisks: significant differences across strains, colored by strain with higher fitness effects (Table S3); ns: not significant. Relative fitness of each strain was calculated with respect to its ancestor (WT or mutator).

Figure 2. Evolutionary consequences of a DFE shifted towards beneficial mutations. (A) Proportion of beneficial, deleterious and neutral mutations in WT and mutator. Median values across environments are indicated in the key; asterisks: significant differences between WT and mutator (Table S5). (B) Schematic showing possible pleiotropic effects of mutations in two environments (W: relative fitness). (C–F) Median proportions of pleiotropic fitness effects of mutations in WT and mutator, and WT DFE artificially shifted towards either beneficial or deleterious mutations (see Methods). Schematics indicate global DFEs (WT=cyan, mutator=pink, WT-shifted=light blue); black lines: relative fitness = 1. (G) Pleiotropic effects of mutations across multiple environments. Frequency distributions show the proportion of mutations with consistent fitness effects (beneficial or deleterious) in a given number of environments. Black lines indicate medians.

Figure 4. Adaptive walk simulations show general benefit of reversing mutation spectra. (A)
Change in mean population fitness, fraction of beneficial and deleterious mutations (f_b and f_d), and
magnitude of beneficial and deleterious effects (s_b and s_d) over the course of adaptive walks of a WT
ancestor. (B, G) Impact of altering the WT mutation bias on f_b at different points along the adaptive walk
(Tv bias = Tv/(Ts+Tv) mutations; GC→AT bias = GC→AT/(GC→AT+AT→GC) mutations; grey dashed
line: WT bias), as a function of the new mutation bias (mutator Tv or GC→AT bias). Line color: WT f_b
at a given time; steeper lines (lighter colors): later evolutionary time, better adapted WT. (C–F) Waterfall
plots show the impact of altering Tv bias (% change, Mutator – WT), as a function of WT (ancestor) and
mutator (new) bias, for a well-adapted WT population ($f_b = 0.04$). (H–I) Impact of introducing a mutation
bias in a hypothetical unbiased ancestor ("anc"), as a function of the new mutation bias (mutator Tv or
GC→AT bias).
575

Figure 5. Phylogenetic analysis shows frequent evolutionary transitions in DNA repair genes and reversals in mutation spectra. (A) Patterns of gene presence/absence in extant taxa were used to map evolutionary transitions in DNA repair genes on the bacterial phylogeny (total transitions are shown in the key), and infer the magnitude of Tv and GC→AT biases relative to WT *E. coli* in each extant taxon (heatmaps). (B) Inferred trajectories of mutational biases across 1093 lineages (C) Number of lineages showing spectrum shifts due to successive evolutionary transitions in repair enzymes (% values calculated as observed fraction across all lineages).

REFERENCES

1. T. Ferenci, Irregularities in genetic variation and mutation rates with environmental stresses. *Environ Microbiol*. **21**, 3979–3988 (2019).
2. V. Katju, U. Bergthorsson, Old trade, new tricks: Insights into the spontaneous mutation process from the partnering of classical mutation accumulation experiments with high-throughput genomic approaches. *Genome Biology and Evolution*. **11**, 136–165 (2018).
3. P. L. Foster, H. Lee, E. Popodi, J. P. Townes, H. Tang, Determinants of spontaneous mutation in the bacterium *Escherichia coli* as revealed by whole-genome sequencing. *Proc. Natl. Acad. Sci. U.S.A.*, 201512136–10 (2015).
4. A. A. M. Al Mamun *et al.*, Identity and function of a large gene network underlying mutagenic repair of DNA breaks. *Science*. **338**, 1344–1348 (2012).
5. R. Maharjan, T. Ferenci, Mutational signatures indicative of environmental stress in bacteria. *Molecular Biology and Evolution*. **32**, 380–391 (2014).
6. L. Y. Yampolsky, A. Stoltzfus, Bias in the introduction of variation as an orienting factor in evolution. *Evol. Dev.* **3**, 73–83 (2001).
7. D. Agashe, N. Shankar, The evolution of bacterial DNA base composition. *J. Exp. Zool.* **322**, 517–528 (2014).
8. H. Long *et al.*, Evolutionary determinants of genome-wide nucleotide composition. *Nat. ecol. evol.* **48**, 582 (2018).
9. J. Bohlin, O. Brynildsrud, T. Vesth, E. Skjerve, D. W. Ussery, Amino acid usage is asymmetrically biased in AT-and GC-rich microbial genomes. *PLoS ONE*. **8**, e69878 (2013).
10. A. Stoltzfus, L. Y. Yampolsky, Climbing Mount probable: Mutation as a cause of nonrandomness in evolution. *Journal of Heredity*. **100**, 637–647 (2009).
11. A. Couce, J. R. Guelfo, J. Blázquez, Mutational spectrum drives the rise of mutator bacteria. *PLoS Genetics*. **9**, e1003167 (2013).
12. A. Couce, O. Tenaillon, Mutation bias and GC content shape antimutator invasions. *Nat Commun*. **10**, 3114–9 (2019).
13. M. A. Streisfeld, M. D. Rausher, Population genetics, pleiotropy, and the preferential fixation of mutations during adaptive evolution. *Evolution*. **65**, 629–642 (2010).
14. D. R. Rokyta, P. Joyce, S. B. Caudle, H. A. Wichman, An empirical test of the mutational landscape model of adaptation using a single-stranded DNA virus. *Nat Genet*. **37**, 441–444 (2005).
15. A. Stoltzfus, D. M. McCandlish, Mutational biases influence parallel adaptation. *Molecular Biology and Evolution*. **34**, 2163–2172 (2017).
16. J. F. Storz *et al.*, The role of mutation bias in adaptive molecular evolution: insights from convergent changes in protein function. *Phil Trans Roy Soc*. **374**, 20180238–10 (2019).
17. J. L. Payne *et al.*, Transition bias influences the evolution of antibiotic resistance in *Mycobacterium tuberculosis*. *Plos Biol*. **17**, e3000265–23 (2019).
18. A. Bachar *et al.*, Point mutations in topoisomerase I alter the mutation spectrum in *E. coli* and impact the emergence of drug resistance genotypes. *Nucleic Acids Research*. **48**, 761–769 (2019).
19. S. M. Leighow, C. Liu, H. Inam, B. Zhao, J. R. Pritchard, Multi-scale predictions of drug resistance epidemiology identify design principles for rational drug design. *CellReports*. **30**, 3951–3963.e4 (2020).
20. C. A. Macken, A. S. Perelson, Protein evolution on rugged landscapes. *Proc. Natl. Acad. Sci. U.S.A.* **86**, 6191–6195 (1989).
21. M. Sane, J. J. Miranda, D. Agashe, Antagonistic pleiotropy for carbon use is rare in new mutations. *Evolution*. **72**, 2202–2213 (2018).
22. A. Eyre-Walker, P. D. Keightley, The distribution of fitness effects of new mutations. *Nat. Rev. Genet.* **8**, 610–618 (2007).
23. E. Firnberg, J. W. Labonte, J. J. Gray, M. Ostermeier, A comprehensive, high-resolution map of a gene's fitness landscape. *Molecular Biology and Evolution*. **31**, 1581–1592 (2014).
24. S. Wielgoss *et al.*, Mutation rate dynamics in a bacterial population reflect tension between adaptation and genetic load. *Proceedings of the National Academy of Sciences*. **110**, 222–227 (2013).
25. H. A. Orr, The population genetics of adaptation: The distribution of factors fixed during adaptive evolution. *Evolution*. **52**, 935–949 (1998).
26. A. Stoltzfus, Mutation-biased adaptation in a protein NK model. *Molecular Biology and Evolution*. **23**, 1852–1862 (2006).
27. G. D. Diwan, D. Agashe, Wobbling forth and drifting back: The evolutionary history and impact of bacterial tRNA modifications. *Molecular Biology and Evolution*. **35**, 2046–2059 (2018).

28. J. A. Eisen, P. C. Hanawalt, A phylogenomic study of DNA repair genes, proteins, and processes.
Mutation Research. **435**, 171–213 (1999).
29. A. V. Cano, J. L. Payne, Mutation bias interacts with composition bias to influence adaptive
evolution. *bioRxiv*. **6** (2020), pp. 2157–48.
30. K. Gomez, J. Bertram, J. Masel, Mutation bias can shape adaptation in large asexual populations
experiencing clonal interference. *bioRxiv*. **17** (2020), pp. e3000265–10.
31. K. Jain, A. James, Fixation probability of a nonmutator in a large population of asexual mutators.
Journal of Theoretical Biology. **433**, 85–93 (2017).
32. Y. Raynes, P. D. Sniegowski, Experimental evolution and the dynamics of genomic mutation rate
modifiers. *Heredity*. **113**, 375–380 (2014).
33. J. A. G. M. de Visser, C. W. Zeyl, P. J. Gerrish, J. L. Blanchard, R. E. Lenski, Diminishing returns
from mutation supply rate in asexual populations. *Science*. **283**, 404–406 (1999).
34. O. Tenaillon, B. Toupance, H. Le Nagard, F. Taddei, B. Godelle, Mutators, population size,
adaptive landscape and the adaptation of asexual populations of bacteria. *Genetics*. **152**, 485–493
(1999).
35. C. F. Gentile, S.-C. Yu, S. A. Serrano, P. J. Gerrish, P. D. Sniegowski, Competition between high-
and higher-mutating strains of *Escherichia coli*. *Biology Letters*. **7**, 422–424 (2010).
36. I. Matic *et al.*, Highly variable mutation rates in commensal and pathogenic *Escherichia coli*.
Science. **277**, 1833–1834 (1997).
37. L. Notley-McRobb, S. Seeto, T. Ferenci, Enrichment and elimination of mutY mutators in
Escherichia coli populations. *Genetics*. **162**, 1055–1062 (2002).
38. R. Hershberg, D. A. Petrov, Evidence that mutation is universally biased towards AT in bacteria.
PLoS Genetics. **6**, e1001115 (2010).
39. A. Stoltzfus, R. W. Norris, On the causes of evolutionary transition:transversion bias. *Molecular
Biology and Evolution*. **33**, 595–602 (2016).
40. E. Denamur *et al.*, Evolutionary implications of the frequent horizontal transfer of mismatch repair
genes. *Cell*. **103**, 711–721 (2000).
41. J. E. LeClerc, B. Li, W. L. Payne, T. A. Cebula, High mutation frequencies among *Escherichia coli*
and *Salmonella* pathogens. *Science*. **274**, 1208–1211 (1996).
42. T. Baba *et al.*, Construction of *Escherichia coli* K-12 in-frame, single-gene knockout mutants: the
Keio collection. *Molecular Systems Biology*. **2**, 2460–11 (2006).
43. L. C. Thomason, N. Costantino, D. L. Court, *E. coli genome manipulation by P1 transduction* (John
Wiley & Sons, Inc., Hoboken, NJ, USA, 2001), vol. 2.
44. H. Li, R. Durbin, Fast and accurate short read alignment with Burrows-Wheeler transform.
Bioinformatics. **25**, 1754–1760 (2009).
45. H. Li *et al.*, The Sequence Alignment/Map format and SAMtools. *Bioinformatics*. **25**, 2078–2079
(2009).
46. D. C. Koboldt *et al.*, VarScan: variant detection in massively parallel sequencing of individual and
pooled samples. *Bioinformatics*. **25**, 2283–2285 (2009).
47. N. F. Delaney *et al.*, Development of an optimized medium, strain and high-throughput culturing
methods for *Methylobacterium extorquens*. *PLoS ONE*. **8**, e62957 (2013).
48. S. R. Eddy, A new generation of homology search tools based on probabilistic inference. *Genome
Inform*. **23**, 205–211 (2009).
49. N. Segata, D. Börnigen, X. C. Morgan, C. Huttenhower, PhyloPhlAn is a new method for improved
phylogenetic and taxonomic placement of microbes. *Nat Commun*. **4**, 2304–11 (2013).
50. J. P. Bollback, SIMMAP: stochastic character mapping of discrete traits on phylogenies. *BMC
Bioinformatics*. **7**, 88 (2006).
51. L. J. Revell, phytools: an R package for phylogenetic comparative biology (and other things).
Methods in Ecology and Evolution. **3**, 217–223 (2011).
693