

Title: Changing fitness effects of mutations through long-term bacterial evolution

Authors: Alejandro Couce^{1,2,3*†}, Anurag Limdi^{4†}, Melanie Magnan¹, Siân V. Owen⁴, Cristina M. Herren^{4,5}, Richard E. Lenski^{6,7}, Olivier Tenaillon^{1,8*††} and Michael Baym^{4*††}

Affiliations:

¹Université Paris Cité and Université Sorbonne Paris Nord, Inserm, IAME, F-75018 Paris, France

²Department of Life Sciences, Imperial College London, SW7 2AZ London, UK

³Centro de Biotecnología y Genómica de Plantas, Universidad Politécnica de Madrid (UPM), 28223 Madrid, Spain

⁴Department of Biomedical Informatics, and Laboratory of Systems Pharmacology, Harvard Medical School, Boston, Massachusetts 02115, USA

⁵Department of Marine and Environmental Sciences, Northeastern University, Boston, Massachusetts 02115, USA

⁶Department of Microbiology and Molecular Genetics, Michigan State University, East Lansing, Michigan 48824, USA

⁷Program in Ecology, Evolution, and Behavior, Michigan State University, East Lansing, Michigan 48824, USA

⁸ Université Paris Cité, Inserm, Institut Cochin, F-75014 Paris, France

†These authors contributed equally to this work

†† These authors contributed equally to this work

*Corresponding authors: a.couce@upm.es, olivier.tenaillon@inserm.fr, baym@hms.harvard.edu


Abstract:

The distribution of fitness effects of new mutations shapes evolution, but it is challenging to observe how it changes as organisms adapt. Using *E. coli* lineages spanning 50,000 generations of evolution, we quantify the fitness effects of insertion mutations in every gene. Macroscopically, the fraction of deleterious mutations hardly changed over time, while the beneficial tail declined sharply, approaching an exponential distribution. Microscopically, changes in individual gene essentiality and deleterious effects often occurred in parallel; altered essentiality is only partly explained by structural variation. The identity and effect sizes of beneficial mutations changed rapidly over time, but many targets of selection remained predictable owing to the importance of loss-of-function mutations. Taken together, these results demonstrate the dynamic, but statistically predictable, nature of mutational fitness effects.

One-Sentence Summary:

Couce, Limdi, *et al.* report predictable and parallel changes in the fitness effects of mutations in replicate lineages of a multi-decades-long *E. coli* evolution experiment.

Summary Figure



Mapping the distribution of fitness effects over evolution. Transposon insertion mutagenesis of *E. coli* strains from the long-term evolution experiment and sequencing-based fitness assays enable characterization of macroscopic and microscopic properties of the distribution of fitness effects (DFE). The overall shape of the DFE is conserved, except for a declining beneficial tail, while the effects of specific mutations and even gene essentiality often change in parallel across replicate evolving populations. The ancestral DFE, combined with target size of genes, is predictive of drivers of adaptation. Schematic illustration created with BioRender ([biorender.com](https://www.biorender.com)).

One-Page Summary

Introduction

Evolution is constrained by the mutations accessible to natural selection. The benefits and costs of these mutations are described by the distribution of fitness effects (DFE). The DFE governs the tempo and mode of adaptation, and it reflects the mutational robustness of genotypes. However, the DFE does not remain static over evolution; with every accumulating mutation, the effects and accessibility of subsequent mutations can change through genetic interactions. Understanding how the DFE changes through adaptation is important for theoretical models that seek to describe the speed of adaptation, maintenance of genetic diversity, and pace of the molecular clock.

Rationale

Here, we quantified the effects of hundreds of thousands of insertion mutations in 12 populations of *Escherichia coli* through 50,000 generations of experimental evolution. We generated high-coverage transposon insertion libraries in the ancestral and evolved strains, and we measured the relative fitness effects of these mutations in sequencing-based bulk competitions. We characterized both the overall statistical properties of the DFEs and the effects of mutations in specific genes.

Results

Macroscopically, we observed little or no change in the deleterious tail of the DFE. By contrast, the fraction of beneficial mutations declined rapidly during early evolution, with its form approaching an exponentially distributed tail. We found no systematic changes in robustness to mutations, measured as the mean of the DFE. At a microscopic scale, we observed frequent changes in the fitness effects of insertion mutations in specific genes. Both the genetic identity and effect sizes of beneficial mutations changed over time. In the deleterious tail, there were frequent changes in the fitness costs of mutations and even in gene essentiality. These changes often evolved in parallel across independent lineages, and the changes in gene essentiality were only partially explained by structural variation. Despite pervasive changes in the fitness effects of particular mutations over time, many genetic targets of selection could still be predicted from the ancestral DFE, owing to the importance of loss-of-function mutations that scale with the length of the target genes.

Conclusions

Overall, we find that the macroscopic features of the fitness landscape remained largely unchanged over this multi-decade evolution experiment, except for truncation of the beneficial tail of the DFE. Over the short term, the genetic drivers of adaptation were often predictable from the microscopic details of the DFE, especially combined with the target size of genes harboring beneficial mutations. As the evolving populations accumulated more mutations over longer timescales, pervasive epistasis led to changes in the magnitude and even the sign of the fitness effects of many mutations. As a consequence, some evolutionary paths that were inaccessible to the ancestor became available to the evolving populations, while others were closed off. Moreover, many of the changes in the fitness effects of particular mutations, both beneficial and deleterious, occurred in parallel across the replicate populations. Thus, some microscopic features of the DFEs changed repeatedly and predictably over time, even as the macroscopic features of the fitness landscape remained largely unchanged.

Main Text:

Evolution in asexual populations is a local process because selection can only act on mutants generated from existing genotypes. Thus, information about the relative fitness of the genotypes that can arise in the mutational neighborhood of the current population is essential for predicting future evolution. The distribution of fitness effects (DFE) captures the properties of an organism's mutational neighborhood: the proportion and magnitude of beneficial mutations determines the tempo and mode of adaptation, whereas the fraction of neutral and deleterious mutations defines the organism's robustness to mutational perturbations. Indeed, the DFE lies at the core of many theories describing fundamental evolutionary phenomena, including the speed of adaptation (1), fitness decay in small populations (2), the maintenance of genetic variation (3), the probability of parallel (4) versus divergent (5) evolution, the pace of the molecular clock (6), and the evolution of sex (7) and mutation rates (8). However, it is unclear whether the general properties of local mutational neighborhood remain static over long periods of evolution because with each successive mutation in a lineage, the accessibility and effect of subsequent mutations can be altered through genetic interactions (i.e., epistasis) (9, 10).

The evolution of the overall shape of the DFE has received much theoretical and empirical attention. The beneficial tail of a DFE is expected to shorten as beneficial substitutions accumulate in an evolving population. Indeed, experiments with microbes show that the speed of adaptation steadily declines during adaptation to a constant environment (11–13), but it is generally unclear whether this deceleration reflects a decline in the availability or magnitude of new beneficial mutations (13). Besides becoming shorter, Extreme Value Theory predicts from simple statistical principles that the beneficial tail should become exponentially distributed as the population approaches a fitness peak (14, 15). While many studies support this model (16, 17), a few studies have reported non-exponential distributions of beneficial effects, and it is unclear whether these exceptions represent populations far away from their fitness peak or, alternatively, the inadequacy of the theory (18, 19). The picture is even more complicated for the deleterious tail of the DFE (20, 21). On the one hand, selection can favor mechanisms conferring increased robustness to mutational perturbations, especially at high mutation rates and in large populations (22–25), an idea with mixed support from studies with viruses and yeast (26–28). On the other hand, recent theoretical work suggests the genetic architecture of complex traits may lead to mutations being on average more detrimental on fitter genetic backgrounds (29), consistent with empirical data from crosses among diverse yeast strains (30).

However, these predictions address only the macroscopic form of the DFE, with little attention to the microscopic processes underlying changes in its overall shape. In the beneficial tail, the microscopic details may determine the extent to which adaptive pathways are predictable (31). For example, in the absence of interactions among mutations, adaptation will shorten the beneficial tail simply by the process of sampling without replacement, and therefore a complete DFE would suffice to specify the probabilities of all possible adaptive pathways in a given environment. By contrast, if each accumulated mutation changes the fitness effects and rank order of the remaining mutations (32), then predicting adaptive pathways would be impossible beyond the very short term.

In the deleterious tail, the microscopic details may reveal which physiological processes and genes are important or essential for fitness, and how those processes and genes might change over time. Further, those details may provide evidence bearing on whether changes in the deleterious tail are the product of natural selection acting directly on mutational robustness or, alternatively, a by-product of selection on related physiological processes (22). Moreover, the extreme of this tail contains the set of essential genes

whose loss would render the organism inviable. Prior work has shown that gene essentiality can vary greatly between species and even strains of the same species (33–35). For instance, about a third of the essential genes in *Escherichia coli* are nonessential in *Bacillus subtilis*, and vice versa (36). Essentiality is also malleable over shorter timescales: in *Saccharomyces cerevisiae* and *Staphylococcus aureus*, many essential genes become nonessential following selection for suppressors (37, 38), and horizontal gene transfer alters the essentiality of some core genes in *E. coli* (35). To what extent gene essentiality remains constant in the absence of direct selection, environmental change, or recombination is unclear. However, that issue has broad fundamental interest (e.g., understanding species’ ecological and geographic ranges) (39) and applied consequences (e.g., the quest for the “minimal genome”) (40).

Empirical studies of the DFE have generally been either small in scale (41) or focused on narrow genomic regions (42), and they typically lack detailed information on the population’s level of adaptation to the test environment. Consequently, it has been difficult to distinguish among competing hypotheses about the evolution of the DFE during the course of adaptation, both at the macroscopic and microscopic levels. To address this challenge, one would ideally like to measure the relative fitness of the complete set of genome-wide mutants at multiple time points along a well-characterized adaptive trajectory. To do so, we turned to the Long-Term Evolution Experiment (LTEE), in which twelve populations of *E. coli* have been serially propagated in a glucose-limited minimal medium (43) for over 75,000 generations.

To quantify changes in the DFE, we generated genome-wide transposon insertion libraries in strains isolated at several time points from the LTEE, and we measured relative fitness values using high-resolution, bulk competitions. Such insertions typically lead to losses of function; by their nature, spontaneous loss-of-function mutations occur readily, and so our approach surveys a large (but not complete) portion of the fitness landscape accessible by single-step mutations. Of note, we also observed two types of more subtle effects. First, an insertion in the C-terminus of a gene may cause only a partial loss of function or even a change in function (44). We observed a number of striking examples of this outcome, including insertions in this region that were not merely tolerated but in fact conferred large fitness benefits (Fig. S1). Second, the positions of many beneficial insertions, including in intergenic regions and genes upstream of known targets of adaptation in the LTEE, suggest effects on gene expression (Fig. S1).

Our experimental system covers a large fitness gradient (>70% gains) that was generated by selection of spontaneous mutations in a constant environment, with no horizontal gene transfer (11). It is therefore well-suited for detecting evolutionary trends in mutational robustness and the size of the essential gene set. Moreover, the most important mutations driving adaptation have been identified from signatures of parallelism in whole-genome sequences (45, 46), allowing predictions based on the DFE at one time point to be compared with the actual fate of mutations observed during later evolution. Lastly, by comparing patterns in changing fitness effects across multiple independently evolving lineages, we can characterize the extent to which changes in the DFE are idiosyncratic or convergent.

Results

High-throughput insertion mutagenesis and fitness measurements

We performed two sets of experiments that analyzed the DFEs of many clones from the LTEE. In one experiment, we focused on changes in mutational robustness and gene essentiality during evolution. To do so, we constructed high-coverage transposon libraries in the LTEE ancestors (REL606 and REL607) and a clone isolated from each population (Ara+1 to Ara+6 and Ara–1 to Ara–6) at 50,000 generations.

In the other experiment, we focused on the early, rapid change in the properties of the beneficial tail. To that end, we made transposon libraries in the ancestor and clones sampled at 2,000 and 15,000 generations from two populations (Ara+2 and Ara-1), when fitness had increased by ~25% and ~50%, respectively (11). In both experiments, we obtained >100,000 unique insertions, disrupting >78% of the genes, with >95% overlap in genes disrupted in the ancestral and evolved libraries (Fig. S1).

We estimated the fitness effects of all these mutants as selection coefficients, which we calculated from the frequency trajectories of every allele based on high-throughput sequencing during bulk competition assays under the same conditions as in the LTEE (Fig. 1, Fig. 2A; Methods). This sequencing-based approach resolves the identity of each mutant at the molecular level; it allows us to interrogate both overall trends and the microscopic details of the locally accessible mutational landscape. We inferred fitness effects relative to a set of reference mutations, which consisted of insertions in known or presumed neutral loci, in the same transposon library (Fig. 1D, Methods). This approach allows relative fitness effects to be compared across the LTEE strains. The resulting fitness estimates were highly reproducible between technical replicates and consistent with independent estimates obtained from pairwise competitions between engineered deletion mutants and their unmutated parents (Fig. S3, Data S1).

No systematic changes in the overall shape of the DFE

To investigate whether the overall form changed over time, we first compared the DFEs of the two LTEE ancestors and a clone from each of the 12 populations evolved independently for 50,000 generations. We excluded two evolved samples from further analyses because their fitness measurements were unreliable for technical reasons and therefore not comparable to the ancestor. In Ara+4, the within-gene measurement variability for fitness was extremely high, and the correlation between technical replicates was poor (Fig. S4A). In Ara-2, a few insertion mutations increased rapidly and outcompeted other mutations (Fig. S4B-C), which made the measurements unreliable and systematically biased (see Supplementary Text 1 for more details). Overall, most mutations are nearly neutral (within ~2-3% of neutrality, depending on the strain), but in all cases with a much heavier tail of deleterious mutations than beneficial mutations (Fig. 2B), consistent with previous results (26–28). The aggregate DFEs for the ancestors and evolved lines were nearly identical, except for an excess of mutations that are beneficial ($s > 0.03$, an effect reliably distinguishable from measurement noise) in the ancestral over the evolved backgrounds (0.9% vs 0.5% of all mutations, respectively; Fig. 2C, note the logarithmic scaling). This difference in the supply of beneficial mutations and its evolutionary significance are examined in depth in our second experiment (see below).

There was no systematic directional trend in how the means of the DFEs changed during evolution (t -test based on population means: $p = 0.37$). While the mean fitness effect differed significantly between the ancestor and several evolved lines considered individually (Fig. 2D), these differences vary in their direction (two evolved clones had higher means than the ancestor, and three had lower means), and they are primarily driven by noisy measurements in the deleterious tail (Fig. S5). Therefore, robustness measured as the overall mean of the DFE of insertion mutations did not systematically change during the 50,000 generations of adaptation.

The constancy of the deleterious tail we observe over time stands in contrast to a study that measured the DFE for 91 insertion mutations in hybrid yeast genotypes with fitness values spanning ~20%, in which deleterious effects were significantly worse in the more-fit backgrounds (30). A potentially important difference is that the fitness variation among the yeast backgrounds was generated by crossing two distantly related strains, whereas we use a series of backgrounds from lineages undergoing adaptation to

the same environment in which we assessed the fitness effects of the new mutations. In any case, theoretical predictions about the tail of deleterious mutations differ substantially and have been guided mostly by plausibility arguments (20, 21), and so these studies collectively should help refine current models by clarifying their assumptions and narrowing the range of parameters.

Parallel changes in fitness effects over evolution

A conserved macroscopic distribution does not preclude microscopic changes in the effects of individual mutations. Therefore, we examined whether and how the fitness effects of the same insertion mutations varied between the ancestor and evolved strains. We restricted this analysis to insertions with fitness effects $s > -0.3$ in both the ancestor and evolved strain, as measurements of extremely deleterious effects have more measurement noise. The fitness effects of many mutations differed between the ancestral and evolved strains, with some becoming more deleterious and others less deleterious (Fig. 3A). Depending on the evolved strain, between 3% and 6% of the mutations had significantly different fitness effects from those in the ancestor (Fig. 3B), and 13% had differential effects in at least one evolved strain.

We observed significant parallelism across the independent lineages in the genes with fitness effects that changed significantly over evolution. We first examined this possibility through hierarchical clustering of mutations that were roughly neutral in the ancestor ($s > -0.05$) and clearly deleterious in an evolved strain ($-0.3 < s < -0.15$), and vice versa (Fig. 3C and D). While many such changes were specific to individual lineages, many others occurred in parallel across multiple lineages. To assess whether the observed parallelism was greater than expected by chance, we compared the two complementary cumulative distributions of differential effects of gene disruptions in multiple lineages against a null distribution, which we generated by shuffling the fitness profiles of each population 10,000 times. Both the neutral-to-deleterious and deleterious-to-neutral transitions occurred in parallel much more often than expected by chance (Fig. 3E). This outcome was insensitive to the chosen cutoff values (Fig. S6). These parallel changes across independent lineages suggest that selection acted, directly or indirectly, to influence those changes.

Parallel changes in gene essentiality over evolution

Moving toward the extreme deleterious tail, we next investigated gene essentiality. Strict lethality or an absolute inability to replicate is often difficult to distinguish from extreme growth defects. For this analysis, we therefore define a gene as differentially essential between the ancestor and an evolved clone if (i) the fitness effect of disruption $s > -0.15$ in one strain and $s < -0.3$ in the other, or (ii) mutants were absent in the library prior to the bulk competition in the LTEE medium DM25, suggesting that the gene was essential in LB (Supplementary Materials). This approach ensured that small changes in fitness effects (say from -0.31 to -0.29) were not counted as changes in essentiality. Also, our choice of $s < -0.3$ emerged from simulated competitions, which indicated that mutations with deleterious effects of this magnitude or larger could not be reliably distinguished from lethality (Fig. S7). Using the cutoff $s < -0.3$, we detected 557 genes that were essential in DM25 in the ancestor (Supplementary Materials).

We found genes that went from nonessential to essential, and vice versa, in all the LTEE lines (Fig. 4A, Data S2). We confirmed two examples of differential gene essentiality in DM25 using clean deletion mutants in the ancestor REL606 and Ara-1 (Fig. S8, Data S1). In total, 77 nonessential genes became essential in at least one evolved lineage, and 97 essential genes became nonessential in at least one lineage, corresponding to ~17% of the essential genes in the ancestor. However, many more genes became nonessential in Ara-6 than in the other evolved lines (Fig. 4C), as a result of gene duplications discussed below. If we exclude Ara-6, then the nonessential-to-essential transition is more common.

Indeed, across the other LTEE populations, we observed a significant tendency for more nonessential genes to become essential than the reverse change ($p = 0.0008$, Mann-Whitney U test). This asymmetry suggests that mutational robustness in terms of gene essentiality typically decreased during the LTEE. Both the essential-to-nonessential and nonessential-to-essential transitions occurred in parallel much more often than expected by chance (Fig. 4D). This outcome was insensitive to the exact cutoff values for essentiality (Fig. S9), and it persisted when we partition essentiality changes by the culture medium (Fig. S10). This parallel evolution in gene essentiality again implies that these changes result from selection. It is unclear how selection would act directly on essentiality; instead, this parallelism is presumably a correlated response to selection on gene expression or other metabolic traits.

Gene essentiality has previously been associated with highly expressed genes (47–49). We therefore examined whether changes in gene essentiality were associated with altered expression levels. We used a recently published RNA-Seq dataset for the LTEE ancestor and evolved strains at 50,000 generations (50). Consistent with previous findings, essential genes have higher expression levels on average than nonessential genes (Fig. S11A). However, for those genes that became essential or nonessential during the LTEE, we find no significant differences in the normalized expression levels in the ancestor and evolved strains (Fig. S11B). This result implies that changes in essentiality are not generally related to altered levels of gene expression.

Changes in gene essentiality could also arise as by-products of other mutations, especially losses or gains of other gene functions. Gene duplications can give rise to robustness by providing functional redundancy (51), while deletions can increase the essentiality of other genes by eliminating existing redundancies. We examined these possibilities by sequencing the ancestors and 50,000-generation clones with high coverage (>60X) to identify all large deletions and duplications in the evolved genomes. We then asked whether changes in gene essentiality were associated with these structural variants and their potential effects on redundancy given homologs in the ancestral genome (Data S3). We found some cases where structural variants were associated with changes in gene essentiality. These cases included parallel deletions in most lineages that spanned the *rfb* operon and caused insertions in some paralogs to become essential in the evolved clones (Figure S12A). For the vast majority of newly essential genes, however, we found no evidence that essentiality was caused by loss of redundant genes. With respect to duplications, the genome from population Ara-6 has two large duplications spanning ~300 and ~25 genes (Figure S12B). Ara-6 alone accounts for the majority of transitions from essential-to-nonessential genes, and most of those transitions are found in the duplicated regions (Figure S12C). Further details and analyses are provided in the Supplement (see “Gains and losses of functional redundancy explain some, but not most, changes in essentiality”).

Rapid contraction of the beneficial tail of the DFE

Our first experiment showed substantial changes in the small, but critically important, beneficial tail of the DFE. We therefore conducted additional experiments focused specifically on this tail and how it changed over evolution. Half of the ~70% gain typically seen at 50,000 generations of the LTEE had already occurred by 5,000 generations (11). We decided therefore to create transposon libraries in clones sampled at 2000 (“2K”) and 15,000 (“15K”) generations, when fitness had increased by ~25% and ~50%, respectively. To increase our resolution near selective neutrality, we divided each locus into 5 segments of equal length and then pooled the insertions within each segment. This approach expands the range of potentially observable beneficial mutations by enabling detection of polar effects within transcription units, effects linked to regulatory intergenic regions, and potentially subtle effects of insertions in the C-termini of protein-coding genes (Fig. S2). As an added benefit, comparing the fitness effects among

segments of the same locus helps identify potential artifacts and provides a within-experiment control to quantify the reproducibility of the fitness estimates (Supplementary Materials, Fig. S13).

We first focused on samples obtained from population Ara+2. Figure 5 (A-C) shows that the fraction of beneficial insertion mutations is substantially larger in the ancestor than in the evolved backgrounds (6.8% for ancestor [Anc] versus 4.3% and 3.2% for 2K and 15K, respectively; $p < 0.044$ both cases, two-sample Kolmogorov–Smirnov [K-S] test). In contrast, and in agreement with what we observed for the 50,000-generation clones, the deleterious fraction is essentially constant across the three backgrounds (20.5% for Anc versus 21.0% and 19.6% for 2K and 15K, respectively; $p > 0.076$ both cases, two-sample K-S test). These patterns are consistent with analyses at the level of individual genes for both beneficial and deleterious mutations (Fig. 5D, Fig. S14). To examine whether these results depend on the particular lineage, we also measured the DFEs for clones sampled at 2000 and 15,000 generations from population Ara–1, which accumulated a different set of beneficial mutations along its independent adaptive trajectory (Supplementary Materials, Table S1). At least two major features distinguish the evolutionary history of this lineage from that of Ara+2. First, Ara–1 fixed a mutation in *topA* early in the LTEE. Mutations in this gene confer among the largest fitness benefits seen in the LTEE for any single substitution (52); they fixed in five of the twelve populations, but they never reached detectable frequency in Ara+2. Second, Ara–1 evolved a mutator phenotype, whereas Ara+2 retained the low ancestral mutation rate throughout the experiment; however, Ara–1 became hypermutable only after ~21,000 generations, and hence poses no added complications to our analysis of the evolved clones from earlier generations. Despite independent histories, we obtained strikingly similar results for these two lineages, at both the macroscopic and microscopic levels (Fig. S15). Our findings demonstrate that the contraction of the beneficial tail of the DFE occurred early and quickly as adaptation proceeded. Specifically, the small number of beneficial mutations that accumulated during the first 2000 generations were sufficient to have a significant impact on the adaptive landscape of the evolving population.

An exponential tail of beneficial mutations emerged during adaptation

Extreme Value Theory predicts on statistical grounds that the effects of beneficial mutations should be exponentially distributed when a population is well-adapted to its environment (1, 14). Despite some empirical support (16, 17), the evidence remains inconclusive owing to a severe limitation of most studies: without detailed knowledge of a population’s evolutionary history, it is difficult to characterize its level of adaptation to a particular environment. Our data, by contrast, provides a uniquely powerful test of these ideas. We find that, indeed, beneficial mutations in the evolved genetic backgrounds are well fit by an exponential distribution, whereas this distribution is decisively rejected for the ancestor ($P < 0.001$ for Anc versus $p = 0.571$ and $p = 0.852$ for Ara+2 clones 2K and 15K, respectively; one-sample K-S test). We also considered alternative distributions, but the exponential provides the best fit for the evolved backgrounds (Supplementary Materials, Table S2). Note that the exponential distribution is a special case of both the Weibull and gamma distributions, so it is not surprising that the data also fit well to them. These two distributions can be thought of as natural transitional shapes before reaching the limiting case of the exponential distribution. Indeed, the beneficial tail for the ancestor was fit to different degrees by both gamma and Weibull distributions ($p = 0.035$ and $p = 0.29$, respectively; one-sample K-S test), consistent with previous studies of viral and bacterial genotypes thought to be poorly adapted to their test environments (18, 19). Overall, our results support the view that, after an early period of rapid adaptation to a new environment, the distribution of beneficial mutations becomes exponential. Thus, by analyzing changes in the DFE in a temporal series of genetic backgrounds becoming better adapted to their environment, we have reconciled otherwise disparate pieces of evidence relevant to general models of adaptation.

Changing identity of beneficial mutations and sign epistasis

We next sought to understand how changes in the DFE's macroscopic structure emerged from changes at the level of genes and mutations. We found that during the early phase of adaptation, deleterious mutations typically exhibit only slight epistasis across the three focal genetic backgrounds of the Ara+2 lineage (Fig. S14). That is, the magnitude of their harmful effects may vary, but deleterious mutations in the ancestor tend to remain deleterious in the evolved backgrounds, consistent with the observed constancy of the deleterious tail (see Fig. S16 for more details). In striking contrast, beneficial mutations are dominated by strong sign-epistatic interactions (Fig. 5D). Only 5.9% of the mutations beneficial in the ancestor are still beneficial at 2,000 generations, with most becoming effectively neutral (76.9%) and some deleterious (17.2%) (Fig. 5E at left). This pattern also holds in the reverse direction: most beneficial mutations at 2,000 generations are neutral (72.5%) or deleterious (17.9%) in the ancestor (Fig. 5E at left). Similar patterns occur when comparing how fitness effects changed between 2,000 and 15,000 generations (Fig. 5E at right). Given the transitory nature of beneficial effects, we asked whether the overall DFE of the initially beneficial mutations retains even a slightly positive tendency at the later time points. In fact, it does not. The DFE of mutations that were beneficial in the ancestor becomes indistinguishable from a random sample of the parent distribution (Fig. 5F at left), and the same holds for the reverse scenario (Fig. 5F at right) ($p > 0.085$ both cases; two-sample K-S test). This regression to the mean persists even when we account for measurement noise around neutrality (Fig. S13C and D).

What explains this turnover in the identity of the beneficial mutations? In a previous study, the first five mutations to fix in one LTEE population were shown to exhibit diminishing-returns epistasis, such that their benefits declined in magnitude as the background fitness increased (52). However, it was unlikely *a priori* that these five mutations would show sign epistasis because they were chosen precisely because their combination was favored by natural selection (53). By contrast, another study analyzed the co-occurrence of fixed mutations across 115 lines of *E. coli* that evolved under thermal stress, and it found that sign epistasis was common (54). Moreover, that study found that the prevalence of different types of epistasis reflected the modular architecture of cellular traits: mutations affecting different modules tended to have additive effects, while those impacting the same module tended to be redundant. We therefore investigated the extent of modularity in our data, and we found that beneficial mutations in the ancestral background often occurred repeatedly in the same operons (Supplementary Materials, $P < 0.01$). Mutations in the same operon typically alter the same cellular process and often in similar ways, and therefore the potential for redundancy at this functional level provides a simple explanation for why large sets of beneficial mutations disappear, and other sets emerge, as adaptation proceeds. More generally, the increased prevalence of sign epistasis with adaptation has also been predicted from general properties of the genotype-to-fitness map (55).

Target size is an important predictor of the genes that accumulate beneficial mutations

We identified a large set of loci that can produce beneficial mutations, including some known targets for adaptation in the LTEE (e.g., *topA*, *pykF*, *nadR*) (45). However, the fate of beneficial mutations in the course of evolution is determined not only by their individual fitness effects, but also by their occurrence rate and the nature of their interactions with other beneficial mutations (31, 32, 56, 57). As a consequence, only a fraction of all possible beneficial mutations will contribute to adaptation in an evolving population. To gain further insight into this issue, we compared our data with metagenomic data previously obtained by sequencing whole-population samples from the 12 LTEE populations over the course of 60,000 generations (46). We see a significant, but fairly weak, correlation between our fitness estimates for mutations in the ancestor and the abundance of corresponding alleles during the LTEE ($r = 0.26$, Fig.

6A), and this correlation largely disappears when using the beneficial effects estimated in the evolved backgrounds. In contrast, the abundance of alleles in the metagenomic data correlates more strongly with the target size of the locus ($r = 0.71$, Fig. 6B-C) (Supplementary Materials). These patterns are consistent with intense competition among independently arising beneficial mutations (i.e., clonal interference), a pervasive phenomenon in the LTEE (46, 58). Under intense clonal interference, the rate at which particular beneficial mutations occur may shape genomic evolution even more than their fitness effects (59). In any case, the best linear model includes target size as the most explanatory single variable, but it also includes significant contributions from the fitness effects in both the ancestral and 2,000-generation genetic backgrounds (Fig. 6C, Table S3). Finally, we note that a potentially important factor contributing to the observed weak correlations is that our methods involve insertion mutations, which usually, but not always (Fig. S2), cause losses of function. While losses of unused functions have contributed to adaptation in the LTEE (45, 60), subtle changes that typically require point mutations have also been important in refining some functions (31, 45, 61).

Predicting future beneficial mutations as adaptation proceeds

Given that sign epistasis is widespread, it is natural to ask for how long the information about the particular loci in the beneficial tail of a DFE can successfully predict the subsequent steps of adaptation. To address this question, we used the metagenomic data to record the alleles nearing fixation through time, and we calculated how many corresponded to loci for which we detected beneficial effects. We found that the ancestral DFE predicted most of the loci where mutations became dominant early in the LTEE populations; the predictive power decays rapidly, but it was still evident for ~15,000 generations (Fig. 6D). This decay was largely driven by lineages that evolved hypermutability early in the LTEE; when these mutator populations are excluded from the analysis, the ancestral DFE retained significant predictive power through 50,000 generations (Fig. S17A). In turn, the DFEs measured in the evolved backgrounds had lower predictive power, and it took longer for their predictions to materialize; the latter effect may reflect the declining rate of adaptation. These patterns corroborate work showing that parallel genomic evolution was more common early in the LTEE than in later generations (45, 62).

Finally, why does the ancestral DFE have such predictive power, when it is based on insertion mutations that represent only a limited set of all possible mutations from a functional standpoint? To address this question, we quantified how many loci with frequent beneficial mutations in the LTEE include mutations with presumed loss-of-function effects. To that end, we assumed that nonsense, frameshift, deletions, and insertions cause losses of function. We find that these presumptive inactivating mutations contribute most (>50%) of the early adaptive mutations in the LTEE, and they continue to be a sizable fraction over the long run (~25%, Fig. S17B). Of note, another study with *Methylobacterium extorquens* adapted to use methanol as the sole carbon source also found that most early beneficial mutations appear to disrupt functions (63). These results support the “coupon-collecting” model of rapid evolution (46, 54), in which “rough-and-ready” loss-of-function mutations dominate the early phase of adaptation to a new environment owing simply to their high rates of occurrence. Under this model, many initially beneficial mutations also become redundant because they inactivate the same functional module. This model implies that fitness effects alone are inadequate for predicting adaptive fixations, but taking target size into account compensates for this uncertainty. This interpretation satisfactorily explains our findings that the initial drivers of adaptation are predictable despite widespread and strong epistasis, and that target size is an important predictor of beneficial alleles that fix early when a population encounters a new environment.

Conclusions and Discussion

This paper began as two separate projects, performed by two different teams, using similar but not identical methods. As we discussed our findings together, we discovered that each project both reinforced and complemented the other. They reinforce one another by finding the same evolution of the overall form of the DFE; they are complementary because one project delved deeply into the fine-scale genetic changes in the deleterious tail, and the other did so for the beneficial tail. Thus, together we have characterized changes in the DFE over the course of long-term evolution in a new environment with unprecedented resolution, including both the distribution's overall form and the effects of specific mutations. At a macroscopic scale, the idiosyncratic shape of the beneficial tail of the DFE became truncated, leading to an exponential distribution as predicted by some models (14, 15). By contrast, there was no discernible change in the deleterious tail of the DFE, and mutational robustness, measured as the mean of the DFE across the replicate populations, was also unchanged over adaptation, suggesting that robustness was not under strong directional selection. With the notable exception of a population that evolved large duplications encompassing many genes, we observed a tendency for more genes to become essential than nonessential, lending some support to the “increasing costs” model of epistasis (29, 30), but this effect disappeared when we examined the entire DFE. Overall, our results paint a complex picture of changing fitness effects that no simple model adequately captures.

At a microscopic scale, we found frequent changes in the fitness effects of particular mutations, even as the overall statistical properties of the DFE remained nearly constant. In the deleterious tail, there were frequent shifts in the effects of specific mutations (~13% of those in nonessential genes) over 50,000 generations, with some mutations becoming more deleterious and others less so. Similarly, we also observed frequent changes in the identity of beneficial mutations over time, even over just 2000 generations. This dynamic pattern implies that the beneficial tail of the DFE is continually replenished by new and functionally different mutations as adaptation proceeds, even as other mutations lose their advantage. This shifting set of beneficial mutations over time helps to explain the sustained gains in fitness observed over tens of thousands of generations in the LTEE .

Prior work has shown that gene essentiality is not a static property of a species, but the rate at which it changes is unknown (33, 37, 38). Here we show that ~3% of the genome had altered essentiality, which is similar to the variation in essentiality across diverse strains of *E. coli* when tested in three environments and often involving horizontally transferred genes (35). In contrast, the changes in gene essentiality that we observe in the LTEE happened over a much shorter evolutionary timescale, in the absence of any horizontal transfer, and without applying direct selection to suppress or enhance essentiality. Our demonstration of the fluid nature of essentiality indicates that the foundation of a minimal autonomous genome should not rely on a static snapshot of essentiality, because deleting genes can impact the potential for further genome reductions.

The ability to predict evolution remains elusive, in part because it requires a deep understanding of fitness landscapes and how they change. We found that the beneficial tail of the ancestral DFE is strongly predictive of the actual targets of selection in the LTEE, as inferred from the mutations nearing fixation in metagenomic data, particularly during early adaptation. This predictability reflects the prominent role that loss-of-function mutations had early in the LTEE, which seems applicable to other model systems (56, 57, 63). Over the long-term, however, pervasive epistasis resulted in declining predictability of these driver mutations, as the fitness effects of many mutations changed in magnitude and even their sign. Consequently, evolutionary paths that were inaccessible to the ancestor became available, while others were closed off, as reported recently in protein evolution (64). Because natural selection has steered most

of the LTEE populations along similar trajectories, the paths that open or close are often the same across independently evolving lineages. Taken together, our results demonstrate the dynamic, but statistically predictable, nature of mutational fitness effects; they show that some features of evolutionary trajectories change repeatedly and predictably over time, even as the macroscopic features of the fitness landscape remain largely unchanged.

Methods

We used two suicide-plasmid delivery systems to construct the transposon libraries in the ancestor and several evolved clones from the *E. coli* long-term evolution experiment (LTEE) (Table S4). We then passaged the transposon libraries in DM25, the medium in which the populations have evolved (43), for 4 to 8 days, and we then isolated genomic DNA from the pool of mutants. In the first set of experiments discussed in the main text, we followed an approach we refer to as UMI-TnSeq that uses the mariner transposon carried by the pSC189 plasmid (65, 66). We used this method to disrupt all genes in 50,000-generation clonal isolates from all 12 LTEE populations. The genomic regions adjacent to the insertion site were captured using a tagmentation-based approach. To control for potential PCR bias, we attached unique molecular identifiers (UMIs) to individual molecules during PCR amplification (see Detailed Experimental Protocols in Supplementary Materials). In the second set of experiments, we used the INSeq methodology (67), focusing on 2000- and 15,000-generation clones from two LTEE populations, called Ara+2 and Ara-1. We chose these populations because they neither evolved hypermutability nor diversified into stably coexisting lineages during the first 15,000 generations. Many other LTEE populations evolved one or both of these features, which would complicate testing our hypotheses (46).

After estimating the frequency of insertion mutants in the transposon libraries from bulk sequencing over the course of the fitness assays, we estimated the relative fitness of each mutant using linear regression of $\ln(\text{frequency})$ of each mutant against the number of generations of selection during the assay. In the UMI-TnSeq analysis, we calculated the fitness effects of disrupting a given gene by averaging over all insertion sites in its interior (excluding the initial 10% and final 25% of the gene). In the INSeq analysis, we calculated fitness effects at the level of sub-genic regions by dividing each locus into five equally sized segments, while requiring a minimum size of 100 bp per segment.

There are two main differences between the UMI-TnSeq and INSeq approaches. First, polar effects within transcription units are expected to be more accentuated with the INSeq approach, because the 1.5-Kb insert carries two transcriptional terminators after the kanamycin resistance gene. The second difference concerns how regions adjacent to the insertion site are identified. The INSeq transposon encodes recognition sequences for the restriction enzyme MmeI, which cuts 20 bp away from its binding site and thus allows the capture of the 14 bp adjacent to the insertion site (Fig. S18). This approach should minimize PCR bias because the genomic fragments are of uniform length, thus reducing the need to add UMIs during PCR. We also performed a replicate experiment with the Ara+2 samples to show that applying the UMI-TnSeq methodology to the INSeq transposon libraries yields essentially the same results (Fig. S19).

Acknowledgments

We thank Milo Johnson, Michael Desai, Andrew Murray, Thomas Bernhardt, Célia Souque, Fernando Rossine, and Daniel Eaton for valuable feedback and discussions; and Adrien Launay, Alexandra Baron, Romain Fernandes, Damien Roux, Thao Truong, Joel Sher, Tanush Jagdish, Karel Brinda, and Natalia Quiñones-Olvera for technical assistance. We also thank the Bauer Core Facility at Harvard University

and the Service of Biochemistry at Bichat Hospital for providing some of the sequencing, and the O2 cluster from the Research Computing Group at Harvard Medical School for supporting some of the computational work.

Funding

M.B. acknowledges support from the NIGMS of the National Institutes of Health (R35GM133700), the David and Lucile Packard Foundation, the Pew Charitable Trusts, and the Alfred P. Sloan Foundation. This work was partially supported by the European Commission under the 7th Framework Program (ERC Grant 310944 to O.T.), *Foundation pour le Recherche Médicale* (EQU201903007848 to O.T.), *Agence Nationale pour la Recherche* ANR GeWiEp (ANR-18-CE35-0005-0 to O.T.), and the Horizon 2020 Framework Programme (MSCA-IF 750129 to A.C.). The work was also partly supported by the *Ministerio de Ciencia, Innovación y Universidades* (PID2019-110992GA-I00 to A.C.), and a *Comunidad de Madrid “Talento”* Fellowship (2019-T1/BIO-12882 to A.C.). A.L. acknowledges support from the Molecules, Cells, and Organisms Graduate Program, Harvard University. R.E.L. acknowledges support from the US National Science Foundation (DEB-1951307) and the John Hannah endowment at Michigan State University.

Competing interests

The authors declare no competing interests.

Data and Code Availability

Raw sequencing reads are available from the NCBI BioProject database (PRJNA814281 and PRJNA979973). Processed data are available from Zenodo (10.5281/zenodo.6547536 and 10.5281/zenodo.7985455); source code for the sequencing pipeline, downstream analyses, and figure generation are available from GitHub (https://github.com/baymlab/2022_Limdi-TnSeq-LTEE and <https://github.com/ACouce/LTEE2022>).

Author Contributions

AC, AL, OT, MB conceived the project; AC, AL, MM, SVO designed and conducted experiments; AC, AL, CMH, REL, OT, MB designed statistical analyses; AC and AL analyzed the data; REL directs the LTEE and provided strains and critical feedback on interpretation; all authors wrote and revised the manuscript.




Fig. 1. Schematic representation of mutagenesis and fitness assay pipeline. (A) The Long-Term Evolution Experiment (LTEE) is an ongoing experiment in which 12 populations of *E. coli* evolve in and adapt to a glucose-limited minimal medium. (B) Here, we created transposon libraries in the LTEE ancestors and clones from the evolving populations. We transferred a mariner transposon with a *kanR* resistance gene, and we selected transconjugants on medium containing kanamycin. (C) We then propagated the resulting insertion libraries for several days in the same minimal medium as used in the LTEE, and we quantified the abundance of mutants over time using sequence data. (D) The abundance trajectories of a set of neutral loci were used to normalize coverage depth across time points, providing an internal reference to estimate selection coefficients of mutations (left). The fitness effects for these neutral loci were closely centered around zero (right). Panels (A,B) created with biorender.com.




Fig. 2. The overall distribution of fitness effects (DFE) is largely unchanged after 50,000 generations. (A) Frequency trajectories of the whole mutant library in the ancestor (left), and mapping of estimated fitness effects along the chromosome (right). Colors indicate fitness effects, from deleterious (red) to beneficial (blue). (B) Ridge plot of the overall DFE in the two LTEE ancestors (Anc: REL606, Anc*: REL607) (gray), which differ by a neutral marker, and 50,000-generation clones sampled from each population (blue). We excluded two strains (Ara-2 and Ara+4) from further analyses (see text; Fig. S4). The histograms were smoothed using kernel density estimation and are shown with a linear y-axis. DFEs are only shown for fitness effects ranging from -0.1 to 0.05 , as the density outside these regions is very low. (C) Comparison of the aggregated DFEs of the ancestral and evolved strains. Here the histograms are plotted with a logarithmic y-axis to show more clearly the deleterious and beneficial tails of the DFEs. (D) Means of the DFEs: error bars indicate the 95% confidence interval in the estimate of means given the associated measurement noise in the bulk fitness assays. Statistically significant differences between the evolved lines and ancestors after Bonferroni correction for multiple tests are indicated (Z-test; *** $p < 0.001$, ** $0.001 < p < 0.005$, * $0.005 < p < 0.05$).






A**B****E****C** Neutral → Deleterious**D** Deleterious → Neutral

Fig. 3. Extensive and parallel changes in fitness effects of insertion mutations over evolution. (A) Pairwise comparison of fitness effects of mutations in nonessential genes ($s > -0.3$) between the ancestor (REL606) and each evolved strain. Purple: more deleterious in the ancestor, green: more deleterious in the evolved strain, Bonferroni corrected p value < 0.05 (two-tailed Z-test). (B) Fraction of mutations (with $s > -0.3$ in both the ancestor and the evolved strain) with significant differences in fitness effects between the ancestor and each evolved clone (Bonferroni corrected p -values < 0.05). (C, D) Clustered heatmaps showing fitness effects (scale at right) of gene disruptions that became roughly neutral ($s > -0.05$) or clearly deleterious ($-0.3 < s < -0.15$) in at least one 50,000-generation strain. Genes that were deleted during evolution are shown in white. Genes with mutations conferring fitness effects below -0.3 (the threshold for essentiality) were set to -0.3 for the clustering and visualization. (E) Parallel changes in fitness effects. We estimated the expected number of parallel changes from chance alone by shuffling the profile of changes in fitness effects 10,000 times and counting how often the same genes had parallel changes (neutral to deleterious or deleterious to neutral) in at least m populations. The expectation is an average over 10,000 simulations, and therefore it can be < 1 .




Fig. 4. Extensive and parallel changes in gene essentiality over evolution. (A) Number of genes that are differentially essential between the ancestor and each evolved strain. (B, C) Clustered heatmaps showing fitness effects (scale at right) of genes that evolved to become essential or nonessential in at least one 50,000-generation strain. Genes that were deleted during evolution are shown in white. Genes with mutations conferring fitness effects below -0.3 (the threshold for essentiality) were set to -0.3 for the clustering and visualization. (D) Parallel changes in gene essentiality. We estimated the expected number of parallel changes from chance alone by shuffling the profiles of changes in gene essentiality 10,000 times and counting how often the same genes had altered essentially in at least m populations. The expectation is an average over 10,000 simulations, and therefore it can be < 1 .




Fig. 5. Rapid contraction of the beneficial fraction over the first 15,000 generations. (A) DFEs in the ancestor (black), 2K (blue), and 15K (red) backgrounds from population Ara+2. Note that the logarithmic scaling of the y-axis exaggerates minor, nonsignificant differences in the extreme deleterious tails. (B) Only the beneficial tails underwent substantial changes during evolution, as indicated by comparing the cumulative fitness distributions for the ancestor and 2K evolved strain (left), and for the ancestor and 15K strain (right). Shaded areas show 90% bootstrapped confidence intervals. (C) Beneficial tails rapidly became exponentially distributed. Histograms show the best fits to exponential distributions (dashed lines) in the ancestor (gray), 2K (blue), and 15K (red) backgrounds. Note that all three x-axes use the same scale. (D) The genes and intergenic regions with the most beneficial alleles in

the ancestral background, and their fitness effects in the 2K (blue) and 15K (red) backgrounds. **(E)** Most of the beneficial mutations available to the ancestor became neutral or deleterious in the 2K background (black arrows), while most beneficial mutations available in the 2K background were neutral or deleterious in the ancestor (red arrows). The same general pattern occurs when comparing beneficial mutations in the 2K and 15K backgrounds (right panel). **(F)** More than 90% of initially beneficial mutations became neutral or deleterious in later generations (left), and >90% of beneficial mutations from later generations were neutral or deleterious in the ancestor.




Fig. 6. Determinants of evolutionary outcomes. (A, C) The prevalence of the observed beneficial mutations in the LTEE is better explained by the mutational target size (A, area and color of dots represent fitness) than by the magnitude of beneficial fitness effects measured in the ancestor (C, area and color of dots represent target size). (B) The best linear model for mutation prevalence includes fitness but is more strongly dependent on the mutational target size (area of dots represents target size, and color represents fitness). (D) The predictive capacity of DFEs as a function of time in the LTEE. Values show the fraction of numerically dominant alleles at each generation that were captured by the DFE measured in the ancestor (black), 2K (red), and 15K (blue) backgrounds. For the ancestor, we measured this fraction across all 12 LTEE populations; for the evolved backgrounds, the fraction includes only the focal population. Shaded areas show the null expectations based on randomly sampling neutral and deleterious mutations.

References

1. C. O. Wilke, The speed of adaptation in large asexual populations. *Genetics* **167**, 2045–2053 (2004).
2. M. Lynch, R. Bürger, D. Butcher, W. Gabriel, The mutational meltdown in asexual populations. *J. Hered.* **84**, 339–344 (1993).
3. D. Charlesworth, B. Charlesworth, M. T. Morgan, The pattern of neutral molecular variation under the background selection model. *Genetics* **141**, 1619–1632 (1995).
4. L.-M. Chevin, G. Martin, T. Lenormand, Fisher’s model and the genomics of adaptation: restricted pleiotropy, heterogenous mutation, and parallel evolution. *Evolution* **64**, 3213–3231 (2010).
5. J. J. Welch, A. Eyre-Walker, D. Waxman, Divergence and polymorphism under the nearly neutral theory of molecular evolution. *J. Mol. Evol.* **67**, 418–426 (2008).
6. T. Ohta, The nearly neutral theory of molecular evolution. *Annu. Rev. Ecol. Syst.* **23**, 263–286 (1992).
7. J. R. Peck, G. Barreau, S. C. Heath, Imperfect genes, Fisherian mutation and the evolution of sex. *Genetics* **145**, 1171–1199 (1997).
8. 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).
9. M. L. M. Salverda, E. Dellus, F. A. Gorter, A. J. M. Debets, J. van der Oost, R. F. Hoekstra, D. S. Tawfik, J. A. G. M. de Visser, Initial mutations direct alternative pathways of protein evolution. *PLOS Genet.* **7**, e1001321 (2011).
10. D. Aggeli, Y. Li, G. Sherlock, Changes in the distribution of fitness effects and adaptive mutational spectra following a single first step towards adaptation. *Nat. Commun.* **12**, 5193 (2021).
11. M. J. Wiser, N. Ribeck, R. E. Lenski, Long-term dynamics of adaptation in asexual populations. *Science* **342**, 1364–1367 (2013).
12. B. H. Good, M. M. Desai, The impact of macroscopic epistasis on long-term evolutionary dynamics. *Genetics* **199**, 177–190 (2015).
13. A. Couce, O. A. Tenaillon, The rule of declining adaptability in microbial evolution experiments. *Front. Genet.* **6**, 99 (2015).
14. J. H. Gillespie, Molecular evolution over the mutational landscape. *Evolution* **38**, 1116–1129 (1984).
15. H. A. Orr, The distribution of fitness effects among beneficial mutations. *Genetics* **163**, 1519–1526 (2003).
16. R. Kassen, T. Bataillon, Distribution of fitness effects among beneficial mutations before selection in experimental populations of bacteria. *Nat. Genet.* **38**, 484–488 (2006).
17. D. R. Rokytá, 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).
18. R. Sanjuán, A. Moya, S. F. Elena, The distribution of fitness effects caused by single-nucleotide substitutions in an RNA virus. *Proc. Natl. Acad. Sci. U. S. A.* **101**, 8396–8401 (2004).
19. R. D. H. Barrett, R. Craig MacLean, G. Bell, Mutations of intermediate effect are responsible for adaptation in evolving *Pseudomonas fluorescens* populations. *Biol. Lett.* **2**, 236–238 (2006).
20. D. P. Rice, B. H. Good, M. M. Desai, The evolutionarily stable distribution of fitness effects. *Genetics* **200**, 321–329 (2015).

21. L. Perfeito, A. Sousa, T. Bataillon, I. Gordo, Rates of fitness decline and rebound suggest pervasive epistasis. *Evol. Int. J. Org. Evol.* **68**, 150–162 (2014).
22. J. A. G. M. de Visser, J. Hermisson, G. P. Wagner, L. Ancel Meyers, H. Bagheri-Chaichian, J. L. Blanchard, L. Chao, J. M. Cheverud, S. F. Elena, W. Fontana, G. Gibson, T. F. Hansen, D. Krakauer, R. C. Lewontin, C. Ofria, S. H. Rice, G. von Dassow, A. Wagner, M. C. Whitlock, Perspective: Evolution and detection of genetic robustness. *Evolution* **57**, 1959–1972 (2003).
23. R. Sanjuán, J. M. Cuevas, V. Furió, E. C. Holmes, A. Moya, Selection for robustness in mutagenized RNA viruses. *PLoS Genet.* **3**, e93 (2007).
24. C. O. Wilke, J. L. Wang, C. Ofria, R. E. Lenski, C. Adami, Evolution of digital organisms at high mutation rates leads to survival of the flattest. *Nature* **412**, 331–333 (2001).
25. J. D. Bloom, Z. Lu, D. Chen, A. Raval, O. S. Venturelli, F. H. Arnold, Evolution favors protein mutational robustness in sufficiently large populations. *BMC Biol.* **5**, 29 (2007).
26. M. S. Johnson, M. M. Desai, Mutational robustness changes during long-term adaptation in laboratory budding yeast populations. *eLife* **11**, e76491 (2022).
27. A. Butković, R. González, I. Cobo, S. F. Elena, Adaptation of turnip mosaic potyvirus to a specific niche reduces its genetic and environmental robustness. *Virus Evol.* **6**, veaa041 (2020).
28. I. S. Novella, J. B. Presluid, C. Beech, C. O. Wilke, Congruent evolution of fitness and genetic robustness in vesicular stomatitis virus. *J. Virol.* **87**, 4923–4928 (2013).
29. G. Reddy, M. M. Desai, Global epistasis emerges from a generic model of a complex trait. *eLife* **10**, e64740 (2021).
30. M. S. Johnson, A. Martsul, S. Kryazhimskiy, M. M. Desai, Higher-fitness yeast genotypes are less robust to deleterious mutations. *Science* **366**, 490–493 (2019).
31. R. J. Woods, J. E. Barrick, T. F. Cooper, U. Shrestha, M. R. Kauth, R. E. Lenski, Second-order selection for evolvability in a large *Escherichia coli* population. *Science* **331**, 1433–1436 (2011).
32. D. M. Weinreich, R. A. Watson, L. Chao, Perspective: Sign epistasis and genetic constraint on evolutionary trajectories. *Evolution* **59**, 1165–1174 (2005).
33. G. Rancati, J. Moffat, A. Typas, N. Pavelka, Emerging and evolving concepts in gene essentiality. *Nat. Rev. Genet.* **19**, 34–49 (2018).
34. L. Parts, A. Batté, M. Lopes, M. W. Yuen, M. Laver, B.-J. San Luis, J.-X. Yue, C. Pons, E. Eray, P. Aloy, G. Liti, J. van Leeuwen, Natural variants suppress mutations in hundreds of essential genes. *Mol. Syst. Biol.* **17**, e10138 (2021).
35. F. Rousset, J. Cabezas-Caballero, F. Piastra-Facon, J. Fernández-Rodríguez, O. Clermont, E. Denamur, E. P. C. Rocha, D. Bikard, The impact of genetic diversity on gene essentiality within the *Escherichia coli* species. *Nat. Microbiol.* **6**, 301–312 (2021).
36. B.-M. Koo, G. Kritikos, J. D. Farelli, H. Todor, K. Tong, H. Kimsey, I. Wapinski, M. Galardini, A. Cabal, J. M. Peters, A.-B. Hachmann, D. Z. Rudner, K. N. Allen, A. Typas, C. A. Gross, Construction and analysis of two genome-scale deletion libraries for *Bacillus subtilis*. *Cell Syst.* **4**, 291–305.e7 (2017).
37. G. Liu, M. Y. J. Yong, M. Yurieva, K. G. Srinivasan, J. Liu, J. S. Y. Lim, M. Poidinger, G. D. Wright, F. Zolezzi, H. Choi, N. Pavelka, G. Rancati, Gene essentiality is a quantitative property linked to cellular evolvability. *Cell* **163**, 1388–1399 (2015).
38. F. Rosconi, E. Rudmann, J. Li, D. Surujon, J. Anthony, M. Frank, D. S. Jones, C. Rock, J. W. Rosch, C. D. Johnston, T. van Opijnen, A bacterial pan-genome makes gene essentiality strain-dependent and evolvable. *Nat. Microbiol.* **7**, 1580–1592 (2022).
39. O. M. Maistrenko, D. R. Mende, M. Luetge, F. Hildebrand, T. S. B. Schmidt, S. S. Li, J. F. M. Rodrigues, C. von Mering, L. Pedro Coelho, J. Huerta-Cepas, S. Sunagawa, P. Bork,

- Disentangling the impact of environmental and phylogenetic constraints on prokaryotic within-species diversity. *ISME J.* **14**, 1247–1259 (2020).
40. E. Martínez-García, V. de Lorenzo, The quest for the minimal bacterial genome. *Curr. Opin. Biotechnol.* **42**, 216–224 (2016).
 41. R. Sanjuán, Mutational fitness effects in RNA and single-stranded DNA viruses: common patterns revealed by site-directed mutagenesis studies. *Philos. Trans. R. Soc. B Biol. Sci.* **365**, 1975–1982 (2010).
 42. H. Kemble, P. Nghe, O. Tenaillon, Recent insights into the genotype–phenotype relationship from massively parallel genetic assays. *Evol. Appl.* **12**, 1721–1742 (2019).
 43. R. E. Lenski, M. R. Rose, S. C. Simpson, S. C. Tadler, Long-term experimental evolution in *Escherichia coli*. I. Adaptation and divergence during 2,000 generations. *Am. Nat.* **138**, 1315–1341 (1991).
 44. E. C. A. Goodall, A. Robinson, I. G. Johnston, S. Jabbari, K. A. Turner, A. F. Cunningham, P. A. Lund, J. A. Cole, I. R. Henderson, The essential genome of *Escherichia coli* K-12. *mBio* **9**, e02096-17 (2018).
 45. O. Tenaillon, J. E. Barrick, N. Ribeck, D. E. Deatherage, J. L. Blanchard, A. Dasgupta, G. C. Wu, S. Wielgoss, S. Cruveiller, C. Médigue, D. Schneider, R. E. Lenski, Tempo and mode of genome evolution in a 50,000-generation experiment. *Nature* **536**, 165–170 (2016).
 46. B. H. Good, M. J. McDonald, J. E. Barrick, R. E. Lenski, M. M. Desai, The dynamics of molecular evolution over 60,000 generations. *Nature* **551**, 45–50 (2017).
 47. J.-F. Gout, D. Kahn, L. Duret, P. P.-G. Consortium, the relationship among gene expression, the evolution of gene dosage, and the rate of protein evolution. *PLOS Genet.* **6**, e1000944 (2010).
 48. J. L. Cherry, Expression level, evolutionary rate, and the cost of expression. *Genome Biol. Evol.* **2**, 757–769 (2010).
 49. J. Zhang, J.-R. Yang, Determinants of the rate of protein sequence evolution. *Nat. Rev. Genet.* **16**, 409–420 (2015).
 50. J. S. Favate, S. Liang, A. L. Cope, S. S. Yadavalli, P. Shah, The landscape of transcriptional and translational changes over 22 years of bacterial adaptation. *eLife* **11**, e81979 (2022).
 51. M. Lynch, J. S. Conery, The evolutionary fate and consequences of duplicate genes. *Science* **290**, 1151–1155 (2000).
 52. A. I. Khan, D. M. Dinh, D. Schneider, R. E. Lenski, T. F. Cooper, Negative epistasis between beneficial mutations in an evolving bacterial population. *Science* **332**, 1193–1196 (2011).
 53. F. Blanquart, G. Achaz, T. Bataillon, O. Tenaillon, Properties of selected mutations and genotypic landscapes under Fisher’s geometric model. *Evolution* **68**, 3537–3554 (2014).
 54. O. Tenaillon, A. Rodríguez-Verdugo, R. L. Gaut, P. McDonald, A. F. Bennett, A. D. Long, B. S. Gaut, The molecular diversity of adaptive convergence. *Science* **335**, 457–461 (2012).
 55. D. Greene, K. Crona, The changing geometry of a fitness landscape along an adaptive walk. *PLoS Comput. Biol.* **10**, e1003520 (2014).
 56. P. A. Lind, E. Libby, J. Herzog, P. B. Rainey, Predicting mutational routes to new adaptive phenotypes. *eLife* **8**, e38822 (2019).
 57. R. L. Moran, E. J. Richards, C. P. Ornelas-García, J. B. Gross, A. Donny, J. Wiese, A. C. Keene, J. E. Kowalko, N. Rohner, S. E. McGaugh, Selection-driven trait loss in independently evolved cavefish populations. *Nat. Commun.* **14**, 2557 (2023).
 58. D. E. Deatherage, J. E. Barrick, High-throughput characterization of mutations in genes that drive clonal evolution using multiplex adaptome capture sequencing. *Cell Syst.* **12**, 1187-1200.e4 (2021).

59. M. M. Desai, D. S. Fisher, A. W. Murray, The speed of evolution and maintenance of variation in asexual populations. *Curr. Biol.* **17**, 385–394 (2007).
60. V. S. Cooper, D. Schneider, M. Blot, R. E. Lenski, Mechanisms causing rapid and parallel losses of ribose catabolism in evolving populations of *Escherichia coli* B. *J. Bacteriol.* **183**, 2834–2841 (2001).
61. R. Maddamsetti, P. J. Hatcher, A. G. Green, B. L. Williams, D. S. Marks, R. E. Lenski, Core Genes Evolve Rapidly in the Long-term evolution experiment with *Escherichia coli*. *Genome Biol. Evol.* **9**, 1072–1083 (2017).
62. R. E. Lenski, Convergence and Divergence in a Long-term experiment with bacteria. *Am. Nat.* **190**, S57–S68 (2017).
63. H.-H. Chou, H.-C. Chiu, N. F. Delaney, D. Segrè, C. J. Marx, Diminishing returns epistasis among beneficial mutations decelerates adaptation. *Science* **332**, 1190–1192 (2011).
64. Y. Park, B. P. H. Metzger, J. W. Thornton, Epistatic drift causes gradual decay of predictability in protein evolution. *Science* **376**, 823–830 (2022).
65. S. L. Chiang, E. J. Rubin, Construction of a mariner-based transposon for epitope-tagging and genomic targeting. *Gene* **296**, 179–185 (2002).
66. L. Ferrières, G. Hémery, T. Nham, A.-M. Guérout, D. Mazel, C. Beloin, J.-M. Ghigo, Silent mischief: bacteriophage Mu insertions contaminate products of *Escherichia coli* random mutagenesis performed using suicidal transposon delivery plasmids mobilized by broad-host-range RP4 conjugative machinery. *J. Bacteriol.* **192**, 6418–6427 (2010).
67. A. L. Goodman, N. P. McNulty, Y. Zhao, D. Leip, R. D. Mitra, C. A. Lozupone, R. Knight, J. I. Gordon, Identifying genetic determinants needed to establish a human gut symbiont in its habitat. *Cell Host Microbe* **6**, 279–289 (2009).