Genomics 104 (2014) 412–416

Contents lists available at ScienceDirect

Genomics

journal homepage: www.elsevier.com/locate/ygeno

The spectrum of adaptive mutations in experimental evolution

Gregory I. Lang [a,*], Michael M. Desai [b,*]

[a] Department of Biological Sciences, Lehigh University, Bethlehem, PA 18015, USA
[b] Department of Organismic and Evolutionary Biology, Department of Physics, FAS Center for Systems Biology, Harvard University, Cambridge, MA 02138, USA

ARTICLE INFO

Article history:
Received 30 August 2014
Accepted 19 September 2014
Available online 28 September 2014

Keywords:
Experimental evolution
Adaptive mutations
Next-generation sequencing

ABSTRACT

A primary goal of recent work in experimental evolution is to probe the molecular basis of adaptation. This requires an understanding of the individual mutations in evolving populations: their identity, their physiological and fitness effects, and the interactions between them. The combination of high-throughput methods for laboratory evolution and next-generation sequencing methods now makes it possible to identify and quantify mutations in hundreds of replicate populations over thousands of generations, and to directly measure fitness effects and epistatic interactions. Many laboratories are now leveraging these tools to study the molecular basis of adaptation and the reproducibility of evolutionary outcomes across a variety of model systems. Genetic analyses on evolved populations are shedding light on the statistics of epistasis between evolved mutations. Here we review the current understanding of the spectrum of mutations observed across these systems, with a focus on epistatic interactions between beneficial mutations and constraints on evolutionary outcomes. We emphasize evolution in asexual microbes, where next generation sequencing methods have been widely applied.

Experimental evolution aims to exploit the advantages of laboratory systems to understand the mechanisms of evolutionary change. One key goal is to analyze the molecular basis of adaptation: which mutations underlie evolutionary adaptation, and why? Laboratory evolution experiments offer many advantages in addressing this question. By initiating a large number of replicate populations from the same starting point and maintaining them in identical conditions, it is possible to "replay the tape of life" to directly observe the role of chance and contingency in determining evolutionary outcomes [26]. The reproducibility and predictability of evolution can be quantified by measuring the degree of parallelism in both the phenotypic and genetic responses to selection across many replicate lines. Experimental conditions can be precisely controlled: population sizes and mutation rates can be adjusted over several orders of magnitude, and (in many systems) populations can be preserved, creating a "frozen fossil record." One can therefore return to any population at any timepoint in an experiment in order to measure fitness, assess phenotype, identify mutations, or replay the tape again. The frozen fossil record allows individuals to be compared to (and competed against) any other timepoint in the evolutionary history. Because most laboratory evolution experiments typically study model organisms such as phage, bacteria, yeast, and drosophila, many genetic tools are available to aid in the downstream genetic analysis of evolved populations. In addition, large communities of researchers focus on these model systems, leading to a deep understanding of their biology, which provides a meaningful context in which to interpret the outcomes of evolutionary experiments. Together, these advantages make laboratory evolution experiments a powerful tool for probing the molecular basis of adaptation.

1. Realizing the promise of experimental evolution

A key difficulty in analyzing evolution experiments is in identifying and tracking evolutionary changes through time. Over the past two decades, the primary observable quantity has been changes in fitness. This is typically assessed using direct competitions of evolved clones or populations against labeled reference strains (typically versions of the ancestral strains labeled with drug, nutrient, or fluorescent markers). These competitive fitness assays make it possible to precisely quantify the fitness difference between two strains; current experiments often measure this competitive fitness to within a fraction of a percent [23]. Frequency-dependent and non-transitive fitness interactions do occasionally appear in some systems [29,30,49,51,59]. However, simple directional selection is more typically observed, and fitness increases steadily through time as populations adapt (though this simple picture may be challenged soon, as more sensitive assays uncover common but subtle frequency-dependent and non-transitive interactions).

Fitness increases indicate that populations are adapting, but parallelism in the rate of adaptation does not necessarily mean that populations are adapting by way of similar mutations. In some cases, phenotypically distinguishable types repeatedly fix in populations. For example, all 12 lines in Richard Lenski's long-term evolution experiment (LTEE) in *Escherichia coli* evolved large cell size [44] and similar patterns of gene expression changes [15], and 6 of 12 lines evolved elevated mutation

* Corresponding authors.
E-mail addresses: glang@lehigh.edu (G.I. Lang), mmdesai@fas.harvard.edu (M.M. Desai).

http://dx.doi.org/10.1016/j.ygeno.2014.09.011
0888-7543/

rates [56,62]. In asexual haploid budding yeast populations, sterility repeatedly evolves [40], and patterns of gene expression change in response to nutrient limitations show some regularities [27]. These examples of parallelism in the evolution of phenotypes other than fitness suggest that a specific class of parallel mutations drive adaptation in these populations. However, in many other cases parallel phenotypic changes occur not because they are advantageous, but because of relaxed selection and mutation accumulation. For example, numerous studies have shown that adaptation in the LTEE leads to a reduction in catabolic breadth and thermal-tolerance in *E. coli* [5,16–18,43,55].

In addition to their implications for the physiological basis of adaptation, spontaneously arising (or pre-existing) phenotypic markers can provide a window into the dynamics of adaptation. For example, periodic accumulation and purging of phenotypic markers has been used to identify potential selective sweeps [1,48], or to analyze clonal interference between multiple competing adaptive mutations [40]. Alternatively, populations can be initiated with multiple phenotypically distinguishable markers [2,22,28,32], to probe the effects of adaptive mutations that alter marker frequencies. These approaches allow us to make a statistical estimate of the rates and fitness effects of beneficial mutations that drive adaptation.

Unfortunately, while all of these approaches provide some insight into the mechanisms of evolutionary change, they do not provide a direct view of the molecular basis of adaptation. They cannot tell us what the adaptive mutations are, why they are beneficial, or how the evolutionary process picks out these specific mutations from amidst a larger set of possibilities. This limits our ability to draw conclusions about the evolutionary process in nature more generally. Fortunately, in recent years, rapid improvements in sequencing technology now make it possible to directly identify the mutations underlying adaptation in laboratory evolution experiments. This has led to a surge of recent studies that attempt to characterize the molecular basis of adaptation in these experiments [19]. In this review, we discuss the implications of these studies for our understanding of the nature of adaptive mutations in experimental evolution, and we survey key current challenges in collecting and interpreting these data.

2. Identifying mutations that arise during experimental evolution

Sanger sequencing methods have long been applied to experimentally evolved bacteriophage populations, with genome sizes on the order of several kilobases [9,10,12,61]. However, until recently, the only way to identify likely beneficial mutations in larger genomes was to target and Sanger sequence candidate genes based on parallel phenotypic evolution. While this approach did help find some mutations underlying adaptation, typically only a tiny fraction of beneficial mutations can be identified in this way. In the last five years, next-generation sequencing (particularly Illumina) has now made it practical and cost-effective to use whole-genome sequencing to analyze genetic changes in evolution experiments in organisms with much larger genome sizes.

The first experimentally-evolved microbial genomes to be sequenced were a set of bacterial clones from six timepoints of a single population from Lenski's long-term evolution experiment [3,4]. These clones provided the first full view of genomic evolution outside of bacteriophage. This study demonstrated the potential of applying full genome sequencing as a direct readout of experimental evolution, highlighting the difference between the rates of genotypic and phenotypic evolutionary changes (and in this case finding a surprising discordance between these rates). Though it is only five years old, the scope of this sequencing project pales in comparison to what is possible today. With current sequencing technologies, several orders of magnitude more data can be produced routinely as a full readout of experimental evolution. A single lane on an Illumina HiSeq 2500 can generate tens of billions of base pairs of sequencing data per lane — roughly 7000-fold coverage of the *E. coli* genome or 3000-fold coverage of the yeast genome. The current challenge is not sequencing capacity, but how to best maximize this powerful tool.

Following the example of Barrick et al. [4], next-generation sequencing technology has been leveraged to identify mutations in evolved genomes from a variety of systems including bacteriophage, *E. coli* [6,57], Methylobacterium [14], *Saccharomyces cerevisiae* [35,36,38,42], and *Drosophila* [13,58]. A number of sequencing strategies are being used, each of which provides a different perspective on evolved populations. The choice of sequencing strategy (clones versus populations, endpoint or timecourse) depends on details of the sequencing technology and on the questions being addressed.

One common approach is to sequence endpoint clones from an evolved population. This makes it relatively straightforward to identify all fixed mutations (compared to population sequencing) since all evolved mutations will be at expected frequencies of 1 for haploid genomes (or 0.5 for heterozygous mutations in diploid genomes). This simplifies downstream processing of the data, but produces a limited picture of genome evolution. Sequencing of endpoint clones can miss much of the genetic variation in the population. In particular, sequencing of clones will fail to detect subpopulations which may be common in experimental evolution due to balancing selection [34,40,46,59]. In addition, sequencing of clones will identify a random subset of low-frequency mutations that happened to be present in the clone that was selected to be representative of the population. It is therefore advisable to sequence more than a single clone from each population and timepoint. Until recently, this has been challenging because a separate sequencing library must be prepared for each clone, and although sequencing costs are low, this library preparation is expensive. Recent work, however, has developed new methods to prepare sequencing libraries appropriate for multiplexed sequencing of many clones in a single Illumina lane at minimal cost [36].

In contrast to sequencing clones, whole-population "metagenomic" sequencing gives a more precise picture of the mutations present in the population at a given timepoint, but requires higher coverage to identify mutations at intermediate frequencies (and to accurately quantify their frequency). This also introduces new challenges in data processing, particularly in distinguishing low-frequency mutations from sequencing or alignment errors. Several strategies have been used to detect low frequency alleles in evolved populations, including sequencing to high coverage, paired-end sequencing of the same short read [38], circular sequencing and related strategies [45], and leveraging the information attained from time-course sequencing [42]. Because of the inherent tradeoffs in sequencing strategies, each of these choices reveals slightly different patterns of genomic evolution, and the optimal sequencing strategy may be a combination of methods, such as whole-population timecourse sequencing combined with sequencing several endpoint clones.

3. Identifying drivers of adaptation

Whole genome sequencing has the potential to reveal all mutations in an evolved genome. Among these are the beneficial mutations, but also neutral or deleterious mutations that happen to arise in the same background and hitchhike to high frequency. This is particularly an issue in asexual populations where the entire genome is one linkage group, but hitchhiking will also occur in sexual populations over genomic length scales where recombination is rare. How then does one distinguish between beneficial driver mutations and neutral or deleterious hitchhikers?

The most obvious approach is to reconstruct each evolved mutation in the ancestral background and directly measure the effect of each on fitness. While this can be done for small numbers of mutations in genetically-tractable organisms, it is not feasible in most situations. Therefore, statistical methods for distinguishing drivers from hitchhikers have become a necessity in analyzing experimental evolution. Common methods include testing for deviations of the expected ratio

of nonsynonymous to synonymous mutations under neutrality or identifying those genes or groups of genes in the same biological processes in which mutations occur more often than expected by chance.

Sequencing of a 20,000-generation evolved clone of *E. coli* revealed 26 coding sequence mutations, none of which were synonymous [4]. Under a neutral model, the probability of observing no synonymous mutations is negligible, suggesting that most if not all of the mutations observed in this experiment are beneficial. Similar underrepresentation of synonymous mutations is observed in other *E. coli* evolution experiments [6,57] as well as in yeast [38]. However, even in rapidly adapting populations, neutral mutations can often hitchhike to high frequency, depending on the relevant population genetic parameters. For example, a different yeast evolution study that was performed at a population size orders of magnitude smaller than Kvitek and Sherlock, found 15% mutations to be synonymous [42]. This percentage of synonymous mutations is similar to what is observed in the Lenski *E. coli* long-term evolution experiment after the emergence of a mutator phenotype [4], though the populations in Lang et al. [42] have a normal mutation rate. The degree of genetic hitchhiking is, therefore, dependent on both population size and mutation rate, as well as the distribution of fitness effects of beneficial mutations.

Even in populations where the majority of mutations are nonsynonymous, it cannot be assumed that all observed mutations are beneficial. In sexual organisms such as yeast, driver and hitchhiker mutations can be separated experimentally by backcrossing evolved clones to the ancestor and selecting for progeny that display the evolved phenotype [35]; related approaches have recently been applied in *E. coli* [50]. Similar bulk-segregant sequencing strategies have been used to map the genetic basis of complex traits [21].

To help differentiate the beneficial mutations from putative hitchhikers, one can take advantage of a central property of evolution experiments: replicate populations. Mutations that appear in multiple evolution experiments are likely to be drivers of adaptation. This parallelism is typically not at the nucleotide level: it is rare to observe the same exact mutation in replicate populations. When this does occur, it is often due to the presence of a mutational hotspot such as a homopolymeric run or due to severe constraint on the type of alteration to protein function. Commonly, at least in haploids, adaptation selects for loss or attenuation of function of target genes, in which case any number of mutations within a given gene will satisfy the selection. For example, we recently used this logic to identify genes that regularly yield drivers and distinguish these from hitchhiker mutations in a whole-population time-course sequencing from a 1000 generation yeast evolution experiment [42]. Across 40 populations, 723 coding mutations were identified, with 23 genes (out of ~5800 genes in the yeast genome) mutated in three or more replicate lines; all but one of these mutations was nonsynonymous. This near lack of synonymous mutations is in line with the spectrum of mutations from the Lenski long-term evolution experiment [4] and provides strong support for the statistical models of determining driver mutations.

Since selection acts on phenotype, and many genotypic changes are phenotypically indistinguishable, additional statistical power can be gained by grouping mutations, not by gene, but by biological process. Taking this more general perspective reveals higher levels of convergence in evolutionary outcomes at higher levels of biological organization. This principle is clearly demonstrated in recent studies that identified 1331 mutations from the sequencing of endpoint clones from 115 populations of *E. coli* that were evolved for 2000 generations at high temperature [52,57]. These independently evolved clones share few identical mutations, but as one changes perspective to higher levels of biological organization—from genes to operons to functional units—more similarities emerge between parallel evolved populations, such that any two evolved populations are likely to share 2.6% of point mutations, but 31.5% of affected functional units.

4. The nature of adaptive mutations

Advancing a full mechanistic model of adaptive evolution requires understanding the effects of individual beneficial mutations—on protein function, gene expression, metabolism, and fitness. Most long-term evolution experiments thus far have been performed in bacteria or haploid yeast populations, where, in most environments, there exist a number of loss-of-function mutations that provide a selective advantage. Given the large target size for these types of mutations, loss-of-function mutations often predominate the spectra of mutations recovered from long-term evolution experiments. Some of these loss events are neutral, attributable to mutation accumulation in the absence of selection for function, such as the reduction of catabolic breadth in *E. coli* [17,18,43]. However, many loss-of-function mutations have been confirmed to provide a selective advantage. For instance sterility in yeast provides a selective advantage by eliminating unnecessary gene expression [41]. The availability of beneficial loss-of-function mutations and the large target size for these events ensure that these mutations will come to dominate experimental evolution over short time scales. Over long time scales or in specialized conditions, mutational spectra may shift towards gain of function mutations. In diploid populations, we may also see a shift in the mutational spectrum away from loss-of-function mutations, towards dominant or overdominant mutations [24,54]. However, there is currently only limited data describing the mutations that occur during experimental evolution in diploids, leaving the exact nature of this shift unclear.

Mutations affecting gene dosage are common in experimental evolution, in particular under nutrient limitation. For example, in chemostat cultures, where growth is strongly limited by a single nutrient, mutations that increase the import of the limiting nutrient are favored. Evolution in sulfur-limited media leads to the specific amplification of the high-affinity sulfur transporter, *SUL1* [27]. Evolution under glucose limitation selects for amplification of the hexose transporters [11,20,32], and nitrogen limitation selects for amplification of nitrogen transporters *PUT4*, *DUR3*, and *DAL4* [31]. The fitness effects of these amplification events are on the order of 10%, roughly an order of magnitude more advantageous than the beneficial mutations observed in rich medium conditions. This existence of these large effect mutations, combined with the large population sizes, may partially explain why some of the smaller effect mutations observed in rich medium conditions are not found in chemostat evolution experiments.

Adaptation to a limiting nutrient also commonly occurs by way of structural rearrangements that generate novel gene fusions. In yeast glucose limited chemostats, an increase in respiration is driven by a recurrent translocation involving the *CIT1* gene, encoding citrate synthase, which catalyzes a rate-limiting step in the TCA cycle [20]. This rearrangement possibly de-represses the expression of *CIT1* in a glucose limited medium, allowing cells to increase respiration under normally fermentative conditions, thus generating more ATP per molecule of limiting glucose.

Structural rearrangements have also been observed in the evolution of aerobic citrate utilization in *E. coli*, which normally only metabolizes citrate in the absence of oxygen. The causative mutation is a two-step mutational event consisting of a novel structural rearrangement that brought a promoter the CitT gene under the control of an aerobically-active promoter followed by a tandem amplification of this region [7,50]. The net effect is that the citrate transporter is expressed at a high level in the presence of oxygen. The use of structural rearrangements to generate novel gain-of-function fusion proteins is analogous to the formation of the Bcr–Abl fusion protein in chronic myelogenous leukemia [53], and serves to highlight the similarity of molecular mechanisms in adaptive evolution across systems.

A key determinant to the success of any mutation is its effect on organismal fitness. For a given evolution experiment, the probability that a beneficial mutation with a given fitness effect will fix depends on underlying population genetic parameters such as mutation rate

and population size [25]. For the same distribution of possible mutational fitness effects, changes to underlying parameters, such as population size, can shift the characteristic fitness effects of fixed beneficial mutations. Thus even in seemingly identical environments that present the same selective pressures, the details of the experimental setup (e.g. the population size) can alter the spectrum of adaptive mutations we expect to observe.

5. Future opportunities and challenges

The decreasing costs of next generation sequencing and the decentralization of sequencing facilities away from large sequencing centers into individual laboratories have opened a new frontier in the study of experimental evolution. Just a decade ago, it was practically impossible to identify all of the mutations in an evolved genome and trace those changes over thousands of generations. Now, this is routine. As sequencing technology continues to improve, over the next decade we are likely to see a dramatic increase in our ability to test fundamental questions in evolutionary biology in the laboratory. With this new technology comes numerous technical challenges, particularly in designing computational tools to identify and track mutations in population sequencing data, and in developing methods to infer their physiological and fitness effects.

In the immediate future, further sequencing studies are likely to provide a dramatically improved view of the reproducibility of evolution. To move towards a full mechanistic understanding of how chance and determinism influence evolutionary dynamics, we need to focus on individual mutations—their immediate effect on fitness, the biological basis of these fitness effects, and their epistatic interactions. The degree of reproducibility of evolutionary outcomes depends on perspective—the degree of parallelism will differ for genotypic and phenotypic evolution. Phenotypic evolution is necessarily more constrained, for two reasons. First, selection acts on phenotype, and many genotypic changes are phenotypically indistinguishable. Second, random processes such as genetic hitchhiking allow for the fixation of neutral (or small effect) mutations, contributing to the uniqueness of genotypic evolution.

Evolutionary reproducibility—both phenotypic and genotypic—particularly depends on the patterns of gene–gene (epistatic) interactions. While there are often substantial epistatic interactions among mutations in a single gene [60] or among deleterious mutations [39], the emerging picture is that overall patterns of epistatic interactions among beneficial mutations are dominated by diminishing returns epistasis [14,33,36], with only a few examples of sign epistasis [37] between beneficial mutations in experimental evolution. This suggests that fitness landscapes are “smooth” in some sense. A second type of generic epistatic interaction is prevalent among beneficial mutations in the same biological process. Once a phenotype is attained, all of the other mutations that produce the same phenotype are selectively neutral. For example, there are at least nine genes in which mutations eliminate signaling through the yeast mating pathway resulting in a selective advantage [41]. These mutations are commonly observed in long-term evolution [40], but never do two sterile mutations fix in the same background [42]. Once a sterile mutation arises, all other mutations producing this phenotype are no longer strongly beneficial. A similar observation has also been reported in *E. coli* [57]. This epistasis between mutations in the same biological process contributes both to parallelism in phenotypic evolution and uniqueness in genotypic evolution.

In the absence of epistasis between mutations in different biological processes, parallel evolution experiments will converge on the same phenotype, each by way of a unique evolutionary trajectory. Phenotypic evolution, therefore, may be predictable despite the inherent randomness of genotypic evolution. It is not clear how much epistasis is required for phenotypic evolutionary trajectories to diverge, nor do we have sufficient empirical data to assess the number of beneficial mutations or the patterns of epistasis in evolving populations.

Although phenotypic evolution may be largely predictable over short time scales, over longer time scales the distribution of possible evolutionary outcomes is conditioned by epistasis, whereby fixed mutations either permit (or prohibit) particular evolutionary paths. Indeed, over very long time scales, chance fixation events can affect subsequent evolutionary change by allowing the future evolution of more complex phenotypes. For example, the emergence of the Cit+ phenotype in the Lenski long-term evolution experiment arose only after 30,000 generations [8]. The full acquisition of strong growth on citrate required potentiating mutations as well as later refining mutations that strengthened this phenotype [50].

The emergence of the Cit+ phenotype is the exception in experimental evolution, where most evolved mutations affect independent genes and biological pathways, driven largely by large-target loss-of-function mutations. As long-term evolution experiments more routinely stretch into the thousands or tens-of-thousands of generations, further instances of multiple mutations affecting the same gene or biological process will likely be observed. In these cases the patterns of epistasis will likely differ from the general diminishing returns epistasis ubiquitously observed in evolution experiments. Retrospective studies of protein evolution such as mammalian hormone receptors [47] and beta lactamase [60] reveal that sign epistasis strongly constrains evolutionary pathways in these systems. As the field of experimental evolution—and long-term experiments themselves—continue to age, a fuller picture of how individual mutations govern evolutionary outcomes will emerge.

References

[1] K.C. Atwood, L.K. Schneider, F.J. Ryan, Periodic selection in *Escherichia coli*, Proc. Natl. Acad. Sci. U. S. A. 37 (1951) 146–155.
[2] J.E. Barrick, M.R. Kauth, C.C. Strelioff, R.E. Lenski, Escherichia coli rpoB mutants have increased evolvability in proportion to their fitness defects, Mol. Biol. Evol. 27 (2010) 1338–1347.
[3] J.E. Barrick, R.E. Lenski, Genome-wide mutational diversity in an evolving population of *Escherichia coli*, Cold Spring Harb. Symp. Quant. Biol. 74 (2009) 119–129.
[4] J.E. Barrick, D.S. Yu, S.H. Yoon, H. Jeong, T.K. Oh, et al., Genome evolution and adaptation in a long-term experiment with *Escherichia coli*, Nature 461 (2009) 1243–1247.
[5] A.F. Bennett, R.E. Lenski, An experimental test of evolutionary trade-offs during temperature adaptation, Proc. Natl. Acad. Sci. U. S. A. 104 (Suppl. 1) (2007) 8649–8654.
[6] D. Blank, L. Wolf, M. Ackermann, O.K. Silander, The predictability of molecular evolution during functional innovation, Proc. Natl. Acad. Sci. U. S. A. 111 (2014) 3044–3049.
[7] Z.D. Blount, J.E. Barrick, C.J. Davidson, R.E. Lenski, Genomic analysis of a key innovation in an experimental *Escherichia coli* population, Nature 489 (2012) 513–518.
[8] Z.D. Blount, C.Z. Borland, R.E. Lenski, Historical contingency and the evolution of a key innovation in an experimental population of *Escherichia coli*, Proc. Natl. Acad. Sci. U. S. A. 105 (2008) 7899–7906.
[9] J.P. Bollback, J.P. Huelsenbeck, Clonal interference is alleviated by high mutation rates in large populations, Mol. Biol. Evol. 24 (2007) 1397–1406.
[10] J.P. Bollback, J.P. Huelsenbeck, Parallel genetic evolution within and between bacteriophage species of varying degrees of divergence, Genetics 181 (2009) 225–234.
[11] C.J. Brown, K.M. Todd, R.F. Rosenzweig, Multiple duplications of yeast hexose transport genes in response to selection in a glucose-limited environment, Mol. Biol. Evol. 15 (1998) 931–942.
[12] J.J. Bull, M.R. Badgett, H.A. Wichman, J.P. Huelsenbeck, D.M. Hillis, et al., Exceptional convergent evolution in a virus, Genetics 147 (1997) 1497–1507.
[13] M.K. Burke, J.P. Dunham, P. Shahrestani, K.R. Thornton, M.R. Rose, et al., Genome-wide analysis of a long-term evolution experiment with *Drosophila*, Nature 467 (2010) 587–590.
[14] H.H. Chou, H.C. Chiu, N.F. Delaney, D. Segre, C.J. Marx, Diminishing returns epistasis among beneficial mutations decelerates adaptation, Science 332 (2011) 1190–1192.
[15] T.F. Cooper, D.E. Rozen, R.E. Lenski, Parallel changes in gene expression after 20,000 generations of evolution in *Escherichia coli*, Proc. Natl. Acad. Sci. U. S. A. 100 (2003) 1072–1077.
[16] V.S. Cooper, A.F. Bennett, R.E. Lenski, Evolution of thermal dependence of growth rate of *Escherichia coli* populations during 20,000 generations in a constant environment, Evolution 55 (2001) 889–896.
[17] V.S. Cooper, R.E. Lenski, The population genetics of ecological specialization in evolving *Escherichia coli* populations, Nature 407 (2000) 736–739.
[18] 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 (2001) 2834–2841.
[19] J.R. Dettman, N. Rodrigue, A.H. Melnyk, A. Wong, S.F. Bailey, et al., Evolutionary insight from whole-genome sequencing of experimentally evolved microbes, Mol. Ecol. 21 (2012) 2058–2077.

[20] M.J. Dunham, H. Badrane, T. Ferea, J. Adams, P.O. Brown, et al., Characteristic genome rearrangements in experimental evolution of *Saccharomyces cerevisiae*, Proc. Natl. Acad. Sci. U. S. A. 99 (2002) 16144–16149.
[21] I.M. Ehrenreich, N. Torabi, Y. Jia, J. Kent, S. Martis, et al., Dissection of genetically complex traits with extremely large pools of yeast segregants, Nature 464 (2010) 1039–1042.
[22] E.M. Frenkel, B.H. Good, M.M. Desai, The fates of mutant lineages and the distribution of fitness effects of beneficial mutations in laboratory budding yeast populations, Genetics 196 (2014) 1217–1226.
[23] R. Gallet, T.F. Cooper, S.F. Elena, T. Lenormand, Measuring selection coefficients below 10–3: method, questions, and prospects, Genetics 190 (2012) 175–186.
[24] A.C. Gerstein, S.P. Otto, Ploidy and the causes of genomic evolution, J. Hered. 100 (2009) 571–581.
[25] B.H. Good, I.M. Rouzine, D.J. Balick, O. Hallatschek, M.M. Desai, Distribution of fixed beneficial mutations and the rate of adaptation in asexual populations, Proc. Natl. Acad. Sci. 109 (2012) 4950–4955.
[26] S.J. Gould, Wonderful Life, Norton, New York, USA, 1989.
[27] D. Gresham, M.M. Desai, C.M. Tucker, H.T. Jenq, D.A. Pai, et al., The repertoire and dynamics of evolutionary adaptations to controlled nutrient-limited environments in yeast, PLoS Genet. 4 (2008) e1000303.
[28] M. Hegreness, N. Shoresh, D. Hartl, R. Kishony, An equivalence principle for the incorporation of favorable mutations in asexual populations, Science 311 (2006) 1615–1617.
[29] R.B. Helling, C.N. Vargas, J. Adams, Evolution of *Escherichia coli* during growth in a constant environment, Genetics 116 (1987) 349–358.
[30] M.D. Herron, M. Doebeli, Parallel evolutionary dynamics of adaptive diversification in *Escherichia coli*, PLoS Biol. 11 (2013) e1001490.
[31] J. Hong, D. Gresham, Molecular specificity, convergence and constraint shape adaptive evolution in nutrient-poor environments, PLoS Genet. 10 (2014) e1004041.
[32] K.C. Kao, G. Sherlock, Molecular characterization of clonal interference during adaptive evolution in asexual populations of *Saccharomyces cerevisiae*, Nat. Genet. 40 (2008) 1499–1504.
[33] 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 (2011) 1193–1196.
[34] M. Kinnersley, J. Wenger, E. Kroll, J. Adams, G. Sherlock, et al., Ex uno plures: clonal reinforcement drives evolution of a simple microbial community, PLoS Genet. 10 (2014) e1004430.
[35] J.H. Koschwanez, K.R. Foster, A.W. Murray, Improved use of a public good selects for the evolution of undifferentiated multicellularity, Elife 2 (2013) e00367.
[36] S. Kryazhimskiy, D.P. Rice, E.R. Jerison, M.M. Desai, Global epistasis makes adaptation predictable despite sequence-level stochasticity, Science 344 (2014) 1519–1522.
[37] D.J. Kvitek, G. Sherlock, Reciprocal sign epistasis between frequently experimentally evolved adaptive mutations causes a rugged fitness landscape, PLoS Genet. 7 (2011) e1002056.
[38] D.J. Kvitek, G. Sherlock, Whole genome, whole population sequencing reveals that loss of signaling networks is the major adaptive strategy in a constant environment, PLoS Genet. 9 (2013) e1003972.
[39] J. Lalic, S.F. Elena, Magnitude and sign epistasis among deleterious mutations in a positive-sense plant RNA virus, Heredity 109 (2012) 71–77.
[40] G.I. Lang, D. Botstein, M.M. Desai, Genetic variation and the fate of beneficial mutations in asexual populations, Genetics 188 (2011) 647–661.
[41] G.I. Lang, A.W. Murray, D. Botstein, The cost of gene expression underlies a fitness trade-off in yeast, Proc. Natl. Acad. Sci. U. S. A. 106 (2009) 5755–5760.
[42] G.I. Lang, D.P. Rice, M.J. Hickman, E. Sodergren, G.M. Weinstock, et al., Pervasive genetic hitchhiking and clonal interference in forty evolving yeast populations, Nature 500 (2013) 571–574.
[43] N. Leiby, C.J. Marx, Metabolic erosion primarily through mutation accumulation, and not tradeoffs, drives limited evolution of substrate specificity in *Escherichia coli*, PLoS Biol. 12 (2014) e1001789.
[44] R.E. Lenski, M. Travisano, Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations, Proc. Natl. Acad. Sci. U. S. A. 91 (1994) 6808–6814.
[45] D.I. Lou, J.A. Hussmann, R.M. McBee, A. Acevedo, R. Andino, et al., High-throughput DNA sequencing errors are reduced by orders of magnitude using circle sequencing, Proc. Natl. Acad. Sci. 110 (2013) 19872–19877.
[46] M.J. McDonald, S.M. Gehrig, P.L. Meintjes, X.X. Zhang, P.B. Rainey, Adaptive divergence in experimental populations of Pseudomonas fluorescens. IV. Genetic constraints guide evolutionary trajectories in a parallel adaptive radiation, Genetics 183 (2009) 1041–1053.
[47] E.A. Ortlund, J.T. Bridgham, M.R. Redinbo, J.W. Thornton, Crystal structure of an ancient protein: evolution by conformational epistasis, Science 317 (2007) 1544–1548.
[48] C. Paquin, J. Adams, Frequency of fixation of adaptive mutations is higher in evolving diploid than haploid yeast populations, Nature 302 (1983) 495–500.
[49] J. Plucain, T. Hindré, M. Le Gac, O. Tenaillon, S. Cruveiller, et al., Epistasis and allele specificity in the emergence of a stable polymorphism in *Escherichia coli*, Science 343 (2014) 1366–1369.
[50] E.M. Quandt, D.E. Deatherage, A.D. Ellington, G. Georgiou, J.E. Barrick, Recursive genomewide recombination and sequencing reveals a key refinement step in the evolution of a metabolic innovation in *Escherichia coli*, Proc. Natl. Acad. Sci. U. S. A. 111 (2014) 2217–2222.
[51] P.B. Rainey, M. Travisano, Adaptive radiation in a heterogeneous environment, Nature 394 (1998) 69–72.
[52] A. Rodríguez-Verdugo, D. Carrillo-Cisneros, A. González-González, B.S. Gaut, A.F. Bennett, Different tradeoffs result from alternate genetic adaptations to a common environment, Proc. Natl. Acad. Sci. 111 (2014) 12121–12126.
[53] J.D. Rowley, Letter: a new consistent chromosomal abnormality in chronic myelogenous leukaemia identified by quinacrine fluorescence and Giemsa staining, Nature 243 (1973) 290–293.
[54] D. Sellis, B.J. Callahan, D.A. Petrov, P.W. Messer, Heterozygote advantage as a natural consequence of adaptation in diploids, Proc. Natl. Acad. Sci. U. S. A. 108 (2011) 20666–20671.
[55] S.C. Sleight, N.S. Wigginton, R.E. Lenski, Increased susceptibility to repeated freeze-thaw cycles in *Escherichia coli* following long-term evolution in a benign environment, BMC Evol. Biol. 6 (2006) 104.
[56] P.D. Sniegowski, P.J. Gerrish, R.E. Lenski, Evolution of high mutation rates in experimental populations of *E. coli*, Nature 387 (1997) 703–705.
[57] O. Tenaillon, A. Rodriguez-Verdugo, R.L. Gaut, P. McDonald, A.F. Bennett, et al., The molecular diversity of adaptive convergence, Science 335 (2012) 457–461.
[58] H. Teotonio, I.M. Chelo, M. Bradic, M.R. Rose, A.D. Long, Experimental evolution reveals natural selection on standing genetic variation, Nat. Genet. 41 (2009) 251–257.
[59] C.C. Traverse, L.M. Mayo-Smith, S.R. Poltak, V.S. Cooper, Tangled bank of experimentally evolved Burkholderia biofilms reflects selection during chronic infections, Proc. Natl. Acad. Sci. U. S. A. 110 (2013) E250–E259.
[60] D.M. Weinreich, N.F. Delaney, M.A. Depristo, D.L. Hartl, Darwinian evolution can follow only very few mutational paths to fitter proteins, Science 312 (2006) 111–114.
[61] H.A. Wichman, M.R. Badgett, L.A. Scott, C.M. Boulianne, J.J. Bull, Different trajectories of parallel evolution during viral adaptation, Science 285 (1999) 422–424.
[62] S. Wielgoss, J.E. Barrick, O. Tenaillon, M.J. Wiser, W.J. Dittmar, et al., Mutation rate dynamics in a bacterial population reflect tension between adaptation and genetic load, Proc. Natl. Acad. Sci. U. S. A. 110 (2013) 222–227.