Idiosyncratic purifying selection on metabolic enzymes in the long-term evolution experiment with *Escherichia coli*

Rohan Maddamsetti[1]

[1]Department of Biomedical Engineering, Duke University, Durham, NC, USA
E-mail: rohan.maddamsetti@duke.edu

Abstract

Bacteria, Archaea, and Eukarya all share a common set of metabolic reactions. This implies that the function and topology of central metabolism has been evolving under purifying selection over deep time. Central metabolism may similarly evolve under purifying selection during long-term evolution experiments, although it is unclear how long such experiments would have to run (decades, centuries, millennia) before signs of purifying selection on metabolism appear. I hypothesized that central and superessential metabolic enzymes would show evidence of purifying selection in the long-term evolution experiment with *Escherichia coli* (LTEE), in which 12 initially identical bacterial populations have been evolving in the laboratory for more than 30 years and 60,000 bacterial generations. I also hypothesized that enzymes that specialize on single substrates would show stronger evidence of purifying selection in the LTEE than generalist enzymes that catalyze multiple reactions. I tested these hypotheses by analyzing metagenomic time series covering 60,000 generations of the LTEE. I find mixed support for these hypotheses: patterns of purifying selection on metabolic enzymes, at least after 60,000 generations of experimental evolution, are largely idiosyncratic and population-specific.

Significance Statement

Purifying selection conserves organismal function over evolutionary time. I looked for evidence of purifying selection on metabolic enzymes in an ongoing long-term evolution experiment with *E. coli*. While some populations show signs of purifying selection, the overall pattern is inconsistent. To explain this finding, I propose that each population's metabolism is evolving in a molecular game of Jenga. In this conceptual model, loss-of-function mutations degrade costly, redundant, and inessential metabolic functions, after which purifying selection begins to dominate. The threshold at which purifying selection activates depends on the idiosyncratic trajectory of lost redundancies in each population.

Introduction

Researchers have learned much about adaptive processes by conducting evolution experiments with large populations of microbes. By contrast, researchers have rarely, if ever, used experimental evolution to study purifying selection, because it may take hundreds or even thousands of years to see evolutionary stasis in action. My colleagues and I have recently started to address this research gap, by examining detailed time series of genetic change in Lenski's long-term evolution experiment with *Escherichia coli* (LTEE).

The LTEE has studied the evolution of 12 initially identical populations of *Escherichia coli* in carbon-limited minimal medium for more than 30 years and 60,000 generations of bacterial evolution (Lenski, et al. 1991; Lenski 2017). Six of the populations can grow on arabinose (populations Ara+1 to Ara+6), while six cannot (populations Ara−1 to Ara−6). This neutral phenotypic marker is used to measure fitness gains over time (Lenski, et al. 1991; Wiser, et al. 2013). The LTEE populations have diverged in their mutation rates and biases, as several LTEE populations have evolved elevated mutation rates due to defects in DNA repair (Sniegowski, et al. 1997; Tenaillon, et al. 2016; Good, et al. 2017; Maddamsetti and Grant 2020a). These hypermutator populations tend to adapt faster than nonmutator populations that have retained the ancestral mutation rate (Wiser, et al. 2013), even though their more rapid genomic evolution largely reflects the accumulation of nearly neutral mutations (Couce, et al. 2017).

Because so many mutations are observed in the hypermutator populations, sets of genes that are significantly depleted of observed mutations can be inferred to be evolving under purifying selection. Based on this idea, my colleagues and I developed a randomization test that we call STIMS (Statistical Test to Infer Mode of Selection) to detect selection on pre-defined sets of genes, using metagenomic time series covering 60,000 generations of the LTEE (Good, et al. 2017). STIMS can detect both positive and purifying selection, and is able to control for both temporal variation in mutation rates over time as well as mutation rate variation over the chromosome (Maddamsetti and Grant 2020a; Maddamsetti and Grant 2020b). We used STIMS to find evidence of purifying selection on aerobic-specific and anaerobic-specific genes (Grant, Maddamsetti, et al. 2020) and essential genes (Maddamsetti and Grant 2020b) in the LTEE. In addition, we found evidence of purifying selection on genes encoding highly expressed and highly interacting proteins, using different methods (Maddamsetti 2020).

Here, I use STIMS to test four sets of metabolic genes for purifying selection in the LTEE: (1) those that encode enzymes that catalyze the core reactions of *E. coli* central metabolism, as curated in the BiGG Models knowledge base (King, et al. 2015); (2) those encoding enzymes that catalyze "superessential" reactions that have been found to be essential in all tested bacterial metabolic network models (Barve, et al. 2012); (3) and (4) those that encode specialist and generalist enzymes, respectively, in the *E. coli* metabolic reaction network, as determined by their substrate and reaction specificity (Nam, et al. 2012).

I tested two specific hypotheses. First, I hypothesized that the BiGG core metabolic enzymes and the superessential metabolic enzymes would show evidence of purifying selection. Second, I hypothesized that specialist enzymes would show stronger purifying selection than generalist enzymes. My results show mixed support for these hypotheses. I find that patterns of purifying selection on metabolism, at least after 60,000 generations of experimental evolution, are largely idiosyncratic and population-specific.

Results

Overlap between examined gene sets

I first examined the overlap between the four gene sets of interest (Figure 1). If they are completely separate, then running STIMS on each would represent four independent statistical tests. At the other extreme, I would be conducting the same test four times.

Specialist and generalist enzymes are mutually exclusive by definition. While I expected the BiGG core and superessential enzymes to overlap, they are mutually exclusive (Figure 1). Out of 136 BiGG core enzymes, 75 are classified as specialist, while 52 are classified as generalists. The remaining 9 were not classified as either. Out of 96 superessential enzymes, 59 are classified as specialists, and 37 are classified as generalists. Therefore, two pairs of comparisons— BiGG core versus superessential, and specialist versus generalist— are statistically independent.

Purifying selection on BiGG core metabolic enzymes

Out of the 6 hypermutator populations, 3 show significant purifying selection on the BiGG core metabolic genes: Ara−1 (STIMS bootstrap: $p = 0.05$), Ara+3 (STIMS bootstrap: $p = 0.035$), and Ara+6 (STIMS bootstrap: $p = 0.003$). In addition, Ara−4 trends toward purifying selection (STIMS bootstrap: $p = 0.066$). Therefore, the BiGG core metabolic genes show purifying selection in some—but not all—hypermutator LTEE populations (Figure 2A). No trend is seen in the nonmutator LTEE populations (Supplementary Figure 1A).

Idiosyncratic purifying selection on superessential metabolic genes

Only two populations show evidence of purifying selection on superessential metabolic genes: Ara−1 (STIMS bootstrap: $p < 0.001$) and Ara+6 (STIMS bootstrap: $p = 0.01$). The strength of purifying selection on superessential genes therefore seems to be population-specific (Figure 2B). No trend is seen in the nonmutator LTEE populations (Supplementary Figure 1B).

Specialist enzymes tend to evolve under stronger purifying selection than generalist enzymes

Two hypermutator populations show evidence of purifying selection on specialist enzymes (Figure 3A): Ara−4 (STIMS bootstrap: $p = 0.049$) and Ara+6 (STIMS bootstrap: $p < 0.001$). Two nonmutator populations are significantly depleted of mutations in specialist enzymes (Supplementary Figure S2A): Ara+1 (STIMS bootstrap: $p = 0.01$) and Ara+2 (STIMS bootstrap: $p = 0.035$). It is unclear whether the patterns in the nonmutator populations reflect relaxed or purifying selection (Grant, Maddamsetti, et al. 2020; Maddamsetti and Grant 2020b).

Ara−1 (STIMS bootstrap: $p = 0.001$) is the only population in which I found evidence of purifying selection on generalist enzymes (Figure 3B, Supplementary Figure S2B). One nonmutator population, Ara+5, shows evidence of positive selection on generalist enzymes (STIMS bootstrap: $p = 0.0277$). Altogether, these results somewhat support the hypothesis that specialist enzymes evolve under stronger purifying selection than generalist enzymes, with Ara−1 serving as a notable exception.

Discussion

I found evidence of purifying selection on metabolic enzymes only in specific populations. This finding indicates that the strength of purifying selection on metabolism in each population depends on its particular evolutionary history during the LTEE.

Two hypermutator populations may be considered as outliers for the purpose of this paper. First, Ara−2 lost its hypermutator phenotype following the fixation of a reversion mutation at 42,250 generations (Maddamsetti and Grant 2020a). It is challenging to detect a cumulative signal of purifying selection in later generations of this population owing to the lower mutation rate (Maddamsetti and Grant 2020b). Second, Ara−3 evolved the ability to use citrate as a carbon source at ~31,000 generations (Blount, et al. 2008; Blount, et al. 2012). Its phenotypic evolution is now on a different trajectory from the other populations (Blount, et al. 2018; Blount, et al. 2020; Grant, Magid, et al. 2020) .

Overall, Ara+6 shows the strongest signal of purifying selection on metabolic enzymes. This finding is consistent with previous results that show strong purifying selection on aerobic and anaerobic-specific genes in Ara+6 (Grant, Maddamsetti, et al. 2020) and strong purifying selection on essential genes in Ara+6 (Maddamsetti and Grant 2020b). Ara+6 has accumulated the most mutations over the course of the LTEE, which may explain its strong signal of purifying selection.

Mutation accumulation does not seem to be sufficient to explain the varying signal of purifying selection across populations. If mutation accumulation were a dominant factor across populations, then the populations with the most mutations should show the strongest signal of purifying selection on metabolic enzymes. However, like Ara+6, Ara+3 has also accumulated many mutations (Tenaillon, et al. 2016; Couce, et al. 2017; Maddamsetti and Grant 2020a), and yet it shows much weaker evidence of purifying selection on metabolic enzymes. Moreover, Ara−1 shows by far the strongest signal of purifying selection on superessential genes, even though it evolved a hypermutator phenotype much later than Ara+3 and Ara+6 (Barrick and Lenski 2009; Maddamsetti, et al. 2015; Tenaillon, et al. 2016; Good, et al. 2017; Maddamsetti and Grant 2020a).

What may account for this idiosyncratic variation? One possibility is that the strength of purifying selection on metabolic enzymes depends on *how* each population has lost metabolic

function over time (Cooper and Lenski 2000; Leiby and Marx 2014). Ara+6 in particular has lost the ability to grow on many sugars that the ancestral LTEE clone can use (Leiby and Marx 2014). Those losses of function might have caused stronger purifying selection on the metabolic pathways that have remained intact. One way to conceptualize this hypothesis is as follows: at first, many metabolic pathways may evolve nearly neutrally, such that many loss-of-function mutations accumulate, especially in the hypermutator LTEE populations. Eventually, the accumulated effects of these losses of function reach a critical point, after which further losses of metabolic function become deleterious. This critical point corresponds to the percolation threshold at which a metabolic network breaks into multiple disconnected components (Smart, et al. 2008; Pearcy, et al. 2016). The time required to reach this critical threshold may depend on the particular metabolic losses of function that happened to occur on the line of descent of each evolving population. In this conceptual model, idiosyncratic differences in the strength of purifying selection on metabolic enzymes across populations reflects idiosyncratic changes in the pathways that metabolic flux can follow in each population.

This conceptual model (which may be playfully called the "Jenga model") could generalize to any phenotype that is mechanistically caused by a molecular network with many redundancies (Albergante, et al. 2014). Such a network may have many nodes or edges that can be lost without any loss of function. At some point, however, the network has lost most of its redundancies, causing strong purifying selection against further losses of function.

The Jenga model is closely related to the hypothesis that interconnections between aerobic and anaerobic metabolism prevent the loss of anaerobic function despite evolving under relaxed selection in LTEE (Grant, Maddamsetti, et al. 2020). The main difference between the buttressing pleiotropy model proposed by Grant et al. (2020) and the model considered here is that I propose that purifying selection on metabolic enzymes may only become evident after passing a critical threshold of network evolution. Future research could study the plausibility of this model using evolution experiments with organisms with minimal genomes, including endosymbionts and engineered organisms (Gil, et al. 2002; Hutchison, et al. 2016). Specifically, organisms with minimal genomes should show strong purifying selection against losses of metabolic function, while adding redundancies should weaken purifying selection.

My hypothesis makes a specific prediction with regard to the LTEE. The function of many metabolic genes, or their regulation, may already be knocked out in many populations

given the extensive losses of metabolic function that have been observed (Leiby and Marx 2014). In most cases, the genetic causes of these losses of function are unknown. If idiosyncratic purifying selection is caused by the progressive loss of metabolic function, then one should be able to predict which enzymes have come under the strongest purifying selection in the network, based on which enzymes have gained the most metabolic flux. Testing this prediction will require metabolomics experiments to describe how metabolic flux has evolved in the LTEE, and targeted knockout experiments to see whether evolved bottlenecks in each population's metabolic network can account for idiosyncratic purifying selection on metabolic enzymes in the LTEE.

Materials and Methods

Data sources

Pre-processed LTEE metagenomic data were downloaded from: https://github.com/benjaminhgood/LTEE-metagenomic. Genes in the BiGG *E. coli* core metabolic model were downloaded from: http://bigg.ucsd.edu/models/e_coli_core. *E. coli* genes encoding superessential metabolic reactions were taken from supplementary table S6 of Barve et al. (2014). *E. coli* genes encoding specialist and generalist metabolic enzymes were taken from supplementary table S1 of Nam et al. (2012). These tables were then merged with NCBI Genbank gene annotation for the ancestral LTEE strain, *Escherichia coli* B str. REL606. The UpSet visualization for the overlap between the four sets of metabolic enzymes was produced using the UpSetR R Package (Conway, et al. 2017).

Data analysis with STIMS

STIMS is fully described in Maddamsetti and Grant (2020). Briefly, STIMS counts the cumulative number of mutations occurring over time in a gene set of interest, and compares that number to a null distribution that is constructed by subsampling random gene sets of equivalent size. The number of observed mutations in a gene set is normalized by the total length of that gene set in nucleotides. Bootstrapped *p*-values were calculated separately for each population. *p*-

values for one-sided tests for purifying selection were calculated as the lower-tail probability, assuming the null distribution, of sampling a normalized cumulative number of mutations that is less than the normalized cumulative number of mutations in the gene set of interest.

In the visualizations shown in the figures, the top 2.5% and bottom 2.5% of points in the null distribution are omitted, such that each panel can be interpreted as a two-sided randomization test with a false-positive (type I error) probability $\alpha = 0.05$. The variance in the bootstrapped null distributions depends on the cardinality of the gene set of interest: larger gene sets have smaller sampling variance, and smaller gene sets have larger sampling variance.

Data Availability Statement

The data and analysis codes underlying this article are available on the Dryad Digital Repository (DOI pending publication). Analysis codes are also available at: https://github.com/rohanmaddamsetti/LTEE-network-analysis.

Acknowledgements

I thank Richard Lenski, Jeffrey Barrick, and Lingchong You for guidance and advice, Nkrumah Grant for valuable discussions, and Zachary Blount for detailed comments and feedback. The LTEE is supported by a grant from the National Science Foundation (currently DEB-1951307) to Richard Lenski and Jeffrey Barrick.

References

Albergante L, Blow JJ, Newman TJ. 2014. Buffered Qualitative Stability explains the robustness and evolvability of transcriptional networks. eLife 3:e02863.

Barrick JE, Lenski RE editors. Cold Spring Harbor symposia on quantitative biology. 2009.

Barve A, Rodrigues JFM, Wagner A. 2012. Superessential reactions in metabolic networks. Proceedings of the National Academy of Sciences 109:E1121-E1130.

Blount ZD, Barrick JE, Davidson CJ, Lenski RE. 2012. Genomic analysis of a key innovation in an experimental Escherichia coli population. Nature 489:513-518.

Blount ZD, Borland CZ, Lenski RE. 2008. Historical contingency and the evolution of a key innovation in an experimental population of Escherichia coli. Proceedings of the National Academy of Sciences 105:7899-7906.

Blount ZD, Lenski RE, Losos JB. 2018. Contingency and determinism in evolution: Replaying life's tape. Science 362.

Blount ZD, Maddamsetti R, Grant NA, Ahmed ST, Jagdish T, Baxter JA, Sommerfeld BA, Tillman A, Moore J, Slonczewski JL, et al. 2020. Genomic and phenotypic evolution of Escherichia coli in a novel citrate-only resource environment. eLife 9:e55414.

Conway JR, Lex A, Gehlenborg N. 2017. UpSetR: an R package for the visualization of intersecting sets and their properties. Bioinformatics 33:2938-2940.

Cooper VS, Lenski RE. 2000. The population genetics of ecological specialization in evolving Escherichia coli populations. Nature 407:736-739.

Couce A, Caudwell LV, Feinauer C, Hindré T, Feugeas J-P, Weigt M, Lenski RE, Schneider D, Tenaillon O. 2017. Mutator genomes decay, despite sustained fitness gains, in a long-term experiment with bacteria. Proceedings of the National Academy of Sciences 114:E9026-E9035.

Gil R, Sabater-Muñoz B, Latorre A, Silva FJ, Moya A. 2002. Extreme genome reduction in Buchnera spp.: toward the minimal genome needed for symbiotic life. Proceedings of the National Academy of Sciences 99:4454-4458.

Good BH, McDonald MJ, Barrick JE, Lenski RE, Desai MM. 2017. The dynamics of molecular evolution over 60,000 generations. Nature 551:45-50.

Grant NA, Maddamsetti R, Lenski RE. 2020. Maintenance of metabolic plasticity despite relaxed selection in a long-term evolution experiment with Escherichia coli. bioRxiv.

Grant NA, Magid AA, Franklin J, Dufour YS, Lenski R. 2020. Changes in Cell Size and Shape During 50,000 Generations of Experimental Evolution with Escherichia coli. bioRxiv.

Hutchison CA, Chuang R-Y, Noskov VN, Assad-Garcia N, Deerinck TJ, Ellisman MH, Gill J, Kannan K, Karas BJ, Ma L. 2016. Design and synthesis of a minimal bacterial genome. Science 351.

King ZA, Lu J, Dräger A, Miller P, Federowicz S, Lerman JA, Ebrahim A, Palsson BO, Lewis NE. 2015. BiGG Models: A platform for integrating, standardizing and sharing genome-scale models. Nucleic Acids Research 44:D515-D522.

Leiby N, Marx CJ. 2014. Metabolic erosion primarily through mutation accumulation, and not tradeoffs, drives limited evolution of substrate specificity in Escherichia coli. PLoS Biol 12:e1001789.

Lenski RE. 2017. Experimental evolution and the dynamics of adaptation and genome evolution in microbial populations. The ISME journal 11:2181-2194.

Lenski RE, Rose MR, Simpson SC, Tadler SC. 1991. Long-term experimental evolution in Escherichia coli. I. Adaptation and divergence during 2,000 generations. The American Naturalist 138:1315-1341.

Maddamsetti R. 2020. Universal constraints on protein evolution in the long-term evolution experiment with Escherichia coli. bioRxiv.

Maddamsetti R, Grant NA. 2020a. Divergent Evolution of Mutation Rates and Biases in the Long-Term Evolution Experiment with Escherichia coli. Genome biology and evolution 12:1591-1603.

Maddamsetti R, Grant NA. 2020b. A simple test to infer mode of selection in metagenomics time series of evolving asexual populations. bioRxiv.

Maddamsetti R, Lenski RE, Barrick JE. 2015. Adaptation, clonal interference, and frequency-dependent interactions in a long-term evolution experiment with Escherichia coli. Genetics 200:619-631.

Nam H, Lewis NE, Lerman JA, Lee D-H, Chang RL, Kim D, Palsson BO. 2012. Network context and selection in the evolution to enzyme specificity. Science 337:1101-1104.

Pearcy N, Chuzhanova N, Crofts JJ. 2016. Complexity and robustness in hypernetwork models of metabolism. Journal of theoretical biology 406:99-104.

Smart AG, Amaral LA, Ottino JM. 2008. Cascading failure and robustness in metabolic networks. Proceedings of the National Academy of Sciences 105:13223-13228.

Sniegowski PD, Gerrish PJ, Lenski RE. 1997. Evolution of high mutation rates in experimental populations of E. coli. Nature 387:703-705.

Tenaillon O, Barrick JE, Ribeck N, Deatherage DE, Blanchard JL, Dasgupta A, Wu GC, Wielgoss S, Cruveiller S, Médigue C. 2016. Tempo and mode of genome evolution in a 50,000-generation experiment. Nature 536:165-170.

Wiser MJ, Ribeck N, Lenski RE. 2013. Long-term dynamics of adaptation in asexual populations. Science 342:1364-1367.

Tables and Figures

Figure 1. Intersections among the sets of metabolic enzymes analyzed in this work. The intersections between four pre-defined sets: superessential enzymes, BiGG core metabolic enzymes, specialist enzymes, and generalist enzymes were examined (see text for further details about these gene sets). The cardinality of each of these sets is shown in the horizontal bar graph on the left. Every non-empty subset of the four sets is visualized as a set of dots connected by lines. The cardinality for each non-empty subset (set intersections) is shown in the vertical bar graph.

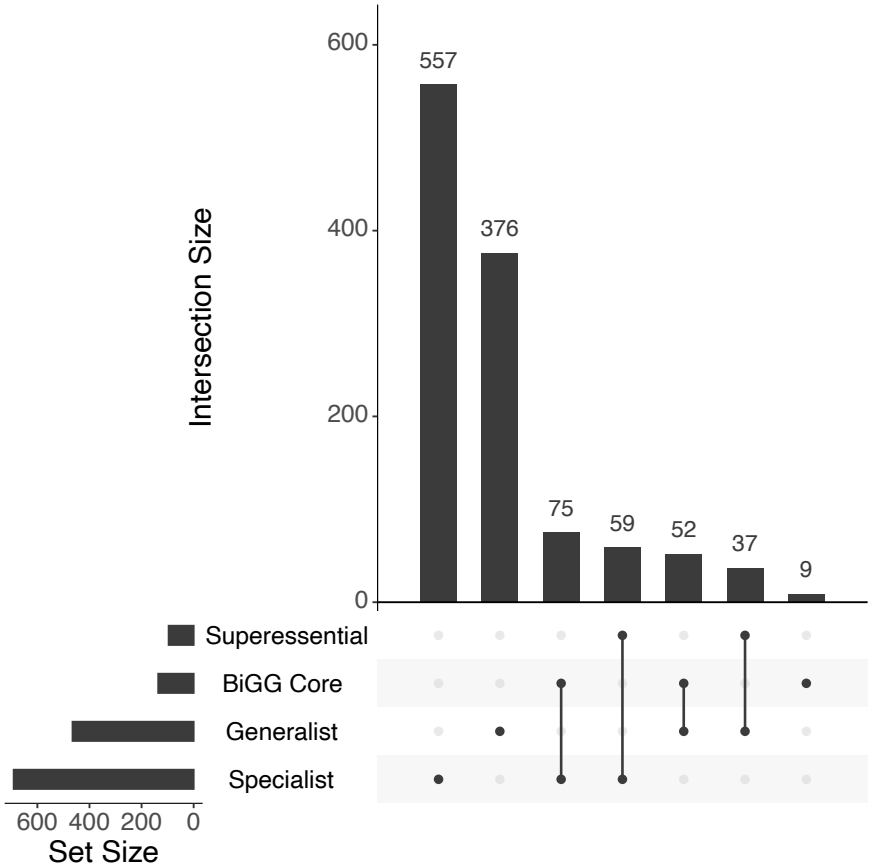

Figure 2. The results of running STIMS on BiGG core enzymes and superessential enzymes (hypermutator populations only). Each panel shows the cumulative number of mutations in the gene set of interest (solid line) in the six hypermutator LTEE populations. For comparison, random sets of genes (with the same cardinality as the gene set of interest) were sampled 1000 times, and the cumulative number of mutations in those random gene sets, normalized by gene length, were calculated. The middle 95% of this null distribution is shown as shaded points. When a solid line falls below the shaded region, then the gene set of interest show a significant signal of purifying selection. A) The results of running STIMS on enzymes in the BiGG *E. coli* core metabolism model. B) The results of running STIMS on enzymes catalyzing superessential metabolic reactions.

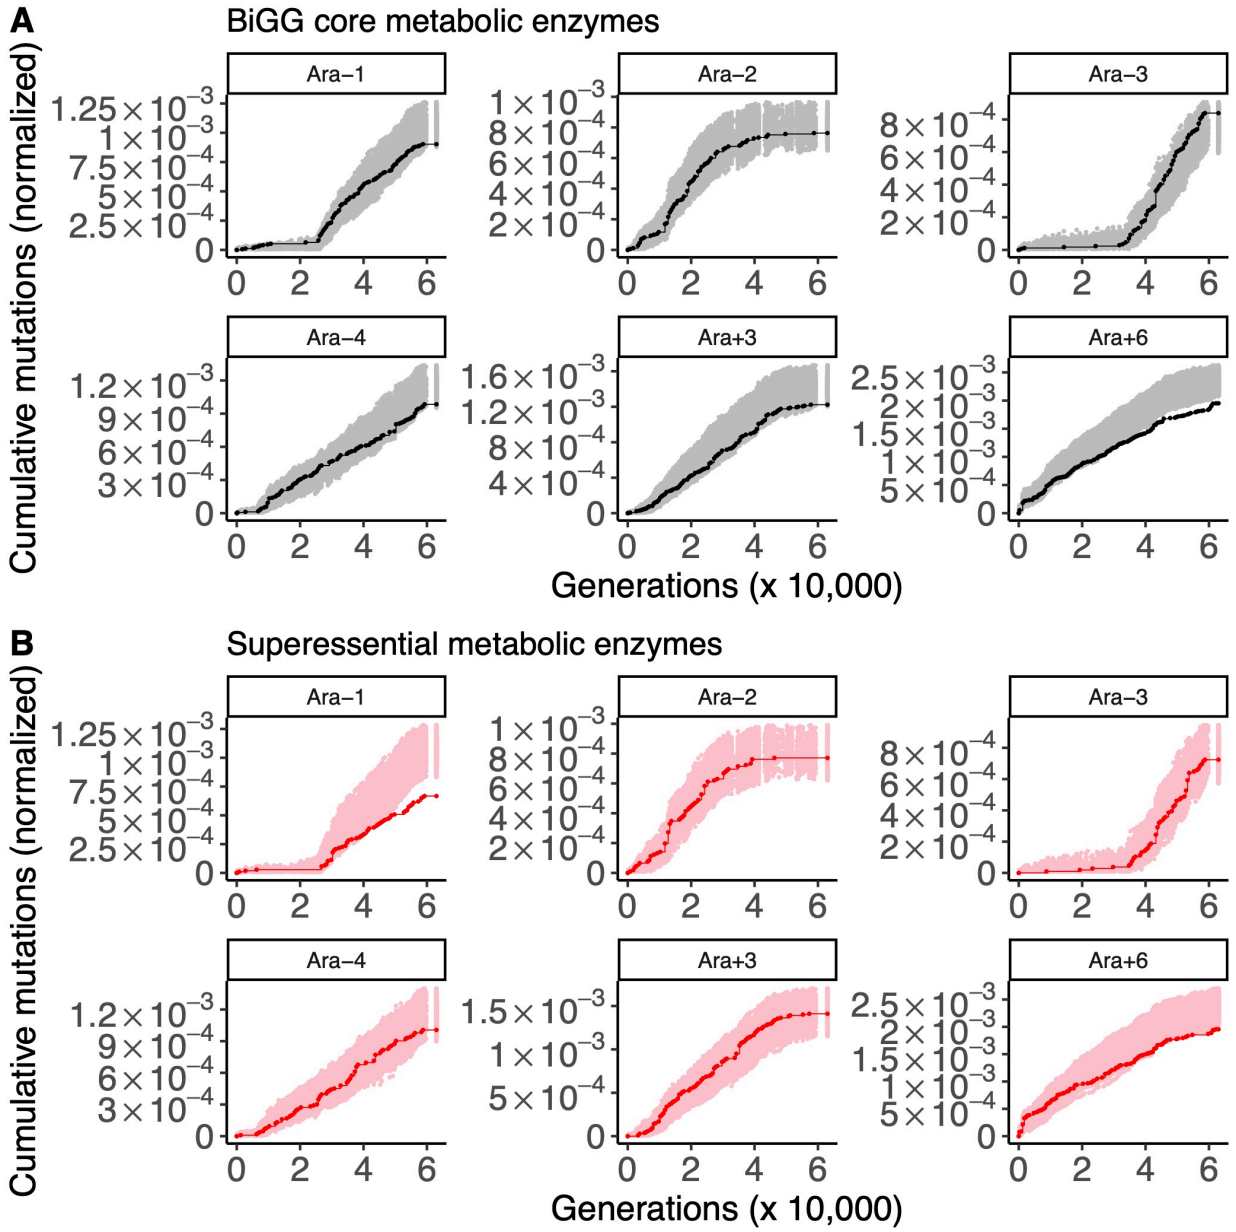

Figure 3. **The results of running STIMS on specialist and generalist metabolic enzymes (hypermutator populations only)**. Each panel shows the cumulative number of mutations in the gene set of interest (solid line) in the six hypermutator LTEE populations. For comparison, random sets of genes (with the same cardinality as the gene set of interest) were sampled 1000 times, and the cumulative number of mutations in those random gene sets, normalized by gene length, were calculated. The middle 95% of this null distribution is shown as shaded points. When a solid line falls below the shaded region, then the gene set of interest show a significant signal of purifying selection. A) The results of running STIMS on specialist enzymes B) The results of running STIMS on generalist enzymes.

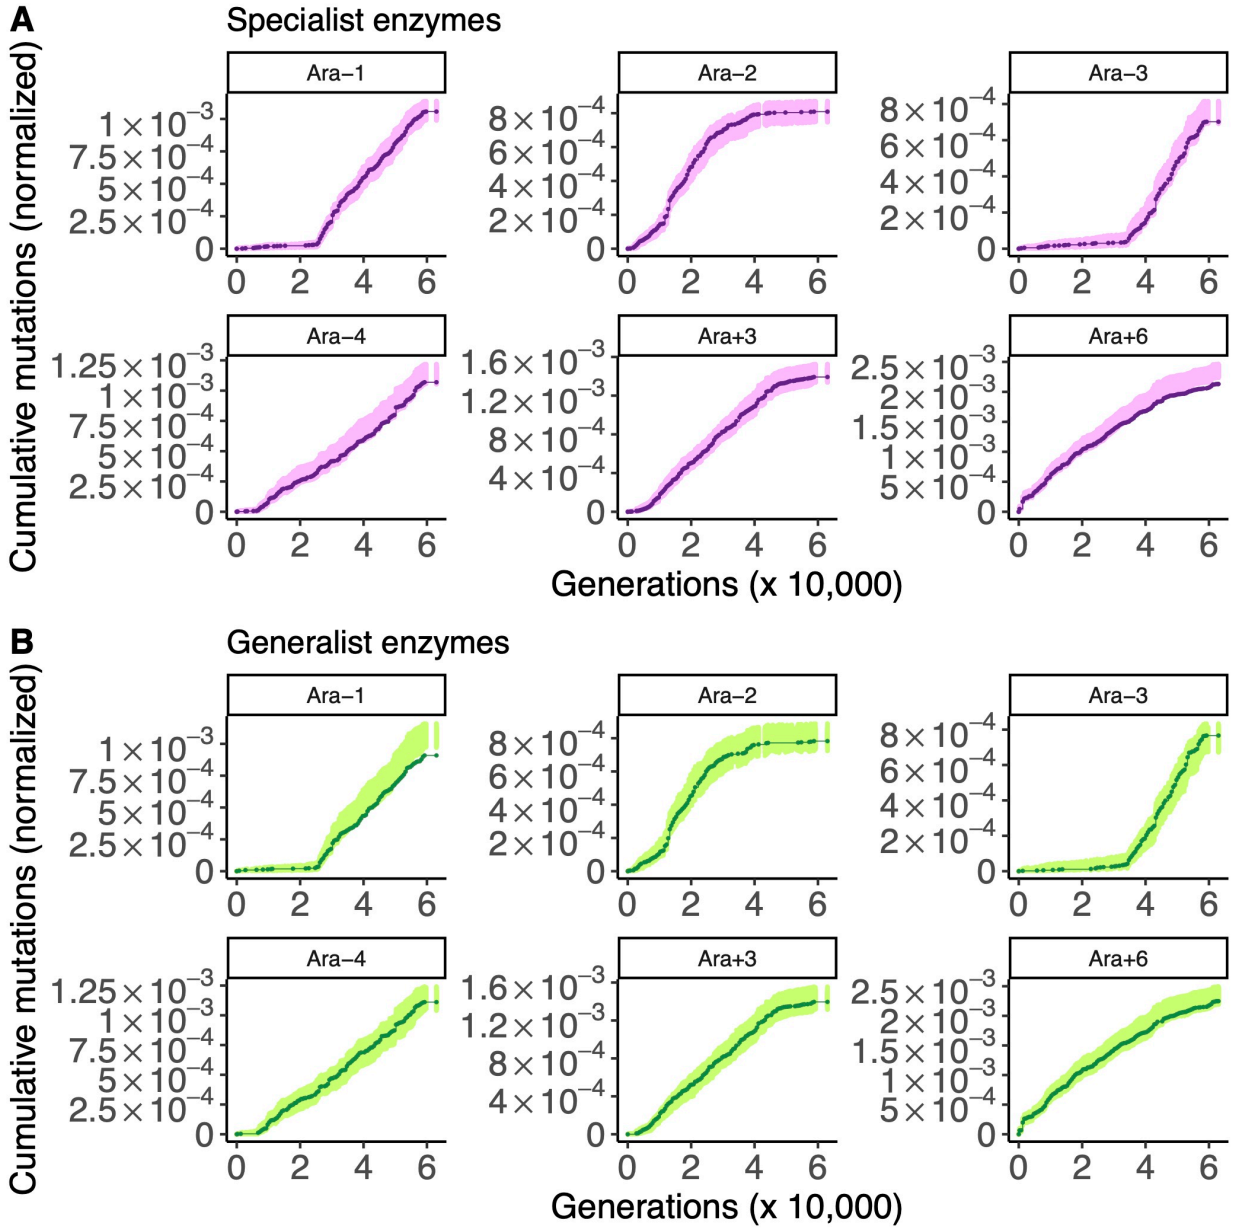

SUPPLEMENTARY INFORMATION

Supplementary File: Mapping of REL606 genes into the four gene sets analyzed in this work (TableS1.csv).

Supplementary Figure 1. The results of running STIMS on BiGG core enzymes and superessential enzymes (nonmutator populations only). Each panel shows the cumulative number of mutations in the gene set of interest (solid line) in the six nonmutator LTEE populations. For comparison, random sets of genes (with the same cardinality as the gene set of interest) were sampled 1000 times, and the cumulative number of mutations in those random gene sets, normalized by gene length, were calculated. The middle 95% of this null distribution is shown as shaded points. When a solid line falls above the shaded region, then the gene set of interest show a significant signal of positive selection. A) The results of running STIMS on enzymes in the BiGG *E. coli* core metabolism model. B) The results of running STIMS on enzymes catalyzing superessential metabolic reactions.

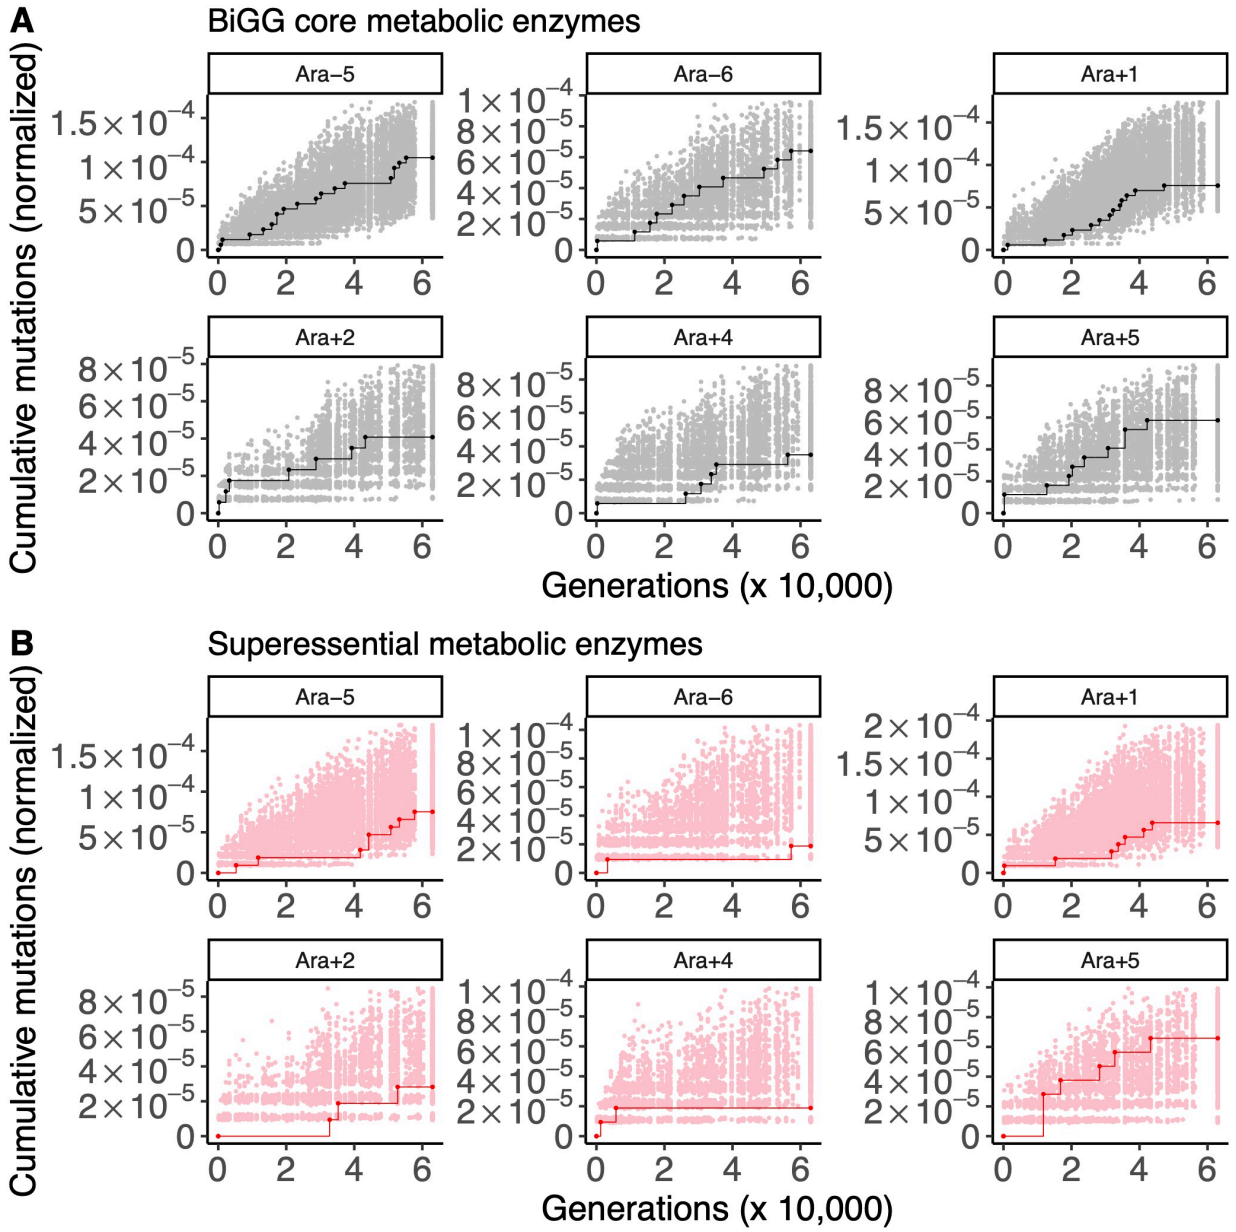
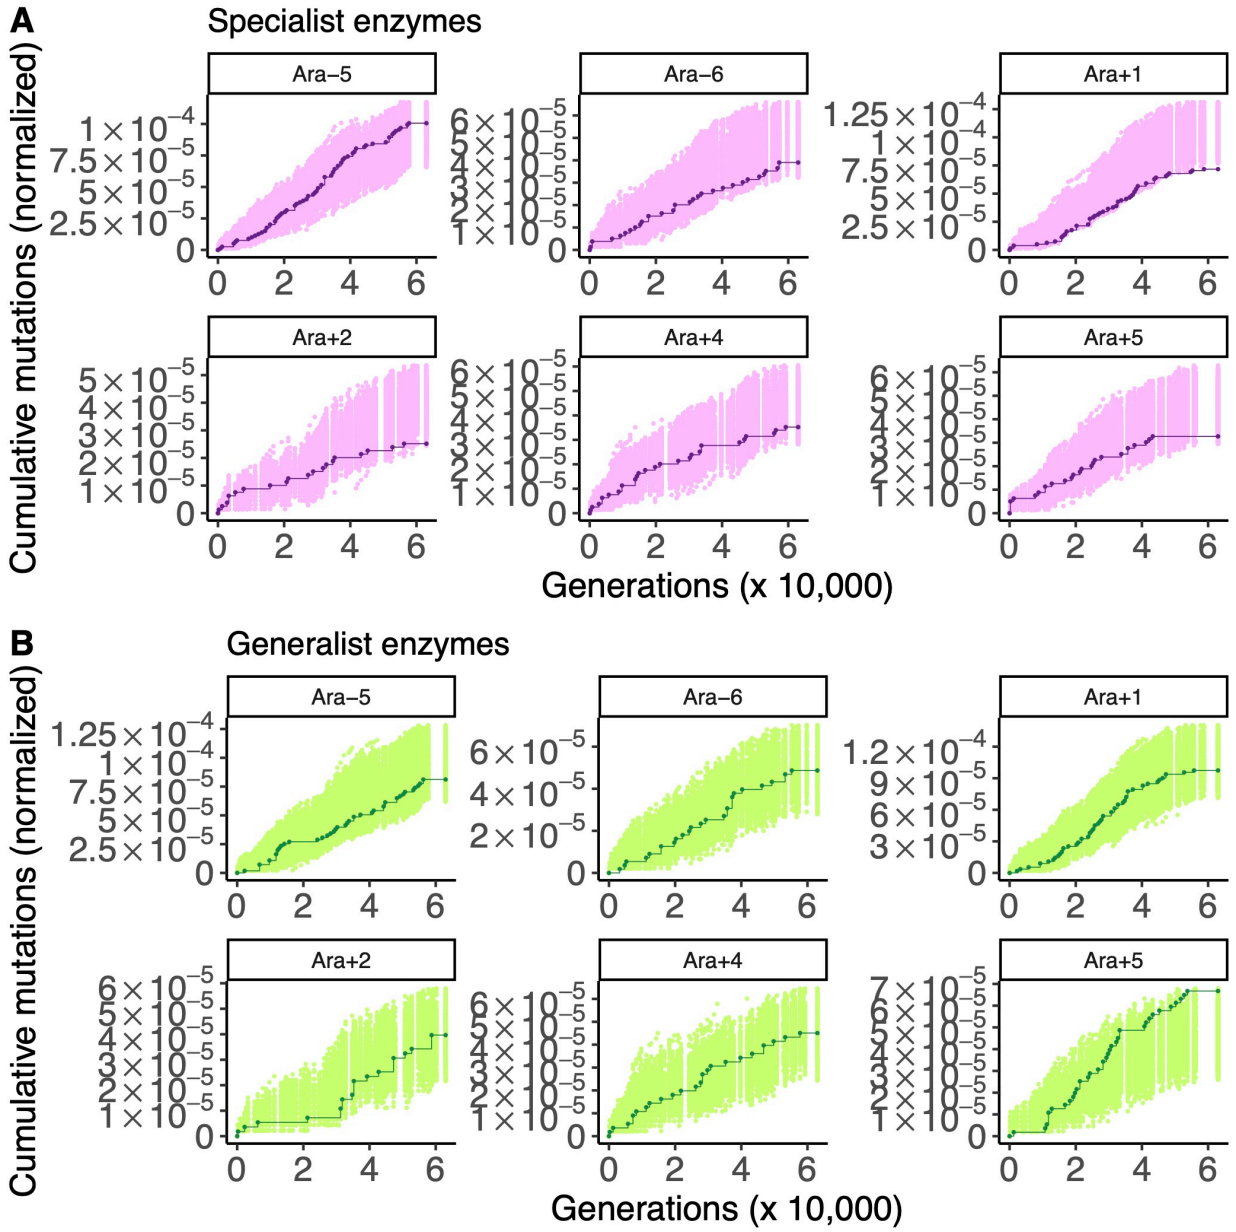
Supplementary Figure 2. The results of running STIMS on specialist enzymes and generalist enzymes (nonmutator populations only). Each panel shows the cumulative number of mutations in the gene set of interest (solid line) in the six nonmutator LTEE populations. For comparison, random sets of genes (with the same cardinality as the gene set of interest) were sampled 1000 times, and the cumulative number of mutations in those random gene sets, normalized by gene length, were calculated. The middle 95% of this null distribution is shown as shaded points. When a solid line falls above the shaded region, then the gene set of interest show a significant signal of positive selection. A) The results of running STIMS on specialist enzymes. B) The results of running STIMS on generalist enzymes.