Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

From Possibility to Probability: Population-Genetic Accessibility of Enzyme Specialization After Gene Duplication Original Article

Jordan A. Dowell

jdowell@lsu.edu

Louisiana State University

Research Article

Keywords: catalytic frustration, enzyme promiscuity, evolutionary accessibility, gene dosage, gene duplication, SLiM

Posted Date: November 16th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11599777/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Promiscuous enzymes catalyze more than one reaction, creating opportunities for new functions but also trade-offs when improving one reaction reduces another. Gene duplication may resolve this conflict by allowing the two copies to specialize, but only if the new copy survives long enough to diverge. We combined a constraint-based model of enzyme activity with forward-time simulations in SLiM to separate what happens after a duplicate is retained from whether a new duplicate can reach that state. Under dosage compensation, persistent specialization evolved in 56 of 300 populations in which duplication was imposed, but in none of 5,700 simulations that began with one new duplicate. Including additive dosage increased establishment from 51 to 2,421 of 5,700 origins, increased near-fixation from 1 to 2,421, and produced 102 persistent specialists. However, dosage did not change specialization after population-wide retention: 61 of 1,140 compensated populations and 60 of 1,140 additive populations specialized. Once every individual carried the duplicate, the shared dosage benefit no longer changed which genotypes were favored. All specialists arising from a single origin occurred under strong catalytic antagonism and high mutation supply. Thus, dosage primarily determined whether a new duplicate survived and spread, whereas catalytic conflict and mutation supply determined whether retained copies later specialized.

Introduction

Across the tree of life, metabolism generates extensive chemical diversity that connects cellular function to organismal interactions and ecosystem processes. Specialized metabolites are particularly diverse in plants and fungi, where lineage-specific compounds mediate defense, communication, competition, and environmental responses (Weng et al. 2012; Keller 2019). This diversity ultimately depends on enzymes, yet many enzymes are not restricted to a single reaction. Many catalyze weak secondary reactions in addition to their primary activity (also known as enzyme promiscuity), providing a standing source of biochemical variation that can be impacted by mutation, altered expression, or environmental change (Aharoni et al. 2005; Khersonsky and Tawfik 2010; Copley 2017). Two reactions catalyzed by the same active site may not improve independently, so a mutation that increases one activity may decrease the other and produce antagonistic pleiotropy, adaptive conflict, or catalytic frustration (Gould and Tawfik 2005; Tokuriki et al. 2012; Ferreiro et al. 2014; Schmutzer et al. 2024). Thus, while enzyme promiscuity creates opportunity, it can also create constraint.

Gene duplication provides one route through this constraint. Additional copies may be lost, partition ancestral functions, acquire new functions, or be retained because additional product is advantageous (Ohno 1970; Force et al. 1999; Lynch and Conery 2000; Lynch and Force 2000; Conant and Wolfe 2008). Under escape from adaptive conflict, a multifunctional ancestor is constrained by opposing selective demands, and duplication permits each copy to improve a different activity (Des Marais and Rausher 2008; Sikosek et al. 2012). Among major hypotheses explaining the origins of diverse enzymes, Innovation–Amplification–Divergence describes a related sequence in which selection first acts on a weak secondary activity, amplification increases dosage and mutational target size, and later mutations

refine function (Francino 2005; Bergthorsson et al. 2007; Näsvall et al. 2012). Both frameworks explain why two copies can support specialization. However, neither a favorable endpoint nor an available mutational path guarantees that a newly arisen copy will persist.

This conflict raises a key question: if selection maintains the performance of an adapted multifunctional enzyme, how does a population reach a two-copy state in which its catalytic functions can separate? Further, while the existence of a weak secondary activity establishes biochemical potential, it does not establish that a population can retain the genetic state needed to realize that potential. For example, a duplicate initially occurs in one genome and is therefore exposed to stochastic loss. Even when later two-copy genotypes are superior, the originating copy may be neutral, deleterious, or beneficial only after additional mutations. Population-genetic models have therefore focused on preserving new duplicates, the waiting time to subfunctionalization, and the many unsuccessful origins that can precede stable incorporation (Lynch et al. 2001; Kondrashov and Kondrashov 2006; Proulx 2012). This distinction separates two questions that are often combined: what can evolve after two copies are retained, and how often does a population reach that retained state?

When enzyme abundance limits flux, an extra expressed copy can provide an immediate benefit and shift a rare duplicate from drift-dominated loss toward establishment (Kondrashov and Kondrashov 2006; Andersson and Hughes 2009). Finite-population models have integrated positive dosage, multifunctionality, duplication, and specialization (Sikosek et al. 2012), and experimental evolution supports amplification as a way to preserve weak activity while later mutations improve it (Näsvall et al. 2012). However, dosage effects depend on biological context. Increased product may be attenuated by expression feedback, constrained by stoichiometric balance, limited by saturation, or made deleterious through toxicity and metabolic cost (Qian and Zhang 2008; Lan and Pritchard 2016; Ascencio et al. 2021). Complete dosage compensation and copy-number-additive expression therefore represent useful mechanistic bounds rather than universal descriptions of duplicated genes.

Recent empirical and theoretical studies further show why retention and divergence should be separated. Mihajlovic et al. (2025) found that two-copy fluorescent-protein systems accumulated greater mutational diversity and robustness but did not evolve phenotypes faster or specialize more frequently, and one copy was often inactivated. While, Schmutzer et al. (2024) constructed feasible spaces for enzymes catalyzing multiple reactions from five classes of pairwise constraints. Their optimization framework compared a multifunctional enzyme with subfunctionalized copies selected on different reactions and identified when specialization could relieve, or fail to relieve, catalytic frustration. However, both approaches largely begin after the critical population-genetic transition has occurred. Importantly, a duplicated system can contain evolutionary opportunities that remain unavailable to a new copy if that copy is lost while rare.

Here, we separate the origin of a duplicate from its evolution after retention using explicit diploid populations. We ask (1) whether catalytic frustration favors specialization once two copies are present, (2) whether a single new copy can survive long enough to access that outcome, and (3) whether an

immediate dosage benefit changes these probabilities. We predict that strong antagonism will favor specialization after retention, that most dosage-compensated origins will be lost, and that additive dosage will act primarily during establishment rather than after population-wide retention. Using matched ancestral populations, catalytic treatments, and random seeds, we extend Schmutzer et al. (2024) 's feasible-space framework to quantify when biochemical possibility becomes population-genetically accessible.

Materials and Methods

Catalytic phenotype and feasible spaces

Each gene copy encoded two continuous catalytic activities, x_1 and x_2, representing two reactions required by the organismal performance function and weighted symmetrically in that function. The two activities therefore represent jointly required reactions rather than a primary reaction paired with a selectively neutral side activity, allowing the model to test whether duplication resolves selection acting simultaneously on both functions. All allelic activities were bounded by $0 \le x_1 \le 1$ and $0 \le x_2 \le 1$. Five catalytic relationships described how the two activities could vary. Strong antagonism was the union of $x_1 + 4x_2 \le 1$ and $4x_1 + x_2 \le 1$. Weak antagonism required $x_1 + 4x_2 \le 4$ and $4x_1 + x_2 \le 4$. The unconstrained relationship imposed only the unit box. Weak synergism required $-x_1 + 5x_2 \le 4$ and $5x_1 - x_2 \le 4$. Strong synergism required $x_2 \le 3x_1$, $x_1 \le 3x_2$, $2x_2 - x_1 \le 1$, and $2x_1 - x_2 \le 1$ (Online Resource 1, Supplementary Table S1). This two-activity design isolates one pairwise slice of the multi-reaction framework of Schmutzer et al. (2024) and adapts their five pairwise constraint classes to a stochastic population-genetic model.

For every inequality $a \cdot x \le c$, boundary violation (D) was $\frac{[a \cdot x - c]_+}{||a||_1}$. Intersections used the maximum normalized violation, whereas the disjunctive strong-antagonism space used the minimum violation to either feasible branch. Unit-box and relationship violations were then combined by their maximum. Individual violation D was the maximum across all active allelic copies. Under the hard boundary, any $D > 10^{-9}$ received a reproductive-weight floor of 10^{-6}. Under the soft boundary, biochemical performance was multiplied by $P = \exp(-100D^2)$. The soft rule represents a gradual penalty outside the declared catalytic space; the hard rule represents an effectively absolute biochemical boundary. All five relationships were evaluated under the hard rule. We also evaluated strong antagonism, the unconstrained relationship, and strong synergism under the soft rule.

Population-genetic model, mutation, and fitness

Simulations were implemented in SLiM 4.0.1 as diploid, panmictic Wright–Fisher populations of N = 1,000 individuals (Haller and Messer 2019, 2023). Each catalytic copy contained 1,000 mutable positions. Before duplication, one copy was active on each haploid genome; after duplication, the second 1,000-position copy became active only on genomes carrying the duplication marker. Per-position

mutation rates were 10^{-7} or 10^{-6} per generation, corresponding to per-copy catalytic-mutation rates of 10^{-4} or 10^{-3}. A duplicated chromosome therefore carried twice the catalytic mutational target of an unduplicated chromosome. Recombination was 10^{-8} per position per generation. Each new catalytic mutation m independently received effects δ_1 and δ_2 from $N(0, 0.025^2)$. Activity of allelic copy c for reaction r was $x_{cr} = 0.1 + \sum_{m \in c} \delta_{mr}$. Mutational effects were additive within copies, and allelic copy activities were combined according to the dosage rule below. The SLiM mutation-type values h = 0.5 and s = 0 were bookkeeping values; phenotype and reproductive weight were calculated by the custom catalytic model rather than SLiM's default mutation-fitness mapping.

For an individual with C active allelic copies, dosage-compensated activity was $A_r = \frac{1}{C}\sum_{c=1}^{C} x_{cr}$, whereas copy-number-additive activity was $A_r = \frac{1}{2}\sum_{c=1}^{C} x_{cr}$. Thus, an ancestral diploid had two active allelic copies and a nominal 1-fold dosage; a heterozygous duplication had three and a nominal 1.5-fold dosage; and a homozygous duplication had four and a nominal twofold dosage. Catalytic performance was the harmonic mean $H = \frac{2}{A_1 + A_2}$, normalized to ancestral base activity as $M = \frac{H}{0.1}$. Duplicate cost was fixed at zero. For feasible individuals, reproductive weight was $W = \max(M^{\lambda}P, 10^{-6})$, where λ = 0.5, 1, or 2 was the selection-strength multiplier; a hard-boundary violation instead set $W = 10^{-6}$. The reported pre-selection metabolic-fitness score was $R = \max(MP, 10^{-6})$. When every individual carried a feasible duplicate, additive dosage doubled both activities and M, multiplying W by the common factor 2^{λ}. Because Wright–Fisher sampling uses relative reproductive weights, this common factor canceled among feasible genotypes. The full factorial design contained 48 treatments: 30 hard-boundary treatments (five relationships × two mutation rates × three selection strengths) and 18 soft-boundary treatments (three relationships × two mutation rates × three selection strengths). Random seeds identified ancestral populations and post-duplication trajectories. Analyses and figures were produced in R 4.6.0 (R Core Team 2026).

Ancestral adaptation, stationarity and retained-duplication experiments

Each population first evolved with one gene copy. Burn-in duration was calibrated separately for each treatment because a single duration could stop slowly adapting populations too early while adding unnecessary computation to others. At 10N intervals, the stationarity analysis compared two consecutive N-generation windows for mean metabolic-fitness score, activity 1, activity 2, and promiscuity and estimated the normalized slope of each metric in the final window. A checkpoint passed only when the 90% confidence intervals for all eight quantities fell within a predeclared ± 1% equivalence region. The accepted burn-in was the earliest passing checkpoint that continued to pass at all later checkpoints. The calibration produced treatment-specific burn-ins from 10N to 60N generations. Of 48 treatments, 44 met sustained ensemble stationarity. Four high-mutation, strong-antagonism treatments reached stable mean fitness while catalytic traits continued to move along a nearly neutral ridge and were classified as ridge-evolving rather than extended indefinitely. These treatments were followed to

80N for a paired timing sensitivity. We saved and reused evolved ancestral populations and mutation-effect sidecar files so repeated duplicate origins and dosage treatments began from identical population states.

We tested specialization after removing duplicate loss as a barrier. A duplicate was introduced into every haploid genome, producing four active allelic copies per diploid, and populations were followed for (10N) post-duplication generations. We initially simulated 20 independently adapted ancestral populations for each of 48 treatment combinations, producing 960 populations. Because specialization was concentrated in six treatments combining strong antagonism with mutation rate (10^{-6}), we added 30 populations to each of these treatments to improve the precision of their specialization estimates. The final dataset therefore contained 50 populations in each focal treatment and 20 populations in every other treatment, for a total of 1,140 independent populations: 300 high-mutation strong-antagonism populations, 120 low-mutation strong-antagonism populations, and 720 populations representing the remaining catalytic relationships. To test whether longer observation changed outcomes in ridge-evolving conditions, we extended 80 existing trajectories from (10N) to (20N). These were paired duration extensions rather than additional independent populations, and each longer trajectory exactly reproduced the corresponding first (10N) generations. The later matched dosage analysis used the final treatment-specific ancestral population bank, including ancestors evolved for (80N) in ridge-evolving treatments.

Single-origin experiment

We then tested accessibility from a new duplication event. One duplicate was introduced into one haploid genome, giving initial frequency 1/(2N) = 0.0005. We launched five independently seeded origins from each ancestor. All 48 treatments contained 100 origins, and the six focal high-mutation strong-antagonism treatments were expanded to 250 origins per cell. In total, 5,700 origins were nested within 1,140 independently evolved ancestral populations. Under dosage compensation, the new copy reproduced one realized haplotypic catalytic state but did not automatically double total output. When homologous alleles differed, copying one haplotype could slightly change the carrier's normalized catalytic phenotype. For 4,740 compensated origins, we measured the immediate change in log metabolic-fitness score and classified it as beneficial, neutral, or deleterious using $|\Delta \log R| \leq 10^{-12}$ as the numerical definition of neutrality.

Finally, we repeated both experiments under copy-number-additive activity. Every compensated trajectory was paired with an additive trajectory using the same ancestor, catalytic treatment, and post-duplication random seed. The matched dataset contained 5,700 single-origin pairs and 1,140 retained-duplication pairs. The two modes differed only in how they translated active allelic copies into total catalytic activity.

Outcomes and statistical analysis

We defined duplicate establishment as a frequency of 5%, and operational near-fixation as 95%. Hereafter, near-fixation refers to this threshold rather than allele frequency 1. For carrier i, complementary specialization was $S_i = \frac{1}{2} |(\bar{x}_{i,0,1} - \bar{x}_{i,0,2}) - (\bar{x}_{i,1,1} - \bar{x}_{i,1,2})|$, where $\bar{x}_{i,0,r}$ and $\bar{x}_{i,1,r}$ were the mean activities of the original and duplicated copies across the homologs on which each copy was present. The population carrier-specialization index was the mean S_i across individuals carrying at least one duplicate. Persistent specialization required duplicate frequency ≥ 5% and carrier-specialization index ≥ 0.5 at every 10-generation event check spanning 100 generations.

Independent ancestral populations were the replicate units. We report binomial proportions with 95% Wilson intervals when trajectories were independent and use the Wilson upper bound for zero observed events (Wilson 1927). Repeated origins were first averaged within an ancestor. Standard errors, 95% t intervals, and paired tests therefore describe variation among ancestral populations rather than treating five origins from one ancestor as five independent ancestral replicates. Dosage effects were paired absolute risk differences, additive minus compensated. For single-origin outcomes, the five paired differences were averaged within each of 1,140 ancestors and evaluated with two-sided one-sample t tests. Attempt-level exact McNemar tests were retained only as descriptive sensitivity analyses because origins shared ancestral populations. Retained effects used 1,140 paired trajectories and exact two-sided McNemar tests for discordant outcomes (McNemar 1947).

We also calculated a stage-by-dosage contrast: the additive-minus-compensated specialization effect from a single origin minus the corresponding effect when retention was imposed. Each stage contained 1,140 ancestors; 1,060 were shared, and 80 were stage-specific. A treatment-stratified partial-pair cluster bootstrap resampled shared ancestors jointly and unmatched ancestors independently while preserving treatment sizes (50,000 resamples). Time to specialization was analyzed with duplicate loss as a competing event using the Aalen–Johansen estimator (Aalen and Johansen 1978). Near-fixed nonspecialists remained under observation to the declared endpoint. Global pooled estimates are replicate-weighted summaries of the adaptive simulation design, not estimates averaged over an empirical distribution of catalytic relationships, mutation rates, or selection regimes. OpenAI Codex assisted with simulation-code development, R scripting, and supplemental materials development. The author reviewed and validated the simulation code and analyses. The SLiM model and portable execution materials used for ancestral adaptation, retained duplication, single-origin accessibility, and matched dosage experiments are provided in Online Resource 2.

Results

Specialization after duplicate retention

Persistent specialization after imposed retention occurred only where catalytic conflict and mutation supply were both high. Under strong antagonism and mutation rate 10^{-6}, 56 of 300 independently adapted compensated populations specialized (18.7%; 95% Wilson interval, 14.7–23.5%; Table 1;

Fig. 1A). Individual parameter combination estimates ranged from 14% to 28%, with the largest point estimate under the soft boundary and 2× selection (14/50; 28.0%; 95% Wilson interval, 17.5–41.7%).

Table 1
Persistent specialization after imposed retention in the six high-mutation, strong-antagonism treatments. Each row contains 50 independently adapted ancestral populations.

Boundary	Selection multiplier	Persistent specialization	Estimate	95% Wilson interval
Hard	0.5	7/50	14%	7.0–26.2%
Hard	1.0	9/50	18%	9.8–30.8%
Hard	2.0	9/50	18%	9.8–30.8%
Soft	0.5	8/50	16%	8.3–28.5%
Soft	1.0	9/50	18%	9.8–30.8%
Soft	2.0	14/50	28%	17.5–41.7%

No specialization occurred in 120 strong-antagonism populations at mutation rate 10^{-7} (descriptive 95% Wilson upper bound, 3.10%) or in 720 populations representing weak antagonism, unconstrained activity, weak synergism, or strong synergism (upper bound, 0.531%). High-mutation strong antagonism differed from the pooled non-antagonistic relationships (Fisher's exact $p < 0.001$) and from low-mutation strong antagonism ($p < 0.001$). Specialization was therefore a restricted response to strong catalytic conflict and high mutation supply, not a general consequence of retaining a duplicate. In contrast, boundary type did not produce a resolved pooled difference. Specialization occurred in 25/150 hard-boundary populations and 31/150 soft-boundary populations (Fisher's exact $p = 0.459$). Although the soft-boundary, 2×-selection condition had the largest estimate, the overlapping intervals do not support a general boundary effect.

Accessibility from a single compensated origin

The same catalytic opportunity was rarely reached when a compensated duplicate began in one genome (Fig. 1B). Only 51/5,700 origins established (0.90%; ancestor-clustered SE, 0.13%; 95% t interval, 0.65–1.14%). One reached near-fixation (0.018%), and none persistently specialized. The attempt-level Wilson interval for near-fixation was 0.003–0.10%, and the zero-event specialization upper bound was 0.067%. Because five origins shared each ancestor, these latter quantities are descriptive attempt-level bounds rather than cluster-robust intervals. Thus, the catalytic landscape that supported specialization after retention was effectively inaccessible from a rare compensated copy. In addition, establishment differed little across the major treatment classes: 0.90% under hard and 0.90% under soft boundaries; 0.92% at mutation rate 10^{-7} and 0.88% at 10^{-6}; and 0.84%, 0.79%, and 1.05% at selection multipliers 0.5, 1, and 2. Relationship-level estimates ranged from 0.67% to 1.00%. At this stage, rarity and drift outweighed the long-term catalytic benefit available after divergence.

Of 4,740 measured dosage-compensated origins, 4,044 (85.3%) were neutral within numerical tolerance, 449 (9.5%) were deleterious, and 247 (5.2%) were beneficial (Fig. 1C). The mean immediate change in log metabolic-fitness score was − 0.0027. Strong antagonism produced the greatest imbalance: 14.3% of origins were deleterious, 4.25% were beneficial, and 81.45% were neutral, with a mean change of − 0.0067. After retention, antagonism supplied the conflict resolved by complementary specialization. At origin, the same geometry made a copied haplotype more likely to perturb an adapted phenotype. Catalytic frustration therefore changed biological role across stages. Initially, frustration could oppose persistence while favoring later functional division.

Among 80 paired ridge-evolving retained populations, persistent specialization increased from 24/80 at 10N to 30/80 at 20N (Fig. 1D). Six pairs gained specialization, and none lost it (exact two-sided McNemar $p = 0.03$). The mean final specialization index increased by 0.094, and the mean metabolic-fitness score increased by 0.25. Thus, fitness could stabilize while catalytic traits continued to explore a trade-off ridge. Therefore, the longer duration did not resolve the compensated single-origin barrier because most rare copies were lost before that later evolutionary period.

Effects of copy-number-additive dosage

The dosage treatments produced the intended immediate effects. After population-wide duplication, additive activity increased mean catalytic output by 2.00-fold (SD, 0.00035) and the unscaled metabolic-fitness score by 2.00-fold (SD, 0.013). Compensation held these ratios near one (activity, 1.00, SD 0.00018; score, 1.00, SD 0.0067; Fig. 2A). The additive twofold score was not a twofold selective advantage after population-wide retention because all feasible individuals received the same multiplier.

From a single origin, however, additive dosage altered copy fate. Establishment increased from 51/5,700 to 2,421/5,700. The paired risk difference was 41.58% (ancestor-clustered SE, 0.73; 95% t interval, 40.14–43.02; $t_{1139} = 56.68$, $p < 0.001$). Near-fixation increased from 1/5,700 to 2,421/5,700, a paired difference of 42.46% (SE, 0.74; 95% interval, 41.01–43.90; $t_{1139} = 57.58$, $p < 0.001$; Table 2; Fig. 2B). Under the additive rule, every origin that established also reached the 95% near-fixation threshold. The immediate dosage effect changed how many duplicate lineages survived long enough to reach a catalytic endpoint.

Table 2

Matched dosage contrasts. Single-origin intervals and p values are two-sided one-sample t results across 1,140 independent ancestors after averaging five attempts within ancestor. The retained interval uses 1,140 paired trajectories; its p value is from an exact two-sided McNemar test.

Experimental stage and outcome	Compensated	Additive	Paired difference (%)	95% interval	Paired p
Single origin: establishment	51/5,700	2,421/5,700	+ 41.58	40.14 to 43.02	< 0.001
Single origin: near-fixation (≥ 95%)	1/5,700	2,421/5,700	+ 42.46	41.01 to 43.90	< 0.001
Single origin: persistent specialization	0/5,700	102/5,700	+ 1.789	1.280 to 2.299	< 0.001
Retention imposed: persistent specialization	61/1,140	60/1,140	−0.088	−0.604 to 0.429	1.000

Stage-specific dosage effects

Additive dosage produced 102 persistent specialists among 5,700 origins, compared with none under compensation. The paired increase was 1.79% (95% t interval, 1.28–2.30; $t_{1139} = 6.89$, $p < 0.001$). Competing-risk analysis showed that dosage first prevented duplicate loss, after which specialization accumulated among the retained lineages (Online Resource 1, Supplementary Fig. S1). All 102 additive single-origin specialists occurred in the six high-mutation, strong-antagonism treatments. Additive-minus-compensated effects were + 4.0%, + 10.4%, and + 9.6% under hard boundaries and + 4.8%, + 5.6%, and + 6.4% points under soft boundaries at 0.5×, 1×, and 2× selection (Fig. 2C). These estimates describe treatment-specific effects; their wide intervals do not support ranking the six regimes.

As predicted, additivity did not change specialization after retention was imposed. In the matched retained grid, 61/1,140 compensated and 60/1,140 additive trajectories specialized. Four pairs specialized only under additivity and five only under compensation, yielding a paired difference of − 0.09% (95% t interval, − 0.60–0.43; exact McNemar p = 1; Fig. 2D). All 61 compensated events occurred in the six focal conditions (61/300), compared with 56/300 in the primary retained analysis; this difference reflects use of the final treatment-specific ancestor bank, including 80N ridge ancestors. Among feasible genotypes with population-wide copy number, additive scaling provided a common factor and therefore did not alter relative selection.

Importantly, the single-origin dosage effect on specialization exceeded the retained effect by 1.88% (treatment-stratified partial-pair cluster-bootstrap 95% interval, 1.25–2.53; bootstrap SE, 0.33). Together, these results indicate that immediate dosage increased access to the retained two-copy state without changing specialization conditional on retention in this implementation.

Discussion

Biochemical opportunity and population-genetic accessibility

Our results separate two evolutionary processes that are often treated as a single consequence of gene duplication. Catalytic frustration created selection for a division of labor between copies, but this opportunity mattered only after the duplicate escaped stochastic loss. Immediate dosage changed that first transition, by increasing establishment and near-fixation from a single origin by more than 40%, yet it had almost no effect on specialization when retention was imposed population-wide. Thus, the catalytic landscape defined what could evolve after retention, while dosage altered the probability that a population reached that landscape with two copies.

Previous theory and experiments show that dosage can preserve duplicate genes and that retained copies can diverge (Kondrashov and Kondrashov 2006; Bergthorsson et al. 2007; Andersson and Hughes 2009; Sikosek et al. 2012; Näsvall et al. 2012). Proulx (2012) further modeled failed origins, incorporation, regulatory change, and coding specialization under a shared selective environment. Our matched simulations incorporate these stages within the same catalytic-frustration landscape and support a more specific role for dosage, acting primarily during establishment and near-fixation rather than during specialization after retention.

These results place a population-genetic boundary between biochemical models of specialization and models of duplicate preservation. Constraint-based models identify catalytic regions in which specialists can outperform a constrained generalist (Schmutzer et al. 2024), while population-genetic models identify conditions under which a duplicate survives (Lynch et al. 2001; Kondrashov and Kondrashov 2006). Innovation–Amplification–Divergence connects these processes by proposing that increased dosage and mutational target size maintain weak activity while later mutations refine function (Näsvall et al. 2012). In the simulations presented here, a favorable catalytic endpoint remained effectively unavailable until a mechanism preserved the new copy.

Dosage effects before and after retention

Additive dosage converted many drift-dominated origins into lineages that established and reached near-fixation. However, these results do not mean that dosage directed the evolution of specialization. Instead, dosage increased the number of lineages that entered and remained in a state where selection could drive specialization. Once copy number was imposed population-wide, catalytic relationship, mutation supply, and observation time determined whether divergence occurred.

Importantly, a higher probability of specialization does not mean that dosage directly caused catalytic divergence. The unconditional probability can increase even when specialization conditional on retention remains unchanged. In the present model, the retained null result is structurally expected because additive output becomes a common multiplier under proportional selection. This result is therefore a validated negative control, rather than evidence that dosage is generally irrelevant after population-wide

retention. Partial compensation, nonlinear flux, toxicity, expression evolution, or genotype-specific cost would break this invariance.

Consistent with the broader observation that an extra copy does not prescribe a specialized endpoint, Mihajlovic et al. (2025) observed greater mutational robustness and diversity in two-copy systems but not faster phenotypic evolution or more specialization. Although their experimental system differs from this model, both results suggest that relaxed constraint and increased copy number are insufficient on their own. A retention advantage can create time for divergence, but catalytic and ecological selection determine how that time is used.

Catalytic conflict and mutation supply

Every specialization event occurred under strong antagonism and high mutation supply. Strong antagonism created a benefit to dividing catalytic labor, whereas weaker antagonism, uncoupled activity, and synergism did not produce the same selective conflict. However, strong antagonism at the lower mutation rate produced no specialists. The landscape must both favor specialization and be sampled during the time that two copies persist.

Mutation supply consequently had different effects across stages. Increasing the catalytic mutation rate did not rescue compensated origins because most copies disappeared before differentiating mutations could accumulate. Once dosage promoted near-fixation, the same high mutation supply became consequential. This result parallels Innovation–Amplification–Divergence, where amplification can enlarge the mutational target, but the amplified state must remain present long enough for that target to matter (Bergthorsson et al. 2007; Näsvall et al. 2012). In contrast, catalytic frustration had opposing effects before and after retention. When a compensated copy arose, strong antagonism made the copied haplotype more likely to reduce performance than improve it. After retention, the same constraint created the selective conflict that complementary specialization resolved. As such, a biochemical architecture can oppose the origin of a duplicate while favoring its later divergence. This stage dependence may be especially relevant to enzymes whose promiscuous activities are embedded in already adapted primary or specialized metabolic networks. Future work should consider how pathway architecture may promote antagonism or maintain promiscuous activities.

Biological implications of the dosage bounds

In plants, fungi, and other systems with a wide array of duplication events of enzymes, whether a new duplicate alters metabolic output depends strongly on its genomic and expression context. A second copy may initially increase flux, be buffered by feedback, or remain phenotypically silent. However, we intentionally defined the additive and compensated treatments tested here as opposing bounds. Additivity assumes that output scales with functional copy number, whereas compensation holds total output approximately constant. Natural genes can fall between or outside these limits because expression attenuation can buffer duplicate effects (Ascencio et al. 2021), dosage balance can constrain duplicate retention, and later expression divergence can change copy fate (Lan and Pritchard 2016).

Product saturation, protein burden, toxicity, and stoichiometric imbalance may also reverse an immediate benefit.

Tandem arrays, biosynthetic gene clusters, tissue-specific expression, environmental induction, and recurrent copy-number variation can each protect an extra copy long enough for catalytic divergence (Osbourn 2010; Weng et al. 2012; Keller 2019). For an enzyme expressed during a high-demand condition, increased flux may initially favor a duplicate, after which coding or regulatory change partitions activity among tissues, environments, or products. Alternatively, rapid dosage compensation may erase the initial advantage and return the copy to drift-dominated loss. Catalytic conflict may therefore explain why retained family members diverge, while dosage and regulation explain how retention was first achieved. This begs the question of how much dosage benefit is required to cross this retention boundary. Partial transcriptional compensation, saturating activity, dosage cost, and toxicity could be varied independently to identify the minimum immediate benefit needed for retention. In addition, environmental fluctuation and regulatory divergence could then test whether the same copy is favored through constant abundance, condition-specific expression, or both.

Equilibrium wandering and competing outcomes

The 10N-to-20N comparison showed that stable mean fitness does not imply a static catalytic phenotype. Six of 80 ridge-evolving populations specialized during the added interval, and none reversed. This pattern is consistent with equilibrium wandering along a nearly flat catalytic ridge. Requiring all trait movement to stop before duplication would therefore remove a biologically meaningful state rather than produce a more neutral ancestral condition.

Duplicate loss also competes directly with specialization. Analyses restricted to successful specialists would ignore the large majority of origins that disappeared. The cumulative-incidence analysis instead retained all origins and showed that the primary additive effect occurred early through prevention of loss, while specialization accumulated later among surviving lineages. Time to specialization should therefore be interpreted jointly with the probability of ever entering the at-risk retained state.

Scope and limitations

These simulations are not predictions for a particular species. The model represents two abstract catalytic reactions rather than nucleotide-resolved proteins or a specific metabolic pathway, and it excludes recurrent duplication, deletion, gene conversion, regulatory evolution, population structure, environmental change, linkage to other selected loci, and explicit metabolic feedback. Population size and mutation rates were selected for comparative forward simulation rather than as estimates for a particular organism. Population-size rescaling can bias fixation times, allele frequencies, linkage disequilibrium, and fixation probability in forward simulations (Dabi and Schrider 2025). Therefore, simulation generations are not calendar time; instead, the results identify comparative changes in evolutionary accessibility across the declared mechanistic design. Further, the additive experiment isolates one immediate copy-number mechanism. It does not establish that positive dosage is common, and its near-fixation probability should not be extrapolated to dosage-sensitive genes.

Finally, the retained and single-origin experiments answer complementary questions. Imposed retention estimates conditional evolutionary potential; a single origin estimates accessibility under one duplication process. Long-term gene-family evolution will additionally depend on recurrent origins, loss, expression evolution, ecology, and changing selection. Extending the biochemical model from two reactions to sampled multi-reaction architectures is an important next step, but the two-reaction design provides the clearest test of the retention boundary addressed here.

Conclusions

Gene duplication did not resolve catalytic conflict simply by creating a second copy; a population first had to retain that copy. In these simulations, copy-number-additive activity increased establishment and near-fixation from rarity, thereby making specialization accessible. Strong antagonism and high mutation supply then determined where specialization occurred. Once copy number was imposed population-wide, the common additive multiplier canceled from relative selection among feasible genotypes, as predicted by the model. Dosage therefore controlled access to the two-copy state, whereas catalytic conflict and mutation supply controlled what evolved after access was achieved.

Declarations

Competing Interests

The author has no relevant financial or non-financial interests to disclose.

Online Resource 1

Supplementary methods, catalytic feasible-space definitions, and competing-outcome cumulative-incidence analysis for the single-origin experiment (PDF)

Online Resource 2

Commented SLiM simulation model, portable execution helper, example inputs, and documentation for the ancestral calibration, retained-duplication, single-origin, and matched-dosage experiments (ZIP)

Funding

Computational resources were provided by the Louisiana State University High Performance Computing Center under allocation hpc_dowelllac.

Author Contribution

J.A.D. designed, prepared, wrote and reviewed the manuscript.

Acknowledgement

The Louisiana State University High Performance Computing Center provided computational resources, along with startup funds allocated to JAD by Louisiana State University.

Data Availability

The commented SLiM model and portable execution materials are supplied as Online Resource 2. Ancestral-state identifiers are retained so that repeated origins can be analyzed as clusters and dosage treatments can be paired exactly.Online Resource 1 Supplementary methods, catalytic feasible-space definitions, and competing-outcome cumulative-incidence analysis for the single-origin experiment (PDF)Online Resource 2 Commented SLiM simulation model, portable execution helper, example inputs, and documentation for the ancestral calibration, retained-duplication, single-origin, and matched-dosage experiments (ZIP)

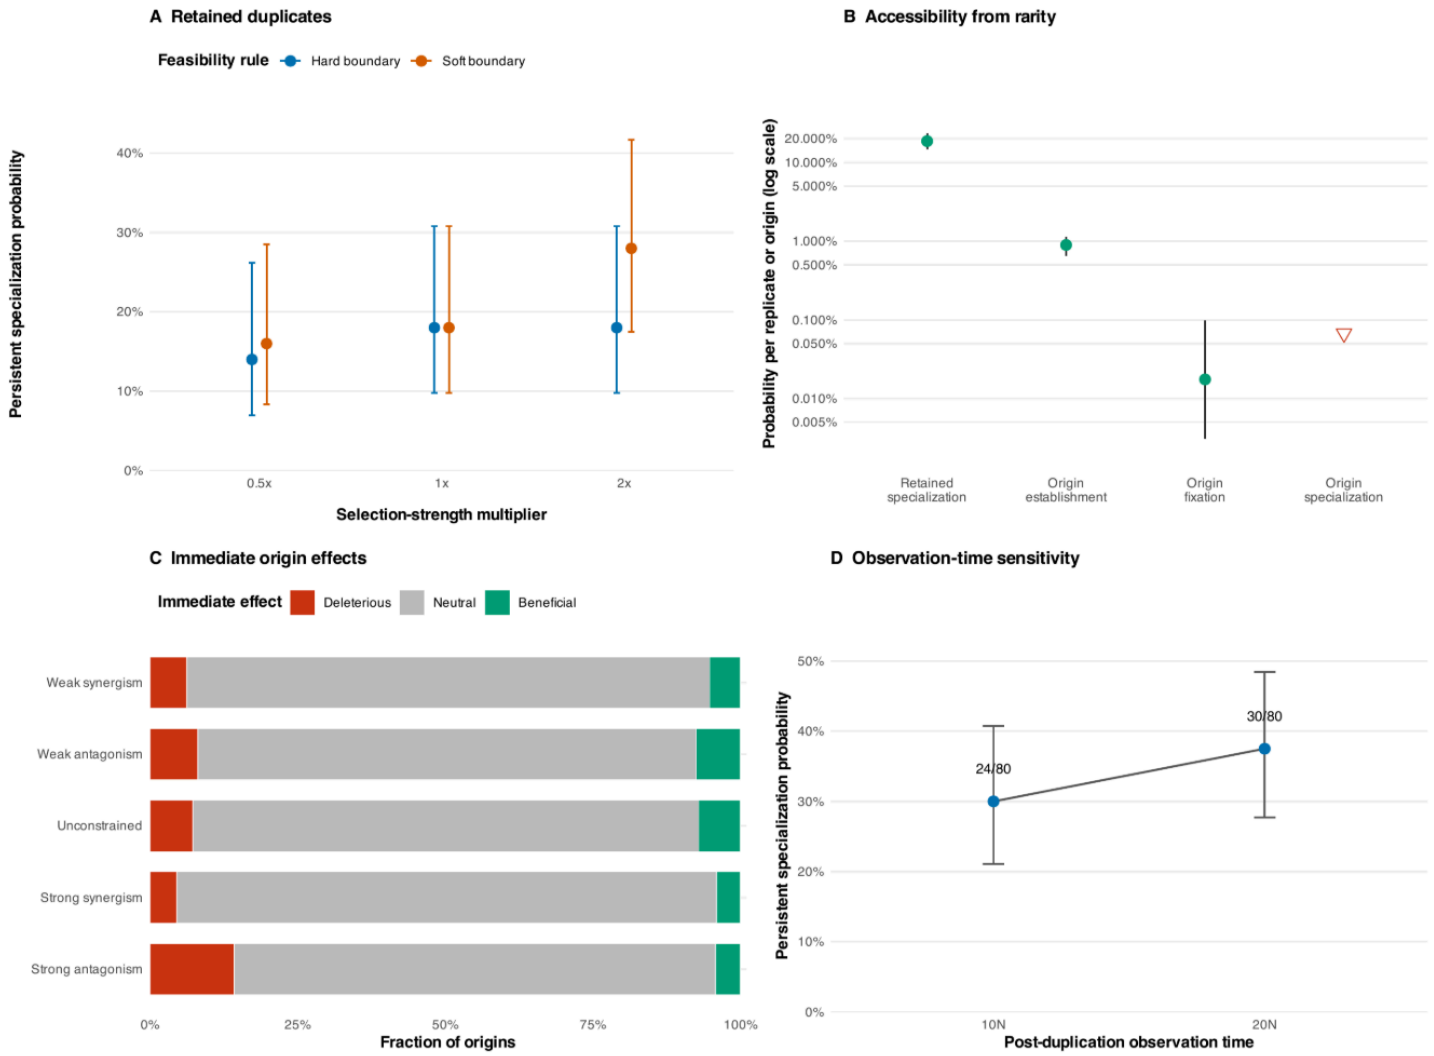
References

1. Aalen OO, Johansen S (1978) An empirical transition matrix for non-homogeneous Markov chains based on censored observations. Scand J Stat 5:141–150
2. Aharoni A, Gaidukov L, Khersonsky O, McQ Gould S, Roodveldt C, Tawfik DS (2005) The ‘evolvability’ of promiscuous protein functions. Nat Genet 37:73–76. https://doi.org/10.1038/ng1482
3. Andersson DI, Hughes D (2009) Gene amplification and adaptive evolution in bacteria. Annu Rev Genet 43:167–195. https://doi.org/10.1146/annurev-genet-102108-134805
4. Ascencio D, Diss G, Gagnon-Arsenault I, Dubé AK, DeLuna A, Landry CR (2021) Expression attenuation as a mechanism of robustness against gene duplication. Proceedings of the National Academy of Sciences USA 118:e2014345118. https://doi.org/10.1073/pnas.2014345118
5. Bergthorsson U, Andersson DI, Roth JR (2007) Ohno’s dilemma: evolution of new genes under continuous selection. Proceedings of the National Academy of Sciences USA 104:17004–17009. https://doi.org/10.1073/pnas.0707158104
6. Conant GC, Wolfe KH (2008) Turning a hobby into a job: how duplicated genes find new functions. Nat Rev Genet 9:938–950. https://doi.org/10.1038/nrg2482
7. Copley SD (2017) Shining a light on enzyme promiscuity. Curr Opin Struct Biol 47:167–175. https://doi.org/10.1016/j.sbi.2017.11.001
8. Dabi A, Schrider DR (2025) Population size rescaling significantly biases outcomes of forward-in-time population genetic simulations. Genetics 229:iyae180. https://doi.org/10.1093/genetics/iyae180

9. Des Marais DL, Rausher MD (2008) Escape from adaptive conflict after duplication in an anthocyanin pathway gene. Nature 454:762–765. https://doi.org/10.1038/nature07092
10. Ferreiro DU, Komives EA, Wolynes PG (2014) Frustration in biomolecules. Q Rev Biophys 47:285–363. https://doi.org/10.1017/S0033583514000092
11. Force A, Lynch M, Pickett FB, Amores A, Yan YL, Postlethwait J (1999) Preservation of duplicate genes by complementary, degenerative mutations. Genetics 151:1531–1545. https://doi.org/10.1093/genetics/151.4.1531
12. Francino MP (2005) An adaptive radiation model for the origin of new gene functions. Nat Genet 37:573–578. https://doi.org/10.1038/ng1579
13. Gould SM, Tawfik DS (2005) Directed evolution of the promiscuous esterase activity of carbonic anhydrase II. Biochemistry 44:5444–5452. https://doi.org/10.1021/bi0475471
14. Haller BC, Messer PW (2019) SLiM 3: forward genetic simulations beyond the Wright–Fisher model. Mol Biol Evol 36:632–637. https://doi.org/10.1093/molbev/msy228
15. Haller BC, Messer PW (2023) SLiM 4: multispecies eco-evolutionary modeling. Am Nat 201:E127–E139. https://doi.org/10.1086/723601
16. Keller NP (2019) Fungal secondary metabolism: regulation, function and drug discovery. Nat Rev Microbiol 17:167–180. https://doi.org/10.1038/s41579-018-0121-1
17. Khersonsky O, Tawfik DS (2010) Enzyme promiscuity: a mechanistic and evolutionary perspective. Annu Rev Biochem 79:471–505. https://doi.org/10.1146/annurev-biochem-030409-143718
18. Kondrashov FA, Kondrashov AS (2006) Role of selection in fixation of gene duplications. J Theor Biol 239:141–151. https://doi.org/10.1016/j.jtbi.2005.08.033
19. Lan X, Pritchard JK (2016) Compensatory drift and the evolutionary dynamics of dosage-sensitive duplicate genes. Genetics 202:765–777. https://doi.org/10.1534/genetics.115.178137
20. Lynch M, Conery JS (2000) The evolutionary fate and consequences of duplicate genes. Science 290:1151–1155. https://doi.org/10.1126/science.290.5494.1151
21. Lynch M, Force A (2000) The probability of duplicate gene preservation by subfunctionalization. Genetics 154:459–473. https://doi.org/10.1093/genetics/154.1.459
22. Lynch M, O'Hely M, Walsh B, Force A (2001) The probability of preservation of a newly arisen gene duplicate. Genetics 159:1789–1804. https://doi.org/10.1093/genetics/159.4.1789
23. McNemar Q (1947) Note on the sampling error of the difference between correlated proportions or percentages. Psychometrika 12:153–157. https://doi.org/10.1007/BF02295996
24. Mihajlovic L, Iyengar BR, Baier F, Barbier I, Iwaszkiewicz J, Zoete V, Wagner A, Schaerli Y (2025) A direct experimental test of Ohno's hypothesis. eLife 13:RP97216. https://doi.org/10.7554/eLife.97216.3
25. Näsvall J, Sun L, Roth JR, Andersson DI (2012) Real-time evolution of new genes by innovation, amplification, and divergence. Science 338:384–387. https://doi.org/10.1126/science.1226521
26. Ohno S (1970) Evolution by Gene Duplication. Springer-, Berlin

27. Osbourn A (2010) Gene clusters for secondary metabolic pathways: an emerging theme in plant biology. Plant Physiol 154:531–535. https://doi.org/10.1104/pp.110.161315
28. Proulx SR (2012) Multiple routes to subfunctionalization and gene duplicate specialization. Genetics 190:737–751. https://doi.org/10.1534/genetics.111.135590
29. Qian W, Zhang J (2008) Gene dosage and gene duplicability. Genetics 179:2319–2324. https://doi.org/10.1534/genetics.108.090936
30. R Core Team (2026) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna, Austria. https://www.R-project.org/
31. Schmutzer M, Dasmeh P, Wagner A (2024) Frustration can limit the adaptation of promiscuous enzymes through gene duplication and specialisation. J Mol Evol 92:104–120. https://doi.org/10.1007/s00239-024-10161-4
32. Sikosek T, Chan HS, Bornberg-Bauer E (2012) Escape from adaptive conflict follows from weak functional trade-offs and mutational robustness. Proceedings of the National Academy of Sciences USA 109:14888–14893. https://doi.org/10.1073/pnas.1115620109
33. Tokuriki N, Jackson CJ, Afriat-Jurnou L, Wyganowski KT, Tang R, Tawfik DS (2012) Diminishing returns and tradeoffs constrain the laboratory optimization of an enzyme. Nat Commun 3:1257. https://doi.org/10.1038/ncomms2246
34. Weng JK, Philippe RN, Noel JP (2012) The rise of chemodiversity in plants. Science 336:1667–1670. https://doi.org/10.1126/science.1217411
35. Wilson EB (1927) Probable inference, the law of succession, and statistical inference. J Am Stat Assoc 22:209–212. https://doi.org/10.1080/01621459.1927.10502953

Figures



Intervals are 95% Wilson intervals except the ancestor-clustered establishment interval in B. Pooled origin estimates are replicate-weighted under the adaptive design; the open triangle is a descriptive zero-event upper bound.

Figure 1

Biochemical opportunity and population-genetic accessibility under dosage compensation. (A) Persistent complementary specialization after imposed retention in the six high-mutation, strong-antagonism treatments. Each estimate represents 50 independently adapted ancestral populations; points and bars show proportions and 95% Wilson intervals. Hard-boundary outcomes were 7/50, 9/50, and 9/50 at 0.5×, 1×, and 2× selection; soft-boundary outcomes were 8/50, 9/50, and 14/50. (B) Retained specialization compared with outcomes from one compensated origin. Retained specialization was 56/300 in the six focal treatments. Across the adaptively replicated 48-cell single-origin design, 51/5,700 origins established, 1/5,700 reached near-fixation, and 0/5,700 specialized. Five origins were nested within each of 1,140 ancestors. The establishment interval is the ancestor-clustered 95% t interval; the near-fixation interval and open specialization bound are descriptive attempt-level Wilson quantities. The y-axis is logarithmic, and denominators differ because the two experiments address different stages. (C) Immediate effect of 4,740 compensated origins on the log metabolic-fitness score: 4,044 neutral, 449 deleterious, and 247 beneficial, with neutrality defined as $|\Delta \log R| \leq 10^{-12}$. (D) Observation-time sensitivity in 80 paired ridge-evolving populations. Persistent specialization increased from 24/80 at 10N to 30/80

at 20N; six pairs gained specialization and none lost it (exact two-sided McNemar p = 0.03125). Bars are marginal 95% Wilson intervals

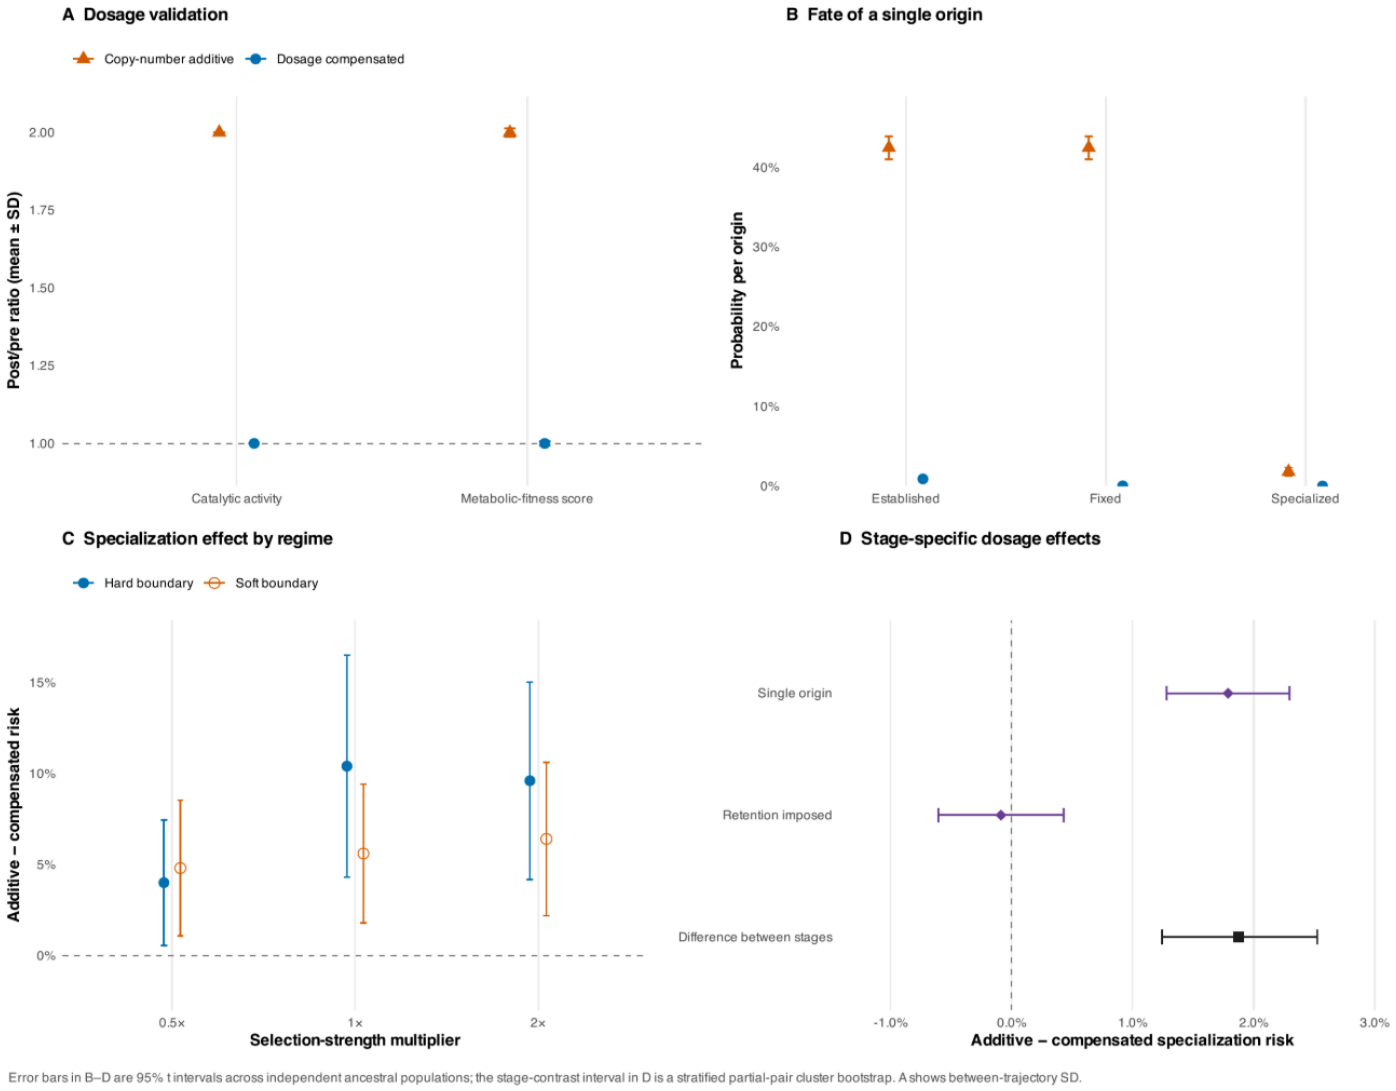


Figure 2

Copy-number-additive dosage increases accessibility from rarity. (A) Validation of dosage rules in 1,140 matched retained trajectories per mode. Additive activity approximately doubled catalytic output and the unscaled metabolic-fitness score, whereas compensation held both near onefold. Values are trajectory means with between-trajectory SDs. (B) Establishment, near-fixation, and persistent specialization among 5,700 matched origins per mode. Each mode contains five attempts from each of 1,140 ancestors; points and 95% t intervals use ancestor-level proportions. Compensated versus additive counts were 51 versus 2,421 for establishment, 1 versus 2,421 for near-fixation, and 0 versus 102 for specialization. (C) Additive-minus-compensated specialization risk in the six focal regimes. Each estimate represents 250 origins nested within 50 ancestors; bars are ancestor-clustered 95% t intervals. Effects at 0.5×, 1×, and 2× selection were +4.0, +10.4, and +9.6 percentage points under hard boundaries and +4.8, +5.6, and +6.4 under soft boundaries. (D) Specialization effects by experimental stage. The

single-origin dosage effect was +1.789 percentage points (95% interval, 1.280–2.299), and the imposed-retention effect was −0.088 points (95% interval, −0.604–0.429; exact McNemar p = 1). Their difference was +1.877 points (treatment-stratified partial-pair cluster-bootstrap 95% interval, 1.246–2.526; 50,000 resamples). Of 1,140 ancestors per stage, 1,060 were shared and resampled jointly; 80 unmatched ancestors per stage were resampled independently

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- ESM1SupplementaryInformation.docx
- ESM2AnnotatedSLIMScripts.zip