Weakest link epistasis and the geometry of genetic load

Florian J.F.Labourel[1,✉], **Florence Bansept**[1], **David M. McCandlish**[2]

[1]Aix Marseille Université, CNRS, Laboratoire de Chimie Bactérienne (UMR7283), Turing Center for Living Systems, Marseille, France, IMM, IM2B, 13402 Marseille, France; [2]Simons Center for Quantitative Biology, Cold Spring Harbor Laboratory, Cold Spring Harbor, NY 11724, USA

✉ **For correspondence:** florian.labourel@univ-amu.fr

Key words: adaptation | complexity | epistasis| Fisher's geometric model | genetic drift | genetic load

Competing interests: The authors declare no competing interests.

Abstract

Because natural selection must optimise multiple traits at once, previous work has suggested that phenotypic dimensionality can substantially worsen the equilibrium fitness defect of a population relative to the phenotypic optimum. However, it remains unclear how conclusions drawn from classical theoretical phenotype–fitness maps extend to models grounded in explicit biological mechanisms. Here we introduce weakest-link epistasis (WLE), a framework in which fitness is determined by the least-fit phenotypic component, an extreme form of diminishing returns epistasis. We show that in this framework, increasing dimensionality amplifies the load in a manner comparable to, but surprisingly not more than, Fisher's geometric model (FGM). Building on this similarity, we demonstrate why genetic load is often invariant across different rules for combining trait-specific fitness components into an overall organismal fitness. We explore these ideas by considering the family of models where the organismal fitness is determined based on the ℓ^p-norm of the vector of trait-specific fitness defects, a framework that includes both FGM and WLE, but also captures a continuum of genetic architectures, ranging from generalist to specialist regimes. Altogether, our approach proposes a new perspective on the geometry of adaptive landscapes, and may help provide quantitative insight into the cost of complexity.

1 Introduction

The difference between the stationary mean fitness of a population and the fittest possible genotype is called the genetic load [1–4]. Since alleles are effectively neutral when their selection coefficient $s << 1/N_e$, where N_e denotes the long-term (effective) population size of a species, the genetic load in a single-trait model is classically predicted to scale as $1/N_e$ [4–7]. However, this apparent predictability of the genetic load is complicated by the genetic architecture underlying fitness (e.g., supply of deleterious mutations [3], degree of epistasis [8]). Notably, increasing the number of fitness-contributing traits can dramatically amplify the supply of deleterious mutations and hence the genetic load [2, 9–11], thereby creating a mutational burden that scales with the complexity of organisms [12–14].

This burden of complexity has most often been studied through the lens of Fisher's geometric model (FGM), under which the fitness landscape contains a single optimum in the phenotypic

space. In this model, the genotypic load carried by an organism increases with the distance to the optimum of its phenotype mapped into the fitness space [10, 14, 15]. In FGM, complexity is thus traditionally captured by the dimensionality of the fitness landscape, with the assumption that each dimension could potentially represent a component (eg., a gene, a trait) contributing to fitness [14]. In this scenario, the genetic load was shown to be substantially exacerbated (relative to the single-trait case), to the extent that it can become proportional to the number of fitness-contributing components [2, 9].

Yet, if FGM has received some theoretical support [14, 16], it remains unknown how accurate and broad a description it is of *in vivo* fitness landscapes [17–20]. Indeed, FGM relies on a very specific assumption about how phenotypic traits combine to produce fitness—namely, that fitness depends on the Euclidean distance to the peak of their respective deviation from their optimum value— which calls into question the universality of its predictions, including those about the genetic load. Accordingly, FGM was extended to demonstrate the genetic load invariance across a much broader range of theoretical fitness models [12]. A central limitation of this finding, however, is that neither fitness nor the underlying definition of complexity is grounded with a clear biological rationale, raising questions about how complexity should be measured [12, 13, 21].

By contrast, a relevant yet unexplored principle governing the combined effect of traits on fitness is Liebig's law of the minimum [22–24]. According to this law, an organism's growth is simply dictated by the most limiting nutrient in its environment, that is its "weakest link", much like the capacity of a barrel composed of staves of unequal length would be limited by its shortest stave. Though originally proposed (and criticised) for the response of organisms to resources [22], this concept conceals a significant potential for generalisation, especially at the cellular scale [24]. In particular, because biosynthetic molecules can either be taken up in the environment or produced through *de novo* synthesis, it is natural to extend this weakest link logic to the mapping of metabolic traits to fitness such that fitness would be dictated by the least-fit trait [25]. Imagine a two-enzyme pathway, for instance, where one enzyme would be inefficient. In this scenario, the other enzyme would have little influence on the metabolic flux (and hence on fitness, provided fitness is positively correlated with the flux), because it has virtually no control on the flux. This epistasis-driven diminishing returns is the classical expectation from metabolic control theory [26, 27], where the flux increase due to enhancing an enzyme's efficiency gets inescapably limited by the presence of other enzymes in the pathway [19, 28–31]. Notably, a similar mapping of traits onto fitness recently helped explain why some metabolic features may be far off their optimum [32]. It is also consistent with the existence of lethal mutations that make fitness equal zero regardless of how fit most of an organism's genes are [33].

Based on this "law of the minimum", we first describe the population genetics consequences of weakest link epistasis (WLE) in terms of fitness landscapes and mutant selection coefficients [25]. We then determine the genetic load under WLE and show that phenotypic dimensionality n amplifies the bias toward maladaptive traits b, which reflects the tendency of mutations to produce deleterious trait changes, yielding an effective maladaptive trait bias nb. To demonstrate the generality of this finding, we prove that a broad class of models for combining trait-specific fitness defects into an overall organismal fitness defect lead to the same genetic load, which we show is a consequence of the geometric relationships between equal fitness contours. We explore these results by considering a class of models where the organismal fitness defect is equal to the ℓ^p norm of the vector of trait-specific fitness defects, and discuss the extent to which these models can capture different phenotypic strategies (e.g. generalists versus specialists) despite the genetic

load being identical.

2 Model & Results

Consider a haploid population of individuals that have a number n of traits (or phenotypic dimensions). For a genotype g, let $\mathbf{F_g} = (F_{g,1}, ..., F_{g,d}, ..., F_{g,n})$, be the vector of corresponding fitness components where $F_{g,d}$ denotes the trait-dependent fitness component associated with the value for the dth phenotypic trait of genotype g. Each such fitness component captures the contribution of a trait to fitness (e.g., the fitness contribution brought by a metabolic reaction, by the shape of an organism, etc.). These fitness components are then combined through a function $f_{org}(\mathbf{F_g})$ to yield the fitness f_g of an organism with genotype g:

$$\underbrace{\mathbf{F_g} = (F_{g,1}, \dots, F_{g,n})}_{\text{phenotype}} \xrightarrow{f_{\text{org}}(\mathbf{F_g})} \underbrace{f_g}_{\text{fitness}}$$

We consider the fitness components $F_{g,d}$ to range from 0 to 1 so that values are easily comparable across traits and can be interpreted as a mapping of trait values into a bounded fitness space. Specifically, we define $F_{g,d} = f_{org}(1, 1, ..., F_{g,d}, ..., 1)$ when all but component d of genotype g has been set equal to 1, i.e. the fitness component $F_{g,d}$ has the interpretation that it is the fitness that would be conferred by the dth trait of genotype g if all its other traits were set to their optimal values.

Based on these premises, we now provide a mathematical formalisation of weakest link epistasis and its direct population genetics consequences.

2.1 Weakest link epistasis from a population genetics perspective

In order to formalize the weakest link logic, we can define the fitness of a genotype $f_g = f_{org}(\mathbf{F_g})$ as the minimum of the trait-dependent fitness components (see Figure 1):

$$f_{\text{WLE}}(\mathbf{F_g}) = \min_{d \in [1,n]} [F_{g,d}] \tag{1}$$

Because $F_{g,d}$ only prescribes fitness when it is the minimum, or limiting factor, it represents the maximum fitness possible given that trait value (e.g., the maximum growth rate that ATP production or nucleotides synthesis can sustain).

2.1.1 Selection coefficients under weakest link epistasis

If mutations are rare enough, a new mutant with genotype g' is either fixed or lost before the next new mutation arises, which implies the absence of clonal interference [5, 7, 34]. We can thus drop genotype indexes and focus on a mutant with trait-dependent fitness components $\mathbf{F}'$ and a resident with trait-dependent fitness components $\mathbf{F}$. Hence, the selection coefficient (such that $f_{g'} = f_g(1 + s)$) reads:

$$s = f_{org}(\mathbf{F}')/f_{org}(\mathbf{F}) - 1 \tag{2}$$

Under WLE, this implies that the selection coefficient of a mutant $\mathbf{F}'$ is thus:

$$s_{\text{WLE}} = \frac{\min \mathbf{F}'}{\min \mathbf{F}} - 1 \tag{3}$$

Whether mutations are deleterious ($s_{\text{WLE}} < 0$), neutral ($s_{\text{WLE}} = 0$) or advantageous ($s_{\text{WLE}} > 0$) therefore simply depends on the sign of the difference between the weakest link of the mutant $\min(\mathbf{F}')$ and the weakest link of the resident $\min(\mathbf{F})$.

Yet, because many mutations are masked by the weakest link effect on fitness, selection is more complex and counterintuitive than it would be in the absence of epistasis. Crucially, this fitness model tends to inflate the supply of purely neutral mutations, since any mutation that preserves $\min(\mathbf{F}') = \min(\mathbf{F})$ leads to $s_{\text{WLE}} = 0$. This happens whenever the weakest link component remains intact after mutation. The resulting extended range of neutral mutations includes any mutation—pleiotropic or not—that does not decrease any of the mutant's fitness contributions below the weakest link value of the resident.

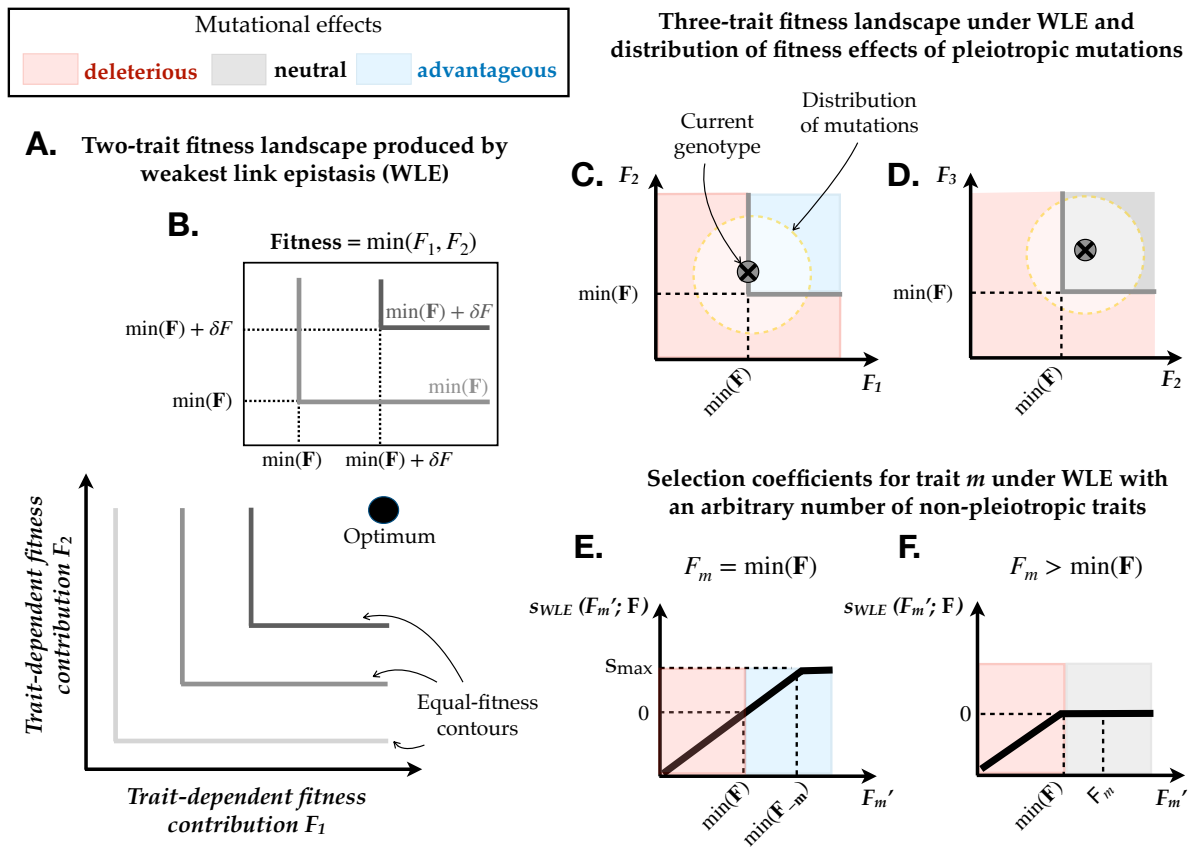

Figure 1. Explanation of the mechanism of weakest-link epistasis (WLE) and its population-genetic consequences when each trait d produces a fitness component F_d. **A–B.** Panels A and B illustrate how WLE operates in a two-dimensional trait space where fitness is bounded (e.g, between 0 and 1) and reaches its optimum at the upper-right corner (black circle). In **B**, the contours of equal fitness are represented in shades of grey, with their values indicated above each contour. Under WLE, fitness equals the minimum of the fitness components, $\min(\mathbf{X})$: it thus remains equal to F_1 as long as $F_2 > F_1$, and symmetrically for F_1. Fitness increases by δF only when both fitness components increase by at least δF. In **A.**, the resulting fitness landscape is composed of quarter-square contours centered on the optimum (in grey). **C–D.** Panels C and D illustrate why WLE produces an excess of neutral and deleterious mutations for the WLE fitness model when pleiotropic mutations affect two traits at a time. As an illustration, the mutational neighbourhood is delimited by circles (with orange-dotted boundaries). In **C**, when one of the mutated trait is the resident's weakest link (trait 1 in this example), the resident genotype (grey cross-in-circle) lies on the equal-fitness contour (grey line). A range of mutations increasing this least-fit trait are advantageous (blue region, upper-right), yet comparatively more mutations are deleterious (red region). The exact balance of advantageous vs. deleterious mutations depends on the exact position of the other trait. In **D**, when mutations affect two traits that are not the resident's weakest link, their effect is either neutral because the weakest link remains unchanged (grey region above the fitness contour) or deleterious otherwise (red region); no mutation can be advantageous in this case. **E–F.** Selective coefficients of non-pleiotropic mutations affecting only trait m. In **E**, for a mutation affecting the weakest link, the selection coefficient s_{WLE} increases (up to s_{max}) until $F'_m = \min(\mathbf{F}_{-\mathbf{m}})$, above which trait m is no longer the weakest link. In **F**, as in panel D, no mutation can be advantageous since the mutated trait m is not the weakest link.

To further understand the selective effect of mutations, we can focus on the non-pleiotropic scenario where mutations only affect one trait at a time. Any mutant genotype g' thus differs from its resident progenitor g by only one phenotypic trait value. We denote this mutant trait value F'_m (in the fitness components space) when it affects trait m, so that $F'_m \neq F_m$. As mutants and residents share the same genotypic (and phenotypic) background $\mathbf{F_{-m}}$, the trait-dependent fitness components are all equal but for the mutant trait:

$$\mathbf{F_{-m}} = \mathbf{F_g} \backslash \{F_m\} = \mathbf{F_{g'}} \backslash \{F'_m\} \tag{4}$$

Weakest link epistasis creates a selective asymmetry between traits that becomes striking in the absence of pleiotropy. In particular, advantageous mutations are fully restricted to cases in which the mutated trait m is the resident's weakest link, that is $F_m = \min(\mathbf{X})$. Provided that F'_m remains the mutant's weakest link ($f_{org}(\mathbf{F'}) = F'_m$), the selection coefficient reduces to:

$$s_m = \frac{F'_m}{\min(\mathbf{F})} - 1. \tag{5}$$

Otherwise, the (unchanged) second-worst trait becomes the mutant's new weakest link. The mutant's selective advantage thereby reaches an upper bound dictated by the second worst trait $\min(\mathbf{F_{-m}})$, above which $f_{org}(\mathbf{F'}) = \min(\mathbf{F_{-m}})$, despite $F'_m > \min(\mathbf{F_{-m}})$. In such cases:

$$s_{max} = \frac{\min(\mathbf{F_{-m}})}{\min(\mathbf{F})} - 1, \text{with } s_{max} \geq 0 \tag{6}$$

By contrast, many mutations reducing fitness components are latent and thus invisible to selection ($s_{\text{WLE}} = 0$); this can happen for any trait except the resident's weakest link, provided that the decrease is small enough to leave the weakest link unchanged – that is, when $F_m > F'_m > \min(\mathbf{F})$. Otherwise, if the mutated trait becomes the new weakest link with $\min(\mathbf{F'}) = F'_{m'}$ the mutation is deleterious, with a selection coefficient given by Eq. (5).

Altogether, the adaptive potential of F'_m can be captured through a selection coefficient $s_{\text{WLE}}(F'_m; \mathbf{F})$ that only depends on itself and the resident set of trait-dependent fitness components $\mathbf{F}$. Based on Eq. (3), this can be formally written as:

$$s_{\text{WLE}}(F'_m; \mathbf{F}) = \begin{cases} \min(s_m, s_{max}), & \text{if } F_m = \min(\mathbf{F}) \text{ is the weakest link,} \\ \min(s_m, 0), & \text{otherwise.} \end{cases} \tag{7}$$

This expression captures both the mutational asymmetry between traits and the intrinsic limitation on advantageous mutations. The same logic extends naturally to pleiotropic mutations, since the selection coefficient depends on a single mutant trait, its weakest link, and how it compares to the resident's one. Hence, mutations that improve several or even many fitness components may still be deleterious if they substantially impair another trait, as often occurs when traits are subject to a trade-off. Due to this limitation on the supply of advantageous mutations, we now consider the genetic load induced by WLE.

2.1.2 Genetic load under weakest link epistasis

By removing unfit genotypes, natural selection biases the observed distribution of fitness, theoretically driving it to the optimum. However, mutations recurrently introduce deleterious genotypes that can segregate and even fix through random drift—despite being disfavoured—thereby limiting fitness improvement. This mutation-selection-drift balance leads to an evolutionarily stationary state for the distribution of fitness $\rho^*(f)$ whose mean deviates from the optimum. The resulting

extent to which the evolutionary process is suboptimal can be quantified through the stationary genetic load L^*, which corresponds to the relative difference $L^* = \frac{f_{opt} - \langle f \rangle^*}{f_{opt}}$ between mean stationary fitness $\langle f \rangle^*$ and the fitness optimum f_{opt}.

One key determinant of the genetic load is the distribution of fitness overall possible genotypes (in the absence of selection), whose density we write as $\rho(f)$. To investigate the impact of WLE on the genetic load, we thus need to evaluate first how it influences $\rho(f)$. For that purpose, we assume trait-dependent fitness components to be independent and identically distributed random variables so that the genotypic density $\rho(F_d)$ is known for each F_d. The genotypic density of fitness $\rho(f)$ can then directly be obtained by combining the respective densities $\rho(F_d)$ with the fitness model $f_{org}(\mathbf{F})$. Under WLE, $f_{org}(\mathbf{F}) = \min(\mathbf{F})$ such that $\rho_{\text{WLE}}(f)$ is given by the distribution of the minimum of the set $\mathbf{F}$. For analytical tractability, we thus choose a distribution for fitness components whose 1st order statistic $F_{(1)} = \min\left\{F_1, F_2, ..., F_n\right\}$ is tractable. The Beta distribution, when applied uniformly across traits such that each F_d obeys $\rho(F_d) \sim \text{Beta}(a, b)$, matches this requirement when its first parameter is set to $a = 1$, which we discuss in more depth in Appendix B.1.1 in relation to the Kumaraswamy distribution (see also [35, 36]). Hence, the density of fitness components reads:

$$\rho(F_d) = \frac{(1 - F_d)^{b-1}}{\int_0^1 (1 - F_d)^{b-1} \, \mathrm{d}F_d} \tag{8}$$

Here, the parameter b creates a bias toward maladaptive phenotypes on each trait by increasing the number of genotypes leading to low-fitness values. From now on, we refer to this as the maladaptive bias. Since $\text{Beta}(1, 1) \overset{d}{=} \text{Uniform}(0, 1)$, the unbiased case occurs when $a = b = 1$, and we can capture the strength of this per-trait maladaptive bias using $\log(b)$ (see Figure 2-A, assuming $n = 1$). Because we regard fitness as being the minimum of the n fitness components, the distribution of the first order statistics $\rho(F_{(1)})$ for a sample of n independent draws from this distribution directly gives the genotypic fitness density under WLE (see Appendix B.1.1 for a proof):

$$\rho_{\text{WLE}}(f) = \frac{(1 - f)^{nb-1}}{\int_0^1 (1 - f)^{nb-1} \, \mathrm{d}f} \sim \text{Beta}(1, nb) \tag{9}$$

Remarkably, this expression highlights how dimensionality n mirrors the role of the maladaptive bias b, which motivates defining nb as an effective maladaptive bias (see Figure 2-A).

To gain insight into the genetic load, we can focus on the stationary state reached under the weak mutation regime, where a mutant genotype either gets fixed or lost [4, 12, 37, 38]. In this scenario, the evolutionary dynamics can be captured by a Markov chain where the transition rates from genotypic state g to g' is simply given by the product of the mutation rate $\mu_{gg'}$ and the probability of fixation $P_{\text{fix}}(g'; g)$ of g' in a population where g is the resident. Since this Markov chain can be shown to be reversible, it obeys detailed balance at stationarity, meaning forward and backward substitutions between two mutationally connected genotypes (from g to g', and in reverse) occur at the same rate [4, 37]. Extending this detailed balance principle to all mutational neighbours (accessible through 1 mutational step), there exists a tractable solution for the stationary density $\rho^*(f)$ of the Markov chain owing to the link between $P_{\text{fix}}(g'; g)$, f_g and $f_{g'}$. Under the haploid Moran scenario and pairwise symmetrical mutations $\mu_{gg'} = \mu_{g'g}$, this stationary fitness distribution has density [4]:

$$\rho^*(f) = \frac{f^{N-1} \, \rho(f)}{\int_0^1 f^{N-1} \, \rho(f) \, \mathrm{d}f}, \tag{10}$$

where N denotes the effective population size, which controls how strongly selection favours higher-fitness genotypes.

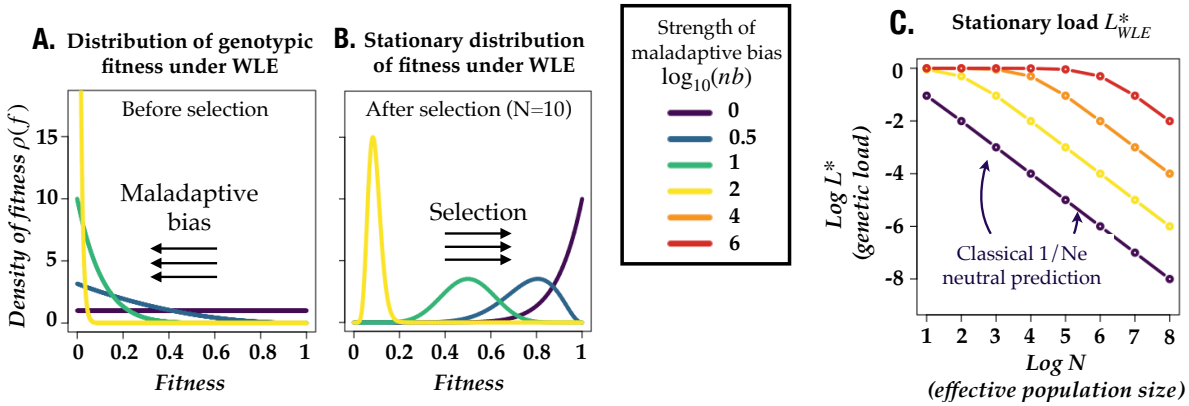

Figure 2. Distribution of fitness and evolution of maladaptation under weakest-link epistasis. Results are shown for varying levels of effective maladaptive biases, ranging from no bias ($log(nb) = 0$, dark blue) to strong bias ($log(nb) = 6$, red). **A.** Distribution of genotypic fitness $\rho_{\text{WLE}}(f)$ expected in the absence of selection. When each of the n trait-dependent fitness components follows Beta$(1, b)$, genotypic fitness is distributed as Beta$(1, nb)$. In this scenario, the per-trait maladaptive bias b and the number of traits n combine symmetrically to push the distribution toward lower fitness values. **B.** Evolutionarily stationary distribution of fitness $\rho^*_{\text{WLE}}(f)$ after selection (illustrated for N=10). Fitness still obeys a Beta distribution, Beta(N, nb), but compared to panel A, the distribution is shifted toward higher fitness by selection. **C.** Evolutionarily stationary load L^* for different effective population sizes N, plotted on log-log scales. The classical population genetics prediction for the genetic load of a single trait scales as $1/N$ and appears as a straight decreasing line with slope -1. Increasing the effective maladaptive bias, either through a larger number of traits n or through a stronger intrinsic bias b, causes substantial maladaptation ($log L^* \approx 0$ so that $L^* \to 1$) even for large population sizes.

Using the density function of beta distributions, we can now plug Eq. (9) into Eq. (10) to get the stationary distribution for fitness under WLE indicating that the stationary distribution of f is distributed as Beta(N, nb):

$$\rho^*_{\text{WLE}}(f) = \frac{f^{N-1}\,(1-f)^{nb-1}}{\int_0^1 f^{N-1}\,(1-f)^{nb-1}\,\mathrm{d}f}. \tag{11}$$

This expression also clearly shows the intuitive antagonistic effect between N and the effective bias $B = nb$ (see Figure 2-B). This can be further formalised by deriving mean stationary fitness $\langle f \rangle^*$ as the first moment of the stationary distribution, which also coincides with the ratio between the (N)th and $(N-1)$th moments of the distribution of genotypic fitness implied by Eq. (9). In our fitness space with upper bound 1, the stationary genetic load reduces to $L^* = 1 - \langle f \rangle^*$, and hence WLE leads to:

$$L^*_{\text{WLE}} = 1 - \left\langle f_{\text{WLE}} \right\rangle^* = \frac{nb}{N + nb} \tag{12}$$

The saturating function in Eq. (12) shows that the genetic load scales linearly with nb/N when $N \gg nb$, while otherwise the expression indicates significant maladaptation, especially in the hypothetical case where nb would be greater than N (see Figure 2-C). Note that, when nb and N are comparable, the stationary fitness distribution is akin to a Gaussian distribution (see Figure 2-B).

Of particular importance is the independence of the genetic load and the mutational scheme (provided that mutation rates are pairwise symmetric)[4, 37]. Consequently, the genetic load of Eq. (12) holds irrespectively of an organism's degree of pleiotropy [37]: for instance, in a purely non-pleiotropic scenario where mutations would always affect one trait (and hence its fitness component) at a time, the genetic load is exactly the same than for any arbitrary complex pleiotropic scenario.

2.2 Unifying invariances in complexity-driven genetic load

In the special case of a uniform distribution of trait-dependent fitness components ($a = b = 1$) the WLE stationary genetic load obtained in Eq. (12) is:

$$L^*_{\text{WLE, unbiased}} = n/(n+N) = L^* \tag{13}$$

If we again assume a uniform distribution of trait-dependent fitness components ($a = b = 1$) and instead let the organismal fitness function be $f_{org}(r) = 1 - r_{\text{FGM}}(\mathbf{X})$, with $r_{\text{FGM}}(\mathbf{X})$ denoting the n-dimensional Euclidean distance between the phenotype $\mathbf{X}$ and the optimum, we recover a version of Fisher's geometric model that has previously appeared in the literature [2, 4, 39], albeit with a more explicit mutational scheme. Interestingly, the load for this version of Fisher's geometric model is identical to the one we obtain here [2, 37]. This echoes previous results showing the invariance of the genetic load in modified versions of FGM with ellipsoidal rather than spherical fitness contours [12]. Together, these findings seem to suggest a broader unifying principle in which the genetic load is invariant across a wide range of fitness landscape geometries.

2.2.1 A sufficient condition for genetic load invariance

To elucidate why seemingly different fitness models yield the same evolutionary outcome, we can return to Eq. (10), in which the only factor that may differ from one fitness model to another is the density of genotypic fitness $\rho(f)$. Consequently, the genetic load obtained in Eq. (13), and indeed the entire stationary fitness distribution $\rho^*(f)$, remains invariant for all scenarios in which:

$$\rho(f) \propto (1-f)^{n-1}. \tag{14}$$

To see how this scaling arises in both WLE and FGM, it is helpful to work with fitness defects rather than fitness components (Figure 3A). To that end, we define the fitness defect associated with trait d as $X_d = 1 - F_d$. Since fitness components cannot exceed their optimum $F_d = 1$, each defect satisfies $0 \leq X_d \leq 1$, and the optimal phenotype corresponds to $(0, \dots, 0)$ in the space of defects. We then combine the trait-dependent fitness defects $\mathbf{X} = (X_1, ..., X_n)$ into the organismal fitness defect r through a rule $r_{org}(\mathbf{X})$, which in turn relates to organismal fitness as $f_{org}(\mathbf{F}) = 1 - r_{org}(\mathbf{X})$. Accordingly, the density of fitness defects $\rho(r)$ is simply related to the density of fitnesses $\rho(f)$ by $\rho(r) = \rho(1-f)$. Through this lens, WLE can be naturally reinterpreted: fitness is fully dictated by the distance between the optimum and the most deleterious fitness component, so that $r_{\text{WLE}}(\mathbf{X}) = \max(\mathbf{X})$, recovering $f_{org}(\mathbf{X}) = 1 - r_{\text{WLE}}(\mathbf{X}) = 1 - \max(\mathbf{X}) = \min(\mathbf{F})$.

Our goal is now to identify a broad class of rules $r(\mathbf{X})$, which, when combined with $f_{org} = 1 - r(\mathbf{X})$, result in $\rho(f) \propto (1-f)^{n-1}$. Working in fitness defects, the corresponding condition is:

$$\rho(r) \propto r^{n-1}. \tag{15}$$

Figure 3B-D show several rules r_{org} that satisfy this scaling for the case $b = 1$, specifically $r_{org}(\mathbf{X}) = \sum_{d=1}^{n} X_d$ in Figure 3B with diagonal equal fitness contours, r_{FGM} in Figure 3C with circular equal fitness contours, and r_{WLE} in Figure 3C with "L"-shaped equal fitness contours. What all of these rules have in common is that they produce contours (i.e. level sets) that are simply rescaled versions of each other. Intuitively, this explains the $n-1$ power scaling of the density, since $\rho(f)$ is proportional to the ($n-1$ dimensional) surface area of the corresponding contour and in n-dimensions surface area scales as the $n-1$-th power of length. Because the relevant rescaling occurs around the optimum rather than the origin, we can also describe the fitness contours as differing by what is known as a homothetic transformation (see Figure 3).

To derive a sufficient condition for the $n-1$-th power scaling when $b = 1$, it is helpful to formulate the analysis in terms of fitness defects, which yields:

$$r_{org}(\lambda \mathbf{X}) = \lambda r_{org}(\mathbf{X}) \tag{16}$$

for $\lambda > 0$. Technically this is known as r_{org} being a homogeneous function of degree 1, but intuitively it states that e.g. halving each individual fitness defect X_d also halves the organismal fitness defect r_{org}. More generally, one can think of this condition as implying that the logic of how trait-dependent fitness defects combine is independent of the overall scale of these defects. This also has specific implications for mutational effects. For example, consider a mutation that always sets the first trait to half its current fitness defect. If introducing this mutation in some background $\mathbf{X}$ makes $r_{opt}(\mathbf{X})$ one third of its current value, then the same mutation introduced in a background $c\mathbf{X}$ will also reduce $r_{opt}(\mathbf{cX})$ to one third its previous value.

Thus far, we have been working in the special case $b = 1$, so that genotypes are assumed to have trait-dependent fitness defects that are jointly uniformly distributed, i.e. $\rho(\mathbf{X}) \propto 1$. We next extend these results to cases with more complex distributions of fitness defects. Intuitively, the issue is that the density of $\rho(r)$ need no longer be proportional to the surface area of the corresponding equal fitness surface because the different points on the same fitness contour have different weights, i.e. different densities with respect to $\rho(\mathbf{X})$. In the Appendix B.1.2 we show that a similar phenomenon arises under a wide range of joint probability distributions $\rho(\mathbf{X})$. Specifically, suppose that:

$$\rho(\lambda \mathbf{X}) = \lambda^{\beta} \rho(\mathbf{X}) \tag{17}$$

for all $\lambda > 0$ and some fixed value β. Technically, this says that we require that the density of fitness defects $\rho(\mathbf{x})$ be a homogeneous function of degree β. More biologically, what this condition says is that the supply of lower fitness defect states changes in a consistent manner as we approach the optimum, so that, for instance, for any $\mathbf{X}$, the density at $\mathbf{X}/2$ is 2^{β} times smaller. Similar to the case for the organismal fitness defect, this means that all of the equal density contours of $\rho(\mathbf{X})$ are scaled versions of each other, but the shape of the generating contour is highly flexible.

What we show in the Appendix B.1.2 is that for any r_{org} that is homogeneous of order 1 and any trait specific defect distribution $\rho(\mathbf{X})$ that is homogeneous of degree β, then we have:

$$\rho(r) \propto r^{\beta+n-1} \tag{18}$$

and hence at stationarity fitness is Beta$(N, \beta + n)$-distributed and in particular has load given by (see section 2.1.2 for more details on the calculation):

$$L^* = \frac{\beta + n}{N + \beta + n}. \tag{19}$$

For instance, when the trait-dependent fitness components are independently Beta$(1, b)$-distributed, $\rho((X))$ is a homogeneous function of degree $\beta = b(n-1)$ and we recover our previous expressions for the stationary fitness distribution and load (see supplemental appendix B.1.2).

This result shows that for fixed population size and a fixed degree of homogeneity for $\rho(\mathbf{X})$, if r_{org} is homogeneous of degree 1, then the whole stationary distribution is invariant to the specific choice of how trait-dependent fitness defects combine and the structure of the neutral joint distribution of trait-dependent defects.

2.2.2 Genetic load across a continuum of adaptation strategies

To more concretely explore the implications of the invariance in genetic load discussed in the previous section, we consider a family of choices for r_{org} that correspond to the ℓ^p-norm of the vector

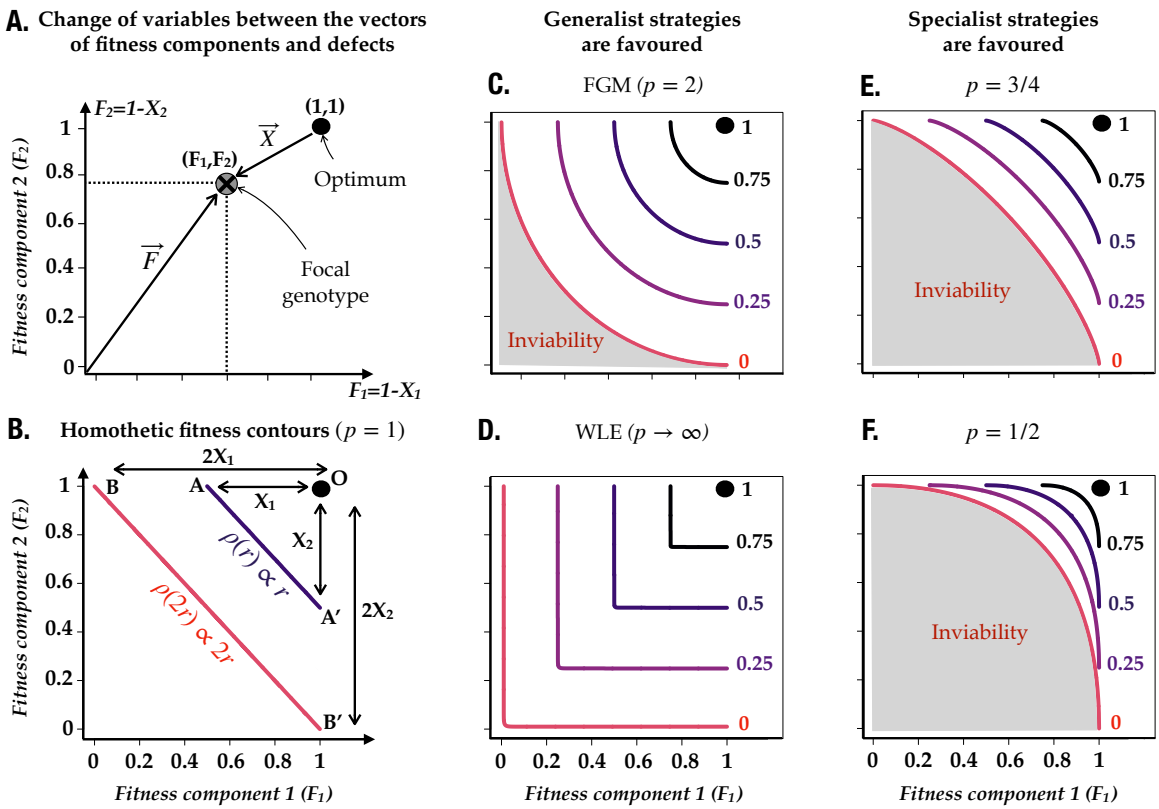

Figure 3. Fitness contours in a two-trait space. A. Change of variables used to interpret weakest link epistasis (WLE) in terms of the vector of fitness defects $\mathbf{X} = \mathbf{1} - \mathbf{F}$. We restrict $\mathbf{F}$ to the unit cube (components $\in [0, 1]$), where the fitness optimum $f = 1$ is reached when all the $F_d = 1$ (here, $F_1 = F_2 = 1$). **B.** Illustration of the homothetic transformation principle behind the genetic load invariance when the organismal fitness defect is defined as $r_{org}(\mathbf{X}) = ||\mathbf{X}||_1 = X_1 + X_2$ and fitness defects are uniformly distributed ($b = 1$). A homothety of ratio k is a transformation that maps any point at distance X_i from a choice of center (here $F_1 = F_2 = 1$) to a new point at distance kX_i (eg., A is sent to B here), preserving shapes and scaling all distances proportionally. Under uniform $\rho(\mathbf{X}\ (b = 1)$, $\rho(r)$ is proportional to the $n - 1$-dimensional surface area of the equal fitness contours and scales as r^{n-1}, where $n = 2$ traits in this example. **C–F.** Fitness contours (from $f = 0$ to $f = 1$, red-black gradient) resulting from fitness defined as $f = 1 - r_p(\mathbf{X})$, with $r_p(\mathbf{X}) = ||\mathbf{X}||_p$ denoting the genotypic fitness defect. Fitness defects are uniformly distributed ($b = 1$). The grey-shaded area represents combinations of fitness defects that are inviable, and expands as p decreases. **C–D.** Two 'generalist' strategies are shown, where improving both traits confers a higher fitness than improving a single one. The limit $p \to \infty$ (panel D) recovers WLE (see Figure 1 for more details), while contours with $p = 2$ retains the weakest-link constraint, albeit in a softened form (panel C, canonical FGM with euclidean distance). **E–F.** Two 'specialist' strategies are shown (panel E: $p = 2/3$; Panel F: $p = 1/2$) where concentrating improvement on one trait is most efficient for increasing fitness.

of fitness defects:

$$r_p(\mathbf{X}) = ||\mathbf{X}||_p = \left(|X_1|^p + \ldots + |X_n|^p\right)^{1/p}, \text{with } p \geq 1 \tag{20}$$

The ℓ^p norm provides a natural formalism for this purpose, since it satisfies the condition of homogeneity of degree 1 and recovers the Euclidean version of FGM ($p = 2$), additive trait specific fitness defects ($p = 1$), and also WLE version as a limit, where fitness is fully determined by the maximum fitness defect as $p \to \infty$ (see Figure 3B-D)

Indeed there is an important ecological meaning behind different values of the parameter p. First, increasing p from the Euclidean version of FGM ($p = 2$) makes the model more and more dependent on the weakest link, which implies that intermediate values reflect a reality in between FGM and WLE: any trait improvement still enhances adaptiveness, albeit with diminishing returns,

as often observed in biological systems [24]. In other words, higher p values reduce the ability to buffer maladaptive traits by limiting the effect (but not the supply) of compensatory mutations. Similarly, we can consider the impact of decreasing p from FGM (Figure 3E,F). Note that, in the cases where $p < 1$, $||X||_p$ is only a quasi-norm, since it fails to satisfy the triangle inequality, but it is still homogeneous of degree 1 and thereby satisfies the condition of Eq. (17). The case where $p = 1$ (also known as the Manhattan distance) represents the "symmetrical" situation, where fitness components/defects can buffer each other equally (e.g., adding 0.1 to to X_1 and X_2 is identical to adding 0.2 to X_1, or to X_n). When p falls under 1, the fitness contours tend to adopt more star-like shapes that favours putting all eggs in a single basket. In this "strongest link epistasis" (SLE) scenario, the fitness reward is higher for organisms that focus on improving one specific trait: because of the diamond shape of contours, the phenotypic distance to the optimum is much lower close to any given trait axis. Under this perspective, SLE can be seen as a model of specialisation, where all but specialist phenotypes would be driven out by natural selection.

A subtle but important implication of this model is that smaller values of p increasingly enlarge the region of trait-dependent fitness component space that results in inviable organisms ($r_p(\mathbf{X}) > 1$, implying negative fitness). For instance, in a two-trait example with $X_1 = 0.9$ and $X_2 = 0.5$, the genotype is inviable under the $r_1(\mathbf{X})$ and $r_2(\mathbf{X})$ rules (because both $r_1 > 1$ and $r_2 > 1$), yet retains viability fitness $f_{org}(\mathbf{X}) = 0.1$ under WLE. In our framework, these inviable combinations are still described by $\rho(\mathbf{X})$ but make no contribution to the stationary distribution, which we consider to have support only on the biologically relevant interval from 0 to 1.

Turning to the issue of genetic load, as in Eq. (8), we assume again the fitness components to be independently Beta-distributed as $F_d \sim \text{Beta}(1, b)$. In this case, regardless of the parameter p, the ℓ^p norms are all homogeneous of degree 1, and hence the distribution of fitnesses at stationarity and corresponding genetic load L^* are identical to that for WLE (Eqs. 9-12) (see APPENDIX B.1.2, as well as APPENDIX B.2.1 and B.2.2 for a more direct proof). The formula for the genetic load can even be further generalised by considering that fitness defects are not identically distributed and obey instead separate $\text{Beta}(b_d, 1)$ for $d = 1$ to n. Denoting mean trait-dependent fitness defect as $\bar{b}$ and assuming again $f(r_p) = 1 - r_p$, this yields (see APPENDIX B.1.2 and APPENDIX B.2.2):

$$L_p^* = \frac{n\bar{b}}{N + n\bar{b}}, \tag{21}$$

where $n\bar{b}$ quantifies the effective maladaptive bias acting on fitness. This result also extends to fitness models with hierarchical genetic architecture, where fitness is determined by the ℓ^p-norm of higher-level traits (e.g., metabolic pathways) and each of these traits is itself given by the ℓ^{p_m}-norm of underlying, lower level trait-dependent fitness defects, with p_m possibly differing across traits (see APPENDIX B.2.3 for the extension to modular inter-dependencies of traits).

Discussion

Diminishing-returns effects are widespread in biology [19, 24, 26, 28]. They arise when improving a phenotypic trait (such as enzyme activity) no longer translates into increased performance of a function (such as metabolic flux) once another component becomes limiting. In some biological systems, phenotypes depend on the efficiency of every underlying component, as in multi-step biosynthetic pathways where failure of a single enzyme can fully prevent the formation of the final product (e.g., synthesis of anthocyanin pigments in plants). From a genetics perspective, such dependencies generate weakest-link epistasis (WLE), which can substantially distort the distribution of fitness effects.

Motivated by these considerations, we first determined the selection coefficient of mutations under WLE. Using this selection coefficient, we show how, in this fitness model, trait dimensionality greatly decreases the supply of advantageous phenotypes, which, in turn, limits the adaptive potential of organisms. Quantitatively, when fitness is determined by the weakest underlying trait, we establish that the the mean stationary genetic load carried by an organism scales as $n\bar{b}/(N + n\bar{b})$, where $n\bar{b}$ represents a compound maladaptive bias. A key feature of this expression is the symmetry between n and $\bar{b}$: phenotypic complexity and intrinsically biased traits (towards low-fitness values) contribute symmetrically to the reduction of fitness.

Intriguingly, the genetic load under WLE can be shown to equal that of Fisher's geometric model (FGM) for unbiased traits ($b = 1$), when fitness decays linearly with the Euclidean distance between the phenotype and the optimum [2, 37]. To elucidate why the complexity scaling of the load is preserved across diverse fitness landscapes despite the seemingly unfavourable genetic architecture imposed by WLE, we show that this arises due to the fact that the equal fitness contours in both cases consist of scaled copies of the same basic shape. We then generalise this scaling to allow for non-uniform joint distributions of trait-dependent fitness components.

To understand the effect of maladaptive biases across more interpretable forms of epistasis, we study the case where fitness defects to combine through arbitrary ℓ^p norms. This generalisation spans a continuum of architectures, including WLE ($p \to \infty$), classical FGM ($p = 2$). This also naturally extends to combinations of ℓ^p-norms that can capture biological modularity and hierarchical trait structure (eg., enzymes forming pathways that themselves combine to determine cell efficiency). Within this broader framework, the genetic load retains the nb scaling, demonstrating that neither the architecture of traits (e.g., modularity) nor the nature of epistasis alters the evolutionary stationary state: what matters is the number of independent phenotypic traits and their average bias toward low-fitness states.

In this work, we have defined fitness as $f_{org}(\mathbf{X}) = 1 - r_{org}(\mathbf{X})$. Earlier work has shown that the invariance of the genetic load can be further generalised by decoupling the rule that specifies the phenotypic 'distance to the optimum' $r_{org}(\mathbf{X})$ from the function $f_{org}(r)$ that maps this distance onto fitness, as is classically done in FGM [12]. For instance, a broad class of definitions for the phenotypic distance to the peak yields the same load $1 - \big((N-1)/N\big)^{n/Q}$ when fitness is given by $f_{org}(r) = \exp(-r^Q)$.

Building on this invariance principle, these earlier approaches attempted top-down inferences of phenotypic dimensionality using the stationary genetic load [12, 21]. However, it remains to be determined whether biologically grounded fitness landscapes satisfy this homogeneity condition (even approximately), for instance in systems where fitness components combine multiplicatively across traits, a situation commonly encountered in the study of life-history traits [40, 41]. No less importantly, even when this condition holds, we have shown throughout this study that the symmetry between phenotypic dimensionality and maladaptive biases obscures the effects of complexity. As a consequence, stationary fitness alone cannot reliably disentangle these two sources of maladaptation unless maladaptive biases would be negligible.

This raises the question of whether there exists reliable empirical estimates of the magnitude of this (maladaptive) bias parameter b. Although the overwhelming predominance of non-functional proteins in random sequences has long been known [42], a more quantitative characterisation of biological quantities related to fitness components (eg., protein binding affinities) has only recently begun to emerge. These empirical distributions appear consistent with b values spanning several orders of magnitude [43, 44]. Crucially, this reinforces the idea that organisms may carry significant genetic loads, and strengthens the paradox of complexity: fitness gains associated with

additional functions should come at the expense of increased maladaptation, that is, a larger difference between mean fitness and the optimal phenotype's fitness for a given number of traits.

This perspective may also help clarify otherwise perplexing empirical findings. For instance, the vast majority of enzymes are only moderately efficient [45, 46], often to the point that their catalytic rates can be improved by several orders of magnitude [47]. Interpreting such patterns will require accounting in greater detail for the influence of genotype-phenotype maps on the distribution of fitnesses that would be realised under neutrality. More broadly, we believe that this framework paves the way for a more quantitative understanding of the outcome of evolution, such as the distribution of fitness effects of mutations. Likewise, it should help predict how biological systems evolve in the face of perturbations, for instance when environmental changes shift the location of the fitness peak, and how this depends on the underlying phenotype-fitness map.

Acknowledgment

We thank Etienne Rajon, Djordje Bajic and Jason Wolf for fruitful discussions that helped conceptualise the framework, and Dmitry Biba for the analogy with Liebig's barrel. DMM acknowledges funding support from NIH grant R35GM133613 and additional funding from the Simons Center for Quantitative Biology at Cold Spring Harbor Laboratory.

References

[1] Motoo Kimura, Takeo Maruyama and James F Crow. 'The mutation load in small populations'. In: *Genetics* 48.10 (Oct. 1963), pp. 1303–1312. ISSN: 1943-2631. DOI: 10.1093/genetics/48.10.1303. eprint: https://academic.oup.com/genetics/article-pdf/48/10/1303/34898065/genetics1303.pdf.

[2] Art Poon and Sarah P. Otto. 'COMPENSATING FOR OUR LOAD OF MUTATIONS: FREEZING THE MELTDOWN OF SMALL POPULATIONS'. In: *Evolution* 54.5 (2021/04/04 2000), pp. 1467–1479. DOI: https://doi.org/10.1111/j.0014-3820.2000.tb00693.x.

[3] Aneil F. Agrawal and Michael C. Whitlock. 'Mutation Load: The Fitness of Individuals in Populations Where Deleterious Alleles Are Abundant'. In: *Annual Review of Ecology, Evolution, and Systematics* 43.1 (2021/04/14 2012), pp. 115–135. DOI: 10.1146/annurev-ecolsys-110411-160257.

[4] Guy Sella and Aaron E. Hirsh. 'The application of statistical physics to evolutionary biology'. In: *Proceedings of the National Academy of Sciences of the United States of America* 102.27 (July 2005), p. 9541. DOI: 10.1073/pnas.0501865102.

[5] M. Kimura. 'On the probability of fixation of mutant genes in a population'. In: *Genetics* 47.6 (June 1962), pp. 713–719.

[6] Michael Lynch. 'Evolution of the mutation rate'. In: *Trends in Genetics* 26.8 (2025/07/09 2010), pp. 345–352. DOI: 10.1016/j.tig.2010.05.003.

[7] David M. McCandlish, Charles L. Epstein and Joshua B. Plotkin. 'Formal properties of the probability of fixation: Identities, inequalities and approximations'. In: *Theoretical Population Biology* 99 (2015), pp. 98–113. DOI: https://doi.org/10.1016/j.tpb.2014.11.004.

[8] Florian Labourel and Etienne Rajon. 'Resource uptake and the evolution of moderately efficient enzymes'. In: *Molecular Biology and Evolution* (May 2021). DOI: 10.1093/molbev/msab132.

[9] Daniel L. Hartl and Clifford H. Taubes. 'Compensatory Nearly Neutral Mutations: Selection without Adaptation'. In: *Journal of Theoretical Biology* 182.3 (1996), pp. 303–309. DOI: https://doi.org/10.1006/jtbi.1996.0168.

[10] H A Orr. 'Adaptation and the cost of complexity.' eng. In: *Evolution* 54.1 (Feb. 2000), pp. 13–20. ISSN: 0014-3820 (Print); 0014-3820 (Linking). DOI: 10.1111/j.0014-3820.2000.tb00002.x.

[11] H Allen Orr. 'The genetic theory of adaptation: a brief history.' eng. In: *Nat Rev Genet* 6.2 (Feb. 2005), pp. 119–127. ISSN: 1471-0056 (Print); 1471-0056 (Linking). DOI: 10.1038/nrg1523.

[12] Olivier Tenaillon et al. 'Quantifying Organismal Complexity using a Population Genetic Approach'. In: *PLOS ONE* 2.2 (Feb. 2007), e217–.

[13] Hervé Le Nagard, Lin Chao and Olivier Tenaillon. 'The emergence of complexity and restricted pleiotropy in adapting networks.' eng. In: *BMC Evol Biol* 11 (Nov. 2011), p. 326. ISSN: 1471-2148 (Electronic); 1471-2148 (Linking). DOI: 10.1186/1471-2148-11-326.

[14] O Tenaillon. 'The Utility of Fisher's Geometric Model in Evolutionary Genetics'. In: *Annual review of ecology, evolution, and systematics* 45 (Nov. 2014), pp. 179–201. DOI: 10.1146/annurev-ecolsys-120213-091846.

[15] R. A. Fisher. *The genetical theory of natural selection.* Oxford, England: Clarendon Press, 1930, pp. xiv, 272–xiv, 272. DOI: 10.5962/bhl.title.27468.

[16] Guillaume Martin. 'Fisher's Geometrical Model Emerges as a Property of Complex Integrated Phenotypic Networks'. In: *Genetics* 197.1 (May 2014), p. 237. DOI: 10.1534/genetics.113.160325.

[17] D. Waxman and J. J. Welch. 'Fisher's Microscope and Haldane's Ellipse.' In: *The American Naturalist* 166.4 (2025/07/11 2005), pp. 447–457. DOI: 10.1086/444404.

[18] D. Waxman. 'Fisher's geometrical model of evolutionary adaptation—Beyond spherical geometry'. In: *Journal of Theoretical Biology* 241.4 (2006), pp. 887–895. DOI: https://doi.org/10.1016/j.jtbi.2006.01.024.

[19] Xiao Yi and Antony M Dean. 'Adaptive Landscapes in the Age of Synthetic Biology'. In: *Molecular biology and evolution* 36.5 (May 2019), pp. 890–907. DOI: 10.1093/molbev/msz004.

[20] Olivier Cotto and Troy Day. 'A null model for the distribution of fitness effects of mutations'. In: *Proceedings of the National Academy of Sciences* 120.23 (2025/07/11 2023), e2218200120. DOI: 10.1073/pnas.2218200120.

[21] Guillaume Martin and Thomas Lenormand. 'A general multivariate extension of Fisher's geometrical model and the distribution of mutation fitness effects across species.' eng. In: *Evolution* 60.5 (May 2006), pp. 893–907. ISSN: 0014-3820 (Print); 0014-3820 (Linking).

[22] Thomas R. Sinclair and Wayne I. Park. 'Inadequacy of the Liebig Limiting-Factor Paradigm for Explaining Varying Crop Yields'. In: *Agronomy Journal* 85.3 (2025/07/11 1993), pp. 742–746. DOI: https://doi.org/10.2134/agronj1993.00021962008500030040x.

[23] Alexander N. Gorban et al. 'Law of the Minimum Paradoxes'. In: *Bulletin of Mathematical Biology* 73.9 (2011), pp. 2013–2044. DOI: 10.1007/s11538-010-9597-1.

[24] Jumpei F. Yamagishi and Tetsuhiro S. Hatakeyama. 'Global constraint principle for microbial growth laws'. In: *Proceedings of the National Academy of Sciences* 122.40 (2025/11/15 2025), e2515031122. DOI: 10.1073/pnas.2515031122.

[25] Florian Labourel. 'A mechanistic view of the evolution of cell diversity : or when constraints meet chance and necessity in life's struggle for multiplication'. 13th Dec. 2021.

[26] H Kacser and J A Burns. 'The control of flux.' eng. In: *Symp Soc Exp Biol* 27 (1973), pp. 65–104. ISSN: 0081-1386 (Print); 0081-1386 (Linking).

[27] R Heinrich and T A Rapoport. 'A linear steady-state treatment of enzymatic chains. General properties, control and effector strength.' eng. In: *Eur J Biochem* 42.1 (Feb. 1974), pp. 89–95. ISSN: 0014-2956 (Print); 0014-2956 (Linking). DOI: 10.1111/j.1432-1033.1974.tb03318.x.

[28] D L Hartl, D E Dykhuizen and A M Dean. 'Limits of adaptation: the evolution of selective neutrality'. In: *Genetics* 111.3 (Nov. 1985), pp. 655–674.

[29] D E Dykhuizen, A M Dean and D L Hartl. 'Metabolic flux and fitness'. In: *Genetics* 115.1 (Jan. 1987), pp. 25–31.

[30] Charlotte Coton et al. 'Evolution of enzyme levels in metabolic pathways: A theoretical approach. Part 1'. In: *Journal of Theoretical Biology* 538 (2022), p. 111015. DOI: https://doi.org/10.1016/j.jtbi.2022.111015.

[31] D. de Vienne, C. Coton and C. Dillmann. 'The genotype–phenotype relationship and evolutionary genetics in the light of the Metabolic Control Analysis'. In: *Biosystems* 232 (2023), p. 105000. DOI: https://doi.org/10.1016/j.biosystems.2023.105000.

[32] David Heckmann, Daniel C. Zielinski and Bernhard O. Palsson. 'Modeling genome-wide enzyme evolution predicts strong epistasis underlying catalytic turnover rates'. In: *Nature Communications* 9.1 (2018), p. 5270. DOI: 10.1038/s41467-018-07649-1.

[33] Lydia Robert et al. 'Mutation dynamics and fitness effects followed in single cells'. In: *Science* 359.6381 (2025/07/28 2018), pp. 1283–1286. DOI: 10.1126/science.aan0797.

[34] Philip J. Gerrish and Richard E. Lenski. 'The fate of competing beneficial mutations in an asexual population'. In: *Mutation and Evolution*. Ed. by Ronny C. Woodruff and James N. Thompson. Dordrecht: Springer Netherlands, 1998, pp. 127–144. ISBN: 978-94-011-5210-5. DOI: 10.1007/978-94-011-5210-5{_}12.

[35] P. Kumaraswamy. 'A generalized probability density function for double-bounded random processes'. In: *Journal of Hydrology* 46.1 (1980), pp. 79–88. DOI: https://doi.org/10.1016/0022-1694(80)90036-0.

[36] M. C. Jones. 'Kumaraswamy's distribution: A beta-type distribution with some tractability advantages'. In: *Statistical Methodology* 6.1 (2009), pp. 70–81. DOI: https://doi.org/10.1016/j.stamet.2008.04.001.

[37] Guy Sella. 'An exact steady state solution of Fisher's geometric model and other models'. In: *Theoretical Population Biology* 75.1 (2009), pp. 30–34. DOI: https://doi.org/10.1016/j.tpb.2008.10.001.

[38] David M McCandlish and Arlin Stoltzfus. 'Modeling evolution using the probability of fixation: history and implications.' eng. In: *Q Rev Biol* 89.3 (Sept. 2014), pp. 225–252. ISSN: 0033-5770 (Print); 0033-5770 (Linking). DOI: 10.1086/677571.

[39] D L Hartl and C H Taubes. 'Towards a theory of evolutionary adaptation.' eng. In: *Genetica* 102-103.1-6 (1998), pp. 525–533. ISSN: 0016-6707 (Print); 0016-6707 (Linking).

[40] Stephen C. Stearns. 'Life-History Tactics: A Review of the Ideas'. In: *The Quarterly Review of Biology* 51.1 (2021/10/19 1976), pp. 3–47. DOI: 10.1086/409052.

[41] D. Roff. *Evolution of life histories: theory and analysis*. Springer Science & Business Media, 1993.

[42] Anthony D. Keefe and Jack W. Szostak. 'Functional proteins from a random-sequence library'. In: *Nature* 410.6829 (2001), pp. 715–718. DOI: 10.1038/35070613.

[43] Mato Lagator et al. 'Predicting bacterial promoter function and evolution from random sequences'. In: *eLife* 11 (2022), e64543. DOI: 10.7554/eLife.64543.

[44] Haiqing Xu et al. 'Chance promoter activities illuminate the origins of eukaryotic intergenic transcriptions'. In: *Nature Communications* 14.1 (2023), p. 1826. DOI: 10.1038/s41467-023-37610-w.

[45] Arren Bar-Even et al. 'The Moderately Efficient Enzyme: Evolutionary and Physicochemical Trends Shaping Enzyme Parameters'. In: *Biochemistry* 50.21 (May 2011), pp. 4402–4410. DOI: 10.1021/bi2002289.

[46] Arren Bar-Even et al. 'The Moderately Efficient Enzyme: Futile Encounters and Enzyme Floppiness'. In: *Biochemistry* 54.32 (Aug. 2015), pp. 4969–4977. DOI: 10.1021/acs.biochem.5b00621.

[47] Moshe Goldsmith and Dan S Tawfik. 'Enzyme engineering: reaching the maximal catalytic efficiency peak'. In: *Current Opinion in Structural Biology* 47 (2017), pp. 140–150. DOI: https://doi.org/10.1016/j.sbi.2017.09.002.

[48] D. Song and A. K. Gupta. 'Lp-Norm Uniform Distribution'. In: *Proceedings of the American Mathematical Society* 125.2 (2025/08/02/ 1997), pp. 595–601.

Supplementary Materials

Supporting text A: Notation

Definition	Parameter, relations & acronyms
Weakest link epistasis	WLE
Fisher's geometric model	FGM
Resident genotype	g
Mutant genotype	g'
Trait-dependent fitness component d of genotype g	$F_{g,d}$
Trait-dependent fitness defect d of genotype g	$X_{g,d}$
Number of traits	n
Population size	N
Maladaptive trait bias towards low-fitness values	b
Vector of fitness components of genotype g	$\mathbf{F_g} = (F_{g,1}, ..., F_{g,n})$
Vector of fitness defects of genotype g	$\mathbf{X_g} = (X_{g,1}, ..., X_{g,n})$
Organismal fitness defect of $\mathbf{X}$	$r_{org}(\mathbf{X})$
ℓ^p-norm of fitness defects	$r_p(\mathbf{X}) = \|\|\mathbf{X}\|\|_p$
Organismal fitness of $\mathbf{X}$ or $\mathbf{F}$	$f_{org}(\mathbf{F}) = 1 - r_{org}(\mathbf{X})$
Effective selection coefficient of a mutant under WLE	s_{WLE}
Joint density of fitness defects	$\rho(\mathbf{X})$
Joint density of fitness components	$\rho(\mathbf{F})$
Genotypic density of organismal fitness defect	$\rho(r)$
Genotypic density of fitness	$\rho(f)$
Stationary density of fitness	$\rho^*(f)$
Expected fitness at stationarity	$\langle f \rangle^*$
Genetic load at stationarity	L^*=1-$\langle f \rangle^*$

Supporting text B: Mathematical appendix

B.1 Genetic load under various fitness models

B.1.1 WLE genetic load for the Kumaraswamy distribution

In this section, we relax the assumption that trait-dependent fitness components F_d are Beta$(1, b)$ distributed. The first order statistics of a set of beta-distributed variables does not generally lead to a closed-form solution (when $a \neq 1$), which precludes further analysis. The Kumaraswamy distribution $\mathcal{K}umar(a, b)$, (with shape parameters a and b) is a Beta-like distribution with compact support ([35]; see Figure 4 for the shape of these distributions) for which the first-order statistic has a closed-form regardless of parameters a and b: if the vector of fitness components $\mathbf{F}$ contains n random variables $F_d \sim \mathcal{K}umar(a, b)$, then the distribution of their minimum is also Kumaraswamy-distributed [35, 36], such that :

$$F_{(1)} \sim \mathcal{K}umar(a, nb) \tag{22}$$

This arises because the influence of the parameter a is slightly modified compared to the Beta-distribution, as reflected by its probability density function. Specifically, the distribution of F_d when $F_d \sim \mathcal{K}umar(a, b)$ is given by:

$$\rho(F_d) = \frac{(F_d)^{a-1}(1-(F_d)^a)^{b-1}}{\int_0^1 (F_d)^{a-1}(1-(F_d)^a)^{b-1}\,\mathrm{d}F_d} \tag{23}$$

where the difference with the Beta-distribution is in the term $(1-(F_d)^a)$, which would simply be $(1-F_d)$ in a Beta-distribution.

Under weakest link epistasis, fitness is determined by the minimum of fitness components. Hence, Eq. (22) corresponds to the genotypic distribution of fitness $\rho_{\text{WLE}}(f)$ in the absence of selection.

Assuming a haploid population and a Moran evolutionary process, we can use Eq. (10) to derive the distribution of fitness at evolutionary stationarity:

$$\rho^*(f) = \frac{f^{N-1}nabf^{a-1}(1-f^a)^{nb-1}}{\int_0^1 f^{N-1}nabf^{a-1}(1-f^a)^{nb-1}\,\mathrm{d}f} \tag{24}$$

Fitness at evolutionary stationarity is thus given by a ratio of raw moments (see main text for explanations) that are analytically tractable for the Kumaraswamy distribution [36]:

$$\langle f\rangle^*_{\text{WLE}} = \frac{\mathcal{B}(1+N/a, nb)}{\mathcal{B}(1+(N-1)/a, nb)} = \frac{\Gamma(1+N/a)\Gamma(1+nb+(N-1)/a)}{\Gamma(1+(N-1)/a)\Gamma(1+nb+N/a)}, \tag{25}$$

where $\mathcal{B}(y)$ and $\Gamma(y)$ denote the Beta and Gamma functions. From Stirling's formula, we have the approximation $\Gamma(y+s) \simeq y^s\Gamma(y)$ when y is large enough, with which we can rewrite fitness at stationary state as:

$$\langle f\rangle^*_{\text{WLE}} \approx \left(\frac{1+(N-1)/a}{1+nb+(N-1)/a}\right)^{1/a}. \tag{26}$$

We can rewrite Eq. (26) as $(1-y)^(1/a)$, where $y = nb/(1+nb+(N-1)/a))$ is small when $nba \ll N$. Using Taylor expansions, we can then simplify Eq. (25) so that in the haploid Moran scenario, we have:

$$L^*_{\text{WLE}} = 1 - \langle f\rangle^*_{\text{WLE}} \tag{27}$$

$$\approx \frac{nb}{nab+a+N}, \text{ if } nba \ll N. \tag{28}$$

As a first order approximation, we could neglect all terms but N in the denominator, but keeping the $nba+a$ term both improves the approximation and helps gain a better intuition of the role played by n,a and b when their product gets close to N.

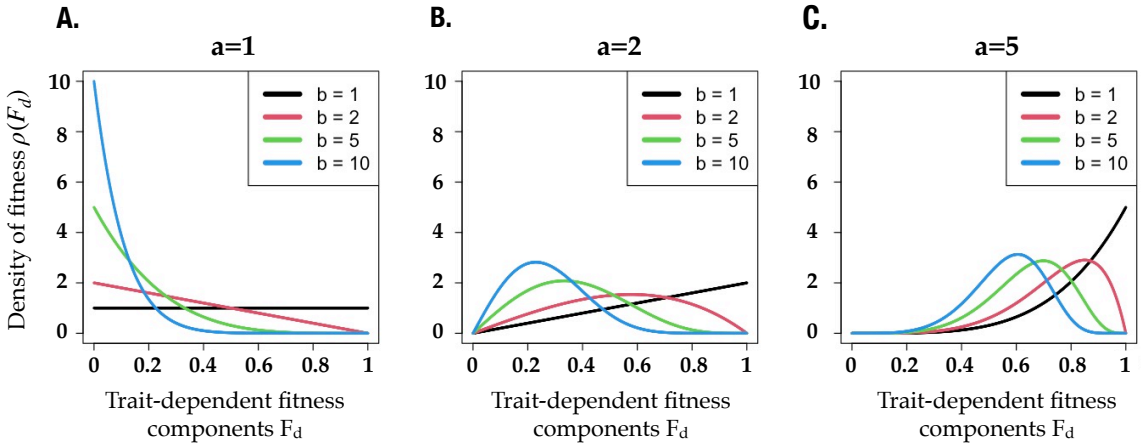

Figure 4. Shape of the genotypic distribution of trait-dependent fitness components when they are identically Kumaraswamy-distributed, $F_d \sim \mathcal{K}umar(a,b)$. Fitness components are bounded between 0 and 1 with their contribution to fitness being maximal at 1. Overall, the parameters a and b modify the genotypic distribution of trait values (i.e., prior to selection), thereby creating a bias toward deleterious traits (when $b > a$) or toward advantageous ones (when $a > b$). **A.** In the case where $a = 1$, F_d follows a Beta-distribution $F_d \sim \mathcal{B}eta(1,b)$ as modeled in the main document. The special case $a = b = 1$ coincides with a uniform distribution, and thus with the absence of bias at the trait level, whereas $b = 2$ leads to a linear decrease of the density of fitness component values (higher weight on low-fitness trait values). For larger b, the density is highly biased towards deleterious values. **B–C.** Increasing the parameter a relaxes the bias towards deleterious traits ($a = 2$, panel B), to the point where high fitness traits become over-represented (especially when $a = 5$, panel C).

From this expression, we first recover that population size N reduces the genetic load. This expression then shows that dimensionality n and the parameter b tend to increase the genetic load in a purely exchangeable way. In the main document, we use the power version of the Beta distribution where $a = 1$, such that b directly represents an intrinsic bias toward low-fitness traits. With the Kumaraswamy distribution (or a Beta-distribution where $a \neq 1$, this bias is better captured by the ratio b/a. Accordingly, while b keeps its influence on inflating low fitness genotypes, the parameter a acts as a counter-bias. Hence, it reduces the genetic load by providing evolution with a large(r) supply of high fitness genotypes. Notably, this result is robust to a broad range of fitness-component distributions, since the Kumaraswamy distribution can approximate diverse shapes, including Gaussian- and exponential-like distributions (see Figure 4). Overall, a has a compound influence that is partly akin population size, partly limiting the influence of the maladaptive bias nb.

Logically, this formula for the genetic load becomes exact when $a = 1$, as $\Gamma(y+1) = y\Gamma(y)$. Starting back from Eq. (26), we can thus write the genetic load exactly as the expression of the main document, because the Kumaraswamy distribution $\mathcal{K}umar(1,b)$ reduces to a Beta-distribution $\mathcal{B}eta(1,b)$ in this case (which also explains why the Beta distribution with parameter $a = 1$ has a tractable first order statistics, as we claim in the main document):

$$L^*_{\text{WLE}} = 1 - \langle f \rangle^*_{\text{WLE}} = \frac{nb}{N + nb}. \tag{29}$$

B.1.2 Invariance of stationary fitness distribution

Here we derive the stationary fitness distribution when both the organismal fitness and the genotypic density are homogeneous functions of the trait specific fitness defects. First, we prove the following technical lemma.

Lemma 1. *Let g be a homogeneous function of degree $\alpha > 0$ and let h be a homogeneous function of degree β, both defined on the positive orthant of $\mathbb{R}^n$. Let*

$$J(x) = \int_{\{\mathbf{Y}>0:\, g(\mathbf{Y})<x\}} h(\mathbf{Y}) d\mathbf{Y}. \tag{30}$$

Then $J(x)$ is a homogeneous function of degree

$$\gamma = \frac{n+\beta}{\alpha}. \tag{31}$$

Proof. Consider

$$J(\lambda x) = \int_{\{\mathbf{Y}>0:\, g(\mathbf{Y})<\lambda x\}} h(\mathbf{Y}) d\mathbf{Y}. \tag{32}$$

We will evaluate this integral by using the change of variables $\mathbf{Y} = \lambda^{1/\alpha}\mathbf{Z}$. This change of variables is an invertible linear transformation defined by the matrix $\lambda^{1/\alpha}\mathtt{I}$ where $\mathtt{I}$ is the identity matrix and hence the determinant of its Jacobian is simply $|\lambda^{1/\alpha}\mathtt{I}| = \lambda^{n/\alpha}$ for all $\mathbf{Z}$. Turning to the region of integration, $g(\mathbf{Y}) < \lambda x$ if and only if $g(\lambda^{1/\alpha}\mathbf{Z}) < \lambda x$, but g is homogeneous of degree α and so $g(\lambda^{1/\alpha}\mathbf{Z}) = \lambda\, g(\mathbf{Z})$. Thus, $g(\mathbf{Y}) < \lambda x$ if and only if $g(\mathbf{Z}) < x$. Likewise, since $\mathbf{Y}$ is a positive multiple $\lambda^{1/\alpha}\mathbf{Z}$ of $\mathbf{Z}$, $\mathbf{Y} > 0$ if and only if $\mathbf{Z} > 0$.

Putting this all together, we have

$$\begin{aligned}
J(\lambda x) &= \int_{\{\mathbf{Y}>0:\, g(\mathbf{Y})<\lambda x\}} h(\mathbf{Y}) d\mathbf{Y} && (33)\\
&= \int_{\{\mathbf{Z}>0:\, g(\mathbf{Z})<x\}} h(\lambda^{1/\alpha}\,\mathbf{Z})\, \lambda^{n/\alpha} d\mathbf{Z} && (34)\\
&= \int_{\{\mathbf{Z}>0:\, g(\mathbf{Z})<x\}} \lambda^{\beta/\alpha} h(\mathbf{Z})\, \lambda^{n/\alpha} d\mathbf{Z} && (35)\\
&= \lambda^{(n+\beta)/\alpha} J(x), && (36)
\end{aligned}$$

where in line 35 we used the fact that h is homogeneous of degree β. Thus $J(\lambda x) = \lambda^{(n+\beta)/\alpha} J(x)$ meaning that J is homogeneous of degree $(n+\beta)/\alpha$ as required. ■

We can now characterise the genotypic density of organismal fitness defect $\rho(r)$ induced by a homogeneous function $r_{org}(\mathbf{X})$ mapping trait-dependent fitness defects to organismal fitness defects when the genotypic density of trait-dependent fitness defects is itself also a homogeneous function.

Proposition 1. *Suppose the organismal fitness defect $r_{org}(\mathbf{X})$ is homogeneous of degree 1, and $\rho(\mathbf{X})$ is homogeneous of degree β. Then $\rho(r) \propto r^{n+\beta-1}$.*

Proof. Let

$$C(x) = \int_{\{\mathbf{X}>0:\, r(\mathbf{X})<x\}} \rho(\mathbf{X}) d\mathbf{X}, \tag{37}$$

so that $\rho(r)$ is proportional to $C'(x)$. By Lemma 1, $C(x)$ is a homogeneous function of degree $n+\beta$, and hence its derivative $C'(x)$ is a homogeneous function of degree $n+\beta-1$, as required. ■

The stationary distribution of fitness itself, and hence the expected load at stationarity, then follow as an immediate consequence of Equation 10.

Corollary 1. *Suppose the organismal fitness defect function $r_{org}(\mathbf{X})$ is homogeneous of degree 1 and $\rho(\mathbf{X})$ is homogeneous of degree β. Then for a Moran population of size N (as we used when deriving the WLE load) the stationary distribution of f is $\mathcal{B}eta(N, \beta + n)$ distributed where n is the number of traits. Moreover, the load is given by:*

$$L^* = \frac{\beta + n}{N + \beta + n}. \tag{38}$$

Proof. By Proposition 1, $\rho(r) \propto r^{n+\beta-1}$ and the density of f is simply a reflected version of the density of r, since $f = 1 - r$, so that we have $\rho(f) \propto (1-f)^{n+\beta-1}$. Plugging these values into Equation 10, we find that the stationary distribution of f has density

$$\rho^*(f) \propto f^{N-1}(1-f)^{n+\beta-1}, \tag{39}$$

which we recognize as the density of a $\mathcal{B}eta(N, \beta+n)$ distribution. Since the mean of a $\mathcal{B}eta(N, \beta+n)$ distribution is $N/(N+\beta+n)$, the stationary load L^* is

$$1 - \frac{N}{N+\beta+n} = \frac{\beta+n}{N+\beta+n}. \tag{40}$$

■

Corollary 2. *Suppose each single-trait defect X_d is homogeneous of degree $b_d - 1$, with b_d denoting the intrinsic maladaptive bias of trait d. Then for a Moran population of size N the stationary distribution of f is $\mathcal{B}eta(N, n\bar{b})$ distributed where n is the number of traits. Moreover, if $\bar{b}$ denotes the mean maladaptive bias across traits, the load is given by:*

$$L^* = \frac{n\bar{b}}{N + n\bar{b}}. \tag{41}$$

Proof. If the probability distribution of each trait $\rho(X_d)$ is homogeneous of degree $b_d - 1$, the joint distribution is given by the product of the marginal distributions. It is thus homogeneous of degree $\sum_{d=1}^{n}(b_d - 1) = n\bar{b} - n$. Replacing by $n\bar{b} - n$ in the previous corollary completes the proof for this second corollary. ■

In the next section, we use the special case of the ℓ^p norm to relate this reasoning to a full probabilistic approach.

B.2 Beta-distributed variables and the ℓ^p norm

B.2.1 Distribution of the ℓ^p norm of Beta-distributed variables

Proposition 2. *Consider a set of n independent variables $X_1, X_2, ..., X_d, ..., X_n$ each distributed as $X_d \sim$ Beta$(a_d, 1)$ with $a_d > 0$. Conditional on $||\mathbf{X}||_p \leq 1$, the distribution obeyed by the ℓ^p norm follows:*

$$||X||_p \sim \mathcal{B}eta(n\bar{a}, 1), \tag{42}$$

with $||X||_p = (|X_1|^p + ... + |X_n|^p)^{1/p}$ and $\bar{a} = \frac{1}{n}\sum_{d=1}^{n} a_d$.

This result is well established for X_d that are uniformly distributed [48]. Below, we extend it to the case of Beta-distributed variables.

Lemma 2. *To prove the proposition in Eq. (42), we first need to show the following lemma. If $X_d \sim \mathcal{B}eta(a_d, 1)$ with $a_d > 0$, then :*

$$(X_d)^p \sim \mathcal{B}eta(a_d/p, 1) \; \forall \; p \in \mathbb{R}_{>0} \tag{43}$$

Proof. To do so, we need to determine the probability density function $f_{Z_d}(z)$ of $Z_d = (X_d)^p$. Using a simple change of variable and the chain rule (for derivatives), it is possible to write the probability density function $f_{Z_d}(z) = f_{X_d}\big(g(z)\big)g'(z)$ provided $x = g(z)$ is a strictly increasing differentiable function. Applying this to the function $x = z^{1/p}$ with $X_d \sim \mathcal{B}eta(a_d, 1)$, we have $g'(z) = z^{1/p-1}/p$. This implies:

$$f_{Z_d}(z) = a_d \; (z^{1/p})^{a_d - 1} \; z^{1/p-1}/p \tag{44}$$

$$= a_d/p \times z^{a_d/p-1}, \tag{45}$$

which completes the proof since this expression is exactly the probability density function of a Beta distribution $\mathcal{B}eta(a_d/p, 1)$. From our premises, it is besides possible to write $q = 1/p$, which immediately yields the useful corollary according to which, if $X_d \sim \mathcal{B}eta(a_d/q, 1)$ with $a_d > 0$:

$$(X_d)^{1/q} \sim \mathcal{B}eta(a_d, 1) \; \forall \; q \in \mathbb{R}_{>0} \tag{46}$$

■

The following lemma can be derived from the aggregation property of the Dirichlet distribution by aggregating the first n variables of $\text{Dir}(\alpha_1, \dots, \alpha_n, 1)$, however here for completeness we provide a direct proof.

Lemma 3. *Consider a set of n variables $Y_d \sim \mathcal{B}eta(\alpha_d, 1)$ with $\alpha_d > 0$. Conditioning on $\sum_{d=1}^{n} Y_d \leq 1$, we can show that:*

$$\sum_{d=1}^{n} Y_d \sim \mathcal{B}eta(\sum_{d=1}^{n} \alpha_d, 1) \tag{47}$$

Proof. We start with two independent variables $Y_1 \sim \mathcal{B}eta(\alpha_1, 1)$ and $Y_2 \sim \mathcal{B}eta(\alpha_2, 1)$ that have compact support on $[0, 1]$, with $\alpha_1, \alpha_2 > 0$. The probability density function for the sum $S = Y_1 + Y_2$ is given by the convolution:

$$f_S(s) = \int_0^1 f_{Y_1}(y_1) \; f_{Y_2}(s - y_1) \; \mathbf{1}_{[0,1]}(s - y_1) \, dy_1, \tag{48}$$

where $\mathbf{1}_{[0,1]}$ denotes the indicator function that enforces variables to live on their support. Because both variables Y_1 and Y_2 take values between 0 and 1, their density vanishes outside this interval. Consequently, $y_2 = s - y_1 \in [0, 1]$ has to be positive. This is enforced by the term $\mathbf{1}_{[0,1]}(s - y_1)$, and implies that the integrand is supported only for $y_1 < s$; otherwise, this would require $y_2 < 0$. We can thus write:

$$f_S(s) = \int_0^s f_{Y_1}(y_1) \; f_{Y_2}(s - y_1) \, dy_1 \tag{49}$$

$$= \int_0^s \alpha_1 y_1^{\alpha_1 - 1} \alpha_2 (s - y_1)^{\alpha_2 - 1} \, dy_1 \tag{50}$$

Exchanging y_1 for us, we have $dy_1 = sdu$ and we can rewrite Eq. (50) as:

$$f_S(s) = \int_0^1 \alpha_1\alpha_2 u^{\alpha_1-1} s^{\alpha_1-1} (s(1-u))^{\alpha_2-1}\, s\, du \tag{51}$$

$$= \int_0^1 \alpha_1\alpha_2 u^{\alpha_1-1}(1-u)^{\alpha_2-1} s^{\alpha_1+\alpha_2-1}\, du \tag{52}$$

We can now include the truncation following which the density of the sum is $f_S(s > 1) = 0$ when the sum of Y_1 and Y_2 is greater than 1. The probability density function must thus be renormalised over the simplex. This can be written as:

$$f_S(s)\mathbf{1}_{[0,1]}(s) = \frac{\int_0^1 \alpha_1\alpha_2 u^{\alpha_1-1}(1-u)^{\alpha_2-1} s^{\alpha_1+\alpha_2-1}\, du}{\int_0^1\int_0^1 \alpha_1\alpha_2 u^{\alpha_1-1}(1-u)^{\alpha_2-1} s^{\alpha_1+\alpha_2-1}\, du\, ds} \tag{53}$$

where normalisation occurs via the denominator. Since u is not a function of s, the integration over u is independent of s. Consequently, all terms that depend on u cancel out in the expression, which reduces to:

$$f_S(s)\mathbf{1}_{[0,1]}(s) = \frac{s^{\alpha_1+\alpha_2-1}}{\int_0^1 s^{\alpha_1+\alpha_2-1}\, ds} = (\alpha_1+\alpha_2)s^{\alpha_1+\alpha_2-1} \tag{54}$$

By definition of the Beta distribution, we have shown through Eq. (54) that the sum of Y_1 and Y_2 obeys $S \sim \mathcal{B}eta(\alpha_1+\alpha_2, 1)$.

Remarking that the second parameter of the distribution is conserved by the summation, we can extend this property by combining the sum S with a third variable Y_3, so that $(S+Y_3) \sim \mathcal{B}eta((\alpha_1+\alpha_2)+\alpha_3, 1)$. Extending it further to n variables demonstrates the above lemma. ■

We can now combine Lemmas (2) and (3) to demonstrate the proposition (2).

Proof. First, including Lemma (2) into Lemma (3) yields:

$$\sum_{d=1}^{n} X_d^p \sim \mathcal{B}eta(\sum_{d=1}^{n} a_d/p, 1) \tag{55}$$

when we condition on $\sum_{d=1}^{n}(X_d)^p \leq 1$. Second, we can notice that this latter requirement is equivalent to $||\mathbf{X}||_p \leq 1$ because variables $X_d \in [0,1]$, and hence $|X_d| = X_d$. Therefore, using Lemma (2) again on the sum $(\sum_{d=1}^{n}|X_d|^p)^{1/p}$ concludes the proof by canceling out the influence of p in the distribution of Eq. (55) and making it equal to Eq. (42). ■

B.2.2 ℓ^p norm and the distribution of fitness defects

We can now apply the property given by condition in Eq. (42) to the vector of fitness defects. We have defined this vector as $\mathbf{X}$, where $X_d \sim \text{Beta}(b_d, 1)$. From Eq. (42), we have:

$$||\mathbf{X}||_p \sim \mathcal{B}eta(n\bar{b}, 1) \tag{56}$$

With $r_p = ||\mathbf{X}||_p$ and $f = 1 - r_p$, we then immediately recover that:

$$f \sim \mathcal{B}eta(1, n\bar{b}) \tag{57}$$

owing to reflection symmetry. This is equivalent to Eq. (9) except it now relies on the average maladaptive bias $\bar{b}$ rather than b. As an immediate consequence, the genetic load takes the form of Eq. (12), with $\bar{b}$ instead of b.

B.2.3 Distribution of the sum of Beta-distributed ℓ^p-norms conditioning on $||\mathbf{X}||_p \leq 1$, and connection to modularity

Distribution of the sum of ℓ^p-norms conditional on $||\mathbf{X}||_p \leq 1$

There is another useful consequence of the fact that the ℓ^p norm of a vector of independent Beta distributed variables $X_d \sim \mathcal{B}eta(a_d, 1)$ remains Beta-distributed when we condition on $||\mathbf{X}||_p \leq 1$: we can now apply this result to determine the distribution followed by the combination of ℓ^p norms themselves.

First, consider M subsets of n_m independent random variables following Beta-distributions:

$$X_d^{(m)} \sim \mathcal{B}eta(a_{d,m}, 1), \text{with } d = 1, \dots, n_m,$$

and the group-level measure of each subset as the ℓ^{p_m} norm:

$$r_m(\mathbf{X^{(m)}}) = ||\mathbf{X^{(m)}}||_{p_m}.$$

Conditional on $||\mathbf{X}^{(m)}||_{p_m} \leq 1$, our prior results imply that:

$$||\mathbf{X^{(m)}}||_{p_m} \sim \mathcal{B}eta(n_m \overline{a}_m, 1), \text{with } \overline{a}_m = \frac{1}{a_m} \sum_{i=1}^{n_m} a_{d,m} \tag{58}$$

From that point, we can now determine the distribution followed by the sum of several ℓ^p-norms, and by taking the ℓ^p-norm of underlying ℓ^p-norms.

Corollary 3. *Consider M subsets of vectors $\mathbf{X^{(m)}}$ where $m = 1, \dots, M$, such that :*

$$||\mathbf{X^{(m)}}||_{p_m} \sim \mathcal{B}eta(n_m \overline{a}_m, 1).$$

Now, consider the set of these M ℓ^{p_m}-norms $\mathbf{R} = (||\mathbf{X^{(1)}}||_{p_1}, \dots, ||\mathbf{X^{(M)}}||_{p_M})$ of order p_m. If the ℓ^{p_m}-norms are combined through a higher-level ℓ^p-norm, the resulting distribution of $||\mathbf{R}||_p$ is still Beta-distributed. Importantly, this property does not depend on the specific values of the exponents p_m, or the number of hierarchical layers.

Proof. Consider a vector containing M group-level norms:

$$\mathbf{R} = (r_1(\mathbf{X^{(1)}}), \dots, r_M(\mathbf{X^{(M)}})), \text{with } r_m(\mathbf{X^{(m)}}) = ||\mathbf{X^{(m)}}||_{p_m}$$

Conditioning on $||\mathbf{R}||_p \leq 1$, we can apply property (2), and the ℓ^p norm $||\mathbf{R}||_p$ is again Beta-distributed according to:

$$||\mathbf{R}||_p \sim \mathcal{B}eta(n\overline{a}, 1) \tag{59}$$

with:

$$n = \sum_{m=1}^{M} n_m \quad \text{and} \quad \overline{a} = \sum_{m=1}^{M} \sum_{d=1}^{n_m} a_{d,m}/n$$

Consequently, the distribution of the combined ℓ^p norm is completely independent of the number of layers, the choice of the orders p_m at each layer, and the final order p. ■

Biological modularity and the genetic load under maladaptive biases

We can use this property to clarify how the genetic load invariance relates to the architecture of traits. An organism may consist of multiple functional modules, within which lower-level traits combine in different ways. For example, the size of an organism may arise as the sum of the sizes of smaller components (eg., bones in a body), which corresponds to the taxicab ℓ^1 norm. conversely, a metabolic rate or the synthesis of a pigment may require all lower-level components to contribute jointly (eg., several glycolytic enzymes are involved to produce ATP), which corresponds to the maximum ℓ^∞ norm in the space of defects. These higher-level traits may then combine to prescribe fitness. Eq. (59) demonstrates that neither modular nor hierarchical architectures of traits modify the genetic load. This latter simply depends on the total number of traits and on the average maladaptive bias of genotypes toward producing low-fitness trait values.