Stabilizing selection on a polygenic trait from the gene's-eye view.

Philibert Courau[1,2,3,*], Amaury Lambert[1,2], Emmanuel Schertzer[3]

November 2025

1. Institut de Biologie de l'ENS (IBENS), École Normale Supérieure, Paris, France 75005
2. Centre Interdisciplinaire de Recherche en Biologie (CIRB), Collège de France, Paris, France 75005
3. Department of Statistics, University of Vienna, Vienna, Austria 1090

∗ Corresponding author: philibert.couraugmail.com.

Abstract

We study a polygenic trait under stabilizing selection at statistical equilibrium, where genetic effect, mutation rate and mutational bias are heterogeneous across loci. The model assumes L biallelic sites subject to reversible mutations, each allele described by its frequency in the population. Using a diffusion approximation, a mean-field approximation and neglecting linkage disequilibrium, we predict consistent phenomena across several regimes of selection: (1) a small deviation Δ^* of the trait mean from its optimal value appears and persists due to genetic mutations not aligned with selection; (2) while this deviation is often undetectable at the trait level, it leaves a substantial signature at the locus level by favoring alleles reducing it, resulting in genic selection with mean coefficient s^* proportional to $-\Delta^*$ acting pervasively; (3) with stronger selection on the trait, (3a) the value of Δ^* is decreased but the intensity of genic selection is increased in inverse proportion, resulting in an essentially constant, non negligible value of s^*. We show how the stationary distribution of allelic frequencies can be obtained from Δ^*. The latter can then be characterized as the solution to a fixed-point equation. Finally, we quantify several macroscopic observables of interest (genetic variance, description of the fluctuations of the trait mean as an Ornstein-Uhlenbeck process). The orders of magnitude of the macroscopic observables can be derived on a wide region of the parameter space. The model shows good fit and can straightforwardly be extended to accommodate pleiotropy, dominance, and some forms of epistasis. We also discuss the different breakdown which may occur (Bulmer effect, Hill-Robertson effect, breakdown of the Ornstein-Uhlenbeck approximation for the dynamics of the trait mean, depletion of genetic variability due to low mutation rates).

Adaptation; Genetic diversity ; Stabilizing selection ; Population genetics ; Genetic architecture ; Mutational bias ; Polygenic adaptation

Historically, the infinitesimal model [1] has been considered to bridge population genetics, which is the field that studies the evolution of the genetic composition of populations, and quantitative genetics, which is the field that studies the evolution of quantitative traits. In the infinitesimal model, (suitably transformed) quantitative traits are determined by the sum of a genetic trait value and an independent environmental effect. The genetic trait value is inherited as follows: if two unrelated organisms with respective genetic trait values z^A and z^B reproduce, their offspring has normally distributed genetic trait value with mean $(z^A + z^B)/2$ and variance V_S called the segregation variance. It was long conjectured [2] and proved in [1] that this model is compatible with Mendelian genetics, provided the quantitative trait is determined by a large number of loci acting additively on the trait, with no major-effect locus.

This allows for an autonomous description of the evolution of the trait, which we call the **trait's-eye view**. This was historically done by Wright [3] and Lande [4], and a wealth of mathematical models extending this approach has been developed since [5] with recent efforts from the partial differential equation community to determine their analytical properties and robustness (see, for instance, [6, 7]).

The segregation variance V_S is the key component of the infinitesimal model. It is supposed to result from a balance between mutations generating genetic variability and selection or genetic drift eroding it. However, in trait-based models, such as Lande's famous study of stabilizing selection [4], the segregation variance is a fixed parameter of the model. A significant improvement to traditional quantitative genetic models should consist in modeling jointly the genetic composition of the population with the evolution of the trait [8].

Recently, some authors have used tools such as Wright's formula [9] to study the dynamics of the frequency of a single allele *conditional on the evolution of the trait.* Assuming linkage equilibrium, this formula gives the effective selection coefficient at a biallelic locus as a function of the gradient of mean logfitness in the population. This approach was used to interpret Genome-Wide Association Studies (GWAS) results in [10], and to model adaptation after a change of environment in [11, 12].

In principle, since the genes determine the trait and only the genes are transmitted, the genetic composition of the population evolves autonomously, without the need to specify any auxiliary variable, and from it the trait distribution should be deduced. In the present article, we seek to fullfil this program and start directly from the allelic frequencies at all loci underlying a polygenic trait. We assume stabilizing selection which, as emphasized in [13], is a realistic assumption for wild populations.

The difficulty of this approach is that genes are coupled in a nontrivial manner by recombination and selection. Now, if we assume that genes are in perfect linkage equilibrium (see SI F.2 for a mathematical discussion about the domain of validity of this assumption), then selection remains as the only coupling force–the selective advantage of an allele at a given locus depends on the entire trait distribution, which itself depends on all other loci.

We show using varied mathematical tools (diffusion approximation, mean field theory, fixed-point equations, slow-fast approximations) that this dependence vanishes with the number of loci.

Following the philosophical literature on population genetics [14], we call the **gene's-eye view** this approach, which studies the distribution of a polygenic trait determined by a large number of unlinked loci, through the dynamics of their allelic frequencies, to which diffusion approximations and mean-field approximations are applied. This approach arose recently, notably for stabilizing selection in the seminal paper by Charlesworth [8]. Concerning directional selection, a similar philosophy is used to describe spatial populations in [15, 16]. It lets us characterize properties of the trait distribution that emerge directly from the sole description of genes, such as the genetic variance, the distance to the selection optimum, and the magnitude of the fluctuations of the population mean trait value at statistical equilibrium–though our model can in principle be extended to describe the dynamics out-of-equilibrium.

Our computations distinguish three regimes which we call weak, moderate and strong selection, based on the value of the **selection-drift ratio**

$$\omega_e^{-2} := 2N(L\bar{\alpha})^2\omega^{-2}$$

compared to powers of L, where N is the (haploid) effective population size, ω^{-2} is the strength of stabilizing selection, $\bar{\alpha}$ is the average effect size of a locus and L is the number of loci.

1 Description of the population model

1.1 The individual-based model

We consider a large, panmictic population composed of N diploid organisms, displaying an additive quantitative trait z controlled by L biallelic loci. We assume that $N \gg 1$ and $L \gg 1$.

A genotype $g = (g_\ell)_{\ell\in[L]}$ determines the value $Z(g)$ of the trait as

$$Z(g) = z_0 + \sum_{\ell\in[L]} \alpha_\ell g_\ell,$$

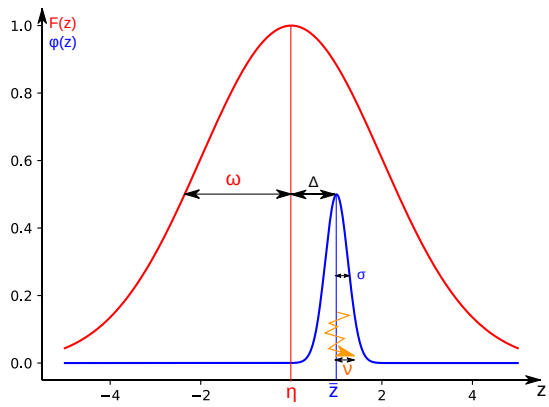

Figure 1: Phenotypic observables. The wide curve is the Gaussian fitness function $F(z) = e^{W(z)}$, which has typical width ω and is centered on η. The narrow curve is the distribution of the trait in the population, which has mean $\bar{z}$ and variance σ^2. The deviation of $\bar{z}$ from the optimum is $\Delta = \bar{z} - \eta$. The broken line represents the fluctuations of $\bar{z}$ through time, which have magnitude ν.

where $g_\ell \in \{0, 1, 2\}$ is the gene content and $\alpha_\ell \geq 0$ is the genetic additive effect of the trait-increasing allele at locus ℓ. We define the **mean additive effect** $\bar{\alpha} := (1/L) \sum_\ell \alpha_\ell$. The model can accomodate extensions to polyploidy, pleiotropy, dominance or some forms of epistasis (see SI I). We assume that $z_0 = 0$ without loss of generality.

We define the probability of mutation per generation from the trait-decreasing allele to the trait-increasing allele at locus ℓ, and back, as $\mu_\ell = (\mu_\ell^+, \mu_\ell^-)$. We assume that the $(\alpha_\ell, \mu_\ell)_{\ell \in [L]}$ are exchangeable, such that (α_ℓ, μ_ℓ) has a prescribed distribution which is not heavy-tailed, so that in particular there are no large-effect mutations.

The logfitness of an organism with trait value z is given by

$$W(z) = -\frac{1}{2\omega^2}(z - \eta)^2 \tag{1}$$

with ω^{-2} the **selection strength** and η the **selection optimum**.

Reproduction occurs via Wright-Fisher sampling: at each generation, organisms independently choose two parents with probability proportional to their fitness and the two parental genomes recombine with one or several crossovers.

Importantly, model details regarding recombination do not matter, but we assume that recombination is strong enough to allow us to neglect linkage disequilibrium (this assumption is discussed in SI F.2).

1.2 Three trait values

A model studying the dynamics of a polygenic trait under stabilizing selection from the gene's-eye view distinguishes three important theoretical trait values.

The first is the **selection optimum** η.

The second is the **mutational optimum** z_M defined as the mean trait value that the population would reach at equilibrium if selection were completely absent

$$z_M := \sum_{\ell=1}^{L} \alpha_\ell \frac{2\mu_\ell^+}{\mu_\ell^+ + \mu_\ell^-}, \tag{2}$$

because $2\mu_\ell^+/(\mu_\ell^+ + \mu_\ell^-)$ is the mean value of g_ℓ at mutational equilibrium. This value has no *a priori* reason to coincide with η, and we call $z_M - \eta$ the **trait mutational bias**.

The third is the **heterozygous trait value** z_H, which is the trait value of an organism heterozygous at every locus

$$z_H := \sum_{\ell=1}^{L} \alpha_\ell = L\bar{\alpha}. \tag{3}$$

We assume that $0 < \eta < 2z_H$, so that many different genotypes can realize the selection optimum and the population close to the optimum is not depleted in genetic variability. We may also assume that the trait is measured in units such that $z_H = 1$ so that in particular $\bar{\alpha} = 1/L$.

Other key ecological observables are represented in Fig. 1 and are explicitly defined in Section 2.

Assumptions. The details of our explicit quantitative assumptions are given in SI B. Some assumptions are made for mathematical convenience, others are necessary to reach the polygenic limit, defined as the regime in which allelic frequencies obey the polygenic equation defined below in (9). Here, we informally explain the underlying philosophy.

The mathematically convenient assumptions are as follows. Define $|\bar{\theta}|$ the mean mutation rate

$$|\bar{\theta}| := \frac{1}{L} \sum_{\ell \in [L]} \theta_\ell^+ + \theta_\ell^- .$$

We assume that (A1) all loci have mutation rates of order $|\bar{\theta}|$ and additive effect $\alpha_\ell \sim 1/L$, so that no one locus disproportionately contributes to the trait or the genetic variance. We assume that (A2) mutations are weaker than genetic drift ($|\bar{\theta}| \lesssim 1$) and that (A3, A6) the mutational bias is not too large or too small. We assume that (A4) the selection optimum η is accessible, meaning many different genotypes can realize the selection optimum (in particular, the population close to the optimum will not be depleted in genetic variability). We assume that (A5) selection is strong enough to have a detectable footprint on the genes, but not so strong as to completely deplete genetic variability in the trait.

We argue in SI F that the equations we obtain further require a sufficiently large population size (N1) and a sufficient mutational input (N2, N3) to provide an accurate description of the system.

1.3 Diffusion approximation

We let P_t^ℓ denote the frequency of the trait-increasing allele at locus ℓ, where time is now measured in units of $2N$ generations; that is, P_t^ℓ represents the frequency at generation $\lfloor 2Nt \rfloor$. We define

$$\theta_\ell := 2N\mu_\ell, \quad \omega_e := \frac{\omega}{\sqrt{2N}},$$

so that θ_ℓ is the effective mutation rate, and $\omega_e^{-2} = 2N\omega^{-2}$ is the **selection-drift ratio**.

As explained in Supplementary Information (SI C), neglecting linkage disequilibrium allows approximating allelic frequency dynamics of the individual loci by Wright-Fisher diffusions which are only coupled by selection

$$\mathrm{d}P_t^\ell = s_\ell(\mathbf{P}_t) P_t^\ell (1 - P_t^\ell)\,\mathrm{d}t \;+\; \left(\theta_\ell^+ (1 - P_t^\ell) - \theta_\ell^- P_t^\ell\right)\mathrm{d}t + \sqrt{P_t^\ell (1 - P_t^\ell)}\,\mathrm{d}B_t^\ell, \tag{4}$$

where $\mathbf{P}_t = (P_t^\ell)_{\ell \in [L]}$ and $(B^\ell)_\ell$ are independent Brownian motions. Classically, the selection coefficient s_ℓ is the regression coefficient of the fitness W on genotype g_ℓ

$$s_\ell(\mathbf{P}_t) := 2N \frac{\mathbf{Cov}_{\mathbf{P}_t}\left[W(Z(g)), g_\ell\right]}{\mathbf{Var}_{\mathbf{P}_t}[g_\ell]},$$

where for $\mathbf{p} = (p_\ell)_{\ell \in [L]} \in [0,1]^L$, $\mathbf{Cov}_{\mathbf{p}}$ and $\mathbf{Var}_{\mathbf{p}}$ are the covariance and variance of a vector of independent random variables $g = (g_\ell)_{\ell \in [L]}$ such that g_ℓ is $Binomial(2, p_\ell)$ (Hardy-Weinberg Linkage Equilibrium). We can show that s_ℓ (see SI C.2) satisfies

$$s_\ell(\mathbf{P}_t) = -\frac{\alpha_\ell}{\omega_e^2}(\bar{z}_t - \eta) + \frac{\alpha_\ell^2}{\omega_e^2}\left(P_t^\ell - \frac{1}{2}\right), \tag{5}$$

where $\bar{z}_t$ is the mean trait value at time t

$$\bar{z}_t = \sum_{\ell=1}^{L} \alpha_\ell \mathbf{E}_{\mathbf{P}_t}(g_\ell) = \sum_{\ell=1}^{L} 2\alpha_\ell P_t^\ell .$$

Here, $\mathbf{E}_{\mathbf{p}}$ is the expectation associated to the random variable g just defined with law determined by $\mathbf{p}$.

As a result, allelic frequencies (P_t^ℓ) are now coupled only by their weighted mean $\bar{z}_t$. This coupling can be further simplified when L is large thanks to a mean-field approximation.

2 Macroscopic observables from gene's-eye view

2.1 Locus dynamics and mean field approximation

Mean-field theory is a powerful approach for analyzing complex systems composed of a large number L of interacting components. When the constituents interact through a weighted mean of their individual characteristics,

as $L \to \infty$ this weighted mean behaves deterministically as by the law of large numbers, and the constituents effectively behave independently [17].

In the case where the genetic effects and mutation rates $(\alpha_\ell, \theta_\ell) \equiv (\alpha, \theta)$ are equal across loci, by replacing the mean with its expectation, we get

$$\bar{z}_t = \sum_{\ell=1}^{L} 2\alpha P_t^\ell = 2L\alpha \times \left(\frac{1}{L} \sum_{\ell=1}^{L} P_t^\ell \right) \approx 2L\alpha \mathbb{E}\left[P_t\right], \tag{6}$$

where P_t has the law of P_t^ℓ (which here does not depend on ℓ). When the $(\alpha_\ell, \theta_\ell)$ are now distributed, we replace the *weighted* mean with an 'augmented' expectation which includes averaging also over the distribution of $(\alpha_\ell, \theta_\ell)$

$$\bar{z}_t \approx 2L\mathbb{E}\left[\alpha P_t\right]. \tag{7}$$

We can now define $\Delta_t := \bar{z}_t - \eta$ as the deviation of the trait mean from the optimum and Δ^* as the **mean deviation at statistical equilibrium**

$$\Delta^* := 2L\mathbb{E}^*[\alpha P_t] - \eta \tag{8}$$

where $\mathbb{E}^*$ is the expectation at equilibrium, which is also the time average

$$\mathbb{E}^*[\alpha P_t] = \lim_{T \to +\infty} \frac{1}{T} \int_{t=0}^{T} \frac{1}{L} \sum_{\ell=1}^{L} \alpha_\ell P_t^\ell.$$

Then, the allelic frequency P_t^ℓ at each locus ℓ independently obeys the following stochastic differential equation which is autonomous in the sense that each term only depends on P_t^ℓ itself

$$\mathrm{d}P_t^\ell = \xi_{\Delta^*,\alpha_\ell}(P_t^\ell) P_t^\ell (1 - P_t^\ell)\,\mathrm{d}t \;+\; \left(\theta_\ell^+ (1 - P_t^\ell) - \theta_\ell^- P_t^\ell\right) \mathrm{d}t + \sqrt{P_t^\ell (1 - P_t^\ell)}\,\mathrm{d}B_t^\ell, \tag{9}$$

where

$$\xi_{\delta,a}(x) := -\frac{a}{\omega_e^2}\delta + \frac{a^2}{\omega_e^2}\left(x - \frac{1}{2}\right). \tag{10}$$

We call (9) the polygenic equation, and the regime in which it is accurate the polygenic limit. A formal justification for this mean-field approximation is spelled out in Section 2.4, where we also describe the dynamics of the trait.

Let us briefly discuss the expression of $\xi_{\Delta^*,\alpha_\ell}$ given in (10). The second term is known as Robertson's underdominant effect [18]. The term $-\alpha_\ell \Delta^*/\omega_e^2$ can be interpreted as a *bias-correcting selection coefficient*. It reflects the selective advantage of trait-increasing alleles when $\Delta^* < 0$ (the trait is on average below the optimum) and their disadvantage when $\Delta^* > 0$ (the trait is on average above the optimum). We define s^* as its mean value

$$s^* = -\frac{\Delta^*}{L\omega_e^2}. \tag{11}$$

In the absence of selection, mutations naturally drive the trait to the mutational optimum z_M defined in (2). In view of (10) and (11), this occurs when $\omega_e^{-2} \ll L$. Then $\Delta^* \approx z_M - \eta$, and selection is negligible at the genetic and trait levels (see SI E.2). As soon as the order of magnitude of ω_e^{-2} is equal to L or higher, what shapes the value of Δ^* is the balance between mutation (pushing $\bar{z}_t$ to z_M) and selection (pushing $\bar{z}_t$ to η).

2.2 Bridging microscopic and macroscopic scales

The consistency between microscopic (locus-level) and macroscopic (trait-level) dynamics is central to the mean-field approach and allows for an explicit determination of Δ^* as the solution to a fixed-point problem.

In (9), take $(\Delta^*, \alpha_\ell, \theta_\ell) = (\delta, \alpha, \theta)$. The stationary density of this diffusion [19] is then given for $p \in (0, 1)$ by

$$\Pi_{\delta,\alpha,\theta}(p) := C_{\delta,\alpha,\theta}\, p^{2\theta^+ - 1}(1 - p)^{2\theta^- - 1} e^{2\int_0^p \xi_{\delta,\alpha}(u)\,\mathrm{d}u}, \tag{12}$$

where $C_{\delta,\alpha,\theta}$ is a normalizing constant. Using the mean-field approximation to average trait contributions across loci, from (8) we deduce that Δ^* satisfies the following fixed-point equation:

$$\Delta^* = 2L\mathbb{E}\left[\alpha \int_0^1 x\, \Pi_{\Delta^*,\alpha,\theta}(x)\mathrm{d}x\right] - \eta, \tag{13}$$

where the expectation is now only taken over the distribution of $(\alpha_\ell, \theta_\ell)$. The latter relation embodies a *self-consistency condition* from the gene's-eye view of quantitative genetics that allows tying the dynamics of allelic

frequencies with macroscopic observables. We deduce from this fixed-point equation in SI E.3 that (except when $\omega_e^{-2} \ll L$, where $\Delta^* \approx z_M - \eta$) the mean trait deviation always scales as

$$\Delta^* \sim \omega_e^2 L = \omega^2 \frac{L}{2N}. \tag{14}$$

In this equation, $L/(2N)$ can be interpreted as the genetic drift accumulated along the genome, while ω^2 is the inverse of the strength of directional selection felt by the trait away from the optimum.

Similarly, we can deduce the trait variance σ^2 as a result of the independent, genetic additive contributions

$$\sigma^2 = 2L\mathbb{E}\left[\alpha^2 \int_0^1 x(1-x)\Pi_{\Delta^*,\alpha,\theta}(x)\mathrm{d}x\right] \sim \frac{|\bar{\theta}|}{L}. \tag{15}$$

We plot in Fig. 3 the result of simulations using the individual-based model against our theoretical predictions for different selection strengths. As expected, stronger selection leads to a smaller distance to the optimum Δ^*. We also see that when selection is sufficiently strong, it depletes the genetic variance σ^2. More precisely, we see that when selection is weak, our theoretical predictions for Δ_t and σ_t match simulation results even when $L = 100$ and $N = 50$. On the other hand, when selection is strong, we see a mismatch in the genetic variance σ^2 due to the build-up of negative Linkage Disequilibrium (LD): this is the so-called Bulmer effect [20]. The mismatch is decreased when the population size is increased ($N = 500$). Furthermore, the decrease in σ^2 due to the Bulmer effect decreases the effectiveness of the bias-correcting coefficient of selection, increasing the equilibrium distance to the optimum Δ^* with respect to neutral expectations. Following the Quasi-Linkage Equilibrium approach [21], we argue in SI F.2 that the Bulmer effect can be neglected if $N \gg |\bar{\theta}| \ln(L)/(L\omega_e^2)$.

2.3 Selection regimes

As noted above, when $\omega_e^{-2} \ll L$, selection has no substantial effect on the system (SI E.2). In contrast, when $\omega_e^{-2} \gg L^2$, Robertson's underdominant effect becomes leading ($\frac{\alpha^2}{\omega_e^2} \gg 1$), resulting in a depletion of genetic variance with a high concentration of allelic frequencies at 0 and 1. In this "ultra-strong" selection regime, the intensity of selection is the main determinant of genetic variance. Inbetween these two extremes, our analysis identifies three distinct regimes, as illustrated in Fig. 2 and summarized in Table 1. Recall from (15) that the trait population variance σ^2 is always of order $|\bar{\theta}|/L$.

1. *Weak selection regime* ($\omega_e^{-2} \sim L$): In this regime, the population is concentrated away from the fitness optimum:
$$\sigma \ll \Delta^* \sim 1.$$

2. *Strong selection regime* ($\omega_e^{-2} \sim L^2$): Here, the population is tightly concentrated around the fitness optimum:
$$\Delta^* \ll \sigma \ll 1.$$
It is only in this regime that Robertson's underdominant effect becomes significant, so that genes also experience disrupting selection in addition to genic selection against mutational bias.

3. *Moderate selection regime* ($L \ll \omega_e^{-2} \ll L^2$): This regime interpolates between the weak and the strong regime, with Δ^* and σ of the same order when $\omega_e^{-2}|\bar{\theta}|^{\frac{1}{2}} \sim L^{3/2}$.

Persistence of s^* across scales. At the locus level, the mean selection coefficient of trait-increasing alleles $s^* = -\Delta^*/(L\omega_e^2)$ is proportional to the product of the trait deviation Δ^* from the optimum and of the strength of selection ω_e^{-2}. At the trait level, we see from (14) that the strength of selection has a direct effect on Δ^* which keeps the product $\Delta^*\omega_e^{-2}$ of order L. As a result, s^* remains of order 1 throughout the three regimes and as L increases. In other words, the mutational bias always has a substantial effect at the locus level. This effect is asymptotically ($L \to \infty$) independent of the strength of selection since s^* approaches some value independent of ω_e as long as $L \ll \omega_e^{-2} \ll L^2$ (SI E.4).

In particular, in this regime, which we called the moderate selection regime, the same genetic architecture, and accordingly similar macroscopic observables, can arise for very different values of ω_e^{-2}. We illustrate this in Figure 4, in which the theoretical prediction for the rescaled genetic variance $L\sigma^2$ is plotted in the limit $L \to +\infty$, as a function of the **selection power** defined with $\omega_e^{-2} = L^b$, equivalent to

$$b := \frac{\ln(\omega_e^{-2})}{\ln(L)}. \tag{16}$$

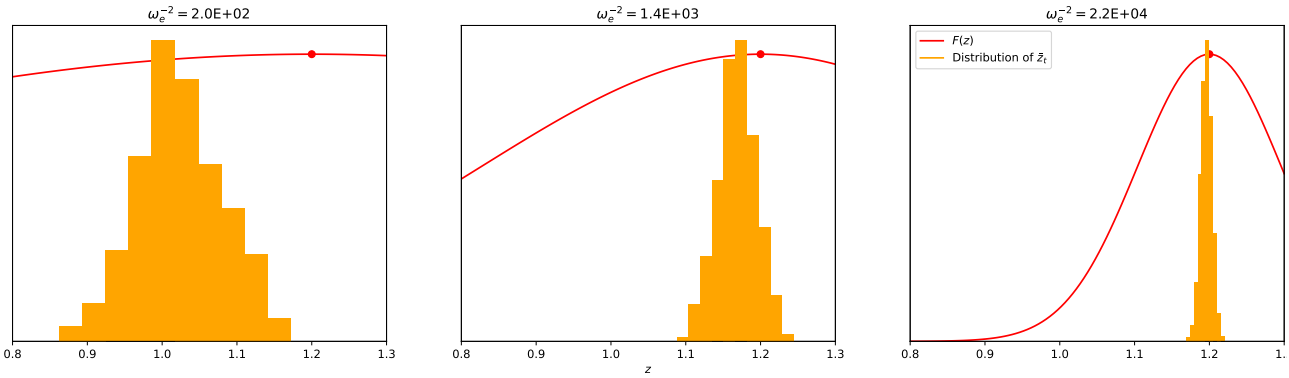

Figure 2: The histogram shows the distribution of the population mean trait $\bar{z}_t$ and not a snapshot of the trait distribution at a given time. For three different values of ω_e corresponding respectively to weak, moderate and strong selection, we show a histogram of 10,000 values of $\bar{z}_t$ after a burn-in of 24,000 generations. The parameters used are $N = 100, L = 100$, $\eta = 1.2, \theta = (0.1, 0.2)$, and the parameters $(\alpha_\ell)_{\ell \in [L]}$ were sampled with distribution $Exponential(L)$. We superimpose in red the fitness function $F(z)$, and the dot corresponds to η. In particular, Δ^* is the distance between the mean of the orange distribution and the red dot, and ν is the width of the orange distribution. As the strength of selection increases, both Δ^* and ν decrease. Only in strong selection (right panel) do we see Δ^* and ν being of the same order of magnitude.

In this figure, it can be seen that $L\sigma^2$ is expected to converge to a step function of b, with a single value corresponding to $1 < b < 2$ and discontinuities at $b = 1$ (weak selection regime) and $b = 2$ (strong selection regime). We also expect $\Delta^* L^{b-1}$ to be independent of $b \in (1, 2)$ (see Table 1). Note that the mutational bias only has a macroscopically detectable effect $\Delta^* \gtrsim 1$ if $\sigma \lesssim \Delta^*$, that is, when $\omega_e^{-2}|\bar{\theta}|^{1/2} \lesssim L^{3/2}$ (see Table 1). This means that the strength of selection ω_e^{-2} cannot be inferred when $|\bar{\theta}|^{-1/2} L^{3/2} \ll \omega_e^{-2} \ll L^2$, neither from macroscopic data, because Δ^* cannot be measured, nor from genomic data, because the distribution of P_t is independent of ω_e^{-2}.

Non-monotonicity of the variance. Another counterintuitive prediction of our model is that σ does not always decrease with higher selection pressure. In the absence of selection, when mutational bias is strong, allelic frequencies tend to accumulate close to 0 or to 1 depending on the sign of mutational bias. Increasing selection at the trait level (parameter b in Fig. 4, panels D, E, F) induces bias-correcting directional selection at the locus level that recenters the frequencies, thus increasing the population variance. We prove in SI H that when the $(\alpha_\ell, \theta_\ell)$ are constant across loci, the criterion for some level of weak selection to increase σ is that going from the mutational optimum z_M to the selection optimum η brings the trait closer to the heterozygous trait value z_H. That is: $\eta, z_H < z_M$ or $\eta, z_H > z_M$. To put it another way, starting from a situation with very weak selection, where $\mathbb{E}^*[\bar{z}] \simeq z_M$, decreasing ω (increasing the strength of selection) brings $\mathbb{E}^*[\bar{z}]$ closer to η. If, in so doing, $\mathbb{E}^*[\bar{z}_t]$ also gets closer to z_H, then the action of selection will lead some loci which the mutational bias kept at one boundary closer to 1/2, and σ will increase for entropic reasons.

Heterogeneity across loci. In Figure 5 we plot the joint distribution of (P_t^ℓ, α_ℓ), assuming that mutation rates are symmetric ($\theta^+ = \theta^- = 1/2$) and that the trait mutational bias $z_M - \eta$ is negative, thereby selecting trait-increasing alleles.

As expected, the distribution of allelic frequencies is observed to be approximately uniform in the absence of selection and to be biased to the right in the presence of substantial selection. In the moderate regime, the bias-correcting selection term indeed shifts the distribution to the right, effectively counteracting the asymmetry of mutation rates. In the strong regime, the disruptive selection driven by Robertson's underdominant effect (10) results in a U-shape distribution, increasing the weight on extreme allelic frequencies.

2.4 Trait dynamics

In the SI D, we show that in the moderate and strong regimes, the fluctuations of the mean deviation Δ_t away from its statistical average Δ^* evolves according to an Ornstein–Uhlenbeck process:

$$\mathrm{d}\Delta_t = \rho(\Delta^* - \Delta_t)\,\mathrm{d}t + \nu\sqrt{2\rho}\,dB_t, \tag{17}$$

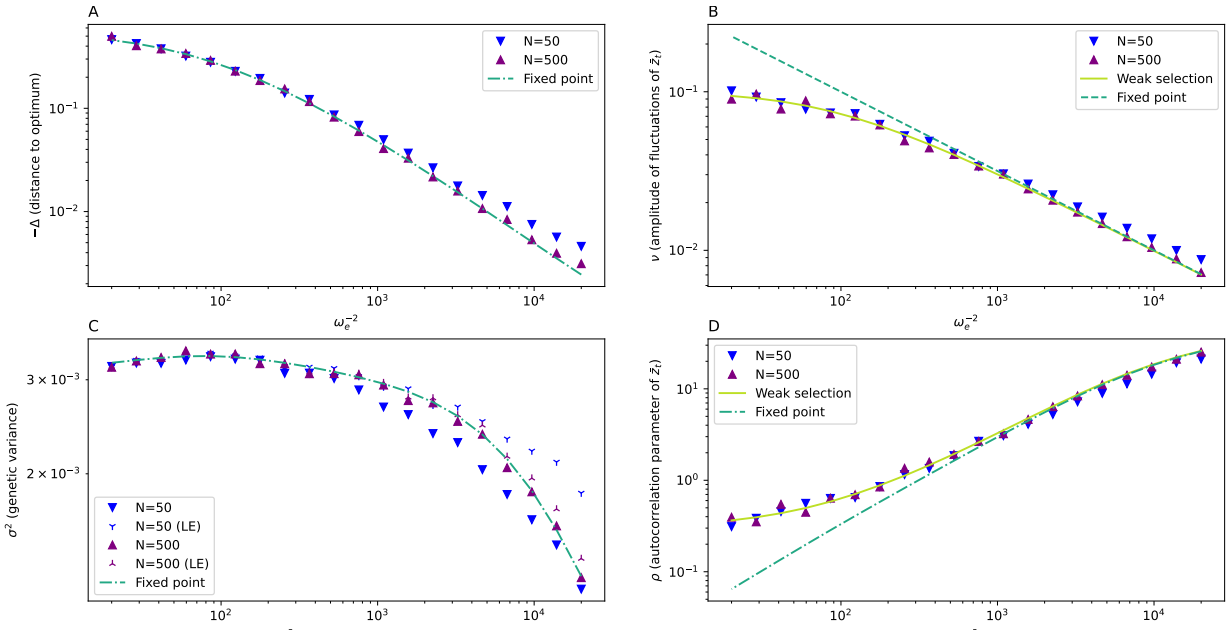

Figure 3: Comparison of theoretical predictions with numerical simulations for the macroscopic observables of Fig. 1 for varying selection regimes, from weak selection (left) to strong selection (right). The simulations were carried out with $L = 100$, the same mutation probabilities for all loci $\theta = (0.1, 0.2)$, and additive effects $(\alpha_\ell)_{\ell\in[L]}$ exponentially distributed with parameter L (in particular $\bar{\alpha} = 1/L$), so that $z_M = 2/3$. The selection optimum is $\eta = 1.2$. The same $(\alpha_\ell)_{\ell\in[L]}$ were used in all simulations, and the predictions were made conditional on the $(\alpha_\ell)_{\ell\in[L]}$. The simulations were run for $T = 500N$ generations, and each observable was measured as an average over the last $250N$ generations. For the magnitude of the fluctuations ν and the autocorrelation parameter ρ, the predictions for weak selection use the corrections derived in SI F.5. The predictions of the fixed point equation are derived in SI B.3. The predictions for moderate selection are derived in SI E.4. For the genetic variance, the simulation results distinguish between the genetic variance in the trait within the population (filled triangles), and the genic variance (three-pointed stars), which is the variance in the trait if the population was in linkage equilibrium (neglecting correlations between pairs of loci).

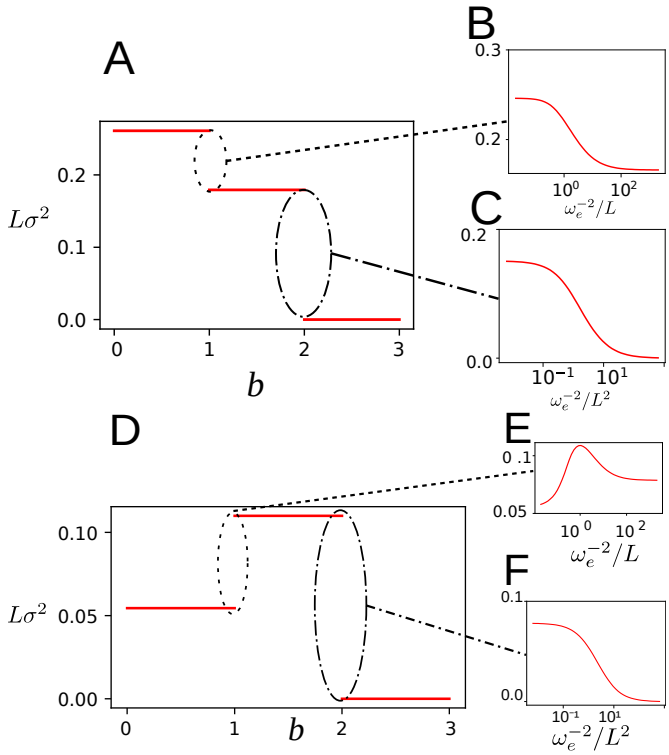

Figure 4: The predicted genetic variance at equilibrium in a limit system with $L \to +\infty, N \to +\infty$ with no mutational bias (**A**,**B**,**C** with $\theta_\ell = (0.1, 0.1)$) or strong mutational bias (**D**,**E**,**F** with $\theta_\ell = (0.01, 0.1)$), as a function of the selection power b (see (16)) and the selection-drift ratio ω_e^{-2}. We set $\eta = 1.5$ and let $(\alpha_\ell)_{\ell \in [L]}$ be exponentially distributed with parameter L (in particular $\bar{\alpha} = \frac{1}{L}$). **A**&**D** We expect the rescaled genetic variance $L\sigma^2$ in the limit to converge to a step function of the parameter b. When $b < 1$ the genetic variance is equivalent to that of a neutral model (no selection), when $b > 2$ we expect the genetic variance to be completely depleted, and when $b \in (1, 2)$, $L\sigma^2$ converge to the variance corresponding to moderate selection. The behavior at the critical points $b = 1$ and $b = 2$ correspond respectively to weak selection (**B** and **E**) and strong selection (**C** and **F**).

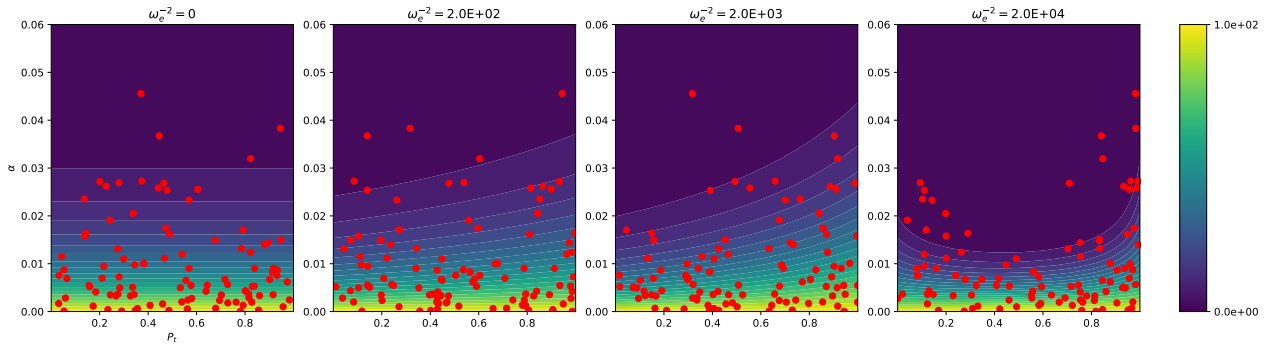

Figure 5: For no selection, weak, moderate and strong selection respectively, we plot the joint distribution of (P_t^ℓ, α_ℓ) after a single run. We chose symmetric mutation rates $\theta = (0.5, 0.5)$ to make the transition between the selection regimes more apparent. The other parameters are $L = 100, N = 500, \eta = 1.2$ and α_ℓ is distributed as $Exponential(L)$. The colour plot correspond to the predicted densities in the corresponding regimes.

Observable	Symbol	Magnitude	Eq.
Mean trait deviation from optimum	Δ^*	$L^{-(b-1)}$	(8)
Std deviation of trait distribution	σ	$\lvert\bar{\theta}\rvert^{1/2}L^{-1/2}$	(15)
Timescale of trait mean fluctuation	τ	$\lvert\bar{\theta}\rvert^{-1}L^{-(b-1)}$	(19)
Amplitude of trait mean fluctuation	ν	$L^{-b/2}$	(19)
Bias-correcting selection coeff.	s^*	1	(11)
Robertson's under-dominant effect	$\frac{\bar{\alpha}^2}{\omega_e^2}$	$L^{-(2-b)}$	(10)

Table 1: Order of magnitude of macroscopic observables as a function of the number L of loci and $b \in [1,2]$, assuming $\omega_e^{-2} \sim L^b$, interpolating between the weak regime ($b = 1$), the moderate regime ($b \in (1,2)$) and the strong regime ($b = 2$). This table is derived in SI E

In this equation, ν captures the magnitude of the stochastic fluctuations and ρ the autocorrelation parameter of $\bar{z}_t$ such that

$$\mathbb{C}\mathbf{ov}[\bar{z}_t, \bar{z}_{t+u}] = \nu^2 e^{-\rho u}. \tag{18}$$

We find in SI E.3 that

$$\nu \sim \omega_e \quad \text{and} \quad \rho \sim \frac{1}{\tau}, \tag{19}$$

where we define the characteristic timescale of the trait $\tau := \frac{L\omega_e^2}{|\bar{\theta}|}$. Under weak selection ($\omega_e^{-2} \sim L$), we show in SI F.5 that (17) is no longer justified mathematically (see Fig. S7). We derive a correction for ν^2 and ρ^2 which, though not mathematically rigorous, yields a good numerical approximation when mutation rates $|\theta_\ell|$ are constant across loci. See the right panels of Fig. 3 for a comparison between theoretical predictions and individual-based simulations for ρ and ν. Increasing the strength of stabilizing selection decreases the magnitude of the fluctuations of $\bar{z}_t$ (ν^2), but accelerates them (increase in ρ). We plot in Fig. S7 the validity of equation (18) for the autocorrelation: it fits well even for weak selection, in which we expect this equation not to be valid.

Justifying the mean-field approximation. In Section 2.1, the locus dynamics was obtained by formally replacing Δ_t with its average Δ^*. How do we justify this approximation? From (5), we have

$$s_\ell(\mathbf{P}_t) = -\frac{\alpha}{\omega_e^2}\left(\Delta^* + \varepsilon_t\right) + \frac{\alpha^2}{\omega_e^2}\left(P_t^\ell - \frac{1}{2}\right), \tag{20}$$

where $\varepsilon_t := \Delta_t - \Delta^*$ is the fluctuation of the trait mean.

In the weak/moderate regimes ($L \lesssim \omega_e^{-2} \ll L^2$), the fluctuations are negligible compared to the mean deviation Δ^* ($\nu \ll \Delta^*$ in Table 1) and a classical mean-field approximation applies.

The strong regime ($\omega_e^{-2} \sim L^2$) is more subtle and can only be justified by a *dynamical* mean-field approach. As already discussed, fluctuations cannot be ignored ($\nu \sim \Delta^*$, see Table 1) and at first sight, it seems that Δ_t is not approximately constant equal to Δ^* in (20). To address this issue, we adopt a *slow-fast approximation* [22]. Since the trait mean Δ_t fluctuates on a much faster timescale than the locus-level dynamics, this separation of timescales leads to a *slow-fast averaging effect*, allowing us to effectively decouple macroscopic (trait-level) and microscopic (locus-level) dynamics. As a result, in (5), we may still, despite non-negligible fluctuations, approximate Δ_t by Δ^*, now seen as a time average.

2.5 Breakdown

In SI F, we discuss the breakdown of our results, which can come from several factors. First, in Section F.2 we discuss how the assumption of HWLE may break down due to two different causes. The Bulmer effect [23] appears if the population size is too small with respect to the strength of selection, as can be seen in Figure 3C where the genic variance (labelled 'σ^2 (LE)') differs from the genetic variance (labelled 'σ') for $N = 50$ under strong selection. Under this effect, selection generates negative LD by favoring genomes in which genes are negatively correlated. The Hill-Robertson effect [24] similarly results in negative LD when the population size is too small, but may occur even if the population is very far from the optimum, in which case selection is well described as directional. It results from the fact that as LD randomly fluctuates, when it is positive selection

and recombination jointly act to deplete it, whereas when it is negative only recombination acts to restore LE (see Figure S3 and S4).

Second, in Section F.3 we argue that the mean-field approximation does not hold if the mutation rate is too small, specifically if $L|\bar{\theta}| \lesssim 1$. This assumption is crucial to ensure that many loci are segregating.

Third, in Section F.5, we claim that the description of the fluctuations of $(\Delta_t)_{t\geq 0}$ as an Ornstein-Uhlenbeck process in (17) breaks down if the error on the mean-field approximations of (17) is of the same order as the fluctuations of $(\Delta_t)_{t\geq 0}$, which occurs in particular for weak selection ($\omega_e^{-2} \sim L$). This breakdown can be corrected in a numerically satisfying way (as in Figure 3B and D) but not in a mathematically rigorous way (see Figure S7).

We also illustrate in Figure S6 the need for the $(\alpha_\ell)_{\ell\in[L]}$ to have well-behaved moments. Finite variance seems sufficient for practical purposes, but we recall [25] found rough estimates of tail exponents in additive effects between 1 and 2.5 (see also [26]).

Of note, we do not observe qualitative breakdown in which the population dynamics completely shift to genotype selection, described as clonal condensation in [27]. We also do not observe the appearance of sweeps at individual loci as was described [28]. It seems both these phenomena (which were to our knowledge not studied under stabilizing selection) can only occur under stabilizing selection in the "ultra-strong" selection regime $\omega_e^{-2} \gg L^2$, when selection completely overwhelms genetic drift, or in the out-of-equilibrium dynamics if the system is originally very far from the optimum [29].

2.6 Extensions

Many extensions are possible. In SI I.1 we mention polyploidy, for which the results are essentially unchanged, replacing $2N_e$ with kN_e where k is the ploidy. In SI I.2 we add pleiotropy: each allele influences d traits. Future work could study this behavior when $d \gg 1$, which should agree with the conclusions of [10, 30].

2.6.1 Dominance

In SI I.3 we add dominance: this means the trait function Z is now of the form

$$Z(g) = \sum_{\ell\in[L]} \alpha_\ell g_\ell + D_\ell \mathbb{1}_{[g_\ell=1]}$$

where D_ℓ is the dominance coefficient at locus ℓ. The selection coefficient on locus ℓ is then $s_\ell(\mathbf{P}_t) = \xi_{\Delta_t,\alpha_\ell,D_\ell}(P_t^\ell)$ where for $\delta \in \mathbb{R}, a \in \mathbb{R}_+, D \in \mathbb{R}$ we define

$$\begin{aligned}
\xi_{\delta,a,D}(p) &:= -\frac{\beta_{a,D}(p)}{\omega_e^2}\delta + \omega_e^{-2}\left(\beta_{a,D}(p)^2\left(p-\frac{1}{2}\right) + 2aDp(1-p)\right)\\
\Delta_t &= \mathbf{E}_{\mathbf{P}_t}[Z(g)] - \eta \;=\; 2\sum_{\ell\in[L]}\left(\alpha_\ell P_t^\ell + D_\ell P_t^\ell(1-P_t^\ell)\right) \;-\; \eta.\\
\beta_{a,D}(p) &:= a + D(1-2p)
\end{aligned}$$

It is remarkable that whereas $\xi_{\Delta^*,\alpha_\ell}(P_t^\ell)$ only was a polynomial of order 1 in P_t^ℓ in the additive case above (10), in the dominance case $\xi_{\Delta_t,\alpha_\ell,D_\ell}(P_t^\ell)$ is a polynomial of order 3 in P_t^ℓ.

Replacing Δ_t with $\Delta^* := \mathbb{E}[\Delta_t]$, we find the following approximate equation for $(P_t^\ell)_{t\geq 0}$

$$\mathrm{d}P_t^\ell = s_\ell(\mathbf{P}_t)P_t^\ell(1-P_t^\ell)\,\mathrm{d}t \;+\; \left(\theta_\ell^+(1-P_t^\ell) - \theta_\ell^- P_t^\ell\right)\mathrm{d}t + \sqrt{P_t^\ell(1-P_t^\ell)}\,\mathrm{d}B_t^\ell,$$

From there, we find the stationary distribution of P_t^ℓ to be $\Pi_{\Delta^*,\alpha_\ell,D_\ell,\theta_\ell}$ where we define

$$\Pi_{\delta,a,D,\theta}(p) := C_{\delta,a,D,\theta}\, p^{2\theta^+-1}(1-p)^{2\theta^--1}e^{2\int_0^p \xi_{\delta,a,D}(u)\,\mathrm{d}u}$$

with $C_{\delta,a,D,\theta}$ a normalization constant.

Finally, the definition of Δ^* lets us obtain the following fixed-point equation

$$\Delta^* \simeq 2I_{dom}(\Delta^*) - \eta$$

where

$$I_{dom}(\delta) := \mathbb{E}^*\left[\int L(\alpha p + Dp(1-p))\Pi_{\delta,\alpha,D,\theta}(p)\mathrm{d}p\right]$$

where the expectation is over (α, D, θ) taken to have the same law as $(\alpha_{\ell_U}, D_{\ell_U}, \theta_{\ell_U})$ for ℓ_U sampled uniformly at random over $[L]$.

2.6.2 Epistasis

In SI J we discuss epistasis. In terms of model, this means the trait function Z may include interaction terms between locus ℓ and locus ℓ' for $\ell' \neq \ell$. We discuss two extreme cases for which our approach can be adapted to account for epistasis.

Small clumps of interacting loci. The first model we discuss for epistasis is the "small clumps" model (Section J.2.1). In this model, the genome $[L]$ can be partitioned into a large number of small clumps $(\mathcal{I}_k)_{k\in[n]}$, such that $Z(g)$ can be obtained as the sum of $Z_k(g_{\mathcal{I}_k})$ where $g_{\mathcal{I}_k} = (g_\ell)_{\ell\in\mathcal{I}_k}$

$$Z(g) = \sum_{k\in[n]} Z_k(g_{\mathcal{I}_k}). \tag{21}$$

For the sake of illustration, we consider a model with clumps of size 2, in which $I_k = \{2k-1, 2k\}$ and (see Section J.2.1)

$$Z(g) = \sum_{k\in[L/2]} \alpha_{2k-1,2k}\, g_{2k-1}\, g_{2k}.$$

More generally for any clump model, we compute the selection coefficient at locus ℓ

$$s_\ell(\mathbf{P}_t) = -\frac{\mathbf{Cov}_{\mathbf{P}_t}\left[\left(Z_k(g_{\mathcal{I}_k}) - \mathbf{E}_{\mathbf{P}_t}\left[Z_k(g_{\mathcal{I}_k})\right]\right)^2, g_\ell\right] + 2\mathbf{Cov}_{\mathbf{P}_t}\left[Z_k(g_{\mathcal{I}_k}), g_\ell\right]\Delta_t}{4\omega_e^2 P_t^\ell(1-P_t^\ell)}.$$

where $\Delta_t := \mathbf{E}_{\mathbf{P}_t}\left[Z(g) \ - \eta\right]$. In other words, the selection coefficient at locus $\ell \in \mathcal{I}_k$ only depends on $(P_t^{\ell'})_{\ell'\in\mathcal{I}_k}$ and on the macroscopic observable Δ_t, and can therefore be rewritten $s_\ell(\mathbf{P}_t) = \hat{s}_{\ell,\Delta_t}((\mathbf{P}_t^\ell)_{\ell\in\mathcal{I}_k})$.

Because we assume that the number of clumps is large ($n \gg 1$), we may use a mean-field approximation to replace Δ_t with $\Delta^* := \mathbb{E}\left[\Delta_t\right]$ (as in Section 2.1). It follows that the dynamics within clump $\mathcal{I}_k$ can be written as an autonomous system, with an explicit stationary distribution which depends on Δ^* (see SI J.2 for details). This lets us finally obtain a fixed-point equation for Δ^*, using that $\Delta^* = \mathbb{E}\left[\mathbf{E}_{\mathbf{P}_t}\left[Z(g)\right]\right] - \eta$ and using the stationary distribution of $\mathbf{P}_t$ as a function of Δ^*.

For instance, in the model with clumps of size 2 (see (21)), the fixed point equation is (see in SI equation (71))

$$\Delta^* = 4L\mathbb{E}\left[\tilde{\alpha}\int\int p_1 p_2 \Pi_{W_{\tilde{\alpha},\Delta^*},\tilde{\Theta}}(p_1,p_2)\mathrm{d}p_1\mathrm{d}p_2\right] \ - \ \eta$$

where $(\tilde{\alpha}, \tilde{\Theta}) = (\alpha_{2k_U-1,2k_U}, (\theta_{2k_U-1}, \theta_{2k_U}))$ for k_U some uniform random variable on $[n]$, $\Pi_{W,\Theta}$ is the stationary solution to a 2-locus Wright-Fisher diffusion with fitness function W and mutation rates Θ (see in SI Equation (70) for details) and $W_{\alpha,\delta}$ is defined with

$$\forall \alpha, \delta \in \mathbb{R}, g_1, g_2 \in \{0,1,2\}, \quad W_{\alpha,\delta}(g_1,g_2) := -\frac{\alpha^2}{2\omega^2} g_1 g_2 \left(g_1 g_2 - 4(g_1-1)(g_2-1)\right) - \frac{\delta g_1 g_2}{\omega^2}.$$

A similar fixed point equation is derived for arbitrary clumps.

Diffuse epistasis across all loci. The other case we briefly discuss is based on Sherrington-Kirkpatrick-like models [31], in which every locus interacts loosely with every other, the typical example (discussed in SI J.3.1) being

$$Z(g) = \sum_{\ell\neq\ell'} \alpha_{\ell,\ell'}\, g_\ell\, g_{\ell'}.$$

For a given $\ell \in [L]$, we can rewrite this as

$$Z(g) = g_\ell \sum_{\ell'\neq\ell} \alpha_{\ell,\ell'}\, g_{\ell'} \quad + \sum_{\substack{\ell',\ell''\in[L]\setminus\{\ell\} \\ \ell'\neq\ell''}} \alpha_{\ell,\ell'}\, g_\ell\, g_{\ell'}.$$

Due to the large value of L, it makes sense to expect $\sum_{\ell'\neq\ell}\alpha_{\ell,\ell'}g_{\ell'}$ to behave as a deterministic macroscopic observable $\hat{\alpha}^* = L\mathbb{E}\left[\alpha 2P_t\right]$ where α, P_t correspond to $(\alpha_{\ell,\ell_U}, P_t^{\ell_U})$ for ℓ_U sampled uniformly at random on $[L]\setminus\{\ell\}$. In this case, from the gene's-eye view of ℓ, Z can be approximated as an additive trait with additive effect $\hat{\alpha}^*$. In turn, this means that we may express the stationary distribution of P_t^ℓ as $\Pi_{\Delta^*,\hat{\alpha}^*,\theta_\ell}$ using the definition (12). This finally lets us find a fixed-point system for $(\Delta^*, \hat{\alpha}^*)$

$$\begin{aligned}\Delta^* =& \frac{(\hat{\alpha}^*)^2}{2\mathbb{E}\left[\alpha\right]} - \eta \\ \hat{\alpha}^* =& 2\mathbb{E}\left[\alpha\right]\mathbb{E}\left[L\int p\Pi_{\hat{\alpha}^*,\Delta^*,\theta}(p)\mathrm{d}p\right]\end{aligned}$$

(see in SI system (75-76)). In Section J.3.2, we give directions on how to generalize this approach to any trait function.

3 Discussion

3.1 Describing a polygenic system from the gene's-eye view

Our paper offers a comprehensive framework to describe a polygenic system from the gene's-eye view. This lets us describe the distribution of a gene at stationarity as well as macroscopic observables. These ideas were already present in the case of constant (α, θ) across loci, small mutation rates ($|\bar{\theta}| \ll 1$) and moderate/strong selection ($\omega_e^{-2} \gg L$) in [8] (see SI E.4). We show in SI G that the gene's-eye view and the traditional trait's-eye view from [3] lead to the same equations.

The equation for the evolution of an allele depends on microscopic, intrinsic properties of the allele (the mutation probability μ, the genetic effect α) and also on an external variable, the bias-correcting coefficient s^*, which is determined by the value of a macroscopic observable (the mean trait value), which is itself a solution to a fixed point equation. This sort of approach from the gene's-eye view has seen increasing use in the past few years: for instance, [15] and [16] have used a similar approach to describe the stationary distribution of a population under directional selection and immigration. Another example is [32] (not yet peer-reviewed), in which the macroscopic observable is the genetic variance σ^2 under fluctuating stabilizing selection, which determines the allelic frequencies, which in turn determine the genetic variance, thus yielding a fixed-point equation (their equation 16).

Taking a different approach based on the stationary solution to the multi-locus Wright-Fisher diffusion (4), [33] has recently found an explicit approximate expression for the solution Δ^* to the fixed-point equation (13) when the mutational bias is small ($|z_M - \eta| \ll 1$) and $\theta^+ = \theta^-$: this approximation is equivalent to computing the variance $\mathbb{V}\mathbf{ar}^*[P_t|\alpha, \theta]$ by approximating the distribution of P_t with the neutral distribution $\Pi_{0,0,\theta}$. Using this idea coupled with the central limit theorem leads to the first-order perturbation to $\Pi_{0,0,\theta}$ due to selection.

3.2 On the importance of trait mutational bias

Trait mutational bias in a quantitative trait describes the situation in which the additive effect of mutations on the trait has a nonzero mean. This should not be confused with mutational bias in fitness. Mutational bias in fitness refers to the fact that if the population is close to the optimum, new mutations will tend to be deleterious (see for instance [34]). Even if there was no trait mutational bias (say $z_H = \eta = z_M$), due to Robertson's underdominant term in the selection coefficient (10), we would still have that, on average, a mutation with large effect α is deleterious. The impact of trait mutational bias on a polygenic trait has been described in [35] numerically, and [8] has provided a theoretical treatment which includes approximate solutions assuming low mutation rates ($|\bar{\theta}| \ll 1$) and constant $(\theta_\ell, \alpha_\ell)$ across loci. [8] suggests two methods to approximate Δ^*, one (Eq. A3.b) which coincides with our own (see our SI E.4) and one which requires approximating $\mathbb{E}\left[P_t(1-P_t)|(\alpha, \theta)\right]$ with its neutral expectation (Eq. S3.3b).

Apart from these two exceptions, trait mutational bias is typically neglected in theoretical studies of polygenic traits under stabilizing selection. For instance, [18, 36, 37, 38, 39, 40, 10] all study a quantitative trait under stabilizing selection, and always assume that at equilibrium the mutation rate does not push the population in any specific direction away from the optimum ([10] consider anisotropic pleiotropic mutations as an extension of their base model, but still assume that the mean effect of any new mutation is zero). The models for quantitative traits from statistical physics such as [21, 40, 41] also typically assume symmetric mutation rates. This seems to be motivated not just by mathematical tractability, but by the idea that when the mutation rate is very low, the directional effect of mutations can be neglected. It would be unreasonable to argue that mutational bias is typically small, as this would imply that mutation favors the same trait values as selection, which would be a strange coincidence, in particular if this were satisfied across distinct environments. In fact, experimental data documenting mutational bias can be obtained from the evolution of morphological traits in mutation accumulation (MA) experiments. Morphological quantitative traits are typically assumed to be under stabilizing selection [13]. In MA lines, such traits generally evolve in a deterministic direction. For instance, table 2 of [42] reviews the case of drosophila, in which the proportion of trait-increasing mutations for sternopleural/abdominal bristle number and wing length range between 0.4 and 0.07. More recently, [43] has reported significant mutational bias in locomotion traits in *Caenorhabditis elegans*. Mutational bias could also account for cases in which directional selection and stabilizing selection are simultaneously observed on phenotypic data [44]. Population genetic analyses also suggest significant mutational bias for codon usage [45, 46, 47, 48]. Finally, [30] has established an empirical test of mutational bias for GWAS in humans, and successfully detected a bias for traits such as height, BMI, and educational attainment.

In the presence of mutational bias, which we expect to be pervasive, weak selection cannot be distinguished from directional selection [35], because the population is too far from the optimum ($\frac{\Delta^*}{\sigma} \gg 1$). Such an

equilibrium population would be in selection/mutation balance for a quantitative trait. While stable quantitative traits under directional selection are sometimes found in natural populations, mutational bias is however rarely mentioned as a putative explanation, presumably due to the relative strength of observed selection with respect to mutation (it is for instance never mentioned on reviews such as [49]). The possibility of mutations and genetic drift preventing a population from reaching the optimal trait value is discussed in the literature as the drift-barrier hypothesis. The typical quantitative trait for which this is deemed plausible is the mutation rate [50]. Our framework could be adapted to model this, following [51].

3.3 Example of a practical application: moderate selection and human height

One convenient aspect of our classification is that the qualitative behavior of the stationary system relies on parameters $\omega_e, |\bar{\theta}|, \tau$, which can be empirically measured.

If we take the example of the human population, the effective population size is of order $N_e \sim 20,000$ [52]. For a trait like height, [53] estimate the mutational target to be of order 10^7 ([54] obtained a saturated GWAS map for height from $12,111$ loci, so a certain lower bound for L is 10,000). The average effect of a locus is taken to be of order $\bar{\alpha} = 10^{-3}m$, in line with previous estimates (the effect sizes vary between $.03cm$ and $1cm$ according to [55]). In [44], the UK biobank was used to estimate the parameter $\frac{\omega^{-2}}{2}$, which was found to be of order 1.5×10^{-2} in units of σ^{-2}, with σ of order $0.1m$. Finally, we take the mean mutation rate $|\bar{\theta}|$ per locus per generation to be $10^{-8}N_e$, where N_e is the effective population size [56]. If we neglect environmental variation for simplicity, we may compute

$$\begin{aligned}\omega_e^{-2} =& N_e(L\bar{\alpha})^2\omega^{-2}\\ =& 20,000 \times (10^7 \times 10^{-3})^2 \times (2 \times 1.5 \times 10^{-2}/0.1^2)\\ =& 6 \times 10^{12}\end{aligned}$$

This is compatible with moderate selection. Of course, applying this crude computation to human data assumes equilibrium and ignores population structure, linkage, the fluctuations of selection, the environmental contribution to height, pleiotropy and the fact that the additive effects α have a heavy-tailed distribution.

3.4 Applicability to GWAS data

Our theoretical model yields a stationary distribution for the joint distribution of allelic frequency and additive effect (P_t, α). This distribution is precisely what GWAS are measuring, and which the works of [30, 10, 53, 57] are interpreting from an evolutionary standpoint. Our contribution to their framework is the bias-correcting selection coefficient s^*. From (10), we expect the distribution of P_t conditional on α to be proportional to $\Pi[\xi_\alpha, \mu, N_e]$ where

$$\xi_\alpha(x) := A\alpha + B\alpha^2\left(\frac{1}{2} - x\right) \tag{22}$$

for some constants A and B which could be estimated empirically by maximum likelihood.

Empirically, [30] validates our core idea by detecting asymmetries in the distribution of (P_t, α). More precisely, their model is concerned with (Y_t, β) where Y_t is the frequency of the derived allele (as opposed to the ancestral allele) and β is its additive effect. Compared to our model, β can be negative, but $|\beta| = \alpha$. When $\beta < 0$, then $Y_t = 1 - P_t$, otherwise $Y_t = P_t$. They consistently reject the null hypothesis of unbiased mutations (same average value of β conditional on $Y_t = p$ or on $Y_t = 1 - p$) in 6 out of 9 GWAS of human traits. One would expect this to mean our model (22) would fare well in practice.

It turns out that model (22) has been tested in the case of quantitative disease GWAS in [57] (not yet peer-reviewed). In this preprint, the authors do not really justify model (22), but test it against the "traditional" model of [10] which corresponds to setting $A = 0$. Model (22) fares rather poorly, and for only 2 of 27 tested traits this model outperforms the $A = 0$ model in terms of restricted maximum likelihood (versus 21 traits that support the $A = 0$ model).

If it was confirmed that model (22) generally underperforms with respect to the model with $A = 0$, this could be interpreted in at least two different directions. First, we assume that genetic diversity is maintained by mutations. It can also be maintained by spatial structure [58]: if a very large population is subdivided into many small demes, then genetic diversity within each deme can be maintained by migration. In our model, if we assume a Levene migration model [59] we could add this effect by replacing the mutation probabilities (μ_ℓ^+, μ_ℓ^-) with $(\mu_\ell^+ + m\mathbb{E}^*[P_t^\ell], \mu_\ell^- + m\mathbb{E}^*[1 - P_t^\ell])$ for some migration parameter $m \gg \mu_\ell^{(\pm)}$. In such a setting, mutational bias should become negligible. Second, it could be seen as a signature of selection on the distribution (α, θ) itself. For instance, the mutation rate is highly variable across the genome [60], shaped by evolvable factors such as chromatin structure [61]. The distribution of α can be shaped through second-order selection favoring robust gene regulatory networks. A quantitative trait under stabilizing selection can evolve towards robustness

in a stable environment, because selection favors genes which reduces the genetic variance in the trait [62]. The process by which stabilizing selection shapes the distribution of α is called genetic canalization [63] (see also [64]).

A framework for a systematic evolutionary analysis GWAS analysis was recently proposed [53] (currently in revision). The model considers that each locus has a highly pleiotropic effect, which determines its selection coefficient in Robertson's underdominant term (the pleiotropic equivalent of the second term of (10), see SI I.2). It is assumed that conditional on this selection coefficient, the additive effect of the alleles at this locus on a given focal trait has a symmetric distribution. This assumption posits that selection and mutation act identically on trait-increasing and trait-decreasing alleles, implying that $z_M = z_H = \eta$ and that the distributions of (α, P_t) and $(\alpha, 1 - P_t)$ are identical. We would argue for a generalization of this approach including a possible bias in mutation and in allelic effects, for example through the addition of a hyperparameter.

3.5 Future directions

On top of all extensions discussed in Section 2.6, it will also be straightforward to add fluctuating environments as in [32] (provided the fluctuations are of the same order as σ or smaller), genes influencing plasticity (G×E interactions, which are studied from the trait's-eye view in [65]), spatial structure as in [58], and more sophisticated selection (for instance if $W(z)$ is skewed).

One could also hope to derive first-order corrections for linkage disequilibrium, following [20, 66]. Another mechanism to consider would be insertions-deletions, which would let the number of loci L evolve [67, 68, 69, 70, 71]. On another level, there is growing evidence for the rôle of introgression and admixture [72] as an alternative to mutation for generating genetic diversity. Theoretical models such as [73] offer powerful descriptions of the long-term effect of an admixture event on a polygenic trait. These models allow for the representation of migration as an input in genetic variability under assumptions of strong recombination and weak selection [74]. Finally, one should account for the evolution of the effective population size N_e, the fluctuations of which can be influenced directly by the effect of selection [5, 58, 75].

Future work will also describe the dynamics of the system out-of-equilibrium, extending the work of [12, 76, 77]. This will allow us to tackle questions such as the response of a population to a decrease in population size or a change in optimum, and the corresponding genetic load [78]. Overall, our framework lets us account for most biologically relevant features of polygenic systems (with the exception of LD) and provides efficient theoretical predictions as well as explicit criteria on the parameters for our predictions to hold.

4 Data availability

Our programs were run on Python 3.10.12. All of our code is available at https://github.com/PhCourau/gene_s_eye_view_of_quantitative_genetics_applications.

Supporting Information

All of our supporting information is organised as a single file, which is structured as follows.

A **Model and notation**

B **Outline of the derivation**

C **The diffusion approximation**

D **The polygenic equation from the gene's eye-view**

E **Observables and scalings**

F **Breakdown of the polygenic limit**

G **Equivalence with the trait's eye-view**

H **Derivation of the criterion for selection to increase genetic variance.**

I **Some extensions**

J **Epistasis**

5 Acknowledgments

We thank Paul Jenkins and Henrique Teotónio for helpful comments and references. We thank Archana Devi, Mathilde André and François Ged for fruitful discussions. Thanks also to Céline Teplitsky who first introduced P. Courau to quantitative genetics.

6 Funding

Philibert Courau and Amaury Lambert were supported by the Center for Interdisciplinary Research in Biology (CIRB, Collège de France) and the Institute of Biology of ENS (IBENS, École Normale Supérieure, Université PSL). Philibert Courau was also supported by the Contrat Doctoral Spécifique Normalien. Emmanuel Schertzer gratefully acknowledges support from the FWF project PAT3816823.

7 Conflicts of interest

The authors declare no conflict of interest.

References

[1] N H Barton, A M Etheridge, and A Véber. The infinitesimal model: Definition, derivation, and implications. *Theoretical Population Biology*, 118:50–73, 12 2017.

[2] R A Fisher. The correlation between relatives on the supposition of mendelian inheritance. *Transaction Royal Society of Edinburgh*, 52:399–433, 1918.

[3] Sewall Wright. The analysis of variance and the correlations between relatives with respect to deviations from an optimum. *Journal of Genetics*, 30(2):243–256, Mar 1935.

[4] Russell Lande. Natural selection and random genetic drift in phenotypic evolution. *Evolution*, 30(2):314–334, 06 1976.

[5] N.H. Barton and A.M. Etheridge. Establishment in a new habitat by polygenic adaptation. *Theoretical Population Biology*, 122:110–127, 2018. Paul Joyce.

[6] Florian Patout. The cauchy problem for the infinitesimal model in the regime of small variance. *Analysis & PDE*, 16(6):1289–1350, 2023.

[7] Vincent Calvez, Jimmy Garnier, and Florian Patout. Asymptotic analysis of a quantitative genetics model with nonlinear integral operator. *Journal de l'École polytechnique-Mathématiques*, 6:537–579, 2019.

[8] Brian Charlesworth. Stabilizing selection, purifying selection, and mutational bias in finite populations. *Genetics*, 194(4):955–971, 08 2013.

[9] S Wright. The distribution of gene frequencies in populations. *Proceedings of the National Academy of Sciences*, 23(6):307–320, 1937.

[10] Y B. Simons, K Bullaughey, R R. Hudson, and G Sella. A population genetic interpretation of gwas findings for human quantitative traits. *PLoS Biology*, 16:e2002985, 4 2017.

[11] Laura Katharine Hayward and Guy Sella. Polygenic adaptation after a sudden change in environment. *eLife*, 11:e66697, sep 2022.

[12] Laura Katharine Hayward and Guy Sella. Polygenic adaptation after a sudden change in environment. *Elife*, 11:e66697, 2022.

[13] Guy Sella and Nicholas H Barton. Thinking about the evolution of complex traits in the era of genome-wide association studies. *Annual review of genomics and human genetics*, 20(1):461–493, 2019.

[14] J Arvid Ågren. *The gene's-eye view of evolution.* Oxford University Press, 2021.

[15] Arthur Zwaenepoel, Himani Sachdeva, and Christelle Fraïsse. The genetic architecture of polygenic local adaptation and its role in shaping barriers to gene flow. *Genetics*, 228(3):iyae140, 08 2024.

[16] Parvathy Surendranadh and Himani Sachdeva. Effect of assortative mating and sexual selection on polygenic barriers to gene flow. *Evolution*, 79(7):1185–1198, 05 2025.

[17] LP Chaintron and A Diez. Propagation of chaos: A review of models, methods and applications. i. models and methods. *Kinetic and Related Models*, 15(6):895–1015, 2022.

[18] Alan Robertson. The effect of selection against extreme deviants based on deviation or on homozygosis. *Journal of Genetics*, 54:236–248, 1956.

[19] Motoo Kimura. Diffusion models in population genetics. *Journal of Applied Probability*, 1(2):177–232, 1964.

[20] M. G. Bulmer. Linkage disequilibrium and genetic variability. *Genetical Research*, 23(3):281–289, 1974.

[21] Richard A. Neher and Boris I. Shraiman. Statistical genetics and evolution of quantitative traits. *Rev. Mod. Phys.*, 83:1283–1300, Nov 2011.

[22] Wei Hong, Shihu Li, and Wei Liu. Strong convergence rates in averaging principle for slow-fast mckean-vlasov spdes. *Journal of Differential Equations*, 316:94–135, 2022.

[23] M. G. Bulmer. The effect of selection on genetic variability. *The American Naturalist*, 105(943):201–211, 1971.

[24] W. G. Hill and Alan Robertson. The effect of linkage on limits to artificial selection. *Genetical Research*, 8(3):269–294, 1966.

[25] Todd L. Parsons and Peter L. Ralph. Large effects and the infinitesimal model. *Theoretical Population Biology*, 156:117–129, 2024.

[26] Jeffrey P. Spence, Nasa Sinnott-Armstrong, Themistocles L. Assimes, and Jonathan K. Pritchard. A flexible modeling and inference framework for estimating variant effect sizes from gwas summary statistics, 2022.

[27] Richard A Neher, Marija Vucelja, Mark Mezard, and Boris I Shraiman. Emergence of clones in sexual populations. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):P01008, jan 2013.

[28] Ilse Höllinger, Pleuni S Pennings, and Joachim Hermisson. Polygenic adaptation: From sweeps to subtle frequency shifts. *PLoS genetics*, 15(3):e1008035, 2019.

[29] Ilse Höllinger, Benjamin Wölfl, and Joachim Hermisson. A theory of oligogenic adaptation of a quantitative trait. *Genetics*, 225(2):iyad139, 08 2023.

[30] Lawrence H. Uricchio, Hugo C. Kitano, Alexander Gusev, and Noah A. Zaitlen. An evolutionary compass for detecting signals of polygenic selection and mutational bias. *Evolution Letters*, 3(1):69–79, 02 2019.

[31] Scott Kirkpatrick and David Sherrington. Infinite-ranged models of spin-glasses. *Phys. Rev. B*, 17:4384–4403, Jun 1978.

[32] Jason Bertram and Zahra Shafiei. Strong amplification of quantitative genetic variation under a balance between mutation and fluctuating stabilizing selection, 2025.

[33] Archana Devi and Kavita Jain. Polygenic adaptation dynamics in large, finite populations, 2023.

[34] Eloïse Vanhoenacker, Linnéa Sandell, and Denis Roze. Stabilizing selection, mutational bias, and the evolution of sex*. *Evolution*, 72(9):1740–1758, 09 2018.

[35] Xu-Sheng Zhang and William G Hill. The anomalous effects of biased mutation revisited: Meanoptimum deviation greatly, however, and that the apparent directional selection thereby caused can be substantial. *Genetics*, 179, 2008.

[36] BDH Latter. Natural selection for an intermediate optimum. *Australian Journal of Biological Sciences*, 13(1):30–35, 1960.

[37] M. G. Bulmer. The genetic variability of polygenic characters under optimizing selection, mutation and drift. *Genetical Research*, 19(1):17–25, 1972.

[38] Harold P. de Vladar and Nick H. Barton. The statistical mechanics of a polygenic character under stabilizing selection, mutation and drift. *Journal of The Royal Society Interface*, 8(58):720–739, 2011.

[39] Katarína Bod'ová, Enikő Szép, and Nicholas H Barton. Dynamic maximum entropy provides accurate approximation of structured population dynamics. *PLOS Computational Biology*, 17(12):e1009661, 2021.

[40] Armita Nourmohammad, Stephan Schiffels, and Michael Lässig. Evolution of molecular phenotypes under stabilizing selection. *Journal of Statistical Mechanics: Theory and Experiment*, 2013(01):P01012, jan 2013.

[41] V Dichio, H-L Zeng, and E Aurell. Statistical genetics in and out of quasi-linkage equilibrium. *Reports on Progress in Physics*, 86(5):052601, apr 2023.

[42] A. García-Dorado, C. López-Fanjul, and A. Caballero. Properties of spontaneous mutations affecting quantitative traits. *Genetical Research*, 74(3):341–350, 1999.

[43] François Mallard, Luke Noble, Charles F Baer, and Henrique Teotónio. Variation in mutational (co)variances. *G3 Genes/Genomes/Genetics*, 13(2):jkac335, 12 2022.

[44] Jaleal S. Sanjak, Julia Sidorenko, Matthew R. Robinson, Kevin R. Thornton, and Peter M. Visscher. Evidence of directional and stabilizing selection in contemporary humans. *Proceedings of the National Academy of Sciences*, 115(1):151–156, 2018.

[45] Paul M. Sharp, Elizabeth Bailes, Russell J. Grocock, John F. Peden, and R. Elizabeth Sockett. Variation in the strength of selected codon usage bias among bacteria. *Nucleic Acids Research*, 33(4):1141–1153, 01 2005.

[46] Ruth Hershberg and Dmitri A. Petrov. Selection on codon bias. *Annual Review of Genetics*, 42(Volume 42, 2008):287–299, 2008.

[47] Kai Zeng and Brian Charlesworth. Estimating selection intensity on synonymous codon usage in a nonequilibrium population. *Genetics*, 183(2):651–662, 10 2009.

[48] Kai Zeng. A simple multiallele model and its application to identifying preferred–unpreferred codons using polymorphism data. *Molecular Biology and Evolution*, 27(6):1327–1337, 01 2010.

[49] Joel G. Kingsolver and Sarah E. Diamond. Phenotypic selection in natural populations: What limits directional selection? *The American Naturalist*, 177(3):346–357, 2011. PMID: 21460543.

[50] Way Sung, Matthew S Ackerman, Samuel F Miller, Thomas G Doak, and Michael Lynch. Drift-barrier hypothesis and mutation-rate evolution. *Proceedings of the National Academy of Sciences*, 109(45):18488–18492, 2012.

[51] Kevin J. Dawson. Evolutionarily stable mutation rates. *Journal of Theoretical Biology*, 194(1):143–157, 1998.

[52] Romain Fournier, Zoi Tsangalidou, David Reich, and Pier Francesco Palamara. Haplotype-based inference of recent effective population size in modern and ancient dna samples. *Nature Communications*, 14(1):7945, Dec 2023.

[53] Yuval B. Simons, Hakhamanesh Mostafavi, Courtney J. Smith, Jonathan K. Pritchard, and Guy Sella. Simple scaling laws control the genetic architectures of human complex traits, 2022.

[54] Loïc Yengo, Sailaja Vedantam, and Eirini et al. Marouli. A saturated map of common genetic variants associated with human height. *Nature*, 610(7933):704–712, Oct 2022.

[55] Peter Visscher. Quantitative genetics isn't dead yet. *International Conference on Quantitative Genetics 6*, March 2021. https://www.youtube.com/watch?v=tDcTjatiWjg.

[56] Laure Ségurel, Minyoung J. Wyman, and Molly Przeworski. Determinants of mutation rate variation in the human germline. *Annual Review of Genomics and Human Genetics*, 15(Volume 15, 2014):47–70, 2014.

[57] E. Koch, N. J. Connally, N. Baya, M. P. Reeve, M. Daly, B. Neale, E. S. Lander, A. Bloemendal, and S. Sunyaev. Genetic association data are broadly consistent with stabilizing selection shaping human common diseases and traits, 2024.

[58] Jitka Polechová and Nicholas H. Barton. Limits to adaptation along environmental gradients. *Proceedings of the National Academy of Sciences*, 112(20):6401–6406, 2015.

[59] Thomas Nagylaki. Evolution under the multilocus levene model without epistasis. *Theoretical Population Biology*, 76(3):197–213, 2009.

[60] Gustavo V Barroso and Julien Y Dutheil. The landscape of nucleotide diversity in *drosophila* melanogaster is shaped by mutation rate variation. *Peer Community Journal*, 3, 2023.

[61] Kateryna D. Makova and Ross C. Hardison. The effects of chromatin organization on variation in mutation rates in the genome. *Nature Reviews Genetics*, 16(4):213–223, Apr 2015.

[62] Günter P. Wagner, Ginger Booth, and Homayoun Bagheri-Chaichian. A population genetic theory of canalization. *Evolution*, 51(2):329–347, 04 1997.

[63] Thomas Flatt. The evolutionary genetics of canalization. *The Quarterly Review of Biology*, 80(3):287–316, 2005. PMID: 16250465.

[64] Kerry Geiler-Samerotte, Federica M.O. Sartori, and Mark L. Siegal. Decanalizing thinking on genetic canalization. *Seminars in Cell & Developmental Biology*, 88:54–66, 2019. Canalization, a central concept in biology.

[65] Sara Via and Russell Lande. Genotype-environment interaction and the evolution of phenotypic plasticity. *Evolution*, 39(3):505–522, 1985.

[66] Sherif Negm and Carl Veller. The effect of long-range linkage disequilibrium on allele-frequency dynamics under stabilizing selection, 2024.

[67] Dmitri A. Petrov, Elena R. Lozovskaya, and Daniel L. Hartl. High intrinsic rate of dna loss in drosophila. *Nature*, 384(6607):346–349, Nov 1996.

[68] D A Petrov, Y C Chao, E C Stephenson, and D L Hartl. Pseudogene evolution in drosophila suggests a high rate of dna loss. *Molecular Biology and Evolution*, 15(11):1562–1567, 11 1998.

[69] Josep M Comeron and Martin Kreitman. The correlation between intron length and recombination in drosophila: Dynamic equilibrium between mutational and selective forces. *Genetics*, 156(3):1175–1190, 11 2000.

[70] Susan E Ptak and Dmitri A Petrov. How intron splicing affects the deletion and insertion profile in drosophila melanogaster. *Genetics*, 162(3):1233–1244, 11 2002.

[71] Evgeny V. Leushkin, Georgii A. Bazykin, and Alexey S. Kondrashov. Strong mutational bias toward deletions in the drosophila melanogaster genome is compensated by selection. *Genome Biology and Evolution*, 5(3):514–524, 02 2013.

[72] Nathaniel B. Edelman and James Mallet. Prevalence and adaptive impact of introgression. *Annual Review of Genetics*, 55(Volume 55, 2021):265–283, 2021.

[73] Himani Sachdeva and Nicholas H Barton. Introgression of a block of genome under infinitesimal selection. *Genetics*, 209(4):1279–1303, 2018.

[74] Himani Sachdeva. Reproductive isolation via polygenic local adaptation in sub-divided populations: Effect of linkage disequilibria and drift. *PLOS Genetics*, 18(9):1–29, 09 2022.

[75] Enikő Szép, Himani Sachdeva, and Nicholas H. Barton. Polygenic local adaptation in metapopulations: A stochastic eco-evolutionary model. *Evolution*, 75(5):1030–1045, 2021.

[76] Nathan W Anderson, Lloyd Kirk, Joshua G Schraiber, and Aaron P Ragsdale. A path integral approach for allele frequency dynamics under polygenic selection. *Genetics*, 229(1):1–63, 11 2024.

[77] William R Milligan, Laura K Hayward, and Guy Sella. When should adaptation arise from a polygenic response versus few large effect changes? *bioRxiv*, page 2025.05.15.654234, 2025. https://www.biorxiv.org/content/biorxiv/early/2025/05/17/2025.05.15.654234.full.pdf.

[78] Brian Charlesworth. Why we are not dead one hundred times over. *Evolution*, 67(11):3354–3361, 2013.