The impact of background selection in mutation-selection-drift balance models of complex trait evolution

Xinyi Li[1] and Jeremy J Berg[1,2]

Affiliations

[1]Committee on Genetics, Genomics, & Systems Biology, University of Chicago, Chicago, IL, USA

[2]Department of Human Genetics, University of Chicago, Chicago, IL, USA

Abstract

Over the past three decades, theoretical population genetic models and empirical analyses have elaborated on how the elimination of deleterious variants reduces variation at linked neutral sites. This process is known as background selection. However, the impact of background selection has been studied almost exclusively without reference to its effect on phenotypic variation. Here, we use population genetic theory and simulation to investigate the consequences of background selection for linked variation impacting complex traits. We find the impact of background selection on a complex trait depends on the type of selection acting on the trait, the effect size distribution, and the background selection intensity. For instance, background selection impacts a much broader range of effect sizes for traits under directional selection via a liability threshold mechanism (i.e. when selection acts on individuals with phenotypes exceeding a threshold) than when fitness decays exponentially with trait value. We also find that background selection can increase the population prevalence of a liability threshold trait (e.g. a complex binary disease trait). Furthermore, while background selection can only reduce the genetic variance of a trait under directional selection (as might be expected from the classical population genetic theory for linked neutral variants) we find that background selection can increase the genetic variance of a trait under stabilizing selection. The results presented here use a synthesis of population genetic and quantitative genetic modeling, and overall, have relevance for understand the genetic architecture and evolution of complex traits, including complex diseases.

Significance

Physical linkage of genetic variants to nearby deleterious mutations shapes their evolution. The effect of this linkage is similar to a reduction in population size, because chromosomes carrying deleterious alleles are quickly removed. This leads to reduced genetic variation for neutral variants, and for variants under weak directional selection. Here, we analyze the impact of linkage in models of selection on complex, polygenic traits. We find that when selection acts on traits through a threshold, linkage reduces genetic variation for much more strongly selected sites than existing theory predicts. In contrast, if selection stabilizes trait values toward an optimum, then linkage may actually increase genetic variation. Understanding the fully effect of linkage therefore requires understand how selection acts on traits.

references.bib

Introduction

Any mutation that impacts fitness alters the fates of other mutations to which it is linked. If the mutation is sufficiently deleterious, it will quickly be purged from the population, carrying with it linked mutations on the same haplotype. To a first approximation, this "background selection" (BGS) acts like a local reduction in the effective population size, reducing the number of mutations that enter the population every generation, and increasing their rate of drift. This, in turn, leads to a local reduction in the level of neutral genetic variation (Charlesworth et al., 1993; Hudson & Kaplan, 1995; Nordborg et al., 1996).

The strength of the reduction in neutral diversity depends largely on the density of deleterious mutations relative to the local recombination rate, and is often measured via a "B value", which estimates the size of the reduction in neutral genetic diversity in a given genomic region, relative to what would be expected in the absence of BGS (McVicker et al., 2009). A recent study of population genetic variation in humans estimated that BGS has reduced neutral diversity by about 17% on average across the genome (i.e. mean $B = 0.83$; Murphy et al. (2022)). In Drosophila, this reduction is plausibly as high as 80% when the effects of recurrent selective sweeps are also included (i.e. mean $B \approx 0.2$; Elyashiv et al. (2016)). The existence of broad-scale correlations between recombination rate and levels of diversity across a wide range of taxa suggests that such linked selection effects are important broadly across the tree of life (Buffalo, 2021; Coop, 2016; Corbett-Detig et al., 2015).

It is less clear how BGS shapes variation for complex traits. At least two empirical studies have invoked BGS to explain observed patterns of genetic variation for complex traits. In *C.elegans*, Rockman et al. (2010) observed that loci associated with gene expression in gene-poor, high recombination chromosome arms explain significantly more variance in transcript abundance than loci in the gene-rich chromosomal centers. They interpreted this pattern as evidence that BGS has played a significant

role in shaping levels of genetic variance for gene expression, and has acted similarly to how it acts on neutral variation. In humans, Pardiñas et al. (2018) found that the per-SNP contribution to heritability due to common variants associated with schizophrenia is greater in regions with lower B values (i.e. stronger background selection). They argue that this is consistent with the effect of BGS, as causal variants under negative selection would be expected to segregate at slightly higher frequencies when the local effect of genetic drift is greater.

The intuition behind these explanations may be correct, but we do not currently have mathematical modeling results that address the impact of BGS on genetic variation underlying complex traits, or its impact on the distribution of complex phenotypes in the population. Previous work has studied the impact of BGS on individual variants with fixed directional selection coefficients (Barton, 1995; Charlesworth, 1994; Stephan et al., 1999). We refer to models of this type "single-site models". These results can be applied to complex traits only if we assume the sites that contribute to variation in the trait evolve independently of one another (see, e.g. McVEAN and Charlesworth (1999)). However, under the most plausible models for complex traits at mutation-selection-drift balance, the strength or even the direction of selection acting on an individual allele depends on the state of all the other loci that contribute to variation in the trait. As we will show, these dependencies in turn modify the impact of BGS on complex trait loci.

Here, we consider three models of selection acting on complex phenotypes, and study how background selection alters their evolution at mutation-selection-drift balance (Figure 1). We first consider two models of directional selection which lie at opposite extremes along a continuum of possible models for directional selection on a complex phenotype. On one end of the continuum is an exponential fitness function, which causes individual variants to evolve as if under independent additive directional selection (this model therefore maps to existing results from McVEAN and Charlesworth (1999) which assume that individual variants are independently selected). At the other end of this continuum is a model in which selection acts against individuals whose phenotype exceeds a physiological threshold (King, 1966), with this selection pressure balanced by a mutational pressure toward the threshold. Threshold selection represents the most extreme form of synergistic epistasis (Crow & Kimura, 1979). Models with other forms of synergistic epistasis (e.g. Charlesworth (1990) and Kimura and Maruyama (1966)), would be expected to have an intermediate behavior between these two extremes. We return to this point in Discussion.

In contrast to directional selection, it is likely that many complex traits evolve under stabilizing selection around an optimum phenotype (Fisher, 1930; Robertson, 1956; Simons et al., 2018). This is supported by evidence for stasis in the fossil record over long evolutionary timescales (Smith, 1983), by direct measurements in both wild (Weber, 1990) and human populations (Karn & Penrose, 1951; Sanjak et al., 2018), and by an emerging argument in human genetics that the genetic architectures of the most complex human phenotypes are best fitted by a model of stabilizing selection (Koch et al., 2024; O'Connor et al., 2019; Simons et al., 2022). Stabilizing selection results in qualitatively different

evolutionary dynamics than directional selection, and to our knowledge, no prior work has considered the impact of background selection in the context of stabilizing selection.

Our paper is organized as follows. We first describe the assumptions of our genotype-to-phenotype model, and our model of BGS. These assumptions are common to all three models of phenotypic selection that we consider. We then reconsider the impact of background selection on individual sites experiencing directional selection, and show how the effect of BGS on heterozygosity at these sites can be understood simply in terms of the relationship between the expected contribution to heterozygosity and the fixation rate. These results, in turn, enable our analysis of both the exponential model and the threshold model for complex traits. We conclude with the model of stabilizing selection (Figure 1).

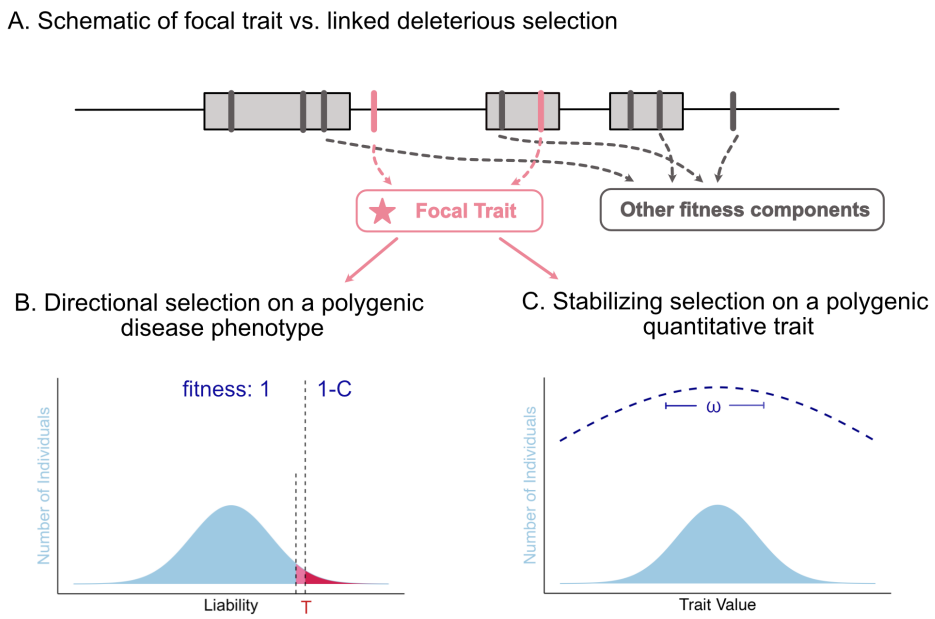

Figure 1: Illustration of focal traits and linked deleterious selection. (A) We study models in which phenotypic variation for a focal trait is caused by a large number of loci that are linked to loci which have deleterious effects via other fitness components. We are primarily interested in models where the focal trait is a disease evolving under the liability threshold model (B) or a quantitative trait under stabilizing selection (C).

Model and Results

Phenotype and mutation model

We assume that an individual's trait (Z_i) arises from the standard additive model of quantitative genetics (Falconer, 1996), with both genetic and environmental components: $Z_i = G_i + E_i$. We assume that the total number of sites that affect the trait is large (i.e. $L \gg 1$), and that at each site there is a trait-increasing and a trait-decreasing allele, each of which mutates into the other type with probability μ per generation. The genetic component of individual i's trait value is thus a weighted sum of contributions from the L loci: $G_i = \sum_{l=1}^{L} a_l g_{il}$, where a_l is the effect on the trait of substituting one additional trait increasing allele at locus l, and g_{il} is the number of trait-increasing alleles carried by individual i at locus l. The minimum theoretically possible genetic value is therefore zero, while the maximum is $2\sum_{l=1}^{L} a_l = 2L\bar{a}$, where $\bar{a}$ is the mean effect size. In the absence of selection, the mean genetic value at mutation-drift equilibrium would be expected to vary around a long-term average of $L\bar{a}$, reflecting the symmetry of both mutation and drift. We assume that the environmental contribution, E_i, is normally distributed with mean zero and variance V_E.

Background selection model

We assume the L focal sites contributing to the trait are embedded in a recombining chromosome of length $M+L$ base pairs (with $M \gg L$), with recombination rate r per base pair, in which background negatively selected mutations with selection coefficient s_{bgs} arise at rate ν per base pair among the M sites that do not contribute to the trait. We assume that selection against these mutations is strong (i.e. $2Ns_{bgs} \gg 1$), so that they are all quickly purged, with no chance of fixing. The impact of this background selection can be approximated as a reduction in the effective population size. Here, we present the most relevant of these results for our work in intuitive and heuristic form.

The expected frequency of the deleterious allele at each background site is $\frac{\nu}{s_{bgs}}$. It follows that we expect there to be a total of approximately $2M\frac{\nu}{s_{bgs}}$ deleterious alleles in the population at any given time. When a new mutation arises on a background carrying one of these deleterious alleles, the probability that it will be purged by linked selection before recombining onto a different haplotype depends on the distance between the deleterious allele and the new mutation. To a first approximation, mutations that arise within $\frac{s_{bgs}}{r}$ base pairs of a deleterious allele will be purged too quickly to recombine, while those at greater distances will recombine before being purged. As a result, only a fraction $\frac{s_{bgs}}{Mr}$ of the $2M\frac{\nu}{s_{bgs}}$ deleterious alleles segregating in the population will be located close enough to the new mutation to impact it before it can recombine. As a result, the probability that a new mutation escapes immediate loss either by arising on a background free of deleterious mutations

or by quickly recombining onto one, is therefore approximately

$$B \approx e^{-\frac{2M\nu}{s_{bgs}} \cdot \frac{s_{bgs}}{Mr}} = e^{-\frac{2\nu}{r}} \tag{1}$$

(Hudson & Kaplan, 1995; Nordborg et al., 1996).

Then, to a first approximation, the effect of background selection is to remove a fraction $1 - B$ of newly arising mutations. Thus, because a fraction $1 - B$ of lineages in the population are always destined for immediate loss, mutations that survive the initial cull evolve as if the population size is reduced by a factor B. For simplicity, going forward we represent the effect of background selection in terms of the reduction in effective size, B.

The effect of background selection on phenotypes under directional selection

As we outlined in the Introduction, we will consider two different models which result in directional selection at the phenotype level. In both cases, we find the expected change for an individual allele currently found at frequency q is

$$\mathbb{E}[\Delta q | s, q] = sq(1 - q), \tag{2}$$

where the selection coefficient s does not change with frequency (Berg et al., 2025; Milkman, 1978). In other words, each individual allele's expected evolution is well described by a model of additive directional selection at a single site. We can therefore rely on standard diffusion results (Kimura, 1964), in which the strength of selection relative to genetic drift is parameterized by the population scaled selection coefficient, $\gamma = 2Ns$. We largely assume that $|s| \ll |s_{bgs}|$, though we comment briefly on violations of this assumption. Finally, we also assume that $2N\mu \ll 1$, so that all segregating variants have a single mutational origin (Kimura, 1969).

Single-site additive directional selection model

As background, we begin by considering the impact of background selection on contributions to heterozygosity at a single site without an explicit complex trait model. The results we discuss here are partly our own framing of existing results due to Charlesworth (1994) and Stephan et al. (1999) and partly an extension of existing results by Cvijović et al. (2018). These results form a helpful preface for the complex trait models we study.

Impact of BGS the on heterozygosity at sites under directional selection

We imagine tracking a single allele under directional selection with selection coefficient s, starting with its introduction at frequency $q = 1/2N$, and recording its total contribution to heterozygosity during its transit through the population. Under our assumption that the population scaled mutation rate is low (i.e. $2N\mu \ll 1$), so that most sites are fixed at any one time, this quantity is proportional to the

expected per site heterozygosity. Summing the per generation change in allele frequency (i.e. Equation (2)) over an allele's transit through the population and taking the expectation, the expected total change in frequency is $\mathbb{E}[\sum_t \mathbb{E}[\Delta q_t|s]|N] = \pi(s,N) - 1/2N$, where $\pi(s,\gamma) \approx 2s/(1-e^{-4Ns})$ is the standard "small s" diffusion approximation for the fixation probability (Kimura, 1962), and $\frac{1}{2N}$ is the fixation probability of a neutral allele. The expected total contribution to heterozygosity for an allele with scaled coefficient $\gamma = 2Ns$, $h(\gamma) = \mathbb{E}[\sum_t q_t(1-q_t)|s,N]$, is therefore closely linked to the fixation processes, as

$$h(\gamma) = \frac{\rho(\gamma) - 1}{\gamma} \tag{3}$$

(Kimura, 1969) where

$$\rho(\gamma) = 2N\pi(s,\gamma) \approx \frac{2\gamma}{1-e^{-2\gamma}} \tag{4}$$

is the total rate of fixation per $2N$ new mutations. We refer to the quantity $\rho(\gamma) - 1$ as the "selective fixation rate": it gives the rate at which derived alleles are driven to (or prevented from) fixation by selection (recall that we expect one fixation per $2N$ mutations due to genetic drift; i.e. $\lim_{\gamma\to 0}\rho(\gamma) = 1$).

When selection is negative, this rate is negative, reflecting the fact that negative selection prevents some or all of these neutral fixations from occurring, thereby balancing the mutational input and ensuring that the net rate of change is zero. Equation (3) conveys the fundamental link between the long-term pace of evolutionary change due to selection (the selective fixation rate in the numerator), the variance created within a population by the sites responsible for that change (the heterozygosity on the left hand of the equation), and the strength of selection (the scaled selection coefficient, the denominator). This relationship will play an important role in our understanding of the phenotypic models below.

To understand the effect of background selection in the single-site model, we apply the reduced effective size approximation to Equation (3). We assume that only a fraction B of new mutations survive long enough to contribute anything to heterozygosity, while those that do evolve with an "effective" scaled coefficient $\gamma_B = \gamma B$, so that selection is less effective due to the increased rate of drift. The reduction in heterozygosity at a linked selected site (B_γ) due to BGS is therefore

$$B_\gamma = B\frac{h(\gamma B)}{h(\gamma)} \tag{5a}$$

$$= \frac{\rho(\gamma B) - 1}{\rho(\gamma) - 1} \tag{5b}$$

(Figure 2). Equation (5) provides two different perspectives on how background selection impacts heterozygosity at directly selected sites. Equation (5a) captures the effects of the two channels through which the population size impacts the heterozygosity (Figure 2A). For sites under weak negative selection (i.e. $\gamma \sim -1$), Equation (5a) says that the heterozygosity reduction due to purging of

new mutations on deleterious backgrounds (i.e. $B < 1$) is partially countered by the fact that each surviving mutation contributes a bit more on average (i.e. $1 < {}^{h(\gamma B)}/_{h(\gamma)} < \frac{1}{B}$), so that the net effect is not as large as in the neutral case. When selection is strong and negative (i.e. $\gamma \ll -1$), the two effects cancel (i.e. ${}^{h(\gamma B)}/_{h(\gamma)} = {}^{1}/_{B}$), so there is no net effect on heterozygosity. For sites under weak positive selection (i.e. $\gamma \sim 1$), each mutation surviving the initial purge contributes slightly less on average (i.e. ${}^{h(\gamma B)}/_{h(\gamma)} < 1$), so the net effect of background selection is to reduce heterozygosity by even more than in the neutral case. When selection is strong and positive (i.e. $\gamma \gg 1$), the contribution to heterozygosity conditional on surviving the purge becomes independent of population size effects (i.e. ${}^{h(\gamma B)}/_{h(\gamma)} = 1$), so the heterozygosity is affected only by the purging of new mutations.

In contrast, Equation (5b) captures the effect of background selection in terms of its ultimate impact on the long-term rate of evolutionary change via the selective fixation rates (Figure 2B). The effect on heterozygosity then simply follows from the relationship between the two quantities. From this perspective, the background selection effect on heterozygosity, B_γ, weakens as direct selection becomes more strongly negative because the selective fixation rate becomes less and less sensitive to differences in (effective) population size as selection gets stronger, plateauing to $\rho(\gamma) - 1 = -1$ in the limit as all fixations are prevented and the BGS effect on heterozygosity vanishes. In contrast, as selection becomes more positive, the effect initially strengthens, as the selective fixation rate initially becomes more sensitive to differences in (effective) population size as selection becomes more positive, before ultimately plateauing again as population size effects become unimportant to the fixation probability (among a reduced number of new mutations) in the strong selection limit. Figure 2C illustrates the diversity reduction predicted by Equation (5) across a range of γ, in comparison to simulations.

The impact of BGS on the frequency spectrum

We can gain some further insight by considering the effect on the frequency spectrum of segregating alleles. In the low mutation rate limit that we employ, the frequency spectrum in the absence of background selection is proportional to the expected sojourn time of an allele at a given frequency, i.e. the number of generations it is expected to spend in the $[q, q + dq]$ frequency interval. In the absence of background selection, this is given by

$$\xi(q|r) = \frac{2}{1 - e^{-2\gamma}} \frac{1 - e^{-2\gamma(1-q)}}{q(1-q)} \tag{6}$$

(Bustamante et al., 2001; Wright, 1938). In the insets of Figure 2C, we illustrate how background selection impacts the frequency spectrum using the reduced effective size approximation (i.e. $\xi_B(q|\gamma) = B\xi(q|\gamma B)$) by plotting the expectation for the frequency spectrum under direct selection (with and without background selection) relative to the neutral expectation (that is, we plot ${}^{\xi_B(q|\gamma)}/_{\xi(q|0)}$ and ${}^{\xi(q|\gamma)}/_{\xi(q|0)}$, where $\xi(q|0) = {}^{2}/_{q}$ is the neutral expectation), and by plotting the expectation for the

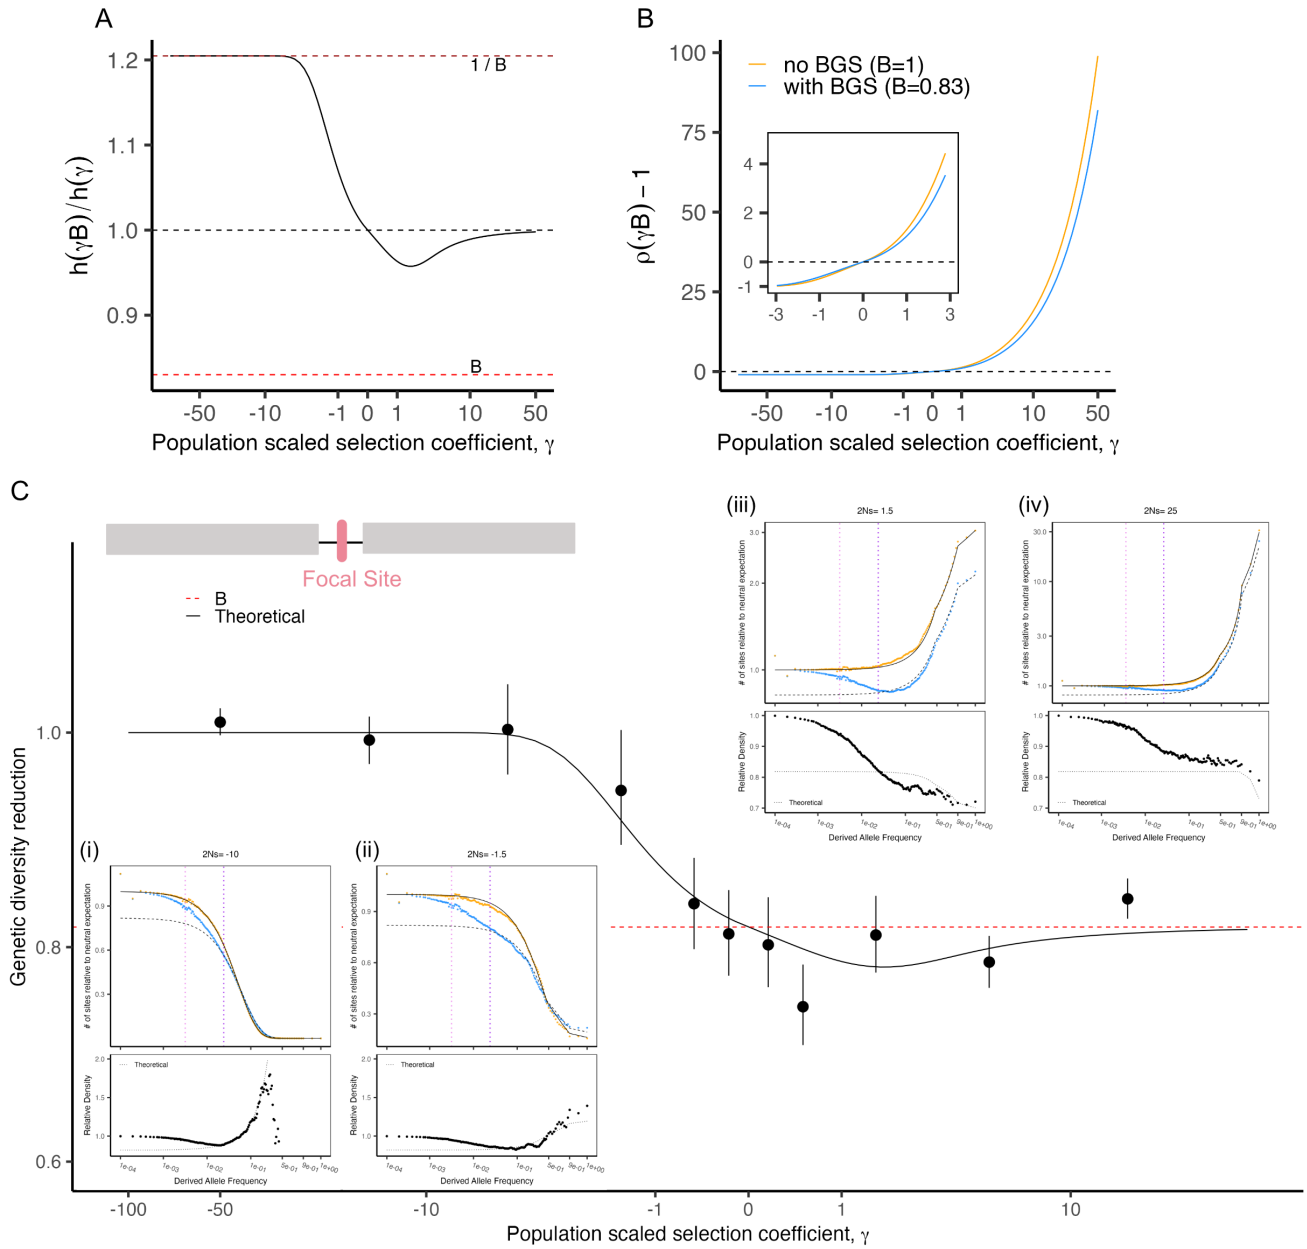

frequency spectrum with background selection relative to that without (i.e. $\xi_B(q|\gamma)/\xi(q|\gamma)$). We compare these predictions to the averages obtained from our Wright-Fisher simulations of recombining chromosomes.

These comparisons help to further illustrate the fundamental results expressed by Equation (5), showing how, for example, alleles under weak negative selection spend slightly more time on average at higher frequencies (and fix at a correspondingly higher rate) when background selection is acting than when it is not (Figure 2C(ii)). The other relationships expressed by Equation (5) across other selection regimes, which we elaborated in the preceding paragraphs, can similarly be visualized in the

Figure 2: Genetic diversity reduction for single-site additive directional selection model under background selection. (A) The predicted change in the contribution to heterozygosity conditional on not being removed via linkage to a BGS site (i.e. the ratio $h(\gamma B)/h(\gamma)$ from equation (5a)). (B) The selective fixation rate, $\rho(\gamma) - 1$, with and without BGS, over the range [-50,50] in the main plot and [-3,3] in the inset. (C) The solid black line shows the diversity reduction, B_γ, predicted by Equation (5), while the horizontal red dashed line shows the diversity reduction expected under neutrality, B. In simulations we used the chromosome structure shown on the top left (gray bars represent regions with BGS sites), varied the scaled selection coefficients of the focal sites and measured the average genetic diversity reduction over many replicates. In the insets (i)- (iv), the upper panels show the site frequency spectrum, scaled relative to its neutral expectation, for $\gamma = -10, -1.5, 1.5, 25$. The color indicates whether background selection is (blue) or is not (orange) operating. The solid line represents the predicted deleterious SFS with effective population size N; while the dashed line represents the predicted deleterious SFS with effective population size NB. The left violet dashed vertical line plots $1/2N\sigma$ where σ is the average genome-wide standard deviation of the background fitness from simulations. The right purple dashed vertical line plots $1/2NBs_{bgs}$. Lower panels plot the density of alleles at each frequency under background selection compared to no background selection.

corresponding Figure 2C insets.

The impact of BGS on low frequency variants

While the reduced effective size approximation performs well at higher frequencies, it is inaccurate at low frequencies. This is because the impact of background selection is not in fact instantaneous but takes several generations to accrue. In a model where the focal sites are neutral and there is no recombination, Cvijović et al. (2018) showed that the effect of this process on the frequency spectrum of neutral alleles can be modeled in terms of two important thresholds. In their model, the frequency spectrum resembles that in the absence of background selection at frequencies less than $1/2N\sigma$ (where σ is the standard deviation of fitness across backgrounds), because alleles found at these frequencies are too young to have experienced much background selection. In contrast, at frequencies greater than $1/2NBs_{bgs}$ the reduced effective size approximation applies, because only alleles on backgrounds free of deleterious mutations can segregate above this threshold. Between these thresholds, there is a transition between the two regimes.

In our simulations, we find that these thresholds accurately describe the transitions in the shape of the site frequency spectrum for directly selected alleles experiencing background selection in the presence of recombination, provided that we use the standard deviation in fitness across the entire recombining chromosome (i.e. $\sigma = \sqrt{2M\nu s_{bgs}}$) to determine the lower threshold, and the local reduction in effective size (i.e. Equation (1)) to determine the upper threshold. The logic behind these choices is that alleles at very low frequencies are very young, and so have not yet had an opportunity to recombine very much. As a result, they are still associated with the original haplotype they arose on over a very long chromosomal scale. The initial onset of background selection effects

in the frequency spectrum therefore depends on the effects of deleterious mutations arising over this very long chromosomal scale. In contrast, alleles that manage to reach the upper threshold $1/2NBs_{bgs}$, have had more time to recombine, and are thus associated with their initial haplotype over a shorter length scale of approximately s_{bgs}/r bases. The alleles whose frequencies exceed this second threshold are those that arose on haplotypes carrying no deleterious mutations within a distance s_{bgs}/r of the focal site, and which happened to recombine away from deleterious mutations at greater distances fast enough to avoid being purged.

These thresholds also help us to understand how and when the reduced effective size approximation begins to fully break down for directly selected loci. Specifically, alleles under very strong negative selection are removed from the population without ever reaching frequency $1/2N\sigma$. For these alleles BGS not only has no impact on the heterozygosity and fixation rate, but also does not impact the SFS (Figure S1A). For very strong positive selection (i.e. stronger than the effect of BGS), positive selection can drive an allele to high frequency even if it is linked to one or more deleterious alleles, so long as the net effect on fitness of the haplotype is still strongly positive. BGS effects become unimportant in this limit, so there is no effect on fixation rates, heterozygosity, or the frequency spectrum (Figure S1B).

Phenotypic models with directional mutation-selection-drift balance

With the insights from the single-site model in mind, we now turn to models of complex trait evolution. We consider two different trait-level fitness functions that result in directional selection on individual loci, and ask how background selection impacts the equilibrium state of the population under each fitness function.

Exponential fitness function

With exponential selection, fitness declines exponentially with increasing phenotype, i.e. individual i's fitness is proportional to:

$$W(Z_i) \propto e^{-\frac{Z_i}{\lambda}}, \tag{7}$$

where λ gives the rate at which fitness declines with increasing phenotype. Assuming the scale of the phenotypic variance is small relative to the rate of fitness decay, alleles contributing to the trait evolve independently under additive selection. Alleles with phenotypic effect a experience a directional selection coefficient of approximately $s \approx a\lambda^{-1}$, so that the scaled selection coefficient is

$$\gamma = 2Ns \approx 2Na\lambda^{-1}. \tag{8}$$

For simplicity, we initially assume that all sites have the same effect size. If the population size, N, and the rate of fitness-decay as a function of the phenotype, λ, remain fixed over long timescales,

then fixations of trait-increasing and trait-decreasing mutations eventually reach a detailed balance. Because we assume the population scaled mutation rate is low (i.e. $2N\mu \ll 1$), most sites are fixed at any one time, and the mean phenotype of the population is well approximated by

$$\bar{G} \approx 2Lap^{+}(\gamma), \tag{9}$$

where

$$p^{+}(\gamma) = \frac{\rho(-\gamma)}{\rho(\gamma) + \rho(-\gamma)} = \frac{1}{1 + e^{2\gamma}} \tag{10}$$

is the fraction of sites fixed for the trait-increasing allele at equilibrium. If selection is absent, $p^{+}(\gamma) = \frac{1}{2}$, as sites fix with equal probability at equilibrium. If selection is weak ($\gamma \approx 1$), some sites are expected to be fixed in both states but with a bias against deleterious alleles fixing (i.e. $0 < p^{+}(\gamma) < \frac{1}{2}$). If selection is strong, $p^{+}(\gamma) = 0$, because deleterious alleles cannot fix.

The total genetic variance is then due to a mixture of contributions from sites where the derived allele increases the trait (these mutations originated from sites previously fixed for the trait-decreasing allele, are deleterious, and contribute $h(-\gamma)$ to heterozygosity), and those where the derived allele decreases the trait (which originated from sites previously fixed for the trait-increasing allele, are beneficial, and contribute $h(\gamma)$), i.e.

$$V_G = \theta \left[(1 - p^{+}(\gamma))h(-\gamma) + p^{+}(\gamma)h(\gamma)\right] \tag{11}$$

After applying equation (3) and simplifying, we find that

$$V_G = \frac{\theta}{\gamma} b(\gamma) \tag{12}$$

where $b(\gamma) = 1 - 2p^{+}(\gamma) = tanh(\gamma)$ measures the strength of the fixation asymmetry toward trait decreasing alleles, and $\frac{\theta}{2} = 2NL\mu$ is the total population scaled mutation rate across all sites contributing to the trait.

If we add BGS to the model, this leads to an increase in the fixation rate for deleterious alleles, and a decrease for beneficial alleles (see Equation (4) above). The mean with BGS is therefore

$$\bar{G}_B \approx 2Lp^{+}(\gamma B) = \frac{2L}{1 + e^{2\gamma B}}. \tag{13}$$

The variance, in turn, is given by

$$\begin{aligned} V_{G,B} &= \theta B[(1 - p^{+}(\gamma B))h(-\gamma B) + p^{+}(\gamma B)h(\gamma B)] \\ &= \frac{\theta}{\gamma} b(\gamma B), \end{aligned} \tag{14}$$

so that the variance is reduced by a factor

$$\frac{V_{G,B}}{V_G} = \frac{b(\gamma B)}{b(\gamma)} = \frac{tanh(\gamma B)}{tanh(\gamma)} \tag{15}$$

(Figure 3A). Equation (15) predicts essentially the same reduction in variance as Equation (5) does for the case when all alleles are deleterious, reflecting the fact that even weak selection creates significant fixation asymmetries, so that most new mutations are deleterious. Thus, for effectively neutral alleles in this model (i.e. those with $\gamma \ll 1$), the variance is rescaled by B, whereas for strongly selected sites (i.e. those with $\gamma \gg 1$) there is no effect. For weakly selected sites, the reduction is intermediate and changes monotonically with the scaled selection coefficient. If there is a distribution of effect sizes across sites, the magnitude of the reduction in variance for sites of each effect size follows from Equation (14) independently, given the scaled selection coefficient γ for sites of that effect size.

Due to the contrast with the threshold model that follows below, it is important to note that for weakly selected sites in the exponential model, the difference between the fraction of sites fixed for the trait-increasing allele with background selection, and that without, will be large, so long as there is a non-trivial reduction in effective size. Returning to the single-effect case for simplicity, the mean shift due to BGS, standardized by the scale of the genetic variance, is

$$\frac{\bar{G}_B - \bar{G}}{\sqrt{V_G}} = 2L\left(p^+(\gamma B) - p^+(\gamma)\right)\sqrt{\frac{\gamma}{\theta b(\gamma)}} \approx \gamma(1-B)\frac{L}{\sqrt{\theta}} \gg 1 \tag{16}$$

Thus, the mean phenotype at equilibrium in the presence of background selection may be well outside the range of phenotypes that are seen when it is absent.

Threshold model

In this model, an individual's fitness depends on whether their phenotype exceeds some threshold value, T, or not. Individual i's fitness is proportional to

$$W(Z_i) \propto 1 - \mathbb{1}[Z_i > T]C \tag{17}$$

where C is the cost of exceeding the threshold, and $\mathbb{1}[Z_i > T]$ is an indicator that is equal to one if $Z_i > T$, and zero otherwise. One way to think of this model is as a model of complex disease, in which case we refer to the phenotype, Z_i, as an individual's liability. Alternatively, threshold models of this form have also been used to study genetic load, in which case the underlying phenotype is sometimes called "fitness potential" (Kondrashov, 2018; Milkman, 1978). A complete analysis of this model at equilibrium is given by Berg et al. (2025) in the complex disease context, and we largely adopt this framing here. As with the exponential model above, it is easiest to begin by considering the case where all sites have the same effect size, which we term the "single-effect-size threshold model".

Before considering the effects of BGS in this model, it is helpful to first consider two features of the equilibrium in the absence of BGS, which were derived by Berg et al. (2025).

First, the emergent selection coefficient acting on each site is approximately $s \approx \delta C$, where δ is the fraction of the population who would be pushed across the threshold by substituting one additional liability-increasing allele for a liability-decreasing one. The value of δ is given approximately by

$$\delta \approx F(T-a) - F(T). \tag{18}$$

Here, $F(T)$ is the proportion of the population whose liability exceeds the threshold, or equivalently, the probability that a random individual will develop the disease, while $F(T-a)$ gives the same for a random individual that has had a single liability-decreasing allele substituted for a liability-increasing one. It follows that δ represents the average effect of a liability-increasing allele on an individual's risk of developing the disease, i.e. its "risk effect" (Berg et al., 2025).

If we assume that the liability effect, a, of an individual variant is small relative to the scale of the total variance in liability (i.e. $a \ll \sqrt{V_P}$), then we can further approximate the risk effect as $\delta \approx af(T)$, where $f(T)$ is the density of liability at the threshold. The scaled selection coefficient is then $\gamma \approx 2Naf(T)C$. Thus, in contrast to the exponential model, the selection coefficient of any given variant in the threshold model is highly sensitive to the state of the other variants contributing to liability, via the threshold density $f(T)$. Relatively small changes in the mean liability can therefore result in relatively large changes in the selection coefficient acting on each site, because they lead to changes in the threshold density $f(T)$. The same is true for changes to the variance. Berg et al. (2025) also showed that for reasonable parameter values, the small effect approximation, $\gamma \approx 2Naf(T)C$, is always valid for sites under weak selection (i.e. $\gamma \sim 1$), and is also expected to hold for sites that belong to the "weak end" of the strongly selected regime (e.g. those with $\gamma \approx 10$).

Second, in the threshold model, the fraction of sites fixed for the liability-increasing allele is determined by the position of the threshold relative to the range of possible genetic liability values, not by the evolutionary dynamics. To see this, consider first that the fraction of sites fixed for the liability-increasing allele is approximately equal to the ratio of the mean liability to the theoretical maximum genetic liability, i.e. $p^+(\gamma) \approx \bar{G}/2La$ (see Equation (9)). Second, when the population is at equilibrium, the mean must necessarily be within a few genetic standard deviations of the threshold (i.e. $\bar{G} - T \sim \sqrt{V_G}$). Because the per-site scaled mutation rate is low (i.e. $2N\mu \ll 1$), the genetic standard deviation is small relative to the range of possible genetic liabilities (i.e. $\sqrt{V_G} \ll 2La$). This implies that the fraction of sites fixed for the liability-increasing allele is approximately

$$p^+(\gamma) \approx p_T = \frac{T}{2La} \tag{19}$$

where p_T measures of the relative position of the threshold along the liability continuum. It then follows from inverting the relationship between $p^+(\gamma)$ and γ (Equation (10)) that the value of the scaled

selection coefficient at equilibrium is determined entirely by the relative position of the threshold, i.e.

$$\gamma \approx \frac{1}{2} \ln \frac{1 - p_T}{p_T}. \tag{20}$$

As we note in the previous paragraph, the scaled selection coefficient of small effect sites can also be expressed in terms of the threshold density, i.e. $\gamma \approx 2Naf(T)C$. It follows that the equilibrium is established by evolving the correct threshold density necessary to ensure that Equation (20) holds, i.e. $f(T) \approx \frac{1}{4NCa} \ln \frac{1-p_T}{p_T}$. As a result, at equilibrium in the threshold model, the genetic variance is the same as it is in the exponential model (i.e. Equation (12)), except that the scaled selection coefficient is determined by the position of the threshold (i.e. Equation (20)), rather than by the rate of fitness decline, as it is in the exponential model (i.e. Equation (8)).

Notably, because these results follow from the fixation dynamics, they apply when $p_T \gg \frac{1}{2L}$, so that at least some sites are fixed for the deleterious allele at equilibrium, and fixations occur in both directions. This corresponds to the scaled selection coefficients of $\gamma < \frac{1}{2} \ln 2L$, and we can think of a single-effect-size model meeting this condition as belonging to a "weak selection" regime. In contrast, $\gamma \gtrsim \frac{1}{2} \ln 2L$ implies that $p_T \sim \frac{1}{2L}$, so that there are no sites fixed for the deleterious allele. This corresponds to a strong selection regime, in which the selection coefficients are determined differently from those suggested by Equation (20). Below, we first describe the impact of background selection in the weak selection regime, before considering the strong selection regime, and ultimately extending our analysis to include models with a distribution of effects.

To understand the impact of BGS, it is helpful to start with the effect on the mean liability, $\bar{G}$. In contrast to the exponential model, where the decreased effective population size due to BGS leads to a change in the selective fixation rates and thus a large change in the mean liability (Equation (16)), in the threshold model, selection against individuals exceeding the threshold prevents such a large shift. Rather, the mean must remain close to the threshold, and it follows that the selective fixation rates must remain approximately the same as they are in the absence of BGS. The increase in the rate of drift by a factor of $\frac{1}{B}$ must therefore be compensated by an increase in the (unscaled) selection coefficient by the same factor, i.e.

$$\frac{s_B}{s} = \frac{1}{B}. \tag{21}$$

The resulting scaled selection coefficient in the presence of BGS is

$$\gamma_B = 2NBs_B = 2NB\frac{s}{B} = 2Ns = \gamma, \tag{22}$$

i.e. unchanged from what it is when background selection is not acting. While the substantial shift in mean observed in the exponential model does not occur, there is still a small change in the mean liability. More precisely, to achieve the increase in the selection coefficient captured by Equation (21),

the risk effect must increase by the same factor. This is accomplished by a small shift in the mean liability toward the threshold, so that

$$\frac{\delta_B}{\delta} = \frac{f_B(T)}{f(T)} = \frac{1}{B}, \tag{23}$$

validated by our simulation results (Figure S2). The compensatory evolution of the mean liability also impacts the variance. It is still the case that a fraction $1 - B$ of all new mutations are removed by BGS without exceeding frequency $1/2NBs_{bgs}$, along with their contributions to the fixation rate and the variance. However, because the scaled selection coefficient is unchanged, the changes to the rate of decay in the frequency spectrum at frequencies greater than $1/2NBs_{bgs}$ do not occur. The selective fixation rates in both directions, and therefore the total heterozygosity, are just rescaled by a factor B , so the variance reduction is simply

$$\frac{V_{G,B}}{V_G} = B, \tag{24}$$

independent of the exact value of the selection coefficient. Notably, because this result follows largely from Equation (20), we expect it to hold for the range of scaled selection coefficients for which Equation (20) holds, that is for $\gamma < \frac{1}{2}\ln 2L$. This prediction is supported by our simulation results (Figure 3A).

Impact on disease prevalence

The changes to the threshold density and variance also translate into an increase in disease prevalence. We can approximate the prevalence by assuming that the distribution of liability is Normal. With this assumption, the equilibrium prevalence can be found by standardizing the threshold density, i.e. $\phi(T^*) = f(T)\sqrt{V_p}$ where ϕ is a standard Normal PDF, and V_P is the total variance and $T^* = {}^{T-\bar{G}}/\sqrt{V_P}$ is the distance between the mean and the threshold in standardized units. The prevalence is then given by $\bar{R} = 1 - \Phi\left(\phi_+^{-1}(f(T)\sqrt{V_P})\right)$, where ϕ_+^{-1} is the positively defined inverse of the Normal PDF, Φ is the normal cumulative distribution function (Berg et al., 2025). The prevalence scales close to linearly with the standardized threshold density as long as it is close to zero (which is a reasonable assumption for realistic parameter values). Thus, accounting for both the increase in the threshold density and the reduction in genetic variance due to background selection, the disease prevalence is increased in the presence of background selection approximately by a factor

$$\frac{\bar{R}_B}{\bar{R}} \approx \frac{\phi(T_B^*)}{\phi(T^*)} = \frac{\sqrt{1 - h^2(1-B)}}{B} \tag{25}$$

(Figure 3B) where h^2 is the heritability in the absence of BGS (see Supplementary section A for a derivation of Equation (25), and Figure S3 for an illustration). Thus, for a disease with a heritability close to zero, background selection would be expected to increase the prevalence by approximately a factor of $1/B$, but this factor is closer to $1/\sqrt{B}$ when the heritability is close to one. These calculations

(and all others in this paper), assume a constant population size. In Figure S8, we conduct a limited simulation of the single-effect-size model using a demographic history inferred by Gravel et al. (2011) for the 1000 Genomes sample of Northern Europeans from Utah, and find that background selection has a similar effect on the prevalence in this case as well.

Large effect sites are unaffected by background selection in the single-effect-size model

The above results depend on the assumption that the threshold is far from zero (i.e. $p_T \gg 1/2L$), so that sites are under relatively weak selection (i.e. $\gamma < \frac{1}{2}\ln 2L$). If instead we assume that the threshold is very close to zero (i.e. $p_T \sim \frac{1}{2L}$), then there must be no sites fixed for the deleterious allele, and all sites will belong to the strongly selected regime (i.e. $\gamma > \frac{1}{2}\ln 2L$). In this case, BGS will have no effect on the genetic variance for liability (and consequently no effect on the disease prevalence), because BGS generally has no effect on the expected heterozygosity at sites under strong negative selection. These predictions are also supported by our simulation results (Figure 3A,C).

Variation in effect sizes

So far, we have assumed that all sites have the same effect size. We now consider the extension to the case with a distribution of effects, which we denote by $g(a)$ (though we still assume all sites have the same B). With a distribution of effect sizes, sites with different effects will have different scaled selection coefficients, $\gamma(a)$, and therefore different fixation asymmetries, $p^+(\gamma(a))$. However, the fixation asymmetry for sites of a given scaled selection coefficient must still obey Equation (10), i.e. $p^+(\gamma(a)) = 1/(1+e^{2\gamma(a)})$, while the mean phenotype is still determined by the relative position of the threshold, p_T. Long term equilibrium is therefore achieved when the distribution of scaled selection coefficients across sites satisfies

$$p_T \approx \int g(a)\frac{a}{\bar{a}}p^+(\gamma(a))da \tag{26}$$

(Berg et al., 2025). In words, the average liability-weighted fixation asymmetry given the scaled selection coefficients must be equal to the relative position of the threshold.

If we first continue to assume that $p_T \gg 1/2L$ and we also assume that the distribution of effects is relatively narrow, then all the effect sizes will be small relative to the scale of variation in liability. This allows us to make the small effect size approximation, $\gamma(a) \approx 2Naf(T)C$, across the full range of effect sizes, and we can rewrite Equation (26) as

$$p_T \approx \int g(a)\frac{a}{\bar{a}}\frac{1}{1+e^{4Naf(T)C}}da \tag{27}$$

(Berg et al., 2025). Equation (27) shows that with a distribution of effects, the long-term equilibrium is achieved in the same way as in the single-effect-size model: by evolving the appropriate threshold density, $f(T)$, given the relative position of the threshold, p_T. It follows that our results from the

single-effect-size case also apply when there is a distribution of effect sizes. That is, the reduction in the effective population size by a factor of B is countered by an increase in the threshold density by a factor of $\frac{1}{B}$, so that the scaled selection coefficients remain unchanged across the range of effect sizes. As a result, the contribution to genetic variance is reduced by a factor of B, as in the single-effect-size case.

Interestingly, although Equation (27) employs the small effect approximation, $\gamma(a) \approx 2Naf(T)C$, the solution for $f(T)$ that it implies is in fact valid even if there are large effect sites, as long as a non-trivial fraction of the variance in liability is due to small effect sites for which the approximation holds. This is because the large effect sites will all be under strong selection (i.e. $\gamma(a) \gg 1$), so that they are all fixed for the protective allele (i.e. $p^+(\gamma(a)) \approx 0$) regardless of the exact value of $\gamma(a)$.

Notably, this implies that when both small and large effect sites are present, background selection will also reduce the variance of large effect sites. This is because the selection coefficients of large effect sites are coupled to those of small effect sites via the negative synergistic epistasis created by the interaction between the liability distribution and the threshold. The size of the reduction in variance at large effect sites depends on the size of the variance reduction at small effect sites, as well as the magnitude of the large effect site's effect relative to the background variance in liability.

For example, first consider sites with scaled selection coefficients at the "weak end" of the strongly selected regime (e.g. sites with $\gamma(a) \approx 10$). Their effects are still small enough (in the sense that $a/\sqrt{V_P} \ll 1$) that the small effect approximation, $\gamma(a) \approx 2Naf(T)C$, applies (Berg et al., 2025). As a result, the increase of the threshold density by a factor of $1/B$ due to the impact of background selection on weakly selected sites also leads to an increase in the selection coefficients of these sites. Consequently, their contributions to variance are reduced by a factor of B, just like for weakly selected sites themselves (Figure 3D(i)). This occurs despite the fact that background selection has no effect on such sites in either the single-site model, or the single-effect-size model.

For sites with larger effect sizes, the small effect approximation does not apply at all, and the impact of BGS depends on exactly how large the liability effect is relative to the scale of the total variance in liability. If the liability effect is similar in size to the standard deviation of liability (i.e. $a/\sqrt{V_P} \sim 1$), then the non-linear relationship between the liability effect, a, and the risk effect, $\delta(a)$, results in an increase in the risk effect, and therefore the selection coefficient, $s(a)$, by a factor that is greater than $1/B$, driving an even larger reduction in variance than that experienced by sites belonging to the small effect regime (Figure 3D(ii)). In contrast, if the liability effect is so large that any individual who carries it is pushed across the threshold (i.e. $a/\sqrt{V_P} \gg 1$), then $\delta(a) \approx 1$, and $s(a) \approx C$. Changes to the mean or variance of the liability distribution therefore have no impact on the risk effect, which implies that there is no impact on the selection coefficient, and consequently no effect on the contribution to genetic variance (Figure 3D(iii)). Thus, in the threshold model, only loci with the very largest effect sizes are free from the impact of background selection.

A. Genetic Variance Reduction for Exponential vs. Threshold model

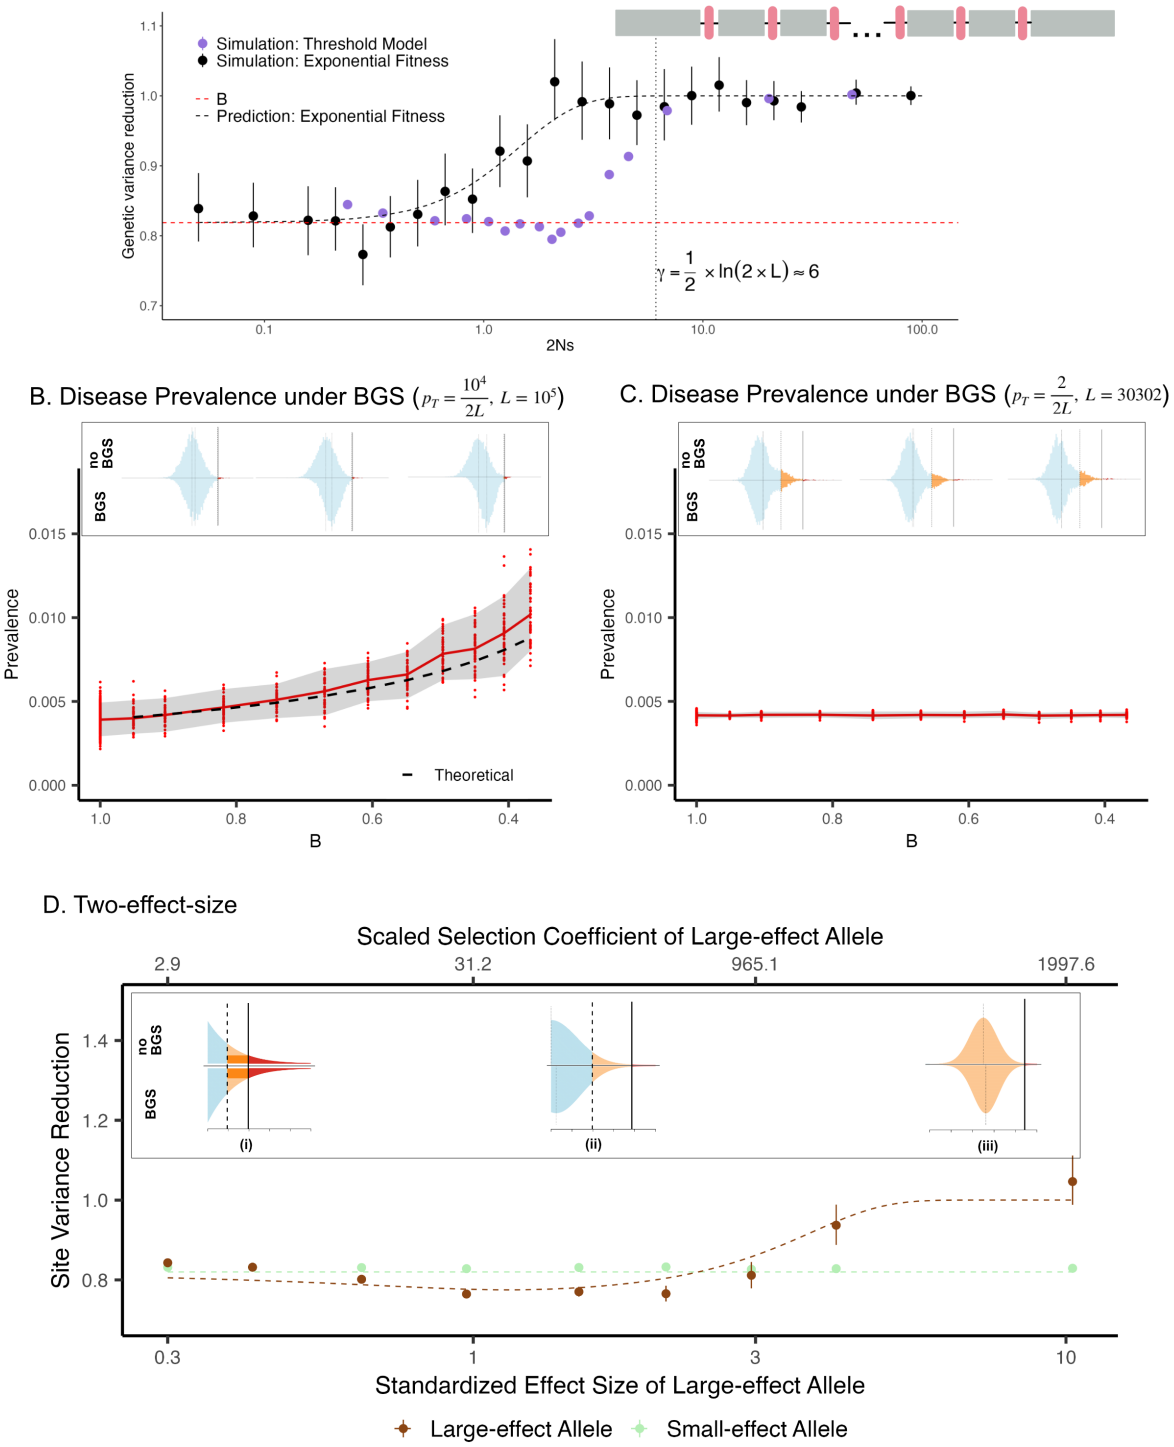

Figure 3: The impact of background selection in the liability threshold model. (A) The genetic variance reduction due to background selection for the exponential fitness function (black) and the threshold model (purple) in simulations. The schematic at top right illustrates the chromosome structure of these simulations, with trait loci (red bars) interspersed among BGS sites (grey regions) (B) The impact of background selection on disease prevalence in the single-effect liability threshold model, when $p_T \gg \frac{1}{2L}$. (B; top inset box) Increasing the strength of BGS shifts the mean toward the threshold and decreases the variance (B; main figure) The disease prevalence increases with stronger background selection. The theoretical prediction of Equation (25) is plotted in a black dashed line. Red dots correspond to averages over individual simulation runs, while the solid line connect means for each value of B. (C) In contrast, when the threshold is close to zero in the single-effect-size model (i.e. $p_T \sim \frac{1}{2L}$), the liability distribution is not affected by background selection, as shown in the box. Consequently, the disease prevalence is not affected by increasing the background selection intensity. In (D), we simulated two-effect traits under the liability threshold model, with one of the effect sizes in the small-effect regime ($a_s = 1$) and the other effect size in the large-effect regime. We varied the effect sizes of the large-effect alleles. (i) – (iii) illustrates how the risk effect size for the large-effect-size allele changes with background selection depending on its effect size. The region in red represents the proportion of diseased individuals no matter whether the individuals carry the large effect allele; orange represents the proportion of individuals that will develop the disease if they carry the large effect allele but will be healthy otherwise (risk effect size); and blue represents the proportion of healthy individuals no matter whether they carry the large effect allele. (i) The effect size is relatively small and the risk effect size can be linearly approximated with the darker orange region. (ii) The linearity does not hold and the risk effect size increase is larger than $\frac{1}{B}$. (iii) When the effect size has a similar scale as the phenotypic variation, the change in liability distribution has a negligible impact on risk effect size. (D) While the small-effect-size alleles (green) have a constant reduction in variance (approximately at a ratio of B), the reduction of variance varies for the large-effect-size allele (brown). With an increased effect size, the ratio drops below B-value and then increases to one. The theoretical predictions (Supplementary section B) of variance reduction for the small-effect and large-effect alleles are plotted in green and red dashed lines.

Stabilizing selection

As we have shown, fitness epistasis in the form of threshold selection results in substantial changes in the way that BGS impacts sites contributing to traits under directional selection. We now consider a different form of fitness epistasis, i.e. stabilizing selection toward an optimum value. As we will show below, this form of epistasis results in BGS effects differ from those described in the two directional selection models we consider above, although in this case, the effects can be understood purely in terms of a corresponding single-site model.

We model stabilizing selection in the standard way by assuming that fitness declines as a Gaussian function of an individual's deviation from an optimum phenotype. We assume this optimum is equal to $L\bar{a}$, and that the population mean is at the optimum, so that an individual's fitness is given by

$$W(Z_i) = e^{\frac{-(Z_i - L\bar{a})^2}{2\omega^2}} \tag{28}$$

(where ω parameterizes the width of the fitness surface), and the mutational input is symmetric at equilibrium.

As is well-known, fitness epistasis due to stabilizing selection induces counterintuitive dynamics for individual alleles contributing to the trait. New mutations are always disfavored, regardless of the direction of their effect, because they push their carriers away from the fitness optimum. If an allele increases in frequency despite this selection against it, fitness epistasis between the focal site and the background results in small, compensatory shifts in frequency at all other sites that contribute to the phenotype, keeping the mean at the optimum. These dynamics lead to changes in the relative fitnesses of the three genotypes at the focal allele, the net result of which is that individual alleles evolve as if experiencing underdominant selection with a selection coefficient $t \approx \frac{a^2}{2\omega^2}$ against heterozygotes (Fisher, 1930; Hayward & Sella, 2022; Negm & Veller, 2024; Robertson, 1956). We denote the corresponding scaled coefficient by $\tau = 2Nt$.

Single-site underdominant selection model

Like the directional selection case, we can begin to understand the impact of BGS on traits under stabilizing selection by understanding the effect on individual variant in the corresponding single-site underdominant model. In the underdominant model, the expected change for an allele at frequency p, is $\mathbb{E}\left[\Delta p | t, p\right] = -t(1-2p)p(1-p)$: selection acts symmetrically against the minor allele, but weakens as it climbs in frequency. Because the relationship between fitness and genotype is a function of the allele frequency in this model, there is no simple analog of Equation (3) relating contributions to heterozygosity to the strength of selection and/or fixation rate. However, we can largely understand the behavior of the single-site underdominant model in terms of its similarities and differences to the single-site directional model.

If selection is either very weak (i.e. $\tau \ll 1$) or very strong (i.e. $\tau \gg 1$), then the changes in the strength and direction of selection during an underdominant allele's transit are unimportant—in the first case because selection is unimportant in general, and in the second case because selection is so strong that the frequency never rises far above zero. In these two cases, the evolution of an underdominant allele is identical to that of a negatively selected allele in the directional model (and we would naturally expect the impact of background selection to be the same as well). However, for alleles that arise under weak selection (roughly, those satisfying $1 \lesssim \tau \lesssim 30$), it is useful to think of them as having a frequency-dependent scaled selection coefficient, which is equal to $\tau(1-2p)$. For these alleles, the initial selection they experience is weak enough that they have an appreciable chance of drifting up to intermediate frequency (Supplementary section C; Figure S4). Once at intermediate frequency, they evolve as if neutral until the minor allele frequency drops back below a threshold of approximately $p \approx \frac{1}{2}(1-\frac{1}{\tau})$, at which point selection again becomes effective and begins to push the allele either to fixation or loss. Such alleles therefore spend more time at the intermediate frequencies that contribute most to heterozygosity than a negatively selected allele with the same scaled selection

coefficient in the directional model.

These dynamics can be summarized (for any value of τ) by the sojourn time of the minor allele at frequency p, which is given approximately by

$$\Psi(p|\tau) \approx 2\frac{e^{-2\tau p(1-p)}}{p(1-p)}, \tag{29}$$

(Hayward & Sella, 2022). The expected lifetime contribution to heterozygosity is therefore

$$h(\tau) \approx 2\int_0^{\frac{1}{2}} x(1-x)\Psi(x|\tau)dx = 2\frac{D_+(\sqrt{\frac{\tau}{2}})}{\sqrt{\frac{\tau}{2}}}, \tag{30}$$

where $D_+(y) = e^{-y^2}\int_0^y e^{u^2}du$ is the Dawson function. Equation (30) coincides with the predictions of the negative directional selection model when selection is either very weak (i.e. $\tau \ll 1$) or very strong (i.e. $\tau \gg 1$), but predicts a greater heterozygosity than the directional model under weak selection (i.e. $1 \lesssim \tau \lesssim 30$), consistent with the dynamics described above (Figure S6).

Applying the reduced effective size approximation, we find that background selection is expected to modify the contribution to heterozygosity of an underdominant site by a factor of

$$B_\tau = B\frac{h(\tau B)}{h(\tau)} = \sqrt{B}\frac{D_+(\sqrt{\frac{\tau B}{2}})}{D_+(\sqrt{\frac{\tau}{2}})} \tag{31}$$

Like the heterozygosity in general, Equation (31) simplifies in the same way as for the directional case when selection is either very weak (i.e. $\frac{h(\tau B)}{h(\tau)} \approx 1$, so that $B_\tau \approx B$) or very strong (i.e. $\frac{h(\tau B)}{h(\tau)} \approx \frac{1}{B}$, so that $B_\tau \approx 1$; see Figure 4A inset). However, for some intermediate values of τ, $\frac{h(\tau B)}{h(\tau)} > \frac{1}{B}$, implying that $B_\tau > 1$. Thus, surprisingly, we find that background selection can increase heterozygosity for some combinations of τ and B (Figure 4A).

This result is closely related to the observation that alleles under strong negative selection in the single-site additive selection model experience no reduction in heterozygosity due to BGS (see above). In that case, each allele that would have exceeded frequency $1/2NBs_{bgs}$ in the absence of background selection does so only if it also initially arises on or quickly recombines onto a background free of deleterious alleles. This happens with probability B. However, the increased rate of drift that these alleles experience means that they attain a higher frequency on average, each contributing more to heterozygosity on average. In the directional selection case with $2Ns_{bgs} < \gamma \ll -1$, we found that $h(\gamma B)/h(\gamma) = 1/B$, so these two effects cancelled out, and there was no net effect on heterozygosity (Figure 2C(i)). However, in the underdominant case, reaching higher frequency is accompanied by even weaker selection. This makes them even more likely to continue drifting to higher frequencies, so that the net result is an increase in heterozygosity.

The scaled selection coefficients that are most strongly impacted by this dynamic, and the size

of their increase in heterozygosity, depend on the magnitude of the reduction in effective population size, B. For example, if the background selection effect is modest, as is the case for the mean effect in humans (where $\bar{B} = 0.83$; Buffalo & Kern, 2024; Murphy et al., 2022), then alleles with scaled selection coefficients satisfying roughly $3 \lesssim \tau \lesssim 10$ will see an increase in their probability of reaching intermediate frequency, with the biggest increase coming around $\tau \approx 6$ (Figure S5). The biggest increase in heterozygosity occurs around $\tau \approx 10$ and is very modest, with $B_\tau \approx 1.03$ (Figure 4A).

In contrast, if we consider a stronger background selection effect of $B = 0.2$ (consistent with the genome-wide average for *Drosophila*; Elyashiv et al. (2016)), then for sites with $\tau \approx 6$ the increase in the rate of drift is not enough to overcome the sharp reduction in the number of alleles exceeding frequency $1/2NBs_{bgs}$, and there are roughly the same number of alleles reaching intermediate frequency as in the absence of BGS (Figure S5). The biggest increase in the number of alleles reaching intermediate frequency instead occurs around $\tau \approx 10$, but the range over which the heterozygosity is increased extends up to $\tau \approx 50$. The biggest increase in heterozygosity occurs at $\tau \approx 25$, and is more significant than the previous case with $B_\tau \approx 1.22$ (Figure S5). Notably, the distortions to the frequency spectrum in this regime can be rather extreme, with significant decreases in the number of low-frequency minor alleles paired with large increases in the number of alleles near $p \approx 0.5$ (Figure S7).

Stabilizing selection trait model

When we turn to considering the effect of background selection in a multi-locus model of a complex trait under stabilizing selection, we find the results for the single-site underdominant model are faithfully replicated at sites contributing to a trait under stabilizing selection. We verified this by simulating a polygenic trait under stabilizing selection in the presence of background selection, using the chromosome configuration described previously to ensure that QTLs experience background selection with $B \approx 0.83$, consistent with the human average. We find good agreement between the single-site predictions of Equation (31) (B_τ) and our simulation results (Figure 4A), validating our theoretical predictions. The applicability of the single-site results to the full multi-locus mode stands in contrast to the threshold model we considered above, and arises because selection does not have to maintain a fixation asymmetry in the stabilizing selection model.

Discussion

In this paper, we have used a combination of analytical theory and simulations to study the impact of background selection on individual variants, and on phenotypes, in three long-studied models of phenotypic mutation-selection-drift balance. Here, we consider 1) what we have learned in general about the impact of background selection in models of phenotype mutation-selection-drift balance, and 2) what our results imply about the likely impact of background selection on phenotypes in

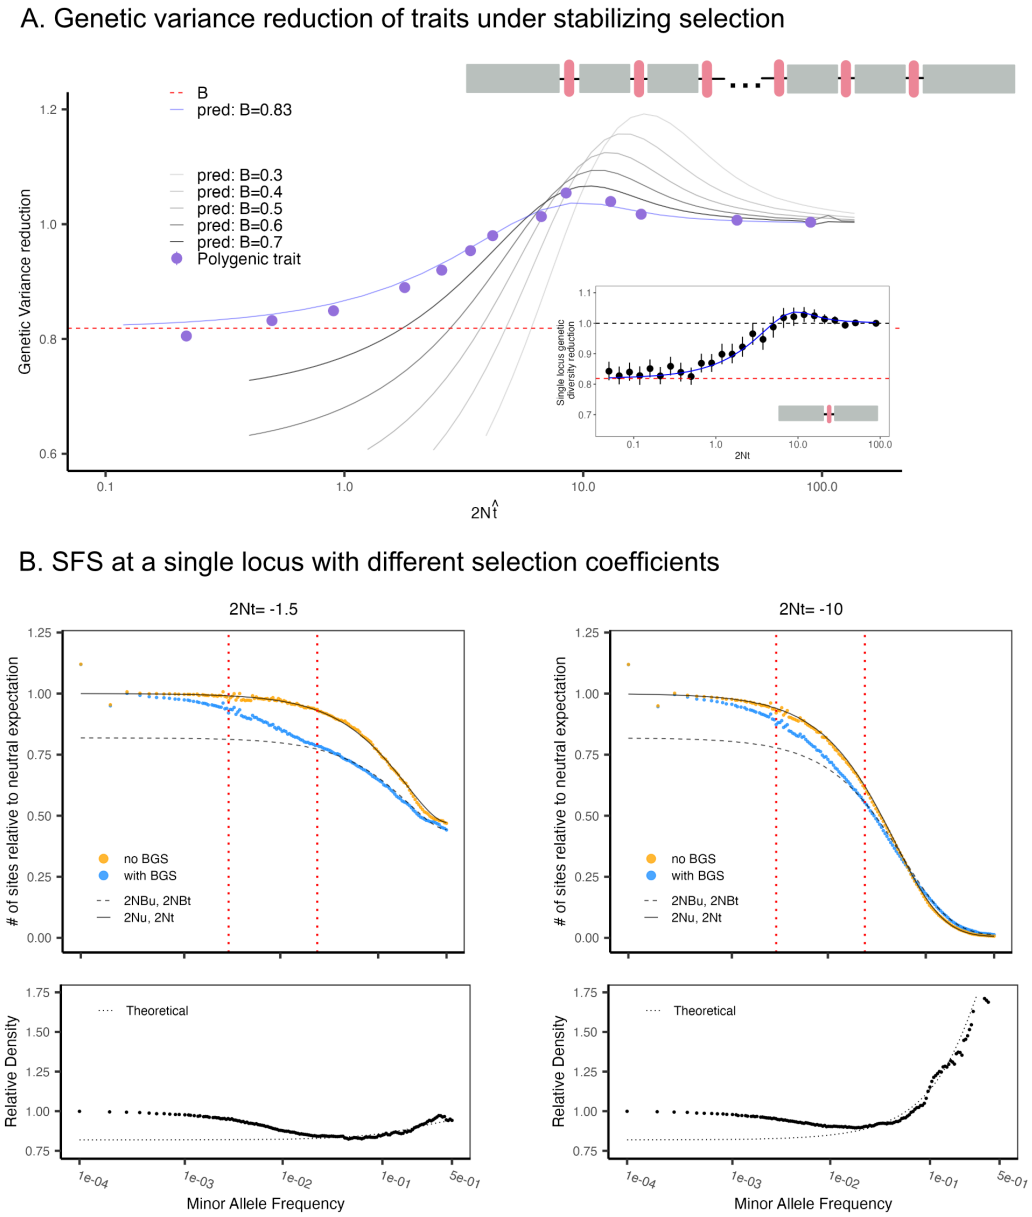

natural populations.

Directional selection models

The directional selection models we consider represent opposite ends of a continuum with respect to the role of synergistic epistasis. In the exponential selection model (i.e. without epistasis), sites

Figure 4: Stabilizing Selection. In (A), we simulated traits with a single constant effect size under stabilizing selection with chromosome structure shown on the top right. Across different settings, we varied effect sizes and measured the reduction in genetic variance due to background selection (y-axis). The x-axis represents the measured average selection coefficient at each causal locus. The method used for these estimations is detailed in Methods. The red dashed line is the expected genetic diversity reduction of a neutral allele at $B = 0.83$. The blue solid line is the theoretical prediction of the genetic reduction of a single selected focal allele with a BGS intensity at $B = 0.83$. The prediction aligns well with the single-locus simulation results, shown in the insect. To demonstrate that the size of increase in heterozygosity due to BGS depends on the magnitude of the BGS effect, we also plot the predicted reduction of genetic diversity with different BGS intensities in grey lines.
In (B) we plot the "relative" site frequency spectrum of the focal selected site with different selection coefficients ($2Nt = 1.5$ on the left panel and $2Nt = 5$ on the right panel) observed in single-locus simulations. The y-axis is the number of sites at each frequency relative to the neutral expectation with $\theta = 4N\mu$. The color represents the scenario whether there is background selection or not. The solid line represents the predicted deleterious SFS with effective population size N; while the dashed line represents the predicted deleterious SFS with effective population size NB. The insets show the relative density of alleles at each frequency bin.

with scaled selection coefficients in the effectively neutral regime experience diversity reductions in the presence of background selection that are similar to those of loci that are truly neutral, but this effect weakens as the strength of direct selection increases, so that background selection has a negligible effect on contributions to variance from loci with $\gamma \gtrsim 3$. For loci with $\gamma \lesssim 3$, background selection has opposing impacts on the fixation rates of alleles under positive and negative selection, and therefore drives an increase in their contribution to the mean phenotype that is large relative to the scale of the standing genetic variation. However, because the rate of fitness decline is independent of the phenotypic value, this change in the position of the mean does not impact the underlying selection coefficients. These effects all follow directly from the local impact of selection against linked backgrounds carrying deleterious alleles, and can therefore be predicted straightforwardly from the classic "reduced effective population size" approximation for a single site.

In contrast, when selection acts through a threshold, there is strong synergistic epistasis among sites. The increase in the mean driven by the impact of background selection on fixation rates quickly causes an increase in the raw selection coefficients (i.e. *unscaled* by population size). This increase in the unscaled selection coefficients offsets the "reduced effective population size" effect on the fixation rates, so that the fixation rates ultimately remain unchanged. Because the selection coefficients of all sites are coupled to one another via the density of the liability distribution at the threshold (and, for loci with large effects, through higher moments of the liability distribution), the increase in unscaled selection coefficients, and therefore the reduction in genetic variance, extends even to strongly selected loci that would otherwise not normally be affected by background selection.

Notably, this is a global effect, which does not arise directly from the effects of linked selection on allelic dynamics, but from the downstream consequences of this change in the dynamics. In

our analyses, we assume that B is the same for all loci. The result is that the *unscaled* selection coefficients of weakly selected sites increase by a factor of $1/B$, exactly canceling the impact of the reduced effective population size on the scaled selection coefficients. If the value of B varies across the genome, then sites in different regions of the genome will experience different reductions in the effective population size, but because the increase in *unscaled* selection coefficients is mediated by the threshold density, $f(T)$, they will all experience an increase in their selection coefficients by a multiplicative factor $1/\tilde{B}$, where $\tilde{B}$ is a weighted average of the B values at sites that contribute to the trait. A complete exploration of this case is beyond our scope, but in Supplementary Section F we consider the case of a single-effect-size threshold model in the weak selection regime, in which there are two different genomic regions with different B values. There, we show how to solve this model numerically, and derive analytical solutions for two limiting cases. Briefly, when $p_T \approx 1/2$, so that there is little mutational pressure and selection is very weak, $\tilde{B}$ is approximately equal to a simple average of B across the two regions, and the variance is ultimately reduced in each region in proportion to the local B value. In contrast, in the limit as the fraction of fixed sites for the liability increasing allele goes to zero (i.e. $p_T \to 0$), $\tilde{B}$ comes to equal the minimum of the B values across regions, and the two regions experience variance reductions of the same size.

Because the two directional models we consider represent the ends of a continuum with respect to the effects of epistasis, our results also provide some insight into the behavior of models that are intermediate (e.g. Charlesworth (1990) and Kimura and Maruyama (1966)), which would be expected to have intermediate behavior. Models with strong synergistic epistasis, in which fitness transitions from high to low over a narrow range of genotypes, would be expected to exhibit relatively small changes in the mean phenotype, and would be expected to exhibit reductions in the contribution to genetic variance across a broad range of selection coefficients. In contrast, models with weaker epistasis, in which fitness declines at a more steady rate with an increasing number of deleterious alleles, would be expected to exhibit relatively large changes in the mean, with reductions in variance mostly confined to weakly selected loci.

Stabilizing selection model

Our results regarding the stabilizing selection model are in some ways simpler, as the effects can be explained entirely in terms of the corresponding underdominant single-site model. Whether background selection increases or decreases a sites's expected contribution to heterozygosity and trait variance (or has no effect) depends on the strength of direct selection acting on the site. This simplicity results largely from the fact that we assume symmetric mutation in our analysis of stabilizing selection. As a result, selection does not need to maintain a directional fixation asymmetry, so the global dynamics that emerge in our analysis of the threshold model do not factor into the stabilizing selection model. However, this would be expected to change in a model with both stabilizing selection and a directional mutational pressure. The impacts of BGS in such a model would be expected to blend features of the

threshold model and the symmetric stabilizing selection model that we consider here.

The linked selection effects we study in this paper are not the only channel through which linkage can impact the evolution of traits under stabilizing selection. This is because by selecting against individuals with extreme phenotypes, stabilizing selection generates negative LD between alleles of similar signs, resulting in a short-term decrease in the genetic variance, a phenomenon known as the Bulmer effect (Bulmer, 1974; Bulmer, 1971). The magnitude of this effect depends on the harmonic mean recombination rate across sites contributing to the trait, so it is stronger if sites are tightly linked. Viewed from the perspective of a single focal site, the Bulmer effect amounts to a negative covariance between the trait increasing allele at that site and the genetic background contribution from all other sites. This results in an attenuation of the marginal phenotypic association, and thus the marginal effect on fitness, of the allele at the focal site (i.e. a decrease in the site's "effective" selection coefficient; Negm & Veller, 2024). Although we did not set out to study this effect, in the process of calibrating our simulations we discovered that this decrease in the strength of selection experienced by individual sites leads to a long-term increase in the *genic* variance that is greater than the short-term reduction due to the negative LD. As a result, the total, long-term impact of the Bulmer effect is actually to increase the genetic variance, rather than to decrease it. We detail this incidental finding briefly in Supplementary section D.

Implications for natural populations

Our study raises the important question of what effect we expect BGS to have on the evolution of complex traits in natural populations. It is useful first to revisit the existing empirical analyses that have invoked BGS effects to explain patterns in the genetic architecture of complex traits.

Specifically, working in *C. elegans*, Rockman et al. (2010) invoked BGS to explain the observation that loci associated with gene expression in gene-poor, high-recombination chromosome arms explain significantly more variance in transcript abundance than loci in the gene-rich chromosomal centers. This is a plausible explanation in light of our results, so long as most of the sites that Rockman et al. (2010) identified as affecting transcript abundance are effectively neutral or very weakly selected (if this is not the case, then BGS does not result in a local decrease in variance).

Alternatively, stratified LD score regression (S-LDSC) analyses of human complex traits and diseases have suggested that contributions to heritability are enriched in regions experiencing stronger background selection (Gazal et al., 2017; Pardiñas et al., 2018). These results are potentially consistent with any of the models that we consider, even though only stabilizing selection can drive an absolute increase in the heritability. This is because S-LDSC measures heritability enrichments on a per-SNP basis. While BGS reduces heritability in models of directional/negative selection by quickly removing a fraction $1-B$ of newly arising mutations, the contribution per segregating site is increased slightly due to the reduced efficacy of selection (see Supplementary section E, Figure S10). Second, because both the contributions to heritability (Finucane et al., 2015; Maurano et al., 2012; Pickrell,

2014) and the signals of BGS are enriched for signals of biological function in general (Murphy et al., 2022), it is likely that causal alleles simply arise at higher rates in regions of stronger BGS. Thus, such enrichments may simply reflect the heterogeneous distribution of functional sites across the genome, rather than any specific effect of background selection.

More generally, as our results show, the exact impact of BGS on complex traits depends on the model of selection that acts on the phenotype. Much of the literature on complex trait evolution assumes that stabilizing selection is the norm (Kingsolver et al., 2001; Sella & Barton, 2019). This assumption has been hard to assess, but GWAS data on human quantitative traits and diseases has recently begun to shed some new light, as analyses of the joint distribution of effect sizes and allele frequencies among genome-wide significant GWAS loci support a model of ubiquitous stabilizing selection across phenotypes (Simons et al., 2022). Notably, this appears to be the case even for GWAS variants associated with disease traits (Berg et al., 2025; Koch et al., 2024), which we might otherwise have expected to fit a model of directional selection. Assuming these results hold, the impact of background selection on human phenotypic variation may be modest, if somewhat complicated. For example, Simons et al. (2022) inferred a distribution of selection coefficients under a model of stabilizing selection for 95 quantitative traits from the UK Biobank, finding that this distribution spans several orders of magnitude. However, most of the mass falls into the weak selection regime. Thus, for the bulk of GWAS loci, the primary effect of background selection may be to marginally increase the number of generations that a typical variant spends drifting at intermediate frequencies before selection is able to push it back out of the population. The net effect on levels of genetic variance might in principle be positive, though the effect would be quite small, given the modest strength of background selection in humans.

It remains unclear to what extent the results of human GWAS should guide expectations about quantitative genetic variation in other species. Assuming stabilizing selection is the norm, then the impact of background selection in other species depends on how the strength of selection at a typical trait associated variant, the effective population size (before factoring in background selection effects), and the strength of background selection, all vary across taxa. For example, *Drosophila melanogaster* is estimated to be impacted by much stronger linked selection than humans (Elyashiv et al., 2016). Thus, if quantitative genetic variation in Drosophila is similarly dominated by weakly selected contributions to phenotypes under stabilizing selection, then the increase in genetic variance might plausibly be non-trivial. However, Drosophila are also estimated to have an effective population size two orders of magnitude larger than humans (Buffalo, 2021). Thus, if the variants responsible for the bulk of quantitative genetic variation in *D. melanogaster* have similar (unscaled) selection coefficients as in humans, then we would expect them to be strongly selected, and we would conclude that background selection would have no impact on their contributions to genetic variance.

Of course, it is plausible that even if stabilizing selection is a sufficient model for GWAS traits in humans, it may not be a sufficient explanation for all quantitative genetic variation, either because the

kinds of phenotypes that human geneticists are interested in and able to measure represent a biased subset of all axes of phenotypic variation, or because humans may not be an ideal model for other species. It would be particularly interesting, vis-à-vis our results, to know whether there exist axes of phenotype variation that clearly correspond to our threshold model, particularly if the threshold is positioned such that only a small number of deleterious fixations are tolerated. To the extent that axes of phenotypic variation fitting this description exist in natural populations, then background selection might ultimately be responsible for a rather large reduction in levels of quantitative genetic variance, in contrast to what we would expect in the case of stabilizing selection.

Methods

Simulation of single-site selection models

To validate the predictions of our single-site models, we performed SLiM simulations (Haller & Messer, 2019) on a single site under directional or underdominant selection, across a range of selection coefficients. In each case, we simulated $5,000$ recombining diploid genomes, with each genome having $M = 4 \times 10^5$ background sites, and a single focal site in the middle of the chromosome. The per-site recombination rate is set to 2.5×10^{-6}, such that the map length is equal to 1 Morgan. The mutation rate at each of the M background sites is 2.5×10^{-7}. When we simulate without background selection, all mutations at the M background sites other than the focal one are neutral. When we simulate with background selection, all mutations at these M background sites have a selection coefficient $s = -0.005$. With $\nu = 2.5 \times 10^{-7}$ and $r = 2.5 \times 10^{-6}$, this produces a B of 0.82. To vary the background selection strength, we tune the mutation rate (ν) with a fixed recombination rate (r) to produce a specified B (Equation (1)).

The mutation rate at the focal site is set to zero, and we introduce mutations at this site "manually". We first let the simulation of the background sites run for $2N$ generations so that the population reaches mutation-selection balance. After this burn-in period, we introduced a single selected allele at the focal site in the center of the genome.

To measure the reduction of heterozygosity by BGS, we adopt the method used by Charlesworth et al. (1993) to measure the expected total contribution to heterozygosity of a new mutation during its transit through the population (H_{tot}). This is given by the sum of heterozygosity at the focal site during the T generations that the mutation segregated in the population,

$$H_{tot} = 2 \sum_{i=1}^{T} x_i(1 - x_i), \tag{32}$$

which has a theoretical expectation of $\mathbb{E}\left[H_{tot}|\gamma\right] = h(\gamma)$, or $\mathbb{E}\left[H_{tot}|\tau\right] = h(\tau)$, depending on the model, and is thus proportional to the expected heterozygosity for a single site at a random sampling

time in the low mutation limit. For each selection coefficient, we average the value of H_{tot} across 10^7 separate focal sites. These mutations were simulated across 10^4 replicate burn-ins, with 10^3 mutations considered per burn-in. We then take the ratio of H_{tot} under BGS relative to no BGS as our estimator of the impact of background selection on the expected heterozygosity. To obtain a standard error, we treat the value of H_{tot} for a single mutation's transit through the population as a single data point, and bootstrap across these values.

Simulation of complex Traits

Exponential selection model

To plot the reduction in variance due to background selection under the exponential model, we take advantage of the theoretical equivalence between the exponential model and the single-site additive model. To this end, we simulate a single site under directional selection with a fixed selection coefficient, as described above, except that each time a mutation ends its transit through the population and we revert the simulation back to the end of the burn-in, we choose the sign of the effect of the next mutation to be opposite to that of the allele that fixed at the end of the previous mutation's transit; i.e. if it was the deleterious allele that fixed, then the next mutation would be beneficial, and vice versa. This scheme simulates the long-term fixation dynamics under the exponential model. Calculations regarding the reduction in heterozygosity due to BGS are then otherwise performed as described above in the single-site case.

Threshold selection model

In this case, we simulate under the full threshold model, again with $5,000$ diploid individuals. For each choice of $p_T = \frac{T}{2La}$ in the single-effect-size model, we simulated $L = 10^5$ causal sites which contribute to phenotypic variation, set $a = 1$, and choose T so as to attain the correct value of p_T. We set the mutation rate for causal sites to $\mu = 10^{-7}$. All causal sites are equally spaced along the chromosome. Between each pair of causal sites, there are 25 BGS sites, as well as a single neutral site (these neutral sites are used to empirically check the strength of the BGS effect, i.e. the B value, against the expectation from Equation (1)). For most of our simulations, the recombination rate among adjacent sites is set to $r = 2 \times 10^{-7}$, and the mutation rate at BGS sites is set to $\nu = 2 \times 10^{-8}$, to yield an expected B value of $e^{-0.2} \approx 0.82$, approximately equal to the estimated human average. To ensure that all causal sites experience a similar background selection effect, we added 6.5×10^5 negatively selected sites on each end of the chromosomes, such that the chromosome has a total map length of 0.8 Morgan. In the simulations where we vary the strength of the BGS effect, we do so by varying the mutation rate for BGS sites, ν, with the recombination rate held constant.

To compare traits with and without background selection, we first simulate a trait with the environmental variance tuned so that the heritability is $h^2 = 0.5$ in the absence of background selection

(i.e. setting $s = 0$ for all BGS sites). We then rerun the same simulation with background selection turned on (i.e. with $s = -0.005$ for all BGS sites), holding the environmental variance constant at the same value used in the no background selection case. In all cases, we set the fitness cost of being across the threshold at $C = 0.2$. An individual's fitness is then computed as a product of their fitness effect due to the trait (Equation (17)) together with independent multiplicative contributions from the BGS sites.

To compare the single-effect-size trait model to the single-site results, we assign a scaled selection coefficient to each set of trait simulations. This quantity measures the average selection coefficient that causal variants experience in the trait model. For each choice of $p^+(a)$ in the simulations, we estimated the unscaled selection coefficient by regressing individual fitness against genotypes in the absence of background selection. This regression was performed for each causal site segregating in a given generation and repeated across multiple generations, with the average slope serving as our estimate. We then obtained the scaled selection coefficient by multiplying this estimate by $2N$. The corresponding simulation run, conducted with the same $p^+(a)$ but with background selection included, was assigned the same scaled selection coefficient. In principle, in the small effect regime we could obtain the scaled selection coefficient analytically using Equation (20), but this method breaks down for large effects, so we use the empirical regression-based estimate across the full range.

Stabilizing Selection

Our simulations under stabilizing selection are similar to those under the threshold model, with a few key differences. Notably, to avoid the Bulmer effect in our simulations of stabilizing selection (see Supplementary section D and Figure S9), we had to increase the harmonic mean recombination rate between pairs of causal sites (Bulmer, 1974; Bulmer, 1971; Negm & Veller, 2024; Veller & Coop, 2024). In these simulations, we set $L = 5.4 \times 10^4$, and split the causal sites across 6 independent chromosomes, so that each chromosome has 9000 causal sites. Between each pair of adjacent causal sites on a given chromosome, there are 98 BGS sites and 1 neutral site. The recombination rate is set to $r = 10^{-6}$ per base pair, while the mutation rate is $\mu = \nu = 10^{-7}$ for both causal and BGS sites. We set the width of the stabilizing selection surface $\omega^2 = 2N$ and vary the effect size for each site a.

To compare the single-effect-size trait model to the single-site results, for each choice of a, we measure the average heterozygosity of all the causal sites in the simulations without background selection over multiple generations and solved for the τ with Equation (30). We then multiply this estimate by $2N$ and assign it to the simulations with the same a with and without background selection.

Resource

The documented code, including simulations, numerical solutions of the model, and figure-geneartion scripts, is publicly available at https://github.com/xinyli/BGS_msdb.git.

Acknowledgements

We would like to thank John Novembre for many helpful conversations, and for reading and commenting extensively on several drafts of the manuscript. We would also like to thank members of the Berg, Novembre and Steinrücken labs for feedback on the work at various stages. XL was supported by NIH grant R35 GM151257 to John Novembre and R01 HG010773 to John Novembre and Xin He. JJB was supported by R35 GM149521.

References

Barton, N. H. (1995). Linkage and the limits to natural selection. *Genetics*, *140*(2), 821–841.

Berg, J. J., Li, X., Riall, K., Hayward, L., & Sella, G. (2025). *Mutation-selection-drift balance models of complex diseases* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/2025.05.18.654722v1].

Buffalo, V. (2021). Quantifying the relationship between genetic diversity and population size suggests natural selection cannot explain lewontin's paradox. *Elife*, *10*, e67509.

Buffalo, V., & Kern, A. D. (2024). A quantitative genetic model of background selection in humans. *Plos Genetics*, *20*(3), e1011144.

Bulmer, M. (1974). Linkage disequilibrium and genetic variability. *Genetics Research*, *23*(3), 281–289.

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

Bustamante, C. D., Wakeley, J., Sawyer, S., & Hartl, D. L. (2001). Directional selection and the site-frequency spectrum. *Genetics*, *159*(4), 1779–1788.

Charlesworth, B. (1990). Mutation-selection balance and the evolutionary advantage of sex and recombination. *Genetics Research*, *55*(3), 199–221.

Charlesworth, B. (1994). The effect of background selection against deleterious mutations on weakly selected, linked variants. *Genetics Research*, *63*(3), 213–227.

Charlesworth, B., Morgan, M., & Charlesworth, D. (1993). The effect of deleterious mutations on neutral molecular variation. *Genetics*, *134*(4), 1289–1303.

Coop, G. (2016). *Does linked selection explain the narrow range of genetic diversity across species?* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/042598v1].

Corbett-Detig, R. B., Hartl, D. L., & Sackton, T. B. (2015). Natural selection constrains neutral diversity across a wide range of species. *PLoS biology*, *13*(4), e1002112.

Crow, J. F., & Kimura, M. (1979). Efficiency of truncation selection. *Proceedings of the National Academy of Sciences*, *76*(1), 396–399.

Cvijović, I., Good, B. H., & Desai, M. M. (2018). The effect of strong purifying selection on genetic diversity. *Genetics*, *209*(4), 1235–1278.

Elyashiv, E., Sattath, S., Hu, T. T., Strutsovsky, A., McVicker, G., Andolfatto, P., Coop, G., & Sella, G. (2016). A genomic map of the effects of linked selection in drosophila. *PLoS genetics*, *12*(8), e1006130.

Falconer, D. S. (1996). *Introduction to quantitative genetics*. Pearson Education India.

Finucane, H. K., Bulik-Sullivan, B., Gusev, A., Trynka, G., Reshef, Y., Loh, P.-R., Anttila, V., Xu, H., Zang, C., Farh, K., et al. (2015). Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nature genetics*, *47*(11), 1228–1235.

Fisher, R. A. (1930). *The genetical theory of natural selection*. Oxford University Press.

Gazal, S., Finucane, H. K., Furlotte, N. A., Loh, P.-R., Palamara, P. F., Liu, X., Schoech, A., Bulik-Sullivan, B., Neale, B. M., Gusev, A., et al. (2017). Linkage disequilibrium–dependent architecture of human complex traits shows action of negative selection. *Nature genetics*, *49*(10), 1421–1427.

Gravel, S., Henn, B. M., Gutenkunst, R. N., Indap, A. R., Marth, G. T., Clark, A. G., Yu, F., Gibbs, R. A., Project, 1. G., Bustamante, C. D., et al. (2011). Demographic history and rare allele sharing among human populations. *Proceedings of the National Academy of Sciences*, *108*(29), 11983–11988.

Haller, B. C., & Messer, P. W. (2019). Slim 3: Forward genetic simulations beyond the wright–fisher model. *Molecular biology and evolution*, *36*(3), 632–637.

Hayward, L. K., & Sella, G. (2022). Polygenic adaptation after a sudden change in environment. *Elife*, *11*, e66697.

Hudson, R. R., & Kaplan, N. L. (1995). Deleterious background selection with recombination. *Genetics*, *141*(4), 1605–1617.

Karn, M. N., & Penrose, L. (1951). Birth weight and gestation time in relation to maternal age, parity and infant survival. *Annals of eugenics*, *16*(1), 147–164.

Kimura, M. (1962). On the probability of fixation of mutant genes in a population. *Genetics*, *47*(6), 713.

Kimura, M. (1964). Diffusion models in population genetics. *Journal of Applied Probability*, *1*(2), 177–232.

Kimura, M. (1969). The number of heterozygous nucleotide sites maintained in a finite population due to steady flux of mutations. *Genetics*, *61*(4), 893.

Kimura, M., & Maruyama, T. (1966). The mutational load with epistatic gene interactions in fitness. *Genetics*, *54*(6), 1337.

King, J. L. (1966). The gene interaction component of the genetic load. *Genetics*, *53*(3), 403.

Kingsolver, J. G., Hoekstra, H. E., Hoekstra, J. M., Berrigan, D., Vignieri, S. N., Hill, C., Hoang, A., Gibert, P., & Beerli, P. (2001). The strength of phenotypic selection in natural populations. *The American Naturalist*, *157*(3), 245–261.

Koch, E., Connally, N., Baya, N., Reeve, M., Daly, M., Neale, B., Lander, E., Bloemendal, A., & Sunyaev, S. (2024). *Genetic association data are broadly consistent with stabilizing selection shaping human common diseases and traits* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/2024.06.19.599789v2].

Kondrashov, A. S. (2018). Through sex, nature is telling us something important. *Trends in Genetics*, *34*(5), 352–361.

Maurano, M. T., Humbert, R., Rynes, E., Thurman, R. E., Haugen, E., Wang, H., Reynolds, A. P., Sandstrom, R., Qu, H., Brody, J., et al. (2012). Systematic localization of common disease-associated variation in regulatory dna. *Science*, *337*(6099), 1190–1195.

McVEAN, G. A., & Charlesworth, B. (1999). A population genetic model for the evolution of synonymous codon usage: Patterns and predictions. *Genetics Research*, *74*(2), 145–158.

McVicker, G., Gordon, D., Davis, C., & Green, P. (2009). Widespread genomic signatures of natural selection in hominid evolution. *PLoS genetics*, *5*(5), e1000471.

Milkman, R. (1978). Selection differentials and selection coefficients. *Genetics*, *88*(2), 391–403.

Murphy, D. A., Elyashiv, E., Amster, G., & Sella, G. (2022). Broad-scale variation in human genetic diversity levels is predicted by purifying selection on coding and non-coding elements. *Elife*, *12*, e76065.

Negm, S., & Veller, C. (2024). *The effect of long-range linkage disequilibrium on allele-frequency dynamics under stabilizing selection* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/2024.06.27.601075v1].

Nordborg, M., Charlesworth, B., & Charlesworth, D. (1996). The effect of recombination on background selection. *Genetics Research*, *67*(2), 159–174.

O'Connor, L. J., Schoech, A. P., Hormozdiari, F., Gazal, S., Patterson, N., & Price, A. L. (2019). Extreme polygenicity of complex traits is explained by negative selection. *The American Journal of Human Genetics*, *105*(3), 456–476.

Pardiñas, A. F., Holmans, P., Pocklington, A. J., Escott-Price, V., Ripke, S., Carrera, N., Legge, S. E., Bishop, S., Cameron, D., Hamshere, M. L., et al. (2018). Common schizophrenia alleles are enriched in mutation-intolerant genes and in regions under strong background selection. *Nature genetics*, *50*(3), 381–389.

Pickrell, J. K. (2014). Joint analysis of functional genomic data and genome-wide association studies of 18 human traits. *The American Journal of Human Genetics*, *94*(4), 559–573.

Robertson, A. (1956). The effect of selection against extreme deviants based on deviation or on homozygosis: With two text-figures. *Journal of Genetics*, *54*, 236–248.

Rockman, M. V., Skrovanek, S. S., & Kruglyak, L. (2010). Selection at linked sites shapes heritable phenotypic variation in c. elegans. *Science*, *330*(6002), 372–376.

Sanjak, J. S., Sidorenko, J., Robinson, M. R., Thornton, K. R., & Visscher, P. M. (2018). Evidence of directional and stabilizing selection in contemporary humans. *Proceedings of the National Academy of Sciences*, *115*(1), 151–156.

Sella, G., & Barton, N. H. (2019). 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.

Simons, Y. B., Bullaughey, K., Hudson, R. R., & Sella, G. (2018). A population genetic interpretation of gwas findings for human quantitative traits. *PLoS biology*, *16*(3), e2002985.

Simons, Y. B., Mostafavi, H., Smith, C. J., Pritchard, J. K., & Sella, G. (2022). *Simple scaling laws control the genetic architectures of human complex traits* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/2022.10.04.509926v1].

Smith, J. M. (1983). The genetics of stasis and punctuation. *Annual review of genetics*, *17*(1), 11–25.

Stephan, W., Charlesworth, B., & McVEAN, G. (1999). The effect of background selection at a single locus on weakly selected, partially linked variants. *Genetics Research*, *73*(2), 133–146.

Veller, C., & Coop, G. M. (2024). Interpreting population-and family-based genome-wide association studies in the presence of confounding. *PLoS Biology*, *22*(4), e3002511.

Weber, K. E. (1990). Selection on wing allometry in drosophila melanogaster. *Genetics*, *126*(4), 975–989.

Wright, S. (1938). The distribution of gene frequencies under irreversible mutation. *Proceedings of the National Academy of Sciences*, *24*(7), 253–259.

A. Derivation of Disease Prevalence Change under BGS

Let V_E environmental variance, and let V_G be the genetic variance in the absence of background selection. The heritability in the absence of background selection is then

$$h^2 = \frac{V_G}{V_G + V_E}, \tag{A.1}$$

and the standardized threshold density is

$$\phi(T^*) = f(T)\sqrt{V_p}, \tag{A.2}$$

where $V_P = V_G + V_E = \frac{V_G}{h^2}$ is the total variance.

In the presence of background selection, the genetic variance is reduced to $V_{G,B} = BV_G$, and the phenotypic variance to $V_{P,B} = V_{G,B} + V_E$, so the heritability becomes

$$h_B^2 = \frac{V_{G,B}}{V_{P,B}} = \frac{V_{G,B}}{V_{G,B} + V_E} = \frac{BV_G}{BV_G + V_E}, \tag{A.3}$$

and we can write the standardized threshold density in the presence of background selection in terms of original quantities in its absence as

$$\phi(T_B^*) = f_B(T)\sqrt{V_{P,B}} = \frac{f(T)}{B}\sqrt{BV_G + V_E}, \tag{A.4}$$

where $T_B^* = \frac{T - \bar{G}_B}{\sqrt{V_{P,B}}}$.

Now, we can eliminate the V_E by solving Equation (A.1) for $V_E = \frac{1-h^2}{h^2}V_G$ and substituting to get

$$\phi(T_B^*) = f_B(T)\sqrt{V_{P,B}} = \frac{f(T)}{B}\sqrt{V_G\left(\frac{1}{h^2} - (1-B)\right)}, \tag{A.5}$$

which is now written entirely in terms of counterfactual quantities from the equilibrium in the absence of background selection, and the B value. The increase in the standardized threshold density due to background selection is thus

$$\frac{\phi(T_B^*)}{\phi(T^*)} = \frac{\sqrt{1 - h^2(1-B)}}{B}, \tag{A.6}$$

i.e. Equation (25) in the main text.

B. Solving the Two Effect-size Model with Background Selection

Here, we detail the method for solving the two-effect-size liability threshold model in the presence of background selection. Berg et al. (2025) describes the method for solving this model in the absence of background selection in their Supplementary Section S7. Here, for completeness, we largely reproduced that description, with the addition of the BGS effect (as well as a few minor differences in notation).

In our two-effect model, we assume that a fraction $p_S = 1 - p_L$ of sites have small effects, while the remaining fraction p_L have large effects. We model the distribution of liability in the population as a convolution of two distributions. The first is a Normal component with variance $V_{G,B,S} + V_E$, where $V_{G,B,S}$ is the variance from small effect sites (in the presence of background selection), and V_E is the environmental variance. The second is a Poisson distribution on the number of large effect alleles that an individual carries. This Poisson distribution has mean

$$\lambda_L = \frac{2Lp_L\mu}{\delta(a_L)C}, \tag{B.1}$$

where

$$\delta(a_L) = F(T^* - a_L) - F(T^*) \tag{B.2}$$

is the risk effect of the large effect sites, $T^* = T - \bar{Z}$ is the distance between the mean and the threshold,

$$F(Z) = \sum_i P(i|\lambda_L)Q(Z|a_L(i - \lambda_L), V_{G,B,S} + V_E) \tag{B.3}$$

is the probability that an individual's liability exceeds a value of Z, where $Q(Z|u, \sigma^2) = 1 - \Phi(Z|u, \sigma^2)$ is the complementary CDF of a Normal distribution with mean u and variance σ^2, and $P(i|\lambda_L)$ is the probability that a Poisson random variable with mean λ_L takes a value of i. The density on total liability, in turn, is

$$f(Z) = \sum_i P(i|\lambda_L)\phi(Z|a_L(i - \lambda_L), V_{G,B,S} + V_E) \tag{B.4}$$

where $\phi(Z|u, \sigma^2)$ is the Normal PDF.

We solve two-effect model in the following way. First, we can solve for the fraction of sites fixed for the liability increasing allele at small effect sites. Because all large effect sites are fixed for the liability decreasing allele, this fraction is equal to

$$p^+(\gamma(a_S)) = \frac{\bar{a}p_T}{a_S p_S} \tag{B.5}$$

where

$$\bar{a} = a_S p_S + a_L p_L \tag{B.6}$$

is the mean effect size. The scaled selection coefficient of the small effect sites is determined entirely by the threshold induced fixation asymmetry and the resulting long-term fixation dynamics, and is therefore given by

$$\gamma(a_S) = \frac{1}{2} \ln \frac{1 - p^+(\gamma(a_S))}{p^+(\gamma(a_S))}, \tag{B.7}$$

and the threshold density by

$$f_B(T) = \frac{\gamma(a_S)}{2NBCa_S} = \frac{1}{4NBCa_S} \ln \frac{1 - p^+(\gamma(a_S))}{p^+(\gamma(a_S))}. \tag{B.8}$$

The genetic variance due to small effect sites is

$$V_{G,B,S} = \frac{\theta_L p_S B a_S b(\gamma(a_S))}{\gamma(a_S)}, \tag{B.9}$$

where $b(\gamma(a_S)) = 1 - 2p^+(\gamma(a_S))$ is the degree of mutational asymmetry at small effect sites, $\theta_L = 4NL\mu$ is the total population scaled mutation rate of the trait, and the B value is given by Equation (1) in the main text.

We can then solve for the value of $\delta(a_L)$ via a line search. First, we plug T^* into the right hand side of Equation (B.4) and set it equal to the right hand side of Equation (B.8), yielding

$$\frac{1}{4NBCa_S} \ln \frac{1 - p^+(\gamma(a_S))}{p^+(\gamma(a_S))} = \sum_i P(i|\lambda_L)\phi(T^*|a_L(i - \lambda_L), V_{G,B,S} + V_E). \tag{B.10}$$

Then, for a proposed value of $\delta(a_L)$, we compute λ_L using Equation (B.1), plug it into the RHS of Equation (B.10), and numerically solve for the value of T^* that satisfies Equation (B.10). We can then conduct a line search for the value of $\delta(a_L)$ that satisfies Equation (B.2), given the value of T^* obtained from solving Equation (B.10).

To constrain the liability-scaled heritability to a specific value, we add one additional step. Given the proposed value of $\delta(a_L)$, we compute the large effect contribution to the genetic variance as

$$V_{G,B,L} = a_L^2 \lambda_L \tag{B.11}$$

and then compute the environmental variance as

$$V_E = \frac{1 - h^2}{h^2}. \tag{B.12}$$

Note that the inclusion of the B value in Equations (B.8) and (B.9) to account for the impact of BGS on weakly selected sites represents the only substantial difference between the algorithm presented in Berg et al. (2025) and the one presented here. The solution in the case of no BGS is therefore obtained simply by setting $B = 1$. To compare cases with and without background selection,

we first solve the above system of equations for a fixed value of h^2 with $B = 1$, to obtain the solution in the absence of background selection. We then solve in the presence of background selection, holding the environmental variance (as opposed to the heritability), constant at the value we obtained in the no BGS case.

C. Hitting Probability at p=0.5 in the Symmetric Underdominant Model

Under background selection, alleles subject to underdominant selection can reach a higher average frequency and contribute more to heterozygosity than they would in the absence of background selection (see main text). The magnitude of this effect depends on both the strength of background selection and the intensity of underdominant selection. To illustrate this effect quantitatively, we consider the probability that a newly introduced allele reaches a frequency of $\frac{1}{2}$. Let $\lambda(q|p,\tau)$ denote the probability that an allele initially at frequency p reaches frequency q under the symmetric underdominant model, given a scaled selection coefficient τ and assuming $p < q$. Of particular interest is $\lambda(\frac{1}{2}|\frac{1}{2N},\tau)$, which represents the probability that a new allele, starting at frequency $\frac{1}{2N}$, eventually reaches $\frac{1}{2}$. We refer to this quantity as the hitting probability at frequency $\frac{1}{2}$.

We can take advantage of the symmetry of the dynamics in the symmetric underdominant model to derive this probability directly from the fixation probability. The fixation probability of a mutation at frequency p in this model is

$$v(\tau,p) = \frac{1}{2} - \frac{\mathrm{Erf}(\sqrt{\frac{\tau}{2}}(1-2p))}{\mathrm{Erf}(\sqrt{\frac{\tau}{2}})}$$

(Ewens, 2004).

The fixation probabilities for frequencies p and q are related via

$$v(\tau,p) = \lambda(q|p,\tau)v(\tau,q).$$

Notably, $v(\tau,\frac{1}{2}) = \frac{1}{2}$ for any value of τ, reflecting the symmetry of selection with respect to minor allele frequency in this model. It follows immediately that

$$\lambda\left(\frac{1}{2}\middle|p,\tau\right) = 2v(\tau,p).$$

It is easier to interpret our hitting probability of interest if we rescale it so that it is independent of population size. We can do this by first approximating $v(\tau,p)$ with a first order Taylor expansion,

which yields

$$v(\tau, p) \approx p \frac{e^{-\frac{\tau}{2}} \sqrt{\frac{2\tau}{\pi}}}{\mathrm{Erf}\left(\sqrt{\frac{\tau}{2}}\right)}$$

Second, we consider a "population scaled hitting rate", i.e. the number of alleles starting at frequency $p = \frac{1}{2N}$ and reaching frequency $q = \frac{1}{2}$, per $2N$ new mutations. This quantity is given by

$$\begin{aligned}\Gamma(\frac{1}{2}|\tau) = 2N\lambda\left(\frac{1}{2}|\frac{1}{2N}, \tau\right) &\approx 2N \frac{2}{2N} \frac{e^{-\frac{\tau}{2}} \sqrt{\frac{2\tau}{\pi}}}{\mathrm{Erf}\left(\sqrt{\frac{\tau}{2}}\right)} \\ &\approx \frac{2e^{-\frac{\tau}{2}} \sqrt{\frac{2\tau}{\pi}}}{\mathrm{Erf}\left(\sqrt{\frac{\tau}{2}}\right)} = 2\sqrt{\tau} \frac{\phi(\sqrt{\tau})}{2\Phi(\sqrt{\tau}) - 1},\end{aligned} \tag{C.1}$$

where ϕ and Φ are the standard normal PDF and CDF respectively.

In the neutral limit, we have $\lim_{\tau \to 0} \Gamma(\frac{1}{2}|\tau) = 2$. We approximate the population scaled hitting rate in the presence of background selection using the reduced effective size approximation,

$$\Gamma_B\left(\frac{1}{2}|\tau\right) \approx B\Gamma\left(\frac{1}{2}|\tau B\right) \approx 2B\sqrt{\tau B} \frac{\phi(\sqrt{\tau B})}{2\Phi(\sqrt{\tau B}) - 1}.$$

Thus, the change in the scaled hitting rate due to background selection is

$$\Gamma_B\left(\frac{1}{2}|\tau\right) - \Gamma\left(\frac{1}{2}|\tau B\right) \approx 2\sqrt{\tau}\left(B\sqrt{B} \frac{\phi(\sqrt{\tau B})}{2\Phi(\sqrt{\tau B}) - 1} - \frac{\phi(\sqrt{\tau})}{2\Phi(\sqrt{\tau}) - 1}\right) \tag{C.2}$$

(Figure S5).

D. An Incidental Finding That the Bulmer Effect Increases the Genetic Variance in the Long Term

In our multi-locus stabilizing selection simulations, when we used the same recombination rate between causal loci as used in the simulation of the liability threshold model (5×10^{-6}), the genetic diversity reduction of causal sites under background selection initially did not match our expectations based on our theory and our single-locus simulations (Figure S9). We hypothesized that this might be related to the Bulmer effect (Bulmer, 1976; Bulmer, 1971), i.e. the accumulation of negative linkage disequilibrium between causal sites due to selection. More precisely, the genetic variance can be composed into two components:

$$Var\left(\sum_j a_j g_j\right) = \sum_j a_j^2 Var(g_j) + \sum_j \sum_{k \neq j} a_j a_k Cov(g_j, g_k), \tag{D.1}$$

where the first term is the additive genic variance and the second term captures the contribution of linkage disequilibrium among causal sites. Under neutrality, the expectation of the LD term is equal to zero. However, because stabilizing selection acts most strongly against individuals carrying combinations of alleles that result in extreme phenotypes, it causes negative covariance among alleles of like sign, leading to a negative expectation for the contribution of the LD term. Thus, while stabilizing selection is acting, it leads to a reduction in the total genetic variance due to this negative LD. The reduction in genetic variance due to this negative LD is what is known as the "Bulmer effect". Because recombination will quickly eliminate this negative LD if stabilizing selection ceases, the Bulmer effect is generally understood as a "short term" phenomenon.

For any given focal allele, the negative LD induced by the Bulmer effect leads to an attenuation of the allele's marginal correlation with the phenotype, relative to its causal effect (Bulmer, 1971; Veller & Coop, 2024). In the modern statistical genetics literature, this effect is sometimes known as "linkage masking" (Brown et al., 2016). Because the expected change in the frequency of an allele from a given generation to the next depends on its marginal correlation with fitness in that generation, rather than its causal effect on fitness, the attenuation of the marginal phenotypic correlation in turn leads to an attenuation of the selection coefficient, and this in turn leads to a reduction in the magnitude of the expected change in frequency due to selection. More precisely, whereas in the absence of any Bulmer effect, the expected change in the minor allele frequency is given by:

$$\mathbb{E}(\Delta x) \approx -t(a)x(1-x)(1-2x), \tag{D.2}$$

with $t(a) = \frac{a^2}{2\omega^2}$, in its presence we expect that $t(a) < \frac{a^2}{2\omega^2}$, where the magnitude of the reduction in $t(a)$ depends on the amount of negative LD. This slowdown in the pace of frequency change is studied in greater depth by Negm and Veller (2024).

This effect has an intriguing implication. That is, while the negative LD induced by stabilizing selection causes a short-term reduction in the genetic variance relative to the genic variance, over long timescales, we would expect it to increase in the genic variance, due to the weaker selection on individual variants. Whether the net effect is to increase or decrease the genetic variance then depends on which effect is larger.

We hypothesized that this phenomenon may be responsible for the difference that we initially observed between our predictions and our simulation results. To investigate this hypothesis, we performed simulations under the stabilizing selection model without background selection and varied the rate of recombination between the neighboring causal sites (Figure S9). When the recombination rate is low, we find that the genetic variance is reduced relative to the genic variance, but is nevertheless increased relative to what we observed with higher recombination rates (which match the theoretical predictions based on unlinked theory; Simons et al. (2018)). Thus, our simulation results suggest that over long time scales, the net impact of the Bulmer effect is to increase the genetic variance, as the increase in the genic variance caused by the slowdown in the pace of allele frequency change

is larger than the reduction in variance due to the negative LD. Moreover, when we increased the recombination rate in our study of background selection to a level at which there was no longer any Bulmer effect or inflation of the genic variance, we found that our theoretical prediction for the impact of background selection then matched our theoretical predictions (Figure 4A), indicating that the Bulmer effect was indeed the culprit.

Our primary focus was on understanding the impact of background selection, so we did not study this effect any further. However, it is worth noting that, in addition to Negm and Veller (2024), at least two prior publications have also studied this effect. Specifically, Lande (1975) found that in his Gaussian allelic model, the decrease in genetic variance due to the Bulmer effect is exactly offset by the increase in genic variance, leading to no net effect. In contrast, Turelli and Barton (1990), employing their "rare alleles" approximation, do predict an increase in the genetic variance when recombination rates are low, but their predictions do not match our simulation results. A more complete analysis of this phenomenon would be valuable.

E. Heterozygosity Per Segregating Variant with Background Selection

For a single allele under directional selection with population scaled selection coefficient γ, we can measure the heterozygosity per segregating site (conditional on exceeding a minor allele frequency threshold p) as

$$\frac{\int_p^{1-p} q(1-q)\xi(q|\gamma)dq}{\int_p^{1-p} \xi(q|\gamma)dq}, \tag{E.1}$$

where $\xi(q|\gamma)$ is the sojourn time of the allele at frequency q. The numerator is the total heterozygosity, the same quantity measured in Equation (6), and the denominator quantifies the number of generation that the site is segregating. As we show in Figure 2, the site frequency with background selection is a transition between $\xi(q|\gamma)$ and $\xi_B(q|\gamma)$ depending on the frequency. In Figure S10, we plot the heterozygosity per segregating site at $\gamma = -1.5$ and $\gamma = -10$, for a range of minor allele frequency thresholds, using the site frequency spectra obtained from our simulations.

When sites are under weak negative selection (i.e. $\gamma = -1.5$, Figure S10A left), the heterozygosity reduction is between B and 1 (Figure 2). The reduction in the number of segregating alleles is slightly smaller than the reduction in heterozygosity because alleles with frequency $f < \frac{1}{2N\sigma}$ are not impacted by background selection. This leads to a reduction of heterozygosity per segregating variants. When we limit to more common variants (i.e. $p \approx 0.01$), these variants arise from backgrounds free of deleterious mutations. We observe a larger reduction in the number of segregating sites than a reduction in heterozygosity, which results in an increase of heterozygosity per segregating variant due to background selection. In other words, under background selection, common variants tend to

segregate with higher frequency in the population for a longer period of time due to an increased rate of drift. However, this is not true if we include rare variants.

When sites are under strong negative selection (Figure S10A right), background selection has a negligible impact on the heterozygosity (Figure 2). However, background selection can still remove new mutations on deleterious backgrounds. It results in a reduction in the number of segregating sites and an increase in heterozygosity per variant.

We performed the same analysis on a single locus under under-dominant selection. The heterozygosity per segregating variant is

$$\frac{\int_p^{\frac{1}{2}-p} q(1-q)\Psi(q|\tau)dq}{\int_p^{\frac{1}{2}-p} \Psi(q|\tau)dq}, \tag{E.2}$$

where we substitute the empirical site frequency spectrum for $\Psi(q|\tau)$. Similar pattern is observed for locus with under-dominant selection (Figure S10B).

F. Variation in B in the Threshold Model

In this section, we explore genomes subject to varying levels of BGS intensity under the threshold model, while assuming a constant effect size across all causal sites. Our primary focus is on cases where causal sites are effectively neutral or weakly selected, as fixation rates of alleles in this regime are highly sensitive to the intensity of BGS. Specifically, variations in BGS intensity across different genomic regions lead to substantial differences in the equilibrium state of fixed alleles between those regions. For example, more alleles will be fixed at a liability-increasing state in a stronger BGS region.

Despite these regional differences, the threshold model imposes a global constraint: the total fraction of sites fixed for liability-increasing alleles across the genome must remain approximately constant at $p_T = \frac{T}{2aL}$. To satisfy this constraint, the selection coefficient across all target sites adjusts by a constant factor. Here, we ask how this selection coefficient adjusts globally to maintain this genome-wide balance?

To address this, we define P_i as the proportion of target sites subject to a specific background selection intensity B_i. We write $\tilde{B}$ for the inverse of the factor by which the unscaled selection coefficients are increased due to the compensatory evolution of the liability distribution. Recalling from the main text that $p^+(\gamma) = \frac{1}{1+e^{2\gamma}}$ gives the fraction of sites with scaled coefficient γ that are fixed for the liability increasing allele, a solution for $\tilde{B}$ can be found by solving:

$$\sum_i P_i \frac{1}{1+e^{2\gamma\frac{B_i}{\tilde{B}}}} = p_T. \tag{F.1}$$

We can understand the implications of Equation (F.1) results by considering a simple model with two regions, each containing an equal proportion of target sites ($P_1 = P_2 = \frac{1}{2}$), under the assumption that sites in region 2 experience no background selection (i.e. $B_1 < 1$ and $B_2 = 1$).

We consider two extremes, starting with the case where $p_T \approx \frac{1}{2}$, so that selection is very weak (i.e. $\gamma \ll 1$). In this case, we can make a first-order approximation of $p^+(\gamma)$ around the point $\gamma = 0$ yielding $p^+(\gamma) \approx \frac{1}{2} - \frac{\gamma}{2}$. In this case, Equation (F.1) becomes

$$\frac{1}{2} \cdot (\frac{1}{2} - \frac{\gamma}{2}\frac{B_1}{\tilde{B}}) + \frac{1}{2} \cdot (\frac{1}{2} - \frac{\gamma}{2}\frac{1}{\tilde{B}}) = p_T. \tag{F.2}$$

Recalling that in the absence of BGS, $\gamma \approx \frac{1}{2} \ln \frac{1-p_T}{p_T}$ (Equation (20)), we can make a first-order approximation of γ around the point $p_T \approx \frac{1}{2}$, which yields $\gamma \approx 1 - 2p_T$. Plugging this into Equation (F.2) allows us to solve explicitly for $\tilde{B}$, which gives

$$\tilde{B} \approx \frac{B_1 + 1}{2}, \tag{F.3}$$

(for arbitrary B_2, $\tilde{B} \approx \frac{B_1+B_2}{2}$). That is, $\tilde{B}$ is approximately equal to the arithmetic mean of the BGS intensity across regions. Intuitively, the fraction of sites fixed for the liability-increasing allele increases in the region with stronger background selection, and decreases by the same amount in the region with weaker background selection, so that the total fraction remains fixed. It follows from the linear relationship between $p^+(\gamma)$ and γ in this regime that $\tilde{B}$ must be equal to the average value of B. More generally, if we imagine a model in which B varies broadly across the genome, and $p_T \approx \frac{1}{2}$, we would expect that $\tilde{B}$ should equal to the arithmetic mean of B across causal sites.

In contrast, if we consider the limit in which $p_T \to 0$, so that selection is stronger (e.g. $\gamma \approx 5$), nearly all sites must already be fixed for the liability-decreasing allele. In this regime, $p^+(\gamma)$ is relatively insensitive to increases in the scaled selection coefficient, because most sites are already fixed for the liability-decreasing allele. Thus, the most that can be accomplished by increasing the scaled selection coefficient from γ to $\frac{\gamma}{\tilde{B}}$ in region 2 (i.e. the one not experiencing the direct effects of background selection) is to switch a few sites that are fixed for the liability-increasing allele to be fixed for the liability-decreasing allele. Thus, we can assume that for region two, $p^+(\frac{\gamma}{\tilde{B}}) \approx Fp_T$, where $0 < F < 1$. As a result, the number of sites in region one that must be fixed for the liability-increasing allele must increase by a factor of $2(1 - \frac{F}{2})$, i.e. $p^+(\gamma\frac{B_1}{\tilde{B}}) \approx 2p_T(1 - \frac{F}{2})$. Plugging in the definition of $p^+(\gamma\frac{B}{\tilde{B}})$ gives

$$\frac{1}{1 + e^{2\gamma\frac{B_1}{\tilde{B}}}} \approx 2p_T(1 - \frac{F}{2}), \tag{F.4}$$

which we can solve for $\tilde{B}$ to obtain

$$\tilde{B} \approx B_1 \frac{2\gamma}{\ln \frac{1-2p_T(1-\frac{F}{2})}{2p_T(1-\frac{F}{2})}}. \tag{F.5}$$

Recognizing that γ represents the scaled selection coefficient that we would obtain in the case without

background selection (i.e. $\gamma \approx \frac{1}{2} \ln \frac{1-p_T}{p_T}$, Equation (20)), we have

$$\tilde{B} \approx B_1 \lim_{p_T \to 0} \left[\frac{\ln \frac{1-p_T}{p_T}}{\ln \frac{1-2p_T(1-\frac{F}{2})}{2p_T(1-\frac{F}{2})}} \right] = B_1. \tag{F.6}$$

That is, when the threshold is close to zero, $\tilde{B}$ is approximately equal to the BGS intensity in region 1 (B_1), indicating that regions with higher BGS intensity dominate the global effect. Intuitively: to the extent that the increase in the scaled selection coefficient from γ to $\frac{\gamma}{\tilde{B}}$ for region 2 results in more of those sites being fixed for the liability decreasing allele, this is offset by an increase in the number of sites in region 1 that are fixed for the liability increasing allele. However, this increase in the fixed fraction for region 1 is accomplished by an essentially negligible decrease in the scaled selection coefficients. It follows that $\tilde{B}$ must be approximately equal to B_1. More generally, if we imagine a model in which B varies broadly across the genome and $p_T \to 0$, we would expect $\tilde{B}$ to be close to the minimum of B across causal sites.

We now consider how these dynamics affect the contribution to genetic variance explained by regions with different BGS intensities. Following Equation (12) from the main text and considering both local and global effects, the genetic variance contributed by each region can be expressed as:

$$V_i = \frac{\theta_i B_i}{\gamma \frac{B_i}{\tilde{B}}} b(\gamma \frac{B_i}{\tilde{B}}) = \frac{\theta_i \tilde{B}}{\gamma} b(\gamma \frac{B_i}{\tilde{B}}), \tag{F.7}$$

where $b(\gamma) = 1 - 2p^+(\gamma) = \frac{e^{2\gamma}-1}{e^{2\gamma}+1}$ measures the asymmetry of the mutational input at equilibrium with scaled coefficient γ, and $\theta_i = 4NLP_i\mu$ is the population scaled mutation rate for region i. For our simplified two region scenario with $B_2 = 1$ and $P_1 = P_2 = \frac{1}{2}$, the relative genetic variance contribution from region 1 is:

$$\frac{V_1}{V_1 + V_2} = \frac{P_1 \cdot \frac{\tilde{B}}{\gamma} \cdot b(\gamma \frac{B_1}{\tilde{B}})}{P_1 \cdot \frac{\tilde{B}}{\gamma} \cdot b(\gamma \frac{B_1}{\tilde{B}}) + P_1 \cdot \frac{\tilde{B}}{\gamma} \cdot b(\gamma \frac{1}{\tilde{B}})} = \frac{b(\gamma \frac{B_1}{\tilde{B}})}{b(\gamma \frac{B_1}{\tilde{B}}) + b(\gamma \frac{1}{\tilde{B}})}. \tag{F.8}$$

When $p_T \approx \frac{1}{2}$, so that $\gamma \ll 1$, Equation (F.8) reduces to

$$\frac{V_1}{V_1 + V_2} = \frac{B_1}{B_1 + 1}, \tag{F.9}$$

i.e. the relative contribution to variance from the region experiencing background selection is reduced in proportion to the strength of the local background selection effect. Intuitively, selection does not have to work very hard to maintain the fixation asymmetry, so these dynamics do not factor into the relative reductions in variance, and we obtain the same result that we would obtain if the trait were

completely neutral. In contrast, in the limit as $p_T \to 0$, Equation (F.8) becomes

$$\frac{V_1}{V_1 + V_2} = \frac{1}{2}, \tag{F.10}$$

i.e. background selection has no impact on the relative contribution to variance across regions. This is because the contributions to variance in the two different regions are both reduced by the same amount, although for slightly different reasons. In region 1, the direct, local impact of background selection effectively reduces the number of alleles entering the population by a factor of $B_1 \approx \tilde{B}$, but the effective scaled selection coefficient is unaltered, as $B_1 \approx \tilde{B}$ implies that $\gamma \frac{B_1}{\tilde{B}} \approx \gamma$ for region 1, resulting in a total variance reduction of $B_1 \approx \tilde{B}$. In region two, there is no local impact on the effective number of mutations entering the population, but the scaled selection coefficient is increased from γ to $\frac{\gamma}{\tilde{B}}$, which results in a reduction of the variance by a factor of $\tilde{B}$, yielding Equation (F.10).

In Figure S11, we plot numerical solutions for $\tilde{B}$ and $\frac{V_1}{V_1+V_2}$ as functions of B_1, for different values of γ. These plots confirm the limiting results given above for $p_T \approx \frac{1}{2}$ and $p_T \to 0$, and indicate a smooth transition between them for intermediate values. In Figure S12, we compare the numerical solutions for $\frac{V_1}{V_1+V_2}$ to simulation results with $B_1 \approx 0.82$, finding good agreement.

Supplementary Figures

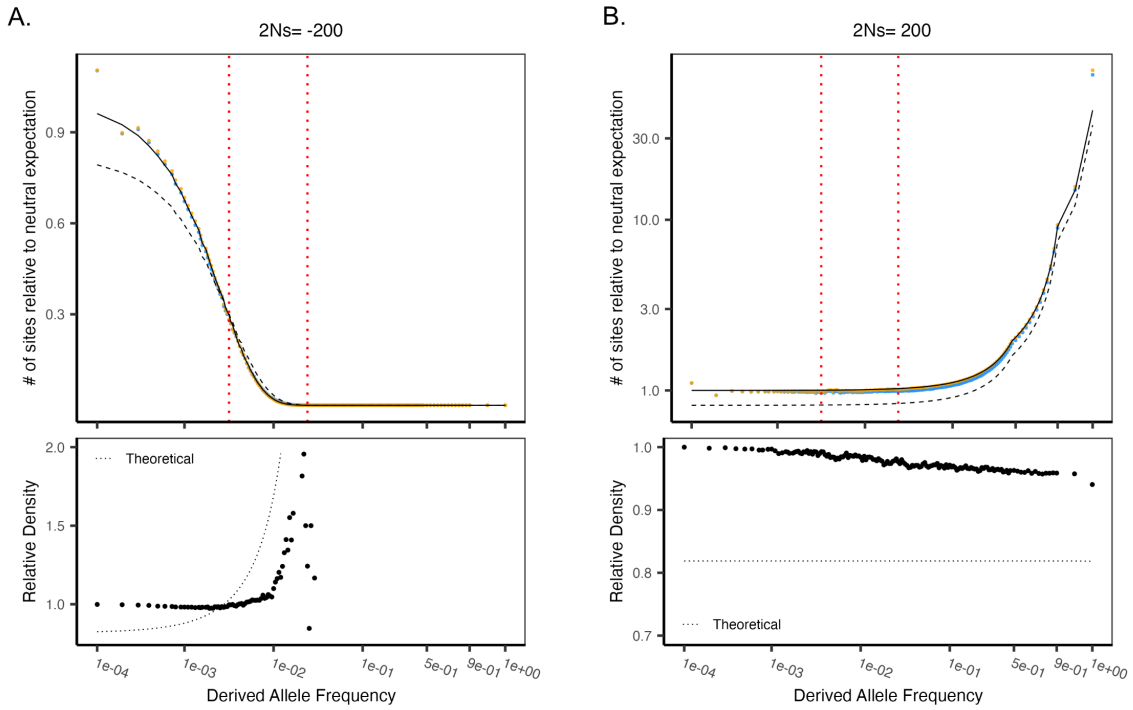

Figure S1: Site frequency spectrum of focal additively selected allele at (A) $2Ns = -200$, and (B) $2Ns = 200$. In the upper panel, we plot the site frequency spectrum of the focal additively selected allele, scaled relative to its neutral expectation, with a background selection intensity $B \approx 0.82$. Color represents the scenario whether there is BGS or not. The solid line represents the predicted deleterious SFS with effective population size N; while the dashed line represents the predicted SFS with effective population size NB. In the lower panel, we plot the expected and empirical ratio of the sojourn time that the focal allele spends at different frequencies with BGS compared to no BGS.

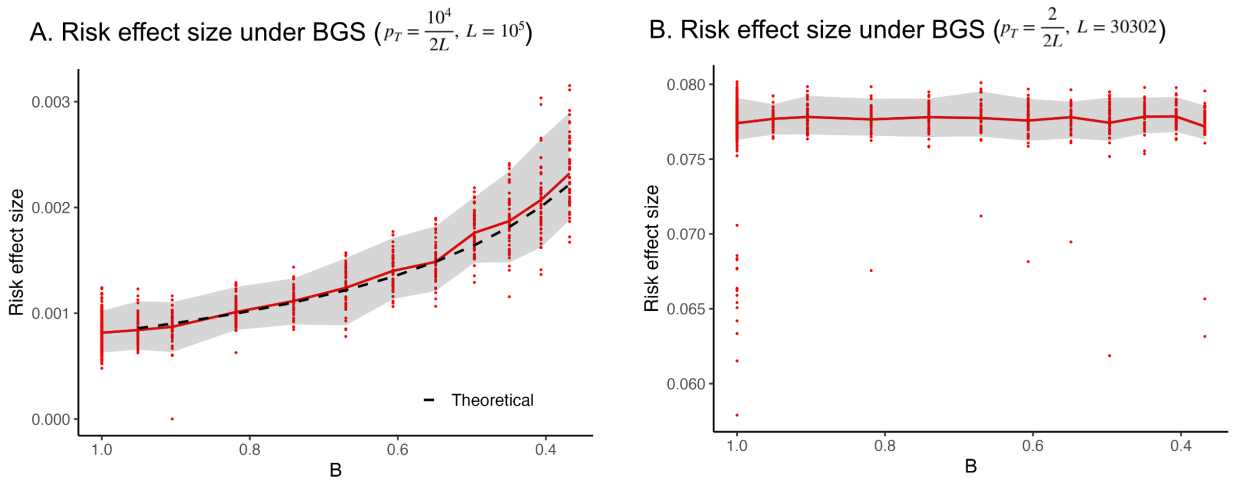

Figure S2: The impact of BGS on risk effect size in the single-effect liability threshold model. (A) When $p_T \gg \frac{1}{2L}$, the risk effect size increases with increasing strength of BGS. The black dashed line shows the theoretical prediction of Equation (23). (B) When $p_T \sim \frac{1}{2L}$, the risk effect size does not change under BGS.

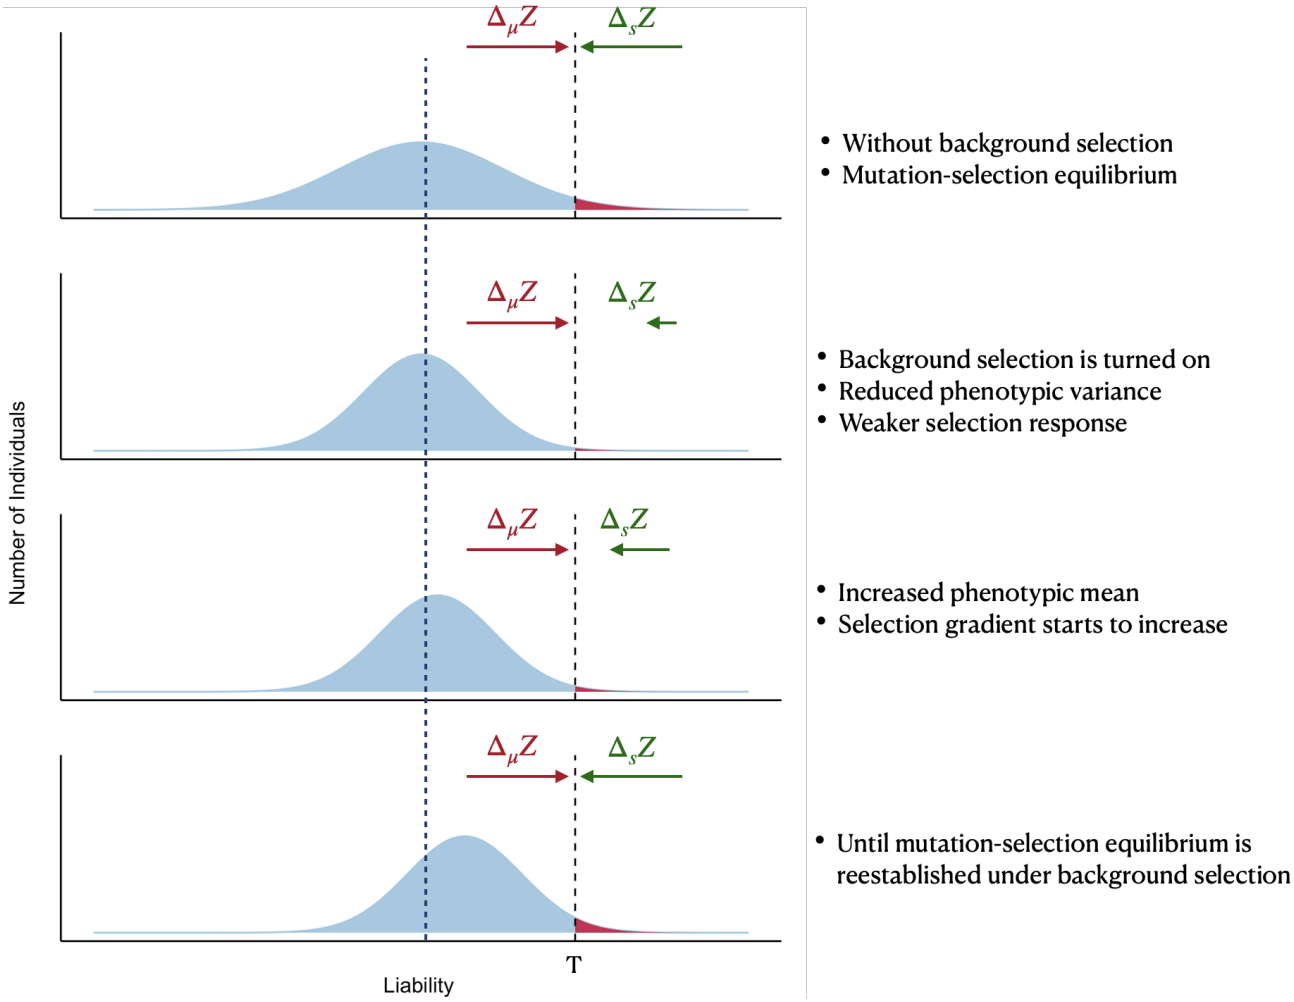

Figure S3: A cartoon illustration that explains how background selection impacts disease prevalence. This illustration imagines a dynamic process where the population is first at equilibrium in the absence of background selection, before it is "turned on", and the population evolves to the new equilibrium. In reality, top and bottom panels should be viewed as counterfactual comparison in a model without vs with background selection.

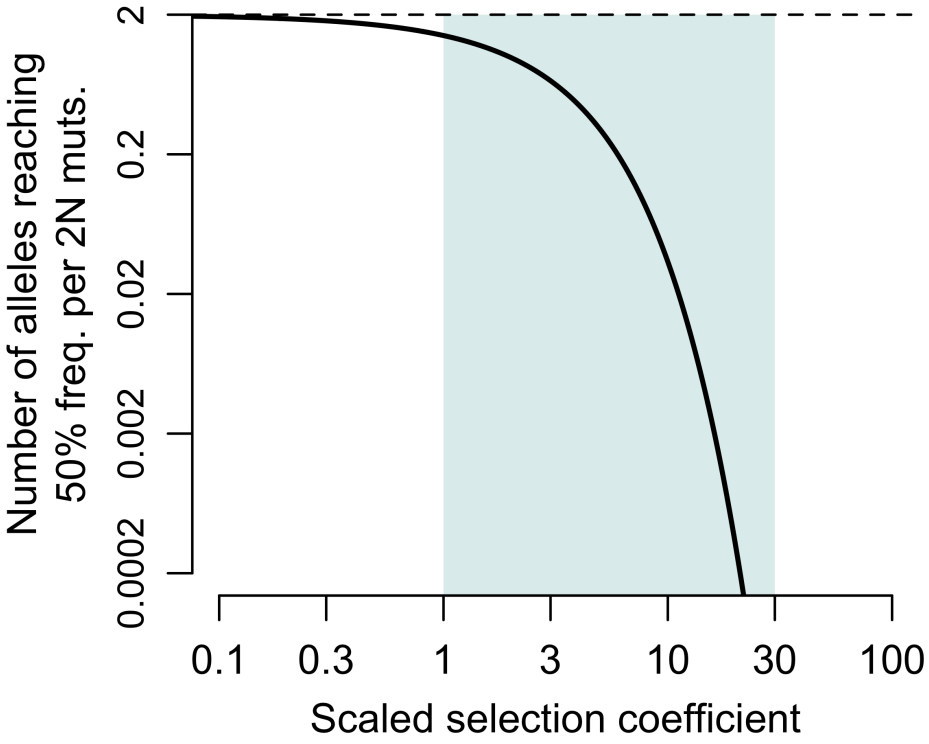

Figure S4: The number of alleles reaching frequency $\frac{1}{2}$ per $2N$ mutations at different scaled selection coefficients with underdominant selection, predicted by Equation (C.1).

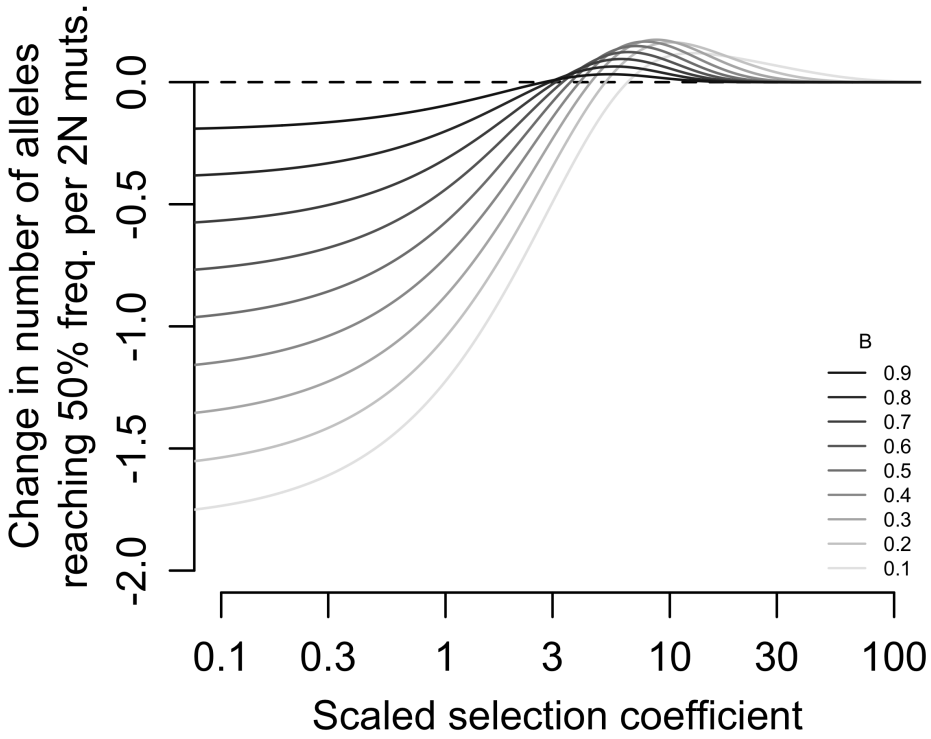

Figure S5: Change in number of alleles that reach frequency 0.5 per $2N$ mutations when BGS is present compared to no BGS (Equation (C.2)). The colors represent different background selection intensities from $B = 0.1$ to $B = 0.9$.

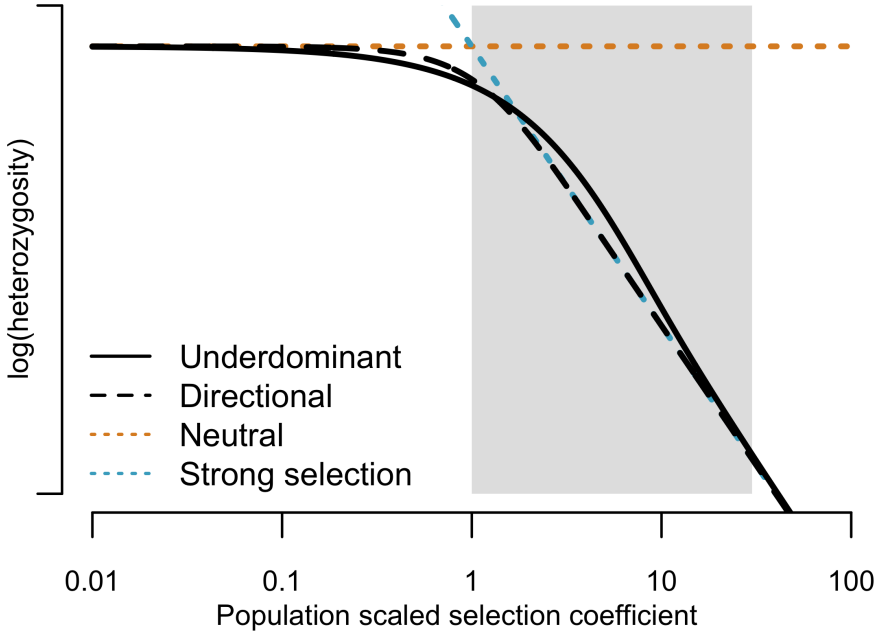

Figure S6: A comparison of expected heterozygosity for alleles with directional selection (Equation (3), dashed black line) vs. under dominant selection (Equation (30), solid black line) at different scaled selection coefficients. The expected neutral heterozygosity is θ, and the expected heterozygosity for strong selection is $\frac{\theta}{\gamma}$.

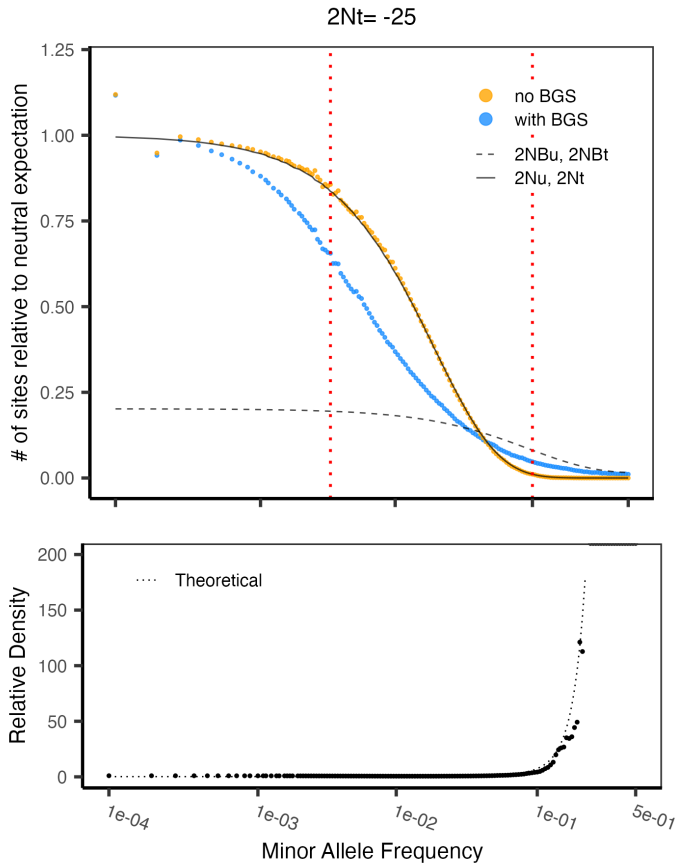

Figure S7: Site frequency spectrum (upper panel) and relative density (lower panel) of a selected allele with under-dominant selection $2Nt = -25$, with a background selection intensity $B \approx 0.2$. The plotting scheme is the same as Figure S1.

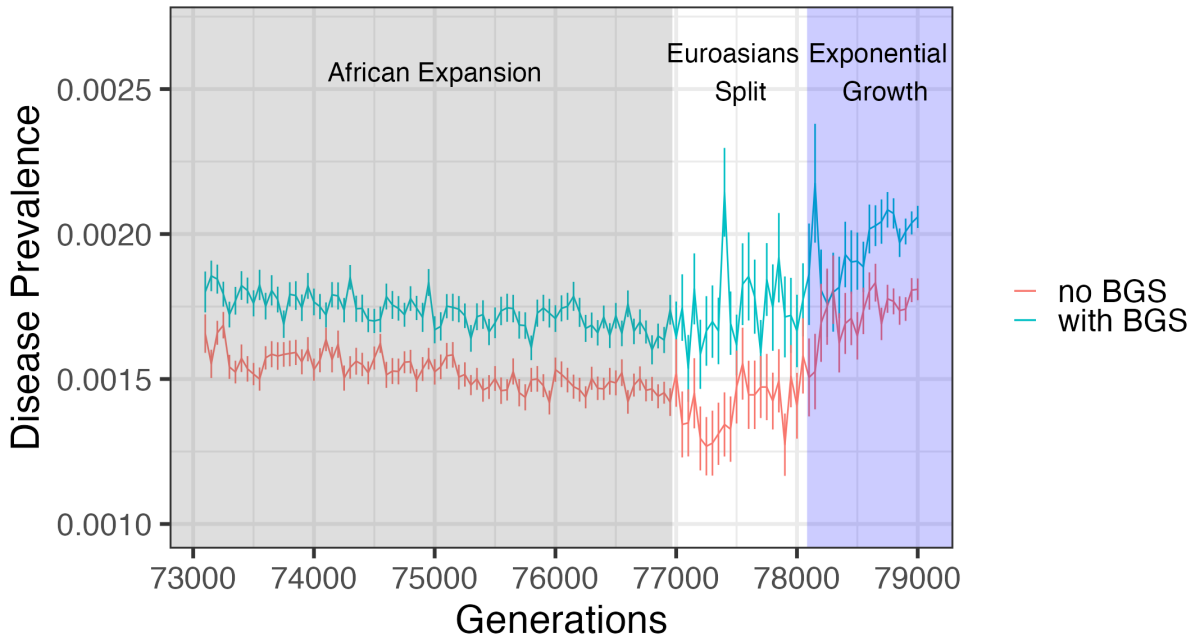

Figure S8: We simulated a "CEU-like" demographic history with the out-of-Africa demographic history (Gravel et al., 2011) with the strength of background selection set to reduce neutral variance by an average 17%, as estimated in humans (Murphy et al., 2022). We simulated a disease under the liability threshold model with heritability $h^2 = 0.5$, and a mutational asymmetry $p^+(a) = 0.05$. We reported the disease prevalence of the African population before the Euroasians split and of the European population after the split. We compared the disease prevalence with no background selection (red) and with background selection (green) intensity $B \approx 0.83$ (average estimated in humans). The error bar represents the standard error of the mean among 100 replications. In these simulations, the prevalence of a polygenic disease with heritability 0.5 and weakly selected causal loci in the presence day (generation 79000) in the absence of background selection was 0.18% (95% CI: 0.176% - 0.191%), and its presence 0.21% (95% CI: 0.203% - 0.219%), an average increase of 14.9%.

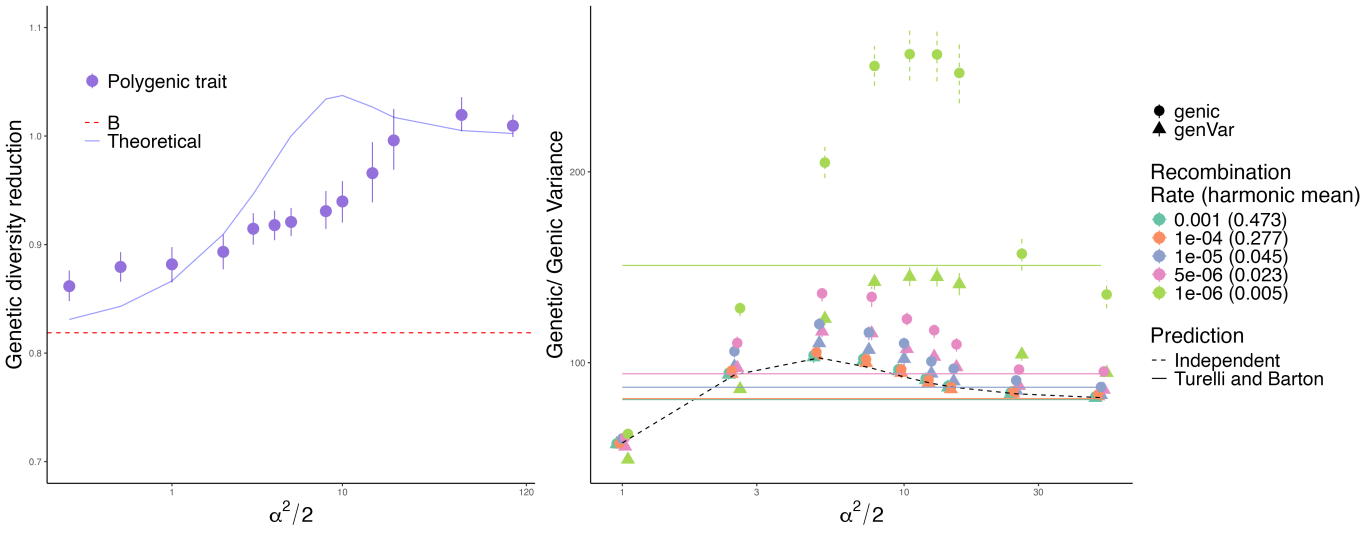

Figure S9: Left panel: we simulated polygenic traits evolving under stabilizing selection and set the recombination rate between causal loci at 5×10^{-6} and mutation rate at 2×10^{-8}. The genetic variance reduction of the trait under background selection deviates from the single-site prediction. Right panel: we simulated traits under stabilizing selection without background selection and varied the recombination rates between causal loci. The mutation rate is fixed at 2×10^{-8}. We plot the genic variance (circle) and genetic variance (triangle) and compare them to the predicted genetic variance that assumes independent causal loci (dashed line), and to the "rare alleles" approximation from Turelli and Barton, 1990.

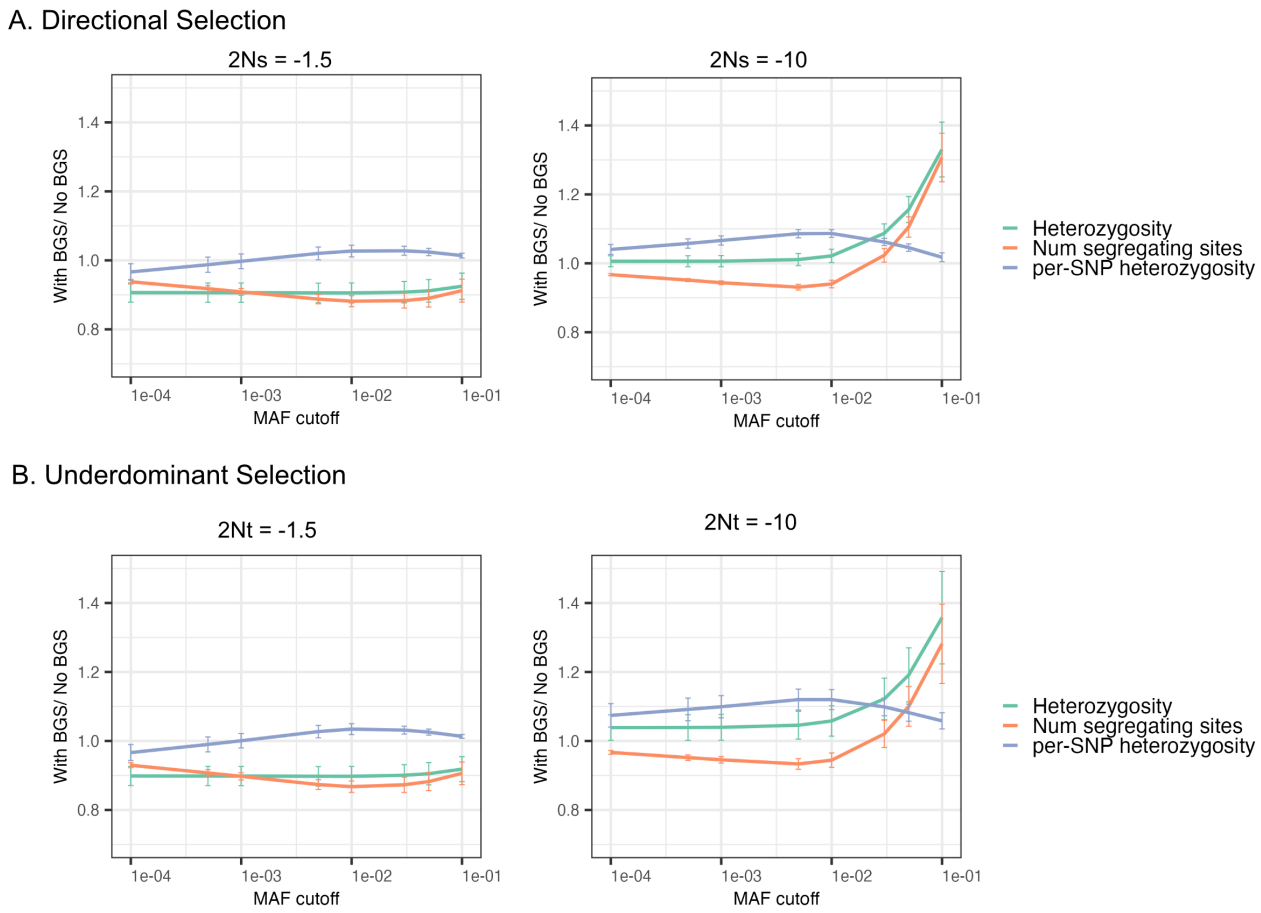

Figure S10: From the site frequency spectrum of alleles under (A) additive selection and (B) underdominant selection, we measured the relative total heterozygosity (red), the relative number of segregating sites (green), and the relative per-SNP heterozygosity (blue) with background selection compared to no background selection using different minor allele frequency cutoffs. The left panels show weakly selected alleles, and the right panels show strongly selected alleles.

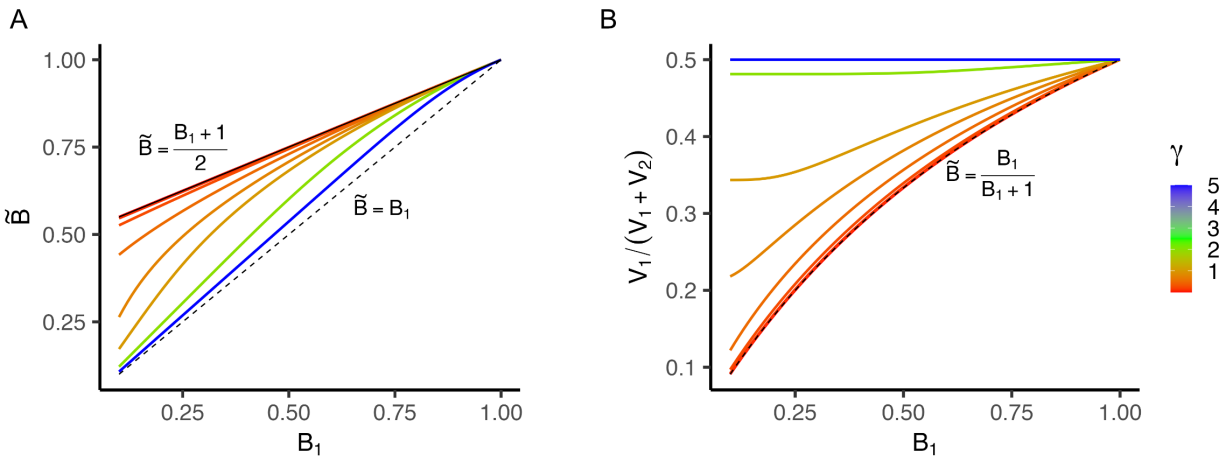

Figure S11: Background selection effect in a scenario where one chromosome is devoid of BGS ($B_2 = 1$) and the other chromosome experiences varying BGS intensities (B_1). (A) We solve for the global background selection effect ($\tilde{B}$) numerically using Equation (F.1). The solid black line represents the average of B_1 and B_2, while the dashed black line indicates the value of B_1. (B) The relative genetic variance contribution from chromosome one is calculated with Equation (F.7).

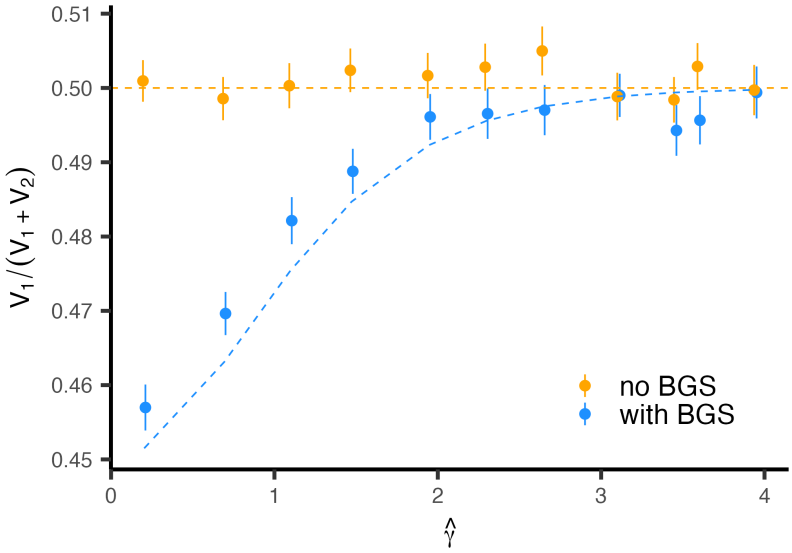

Figure S12: We simulated traits using the liability threshold model under varying thresholds (corresponding to different γ on the x-axis), with causal sites evenly distributed across two chromosomes. We reported the genetic variance ratio contributed by chromosome one in the presence and absence of background selection. When background selection is applied, chromosome one experiences a background selection intensity of $B \approx 0.82$, and chromosome two experiences no background selection. The theoretical expectations, shown as dashed lines, are calculated using Equation (F.7).

Supplementary References

Berg, J. J., Li, X., Riall, K., Hayward, L., & Sella, G. (2025). *Mutation-selection-drift balance models of complex diseases* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/2025.05.18.654722v1].

Brown, B. C., Price, A. L., Patsopoulos, N. A., & Zaitlen, N. (2016). Local Joint Testing Improves Power and Identifies Hidden Heritability in Association Studies. *Genetics*, *203*(3), 1105–1116. https://doi.org/10.1534/genetics.116.188292

Bulmer, M. (1976). The effect of selection on genetic variability: A simulation study. *Genetics Research*, *28*(2), 101–117.

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

Ewens, W. J. (2004). *Mathematical population genetics: Theoretical introduction* (Vol. 27). Springer.

Gravel, S., Henn, B. M., Gutenkunst, R. N., Indap, A. R., Marth, G. T., Clark, A. G., Yu, F., Gibbs, R. A., Project, 1. G., Bustamante, C. D., et al. (2011). Demographic history and rare allele sharing among human populations. *Proceedings of the National Academy of Sciences*, *108*(29), 11983–11988.

Lande, R. (1975). The maintenance of genetic variability by mutation in a polygenic character with linked loci. *Genetics Research*, *26*(3), 221–235.

Murphy, D. A., Elyashiv, E., Amster, G., & Sella, G. (2022). Broad-scale variation in human genetic diversity levels is predicted by purifying selection on coding and non-coding elements. *Elife*, *12*, e76065.

Negm, S., & Veller, C. (2024). *The effect of long-range linkage disequilibrium on allele-frequency dynamics under stabilizing selection* [Preprint at bioRxiv https://www.biorxiv.org/content/10.1101/2024.06.27.601075v1].

Simons, Y. B., Bullaughey, K., Hudson, R. R., & Sella, G. (2018). A population genetic interpretation of gwas findings for human quantitative traits. *PLoS biology*, *16*(3), e2002985.

Turelli, M., & Barton, N. H. (1990). Dynamics of polygenic characters under selection. *Theoretical Population Biology*, *38*(1), 1–57.

Veller, C., & Coop, G. M. (2024). Interpreting population-and family-based genome-wide association studies in the presence of confounding. *PLoS Biology*, *22*(4), e3002511.