A model for background selection in non-equilibrium populations

Gustavo V. Barroso[1,*] and Aaron P. Ragsdale[1,*]

[1]Department of Integrative Biology, University of Wisconsin-Madison, USA, 53706
[*]gvbarroso@gmail.com, apragsdale@wisc.edu

September 2, 2025

Abstract

In many taxa, levels of genetic diversity are observed to vary along their genome. The framework of background selection models this variation in terms of linkage to constrained sites, and recent applications have been able to explain a large portion of the variation in human genomes. However, these studies have also yielded conflicting results, stemming from two key limitations. First, existing models are inaccurate in a critical region of parameter space ($N_e s \sim -1$), where the local reduction in diversity is sharpest. Second, they assume a constant population size over time. Here, we develop predictions for diversity under background selection based on the Hill-Robertson system of two-locus statistics, which allows for population size changes. We treat the joint effect of multiple selected loci independently, but we show that interference among them is well captured through local rescaling of mutation, recombination and selection in an iterative procedure that converges quickly. We further accommodate existing background selection theory to non-equilibrium demography, bridging the gap between weak and strong selection. Simulations show that our predictions are accurate across the entire range of selection coefficients. We characterize the temporal dynamics of linked selection under population size changes and demonstrate that patterns of diversity can be misinterpreted by other models. Specifically, biases due to the incorrect assumption of equilibrium carry over to downstream inferences of the distribution of fitness effects and deleterious mutation rate. Jointly modeling demography and linked selection therefore improves our understanding of the genomic landscape of diversity, which will help refine inferences of linked selection in humans and other species.

Introduction

Patterns of genetic variation reflect ancestral events such as historical population size changes, population splits, migrations, and episodes of natural selection. Whereas demographic processes affect the whole genome, selection operates on functionally constrained sites, sorting polymorphisms based on their phenotypic effect. The interplay between neutral stochastic processes and the deterministic force of negative selection is central to evolutionary biology (Kimura and Ohta, 1971; Cutter and Payseur, 2013; Pouyet and Gilbert, 2021; Barroso and Dutheil, 2023). Early work in population genetics theory has characterized how single-site summary statistics behave under different demographic histories and degrees of direct negative selection (Ewens, 1964; Sawyer and Hartl, 1992; Marth et al., 2011).

The effects of negative selection extend beyond constrained sites. This is because physical linkage connects the evolutionary histories of sites along the same chromosome, such that linked loci tend to share coalescence events en route to their most recent common ancestors (Hudson and Kaplan, 1988, 1995b). On the other hand, recombination weakens these correlations by separating lines of descent at crossing-over breakpoints during meiosis. Consequently, since the rate of recombination between loci increases with their physical distance, nearby sites share a larger fraction of their histories than sites that are farther apart (Hill and Robertson, 1966, 1968; Ohta and Kimura, 1971; Cutter and Payseur, 2013). The distortion in genealogies of neutral sites caused by the presence of negatively constrained loci is termed background selection and

has received much attention in the past three decades (e.g., Charlesworth et al., 1993; Hudson and Kaplan, 1995b; Nordborg et al., 1996; McVicker et al., 2009; Good et al., 2014; Cvijović et al., 2018; Elyashiv et al., 2016; Ewing and Jensen, 2016; Torres et al., 2020; Murphy et al., 2022; Buffalo and Kern, 2024). In genomic data, a well-documented signature of background selection is the spatial covariance between genetic diversity and features like gene density and recombination rate (Aguade et al., 1989; Begun and Aquadro, 1992; Andolfatto and Przeworski, 2001; Lohmueller et al., 2011; Castellano et al., 2018). Accurately characterizing the evolutionary mechanisms that shape these patterns is key to establishing a proper baseline model of genome evolution (Comeron, 2017; Johri et al., 2020).

The most well-studied measure of background selection is the reduction in pairwise nucleotide diversity, π. Classic background selection theory (cBGS) describes how the ratio π/π_0 (denoted B, with π_0 being the expected diversity in the absence of linked selection) is affected by the continual influx and removal of strongly deleterious variants (Charlesworth et al., 1993; Hudson and Kaplan, 1995a; McVicker et al., 2009). The assumption of mutation-selection balance employed by cBGS offers elegant formulas which incorporate multiple constrained sites through simple multiplication. At the time, the approximations of cBGS were justified by the prevailing view that the strength of selection in nature should be quite high, with selection coefficients against heterozygous genotypes on the order of -0.01 and stronger. Under this regime, selection is strong enough that its effect can indeed be treated deterministically (Haldane, 1927), even in relatively small populations, and equilibrium diversity at constrained sites is low enough that they do not substantially interfere with each other, even at short genetic distances (Hill and Robertson, 1966). Thirty years on, however, investigation of genomic data from several species has revealed wide distributions of fitness effects, with mean selection coefficients on the order of -0.001 and typically long tails spanning from moderately selected to nearly-neutral mutations (Eyre-Walker et al., 2006; Eyre-Walker and Keightley, 2007; Huber et al., 2017). Furthermore, other empirical work has found evidence of pervasive interference among selected sites (Castellano et al., 2016).

Yet predictions of B-value maps (B-maps) in humans and fruit flies have relied on the analytical approximations of cBGS (McVicker et al., 2009; Elyashiv et al., 2016; Murphy et al., 2022). Only recently has a different theoretical framework (Santiago and Caballero, 1995, 1998, 2016) been translated into the computational machinery required for statistical inference (Buffalo and Kern, 2024). Unlike cBGS, the approach from Santiago and Caballero (2016) (SC16) is a quantitative model based on the variance in fitness within a population, which is a function of the influx of deleterious variation and rate of fixation of slightly deleterious mutations (i.e., Muller's Ratchet, Muller, 1964). In practice, this has produced accurate predictions of B-values in the weak selection regime (Buffalo and Kern, 2024). However, the model (as currently implemented) breaks down for scaled selection coefficients ($N_e s$) in the broad vicinity of -1. This is a particularly compelling region of parameter space, where interference among constrained sites is highest and the local reduction in diversity is strongest (McVean and Charlesworth, 2000; Comeron and Kreitman, 2002; Good et al., 2014; Ragsdale, 2022). More realistic models of linked selection rely on accurate prediction of B-values in this regime.

Current models have other limitations besides the range of selection regimes where they apply. A critical shortcoming shared by cBGS and SC16 theory is their restriction to equilibrium demography. Fluctuating population sizes impact the segregation trajectories of deleterious variants, which in turn influence diversity at linked neutral sites. Predicting these dynamics is challenging because the drift-effective population size (N_e) simultaneously dictates the overall level of (deleterious) genetic variation maintained in the population ($\mathcal{O}(N_e\mu)$), the relative strength of selection versus drift ($\mathcal{O}(N_e s)$), and the time available for recombination to break down correlations among genealogies ($\mathcal{O}(N_e r)$). For example, deleterious variants subject to efficient selection in a large ancestral population may transition to the "interference" regime ($N_e s \sim -1$) after a sharp bottleneck, where drift plays a pivotal role. Meanwhile, concerted decreases in $N_e\mu$ and $N_e r$ will lead to lower deleterious diversity and tighter linkage, respectively, with opposing effects on the extent of linked selection. What, then, is the net outcome of these perturbed forces on B-values?

The distinct relationships of B against $N_e s$, $N_e r$ or $N_e\mu$ would be enough to challenge intuition, but the situation is further complicated because the drift-effective population size, as recorded in genetic diversity, represents the harmonic mean between the sizes in consecutive epochs (Charlesworth, 2009). Therefore, after

a sudden demographic shift and until the new equilibrium is reached, N_e itself changes gradually over time, reflecting the proportion of uncoalesced ancestry within each epoch, and with it the compound parameters $N_e s$, $N_e r$ and $N_e \mu$. In general, if the time-scale of consecutive demographic changes is short enough to prevent the population from reaching stationarity, as is the case in natural populations, then patterns of diversity observed at any time should arise from a complex combination of past dynamics and not simply their average. While previous work used non-equilibrium simulations to describe the temporal dynamics of B-values (Torres et al., 2020), incorporating them into theoretical models is important for the development of inference tools that can be applied to data.

In this article we introduce a new model (`moments++`) that can accurately predict B-values under a wide range of evolutionary scenarios, from weak to strong selection, passing through the interference selection regime, and under non-equilibrium single-population demography. Our method is based on the Hill and Robertson (1968) system of two-locus statistics that has recently been generalized (Ragsdale and Gravel, 2019). We incorporate selection into the two-locus system and model its effect on a neutral locus located at an arbitrary recombination distance. Inspired by the pioneering work of Hudson and Kaplan (1995a,b), we treat the joint effect of multiple selected loci independently, but we show that under equilibrium, interference among them is well captured through iterative rescaling of mutation, recombination and selection. We explore the dynamics of linked selection under non-equilibrium demography and discuss how patterns of diversity can be misinterpreted by other models.

Results and Discussion

Predicting diversity under linked selection in a two-locus framework

Our objective is to model the effect of linked selection on the expected diversity of a focal site. We begin our approach with a pure two-locus model, where a single constrained site stands a recombination distance r away from the focal neutral site (Appendix A, Figure 1A). Although B-values are agnostic to possible selection on the focal site, we develop the two-locus model under the assumption that the focal site evolves neutrally to avoid formalizing its reciprocal linked selection effect. Later, using a heuristic correction for multi-locus interference, we will relax this assumption and predict diversity for neutral and constrained sites across the genome.

Negative selection in the constrained locus interferes with the expected sojourn time at the focal neutral locus. Briefly, the presence of a deleterious variant introduces fitness variance in the population, such that neutral polymorphisms associated with the deleterious background have increased extinction probability whereas those associated with the ancestral background have increased chance of fixation. The magnitude of this effect depends on the extent of linkage disequilibrium (LD) between the constrained and neutral loci, suggesting the two-locus model of Hill and Robertson (1968) as a natural framework for our endeavor.

In the absence of linked selection, expected pairwise diversity at a neutral locus, $\mathbb{E}[\pi_0]$, is governed by a balance between mutation and drift (Wright, 1931). This follows the familiar expressions for the expected changes in diversity in one generation:

$$\Delta_u \mathbb{E}[\pi] = 2u,$$
$$\Delta_{N_e} \mathbb{E}[\pi] = -\frac{1}{2N_e}\mathbb{E}[\pi].$$

The dynamics of $\mathbb{E}[\pi]$ under neutrality depend only on itself, and this simple system closes. With selection at a linked locus (left, arbitrarily), allele frequencies and pairwise diversity at the neutral (right) locus change depending on its association with the selected allele (Appendix B). For example, the expected allele frequency (q) at the right locus changes due to selection at the left locus as $\Delta_s \mathbb{E}[q] = s\mathbb{E}[D]$, where D is the standard covariance measure of LD. For pairwise diversity at that locus,

$$\Delta_s \mathbb{E}[\pi_R] = 2s\mathbb{E}[D(1-2q)]. \tag{1}$$

Here, $\mathbb{E}[D(1-2q)]$ is positive, and the decay of π_R due to LD with the constrained locus becomes explicit.

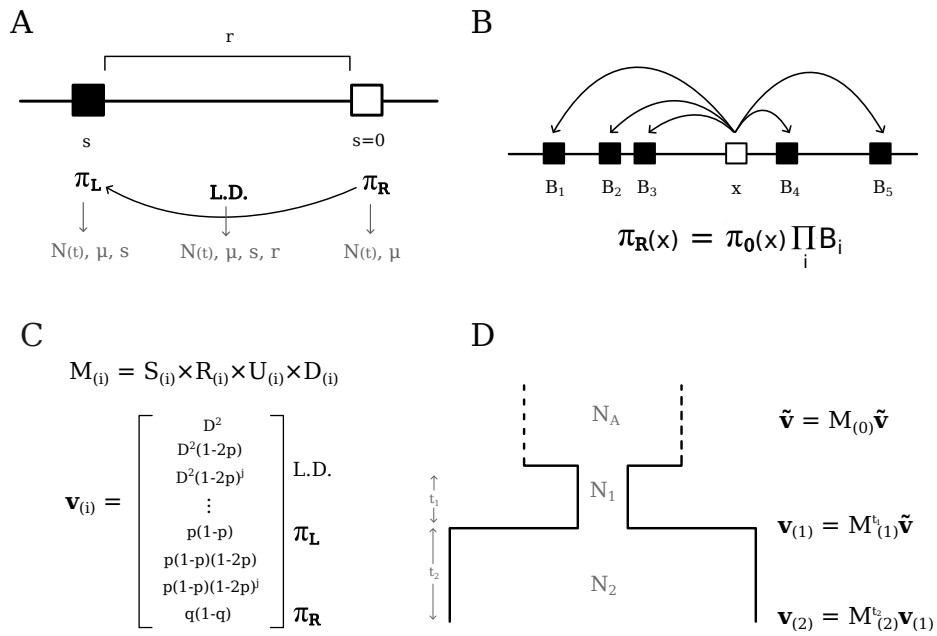

Figure 1: Schematic representation of the Hill-Robertson system with negative selection against the derived allele at the left locus. A) In a pure two-locus model, neutral pairwise diversity (π_R) is modulated by LD with a single constrained site located at recombination distance r that experiences selection strength s. B) The multi-locus extension of the pure two-locus model, where the total reduction in diversity, arising from many pairwise interactions, is multiplicative. C) A simplified representation of the system components, where M is the transition matrix that embodies the evolutionary parameters and operates on $\mathbf{v}$ (i indexes the epoch). D) Obtaining expectations under a three-epoch model, starting from equilibrium in the ancestral population (top) and finishing at present time (t_i is the number of generations spent on epoch i).

With the inclusion of linked selection, the number of statistics needed to model $\mathbb{E}[\pi_R]$ grows and the system and no longer closes (Appendix B). This happens because the addition of $\mathbb{E}[D(1-2q)]$ also requires understanding its dynamics under mutation, drift, recombination and selection (Ragsdale and Gravel, 2019). In computing these terms, the system of recursions on these statistics (moments of the full two-locus haplotype frequency distribution) accumulates an ever-increasing number of moments, with the presence of selection consistently demanding moments of slightly higher order than those previously found (Appendix A, and see Jouganous et al. (2017); Ragsdale and Gravel (2019)). However, these take a familiar form, as we soon find moments that were originally studied by Hill and Robertson (1968); Ohta and Kimura (1969) to compute the variance of D under neutrality: $\mathbb{E}[D^2]$, $\mathbb{E}[D(1-2p)(1-2q)]$, and $\mathbb{E}[p(1-p)q(1-q)]$ (where p and q are the derived allele frequencies at the left and right loci, respectively). Therefore, the moments accumulated due to the inclusion of selection resemble those in the original Hill-Robertson basis, augmented by factors of $(1-2p)^j$, with $j \geq 0$. That is, the system includes all terms $\mathbb{E}[D^2(1-2p)^j]$, $\mathbb{E}[D(1-2p)^j(1-2q)]$, and $\mathbb{E}[p(1-p)(1-2p)^j q(1-q)]$, as well as the single-locus terms $\mathbb{E}[\pi_L(1-2p)^j]$, and we label this basis $\mathbf{v}$. Within $\mathbf{v}$, each moment with a given factor $(1-2p)^j$ depends only on its close neighbors (moments with $1-2p$ factors of $j \pm 1$ or 2) so that the system remains sparse. This suggests that a treatment of the two-locus system with selection is manageable, but a moment-closure approximation is required to obtain a concise matrix representation of the recursions.

Surprisingly, naive truncation of the system after a sufficient number of $1-2p$ factors produces accurate results (Figure S7). Since stronger selection demands faster replacement of eliminated haplotypes (Appendix A, the required maximum order of $(1-2p)^j$ is an increasing function of $|N_e s|$. Since the extra information

brought by the inclusion of a new haplotype in a small sample is higher than in a sample that is already large, this function is supra-linear (Jouganous et al., 2017). Throughout the remainder of this article, when modeling different strengths of selection, we compute statistics using the appropriate truncation level of $(1-2p)^j$ factors needed for accurate prediction, but omit this information for ease of exposition.

Modeling strong linked selection under non-equilibrium demography

The demanding nature of the selection operator imposes practical problems for $|s| \gtrsim \mathcal{O}(0.01)$ (see Conclusions). Fortunately, the strong-selection regime is where cBGS approximations can take over without losing accuracy (Charlesworth et al., 1993; Hudson and Kaplan, 1995b; Nordborg et al., 1996). In particular, Nordborg (1997) showed that the structured coalescent framework can be used to model background selection through a separation of time scales. In this approach, coalescence at the neutral locus occurs within allelic classes defined by the number of haplotypes in a sample that carry the deleterious variant at the constrained locus, and transitions between classes due to recombination and mutation are rapid relative to coalescence. With strong selection, mutation-selection balance applies, maintaining the frequency of the deleterious class at $p = u/s$, and the reduction in coalescence times was found to be independent of N_e (Nordborg et al., 1996; Nordborg, 1997):

$$B = 1 - \frac{u}{s(1+r/s)^2}.$$

To incorporate population size changes, we assume a piecewise-constant population size history. Using phase-type theory (Hobolth et al., 2019), we can efficiently calculate probabilities of coalescence within each epoch and the expected T_{MRCA} conditional on coalescing within a given epoch (Methods). This is done for both the standard neutral coalescent and structured coalescent (Nordborg, 1997) models to obtain the expected reduction in diversity.

On the surface, the independence of B-values from N_e at steady state could suggest that they would be invariant to changes in population size. Indeed, size changes have little effect on deleterious diversity, so that π_L remains approximately constant (Figure S8, top panels). However, B-values are transiently affected, as both π_0 and π_R depend more sensitively on N_e – and to different degrees. To see this, consider the distribution of genealogies under a two-epoch demographic model. After the size change, and until steady state is re-established, the probability of coalescence within the most recent epoch disproportionally differs whether the neutral locus is linked to a selected locus or not. This distorts the reduction in the expected coalescence times of linked neutral sites relative to unlinked sites, so that the ratio π_R/π_0 varies temporally. Another interpretation of this phenomenon is that π_R reaches the new steady-state at a faster rate than π_0 due to its increased coalescence rate, since to a first approximation, linked selection incurs a reduction in N_e (Charlesworth et al., 1993; Charlesworth, 2009). Figure S3 demonstrates this effect, and our theoretical results confirm observations from simulations carried out in Torres et al. (2020).

The extension of the Nordborg (1997) cBGS framework to non-equilibrium demography completes the bridge between weak and strong selection. By joining the numerical solutions of `moments++` and the analytical predictions extended from Nordborg's model, we can predict B-values under arbitrary single-population demographic histories and distributions of fitness effects. While we focus here on changes in population size, this approach based on phase-type theory is quite flexible. It may be extended to accommodate piecewise-constant changes in the mutation rate at the selected site, recombination rate, the strength of selection (as long as it remains strong, relative to drift), or a combination of these.

B-value behavior under equilibrium demography

We first benchmarked our model against equilibrium ($N_e = 10,000$) simulations of a 100 kb segment with uniform recombination and mutation rates (Methods). In multi-locus settings, the B-value of each site is obtained by multiplying together the B-values resulting from every pairwise interaction. Therefore, our approach follows the formulation of Hudson and Kaplan (1995a,b), who started from two-locus reasoning to arrive at analytical approximations of multi-locus effects. We observe close agreement between our predictions

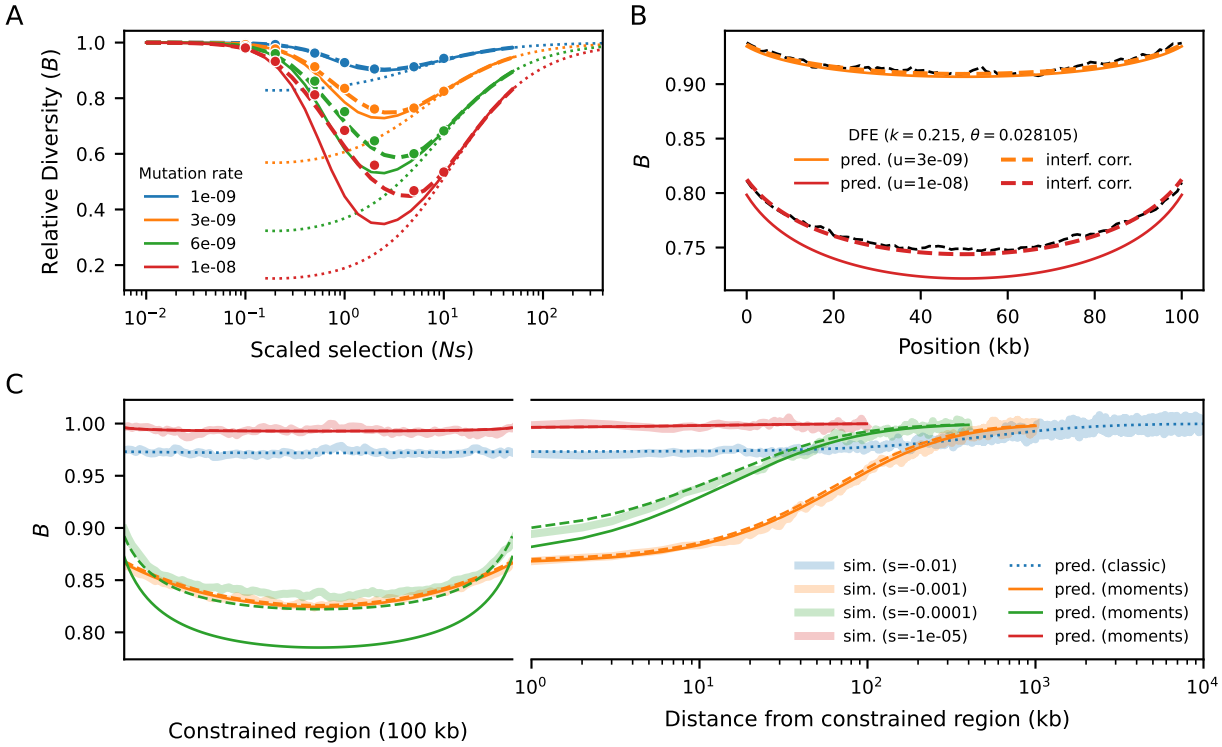

Figure 2: Benchmarking against equilibrium simulations of a 100 kb segment ($N_e = 10,000$). A) B-value in the middle of the segment, for different mutation rates and strengths of selection. Dots are simulated B-values (averaged over 10,000 replicates). Solid lines show naive predictions, dashed lines incorporate interference correction and dotted lines mark cBGS. B) B-map predictions remain accurate when s follows a distribution of fitness effects (black dashed lines show simulated data). C) B-values along the simulated segment. A concave pattern emerges because loci near the edges are on average further from constrained elements, and is counteracted by stronger interference in the center, with flattens the B-map. As the distance from the selected segment increases moving into a neutral flanking region, linked selection relaxes at a pace that depends on s and r. Weak-to-moderate selection ($N_e s \sim -1$) distorts patterns of diversity sharply albeit locally, whereas stronger selection ($N_e s \leq -10$) imposes a small but still relevant reduction throughout the entire region. Solid lines show naive predictions, dashed lines incorporate interference correction, and shaded envelops mark simulated data.

and simulations, from weak to strong selection, with substantial overlap in the regime where cBGS can take over (Figure 2). Moreover, unlike the implementation of Buffalo and Kern (2024), there is no discontinuity in the prediction around $N_e s$ of -1. This is precisely the peak of linked selection, where the combination of intermediate deleterious diversity and intermediate rates of selective elimination leads to a sharp reduction in diversity at nearby linked sites.

Because our multiplicative model assumes that constrained loci evolve independently, ignoring interference leads to artificially deflated predictions at higher mutation rates. To alleviate this bias, we draw inspiration from cBGS (Charlesworth et al., 1993; Charlesworth, 2009) and model interference as a reduction in N_e at constrained loci, which varies along the genome. In practice, we simultaneously scale μ, r and s by the corresponding B-values at each locus (Good et al., 2014), such that the effective rates of mutation, recombination and the efficiency of selection reflect local N_e across the genome. This changes B-value predictions all around, so we repeat the process until convergence (Methods, Figure S1). This procedure restores accuracy, and we conclude that `moments++` can predict B-values over a wider range of selection strengths than previous methods.

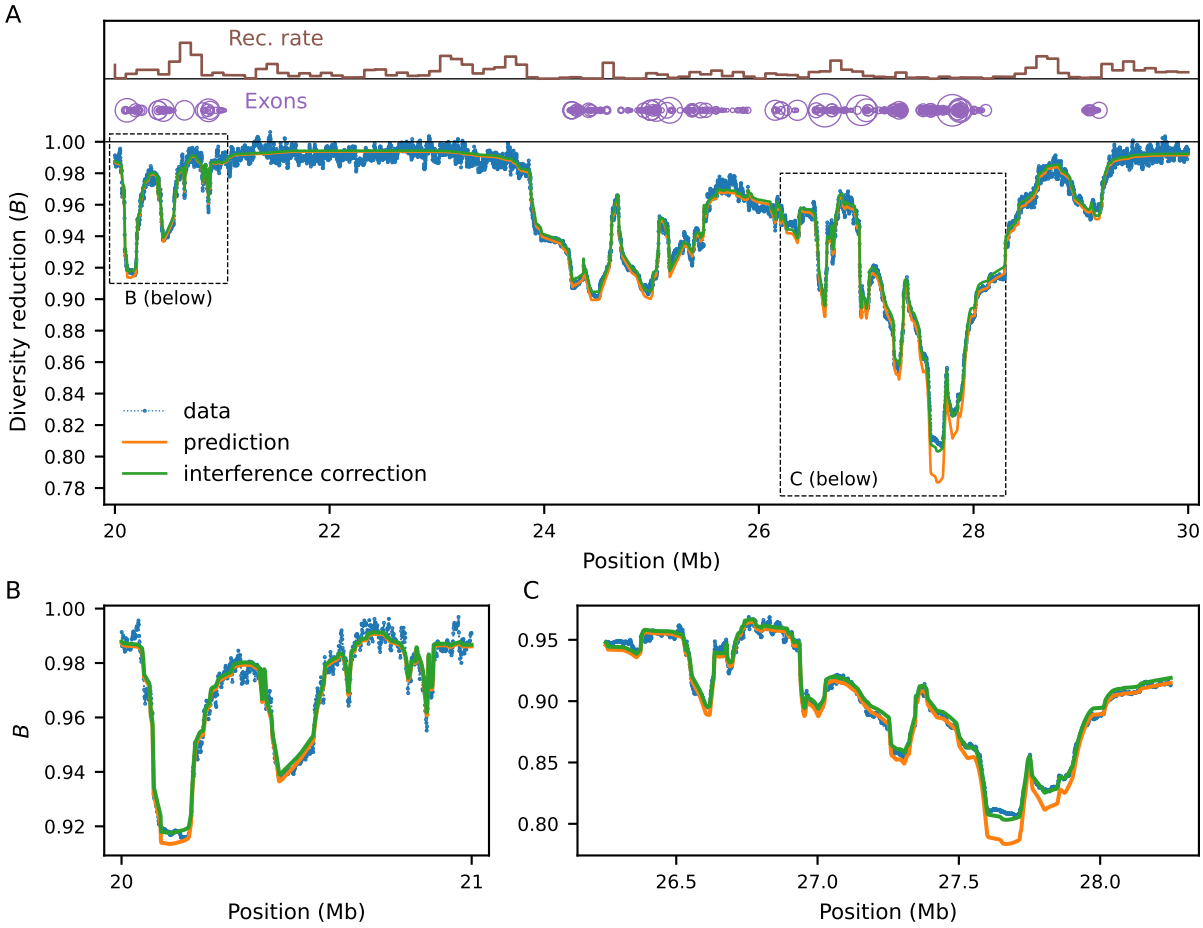

Figure 3: Benchmarking against equilibrium simulations ($N_e = 10,000$) mimicking the structure of human chromosome 2, and zooming in on a 10 Mb segment (A). Local recombination rates and exon density are shown at the top. Simulated data (averaged over 10,000 replicates) is shown by blue dots. Green and orange lines show prediction with and without interference correction, respectively. Regions distant from exons experience weaker reduction in diversity, whereas B-values dip in exon-rich regions, especially if the local recombination rate is low (cf. panels B and C). Here the mutation rate is 10^{-8} and s follows the Kim et al. (2017) DFE.

We now increase the complexity of the simulations by partially mimicking the first 30 Mb of human chromosome 2 (Methods). The heterogeneous distributions of constrained elements (here, annotated exons (Schneider et al., 2017)) and recombination rates (Spence and Song, 2019) lead to intricate patterns of linked selection along this segment, which our model predicts remarkably well (Figure 3). Regions that are on average distant from exons experience weaker reduction in diversity and negligible interference. However, even in regions largely devoid of exons, B-values stand unequivocally below one. This happens because they are affected by strongly selected sites whose effect is less sensitive to genetic distance (Figure 2C). Conversely, B-values dip in exon-rich regions, especially if the local recombination rate is low. This is driven by the strong narrow-range distortions caused by weak-to-moderate selection ($N_e s \sim -1$). These results suggest that predicting B-values across the full range of selection coefficients – rather than truncating the distribution of fitness effects (DFE) to exclusively consider strong selection (McVicker et al., 2009; Murphy et al., 2022) – can substantially improve models of genetic diversity (Figure S6).

B-value dynamics in non-equilibrium populations

At steady-state, genetic diversity is balanced among drift, mutation, recombination and (linked) selection. Changes in population size lead to transient dynamics of linked selection, until the new equilibrium is reached (Torres et al., 2020). Since B-values represent the accumulated reduction in diversity at a focal site due to $\mathcal{O}(N_e)$ generations of negative selection on linked sites, and since $N(t)$ concertedly modulates the efficiency of selection, deleterious diversity and effective recombination, we find that population fluctuations can have a substantial impact on the evolution of this statistic. B-value dynamics are often non-monotonic over time, depending on the strength of selection and demographic history (Figures S9, S10, Appendix C).

Here we focus on multi-locus models, in which we couple the chromosome layout from Figure 3 (including a DFE inferred from human data by Kim et al. (2017)) with three demographic models: 1) a 10-fold bottleneck ($N = 10\,000$ to $1\,000$, tracked for 25,000 generations after size change); 2) a 10-fold expansion ($N = 10\,000$ to $100\,000$, also for 25,000 generations); and 3) a four-epoch model that loosely resembles human effective population size history inferred by Cousins et al. (2024). By sampling the expected B-map at different time points after the first population size change, we show that neglecting demography can bias predictions and subsequent downstream inferences.

To assess the bias introduced by incorrectly assuming equilibrium, we used the drift-effective population size at each sampling generation ($N_e(t) = \pi_0(t)/4\mu$) to predict a series of steady-state B-maps. This approximation assumes that deleterious diversity and selective efficiency have remained constant over time, as well as the population-size scaled recombination and mutation rates. Therefore, they contrast with the demography-aware B-map which reflects genealogical history and evolves over time.

In the bottleneck scenario, the discrepancy between equilibrium and demography-aware B-maps initially grows as the ancestral size contributes less to drift-effective N_e and then shrinks again as N_e approaches the current size (Figure S4, blue). The bias peaks $\sim 5,000$ generations after the bottleneck, when most of the change in π_0 has already occurred (drift-effective $N_e \sim 2,000$), suggesting an asymmetric effect of historical population sizes in shaping B-values. Conversely, the expansion scenario sees the bias grow throughout the 25,000 generations of evolution (Figure S4, orange), whereas the Cousins et al. (2024) model shows more subtle deviations from equilibrium B-maps (Figure S5). Since in the latter model the ratios between consecutive population sizes do not differ as dramatically as a 10-fold bottleneck or expansion, discrepancies in predicted B-maps are not as pronounced. Thus, assuming equilibrium will bias predictions in general, in agreement with simulations (Torres et al., 2020). However, the magnitude of the bias depends on the precise population history and DFE, such that equilibrium B-maps may be a reasonable approximation in some systems.

Implications for DFE inference

What are the practical consequences of neglecting demography when predicting B-values? It seems reasonable that the ensuing biases should carry over to downstream analyses. In this regard, previous applications of equilibrium models of linked selection have also estimated selection parameters, including the DFE and deleterious mutation rate, by minimizing the distance between the predicted landscape of diversity and that observed in human data (McVicker et al., 2009; Murphy et al., 2022; Buffalo and Kern, 2024).

These estimates have been subject to debate. Specifically, the inferred (deleterious) mutation rate is much larger than expected, incompatible with more direct estimates of the human mutation rate (Ségurel et al., 2014; Tian et al., 2022)). Furthermore, the DFE of functional elements such as exons conflicts with those obtained from the site frequency spectrum (Williamson et al., 2005; Keightley and Eyre-Walker, 2007; Boyko et al., 2008; Kim et al., 2017). For some classes of constrained elements the inferred DFE is strongly bimodal, implying a combination of linked selection that is either weak and very localized ($|N_e s| \ll 1$) or long range ($|N_e s| \gg 1$). This suggests problems with identifiability: the elevated mutation rate increases deleterious diversity, reducing B and partially compensating for the low density of mutations around $N_e s = -1$.

To investigate the bias introduced by assuming equilibrium, we examined scenarios of fluctuating demography and several DFE parameterizations. We predicted B-maps in the ancestral population and at every 500 generations after the first population size change, which we then used to infer the deleterious

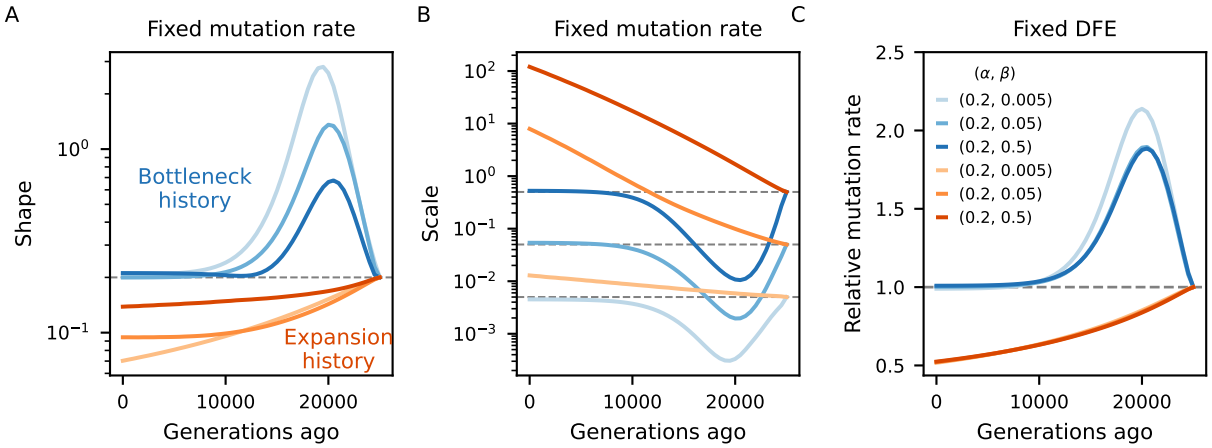

Figure 4: Biases in the inference of selection parameters due to the incorrect assumption of equilibrium. An ancestral population experiences either a 10-fold bottleneck (Blues) or expansion (Reds) and is followed for 25,000 generations after the size change. Solid lines in A-B show the inferred shape and scale of the DFE at different points in time ($N_e(t) = \pi_0(t)/4\mu$), which flows from right to left (dashed lines denote simulated value). Solid lines in C show the relative inferred (deleterious) mutation rate. We either assume the mutation rate is known and fit the shape and scale of the DFE (A-B), or vice-versa (C).

mutation rate and the shape and scale of the Gamma-distributed DFEs that best fit those maps under the assumption of steady-state size history. At each sampling time point, the maximum-likelihood estimates of the aforementioned selection parameters will be those which minimize the distance between equilibrium and demography-aware B-maps. This setup provides a best-case scenario for parameter inference by removing any other source of error, technical or biological (including genealogical noise), that would be present in an analysis of real data, and in this sense it offers a conservative assessment of the bias.

Figure 4 compares simulated and inferred parameter values for the two-epoch scenarios. As expected, they are well recovered in the ancestral population which finds itself at steady-state. As time moves forward after the size change, a strong departure develops under both bottleneck and expansion. The direction of the bias is opposite between demographic scenarios, and within each it is again opposite between shape and scale of the DFE. Mirroring the patterns observed in the previous subsection, the bias grows as the ancestral size contributes less to drift-effective N_e. Under the bottleneck, it then shrinks as N_e approaches the current size (Figure 4, blue curves). Conversely, in case of expansion, it grows throughout the 25,000 generations observed, and predictions from the pure two-locus model (Appendix C, Figure S10) suggest that it would take a long time to restore accuracy using equilibrium assumptions. We also see that the magnitude of the bias depends on the selection regime, with the shape being the most sensitive parameter. Translating inferences to the mean and variance of the DFE, stronger selection leads to stronger absolute deviations (Figure S11).

The behavior is markedly different in the human-like history (Cousins et al., 2024), where consecutive population sizes move the target equilibrium point in opposite directions. Here we tested more dissimilar DFEs and found that the shape parameter is underestimated under strong selection but overestimated otherwise (Figure 5B). The mean selection coefficient is underestimated, but the precise temporal trajectory of the bias depends on the DFE (Figure S12). Since the out-of-Africa event is relatively recent (Gutenkunst et al., 2009; Gravel et al., 2011), we expect that B-maps from distinct human populations will be fairly similar, even after accounting for demography, in agreement with what has been found under equilibrium (Murphy et al., 2022; Buffalo and Kern, 2024).

Taken together, these results highlight the importance of jointly modeling demography and linked selection (Johri et al., 2020, 2023). The extent of the bias in B-maps and downstream inference will depend on the precise population history and DFE, but it can potentially be large. Previous studies may have misinterpreted

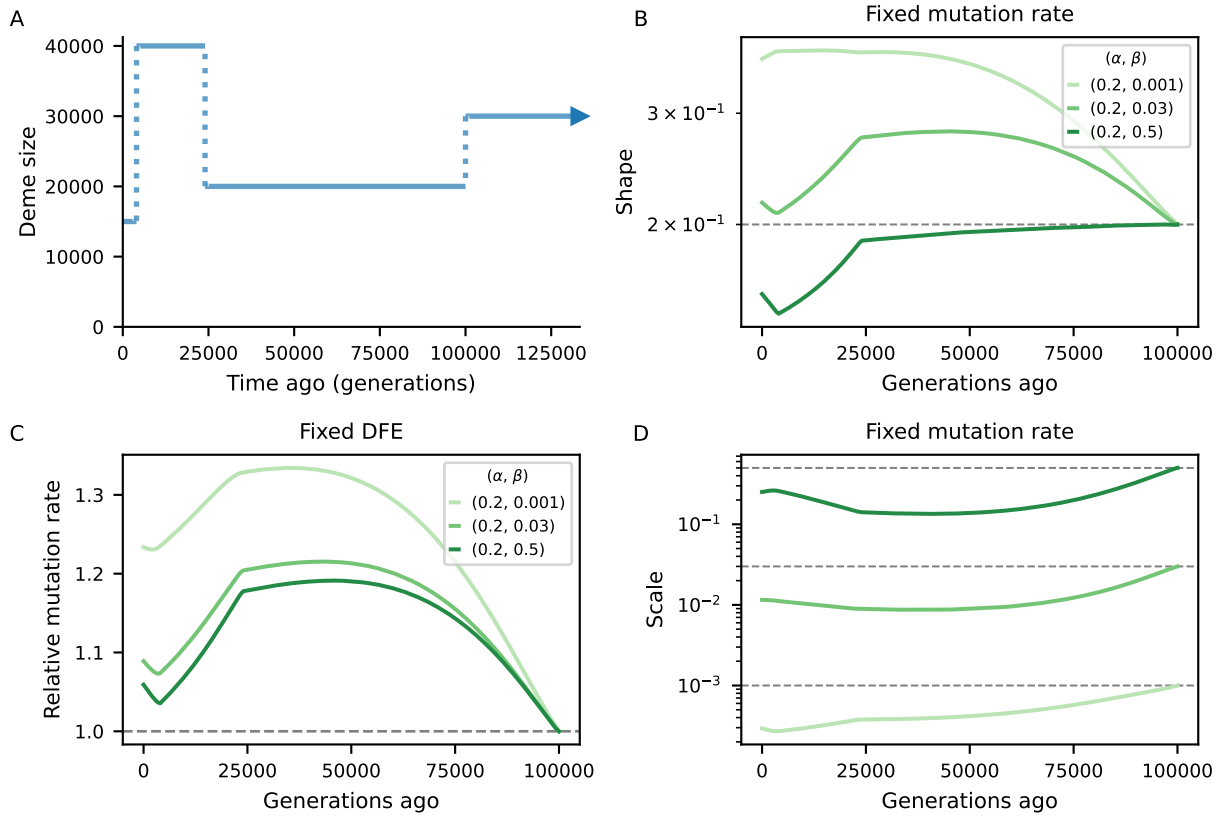

Figure 5: Biases in the inference of selection parameters due to the incorrect assumption of equilibrium. Population size trajectory (top left) is a rough piecewise constant representation of the Cousins et al. (2024) model. Other panels mirror Figure 4, except that shape and scale here combine for less deleterious DFEs. We either assume the mutation rate is known and fit the shape and scale of the DFE (panels B and D), or vice-versa (C).

patterns of diversity, potentially due to fitting probability weights to a discretized DFE which allowed too many degrees of freedom. Moreover, since both cBGS and the Buffalo and Kern (2024) implementation of SC16 show reduced accuracy around $N_e s = -1$, our assessment of the temporal bias – arriving from `moments++` predictions – is conservative relative to what we may expect in previous studies.

Conclusions

Accurately characterizing the effect of linked (negative) selection has been a goal of population genetics for decades (Charlesworth et al., 1993; Charlesworth, 2013; Cutter and Payseur, 2013; Charlesworth and Jensen, 2021). Pioneering studies focused on strongly selected sites (Charlesworth et al., 1993; Hudson and Kaplan, 1995b,a; Nordborg et al., 1996) while a subsequent framework managed to incorporate weak selection (Santiago and Caballero, 1995, 1998, 2016; Buffalo and Kern, 2024). These models have enjoyed success in predicting π across the genome, but their limitations in the interference selection regime and their assumption of equilibrium demography raise questions about their validity. Here we incorporated negative selection in the two-locus system (Hill and Robertson, 1968; Ragsdale and Gravel, 2019), which jointly models the evolutionary forces of drift, mutation, recombination and selection. The flexibility of such two-locus methods supports population-size history and a wide range of selection coefficients, which critically

includes the region around $N_e s = -1$, improving accuracy in several evolutionary scenarios while remaining computationally tractable. We also accommodated the Nordborg (1997) model to historical changes in population size, bridging the gap in B-value prediction, from weak to strong selection, for non-equilibrium populations. Exploring our new model, we found rich temporal dynamics of linked selection that depend on the DFE, genome structure, and precise demographic history.

Being able to readily compute expectations under arbitrary single-population scenario brings several benefits. For example, we showed that the temporal dynamics of B-values is highly dependent on the strength of selection and the prescribed demography (Figure S9). This clarifies the patterns highlighted by Figure S7 in Torres et al. (2020), where the trend may sometimes be hidden behind simulation noise. Instead, our approach provides expectations for diversity statistics in scenarios that would be computationally burdensome to extract from forward-in-time simulation (Figure S10). We note that Johri et al. (2021) derived expectations of B-values under two-epoch scenarios and assuming that $N_e s$, $N_e \mu$ and $N_e r$ remain constant (which holds as long as sampling takes place shortly after the population size change). They used these expectations to investigate the bias in demographic inference incurred by linked selection. Our more general model can foster the development of joint-inference approaches to address this important issue (Johri et al., 2020, 2023; Mackintosh et al., 2025).

Still there are caveats to our implementation. One technicality is that strong selection requires high orders of $(1-2p)^j$ factors, which places the transition matrix close to singularity such that numerical instability may prevent the linear algebra algorithm from finding the steady-state of $\mathbf{v}$ (Figure 1). The critical point of instability is a function of $N(t)$, s and the order of $1-2p$, which makes it difficult to propose a useful rule of thumb to circumvent it. In limited testing, we were always able to find an order of $1-2p$ that both displays high accuracy and prevents descent into numerical chaos. In practice, however, it is easier to simply use `moments++` in scenarios where $|s| \leq 0.005$ and employ our extension of cBGS in the strong selection regime. Similarly, drastic demographic fluctuations can generate expectations of high-order moments of $1-2p$ that are numerically indistinguishable from zero in epochs of low N_e. In such cases, scaling up the mutation rate restores numerical precision and is inconsequential to the prediction of B-values, which measure relative diversity.

A more conceptual challenge is in working out the interference correction under non-equilibrium demography, because population sizes fluctuations cause the dynamics of linked selection to change over time and thus the level of interference. Therefore, no single rescaling of local N_e can properly represent the historical linked selection exerted on a constrained site (analogously to neutral sites) and our approximate correction becomes improper. To circumvent this problem, in this article we focused on the temporal dynamics of B-maps in scenarios with relatively weak interference. We caution that assembling B-maps in species that experience severe size changes and widespread interference should require adjustments to take both into account.

Another current limitation of `moments++` is the restriction to single-population models. If non-equilibrium demography already creates complex temporal dynamics under panmixia, scenarios involving structure, migration and admixture may further exacerbate them (Hasan and Whitlock, 2024). Work is ongoing to assimilate gene-flow into the system and characterize these patterns. Finally, although `moments++` incorporates all two-locus statistics from the original basis $\mathbf{v}$, here we only examined π_R (and B-values). Describing the behavior of LD statistics under linked selection is of interest, but we leave it for a future treatment.

Our main practical result is that inaccurate prediction of B-maps (either by neglecting a critical part of the DFE, assuming equilibrium demography, or both) biases downstream analyses (Figures 4, 5, S6). This helps explain why B-map based inferences of DFEs and deleterious mutation rates (McVicker et al., 2009; Murphy et al., 2022; Buffalo and Kern, 2024) produced estimates that are incompatible with other studies, although other factors could also be in play. These models were justified by their high predictive power, as measured by the variance in observed diversity explained by their B-maps, but aiming highly parameterized DFE models at the B-value landscape alone likely led to over-fitting. Although beyond the scope of this paper, work is underway to investigate models of human genetic diversity with `moments++`.

Besides DFE inference, B-values have also been used to investigate the impact of linked selection on the heritability of complex diseases, following the rationale that causal variants segregate more freely where the local effect of genetic drift is greater (Pardiñas et al., 2018; Li and Berg, 2025). In a similar vein,

Simon and Coop (2024) showed that in recent European history, most of the temporal covariance in allele frequency change that would be attributable to linked selection can instead be explained by gene-flow, except in the lowest B-value bins. B-values are also routinely adopted in genomic masks, as a way to mitigate the background selection bias in demographic inference, e.g. Medina-Muñoz et al. (2023). We expect that our model will help to refine future analyses of this kind, especially when interest lies in close proximity to (weakly) constrained elements, where the cBGS theory employed so far struggles the most (Figure S6).

Methods

Model implementation

Due to the computational burden imposed by selection, we re-implemented the `moments.LD` framework (Jouganous et al., 2017; Ragsdale and Gravel, 2019) in `C++`. In doing so, we made two changes to the basic layout of the model.

First, we obtain the combined transition matrix M by multiplying together (rather than summing) the matrices that represent the individual operators (Selection, Recombination, Mutation, Drift, Figure 1C): $M = S \times R \times U \times D$. This means that terms as small as $\mathcal{O}(\mu rs/2N)$ can be retained within M, rendering it slightly denser than the transition matrix in `moments.LD`. In practice, these inclusions have a negligible effect on the expectations such that predictions from `moments++` and `moments.LD` are indistinguishable under neutrality. Related, note that re-arranging the multiplication above amounts to changing the order of events in the life-cycle (drift represents reproduction), but this also has minimal impact on the predictions.

Second, `moments++` approaches demography as a collection of epochs with piecewise constant population sizes (Figure 1D), instead of the continuous-time treatment of `moments.LD`. This simplification brings about further computational efficiency and streamlines the design of evolutionary models to be explored in future studies, e.g. when selection coefficients, mutation rates and/or recombination rates change over time.

We benchmarked our model against `moments.TwoLocus` (Ragsdale, 2022) for different selection coefficients and orders of $1-2p$ factors (Figure S7). As mentioned in the concluding remarks, some parameter combinations may introduce numerical instability if the order of $1-2p$ factors is exceedingly high (e.g., when over-shooting it to guarantee accuracy in strong selection scenarios). In such cases, `moments++` will throw an error and suggest adjusting this number.

Extending cBGS to piecewise-constant histories

The expected reduction in diversity due to linked selection can be expressed as the expected reduction in the expected coalescence time (T_{MRCA}) of a pair of samples experiencing linked selection relative to that under neutrality. In a constant-sized population, $\mathbb{E}[T_{MRCA}] = 2N_e$ under neutrality. Nordborg (1997) developed a structured coalescence model to obtain the expected T_{MRCA} given strong purifying selection at a linked locus, with mutation rate u, selection coefficient s, and genetic distance r between the selected and neutral loci.

We consider sampling two haplotypes, which could be found in the following states: both haplotypes are free of the deleterious allele (denoted $(2,0)$), one carries the deleterious allele (denoted $(1,1)$), or both carry the deleterious allele (denoted $(0,2)$). The absorbing states are reached by coalescence from states $(2,0)$ and $(0,2)$, and are denoted $(1,0)$ and $(0,1)$. The transition matrix between states is given by Nordborg (1997) (Equation 2) as

$$T = \begin{bmatrix} 1-b_{12} & 0 & 0 & 0 & b_{12} \\ \frac{1}{1-p} & 1-2b_{12}-\frac{1}{1-p} & 2b_{12} & 0 & 0 \\ 0 & b_{21} & 1-b_{12}-b_{21} & b_{12} & 0 \\ 0 & 0 & 2b_{21} & 1-2b_{21}-\frac{1}{p} & \frac{1}{p} \\ b_{21} & 0 & 0 & 0 & 1-b_{21} \end{bmatrix}, \tag{2}$$

where $p = u/s$ is given by mutation-selection balance, and the coefficients b_{12} and b_{21} are

$$b_{12} = pr,$$

$$b_{21} = (1-p)(s+r).$$

Nordborg (1997) showed that the total coalescence rate in this model is given by

$$\lambda = \frac{1}{(b_{12}+b_{21})^2}\left(\frac{b_{21}^2}{1-p} + \frac{b_{12}^2}{p}\right), \tag{3}$$

which is greater than one, so that

$$\mathbb{E}[T_{MRCA}|\text{sel.}] = \frac{2N_e}{\lambda}. \tag{4}$$

To account for piecewise-constant population size changes, we consider the probability that coalescence occurs within a given epoch and compute the T_{MRCA} conditioned on coalescence occurring within that epoch. In this case,

$$\mathbb{E}[T_{MRCA}] = \sum_{\text{epoch}} P(\text{coal. in epoch})\mathbb{E}[T_{MRCA}|\text{coal. in epoch}]. \tag{5}$$

For the neutral coalescent,

$$P(\text{coal. within epoch}) = P(\text{no coal. before epoch}) \times (1 - \exp(-t/2N_e)),$$

where the length of the epoch t is measured in generations. The probability that no coalescence occurs before a given epoch can be iteratively obtained over the preceding epochs, looking backward in time. In the model with selection, this can be found using the sub-intensity matrix of the transition matrix (Equation 2),

$$S = \begin{bmatrix} 1 - 2b_{12} - \frac{1}{\nu(1-p)} & 2b_{12} & 0 \\ b_{21} & 1 - b_{12} - b_{21} & b_{12} \\ 0 & 2b_{21} & 1 - 2b_{21} - \frac{1}{\nu p} \end{bmatrix},$$

in which $\nu = \frac{N_{\text{epoch}}}{N_e}$. Here, N_e is taken to be some reference population size, often arbitrarily chosen as the size in the most ancient epoch. Using the expected relative frequencies of occupying each transient state under mutation-selection balance, we then have

$$P(\text{coal. within epoch}|\text{sel.}) = P(\text{no coal. before epoch}) \times [p^2, 2p(1-p), (1-p)^2] \cdot S \cdot [1,1,1]^T. \tag{6}$$

The expected coalescence time, conditioned on coalescing within an epoch, can be found using the coalescence rate λ (Equation 3). In an epoch spanning time $[0, t)$,

$$\mathbb{E}[T_{MRCA}] = P(\text{coal. in epoch})\mathbb{E}[T_{MRCA}|\text{coal. in epoch}] + (1 - P(\text{coal. in epoch}))(t + \mathbb{E}[T_{MRCA}]). \tag{7}$$

Noting that $\mathbb{E}[T_{MRCA}] = 2N_e/\lambda$ and rearranging, we find

$$\mathbb{E}[T_{MRCA}|\text{coal. in epoch}] = \frac{2N_e}{\lambda} - \frac{1 - P(\text{coal. in epoch})}{P(\text{coal. in epoch})} \times t. \tag{8}$$

Then for any epoch spanning time $[t_0, t_f)$, we can use this same formula, with $t = t_f - f_0$, and simply add t_0 to account for the translation of time to the start of the given epoch.

Multi-locus prediction

At its heart, `moments++` is a pure two-locus model in which the focal neutral site is linked to a single constrained site at arbitrary genetic distance (Figure 1A). Our approach to predicting B-values in the presence of several constrained elements involves pre-computing a lookup table that stores, for a prescribed demographic model, pure two-locus predictions on a parameter grid. In this article, each of the lookup

tables built with `moments++` embodied 35 selection coefficients (from 0 to -0.001) and 72 recombination rates (from 0 to 0.5), roughly linearly spaced in logarithmic scale. We then used our extension of the Nordborg (1997) model to incorporate more deleterious selection coefficients (down to -1), and performed cubic splines interpolation to have a smooth function of B against r that is easily applicable along the chromosome. When elements were constituted by contiguous sites, we computed the B-value reduction by multiplying together the effect of each individual site (i.e., we assume no within-element interference). In doing so, we considered the genetic distance between our focal site and the midpoint of the element. Both approximations (no interference and midpoint recombination distance) are fairly accurate for elements with ~ 1000 sites or less. When dealing with chromosome layouts that mimic the position and lengths of exons in humans, we divided elements with > 1000 sites into smaller adjacent elements. In models where selection coefficients follow a continuous DFE, we used numerical integration to obtain probability weights at each of the selection coefficients stored in the lookup table. In such cases, individual exons were assumed to follow the same DFE and therefore shared probability weights.

Lookup tables contain a single mutation rate (in this article, $\mu = 10^{-8}$ per site per generation). When relevant (e.g., performing interference correction), we applied a linear adjustment of the B-value with respect to element-specific mutation rates. Incorporating a heterogeneous mutation landscape becomes straightforward in this setting, although we did not explore such scenarios here.

Interference correction

Selective interference can be thought of as linked selection among constrained sites. Borrowing from cBGS (Charlesworth et al., 1993; Charlesworth, 2009), we developed a correction scheme where constrained elements impose reductions in each other's effective population size (Good et al., 2014). The reduction in N_e of a given element is then equivalent to its B-value. To mimic this effect, at each round of prediction we scale μ, r and s of every element by its corresponding B-value. The justification is that the dynamics are determined by the scaled rates $N_e\mu$, $N_e r$ and $N_e s$. In turn, this scaling changes the linked selection effect that the elements exert on each other (and on neutral sites), thus we re-compute B-values all around. By iterating this procedure, we can obtain a B-map that approximately accounts for interference among constrained sites. In practice, this converges in < 10 iterations (often fewer), even for dense regions with strong linked selection effects (Figure 2, S1).

B-map predictions

The chromosome layout used to assess the bias in DFE inference under non-equilibrium (Figure 4, 5) used $\mu = 10^{-8}$, $r = 10^{-8}$ as well as an arbitrary 10 Mb chromosome where 20 functional elements, each with length 1 kb, are evenly distributed. Within each element, the nonsynonymous (constrained) to synonymous (neutral) ratio was set to $2.31 : 1$ (Kim et al., 2017). The DFE parameters varied and are depicted in Figures 4, 5.

When assessing the bias incurred by incorrectly assuming equilibrium, the N_e used to predict equilibrium B-maps was based on the expected diversity coming from a purely neutral model ($N_e(t) = \pi_0(t)/4\mu$). Here, $\pi_0(t)$ was extracted from `moments++` predictions of a neutral model that follows the same demography as the linked selection model to which it is being contrasted.

Forward simulations

Multi-locus forward-in-time simulations (Figure 2, 3) were conducted using `fwdpy11` (Thornton, 2014, 2019). We computed pairwise diversity at each site (and by extension, B-values) from the tree sequences using the expected number of mutations given the realized branch lengths (removing mutational noise) and averaged over 10,000 replicates to reduce genealogical noise. Thus, our goal was not to ask how much of the variance in diversity is captured by the predictions, but to quantify the bias in `moments++` using a low-variance estimate of π_R.

In benchmarking simulations (Figure 2), per-base recombination and mutation rates were uniform, with $r = 10^{-8}$ and the deleterious rate μ taking values between 10^{-9} and 10^{-8}. Population sizes were held constant

at 10^4 diploid individuals. In simulations following human genomic annotations, the recombination rates were determined by the OMNI recombination map inferred from YRI individuals in the Thousand Genomes Project (Consortium et al., 2012). Mutation rates were held constant at $\mu = 1.5 \times 10^{-8}$, with the nonsynonymous (selected) mutation rate within annotated exons equal to $\frac{2.31}{1+2.31}\mu$ (as the ratio of new mutations within exons being nonsynonymous:synonymous has been estimated to be $2.31 : 1$ in humans (Huber et al., 2017)). Coding regions were specified by annotated exons in the human genome, using human genome build 37. Selection coefficients were drawn from a gamma-distributed distribution of fitness effects, with shape parameter 0.215 and scale parameter 0.028105 (Kim et al., 2017).

We also developed `twoLocusSim`, a forward-in-time simulator of independent two-locus systems written in `C++`. Briefly, it employs an infinite-sites mutation model to spawn one- and two-locus diversity with a pseudo-random number generator. Recombination and selection operate deterministically on haplotype frequencies, and drift is implemented through multinomial sampling of haplotypes every generation. The software can output statistics in time-series fashion, and we used `twoLocusSim` to understand the mechanism of selection in the Hill-Robertson system (Appendix B).

Data and software availability

Source code for `moments++` and `twoLocusSim` is available at https://github.com/gvbarroso/momentspp. The version of the `moments++` source code used to generate the results in the present article has been frozen and stored in https://doi.org/10.5281/zenodo.15824891. High-level tools for predicting B-maps (using a lookup table that stores `moments++` predictions) are available as a python package called `bgshr`, available at https://github.com/apragsdale/bgshr. The version of `bgshr` used to generate the results in the present article has been frozen and stored in https://doi.org/10.5281/zenodo.15802044. Jupyter notebooks used to generate figures, as well as the associated metadata, can be found at https://github.com/gvbarroso/momentspp/tree/main/doc/examples/figures_paper1. Scripts used to perform simulations are provided in https://github.com/gvbarroso/momentspp/tree/main/doc/examples/figures_paper1/fwdpy11_simulations. The Jupyter notebooks and `fwdpy11` simulation scripts are packaged together with the `moments++` source code (https://doi.org/10.5281/zenodo.15824891).

Acknowledgments

We thank members of the Ragsdale lab for helpful discussions on this work, Nick Collier for careful comments on the manuscript, and Lloyd Kirk for implementing the Docker container build. This work was supported by NIH Award R35-GM154962 to APR.

References

Montserrat Aguade, Naohiko Miyashita, and Charles H Langley. Reduced variation in the yellow-achaete-scute region in natural populations of drosophila melanogaster. *Genetics*, 122(3):607–615, 1989.

Peter Andolfatto and Molly Przeworski. Regions of lower crossing over harbor more rare variants in african populations of drosophila melanogaster. *Genetics*, 158(2):657–665, 2001.

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

David J Begun and Charles F Aquadro. Levels of naturally occurring dna polymorphism correlate with recombination rates in d. melanogaster. *Nature*, 356(6369):519–520, 1992.

CW Birky Jr and J Bruce Walsh. Effects of linkage on rates of molecular evolution. *Proceedings of the National Academy of Sciences*, 85(17):6414–6418, 1988.

Adam R Boyko, Scott H Williamson, Amit R Indap, Jeremiah D Degenhardt, Ryan D Hernandez, Kirk E Lohmueller, Mark D Adams, Steffen Schmidt, John J Sninsky, Shamil R Sunyaev, et al. Assessing the evolutionary impact of amino acid mutations in the human genome. *PLoS genetics*, 4(5):e1000083, 2008.

Vince Buffalo and Andrew D Kern. A quantitative genetic model of background selection in humans. *Plos Genetics*, 20(3):e1011144, 2024.

David Castellano, Marta Coronado-Zamora, Jose L Campos, Antonio Barbadilla, and Adam Eyre-Walker. Adaptive evolution is substantially impeded by hill–robertson interference in drosophila. *Molecular biology and evolution*, 33(2):442–455, 2016.

David Castellano, Jennifer James, and Adam Eyre-Walker. Nearly neutral evolution across the drosophila melanogaster genome. *Molecular Biology and Evolution*, 35(11):2685–2694, 2018.

B Charlesworth. Fundamental concepts in genetics: effective population size and patterns of. *Proc. R. Soc. B*, 281(20132858):9, 2009.

Brian Charlesworth. Background selection 20 years on: the wilhelmine e. key 2012 invitational lecture. *Journal of Heredity*, 104(2):161–171, 2013.

Brian Charlesworth and Jeffrey D Jensen. Effects of selection at linked sites on patterns of genetic variability. *Annual review of ecology, evolution, and systematics*, 52(1):177–197, 2021.

Brian Charlesworth, MT Morgan, and Deborah Charlesworth. The effect of deleterious mutations on neutral molecular variation. *Genetics*, 134(4):1289–1303, 1993.

Josep M Comeron. Background selection as null hypothesis in population genomics: insights and challenges from drosophila studies. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1736): 20160471, 2017.

Josep M Comeron and Martin Kreitman. Population, evolutionary and genomic consequences of interference selection. *Genetics*, 161(1):389–410, 2002.

1000 Genomes Project Consortium et al. An integrated map of genetic variation from 1,092 human genomes. *Nature*, 491(7422):56, 2012.

Trevor Cousins, Daniel Tabin, Nick Patterson, David Reich, and Arun Durvasula. Accurate inference of population history in the presence of background selection. *bioRxiv*, 2024.

Asher D Cutter and Bret A Payseur. Genomic signatures of selection at linked sites: unifying the disparity among species. *Nature Reviews Genetics*, 14(4):262–274, 2013.

Ivana Cvijović, Benjamin H Good, and Michael M Desai. The effect of strong purifying selection on genetic diversity. *Genetics*, 209(4):1235–1278, 2018.

Eyal Elyashiv, Shmuel Sattath, Tina T Hu, Alon Strutsovsky, Graham McVicker, Peter Andolfatto, Graham Coop, and Guy Sella. A genomic map of the effects of linked selection in drosophila. *PLoS genetics*, 12(8): e1006130, 2016.

Warren J Ewens. The pseudo-transient distribution and its uses in genetics. *Journal of Applied Probability*, 1 (1):141–156, 1964.

Gregory B Ewing and Jeffrey D Jensen. The consequences of not accounting for background selection in demographic inference. *Molecular ecology*, 25(1):135–141, 2016.

Adam Eyre-Walker and Peter D Keightley. The distribution of fitness effects of new mutations. *Nature Reviews Genetics*, 8(8):610–618, 2007.

Adam Eyre-Walker, Megan Woolfit, and Ted Phelps. The distribution of fitness effects of new deleterious amino acid mutations in humans. *Genetics*, 173(2):891–900, 2006.

Benjamin H Good, Aleksandra M Walczak, Richard A Neher, and Michael M Desai. Genetic diversity in the interference selection limit. *PLoS genetics*, 10(3):e1004222, 2014.

Simon Gravel, Brenna M Henn, Ryan N Gutenkunst, Amit R Indap, Gabor T Marth, Andrew G Clark, Fuli Yu, Richard A Gibbs, 1000 Genomes Project, Carlos D Bustamante, et al. Demographic history and rare allele sharing among human populations. *Proceedings of the National Academy of Sciences*, 108(29): 11983–11988, 2011.

Ryan N Gutenkunst, Ryan D Hernandez, Scott H Williamson, and Carlos D Bustamante. Inferring the joint demographic history of multiple populations from multidimensional snp frequency data. *PLoS genetics*, 5 (10):e1000695, 2009.

John Burdon Sanderson Haldane. A mathematical theory of natural and artificial selection, part v: selection and mutation. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 23, pages 838–844. Cambridge University Press, 1927.

Asad Hasan and Michael C Whitlock. F st and genetic diversity in an island model with background selection. *PLoS genetics*, 20(12):e1011225, 2024.

WG Hill and Alan Robertson. Linkage disequilibrium in finite populations. *Theoretical and applied genetics*, 38:226–231, 1968.

William G Hill and Alan Robertson. The effect of linkage on limits to artificial selection. *Genetics Research*, 8(3):269–294, 1966.

Asger Hobolth, Arno Siri-Jegousse, and Mogens Bladt. Phase-type distributions in population genetics. *Theoretical population biology*, 127:16–32, 2019.

Christian D Huber, Bernard Y Kim, Clare D Marsden, and Kirk E Lohmueller. Determining the factors driving selective effects of new nonsynonymous mutations. *Proceedings of the National Academy of Sciences*, 114(17):4465–4470, 2017.

Richard R Hudson and Norman L Kaplan. The coalescent process in models with selection and recombination. *Genetics*, 120(3):831–840, 1988.

Richard R Hudson and Norman L Kaplan. Deleterious background selection with recombination. *Genetics*, 141(4):1605–1617, 1995a.

RR Hudson and NL Kaplan. The coalescent process and background selection. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 349(1327):19–23, 1995b.

Parul Johri, Brian Charlesworth, and Jeffrey D Jensen. Toward an evolutionarily appropriate null model: jointly inferring demography and purifying selection. *Genetics*, 215(1):173–192, 2020.

Parul Johri, Kellen Riall, Hannes Becher, Laurent Excoffier, Brian Charlesworth, and Jeffrey D Jensen. The impact of purifying and background selection on the inference of population history: problems and prospects. *Molecular biology and evolution*, 38(7):2986–3003, 2021.

Parul Johri, Susanne P Pfeifer, and Jeffrey D Jensen. Developing an evolutionary baseline model for humans: jointly inferring purifying selection with population history. *Molecular biology and evolution*, 40(5):msad100, 2023.

Julien Jouganous, Will Long, Aaron P Ragsdale, and Simon Gravel. Inferring the joint demographic history of multiple populations: beyond the diffusion approximation. *Genetics*, 206(3):1549–1567, 2017.

Peter D Keightley and Adam Eyre-Walker. Joint inference of the distribution of fitness effects of deleterious mutations and population demography based on nucleotide polymorphism frequencies. *Genetics*, 177(4): 2251–2261, 2007.

Bernard Y Kim, Christian D Huber, and Kirk E Lohmueller. Inference of the distribution of selection coefficients for new nonsynonymous mutations using large samples. *Genetics*, 206(1):345–361, 2017.

Motoo Kimura and Tomoko Ohta. On the rate of molecular evolution. *Journal of Molecular Evolution*, 1: 1–17, 1971.

Xinyi Li and Jeremy J Berg. The impact of background selection in mutation-selection-drift balance models of complex trait evolution. *bioRxiv*, pages 2025–05, 2025.

Kirk E Lohmueller, Anders Albrechtsen, Yingrui Li, Su Yeon Kim, Thorfinn Korneliussen, Nicolas Vinckenbosch, Geng Tian, Emilia Huerta-Sanchez, Alison F Feder, Niels Grarup, et al. Natural selection affects multiple aspects of genetic variation at putatively neutral sites across the human genome. *PLoS genetics*, 7 (10):e1002326, 2011.

Alexander Mackintosh, Maxence Brault, Denis Roze, Martin Lascoux, and Sylvain Glemin. Estimating the reduction in genetic diversity from background selection under non-equilibrium demography and partial selfing. *bioRxiv*, pages 2025–06, 2025.

Gabor T Marth, Fuli Yu, Amit R Indap, Kiran Garimella, Simon Gravel, Wen Fung Leong, Chris Tyler-Smith, Matthew Bainbridge, Tom Blackwell, Xiangqun Zheng-Bradley, et al. The functional spectrum of low-frequency coding variation. *Genome biology*, 12:1–17, 2011.

Gilean AT McVean and Brian Charlesworth. The effects of hill-robertson interference between weakly selected mutations on patterns of molecular evolution and variation. *Genetics*, 155(2):929–944, 2000.

Graham McVicker, David Gordon, Colleen Davis, and Phil Green. Widespread genomic signatures of natural selection in hominid evolution. *PLoS genetics*, 5(5):e1000471, 2009.

Santiago G Medina-Muñoz, Diego Ortega-Del Vecchyo, Luis Pablo Cruz-Hervert, Leticia Ferreyra-Reyes, Lourdes García-García, Andrés Moreno-Estrada, and Aaron P Ragsdale. Demographic modeling of admixed latin american populations from whole genomes. *The American Journal of Human Genetics*, 110(10): 1804–1816, 2023.

Hermann Joseph Muller. The relation of recombination to mutational advance. *Mutation Research/Fundamental and Molecular Mechanisms of Mutagenesis*, 1(1):2–9, 1964.

David A Murphy, Eyal Elyashiv, Guy Amster, and Guy Sella. Broad-scale variation in human genetic diversity levels is predicted by purifying selection on coding and non-coding elements. *Elife*, 12:e76065, 2022.

Magnus Nordborg. Structured coalescent processes on different time scales. *Genetics*, 146(4):1501–1514, 1997.

Magnus Nordborg, Brian Charlesworth, and Deborah Charlesworth. The effect of recombination on background selection. *Genetics Research*, 67(2):159–174, 1996.

Tomoko Ohta and Motoo Kimura. Linkage disequilibrium due to random genetic drift. *Genetics Research*, 13 (1):47–55, 1969.

Tomoko Ohta and Motoo Kimura. Linkage disequilibrium between two segregating nucleotide sites under the steady flux of mutations in a finite population. *Genetics*, 68(4):571, 1971.

Sarah P Otto and Michael C Whitlock. The probability of fixation in populations of changing size. *Genetics*, 146(2):723–733, 1997.

Antonio F Pardiñas, Peter Holmans, Andrew J Pocklington, Valentina Escott-Price, Stephan Ripke, Noa Carrera, Sophie E Legge, Sophie Bishop, Darren Cameron, Marian L Hamshere, et al. Common schizophrenia alleles are enriched in mutation-intolerant genes and in regions under strong background selection. *Nature genetics*, 50(3):381–389, 2018.

Fanny Pouyet and Kimberbly J Gilbert. Towards an improved understanding of molecular evolution: the relative roles of selection, drift, and everything in between. *Peer Community Journal*, 1, 2021.

Aaron P Ragsdale. Local fitness and epistatic effects lead to distinct patterns of linkage disequilibrium in protein-coding genes. *Genetics*, 221(4):iyac097, 2022.

Aaron P Ragsdale and Simon Gravel. Models of archaic admixture and recent history from two-locus statistics. *PLoS genetics*, 15(6):e1008204, 2019.

Aaron P Ragsdale, Timothy D Weaver, Elizabeth G Atkinson, Eileen G Hoal, Marlo Möller, Brenna M Henn, and Simon Gravel. A weakly structured stem for human origins in africa. *Nature*, 617(7962):755–763, 2023.

Enrique Santiago and Armando Caballero. Effective size of populations under selection. *Genetics*, 139(2): 1013–1030, 1995.

Enrique Santiago and Armando Caballero. Effective size and polymorphism of linked neutral loci in populations under directional selection. *Genetics*, 149(4):2105–2117, 1998.

Enrique Santiago and Armando Caballero. Joint prediction of the effective population size and the rate of fixation of deleterious mutations. *Genetics*, 204(3):1267–1279, 2016.

Stanley A Sawyer and Daniel L Hartl. Population genetics of polymorphism and divergence. *Genetics*, 132 (4):1161–1176, 1992.

Valerie A Schneider, Tina Graves-Lindsay, Kerstin Howe, Nathan Bouk, Hsiu-Chuan Chen, Paul A Kitts, Terence D Murphy, Kim D Pruitt, Françoise Thibaud-Nissen, Derek Albracht, et al. Evaluation of grch38 and de novo haploid genome assemblies demonstrates the enduring quality of the reference assembly. *Genome research*, 27(5):849–864, 2017.

Laure Ségurel, Minyoung J Wyman, and Molly Przeworski. Determinants of mutation rate variation in the human germline. *Annual review of genomics and human genetics*, 15(1):47–70, 2014.

Alexis Simon and Graham Coop. The contribution of gene flow, selection, and genetic drift to five thousand years of human allele frequency change. *Proceedings of the National Academy of Sciences*, 121(9): e2312377121, 2024.

Jeffrey P Spence and Yun S Song. Inference and analysis of population-specific fine-scale recombination maps across 26 diverse human populations. *Science Advances*, 5(10):eaaw9206, 2019.

Kevin R Thornton. A c++ template library for efficient forward-time population genetic simulation of large populations. *Genetics*, 198(1):157–166, 2014.

Kevin R Thornton. Polygenic adaptation to an environmental shift: temporal dynamics of variation under gaussian stabilizing selection and additive effects on a single trait. *Genetics*, 213(4):1513–1530, 2019.

Xiaowen Tian, Ruoyi Cai, and Sharon R Browning. Estimating the genome-wide mutation rate from thousands of unrelated individuals. *The American Journal of Human Genetics*, 109(12):2178–2184, 2022.

Raul Torres, Markus G Stetter, Ryan D Hernandez, and Jeffrey Ross-Ibarra. The temporal dynamics of background selection in nonequilibrium populations. *Genetics*, 214(4):1019–1030, 2020.

Scott H Williamson, Ryan Hernandez, Adi Fledel-Alon, Lan Zhu, Rasmus Nielsen, and Carlos D Bustamante. Simultaneous inference of selection and population growth from patterns of variation in the human genome. *Proceedings of the National Academy of Sciences*, 102(22):7882–7887, 2005.

Sewall Wright. Evolution in mendelian populations. *Genetics*, 16(2):97, 1931.

Appendix

A The two-locus system with selection

By relating two-locus statistics to the evolutionary processes that affect them, Hill and Robertson (1968); Ohta and Kimura (1969) developed a system of equations for studying LD under neutrality. They found that solving for the evolution of $\mathbb{E}[D^2]$ under drift and recombination (where D is the standard covariance measure of LD) requires the inclusion of $\mathbb{E}[D(1-2p)(1-2q)]$ and $\mathbb{E}[p(1-p)q(1-q)]$ (p and q being the derived allele frequencies at the left and right loci, respectively). With mutation, the expected pairwise diversity at each locus is needed, that is, $\mathbb{E}[2p(1-p)]$ and $\mathbb{E}[2q(1-q)]$. Together, this set of summary statistics, which we denote $\mathbf{v}$, is closed under neutrality. That is, a full description of their dynamics depends only on other statistics already present in $\mathbf{v}$, leading to a finite number of recursions (Figure 1).

We may think of these statistics (or moments) as representing sampling configurations of four two-locus haplotypes (Ragsdale and Gravel, 2019; Ragsdale, 2022). From this angle, the individual effects of evolutionary processes become apparent. Mutation introduces one-locus diversity in the system by creating a derived allele in a previously monomorphic sample; similarly, it introduces two-locus diversity by creating a derived allele in the monomorphic locus of a sample that already segregates at the other. Recombination shuffles alleles between the two haplotypes, whereas drift reduces both one- and two-locus diversity when a haplotype copies over another (backward-in-time, a coalescence event). This perspective aids modeling other processes such as migration and admixture (Ragsdale and Gravel, 2019; Ragsdale et al., 2023). We now want to incorporate selection into the Hill-Robertson system.

The effect of negative selection on the two-locus sampling configurations is to remove a haplotype carrying the deleterious variant with rate proportional to its fitness effect (in diploids, this is equivalent to semi-dominant selection acting at the gametic stage). Under a soft selection model (i.e., selection does not alter population sizes), the eliminated haplotype must be replaced by a haplotype randomly chosen from the rest of the population. We derive the selection dynamics in expectation, reaching a deterministic model where selection steadily removes a fraction of haplotypes carrying deleterious variants from the population at each generation. Denoting the pair of ancestral and derived alleles at the left locus (constrained) by a, A and similarly b, B for the right locus (neutral), the expected haplotype frequencies after one generation of selection are (to first order in s):

$$\begin{aligned} f'_{AB} &= f_{AB}(1+s)/\bar{w} \\ f'_{Ab} &= f_{Ab}(1+s)/\bar{w} \\ f'_{aB} &= f_{aB}/\bar{w} \\ f'_{ab} &= f_{ab}/\bar{w} \end{aligned}$$

with $\bar{w} = [(1+s)(f_{AB}+f_{Ab})+f_{aB}+f_{ab}]$ the mean fitness of the population (Ragsdale and Gravel, 2019, see section S1.1.4 of the Supplement in).

Converting from two-locus haplotypes to the space of statistics where the Hill-Robertson basis resides, negative selection directly reduces diversity at the left (constrained) locus, as

$$\Delta_s \mathbb{E}[\pi_L] = s\mathbb{E}[2p(1-p)(1-2p)], \tag{A1}$$

where $\pi_L = 2p(1-p)$. Since s is negative and $\mathbb{E}[2p(1-p)(1-2p)]$ is positive, $\mathbb{E}[\pi_L]$ decreases due to direct selection. Moreover, the dependence of π_L on an additional statistic ($\pi_L(1-2p)$) immediately hints at the

non-closure of the system under selection (Jouganous et al., 2017). On the other hand, selection *indirectly* reduces diversity at the neutral locus (equation 1, Figure 1).

The mechanism of haplotype elimination and replacement creates an intricate pattern of dependence among the terms in $\mathbf{v}$. When one of the two loci is constrained while the other remains neutral, the ancestral-derived symmetry in the constrained locus is broken, as well as the left-right symmetry between loci, and the system no longer closes. This means that in the presence of selection, expectations of two-locus statistics in the next generation depend on expectations of statistics that are absent from the original basis. These, in turn, need additional moments themselves, so the system grows indefinitely. Intuitively, such growth could be anticipated from the way selection operates: to replace the eliminated haplotype, we need to know the expected frequency of haplotypes in a larger sample (Jouganous et al., 2017). Thus the expectation of, e.g., D after one generation of selection requires contributions of statistics of the form $D(1-2p)^j$ and so on, recursively, where the order j is unbounded (Figure S2). A graphical representation of the selection operator acting on $\mathbb{E}[D]$ and extending to the first three layers of extra moments is given in Figure S2.

B Proximate and ultimate explanations for the decay of π_R in the Hill-Robertson system

The goal of this section is to understand how negative selection operates in (p, q, D)-space (Ragsdale and Gravel, 2019). To this end, it is instructive to discuss why $\mathbb{E}[D(1-2q)]$ should be positive in the presence of selection. With hindsight, this must be true otherwise π_R would not decay through linkage with a selected locus (Equation 1). Without this retroactive argument, however, the reason is not obvious. $D(1-2q)$ is a LD statistic that receives greater contribution from variants at the tails of the neutral locus' frequency spectrum, with change between consecutive generations given by

$$\Delta\mathbb{E}[D(1-2q)] = -2s\mathbb{E}[D^2] + s\mathbb{E}[D(1-2q)(1-2p)] - r\mathbb{E}[D(1-2q)] - \frac{3}{2N}\mathbb{E}[D(1-2q)],$$

where the first two terms on the right-hand side embody the selection operator ($s < 0$). Its neutral expectation is zero, simply decaying by recombination and drift, but under negative selection $\mathbb{E}[D(1-2q)] > 0$ indicates that the buildup through D^2 is higher than the decay through $\mathbb{E}[D(1-2p)(1-2q)]$. Why should this be?

Taking a step back, we note that $\mathbb{E}[D]$ equals zero under both neutrality and selection. To understand why, first note that $D = 0$ until two-locus diversity emerges. This second mutation event will either create "coupling" (AB, $\Rightarrow D > 0$) or "repulsion" (aB or Ab, $\Rightarrow D < 0$) haplotypes. From the infinite sites model and the assumption of weak mutation ($N_e\mu \ll 1$), it follows that $\mathbb{E}[p] < 1/2$ and $\mathbb{E}[q] < 1/2$. That is, single-locus diversity implies a derived allele is expected to segregate at low frequency, regardless of which locus mutates first. Therefore, the second mutation event in the two-locus system is more likely to produce "repulsion" haplotypes. Naively, this would lead us to believe that mutation tends to push $D < 0$, but this higher probability is balanced by the increased value of $D > 0$ whenever mutation does generate "coupling" haplotypes. Hence mutation does not affect D in expectation (Ragsdale and Gravel, 2019). Meanwhile, recombination and drift lead to its decay:

$$\Delta\mathbb{E}[D] = s\mathbb{E}[D(1-2p)] - r\mathbb{E}[D] - \frac{1}{2N}\mathbb{E}[D]$$

Here we see that selection on D collects only from $D(1-2p)$, which itself evolves according to

$$\Delta\mathbb{E}[D(1-2p)] = \frac{3s}{2}\mathbb{E}[D(1-2p)^2] - \frac{s}{2}\mathbb{E}[D] - r\mathbb{E}[D(1-2p)] - \frac{3}{2N}\mathbb{E}[D(1-2p)]$$

and the recursions for higher-order $\mathbb{E}[D(1-2p)^j]$ moments follow the same pattern of exclusive contributions

from the adjacent $\mathbb{E}[D(1-2p)^{j-1}]$ and $\mathbb{E}[D(1-2p)^{j+1}]$ through selection:

$$\begin{aligned}
\Delta\mathbb{E}[D(1-2p)^2] &= 2s\mathbb{E}[D(1-2p)^3] - s\mathbb{E}[D(1-2p)] - \text{rec.} - \text{drift} \\
\Delta\mathbb{E}[D(1-2p)^3] &= \frac{5s}{2}\mathbb{E}[D(1-2p)^4] - \frac{3s}{2}\mathbb{E}[D(1-2p)^2] - \text{rec.} - \text{drift} \\
&\vdots \\
\Delta\mathbb{E}[D(1-2p)^j] &= \frac{(j+2)s}{2}\mathbb{E}[D(1-2p)^{j+1}] - \frac{js}{2}\mathbb{E}[D(1-2p)^{j-1}] - \text{rec.} - \text{drift}
\end{aligned}$$

`moments++` reveals that all $\mathbb{E}[D(1-2p)^j]$ moments are zero regardless of selective constraint. When $s < 0$, each collects from their adjacent neighbors with opposite signs and in the end these contributions cancel out. This means that $\mathbb{E}[D]$ is flat over the distribution of derived allele frequencies at the constrained locus (Figure A1A). Indeed, the change in p after one generation of selection,

$$\Delta_s\mathbb{E}[p] = s\mathbb{E}[p(1-p)],$$

is independent of the neutral locus. Therefore conditioning on p (via $1-2p$) should not affect the covariance between allele frequencies. Conversely, $\mathbb{E}[D(1-2q)] > 0$ suggests that $\mathbb{E}[D]$ is *not* flat over the distribution of derived allele frequencies at the neutral locus (Figure A1B). This happens because the expected change in q after one generation of selection,

$$\Delta_s\mathbb{E}[q] = s\mathbb{E}[D],$$

is entirely attributed to linkage with the constrained locus. (Hence conditioning on q (via $1-2q$) does affect the expectation of the covariance term D.) In fact, since $\mathbb{E}[D]$ is zero, the two-locus system predicts that selection does not affect q in expectation. This happens because the decrease in frequency when the neutral mutation falls on the derived (deleterious) background is compensated by its increase when it falls on the ancestral background, which is relatively more fit. (A similar logic applies when the neutral locus segregates first.) In other words, linked selection should not change the mean of the neutral allele frequency spectrum, instead distorting its outline by increasing mass at both tails (Ewing and Jensen, 2016; Cvijović et al., 2018) – in line with the classic result that linked selection does not alter the rate of fixation of neutral alleles (Birky Jr and Walsh, 1988). This formulation is diametrically opposed from cBGS models which postulate that the reduction in π_R comes from a decrease in N_e, changing the mean of the neutral allele frequency spectrum without touching its outline.

Having built some intuition, we come back to the original question: how does negative selection cause $\mathbb{E}[D(1-2q)] > 0$? Two alternative mechanisms could be in play. First, this distortion could be manifested immediately upon the mutation event that generates two-locus diversity. In this case, selection would influence $\mathbb{E}[D(1-2q)]$ through its "stationary" effect on the site frequency spectrum. Second, even if $\mathbb{E}[D(1-2q)] = 0$ immediately after the second mutation, selection could distort its expectation over the course of time. This could happen if elimination of haplotypes carrying the deleterious allele does not preserve the symmetry of D (the canceling out of negative and positive instances) over the distribution $f(q)$.

We tested whether $\mathbb{E}[D(1-2q)] > 0$ upon the the second mutation event (first hypothesis). Conditioning on the right locus segregating first, then prior to the second mutation we have $f_{aB} = q$ and $f_{ab} = 1-q$. With probability q, the second mutation creates the AB haplotype at frequency $1/2N$, in which case $D = f_{ab}f_{AB} - f_{aB}f_{Ab} = (1-q)/2N$. Otherwise, (probability $1-q$) it creates the Ab haplotype and $D = -q/2N$. Weighing $D(1-2q)$ in both cases by their respective probabilities and integrating over the neutral frequency spectrum $f(q) \propto 1/q$, we have

$$\int_0^1 (q(1-2q)(1-q)/2N + (1-q)(1-2q)(-q/2N))q^{-1}\,dx$$

which yields zero, therefore, immediately after the second mutation event and if the right locus segregate first, $\mathbb{E}[D(1-2q)] = 0$. Similarly, if the left locus segregates first we integrate over the deleterious site frequency

spectrum $f(p)$ (Ewens, 1964) to obtain

$$\int_0^1 (p(1-2q)(1-p)/2N + (1-p)(1-2q)(-p/2N))((1-e^{-4N_e s(1-p)})/((1-e^{-4N_e s})p(1-p)))\,dx$$

which likewise yields zero, thus we reject the first hypothesis. We conclude that selection builds up $\mathbb{E}[D(1-2q)]$ over the course of time.

To understand why, consider again the situation where the right locus segregates first. We have seen that, immediately after the left mutation, $\mathbb{E}[D] = 0$ regardless of q. However, if $q < 1/2$, then selection will mostly act against Ab haplotypes, such that over time D will tend to be positive on the first half of the neutral frequency spectrum. Conversely, if $q > 1/2$, then selection will mostly act against AB haplotypes, such that over time D will tend to be negative on the second half of the neutral frequency spectrum. Both situations lead to $\mathbb{E}[D(1-2q)] > 0$ and we conclude that selection distorts the distribution of D over $f(q)$ by virtue of its effect after two-locus diversity is established.

We confirm these trends in forward-in-time simulations with `twoLocusSim` (Figure A1). These simulations further allow us to pinpoint thar D increases at first, then decreases, whereas $\mathbb{E}[D(1-2q)]$ tends to remain above zero over time (Figure A2).

C Temporal dynamics in the pure two-locus model

We now discuss the behavior of π_R when the focal neutral site is linked to a single constrained site. For simplicity, we describe how B-values evolve in two-epoch demographic models (10-fold bottleneck or expansion). Patterns from this pure two-locus model will help us interpret multi-locus scenarios (Figure 2, 3).

A comprehensive description of the temporal dynamics at the neutral locus requires understanding the behavior of the linked constrained locus. Therefore, we start by analyzing the trajectory of π_L as a function of its selection coefficient and demographic model. This corresponds to a single-site model whose behavior was described by Otto and Whitlock (1997) and has led to DFE inference methods based on the frequency spectrum of (independent) constrained sites (Boyko et al., 2008; Keightley and Eyre-Walker, 2007; Kim et al., 2017). Although not novel, the prediction of π_L is part of the `moments++` framework and we revisit it here to aid the interpretation of B-value dynamics.

We track the expected value of π_L for 120,000 generations following a population size change. As expected, the patterns are markedly different between bottleneck and expansion scenarios (Figure S8). Focusing on the expansion first, we see that if selection in the ancestral population is strong ($2N_e s = -20$, Figure S8A), then π_L increases ever so slightly at first but pulls back to the previous equilibrium point. This happens because the steady-state in the ancestral population is already dominated by mutation-selection balance ($\pi_L \approx u/s$) such that a 10-fold increase in N_e can only move the population deeper into the same regime. On the other hand, if the ancestral population experiences weak-to-moderate selection (Figure S8B, $2N_e s = -2$), a 10-fold expansion brings it close to mutation-selection-balance. After an initial uptick in π_L caused by the reduced coalescence rate, selection starts purging the excess of deleterious variants. Steady-state lies intermediate between ancestral π_L and such peak. Finally, if the ancestral population finds itself in the nearly-neutral regime (Figure S8F, $2N_e s = -0.2$), where the equilibrium π_L is relatively close to the neutral expectation ($\pi_0 = 4N_e\mu$), a 10-fold expansion will introduce weak-to-moderate selection. π_L will rise smoothly, driven by the 10-fold increase in the scaled mutation rate, but eventually plateau at a point distinguishably different from π_0, as selection becomes a relevant force.

The temporal dynamics of π_L are very different in the bottleneck case. If selection in the ancestral population is not strong ($2N_e s = -2$ or $2N_e s = -0.2$, Figure S8B, C), π_L decays monotonically and quickly reaches the new equilibrium. However, if selection is strong enough ($2N_e s = -20$, Figure S8A), a quick drop in π_L is followed by partial recovery towards an intermediate equilibrium point. This happens because the population suddenly leaves mutation-selection balance to enter the mutation-drift-selection regime. Deleterious mutations are now met with selection that is effectively weaker than in the ancestral population, thus π_L upticks after the initial drop. Note that in both demographic models the absolute change in π_L under strong selection (y-axes in Figure S8A) is quite low. In fact, in the expansion case the changes are negligible (as

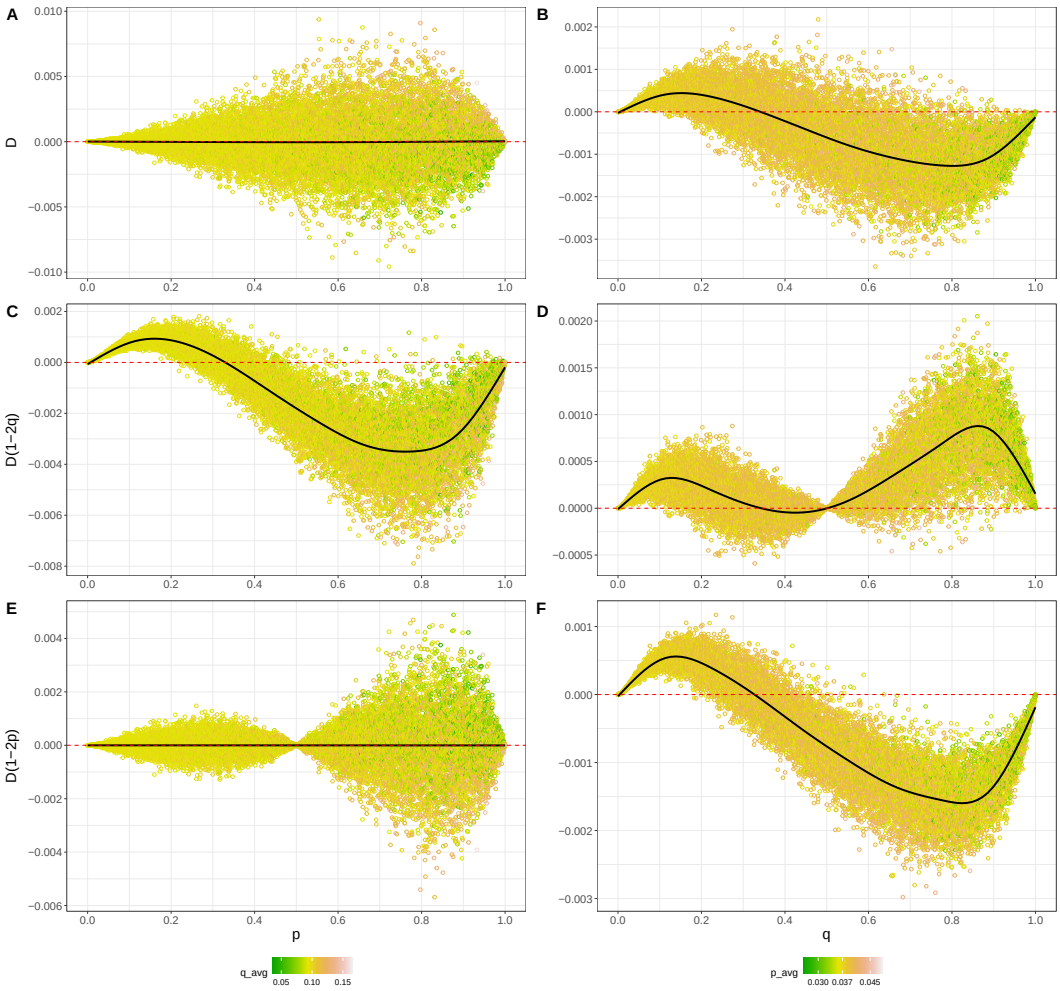

Figure A1: D statistics simulated with `twoLocusSim`. In each of 12,000 replicates, independent two-locus systems evolved for 10^7 generations (sequence length $L = 10$) with $\mu = 10^{-6}$, $r = 10^{-4}$ and $s = -10^{-4}$. Each point represents the average over all generations (after a burn-in period of 400,000 generations), conditioning on a particular frequency of p (A, C, E) or q (B, D, F). Black line shows the loess smoothing curve. Note that D is flat (in expectation) over $f(p)$ (A), but not over $f(q)$ (B).

predicted by theory), however, in the bottleneck case the modest non-monotonic behavior of π_L promotes interesting dynamics of diversity at the linked neutral site, which we describe next.

Population expansions lead to lower equilibrium B-values relative to the ancestral population (Figure S9, S10). When selection in the ancestral population is already strong ($2N_A s = -20$), the B-value increases for many generations before starting to decline. This initial uptick means that the rate of change of π_R is, at first, faster than that of π_0. Eventually, the B-value returns to the previous steady-state, as expected from cBGS predictions. On the other extreme, if deleterious variants behave nearly-neutrally in the ancestral population ($2N_e s = -0.2$), a 10-fold expansion will send them to the weak-to-moderate regime. Driven by the increase in N_e, overall levels of diversity increase, but the higher π_L renders the increase in π_R smaller than that of π_0 and the B-value goes straight down. Finally, when $2N_e s = -2$ a 10-fold expansion imposes efficient selection. For the first $\sim N_e$ generations, the B-value curve resembles the mirror image of π_L, decreasing at first and then creeping back up. Due to the increased deleterious diversity, the B-value then slowly drops below its

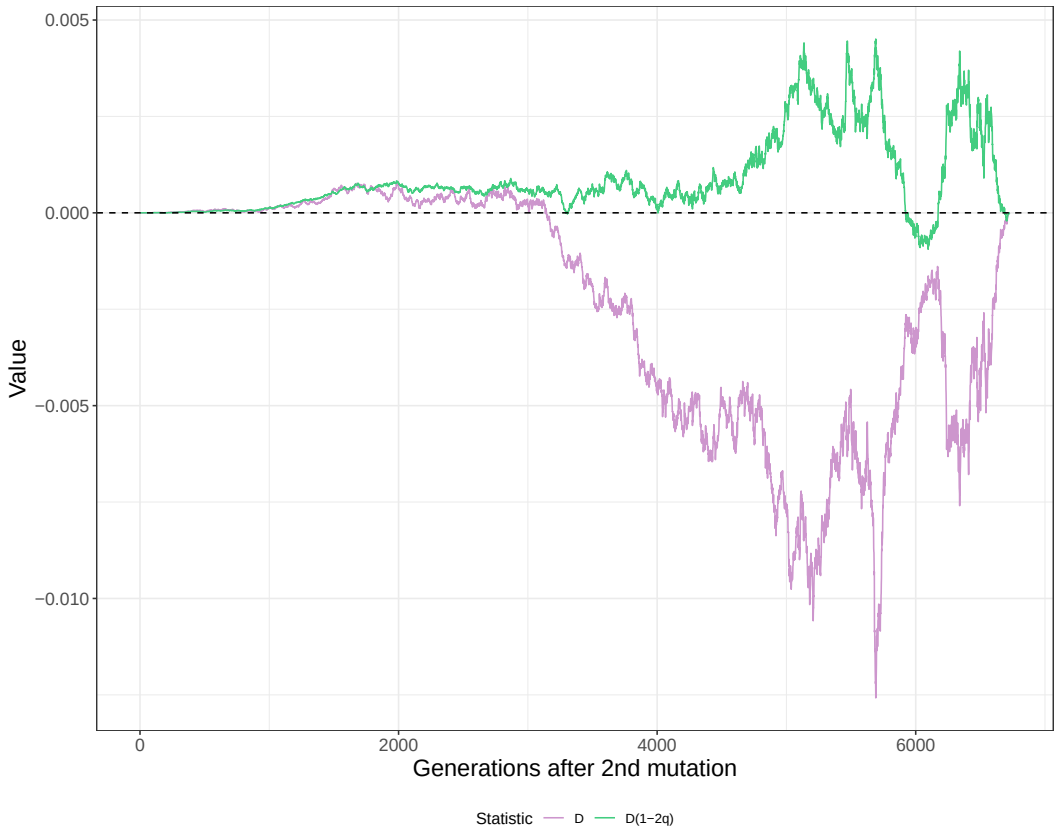

Figure A2: The temporal dynamics of D statistics after two-locus diversity is created. In each of 12,000 replicates, independent two-locus systems evolved for 10^9 generations (sequence length $L = 1$) with $N_e = 10,000$, $\mu = 10^{-8}$, $r = 10^{-8}$ and $s = -10^{-3}$. Shown are the average statistics over all replicates, conditioning on two-locus diversity reaching a particular generation (i.e., no fixation or loss). Therefore, since the number of simulations still segregating decreases over time, noise in these estimates increases and becomes substantial after $\sim 5,000$ generations.

ancestral steady-state. Overall, these dynamics are quantitatively dependent on the recombination rate, where strongly deleterious mutations reduce B-values over long genetic distances while weakly deleterious mutations have a much more localized effect.

Population bottlenecks lead to higher equilibrium B-values relative to the ancestral population (Figure S9). Since these new equilibria do not depend on the slow supply of new mutations, they are attained much faster than in the expansion scenarios. When the ancestral population is under weak and moderate selection the increase is monotonic whereas the strong selection scenario reveals an initial drop in B, followed by recovery and beyond. In the latter case, the drop in π_R is initially more pronounced than in π_0 because the approach to the new equilibrium happens faster in the presence of selection (see subsection "Modeling strong selection under non-equilibrium demography" and Torres et al. (2020)). Then, as π_L recovers but settles at an intermediate point (Figure S8), linked selection is relaxed relative to the ancestral population, and the B-value finds an elevated steady-state. In all selection regimes, the relationship between the equilibria points for B and π_L is intuitive: the first rises whenever the second drops. Under a 10-fold bottleneck, it is the stronger selection regime that causes non-monotonic dynamics. As in the expansion scenarios, recombination primarily modulates the magnitude of linked selection while preserving the shape of temporal dynamics.

The results from two-epoch models support that equilibrium B-values correlate negatively with N_e (Nordborg et al., 1996; Torres et al., 2020). (This explains why interference – approximated as a reduction in

the N_e of constrained elements – always weakens linked selection effects). Indeed, this trend holds as long as the population is not under mutation-selection balance, in which case π_L is dictated by the ratio μ/s and independent of N_e, as predicted by theory. The same logic applies under non-equilibrium demography, but the distinct evolutionary pace of π_0, π_R and π_L shapes the temporal trajectory of B-values. Their transient behavior can be non-monotonic and last for several hundred generations, longer than most populations remain static in nature.

Supplemental Figures

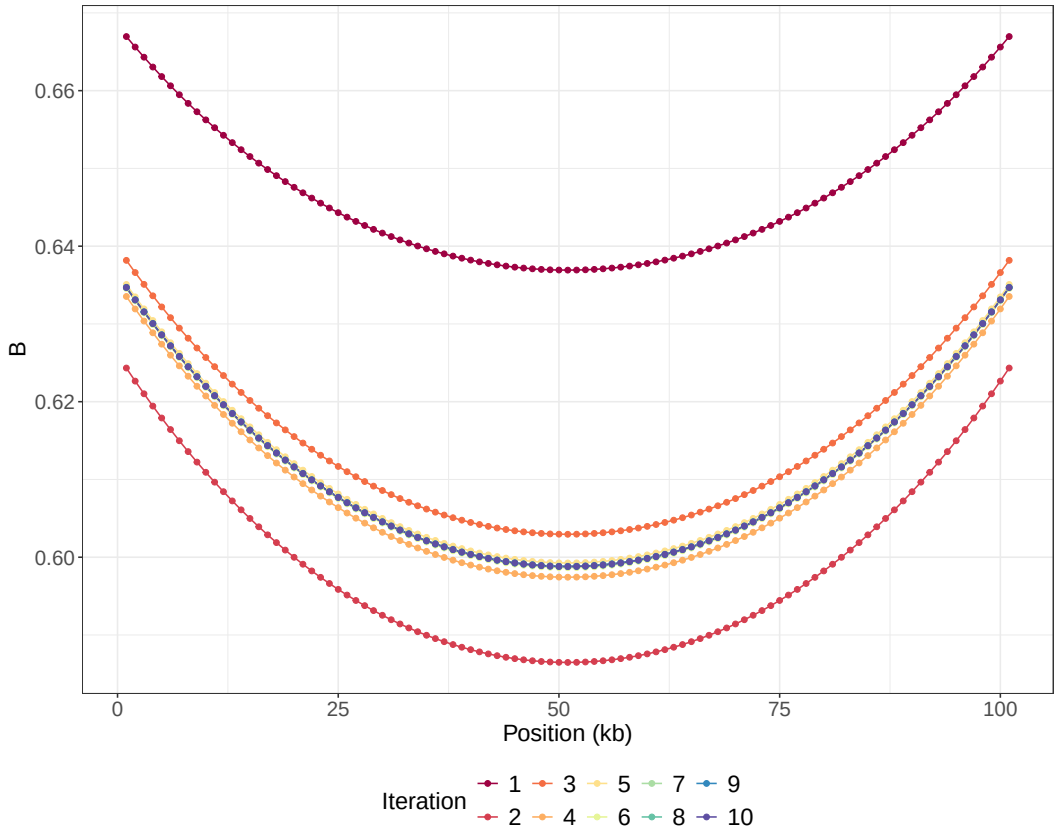

Figure S1: Predicted B-maps along a 100 kb segment, after each round of interference correction (color). Note how the iterative re-scaling of μ, r, and s leads to B-maps bouncing around the final values (convergence happens visually after the 6th iteration). Here the 100 kb segment layout follows that of Figure 2

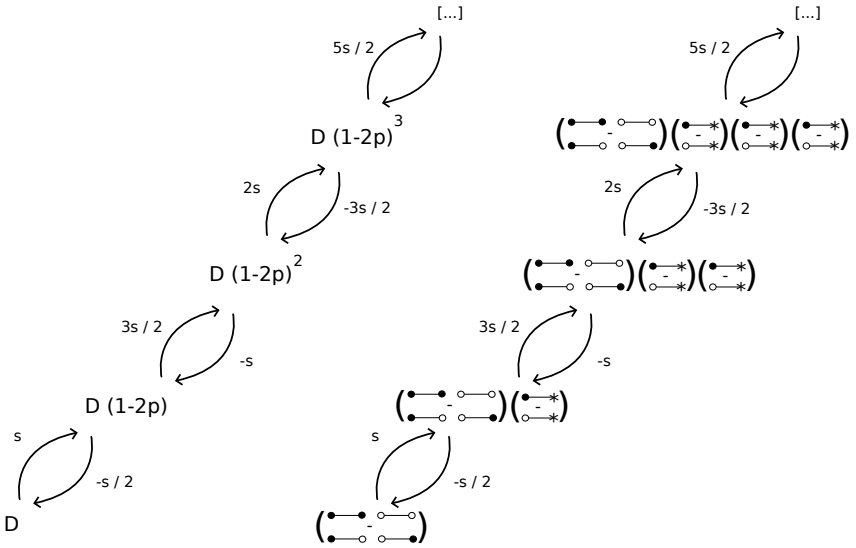

Figure S2: Representation of the dependencies among $D(1-2p)^j$ statistics imposed by the selection operator in (p, q, D)-space (left) and the equivalent haplotype space (right). Arrows denote "collects from", with the respective entry in the selection matrix shown above ($s < 0$). Haplotype subtractions proceed from top to bottom of the configuration within parentheses. Asterisks denote invariance to the allelic state at the right locus $(1-2p)$.

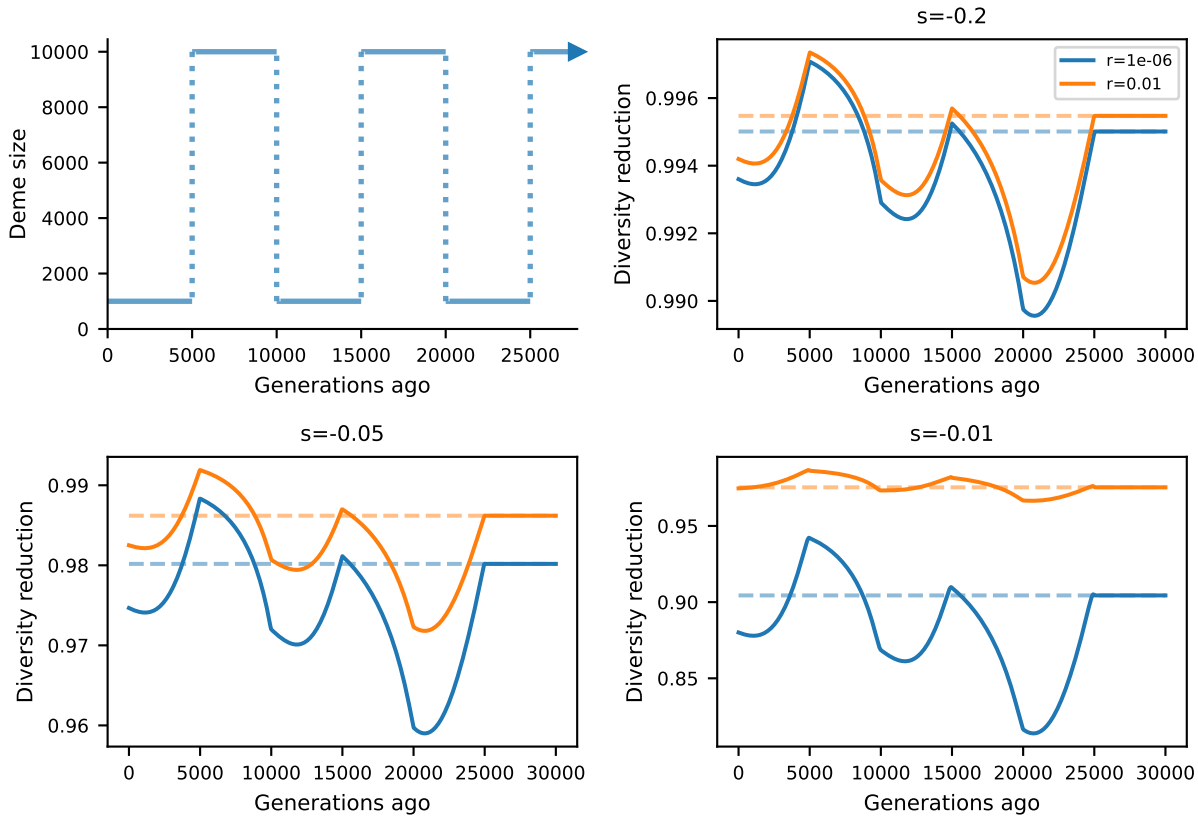

Figure S3: Temporal dynamics of BGS due to strong selection, as predicted using our extension to Nordborg (1997). Top left panel shows the demographic history considered. Each other panel shows, for a given selection coefficient, the predicted B-value trajectory for either tightly linked ($r = 10^{-6}$, blue lines) and loosely linked loci ($r = 10^{-2}$, orange lines). Dashed lines indicate the respective steady state solutions. Here the mutation rate is set to $u = 10^{-3}$.

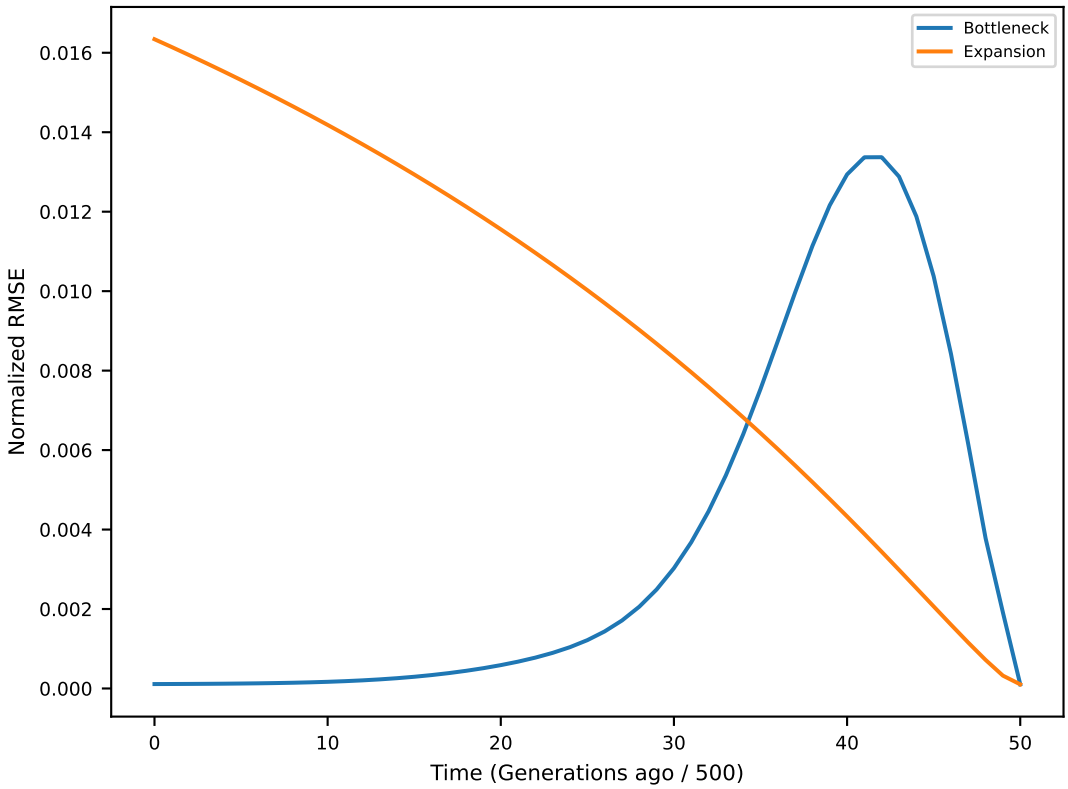

Figure S4: Temporal evolution of the the Normalized Rooted Mean Squared Error between the demography-aware B-map and steady-state B-maps, predicted every 500 generations after the size change. Blue and orange lines show results from a 10-fold bottleneck and expansion, respectively. Steady-state B-maps are obtained using the equilibrium N_e value calculated from neutral genetic diversity in the absence of linked selection (equivalent to the harmonic mean of population sizes in the relevant time-frame). Here the chromosome layout follows that of Figure 3.

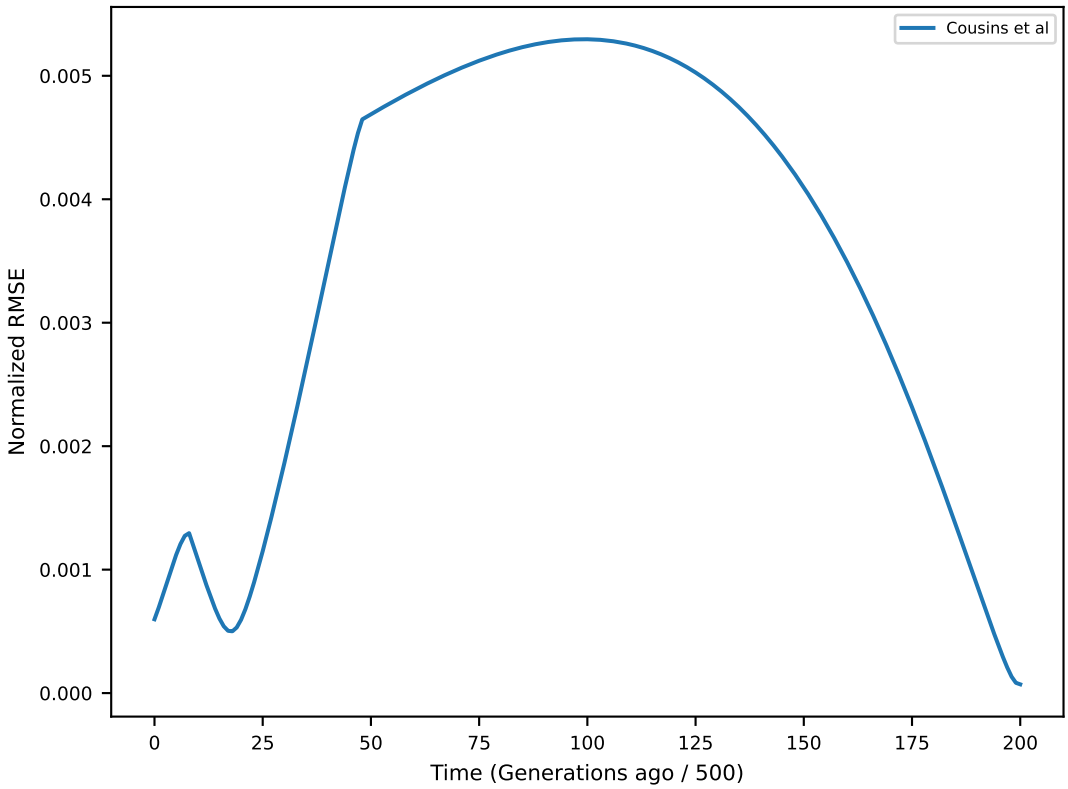

Figure S5: Temporal evolution of the the Normalized Rooted Mean Squared Error between the demography-aware B-map and steady-state B-maps, predicted every 500 generations after the first size change in the Cousins et al. (2024) demography. Steady-state B-maps are obtained using the equilibrium N_e value calculated from neutral genetic diversity in the absence of linked selection (equivalent to the harmonic mean of population sizes in the relevant time-frame). Here the chromosome layout follows that of Figure 3.

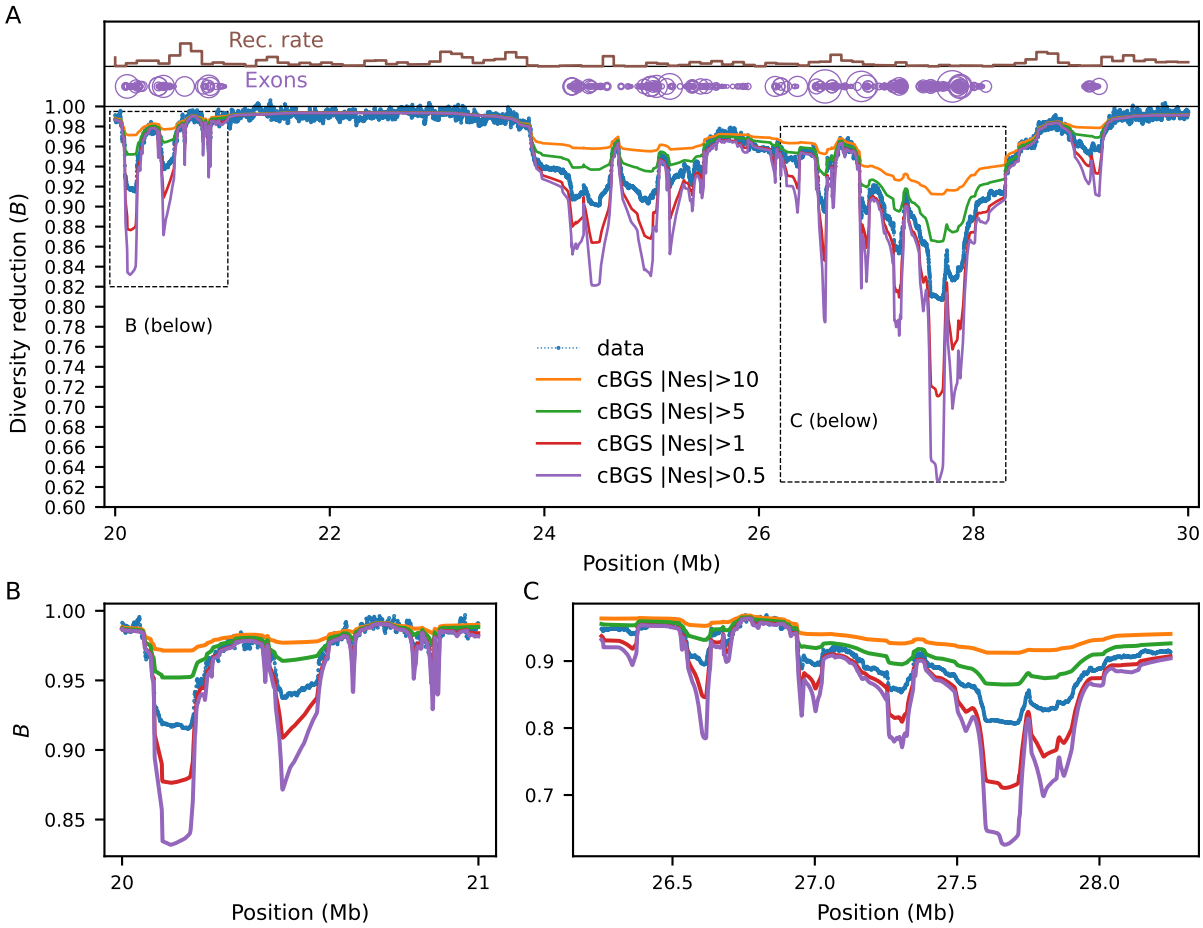

Figure S6: Predictions from cBGS, truncating the DFE at different points. Notice the overestimation of B-values caused by excluding a large portion of the DFE (orange and green lines) as well as the underestimation caused by the inappropriateness of the cBGS model at weaker selection coefficients (red and purple lines). Here the chromosome layout follows that of Figure 3.

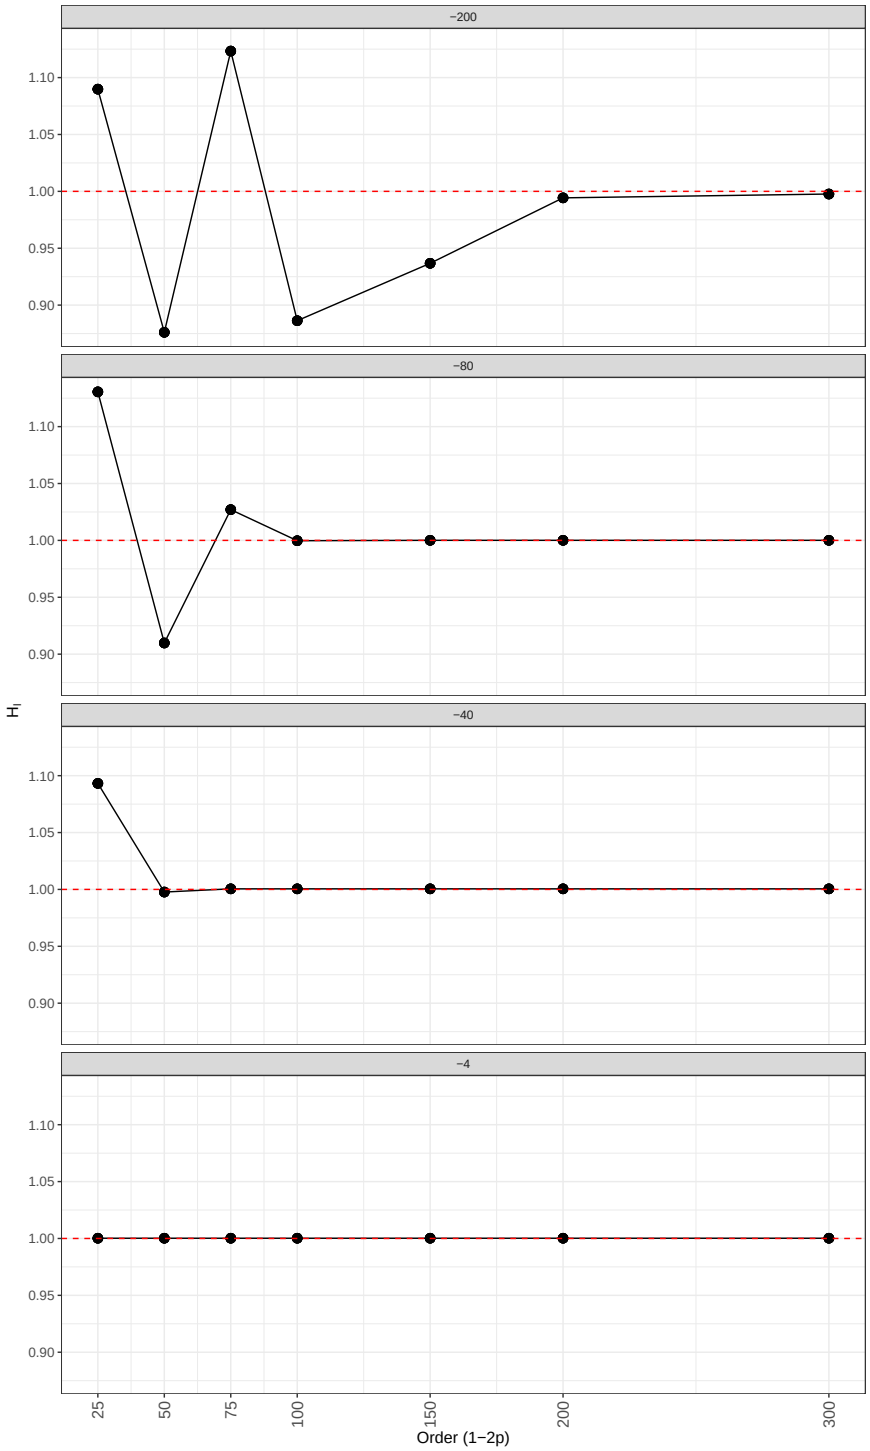

Figure S7: Benchmarking of the implementation of the selection operator in `moments++`, with the truncation strategy to close the system. Shown are ratios of π_l predicted with `moments++` to π_l predicted with `moments.TwoLocus`, as a function of the Order of $1-2p$ factors included in the `moments++` model (panels show $N_e s$). Dashed red lines denote the target line indicating good agreement with our gold-standard.

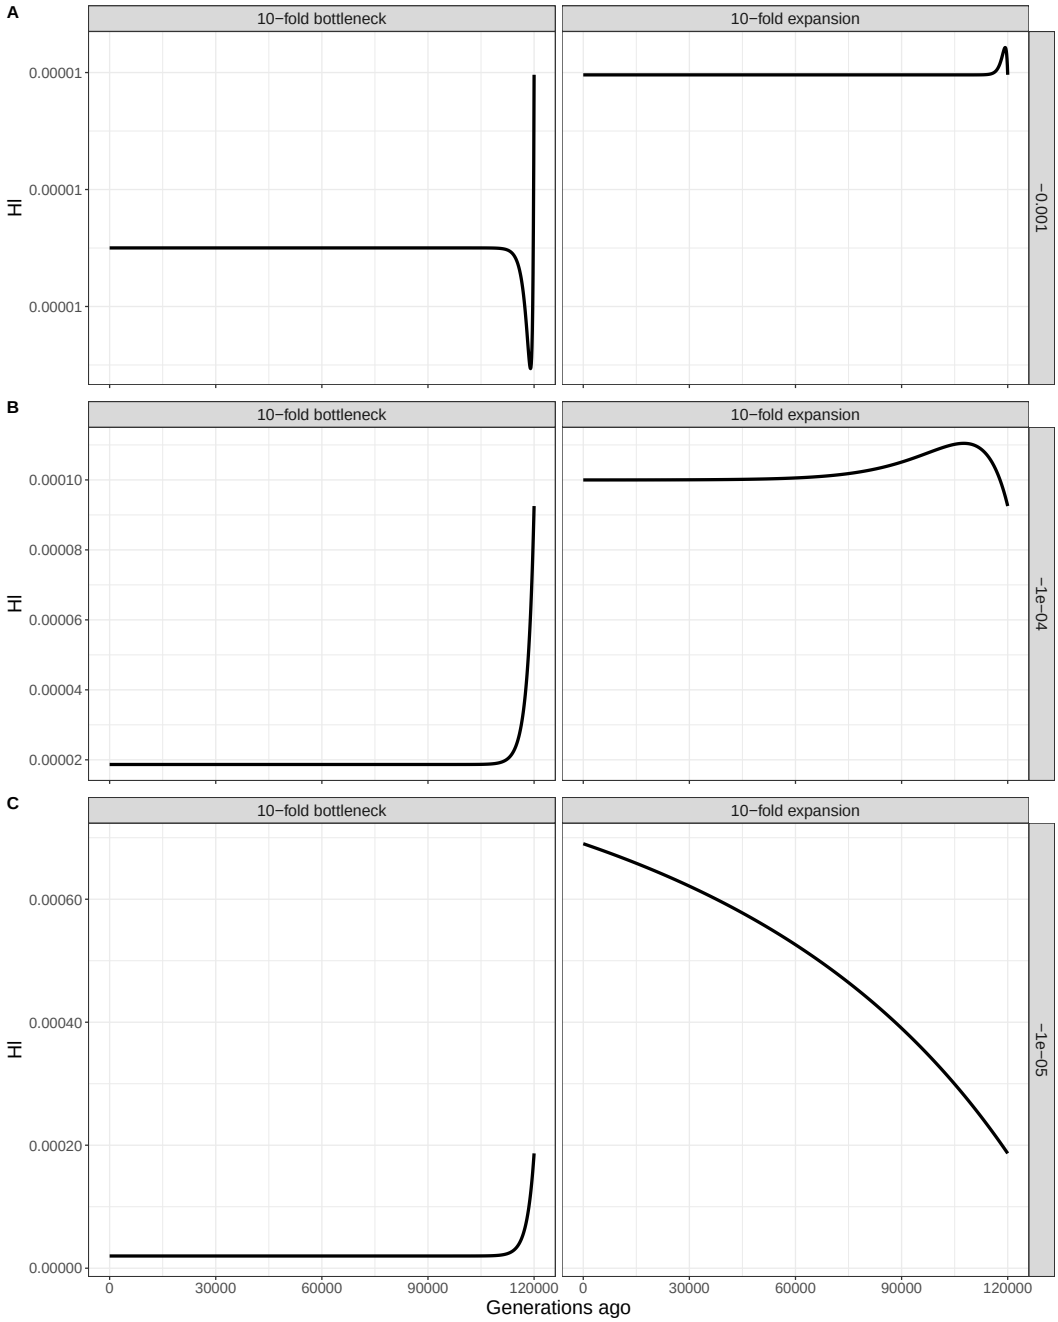

Figure S8: Temporal trajectories of π_L following a 10-fold bottleneck (left panels) or expansion (right panels). A) $s = -0.001$. B) $s = -0.0001$. C) $s = -0.00001$. Here $\mu = 10^{-8}$.

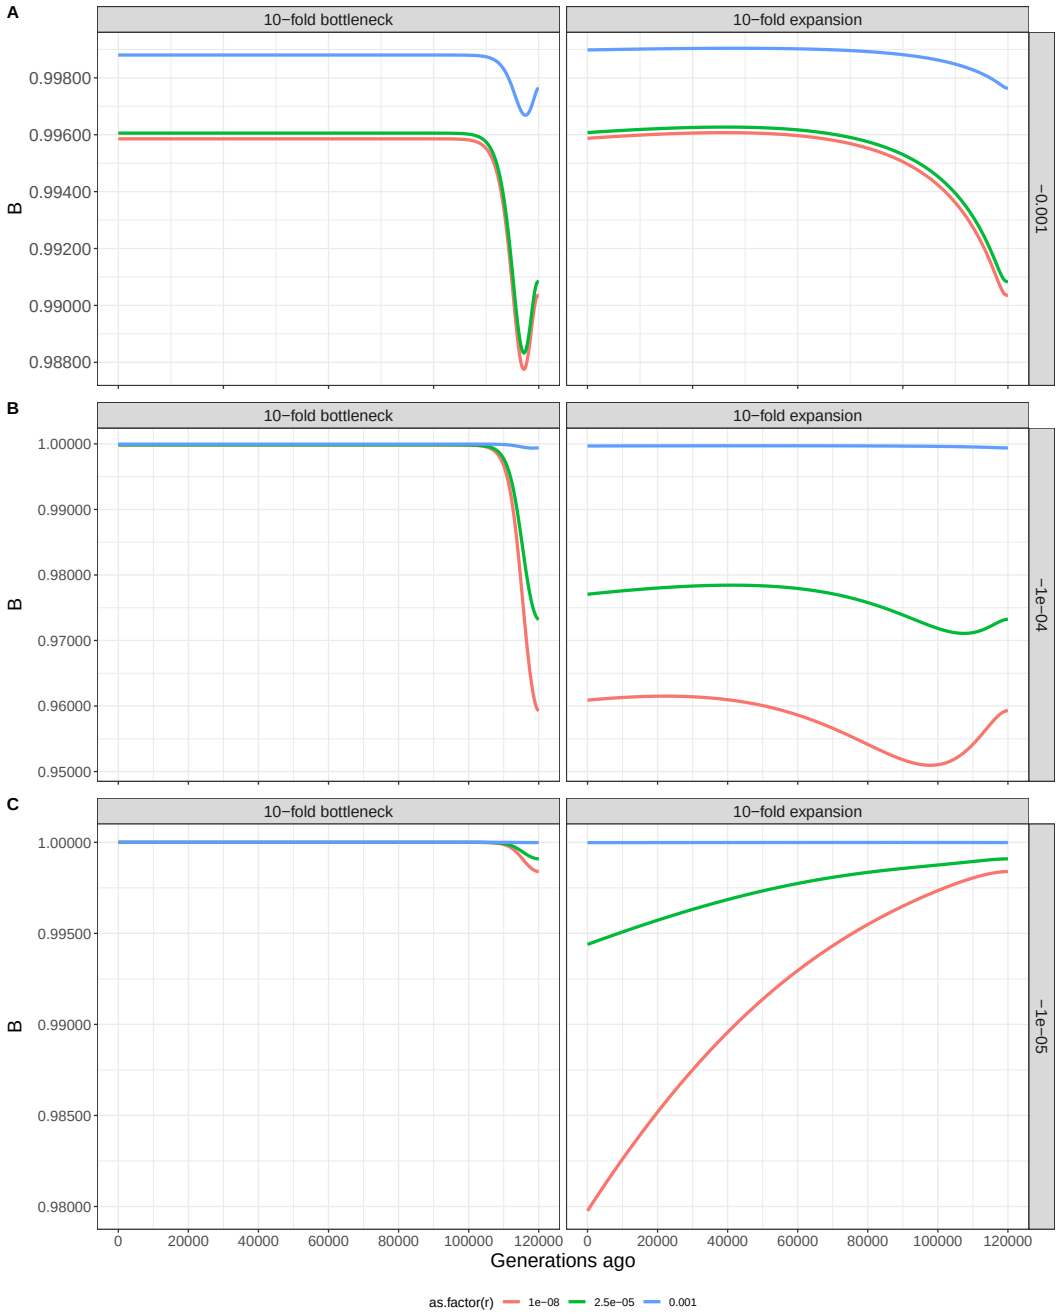

Figure S9: Temporal trajectories of B-values following a 10-fold bottleneck (left panels) or expansion (right panels), for a pure two-locus model. Note how the expansion scenario does not reach the new equilibrium after 120,000 generations. A) $s = -0.001$. B) $s = -0.0001$. C) $s = -0.00001$. Colors denote the rate of recombination. Here $\mu = 10^{-8}$ and B-values are raised to the 1000th power to facilitate visualization.

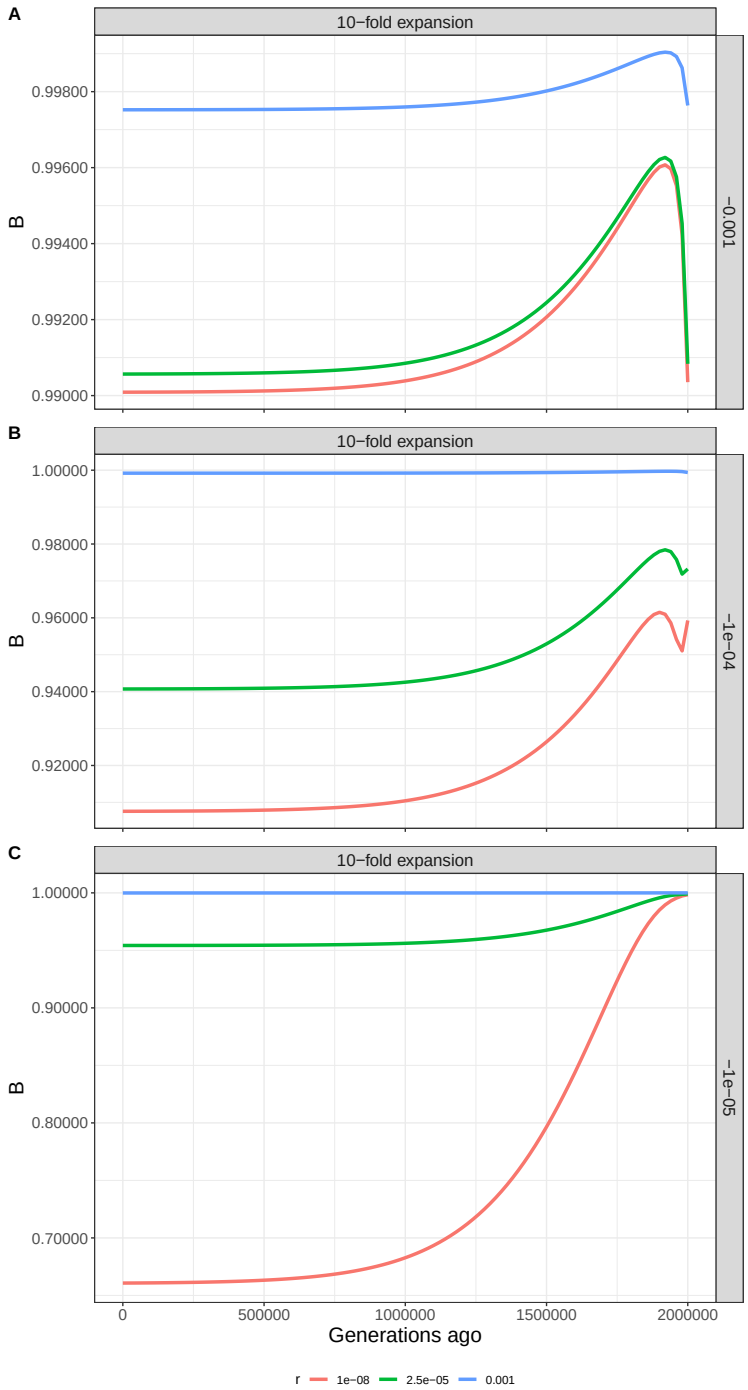

Figure S10: Temporal trajectories of B-values under the 10-fold expansion model from Figure S9 where we track statistics for 2,000,000 generations to ensure that the new equilibrium point is found. A) $s = -0.001$. B) $s = -0.0001$. C) $s = -0.00001$. Colors denote the rate of recombination. Here $\mu = 10^{-8}$ and B-values are raised to the 1000th power to facilitate visualization.

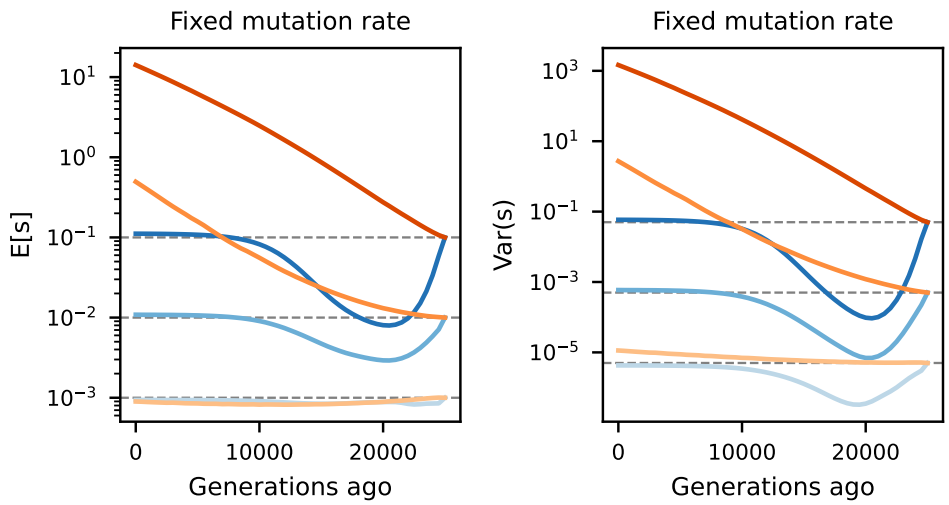

Figure S11: Mean (left) and variance (right) of the inferred DFEs depicted in Figure 4.

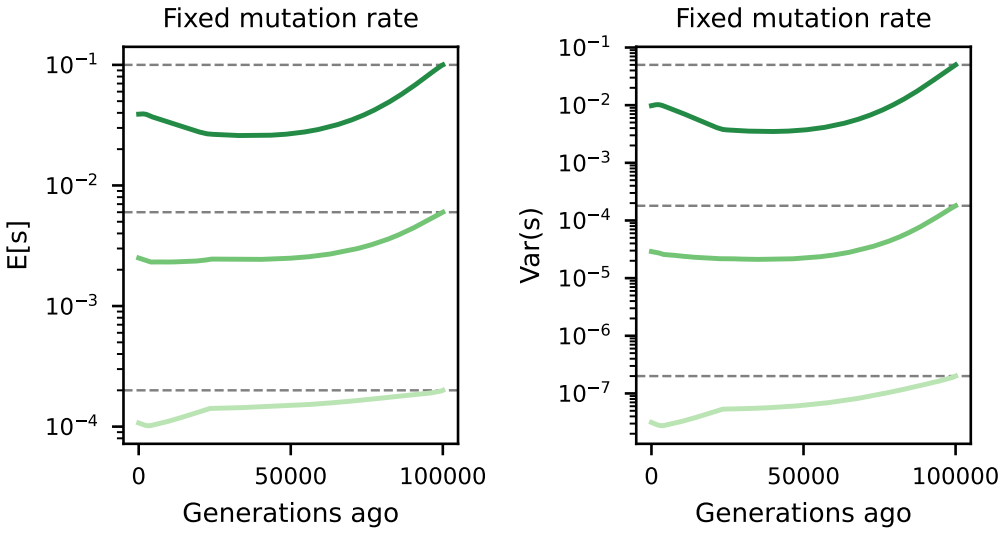

Figure S12: Mean (left) and variance (right) of the inferred DFEs depicted in Figure 5.