GENETICS, 2025, 229(4), iyae210
https://doi.org/10.1093/genetics/iyae210
Advance Access Publication Date: 18 December 2024
Genomic Data Analyses in Biobanks

OXFORD GENETICS

Characterizing selection on complex traits through conditional frequency spectra

Roshni A. Patel [1,2,*], Clemens L. Weiß [3], Huisheng Zhu [4], Hakhamanesh Mostafavi [5,6], Yuval B. Simons [7], Jeffrey P. Spence [1,*], Jonathan K. Pritchard [1,4,*]

[1]Department of Genetics, Stanford University, Stanford, CA 94305, USA
[2]Department of Quantitative and Computational Biology, University of Southern California, Los Angeles, CA 90089, USA
[3]Stanford Cancer Institute Core, Stanford University, Stanford, CA 94305, USA
[4]Department of Biology, Stanford University, Stanford, CA 94305, USA
[5]Center for Human Genetics and Genomics, New York University School of Medicine, New York, NY 10016, USA
[6]Division of Biostatistics, Department of Population Health, New York University School of Medicine, New York, NY 10016, USA
[7]Department of Medicine, University of Chicago, Chicago, IL 60637, USA

*Corresponding author: Department of Quantitative and Computational Biology, University of Southern California, Los Angeles, CA 90089, USA. Email: roshni@usc.edu; *Corresponding author: Department of Genetics, Stanford University, Stanford, CA 94305, USA. Email: jspence@stanford.edu; *Corresponding author: Departments of Genetics and Biology, Stanford University, Stanford, CA 94305, USA. Email: pritch@stanford.edu

Natural selection on complex traits is difficult to study in part due to the ascertainment inherent to genome-wide association studies (GWAS). The power to detect a trait-associated variant in GWAS is a function of its frequency and effect size – but for traits under selection, the effect size of a variant determines the strength of selection against it, constraining its frequency. Recognizing the biases inherent to GWAS ascertainment, we propose studying the joint distribution of allele frequencies across populations, conditional on the frequencies in the GWAS cohort. Before considering these conditional frequency spectra, we first characterized the impact of selection and non-equilibrium demography on allele frequency dynamics forwards and backwards in time. We then used these results to understand conditional frequency spectra under realistic human demography. Finally, we investigated empirical conditional frequency spectra for GWAS variants associated with 106 complex traits, finding compelling evidence for either stabilizing or purifying selection. Our results provide insights into polygenic score portability and other properties of variants ascertained with GWAS, highlighting the utility of conditional frequency spectra.

Keywords: allele frequency spectra; selection; complex traits; GWAS; polygenic scores

Introduction

Over the last two decades, genome-wide association studies (GWAS) have uncovered countless associations between genetic variants and complex traits. In the process, it has become apparent that the effect sizes of trait-associated variants are inversely correlated with their frequencies (Yang J *et al.* 2010, 2015; Speed *et al.* 2012). This relationship has been interpreted as a sign that natural selection acts on complex traits (Zeng *et al.* 2018, 2021; Schoech *et al.* 2019). Furthermore, patterns of diversity around trait-associated variants appear to deviate from neutral expectations (Gazal *et al.* 2017; Speed *et al.* 2020).

To explain these observations, different models of selection on complex traits have been proposed. Several groups have proposed models of stabilizing selection, wherein intermediate trait values are favored (Lande 1976; Turelli 1984; Simons *et al.* 2018; Zhang *et al.* 2023). Stabilizing selection on a trait results in selection against the minor allele at independent genetic variants underlying the trait (Robertson 1956; Simons *et al.* 2018; Walsh and Lynch 2018). In other words, the direction of selection at individual alleles depends on the frequency of the allele. In contrast, others have proposed a model of purifying selection, which entails selection against all new mutations (Keightley and Hill 1990; Eyre-Walker 2010; Caballero *et al.* 2015; Zeng *et al.* 2021). Under such models, the direction of selection is independent of allele frequency. Finally, some have proposed models of directional selection on complex traits, where the direction of selection at individual alleles depends on the sign of their effects (Guo *et al.* 2018).

Distinguishing between competing models of selection is essential for interpreting the results of GWAS. The mode and strength of selection acting on traits determines whether GWAS tend to ascertain variants with large or small effects. Moreover, because selection influences the genetic architecture of traits at the molecular level, selection can also shape the extent to which discovered variants are pleiotropic and further upstream in regulatory networks, or more specific and further downstream in regulatory networks (Mostafavi *et al.* 2023). Naturally, this has consequences for the proportion of heritability explained by discovered variants (O'Connor *et al.* 2019; Weiner *et al.* 2023), as well as the portability of polygenic scores, or the extent to which a polygenic score derived from a particular GWAS cohort is predictive in genetically dissimilar individuals (Wang *et al.* 2020; Durvasula and Lohmueller 2021; Yair and Coop 2022). Thus, a clear understanding of selection on complex traits is not only essential for studying complex trait evolution, but also for understanding complex trait architecture more broadly.

Received on 24 September 2024; accepted on 03 December 2024

Selection on complex traits has been challenging to investigate in part because of the technical limitations of GWAS. First, the power to ascertain a trait-associated variant in GWAS depends on the variant's frequency and effect size; GWAS are more powered to ascertain variants with large minor allele frequencies and large effect sizes. For traits under selection, the effect size of a variant determines the strength of selection on the variant, which in turn constrains its frequency. This logic also applies to traits that are genetically correlated with traits under selection. Trait-associated variants ascertained in GWAS are therefore enriched for common variants and thought to be depleted for the most strongly selected variants (Pritchard 2001; Manolio et al. 2009). Second, GWAS generally rely on genotype arrays and imputation panels to infer genotypes genome-wide. By implicitly prioritizing a subset of genetic variation, arrays and imputation algorithms will bias discovered variants in a manner that is particularly difficult to model (Clark et al. 2005; Lachance and Tishkoff 2013; Kim et al. 2018). As a result, the process of GWAS ascertainment can obscure the hallmark signatures of selection.

To avoid certain biases induced by GWAS ascertainment, we propose conditioning on the allele frequency in the GWAS cohort by leveraging additional, diverged populations. These additional populations provide orthogonal sources of information about selection. Moreover, conditioning on the frequency in the GWAS cohort effectively conditions on GWAS ascertainment itself. Yet, in order to combine information across populations, we need to develop a framework for describing how selection and non-equilibrium demography impact the distribution of allele frequencies in one population conditional on another.

A wealth of previous work has used population genetics theory to answer questions at the intersection of selection, demography, and (joint) allele frequency trajectories. Maruyama (1974) derived theoretical expectations for the impact of selection on allele frequencies within a single population under equilibrium demography. More recently, others have characterized the impact of non-equilibrium demographies on both neutral variation (Myers et al. 2008; Bhaskar and Song 2014; Bhaskar et al. 2015; Terhorst and Song 2015; Spence et al. 2016; Baharian and Gravel 2018; Ragsdale et al. 2018; Rosen et al. 2018; Schraiber 2018) and selected alleles within a single population (Slatkin 2001; Evans et al. 2007; Živković and Stephan 2011; Song and Steinrücken 2012; Chen and Slatkin 2013; Schraiber et al. 2016; Ortega-Del Vecchyo et al. 2022). Consequently, we have a solid theoretical expectation for how selection and demography impact the allele frequency spectrum within one population. There has also been a line of work exploring how these two evolutionary forces impact the joint frequency spectrum for two or more populations (Gutenkunst et al. 2009; Lukić et al. 2011; Lukić and Hey 2012; Yang MA et al. 2014; Jouganous et al. 2017; Kamm JA et al. 2017; Kern and Hey 2017; Kamm J et al. 2020; Dilber and Terhorst 2024).

However, the effect of conditioning on frequencies in one population has been under-explored. We propose conditioning on frequencies in a GWAS cohort as a means of studying selection on complex traits, but previous work on conditional frequency spectra has focused on neutral variants, not selected variants (Chen et al. 2007; Harpak et al. 2016; Durvasula and Sankararaman 2020). While it is known that different models of selection leave distinct signatures on joint frequency spectra, it is unclear if this also holds for conditional frequency spectra. Conditional frequency spectra can be understood by considering selection both forwards and backwards in time: first backwards from the conditional population to a shared ancestor, then forwards to the second, unconditioned population. The impact of selection on allele frequency trajectories backwards in time is poorly understood, especially under non-equilibrium demographies.

Thus, we sought to characterize conditional frequency spectra across modern human populations under different models of selection. Through our theoretical analyses of conditional frequency spectra, we identified strong qualitative signatures associated with each model of selection. We then used 106 complex trait GWAS to investigate empirical conditional frequency spectra at trait-associated variants and assess evidence for different models of selection on complex traits. Finally, we used this framework to understand how selection and non-equilibrium demography impact the portability of polygenic scores. Our results highlight the utility of conditional frequency spectra as a tool for studying selection on complex traits.

Results

GWAS ascertainment biases allele frequency spectra

We first sought to demonstrate the impact of GWAS ascertainment on the allele frequency spectrum. We began by curating a set of variants associated with 106 complex traits and diseases. We obtained 18,229 trait-associated variants from published GWAS of 94 quantitative traits in the UK Biobank "White British" cohort, approximately 337,000 individuals (Simons et al. 2022). We additionally obtained a set of 1,040 trait-associated variants from published GWAS of 12 complex diseases in various cohorts with European ancestries (Supplementary Table S1).

We compared the derived allele frequencies of trait-associated variants to those of approximately 11 million non-associated variants imputed in the UK Biobank. As expected, we found that the frequencies of trait-associated variants are enriched for common variation, reflecting the bias induced by the process of ascertainment (Fig. 1a). The tremendous discrepancy between the frequencies of trait-associated variants and the true allele frequency spectrum emphasizes the utility of conditional frequency spectra (Fig. 1b). Similar to joint frequency spectra, conditional frequency spectra can encode information about demography and natural selection. Unlike joint frequency spectra, conditional frequency spectra can circumvent the biases induced by GWAS ascertainment by conditioning on the frequencies of trait-associated variants in the GWAS cohort.

Characterizing allele frequency dynamics over time

To build intuition for the allelic dynamics captured by conditional frequency spectra, we characterized the impact of demography and selection on allele frequencies over time. We considered four different demographic models relevant for human populations: equilibrium demography (i.e. constant population size), bottleneck, exponential population growth, and bottleneck followed by exponential growth. We also considered the following three types of selection: purifying selection against new alleles, directional selection on traits, and stabilizing selection on traits. These three modes of selection can be well-approximated by different allelic dynamics. Purifying selection acts as negative genic selection against derived alleles; directional selection at the trait level acts as positive genic selection for trait-improving alleles; and stabilizing selection can be approximated by perfect underdominance where only heterozygotes have reduced fitness (Robertson 1956; Walsh and Lynch 2018).

We first recapitulated well-known results about how selection impacts the trajectory of an allele beginning at a particular frequency. We simulated allele frequencies under a Wright–Fisher

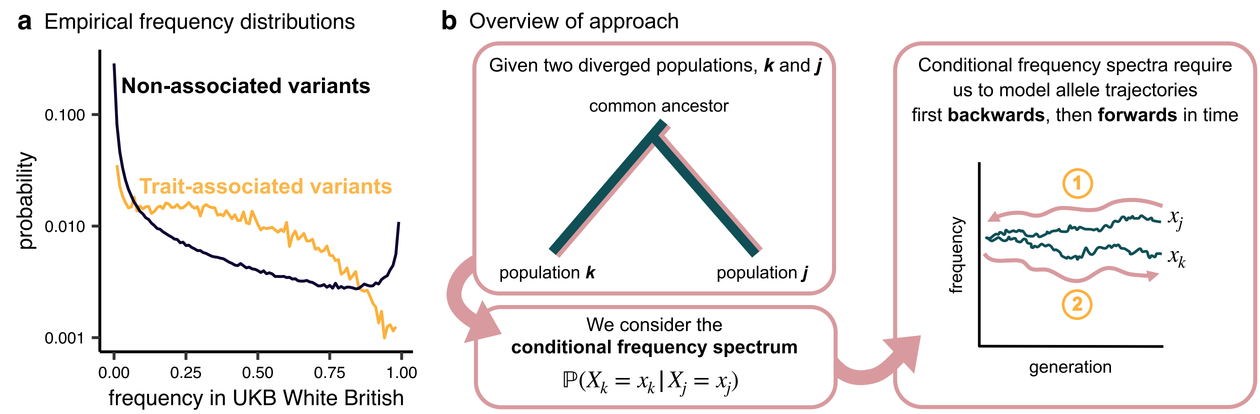

Fig. 1. Overview. a) Derived allele frequency distribution of 19,269 trait-associated variants and approximately 11 million non-associated variants in the UK Biobank White British cohort. b) Overview of conditional frequency spectra. Given two populations k and j, the conditional frequency spectrum $\mathbb{P}(X_k = x_k \mid X_j = x_j)$ requires us to first consider allele trajectories backwards in time from x_j to the frequency in the common ancestor of populations k and j; then ultimately forwards in time to x_k.

model for 2,000 generations and a constant population size with $N_e = 10{,}000$. Regardless of an allele's frequency, positive selection and negative selection consistently act to increase and decrease frequencies, respectively. In contrast, allelic dynamics under stabilizing selection exhibit a frequency dependence because stabilizing selection results in selection against the *minor* allele. Alleles at low frequencies (e.g. 0.1) have nearly indistinguishable dynamics under stabilizing selection and negative selection (Fig. 2a), but as the derived allele frequency increases, the dynamics under stabilizing selection gradually diverge from those under negative selection. When alleles are at a frequency of 0.5, dynamics under stabilizing selection resemble those of neutrality (Fig. 2b), and at frequencies well above 0.5, stabilizing selection begins to look similar to positive selection (Fig. 2c).

We next summarized allelic dynamics by computing the probability of an allele transitioning from a given ancestral frequency, x_{anc}, to a given descendant frequency, x_{desc}, $P(X_{desc} = x_{desc} \mid X_{anc} = x_{anc}; s, \eta)$, where s represents the selection scenario and η the demography. We computed these forward-in-time transition probabilities using two different approaches: Wright–Fisher simulations conducted with SLiM (Haller and Messer 2019), and numerically with fastDTWF (Spence *et al.* 2023).

Using these transition probabilities, we computed the expected frequency in the descendant population conditional on the ancestral frequency. Results for our equilibrium model are presented in Fig. 2d, with expected allele frequencies remaining unchanged over time under a neutral model; increasing under positive selection; decreasing under negative selection; and behaving in a frequency-dependent manner under stabilizing selection. For non-equilibrium models, the results are qualitatively consistent (Supplementary Fig. S1), but by shrinking population sizes, bottlenecks increase the effect of genetic drift, thereby reducing the differences between selection and neutrality (Supplementary Fig. S1a, S1c, S1d). Transition probabilities were concordant between SLiM simulations and fastDTWF, though there were some discrepancies at low allele frequencies due to the different assumptions of both methods (Methods, Supplementary Fig. S1a–c). We next considered the impact of selection and demography backwards in time. In other words, given an allele segregating in a descendant population, what do we expect its frequency to have been in an ancestral population? Under equilibrium demography, we found that selected alleles are expected to be at lower frequencies in an ancestral population regardless of the direction of selection. Moreover, conditional on the frequency in the descendant population, the distribution of frequencies in the ancestral population is identical for positive and negative selection (Fig. 3b), consistent with a classic result from Maruyama (1974).

This can be rationalized by thinking of the backward transition probability $P(X_{anc} = x_{anc} \mid X_{desc} = x_{desc}; s, \eta)$ as a posterior probability proportional to the likelihood of transitioning from some ancestral frequency to the descendant frequency, $P(X_{desc} = x_{desc} \mid X_{anc} = x_{anc}; s, \eta)$, weighted by the prior probability of being at that ancestral frequency, $P(X_{anc} = x_{anc}; s, \eta)$ (Fig. 3a; Equation 2). Under equilibrium demography and negative selection, the likelihood is maximized by values of x_{anc} greater than x_{desc} because negative selection decreases frequencies over time. Conversely, under positive selection, the likelihood is maximized by values of x_{anc} less than x_{desc} because positive selection increases frequencies over time. However, the prior distribution differs under negative and positive selection in the opposite way: high derived allele frequencies in the ancestor are much less probable under negative selection compared to positive selection. For equilibrium demographies, these two forces exactly balance, resulting in negative and positive selection having identical posterior distributions (Fig. 3a).

Under non-equilibrium demography, different trends emerged (Supplementary Fig. S2). The demographic models we considered only differ from each other following the split from the ancestral population; thus, while selection impacts both the prior $P(X_{anc} = x_{anc}; s, \eta)$ and the likelihood $P(X_{desc} = x_{desc} \mid X_{anc} = x_{anc}; s, \eta)$, our demographic models only change the likelihood.

In the presence of a bottleneck, the increased strength of genetic drift makes the likelihood flatter, and as a result, the posterior is dominated by the prior, such that the expected frequency in the ancestor does not vary dramatically across descendant frequencies (Fig. 3c, Supplementary Fig. S2c). Specifically, we found that alleles subject to negative and stabilizing selection are likely to be rare in the ancestor regardless of their frequency in the present. We found that a bottleneck followed by exponential growth reverses some of the observed flattening because the increase in population size decreases the strength of genetic drift (Fig. 3d, Supplementary Fig. S2d).

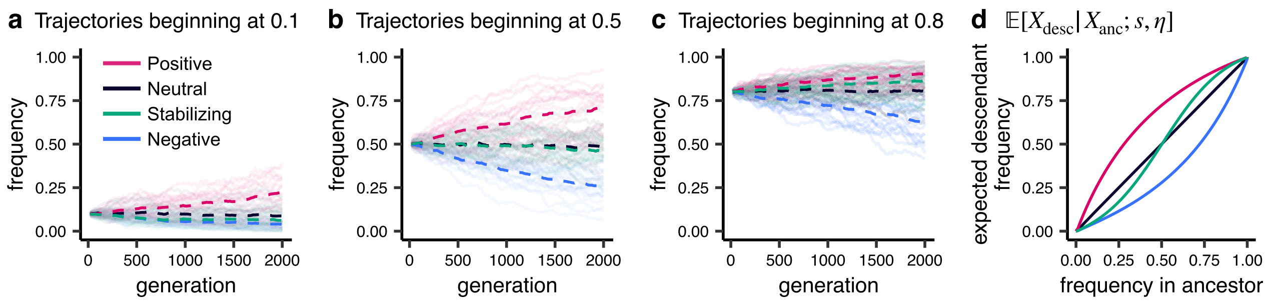

Fig. 2. Forward transitions. a, b, c) Allele frequency trajectories simulated under a Wright–Fisher model. Dashed lines depict the mean frequency across 20 trajectories for each model of selection. d) Expected frequency in a descendant population conditional on the frequency in the ancestral population, computed with fastDTWF. For all panels, the demographic model consists of a constant population size with $N_e = 10{,}000$ for 2,000 generations, and selection coefficients correspond to $|hs| = 5.0 \times 10^{-4}$.

Under exponential growth alone, genetic drift is much weaker, meaning the likelihood is sharper than it is at equilibrium, and the posterior is dominated by the likelihood instead of the prior. As a result, alleles subject to negative selection are expected to be at higher frequencies in the ancestral population relative to alleles evolving neutrally, while alleles subject to positive selection are expected to be at lower frequencies (Fig. 3e, Supplementary Fig. S2e). We also varied the strength of the bottleneck and the degree of exponential growth, finding qualitatively similar results (Supplementary Fig. S3).

Selection and the out-of-Africa demographic model

Having developed intuition for forward and backward allelic dynamics under simple demographies, we next considered their implications for realistic human demographies. We generated allele frequency distributions under a demographic model inferred from Han Chinese in Beijing, China (CHB); Yoruba in Ibadan, Nigeria (YRI); and Northern Europeans from Utah (CEU) in the 1,000 Genomes Project (Fig. 4a) (Auton *et al.* 2015; Jouganous *et al.* 2017). Briefly, this model consists of an ancestral human population under equilibrium demography with 24,000 individuals approximately 4,100 generations ago. At this time, the out-of-Africa migration occurs, and the population of the out-of-Africa branch shrinks by almost 90%, reflecting a severe bottleneck. After an additional 2,500 generations, the out-of-Africa branch splits into the CEU and CHB populations, both of which undergo a modest bottleneck and then experience exponential growth until the present day.

We first sought to understand the impact of selection on frequency spectra in ancestral populations conditioned on the present-day frequency in CEU. Surprisingly, we found that the expected frequency in the ancestor of CEU and CHB, $\mathbb{E}[X_{\text{CEU–CHB ancestor}} | X_{\text{CEU}}; s, \eta]$, is similar for alleles evolving neutrally and those under strong negative selection ($s = 10^{-3}$) (Supplementary Fig. S4a).

One might suppose that selection resembles neutrality because the 1,600 generations separating the CEU–CHB ancestor from CEU does not provide enough time for selection to act. However, this explanation can be ruled out by considering allele frequency dynamics forwards in time along this branch: it is apparent that selection substantially impacts allele frequencies over the 1,600 generations between the CEU–CHB ancestor and CEU (Supplementary Fig. S4b). Instead, our observations can be explained by non-equilibrium demography, which results in offsetting effects between the likelihood of CEU frequencies and the prior distribution in the CEU–CHB ancestor, similar to what we saw previously (Fig. 3a).

In contrast, we found that the expected frequency in the ancestor of CEU and YRI, $\mathbb{E}[X_{\text{CEU–YRI ancestor}} | X_{\text{CEU}}; s, \eta]$, is similar for alleles under strong stabilizing selection and alleles under strong negative selection (Supplementary Fig. S4c). Regardless of the frequency in CEU, alleles under strong stabilizing or negative selection are likely to be maintained at low frequencies in the ancestor of CEU and YRI. This is concordant with the results we saw for a pure bottleneck scenario (Fig. 3c), suggesting that the strength of the out-of-Africa bottleneck outweighs the subsequent exponential growth.

The stark difference in the backward transition probabilities for the CEU–YRI ancestor and the CEU–CHB ancestor is particularly notable given the qualitative similarity of the forward transition probabilities (Supplementary Fig. S4b, S4d). Ultimately, this underscores our finding that backward transitions are more sensitive to demography than forward transitions.

We next turned to understanding the allele frequency spectrum in one modern human population conditional on the frequency in another modern population. We focused on the distribution of frequencies in YRI and CHB conditional on the frequency in CEU, given that the majority of GWAS participants are highly genetically similar to CEU (Mills and Rahal 2019).

We observed that conditional on the frequency in CEU, stabilizing and negative selection both result in lower frequencies in CHB and YRI relative to neutrality, regardless of the CEU frequency (Fig. 4c, 4d). At low frequencies in CEU, positive selection and neutrality result in similar expected frequencies in CHB and YRI. However, as the CEU frequency increases, positively selected alleles are expected to be at much higher frequencies in CHB and YRI relative to neutral alleles. To understand this result, note that low frequency alleles tend to be younger on average and less likely to be segregating in the ancestral populations, especially when subject to selection (Supplementary Fig. S5a, S5d) (Maruyama 1974). As the CEU frequency increases, the likelihood of a positively selected allele segregating in the ancestral population increases, such that the allele can experience positive selection along the other branches, resulting in higher frequencies in CHB and YRI relative to neutrality (Supplementary Fig. S5c, S5f).

Finally, we considered which conditional frequency spectra would be most informative for identifying the mode of selection acting on complex traits. We found that comparing YRI and CEU is more informative than comparing CHB and CEU, especially when selection is weak (Supplementary Fig. S6). We also found that it is generally much easier to distinguish neutrality from

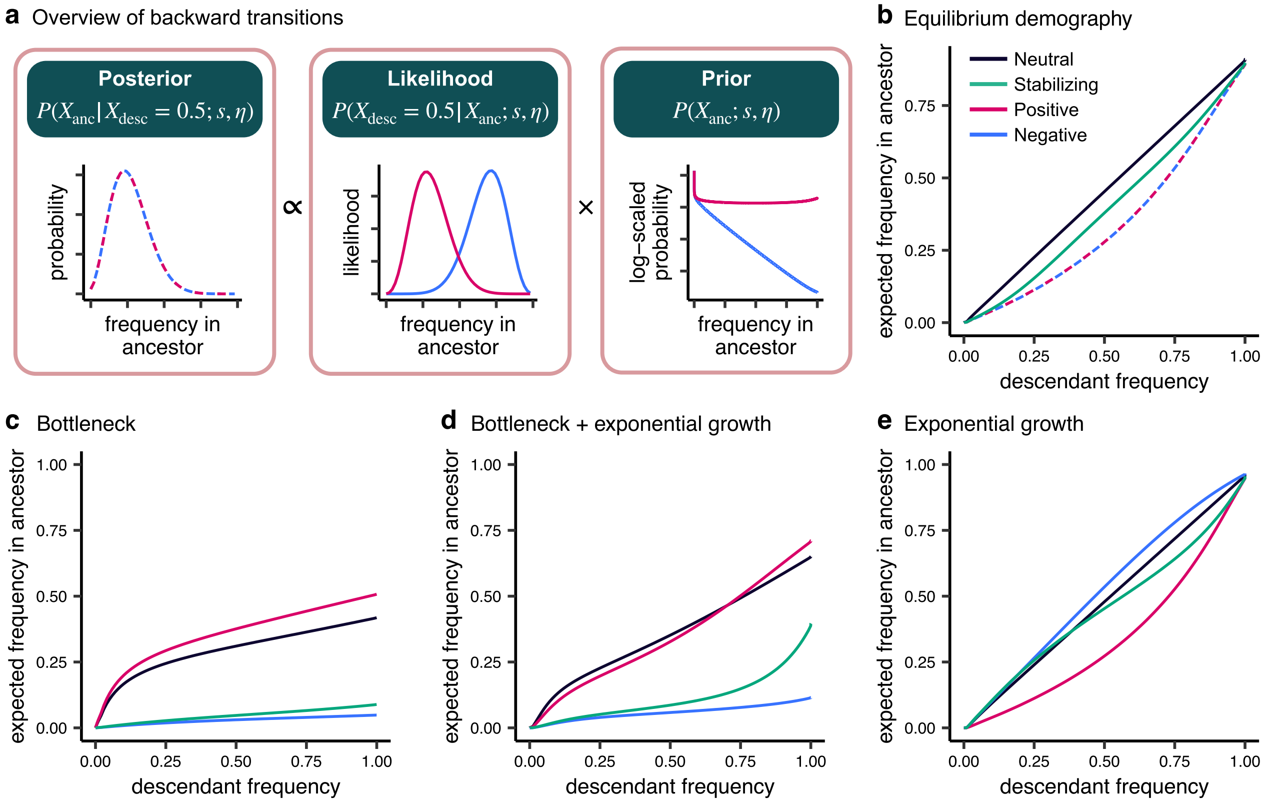

Fig. 3. Backward transitions. a) Backward transition probabilities can be interpreted as a posterior consisting of a likelihood (i.e. forward transition probabilities), multiplied by a prior (i.e. the marginal distribution in the ancestor). b–e) Expected frequency in an ancestral population conditional on the frequency in the descendant population, computed with fastDTWF. Selection coefficients correspond to $|hs| = 5.0 \times 10^{-4}$. Demographic models consist of an ancestral population with $N_e = 10{,}000$ and b) constant population size for 2,000 generations; c) $0.1N_e$ bottleneck; d) $0.1N_e$ bottleneck and exponential growth at a rate of 0.1%; e) exponential growth at a rate of 0.1% each generation. To enhance visibility, overlapping distributions are represented with dashed lines.

modes of selection than it is to distinguish *among* the different modes of selection. Specifically, frequency spectra in YRI conditional on CEU are similar under strong negative selection and strong stabilizing selection. At high (derived) allele frequencies in CEU, stabilizing selection and negative selection become easier to distinguish (Supplementary Fig. S5), but alleles at such high frequencies are discovered infrequently, particularly under these modes of selection (Fig. 1a, 4b). Interestingly, stabilizing and negative selection are easier to distinguish when selection is weaker – in this regime, derived alleles are more likely to reach frequencies greater than 0.5, where the two modes of selection exhibit different allelic dynamics (Supplementary Fig. S7a–c, Fig. 2c).

One consideration is that our approach inherently assumes a tree-like demographic model: when two populations split, allele frequencies in both populations must evolve independently. However, it is well-known that human demographic history featured ongoing migration between populations (Gutenkunst *et al.* 2009; Ragsdale and Gravel 2019). Such migration would generate a correlation between population allele frequencies, potentially impacting the conditional frequency spectrum. To understand the impact of migration on conditional frequency spectra, we used the software package *moments*, which is capable of modeling the joint frequency spectrum under more complex and non-tree-like demographies (Jouganous *et al.* 2017). Under weak selection, we found that conditional frequency spectra are qualitatively consistent in the presence of migration. Conditional on the frequency in CEU, alleles subject to negative or stabilizing selection are expected to be at lower frequencies in CHB and YRI relative to neutral alleles, while alleles subject to positive selection are expected to be at higher frequencies (Supplementary Fig. S8). Interestingly, we also observed some notable differences in the presence of migration: at frequencies upwards of 90% in CEU, alleles subject to stabilizing selection are likely to be at higher frequencies in CHB and YRI relative to neutral alleles.

Empirical analysis of trait-associated variants

To understand the mode of selection acting on human complex traits, we investigated empirical conditional frequency spectra. There are two primary challenges in using our theoretical work on conditional frequency spectra to interpret empirical data. First, we expect trait-associated variants ascertained with GWAS to represent a mixture of different selection coefficients. Second, as we observed in our theoretical work, conditional frequency spectra are sensitive to demography, meaning that our quantitative theoretical results are sensitive to the exact demographic parameters in the Jouganous *et al.* (2017) out-of-Africa model. However, our qualitative theoretical results are robust to small changes in demographic parameters along the branches of the tree (Supplementary Fig. S9), as well as in the ancestral population (Supplementary Fig. S10). This is especially true when conditioning on a modern population that has experienced a bottleneck: a recent bottleneck generates a strong qualitative signal that is recapitulated across many demographic scenarios.

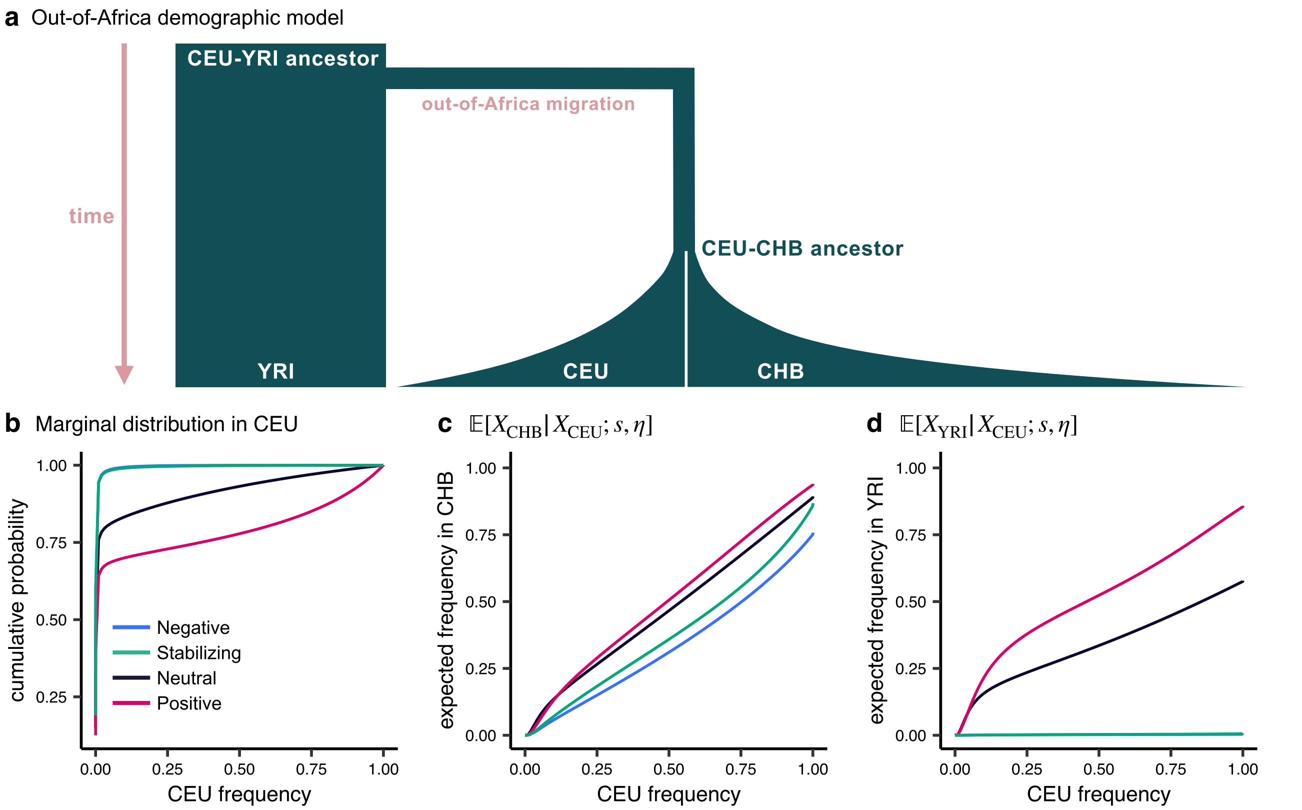

Fig. 4. Out-of-Africa demography. a) Overview of out-of-Africa demographic model inferred from YRI, CEU, and CHB by Jouganous et al. (2017). Widths and lengths of branches are approximately proportional to population sizes and split times. b) Marginal distribution in CEU predicted by our theoretical results. c, d) Expected frequency in CHB and YRI conditional on the frequency in CEU. For all panels, selection coefficients correspond to $|hs| = 5.0 \times 10^{-4}$, computed with fastDTWF.

Thus, given the challenges associated with drawing quantitative conclusions from conditional frequency spectra, we instead qualitatively investigate the mode of selection acting on complex traits. To do so, we compare conditional frequency spectra for trait-associated variants to conditional frequency spectra for "matched variants," a set of putatively neutral, non-associated variants with similar genomic properties. We generated a set of matched variants by matching trait-associated variants to non-associated variants on two metrics: derived allele frequency in the UK Biobank White British cohort, and B-value, a measure of background selection (Murphy et al. 2022). Matching on B-value accounts for differences in allele frequency spectra that are due to selection at linked sites, as opposed to selection at the focal variant. Though some fraction of these matched variants may be functional and under selection, this only makes our comparisons more conservative.

We compared conditional frequency spectra in YRI and CHB between trait-associated variants and matched variants. We conditioned on frequencies in the UK Biobank White British cohort because it is the GWAS cohort for 94 of the 106 complex traits we analyzed, and because the individuals in this cohort are generally genetically similar to the individuals in the remaining GWAS cohorts. To compare conditional frequency spectra between trait-associated variants and matched variants, we first grouped variants into 10 deciles based on their frequency in the UK Biobank White British. Within each decile, we computed the mean frequency in YRI and CHB for trait-associated variants and matched variants.

Across all deciles, we found that variants associated with quantitative traits have a systematically lower mean frequency in YRI relative to matched variants (Fig. 5a; $p = 2.4 \times 10^{-93}$; unpaired two-sided t-test). The same pattern was broadly observed for CHB, albeit to a lesser extent (Fig. 5b; $p = 9.8 \times 10^{-22}$; unpaired two-sided t-test). Much of this signal is driven by the fact that trait-associated variants are less likely to be polymorphic in YRI and CHB (Supplementary Fig. S11, S12). (Notably, given the empirical sample sizes of YRI and CHB, it is plausible that trait-associated variants are still segregating in the larger populations.)

Having established that trait-associated variants appear non-neutral, we aimed to understand which mode of selection explains the signal. We considered three models of selection on complex traits: stabilizing selection, directional selection, and purifying selection (i.e. negative selection against new mutations). Based on our theoretical results, stabilizing and negative selection on alleles produce similar conditional frequency spectra (Fig. 4d), making it difficult to distinguish between stabilizing and purifying selection at the trait level. However, these two models are easy to distinguish from directional selection: if a trait is under directional selection, then alleles with positive and negative effects on the trait should experience directional selection in opposite directions, and hence have different conditional frequency spectra. This means that under directional selection, *either* trait-associated variants with a positive effect *or* trait-associated variants with a negative effect should be at higher frequencies in YRI and CHB relative to matched variants. Under stabilizing or purifying selection, however, we always expect to see trait-associated variants at lower frequencies in YRI and CHB relative to matched variants, regardless of the sign of their effect.

For each trait, we divided trait-associated variants into two groups based on the sign of their effect, and compared their

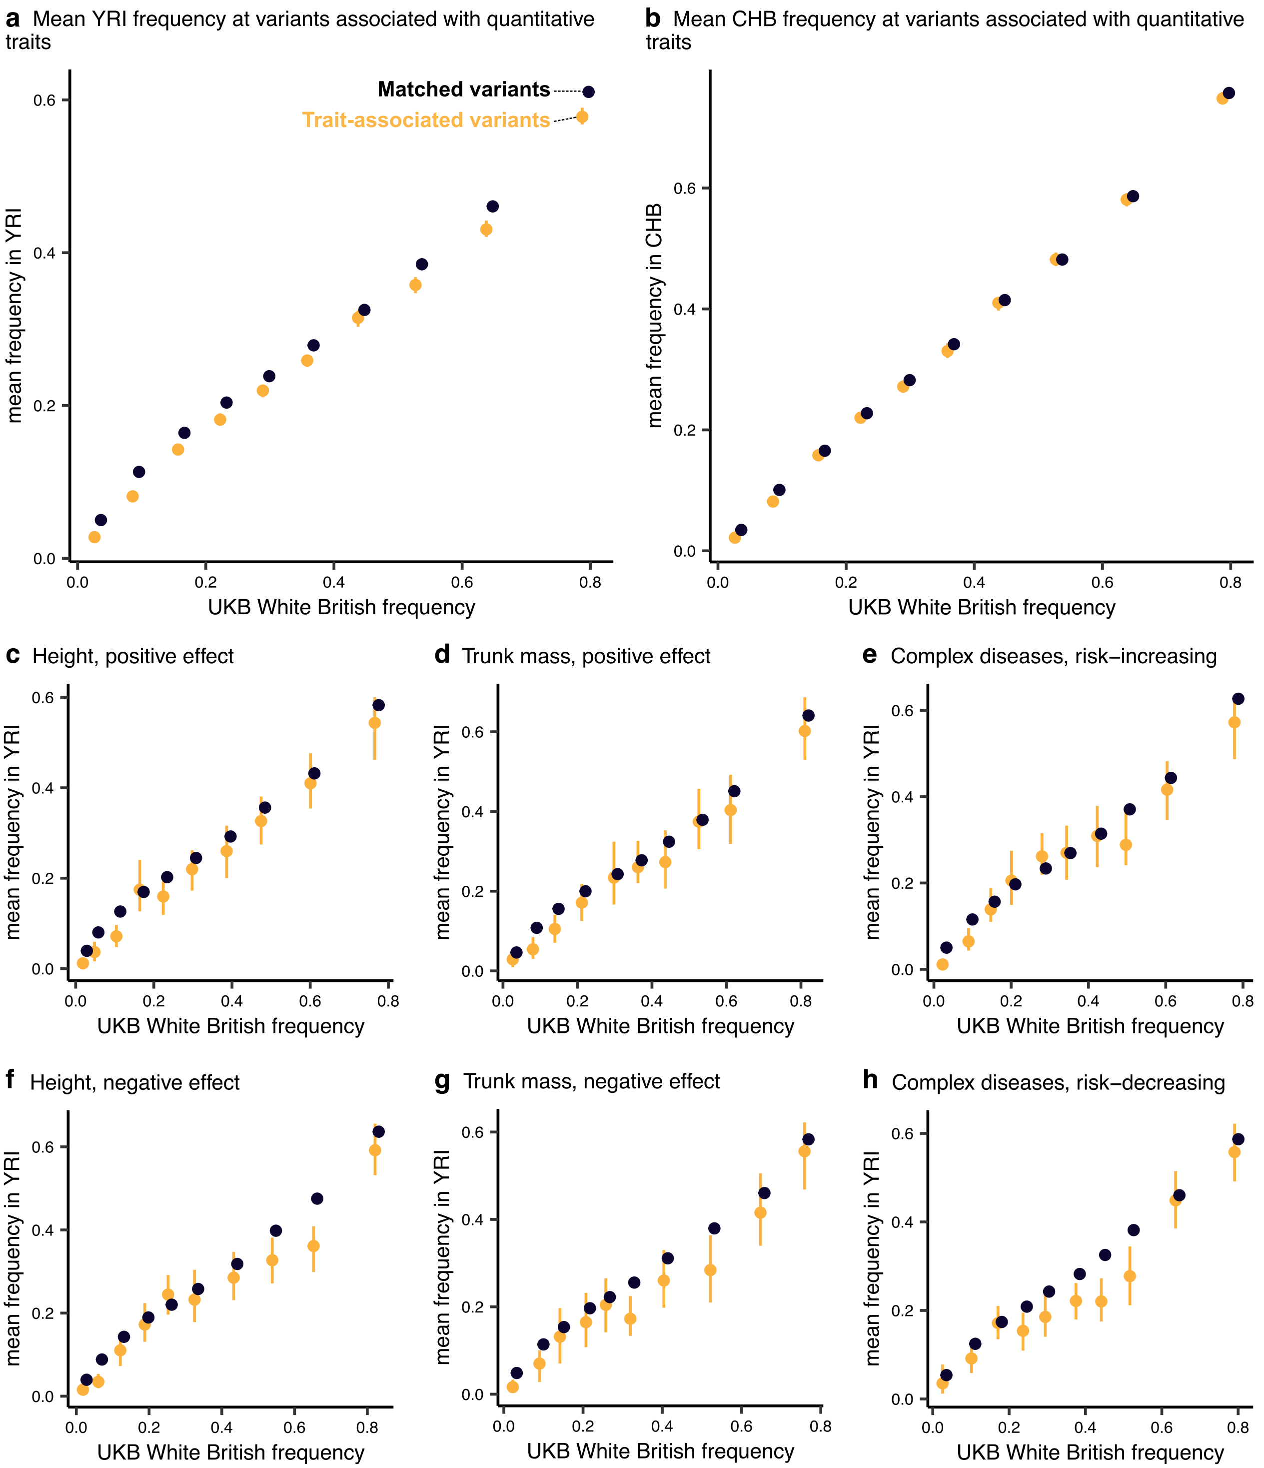

Fig. 5. Conditional frequency spectra for trait-associated variants. a, b) Mean frequency in YRI and CHB conditional on UK Biobank White British frequency decile for quantitative trait-associated variants and matched variants. c–h) Mean frequency in YRI conditional on UK Biobank White British frequency decile for variants associated with height, trunk mass, and complex diseases. For all panels, error bars depict the 95% confidence interval for the mean, calculated from 100 bootstrap samples. Points are jittered along the x-axis (UK Biobank White British frequency) for better visibility.

conditional frequency spectra to matched variants. Specifically, we performed one-sided tests of the alternate hypothesis that trait-associated variants have lower frequencies in YRI relative to matched variants. Across quantitative traits, we found that trait-associated variants have significantly lower frequencies in YRI relative to matched variants, regardless of the sign of their effect (Fig. 5c, 5d, 5f, 5g; $p = 3.2 \times 10^{-10}$ positive effect on height, $p = 2.3 \times 10^{-8}$ negative effect on height, $p = 4.6 \times 10^{-4}$ positive effect on trunk mass, $p = 7.6 \times 10^{-3}$ negative effect on trunk mass; unpaired one-sided t-test; see also Supplementary Fig. S13, S14). For complex diseases, we also found that trait-associated variants have lower frequencies in YRI relative to matched variants, regardless of whether they are risk-increasing or risk-decreasing (Fig. 5e, 5h; $p = 1.4 \times 10^{-6}$ risk-increasing, $p = 1.5 \times 10^{-11}$ risk-decreasing; unpaired one-sided t-test). Thus, we do not find evidence for directional selection, and instead find compelling evidence that stabilizing or purifying selection is the predominant mode of selection acting on complex traits and diseases.

Given that our results depend on knowing which allele is ancestral or derived, one consideration is the potential mispolarization of variants. In particular, variants at CpG sites are more difficult to polarize because of the higher mutation rate at these sites (Jónsson et al. 2017; Keightley and Jackson 2018). To account for this, we repeated our analyses excluding all transition mutations, and found that our results were qualitatively unchanged (Supplementary Fig. S15).

Another potential consideration is that our set of trait-associated variants likely contains some non-functional "tag" variants in linkage disequilibrium (LD) with the true causal variants: if two variants are in LD, their association signals are also highly correlated, making it difficult to determine which variant is causal. Ascertainment of non-functional tag variants can affect our results in two ways. First, tag variants and causal variants could be at different frequencies, distorting our empirical conditional frequency spectra. Second, conditional frequency spectra encode information about *derived* allele frequencies. If the derived allele at the tag variant is in negative LD with the derived allele at the causal variant, the trajectory of the derived allele at the tag variant will be opposite of that at the causal variant. Then, in conditional frequency spectra, the true direction of selection would appear to be reversed. While this would make our qualitative test of neutrality against selection more conservative, it could bias our results when trying to determine the mode of selection.

We used coalescent simulations to better understand these potential sources of bias. First, we investigated whether variants in LD could be at dramatically different frequencies. We found that for variants with $r^2 \geq 0.5$, the corresponding minor allele frequencies are essentially identical with high probability (Supplementary Fig. S16a). Next, we estimated the probability that the derived allele at a tag variant is in negative LD with the derived allele at a causal variant. We found that negative LD is uncommon when the derived allele at the tag variant is rare, but negative LD between derived alleles happens with ~50% probability when the derived allele at the tag variant reaches higher frequencies (Supplementary Fig. S16b).

Given that rare derived alleles more reliably tag derived causal variants, we repeated our tests for directional selection using trait-associated variants in the lowest quintile of UK Biobank White British derived allele frequency. For quantitative traits, we found that our results were qualitatively unchanged; we find that both trait-increasing and trait-decreasing variants have significantly lower frequencies in YRI relative to matched variants, ruling out directional selection ($p = 2.4 \times 10^{-9}$ positive effect on height, $p = 1.2 \times 10^{-8}$ negative effect on height, $p = 5.7 \times 10^{-4}$ positive effect on trunk mass, $p = 7.4 \times 10^{-4}$ negative effect on trunk mass; unpaired one-sided t-test). For complex diseases, we found that risk-increasing variants had significantly lower frequencies than matched variants but did not find significant evidence for risk-decreasing variants ($p = 1.7 \times 10^{-15}$ risk-increasing, $p = 0.09$ risk-decreasing; unpaired one-sided t-test). This could be due to the reduction in power when excluding 80% of trait-associated variants, or could be compatible with directional selection against variants increasing disease risk.

Impact of selection and demography on polygenic score portability

Our results have broader implications for applications of GWAS data, particularly in the context of polygenic scores. Polygenic scores estimate the genome-wide genetic contribution to a trait or disease using trait-associated variants ascertained in a GWAS. The phenotypic variance explained by polygenic scores is known to decrease in populations with less genetic similarity to the GWAS cohort – commonly referred to as a lack of "portability" (Martin AR et al. 2019; Ding et al. 2023). The portability of a polygenic score is influenced by many factors including differing patterns of linkage between ascertained variants and causal variants across populations; bias in effect sizes due to population structure; and environmental differences across populations (Vilhjálmsson et al. 2015; Wojcik et al. 2019; Mostafavi et al. 2020; Wang et al. 2020; Patel et al. 2022; Privé et al. 2022). Here, we focus on understanding the impact of allele frequency differences on portability, particularly in the context of different modes of selection.

For quantitative traits, the phenotypic variance explained by a polygenic score is proportional to $\sum_i 2x_i(1 - x_i)\beta_i^2$: the sum of squared effect sizes, β_i^2, at trait-associated variants, weighted by heterozygosity, $2x_i(1 - x_i)$. Thus, systematic differences in heterozygosity across populations will drive differences in the variance explained by the polygenic score, affecting portability.

Using theoretical conditional frequency spectra, we computed the expected heterozygosity of variants in CHB and YRI conditional on their frequency in CEU. We found that in both CHB and YRI, stabilizing and negative selection reduce heterozygosity relative to neutral alleles for low to moderate CEU frequencies (Fig. 6a, 6b). At CEU frequencies greater than 0.75, however, we found that negative and stabilizing selection can actually increase heterozygosity in CHB above neutral levels.

Relative to CHB, selection has a much more dramatic impact on the heterozygosity of alleles in YRI (Fig. 6a, 6b). Regardless of CEU frequency, negative and stabilizing selection result in extremely low heterozygosity in YRI. Even weakly selected alleles (corresponding to $Ns \approx 1$) have a substantially reduced heterozygosity in YRI, almost 10% lower than neutral alleles. In CHB, however, the heterozygosity under weak selection is nearly indistinguishable from neutrality.

This suggests that selection acting on complex traits (or correlated traits) will strongly impact polygenic score portability from CEU–like populations to YRI-like populations. To understand the role of demography, we considered a simpler demographic model consisting of two populations that split 2,000 generations ago and maintained a constant population size of $N_e = 10{,}000$. We computed the expected heterozygosity in one population conditional on the other and again found that selected alleles have reduced heterozygosity in comparison to neutral alleles (Supplementary Fig. S17). However, the magnitude of this difference is much smaller than what our theoretical results predict for CEU and YRI. It is therefore evident that while selection reduces polygenic score portability across all demographic scenarios, certain demographic scenarios can exacerbate the impacts of selection (Durvasula and Lohmueller 2021). Indeed, much of the reduction in polygenic score portability observed in empirical data can likely be attributed to the out-of-Africa bottleneck experienced by CEU-like populations (Martin AR et al. 2019; Ding et al. 2023).

To understand this phenomenon in empirical data, we compared the mean heterozygosity of trait-associated variants and matched variants in CHB and YRI across each decile of UK Biobank White British frequency. We found that trait-associated variants systematically have less heterozygosity in CHB and YRI, relative to matched variants (Fig. 6c, 6d; $p = 9.9 \times 10^{-39}$ CHB, $p = 5.0 \times 10^{-123}$ YRI; unpaired

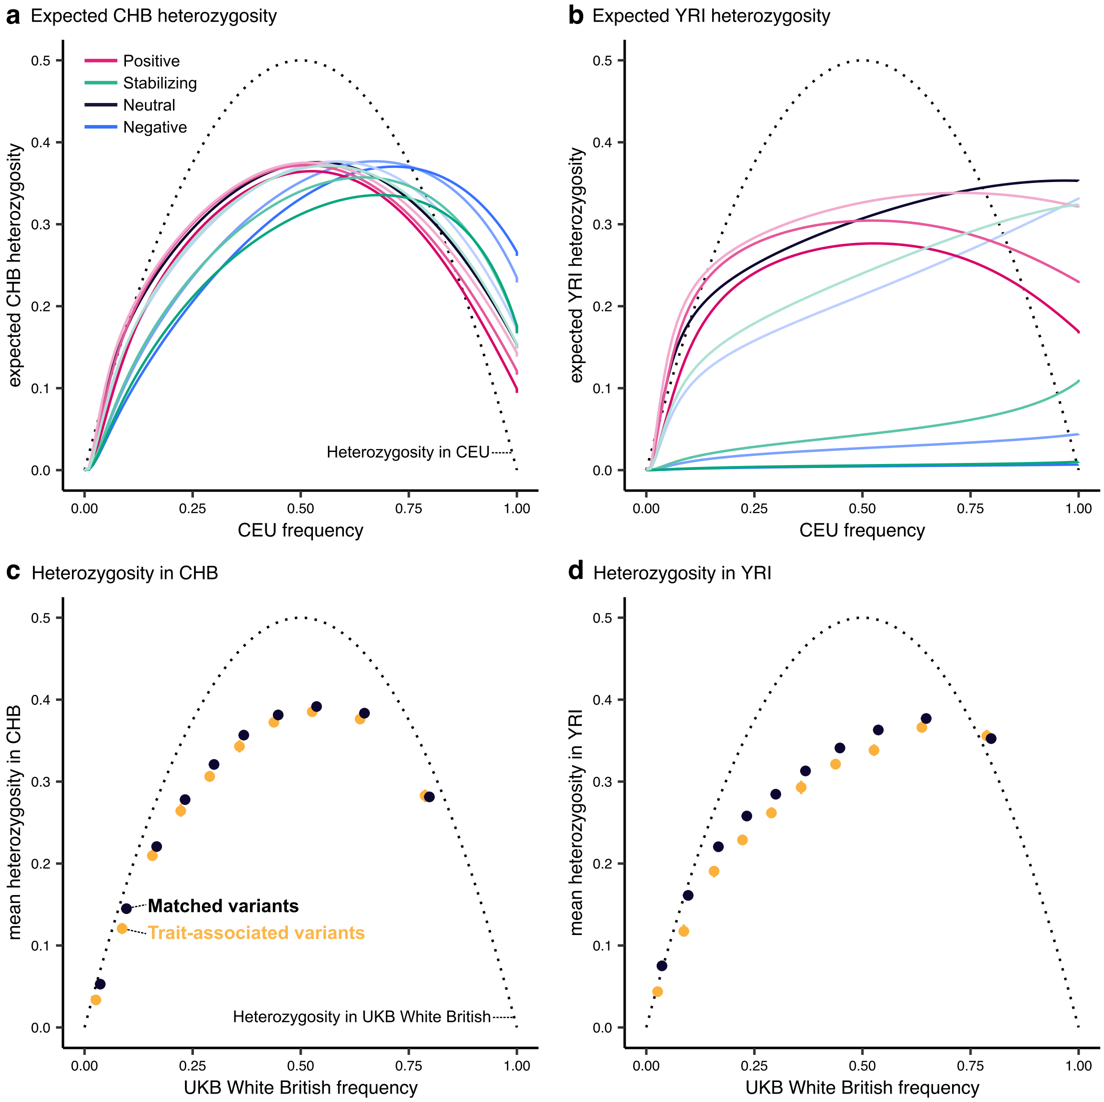

Fig. 6. Implications for polygenic score portability. a, b) Expected heterozygosity in CHB and YRI conditional on the frequency in CEU, computed with fastDTWF. Selection coefficients range from $|hs| = 5.0 \times 10^{-5}$, shown in the lightest shades, to $|hs| = 5.0 \times 10^{-4}$, shown in the darkest shades. c, d) Mean heterozygosity in CHB and YRI conditional on UK Biobank White British frequency decile for all trait-associated variants and matched variants. Error bars depict the 95% confidence interval for the mean, calculated from 100 bootstrap samples, and points are jittered along the x-axis (UK Biobank White British frequency) for better visibility. For all panels, dotted line corresponds to the heterozygosity in the conditional population.

two-sided t-test). As suggested by our theoretical results, the reduction in heterozygosity at trait-associated variants is stronger in YRI compared to CHB. These results illustrate how differences in allele frequencies driven by selection can contribute to the poor portability of polygenic scores.

Discussion

We presented conditional frequency spectra as a tool for studying selection on complex traits. The utility of conditional frequency spectra stems from the peculiarities of GWAS: GWAS ascertainment implicitly prioritizes variants with large minor allele frequencies and large effect sizes, both of which are related to the strength of selection on a variant. In doing so, GWAS ascertainment obscures much of the information one actually wants to learn from GWAS. Here, we used conditional frequency spectra to circumvent the issues caused by GWAS ascertainment to study selection on complex traits – but conditional frequency spectra should be broadly useful for other statistical genetics applications as well.

Our theoretical analyses of conditional frequency spectra identified robust qualitative signatures for each model of selection. Applying this intuition to empirical data, we found significant evidence for stabilizing or purifying selection acting on trait-associated variants, but no evidence for directional selection. We note that our approach only enables us to detect sustained directional selection shared by all branches in the tree. Others have previously proposed that selection acts in divergent directions across human populations, structuring many complex trait phenotypes (Guo *et al.* 2018); our analyses cannot rule out this possibility. However, our work does highlight that certain empirical observations can be compatible with multiple models of selection on complex traits. This emphasizes the benefit of invoking explicit population genetic models: models are crucial for interpreting what observations can or cannot tell us about selection, even if inference under such models is difficult.

Though empirical conditional frequency spectra are informative about the mode of selection acting on complex traits, we caution against using them to interpret the strength of selection.

First, trait-associated variants represent a mixture of selection coefficients – and moreover, trait-associated variants at intermediate frequencies in the GWAS cohort are likely to represent a different mixture of selection coefficients relative to trait-associated variants that are rare. Second, trait-associated variants in UK Biobank will not always tag the causal variant in other populations. Imperfect tagging will increase the similarity between conditional frequency spectra for trait-associated variants and matched variants, such that the strength of selection appears weaker.

Drawing quantitative conclusions about the strength of selection is also challenging because of the uncertainty in demographic models. Given how sensitive backward transitions are to demography, demographic misspecification can hinder the inference of selection coefficients from conditional frequency spectra. In our case, demographic misspecification arises from the fact that we use 1,000 Genomes CEU to construct our theoretical predictions, but we use variants ascertained in UK Biobank White British and other European cohorts for our empirical analyses. These subtle differences in allele frequency between European populations could bias quantitative inference of selection coefficients. To account for this, one could use putatively neutral variation to infer a demographic model relating the specific cohorts of individuals represented in empirical conditional frequency spectra. This demographic model could then be used to generate theoretical expectations of conditional frequency spectra and ultimately obtain more robust estimates of selection coefficients from data. In any case, care would need to be taken such that errors in the inferred demographic model do not affect estimates of selection.

Finally, we examined the consequences of selection on complex traits for polygenic score portability. Though it is well-known that selection on complex traits can affect the portability of polygenic scores (Wang et al. 2020; Durvasula and Lohmueller 2021; Yair and Coop 2022), our approach enables us to contrast different models of selection, as well as understand the respective contributions of demography and selection. We find that while selection reduces portability, the out-of-Africa bottleneck and subsequent drift likely plays an outsized role in the decreased portability from CEU-like populations to YRI-like populations. Here, we only modeled alleles, as opposed to traits, and so our results cannot speak to the effects of environmental heterogeneity, different contributions of the additive genetic component across populations, gene–gene interactions, or gene–environment interactions. Investigating these aspects further is likely to be informative for improving polygenic score portability across groups.

In conclusion, we characterize the conditional frequency spectrum under different models of selection, underscoring the value of conditioning on GWAS ascertainment. Our work parallels recent findings by Koch et al. (2024), who use the joint distribution of effect size and allele frequency to compare likelihoods for different models of selection. Together, these complementary approaches paint a compelling picture of stabilizing selection shaping the architecture of human complex traits.

Methods

Theoretical analysis of allele frequency spectra

We characterized the impact of demography and selection on the distribution of allele frequencies across populations by treating demography, η, and selection, s, as fixed quantities in a discrete probability distribution over allele frequencies X. We considered four distinct types of allele frequency distributions:

i) Forward transitions $P(X_{\text{desc}} = x_{\text{desc}} \mid X_{\text{anc}} = x_{\text{anc}}; s, \eta)$: the distribution of frequencies in a descendant population, conditional on the frequency in its ancestral population
ii) Marginals $P(X_k = x_k; s, \eta)$: the marginal distribution of frequencies in population k
iii) Backward transitions $P(X_{\text{anc}} = x_{\text{anc}} \mid X_{\text{desc}} = x_{\text{desc}}; s, \eta)$: the distribution of frequencies in an ancestral population, conditional on the frequency in its descendant population
iv) Conditional frequency spectra $P(X_j = x_j \mid X_k = x_k; s, \eta)$: the distribution of frequencies X in population j, conditional on the frequencies in population k, where j and k share a common ancestor.

For each branch in a particular demographic model, we obtained forward transitions $P(X_{\text{desc}} = x_{\text{desc}} \mid X_{\text{anc}} = x_{\text{anc}}; s, \eta)$ using two different implementations of the Wright–Fisher model. First, we used fastDTWF to numerically compute likelihoods under the discrete-time Wright–Fisher model (Spence et al. 2023). fastDTWF is restricted to modeling demographies that are piecewise constant. As such, we approximated exponential population growth by a piecewise constant model by updating the population size every 20 generations. Second, we used forward-time SLiM simulations, which are capable of modeling more flexible demographies, to obtain Monte Carlo approximations of the forward transitions (Haller and Messer 2019). We started at an initial ancestral frequency x_{anc} and simulated forward-in-time, recording the frequency in every downstream population in the demographic model. We then approximated the forward transition $P(X_{\text{desc}} = x_{\text{desc}} \mid X_{\text{anc}} = x_{\text{anc}}; s, \eta)$ with the empirical distribution across replicate simulations. We performed 1,000 replicates for each initial frequency for simple two population demographic models and 2,000 replicates for the out-of-Africa demographic model (see *Modeling demography* below for details). The forward transitions generated by fastDTWF and SLiM simulations differed in that while fastDTWF incorporates both recurrent mutations and new mutations, our particular implementation of SLiM simulations requires alleles to be segregating in the ancestral population of the demographic model.

To compute a marginal distribution for a descendant population, we first obtained the marginal distribution of the ancestral population in the demographic model, $P(X_{\text{anc}}; s, \eta)$, using fastDTWF. Unless explicitly stated otherwise, we assumed the ancestral population is at mutation–selection–drift balance conditioned on non-fixation of the derived allele (see Spence et al. 2023 for more details). We then computed the marginal distribution for the descendant population as follows:

$$P(X_{\text{desc}} = x_{\text{desc}}; s, \eta) = \sum_{x_{\text{anc}}} \Big(P(X_{\text{desc}} = x_{\text{desc}} \mid X_{\text{anc}} = x_{\text{anc}}; s, \eta) \times P(X_{\text{anc}} = x_{\text{anc}}; s, \eta) \Big) \tag{1}$$

Using the marginal probability $P(X_{\text{desc}} = x_{\text{desc}}; s, \eta)$, we computed the backward transition as follows:

$$P(X_{\text{anc}} = x_{\text{anc}} \mid X_{\text{desc}} = x_{\text{desc}}; s, \eta) = \frac{1}{P(X_{\text{desc}} = x_{\text{desc}}; s, \eta)} \Big(P(X_{\text{desc}} = x_{\text{desc}} \mid X_{\text{anc}} = x_{\text{anc}}; s, \eta) \times P(X_{\text{anc}} = x_{\text{anc}}; s, \eta) \Big) \tag{2}$$

Similarly, we computed a conditional frequency distribution $P(X_j = x_j \mid X_k = x_k; s, \eta)$ in which populations j and k share a common ancestor:

$$P(X_j = x_j \mid X_k = x_k; s, \eta) = \frac{1}{P(X_k = x_k; s, \eta)} \sum_{x_{anc}} \Big(P(X_j = x_j \mid X_{anc} = x_{anc}; s, \eta) \times P(X_k = x_k \mid X_{anc} = x_{anc}; s, \eta) \times P(X_{anc} = x_{anc}; s, \eta) \Big) \quad (3)$$

In practice, we performed these computations on frequency distributions by first aggregating allele frequencies into bins. A distribution over allele frequencies is a discrete probability distribution with $2N$ possible outcomes, where N is the diploid population size. Thus, without binning frequencies, the number of computations required for basic operations quickly becomes cumbersome. Moreover, adjacent conditional distributions will be trivially similar (e.g. the probability distribution $P(X_j \mid X_k = 10{,}000 \text{ counts}; s, \eta)$ is nearly identical to the probability distribution $P(X_j \mid X_k = 10{,}001 \text{ counts}; s, \eta)$) (Spence *et al.* 2023). To account for the fact that adjacent conditional distributions are more dissimilar when alleles are close to loss or fixation, we used denser binning for allele frequencies close to 0 and 1. Specifically, we created bins corresponding to 0 derived allele counts, 1 count, 2 counts, (2, 5] counts, (5, 10] counts, and (10, 20] counts to cover the space of allele frequencies close to 0; and bins corresponding to $2N$ counts, $2N-1$ counts, and $(2N-5, 2N-2]$ counts to cover the space of allele frequencies close to 1. To cover the remaining space of allele frequencies, we created frequency bins with a width of 0.01. For each population this procedure generated roughly 100 allele frequency bins, regardless of the population size N.

Modeling demography

We relied on two different types of demographic models. First, we considered a simple two population model consisting of one ancestral population and one descendant population separated by 2,000 generations. We used this model to understand how selection impacts allele frequency dynamics under different demographic conditions: equilibrium demography (i.e. constant population size), a bottleneck (i.e. a sudden decrease in population size), exponential population growth, and a bottleneck followed by exponential growth. We focused on parameter values relevant to human demographic history, modeling an ancestral population size with $N_e = 10{,}000$; a bottleneck of $0.1N_e$ and $0.3N_e$; and an exponential growth rate of 0.05% and 0.1% (Gutenkunst *et al.* 2009; Jouganous *et al.* 2017; Ragsdale and Gravel 2019).

The second demographic model we considered was the out-of-Africa model inferred by Jouganous *et al.* (2017). Jouganous *et al.* (2017) inferred demographic parameters relating three modern-day populations in the 1000 Genomes Project (Auton *et al.* 2015): Han Chinese in Beijing, China (CHB); Yoruba in Ibadan, Nigeria (YRI); and Northern Europeans from Utah (CEU). In brief, this model consists of a common ancestor for CHB, YRI, and CEU 4,100 generations ago, followed by a strong out-of-Africa bottleneck and subsequent split between CHB and CEU 1,600 generations ago. CHB experiences a modest bottleneck after splitting from CEU, and the CHB and CEU branches both experience exponential growth following their split (Fig. 3a). Our model only differs from that of Jouganous *et al.* (2017) in that we did not model migration between branches; our probability computations assume that populations are independent after branching (but see *Conditional frequency spectra with migration* below for a different approach that incorporates migration between branches).

For all theoretical results, we report the exact frequency spectrum of the entire population in the demographic model (as opposed to an approximate frequency spectrum estimated from a smaller sample).

Modeling selection

We considered three types of selection on alleles: negative genic selection, positive genic selection, and stabilizing selection. Using common notational convention, we can represent the fitness of the *AA* genotype as 1, the *Aa* genotype as $1 + hs$, and the *aa* genotype as $1 + s$, where *A* is the ancestral allele and *a* is the derived allele (Gillespie 2004). Under negative genic selection, $h = 0.5$ and $s < 0$. We considered three values of s: -1×10^{-4}, -5×10^{-4}, and -1×10^{-3}. These values range from extremely weak (corresponding to $Ns \approx 1$ in the ancestral population) to the strongest selection coefficients inferred for trait-associated variants ascertained in complex trait GWAS (Simons *et al.* 2022). Under positive genic selection, $h = 0.5$ and $s > 0$. We modeled values of s ranging from $+1 \times 10^{-4}$ to $+1 \times 10^{-3}$, analogous to negative selection. Under stabilizing selection, allele frequency dynamics resemble a scenario where $hs < 0$ but $s = 0$, such that only heterozygotes experience a fitness cost (Robertson 1956; Walsh and Lynch 2018). We modeled values of hs ranging from -5×10^{-4} to -5×10^{-5}, analogous to negative selection.

Conditional frequency spectra with migration

To generate conditional frequency spectra that incorporate migration, we used *moments* (Jouganous *et al.* 2017). *moments* is able to model the joint frequency spectrum between populations under complex demographic scenarios (albeit with weaker selection regimes). We obtained the joint frequency spectrum between CEU, CHB, and YRI under the demographic model inferred by Jouganous *et al.* (2017) and weak selection ($|hs| = 5.0 \times 10^{-5}$). As before, we used the joint frequency spectrum to compute conditional frequency spectra by marginalizing out one population and conditioning on another.

Empirical analysis of trait-associated variants

We obtained trait-associated variants from published GWAS for 94 quantitative complex traits and 12 (binary) complex diseases. For quantitative traits, we obtained a set of 18,229 variants ascertained in the UK Biobank "White British" cohort with a minor allele frequency of at least 0.01 (Simons *et al.* 2022) (http://www.nealelab.is/uk-biobank). The White British cohort consists of approximately 337,000 unrelated individuals in the UK Biobank. We restricted our analyses to this subset of the UK Biobank because our approach relies on performing GWAS in a cohort that is relatively homogeneous and has reasonably high genetic similarity to one of the populations in our out-of-Africa demographic model. Moreover, the size of the White British cohort enables less noisy estimation of allele frequencies.

For disease traits, we obtained a total of 1,040 variants ascertained in cohorts with European ancestries using individual case/control GWAS (Supplementary Table S1) (International IBD Genetics Consortium (IIBDGC) *et al.* 2013; Michailidou *et al.* 2015; De Lange *et al.* 2017; Scott *et al.* 2017; Pardiñas *et al.* 2018; Schumacher *et al.* 2018; Nalls *et al.* 2019; Mullins *et al.* 2021; Aragam *et al.* 2022; Bellenguez *et al.* 2022; Ishigaki *et al.* 2022; Demontis *et al.* 2023). Paralleling our quantitative trait analyses, we filtered out variants with a minor allele frequency less than 0.01 in the UK Biobank White British.

For each trait-associated variant, we generated a set of "matched variants" that share similar properties. We started with approximately 11 million (imputed) variants in the UK Biobank that are not found to be trait-associated in our GWAS datasets. For a given trait-associated variant, we identified matched variants by selecting for two different criteria. First, we matched on the derived allele frequency in the UK Biobank White British. To compute derived allele frequencies, we obtained the ancestral allele state inferred by Auton et al. (2015) and stored in the Ensembl variation database for all trait-associated variants and all other imputed variants in the UK Biobank (Martin FJ et al. 2023). Briefly, Auton et al. (2015) identified ancestral allele states using a multiple genome alignment between human, chimp, orangutan, and rhesus macaque to infer ancestral sequences and annotate ancestral allele states. We matched on UK Biobank White British derived allele frequency by binning continuous frequencies into 100 evenly spaced bins between 0 and 1; two derived allele frequencies were deemed to match if they belonged to the same bin.

Second, we matched variants on their B-value, a background selection statistic (Murphy et al. 2022). Background selection impacts allele frequency spectra by decreasing genetic diversity, and varies throughout the genome as a function of recombination rate, gene density, and other genomic features. Thus, matching on B-values accounts for differences in allele frequency spectra between trait-associated variants and matched variants that arise due to varying levels of background selection throughout the genome. We matched on B-value by binning the B-values for the 11 million variants imputed by UKB into 15 quantiles; two B-values were deemed to match if they belonged to the same bin. We discarded trait-associated variants with fewer than 500 matched variants, generating an average of 11,547 matched variants for each trait-associated variant.

We generated empirical conditional frequency spectra for both trait-associated variants and matched variants. We considered two types of conditional frequency spectra: the distribution of frequencies in CHB (Han Chinese from Beijing, China) conditional on the frequency in UK Biobank White British, and the distribution of frequencies in YRI (Yoruba in Ibadan, Nigeria) conditional on the frequency in UK Biobank White British. To compute allele frequencies in CHB and YRI, we used 103 CHB individuals and 108 YRI individuals from 1,000 Genomes Phase 3, respectively.

To pool information across the relatively small number of trait-associated variants, we split them into ten deciles based on their frequency in UK Biobank White British. Within each UK Biobank frequency decile, we computed the mean frequency in CHB and YRI across all trait-associated variants in the decile. This approximates the expectation over the conditional frequency spectrum, $\mathbb{E}[X_{\mathrm{CHB}} \mid X_{\mathrm{UKB}}]$ and $\mathbb{E}[X_{\mathrm{YRI}} \mid X_{\mathrm{UKB}}]$, for trait-associated variants. We next computed the analogous quantity for matched variants by aggregating across the matched variants for each trait-associated variant in a particular UK Biobank frequency decile. To compute the 95% confidence interval for the mean(s), we separately bootstrapped over trait-associated and matched variants, and reported the 0.025 and 0.975 quantiles across 100 replicates.

To test for selection, we compared trait-associated and matched variants by performing an unpaired two-sided t-test for unequal sample variances within each UK Biobank frequency decile. We combined p-values across deciles using Fisher's method.

To specifically test for directional selection, we split trait-associated variants into two groups based on the sign of their effect (i.e. trait-increasing and trait-decreasing). For each group, we tested the alternative hypothesis that trait-associated variants have a lower frequency in YRI and CHB relative to matched variants by performing an unpaired one-sided t-test for unequal sample variances within each UK Biobank frequency decile. We again combined p-values across deciles using Fisher's method. When the null hypothesis was rejected for both trait-increasing and trait-decreasing variants, we interpreted this as evidence against directional selection.

Analysis of imperfect tagging

We aimed to understand the probability that a trait-associated variant accurately tags the "correct" allele at the true causal variant – in other words, the probability that derived alleles at a pair of linked variants are positively correlated. Using msprime (Baumdicker et al. 2022), we performed coalescent simulations for 100 diploid individuals. Specifically, we simulated 100,000 coalescent trees with exactly two mutations. For each tree, we calculated the linkage disequilibrium (LD) between the pair of mutations as the squared correlation coefficient of their genotypes. We then computed the probability that derived alleles are linked at pairs of mutations meeting various LD thresholds.

Data availability

Conditional frequency spectra and other allele frequency distributions can be found at https://github.com/roshnipatel/conditional-frequency-spectra, along with code for performing computations on frequency distributions and analyzing empirical data. Supplemental material available at GENETICS online.

Acknowledgments

We thank Alyssa Lyn Fortier and Courtney Smith for valuable feedback on the manuscript, and members of the Pritchard lab for helpful conversations about this work.

Funding

Research reported in this publication was supported by NIH grants R01HG011432 and R01HG008140.

Conflicts of interest

The author(s) declare no conflicts of interest.

Literature cited

International IBD Genetics Consortium (IIBDGC), Agliardi C, Alfredsson L, Alizadeh M, Anderson C, Andrews R, Søndergaard HB, Baker A, Band G, Baranzini SE. 2013. Analysis of immune-related loci identifies 48 new susceptibility variants for multiple sclerosis. Nat Genet. 45:1353–1360. doi:10.1038/ng.2770.

Aragam KG, Jiang T, Goel A, Kanoni S, Wolford BN, Atri DS, Weeks EM, Wang M, Hindy G, Zhou W, et al. 2022. Discovery and systematic characterization of risk variants and genes for coronary artery disease in over a million participants. Nat Genet. 54: 1803–1815. doi:10.1038/s41588-022-01233-6.

Auton A, Abecasis GR, Altshuler DM, Durbin RM, Abecasis GR, Bentley DR, Chakravarti A, Clark AG, Donnelly P, Eichler EE, *et al.* 2015. A global reference for human genetic variation. Nature. 526:68–74. doi:10.1038/nature15393.

Baharian S, Gravel S. 2018. On the decidability of population size histories from finite allele frequency spectra. Theor Popul Biol. 120: 42–51. doi:10.1016/j.tpb.2017.12.008.

Baumdicker F, Bisschop G, Goldstein D, Gower G, Ragsdale AP, Tsambos G, Zhu S, Eldon B, Ellerman EC, Galloway JG, *et al.* 2022. Efficient ancestry and mutation simulation with msprime 1.0. Genetics. 220:iyab229. doi:10.1093/genetics/iyab229.

Bellenguez C, Küçükali F, Jansen IE, Kleineidam L, Moreno-Grau S, Amin N, Naj AC, Campos-Martin R, Grenier-Boley B, Andrade V, *et al.* 2022. New insights into the genetic etiology of Alzheimer's disease and related dementias. Nat Genet. 54:412–436. doi:10.1038/s41588-022-01024-z.

Bhaskar A, Song YS. 2014. Descartes' rule of signs and the identifiability of population demographic models from genomic variation data. Ann Stat. 42:2469–2493. doi:10.1214/14-AOS1264.

Bhaskar A, Wang YXR, Song YS. 2015. Efficient inference of population size histories and locus-specific mutation rates from large-sample genomic variation data. Genome Res. 25:268–279. doi:10.1101/gr.178756.114.

Caballero A, Tenesa A, Keightley PD. 2015. The nature of genetic variation for complex traits revealed by GWAS and regional heritability mapping analyses. Genetics. 201:1601–1613. doi:10.1534/genetics.115.177220.

Chen H, Green RE, Pääbo S, Slatkin M. 2007. The joint allele-frequency spectrum in closely related species. Genetics. 177: 387–398. doi:10.1534/genetics.107.070730.

Chen H, Slatkin M. 2013. Inferring selection intensity and allele age from multilocus haplotype structure. G3 (Bethesda). 3: 1429–1442. doi:10.1534/g3.113.006197.

Clark AG, Hubisz MJ, Bustamante CD, Williamson SH, Nielsen R. 2005. Ascertainment bias in studies of human genome-wide polymorphism. Genome Res. 15:1496–1502. doi:10.1101/gr.4107905.

De Lange KM, Moutsianas L, Lee JC, Lamb CA, Luo Y, Kennedy NA, Jostins L, Rice DL, Gutierrez-Achury J, Ji S-G, *et al.* 2017. Genome-wide association study implicates immune activation of multiple integrin genes in inflammatory bowel disease. Nat Genet. 49:256–261. doi:10.1038/ng.3760.

Demontis D, Walters GB, Athanasiadis G, Walters R, Therrien K, Nielsen TT, Farajzadeh L, Voloudakis G, Bendl J, Zeng B, *et al.* 2023. Genome-wide analyses of ADHD identify 27 risk loci, refine the genetic architecture and implicate several cognitive domains. Nat Genet. 55:198–208. doi:10.1038/s41588-022-01285-8.

Dilber E, Terhorst J. 2024. Faster inference of complex demographic models from large allele frequency spectra.

Ding Y, Hou K, Xu Z, Pimplaskar A, Petter E, Boulier K, Privé F, Vilhjálmsson BJ, Olde Loohuis LM, Pasaniuc B. 2023. Polygenic scoring accuracy varies across the genetic ancestry continuum. Nature. 618:774–781. doi:10.1038/s41586-023-06079-4.

Durvasula A, Lohmueller KE. 2021. Negative selection on complex traits limits phenotype prediction accuracy between populations. Am J Hum Genet. 108:620–631. doi:10.1016/j.ajhg.2021.02.013.

Durvasula A, Sankararaman S. 2020. Recovering signals of ghost archaic introgression in African populations. Sci Adv. 6:eaax5097. doi:10.1126/sciadv.aax5097.

Evans SN, Shvets Y, Slatkin M. 2007. Non-equilibrium theory of the allele frequency spectrum. Theor Popul Biol. 71:109–119. doi:10.1016/j.tpb.2006.06.005.

Eyre-Walker A. 2010. Genetic architecture of a complex trait and its implications for fitness and genome-wide association studies. Proc Natl Acad Sci U S A. 107:1752–1756. doi:10.1073/pnas.0906182107.

Gazal S, Finucane HK, Furlotte NA, Loh P-R, Palamara PF, Liu X, Schoech A, Bulik-Sullivan B, Neale BM, Gusev A, *et al.* 2017. Linkage disequilibrium-dependent architecture of human complex traits shows action of negative selection. Nat Genet. 49: 1421–1427. doi:10.1038/ng.3954.

Gillespie JH. 2004. Population genetics. Johns Hopkins University Press.

Guo J, Wu Y, Zhu Z, Zheng Z, Trzaskowski M, Zeng J, Robinson MR, Visscher PM, Yang J. 2018. Global genetic differentiation of complex traits shaped by natural selection in humans. Nat Commun. 9:1865. doi:10.1038/s41467-018-04191-y.

Gutenkunst RN, Hernandez RD, Williamson SH, Bustamante CD. 2009. Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. PLoS Genet. 5: e1000695. doi:10.1371/journal.pgen.1000695.

Haller BC, Messer PW. 2019. SLiM 3: Forward genetic simulations beyond the Wright–Fisher model. Mol Biol Evol. 36:632–637. doi:10.1093/molbev/msy228.

Harpak A, Bhaskar A, Pritchard JK. 2016. Mutation rate variation is a primary determinant of the distribution of allele frequencies in humans. PLoS Genet. 12:e1006489. doi:10.1371/journal.pgen.1006489.

Ishigaki K, Sakaue S, Terao C, Luo Y, Sonehara K, Yamaguchi K, Amariuta T, Too CL, Laufer VA, Scott IC, *et al.* 2022. Multi-ancestry genome-wide association analyses identify novel genetic mechanisms in rheumatoid arthritis. Nat Genet. 54: 1640–1651. doi:10.1038/s41588-022-01213-w.

Jónsson H, Sulem P, Kehr B, Kristmundsdottir S, Zink F, Hjartarson E, Hardarson MT, Hjorleifsson KE, Eggertsson HP, Gudjonsson SA, *et al.* 2017. Parental influence on human germline de novo mutations in 1,548 trios from Iceland. Nature. 549:519–522. doi:10.1038/nature24018.

Jouganous J, Long W, Ragsdale AP, Gravel S. 2017. Inferring the joint demographic history of multiple populations: Beyond the diffusion approximation. Genetics. 206:1549–1567. doi:10.1534/genetics.117.200493.

Kamm J, Terhorst J, Durbin R, Song YS. 2020. Efficiently inferring the demographic history of many populations with allele count data. J Am Stat Assoc. 115:1472–1487. doi:10.1080/01621459.2019.1635482.

Kamm JA, Terhorst J, Song YS. 2017. Efficient computation of the joint sample frequency spectra for multiple populations. J Comput Graph Stat. 26:182–194. doi:10.1080/10618600.2016.1159212.

Keightley PD, Hill WG. 1990. Variation maintained in quantitative traits with mutation-selection balance: pleiotropic side-effects on fitness traits. Proc Biol Sci. 242:95–100. doi:10.1098/rspb.1990.0110.

Keightley PD, Jackson BC. 2018. Inferring the probability of the derived vs the ancestral allelic state at a polymorphic site. Genetics. 209:897–906. doi:10.1534/genetics.118.301120.

Kern AD, Hey J. 2017. Exact calculation of the joint allele frequency spectrum for isolation with migration models. Genetics. 207: 241–253. doi:10.1534/genetics.116.194019.

Kim MS, Patel KP, Teng AK, Berens AJ, Lachance J. 2018. Genetic disease risks can be misestimated across global populations. Genome Biol. 19:179. doi:10.1186/s13059-018-1561-7.

Koch EM, Connally NJ, Baya NA, Reeve MP, Daly MJ, Neale BM, Lander ES, Bloemendhal A, Sunyaev SR. 2024. Genetic association data are broadly consistent with stabilizing selection shaping human common diseases and traits. Pages: 2024.06.19.599789 Section: New Results.

Lachance J, Tishkoff SA. 2013. SNP ascertainment bias in population genetic analyses: Why it is important, and how to correct it. Bioessays. 35:780–786. doi:10.1002/bies.v35.9.

Lande R. 1976. Natural selection and random genetic drift in phenotypic evolution. Evolution. 30:314–334. doi:10.2307/2407703.

Lukić S, Hey J. 2012. Demographic inference using spectral methods on SNP data, with an analysis of the human out-of-Africa expansion. Genetics. 192:619–639. doi:10.1534/genetics.112.141846.

Lukić S, Hey J, Chen K. 2011. Non-equilibrium allele frequency spectra via spectral methods. Theor Popul Biol. 79:203–219. doi:10.1016/j.tpb.2011.02.003.

Manolio TA, Collins FS, Cox NJ, Goldstein DB, Hindorff LA, Hunter DJ, McCarthy MI, Ramos EM, Cardon LR, Chakravarti A, et al. 2009. Finding the missing heritability of complex diseases. Nature. 461:747–753. doi:10.1038/nature08494.

Martin FJ, Amode MR, Aneja A, Austine-Orimoloye O, Azov A, Barnes I, Becker A, Bennett R, Berry A, Bhai J, et al. 2023. Ensembl 2023. Nucleic Acids Res. 51:D933–D941. doi:10.1093/nar/gkac958.

Martin AR, Kanai M, Kamatani Y, Okada Y, Neale BM, Daly MJ. 2019. Clinical use of current polygenic risk scores may exacerbate health disparities. Nat Genet. 51:584–591. doi:10.1038/s41588-019-0379-x.

Maruyama T. 1974. The age of an allele in a finite population. Genet Res. 23:137–143. doi:10.1017/S0016672300014750.

Michailidou K, Beesley J, Lindstrom S, Canisius S, Dennis J, Lush MJ, Maranian MJ, Bolla MK, Wang Q, Shah M, et al. 2015. Genome-wide association analysis of more than 120,000 individuals identifies 15 new susceptibility loci for breast cancer. Nat Genet. 47:373–380. doi:10.1038/ng.3242.

Mills MC, Rahal C. 2019. A scientometric review of genome-wide association studies. Commun Biol. 2:1–11. doi:10.1038/s42003-018-0261-x.

Mostafavi H, Harpak A, Agarwal I, Conley D, Pritchard JK, Przeworski M. 2020. Variable prediction accuracy of polygenic scores within an ancestry group. Elife. 9:e48376. doi:10.7554/eLife.48376.

Mostafavi H, Spence JP, Naqvi S, Pritchard JK. 2023. Systematic differences in discovery of genetic effects on gene expression and complex traits. Nat Genet. 55:1866–1875. doi:10.1038/s41588-023-01529-1.

Mullins N, Forstner AJ, O'Connell KS, Coombes B, Coleman JR, Qiao Z, Als TD, Bigdeli TB, Børte S, Bryois J, et al. 2021. Genome-wide association study of more than 40,000 bipolar disorder cases provides new insights into the underlying biology. Nat Genet. 53: 817–829. doi:10.1038/s41588-021-00857-4.

Murphy DA, Elyashiv E, Amster G, Sella G. 2022. Broad-scale variation in human genetic diversity levels is predicted by purifying selection on coding and non-coding elements. Elife. 12:e76065. doi:10.7554/eLife.76065.

Myers S, Fefferman C, Patterson N. 2008. Can one learn history from the allelic spectrum? Theor Popul Biol. 73:342–348. doi:10.1016/j.tpb.2008.01.001.

Nalls MA, Blauwendraat C, Vallerga CL, Heilbron K, Bandres-Ciga S, Chang D, Tan M, Kia DA, Noyce AJ, Xue A, et al. 2019. Identification of novel risk loci, causal insights, and heritable risk for Parkinson's disease: A meta-analysis of genome-wide association studies. Lancet Neurol. 18:1091–1102. doi:10.1016/S1474-4422(19)30320-5.

O'Connor LJ, Schoech AP, Hormozdiari F, Gazal S, Patterson N, Price AL. 2019. Extreme polygenicity of complex traits is explained by negative selection. Am J Hum Genet. 105:456–476. doi:10.1016/j.ajhg.2019.07.003.

Ortega-Del Vecchyo D, Lohmueller KE, Novembre J. 2022. Haplotype-based inference of the distribution of fitness effects. Genetics. 220:iyac002. doi:10.1093/genetics/iyac002.

Pardiñas AF, Holmans P, Pocklington AJ, Escott-Price V, Ripke S, Carrera N, Legge SE, Bishop S, Cameron D, Hamshere ML, et al. 2018. Common schizophrenia alleles are enriched in mutation-intolerant genes and in regions under strong background selection. Nat Genet. 50:381–389. doi:10.1038/s41588-018-0059-2.

Patel RA, Musharoff SA, Spence JP, Pimentel H, Tcheandjieu C, Mostafavi H, Sinnott-Armstrong N, Clarke SL, Smith CJ, V.A. Million Veteran Program, et al. 2022. Genetic interactions drive heterogeneity in causal variant effect sizes for gene expression and complex traits. Am J Hum Genet. 109:1286–1297. doi:10.1016/j.ajhg.2022.05.014.

Pritchard JK. 2001. Are rare variants responsible for susceptibility to complex diseases? Am J Hum Genet. 69:124–137. doi:10.1086/321272.

Privé F, Aschard H, Carmi S, Folkersen L, Hoggart C, O'Reilly PF, Vilhjálmsson BJ. 2022. Portability of 245 polygenic scores when derived from the UK Biobank and applied to 9 ancestry groups from the same cohort. Am J Hum Genet. 109:12–23. doi:10.1016/j.ajhg.2021.11.008.

Ragsdale AP, Gravel S. 2019. Models of archaic admixture and recent history from two-locus statistics. PLoS Genet. 15:e1008204. doi:10.1371/journal.pgen.1008204.

Ragsdale AP, Moreau C, Gravel S. 2018. Genomic inference using diffusion models and the allele frequency spectrum. Curr Opin Genet Dev. 53:140–147. doi:10.1016/j.gde.2018.10.001.

Robertson A. 1956. The effect of selection against extreme deviants based on deviation or on homozygosis. J Genet. 54:236–248. doi:10.1007/BF02982779.

Rosen Z, Bhaskar A, Roch S, Song YS. 2018. Geometry of the sample frequency spectrum and the perils of demographic inference. Genetics. 210:665–682. doi:10.1534/genetics.118.300733.

Schoech AP, Jordan DM, Loh P-R, Gazal S, O'Connor LJ, Balick DJ, Palamara PF, Finucane HK, Sunyaev SR, Price AL. 2019. Quantification of frequency-dependent genetic architectures in 25 UK Biobank traits reveals action of negative selection. Nat Commun. 10:790. doi:10.1038/s41467-019-08424-6.

Schraiber JG. 2018. Assessing the relationship of ancient and modern populations. Genetics. 208:383–398. doi:10.1534/genetics.117.300448.

Schraiber JG, Evans SN, Slatkin M. 2016. Bayesian inference of natural selection from allele frequency time series. Genetics. 203: 493–511. doi:10.1534/genetics.116.187278.

Schumacher FR, Olama AAAl, Berndt SI, Benlloch S, Ahmed M, Saunders EJ, Dadaev T, Leongamornlert D, Anokian E, Cieza-Borrella C, et al. 2018. Association analyses of more than 140,000 men identify 63 new prostate cancer susceptibility loci. Nat Genet. 50:928–936. doi:10.1038/s41588-018-0142-8.

Scott RA, Scott LJ, Mägi R, Marullo L, Gaulton KJ, Kaakinen M, Pervjakova N, Pers TH, Johnson AD, Eicher JD, et al. 2017. An expanded genome-wide association study of type 2 diabetes in Europeans. Diabetes. 66:2888–2902. doi:10.2337/db16-1253.

Simons YB, Bullaughey K, Hudson RR, Sella G. 2018. A population genetic interpretation of GWAS findings for human quantitative traits. PLoS Biol. 16:e2002985. doi:10.1371/journal.pbio.2002985.

Simons YB, Mostafavi H, Smith CJ, Pritchard JK, Sella G. 2022. Simple scaling laws control the genetic architectures of human complex traits.

Slatkin M. 2001. Simulating genealogies of selected alleles in a population of variable size. Genet Res. 78:49–57. doi:10.1017/S0016672301005183.

Song YS, Steinrücken M. 2012. A simple method for finding explicit analytic transition densities of diffusion processes with general diploid selection. Genetics. 190:1117–1129. doi:10.1534/genetics.111.136929.

Speed D, Hemani G, Johnson M, Balding D. 2012. Improved heritability estimation from genome-wide SNPs. Am J Hum Genet. 91:1011–1021. doi:10.1016/j.ajhg.2012.10.010.

Speed D, Holmes J, Balding DJ. 2020. Evaluating and improving heritability models using summary statistics. Nat Genet. 52:458–462. doi:10.1038/s41588-020-0600-y.

Spence JP, Kamm JA, Song YS. 2016. The site frequency spectrum for general coalescents. Genetics. 202:1549–1561. doi:10.1534/genetics.115.184101.

Spence JP, Zeng T, Mostafavi H, Pritchard JK. 2023. Scaling the discrete-time Wright–Fisher model to Biobank-scale datasets. Genetics. 225:iyad168. doi:10.1093/genetics/iyad168.

Terhorst J, Song YS. 2015. Fundamental limits on the accuracy of demographic inference based on the sample frequency spectrum. Proc Natl Acad Sci U S A. 112:7677–7682. doi:10.1073/pnas.1503717112.

Turelli M. 1984. Heritable genetic variation via mutation-selection balance: Lerch's zeta meets the abdominal bristle. Theor Popul Biol. 25:138–193. doi:10.1016/0040-5809(84)90017-0.

Vilhjálmsson BJ, Yang J, Finucane HK, Gusev A, Lindström S, Ripke S, Genovese G, Loh P-R, Bhatia G, Do R, *et al.* 2015. Modeling linkage disequilibrium increases accuracy of polygenic risk scores. Am J Hum Genet. 97:576–592. doi:10.1016/j.ajhg.2015.09.001.

Walsh B, Lynch M. 2018. Evolution and selection of quantitative traits. Oxford University Press.

Wang Y, Guo J, Ni G, Yang J, Visscher PM, Yengo L. 2020. Theoretical and empirical quantification of the accuracy of polygenic scores in ancestry divergent populations. Nat Commun. 11:3865. doi:10.1038/s41467-020-17719-y.

Weiner DJ, Nadig A, Jagadeesh KA, Dey KK, Neale BM, Robinson EB, Karczewski KJ, O'Connor LJ. 2023. Polygenic architecture of rare coding variation across 394,783 exomes. Nature. 614:492–499. doi:10.1038/s41586-022-05684-z.

Wojcik GL, Graff M, Nishimura KK, Tao R, Haessler J, Gignoux CR, Highland HM, Patel YM, Sorokin EP, Avery CL, *et al.* 2019. Genetic analyses of diverse populations improves discovery for complex traits. Nature. 570:514–518. doi:10.1038/s41586-019-1310-4.

Yair S, Coop G. 2022. Population differentiation of polygenic score predictions under stabilizing selection. Philos Trans R Soc B: Biol Sci. 377:20200416. doi:10.1098/rstb.2020.0416.

Yang J, Bakshi A, Zhu Z, Hemani G, Vinkhuyzen AAE, Lee SH, Robinson MR, Perry JRB, Nolte IM, van Vliet-Ostaptchouk JV, *et al.* 2015. Genetic variance estimation with imputed variants finds negligible missing heritability for human height and body mass index. Nat Genet. 47:1114–1120. doi:10.1038/ng.3390.

Yang J, Benyamin B, McEvoy BP, Gordon S, Henders AK, Nyholt DR, Madden PA, Heath AC, Martin NG, Montgomery GW, *et al.* 2010. Common SNPs explain a large proportion of the heritability for human height. Nat Genet. 42:565–569. doi:10.1038/ng.608.

Yang MA, Harris K, Slatkin M. 2014. The projection of a test genome onto a reference population and applications to humans and archaic hominins. Genetics. 198:1655–1670. doi:10.1534/genetics.112.145359.

Zeng J, de Vlaming R, Wu Y, Robinson MR, Lloyd-Jones LR, Yengo L, Yap CX, Xue A, Sidorenko J, McRae AF, *et al.* 2018. Signatures of negative selection in the genetic architecture of human complex traits. Nat Genet. 50:746–753. doi:10.1038/s41588-018-0101-4.

Zeng J, Xue A, Jiang L, Lloyd-Jones LR, Wu Y, Wang H, Zheng Z, Yengo L, Kemper KE, Goddard ME, *et al.* 2021. Widespread signatures of natural selection across human complex traits and functional genomic categories. Nat Commun. 12:1–12. doi:10.1038/s41467-021-21446-3.

Zhang MJ, Durvasula A, Chiang C, Koch EM, Strober BJ, Shi H, Barton AR, Kim SS, Weissbrod O, Loh P-R, *et al.* 2023. Pervasive correlations between causal disease effects of proximal SNPs vary with functional annotations and implicate stabilizing selection. Pages: 2023.12.04.23299391.

Živković D, Stephan W. 2011. Analytical results on the neutral non-equilibrium allele frequency spectrum based on diffusion theory. Theor Popul Biol. 79:184–191. doi:10.1016/j.tpb.2011.03.003.

Editor: S. Gravel