PNAS RESEARCH ARTICLE | GENETICS OPEN ACCESS

Observational epidemiological studies can mitigate genetic confounding with a genetic relatedness matrix

Roshni A. Patel[a,b,c,1], Joshua G. Schraiber[c], Matt Pennell[d,1], and Michael D. Edge[c,1]

Affiliations are included on p. 10.

Edited by Marcus W. Feldman, Stanford University, Stanford, CA; received November 24, 2025; accepted April 6, 2026

Observational studies are commonly used in psychology and epidemiology to identify risk factors correlated with health outcomes. However, these studies are vulnerable to confounding when shared genetic variation influences both the putative risk factor and outcome. Researchers have often controlled for this type of genetic confounding using polygenic scores, but these scores are noisy and biased estimators of a trait's genetic component. While some newer methods offer significant improvements over polygenic scores, they still rely on genome-wide association studies (GWAS) summary statistics, which may be untenable for certain datasets. Here, we develop an analogous method that leverages a genetic relatedness matrix to control genetic confounding when testing for nongenetic risk factors. In simulations, we find that our method outperforms existing approaches, particularly at sample sizes that are large by the standards of much human research but smaller than datasets often used in human genetics. We also demonstrate that existing methods are susceptible to poor GWAS portability, whereas our method is inherently robust to such concerns, conditional on the availability of individual genotype data. Finally, we apply our method to the UK Biobank to reanalyze social risk factors for health outcomes in previously understudied cohorts.

genetic confounding | epidemiological association study | genetic relatedness matrix | PGS portability | genetic correlation

Significance

Large observational datasets are often used in epidemiology or social sciences to identify risk factors associated with health outcomes. However, these studies can be misleading when the putative risk factor and outcome are both complex genetic traits: a correlated genetic basis can create spurious associations between a putative risk factor and outcome. Existing approaches to address this problem typically rely on summary statistics from genome-wide association studies, which require biobank-scale datasets. In this work, we introduce an alternative method that uses a genetic relatedness matrix to control for genetic confounding. We show that our approach is closely related to existing methods but outperforms them in smaller datasets, making it a particularly valuable tool for research on understudied traits and populations.

Empirical research across the social, behavioral, and health sciences often relies on observational studies to test for relationships between putative exposures and outcomes. However, these observational studies can be subject to confounding when the exposure and outcome are both complex traits with a genetic component. Complex traits are influenced by many genetic variants, with estimates for some traits reaching into the tens or even hundreds of thousands (1–3). These variants also tend to be pleiotropic, influencing multiple traits simultaneously, which raises the possibility that both exposure and outcome are influenced by shared genetic variation (4).

Shared genetic variation poses a problem for causal inference in observational epidemiological studies. In many instances, researchers are interested in understanding whether intervening on a particular exposure (e.g. BMI) could modify someone's risk for a particular outcome (e.g. high cholesterol). In other words, to understand the effect of environmental interventions, researchers seek to understand whether an observed correlation between exposure and outcome manifests independently of genetics: Correlations that arise solely due to shared genetic variation may not readily translate to effective interventions. For example, previous work has found that associations between developmental risk factors and later-in-life psychopathology are attenuated after accounting for shared genetics and other factors (5). Consequently, there is a growing awareness of the importance of controlling for shared genetic variation that affects the putative exposure and outcome of interest, both in studies of population samples, as we focus on here, as well as in twin studies (6–9). Within the social sciences, recognition of this problem has motivated a surge of interest in behavioral genetics and the use of genetic information to better understand the effects of behavioral interventions (10).

Historically, efforts to control genetic confounding—specifically, the genetic confounding that results from genetic correlation between exposure and outcome—have relied on including covariates that capture the genetic variation underlying the exposure. One approach conditions on a small number of genetic variants that have a large effect on the exposure trait (11–14). However, this approach relies on the genetic architecture of the exposure trait being dominated by common, large-effect variants, which is not always true for complex traits (2). Another line of work has sought to control for genetic confounding by including the polygenic score for the exposure as a covariate

Author contributions: R.A.P., J.G.S., M.P., and M.D.E. designed research; R.A.P. performed research; R.A.P. analyzed data; and R.A.P., J.G.S., M.P., and M.D.E. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

[1]To whom correspondence may be addressed. Email: roshnip@uoregon.edu, mpennell@cornell.edu, or edgem@usc.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.2533909123/-/DCSupplemental.

Published May 6, 2026.

when regressing outcome on exposure (6, 15–20). Polygenic scores (PGS) estimate the genetic component of a trait by combining information across hundreds or thousands of genetic variants ascertained from genome-wide association studies (GWAS). In theory, one could control genetic confounding by perfectly estimating the true value of the genetic component of the exposure, and including it as a covariate when regressing outcome on exposure. But in practice, PGS estimated from data do not capture all causal variants, and the effects of these causal variants are estimated with bias and noise. In response, a considerable body of work has sought to improve control of genetic confounding by improving PGS estimation and carefully modeling measurement error (21–24). However, less progress has been made in improving the "portability" of PGS. Due to population structure confounding and challenges with fine-mapping causal variants, PGS are known to be less accurate when estimated from and applied to groups with different genetic ancestries or environments (25–27). As a result, PGS-corrected estimates of the exposure effect are likely to be unreliable when applied to populations that are underrepresented in the cohorts from which that PGS was estimated.

Recently, ref. 28 introduced an alternative approach to control genetic confounding named PENGUIN. PENGUIN-corrected estimates of the exposure effect are a tremendous improvement over previous methods that rely on polygenic scores. Whereas earlier PGS-based methods sought to control genetic confounding by noisily estimating individual-level genetic values, PENGUIN uses LD score regression (LDSC) (1) to directly estimate the genetic correlation between exposure and outcome from GWAS summary statistics. This enables PENGUIN to estimate the exposure effect by subtracting off the correlation between exposure and outcome that is mediated through genetics.

However, at least two problems remain. First, both PENGUIN and PGS-based methods presuppose that data on exposure and outcome have been collected in samples large enough to allow very precise GWAS estimates, whether in the study at hand or in an external dataset. In practice, this may require sample sizes of hundreds of thousands of individuals. This requirement limits the utility of existing datasets where valuable phenotypic information may have been collected in smaller sample sizes. Moreover, biobank-scale measurements may be fundamentally untenable for many variables analyzed in observational epidemiological studies: Psychologists commonly study traits that are time-consuming and labor-intensive to measure, such as longitudinal phenotypes or scales computed from multiquestion surveys. Second, even when traits have been measured in large cohorts, methods that rely on GWAS summary statistics are susceptible to poor portability across cohorts with differences in genetic ancestries or environmental backgrounds (26, 27, 29). Despite growing recognition of this problem, progress in expanding the diversity of biobank participants remains slow (30). Thus, there remains a need to develop a method that controls genetic confounding when studying traits that are hard to measure, or populations underrepresented in existing GWAS.

To develop a flexible method to control genetic confounding in observational epidemiological studies, we propose using a genetic relatedness matrix (GRM) to fit a bivariate linear mixed model. GRMs are constructed from individual-level genotype data, often as the variance–covariance matrices of standardized, additive genotypes, and they encode information on genetic relationships among individuals (31). Thus, the GRM can be used to partition the naive correlation between exposure and outcome into a component mediated through genetics and a component independent of genetics (32, 33). Here, we demonstrate the conceptual similarity between GRM-corrected estimates and estimates derived from earlier GWAS-based methods. We show that GRM-corrected estimates largely resolve the two problems associated with GWAS-corrected estimates: They can be applied to small datasets and computed within-sample, mitigating portability concerns. We conclude by using our method to reanalyze social risk factors for health outcomes in 6,104 Indian and 8,483 Black British individuals in the UK Biobank.

Results

Statistical Model. We begin by introducing a generative statistical model of two phenotypes that we refer to as the exposure, X, and the outcome, Y (Fig. 1*A*):

$$\begin{aligned} X &= G\beta_X + e_X \\ Y &= G\beta_Y + e_Y + be_X, \end{aligned} \qquad [1]$$

where, for n individuals and m genetic variants, G is the centered $n \times m$ genotype matrix; β_X and β_Y are $m \times 1$ vectors of additive SNP effect sizes; and e_X and e_Y are $n \times 1$ vectors of nongenetic effects on the traits. We assume the entries of e_X and e_Y are independently and identically distributed across individuals with expectations of 0 and that e_X and e_Y are uncorrelated. We further assume that nongenetic effects e_X and e_Y are independent of G. We assume β_X and β_Y have expected values of 0, and we allow that β_X and β_Y could be correlated.

Under this generative model, the outcome is correlated with the exposure through both genetic and nongenetic mechanisms, and the coefficient b measures the effect of the nongenetic component of the exposure. In other words, b is the quantity we wish to estimate in observational epidemiological studies: It quantifies the extent to which intervening on the exposure will affect the outcome variable. In contrast, the naive estimate of the exposure effect—obtained by simply regressing the outcome on the exposure—is $\frac{\mathrm{Cov}(G\beta_X, G\beta_Y)+b\mathrm{Var}(e_X)}{\mathrm{Var}(G\beta_X)+\mathrm{Var}(e_X)}$ in expectation. In other words, the naive estimate of the exposure effect is subject to both regression attenuation and bias from genetic correlation between traits (see *SI Appendix, Supplementary Note* 1 for details).

Although this model has not been explicitly formalized in prior literature (to our knowledge), we show that it can be understood as the basis for including the exposure PGS as a covariate and other existing approaches to control genetic confounding (*SI Appendix, Supplementary Note* 1). Given that e_X is unobserved and unmeasurable, estimating b requires an indirect approach. PENGUIN approaches this problem by modeling $\mathrm{Cov}(G\beta_X, G\beta_Y)$, whereas PGS-based methods seek to estimate $G\beta_X$. Here, we propose estimating the interventional effect of the exposure by using a GRM to partition the covariance between exposure and outcome into a genetic component and a nongenetic component.

If we assume that both the effects of genetic variants and the nongenetic effects on traits are normally distributed, the exposure and outcome can be modeled with a bivariate normal distribution:

$$\begin{pmatrix} X \\ Y \end{pmatrix} = N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \Sigma_G \otimes R + \Sigma_E \otimes I \right), \qquad [2]$$

where R is an $n \times n$ GRM, I is the n-dimensional identity matrix, and Σ_G and Σ_E represent the 2×2 genetic and nongenetic variance–covariance matrices between exposure and outcome.

Given a set of genotypes from which to construct the GRM R, and phenotypes for the exposure X and outcome Y, we can estimate the genetic and nongenetic variance components Σ_G and Σ_E by using variance component analysis to model the covariance of X and Y across individuals. In fact, the model in Eq. **2** may be familiar as bivariate GREML, which is commonly used to estimate heritability and genetic correlation using the entries of Σ_G (32, 33). In contrast, we propose using the entries of Σ_E to estimate the interventional effect of an exposure in the context of observational epidemiological studies. Specifically, let the entries of Σ_E be denoted as follows:

$$\Sigma_E = \begin{pmatrix} \sigma^2_{e_X} & \sigma_{e_{XY}} \\ \sigma_{e_{XY}} & \sigma^2_{e_Y} \end{pmatrix},$$

where $\sigma_{e_{XY}}$ is the covariance between the nongenetic components of the exposure X and outcome Y; $\sigma^2_{e_X}$ is the variance of the nongenetic component of the exposure X; and $\sigma^2_{e_Y}$ is the variance of the nongenetic component of the exposure Y. As implied by Eq. **1**, the covariance between the nongenetic component of the exposure X (e_X) and the nongenetic component of the outcome Y ($e_Y + be_X$) is $\sigma_{e_{XY}} = b\sigma^2_{e_X}$. Thus, an estimator for the interventional effect of the exposure X is given by

$$\hat{b} = \frac{\hat{\sigma}_{e_{XY}}}{\hat{\sigma}^2_{e_X}},$$

where $\hat{\sigma}_{e_{XY}}$ and $\hat{\sigma}^2_{e_X}$ are estimators of the corresponding quantities—in this paper, the ones provided by GREML.

Assessing the Utility of GRM-Corrected Estimates. Though PGS-, PENGUIN-, and GRM-corrected estimates share the same underlying model of trait variation, these approaches impose different data requirements (Fig. 1*B*). PGS-corrected estimates use individual-level genotype and phenotype data to regress outcome on exposure, and the PGS itself is typically constructed using either in-sample or out-of-sample GWAS summary statistics. In contrast, PENGUIN-corrected estimates do not require individual-level genotypes. Instead, PENGUIN uses GWAS summary statistics to estimate the genetic correlation between traits, and individual-level phenotypes to estimate the overall phenotypic correlation between exposure and outcome. Finally, GRM-corrected estimates require no GWAS summary statistics at all and instead partition the covariance between exposure and outcome strictly using individual-level genotypes and phenotypes.

Thus, we sought to understand the utility of different methods for controlling genetic confounding. We simulated genotypes for individuals under a realistic model of human demography (34, 35). Conditional on these genotypes, we simulated phenotypes for exposure and outcome traits under the generative model described in Eq. **1** and an additive genetic architecture with frequency-dependent effect sizes (*Materials and Methods*).

In simulations, we find that the perfect PGS—that is, the true value of the genetic component of the exposure—would control genetic confounding so successfully that the resulting estimates lie on the diagonal $y = x$ line (Fig. 2*A*). However, as the correlation between the estimated PGS and the true genetic value decreases, PGS-corrected estimates of the exposure effect become increasingly susceptible to genetic confounding. Any genetic effects not captured by the estimated PGS are absorbed into the coefficient of the exposure, which is why the resulting estimate of exposure effect is biased toward the genetic correlation between exposure and outcome. Strikingly, even when a PGS is constructed with 100% of the true causal variants, noisy estimation of variant effect sizes will result in biased estimates of exposure effect. Given that current GWAS are not expected to identify 100% of the causal variants for a trait, this represents a relatively optimistic scenario. When only 20% of causal variants are captured, the PGS-corrected estimate of exposure effect barely performs better than the estimate from the naive regression of outcome on exposure. For the remainder of this manuscript, we therefore prioritize comparing GRM-corrected estimates with PENGUIN-corrected estimates.

We next compared the performance of PENGUIN- and GRM-corrected estimates in realistic scenarios (Fig. 2*B*). Some observational epidemiological studies are conducted in biobanks with information on hundreds of thousands of individuals, whereas others rely on datasets with a few thousand individuals (6, 16). We find that GRM-corrected estimates derived from sample sizes of just 4,000 individuals perform as well as PENGUIN-corrected estimates derived from biobank-scale data. In smaller sample sizes, we find that GRM-corrected estimates substantially outperform PENGUIN-corrected estimates, reflecting the loss of information when using GWAS summary statistics rather than individual genotypes in modest sample sizes. At a sample size of 4,000 individuals, GWAS summary statistics are too noisy for accurate estimation of heritability and genetic correlation via LD score regression. This highlights the particular utility of GRM-corrected estimates in studies with sample sizes that are too small for reliable GWAS summary statistics. On the other hand, the utility of PENGUIN-corrected estimates becomes apparent as sample sizes approach hundreds of thousands of individuals. In this setting, constructing a GRM becomes computationally infeasible, but GWAS summary statistics become less noisy and more accurate.

Interrogating Assumptions About Trait Heritability. To identify the potential limitations of GRM- and PENGUIN-corrected estimates, we next examined their robustness to assumptions about trait heritability. Both approaches estimate the interventional effect of the exposure by partitioning the genetic correlation between traits from the nongenetic correlation between them, which requires making assumptions about how trait heritability is distributed across individuals and variants. Thus, we sought to understand how violations of these assumptions affect our ability to recover the true interventional value of the exposure effect.

We first evaluated the effect of violating assumptions about random mating. Population structure and cross-trait assortative mating can both induce spurious correlations between individuals' genotypes and are known to bias traditional heritability and genetic correlation estimators (36, 37). To characterize the effect of population structure, we simulated two subpopulations with an F_{ST} of 0.005, a value comparable to those observed between populations within continental Europe. We find that when the mean exposure phenotype is shared between subpopulations, the CIs for GRM- and PENGUIN-corrected estimates generally include the true exposure effect, though PENGUIN-corrected estimates have much wider CIs (*SI Appendix*, Fig. S1*A*). However, if the mean environmental contributions to the exposure and outcome (that is, the subpopulation-mean values of e_X and e_Y) differ between subpopulations (e.g., due to environmental differences), we find that PENGUIN-corrected estimates exhibit greater bias than GRM-corrected estimates (*SI Appendix*, Fig. S1*B*). In simulations of cross-trait assortative mating, we find that CIs for GRM-corrected estimates generally recover the true exposure effect, but estimates can be prone to bias as the strength of cross-trait assortative mating increases

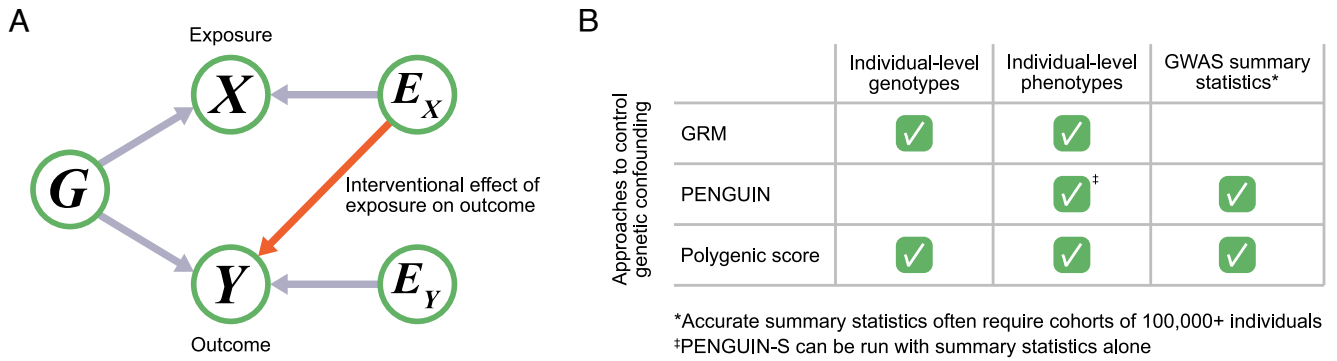

Fig. 1. Overview. (*A*) Causal graph for observational epidemiological studies where both exposure and outcome are genetic traits. The orange arrow illustrates the quantity of interest: the interventional effect of the exposure on the outcome. (*B*) Data requirements for GRM-, PENGUIN-, and PGS-corrected estimates of exposure effect.

(*SI Appendix*, Fig. S1*C*). We were unable to evaluate PENGUIN-corrected estimates under assortative mating given the large sample sizes required for PENGUIN and the computational complexity of assortative mating simulations (37).

Last, we considered assumptions about how heritability is distributed across variants. We find that if the exposure and outcome are simulated with a heritability model that violates method assumptions, both GRM- and LDSC-based estimates of heritability will be biased, consistent with previous findings (*SI Appendix*, Fig. S2*A*) (38). Nevertheless, we find that GRM-corrected estimates of exposure effect are typically more successful at recovering the true effect (*SI Appendix*, Fig. S2*B*).

Portability and Methods to Control Genetic Confounding. Finally, we sought to characterize the portability of methods to control genetic confounding. Whereas GRM-corrected estimates can only be obtained from a single dataset with individual-level genotype and phenotype data, both PGS- and PENGUIN-corrected estimates can be obtained by combining information across multiple cohorts. In the case of PENGUIN-corrected estimates, this feature presents a considerable methodological advantage: By leveraging external GWAS summary statistics, researchers can use PENGUIN to estimate exposure effect even in the absence of individual-level genotype data (Fig. 1*B*). At the same time, combining information across cohorts makes both PENGUIN- and PGS-corrected estimates susceptible to concerns about poor portability across cohorts.

To characterize the portability of PGS- and PENGUIN-corrected estimates, we first distinguish between the types of portability concerns relevant to both methods (Fig. 3*A*). The portability of PGS has been written about extensively (25–27, 39, 40), and is defined as a loss of accuracy when GWAS are conducted in one cohort and the resulting PGS is applied to a second cohort. This loss of accuracy has many causes, including differences in allele frequency of causal variants, differences in LD between causal variants and tag variants identified in GWAS, differences in the role of the environment, and gene-by-gene or gene-by-environment interactions (26, 29, 41–43).

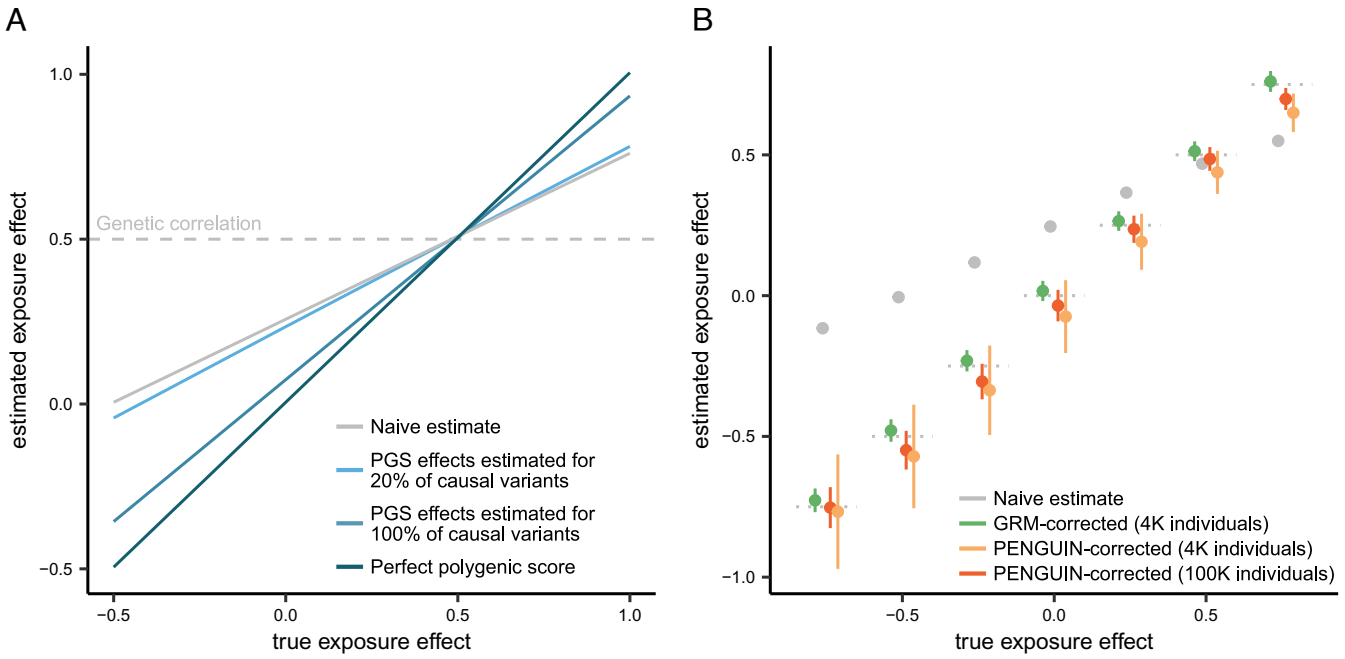

Fig. 2. Utility of GRM-corrected estimates. (*A*) PGS-corrected estimates of exposure effect, for i) PGS constructed from 20% of causal variants and effects estimated with noise; ii) PGS constructed from 100% of causal variants and effects estimated with noise; and iii) the perfect PGS, which exactly overlaps the diagonal, i.e. the $y = x$ line. (*B*) GRM- and PENGUIN-corrected estimates of exposure effect in simulated data. Horizontal dotted lines correspond to the diagonal, i.e. the $y = x$ line. Error bars correspond to mean ± 1 SD. In both panels, the naive estimate is obtained by directly regressing outcome on exposure.

A

	Concerns about portability of GWAS summary statistics in test cohort
GRM	**None**; exposure effect can be estimated entirely within-sample for **test** cohort if genotypes available
PENGUIN	Per-SNP heritability is assumed to be identical between **GWAS** and **test** cohort, but systematic differences can arise due to demographic history
Polygenic score	Variants and effect sizes ascertained from **GWAS** cohort may not be useful predictors in **test** cohort when LD patterns or allele frequencies differ

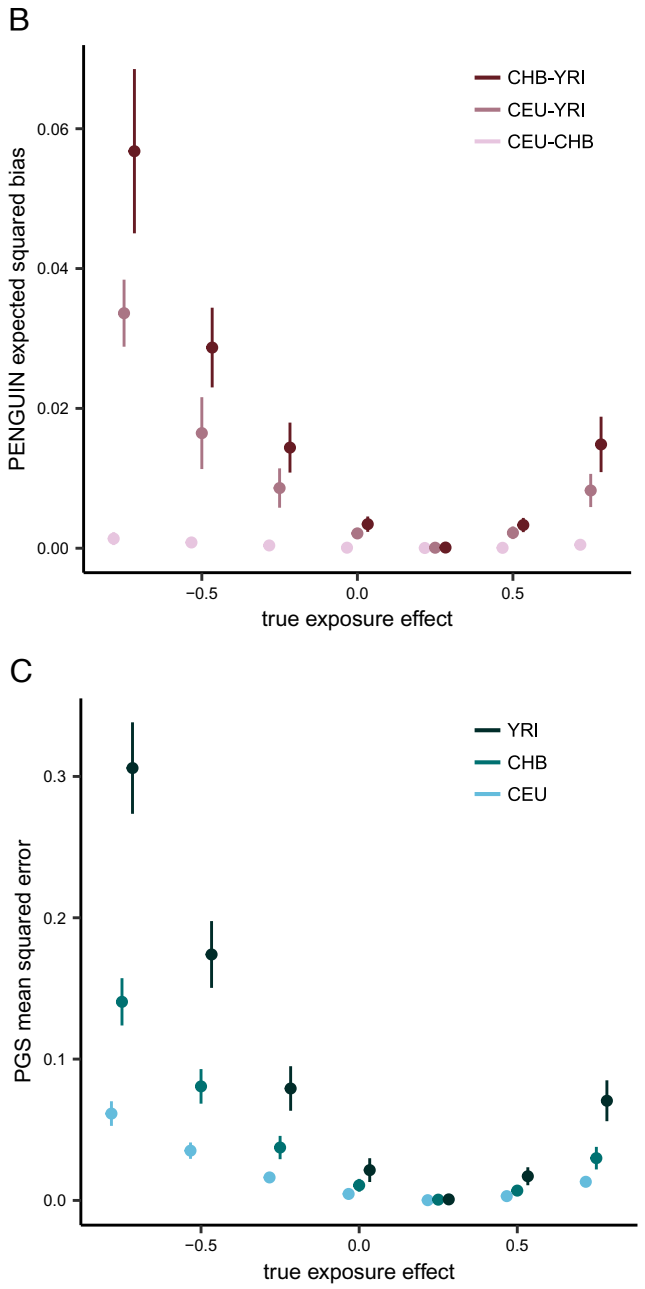

Fig. 3. Portability of methods. (*A*) Overview of portability concerns associated with GRM-, PENGUIN-, and PGS-corrected estimates of exposure effect. (*B*) Squared bias in PENGUIN-corrected estimates of exposure effect when GWAS summary statistics and individual-level data are obtained from two different populations. (*C*) Mean squared error in PGS-corrected estimates of exposure effect when GWAS summary statistics are obtained from simulated CEU individuals. Error bars correspond to mean ± 1 SD.

In contrast, the PENGUIN estimator faces a different kind of portability challenge. In the context of the PENGUIN estimator, "poor portability" means that PENGUIN-corrected estimates are inaccurate when GWAS summary statistics and individual-level phenotypes are obtained from two different cohorts. We show analytically that the PENGUIN estimator is no longer guaranteed to recover the true exposure effect when per-SNP heritability differs between the two cohorts due to factors such as demographic history or selection shaping allele frequencies, or cohort-specific gene-by-gene and gene-by-environment interactions modifying effect sizes (*SI Appendix, Supplementary Note 2*). Specifically, we find that the expected value of the PENGUIN estimator is

$$\mathbb{E}[\hat{b}_{\mathrm{PENG}}] \approx \frac{b\sigma_e^2 + \sum_{k=1}^{M} 2\beta_{X,k}\beta_{Y,k}\left(p_{k,\mathrm{ind}}(1-p_{k,\mathrm{ind}}) - p_{k,\mathrm{GWAS}}(1-p_{k,\mathrm{GWAS}})\right)}{\sigma_e^2 + \sum_{k=1}^{M} 2\beta_{X,k}^2\left(p_{k,\mathrm{ind}}(1-p_{k,\mathrm{ind}}) - p_{k,\mathrm{GWAS}}(1-p_{k,\mathrm{GWAS}})\right)}, \quad [3]$$

where $\beta_{X,k}$ and $\beta_{Y,k}$ are the effect of the kth variant on the exposure X and the outcome Y, and $p_{k,\mathrm{ind}}$ and $p_{k,\mathrm{GWAS}}$ are the frequency of the kth variant in the individual-level and GWAS cohorts respectively.

To clarify the magnitude of error associated with these two types of portability concerns, we simulated populations under a previously published demographic model inferred from 1,000 Genomes CEU (Utah residents with Northern and Western European ancestry), CHB (Han Chinese in Beijing, China), and YRI (Yoruba in Ibadan, Nigeria). Importantly, the portability concerns we modeled are purely due to differences in allele frequency and linkage disequilibrium arising from demographic history. We do not consider gene-by-environment interactions, population structure, or population-specific exposure effects, all of which could further reduce the portability of methods to control genetic confounding.

To characterize PENGUIN portability, we computed the squared bias of the estimator as $(\mathbb{E}[\hat{b}_{\mathrm{PENG}}] - b)^2$. As before, we find that the squared bias increases as the true exposure effect becomes more dissimilar to the genetic correlation (Fig. 3*B*). We also find that the error is most pronounced for comparisons between CEU and YRI, or CHB and YRI. This reflects systematically different per-SNP heritabilities in CEU and CHB relative to YRI, and can be explained by the different amounts of genetic drift in their respective demographic histories.

We next characterized PGS portability by comparing the mean squared error associated with out-of-sample GWAS and in-sample GWAS. Briefly, we performed GWAS on the exposure trait in simulated CEU individuals, clumped trait-associated variants to construct a PGS, and applied that PGS to control genetic confounding in CHB and YRI. Similar to PENGUIN, we find that PGS-corrected estimates exhibit poor portability between groups (Fig. 3*C*). However, we reiterate that unlike PENGUIN, PGS-corrected estimates can be biased even with in-sample GWAS (Fig. 2*A*).

Empirical Data Analyses. Having demonstrated the utility of the GRM in controlling genetic confounding, as well as its robustness to assumptions about trait heritability, we turned our attention to empirical data. Specifically, we sought to understand the effect of loneliness on health outcomes in the UK Biobank. Chronic

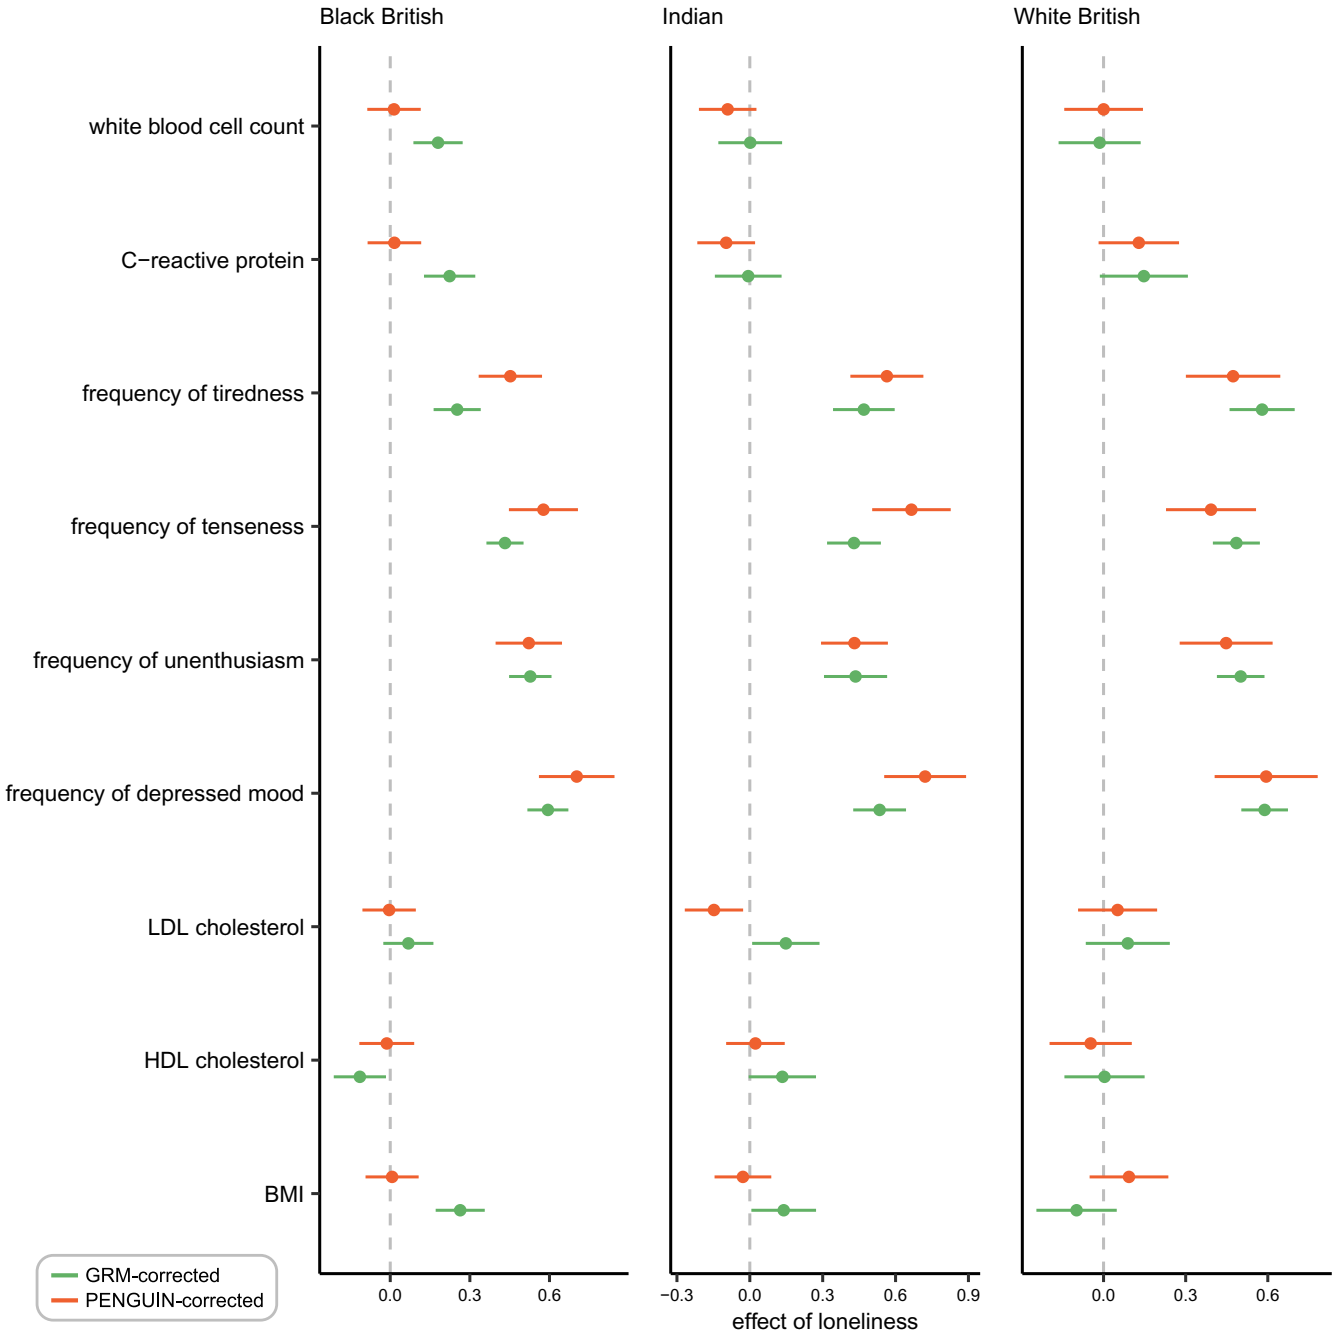

Fig. 4. Effect of loneliness on health outcomes in UK Biobank. GRM- and PENGUIN-corrected estimates of the effect of loneliness on inflammatory markers, mental health traits, and cardiovascular risk factors in the UK Biobank Black British, Indian, and White British. PENGUIN-corrected estimates use individual-level data from the specified group and GWAS summary statistics from the UK Biobank White British. Error bars correspond to mean ± 1 SD. Quantitative traits were inverse normal transformed, such that the x-axis scale has the same interpretation for all quantitative traits, and a different interpretation for categorical mood frequency phenotypes.

loneliness has an incidence of 5 to 10% in modern societies (44, 45), and is recognized as a significant public health concern (46). Using observational epidemiological studies, researchers have found that loneliness is associated with numerous health and social outcomes, including increased risk for cardiovascular disease (45, 47–50), increased inflammation (51, 52), and poor mental health (18, 49, 53, 54). At the same time, loneliness seems to be genetically correlated with some of these same health outcomes (9, 53), raising the possibility that it could be useful to control for potential genetic confounding when testing for loneliness as a risk factor.

We estimated the effect of a binary loneliness variable on three cardiovascular traits, four mental health traits, and two inflammation markers in the UK Biobank. Given that risk factors can manifest differently across social and environmental contexts (55–57), we separately analyzed the effect of loneliness in three subsets of the UK Biobank: a random sample of 7,000 individuals who identified as White and British (hereafter "White British"), 6,104 individuals who identified as Indian, and 8,483 individuals who identified as Black or Black British (hereafter "Black British"). For each subset, we computed both GRM- and PENGUIN-corrected estimates of the effect of

loneliness. To compute PENGUIN-corrected estimates, we used GWAS summary statistics generated from 361,194 White British individuals.

Across all three groups, we generally find that loneliness either has no effect or is associated with poorer health outcomes, corroborating existing literature (Fig. 4). We find that GRM- and PENGUIN-corrected estimates are most concordant for White British individuals ($r^2 = 0.96$, compared with 0.85 for Black British and 0.92 for Indian). This likely reflects the fact that PENGUIN-corrected estimates face portability challenges when individual-level phenotypes and GWAS summary statistics are obtained from two different populations.

This has implications for interpreting the effect of loneliness on health outcomes. For example, PENGUIN trained on a White British GWAS finds no effect of loneliness on inflammation markers in Black British individuals. However, GRM-corrected estimates indicate that loneliness is associated with both increased white blood cell count and increased C-reactive protein. This is particularly notable because loneliness does not appear to have any effect on inflammation in the Indian or White British subsets—thus, the association between loneliness and inflammation is only evident with GRM-corrected estimates, and only in the Black British subset. In some rare instances, we also find that PENGUIN-corrected estimates and GRM-corrected estimates have opposite signs, though we caution that these differences are not significant.

Discussion

We presented an approach that uses the genetic relatedness matrix (GRM) to control genetic confounding in observational epidemiological studies. Echoing earlier work (58), our results reiterate the conceptual similarity between LDSC- and GRM-based estimates of heritability. We show that the estimator obtained by the LDSC-based method PENGUIN is equivalent to the estimator obtained when the GRM is used to control genetic confounding. In parallel, our work also highlights the deep connections between statistical genetics and phylogenetics (59). When studying ecological and evolutionary relationships between traits, phylogeneticists often use interspecific data to test for an association between two traits. Crucially, researchers must account for confounding that results from species' shared ancestry via their phylogenetic relationships (60–63). This problem in phylogenetics bears a strong resemblance to genetic confounding in observational epidemiological studies, and in fact, the model we propose is closely analogous to one recently described by ref. 64 in the context of controlling phylogenetic confounding.

Our results provide several lessons for empirical studies. First, our results showcase the value of the GRM in settings where large, accurate GWAS are not available for a given trait or population. When biobank-scale data or GWAS summary statistics for the specific trait and population are readily available, existing methods such as PENGUIN are well suited to the task. Indeed, the proliferation of statistical methods based on summary statistics has been an important advance in statistical genetics, allowing the application of GWAS information out-of-sample and also decreasing computational burden (1). For understudied traits or populations, however, in-sample GWAS will be too noisy for accurate estimation, and out-of-sample GWAS risk portability concerns. In this setting, observational epidemiological studies can benefit from using GRM-corrected estimates of exposure effect.

Second, our results highlight the potential value of cohort-specific analyses in epidemiological research aimed at identifying effective interventions. In the UK Biobank, we observed differences in the effect of loneliness on multiple health and social outcomes across Black British, Indian, and White British individuals. The differences we observe could represent either gene-by-environment interactions, environment-by-environment interactions, or both, given that both genetic and nongenetic factors could vary among groups. Leveraging the GRM to characterize such interactions is an important direction for future research.

One important consideration in applying any data-analytic approach is computational tractability. For the results presented here, we performed variance-components analysis using GCTA-GREML, a method that partitions trait variance using restricted maximum likelihood. GCTA-GREML and similar methods become increasingly slow and computationally expensive as sample sizes increase into the tens of thousands, largely because of the difficulty of forming and storing the GRM, which has n^2 entries for a sample of size n. In this setting, a promising alternative is SCORE, a recent method developed by ref. 65. Much like GCTA-GREML, SCORE incorporates individual-level genotype data to partition variance for two traits. Unlike GCTA-GREML, however, SCORE uses stochastic trace estimation to avoid explicitly forming a GRM, allowing it to scale to hundreds of thousands of individuals. Although we do not evaluate SCORE in this study, it appears well suited to sample sizes that are too small for GWAS yet too large for GCTA-GREML.

Another important consideration is the construction of the GRM, which is central to our approach. There are different ways to construct GRMs, corresponding to different assumptions about how genetic similarity between individuals translates into phenotypic similarity. In this work, we use what is sometimes called the "canonical" GRM, defined as the variance–covariance matrix of additive genotype values that are standardized by the estimated SD of the genotype assuming Hardy–Weinberg equilibrium. Although widely used, this construction is not necessarily optimal: It measures similarity at observed variants, which may differ from genome-wide relatedness, and it corresponds to an implicit assumption that $\mathrm{E}(\beta^2) \propto \frac{1}{p(1-p)}$, such that the squared effect sizes of variants are inversely proportional to allelic heterozygosity. In our simulations, the canonical GRM generally provides adequate control of genetic confounding, even when the data are generated under alternative models. Nevertheless, alternative GRMs could readily be incorporated into observational epidemiological studies. In particular, recent work has highlighted the advantages of GRMs computed on the basis of estimated ancestral recombination graphs, especially when not all variants are genotyped (66–69). Other alternatives include GRMs that encode different assumptions about the effect size-allele frequency relationship, or explicitly incorporate information on linkage disequilibrium (LD) or heritability (38). The performance of the GRM under different generative models also remains to be characterized: Although our simulations incorporated several complexities, including population structure and assortative mating, many alternative data-generating mechanisms remain unexplored.

Our results present an opportunity to reconsider the causal model underlying observational epidemiological studies. In this work, we focused on estimating the effect of the exposure under a specific causal model where potentially correlated genetic variation affects both exposure and outcome, and researchers are interested in the interventional effect of the exposure,

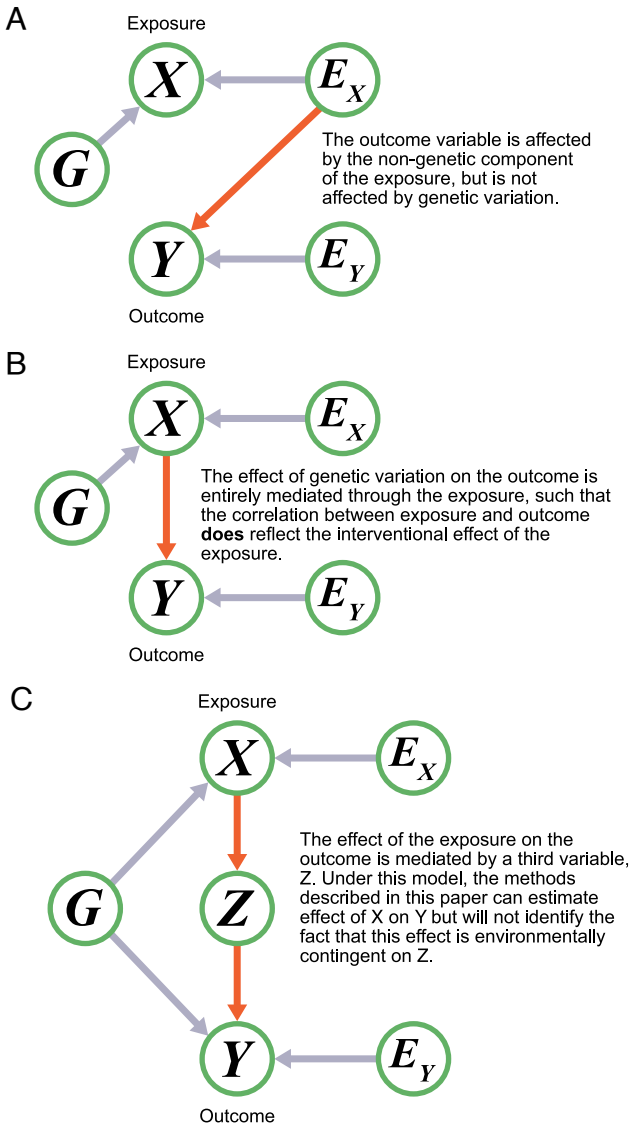

Fig. 5. Alternative causal models. (*A*) There is no effect of genetic variation on the outcome trait. (*B*) The effect of genetic variation is entirely mediated through the exposure trait. (*C*) The effect of genetic variation is mediated through the exposure trait, and the effect of the exposure is itself mediated by a third environmental variable.

i.e. the effect of the nongenetic component (Fig. 1*A*). We show in *SI Appendix*, *Supplementary Note* 1 that this model underlies a large body of existing work in epidemiology and social science that seeks to control genetic confounding by using the polygenic score for the exposure. Nevertheless, we will briefly discuss several alternative causal models that might occur to readers (Fig. 5). Two of these models, though apparently different from the generative model displayed in Fig. 1, are in fact special cases of it, and estimates derived from our approach (or, given the appropriate conditions, from PENGUIN) will provide germane information about causal relationships. However, the third is an extension of the model that requires a different interpretation.

First, suppose that there is no effect of genetic variation on the outcome trait (Fig. 5*A*). This can be viewed a special case of the generative model in Fig. 1*A* in which $\beta_Y = 0$, and therefore the genetic correlation between exposure and outcome is 0. In this setting, the naive regression of outcome on exposure is not subject to genetic confounding, but as we see in *SI Appendix*, *Supplementary Note* 1, the naive estimator is still subject to regression attenuation. This motivates the use of GRM- or PENGUIN-corrected estimates of exposure effect even when researchers believe the causal model at hand is Fig. 5*A*. In principle, both methods should successfully estimate the exposure effect even if the heritability of the outcome trait is 0.

A second possible scenario is one in which the effect of genetic variation on the outcome is entirely mediated through the exposure (Fig. 5*B*). In this setting, it seems plausible that researchers may want to operationalize a different definition of "interventional effect." In keeping with the existing literature, our original model defines the interventional effect of the exposure as the effect of the nongenetic component of the exposure. In contrast, in Fig. 5*B*, it may appear more suitable to consider the interventional effect as the effect of the total exposure on the outcome, rather than simply the effect of the nongenetic component of the exposure. However, Fig. 5*B* can be viewed as a special case of our generative model where exposure and outcome are perfectly genetically correlated, i.e. $\beta_Y = c\beta_X$ for some scalar c. In this setting, it turns out that the naive regression of outcome on exposure is statistically unbiased and equal in expectation to GRM- and PENGUIN-corrected estimates.

In contrast to the models in Fig. 5 *A* and *B*, the model in Fig. 5*C* is not a special case of generative model in Fig. 1*A*. In particular, now the effect of the exposure on the outcome is mediated by another variable, Z, which might be environmentally contingent. This model is motivated by Jencks' famous "red-headed child" thought experiment, which he pithily summarized, "If, for example, a nation refuses to send children with red hair to school, the genes that cause red hair can be said to lower reading scores. This does not tell us that children with red hair cannot learn to read" (70). In Jencks' example, there is a "genetic" effect of red hair on educational attainment, but it is transmitted via a societal decision to discriminate against red-haired children. A model like ours might indeed suggest, in Jencks' world, that an environmental modification of hair color—perhaps via sun exposure, bleach, or dye—might affect a child's educational attainment. But the data analysis suggested here would not, on its own, reveal the environmental contingency of the effect of red hair on educational attainment, which is the crucial consideration in assessing its modifiability. We raise this case to note that the ability to draw causal conclusions from variance decompositions rests on simplifying and sometimes restrictive assumptions (71). Thus, the outputs of our approach, despite their favorable properties under the model assumed here and in previous work, should always be treated with caution.

Ultimately, our results position the GRM as a flexible tool for addressing genetic confounding in observational epidemiological studies, particularly in settings where large, well-powered GWAS are unavailable. More broadly, our findings underscore the methodological similarities between disparate lines of work in epidemiology, statistical genetics, and phylogenetics—and the ways in which these connections can be leveraged to improve estimation across domains.

Materials and Methods

Estimating Exposure Effect.

GRM-corrected estimates. To construct a GRM for a particular set of individuals, we used all biallelic variants with minor allele frequency greater than 0.01. Unless specified otherwise, the GRM was constructed with an algorithm corresponding to the assumption that all variants contribute equally to heritability (GCTA flag --make-grm-alg 0).

To compute GRM-corrected estimates of the interventional exposure effect, we used bivariate GREML (32) to decompose the variance of the exposure and outcome into genetic and nongenetic components. We estimated the exposure effect using the estimated variance–covariance matrix for the nongenetic components of the exposure and outcome. Specifically, we estimated b as:

$$\hat{b} = \frac{\hat{\sigma}_{e_{XY}}}{\hat{\sigma}^2_{e_X}},$$

where $\hat{\sigma}_{e_{XY}}$ is the estimated covariance between the nongenetic components of the exposure and outcome, and $\hat{\sigma}^2_{e_X}$ is the estimated variance of the nongenetic component of the exposure.

To obtain SEs of the GRM-corrected estimator, we used two approaches. First, we computed the SD of the estimates obtained across multiple replicates in simulations. Second, we approximated the SE of the estimator using the delta method:

$$\mathrm{Var}\left(\frac{\hat{\sigma}_{e_{XY}}}{\hat{\sigma}^2_{e_X}}\right) = \left(\frac{\hat{\sigma}_{e_{XY}}}{\hat{\sigma}^2_{e_X}}\right)^2 \cdot \left(\frac{\mathrm{Var}(\hat{\sigma}_{e_{XY}})}{(\hat{\sigma}_{e_{XY}})^2} + \frac{\mathrm{Var}(\hat{\sigma}^2_{e_X})}{(\hat{\sigma}^2_{e_X})^2} - 2\frac{\mathrm{Cov}(\hat{\sigma}_{e_{XY}}, \hat{\sigma}^2_{e_X})}{\hat{\sigma}_{e_{XY}}\hat{\sigma}^2_{e_X}}\right) \quad [4]$$

To obtain the variance and covariance of the estimated components, we lightly modified GCTA (https://github.com/roshnipatel/GCTA) to report the Fisher information matrix, which is already computed by the software during each REML iteration. We verified in simulations that both SEs were comparable (*SI Appendix*, Fig. S3).

PENGUIN-corrected estimates. To compute PENGUIN-corrected estimates in simulated data, we first generated GWAS summary statistics for biallelic variants with minor allele frequency greater than 0.01 using plink2 (72). We next generated LD scores for our simulated genotypes using LDSC (1) with a window size of 100 kb. Using our GWAS summary statistics and LD scores, we ran individual-level PENGUIN with default parameters. In simulations of random mating, we did not include additional covariates when running GWAS or PENGUIN. In simulations of population structure or assortative mating, we included one principal component when generating GWAS summary statistics and subsequently computing PENGUIN-corrected estimates of exposure effect.

PGS-corrected estimates. Due to the many technical decisions necessary to construct polygenic scores (73), we sought to place an upper bound on the utility of PGS-corrected estimates through semianalytical means. We simulated trait architecture and genotypes and phenotypes for 100,000 individuals as described below (see *Random mating*). We then simulated the performance of PGS-corrected estimates in scenarios in which 20% and 100% of causal variants are identified in GWAS and included in the polygenic score. To do so, we started by randomly sampling the aforementioned proportion of causal variants, given that all variants contribute equally to heritability in expectation under the model we assume (see *Random mating*). In particular, for the subset of causal variants identified in GWAS, we simulated effect sizes estimated with noise using a standard approach (74). Estimated effect sizes $\hat{\beta}$ were drawn from a normal distribution centered on the true effect size, with a variance that depends on the true effect size, the allele frequency of the variant, f; and the number of individuals in the GWAS, n.

$$\hat{\beta} \sim N\left(\beta, \frac{1 - 2f(1-f)\beta^2}{2nf(1-f)}\right)$$

Using the estimated effect sizes $\hat{\beta}$, we then simulated an estimated polygenic score for each individual. Using this estimated polygenic score, we computed the PGS-corrected estimate of the exposure effect for ten replicates. We reported the mean of the estimated exposure effect across all ten replicates.

Simulating Genotypes and Phenotypes.

Random mating. Given that our empirical analyses rely on GWAS summary statistics obtained from the UK Biobank White British cohort, we conducted the majority of simulation analyses in simulated CEU individuals, except when analyzing portability (see *Analyzing portability* below). We simulated genotypes on chromosome 22 for CEU individuals under a published demographic model using the msprime engine in stdpopsim (34, 35, 75). For downstream analyses, we centered but did not standardize genotypes. To compare the utility of GRM- and PENGUIN-corrected estimates of exposure effect across different sample sizes, we randomly partitioned the simulated individuals into a GWAS cohort of 100,000 individuals and a test cohort of 4,000 individuals.

We simulated genetic values for exposure and outcome under an additive genetic architecture. To simulate causal variants, we first filtered variants for biallelic sites with a minor allele frequency greater than 0.01 in the test cohort. Next, we randomly sampled 10,000 variants to be the causal variants. We simulated effect sizes for the ith variant using a bivariate normal distribution

$$\begin{pmatrix}\beta_{X,i}\\ \beta_{Y,i}\end{pmatrix} \sim N\left(\begin{pmatrix}0\\0\end{pmatrix}, \begin{pmatrix}\frac{h^2}{2mf_i(1-f_i)} & \frac{\rho h^2}{2mf_i(1-f_i)}\\ \frac{\rho h^2}{2mf_i(1-f_i)} & \frac{h^2}{2mf_i(1-f_i)},\end{pmatrix}\right),$$

where h^2 is the heritability of the exposure trait, ρ is the effect size correlation between exposure and outcome, f_i is the allele frequency of the ith variant in the test cohort, and m is the number of causal variants. This distribution corresponds to the assumption that each causal variant contributes equally to heritability [often known in statistical genetics as "the alpha model" (38)]. For a subset of results, we also analyzed alternative models with a weaker dependence between effect size and frequency, such that $\mathrm{Var}(\beta_i) \propto (f_i(1-f_i))^\alpha$ for $\alpha \in [0, -0.3, -0.5]$.

We simulated trait-specific environmental effects as $N(0, 1-h^2)$. Conditional on genetic values and environmental effects, we simulated phenotypes for exposure and outcome under the generative model we introduce in the Results:

$$X = G\beta_X + e_X$$
$$Y = G\beta_Y + e_Y + be_X$$

Note that this model ensures that both the variance of genetic values and the heritability of the exposure, X, is equal to h^2, while only the variance of genetic values for the outcome is equal to h^2; the heritability will vary with the magnitude of b.

Population structure. To model population structure, we simulated a population split followed by no migration between subpopulations. Specifically, we simulated a split time of 200 generations ago, and a constant population size of 10,000 individuals for each subpopulation. For downstream analysis, we sampled 52,000 individuals from each subpopulation. We simulated both stratified and unstratified phenotypes, and stratified phenotypes were simulated such that the environmental effect of both exposure and outcome was shifted by 0.5 units.

To compute GRM-corrected estimates, we analyzed a joint sample of 2,000 individuals from each subpopulation. To compute PENGUIN-corrected estimates, we used a joint sample of 50,000 individuals from each subpopulation to perform a GWAS. We then used a (nonoverlapping) joint sample of 2,000 individuals from each subpopulation to comprise individual-level data for PENGUIN. We computed LD scores in a random subset of 4,000 individuals from one population, and incorporated principal components as covariates to correct population structure for both the GWAS and PENGUIN itself.

Assortative mating. We simulated genotypes and phenotypes under assortative mating using the software package xftsim (37). Given the computational complexity of assortative mating simulations, we were unable to simulate haplotypes at a sample size capable of computing PENGUIN-corrected estimates. To evaluate GRM-corrected estimates in simulations of assortative mating, we simulated founder haplotypes of 4,000 CEU individuals on chromosome 22 using stdpopsim (35). We seeded our simulations with these founder haplotypes and simulated 5 generations of cross-trait assortative mating with correlations ranging from -0.6 to $+0.6$. We simulated causal variants, effect sizes, and phenotypes under the models described previously (see *Random mating* above for more details).

Analyzing Portability.

Portability of PENGUIN-corrected estimates. To characterize the portability of PENGUIN-corrected estimates, we simulated 10,000 CEU, CHB, and YRI individuals under the demographic model inferred by ref. 34. We randomly

sampled 10,000 causal variants from the pool of biallelic sites that were segregating at a frequency of 1×10^{-4} in all populations. We simulated multivariate effect sizes independent of frequency (i.e. with $\alpha = 0$), given the impossibility of simulating frequency-dependent effects in a multipopulation model. We simulated a genetic correlation of 0.25 between exposure and outcome and an exposure heritability of 0.5 in CEU. Using simulated frequencies and effect sizes, we computed the analytical expectation of the PENGUIN estimator as follows, where $\beta_{X,k}$ and $\beta_{Y,k}$ are the effect of the kth variant on the exposure X and the outcome Y, and $p_{k,\text{ind}}$ and $p_{k,\text{GWAS}}$ are the frequency of the kth variant in the individual-level and GWAS cohorts respectively. (For a full derivation, see *SI Appendix, Supplementary Note* 2.)

$$\mathbb{E}[\hat{b}_{\text{PENG}}] \approx \frac{b\sigma_e^2 + \sum_{k=1}^{M} 2\beta_{X,k}\beta_{Y,k}\left(p_{k,\text{ind}}(1-p_{k,\text{ind}}) - p_{k,\text{GWAS}}(1-p_{k,\text{GWAS}})\right)}{\sigma_e^2 + \sum_{k=1}^{M} 2\beta_{X,k}^2\left(p_{k,\text{ind}}(1-p_{k,\text{ind}}) - p_{k,\text{GWAS}}(1-p_{k,\text{GWAS}})\right)}, \quad [5]$$

Across each of 10 simulation replicates, we computed the squared bias as $(\mathbb{E}[\hat{b}_{\text{PENG}}] - b)^2$ and reported the mean squared bias across simulations.

Portability of PGS-corrected estimates. To characterize the portability of PGS-corrected estimates, we simulated 100,000 CEU and 10,000 CHB and YRI individuals under the demographic model inferred by ref. 34. We simulated causal variants and their effect sizes as described in *Portability of PENGUIN-corrected estimates*. We generated GWAS summary statistics with plink2 (72) in the set of variants with MAF $\geq$0.01 in CEU, and clumped variants with $p \leq 1 \times 10^{-5}$, $r^2 \leq 0.5$, and a window size of 250 kb. We computed PGS by summing over the product of clumped variants and their estimated effect sizes, and reported the mean squared error in each population.

UK Biobank Analyses. We analyzed the effect of loneliness on three cardiovascular traits (BMI, HDL cholesterol, and LDL cholesterol); four mental health traits (frequency of depressed mood, unenthusiasm, tiredness, and tenseness); and two immune-related traits (C-reactive protein and white blood cell count). To encode loneliness phenotypes, we used responses to the question "Do you often feel lonely?", where "No" was encoded as 0, and "Yes" was encoded as 1. To encode categorical mood frequency phenotypes, we used the default encoding provided by the UK Biobank, where 1 corresponds to "Not at all," 2 corresponds to "Several days," 3 corresponds to "More than half the days," and 4 corresponds to "Nearly every day." To increase the interpretability of exposure effect, we inverse normal transformed phenotypes for quantitative traits (white blood cell count, C-reactive protein, LDL cholesterol, HDL cholesterol, and BMI) and removed outliers with a z-score greater than 3. We separately computed GRM- and PENGUIN-corrected estimates in three subsets of the UK Biobank: individuals who self-identified their ethnic background as Black or Black British (hereforth "Black British"); Indian; and both White and British (hereforth "White British").

To compute GRM-corrected estimates, we constructed a GRM in each subset using all genotyped variants with minor allele frequency greater than 0.01. Given that there were 442,719 individuals who identified as White British, we randomly sampled 7,000 individuals to obtain a tractable sample size for computing GRM-corrected estimates. To compute GRM-corrected estimates in the Black British and Indian subsets, we used data from all 8,483 Black British and 6,104 Indian individuals. We performed variance components analysis using bivariate GREML and included covariates that captured age, sex, assessment center, and 10 principal components.

We computed PENGUIN-corrected estimates using GWAS summary statistics obtained from all UK Biobank White British individuals (http://www.nealelab.is/uk-biobank/). We used individual-level phenotype data from the random sample of 7,000 White British individuals and all Black British and Indian individuals. As before, we included covariates that captured age, sex, and 10 principal components.

Data, Materials, and Software Availability. Software for generating GRM-corrected estimates of exposure effect and code for replicating results in this manuscript is available at https://github.com/roshnipatel/grm-corrected-exposure (76).

ACKNOWLEDGMENTS. We thank members of the Edge, Pennell, and Mooney labs for helpful conversations about this work. We are also grateful to Graham Coop for thoughtful comments on an earlier draft. R.A.P., J.G.S., and M.D.E. were supported by NIH grant R35GM137758. M.P. was supported by NIH grant R35GM151348.

Author affiliations: [a] Department of Data Science, University of Oregon, Eugene, OR 97403; [b] Institute of Ecology and Evolution, University of Oregon, Eugene, OR 97403; [c] Department of Quantitative and Computational Biology, University of Southern California, Los Angeles, CA 90089; and [d] Department of Computational Biology, Cornell University, Ithaca, NY 14853

1. B. K. Bulik-Sullivan *et al.*, LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nat. Genet.* **47**, 291–295 (2015).
2. N. Sinnott-Armstrong, S. Naqvi, M. Rivas, J. K. Pritchard, GWAS of three molecular traits highlights core genes and pathways alongside a highly polygenic background. *eLife* **10**, e58615 (2021).
3. L. Yengo *et al.*, A saturated map of common genetic variants associated with human height. *Nature* **610**, 704–712 (2022).
4. B. Bulik-Sullivan *et al.*, An atlas of genetic correlations across human diseases and traits. *Nat. Genet.* **47**, 1236–1241 (2015).
5. B. M. D'Onofrio, Q. A. Class, B. B. Lahey, H. Larsson, Testing the developmental origins of health and disease hypothesis for psychopathology using family-based quasi-experimental designs. *Child Dev. Perspect.* **8**, 151–157 (2014).
6. B. M. D'Onofrio, A. Sjölander, B. B. Lahey, P. Lichtenstein, A. S. Öberg, Accounting for confounding in observational studies. *Annu. Rev. Clin. Psychol.* **16**, 25–48 (2020).
7. A. E. Eney, S. Tsang, J. A. Delaney, E. Turkheimer, G. E. Duncan, Cross-sectional association between soda consumption and body mass index in a community-based sample of twins. *Nutr. J.* **16**, 48 (2017).
8. S. R. Jaffee, T. S. Price, The implications of genotype–environment correlation for establishing causal processes in psychopathology. *Dev. Psychopathol.* **24**, 1253–1264 (2012).
9. R. Moshtael, M. E. Lynch, G. E. Duncan, C. R. Beam, A genetically informed study of the association between perceived stress and loneliness. *Behav. Genet.* **54**, 268–277 (2024).
10. K. P. Harden, P. D. Koellinger, Using genetics for social science. *Nat. Hum. Behav.* **4**, 567–576 (2020).
11. D. Choi, K. J. Tsuchiya, N. Takei, Interaction effect of oxytocin receptor (OXTR) rs53576 genotype and maternal postpartum depression on child behavioural problems. *Sci. Rep.* **9**, 7685 (2019).
12. D. Manousaki *et al.*, Vitamin D levels and risk of type 1 diabetes: A Mendelian randomization study. *PLoS Med.* **18**, e1003536 (2021).
13. G. Davey Smith, S. Ebrahim, "Mendelian randomization": Can genetic epidemiology contribute to understanding environmental determinants of disease? *Int. J. Epidemiol.* **32**, 1–22 (2003).
14. G. Davey Smith, G. Hemani, Mendelian randomization: Genetic anchors for causal inference in epidemiological studies. *Hum. Mol. Genet.* **23**, R89–R98 (2014).
15. D. W. Belsky *et al.*, The genetics of success: How single-nucleotide polymorphisms associated with educational attainment relate to life-course development. *Psychol. Sci.* **27**, 957–972 (2016).
16. D. Conley, Socio-genomic research using genome-wide molecular data. *Annu. Rev. Sociol.* **42**, 275–299 (2016).
17. A. G. Cuevas, F. D. Mann, D. R. Williams, R. F. Krueger, Discrimination and anxiety: Using multiple polygenic scores to control for genetic liability. *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2017224118 (2021).
18. S. L. Lee *et al.*, The association between loneliness and depressive symptoms among adults aged 50 years and older: A 12-year population-based cohort study. *Lancet Psychiatry* **8**, 48–57 (2021).
19. L. Schmitz, D. Conley, Modeling gene-environment interactions with quasi-natural experiments. *J. Pers.* **85**, 10–21 (2017), https://onlinelibrary.wiley.com/doi/pdf/10.1111/jopy.12227.
20. B. Zheng, Q. Lu, J. Fletcher, *Estimating Causal Effects of Fertility on Life Course Outcomes: Evidence Using A Dyadic Genetic Instrumental Variable Approach* (National Bureau of Economic Research, Cambridge, MA, Working Paper w30955, 2023).
21. T. A. DiPrete, C. A. P. Burik, P. D. Koellinger, Genetic instrumental variable regression: Explaining socioeconomic and health outcomes in nonexperimental data. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E4970–E4979 (2018).
22. J. B. Pingault *et al.*, Genetic sensitivity analysis: Adjusting for genetic confounding in epidemiological associations. *PLoS Genet.* **17**, e1009590 (2021).
23. J. B. Pingault *et al.*, Research Review: How to interpret associations between polygenic scores, environmental risks, and phenotypes. *J. Child Psychol. Psychiatry* **63**, 1125–1139 (2022), https://onlinelibrary.wiley.com/doi/pdf/10.1111/jcpp.13607.
24. M. J. Uddin, C. Hjorthøj, T. Ahammed, M. Nordentoft, C. T. Ekstrøm, The use of polygenic risk scores as a covariate in psychological studies. *Methods Psychol.* **7**, 100099 (2022).
25. A. R. Martin *et al.*, Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat Genet.* **51**, 584–591 (2019).
26. H. Mostafavi *et al.*, Variable prediction accuracy of polygenic scores within an ancestry group. *eLife* **9**, e48376 (2020).
27. J. Y. Wang *et al.*, Three open questions in polygenic score portability. *Nat. Commun.* **17**, 942 (2026).

28. Z. Zhao *et al*., Controlling for polygenic genetic confounding in epidemiologic association studies. *Proc. Natl. Acad. Sci. U.S.A.* **121**, e2408715121 (2024).
29. R. A. Patel *et al*., Characterizing selection on complex traits through conditional frequency spectra. *Genetics* **229**, iyae210 (2025).
30. M. C. Mills, C. Rahal, A scientometric review of genome-wide association studies. *Commun. Biol.* **2**, 1–11 (2019).
31. D. Speed, D. J. Balding, Relatedness in the post-genomic era: Is it still useful? *Nat. Rev. Genet.* **16**, 33–44 (2015).
32. S. Lee, J. Yang, M. Goddard, P. Visscher, N. Wray, Estimation of pleiotropy between complex diseases using single-nucleotide polymorphism-derived genomic relationships and restricted maximum likelihood. *Bioinformatics* **28**, 2540–2542 (2012).
33. J. Yang *et al*., Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* **42**, 565–569 (2010).
34. R. N. Gutenkunst, R. D. Hernandez, S. H. Williamson, C. D. Bustamante, Inferring the joint demographic history of multiple populations from multidimensional SNP frequency data. *PLoS Genet.* **5**, e1000695 (2009).
35. J. R. Adrion *et al*., A community-maintained standard library of population genetic models. *eLife* **9**, e54967 (2020).
36. R. Border *et al*., Assortative mating biases marker-based heritability estimators. *Nat. Commun.* **13**, 660 (2022).
37. R. Border *et al*., Simple models of non-random mating and environmental transmission bias standard human genetics statistical methods. bioRxiv [Preprint] (2024), https://doi.org/10.1101/2024.10.16.618755 (Accessed 10 October 2024).
38. D. Speed *et al*., Reevaluation of SNP heritability in complex human traits. *Nat Genet.* **49**, 986–992 (2017).
39. J. J. Berg *et al*., Reduced signal for polygenic adaptation of height in UK Biobank. *eLife* **8**, e39725 (2019).
40. M. Sohail *et al*., Polygenic adaptation on height is overestimated due to uncorrected stratification in genome-wide association studies. *eLife* **8**, e39702 (2019).
41. R. A. Patel *et al*., Genetic interactions drive heterogeneity in causal variant effect sizes for gene expression and complex traits. *Am. J. Hum. Genet.* **109**, 1286–1297 (2022).
42. Y. Wang *et al*., Theoretical and empirical quantification of the accuracy of polygenic scores in ancestry divergent populations. *Nat. Commun.* **11**, 3865 (2020).
43. S. Yair, G. Coop, Population differentiation of polygenic score predictions under stabilizing selection. *Philo. Trans. R. Soc. B* **377**, 20200416 (2022).
44. M. E. Beutel *et al*., Loneliness in the general population: Prevalence, determinants and relations to mental health. *BMC Psychiatry* **17**, 97 (2017).
45. C. Hakulinen *et al*., Social isolation and loneliness as risk factors for myocardial infarction, stroke and mortality: UK Biobank cohort study of 479 054 men and women. *Heart* **104**, 1536–1542 (2018).
46. Office of the Surgeon General, *Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community* (US Department of Health and Human Services, Washington, D.C., 2023).
47. Z. Liu *et al*., Association between loneliness and incident atrial fibrillation across different genetic predisposition: Findings from the UK Biobank. *Eur. J. Med. Res.* **29**, 1–11 (2024).
48. E. Paul, F. Bu, D. Fancourt, Loneliness and risk for cardiovascular disease: Mechanisms and future directions. *Curr. Cardiol. Rep.* **23**, 1–7 (2021).
49. L. Rødevand *et al*., Polygenic overlap and shared genetic loci between loneliness, severe mental disorders, and cardiovascular disease risk factors suggest shared molecular mechanisms. *Transl. Psychiatry* **11**, 3 (2021).
50. N. Xia, H. Li, Loneliness, social isolation, and cardiovascular health. *Antioxid. Redox Signal.* **28**, 837–851 (2018).
51. T. Matthews *et al*., Social isolation, loneliness, and inflammation: A multi-cohort investigation in early and mid-adulthood. *Brain Behav. Immun.* **115**, 727–736 (2024).
52. K. Van Bogart *et al*., The association between loneliness and inflammation: Findings from an older adult sample. *Front. Behav. Neurosci.* **15**, 801746 (2022).
53. A. Abdellaoui *et al*., Phenome-wide investigation of health outcomes associated with genetic predisposition to loneliness. *Hum. Mol. Genet.* **28**, 3853–3865 (2019).
54. J. Dennis *et al*., Genetic risk for major depressive disorder and loneliness in sex-specific associations with coronary artery disease. *Mol. Psychiatry* **26**, 4254–4264 (2021).
55. R. A. Nianogo *et al*., Risk factors associated with Alzheimer disease and related dementias by sex and race and ethnicity in the US. *JAMA Neurol.* **79**, 584–591 (2022).
56. E. Proto, C. Quintana-Domeque, COVID-19 and mental health deterioration by ethnicity and gender in the UK. *PLoS One* **16**, e0244419 (2021).
57. M. T. Williams, L. K. Chapman, J. Wong, E. Turkheimer, The role of ethnic identity in symptoms of anxiety and depression in African Americans. *Psychiatry Res.* **199**, 31–36 (2012).
58. B. Bulik-Sullivan, Relationship between LD Score and Haseman-Elston Regression (2015).
59. J. G. Schraiber, M. D. Edge, M. Pennell, Unifying approaches from statistical genetics and phylogenetics for mapping phenotypes in structured populations. *PLoS Biol.* **22**, e3002847 (2024).
60. J. Felsenstein, Phylogenies and the comparative method. *Am. Nat.* **125**, 1–15 (1985).
61. M. Lynch, Methods for the analysis of comparative data in evolutionary biology. *Evolution* **45**, 1065–1080 (1991), https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1558-5646.1991.tb04375.x.
62. E. P. Martins, T. F. Hansen, Phylogenies and the comparative method: A general approach to incorporating phylogenetic information into the analysis of interspecific data. *Am. Nat.* **149**, 646–667 (1997).
63. J. C. Uyeda, R. Zenil-Ferguson, M. W. Pennell, Rethinking phylogenetic comparative methods. *Syst. Biol.* **67**, 1091–1109 (2018).
64. M. Westoby, L. Yates, B. Holland, B. Halliwell, Phylogenetically conservative trait correlation: Quantification and interpretation. *J. Ecol.* **111**, 2105–2117 (2023), https://onlinelibrary.wiley.com/doi/pdf/10.1111/1365-2745.14150.
65. Y. Wu *et al*., Fast estimation of genetic correlation for biobank-scale data. *Am. J. Hum. Genet.* **109**, 24–32 (2022).
66. C. Fan, N. Mancuso, C. W. Chiang, A genealogical estimate of genetic relationships. *Am. J. Hum. Genet.* **109**, 812–824 (2022).
67. V. Link *et al*., Tree-based QTL mapping with expected local genetic relatedness matrices. *Am. J. Hum. Genet.* **110**, 2077–2091 (2023).
68. B. C. Zhang, A. Biddanda, A. F. Gunnarsson, F. Cooper, P. F. Palamara, Biobank-scale inference of ancestral recombination graphs enables genealogical analysis of complex traits. *Nat Genet.* **55**, 768–776 (2023).
69. B. Lehmann *et al*., On args, pedigrees, and genetic relatedness matrices. *Genetics* **232**, iyaf219 (2025).
70. C. Jencks *et al*., *Inequality: A Reassessment of the Effect of Family and Schooling in America* (Basic Books, New York, 1972).
71. R. Lewontin, Annotation: the analysis of variance and the analysis of causes. *Am. J. Hum. Genet.* **26**, 400–411 (1974).
72. C. C. Chang *et al*., Second-generation PLINK: Rising to the challenge of larger and richer datasets. *Gigascience* **4**, s13742–015-0047-8 (2015).
73. Y. S. Lin *et al*., Causal variant capture in genotype discovery approaches drives polygenic prediction performance across traits and populations. bioRxiv [Preprint] (2026). https://doi.org/10.1101/2025.03.18.644029 (Accessed 19 February 2026).
74. B. Vilhjálmsson *et al*., Modeling linkage disequilibrium increases accuracy of polygenic risk scores. *Am. J. Hum. Genet.* **97**, 576–592 (2015).
75. F. Baumdicker *et al*., Efficient ancestry and mutation simulation with msprime 1.0. *Genetics* **220**, iyab229 (2022).
76. R. A. Patel, grm-corrected-exposure. GitHub. https://github.com/roshnipatel/grm-corrected-exposure. Accessed 9 December 2025.