Assumptions about frequency-dependent architectures of complex traits bias measures of functional enrichment

Shadi Zabad¹, Aaron P. Ragsdale², Rosie Sun², Yue Li^{*,1,3}, and Simon Gravel^{**,2,3}

¹Department of Computer Science, McGill University, Montreal, QC, Canada

²Department of Human Genetics, McGill University, Montreal, QC, Canada

³Quantitative Life Science, McGill University, Montreal, QC, Canada

Abstract

Linkage-Disequilibrium Score Regression (LDSC) is a popular framework for analyzing GWAS summary statistics that allows for estimating SNP heritability, confounding, and functional enrichment of genetic variants with different annotations. Recent work has highlighted the influence of implicit and explicit assumptions of the model on the biological interpretation of the results. In this work, we explored a formulation of LDSC that replaces the r^2 measure of LD with a recently-proposed unbiased estimator of the D^2 statistic. In addition to modest statistical difference across estimators, this derivation highlighted implicit and unrealistic assumptions about the relationship between allele frequency, effect size, and annotation status. We carry out a systematic comparison of alternative LDSC formulations by applying them to summary statistics from 47 GWAS traits. Our results show that commonly used models likely underestimate functional enrichment. These results highlight the importance of calibrating the LDSC model to achieve a more robust understanding of polygenic traits.

Keywords: LD Score Regression; GWAS; Functional Enrichment; Heritability Estimation; Genetic Architecture

*Corresponding author. E-mail: yueli@cs.mcgill.ca

**Corresponding author. E-mail: simon.gravel@mcgill.ca

1 Introduction

Linkage-Disequilibrium Score Regression (LDSC) provides a general framework for understanding the architecture of polygenic traits from GWAS summary statistics (Bulik-Sullivan, Loh, et al., 2015; Pasaniuc and Price, 2017). Because of its tractability and computational efficiency, it has become a central tool to estimate statistical and population genetic quantities, such as SNP heritability (h_{SNP}^2) and confounding in GWAS results (Bulik-Sullivan, Loh, et al., 2015), functional enrichment (Finucane et al., 2015), cross-trait genetic correlations (Bulik-Sullivan, Finucane, et al., 2015) as well as genetic correlations across ethnic groups for the same trait (Shi et al., 2019). This diversity of applications stems from the ease of incorporating various assumptions about the architecture of a trait in the LDSC framework. In its original formulation, the LDSC model made the implicit assumption that the expected phenotypic variance explained by each genetic variant was independent of allele frequency. However, subsequent work has shown that this assumption is restrictive and results in downwardly biased estimates of heritability (Speed, Cai, Johnson, Nejentsev, and Balding, 2017).

To guard against these potential sources of bias, the trend in the field has been towards incorporating more expressive models that describe the expected squared effect of a variant as a function of different genetic annotations (Finucane et al., 2015; Gazal et al., 2017; Gazal et al., 2018; Hormozdiari et al., 2018; Hujoel, Gazal, Hormozdiari, van de Geijn, and Price, 2019), such as MAF decile bins and whether a SNP occurs in an annotated enhancer region. An advantage of this approach, commonly referred to as Stratified LDSC (S-LDSC), is that it allows for arbitrary relationships between allele frequency and mean squared effect sizes, thus reducing the need for a priori assumptions that can bias results. It also provides measures of functional enrichment for the categories incorporated into the model. Despite these many advantages, previous work has shown that the way we model the dependence of effect size on functional annotations can result in substantially different estimates of partitioned heritability and functional enrichment (Gazal, Marquez-Luna, Finucane, and Price, 2019; Speed et al., 2017; Speed, Hemani, Johnson, and Balding, 2012).

In this work, we reformulate the LD Score regression framework by replacing the r^2 measure of LD with an estimate based on the D^2 statistic, which has desirable statistical properties

(Ragsdale and Gravel, 2020). Even though the choice of the LD statistic itself has little effect on heritability estimates, this formulation highlights implicit assumptions of commonly-used stratified models. Specifically, these models imply relationships between allele frequency, effect size, and annotation status that do not agree with empirical observations. While these assumptions have a small impact on global estimates of SNP heritability, we show that they can result in systematic biases in estimates of functional enrichment. To correct for these biases, we propose using modified S-LDSC models that better capture the empirical relationship between allele frequency and effect size for annotated variants.

2 Materials and Methods

2.1 LD Score Regression as a function of D^2

In our formulation of LD Score Regression, we assume a standard linear polygenic model $Y = X\beta + \varepsilon$, where Y is a vector of phenotypes for N individuals, X is a $N \times M$ mean-centered (but not variance-normalized) matrix that encodes the genotype at M SNPs, β is a vector of random and independent effect sizes, and ε is a vector of random and independent environmental and other effects. Under this model, assuming no confounding, the expectation of the χ^2 association statistic at locus j can be expressed as a function of the squared covariance in allele frequency D_{jk}^2 at variants j and k as:

$$\mathbb{E}[\chi_j^2] \approx N \sum_k \frac{4}{\text{Var}(X_j)} D_{jk}^2 \text{Var}(\beta_k) + 1. \quad (1)$$

Here, $\text{Var}(\beta_k)$ is the variance of causal effect size at SNP k and $\text{Var}(X_j) = 2p_j(1 - p_j)$ is the variance in genotype across individuals at the focal SNP j , where p_j is its minor allele frequency in the population. In LDSC, we seek to build a parameterized model for $\text{Var}(\beta_k)$ that can be fitted to the observed χ^2 . Since $\text{Var}(\beta_k)$ is the mean squared effect size for variant k , we can then sum over all SNPs to obtain an estimate of SNP heritability. While the simplest model takes $\text{Var}(\beta_k)$ to be a constant across SNPs, various models have been proposed that take into account allele frequencies as well as functional annotations. Here we consider models of the

form

$$\text{Var}(\beta_k) = \text{Var}(X_k)^{-\alpha} \left[\tau_0 + \theta \left(\sum_{c=1}^{96} \tau_c a_c(k) \right) \right] \quad (2)$$

$$\alpha \in \{0, 0.25, 0.5, 0.75, 1\}, \quad \theta \in \{0, 1\}.$$

The mean squared effect size $\text{Var}(\beta_k)$ of SNP k is defined in terms of two components: a dependence on minor allele frequency (MAF) through the parameter α , and a dependence on a set of 96 functional annotations $a_c(k)$ that were incorporated into the baseline-LD model (v2.2) (Gazal et al., 2017; Gazal et al., 2019, see Web Resources). This dependence is controlled by the parameter θ . If $\theta = 0$, we recover the univariate models introduced in Bulik-Sullivan, Loh, et al., 2015. If $\theta = 1$ we obtain the stratified models that are most commonly used in more recent work (Finucane et al., 2015; Gazal et al., 2017; Gazal et al., 2018; Hormozdiari et al., 2018; Hujoel et al., 2019).

The contribution of each annotation a_c to SNP heritability is modulated by the corresponding parameter τ_c , which is estimated in the LDSC regression procedure. In this formulation, τ_0 corresponds to the “base” annotation that includes all SNPs (e.g. $a_0(k) = 1$ for all k). Overall, in this framework, the expected heritability contributed by a SNP is defined as $\mathbb{E}[h_j] = \text{Var}(X_j)\text{Var}(\beta_j)$ (Speed and Balding, 2019; Speed, Holmes, and Balding, 2020). Table 1 summarizes some of the relevant models that will be analyzed in this paper.

Heritability model	Parameters	Comments
Constant effect size model	$\alpha = 0; \theta = 0$	The variance in effect size $\text{Var}(\beta_j)$ is independent of allele frequency.
GCTA model (Yang et al., 2010, Bulik-Sullivan, Loh, et al., 2015)	$\alpha = 1; \theta = 0$	The expected heritability $\mathbb{E}[h_j]$ contributed by a variant is constant.
Baseline-LD model (v2.2) (Gazal, Marquez-Luna, Finucane, and Price, 2019)	$\alpha = 1; \theta = 1$	A stratified LDSC model that includes a total of 96 functional annotations, including continuous LD-related annotations and MAF bins. This model gives more weight to annotations found in rare variants.
MAF-independent baseline-LD model	$\alpha = 0; \theta = 1$	Similar to the baseline-LD model, though the effect of each annotation is independent of allele frequency.

Table 1: The SNP-heritability models analyzed in this paper.

In previous work, α has often been treated as a fixed parameter, usually set to $\alpha = 1$. In the univariate case ($\theta = 0$), this is a mathematically convenient choice that implies constant variance explained by all SNPs and leads to simple expressions of LD scores in terms of the LD statistic r^2 (as detailed below). In turn, this implies that rare variants have much larger per-allele biological effect than common variants. In other research settings, α is treated as a continuous parameter that can be fit to data to measure the strength of negative selection for a given trait. Average inferred values of α for a number of UK Biobank traits ranged from 0.2 to 0.5, with a mean value of roughly 0.38 (Schoech et al., 2019; Speed et al., 2017; Zeng et al., 2018). However, these values of α were inferred for the univariate setting and it is not clear if the interaction between MAF and the various functional annotations can still be characterized by these global estimates. In fact, a recent report using an updated version of the LDAK and SumHer models (Speed and Balding, 2019; Speed et al., 2012; Speed et al., 2020) found variations in the α values ranging from 0.23 to 0.67 for SNPs in all but one functional category, and 0.25 for all variants.

Despite these observations, in standard applications of the stratified LDSC models ($\theta = 1$), it is still commonly assumed that $\alpha = 1$ (Gazal et al., 2017; Gazal et al., 2018; Hormozdiari et al., 2018; Hujoel et al., 2019). The choice $\alpha = 1$ is not particularly biologically plausible — it is a leftover from the mathematically convenient formulation of LDSC in terms of the r^2 statistic. Equation (2) shows that the assumption $\alpha = 1$ in the stratified setting implies that the contribution of each annotation to predicted effect size is much larger for rare variants. Put another way, under the $\alpha = 1$ model, rare variants in a given functionally enriched category c are predicted to have much larger per-allele biological effect than common variants in the same category. Since this assumption has been found to not be valid in genome-wide analyses (Schoech et al., 2019; Zeng et al., 2018), there is no compelling reason to think that it holds within specific categories of variants.

3 Results

3.1 Comparing LD scores across models and estimators

The family of models outlined in Equations (1) and (2) may be equivalently expressed in terms of LD scores:

$$\mathbb{E}[\chi_j^2] \approx N \left[\tau_0 \ell(j, 0) + \theta \sum_{c=1}^{96} \tau_c \ell(j, c) \right] + 1, \quad (3)$$

where the LD score $\ell(j, c)$ of variant j and category c is defined as the sum of the annotation value $a_c(k)$ multiplied by a MAF-weighted linkage disequilibrium term $L_{jk}(\alpha)$ over all neighbouring SNPs k :

$$\begin{aligned} \ell(j, c) &:= \sum_k L_{jk}(\alpha) a_c(k) \\ L_{jk}(\alpha) &= \frac{4D_{jk}^2}{\text{Var}(X_j)} \text{Var}(X_k)^{-\alpha} = r_{jk}^2 \text{Var}(X_k)^{1-\alpha}. \end{aligned} \quad (4)$$

In standard applications of LDSC, the MAF-weighted LD measure $L_{jk}(\alpha)$ is typically taken to be r_{jk}^2 , which implies $\alpha = 1$. However, for arbitrary α , it can be expressed more generally in terms of a product of the variances at SNPs j and k with the squared covariance D_{jk}^2 or the squared correlation r_{jk}^2 . This generalized measure can be estimated in-sample or from a reference panel using, e.g., the following two estimators:

$$\begin{aligned} \hat{L}_{jk,D} &= \frac{\widehat{4D_{jk}^2}}{\widehat{\text{Var}(X_j)}} \widehat{\text{Var}(X_k)^{-\alpha}} \\ \hat{L}_{jk,r} &= \widehat{r_{jk}^2} \widehat{\text{Var}(X_k)^{1-\alpha}} \end{aligned}$$

where $\widehat{D_{jk}^2}$ is the unbiased estimator of Ragsdale and Gravel, 2020 and $\widehat{r_{jk}^2}$ is the bias-corrected estimator used in the original LDSC model (Bulik-Sullivan, Loh, et al., 2015; Yin and Fan, 2001). The estimator for $\hat{L}_{jk,r}$ has a particularly simple form for $\alpha = 1$, which is probably a reason for the popularity of this parameter choice. To evaluate whether choices of estimators affect LDSC regression results, we analyzed the differences between the two statistical estima-

tors across different values of α . To perform this analysis, we computed LD scores according to Equation (4) using both estimators and α values of $\{0, 0.25, 0.5, 0.75, 1\}$ for three of the super populations in the 1000 Genomes Project (Africans, Asians and Europeans) (Altshuler et al., 2010; The 1000 Genomes Project Consortium, 2012).

Both estimators produce highly correlated estimates ($R^2 = 0.99$) (Figure 1 (a)), and their bias is comparable even for small panel sample size (Supplementary Material, Supplementary Figure 1). The consistency of these two estimators suggests that they perform equally well in estimating LD scores irrespective of the value of α . Thus, biological realism might be a more relevant criterion when choosing α than statistical convenience. For the remainder of the discussion in the main text, we will use the $\hat{L}_{jk,D}$ estimator.

Next, we sought to highlight the influence of model assumptions on the LD scores computed from the 1000 Genomes data. For clarity of exposition, we focus on the univariate case where the LDSC model can be written as: $\mathbb{E}[\chi_j^2] \approx N\tau_0\ell(j, 0) + 1$ and the LD score is simply the sum of the generalized LD measures $L_{jk}(\alpha)$. To qualitatively compare the LD scores with different values of α in a consistent manner, we normalize them such that the slope of the univariate regression becomes $\frac{N}{M}h_{SNP}^2$, where h_{SNP}^2 is the total heritability (see Appendix C.1).

As expected, Figure 1(b) shows that common variants have higher LD scores under the $\alpha = 0$ model compared to the $\alpha = 1$ model, and the opposite trend is observed for rare variants. This effect is less stark for intermediate values of α (Supplementary Figure 2). In practical terms, the choice of α determines which categories of SNPs will have higher weight in the LDSC regression and, as has been documented before (Speed et al., 2017), this will in turn influence the global estimates of SNP heritability.

3.2 Estimates of SNP-heritability and Confounding for 47 GWAS Traits

To understand the implications of choosing different values of α on global estimates of SNP heritability and confounding, we applied the LD Score regression framework to a total of 47 GWAS traits (41 independent traits), for which summary statistics have been previously analyzed (Gazal et al., 2019; Hujoel et al., 2019) (Supplementary Table S1, see Web Resources). We confirmed that the choice of statistical estimator has a modest impact on estimates of global

heritability and confounding (Supplementary Table S7): for example, the $\hat{L}_{jk,D}$ estimator in the $\alpha = 1$ stratified model gave lower estimates of heritability than the $\hat{L}_{jk,r}$ estimator for all 47 traits, but the relative differences were always smaller than 6% (Supplementary Table S3), and always smaller than twice the block-jackknife standard error of the estimator. By contrast, the choice of α can produce large differences.

For the univariate models ($\theta = 0$), we find that estimates of heritability with $\alpha = 0$ are different by more than two standard errors to the $\alpha = 1$ estimate for 29 out of the 47 traits considered (Figure 2, Supplementary Figure 6, Supplementary Tables S5-6). In general, we observe that the univariate models with $\alpha = 0$ produce heritability estimates that are closer to the estimates from the stratified models for 44 out of 47 traits (Figure 2(a), Supplementary Tables S2-4), which is consistent with previous estimates of α values that are closer to zero ($\alpha \approx 0.38$, Schoech et al., 2019; Zeng et al., 2018). This trend is reversed for the intercept, with the $\alpha = 1$ model producing estimates that are closer to the stratified models for most traits. Supplementary Figure 6(b) shows that intermediate values of α , such as $\alpha = 0.25$, produce estimates of confounding and global heritability that are closer to the stratified models than both $\alpha = 0$ and $\alpha = 1$.

In the case of the stratified models ($\theta = 1$), the differences are more subtle. Out of the 47 traits analyzed, none showed a significant difference in the global estimates of SNP heritability across the models with $\alpha = 0$ and $\alpha = 1$ (Supplementary Tables S2-4). At the same time, small but significant differences are observed in the estimates of the intercept for 3 of the traits analyzed (Eosinophil Count, Tanning, and White Blood Cell Count) (Supplementary Tables S2-4). Overall, our analysis confirms that the stratified models (with the MAF decile bins included) successfully counteract the bias induced by an arbitrary choice of α , producing global estimates of SNP heritability and confounding that are largely concordant across different choices of α (Figure 2, Supplementary Figure 7).

3.3 Examining Estimates of Coefficients and Functional Enrichment

The preceding analyses showed that the stratified models generally produce concordant estimates of global parameters such as SNP heritability and confounding, with the choice of α

having only a minor impact. However, as has been noted above, different choices of α can still result in different estimates for quantities associated with partitioned heritability, such as standardized heritability coefficients and functional enrichment. To examine the influence of α on these partitioned heritability metrics, we first focus on the per-standardized annotation coefficients, as defined by Gazal et al., 2017:

$$\tau_c^* = \frac{M \cdot sd(a_c)}{h_g^2} \hat{\tau}_c \quad (5)$$

The quantity τ_c^* is defined as additive change in per-SNP heritability associated to a 1 standard deviation increase in the value of the annotation ($sd(a_c)$), normalized by the average per-SNP heritability over all SNPs for the trait (h_{SNP}^2/M). Despite its usefulness for examining contributions to heritability across traits, the τ^* metric as defined above is not suitable for comparing models with different values of α , since the $\hat{\tau}_c$ in these models are on different scales (see Appendix B). Here, we propose a modified metric that captures the contributions of the different annotations to heritability as well as the influence of α :

$$\tau^*(w) = \frac{M \cdot sd(\text{Var}(X)^{1-\alpha} a_c)}{h_g^2} \hat{\tau}_c \quad (6)$$

The modified metric $\tau^*(w)$ has the same overall interpretation as the metric proposed by Gazal et al., 2017, with the main difference being that we multiply the value of the annotation by the variance in allele frequency $\text{Var}(X)^{1-\alpha}$ when computing its standard deviation.

Figure 3(a) shows the estimates for the standardized coefficients for all the functional annotations that we analyzed, meta-analyzed across the 47 GWAS traits. We find that, in general, models with $\alpha = 0$ and $\alpha = 1$ produce comparable estimates of $\tau^*(w)$, with some notable exceptions (Supplementary Table S8). In particular, the models tend to diverge in the significance they assign to annotations associated with QTLs (e.g. the set of MaxCPP annotations introduced in Hormozdiari et al., 2018, with the $\alpha = 0$ models often failing to reach the Bonferroni significance threshold for those annotations (3(a), Supplementary Table S8). Comparable differences between the models are also observed for some of the LD-related annotations (e.g. CpG Content and Recombination Rate) that were introduced by Gazal et al., 2017 (Figure 3(c),

Supplementary Table S8). Note that some of these differences persist for intermediate values of α , though at a much smaller scale (Supplementary Figures 8-9). Thus the choice of α can influence the predicted biological relevance of different annotations even when frequency bins are used to model frequency-dependent effects.

Figure 3(a) also has a slope higher than one, indicating that the choice $\alpha = 0$ leads to stronger estimates of enrichment for highly enriched categories. We see a similar effect when examining heritability enrichment, a measure of enrichment commonly-used for binary annotations. Finucane et al., 2015 (Figure 3(b), Supplementary Table S9). Heritability enrichment is the ratio of the proportion of heritability explained by SNPs in a given functional category to the proportion of SNPs in that category. As expected, intermediate values of α produce less systematic shift in the estimates of enrichment for both statistics (Supplementary Figure 10).

To test the significance of the enrichment estimates, we compute p -values using the differential enrichment metric as defined by Hujoel et al., 2019 (see Appendix B). Even though $\alpha = 0$ produce higher enrichment results for strongly enriched categories, two functional categories that are deemed highly significant under the $\alpha = 1$ model fail to reach the significance threshold under the $\alpha = 0$ model (Enhancer (Hoffman) and TSS (Hoffman), Supplementary Table S9). For example, Enhancer (Hoffman) has unadjusted p -values of 0.229 under $\alpha = 0$ and 0.0005 under $\alpha = 1$. The opposite effect is observed for the Weak Enhancer (Hoffman) functional category, where the $\alpha = 0$ model reports it to be highly significant (Supplementary Table S9). These results again highlight the important role that the α parameter plays in the stratified LDSC framework, with different values of α potentially leading to different interpretations about the genetic architecture of complex traits.

3.4 The influence of α on models of SNP heritability

To explain the systematic shifts in enrichment as a function of α in the stratified models, here we empirically explore predicted mean squared effect size $\text{Var}(\beta_j)$ and corresponding association statistic χ_j^2 as a function of allele frequency under the LDSC models.

In the standard univariate model ($\theta = 0$, $\alpha = 1$), according to Equation 2, rare variants are predicted to have very large effect relative to common variants (Figure 4(a)). As a result, the

model tends to strongly overestimate the chi-square statistics for rare variants, and underestimate them for common variants (Figure 4(b)). This effect is more pronounced for $\alpha = 1$ than for any other choice of α (Figure 4(b)) and it helps explain the general trend of the standard univariate model producing downwardly biased estimates of SNP heritability (Speed et al., 2017): to avoid having large squared errors for rare variants in the regression, the model must underestimate the effect sizes of common variants, thereby underestimating the total heritability. This bias can be reduced by choosing a more biologically plausible value of α , or by using stratified regression.

We expect the stratified models to better describe the relationship between allele frequency and effect size, since these models have the freedom to assign arbitrary mean predicted effect sizes for each frequency bin. Indeed, all stratified models correctly overcome the bias in predicting the χ^2 statistic (Figure 4(d)). However, they do so while predicting different effect sizes for rare and common variants (Figure 4(c)). Consequently, the weight of rare variants in the heritability of a trait can vary even in stratified models that agree on estimates of overall heritability. This, in turn, implies that the choice of α would lead to different proportions of heritability explained by each frequency bin and thus impact the predicted functional enrichment of variants in those bins (Figure 5(a)).

To explain the systematic differences in enrichment affecting highly enriched annotations, we now focus on the SNPs in the top 10 highly enriched functional categories (HECs, Supplementary Table S11), which comprise roughly 7% of all SNPs. The predicted mean effect size for rare functional variants is again very high for rare variants when $\alpha = 1$ (Figure 4(e)). And, as in the univariate models, this results in correspondingly high weighted mean errors in estimating the χ^2 s for those SNPs (Figure 4(f)). This makes sense given the linear model for $\text{Var}(\beta)$: for variants in strongly enriched categories, the annotation term overwhelms the frequency corrections and the distribution of effect sizes again follows the implicit assumptions induced by our choice of α (Figure 4(e)). Because the frequency bins fail to capture the frequency dependence for variants in highly enriched categories, the LDSC model for these variants is similar to a univariate model in that frequency effects are dictated by the choice of α . As in the univariate case, effects of rare variants are overestimated and, as argued above, the

heritability contributed by these SNPs is underestimated. Figure 4 (b), (d), and (f) also show that in all cases, intermediate α produce less biased predicted chi-squared, as measured by the weighted mean error, than either $\alpha = 0$ and $\alpha = 1$.

Finally, the choice of α also has a large impact on the inferred architecture of complex traits as measured by the proportion of heritability attributed to common and rare variants. Figure 5(a) shows markedly different estimates of the heritability enrichment metric among rare variants. Figure 5(a) also shows that heritability enrichment should not be used as a proxy for functional enrichment: Even though common variants have lower per-allele effect size (e.g., Figure 4c), they are highly enriched for heritability relative to expectations under the (unrealistic) $\alpha = 1$ model (Figure 5(a)). If the goal is to identify classes of variants with high per-allele effect size, we advocate for a different measure of functional enrichment, where contribution to heritability is measured relative to the expectation under a constant per-allele effect size ($\alpha = 0$) (see Appendix B). Figure 5(b) shows the estimated functional enrichment across frequency bins, revealing the expected behaviour of slight depletion for common variants and slight enrichment for rare variants, with substantial differences depending on α . The 0th bin shows the largest difference across choices of α , and also an unexpected functional depletion for the rarest variants and $\alpha = 0$. This 0th bin, which includes variants with allele frequency less than 5%, is often excluded from heritability estimates because of sensitivity to model assumptions and of other technical reasons (Bulik-Sullivan, Loh, et al., 2015).

4 Discussion

In this work, we proposed a reformulation of the LD Score Regression framework that replaces the r^2 measure of LD with the D^2 statistic, leveraging recently-proposed unbiased estimators (Ragsdale and Gravel, 2020). The reformulation highlighted the implicit assumptions about the relationship between a variant’s minor allele frequency and its effect size in commonly-used stratified LDSC models. This relationship, characterized by the parameter α , has been the subject of recent work on the frequency-dependent architecture of complex traits (Schoech et al., 2019; Speed et al., 2017; Speed et al., 2020; Zeng et al., 2018).

In the initial LDSC model, this parameter was set to $\alpha = 1$ for mathematical convenience. Over the years, this assumption has been shown to be biologically implausible and to introduce biases in heritability estimates (Speed et al., 2017). These biases have been primarily addressed in the literature through stratified LDSC, giving the model more parameters and thus more flexibility to overcome the unrealistic assumption. Despite the additional parameters, we have shown that the choice of α still has a substantial effect on partitioned heritability metrics, such as standardized heritability coefficients and functional enrichment.

More recently, Speed et al., 2020 introduced a new heritability model, dubbed BLD-LDAK-alpha, that combines features of the Baseline LD and the LDAK models of SNP heritability and treats alpha as a free parameter that can be estimated from data. This is a promising approach that may alleviate the issues we report here, and the BLD-LDAK-alpha model results are generally consistent with our findings (Supplementary Figure 11). A notable difference between the two approaches is that the models presented here allow for a more flexible dependence of effect sizes on allele frequencies, both by conditioning on MAF bins and fitting the model along a coarse grid of alpha values, while the BLD-LDAK-alpha model allows for more systematic inference of the parametric dependence. The approach taken here allows us to investigate the effect of implicit assumptions on the commonly-used stratified LDSC model and to propose a minimal modification that largely resolves the issue, and to gain further insight into the interpretation of enrichment results.

These biases in partitioned heritability metrics exist because the choice of $\alpha = 1$ implies rigid assumptions about the relationship between allele frequencies and effect sizes for variants in highly enriched categories. Under genome-wide empirical estimates of $\alpha = 0.38$ (Schoech et al., 2019) and no annotation ($\theta = 0$), Equation 2 predicts that the expected squared effect size of an uncommon variant at frequency $p = 0.05$ is 88% higher than a common variant at $p = 0.5$. Under the $\alpha = 1$ model, that ratio is expected to be at 526%. Because the contributions of highly enriched functional annotations dominate the frequency corrections in Equation 2, this effect is still present for highly enriched variants in the model with annotations ($\theta = 1$). These predictions are indeed borne out in the empirical results in Panels (a), (c), and (e) of Figure 4. Importantly, these panels show that conditioning on allele frequency bins is not sufficient to

counteract the effect of setting α to an unrealistic value.

Even though such a synergistic effect is possible in theory, the analyses presented above and recent empirical estimates (Speed et al., 2020) suggest that values in the range $\alpha = 0.2 - 0.5$ are much more plausible across all functional annotations. Even though uncertainty remains about the ‘best’ choice for α , we find much more modest differences in inference results if α is chosen within this range. Therefore, to avoid letting the best be the enemy of the good, we recommend using the genome-wide average $\alpha = 0.38$ as a starting point for most stratified analyses.

Appendix A Derivation of LD Score Regression with the D^2 Statistic

Our re-formulation of the LD Score Regression model starts with a standard polygenic model where the phenotype of N individuals is assumed to linearly depend on M SNPs:

$$Y = X\beta + \varepsilon \quad (7)$$

In this context, we assume that Y is a standardized vector of measured phenotypes and X is a $N \times M$ mean-centered (but not variance-normalized) genotype matrix. Following the derivation in Finucane et al., 2015, we express the Ordinary Least Squares (OLS) solution for the marginal statistic of SNP j as a function of Y :

$$\begin{aligned} \hat{\beta}_j &= \frac{X_j^\top Y}{X_j^\top X_j} = \frac{1}{X_j^\top X_j} X_j^\top (X\beta + \varepsilon) \\ \hat{\beta}_j &= \frac{1}{X_j^\top X_j} \left(\sum_k (X_j^\top X_k) \beta_k + X_j^\top \varepsilon \right) \\ \hat{\beta}_j &= \frac{1}{X_j^\top X_j} \left(2N \sum_k D_{jk} \beta_k + X_j^\top \varepsilon \right) \end{aligned} \quad (8)$$

Where the third line follows from the definition of the D statistic ($D_{jk} := \frac{1}{2N} X_j^\top X_k$). If we define the χ^2 association statistic for SNP j as $\chi_j^2 := N \left(X_j^\top X_j \right) \hat{\beta}_j^2$ and take the expectation over all the random components, we obtain:

$$\mathbb{E}[\chi_j^2] = \frac{4N^2}{X_j^\top X_j} \sum_k D_{jk}^2 \mathbb{E}[\beta_k^2] + \sigma_\varepsilon^2 \quad (9)$$

Under this model, the total narrow-sense SNP heritability is defined as $h_{SNP}^2 = \sum_k \text{Var}(X_k) \text{Var}(\beta_k)$. Given this, coupled with the assumption that $h_g^2 + \sigma_\varepsilon^2 = 1$, we obtain our general model for LD Score Regression with the D^2 statistic:

$$\mathbb{E}[\chi_j^2] = \frac{4N}{\text{Var}(X_j)} \sum_k \text{Var}(\beta_k) \left(D_{jk}^2 - \frac{\text{Var}(X_k) \text{Var}(X_j)}{4N} \right) + 1 \quad (10)$$

Assuming a large sample size for GWAS, this can be approximated by:

$$\mathbb{E}[\chi_j^2] \approx \frac{4N}{\text{Var}(X_j)} \sum_k \text{Var}(\beta_k) D_{jk}^2 + 1. \quad (11)$$

Finally, using the SNP heritability models outlined in Equation (2), the expression above can be equivalently formulated in terms of LD scores:

$$\begin{aligned} \mathbb{E}[\chi_j^2] &\approx N \left[\tau_0 \ell(j, 0) + \theta \sum_{c=1}^{96} \tau_c \ell(j, c) \right] + 1 \\ \ell(j, c) &:= \frac{4}{\text{Var}(X_j)} \sum_k D_{jk}^2 \text{Var}(X_k)^{-\alpha} a_c(k). \end{aligned} \quad (12)$$

Appendix B Definitions of SNP Heritability and Functional Enrichment

B.1 SNP Heritability

SNP heritability quantifies the amount of phenotypic variance explained by a set of SNPs, genotyped or imputed (Yang, Zeng, Goddard, Wray, and Visscher, 2017). If we assume that the phenotype is standardized and the genotype matrix is mean-centered (but not variance-normalized), we can write our model of SNP heritability as:

$$h_{SNP}^2 = \sum_k \text{Var}(X_k) \text{Var}(\beta_k) \quad (13)$$

Using the above formulation, the models of SNP heritability outlined in Equation (2) can be expressed as:

$$h_{SNP}^2 = \sum_k \text{Var}(X_k)^{1-\alpha} \left(\tau_0 + \theta \sum_{c=1}^{96} \tau_c a_c(k) \right) \quad (14)$$

From the above expression, we can see that the relationship between partitioned heritability coefficients τ_c and the total SNP heritability h_{SNP}^2 depends on α . For instance, in the univariate case ($\theta = 0$), the coefficient τ_0 is proportional to per-SNP heritability and is related to the total

SNP heritability as

$$\tau_0 = \frac{h_{SNP}^2}{\sum_k \text{Var}(X_k)^{1-\alpha}} \quad (15)$$

When $\alpha = 1$, we recover the original LDSC formulation where the per-SNP heritability was assumed to be proportional to $1/M$ (Bulik-Sullivan, Loh, et al., 2015). On the other hand, when $\alpha = 0$ the per-SNP heritability is inversely proportional to the average allele frequency variance across all SNPs.

B.2 Heritability and Functional Enrichment

Heritability enrichment is defined as the ratio of the proportion of heritability explained by SNPs in a given category c relative to the proportion of SNPs in that category (Finucane et al., 2015):

$$\text{Heritability Enrichment}(c) = \frac{h_{SNP}^2(c)/h_{SNP}^2}{M(c)/M} \quad (16)$$

Where $h_{SNP}^2(c)$ is the heritability explained by SNPs in category c and $M(c)$ is the number of SNPs in that category. To test for significance of the functional enrichment metric, we use the differential enrichment metric as defined in Hujoel et al., 2019:

$$\text{Differential Enrichment}(c) = \frac{h_{SNP}^2(c)}{M(c)} - \frac{h_{SNP}^2 - h_{SNP}^2(c)}{M - M(c)} \quad (17)$$

As outlined in the main text (Section 2.5), a better measure of functional enrichment would be measuring contributions to heritability relative to expectation under a constant per-allele effect size ($\alpha = 0$) model. This implies that instead of dividing by the proportion of SNPs, we divide by the proportion of genotypic variances:

$$\text{Functional Enrichment}(c) = \frac{h_{SNP}^2(c)/h_{SNP}^2}{\sum_{j \in c} \text{Var}(X_j) / \sum_j \text{Var}(X_j)} \quad (18)$$

Appendix C Analysis Procedures and Evaluation Metrics

C.1 Computing and Comparing LD Scores from 1000 Genomes Data

To conduct the analyses discussed in this paper, we computed LD scores using three different estimators of LD (naive r^2 , corrected r^2 (Bulik-Sullivan, Loh, et al., 2015; Yin and Fan, 2001) and D^2 (Ragsdale and Gravel, 2020)) for $\alpha \in \{0.0, 0.25, 0.5, 0.75, 1.0\}$ and for 3 of the super populations in Phase III of the 1000 Genomes Project (Africans, Asians, and Europeans) (Altshuler et al., 2010; The 1000 Genomes Project Consortium, 2012), excluding the sets of closely related individuals identified in Gazal, Sahbatou, Babron, Genin, and Leutenegger, 2015. We followed the same quality control procedures as in the original LDSC software documentation (Bulik-Sullivan, Loh, et al., 2015) (e.g. discarding singleton SNPs). All of these quality control steps were run with PLINK v1.9 (Purcell et al., 2007).

To qualitatively compare LD Scores from the univariate model with different values of α , we normalized them such that the slope of the univariate regression becomes $\frac{N}{M}h_{SNP}^2$ instead of $N\tau_0$. Given our definition of τ_0 in terms of h_{SNP}^2 in Equation (15), we see that this can be achieved by dividing the LD score for each model by the following quantity:

$$\frac{1}{M} \sum_k \text{Var}(X_k)^{1-\alpha}$$

By construction, when $\alpha = 1$, as in the standard LDSC models, the normalization factor is 1. When $\alpha = 0$, the normalization factor becomes the average genome-wide variance in allele frequency.

C.2 Parameter Estimation and Meta-analysis

To estimate the parameters of the various heritability models, we used the Iteratively re-weighted least squares (IRLS) jackknife estimator from Bulik-Sullivan, Loh, et al., 2015 with some modifications to account for differences in the heritability models. Primarily, since our general model starts with an unnormalized genotype matrix, the total SNP heritability is given by:

$$h_{SNP}^2 = \tau_0 \sum_k \text{Var}(X_k)^{1-\alpha} + \theta \sum_{c=1}^{96} \tau_c \sum_k \text{Var}(X_k)^{1-\alpha} a_c(k)$$

Once the coefficients $\hat{\tau}_c$ have been estimated by the `ldsc` software, to obtain the total heritability, we have to multiply them by different factors. For example, in the case of the $\alpha = 1$ models, we multiply by the sum of the annotations $\sum_j a_c(k)$, whereas for the $\alpha = 0$ models, we multiply by the sum of the annotations weighted by MAF: $\sum_k \text{Var}(X_k) a_c(k)$.

To meta-analyze the coefficients as well as measures of functional enrichment across the 47 GWAS traits, we used a random effects model as implemented in the R package `meta` (Balduzzi, Rücker, and Schwarzer, 2019).

C.3 Empirical Evaluation of SNP heritability models

To empirically evaluate the stratified SNP heritability models as in Section 2.5, we fit each model to the summary statistics from 47 GWAS traits (Supplementary Table S1) using the `ldsc` software (Bulik-Sullivan, Loh, et al., 2015) with standard configurations. For each trait and model, we obtained the mean parameter estimates $\hat{\tau}_c$ and used them to compute the predicted $\mathbb{E}[\hat{\chi}_j^2]$ for all SNPs that were used in the regression. Then, for each MAF bin, we computed the weighted mean error using the following equation:

$$\text{Weighted Mean Error} = \frac{1}{\sum_{j \in \text{MAFbin}} w_j} \sum_{j \in \text{MAFbin}} w_j (\mathbb{E}[\hat{\chi}_j^2] - \chi_j^2) \quad (19)$$

Where the $\mathbb{E}[\hat{\chi}_j^2]$ is the predicted association statistic under the model and χ_j^2 is the observed statistic. The weight w_j is simply the LD score weight that we employed in the regression Bulik-Sullivan, Loh, et al., 2015. This strategy of weighing the performance statistics by the LD score weights has been explored in previous work Speed et al., 2020. Since, in our case, each model has its own set of LD score weights that match its α value, here we used the LD score weights for the $\alpha = 1$ model throughout. The weighted mean errors per MAF bin that we report in the main text are averaged across the 47 GWAS traits.

Supplemental Data

Supplemental Data include ten figures (Supplementary Figures S1-10) and eleven tables (Supplementary Tables S1-11).

Declaration of Interests

The authors declare no conflict of interest.

Acknowledgements

We thank Chris Gignoux and members of the Gravel Lab for useful discussions.

Web Resources

Baseline-LD model version 2.2,

data.broadinstitute.org/alkesgroup/LDSCORE/

Summary statistics for 47 GWAS traits,

data.broadinstitute.org/alkesgroup/LDSCORE/independent_sumstats/

Data and Code Availability

Code to compute the LD scores with the unbiased estimator of D^2 and carry out the analyses discussed in this paper is available on github:

<https://github.com/shz9/unbiased-ldsc>.

References

Altshuler, D. L., Durbin, R. M., Abecasis, G. R., Bentley, D. R., Chakravarti, A., Clark, A. G., ... Peterson, J. L. (2010). A map of human genome variation from population-scale sequencing. *Nature*. doi:10.1038/nature09534

- Balduzzi, S., Rücker, G., & Schwarzer, G. (2019). How to perform a meta-analysis with R: A practical tutorial. *Evidence-Based Mental Health*, (22), 153–160.
- Bulik-Sullivan, B., Finucane, H. K., Anttila, V., Gusev, A., Day, F. R., Loh, P. R., ... Neale, B. M. (2015). An atlas of genetic correlations across human diseases and traits. *Nature Genetics*. doi:10.1038/ng.3406
- Bulik-Sullivan, B., Loh, P. R., Finucane, H. K., Ripke, S., Yang, J., Patterson, N., ... O'Donovan, M. C. (2015). LD score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nature Genetics*. doi:10.1038/ng.3211
- Finucane, H. K., Bulik-Sullivan, B., Gusev, A., Trynka, G., Reshef, Y., Loh, P. R., ... Price, A. L. (2015). Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nature Genetics*. doi:10.1038/ng.3404
- Gazal, S., Finucane, H. K., Furlotte, N. A., Loh, P. R., Palamara, P. F., Liu, X., ... Price, A. L. (2017). Linkage disequilibrium-dependent architecture of human complex traits shows action of negative selection. *Nature Genetics*. doi:10.1038/ng.3954
- Gazal, S., Loh, P. R., Finucane, H. K., Ganna, A., Schoech, A., Sunyaev, S., & Price, A. L. (2018). Functional architecture of low-frequency variants highlights strength of negative selection across coding and non-coding annotations. *Nature Genetics*. doi:10.1038/s41588-018-0231-8
- Gazal, S., Marquez-Luna, C., Finucane, H. K., & Price, A. L. (2019). Reconciling S-LDSC and LDK functional enrichment estimates. doi:10.1038/s41588-019-0464-1
- Gazal, S., Sahbatou, M., Babron, M. C., Genin, E., & Leutenegger, A. L. (2015). High level of inbreeding in final phase of 1000 Genomes Project. *Scientific Reports*. doi:10.1038/srep17453
- Hormozdiari, F., Gazal, S., Van De Geijn, B., Finucane, H. K., Ju, C. J., Loh, P. R., ... Price, A. L. (2018). Leveraging molecular quantitative trait loci to understand the genetic architecture of diseases and complex traits. *Nature Genetics*. doi:10.1038/s41588-018-0148-2
- Hujoel, M. L., Gazal, S., Hormozdiari, F., van de Geijn, B., & Price, A. L. (2019). Disease Heritability Enrichment of Regulatory Elements Is Concentrated in Elements with Ancient

- Sequence Age and Conserved Function across Species. *American Journal of Human Genetics*. doi:10.1016/j.ajhg.2019.02.008
- Pasaniuc, B., & Price, A. L. (2017). Dissecting the genetics of complex traits using summary association statistics. doi:10.1038/nrg.2016.142
- Purcell, S., Neale, B., Todd-Brown, K., Thomas, L., Ferreira, M. A., Bender, D., . . . Sham, P. C. (2007). PLINK: A tool set for whole-genome association and population-based linkage analyses. *American Journal of Human Genetics*. doi:10.1086/519795
- Ragsdale, A. P., & Gravel, S. (2020). Unbiased Estimation of Linkage Disequilibrium from Unphased Data. *Molecular Biology and Evolution*. doi:10.1093/molbev/msz265
- Schoech, A. P., Jordan, D. M., Loh, P. R., Gazal, S., O'Connor, L. J., Balick, D. J., . . . Price, A. L. (2019). Quantification of frequency-dependent genetic architectures in 25 UK Biobank traits reveals action of negative selection. *Nature Communications*. doi:10.1038/s41467-019-08424-6
- Shi, H., Gazal, S., Kanai, M., Koch, E. M., Schoech, A. P., Kim, S. S., . . . Price, A. L. (2019). Population-specific causal disease effect sizes in functionally important regions impacted by selection. *bioRxiv*. doi:10.1101/803452
- Speed, D., & Balding, D. J. (2019). SumHer better estimates the SNP heritability of complex traits from summary statistics. *Nature Genetics*. doi:10.1038/s41588-018-0279-5
- Speed, D., Cai, N., Johnson, M. R., Nejentsev, S., & Balding, D. J. (2017). Reevaluation of SNP heritability in complex human traits. *Nature Genetics*. doi:10.1038/ng.3865
- Speed, D., Hemani, G., Johnson, M. R., & Balding, D. J. (2012). Improved heritability estimation from genome-wide SNPs. *American Journal of Human Genetics*. doi:10.1016/j.ajhg.2012.10.010
- Speed, D., Holmes, J., & Balding, D. J. (2020). Evaluating and improving heritability models using summary statistics. *Nature Genetics*. doi:10.1038/s41588-020-0600-y
- The 1000 Genomes Project Consortium. (2012). An integrated map of genetic variation. *Nature*.

- Yang, J., Benyamin, B., McEvoy, B. P., Gordon, S., Henders, A. K., Nyholt, D. R., ... Visscher, P. M. (2010). Common SNPs explain a large proportion of the heritability for human height. *Nature Genetics*. doi:10.1038/ng.608
- Yang, J., Zeng, J., Goddard, M. E., Wray, N. R., & Visscher, P. M. (2017). Concepts, estimation and interpretation of SNP-based heritability. doi:10.1038/ng.3941
- Yin, P., & Fan, X. (2001). Estimating r^2 shrinkage in multiple regression: A comparison of different analytical methods. *Journal of Experimental Education*. doi:10.1080/00220970109600656
- Zeng, J., De Vlaming, R., Wu, Y., Robinson, M. R., Lloyd-Jones, L. R., Yengo, L., ... Yang, J. (2018). Signatures of negative selection in the genetic architecture of human complex traits. *Nature Genetics*. doi:10.1038/s41588-018-0101-4

Figures

- **Figure 1. Genome-wide comparison of LD Scores in the European samples ($N = 489$) in the 1000 Genomes Project.**

Genome-wide comparison of LD Scores in the European samples ($N = 489$) in the 1000 Genomes Project. (a) shows the distribution of LD Scores obtained from the r^2 (x-axis) and D^2 (y-axis) estimators. In (b) we show the normalized LD Scores from the D^2 estimator but computed according to different models of SNP heritability ($\alpha = 1$ on the x-axis and the $\alpha = 0$ on the y-axis). Points are colored by the focal SNP's minor allele frequency (MAF).

- **Figure 2. Estimates of SNP heritability and confounding for 15 GWAS traits using 4 different models of heritability.**

Estimates of SNP heritability and confounding for 15 GWAS traits using 4 different models of heritability. Dashed bars show estimates for univariate models. Solid bars show estimates for stratified models. Error bars correspond to jackknife standard errors. Color code: turquoise ($\alpha = 1$), red ($\alpha = 0$).

- **Figure 3. Comparison of meta-analyzed standardized coefficients and functional enrichment across different models of SNP heritability**

Comparison of meta-analyzed standardized coefficients and functional enrichment across different models of SNP heritability ($\alpha = 1$ on the x-axis and $\alpha = 0$ on the y-axis). The estimates are meta-analyzed over 47 GWAS traits using a random-effects model. Estimates are shown with standard errors. **(a)** shows standardized coefficients $\tau^*(w)$ and **(b)** shows estimated heritability enrichment.

- **Figure 4. The predicted mean squared effect size and the mean error between predicted and observed χ^2 for SNPs in different MAF bins with different values of α**

The predicted mean squared effect size, computed according to Equation 2, and the mean error between predicted and observed χ^2 for SNPs in different MAF bins with different values of α . Estimates are averaged across 47 GWAS traits. **(a, c)** show predicted mean squared effect size for SNPs under **(a)** univariate models and **(c)** stratified models. **(b, d)** show mean error between predicted and observed χ^2 for SNPs under **(b)** univariate and **(d)** stratified models. Panels **(e, f)** show predicted mean squared effect size and mean error for SNPs in the top 10 Highly Enriched Categories (HECs) under the stratified models.

- **Figure 5. Estimates of enrichment for the MAF bins for stratified LDSC model with five different values of α**

Estimates of enrichment for the MAF bins, computed according to Equation 14, for stratified LDSC model with five different values of α . Estimates are averaged across 47 GWAS traits. **(a)** shows the predicted heritability enrichment metric for 11 MAF bins and **(b)** shows the functional enrichment metric for the same categories.

Figure 1: Genome-wide comparison of LD Scores in the European samples ($N = 489$) in the 1000 Genomes Project. (a) shows the distribution of LD Scores obtained from the r^2 (x-axis) and D^2 (y-axis) estimators. In (b) we show the normalized LD Scores from the D^2 estimator but computed according to different models of SNP heritability ($\alpha = 1$ on the x-axis and the $\alpha = 0$ on the y-axis). Points are colored by the focal SNP's minor allele frequency (MAF).

(a) Estimated SNP-heritability h_g^2

(b) Estimated confounding (Intercept - 1)

Figure 2: Estimates of SNP heritability and confounding for 15 GWAS traits using 4 different models of heritability. Dashed bars show estimates for univariate models. Solid bars show estimates for stratified models. Error bars correspond to jackknife standard errors. Color code: turquoise ($\alpha = 1$), red ($\alpha = 0$).

Figure 3: Comparison of meta-analyzed standardized coefficients and functional enrichment across different models of SNP heritability ($\alpha = 1$ on the x-axis and $\alpha = 0$ on the y-axis). The estimates are meta-analyzed over 47 GWAS traits using a random-effects model. Estimates are shown with standard errors. (a) shows standardized coefficients $\tau^*(w)$ and (b) shows estimated heritability enrichment.

(a) Predicted $\text{Var}(\beta)$ ($\theta = 0$)

(b) Weighted Mean Error ($\theta = 0$)

(c) Predicted $\text{Var}(\beta)$ ($\theta = 1$)

(d) Weighted Mean Error ($\theta = 1$)

(e) Predicted $\text{Var}(\beta)$
($\theta = 1$, HEC SNPs)

(f) Weighted Mean Error
($\theta = 1$, HEC SNPs)

Figure 4: The predicted mean squared effect size and the mean error between predicted and observed χ^2 for SNPs in different MAF bins with different values of α . Estimates are averaged across 47 GWAS traits. (a, c) show predicted mean squared effect size for SNPs under (a) univariate models and (c) stratified models. (b, d) show mean error between predicted and observed χ^2 for SNPs under (b) univariate and (d) stratified models. Panels (e, f) show predicted mean squared effect size and mean error for SNPs in the top 10 Highly Enriched Categories (HECs) under the stratified models.

(a) Heritability enrichment

(b) Functional enrichment

Figure 5: Estimates of enrichment for the MAF bins for stratified LDSC model with five different values of α . Estimates are averaged across 47 GWAS traits. (a) shows the predicted heritability enrichment metric for 11 MAF bins and (b) shows the functional enrichment metric for the same categories.