

Linkage disequilibrium–dependent architecture of human complex traits shows action of negative selection

Steven Gazal^{1,2} , Hilary K Finucane^{1–3}, Nicholas A Furlotte⁴, Po-Ru Loh^{1,2}, Pier Francesco Palamara^{1,2}, Xuanyao Liu^{1,2}, Armin Schoech^{1,2,5}, Brendan Bulik-Sullivan^{2,6}, Benjamin M Neale^{2,6,7} , Alexander Gusev^{1,2} & Alkes L Price^{1,2,8} 

Recent work has hinted at the linkage disequilibrium (LD)-dependent architecture of human complex traits, where SNPs with low levels of LD (LLD) have larger per-SNP heritability. Here we analyzed summary statistics from 56 complex traits (average $N = 101,401$) by extending stratified LD score regression to continuous annotations. We determined that SNPs with low LLD have significantly larger per-SNP heritability and that roughly half of this effect can be explained by functional annotations negatively correlated with LLD, such as DNase I hypersensitivity sites (DHSs). The remaining signal is largely driven by our finding that more recent common variants tend to have lower LLD and to explain more heritability ($P = 2.38 \times 10^{-104}$); the youngest 20% of common SNPs explain 3.9 times more heritability than the oldest 20%, consistent with the action of negative selection. We also inferred jointly significant effects of other LD-related annotations and confirmed via forward simulations that they jointly predict deleterious effects.

Estimating the heritability explained by SNPs^{1,2}, and its distribution across chromosomes^{3,4}, allele frequencies⁵, and functional regions^{6–10}, has yielded rich insights into the polygenic architecture of human complex traits. Recent work has hinted at LD-dependent architectures, defined as a dependence of causal effect sizes on LLD after conditioning on minor allele frequency (MAF), for several complex traits. LD-dependent architectures bias SNP heritability estimates¹¹ and downward biases have been observed for several traits^{11–13}, suggesting larger causal effect sizes for genetic variants with low LLD.

Indeed, heritability is enriched in functional annotations such as DHSs⁷, histone marks^{8,10}, and regions with high GC content⁹, which all have low LLD^{7,14,15}. On the other hand, regions of low recombination rate, which have high LLD, are enriched for exonic deleterious and disease-associated variants¹⁶, suggesting an LD-dependent architecture of opposite effect.

Despite these observations, LD-dependent architectures have not been formally assessed, quantified, or biologically interpreted. Understanding which biological processes shaping the LD patterns of the genome are most directly linked to complex traits is challenging, as many of the corresponding annotations are correlated with each other. To investigate LD-dependent architectures, we extended stratified LD score regression⁸, a method that partitions the heritability of a set of binary genomic annotations using genome-wide association study (GWAS) summary statistics, to continuous-valued annotations; our method produces robust results in simulations. We applied our method to a broad set of LD-related annotations, including LLD, predicted allele age, and recombination rate, to analyze summary statistics from 56 complex traits and diseases (average $N = 101,401$), including 18 traits from the 23andMe research database and 15 traits from the UK Biobank. We inferred jointly significant effects for several LD-related annotations on per-SNP heritability, including predicted allele age: common variants that are more recent tend to have lower LLD and to explain more heritability, which is consistent with the action of negative selection, as selection has had less time to eliminate recent weakly deleterious variants. We confirmed via forward simulations that allele age and other LD-related annotations associated with per-SNP heritability jointly predict the deleterious effects of a variant. Our results implicate the action of negative selection on deleterious variants that affect complex traits.

RESULTS

Overview of methods

Stratified LD score regression⁸ is a method for partitioning heritability across overlapping binary annotations using GWAS summary statistics. The idea of this method is that, for a polygenic trait, LD with an annotation that is enriched for heritability will increase the χ^2 statistic of a SNP more than LD with an annotation that is not enriched for heritability. We extended stratified LD score regression to quantify the effects on heritability of continuous-valued (and/or binary) annotations. Here the idea is that, if a continuous annotation

¹Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, Massachusetts, USA. ²Program in Medical and Population Genetics, Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA. ³Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts, USA. ⁴23andMe, Inc., Mountain View, California, USA. ⁵Department of Systems Biology, Harvard Medical School, Boston, Massachusetts, USA. ⁶Stanley Center for Psychiatric Research, Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA. ⁷Analytic and Translational Genetics Unit, Department of Medicine, Massachusetts General Hospital, Boston, Massachusetts, USA. ⁸Department of Biostatistics, Harvard T.H. Chan School of Public Health, Boston, Massachusetts, USA. Correspondence should be addressed to S.G. (sgazal@hsph.harvard.edu) or A.L.P. (aprice@hsph.harvard.edu).

Received 19 October 2016; accepted 16 August 2017; published online 11 September 2017; doi:10.1038/ng.3954

a is associated with increased heritability, LD to SNPs with large values of a will increase the χ^2 statistic of a SNP more than LD to SNPs with small values of a .

More precisely, the expected χ^2 statistic of SNP j can be written as

$$E[\chi_j^2] = N \sum_c \tau_c l(j, c) + Nb + 1 \quad (1)$$

where

$$l(j, c) = \sum_k a_c(k) r_{jk}^2$$

is the LD score of SNP j with respect to continuous values $a_c(k)$ of annotation a_c , r_{jk} is the correlation between SNPs j and k in a reference panel (for example, Europeans from the 1000 Genomes Project¹⁷), N is the sample size of the GWAS, τ_c is the effect size of annotation a_c on per-SNP heritability (conditioned on all other annotations), and b is a term that measures the contribution of confounding biases¹⁸. We standardize estimated effect sizes $\hat{\tau}$ to report per-standardized-annotation effect sizes $\hat{\tau}^*$, defined as the proportionate change in per-SNP heritability associated with a 1 s.d. increase in the value of the annotation; we note that $\hat{\tau}^*$ can be compared across annotations and across traits. Analogous to ref. 8, standard errors on estimates of $\hat{\tau}^*$ are computed using a block jackknife (Online Methods). We have released open-source software implementing the method (see URLs).

We applied our extension of stratified LD score regression to LLD annotations, adjusted for MAF via MAF-stratified quantile-normalized LD score, as well as other LD-related annotations, including predicted allele age and recombination rate; we included ten MAF bins as additional annotations in all analyses to model MAF-dependent architectures. We also considered functional annotations from a 'baseline model' (refs. 8,19) including 28 main annotations such as coding, conserved, DHSs, and histone marks (59 total annotations; Online Methods).

Although stratified LD score regression has previously been shown to produce robust results using binary annotations⁸, we performed additional simulations to confirm that our extension of stratified LD score regression produces robust results using continuous-valued LD-related annotations and specifically that analyzing LD-related annotations using an LD-based method is appropriate (Online Methods).

SNPs with low LLD have larger per-SNP heritability

We applied our extension of stratified LD score regression to GWAS summary statistics from 56 complex traits and diseases, including 18 traits from 23andMe and 15 traits from UK BioBank (average $N = 101,401$); for 5 traits we analyzed multiple data sets, leading to a total of 62 data sets analyzed (Supplementary Table 1). The standardized effect sizes $\hat{\tau}^*$ for the LLD annotation were consistently negative in all 62 data sets analyzed (Fig. 1 and Supplementary Table 2). In a meta-analysis across 31 independent traits, excluding genetically correlated traits²⁰ in overlapping samples (Supplementary Table 3; average $N = 84,686$; Online Methods), the LLD annotation was highly statistically significant ($\hat{\tau}^* = -0.30$, standard error (s.e.) = 0.02; $P = 2.42 \times 10^{-80}$), confirming that SNPs with a low MAF-adjusted level of LD have larger per-SNP heritability. We also investigated two alternative MAF-adjusted measures of the level of LD, using a sliding-window approach to quantify the level of LD in a genomic region (LLD-REG)¹³ and using the D' coefficient instead of the squared correlation to compute LD scores (LLD- D'); we observed smaller but still significant effects for LLD-REG ($\hat{\tau}^* = -0.22$, s.e. = 0.02; $P = 2.86 \times 10^{-44}$) and LLD- D' ($\hat{\tau}^* = -0.15$, s.e. = 0.02; $P = 2.22 \times 10^{-12}$).




Figure 1 Effect sizes of MAF-adjusted LLD on 20 highly heritable complex traits. Results are displayed for the 20 traits with the highest SNP heritability (subject to low genetic correlation²⁰ between traits). Numerical results for all 56 complex traits are reported in Supplementary Table 2. Error bars represent jackknife 95% confidence intervals. BMI, body mass index; FEV₁, forced expiratory volume in 1 s; FVC, forced vital capacity.

Correlations between LLD and other LD-related annotations

We investigated other LD-related annotations, including MAF-adjusted allele age as predicted using ARGweaver²¹, MAF-adjusted LLD measured in African populations (LLD-AFR), recombination rate^{22,23}, nucleotide diversity¹⁵, a background selection statistic (McVicker B statistic)²⁴, GC content¹⁵, CpG dinucleotide content, replication timing²⁵, and centromeres and telomeres¹⁵. We used the Oxford recombination map²³ and a window size of ± 10 kb for recombination rate and a window sizes of ± 10 kb for nucleotide diversity, ± 1 Mb for GC content, ± 50 kb for CpG content, ± 5 Mb for centromeres and the first/last 10 Mb for telomeres, as these choices produced the most significant signals, although other choices produced similar results (see below and Online Methods). We also considered the 28 main functional annotations from our baseline model. Many of these annotations are highly correlated with LLD and with each other (Fig. 2 and Supplementary Table 4); these correlation patterns inform the interpretation of our heritability results below. In particular, nearly all of the functional annotations from the baseline model are negatively correlated with LLD, with the strongest negative correlations ($-0.20 < r < -0.10$) for histone marks (H3K27ac, H3K4me1, and H3K9ac), conserved regions (GERP NS), and super-enhancers; only repressed regions ($r = 0.05$; depleted for heritability⁸) and transcribed regions ($r = 0.02$) exhibited positive correlations.

One surprising observation was that predicted allele age was positively correlated with LLD ($r = 0.22$; more recent SNPs had lower LLD), whereas a negative correlation might be expected because the LD between two SNPs decays with time. To confirm this observation, we performed coalescent simulations²⁶ using a realistic demographic model for African and European populations²⁷ (Supplementary Note). We observed that, while the LLD of a SNP defined using a fixed set of older SNPs decreased with allele age, older SNPs acquired additional LD with more recent SNPs; the latter effect leads to a positive correlation between predicted allele age and LLD (Supplementary Figs. 1 and 2). We also observed, in both real data and simulations, that allele age was more strongly correlated with LLD-AFR than LLD, as demographic events (for example, bottlenecks) that occurred in European populations distort the relationship between LLD and allele age.

Multiple LD-related annotations impact complex trait architectures

We applied our extension of stratified LD score regression to each of the 13 LD-related annotations defined above, analyzing each




Figure 2 Correlations between LD-related and functional annotations. We report correlations computed on common SNPs (MAF $\geq 5\%$). LLD, LLD-D', LLD-REG, predicted allele age, and LLD-AFR annotations are adjusted for MAF. Numerical results are reported in **Supplementary Table 4**. TFBS, transcription factor binding site.

annotation in turn. We performed meta-analysis of the results across 31 independent traits (**Fig. 3a** and **Supplementary Tables 2** and **3**). All annotations except telomeres were highly significant after correction for multiple testing (**Supplementary Table 3**), and 8 of the remaining 12 annotations remained significant when fitted jointly (**Supplementary Tables 5** and **6**). The predicted allele age ($\tau^* = -0.78$, s.e. = 0.03; $P = 6.27 \times 10^{-175}$) and nucleotide diversity ($\tau^* = -0.78$, s.e. = 0.04; $P = 1.79 \times 10^{-79}$) annotations produced the largest absolute standardized effect size. Interestingly, SNPs in regions with high recombination rates (corresponding to low LLD; $r = -0.49$) had smaller per-SNP heritability ($\tau^* = -0.54$, s.e. = 0.06; $P = 2.39 \times 10^{-18}$), which was inconsistent with the direction of the LLD effect but consistent with the fact that negative selection is more effective in regions with high recombination rates as a consequence of the Hill–Robertson effect²⁸. Thus, per-SNP heritability was most enriched in SNPs with low LLD in regions with low recombination rates, and the opposing effects of these two annotations were stronger when they were conditioned on each other (**Supplementary Fig. 3**). Opposing effects were also observed for the background selection statistic annotation, which was positively correlated with LLD ($r = 0.35$) but had the opposite direction of effect ($\tau^* = 0.51$, s.e. = 0.05; $P = 5.06 \times 10^{-26}$).

To assess how much of the LLD effect is explained by known functional annotations (and because results of stratified LD score regression may be biased in the presence of unmodeled functional annotations⁸), we analyzed each of the 13 LD-related annotations while conditioning on the 59 functional annotations of the baseline model (**Fig. 3b** and **Supplementary Tables 3** and **7**). The effect size of the LLD annotation remained highly significant but was smaller in magnitude ($\tau^* = -0.11$, s.e. = 0.02; $P = 2.57 \times 10^{-11}$), primarily owing to its correlation with DHSs (**Supplementary Fig. 4**). Thus, more than half of the initial LLD signal is explained by known functional annotations. The LLD-REG annotation¹³ was no longer significant in this analysis ($P = 0.19$), indicating that the regional LLD signal is entirely explained by known functional annotations. Predicted allele age produced the largest absolute standardized effect size and the most significant signal ($\tau^* = -0.46$, s.e. = 0.02; $P = 2.38 \times 10^{-104}$); the sign of this effect was consistent across 55 of the 56 traits (positive but not significantly different from zero for glycated hemoglobin concentration (Hb1AC); **Supplementary Table 7**). This indicates that more recent alleles have

larger per-SNP heritability after conditioning on both MAF and known functional annotations. Many other LD-related annotations remained significant (after correction for multiple testing) in the conditional analysis, although LLD-D', replication timing, and centromeres were no longer significant (**Supplementary Fig. 4** and **Supplementary Table 3**).

Finally, we built a model consisting of the 59 functional annotations from the baseline model and the 6 LD-related annotations that remained significant (after correction for multiple testing) when conditioned on each other as well as the baseline model (**Fig. 3c** and **Supplementary Tables 8** and **9**); we call this model the baseline-LD model (Online Methods). We determined that this model produced similar results when using different window sizes for windows-based annotations (for example, recombination rate, nucleotide diversity, and CpG content) or different data sources for recombination rate (**Supplementary Fig. 5**), when performing derived allele frequency (DAF) adjustment instead of MAF adjustment, when using the UK10K Project²⁹ (instead of the 1000 Genomes Project) as the reference panel (**Supplementary Fig. 6**), and across different data sets for the same trait (**Supplementary Fig. 7**). Predicted allele age remained the annotation with the largest absolute standardized effect size ($\tau^* = -0.24$, s.e. = 0.02; $P = 1.08 \times 10^{-23}$), but its effect size decreased owing to its high correlation with the LLD-AFR annotation (**Supplementary Fig. 8**). The effect sizes of LLD-AFR and CpG content increased, owing to opposing effects with the recombination rate and background selection statistic annotations. Effect sizes of the recombination rate, nucleotide diversity, and background selection statistic annotations decreased because they compete with each other, and LLD and GC content were no longer significant (after correction for multiple testing) because of their high correlation with LLD-AFR and CpG content, respectively (**Supplementary Table 10**). Psychiatric diseases and autoimmune diseases exhibited significantly stronger effects for the predicted allele age and background selection statistic annotations, respectively (**Supplementary Table 11** and **Supplementary Note**), possibly owing to the role of selection at different time scales in shaping the genetic architecture of these diseases^{30,31}.

To provide a more intuitive interpretation of the magnitude of the LD-related annotation effects, we computed the proportion of heritability explained by each quintile of each annotation in the baseline-LD model and by each quintile of MAF for comparison

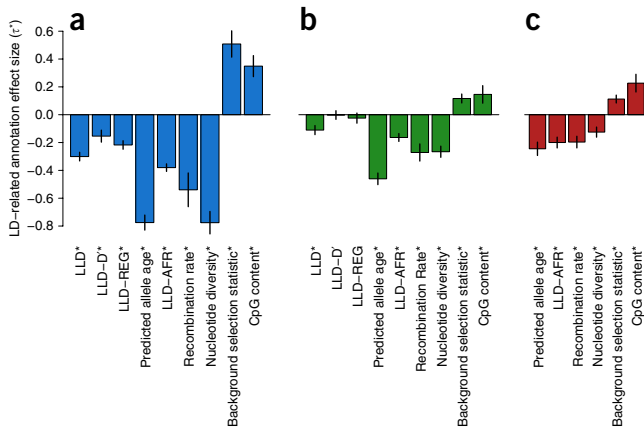


Figure 3 Effect sizes of LD-related annotations subjected to meta-analysis over 31 independent traits. **(a)** Meta-analysis results for nine LD-related annotations. **(b)** Meta-analysis results for nine LD-related annotations, conditioned on the baseline model. **(c)** Meta-analysis results for six LD-related annotations, conditioned on each other and on the baseline model. Results are displayed for the six LD-related annotations that are jointly significant when conditioned on each other and on the baseline model (in **c**). In **a** and **b** only, results are also displayed for the remaining LLD annotations. Numerical results for all annotations analyzed are reported in **Supplementary Table 3** (for **a** and **b**) and **Supplementary Table 8** (for **c**). Numerical results for all 56 complex traits are reported in **Supplementary Table 2** (for **a**), **Supplementary Table 7** (for **b**), and **Supplementary Table 9** (for **c**). Asterisks indicate significance at $P < 0.05$ after Bonferroni correction (0.05/43, 0.05/43, and 0.05/6 for **a–c**, respectively). Error bars represent 95% confidence intervals.

purposes (**Fig. 4**, **Supplementary Fig. 9**, and **Supplementary Tables 9** and **12**). These proportions are computed on the basis of a joint fit of the baseline-LD model but measure the heritability explained by each quintile of each annotation while including the effects of other annotations—in contrast to standardized effect sizes, τ^* , which are conditioned on all other annotations and measure the additional contribution of one annotation to the model. The youngest 20% of common SNPs (based on MAF-adjusted predicted allele age) explained 3.9 times more heritability than the oldest 20%. This is even larger than MAF-dependent effects, in which the 20% of common SNPs with the largest MAFs (>38%) explained 1.8 times more heritability than the 20% with the smallest MAFs (<10%). (We note that slightly smaller per-SNP heritability for less common variants is consistent with somewhat larger per-allele effect sizes for less common variants as expected under realistic levels of negative selection⁹, because less common variants with the same per-allele effect size explain less heritability in proportion to $p(1-p)$, although larger per-SNP heritability for less common variants is theoretically possible under exceedingly strong levels of negative selection; see **Supplementary Fig. 10** and below for additional discussion of MAF-dependent effects.) The heritability explained by quintiles of recombination rate was roughly flat (in contrast to τ^* , which conditions on effects of other annotations; **Fig. 3c**) owing to the inclusion of opposing effects of the LLD-AFR (correlation of -0.39 versus same signs of τ^* ; **Supplementary Table 13**) and CpG content (correlation of 0.18 versus opposite signs of τ^* ; **Supplementary Table 14**) annotations; we note that the effect of recombination rate is dominated by its largest (fifth) quintile (that is, recombination rate hotspots; **Supplementary Fig. 11**), explaining the significant decrease in heritability explained between the fourth and fifth quintiles (**Fig. 4**).

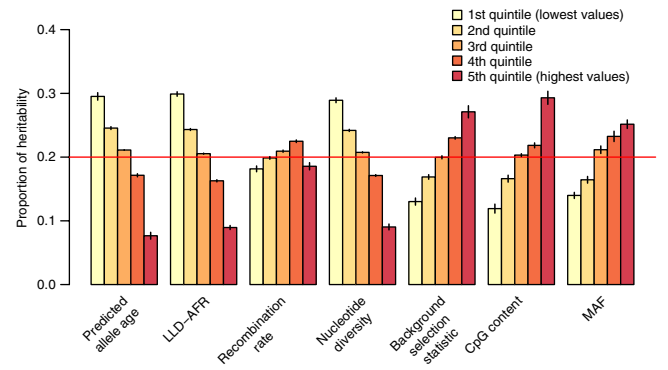


Figure 4 Proportion of heritability explained by the quintiles of each LD-related annotation, subjected to meta-analysis over 31 independent traits. We report results for each LD-related annotation of the baseline-LD model and for MAF for comparison purposes. Numerical results are reported in **Supplementary Table 12**. Results for all 56 complex traits are reported in **Supplementary Figure 9** and **Supplementary Table 9**. Error bars represent jackknife standard errors around the enrichment estimates. The red line indicates the proportion of heritability when there is no enrichment (where 20% of SNPs explain 20% of heritability).

LD-related annotations predict deleterious effects

Our finding that common variants that are more recent tend to explain more complex trait heritability is potentially consistent with the action of negative selection on variants affecting complex traits, as selection has had less time to eliminate recent weakly deleterious variants. We hypothesized that our results for other LD-related annotations might also be explained by the action of negative selection. To investigate this hypothesis, we performed forward simulations³² using a demographic model for African and European populations²⁷ and a range of selection coefficients for deleterious variants (Online Methods). We jointly regressed the absolute value of the selection coefficient against the allele age (now using true allele age instead of predicted allele age), LLD-AFR, recombination rate, and nucleotide diversity annotations from the baseline-LD model to assess whether these annotations are jointly predictive of deleterious effects (the background selection statistic and CpG content annotations could not be investigated as they rely on empirical data). We observed that these four annotations were all significant in the joint analysis (**Fig. 5** and **Supplementary Table 15**), with effect sizes roughly proportional to the standardized effect sizes for trait heritability reported in **Figure 3c**. This suggests that the joint impact of each of these annotations on trait heritability is a consequence of their predictive value for deleterious effects. Indeed, consistent with theory, recent variants are more likely to be deleterious because selection has had less time to remove them³³, variants in regions with low recombination rates are more likely to be deleterious owing to reduced efficiency of selection (Hill–Robertson effect²⁸), and variants in regions with low nucleotide diversity are more likely to be deleterious because these regions tend to have a greater fraction of sites under purifying selection³⁴. In addition, the LLD-AFR annotation contains information complementary to allele age, recombination rate, and nucleotide diversity; we note that LLD-AFR contains roughly the same amount of information (that is, the same effect) as LLD measured in an ancestral population sampled just before the out-of-Africa event (**Supplementary Fig. 12**). We further determined that the predictive value of the nucleotide diversity annotation is contingent on the non-homogeneous distribution of selection coefficients and that the predictive value of the LLD-AFR annotation is largely contingent on the out-of-Africa bottleneck, as the LLD effect




Figure 5 Forward simulations confirm that LD-related annotations predict deleterious effects. We report standardized coefficients for each of four LD-related annotations in a joint regression of absolute selection coefficients against these annotations in data from forward simulations (see text). Numerical results are reported in **Supplementary Table 15**. Error bars represent 95% confidence intervals around the regression coefficient estimates.

disappears in a model with constant population size with a homogeneous distribution of selection coefficients (**Supplementary Fig. 13**). We finally note that we did not expect our results for LD-related annotations to be a signature of positive selection on variants affecting complex traits, as beneficial alleles tend to have increased LD³⁵ and more efficient selection in regions with high recombination rates²⁸, each of which would be inconsistent with the results in **Figure 3c**; indeed, forward simulations involving beneficial mutations confirmed that the LD-related annotations associated with per-SNP heritability do not predict beneficial effects (**Supplementary Fig. 14**).

DISCUSSION

In this study, we assessed the LD-dependent architecture of human complex traits by extending stratified LD score regression⁸ from binary to continuous annotations, an approach that produces robust results in simulations (**Fig. 6**, Online Methods, and **Supplementary Table 16**). We determined that SNPs with low LLD have larger per-SNP heritability across all 56 complex traits analyzed. More than half of this signal can be explained by functional annotations that are negatively correlated with LLD and enriched for heritability, such as DHSs and histone marks. The remaining signal is largely driven by MAF-adjusted predicted allele age, as more recent alleles have larger per-SNP heritability in 55 of the 56 complex traits analyzed, but we also observed multiple jointly significant effects of other LD-related annotations. We showed via forward simulations that all of these jointly significant effects are consistent with the action of negative selection on deleterious variants. As noted above, recent variants are more likely to be deleterious since selection has had less time to remove them³³, variants in regions with low recombination rates are more likely to be deleterious owing to reduced efficiency of negative selection (Hill–Robertson effect²⁸), and variants in regions with low nucleotide diversity are more likely to be deleterious because these regions tend to have a greater fraction of sites under purifying selection³⁴; we also observed higher per-SNP heritability for SNPs with low values of LLD-AFR, capturing a property of variant history that is currently unknown. We note that our genome-wide results on recombination rate differ from the results of Hussin *et al.*¹⁶, who determined that regions of low recombination rate are enriched for exonic deleterious and disease-associated variants: although we do observe a similar recombination rate effect (consistent with the Hill–Robertson effect²⁸) for jointly estimated effect sizes τ^* ,

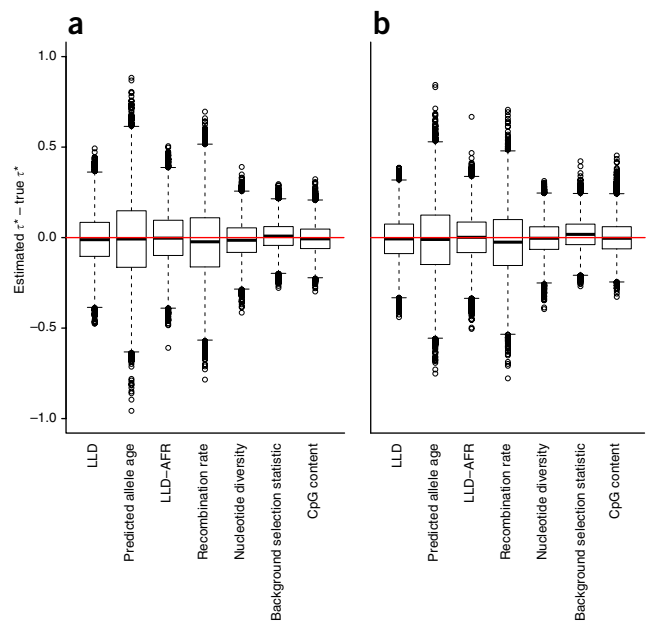


Figure 6 Simulations to assess the extension of stratified LD score regression to continuous LD-related annotations. (**a,b**) We report bias (estimated versus true τ^*) across 10,000 simulations for null simulations with MAF-dependent architecture (**a**) and causal simulations with MAF + LD-dependent architecture (**b**). Results for null simulations with MAF-independent architecture are reported in **Supplementary Figure 15**. The median value of each bias is displayed as a band inside each box. Boxes denote values in the second and third quartiles. The length of each whisker is 1.5 times the interquartile range (defined as the height of each box). All values lying outside the whiskers are considered to be outliers. The red line indicates no bias. Numerical results are reported in **Supplementary Table 16**.

which are conditioned on other annotations and measure the additional contribution of one annotation to the model, this effect is largely canceled out when including the opposing effects of other annotations (**Fig. 4** and **Supplementary Tables 13** and **14**).

While negative selection has long been hypothesized to shape genetic diversity²⁴ and previous studies have emphasized the importance of allele age^{21,36,37} and recombination rate¹⁶, our study demonstrates the impact of negative selection on complex traits on a polygenic, genome-wide scale. Specifically, our results demonstrate that common variants associated with complex traits are weakly deleterious, confirming a hypothesized relationship between the effect size of a variant and its selection coefficient s (refs. 38–41). One of the implications of this finding is that we expect larger per-allele effect sizes for less common variants, consistent with only slightly smaller (per-SNP) heritability explained (**Fig. 4**); this expectation also applies to rare variants, which we do not analyze here.

Our results on LD-dependent architectures have several implications for downstream analyses. First, recent work has suggested that the problem of LD-related bias in SNP heritability estimates^{11,12} could be addressed by modeling regional LD (LD-REG) in addition to MAF¹³. On the other hand, our baseline-LD model contains a considerably larger number of parameters, increasing model complexity but more accurately resolving the underlying signal; in particular, our results suggest that modeling predicted allele age may be more informative than modeling regional LD (**Fig. 4**). Second, previous studies have shown limited improvements in polygenic prediction accuracy^{7,42} and association power^{43,44} using functional annotations,

perhaps because the annotations analyzed in those studies have pervasive LD between in-annotation and out-of-annotation SNPs⁷; however, our LD-related annotations by definition should not have this limitation, making them potentially more useful in those contexts. Third, although SNPs with low LLD have larger causal effect sizes, SNPs with high LLD may have larger χ^2 statistics if they tag multiple causal variants. In the presence of multiple causal variants, fine-mapping strategies based on ranking *P* values⁴⁵ might thus favor high-LLD non-causal variants over causal low-LLD variants. For this reason, approaches that explicitly model multiple causal variants while incorporating LD-dependent architectures using integrative methods⁴⁶ might improve fine-mapping accuracy. Fourth, we observed that predicted allele age is substantially smaller (>0.1 s.d. below average) in transcription start site (TSS), coding, conserved, and UTR regions and below average for all functional annotations except repressed regions (Supplementary Table 17), consistent with stronger selection. The identification of functional noncoding regions under strong selective constraint could be used to improve variant prioritization in whole-genome sequencing studies^{40,47}.

Although our work has provided insights on the genetic architecture of human complex traits, it has several limitations. First, our extension of stratified LD score regression assumes a linear effect of each continuous annotation (equations (1) and (3); Online Methods), which may not always hold; however, this assumption appears reasonable in the continuous annotations that we analyzed (Fig. 4). Second, we restricted all of our analyses to common variants (Online Methods and Supplementary Note), as stratified LD score regression has several limitations when applied to rare variants⁸. Third, while the allele age predictions produced by ARGweaver²¹ were of critical value to this study, they have $>10\%$ missing data, were computed on only 54 sequenced individuals (including only 13 Europeans), and rely on a demographic model with constant population size; the development of computationally tractable methods for predicting allele age remains a research direction of high interest. Fourth, while the effect directions of the LD-related annotations we analyzed were remarkably consistent across all 56 complex traits analyzed, this result does not imply that negative selection acts directly on each of these traits, as selection may be acting on pleiotropic traits⁴⁸. Fifth, while our results suggest that negative selection has a greater impact than positive selection on the genetic architecture of human complex traits, we cannot draw broader conclusions about the roles of negative and positive selection in shaping the human genome^{49,50}. In addition, our forward simulations did not include balancing selection, whose main genomic signature (in contrast to negative selection) is increased nucleotide diversity⁵¹ and would not explain the results in Figure 3c, or stabilizing selection, which uses negative selection to favor intermediate values of phenotypes over extreme values. Sixth, the interpretation of some LD-related annotations remains unclear. The LLD-AFR annotation captures a property of variant history that is currently unknown. The CpG content annotation is highly correlated to the GoNL local mutation rate map annotation⁵² ($r = 0.86$; Supplementary Table 18), but that annotation does not have a significant effect on trait heritability when conditioned on the baseline model (Supplementary Table 19), suggesting that the CpG content annotation might instead tag some functional process absent from the baseline model; indeed, some of our LD-related annotations could be viewed as proxies for currently unknown functional annotations. Despite all of these limitations, our results convincingly demonstrate the action of negative selection on deleterious variants that affect complex traits, complementing efforts to learn about negative selection by analyzing much smaller rare variant data sets.

URLs. ldsc software and tutorials, <http://www.github.com/bulik/ldsc>; baseline-LD annotations, <https://data.broadinstitute.org/alkesgroup/LDSCORE/>; 1000 Genomes Project Phase 3 data, <ftp://ftp.1000genomes.ebi.ac.uk/vol1/ftp/release/20130502>; PLINK software, <https://www.cog-genomics.org/plink2>; ARGweaver allele ages, http://compngen.cshl.edu/ARGweaver/CG_results/download; Oxford recombination map, http://www.shapeit.fr/files/genetic_map_b37.tar.gz; African-American and deCODE recombination maps, http://www.well.ox.ac.uk/~anjali/AAMap/maps_b37.tar.gz; BEDTools software, <http://bedtools.readthedocs.io/en/latest>; human reference sequence, ftp://ftp-trace.ncbi.nih.gov/1000genomes/ftp/technical/reference/human_g1k_v37.fasta.gz; GCTA software, <http://cns.genomics.com/software/gcta/#Download>; ARGON software, <https://github.com/pierpal/ARGON>; BOLT-LMM software, <https://data.broadinstitute.org/alkesgroup/BOLT-LMM>; UK Biobank, <http://www.ukbiobank.ac.uk/>; UK Biobank Genotyping and QC Documentation, http://www.ukbiobank.ac.uk/wp-content/uploads/2014/04/UKBiobank_genotyping_QC_documentation-web.pdf; SLiM software, <https://messengerlab.org/slim/>.

METHODS

Methods, including statements of data availability and any associated accession codes and references, are available in the [online version of the paper](#).

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank the research participants and employees of 23andMe for making this work possible. We thank S. Sunyaev, Y. Reshef, G. Kichaev, D. Speed, and F. Day for helpful discussions. This research has been conducted using the UK Biobank Resource (application number 16549). This research was funded by NIH grants R01 MH101244, R01 MH107649, and U01 HG009088.

AUTHOR CONTRIBUTIONS

S.G. and A.L.P. designed experiments. S.G. performed experiments. S.G., H.K.F., N.A.F., and P.-R.L. analyzed data. S.G. and A.L.P. wrote the manuscript with assistance from H.K.F., N.A.F., P.-R.L., P.F.P., X.L., A.S., B.B.-S., B.M.N., and A.G.

COMPETING FINANCIAL INTERESTS

The authors declare competing financial interests: details are available in the [online version of the paper](#).

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>. Publisher's note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Yang, J. *et al.* Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* **42**, 565–569 (2010).
2. Yang, J., Lee, S.H., Goddard, M.E. & Visscher, P.M. GCTA: a tool for genome-wide complex trait analysis. *Am. J. Hum. Genet.* **88**, 76–82 (2011).
3. Yang, J. *et al.* Genome partitioning of genetic variation for complex traits using common SNPs. *Nat. Genet.* **43**, 519–525 (2011).
4. Shi, H., Kichaev, G. & Pasaniuc, B. Contrasting the genetic architecture of 30 complex traits from summary association data. *Am. J. Hum. Genet.* **99**, 139–153 (2016).
5. Lee, S.H. *et al.* Estimating the proportion of variation in susceptibility to schizophrenia captured by common SNPs. *Nat. Genet.* **44**, 247–250 (2012).
6. Davis, L.K. *et al.* Partitioning the heritability of Tourette syndrome and obsessive compulsive disorder reveals differences in genetic architecture. *PLoS Genet.* **9**, e1003864 (2013).
7. Gusev, A. *et al.* Partitioning heritability of regulatory and cell-type-specific variants across 11 common diseases. *Am. J. Hum. Genet.* **95**, 535–552 (2014).
8. Finucane, H.K. *et al.* Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nat. Genet.* **47**, 1228–1235 (2015).
9. Loh, P.-R. *et al.* Contrasting genetic architectures of schizophrenia and other complex diseases using fast variance-components analysis. *Nat. Genet.* **47**, 1385–1392 (2015).
10. Won, H.-H. *et al.* Disproportionate contributions of select genomic compartments and cell types to genetic risk for coronary artery disease. *PLoS Genet.* **11**, e1005622 (2015).
11. Speed, D., Hemani, G., Johnson, M.R. & Balding, D.J. Improved heritability estimation from genome-wide SNPs. *Am. J. Hum. Genet.* **91**, 1011–1021 (2012).

12. Gusev, A. *et al.* Quantifying missing heritability at known GWAS loci. *PLoS Genet.* **9**, e1003993 (2013).
13. Yang, J. *et al.* Genetic variance estimation with imputed variants finds negligible missing heritability for human height and body mass index. *Nat. Genet.* **47**, 1114–1120 (2015).
14. Koch, C.M. *et al.* The landscape of histone modifications across 1% of the human genome in five human cell lines. *Genome Res.* **17**, 691–707 (2007).
15. Smith, A.V., Thomas, D.J., Munro, H.M. & Abecasis, G.R. Sequence features in regions of weak and strong linkage disequilibrium. *Genome Res.* **15**, 1519–1534 (2005).
16. Hussin, J.G. *et al.* Recombination affects accumulation of damaging and disease-associated mutations in human populations. *Nat. Genet.* **47**, 400–404 (2015).
17. 1000 Genomes Project Consortium. A global reference for human genetic variation. *Nature* **526**, 68–74 (2015).
18. Bulik-Sullivan, B.K. *et al.* LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nat. Genet.* **47**, 291–295 (2015).
19. Liu, X. *et al.* Functional architectures of local and distal regulation of gene expression in multiple human tissues. *Am. J. Hum. Genet.* **100**, 605–616 (2017).
20. Bulik-Sullivan, B. *et al.* An atlas of genetic correlations across human diseases and traits. *Nat. Genet.* **47**, 1236–1241 (2015).
21. Rasmussen, M.D., Hubisz, M.J., Gronau, I. & Siepel, A. Genome-wide inference of ancestral recombination graphs. *PLoS Genet.* **10**, e1004342 (2014).
22. McVean, G.A.T. *et al.* The fine-scale structure of recombination rate variation in the human genome. *Science* **304**, 581–584 (2004).
23. Myers, S., Bottolo, L., Freeman, C., McVean, G. & Donnelly, P. A fine-scale map of recombination rates and hotspots across the human genome. *Science* **310**, 321–324 (2005).
24. McVicker, G., Gordon, D., Davis, C. & Green, P. Widespread genomic signatures of natural selection in hominid evolution. *PLoS Genet.* **5**, e1000471 (2009).
25. Koren, A. *et al.* Differential relationship of DNA replication timing to different forms of human mutation and variation. *Am. J. Hum. Genet.* **91**, 1033–1040 (2012).
26. Palamara, P.F. ARGON: fast, whole-genome simulation of the discrete time Wright-Fisher process. *Bioinformatics* **32**, 3032–3034 (2016).
27. Gravel, S. *et al.* Demographic history and rare allele sharing among human populations. *Proc. Natl. Acad. Sci. USA* **108**, 11983–11988 (2011).
28. Hill, W.G. & Robertson, A. The effect of linkage on limits to artificial selection. *Genet. Res.* **8**, 269–294 (1966).
29. UK10K Consortium. The UK10K project identifies rare variants in health and disease. *Nature* **526**, 82–90 (2015).
30. International Schizophrenia Consortium. Common polygenic variation contributes to risk of schizophrenia and bipolar disorder. *Nature* **460**, 748–752 (2009).
31. Karlsson, E.K., Kwiatkowski, D.P. & Sabeti, P.C. Natural selection and infectious disease in human populations. *Nat. Rev. Genet.* **15**, 379–393 (2014).
32. Messer, P.W. SLiM: simulating evolution with selection and linkage. *Genetics* **194**, 1037–1039 (2013).
33. Maruyama, T. The age of a rare mutant gene in a large population. *Am. J. Hum. Genet.* **26**, 669–673 (1974).
34. Charlesworth, B., Morgan, M.T. & Charlesworth, D. The effect of deleterious mutations on neutral molecular variation. *Genetics* **134**, 1289–1303 (1993).
35. Sabeti, P.C. *et al.* Detecting recent positive selection in the human genome from haplotype structure. *Nature* **419**, 832–837 (2002).
36. Fu, W. *et al.* Analysis of 6,515 exomes reveals the recent origin of most human protein-coding variants. *Nature* **493**, 216–220 (2013).
37. Kiezun, A. *et al.* Deleterious alleles in the human genome are on average younger than neutral alleles of the same frequency. *PLoS Genet.* **9**, e1003301 (2013).
38. Eyre-Walker, A. Genetic architecture of a complex trait and its implications for fitness and genome-wide association studies. *Proc. Natl. Acad. Sci. USA* **107** (Suppl. 1), 1752–1756 (2010).
39. Agarwala, V., Flannick, J., Sunyaev, S., GoT2D Consortium & Altshuler, D. Evaluating empirical bounds on complex disease genetic architecture. *Nat. Genet.* **45**, 1418–1427 (2013).
40. Zuk, O. *et al.* Searching for missing heritability: designing rare variant association studies. *Proc. Natl. Acad. Sci. USA* **111**, E455–E464 (2014).
41. Fuchsberger, C. *et al.* The genetic architecture of type 2 diabetes. *Nature* **536**, 41–47 (2016).
42. Shi, J. *et al.* Winner's curse correction and variable thresholding improve performance of polygenic risk modeling based on genome-wide association study summary-level data. *PLoS Genet.* **12**, e1006493 (2016).
43. Pickrell, J.K. Joint analysis of functional genomic data and genome-wide association studies of 18 human traits. *Am. J. Hum. Genet.* **94**, 559–573 (2014).
44. Sveinbjornsson, G. *et al.* Weighting sequence variants based on their annotation increases power of whole-genome association studies. *Nat. Genet.* **48**, 314–317 (2016).
45. Faye, L.L., Machiela, M.J., Kraft, P., Bull, S.B. & Sun, L. Re-ranking sequencing variants in the post-GWAS era for accurate causal variant identification. *PLoS Genet.* **9**, e1003609 (2013).
46. Kichaev, G. *et al.* Integrating functional data to prioritize causal variants in statistical fine-mapping studies. *PLoS Genet.* **10**, e1004722 (2014).
47. Lee, S., Abecasis, G.R., Boehnke, M. & Lin, X. Rare-variant association analysis: study designs and statistical tests. *Am. J. Hum. Genet.* **95**, 5–23 (2014).
48. Johnson, T. & Barton, N. Theoretical models of selection and mutation on quantitative traits. *Philos. Trans. R. Soc. B Biol.* **360**, 1411–1425 (2005).
49. Hernandez, R.D. *et al.* Classic selective sweeps were rare in recent human evolution. *Science* **331**, 920–924 (2011).
50. Enard, D., Messer, P.W. & Petrov, D.A. Genome-wide signals of positive selection in human evolution. *Genome Res.* **24**, 885–895 (2014).
51. Charlesworth, D. Balancing selection and its effects on sequences in nearby genome regions. *PLoS Genet.* **2**, e64 (2006).
52. Francioli, L.C. *et al.* Genome-wide patterns and properties of *de novo* mutations in humans. *Nat. Genet.* **47**, 822–826 (2015).

ONLINE METHODS

Extension of stratified LD score regression to continuous annotations. The derivation of stratified LD score regression using binary annotations has previously been described⁸. Here we extend the method to continuous-valued annotations.

Suppose that we have a sample of N individuals and a vector $y = (y_1, \dots, y_N)$ of quantitative phenotypes, standardized to mean 0 and variance 1. We assume the infinitesimal linear model

$$y = X\beta + \varepsilon \quad (2)$$

where X is an $N \times M$ matrix of standardized genotypes, $\beta = (\beta_1, \dots, \beta_M)$ is the vector of per-normalized-genotype effect size, and $\varepsilon = (\varepsilon_1, \dots, \varepsilon_N)$ is a mean 0 vector of residuals with variance σ_ε^2 . Here we are interested in modeling β as a mean 0 vector whose variance depends on C continuous-valued annotations, a_1, \dots, a_c

$$\text{var}(\beta_j) = \sum_c a_c(j)\tau_c \quad (3)$$

where $a_c(j)$ is the value of annotation a_c at SNP j and τ_c represents the per-SNP contribution of one unit of the annotation a_c to heritability. This is a generalization of stratified LD score regression⁸, with $a_c(j) \in \{0,1\}$ if annotation a_c has binary values.

Under this model

$$E[\chi_j^2] = N \sum_c \tau_c l(j,c) + 1 \quad (4)$$

where

$$l(j,c) = \sum_k a_c(k)r_{jk}^2$$

is the LD score of SNP j with respect to annotations a_c and r_{jk} is the correlation between SNPs j and k (**Supplementary Note**). Given a vector of χ^2 statistics and LD scores computed from a reference sample, this equation allows us to obtain estimates $\hat{\tau}_c$ of τ_c (ref. 8).

We were interested in both comparing the estimated effect sizes of the different annotations and performing meta-analysis of them across different traits. For this reason, we focused on per-standardized-annotation effect sizes τ_c^* , defined as the additive change in per-SNP heritability associated with a 1 s.d. increase in the value of the annotation, divided by the average per-SNP heritability over all SNPs for the trait, and computed as

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

where h_g^2 the estimated SNP heritability of the trait computed as

$$h_g^2 = \sum_j \text{var}(\beta_j) = \sum_j \sum_c a_c(j)\tau_c$$

$M_{h_g^2}$ is the number of SNPs used to compute h_g^2 , and sd_c is the standard deviation of the annotation a_c . The standard error of $\hat{\tau}_c$ was computed using a block jackknife over SNPs with 200 equally sized blocks of adjacent SNPs⁸ and was used to compute a z score to test for the significance of τ_c^* .

To interpret the heritability explained by a continuous-valued annotation a_c , we computed the expected heritability of each quintile of its annotations. Let $C_{c,q}$ denote the q th quintile of annotation a_c , so that $h_g^2(C_{c,1})$ and $h_g^2(C_{c,5})$ represent the heritability explained by the 20% of SNPs with the lowest and highest values of a_c , respectively. We used the equation

$$h_g^2(C_{c,q}) = \sum_{j \in C_{c,q}} \sum_{c'} a_{c'}(j)\tau_{c'}$$

to estimate $h_g^2(C_{c,q})$.

Application of stratified LD score regression was performed using the guidelines from Finucane *et al.*⁸ and was restricted to data sets of European ancestry (**Supplementary Note**).

Baseline model and functional annotations. The 59 functional annotations that we used to define the baseline model consist of the 53 binary annotations from ref. 8 and an additional 6 annotations. The 53 annotations are derived from 24 main annotations, including coding, UTR, promoter, and intronic regions, the histone marks monomethylation (H3K4me1) and trimethylation (H3K4me3) of histone H3 at lysine 4, acetylation of histone H3 at lysine 9 (H3K9ac), and two versions of acetylation of histone H3 at lysine 27 (H3K27ac), open chromatin as reflected by DHSs, combined chromHMM and Segway predictions (which make use of many Encyclopedia of DNA Elements (ENCODE) annotations to produce a single partition of the genome into seven underlying chromatin states), regions that are conserved in mammals, super-enhancers, and FANTOM5 enhancers. The 53 annotations also include 500-bp windows around each of the 24 main annotations, 100-bp windows around 4 of the main annotations, and an annotation containing all SNPs. We added four binary annotations based on super-enhancers and typical enhancers⁵³, as previously described¹⁹. We also added two conservation annotations based on GERP++ scores⁵⁴, including one continuous annotation based on the neutral rate score (NS) and one binary annotation based on a rejected substitutions (RS) score ≥ 4 , as we observed significant effects for these annotations (**Supplementary Table 20**). We did not include 500-bp windows around the GERP-NS annotation (which is a continuous annotation) or the GERP-RS annotation (which is defined separately for each base pair).

MAF adjustment and LLD annotations. To investigate the LD-dependent architecture of human complex traits, it is essential to account for the relationship between MAF and LD. Indeed, common variants have both higher LD scores and higher per-SNP heritability^{5,9}. For this reason, all of our stratified LD score regression analyses included ten MAF bins coded as ten binary annotations (all with $\text{MAF} \geq 0.05$; **Supplementary Table 21**) in addition to an annotation containing all SNPs.

To quantify the LLD of reference SNPs, we first computed LD scores, defined as the sum of the squared correlations of each SNP with all nearby SNPs in a 1-cM window, using the *ldsc* software. Then, we adjusted these values for MAF via MAF-stratified quantile normalization: for each MAF bin, LD scores were quantile normalized to a normal distribution of mean 0 and variance 1. The LLD of rare variants ($\text{MAF} < 0.05$) was fixed to 0. Because stratified LD score regression is designed to quantify the heritability explained by common SNPs and the heritability explained by rare variants is hypothesized to be relatively low^{1,5,55}, we excluded rare variants from all MAF-adjusted annotations. For the LLD model (**Fig. 1**), we thus modeled the variance of the per-normalized genotype effect size of SNP j as

$$\text{var}(\beta_j) = \tau_0 + \sum_{m=1}^{10} 1_{j \in \text{MAF bin } m} \tau_m + \text{LLD}_j \tau_{\text{LLD}} \quad (6)$$

where τ_0 is an intercept term modeling the per-SNP contribution of each SNP to heritability, $1_{j \in \text{MAF bin } m}$ is an indicator function with value 1 if SNP j belongs to MAF bin m and 0 otherwise, τ_m is the per-SNP contribution of a SNP in MAF bin m to heritability, and τ_{LLD} is the contribution of one unit of the annotation LLD to heritability.

The LLD- D' annotation of a SNP was measured by summing the D' coefficients of that SNP with all nearby SNPs in a ± 0.5 -cM window. Version 1.90b3 of PLINK 2 software⁵⁶ (see URLs) was used to compute D' coefficients for each pair of SNPs. The LLD of a genomic region (LLD-REG) was measured by averaging in 100-kb windows the LD scores computed in 20-Mb regions (ignoring LD $r^2 < 0.01$), as previously described¹³, using the $-ld$ -score-region option of version 1.25.1 of GCTA software². LLD- D' was adjusted for MAF via MAF-stratified quantile normalization. LLD-REG was quantile normalized without MAF adjustment because it is a regional annotation.

The LLD-AFR annotation was measured by computing the LD scores of reference SNPs in 440 unrelated African samples from Phase 3 of the 1000 Genomes Project (the ACB and ASW populations were removed because of the presence of European admixture). LD scores for reference SNPs that were absent in African samples were set to 1. LLD-AFR was also adjusted for MAF via MAF-stratified quantile normalization, using the same European MAF bins.

Other LD-related annotations. We used allele age as predicted by the ARGweaver²¹ method, estimated using 54 unrelated sequenced individuals (including 13 Europeans; see URLs). This annotation was also adjusted for MAF via MAF-stratified quantile normalization, as common variants tend to be older (the correlation for common reference SNPs between available ARGweaver allele ages and MAF is 0.16). 10.2% of common reference SNPs had missing values for predicted allele age; these values were excluded during the MAF-stratified quantile normalization process, and corresponding MAF-adjusted predicted allele ages were set to 0. Adding a binary annotation indicating missing allele age information for common reference SNPs did not change the effect size estimates for predicted allele age (Supplementary Table 22).

Recombination rate, nucleotide diversity, GC content, and CpG content were computed using windows of different sizes around each SNP: ± 10 kb, ± 50 kb, ± 100 kb, ± 500 kb, and $\pm 1,000$ kb. Recombination rates (measured in cM/Mb) were computed from three recombination maps (see URLs): the Oxford map, which estimates recombination rates from LD patterns in African, European, and Asian populations from HapMap 2 (refs. 22,23); the African-American map, which estimates recombination rates from admixture patterns in African-American individuals⁵⁷; and the deCODE map, which estimates recombination rates from Icelandic parent-offspring pairs⁵⁸. These recombination maps measure recombination rates at different time scales: the deCODE map measures recombination that occurred in recent generations, the African-American map measures recombination that occurred in the past ~ 20 generations, and the Oxford map measures recombination that occurred further back in time. The genetic positions of surrounding windows were interpolated linearly from recombination maps using PLINK. We determined that the Oxford map provided the most significant results (Supplementary Table 3), suggesting that the impact of recombination rate on trait heritability operates over a long time scale; we thus used the Oxford map in all primary analyses. Nucleotide diversity was measured as the number of reference SNPs (with minor allele count ≥ 5) per kilobase. Measuring diversity on all 1000 Genomes Project SNPs (down to singletons or doubletons) or the fraction of rare variants⁵⁹ (that is, the diversity of rare variants with allele count < 5) did not furnish more significant results (data not shown). GC content and CpG content were measured using version 2.17.0 of BEDTools software⁶⁰ and the human reference sequence used for the 1000 Genomes Project (see URLs).

The background selection statistic was computed as $1 - \text{McVicker } B$ statistic²⁴ to facilitate interpretation of the results. Background selection statistic values close to 1 represent nearly complete removal of diversity as a result of background selection, and values near 0 indicate little effect. Replication timing was based on the annotation of Koren *et al.*²⁵. 0.19% and 0.27% of reference SNPs had missing values for background selection and replication timing, respectively; these were replaced by the median annotation value based on the remaining reference SNPs.

Finally, telomeres and centromeres were defined using window sizes of 5, 10, and 15 Mb, as described by Smith *et al.*¹⁵.

We thus created 43 LD-related annotations in total (Supplementary Table 23). For annotations computed with different window sizes or using different data sources for recombination rate, the one producing the most significant P value after conditioning on the baseline model was selected as the primary annotation (Supplementary Table 3). Except for telomere and centromere annotations that were not significant in this analysis, other annotations had consistent results with adjacent window sizes. To overcome overfitting, we used a Bonferroni threshold of $0.05/43 = 1.16 \times 10^{-3}$ to assess statistical significance when analyzing one LD-related annotation at a time. We note that this procedure did not affect our final conclusions (Supplementary Fig. 5).

Choice of traits for main analyses and meta-analysis. Stratified LD score regression was applied to 29 publicly available GWAS summary statistic data sets^{61–81} (for age at menopause⁷⁹, effect sizes are publicly available but sample sizes for each SNP were obtained through collaboration), 18 summary statistic data sets from 23andMe, and summary statistics of 15 traits from UK Biobank (Supplementary Note). This led to a total of 62 summary statistic data sets spanning 56 traits (5 traits were represented in multiple data sets) with an average sample size of 101,401 (computed using the largest single data set for each trait; the average sample size of the 62 data sets was 101,989). Analyses

were restricted to traits for which the z score of total SNP heritability computed using the baseline model was at least 6 (Supplementary Table 1). The traits displayed in Figure 1 were selected by prioritizing according to the total SNP heritability, excluding traits with absolute genetic correlation > 0.50 (ref. 20). Traits included in the meta-analyses were selected by prioritizing according to the z score of total SNP heritability and excluding genetically correlated traits in overlapping samples by measuring the intercept of cross-trait LD score regression²⁰ as previously described⁸. We retained 31 independent traits (average $N = 84,686$; Supplementary Table 1) and performed random-effects meta-analyses using the R package *rmeta*.

Construction of the baseline-LD model. We first considered a model including the eight LD-related annotations that were significant after being conditioned on ten MAF bins and the baseline model (LLD, predicted allele age, LLD-AFR, recombination rate, nucleotide diversity, the background selection statistic, GC content, and CpG content) and also including ten MAF bins and the baseline model. We removed LD-related annotations that were not significant (in the meta-analysis of 31 independent traits) one at a time on the basis of the least significant P value (GC content was removed first, then LLD). This procedure produced a baseline-LD model with the 59 annotations of the baseline model, the 10 MAF bins, and the 6 remaining LD-related annotations, leading to a total of 75 annotations. We have made these annotations publicly available (see URLs).

Simulations to assess extension of stratified LD score regression to continuous LD-related annotations. To ensure that applying our extension of stratified LD score regression to continuous LD-related annotations does not produce false positive signals or biased results, we simulated quantitative phenotypes from chromosome 1 UK10K data²⁹ (3,567 individuals and 1,041,378 SNPs). In each simulation, we used the 1000 Genomes Project as the reference panel and evaluated all six LD-related annotations of the baseline-LD model (Fig. 3c), as well as the LLD annotation. We also included an annotation containing all SNPs and annotations for ten MAF bins. In each simulation, we set trait heritability to $h^2 = 0.5$ and selected $M = 100,000$ causal SNPs. Causal SNPs were selected randomly from the 673,779 SNPs present in both the UK10K Project and 1000 Genomes Project, such that all causal SNPs were represented in the reference panel. In null simulations, the variance of per-normalized-genotype effect sizes was simulated proportional to $(p(1-p))^{1+\alpha}$ for a variant of frequency p . We considered simulations with both MAF-independent ($\alpha = -1$; where all SNPs have the same contribution to variance) and MAF-dependent ($\alpha = -0.28$; as previously estimated⁹) architectures. In causal simulations (MAF- + LD-dependent architecture), we used the τ coefficients estimated from the meta-analyses reported in Figure 3a and set the variance of per-normalized-genotype effect sizes using the additive model of equation (6), replacing the LLD annotation with the LD-related annotation of interest. These coefficients were rescaled to constrain the variance of each SNP to be positive and the total h^2 of the 100,000 causal SNPs to be 0.5. Phenotypes were simulated with GCTA² (see URLs). We performed 10,000 simulations for each of the three simulation scenarios (null MAF independent, null MAF dependent, and causal MAF + LD dependent). In each simulation, we estimated the effect size $\hat{\tau}$ using HapMap 3 SNPs as regression SNPs, to account for the possibility that causal SNPs were not included in the set of regression SNPs. Corresponding $\hat{\tau}^*$ values were computed using the simulated h^2 value. (We were interested in the bias of the τ parameter, and not in the h_g^2 parameter, which might be underestimated in simulation scenarios where rare variants have large effect sizes. We note that estimates of τ in real phenotypes may be slightly biased by inaccurate h_g^2 estimates, but that this will not lead to false positive nonzero τ^* estimates.) We observed unbiased estimates of τ^* for most annotations in both null simulations (Fig. 6a and Supplementary Fig. 15) and causal simulations (Fig. 6b) (numerical results in Supplementary Table 16). Only the recombination rate annotation exhibited very slight biases (between -0.028 and -0.025) that are nevertheless far from the estimates observed on real data (-0.540 ; Fig. 3a). We also confirmed accurate calibration of standard errors in both null and causal simulations (Supplementary Table 24). We repeated each of these simulations drawing causal SNPs from all UK10K SNPs (to simulate a scenario where causal SNPs are not represented in the reference panel). Results for null simulations were similar to those above, and results

for causal simulations produced slight biases opposite to (that is, slightly underestimating) true effects (**Supplementary Fig. 16**; numerical results in **Supplementary Table 25**).

Forward simulations to assess the connection between LD-related annotations and negative selection. To investigate the connection between the LD-related annotations of the baseline-LD model (predicted allele age, LLD-AFR, recombination rate, and nucleotide diversity; note that the background selection statistic and CpG content cannot be assessed in simulations as they rely on empirical data and that these simulations used true allele age instead of predicted allele age) and the selection coefficient s , we performed forward simulations under a Wright–Fisher model with selection using version 1.8 of SLiM software³² (see URLs). We simulated 1-Mb regions of genetic length 1 cM. To ensure realistic recombination rate patterns, we divided the 1-Mb regions into three recombination environments¹⁶, including a 475-kb coldspot region containing 4.1% of recombination events (0.08 cM/Mb) and a 140-kb region with high recombination rate containing 58.6% of recombination events (4.18 cM/Mb). The mutation rate was again set to 1.65×10^{-8} mutations per base per generation⁸². New mutations had probabilities d of being deleterious with a dominance coefficient of 0.5 and a selection coefficient s drawn from a gamma distribution with parameters -0.05 and 0.2 (as suggested in the SLiM manual) and $1 - d$ of being neutral ($s = 0$). To study the impact of a non-homogeneous distribution of d across the genome, we divided each recombination environment into two subregions and assigned alternate probabilities d_1 and d_2 of being deleterious in these subregions (results reported in **Fig. 5** used $d_1 = 0.60$ and $d_2 = 0.90$). We performed simulations spanning 100,000 generations under two different demographic scenarios. First, we started from a fixed population size of 7,300 individuals, used the realistic demographic model of Gravel *et al.*²⁷ for the last 5,920 generations, and output 500 European genomes and 500 African genomes. Second, we considered a fixed population size of 10,000 individuals and output 500 individual genomes at the last generation. We simulated 200 1-Mb regions in each demographic scenario. LD scores were computed independently in each 1-Mb fragment on the basis of SNPs with minor allele count ≥ 5 ; allele age and LLD-AFR or LLD (depending on the demographic scenario) were adjusted for MAF via MAF-stratified quantile normalization after merging the 200 1-Mb regions. We performed a multivariate linear regression of the absolute value of the (known) selection coefficients $|s|$ against the ten MAF bin annotations, the MAF-adjusted allele age, the MAF-adjusted LLD-AFR or LLD (depending on the demographic scenario), the true recombination rate, and the nucleotide diversity measured in a ± 10 -kb window around each SNP. The above simulations did not include beneficial mutations, but we also performed simulations with beneficial and neutral mutations only to confirm that positive selection cannot explain the observed results of the LD-related annotations of the baseline-LD model (**Supplementary Fig. 13**).

Data availability. The ldsc software, baseline-LD annotations, and a tutorial on how to use ldsc with continuous-valued annotations are available online (see URLs). Links to all publicly available summary statistics are provided in **Supplementary Table 1**. UK Biobank summary statistics have been submitted to the UK Biobank data showcase. A **Life Sciences Reporting Summary** is available.

53. Vahedi, G. *et al.* Super-enhancers delineate disease-associated regulatory nodes in T cells. *Nature* **520**, 558–562 (2015).

54. Davydov, E.V. *et al.* Identifying a high fraction of the human genome to be under selective constraint using GERP++. *PLoS Comput. Biol.* **6**, e1001025 (2010).
55. Lee, S.H. *et al.* Estimation of SNP heritability from dense genotype data. *Am. J. Hum. Genet.* **93**, 1151–1155 (2013).
56. Chang, C.C. *et al.* Second-generation PLINK: rising to the challenge of larger and richer datasets. *Gigascience* **4**, 7 (2015).
57. Hinch, A.G. *et al.* The landscape of recombination in African Americans. *Nature* **476**, 170–175 (2011).
58. Kong, A. *et al.* Fine-scale recombination rate differences between sexes, populations and individuals. *Nature* **467**, 1099–1103 (2010).
59. Gudbjartsson, D.F. *et al.* Large-scale whole-genome sequencing of the Icelandic population. *Nat. Genet.* **47**, 435–444 (2015).
60. Quinlan, A.R. & Hall, I.M. BEDTools: a flexible suite of utilities for comparing genomic features. *Bioinformatics* **26**, 841–842 (2010).
61. Dubois, P.C.A. *et al.* Multiple common variants for celiac disease influencing immune gene expression. *Nat. Genet.* **42**, 295–302 (2010).
62. Lango Allen, H. *et al.* Hundreds of variants clustered in genomic loci and biological pathways affect human height. *Nature* **467**, 832–838 (2010).
63. Speliotes, E.K. *et al.* Association analyses of 249,796 individuals reveal 18 new loci associated with body mass index. *Nat. Genet.* **42**, 937–948 (2010).
64. Teslovich, T.M. *et al.* Biological, clinical and population relevance of 95 loci for blood lipids. *Nature* **466**, 707–713 (2010).
65. Tobacco and Genetics Consortium. Genome-wide meta-analyses identify multiple loci associated with smoking behavior. *Nat. Genet.* **42**, 441–447 (2010).
66. Psychiatric GWAS Consortium Bipolar Disorder Working Group. Large-scale genome-wide association analysis of bipolar disorder identifies a new susceptibility locus near *ODZ4*. *Nat. Genet.* **43**, 977–983 (2011).
67. Schunkert, H. *et al.* Large-scale association analysis identifies 13 new susceptibility loci for coronary artery disease. *Nat. Genet.* **43**, 333–338 (2011).
68. Jostins, L. *et al.* Host–microbe interactions have shaped the genetic architecture of inflammatory bowel disease. *Nature* **491**, 119–124 (2012).
69. Manning, A.K. *et al.* A genome-wide approach accounting for body mass index identifies genetic variants influencing fasting glycemic traits and insulin resistance. *Nat. Genet.* **44**, 659–669 (2012).
70. Morris, A.P. *et al.* Large-scale association analysis provides insights into the genetic architecture and pathophysiology of type 2 diabetes. *Nat. Genet.* **44**, 981–990 (2012).
71. Cross-Disorder Group of the Psychiatric Genomics Consortium. Identification of risk loci with shared effects on five major psychiatric disorders: a genome-wide analysis. *Lancet* **381**, 1371–1379 (2013).
72. Rietveld, C.A. *et al.* GWAS of 126,559 individuals identifies genetic variants associated with educational attainment. *Science* **340**, 1467–1471 (2013).
73. Boraska, V. *et al.* A genome-wide association study of anorexia nervosa. *Mol. Psychiatry* **19**, 1085–1094 (2014).
74. Okada, Y. *et al.* Genetics of rheumatoid arthritis contributes to biology and drug discovery. *Nature* **506**, 376–381 (2014).
75. Perry, J.R.B. *et al.* Parent-of-origin-specific allelic associations among 106 genomic loci for age at menarche. *Nature* **514**, 92–97 (2014).
76. Schizophrenia Working Group of the Psychiatric Genomics Consortium. Biological insights from 108 schizophrenia-associated genetic loci. *Nature* **511**, 421–427 (2014).
77. Benthall, J. *et al.* Genetic association analyses implicate aberrant regulation of innate and adaptive immunity genes in the pathogenesis of systemic lupus erythematosus. *Nat. Genet.* **47**, 1457–1464 (2015).
78. Cordell, H.J. *et al.* International genome-wide meta-analysis identifies new primary biliary cirrhosis risk loci and targetable pathogenic pathways. *Nat. Commun.* **6**, 8019 (2015).
79. Day, F.R. *et al.* Large-scale genomic analyses link reproductive aging to hypothalamic signaling, breast cancer susceptibility and *BRCA1*-mediated DNA repair. *Nat. Genet.* **47**, 1294–1303 (2015).
80. Okbay, A. *et al.* Genome-wide association study identifies 74 loci associated with educational attainment. *Nature* **533**, 539–542 (2016).
81. Okbay, A. *et al.* Genetic variants associated with subjective well-being, depressive symptoms, and neuroticism identified through genome-wide analyses. *Nat. Genet.* **48**, 624–633 (2016).
82. Palamara, P.F. *et al.* Leveraging distant relatedness to quantify human mutation and gene-conversion rates. *Am. J. Hum. Genet.* **97**, 775–789 (2015).

Life Sciences Reporting Summary

Nature Research wishes to improve the reproducibility of the work we publish. This form is published with all life science papers and is intended to promote consistency and transparency in reporting. All life sciences submissions use this form; while some list items might not apply to an individual manuscript, all fields must be completed for clarity.

For further information on the points included in this form, see [Reporting Life Sciences Research](#). For further information on Nature Research policies, including our [data availability policy](#), see [Authors & Referees](#) and the [Editorial Policy Checklist](#).

▶ Experimental design

1. Sample size

Describe how sample size was determined.

Our work is based on summary statistics. We restricted ourself to datasets with sufficient sample size, SNP heritability and polygenicity by measuring the z score of total SNP heritability estimated by stratified LD score regression, such as recommended in Finucane et al. 2015 Nat Genet. (See "Choice of traits for main analyses and meta-analysis." online method section)

2. Data exclusions

Describe any data exclusions.

Our study was restricted to data sets of European ancestry (see last line of first Online Method section "Extension of stratified LD score regression to continuous annotations" + Supplementary Notes "23andMe dataset." and "UKBiobank dataset." for specific criteria applied on these 2 datasets)

3. Replication

Describe whether the experimental findings were reliably reproduced.

No overall replication was performed as we meta-analyzed data from more than 30 studies and we obtained strongly significant Pvalues (from $P=10^{-100}$ to $P<10^{-20}$). However, Supplementary Figures 5, 6, and 7 show that our results across different definition of the analyzed annotations, and accross traits with different summary statistics available.

4. Randomization

Describe how samples/organisms/participants were allocated into experimental groups.

NA. We only analyzed summary statistics in this study.

5. Blinding

Describe whether the investigators were blinded to group allocation during data collection and/or analysis.

We did not collect data for this study, but analyzed publicly available summary statistics, summary statistics from 23andMe, and summary statistics from UK Biobank.

Note: all studies involving animals and/or human research participants must disclose whether blinding and randomization were used.

6. Statistical parameters

For all figures and tables that use statistical methods, confirm that the following items are present in relevant figure legends (or the Methods section if additional space is needed).

n/a Confirmed

- The exact sample size (n) for each experimental group/condition, given as a discrete number and unit of measurement (animals, litters, cultures, etc.)
- A description of how samples were collected, noting whether measurements were taken from distinct samples or whether the same sample was measured repeatedly.
- A statement indicating how many times each experiment was replicated
- The statistical test(s) used and whether they are one- or two-sided (note: only common tests should be described solely by name; more complex techniques should be described in the Methods section)
- A description of any assumptions or corrections, such as an adjustment for multiple comparisons
- The test results (e.g. p values) given as exact values whenever possible and with confidence intervals noted
- A summary of the descriptive statistics, including central tendency (e.g. median, mean) and variation (e.g. standard deviation, interquartile range)
- Clearly defined error bars

See the web collection on [statistics for biologists](#) for further resources and guidance.

► Software

Policy information about [availability of computer code](#)

7. Software

Describe the software used to analyze the data in this study.

The principal software used for this study is Idsc, developed in our lab. It is publicly available and the link is referenced in the paper. We also put online a tutorial explaining how to use Idsc with continuous annotation (<https://github.com/bulik/ldsc/wiki/Partitioned-Heritability-from-Continuous-Annotations>).

For all studies, we encourage code deposition in a community repository (e.g. GitHub). Authors must make computer code available to editors and reviewers upon request. The *Nature Methods* [guidance for providing algorithms and software for publication](#) may be useful for any submission.

► Materials and reagents

Policy information about [availability of materials](#)

8. Materials availability

Indicate whether there are restrictions on availability of unique materials or if these materials are only available for distribution by a for-profit company.

23andMe Inc summary statistics used in this study are not publicly available. However, all the results derived from these data have no restriction.

9. Antibodies

Describe the antibodies used and how they were validated for use in the system under study (i.e. assay and species).

No antibodies used.

10. Eukaryotic cell lines

a. State the source of each eukaryotic cell line used.

No eukaryotic cell line used.

b. Describe the method of cell line authentication used.

No eukaryotic cell line used.

c. Report whether the cell lines were tested for mycoplasma contamination.

No eukaryotic cell line used.

d. If any of the cell lines used in the paper are listed in the database of commonly misidentified cell lines maintained by [ICLAC](#), provide a scientific rationale for their use.

No eukaryotic cell line used.

► Animals and human research participants

Policy information about [studies involving animals](#); when reporting animal research, follow the [ARRIVE guidelines](#)

11. Description of research animals

Provide details on animals and/or animal-derived materials used in the study.

No animal used.

Policy information about [studies involving human research participants](#)

12. Description of human research participants

Describe the covariate-relevant population characteristics of the human research participants.

This study used human genetic data generated by different sources. We did not directly perform research on human participants.