Pervasive horizontal pleiotropy in human genetic variation is driven by extreme polygenicity of human traits and diseases

Daniel M. Jordan, Ph.D.[1,2,3*], Marie Verbanck, Ph.D.[1,2,3*], Ron Do, Ph.D.[1,2,3]

[1]The Charles Bronfman Institute for Personalized Medicine, Icahn School of Medicine at Mount Sinai, 1468 Madison Avenue, New York, NY, USA

[2]The Icahn Institute for Genomics and Multiscale Biology, Icahn School of Medicine at Mount Sinai, 1425 Madison Ave, New York, NY, USA

[3]Department of Genetics and Genomic Sciences, Icahn School of Medicine at Mount Sinai, 1425 Madison Avenue, New York, NY, USA

*Dr. Jordan and Dr. Verbanck contributed equally to this work.

Addresses for correspondence:

Ron Do, Ph.D.

The Charles Bronfman Institute for Personalized Medicine

Department of Genetics and Genomic Sciences

Icahn School of Medicine at Mount Sinai

One Gustave L. Levy Pl., Box 1003

New York, NY 10029

Email: ron.do@mssm.edu

Tel: 212-241-6206

Abstract

Horizontal pleiotropy, where one variant has independent effects on multiple traits, is important for our understanding of the genetic architecture of human phenotypes. We developed a method to quantify horizontal pleiotropy using genome-wide association summary statistics and applied it to 372 heritable phenotypes measured in 361,194 UK Biobank individuals. We observed horizontal pleiotropy is: 1) pervasive throughout the human genome; 2) especially prominent among highly polygenic phenotypes; 3) detected in 24,968 variants in 7,831 loci; and 4) enriched in active regulatory regions. Our results highlight the central role horizontal pleiotropy plays in the genetic architecture of human phenotypes.

Keywords

Pleiotropy, Polygenicity, Genetic Architecture, GWAS, Statistical Method, R Package

Background

The term “pleiotropy” refers to a single genetic variant having multiple distinct phenotypic effects. In general terms, the existence and extent of pleiotropy has far-reaching implications on our understanding of how genotypes map to phenotypes (1), of the genetic architectures of traits (2,3), of the biology underlying common diseases (4) and of the dynamics of natural selection (5). However, beyond this general idea of the importance of pleiotropy, it quickly becomes difficult to discuss in specifics, because of the difficulty in defining what counts as a direct causal effect and what counts as a separate phenotypic effect.

One particularly important dividing line in these conflicting definitions is the distinction between vertical pleiotropy and horizontal pleiotropy (6,7). When a genetic variant has a phenotypic effect that then has its own downstream effects in turn, that variant exhibits “vertical” pleiotropy.

For example, a variant that increases low density lipoprotein (LDL) cholesterol might also have an additional corresponding effect on coronary artery disease risk due to the causal relationship between these two traits, thus exhibiting vertical pleiotropy. Vertical pleiotropy has been conceptualized and measured by explicit genetic methods like Mendelian randomization.

In contrast, a genetic cause that directly influences multiple traits, without one trait being mediated by another, would exhibit “horizontal” pleiotropy. Horizontal pleiotropy contains some conceptual difficulties, and consequently can be difficult to measure. In principle, we might imagine selecting a variant and counting how many phenotypes are associated with it. Indeed, several versions of this analysis have been performed for different lists of traits (8,2,3,9). However, the results of these analyses are highly dependent on the exact list of traits used, and traits of interest to researchers previously tend to involve only a small number of phenotypes and/or be heavily biased towards a small set of disease-relevant biological systems and processes. Due to these limitations, it is unknown to what extent horizontal pleiotropy affects genetic variation in the human genome at the genome-wide level.

The proliferation of data sources like large-scale biobanks and metabolomics data that include a wide array of phenotypes in one dataset, combined with the growing public availability of genome-wide association studies (GWASs) summary statistic data, especially for extremely large meta-analyses, has allowed the development of methods that use these summary statistics to gain insight into human biology, and particularly into the genetic architecture of complex traits and diseases (10).

Here, we present a method to measure horizontal pleiotropy using publicly available GWAS summary statistics. We focus on measuring horizontal pleiotropy of SNVs on observable traits, meaning a scenario where a single SNV affects multiple phenotypes without relying on a detectable causal relationship between those phenotypes. Using this framework, we are able to

score each SNV in the human genome for horizontal pleiotropy, giving us broad insight into the genetic architecture of pleiotropy. Because our framework explicitly removes correlations between the input phenotypes, and because these phenotypes represent a diverse array of traits and diseases, these insights are largely robust to the specific list of traits studied, and pertain to human biology overall rather than relationships between specific traits.

Results

Defining pleiotropy

We narrowly define the scope of pleiotropy as applying only to genetic variants, and particularly variants investigated as part of GWASs. As effects, we are considering phenotypic outcomes measured by GWASs. By our definition, then, pleiotropy means that one variant shows significant associations across GWASs of multiple traits. We additionally restrict the scope of pleiotropy we are considering to include only horizontal pleiotropy, and to exclude vertical pleiotropy (**Figure 1**). To elaborate on this distinction, suppose we have identified a variant that influences two different traits, trait A and trait B. In vertical pleiotropy, the traits themselves are biologically related, so that the variant's effect on trait A actually causes the effect on trait B. A key feature of vertical pleiotropy is that two traits that are biologically related should be related regardless of which specific gene or variant is causing the effect. This induces correlation between GWAS effect sizes on the two traits across an entire set of variants. For example, we expect that *any* variant that increases LDL cholesterol also increases risk of coronary artery disease, because we suspect that it is the increase in LDL cholesterol itself that causes increased disease risk. This results in a correlation between variant effect sizes for LDL cholesterol and coronary artery disease, which has been detected in multiple studies (11–13). The methodology of Mendelian Randomization uses this predicted correlation within a given set

of variants to formulate a statistical test for causal relationships among traits, which is now widely used for biological discovery (14,15). We extend this methodology to use the entire set of SNVs evaluated by GWAS, treating a GWAS-wide correlation between two traits as evidence of a vertical pleiotropic relationship between these traits.

In the case of horizontal pleiotropy, an individual variant acts on traits A and B without mirroring any trait-level relationship between them. Unlike vertical pleiotropy, since we are not considering the variant-level effect as evidence of a relationship between the two traits, we cannot detect horizontal pleiotropy by detecting correlations between traits. Instead, each horizontally pleiotropic variant acts by its own unique mechanism. These particular pleiotropic variants, therefore, should show a relationship between the two traits that deviates from the relationship we would infer from the genome-wide correlation of effect sizes between them. This deviation from the correlation between traits is not a prediction of any kind of model of pleiotropy, but simply follows from our definition of the term “horizontal pleiotropy”: any pair of traits whose effect sizes are correlated across all variants is by definition related by vertical pleiotropy, while any variant whose effects on two traits substantially deviate from the trait-level relationship between those traits is by definition exhibiting horizontal pleiotropy.

A quantitative score for pleiotropy

We have developed a method to measure horizontal pleiotropy using summary statistics data from GWASs on multiple traits. Our method relies on applying a statistical whitening procedure to a set of input variant-trait associations, which removes correlations between traits caused by vertical pleiotropy and normalizes effect sizes across all traits. Using the decorrelated association Z-scores, we measure two related but distinct components of pleiotropy: the total magnitude of effect on whitened traits (“magnitude” score, denoted P_m) and the total number of whitened traits affected by a variant (“number of traits” score, denoted P_n). Both scores are then

scaled by the number of traits and multiplied by 100, so that the final score represents the value as it would be measured in a dataset of 100 traits. This two-component quantitative pleiotropy score allows us to measure both the magnitude (pleiotropy magnitude score P_m) and quantity (pleiotropy number of traits score P_n) of horizontal pleiotropy for all SNVs in the human genome. In principle these are distinct quantities: the magnitude score P_m measures the total pleiotropic effect size of a variant across all traits, while the number of traits score P_n measures the number of distinct pleiotropic effects a variant has. A variant with a high P_m score and a low P_n score has a large effect spread over a small number of traits; a variant with a low P_m score and a high P_n score has only a minor effect overall, but that effect is spread out across a large number of traits; and a variant with high scores on both components has a large effect that is spread across a large number of traits. Since we expect these scores to be heavily influenced by linkage disequilibrium (LD), we regress P_m and P_n against LD scores to produce an LD-corrected score (P_m^{LD} and P_m^{LD}) (**Figures 2, 3**; **Methods**).

Calculating significance of pleiotropy

We compute *P-values* for the two components of our pleiotropy score using two different procedures, corresponding to two different null expectations.

1. Theoretical *P*-values (Raw pleiotropy score [P_m and P_n] or LD-corrected pleiotropy score [P_m^{LD} and P_n^{LD}]), calculated analogously to *P*-values for genetic association studies including GWAS, based on a null scenario where variants do not exhibit pleiotropic effects on observed traits.
2. Empirical *P*-values (Polygenicity/LD-corrected pleiotropy score [P_m^P and P_n^P]), calculated by permutation of the observed distributions of whitened traits. These *P*-values are based on a null scenario where variants may have significant effects on one or more

traits, but the effects of each variant on each trait are independent and the number of variants with effects on multiple traits is no more than would be expected by chance.

This empirical correction for polygenicity is required because polygenicity is a major factor that can produce pleiotropy. For example, it has been estimated that approximately 100,000 independent loci are causal for height in humans (16). If the total number of independent loci in the human genome is approximately 1 million, this corresponds to about 10% of the human genome having an effect on height. If we imagine multiple phenotypes with this same highly polygenic genetic architecture, we should expect substantial overlap between causal loci for multiple different traits, even in the absence of any true causal relationship between the traits, resulting in horizontal pleiotropy **(Figure 2**).

Power to detect pleiotropy in simulations

We conducted a simulation study to evaluate the performance of our two-component pleiotropy score. We simulated 800,000 variants controlling 100 traits, varying the per-trait liability scale heritability of all traits h^2 and the proportion of pleiotropic and non-pleiotropic causal variants. To introduce LD in the simulations, we used real LD architecture from 800000 SNVs from 1000 Genomes European population. We simulated Z-scores independently for each SNV and then propagate LD for a given SNV by "contaminating" its Z-score according to the Z-scores of the SNVs in LD with it. Under the null model, all trait-variant associations were independent, and no horizontal pleiotropy was added. Under the added-pleiotropy models, we randomly chose a fraction of causal variants and forced them to have simultaneous associations with multiple traits. The simulation study showed that both components of the pleiotropy score were well-powered to detect horizontal pleiotropy (**Figure 4**), and that the LD correction dramatically reduces the dependence of the pleiotropy score on LD (**Supplementary Figure 1**). Under the null hypothesis of no added horizontal pleiotropy, the false positive rate was well controlled for

both scores when there was low heritability or few causal variants. However, when there are many causal variants and high per-variant heritability, the LD-corrected pleiotropy score (P_m^{LD} and P_n^{LD}) detects a large excess of pleiotropic variants, due to serendipitous overlap between causal variants without explicitly induced pleiotropy. The LD/polygenicity-corrected empirical *P*-value (P_m^P and P_n^P) does not detect this serendipitous pleiotropy at the same high rate.

In the presence of added horizontal pleiotropy, our approach was powered to detect pleiotropy with per-variant heritability h^2 as small as 0.002 if there are no non-pleiotropic causal variants. In the presence of both pleiotropic and non-pleiotropic causal variants, detecting pleiotropy was more difficult, but our approach still had appreciable power to detect pleiotropic variants, which increased with increasing per-variant heritability and decreased with increasing numbers of non-pleiotropic causal variants. Adding the correction for polygenic architecture (P_m^P and P_n^P) reduced this power only slightly. The power of our method was not substantially reduced by increasing the number of traits affected by pleiotropic variants (**Supplementary Figure 2**) or by adding a realistic correlation structure between the traits (**Supplementary Figure 3**).

Genome-Wide Pleiotropy Study (GWPS) reveals pervasive pleiotropy

To apply our method to real human association data, we used GWAS association statistics for 372 heritable medical traits measured in 337,119 individuals from the UK Biobank (17–19) . We successfully computed our two-component pleiotropy score for 767,057 variants genome-wide and conducted a genome-wide pleiotropy study (GWPS), by analogy to a standard GWAS (**Figure 3**; **Methods**). **Supplementary Figure 4** shows the resulting quantile-quantile plots (Q-Q plots). We observed significant inflation for both the LD-corrected magnitude score P_m^{LD} and number of traits score P_n^{LD} (Mann-Whitney U test $P < 10^{-300}$ for both). Furthermore, we

observed across both scores that horizontal pleiotropy was widely distributed across the genome, rather than being localized to a few specific loci (**Supplementary Figure 5**). Testing an alternative strategy for computing the phenotype-correlation matrix using all SNVs produced comparable results (Pearson r = 0.995 and 0.964 for P_m^{LD} and P_n^{LD} respectively) to our strategy of using a pruned set of SNVs to account for LD (r^2 < 0.1) (**Supplementary Figure 6**).

Pleiotropy is driven by polygenicity

We applied the permutation-based empirical *P*-value calculation (Polygenicity/LD-corrected pleiotropy score: P_m^P and P_n^P) to correct for the known polygenic architecture of traits and test whether any loci are pleiotropic to a greater extent than would be expected due to polygenicity. **Supplementary Figures 7 and 8** show the resulting Q-Q plots and Manhattan plots. In contrast to the results from the LD-corrected pleiotropy score (P_m^{LD} and P_n^{LD}), we do not find pleiotropy significantly in excess of what would be expected from the known polygenic architecture of traits: there are dramatically fewer loci with genome-wide significant levels of pleiotropy after correcting for polygenic architecture, and the genome-wide distribution of pleiotropy score shows less pleiotropy than expected (Mann-Whitney U test $P < 10^{-300}$ for both P_m^P and P_n^P).

As an additional test of whether the pleiotropy we observe is driven by polygenicity, we calculated the polygenicity of the same 372 heritable traits from the UK Biobank. We measured polygenicity using a version of the genomic inflation factor corrected using LD score λ_{GC}^c (20). We then stratified these traits by λ_{GC}^c after controlling for heritability (**Methods**), and calculated the two-component LD-corrected pleiotropy score [P_m^{LD} and P_n^{LD}]) and *P-values* for each component independently for every variant in the genome using each of these bins of traits. We observed that both scores are highly dependent on polygenicity, with the lowest-polygenicity bins in each heritability class showing very little inflation. (**Figure 5**; **Supplementary Table 1**).

Taken together, these results suggest that extreme polygenicity drives horizontal pleiotropy, and that this has an extremely large effect on the genetic architecture of human phenotypes.

Genome-wide distribution of pleiotropy score gives insight into genetic architecture

In addition to observing genome-wide inflation of the pleiotropy score, we can also gain insight from the distribution of the pleiotropy score on a more granular level.

Figure 6a shows the distribution of pleiotropy score for independent SNVs (LD pruned to a threshold of $r^2 < 0.1$) compared to the expectation under the null hypothesis of no pleiotropic effect. We observe a large excess in the number of traits score P_n^{LD}, and a smaller but still highly significant excess in total magnitude of pleiotropic effect P_m^{LD}. This excess comes in part from a long tail of highly pleiotropic loci that pass the threshold of genome-wide significance (dashed line in **Figure 6a**), but is primarily driven by weak pleiotropy among loci that do not reach genome-wide significance.

Pleiotropy score is correlated with molecular and biological function

To further investigate the properties of pleiotropic variants, we examined the effects of various functional and biochemical annotations on our LD-corrected pleiotropy score (P_m^{LD} and P_n^{LD}) **(Table 1; Methods)**. Using annotations from Ensembl Variant Effect Predictor (21), we observed that both components of the pleiotropy score are higher on average in transcribed regions (coding and UTR) than in intergenic noncoding regions. This result was confirmed and expanded by annotations from Roadmap Epigenomics (22), which showed that regions whose chromatin configurations were associated with actively transcribed regions, promoters, enhancers, and transcription factor binding sites had significantly higher levels of both components of the pleiotropy score, while heterochromatin and quiescent chromatin states had

significantly lower levels. Investigating individual histone marks, we found that both the repressive histone mark H3K27me3 and the activating histone mark H3K27ac were associated with elevated levels of pleiotropy, although the activating mark H3K27ac had a larger effect. This may indicate that being under active regulation at all produces higher levels of pleiotropy, whether that regulation is repressive or activating.

We also used data from the Genotype-Tissue Expression (23) project to measure the connection between transcriptional effects and our pleiotropy score **(Table 1)**. Consistent with the previous observation that functional regions had higher pleiotropy scores, we found that variants that were identified as *cis*-eQTLs for any gene in any tissue had higher pleiotropy scores on average. Within eQTLs, we also observed significant correlations between our pleiotropy score and the numbers of genes (P_m^{LD}: r = 0.036, $P < 2.2 \times 10^{-16}$; P_n^{LD}: r = 0.035, $P < 2.2 \times 10^{-16}$) and tissues (P_m^{LD}: r = 0.062, $P < 2.2 \times 10^{-16}$; P_n^{LD}: r = 0.059, $P < 2.2 \times 10^{-16}$) where the variant was annotated as an eQTL, showing that our pleiotropy score is related to transcriptional measures of pleiotropy.

Finally, we found that variants that are eQTLs for genes whose orthologs are associated with multiple measurable phenotypes in mice or yeast have higher pleiotropy scores, demonstrating that our pleiotropy score is also related to pleiotropy in model organisms.

All these results are consistent when using the Polygenicity/LD-corrected pleiotropy score (P_m^P and P_n^P), indicating that the association of pleiotropy with molecular and biological function is not exclusively driven by highly polygenic architecture (**Additional File 1**).

Genome-wide pleiotropy study identifies novel biological loci

By analogy to standard GWAS, our GWPS methodology can identify individual variants that have a genome-wide significant level of horizontal pleiotropy. Using the LD-corrected magnitude

score P_m^{LD}, we identified 74,335 variants in 8,093 independent loci with a genome-wide significant level of horizontal pleiotropy, while using the LD-corrected number of traits score P_n^{LD} identified 18,393 variants in 2,859 independent loci with a genome-wide significant level of horizontal pleiotropy, all of which are also identified by the LD-corrected magnitude score P_m^{LD} (**Methods**, **Supplementary Table 2**). Applying the same analysis to the Polygenicity/LD-corrected pleiotropy score, using the Polygenicity/LD-corrected magnitude score P_m^P identified no genome-wide significant loci, but using the Polygenicity/LD-corrected number of traits score P_n^P identified 2,674 variants in 432 loci. Strikingly, a majority of loci significant in P_n^{LD} (1,519 of 2,859) or P_n^P (294 of 432), along with a sizeable minority of loci significant in P_m^{LD} (2,934 of 8,093), have no entry in the NHGRI-EBI GWAS catalog, meaning that they have never been reported as an associated locus in any published GWAS. These loci represent an under-recognized class of genetic variation that has multiple weak to intermediate effects that are collectively significant, but no specific strong effect on any one particular trait. Functional enrichment analysis on genes near these genome-wide significant loci implicates a wide range of biological functions, including cell adhesion, post-translational modification of proteins, cytoskeleton, transcription factors, and intracellular signaling cascades (**Additional File 2**). Loci significant in P_n^P show a more focused subset of functions, with a greater role for nuclear proteins regulating transcription and chromatin state, suggesting that these are the functions that exhibit horizontal pleiotropy beyond the baseline level induced by polygenicity. The role of these novel loci and these biological processes in human genetics and biology may be a fruitful area for future study, with the potential for biological discovery.

Pleiotropic loci replicate in independent GWAS datasets

As replication datasets, we used two additional sources of GWAS summary statistics to calculate our LD-corrected pleiotropy score (P_m^{LD} and P_n^{LD}): previously published GWASs and

meta-analyses for 73 human complex traits and diseases, which we collected and curated manually from the literature (**Methods, Supplementary Table 3**) (24); and a previously published study of 430 blood metabolites measured in 7,824 European adults (25). For all variants covered by the UK Biobank, we were able to compute our pleiotropy score independently using these two datasets (**Figure 7**). In the traits and diseases dataset, we observed that 57% of P_m^{LD} loci and 38% of P_n^{LD} loci replicated, while in the blood metabolites dataset, we observed that 17% of P_m^{LD} loci and 12% of P_n^{LD} loci replicated, compared to 5% of P_m^{LD} loci and 6% of P_n^{LD} loci expected by chance according to a permutation-based null model. This high level of replication using independent sets of GWAS summary statistics suggests that our pleiotropy score is capturing an underlying biological property, rather than an artifact of the UK Biobank study.

Pleiotropy is correlated with specific complex traits and diseases

To characterize the phenotypic associations of these loci, we used our replication dataset of published GWAS summary statistics for 73 human quantitative traits and diseases, plus nine additional traits we excluded from our replication dataset for a total of 82 (**Methods**). We are not able to compute directly the degree of pleiotropy exhibited by these traits, since our definition of horizontal pleiotropy applies only to individual variants and does not apply to traits. However, we can identify traits whose GWAS variant associations are correlated to our pleiotropy score, which in some sense represents the traits that contribute most to our signal of pervasive horizontal pleiotropy. **Figure 6c** shows the correlations between our LD-corrected pleiotropy score (P_m^{LD} and P_n^{LD}) and the association statistics for these 82 traits and diseases. The most strongly correlated traits were anthropometric traits like body mass index, waist and hip circumference, and height; certain blood lipid levels, including total cholesterol and triglycerides; and schizophrenia. These are all known to be highly polygenic and heterogeneous traits. The least correlated traits include several measurements of insulin sensitivity and glucose response,

304 such as the insulin sensitivity index (ISI), certain features of brain morphology, and the
305 inflammatory biomarker lipoprotein(a). This may be partly due to low sample size of the
306 corresponding GWASs. However, these correlations do not appear to be driven exclusively by
307 sample size: in cases where multiple GWASs for the same trait have been performed on
308 subsamples of the population (for example, males only, female only, and combined), the sample
309 size only marginally affects the correlation (**Supplementary Table 4**). Another contributing
310 factor may be heritability: height, in particular, is among the most heritable traits we examined,
311 while ISI and the brain morphology features are among the least.

312 # Discussion

313 We have presented a framework for scoring horizontal pleiotropy across human genetic
314 variation. In contrast to previous analyses, our framework explicitly distinguishes between
315 horizontal pleiotropy and vertical pleiotropy or biological causation. After applying both
316 components of our pleiotropy score to 372 heritable medical traits from the UK Biobank, we
317 made the following observations: 1) horizontal pleiotropy is pervasive and widely distributed
318 across the genome; 2)) horizontal pleiotropy is driven by extreme polygenicity of traits; 3)
319 horizontal pleiotropy is significantly enriched in actively transcribed regions and active regulatory
320 regions, and is correlated with the number of genes and tissues for which the variant is an
321 eQTL; 4) there are thousands of loci that exhibit extreme levels of horizontal pleiotropy, a
322 majority of which have no previously reported associations; and 5) pleiotropic loci are enriched
323 in specific complex traits including body mass index, height, and schizophrenia. These findings
324 are largely consistent between the magnitude of pleiotropy score P_m and the number of traits
325 score P_n, although we note some differences where some variants are primarily associated with
326 P_m^{LD} but not P_n^{LD}. This indicates that these signals are driven by loci that both influence a large
327 number of traits and have relatively large combined effects, and secondarily by loci that have

large combined effects but only influence a handful of traits each, with minimal contribution from loci that influence a large number of traits but have small combined effects. Conversely, after applying the correction for polygenicity, we only observe variants that are significant for P_n^P, but not for P_m^P. This indicates that, while there do exist horizontal pleiotropic master control loci that affect more traits than we would expect from the random overlap of multiple highly polygenic traits, the overall effect of these loci is not noticeably larger than we would expect.

This analysis is enabled by the technique of whitening trait associations to remove correlations between traits. This lets us count pleiotropic effects in a more objective and systematic way, as opposed to manually selecting putatively independent traits to count, or manually grouping traits into independent blocks. However, it does come with three major limitations compared to these approaches. First, it is somewhat more difficult to tell which specific traits are driving a signal of pleiotropy at a particular locus. Our whitened traits are combinations of real observed traits, and do not necessarily correspond to any specific biological traits of interest. However, it is relatively easy to inspect the input GWAS summary statistics for a particular variant of interest to see which traits it is associated with. Furthermore, since pleiotropic loci are by definition associated with a large cross-section of traits, this kind of inspection is not likely to be very informative about specific traits. Second, the whitening procedure has the counterintuitive property that a variant that has a narrow effect on a single trait without also affecting correlated traits can appear to be highly pleiotropic. For example, if a variant had a strong risk-increasing effect on coronary artery disease (CAD), but no effect on any of the known upstream risk factors of CAD (such as blood lipid levels or adiposity) or any of the known downstream consequences of CAD (such as inflammatory biomarkers or increased mortality), such a variant would appear as highly pleiotropic in our analysis. Our analysis would interpret the variant as increasing the risk of CAD while suppressing these upstream and downstream factors. We believe this treatment is appropriate, however counterintuitive. Regardless, these kinds of isolated effects are fairly rare:

in our dataset of 372 heritable traits from UK Biobank, only 6% of variants (42,684 of 767,057) reach genome-wide significance for only a single trait. Indeed, it is unlikely by definition that a variant is associated with only one trait from a set of correlated traits, since we compute our correlations from observed association statistics. Third, we assume all genetic effects are additive and independent, and we do not model epistasis or other more complex genetic architectures.

Our findings are in keeping with several recent studies that have found abundant pleiotropy in the genome (26,27,8,2,9). Our pleiotropy score goes a step further than many of these studies by explicitly removing vertical pleiotropy between traits, which are indicative of fundamental biological relationships between traits (8,24,28). Furthermore, the current study has evaluated horizontal pleiotropy in human genetic variation genome-wide, whereas previous studies have focused on only a small subset of disease-associated variants identified from GWAS. Our results therefore suggest that there is substantial complexity and heterogeneity not only in causal relationships between human traits, but also in the genetic architecture of individual traits.

Our findings have several important implications for the field of human genetics. First, our observation of ubiquitous horizontal pleiotropy is problematic for Mendelian Randomization (MR) methods, which assumes horizontal pleiotropy to be absent. Recent developments in the field of MR include methods that account for horizontal pleiotropy explicitly (24,28,29); our results reinforce the importance of these methods. The presence of widespread horizontal pleiotropy suggests that single-instrument methods that independently account for every variant, each of which presumably has pleiotropic effects on many different distinct traits, should be considered in addition to multi-instrument methods for MR, which collapse many variants into a single polygenic score for analysis, and therefore treat all variants equivalently.

Second, our results appear to support the "network pleiotropy" hypothesis of Boyle, Li, and Pritchard (16), which proposes widespread pleiotropy driven by small perturbations of densely connected functional networks, where any perturbation in a relevant cell type will have at least a small effect on all phenotypes affected by that cell type. A subsequent paper detailed a more specific mechanism, where causal effects are driven by many biological components that are only indirectly related to the phenotype itself (30). Many of the functional enrichments we observe, including transcription factors, cytoskeleton, and intracellular signaling cascades, represent components that can plausibly influence a wide variety of cell types and processes, providing evidence for this model over one where a specific biological component is largely responsible for pleiotropy. The fact that the magnitude of pleiotropy score P_m and the number of traits score P_n give largely consistent results also supports this model, where a larger biological effect in a given tissue will perturb a greater number of phenotypes relevant to that tissue, although we note that some variants have high magnitude of pleiotropy score P_m and low number of traits score P_n, which may represent a small class of variants that has large biological effects without perturbing a large number of phenotypes.

While our results largely support this network pleiotropy hypothesis, we have also demonstrated an alternate view of horizontal pleiotropy in the context of highly polygenic causation. In our simulations, introducing extreme polygenicity at the levels suggested by these papers inherently results in high levels of horizontal pleiotropy detectable by our score, independent of any assumptions about the mechanism of pleiotropy or of polygenicity. Indeed, our null hypothesis of no horizontal pleiotropy, that 5% of the genome is independently causal to each trait with $P < 0.05$, is trivially rejected when a single trait is influenced by an unexpectedly large fraction of the genome. This means that, on some level, widespread horizontal pleiotropy in human genetic variation is simply a logical consequence of widespread polygenicity of human traits, regardless of the specific mechanism of either. In simple terms, the more loci are associated with each trait,

the more chances there are for associations with multiple traits to overlap. Supporting this result, we find that controlling for the polygenic architecture of the input traits significantly attenuates our signal of pleiotropy, as does restricting to oligogenic traits. It may be the case that horizontal pleiotropy is only truly widespread among the most complex and polygenic subset of human traits.

Conclusions

In this study, we have presented a quantitative score for horizontal pleiotropy in human genome variation. Using this score, we have identified a genome-wide trend of highly inflated levels of horizontal pleiotropy, an underappreciated relationship between horizontal pleiotropy with polygenicity and functional biology, and a large number of specific novel loci with high levels of horizontal pleiotropy. We expect further investigations using this score to yield deep insights into the genetic architecture of human traits and to uncover important novel biology.

Methods

We developed a statistical method to measure horizontal pleiotropy using a two-component pleiotropy score. For a given variant, we measured 1) the total magnitude of pleiotropic effect the variant has and 2) the number of whitened traits affected by the variant.

Z-scores decorrelation strategy

Observable traits and diseases can be highly correlated, which can lead to inflation of our pleiotropy score if the correlation is not properly accounted for. Therefore, we developed an efficient strategy to remove this correlation and obtain decorrelated traits. Let Z^{raw} denote the matrix of raw Z-scores, with variants in columns and traits in rows, and Σ denote the

corresponding correlation matrix between the Z-scores. Under the null hypothesis of no horizontal pleiotropy, Z-scores for each trait are assumed to follow a Gaussian distribution $N(0,1)$, and the columns of Z^{raw} collectively follow a multivariate Gaussian distribution $N(0, \Sigma)$. Our goal is to eliminate the extra-diagonal terms of the correlation matrix Σ. To achieve this, we use a Mahalanobis whitening transformation on the matrix Z^{raw} to obtain a whitened Z-score matrix Z. The procedure to obtain Z can be formally expressed as:

$$Z = \Sigma^{\frac{-1}{2}} Z^{raw}$$

Under the null hypothesis of no horizontal pleiotropy, we expect Z to follow a multivariate Gaussian distribution $N(0, Id_l)$, where Id_l is the identity matrix of size l, l being the number of traits.

In reality, the true correlation matrix Σ is unknown, and we must use an estimated correlation matrix $\hat{\Sigma}$ obtained by measuring the genome-wide correlation between actual Z-scores. We tested two approaches to obtain $\hat{\Sigma}$, either using all genotyped variants genome-wide or using a subset of variants pruned to $r^2 < 0.1$ in the 1000 Genomes European population to account for the effects of linkage disequilibrium (LD). Both approaches produced similar results (See **Supplementary Figure 6**). In all subsequent analysis, we used covariance matrices estimated from pruned variants.

Computation of the pleiotropy score

We computed two different scores to capture both the magnitude and number of traits of pleiotropy. First, we quantify the total pleiotropic magnitude of effect a variant using the magnitude pleiotropy score P_m:

$$P_m = \frac{100}{l}\sqrt{\sum_{1}^{l} {z_i}^2}$$

where z_i is the whitened Z-score for trait i for a given variant. Second, we quantify the number of whitened traits affected by a variant using the number of pleiotropic traits score P_n:

$$P_n = \frac{100}{l}\sum_{1}^{l} H(z_i - 2)$$

where z_i is the whitened Z-score for trait i for the tested variant and $H()$ is the Heaviside step function which equals 1 if $|z_i| > 2$ and 0 otherwise. 2 represents a standard value of the Z-score which represents the normal threshold for nominal significance ($P < 0.05$).

LD-corrected pleiotropy score

Similarly to LD score regression, each component of the pleiotropy score was regressed on the LD scores for all variants. Then, we regressed out the effect of LD on each component of the pleiotropy score independently to obtain an LD-corrected pleiotropy score. The LD-corrected pleiotropy score components P_m^{LD} and P_n^{LD} are given by:

$$P_m^{LD} = P_m - \beta_m l$$

$$P_n^{LD} = P_n - \beta_n l$$

where l is the LD score of the variant site, and β_m and β_n are the regression coefficients for LD score on P_m and P_n, respectively.

Computation of theoretical *P-values* for the pleiotropy score

Based on the observation that Z follows a multivariate standard Gaussian distribution $N(0, Id_l)$ under the null hypothesis of no pleiotropy, *P-values* can easily be computed for P_m and P_n.

Under the null hypothesis, the square of P_m (or P_m^{LD}) follows a chi-square distribution $\chi^2(l)$ where l is the total number of traits. Likewise, P_n (or P_n^{LD}) follows a binomial distribution $B(l, p)$ where l is the total number of traits and *p* the probability to get a Z-score greater than 2 under the standard Gaussian distribution ($P \approx 0.045$).

Computation of empirical (polygenicity/LD-corrected) *P-values* for the pleiotropy score

To correct for the known polygenic architecture of traits in addition to LD, we additionally computed empirical permutation-based *P*-values for both P_m^{LD} and P_n^{LD}. We performed 25 random permutations of the input Z-scores for each observable trait, producing millions of permuted variants. We calculated P_m and P_n for each of these permuted variants, and then rank ordered the resulting scores. The empirical *P*-value corresponding to a value of P_m^{LD} or P_n^{LD} is given by the fraction of permuted variants with higher scores than the given value. We converted these *P*-values into polygenicity/LD-corrected P_m^P and P_n^P scores by converting each *P*-value into the score it would correspond to under the expected (theoretical) distributions described above.

Simulation framework

We simulated a realistic matrix of Z-scores Z with 100 traits and 800,000 genotyped variants. For non-causal variants, Z-scores for each trait were drawn from an independent Gaussian distribution $N(0,1)$. A subset of variants was designated as causal, either pleiotropically or non-pleiotropically. For these causal variants, Z-scores were drawn from a different Gaussian distribution $N(0, h^2)$,where h^2 is a parameter representing the per-variant heritability of each trait. Non-pleiotropic variants were selected independently for each trait, while pleiotropic variants were selected globally and each forced to be causal for a specified number of traits ν.

Simulations were run for all combinations of the following parameters: 1) correlation structure: absent or present; 2) proportion of pleiotropic causal variants: 0.1% (800/800,000 variants) or 1% (8,000/800,000 variants); 3) proportion of non-pleiotropic causal variants: 0 (0/800,000 variants), 0.1% (800/800,000 variants), or 1% (8,000/800,000 variants); 4) number of traits involved in horizontal pleiotropy ν: 10 or 20; 5) per-variant heritability of traits h^2: 0.0002, 0.002, 0.02, or 0.2. Each scenario was replicated 10,000 times.

Collection of genome-wide association (GWA) summary statistics datasets

First, we retrieved GWA publicly available summary statistics from 545 continuous traits in 361,194 samples from the UK Biobank (17), and 1,403 binary traits from the same dataset calculated using SAIGE (18,19). We used LD score regression to calculate heritability for each trait, using the liability scale for binary traits, and restricted the sample to only traits with a significant *P*-value for nonzero heritability after Bonferroni correction. For every pair of traits with correlation coefficient between Z-scores $r^2 > 0.8$, we additionally removed the member of the pair with lower heritability. This left a total of 372 traits.

Second, we retrieved publicly available genome-wide association (GWA) summary statistics data for 82 complex traits and diseases (31–66) (**Table S9**). For each dataset, we retrieved the appropriate variant annotation (build, rsid, chromosome, position, reference and alternate alleles) and summary statistics (effect size, standard errors, *P-values* and sample size of the study). All variant coordinates (chr, pos) were lifted over to hg19 using the UCSC Genome Browser LiftOver Tool and aligned to the reference and alternate alleles retrieved from the Ensembl variation database. After performing the same pruning of highly correlated phenotypes, we were left with 73 traits and diseases.

Third, we retrieved GWA summary statistics datasets from a GWAS of 453 blood metabolites in 7,824 individuals (67). After performing the same pruning of highly correlated phenotypes, we were left with 430 metabolites.

Genome-wide pleiotropy study (GWPS)

Using the two components of the pleiotropy score, we can run a genome-wide pleiotropy study (GWPS) which consists of computing two *P-values* for each component of the score (P_m^{LD} and P_n^{LD}) and for all variants genome-wide. We computed the pleiotropy score separately for each of the three datasets described above (372 UK Biobank phenotypes, 73 traits and diseases, and 430 blood metabolites). Additionally, we computed the pleiotropy score on a subset of 372 traits with genome-wide significant heritability as calculated by LD Score Regression (20) (univariate heritability significant after Bonferroni correction). The 372 UK Biobank heritable traits were used for discovery, and the 73 traits and diseases and 430 blood metabolites datasets were used for replication. There was a total of 768,756 variants genotyped across all three datasets.

Study of polygenicity on horizontal pleiotropy

To study the effect of polygenicity on horizontal pleiotropy, we first estimated the liability-scale heritability of all 372 traits in our UK Biobank dataset using LD score regression, and stratified all traits into four equally-sized classes of heritability, in order to control for the effect of high heritability separate from the effect of high polygenicity. Next, we estimated the polygenicity of the 372 traits using a corrected version of the genomic inflation factor λ_{CG}^{c} (20). The intercept of LD score regression minus one is an estimator of the mean contribution of confounding bias to the inflation in the test statistics. Therefore, we computed a corrected version of the genomic inflation factor by subtracting the quantity (intercept of LD score regression - 1) from λ_{GC}. The 372 phenotypes were then ranked according to λ_{CG}^{c} within each heritability class, and grouped into 5 equal-sized bins of about 20 phenotypes each. We then recomputed the LD-corrected

pleiotropy score components (P_m^{LD} and P_n^{LD}) for the subset of phenotypes in each bin. The inflation of the pleiotropy score was measured per bin to represent pleiotropy score inflation as a function of polygenicity.

Characterization of the pleiotropic variants

We performed various enrichment analyses for the pleiotropy score to characterize the pleiotropic variants using a variety of annotations that could be a direct consequence of horizontal pleiotropy. Each analysis uses the principle of assigning each variant an annotation category and selecting one category as the reference category. Then, for each category, we selected a set of variants from the corresponding reference category with minor allele frequencies matched to those in the query category, and performed a Student's t-test to test whether the average LD-corrected pleiotropy score (P_m^{LD} and P_n^{LD}) of the variants in each given category is different from the average LD-corrected pleiotropy score of the selected reference variants.

First, we used Ensembl Variant Effect Predictor (21) to classify each variant as noncoding, UTR, nonsynonymous, or coding synonymous, treating noncoding as the reference class. These were complemented by annotations from Roadmap Epigenomics (22). We used the 25-state chromatin state model published by Roadmap Epigenomics to assign each variant 25 scores from 0 to 127, where each score represents the number of epigenomes for which that site is assigned to the corresponding category. We did the same for two specific chromatin marks: the activating mark H3K27ac and the repressive mark H3K27me3. For these annotations, we used a combination of all other categories as a reference set. In other words, the reference set for each category is all variants that are not in that category.

In addition to these molecular annotations, we used expression-related annotations from the Genotype-Tissue Expression project (23). For each variant, we retrieved the number of genes

552 for which the variant is referenced as a *cis* eQTL (expression quantitative trait loci) in any tissue
553 (eGenes), and the number of tissues where the variant is annotated as a *cis* eQTL for any gene
554 (eTissues). We divided variants into bins by number of eGenes (below 10, between 10 and 15,
555 between 15 and 20, and over 20) and eTissues (below 30, between 30 and 35, between 35 and
556 40, and above 40). The reference set used for these analyses were variants that are not
557 annotated as eQTLs in any gene or tissue.

558 Finally, we used model organism phenotypes measured by the International Mouse
559 Phenotyping Consortium (IMPC) (68) and the Saccharomyces Cerevisiae Morphological
560 Database (SCMD) (69). To map ortholog genes from IMPC and SCMD to human variants, we
561 used orthology annotations of gene orthologs, and eQTLs from GTEx. Thus, variants annotated
562 as associated with a mouse or yeast phenotype are those that are annotated as *cis* eQTLs in
563 any tissue for a gene whose ortholog in mouse or yeast is associated with that phenotype. The
564 reference set for this analysis was variants annotated as *cis* eQTLs for genes that are not
565 associated with mouse or yeast phenotypes.

566 Genome-wide significant pleiotropy loci

567 To detect loci with a genome-wide significant pleiotropy, we used the LD-corrected two-
568 component pleiotropy score (P_m^{LD} and P_n^{LD}) computed on the dataset of 372 heritable traits from
569 UK Biobank described above. We used LD clumping as implemented in PLINK to cluster linked
570 variants, with an r^2 threshold of 0.1, a distance threshold of 100 kb, and *P-value* thresholds of 5
571 x 10^{-8} for genome-wide significance and 0.05 for nominal significance. The resulting loci were
572 assigned to genes using 1) localization of variants within a gene, as annotated by Ensembl
573 Variant Effect predictor, and 2) annotation as a *cis* eQTL in any tissue, as annotated by GTEx.
574 We submitted the resulting list of genome-wide significant genes to DAVID for enrichment
575 analysis (70–72).

Permutation-based null model for replication analysis

In general, we should expect only 5% of loci to replicate by chance in each replication dataset; however, it is possible that this number might increase because of polygenicity in the underlying GWAS statistics and the resulting inflation in our pleiotropy score, which may cause substantially more than 5% of the genome to be assigned $P < 0.05$. To correct for this, we performed random permutations of the whitened Z-scores independently for each trait and used these permuted Z-scores to compute our LD-corrected pleiotropy score components (P_m^{LD} and P_n^{LD}). This generates a null expectation that preserves the polygenicity and inflation within each dataset. For both datasets, our null model expected that 5% of loci for P_m^{LD} loci and 6% of loci for P_n^{LD} should replicate. The fraction that replicated in the actual data was substantially higher (**Figure 7**).

Ethics approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Availability of data and material

An R package implementing this pleiotropy score method is available on GitHub at https://github.com/rondolab/PleiotropyScore. The dataset of summary statistics for the 372 medical traits from the UK Biobank and the pleiotropy scores computed from these summary statistics are also available in the same GitHub project at https://github.com/rondolab/PleiotropyScore/tree/master/data. The summary statistics for 430 blood metabolites are available from the original publication where this dataset was reported

599 (61), and the summary statistics for 73 human traits and diseases are available from the original
600 publications where they were reported, as cited in **Supplementary Table 3**.

601

602

603 Competing interests

604 R.D has received research support from AstraZeneca and Goldfinch Bio.

605 Funding

606 R.D is supported by R35GM124836 from the National Institute of General Medical Sciences of
607 the National Institutes of Health, R01HL139865 from the National Heart, Lung, Blood Institute of
608 the National Institutes of Health and previously an American Heart Association Cardiovascular
609 Genome-Phenome Discovery grant (15CVGPSD27130014). D.M.J. is supported by
610 T32HL00782 from the National Heart, Lung, and Blood Institute of the National Institutes of
611 Health. The content is solely the responsibility of the authors and does not necessarily represent
612 the official views of the National Institutes of Health.

613 Authors' contributions

614 D.M.J. and M.V. contributed to study conception, data analysis, interpretation of the results and
615 drafting of the manuscript. R.D. contributed to study conception, interpretation of the results and
616 critical revision of the manuscript.

617 Acknowledgments

We thank the various genome-wide association consortia as well as Dr. Benjamin Neale's and Dr. Seunggeun Lee's group for generously sharing genome-wide association summary statistics for phenome-wide association scan in UK biobank.

References

1. Zhan J, Arking DE, Bader JS. Discovering patterns of pleiotropy in genome-wide association studies. bioRxiv. 2018 Feb 28;273540.

2. Chesmore K, Bartlett J, Williams SM. The ubiquity of pleiotropy in human disease. Hum Genet. 2017 Nov 21;1–6.

3. Socrates A, Bond T, Karhunen V, Auvinen J, Rietveld C, Veijola J, et al. Polygenic risk scores applied to a single cohort reveal pleiotropy among hundreds of human phenotypes. bioRxiv. 2017 Oct 14;203257.

4. Solovieff N, Cotsapas C, Lee PH, Purcell SM, Smoller JW. Pleiotropy in complex traits: challenges and strategies. Nat Rev Genet. 2013 Jul;14(7):483–95.

5. Keightley PD, Hill WG. Variation maintained in quantitative traits with mutation–selection balance: pleiotropic side-effects on fitness traits. Proc R Soc Lond B. 1990 Nov 22;242(1304):95–100.

6. Paaby AB, Rockman MV. The many faces of pleiotropy. Trends in Genetics. 2013 Feb 1;29(2):66–73.

7. Tyler AL, Asselbergs FW, Williams SM, Moore JH. Shadows of complexity: what biological networks reveal about epistasis and pleiotropy. BioEssays. 2009 Feb 1;31(2):220–7.

8. Pickrell JK, Berisa T, Liu JZ, Ségurel L, Tung JY, Hinds DA. Detection and interpretation of shared genetic influences on 42 human traits. Nat Genet. 2016 Jul;48(7):709–17.

9. Kanai M, Akiyama M, Takahashi A, Matoba N, Momozawa Y, Ikeda M, et al. Genetic analysis of quantitative traits in the Japanese population links cell types to complex human diseases. Nat Genet. 2018 Mar;50(3):390–400.

10. Pasaniuc B, Price AL. Dissecting the genetics of complex traits using summary association statistics. Nature Reviews Genetics. 2017 Feb;18(2):117–27.

11. Do R, Willer CJ, Schmidt EM, Sengupta S, Gao C, Peloso GM, et al. Common variants associated with plasma triglycerides and risk for coronary artery disease. Nat Genet. 2013 Nov;45(11):1345–52.

12. Ference BA, Yoo W, Alesh I, Mahajan N, Mirowska KK, Mewada A, et al. Effect of long-term exposure to lower low-density lipoprotein cholesterol beginning early in life on the risk

of coronary heart disease: a Mendelian randomization analysis. J Am Coll Cardiol. 2012 Dec 25;60(25):2631–9.

13. Burgess S, Freitag DF, Khan H, Gorman DN, Thompson SG. Using multivariable Mendelian randomization to disentangle the causal effects of lipid fractions. PLoS ONE. 2014;9(10):e108891.

14. Burgess S, Bowden J, Fall T, Ingelsson E, Thompson SG. Sensitivity Analyses for Robust Causal Inference from Mendelian Randomization Analyses with Multiple Genetic Variants. Epidemiology. 2017;28(1):30–42.

15. Pierce BL, Ahsan H, Vanderweele TJ. Power and instrument strength requirements for Mendelian randomization studies using multiple genetic variants. Int J Epidemiol. 2011 Jun;40(3):740–52.

16. Boyle EA, Li YI, Pritchard JK. An Expanded View of Complex Traits: From Polygenic to Omnigenic. Cell. 2017 Jun 15;169(7):1177–86.

17. Abbott L, Bryant S, Churchhouse C, Ganna A, Howrigan D, Palmer D, et al. UK Biobank. [cited 2018 Aug 7]. Available from: http://www.nealelab.is/uk-biobank

18. Zhou W, Nielsen JB, Fritsche LG, Dey R, Gabrielsen ME, Wolford BN, et al. Efficiently controlling for case-control imbalance and sample relatedness in large-scale genetic association studies. Nature Genetics. 2018 Sep 1;50(9):1335–41.

19. Zhou W, Nielsen JB, Fritsche LG, Elvestad ME, Wolford BN, Lin M, et al. SAIGE [Internet]. [cited 2018 Sep 20]. Available from: https://github.com/weizhouUMICH/SAIGE#uk-biobank-gwas-results

20. Bulik-Sullivan BK, Loh P-R, Finucane HK, Ripke S, Yang J, Consortium SWG of the PG, et al. LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. Nature Genetics. 2015 Mar;47(3):291.

21. McLaren W, Gil L, Hunt SE, Riat HS, Ritchie GRS, Thormann A, et al. The Ensembl Variant Effect Predictor. Genome Biology. 2016 Jun 6;17(1):122.

22. Bernstein BE, Stamatoyannopoulos JA, Costello JF, Ren B, Milosavljevic A, Meissner A, et al. The NIH Roadmap Epigenomics Mapping Consortium. Nature Biotechnology. 2010 Oct 13;28:1045.

23. GTEx Consortium, Laboratory, Data Analysis &Coordinating Center (LDACC)—Analysis Working Group, Statistical Methods groups—Analysis Working Group, Enhancing GTEx (eGTEx) groups, NIH Common Fund, NIH/NCI, et al. Genetic effects on gene expression across human tissues. Nature. 2017 Oct 11;550(7675):204–13.

24. Verbanck M, Chen C-Y, Neale B, Do R. Detection of widespread horizontal pleiotropy in causal relationships inferred from Mendelian randomization between complex traits and diseases. Nature Genetics. 2018 Apr 23;1.

25. Shin S-Y, Fauman EB, Petersen A-K, Krumsiek J, Santos R, Huang J, et al. An atlas of genetic influences on human blood metabolites. Nat Genet. 2014 Jun;46(6):543–50.

688 26. Wang Z, Liao B-Y, Zhang J. Genomic patterns of pleiotropy and the evolution of
689 complexity. PNAS. 2010 Oct 19;107(42):18034–9.

690 27. Sivakumaran S, Agakov F, Theodoratou E, Prendergast JG, Zgaga L, Manolio T, et al.
691 Abundant pleiotropy in human complex diseases and traits. Am J Hum Genet. 2011 Nov
692 11;89(5):607–18.

693 28. Bowden J, Del Greco M F, Minelli C, Davey Smith G, Sheehan N, Thompson J. A
694 framework for the investigation of pleiotropy in two-sample summary data Mendelian
695 randomization. Statist Med. 2017 Jan 1;n/a-n/a.

696 29. Bowden J, Smith GD, Burgess S. Mendelian randomization with invalid instruments: effect
697 estimation and bias detection through Egger regression. Int J Epidemiol. 2015 Apr
698 1;44(2):512–25.

699 30. Liu X, Li YI, Pritchard JK. Trans effects on gene expression can drive omnigenic
700 inheritance. bioRxiv. 2018 Jan 1;425108.

701 31. Prokopenko I, Poon W, Mägi R, Prasad B R, Salehi SA, Almgren P, et al. A central role for
702 GRB10 in regulation of islet function in man. PLoS Genet. 2014 Apr;10(4):e1004235.

703 32. Bradfield JP, Taal HR, Timpson NJ, Scherag A, Lecoeur C, Warrington NM, et al. A
704 genome-wide association meta-analysis identifies new childhood obesity loci. Nat Genet.
705 2012 May;44(5):526–31.

706 33. Major Depressive Disorder Working Group of the Psychiatric GWAS Consortium, Ripke S,
707 Wray NR, Lewis CM, Hamilton SP, Weissman MM, et al. A mega-analysis of genome-wide
708 association studies for major depressive disorder. Mol Psychiatry. 2013 Apr;18(4):497–
709 511.

710 34. van der Valk RJP, Kreiner-Møller E, Kooijman MN, Guxens M, Stergiakouli E, Sääf A, et al.
711 A novel common variant in DCST2 is associated with length in early life and height in
712 adulthood. Hum Mol Genet. 2015 Feb 15;24(4):1155–68.

713 35. Liu JZ, van Sommeren S, Huang H, Ng SC, Alberts R, Takahashi A, et al. Association
714 analyses identify 38 susceptibility loci for inflammatory bowel disease and highlight shared
715 genetic risk across populations. Nat Genet. 2015 Sep;47(9):979–86.

716 36. Schizophrenia Working Group of the Psychiatric Genomics Consortium. Biological insights
717 from 108 schizophrenia-associated genetic loci. Nature. 2014 Jul 24;511(7510):421–7.

718 37. Hibar DP, Stein JL, Renteria ME, Arias-Vasquez A, Desrivières S, Jahanshad N, et al.
719 Common genetic variants influence human subcortical brain structures. Nature. 2015 Apr
720 9;520(7546):224–9.

721 38. Taal HR, St Pourcain B, Thiering E, Das S, Mook-Kanamori DO, Warrington NM, et al.
722 Common variants at 12q15 and 12q24 are associated with infant head circumference. Nat
723 Genet. 2012 Apr 15;44(5):532–8.

724 39. Wood AR, Esko T, Yang J, Vedantam S, Pers TH, Gustafsson S, et al. Defining the role of
725 common variation in the genomic and biological architecture of adult human height. Nat
726 Genet. 2014 Nov;46(11):1173–86.

727 40. Global Lipids Genetics Consortium, Willer CJ, Schmidt EM, Sengupta S, Peloso GM,
728 Gustafsson S, et al. Discovery and refinement of loci associated with lipid levels. Nat
729 Genet. 2013 Nov;45(11):1274–83.

730 41. Pattaro C, Teumer A, Gorski M, Chu AY, Li M, Mijatovic V, et al. Genetic associations at
731 53 loci highlight cell types and biological pathways relevant for kidney function. Nat
732 Commun. 2016 Jan 21;7:10023.

733 42. Locke AE, Kahali B, Berndt SI, Justice AE, Pers TH, Day FR, et al. Genetic studies of body
734 mass index yield new insights for obesity biology. Nature. 2015 Feb 12;518(7538):197–
735 206.

736 43. International Consortium for Blood Pressure Genome-Wide Association Studies, Ehret GB,
737 Munroe PB, Rice KM, Bochud M, Johnson AD, et al. Genetic variants in novel pathways
738 influence blood pressure and cardiovascular disease risk. Nature. 2011 Sep
739 11;478(7367):103–9.

740 44. Köttgen A, Albrecht E, Teumer A, Vitart V, Krumsiek J, Hundertmark C, et al. Genome-
741 wide association analyses identify 18 new loci associated with serum urate concentrations.
742 Nat Genet. 2013 Feb;45(2):145–54.

743 45. Felix JF, Bradfield JP, Monnereau C, van der Valk RJP, Stergiakouli E, Chesi A, et al.
744 Genome-wide association analysis identifies three new susceptibility loci for childhood
745 body mass index. Hum Mol Genet. 2016 Jan 15;25(2):389–403.

746 46. Cousminer DL, Berry DJ, Timpson NJ, Ang W, Thiering E, Byrne EM, et al. Genome-wide
747 association and longitudinal analyses reveal genetic loci linking pubertal height growth,
748 pubertal timing and childhood adiposity. Hum Mol Genet. 2013 Jul 1;22(13):2735–47.

749 47. Horikoshi M, Beaumont RN, Day FR, Warrington NM, Kooijman MN, Fernandez-Tajes J, et
750 al. Genome-wide associations for birth weight and correlations with adult disease. Nature.
751 2016 Oct 13;538(7624):248–52.

752 48. Okbay A, Beauchamp JP, Fontana MA, Lee JJ, Pers TH, Rietveld CA, et al. Genome-wide
753 association study identifies 74 loci associated with educational attainment. Nature. 2016
754 26;533(7604):539–42.

755 49. Okada Y, Wu D, Trynka G, Raj T, Terao C, Ikari K, et al. Genetics of rheumatoid arthritis
756 contributes to biology and drug discovery. Nature. 2014 Feb 20;506(7488):376–81.

757 50. Cousminer DL, Stergiakouli E, Berry DJ, Ang W, Groen-Blokhuis MM, Körner A, et al.
758 Genome-wide association study of sexual maturation in males and females highlights a
759 role for body mass and menarche loci in male puberty. Hum Mol Genet. 2014 Aug
760 15;23(16):4452–64.

761 51. Tobacco and Genetics Consortium. Genome-wide meta-analyses identify multiple loci
762 associated with smoking behavior. Nat Genet. 2010 May;42(5):441–7.

763 52. Cross-Disorder Group of the Psychiatric Genomics Consortium. Identification of risk loci
764 with shared effects on five major psychiatric disorders: a genome-wide analysis. Lancet.
765 2013 Apr 20;381(9875):1371–9.

766 53. Scott RA, Lagou V, Welch RP, Wheeler E, Montasser ME, Luan J, et al. Large-scale
767 association analyses identify new loci influencing glycemic traits and provide insight into
768 the underlying biological pathways. Nat Genet. 2012 Sep;44(9):991–1005.

769 54. Manning AK, LaValley M, Liu C-T, Rice K, An P, Liu Y, et al. Meta-analysis of Gene-
770 Environment interaction: joint estimation of SNP and SNP×Environment regression
771 coefficients. Genet Epidemiol. 2011 Jan;35(1):11–8.

772 55. the CARDIoGRAMplusC4D Consortium. A comprehensive 1000 Genomes-based genome-
773 wide association meta-analysis of coronary artery disease. Nat Genet. 2015
774 Oct;47(10):1121–30.

775 56. Morris AP, Voight BF, Teslovich TM, Ferreira T, Segrè AV, Steinthorsdottir V, et al. Large-
776 scale association analysis provides insights into the genetic architecture and
777 pathophysiology of type 2 diabetes. Nat Genet. 2012 Sep;44(9):981–90.

778 57. Psychiatric GWAS Consortium Bipolar Disorder Working Group. Large-scale genome-wide
779 association analysis of bipolar disorder identifies a new susceptibility locus near ODZ4. Nat
780 Genet. 2011 Sep 18;43(10):977–83.

781 58. Stolk L, Perry JRB, Chasman DI, He C, Mangino M, Sulem P, et al. Meta-analyses identify
782 13 loci associated with age at menopause and highlight DNA repair and immune
783 pathways. Nat Genet. 2012 Jan 22;44(3):260–8.

784 59. Lambert JC, Ibrahim-Verbaas CA, Harold D, Naj AC, Sims R, Bellenguez C, et al. Meta-
785 analysis of 74,046 individuals identifies 11 new susceptibility loci for Alzheimer's disease.
786 Nat Genet. 2013 Dec;45(12):1452–8.

787 60. Dehghan A, Dupuis J, Barbalic M, Bis JC, Eiriksdottir G, Lu C, et al. Meta-analysis of
788 genome-wide association studies in >80 000 subjects identifies multiple loci for C-reactive
789 protein levels. Circulation. 2011 Feb 22;123(7):731–8.

790 61. Neale BM, Medland SE, Ripke S, Asherson P, Franke B, Lesch K-P, et al. Meta-analysis of
791 genome-wide association studies of attention-deficit/hyperactivity disorder. J Am Acad
792 Child Adolesc Psychiatry. 2010 Sep;49(9):884–97.

793 62. Dupuis J, Langenberg C, Prokopenko I, Saxena R, Soranzo N, Jackson AU, et al. New
794 genetic loci implicated in fasting glucose homeostasis and their impact on type 2 diabetes
795 risk. Nat Genet. 2010 Feb;42(2):105–16.

796 63. Shungin D, Winkler TW, Croteau-Chonka DC, Ferreira T, Locke AE, Mägi R, et al. New
797 genetic loci link adipose and insulin biology to body fat distribution. Nature. 2015 Feb
798 12;518(7538):187–96.

799 64. Dastani Z, Hivert M-F, Timpson N, Perry JRB, Yuan X, Scott RA, et al. Novel loci for
800 adiponectin levels and their influence on type 2 diabetes and metabolic traits: a multi-
801 ethnic meta-analysis of 45,891 individuals. PLoS Genet. 2012;8(3):e1002607.

65. Perry JRB, Day F, Elks CE, Sulem P, Thompson DJ, Ferreira T, et al. Parent-of-origin-specific allelic associations among 106 genomic loci for age at menarche. Nature. 2014 Oct 2;514(7520):92–7.

66. van der Harst P, Zhang W, Mateo Leach I, Rendon A, Verweij N, Sehmi J, et al. Seventy-five genetic loci influencing the human red blood cell. Nature. 2012 Dec 20;492(7429):369–75.

67. Shin S-Y, Fauman EB, Petersen A-K, Krumsiek J, Santos R, Huang J, et al. An atlas of genetic influences on human blood metabolites. Nat Genet. 2014 Jun;46(6):543–50.

68. Dickinson ME, Flenniken AM, Ji X, Teboul L, Wong MD, White JK, et al. High-throughput discovery of novel developmental phenotypes. Nature. 2016 22;537(7621):508–14.

69. Saito TL, Ohtani M, Sawai H, Sano F, Saka A, Watanabe D, et al. SCMD: Saccharomyces cerevisiae Morphological Database. Nucleic Acids Res. 2004 Jan 1;32(Database issue):D319-322.

70. Huang DW, Sherman BT, Lempicki RA. Bioinformatics enrichment tools: paths toward the comprehensive functional analysis of large gene lists. Nucleic Acids Res. 2009 Jan;37(1):1–13.

71. Huang DW, Sherman BT, Lempicki RA. Systematic and integrative analysis of large gene lists using DAVID bioinformatics resources. Nat Protoc. 2009;4(1):44–57.

72. DAVID Bioinformatics Resources v.6.8 [Internet]. [cited 2019 Jan 14]. Available from: https://david.ncifcrf.gov/

Figure Titles and Legends

Figure 1: Schematic of different types of pleiotropy.

Previous studies distinguish between vertical pleiotropy, where effects on one trait are mediated through effects on another trait, and horizontal pleiotropy, where effects on multiple traits are independent.

Figure 2: Contributions of linkage disequilibrium (LD) and polygenicity to horizontal pleiotropy.

In addition to the normal sense of horizontal pleiotropy, both linkage disequilibrium (LD) and polygenicity are expected to contribute to horizontal pleiotropy. In the case of LD-induced horizontal pleiotropy, two linked SNVs have independent effects on different traits which appear pleiotropic because of the linkage between the SNVs. In the case of polygenicity-induced horizontal pleiotropy, two highly polygenic traits have an overlap in their polygenic footprint.

Figure 3: Two component pleiotropy score method.

We (i) collect association statistics from the UK Biobank, (ii) process them using Mahalanobis whitening, (iii) compute the two components of our pleiotropy score (P_m and P_n) based on the whitened association statistics, (iv) use LD scores to correct for LD-induced pleiotropy (P_m^{LD} and P_n^{LD}), and (v) use permutation-based *P*-values to correct for polygenic architecture (P_m^P and P_n^P).

Figure 4: Simulation study showing false positive rate (a,b,c,d) and power (e,f,g,h) of two-component pleiotropy score.

Top row shows performance on non-pleiotropic simulated variants (black line shows 5% false positive rate); bottom row shows performance on pleiotropic variants (black line shows 80% power). Simulations were run for both P_m^{LD} (left) and P_n^{LD} (right), and both without correction for

846 polygenicity (a,c,e,g) and with the correction (b,f,d,h), with per-variant heritability ranging from
847 0.0002 to 0.2, proportion of non-pleiotropic causal loci ranging from 0 to 1%, and proportion of
848 pleiotropic causal loci ranging from 0.1% to 1%. Our method has good power to detect
849 pleiotropy for highly heritable traits, though its power is reduced by extreme polygenicity.
850 Extreme polygenicity also increases the false positive rate, though this effect is corrected by our
851 polygenicity correction.

852 **Figure 5: Quantile-quantile (Q-Q) plots showing the inflation of the pleiotropy score as a**
853 **function of polygenicity.**

854 Variants are stratified into 4 batches of about 80 traits each by heritability, and then subdivided
855 into 5 batches of about 20 traits each by polygenicity, as measured by corrected genomic
856 inflation factor λ_{GC}^{c}. Darker shades represent low polygenicity and lighter shades represent high
857 polygenicity. All panels show -$\log_{10}$ transformed *P-values*. The black lines show the expected
858 value under the null hypothesis.

859 **Figure 6: Distribution of the pleiotropy score among variants (a), genes (b), and traits (c).**

860 Panel a shows the global distribution of P_m^{LD} (left) and P_n^{LD} (right) for the 767,057 tested variants.
861 The expected distribution under the null hypothesis of no pleiotropy is shown in red and the
862 observed distribution is shown in blue. The vertical line represents the value of the pleiotropy
863 score corresponding to genome-wide significance ($P < 5 \times 10^{-8}$). 1,769 (P_m^{LD}) and 643 (P_n^{LD})
864 variants are not represented for the sake of clarity, because they have extreme values for the
865 pleiotropy score. Panel b shows the distribution of the average pleiotropy score for coding
866 variants in each gene for P_m^{LD} (left) and P_n^{LD} (right). The top ten genes are represented on the
867 right side of the plots, whereas genes with a pleiotropy score of 0 are represented on the left
868 side of the plots. Panel c shows the contribution of pleiotropic variants to 82 complex traits and
869 diseases. Contribution of pleiotropic variants is calculated as the correlation coefficient between

the absolute value of Z-scores and the pleiotropy score among variants that are genome-wide significant for the pleiotropy score ($P < 5 \times 10^{-8}$ for P_m^{LD} and P_n^{LD} respectively).

Figure 7: Replication analysis for the genome-wide pleiotropy study.

We used 372 UK Biobank heritable medical traits as our discovery dataset, and independent datasets of 73 complex traits and diseases and 430 blood metabolites as replication datasets. In each case, expected fraction of replication was empirically determined using a permutation analysis.

Tables

Table 1: Functional enrichment analysis of pleiotropy score.

			P_m^{LD}	P_n^{LD}
Variant effect predictor		UTR	**+0.24 (±0.01); P = 1.72x10^{-234}**	**+0.69 (±0.02); P = 2.16x10^{-236}**
		coding synonymous	**+0.24 (±0.01); P = 2.49x10^{-99}**	**+0.61 (±0.03); P = 1.92x10^{-76}**
		non synonymous	**+0.19 (±0.01); P = 3.82x10^{-82}**	**+0.48 (±0.03); P = 3.62x10^{-62}**
Roadmap Epigenomics	Promoter	H327ac	**+0.20 (±0.01); $P < 10^{-308}$**	**+0.54 (±0.01); $P < 10^{-308}$**
		H3K27me3	**+0.02 (±0.01); P = 1.40x10^{-18}**	+0.01 (±0.01); P = 0.4
		Active TSS	**+0.20 (±0.02); P = 1.42x10^{-36}**	**+0.54 (±0.04); P = 8.56x10^{-34}**
		Promoter Upstream TSS	**+0.16 (±0.01); P = 4.44x10^{-130}**	**+0.43 (±0.02); P = 4.33x10^{-103}**
		Promoter Downstream TSS 1	**+0.35 (±0.01); P = 1.87x10^{-220}**	**+0.92 (±0.03); P = 3.59x10^{-197}**
		Promoter Downstream TSS 2	**+0.30 (±0.01); P = 2.70x10^{-203}**	**+0.86 (±0.03); P = 3.44x10^{-210}**
	Transcription	Transcribed - 5' preferential	**+0.29 (±0.01); $P < 10^{-308}$**	**+0.88 (±0.01); $P < 10^{-308}$**
		Strong transcription	**+0.38 (±0.01); $P < 10^{-308}$**	**+1.10 (±0.01); $P < 10^{-308}$**
		Transcribed - 3' preferential	**+0.29 (±0.01); $P < 10^{-308}$**	**+0.82 (±0.01); $P < 10^{-308}$**

Transcription & regulation	Weak transcription	+0.21 (±0.01); $P < 10^{-308}$	+0.60 (±0.01); $P < 10^{-308}$
	Transcribed & regulatory (Prom/Enh)	+0.36 (±0.01); $P < 10^{-308}$	+1.00 (±0.02); $P < 10^{-308}$
	Transcribed 5' preferential and Enh	+0.35 (±0.01); $P < 10^{-308}$	+1.00 (±0.01); $P < 10^{-308}$
	Transcribed 3' preferential and Enh	+0.33 (±0.01); $P < 10^{-308}$	+0.92 (±0.02); $P < 10^{-308}$
	Transcribed and Weak Enhancer	+0.32 (±0.01); $P < 10^{-308}$	+0.97 (±0.01); $P < 10^{-308}$
Active enhancer	Active Enhancer 1	+0.13 (±0.01); $P = 4.54 \times 10^{-295}$	+0.32 (±0.01); $P = 5.1 \times 10^{-216}$
	Active Enhancer 2	+0.11 (±0.01); $P = 2.64 \times 10^{-294}$	+0.28 (±0.01); $P = 5.63 \times 10^{-238}$
	Active Enhancer Flank	+0.11 (±0.01); $P < 10^{-308}$	+0.29 (±0.01); $P = 6.06 \times 10^{-270}$
Weak enhancer	Weak Enhancer 1	+0.07 (±0.01); $P = 2.79 \times 10^{-89}$	+0.16 (±0.01); $P = 6.89 \times 10^{-60}$
	Weak Enhancer 2	+0.08 (±0.01); $P < 10^{-308}$	+0.23 (±0.01); $P = 6.52 \times 10^{-291}$
	Primary H3K27ac possible Enhancer	+0.09 (±0.01); $P = 2.72 \times 10^{-259}$	+0.24 (±0.01); $P = 1.53 \times 10^{-187}$
	Primary DNase	+0.03 (±0.01); $P = 3.83 \times 10^{-21}$	+0.05 (±0.01); $P = 1.11 \times 10^{-7}$
	ZNF genes & repeats	+0.08 (±0.01); $P = 1.29 \times 10^{-7}$	+0.20 (±0.04); $P = 6.9 \times 10^{-7}$
	Heterochromatin	-0. 20 (±0.01); $P < 10^{-308}$	-0.61 (±0.01); $P < 10^{-308}$
	Poised Promoter	+0.05 (±0.01); $P = 1.03 \times 10^{-35}$	+0.09 (±0.01); $P = 2.27 \times 10^{-16}$
	Bivalent Promoter	+0.17 (±0.01); $P = 1.28 \times 10^{-93}$	+0.51 (±0.03); $P = 6.29 \times 10^{-88}$

	Repressed Polycomb	**+0.04 (±0.01); $P = 5.77 \times 10^{-42}$**	**+0.06 (±0.01); $P = 1.48 \times 10^{-11}$**
	Quiescent/Low	**-0.41 (±0.01); $P < 10^{-308}$**	**-1.20 (±0.01); $P < 10^{-308}$**
GTEx - number of genes the variant is an eQTL for	eGenes<10	**+0.11 (±0.01); $P = 6.78 \times 10^{-186}$**	**+0.28 (±0.01); $P = 1.04 \times 10^{-140}$**
	eGenes>10 & <15	**+0.19 (±0.01); $P = 4.72 \times 10^{-114}$**	**+0.52 (±0.02); $P = 6.84 \times 10^{-99}$**
	eGenes>15 & <20	**+0.31 (±0.02); $P = 7.98 \times 10^{-52}$**	**+0.88 (±0.06); $P = 5.38 \times 10^{-47}$**
	eGenes>20	**+0.66 (±0.06); $P = 3.40 \times 10^{-27}$**	**+2.07 (±0.18); $P = 1.35 \times 10^{-30}$**
GTEx - number of tissues the variant is an eQTL for	eTissue<30	**+0.10 (±0.01); $P = 1.84 \times 10^{-151}$**	**+0.26 (±0.01); $P = 1.26 \times 10^{-114}$**
	eTissue>30 & <35	**+0.21 (±0.01); $P = 3.70 \times 10^{-187}$**	**+0.54 (±0.02); $P = 6.80 \times 10^{-147}$**
	eTissue>35 & <40	**+0.36 (±0.02); $P = 1.11 \times 10^{-82}$**	**+1.13 (±0.06); $P = 4.24 \times 10^{-92}$**
	eTissue>40	**+0.35 (±0.05); $P = 2{,}42 \times 10^{-13}$**	**+0.97 (±0.14); $P = 7.08 \times 10^{-12}$**
International Mouse Phenotyping Consortium	Phenotypes > 1	**+0.06 (±0.01); $P = 1.91 \times 10^{-6}$**	**+0.19 (±0.04); $P = 2.70 \times 10^{-7}$**
Saccharomyces cerevisiae Morphological Database	Phenotypes > 1	**+0.09 (±0.01); $P = 4.48 \times 10^{-17}$**	**+0.26 (±0.03); $P = 1.53 \times 10^{-18}$**

879

880 We grouped variants by (i) molecular function as annotated by Ensembl, (ii) predicted chromatin state as annotated by the NIH
881 Roadmap Epigenomics Project, (iii) transcriptional effects as annotated by the NIH Genotype-Tissue Expression (GTex) Project, and
882 (iv) effects on model organism phenotypes as annotated by the International Mouse Phenotyping Consortium (IMPC) and

Saccharomyces Cerevisiae Morphological Database (SCMD). For each grouping, we computed the mean LD-corrected pleiotropy score and used a two-sample Student's t-test to determine whether the mean was significantly different from the baseline. We found (i) that coding regions have higher pleiotropy scores than noncoding regions, (ii) that active promoters and enhancers have the highest pleiotropy scores and quiescent and heterochromatin have the lowest, (iii) that variants that control expression of more genes in more tissues have higher pleiotropy scores, and (iv) that genes associated with more than one model organism phenotype have higher pleiotropy scores.

Additional Files

Additional File 1. Functional enrichment analysis of pleiotropy score after applying polygenicity correction.

Excel spreadsheet (.xlsx) showing the equivalent of Table 1 using the LD/polygenicity-corrected scores (P_m^P and P_n^P) instead of the LD-corrected scores (P_m^{LD} and P_n^{LD})

Additional File 2. DAVID enrichment analysis of high-pleiotropy genes.

Excel spreadsheet (.xlsx) showing the results of the DAVID enrichment analysis described in the text.

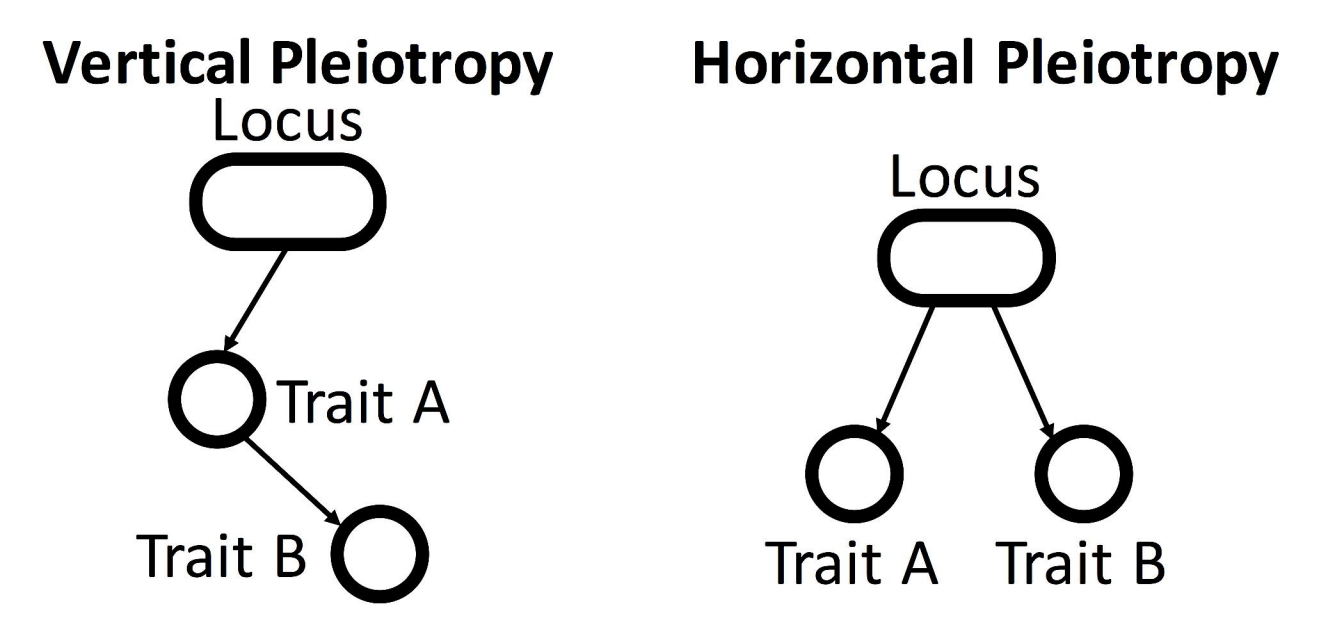
Vertical Pleiotropy
Locus
Trait A
Trait B
Horizontal Pleiotropy
Locus
Trait A
Trait B

Biological (single-locus) Horizontal Pleiotropy

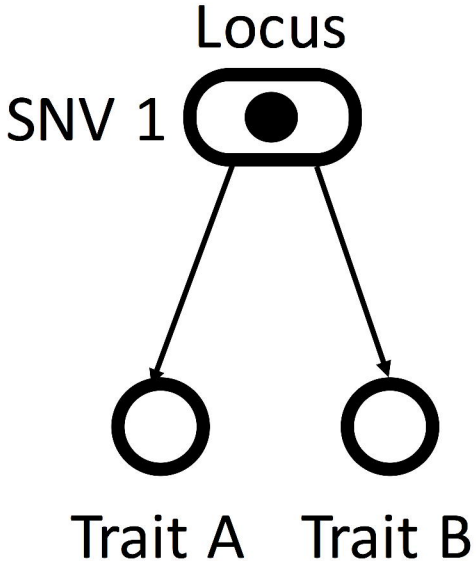

LD-Induced Horizontal Pleiotropy

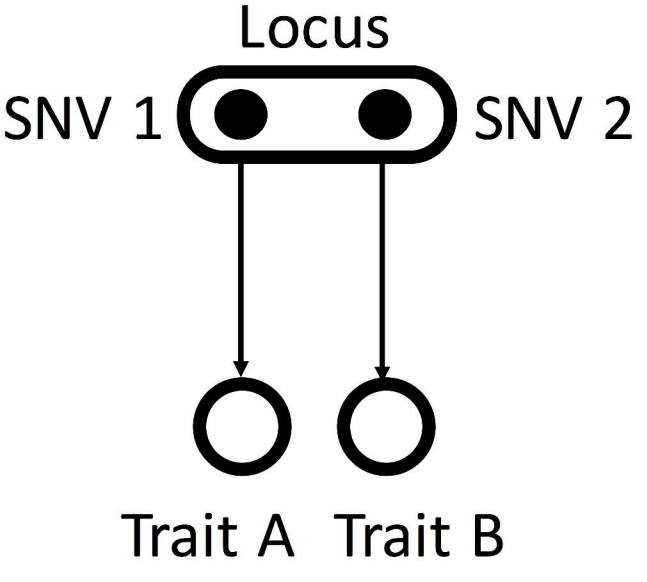

Polygenicity-Induced Horizontal Pleiotropy

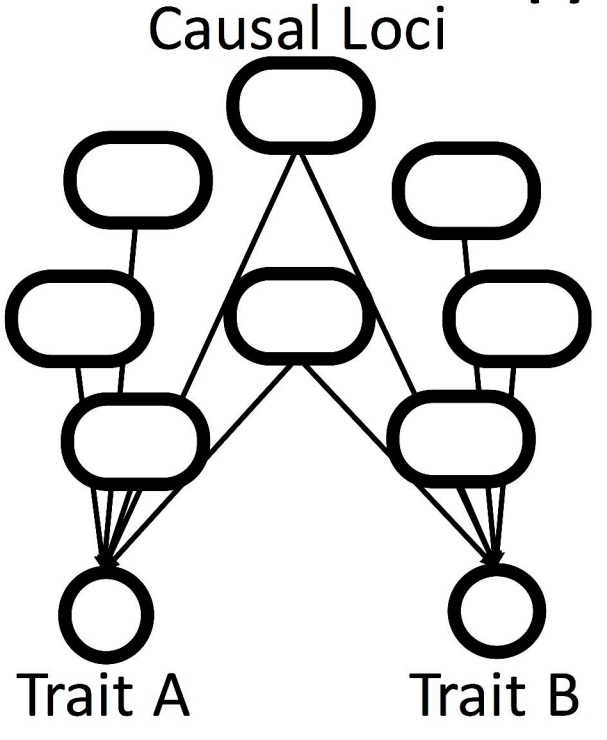

Collection of UK Biobank association summary statistics

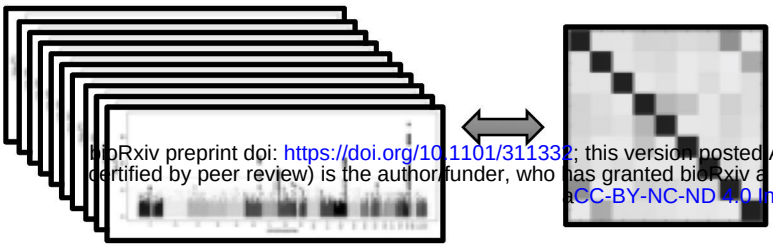

Mahalanobis whitening
(remove trait correlations)

Whitened (uncorrelated) association statistics

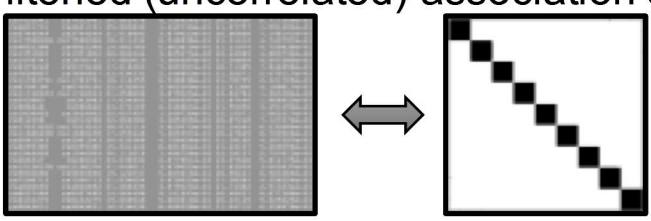

Sum squared association statistics

Counts of nominally significant whitened traits

Magnitude pleiotropy score (P_m)

Number of traits pleiotropy score (P_n)

Correction for LD

LD-corrected Magnitude pleiotropy score (P_m^{LD})

LD-corrected Number of traits pleiotropy score (P_n^{LD})

Correction for polygenicity

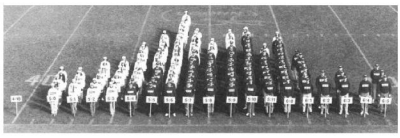

Polygenicity-corrected Magnitude pleiotropy score (P_m^P)

Polygenicity-corrected Number of traits pleiotropy score (P_n^P)

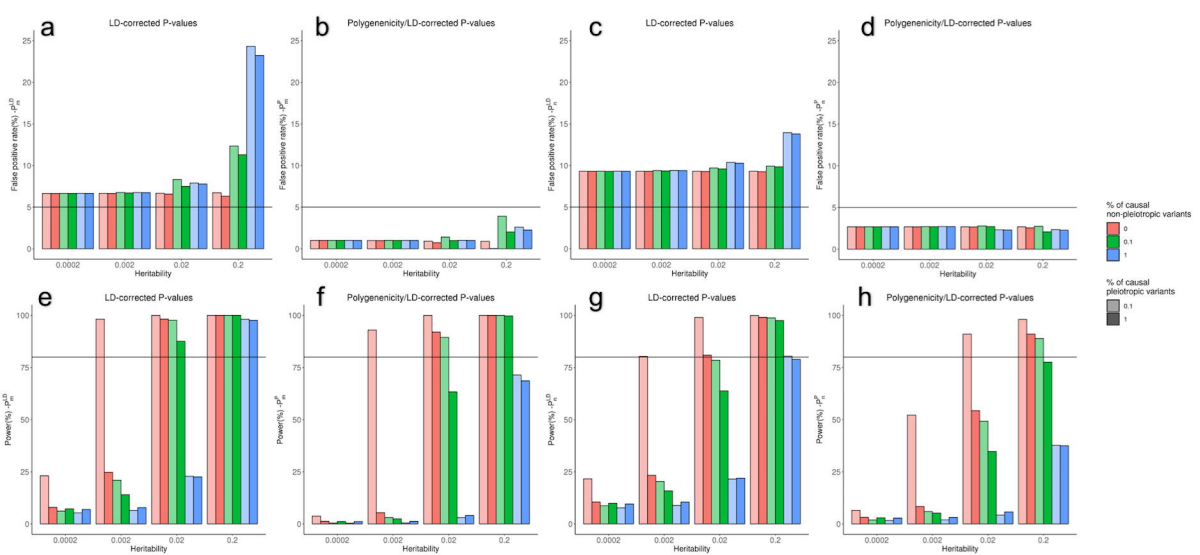
a
LD-corrected P-values
b
Polygenenicity/LD-corrected P-values
c
LD-corrected P-values
d
Polygenenicity/LD-corrected P-values
False positive rate(%) p^{LD}_{nn}
False positive rate(%) p^{B}_{nn}
False positive rate(%) p^{LD}_{n}
False positive rate(%) p^{B}_{n}
0
5
10
15
20
25
0.0002
0.002
0.02
0.2
Heritability
e
LD-corrected P-values
f
Polygenenicity/LD-corrected P-values
g
LD-corrected P-values
h
Polygenenicity/LD-corrected P-values
Power(%) p^{LD}_{nn}
Power(%) p^{B}_{nn}
Power(%) p^{LD}_{n}
Power(%) p^{B}_{n}
0
25
50
75
80
100
0.0002
0.002
0.02
0.2
Heritability
% of causal non-pleiotropic variants
0
0.1
1
% of causal pleiotropic variants
0.1
1

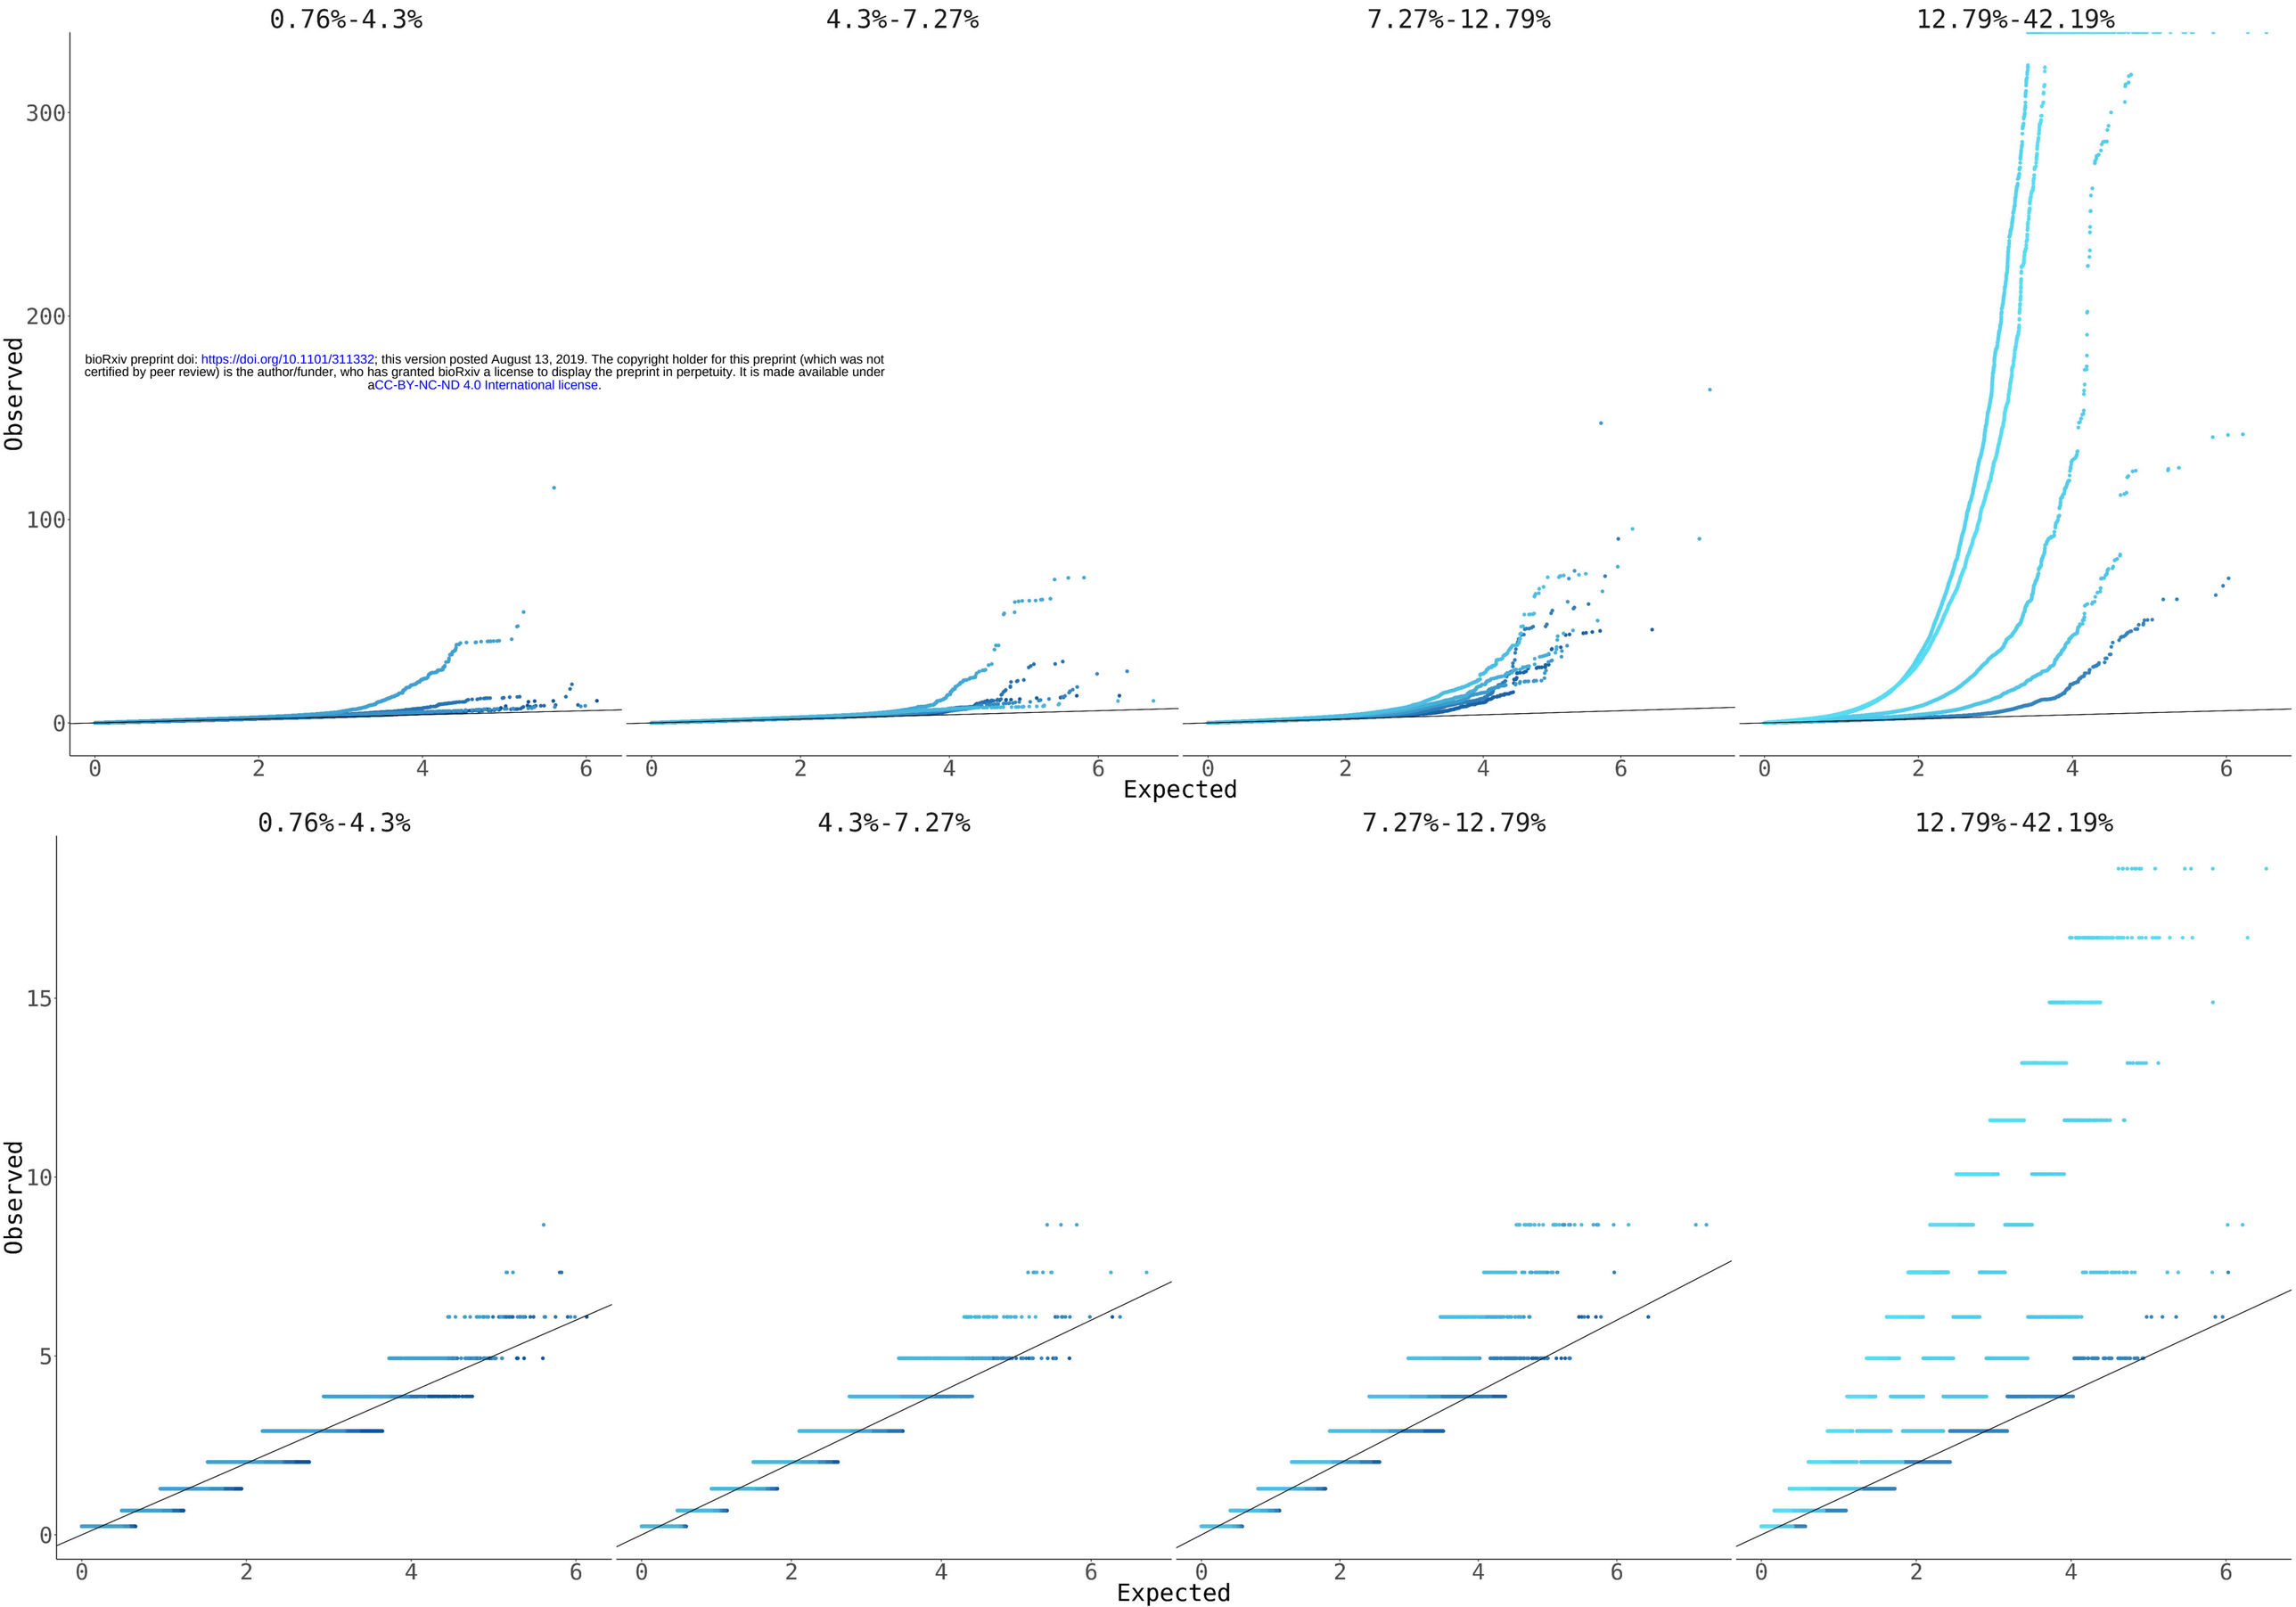

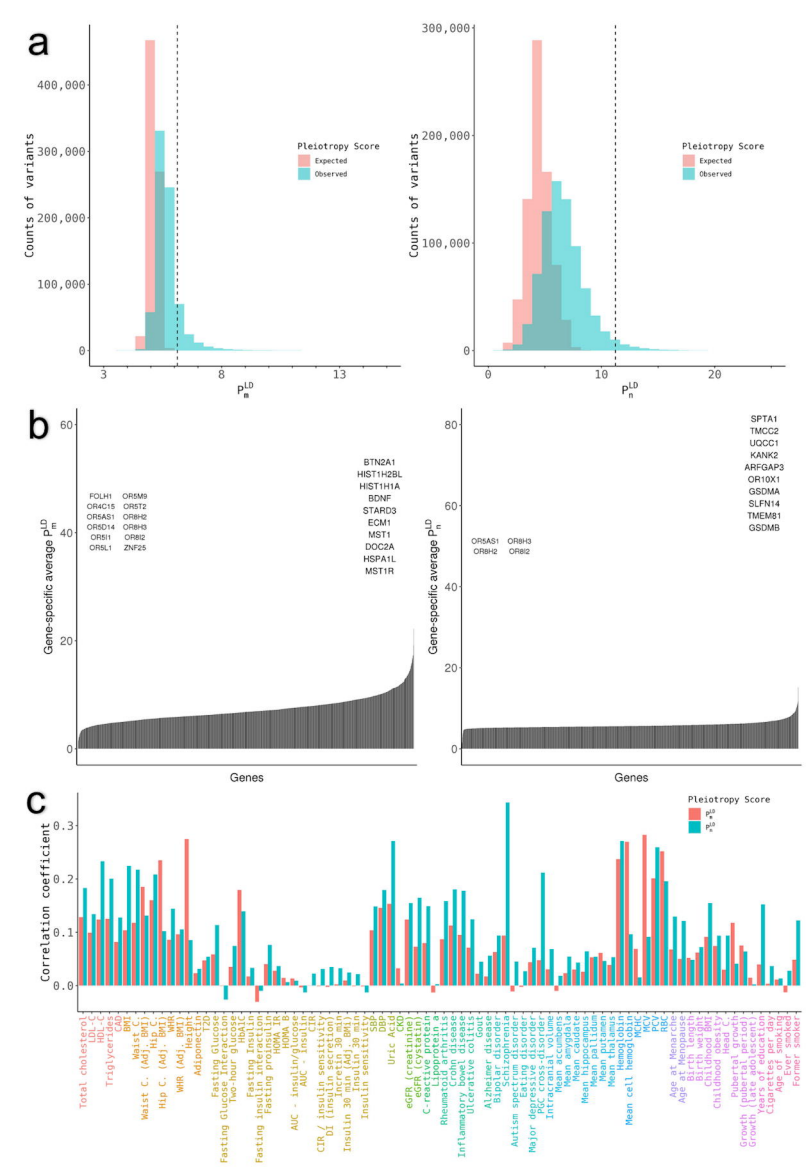
a
Counts of variants
Pleiotropy Score
Expected
Observed
P_m^{LD}
P_n^{LD}
b
Gene-specific average P_E^{LD}
Gene-specific average P_n^{LD}
FOLH1 OR5M9
OR4C15 OR5T2
OR5AS1 OR8H2
OR5D14 OR8H3
OR5I1 OR8I2
OR5L1 ZNF25
BTN2A1
HIST1H2BL
HIST1H1A
BDNF
STARD3
ECM1
MST1
DOC2A
HSPA1L
MST1R
OR5AS1 OR8H3
OR8H2 OR8I2
SPTA1
TMCC2
UQCC1
KANK2
ARFGAP3
OR10X1
GSDMA
SLFN14
TMEM81
GSDMB
Genes
c
Correlation coefficient
Total cholesterol
LDL-C
HDL-C
Triglycerides
CAD
MI
Waist C. (Adj. BMI)
Waist C.
Hip C. (Adj. BMI)
Hip C.
WHR
WHR (Adj. BMI)
BMI
Height
Adiponectin
Fasting glucose
Fasting glucose
Fasting Glucose Increase
Two-hour glucose
HbA1c
Fasting Insulin
Fasting insulin
Fasting insulin interaction
Fasting proinsulin
HOMA-B
HOMA-IR
AUC - insulin/glucose
AUC - insulin
CIR
CIR_i (insulin sensitivity)
DI (insulin secretion)
Incretin 30 min
Insulin 30 min adjusted
Insulin30_min
Insulin sensitivity
SBP
DBP
Uric acid
CKD
eGFR (creatinine)
eGFR (cystatin)
C-reactive protein
Lipoprotein A
Rheumatoid arthritis
Crohn disease
IBD
Ulcerative colitis
Inflammatory bowel disease
Gout
Alzheimer disease
Schizophrenia
Bipolar disorder
Autism spectrum disorder
Eating disorder
Major depressive disorder
PGC cross-disorder
Intracranial volume
White matter
Mean amygdala
Mean caudate
Mean putamen
Mean pallidum
Mean hippocampus
Mean thalamus
Hemoglobin
MCH
Mean cell hemoglobin
MCHC
MCV
PLT
RBC
Age at Menopause
Age at Menopause
Birth length
Birth weight
Childhood BMI
Childhood obesity
Puberty BMI
Pubertal growth
Growth (puberal period)
Growth (puberal period)
Years of education
Cigarettes per day
Age of smoking
Ever smoked
Former smoked

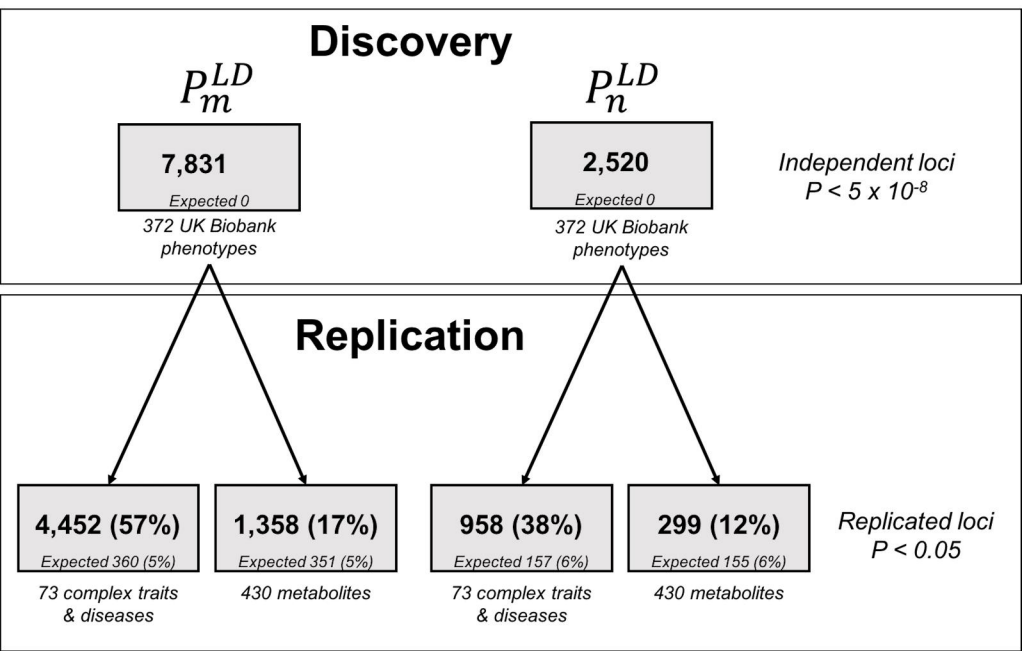
Discovery
P_m^{LD}
7,831
Expected 0
372 UK Biobank phenotypes
P_n^{LD}
2,520
Expected 0
372 UK Biobank phenotypes
Independent loci
P < 5 x 10-8
Replication
4,452 (57%)
Expected 360 (5%)
73 complex traits & diseases
1,358 (17%)
Expected 351 (5%)
430 metabolites
958 (38%)
Expected 157 (6%)
73 complex traits & diseases
299 (12%)
Expected 155 (6%)
430 metabolites
Replicated loci
P < 0.05