

UCSF

UC San Francisco Previously Published Works

Title

A model and test for coordinated polygenic epistasis in complex traits

Permalink

<https://escholarship.org/uc/item/20z6j00v>

Journal

Proceedings of the National Academy of Sciences of the United States of America, 118(15)

ISSN

0027-8424

Authors

Sheppard, Brooke
Rappoport, Nadav
Loh, Po-Ru
[et al.](#)

Publication Date

2021-04-13

DOI

10.1073/pnas.1922305118

Peer reviewed



A model and test for coordinated polygenic epistasis in complex traits

Brooke Sheppard^{a,1}, Nadav Rappoport^{a,b,1}, Po-Ru Loh^{c,d}, Stephan J. Sanders^a, Noah Zaitlen^{a,e,f,1,2}, and Andy Dahl^{e,f,g,1,2}

^aDepartment of Psychiatry and Behavioral Sciences, Weill Institute for Neurosciences, University of California San Francisco, San Francisco, CA 94143; ^bBakar Computational Health Sciences Institute, University of California San Francisco, San Francisco, CA 94143; ^cProgram in Medical and Population Genetics, Broad Institute of Massachusetts Institute of Technology and Harvard, Cambridge, MA 02142; ^dDivision of Genetics, Department of Medicine, Brigham and Women's Hospital and Harvard Medical School, Boston, MA 02115; ^eDepartment of Neurology, University of California Los Angeles, Los Angeles, CA 90095; ^fDepartment of Computational Medicine, University of California Los Angeles, Los Angeles, CA 90095; and ^gSection of Genetic Medicine, University of Chicago, Chicago, IL 60637

Edited by Andrew G. Clark, Cornell University, Ithaca, NY, and approved December 2, 2020 (received for review February 19, 2020)

Interactions between genetic variants—epistasis—is pervasive in model systems and can profoundly impact evolutionary adaptation, population disease dynamics, genetic mapping, and precision medicine efforts. In this work, we develop a model for structured polygenic epistasis, called coordinated epistasis (CE), and prove that several recent theories of genetic architecture fall under the formal umbrella of CE. Unlike standard epistasis models that assume epistasis and main effects are independent, CE captures systematic correlations between epistasis and main effects that result from pathway-level epistasis, on balance skewing the penetrance of genetic effects. To test for the existence of CE, we propose the even-odd (EO) test and prove it is calibrated in a range of realistic biological models. Applying the EO test in the UK Biobank, we find evidence of CE in 18 of 26 traits spanning disease, anthropometric, and blood categories. Finally, we extend the EO test to tissue-specific enrichment and identify several plausible tissue–trait pairs. Overall, CE is a dimension of genetic architecture that can capture structured, systemic forms of epistasis in complex human traits.

genetics | epistasis | polygenic risk

Interaction between the phenotypic effects of genetic variants, or epistasis, is an essential component of biology with important consequences across multiple scientific domains. For example, epistasis significantly impacts evolutionary models, including response to selection or changing environment (1, 2). Epistasis also fundamentally shapes genetic architecture, as the direction of an allele's effect can change based on genetic background (3–6). Epistatic interactions are pervasive in model systems, including in model organisms (2, 7–11) and in recent mammalian gene-level interaction screens (12–15). Moreover, these interactions often represent core biological functions. For example, the molecular chaperone HSP90 modifies diverse disease-model-specific proteins (16–18), and conceptually similar mechanisms protect the cell from aberrant translation at ribosomes (19). Such highly structured forms of epistasis are a core focus of systems biology, and their genetic bases have long been hypothesized to influence complex human disease.

While the observations in model organisms suggest that epistasis is also fundamental in humans (11, 20), it remains poorly understood. This is largely because powerful, interpretable modeling tools are nascent. Genome-wide searches for interacting single-nucleotide polymorphism (SNP) pairs are computationally and statistically onerous, despite some success (21). Additionally, recent and classical methods for epistasis increase power by aggregating interactions across the whole genome (3, 22–28), and their results further support the potential importance of epistasis in complex traits. Interestingly, all of these approaches assume that interaction effect sizes and directions are independent of main effects or, in other words, that epistasis is uncoordinated.

This is in contrast with recent conceptual models of complex human traits that imply or are consistent with coordinated forms

of epistasis. For example, Zuk et al. describe a limiting pathway model of human disease that directly implies negative interactions between SNPs contributing to different pathways (29). Also, the HSP90 community has discussed the possibility of a polygenic version of chaperon function (30), which would induce coordinated interactions between HSP90 buffer SNPs and exonic missense SNPs. Another example is the omnigenic model, which suggests that traits are determined by a few trait-specific “core” genes that are modified by many nonspecific “peripheral” genes (31, 32), implying coordinated epistasis (CE) between the SNPs contributing to “core” and “peripheral” genes. Monogenic disorders with polygenic modifiers are another instance of structured epistasis, where the expressivity and penetrance of mutations in a core gene are modified by many variants distributed among the genome (33–35).

Motivated by these concerns, we formally propose a concrete parameter to capture a specific form of structured epistasis, which we call CE and define formally below. Conceptually, CE measures concerted interactions between pathways that are themselves additively heritable traits. For example, one pathway might capture additive genetic effects on a polygenic buffering pathway, similar to HSP90, and another pathway might capture genetic effects on

Significance

Systems-level interactions across physiological pathways, cell types, and tissues are core biological elements widely studied across diverse fields including evolution, systems biology, and model-organism genetics. However, they are essentially ignored in human genetics, and existing approaches fail to interpretably explain substantial complex trait heritability. Here, we propose the coordinated epistasis model of complex phenotypes that generalizes several recently proposed theoretical epistatic architectures of human traits. Broadly, coordination measures the degree to which epistasis effects act in concert with respect to marginal effects. It summarizes a dimension of polygenic effects orthogonal to parameters like heritability and standard estimates of epistasis.

Author contributions: N.Z. and A.D. designed research; B.S., N.R., N.Z., and A.D. performed research; B.S., N.R., P.-R.L., and A.D. analyzed data; and B.S., N.R., P.-R.L., S.J.S., N.Z., and A.D. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

¹B.S., N.R., N.Z., and A.D. contributed equally to this work.

²To whom correspondence may be addressed. Email: nzaitlen@ucla.edu or andywdahl@uchicago.edu.

This article contains supporting information online at <https://www.pnas.org/lookup/suppl/doi:10.1073/pnas.1922305118/-DCSupplemental>.

Published April 8, 2021.

individual proteins that contribute to disease. In another example, type 2 diabetes (T2D) involves interactions between insulin-sensing and -producing systems, which likely have distinct genetic architectures and operate through different tissues (36–38). We quantify CE with the parameter γ , where $\gamma < 0$ indicates negative epistasis between genetic pathways, on average, dampening the marginal effects of trait-increasing variants; conversely, $\gamma > 0$ indicates positive epistasis, where trait-increasing SNP effects are actually magnified in the full organismal context relative to their genome-wide association study (GWAS) estimates. Importantly, CE is a test for general statistical epistasis and does not generally indicate the presence of biological interactions as in a protein complex (35).

For simplicity, we usually imagine CE in the context of positive SNP effects, where positive/negative epistasis coincide with the more interpretable notions of synergy/antagonism. For example, for a disease trait, positive epistasis and synergy both imply that the joint effect of risk alleles exceeds the sum of their marginal effects. In general, however, CE measures the skew in genetic risks, and $\gamma > 0$ may either dampen or exacerbate marginal genetic effects (39). In Fig. 1, we illustrate γ in the case of two latent interacting risk pathways. The key is that $\gamma \neq 0$ if there are heritable, positively (synergistically) or negatively (antagonistically) interacting pathways (Fig. 1B), while $\gamma = 0$ if the interactions are purely random or if the interactions are absent (Fig. 1A).

Testing for the existence of CE would be straightforward if we knew or could accurately estimate the SNP effects on causal, latent pathways: we could simply build genetic predictors of each pathway and test their interaction. However, these pathways are generally unknown and high dimensional, making CE estimation seem impossible. Surprisingly, we prove that testing for CE can be accomplished by randomly assigning independent sets of SNPs to arbitrary proxy pathways. Concretely, we build polygenic risk scores (PRS) specifically for the even and odd chromosomes and then test their interaction in a linear regression on phenotype (Fig. 1C). We call this the even-odd (EO) test, and we prove analytically that EO reliably estimates CE under polygenicity. We also prove that under a polygenic model, any chromosome split, not just EO, gives identical results in expectation. Importantly, however, partitions aligned closely with the true latent pathways will have greater power. We leverage this fact to test for CE enrichment in specific genomic annotations. Specifically, we focus on tissue-specific annotations to ask whether complex traits are enriched for interactions between tissue-specific pathways.

In this paper, we first introduce the concept of CE formally and define the EO test. We then examine simulated data sets to show that EO is robust to confounding under plausible models of assortative mating and population structure. Then, we perform the EO test in the UK Biobank (UKBB) across 26 traits spanning multiple domains and find 18 with significant CE. We validate the approach, which uses internally cross-validated PRS, by using PRS

constructed from external data sources for 17 of the 26 traits. Finally, we estimate tissue-specific CE across 13 tissues in the UKBB and find several biologically plausible tissue-trait pairs, including several that replicate, as well as enrichment for interacting tissue pairs. We conclude with a discussion of limitations to our approach, implications for association testing and genetic architecture, and possible future extensions.

Results

Coordinated Polygenic Epistasis. Throughout the paper, we assume a polygenic pairwise epistasis model:

$$y_i = \sum_{j=1}^M G_{ij}\beta_j + \sum_{j < j'} G_{ij}G_{ij'} \Omega_{jj'} + \varepsilon_i, \quad [1]$$

where y_i is the phenotype for individual i , and G_{ij} is the genotype for individual i at marker j , and ε is the residual error and assumed to be independent and identically distributed (i.i.d.) Gaussian. The vector β contains the marginal polygenic effects. $\Omega_{jj'}$ is the pairwise interaction effect of SNPs j not equal to j' , so Ω is the matrix of all pairwise SNP epistasis effects in the genome.

The standard additive model assumes no epistasis, that is, $\Omega = 0$. In this model, SNP j always has the same effect β_j on the phenotype, regardless of genetic background or environment. In the polygenic setting, where there are many more SNPs than individuals ($M > N$), total heritability can still be reliably estimated by the random-effect model Genome-based restricted maximum likelihood (GREML), which models β_j as i.i.d. Gaussian (40).

Epistatic models go further by allowing nonzero Ω . To date, epistatic tests have focused either on candidate SNPs or genome-wide screens for SNP pairs, which reduce $M < N$ and facilitate simple fixed-effect models (21). More recently, random-effect models akin to GREML have become popular for estimating the total size of Ω (i.e., the heritability from pairwise epistasis) (25–28). Another recent approach tests for interaction between a single SNP and a genome-wide kinship matrix, a useful compromise that provides SNP-level resolution and also aggregates genome-wide signal (22).

While these methods are useful for characterizing the existence and impact of epistasis, all are limited by the assumption that β and Ω are independent. We are interested in an orthogonal question: when are β and Ω deeply intertwined by latent interacting pathways? Conceptually, β and Ω encode all relevant information, so the goal is to decode the presence of interacting pathways from these parameters. Concretely, we prove that these pathways exist if, and only if, the CE γ is nonzero, where γ is defined as:




Fig. 1. CE and the EO test with two chromosomes. (A) In the additive model, SNP effects from the two phenotype-increasing pathways are summed to produce the phenotype ($\gamma = 0$). (B) Same as A, except the pathways interact either positively (synergistically, \times , $\gamma > 0$) or negatively (antagonistically, \div , $\gamma < 0$). (C) The EO test considers interaction from traits derived from the even and odd chromosomes in place of the unknown pathways truly driving the interaction.

$$\gamma = \text{Cov}_{j < j'} \left(\beta_j \beta_{j'}, \Omega_{jj'} \right). \quad [2]$$

$$y \sim \alpha_o \text{PRS}_o + \alpha_e \text{PRS}_e + \gamma_{eo} \text{PRS}_o * \text{PRS}_e. \quad [3]$$

Intuitively, $\gamma < 0$ negatively skews the total polygenic effect relative to additivity. This is equivalent to dampening the positive marginal phenotypic effects in aggregate, that is, antagonism between positive main effects. Conversely, $\gamma > 0$ positively skews the population, increasing the probability of extremely high phenotypes. Note that $\gamma = 0$ does not imply that epistasis is absent; rather, it implies epistasis effects do not necessarily systematically align with main effects, as has been implicitly assumed by prior polygenic models of epistasis (25–28).

As a stylistic example, imagine there are two genetically independent pathways that are each sufficient for T2D, one based on body mass index (BMI) and one based on pancreas function. Then $\gamma < 0$ is expected, because high-BMI cases are not likely to also have high pancreas risk, which is rare. On the other hand, if a disease requires high risk across distinct systems, then $\gamma > 0$ is expected. As another stylistic example, if asthma requires both immune components and tobacco exposures, then the impact of a smoking-risk SNP on asthma is greater in the presence of immune risk factors.

We provide a more rigorous exploration of coordination in *SI Appendix*. In particular, we prove that several biologically plausible epistasis models induce CE, including a trans genetic regulation model (32) (*SI Appendix*, section 4.1), a polygenic generalization of molecular buffers (16) (*SI Appendix*, section 4.2), and gene–environment interaction with heritable environment (41) (*SI Appendix*, section 4.3 and Fig. S1). We also show that within-pathway coordination contributes to CE and can cause $\gamma \neq 0$ (*SI Appendix*, Example 3 and Fig. S2). We also study CE as a function of the number of latent pathways; in particular, $\gamma = 0$ results in the mathematically natural many-pathway limit, hence CE typically indicates the existence of a parsimonious set of interacting pathways (*SI Appendix*, Example 4).

The EO Estimator for CE. We have defined our target, the CE γ , as a function of the true genetic effects β and Ω . However, these parameters are not known; even worse, they are high dimensional and cannot be accurately estimated. Nonetheless, we develop a simple and powerful method to test for nonzero γ .

The key idea in the EO test is that randomly defined pathways can act as proxies for true pathways. These random pathways will have a significant interaction if, and only if, there truly are latent pathways that interact. Intuitively, this randomized approach has power to detect true signals because many interacting SNP pairs are appropriately assigned under the random partition—the probability of incorrectly assigning all interacting pairs is negligible. Furthermore, the EO test is calibrated because incorrectly assigned SNP pairs do not interact and hence do not cause bias. In other words, correctly assigned SNP pairs will exist and will suffice to drive interactions between the even and odd proxy pathways. Although the incorrectly assigned SNPs will cause power loss and downward bias for $|\gamma|$, this bias can be corrected post hoc if we assume an infinitesimal epistatic model (*SI Appendix*, Proposition 1). In this case, exactly half of all interacting SNP pairs will be captured by the random SNP partition, which allows us to perfectly estimate the aggregate contribution of SNPs whose pairwise interactions are not captured by the EO partition. Mathematically, these facts are related to other randomized approaches for high-dimensional estimation, including GREML, randomized linear algebra, and compressed sensing (42–46).

We propose estimating γ by regressing on the interaction between PRS built specifically from even-indexed and odd-indexed chromosomes (PRS_e and PRS_o , respectively):

The ordinary least squares estimate $\hat{\gamma}_{eo}$ is the EO estimator of the coordination γ . We prove that $\gamma_{eo} = 0$ if, and only if, $\gamma = 0$ assuming the model in Eq. 1 and that there are many causal SNPs (*SI Appendix*, section 3). Therefore, we can simply use a regression test for $\hat{\gamma}_{eo}$ to test for the existence of CE. Furthermore, we prove that the EO estimate is unbiased and consistent when causal SNPs are in linkage equilibrium (*SI Appendix*, Proposition 1). Full details, assumptions, and guarantees of the EO test are provided in *SI Appendix*.

Our choice of even and odd chromosomes is clearly arbitrary. Crucially, we prove that any partition of chromosomes into two groups without linkage disequilibrium (LD) gives identical results under a perfectly infinitesimal model (*SI Appendix*, section 3.2). Therefore, for concreteness, simplicity, and to emphasize the arbitrariness and flexibility of the EO test, we have chosen to define it by the case where one PRS is built specifically from SNPs on even chromosomes and the other is built specifically from SNPs on odd chromosomes.

With finite genomes, however, results may depend heavily on the precise choice of chromosomes used to build PRS_o and PRS_e . Nonetheless, it is important that SNPs are independent across splits to avoid false positives from nonlinear single-variant effects, that is, dominant or recessive variants. Hence, we always evaluate chromosome-level genome partitions. This is analogous to our choice to exclude the Ω_{ij} terms from CE; in both cases, CE implies nonlinear effects of each individual SNP and the aggregate PRS, but the converse is not true. In practice, we use a more flexible form for the EO test, where we jointly fit and test the interactions between all (22 choose 2) chromosome-specific PRS, which simultaneously obviates LD concerns and the arbitrariness of assigning chromosomes based on their canonical indices. Nonetheless, our method does not assume the true causal pathways are located in contiguous or independent genomic regions; rather, this is a condition on our chosen partition of the genome that is used in the EO test. In particular, both CE and EO remain valid when causal SNPs for each pathway are distributed across the genome randomly, and/or when SNPs contribute to multiple pathways (*SI Appendix*, section 3.3).

EO Test Can Distinguish CE from Population Structure and Uncoordinated Epistasis. To examine the properties of the EO test, we simulate data under three biologically plausible genetic architectures: additivity, isotropic (i.e., uncoordinated) epistasis, and CE. For each architecture, we also test the impact of population structure, assortative mating, and adjustment with principal components (see *Materials and Methods*). In each case, we perform a GWAS and then build ordinary PRS as well as PRS based only on variants from odd chromosomes and even chromosomes (PRS_o and PRS_e , respectively). We then test PRS in a second dataset.

We first tested the correlation between PRS_o and PRS_e , called θ_{eo} , which is related to an existing method of assortative mating (47). We found that the test for $\theta_{eo} \neq 0$ reliably indicated the presence of assortative mating or uncorrected population structure. Note that after adjusting for PCs, the test for $\theta_{eo} \neq 0$ was roughly null in the presence of population structure (Table 1).

However, our main focus is on the EO test for $\gamma_{eo} \neq 0$. We found that the EO test was calibrated under both the additive and isotropic interaction models, as expected, with false-positive rates near 0.05 at a nominal $P < 0.05$ threshold. In particular, this is a clear demonstration of a setting where significant pairwise epistasis exists and yet CE does not exist, thereby drawing a clear distinction between CE and the more general concept of epistasis. Further, the EO test has power of roughly 95% under the CE model (Table 1).

Table 1. Polygenic simulations under additivity, isotropic epistasis, or CE assuming random mating, population structure, or assortative mating

	Simulation characteristics			$\theta_{eo} : PRS_e \sim PRS_o$		$\gamma_{eo} : Y \sim PRS_e * PRS_o$	
	Population structure	PC adjusted	Assortative mating	χ^2 (SD)	PR ($\alpha = 0.05$)	χ^2 (SD)	PR ($\alpha = 0.05$)
Additive	N	N	N	1.0 (1.4)	0.05	1.0 (1.4)	0.05
	Y	N	N	81,789.4 (35,386.7)	0.99	101.1 (50.5)	0.95
	Y	Y	N	1 (1.4)	0.05	1.0 (1.4)	0.05
	N	N	Y	40.7 (13.0)	1.00	1.1 (1.6)	0.06
Isotropic epistasis	N	N	N	1.0 (1.4)	0.05	1.0 (1.4)	0.05
	Y	N	N	81,698.3 (35,422.0)	0.99	101.0 (51.1)	0.95
	Y	Y	N	1.0 (1.4)	0.05	1.0 (1.4)	0.05
	N	N	Y	40.6 (12.9)	1.00	1.2 (1.7)	0.08
CE	N	N	N	1 (1.4)	0.05	15.4 (8.3)	0.96
	Y	N	N	81,717.4 (35,394.7)	0.99	109.6 (57.1)	0.96
	Y	Y	N	1.0 (1.4)	0.05	14.2 (8.2)	0.93
	N	N	Y	40.8 (12.9)	1.00	43.4 (14.9)	1.00
Oracle EO partition	N	N	N	1.0 (1.4)	0.05	132.6 (25.2)	1.00

Mean, standard deviation (SD), and positive rate (PR) at nominal $P < 0.05$ are shown for estimates of θ_{eo} and γ_{eo} . We display results averaged over 10,000 simulations using the baseline parameters $n = 10,000$, $h^2 = 0.5$, and $\gamma = 0.2$. For all models, the ordinary EO test was performed by randomly assigning SNPs to either the “Even” or “Odd” group, except in the “Oracle EO Partition” setting where SNPs were grouped according to the true, generative pathways they causally effect.

We did not observe substantial false positives in the EO test under assortative mating. However, we did observe high false-positive rates for the EO test under uncorrected population structure (0.95), although after correcting for PCs, the test is calibrated with an empirical false-positive rate near 0.05 (Table 1). In practice, we recommend adjusting for population-structure proxies when running the EO test, as is standard in human genetics.

We then sought to more comprehensively profile the behavior of the EO test across a range of simulation parameters. First, we varied the sample size, n (SI Appendix, Fig. S3A). As expected, power at nominal $P < 0.05$ grew with sample size, from roughly 5% at $n = 1,000$, up to roughly 95% at $n = 10,000$ (our baseline), and reaching nearly 100% at $n = 20,000$. Next, we varied the additive heritability, which affects the EO test power because it governs the accuracy of the PRS estimates (SI Appendix, Fig. S3B). As expected, power increased with heritability, growing from 35% at $h^2 = 0.1$ to nearly 100% at $h^2 = 0.7$ (our baseline is $h^2 = 0.5$). Finally, we varied the strength of CE parameter, γ . As expected, power grew with $|\gamma|$, reaching 100% when we doubled the parameter relative to our baseline (SI Appendix, Fig. S3C). Furthermore, when $\gamma = 0$, we recovered calibrated tests with roughly 5% false-positive rate. Interestingly, our simulations had similar power regardless the sign of γ , as expected because the sign of the phenotype is arbitrary in these simulations.

An essential feature of the EO test is that it is unbiased even when the even/odd partition of SNPs is chosen randomly with respect to partition of SNPs into the causal interacting pathways (under conditions stated above). Nonetheless, the power to detect CE increases when the even/odd partition more closely aligns with the causal partition; we demonstrate this in our simulations by directly matching the two partitions (see Materials and Methods). As expected, the EO test power increases from 95 to 100% when we use this more precise SNP partition (Table 1).

Our above simulations used minor allele frequencies (MAFs) drawn i.i.d. from a uniform (0.01,0.5) distribution. To assess a more realistic MAF spectrum, we repeated our baseline simulations using MAFs drawn randomly from the MAF spectrum we observed in UKBB, excluding rare variants with $MAF < 0.01$ for simplicity. Nonetheless, the simulation results were nearly identical (SI Appendix, Table S1).

Finally, we assessed the EO test under third-order epistasis (SI Appendix, Table S2). First, when interacting SNP triples were

distributed uniformly at random across the genome, as in the isotropic model of pairwise epistasis, the EO test had positive rate near 0.05. This is appropriately calibrated behavior in the informal sense that there is no relationship between main effects and epistasis effects and in the formal sense of the definition of CE (Eq. 1, above). Second, we simulated a three-pathway model where the product of all three pathways, but not any two, had an effect on the phenotype. In this setting, the main and interaction effects are related, but the formal CE definition is still 0. In the simulation, we observe power near 8% at $P < 0.05$ significance. This is expected in our simulations that use finite genomes; for any simulated dataset, the pathways will have some nonzero mean that induces pairwise interaction, resulting in an average γ estimate that is zero but overdispersed (this is related to similar issues in GREML with finite numbers of causal SNPs) (48). Overall, the EO test is calibrated under isotropic third-order epistasis and has some signal under higher-order coordinated epistasis.

Testing for CE with the EO Test in the UKBB. We next tested for CE with the EO test in the UKBB (49). We studied 21 quantitative and 5 binary traits (SI Appendix, Table S3) chosen to represent a range of trait classes including anthropometric, disease, and blood traits. We specifically analyzed the subjects classified as “White British” and filtered out related individuals to minimize population structure bias while retaining large sample sizes (max $n = 342,816$). We calculated PRS for each chromosome for each sample using a standard thresholding approach (50) (see Materials and Methods). The total PRS is then the sum of the individual chromosomes’ PRS. We evaluated a range of PRS P value thresholds and corrected for this source of multiple testing using a Benjamini–Hochberg false discovery rate (FDR) threshold of 10%. We used 10-fold cross-validation across individuals to minimize bias from in-sample overfitting.

Unlike in perfectly infinitesimal models, in real data, the EO test depends on the specific partition of chromosomes used to estimate γ . To minimize bias from choosing a single split, for example, even versus odd, in practice we jointly test interactions across all distinct pairs of chromosome-specific PRS with an F-test (see Materials and Methods). Following common practice (51–56), we report the significant results per phenotype.

We discovered 18 traits with significant CE (Table 2), including the complex diseases asthma, cardiovascular disease (CVD), eczema, and T2D. We also detected CE for the complex quantitative traits basal metabolic rate, bone mineral density (BMD), lung function (forced expiratory volume [FEV]/ forced vital capacity [FVC]), and height, as well as nine blood traits: glucose, low density lipoprotein (LDL), platelet distribution width, mean platelet volume (MPV), red blood cell distribution width, spheroid cell volume, triglycerides (TG), and counts of monocytes, lymphocytes, and platelets (PLAT). Taken together, these results show that CE contributes to the genetic architecture of multiple complex traits.

To assess the potential for impact by population structure and assortative mating, we calculate the correlation between each pair of chromosome-specific PRS (θ_{eo}). In simulations, we found that $\theta_{eo} > 0$ indicates the presence of structured genotypes, which may be due to assortative mating and/or uncorrected population stratification. Although correction for PCs was sufficient for calibrated inference in simulations (Table 1), we nonetheless studied the relationship between θ_{eo} and γ_{eo} as a metric for possible confounding of the EO test by stratification. This caution is motivated by the complexities of population structure in massive datasets like the UKBB (57–60). However, we found no evidence that population stratification drove our CE inference; larger values of θ did not correlate with γ values or their corresponding P or $\log(P)$ values (SI Appendix, Fig. S4). Overall, even with these conservative extra tests we did not identify evidence that our EO test for CE was substantially confounded by uncorrected structure or assortative mating.

Table 2. EO test results for 18 traits with significant CE in the UKBB

Phenotype	PRS type	τ	%VE	FDR
Basal metabolic rate	Intern quant	1.00E-08	0.74	0.02
BMD	Intern quant	0.001	3.39	2.30E-03
FEV/FVC	Intern quant	1.00E-07	0.52	0.02
Glucose	Intern quant	1.00E-05	2.28	0.02
Height	Intern quant	1.00E-08	2.80	4.00E-03
LDL	Intern quant	1.00E-07	1.30	7.30E-04
Lymphocyte no.	Intern quant	1.00E-05	1.05	2.00E-03
Monocyte no.	Intern quant	1.00E-08	0.95	8.60E-04
PLAT no.	Intern quant	0.0001	5.81	0.08
PLAT distribution width	Intern quant	1.00E-08	3.99	2.60E-04
PLAT volume	Intern quant	0.0001	9.06	2.90E-11
RBC distribution width	Intern quant	1.00E-06	2.04	0.01
Spheroid cell volume	Intern quant	1.00E-06	1.73	0.07
Triglycerides	Intern quant	0.0001	1.55	9.70E-03
Asthma	Intern bin	0.0001	2.79	6.60E-04
Cardiovascular	Intern bin	1.00E-08	3.90	6.50E-05
Eczema	Intern bin	1.00E-08	2.18	8.10E-05
T2D	Intern bin	0.001	3.64	9.40E-08
Corp. hemoglobin	Extern quant	—	19.17	0.04
LDL	Extern quant	1.00E-05	2.19	0.01
PLAT no.	Extern quant	—	16.90	5.20E-03
PLAT volume	Extern quant	—	29.41	1.10E-25

Significant traits based on either internal or external PRS are shown (FDR < 0.1); results for all traits are shown in SI Appendix, Table S3. We summarize evidence for CE per trait using Benjamini–Hochberg FDR across nine PRS P -value thresholds; for single-threshold external PRS, FDR is just an ordinary P value. “Intern” indicates that the PRS was calculated using the cross-validation method with UKBB data (Materials and Methods). “Extern” indicates that the PRS was calculated using GWAS summary statistics from external datasets (Materials and Methods). τ is the marginal SNP P -value threshold used to construct the PRS and that minimizes the EO test P value. %VE describes the variance explained by the chosen PRS across all samples, quant, quantitative.

As a final assessment, we repeated our EO test for 100 random permutations of the PRS among samples. The resulting EO-test FDR was null, as expected (SI Appendix, Fig. S5). Furthermore, we used these permutations to construct empirical P values for the minimum EO test across all PRS thresholds as an alternative to FDR correction. This gives qualitatively similar results to our primary FDR-based analysis, although computational costs bound the minimum attainable empirical P values (SI Appendix, Fig. S6). These analyses do not rule out the possibility of subtle confounding, but they do provide further support for the calibration of the EO test in practice.

Replicating EO Test Results with External PRS. Although we have cross-validated our PRS, we wished to additionally check that our results are not specific to a single population by overfitting dataset-specific confounders (57, 58, 60). We investigated this by constructing PRS using external summary statistics to remove the potential impact of cohort-specific artifacts. We selected eight traits that have large external GWAS summary statistics: asthma, T2D, CVD, height, BMI, TG, educational attainment, and LDL; additionally, we constructed PRS for nine other blood traits that were only available for a specific P -value threshold (see Materials and Methods and SI Appendix, Table S3). We tested replication for the 10 traits in this list with internally significant CE.

We applied the EO test as described above and found that CE replicated for LDL, PLAT, and MPV (Table 2). This validates CE in the sense that our results are not specific to the main effect–size estimates derived from the UKBB dataset. However, significant CE did not replicate for several traits. Broadly, the CE P values were less significant using external PRS, likely due to a combination of winner’s curse and the fact that external studies generally analyze datasets smaller than UKBB and do not perfectly match UKBB in terms of environment and genetic background.

To further strengthen confidence in CE, we performed an alternate replication analysis by directly comparing the internal and external PRS–interaction effect estimates across all pairs of chromosomes (Fig. 2). This test identified highly significant CE replication for 7/10 tested traits. This test is more powerful than simply applying the EO test to the external PRS because it tests a narrower alternative hypothesis. Specifically, the EO test asks whether any chromosome pairs interact, whereas the direct replication test asks whether the external interaction effect estimates correlate with their internal counterparts. Conceptually, this is related to one- versus two-sided tests, as the former add power when prior knowledge of the estimand’s sign is available.

In addition to strengthening confidence in the existence of CE, these replication analyses also suggest that CE can be reliably tested using internally constructed PRS in the future. This can be essential for applications to under-studied traits, populations, or environmental contexts (61–65).

Tissue-Specific Coordination in UKBB. Having demonstrated the existence of CE for several traits, we now consider the possibility that interacting pathways are enriched in trait-relevant tissues. Specifically, we test for tissue-specific enrichment of CE across the above 26 traits and 13 tissue-specific genomic annotations: 7 based on specifically expressed genes (adipocytes, blood cells, brain, hippocampus, liver, muscles, and pancreas) (66, 67) and 6 based on tissue-specific chromatin-marker patterns (adipose, brain, hippocampus, liver, pancreas, and skeletal muscle) (68, 69) as used in ref. 51.

For each tissue–trait pair, the tissue-specific EO test asks whether the coordination mediated through a specific tissue exceeds the genome-wide average. This test is conducted by first creating a tissue-specific PRS for each chromosome (TPRS_{*i*}) and testing it for interaction with the standard PRS on other chromosomes (PRS_{*j*}, $i \neq j$). To focus our test on truly tissue-specific




Fig. 2. CE replication in UKBB. For each of the 10 phenotypes with internally significant CE and with available external PRS, we plot internal interaction estimates (x-axis) against external estimates (y-axis) for each pair of chromosomes. Chromosomes that do not contribute to a particular PRS are excluded. We use the PRS P -value threshold for each PRS to minimize its CE P value; this does not cause inflation because we are not testing the size of these interaction estimates but rather their correlation. This indirect replication test is Bonferroni significant for 7/10 traits when correcting for the number of tested traits. Effect sizes are displayed after each chromosome-specific PRS is centered and scaled, and the P values and red lines correspond to regressions with intercepts constrained to 0.

effects, rather than global effects that happen to be tagged by tissue-specific regions, we adjust for the ordinary chromosome-level PRS interactions ($PRS_i * PRS_j$) as well as the main effects of each chromosome's PRS and TPRS. We use a $P = 0.05$ threshold with a conservative Bonferroni correction to account for the multiple tested tissues. We only evaluate phenotypes with significant ordinary CE (Table 2).

We found 53 instances of tissue-specific CE enrichment using internal PRS (Table 3 and *SI Appendix, Table S4*). This includes nine tissues enriched for LDL CE, several of which are essentially positive controls: the liver is a key regulator of LDL metabolism and adipocytes and muscles are primary sinks for the triglycerides carried by LDL (70). Further confirming these results, adipocyte CE-enrichment replicates when using external LDL PRS.

Another biologically plausible set of tissue-trait pairs is for MPV, which also had the strongest ordinary CE of any trait. First, we find significant enrichment for liver, consistent with its known role as the main producer of the main regulator of PLAT production, thrombopoietin (TPO) (71). Second, we find significant CE enrichment for muscles, which also modestly produce TPO (71); furthermore, PLATs are important in healing muscles. Third, we found suggestive CE with brain tissues—while the underlying biology is less obvious, there is evidence that TPO affects brain development (72). Additionally, recent complementary studies have found evidence that brain tissue plays a role in MPV (73). All three of these tissues replicate for CE enrichment when using external PRS.

We also found tissue-specific CE for complex disease. For example, CVD has CE enrichment for brain, liver, muscle, and adipose and has a clear link to CVD; several brain regions have been implicated in the genetic basis of BMI, a CVD risk factor (74, 75), and liver, muscle, and adipose all are deeply involved in energy homeostasis (76). While none of these tissues replicate, this likely reflects lack of power as even ordinary CE was not significant in either of our external replication tests.

To add further confidence in tissue-specific CE, we tested five permutations of the TPRS across samples. We observed that the resulting P values were appropriately null (*SI Appendix, Fig. S7*).

In general, the tissue-specific EO test can improve power over the ordinary EO test when the correct tissue is identified, despite the fact that the TPRS almost necessarily explain less trait variation than the overall PRS. This is consistent with our simulation result where the EO test power increased under the Oracle approach that partitions SNPs into the true pathways, even though this yields PRS with weaker predictive power. For example, the empirical EO test P value for spheroid cell volume is 0.03 (*SI Appendix, Table S3*), but its P value for blood-cell CE is 0.003; however, the overall PRS explains 1.5% of trait variance while the blood-cell TPRS explains only 0.5%. Similar properties can be observed for many plausible tissue-trait pairs, for example, in the liver-specific CE for glucose and triglycerides. This is further evidence that CE is truly enriched in specific tissues for some traits and illustrates that CE captures a genetic axis that is partially distinct from additive variance.

Tissue-Pair CE in UKBB. We next ask whether CE can be detected between pairs of tissues based on the hypothesis that pathways may be tissue specific. Rather than test the interaction between tissue-specific PRS_i and global PRS_j , we now test for interaction between tissue 1-specific PRS_i and tissue 2-specific PRS_j . For tissue-specific CE, we adjust for the same covariates as above and now additionally adjust for $PRS_i * TPRS_j$ for both tested TPRS and all $i \neq j$. In particular, these tissue-tissue interaction tests are statistically independent of the EO tests and the tissue-specific tests in the above sections under the null hypothesis for tissue-tissue interaction.

If tissue-tissue CE exists, it will likely cause tissue-specific CE; hence, to improve power, we only test tissue-trait pairs that are significant in the above TPRS test (internal or external, Table 3). This reduction in testing dimension is particularly important here because the number of tests scales by the number of evaluated tissues squared. As each phenotype now has a different number of tests, and this test is the most complex and high dimensional we consider, we use an aggressive Bonferroni correction and adjust for all tested trait-tissue-tissue triples, performed separately for internal and external PRS.

Table 3. Tissue-specific CE in the UKBB

Phenotype	Tissue	PRS type	TPRS %VE	P value
BMD	Liver	Intern quant	1.07	1.80E-03
FEV/FVC	FrankeLiver	Intern quant	0.28	1.30E-06
Glucose	Franke.muscles	Intern quant	0.07	2.40E-04
Height	Franke.blood.cells	Intern quant	0.03	4.00E-12
LDL	Franke.blood.cells	Intern quant	0.10	1.10E-04
LDL	Franke.adipocytes	Intern quant	1.77	2.60E-04
LDL	Franke.brain	Intern quant	1.10	6.60E-04
LDL	Franke.hippocampus	Intern quant	0.38	6.50E-06
LDL	Franke.liver	Intern quant	3.02	1.60E-03
LDL	Franke.muscles	Intern quant	0.26	1.30E-03
LDL	Franke.pancreas	Intern quant	1.95	2.00E-03
LDL	Brain	Intern quant	4.71	2.90E-04
LDL	Hippocampus	Intern quant	4.24	1.60E-05
Lymphocyte no.	Franke.hippocampus	Intern quant	0.90	2.10E-03
Monocyte no.	Franke.hippocampus	Intern quant	0.15	4.30E-04
PLAT no.	Franke.adipocytes	Intern quant	1.03	1.10E-03
PLAT no.	Franke.hippocampus	Intern quant	0.39	7.60E-05
PLAT no.	Franke.muscles	Intern quant	0.17	1.20E-03
PLAT no.	Skeletal_muscle	Intern quant	2.92	1.80E-03
PLAT distribution width	Franke.brain	Intern quant	0.04	5.30E-04
PLAT volume	Franke.blood.cells	Intern quant	0.36	1.80E-04
PLAT volume	Franke.adipocytes	Intern quant	1.83	9.60E-07
PLAT volume	Franke.brain	Intern quant	0.00	1.40E-09
PLAT volume	Franke.hippocampus	Intern quant	0.02	6.80E-13
PLAT volume	Franke.liver	Intern quant	1.27	1.30E-03
PLAT volume	Franke.pancreas	Intern quant	0.62	1.30E-04
PLAT volume	Skeletal_muscle	Intern quant	0.62	3.60E-03
Sphered cell volume	Franke.blood.cells	Intern quant	0.51	3.00E-03
Triglycerides	Franke.pancreas	Intern quant	0.49	6.60E-04
Cardiovascular	Franke.brain	Intern bin	0.01	8.20E-05
Cardiovascular	Franke.liver	Intern bin	0.24	1.00E-04
Cardiovascular	Franke.muscles	Intern bin	0.04	2.10E-04
Cardiovascular	Adipose	Intern bin	0.26	1.70E-03
Eczema	Franke.muscles	Intern bin	0.07	2.10E-04
LDL	Franke.adipocytes	Extern quant	0.44	2.30E-03
PLAT volume	Franke.adipocytes	Extern quant	0.08	1.60E-08
PLAT volume	Franke.brain	Extern quant	0.02	2.90E-03
PLAT volume	Franke.muscles	Extern quant	0.06	4.70E-05
PLAT volume	Franke.pancreas	Extern quant	0.00	7.40E-05
PLAT volume	Adipose	Extern quant	0.13	1.10E-03
PLAT volume	Brain	Extern quant	0.02	3.30E-03
PLAT volume	Liver	Extern quant	0.00	4.60E-04

Trait–tissue pairs with significant tissue-specific CE are shown; results for all tested trait–tissue pairs are in [SI Appendix, Table S4](#). For parsimony, we list here only the most significant tissue per trait as well as all significant tissues for traits discussed in the main text. The prefix “Franke.” indicates that the tissue annotation is derived from Franke laboratory (66, 67). TPRS %VE is percent phenotype variance explained by the tissue-specific PRS. *P* values for the tissue-specific CE estimate are Bonferroni adjusted for testing across multiple tissue annotations. “Intern” indicates that the PRS was calculated using the cross-validation method with UKBB data (*Materials and Methods*). “Extern” indicates that the PRS was calculated using GWAS summary statistics from external datasets (*Materials and Methods*).

We find five Bonferroni-significant examples of tissue–tissue CE (Table 4 and [SI Appendix, Table S5](#)). Two highlight the role of muscles in PLAT, one interacting with adipocytes (internal is significant; external $P = 0.02$), and another with hippocampus (external is significant; internal $P = 0.01$); despite the nominal internal/external replication, the biological interpretation of these tissue pairs is unclear. Two other tissue pairs highlight brain interaction with blood cells and liver for MPV; again, the role of the brain is not clear, but both blood cells and liver have strong connection to MPV biology. Finally, we find evidence for liver–pancreas CE driving FEV/FVC1, which is particularly striking as the *P* value for this tissue–tissue CE exceeds either *P* value for these tissues’ individual CE.

Discussion

Epistatic interactions between genetic variants can have profound influences on phenotypic distributions and evolutionary landscapes (39). While there are numerous examples of epistasis in model organisms and in-vitro studies (2, 7–11), especially between pairs of genes, relatively few examples have been observed for complex polygenic traits in humans. In this work, we propose a specific form of epistasis called CE, in which many SNPs in one pathway jointly skew the effects of SNPs in other pathways. This model is inspired by recent theories of human genetic architecture (29, 31). These pathways could represent distinct biological concepts, including different tissues (e.g., neurons and astrocytes in

Table 4. Tissue-Pair CE in the UKBB

Phenotype	Tissue 1	Tissue 2	PRS type	<i>P</i> value
FEV/FVC	Franke.liver	Franke.pancreas	Intern quant	2.20E-05
PLAT no.	Franke.adipocytes	Franke.muscles	Intern quant	1.50E-04
PLAT volume	Franke.blood.cells	Franke.brain	Intern quant	1.80E-05
PLAT volume	Franke.hippocampus	Franke.liver	Intern quant	2.30E-04
PLAT no.	Franke.hippocampus	Franke.muscles	Extern quant	2.00E-04

Trait-tissue-tissue triples with significant tissue-tissue CE are shown; results for all tested triples are shown in [SI Appendix, Table S5](#). We use a $P = 0.05$ significance threshold after Bonferroni adjustment for the total number of tested trait-tissue-tissue triples. The prefix “Franke.” indicates that the tissue annotation is derived from Franke laboratory (66, 67).

Alzheimer’s disease) (77), different systems within a tissue [e.g., stress and immune response in amyotrophic lateral sclerosis (78)], or unknown subtypes of a disease with partially distinct genetic risk factors (29, 79). Importantly, CE is fundamentally different from previous polygenic epistatic models; in particular, existing models assume a purely uncoordinated form of epistasis.

We next developed a test, called the EO test for CE, based on testing the interaction between PRS constructed from independent chromosomes. We showed via extensive analytical examination and empirical simulation that our test will detect CE even though the pathways are unknown. We further show that our test is robust to (at least) moderate population structure and assortative mating. The sign of γ determines which direction CE skews the polygenic contribution, which is closely connected to the probability of observing extreme phenotypes in either direction. In the case where main SNP effects are positive, positive/negative epistasis corresponds to synergy/antagonism between SNPs.

We observed evidence of CE from the EO test for many phenotypes in the UKBB including blood, anthropometric, and disease phenotypes. We interpret this as evidence that CE may affect a range of phenotypes across many domains. To further investigate CE, we examined PRS restricted to genomic regions annotated to specific tissues and found biologically plausible results. Interestingly, for several phenotypes, the tissue-specific EO *P* value increased in significance despite the fact that the total variance explained decreased. This suggests that tissue-specific annotations are enriched for true latent pathways in a trait-relevant way.

There are several limitations to the EO test and potential avenues for future improvements. First, more accurate PRS will improve the power of the test, either through more sophisticated models for PRS or through larger external reference datasets. Second, the EO test uses a random split of SNPs into even and odd PRS, which only partially tags the true latent pathways and in turn reduces power. Using tissue-specific PRS is a step toward identifying which genomic regions harbor the greatest CE. Nonetheless, current variant-level annotations are imprecise, and other annotations may be more relevant depending on researcher interest and underlying trait biology. For example, annotations for cell-type specificity, gene ontology, inferred coexpression networks, or other functional categories related to methylation and chromatin state may be worth exploring. Third, the Bonferroni correction we use is overly conservative due to the strong correlation between splits (and tissues); some recent theoretical progress on testing multiple exchangeable hypotheses may enable power gains in the future (80). Fourth, our random chromosome-splitting procedure itself has noise, and jointly testing all 210 pairs of chromosome-specific PRS interactions may add power; however, this is infeasible for modest data sets or for the extensions to test annotation-specific CE enrichment. Fifth, as with any test of intInteraction, phenotypic scale can impact results. However, we show that the sign and existence of CE is preserved under rescaling, either on average or under smooth transformation, and

hence CE is (approximately) “essential” (81). Sixth, we have not yet directly connected CE estimates to their impact on genetic architecture and heritability, though we do outline a path forward in [SI Appendix](#). Seventh, the possibility of subtle bias from population structure can never be fully ruled out; nonetheless, the EO test was calibrated in simulations with population structure (Table 1), we utilized standard population structure corrections, and we performed several robustness checks that add confidence in our CE inferences ([SI Appendix, Figs. S4–S7](#)). Eighth, the chromosome-level partitions we use will yield conservative EO estimates for CE when epistasis effects are enriched within chromosome, as indicated by prior studies in model organisms (5, 82); formally, this model violates the assumption of uniformly distributed epistasis that is sufficient for unbiased CE estimation ([SI Appendix, Proposition 1](#)). Finally, the EO test only detects statistical CE, not biological CE. For example, a disease that has two subtypes with distinct, purely additive genetic bases will exhibit CE even though neither subtype has any epistasis at all. On the other side of the coin, though, CE can be a useful way to demonstrate the existence of unmodeled subtypes.

Going forward, there are several straightforward applications of CE that may inform novel dimensions of complex trait biology. One could examine SNP×PRS interaction or gene–PRS interaction as in TWAS (83, 84). This will add substantial power when the main effect of a highly penetrant gene or SNP depends heavily on genetic background. Second, expanding to other traits and annotations could help refine the pathways involved and increasingly home in on causal mechanisms. For example, annotation of SNPs based on association with environmental factors, such as diet, exercise, smoking, or stress, can inform the nature of G×E interactions. Third, SNPs in genes known to influence specific pathways from in-vitro molecular assays could be examined at the population scale in the full organismal context. For example, the genes in the ribosomal quality-control (QC) pathway are known to influence degradation of misfolded proteins (19, 85), and this pathway may have important interactions with disease- or trait-specific coding variants segregating in the population. Fourth, we chose to develop CE and the EO test in UKBB because we are motivated by recent theories on genetic architectures of complex human traits, but we hope these ideas will also be applied to model systems to move toward tests of biological epistasis and to leverage prior knowledge. Finally, CE may have implications for medical interventions, which are generally pathway specific. For example, if asthma exhibits strong positive CE between immune and lung pathways, intervening on either pathway will be sufficient to reduce disease burden; conversely, if LDL exhibits negative CE across liver- and BMI-driven pathways, it may be important to simultaneously address both pathways in order to control blood cholesterol.

Materials and Methods

Simulation Description. Phenotypes and genotypes were generated under three phenotypic models—additive, isotropic (uncoordinated) epistasis, and

CE—and three genotypic models— independent samples, assortative mating, and two populations with $F_{ST} = 0.1$ (SI Appendix, section 6). In brief, the additive model had no epistasis. The isotropic model had epistasis heritability of 10%, contributed by 0.1% of SNP pairs having interaction effects drawn independently of the main effects. The CE model had two additive pathways that interact to explain 10% of the phenotypic variation, where each pathway is driven by a different subset of causal SNPs. The genotypes were either drawn i.i.d. over samples, or were drawn after one generation of assortative mating, or were drawn from a Balding–Nichols model (86).

For each simulation, genotypes were generated for 10,000 individuals at 1,000 SNPs. Half of the individuals were randomly selected for the training set, in which the phenotype y was regressed on each genotype g_j to calculate the estimated effect, $\hat{\beta}_j$, of each SNP. We then constructed PRS in the remaining individuals ($PRS = \sum_{j=1}^{1000} g_j \hat{\beta}_j$). PRS were also calculated separately for “Even” and “Odd” SNPs, sets that were defined arbitrarily as there are no chromosomes in these simulations.

We then tested for EO correlation by regressing odd PRS on the even PRS and tested CE by regressing the phenotype on the interaction between the odd PRS and the even PRS, which is the EO test. The test was also performed conditionally on the first genetic PC in the case where population structure is simulated. Finally, we evaluated the EO test when the “Even” and “Odd” SNP sets were chosen to perfectly coincide with SNPs corresponding to the causal pathways one and two, respectively, which uses oracle information that is not known in practice.

A full description of the simulations is provided in SI Appendix, section 6, and all code used is provided at <https://github.com/nadavrap/CoordinatedInteractions>.

UKBB Data Description. Data were obtained from the UKBB project (49). We used imputed genomic data version 3. We used the same phenotypic transformations as defined in ref. 51. Education-level categories (data field no. 6138) were transformed to numerical values according to ref. 87 as follows: “College or University degree” = 20 y; “Advanced level/Advanced Subsidiary level (A levels/AS levels) or equivalent” = 13 y; “Ordinary level/General Certificate of Secondary Education (O levels/GCSEs) or equivalent” = 10 y; “Certificate of Secondary Education (CSEs) or equivalent” = 10 y; “National Vocational Qualification (NVQ) or Higher National Diploma (HND) or Higher National Certificate (HNC) or equivalent” = 19 y; “Other professional qualifications, e.g., nursing or teaching” = 15 y; “None of the above” = 7 y; “Prefer not to answer” = missing. BMI was used as computed by UKBB (data field no. 21001). Heel BMD T-score was computed as the sum of the left and right heel BMD T-score (data fields no. 4106 and 4125). Waist–hip ratio was computed as the ratio of waist circumference and hip circumference (data fields no. 48 and 49). FEV1–FVC ratio phenotype was computed as the ratio between data field FVC (no. 3062) and “FEV1 in 1 second” (FEV1) (no. 3063). T2D phenotype was based on data field “Diabetes diagnosed by doctor” (no. 2443) where values such as “Prefer not to answer” and “Do not know” were considered as missing data. Eczema allergy was based on data field “Blood clot, DVT, bronchitis, emphysema, asthma, rhinitis, eczema, allergy diagnosed by doctor” (no. 6152) where the value “Hay fever, allergic rhinitis, or eczema” was considered to define cases, “Prefer not to answer” as missing, and the rest as controls. Asthma was based on the same data field, where only “Asthma” values were considered as cases. High cholesterol was determined based on samples with data code 1,473 (“high cholesterol”) in field “Noncancer illness code, self-reported” (no. 20002), all other samples were defined as controls. Cardiovascular cases were defined according to ref. 88 as samples having any of the following codes in field “Noncancer illness code, self-reported” (no. 20002): 1,065 (hypertension); 1,066 (heart/cardiac problem); 1,067 (peripheral vascular disease); 1,068 (venous thromboembolic disease); 1,081 (stroke); 1,082 (transient ischemic attack [tia]); 1,083 (subdural hemorrhage/hematoma); 1,425 (cerebral aneurysm); 1,473 (high cholesterol); and 1,493 (other venous/lymphatic disease).

Sample Selection. For sample QC, we filtered out samples who were not classified as “white British” (data field 21,000), or with discordance of genetic sex with declared sex (data fields 22,001 and 31). We filtered out 74,634 samples due to relatedness based on the .dat file provided by the UKBB. Using a greedy approach, samples were removed by number of remaining related samples from largest to lowest until no more related samples were left. In total, 342,815 subjects (158,561 males and 184,254 females) were included in further analysis.

For genotype QC, we filtered out variants with MAF lower than 0.1%. In total, 16,427,045 variants were included. These analyses were performed in plink version 2.00a (89).

Internal PRS.

Cross-validation by phenotype. For each phenotype, we split subjects into 10 folds. For each fold, we estimated effect size of each variant using the remaining nine folds, and we then use these effect sizes to construct a PRS for the held-out fold. Quantitative phenotypes were normalized by first removing outliers defined as samples outside 1.5 times the interquartile range above the upper quartile or below the lower quartile and then performing quantile normalization. Although such normalization will inevitably shrink the non-Gaussianity induced by true CE signal, we view this as an important normalization step to remain conservative. Moreover, we prove smooth monotone transforms cannot create, remove, or change the sign of CE in SI Appendix.

PRS estimation consisted of two steps: 1) effect size estimation; 2) compute PRS using selected variants. The first step is equivalent to performing a standard GWAS. For that purpose, we add as covariates the sex, age, center of assessment, genotyping batch, first 40 PCs as provided by UKBB, and BMI (except when BMI was the target phenotype). Effect sizes were estimated and tested with linear regression for quantitative traits. For binary traits, we used logistic regression and a computationally efficient approximate Firth correction (90, 91) using “cc-residualize no-firth” in plink2. In the second step, PRS were computed for each fold for the 10% of samples that were not used to estimate the effect sizes. We computed PRS using a range of 9 P -value thresholds: 1.0, 0.1, 0.01, 0.001, 0.0001, 1e-05, 1e-06, 1e-07, and 1e-08.

External PRS. An alternative approach is to use effect sizes estimated from external datasets. This must be a cohort with similar ancestry in which the UKBB subjects were not included. We used external summary statistics included in LD Hub (92) from the following studies: height [GIANT consortium (93)]; LDL and triglycerides [Global Lipids Genetics Consortium (94)]; BMI [GIANT consortium (95)]; educational attainment [Social Science Genetic Association Consortium (87)]; T2D [DIAGRAM Consortium (96)]; Asthma [A GABRIEL Consortium (97)]; and CVD [CARDIoGRAMplusC4D Consortium (98)]. For nine additional phenotypes, summary statistics were available for only a single P -value threshold and downloaded from the Polygenic Score Catalog (99): eosinophil counts, reticulocyte counts, PLAT counts, lymphocyte counts, red cell counts, mean corpuscular hemoglobin, white cell counts, PLAT volume, and monocyte counts (100).

Tissue-Specific Variants. Variants based on tissue-specific gene expression were defined as variants that fall within a 100 kb of the genes that were classified as differentially expressed genes by refs. 66 and 67, similarly to ref. 51. Variants were based on chromatin markers from the Roadmap project (68, 69). Any genomic location annotated with any of the DNase, H3K27ac, H3K36me3, H3K4me1, H3K4me3, and H3K9ac was considered as “open chromatin” region. To evaluate PRS based on tissue-specific variants, we used variants in the intersection of imputed variants from UKBB with MAF > 0.01% and variants that were classified as tissue specific. We did not use the clumped set of variants, as the intersection of the two is small.

Data Availability. The code to generate internal and external PRS, to perform EO CE and tissue-specific tests, and to reproduce simulations can be found at <https://github.com/nadavrap/CoordinatedInteractions>. R Markdown notebook to interactively explore CE and the EO test in a range of simulation settings can be found at <https://github.com/nadavrap/CoordinatedInteractions/tree/master/simulations>. UKBB data are available through the UK Biobank Access Management System. Detailed instructions can be found at <https://www.ukbiobank.ac.uk/register-apply/>.

ACKNOWLEDGMENTS. N.Z. is supported by NIH grants K25HL121295, U01HG009080, R01HG006399, R01CA227237, R03DE025665, and Department of Defense W81XWH-16-2-0018. A.D. is supported by NIH grants U01HG009080 and R01HG006399. B.S. is supported by NIH Grant R01CA227237. N.R. is supported by NIH grants K25HL121295 and U01HG009080. S.J.S. is supported by NIH grants R01MH110928 and R01MH111662. This research has been conducted using the UKBB resource under application no. 30397. We thank Erez Dor for his help in removing related samples.

1. N. H. Barton, How does epistasis influence the response to selection? *Heredity* **118**, 96–109 (2017).
2. R. B. Corbett-Detig, J. Zhou, A. G. Clark, D. L. Hartl, J. F. Ayroles, Genetic incompatibilities are widespread within species. *Nature* **504**, 135–137 (2013).

3. D. S. Park *et al.*, An ancestry-based approach for detecting interactions. *Genet. Epidemiol.* **42**, 49–63 (2018).
4. L. J. Sittig *et al.*, Genetic background limits generalizability of genotype-phenotype relationships. *Neuron* **91**, 1253–1259 (2016).

5. H. Shao *et al.*, Genetic architecture of complex traits: Large phenotypic effects and pervasive epistasis. *Proc. Natl. Acad. Sci. U.S.A.* **105**, 19910–19914 (2008).
6. A. Chen, Y. Liu, S. M. Williams, N. Morris, D. A. Buchner, Widespread epistasis regulates glucose homeostasis and gene expression. *PLoS Genet.* **13**, e1007025 (2017).
7. J. S. Bloom *et al.*, Genetic interactions contribute less than additive effects to quantitative trait variation in yeast. *Nat. Commun.* **6**, 8712 (2015).
8. R. B. Brem, J. D. Storey, J. Whittle, L. Kruglyak, Genetic interactions between polymorphisms that affect gene expression in yeast. *Nature* **436**, 701–703 (2005).
9. S. K. G. Forsberg, J. S. Bloom, M. J. Sadhu, L. Kruglyak, Ö. Carlborg, Accounting for genetic interactions improves modeling of individual quantitative trait phenotypes in yeast. *Nat. Genet.* **49**, 497–503 (2017).
10. W. Huang *et al.*, Epistasis dominates the genetic architecture of *Drosophila* quantitative traits. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 15553–15559 (2012).
11. T. F. C. Mackay, Epistasis and quantitative traits: Using model organisms to study gene-gene interactions. *Nat. Rev. Genet.* **15**, 22–33 (2014).
12. A. Dixit *et al.*, Perturb-Seq: Dissecting molecular circuits with scalable single-cell RNA profiling of pooled genetic screens. *Cell* **167**, 1853–1866.e17 (2016).
13. M. A. Horlbeck *et al.*, Mapping the genetic landscape of human cells. *Cell* **174**, 953–967.e22 (2018).
14. T. M. Norman *et al.*, Exploring genetic interaction manifolds constructed from rich single-cell phenotypes. *Science* **365**, 786–793 (2019).
15. C. D. Rau *et al.*, Modeling epistasis in mice and yeast using the proportion of two or more distinct genetic backgrounds: Evidence for “polygenic epistasis”. *PLoS Genet.* **16**, e1009165 (2020).
16. B. Chen, A. Wagner, Hsp90 is important for fecundity, longevity, and buffering of cryptic deleterious variation in wild fly populations. *BMC Evol. Biol.* **12**, 25 (2012).
17. G. P. Meares, A. A. Zmijewska, R. S. Jope, Heat shock protein-90 dampens and directs signaling stimulated by insulin-like growth factor-1 and insulin. *FEBS Lett.* **574**, 181–186 (2004).
18. C. Queitsch, T. A. Sangster, S. Lindquist, Hsp90 as a capacitor of phenotypic variation. *Nature* **417**, 618–624 (2002).
19. K. L. Hickey *et al.*, GIGYF2 and 4EHP inhibit translation initiation of defective messenger RNAs to assist ribosome-associated quality control. *bioRxiv* [Preprint] (2019). <https://doi.org/10.1101/792994> (Accessed 4 October 2019).
20. T. F. Mackay, J. H. Moore, Why epistasis is important for tackling complex human disease genetics. *Genome Med.* **6**, 124 (2014).
21. J. Marchini, P. Donnelly, L. R. Cardon, Genome-wide strategies for detecting multiple loci that influence complex diseases. *Nat. Genet.* **37**, 413–417 (2005).
22. L. Crawford, P. Zeng, S. Mukherjee, X. Zhou, Detecting epistasis with the marginal epistasis test in genetic mapping studies of quantitative traits. *PLoS Genet.* **13**, e1006869 (2017).
23. J.-L. Jannink, Identifying quantitative trait locus by genetic background interactions in association studies. *Genetics* **176**, 553–561 (2007).
24. C. D. Rau *et al.*, The effects of mutations are modified by genetic background in mice. *bioRxiv* [Preprint] (2019). <https://doi.org/10.1101/555383>. Accessed 7 November 2019.
25. C. C. Cockerham, An extension of the concept of partitioning hereditary variance for analysis of covariances among relatives when epistasis is present. *Genetics* **39**, 859–882 (1954).
26. C. R. Henderson, Best linear unbiased prediction of nonadditive genetic merits in noninbred populations. *J. Anim. Sci.* **60**, 111–117 (1985).
27. A. I. Young, R. Durbin, Estimation of epistatic variance components and heritability in founder populations and crosses. *Genetics* **198**, 1405–1416 (2014).
28. Y. Jiang, J. C. Reif, Modeling epistasis in genomic selection. *Genetics* **201**, 759–768 (2015).
29. O. Zuk, E. Hechter, S. R. Sunyaev, E. S. Lander, The mystery of missing heritability: Genetic interactions create phantom heritability. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 1193–1198 (2012).
30. C. C. Milton, B. Huynh, P. Batterham, S. L. Rutherford, A. A. Hoffmann, Quantitative trait symmetry independent of Hsp90 buffering: Distinct modes of genetic canalization and developmental stability. *Proc. Natl. Acad. Sci. U.S.A.* **100**, 13396–13401 (2003).
31. E. A. Boyle, Y. I. Li, J. K. Pritchard, An expanded view of complex traits: From polygenic to omnigenic. *Cell* **169**, 1177–1186 (2017).
32. X. Liu, Y. I. Li, J. K. Pritchard, Trans effects on gene expression can drive omnigenic inheritance. *Cell* **177**, 1022–1034.e6 (2019).
33. A. T. Timberlake *et al.*, Two locus inheritance of non-syndromic midline craniosynostosis via rare *SMAD6* and common *BMP2* alleles. *eLife* **5**, e20125 (2016).
34. J. Goodrich *et al.*, Determinants of penetrance and variable expressivity in monogenic metabolic conditions across 77,184 exomes. *medRxiv* [Preprint] (2020). <https://doi.org/10.1101/2020.09.22.20195529>. Accessed 21 October 2020.
35. J. H. Moore, S. M. Williams, Traversing the conceptual divide between biological and statistical epistasis: Systems biology and a more modern synthesis. *BioEssays* **27**, 637–646 (2005).
36. K. S. Small *et al.*, Regulatory variants at *KLF14* influence type 2 diabetes risk via a female-specific effect on adipocyte size and body composition. *Nat. Genet.* **50**, 572–580 (2018).
37. S. Smemo *et al.*, Obesity-associated variants within *FTO* form long-range functional connections with *IRX3*. *Nature* **507**, 371–375 (2014).
38. K. A. Bailey *et al.*, Evidence of non-pancreatic beta cell-dependent roles of *Tcf7l2* in the regulation of glucose metabolism in mice. *Hum. Mol. Genet.* **24**, 1646–1654 (2015).
39. P. C. Phillips, Epistasis—The essential role of gene interactions in the structure and evolution of genetic systems. *Nat. Rev. Genet.* **9**, 855–867 (2008).
40. J. Yang *et al.*, Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* **42**, 565–569 (2010).
41. A. Dahl *et al.*, A robust method uncovers significant context-specific heritability in diverse complex traits. *Am. J. Hum. Genet.* **106**, 71–91 (2020).
42. P. Drineas, M. W. Mahoney, RandNLA: Randomized numerical linear algebra. *Commun. ACM* **59**, 80–90 (2016).
43. Y. Wu, S. Sankararaman, A scalable estimator of SNP heritability for biobank-scale data. *Bioinformatics* **34**, i187–i194 (2018).
44. E. J. Candes, M. B. Wakin, An introduction to compressive sampling. *IEEE Signal Process. Mag.* **25**, 21–30 (2008).
45. J. Jiang, C. Li, D. Paul, C. Yang, H. Zhao, On high-dimensional misspecified mixed model analysis in genome-wide association study. *Ann. Stat.* **44**, 2127–2160 (2016).
46. A. Pazokitoroudi *et al.*, Efficient variance components analysis across millions of genomes. *Nat. Commun.* **11**, 4020 (2020).
47. L. Yengo *et al.*, Imprint of assortative mating on the human genome. *Nat. Hum. Behav.* **2**, 948–954 (2018).
48. D. Steinsaltz, A. Dahl, K. W. Wachter, Statistical properties of simple random-effects models for genetic heritability. *Electron. J. Stat.* **12**, 321–356 (2018).
49. C. Bycroft *et al.*, The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**, 203–209 (2018).
50. J. Euesden, C. M. Lewis, P. F. O’Reilly, PRSice: Polygenic risk score software. *Bioinformatics* **31**, 1466–1468 (2015).
51. H. K. Finucane *et al.*, Brainstorm Consortium, Heritability enrichment of specifically expressed genes identifies disease-relevant tissues and cell types. *Nat. Genet.* **50**, 621–629 (2018).
52. H. K. Finucane *et al.*, ReproGen Consortium; Schizophrenia Working Group of the Psychiatric Genomics Consortium; RACI Consortium, Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nat. Genet.* **47**, 1228–1235 (2015).
53. A. E. Justice *et al.*, CHD Exome+ Consortium; Cohorts for Heart and Aging Research in Genomic Epidemiology (CHARGE) Consortium; EPIC-CVD Consortium; ExomeBP Consortium; Global Lipids Genetic Consortium; GoT2D Genes Consortium; InterAct; ReproGen Consortium; T2D-Genes Consortium; MAGIC Investigators, Protein-coding variants implicate novel genes related to lipid homeostasis contributing to body-fat distribution. *Nat. Genet.* **51**, 452–469 (2019).
54. P.-R. Loh, G. Kichaev, S. Gazal, A. P. Schoech, A. L. Price, Mixed-model association for biobank-scale datasets. *Nat. Genet.* **50**, 906–908 (2018).
55. A. P. Schoech *et al.*, Quantification of frequency-dependent genetic architectures in 25 UK Biobank traits reveals action of negative selection. *Nat. Commun.* **10**, 790 (2019).
56. Z. Zhu *et al.*, A genome-wide cross-trait analysis from UK Biobank highlights the shared genetic architecture of asthma and allergic diseases. *Nat. Genet.* **50**, 857–864 (2018).
57. J. J. Berg *et al.*, Reduced signal for polygenic adaptation of height in UK Biobank. *eLife* **8**, e39725 (2019).
58. M. Sohail *et al.*, Polygenic adaptation on height is overestimated due to uncorrected stratification in genome-wide association studies. *eLife* **8**, e39702 (2019).
59. A. Abdellaoui *et al.*, Genetic correlates of social stratification in Great Britain. *Nat. Hum. Behav.* **3**, 1332–1342 (2019).
60. S. Kerminen *et al.*, Geographic variation and bias in the polygenic scores of complex diseases and traits in Finland. *Am. J. Hum. Genet.* **104**, 1169–1181 (2019).
61. X. Liu *et al.*, GBAT: A gene-based association method for robust trans-gene regulation detection. *bioRxiv* [Preprint] (2018).
62. A. R. Martin *et al.*, Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat. Genet.* **51**, 584–591 (2019).
63. A. R. Martin *et al.*, Human demographic history impacts genetic risk prediction across diverse populations. *Am. J. Hum. Genet.* **100**, 635–649 (2017).
64. H. Mostafavi, A. Harpak, D. Conley, J. K. Pritchard, M. Przeworski, Variable prediction accuracy of polygenic scores within an ancestry group. *bioRxiv* [Preprint] (2019). doi.org/10.7554/eLife.48376. Accessed 4 October 2019.
65. J. Mefford *et al.*, Efficient estimation and applications of cross-validated genetic predictions to polygenic risk scores and linear mixed models. *J. Comp. Biol.* **27**, 599–612 (2020).
66. R. S. N. Fehrmann *et al.*, Gene expression analysis identifies global gene dosage sensitivity in cancer. *Nat. Genet.* **47**, 115–125 (2015).
67. T. H. Pers *et al.*, Genetic Investigation of Anthropometric Traits (GIANT) Consortium, Biological interpretation of genome-wide association studies using predicted gene functions. *Nat. Commun.* **6**, 5890 (2015).
68. A. Kundaje *et al.*, Roadmap Epigenomics Consortium, Integrative analysis of 111 reference human epigenomes. *Nature* **518**, 317–330 (2015).
69. ENCODE Project Consortium, An integrated encyclopedia of DNA elements in the human genome. *Nature* **489**, 57–74 (2012).
70. K. R. Feingold, C. Grunfeld, “Introduction to lipids and lipoproteins” in *Endotext*, K. R. Feingold *et al.*, Eds. (MDText.com, Inc., 2000).
71. K. Kaushansky, Lineage-specific hematopoietic growth factors. *N. Engl. J. Med.* **354**, 2034–2045 (2006).
72. H. Ehrenreich *et al.*, A hematopoietic growth factor, thrombopoietin, has a proapoptotic role in the brain. *Proc. Natl. Acad. Sci. U.S.A.* **102**, 862–867 (2005).
73. A. N. Barbeira *et al.*, Widespread dose-dependent effects of RNA expression and splicing on complex diseases and traits. *Genetics* [Preprint] (2019). <https://doi.org/10.1101/814350> (Accessed 13 December 2019).
74. N. Cherbuin, K. Sargent-Cox, M. Fraser, P. Sachdev, K. J. Anstey, Being overweight is associated with hippocampal atrophy: The PATH through life study. *Int. J. Obes.* **39**, 1509–1514 (2015).
75. C. A. Raji *et al.*, Brain structure and obesity. *Hum. Brain Mapp.* **31**, 353–364 (2010).

76. J. H. Stern, J. M. Rutkowski, P. E. Scherer, Adiponectin, leptin, and fatty acids in the maintenance of metabolic homeostasis through adipose tissue crosstalk. *Cell Metab.* **23**, 770–784 (2016).
77. C. L. Masters *et al.*, Alzheimer's disease. *Nat. Rev. Dis. Primers* **1**, 15056 (2015).
78. S. Phani, D. B. Re, S. Przedborski, The role of the innate immune system in ALS. *Front. Pharmacol.* **3**, 150 (2012).
79. A. Dahl *et al.*, Reverse GWAS: Using genetics to identify and model phenotypic subtypes. *PLoS Genet.* **15**, e1008009 (2019).
80. S. Rosset, R. Heller, A. Painsky, E. Aharoni, Optimal and maximin procedures for multiple testing problems. [Preprint] (2018). (Accessed 4 October 2019).
81. S. Sverdlov, E. A. Thompson, The epistasis boundary: Linear vs. nonlinear genotype-phenotype relationships. *bioRxiv*, 503466 (2018). <https://doi.org/10.1101/503466>. Accessed 4 October 2019.
82. S. N. Yazbek *et al.*, Deep congenic analysis identifies many strong, context-dependent QTLs, one of which, Slc35b4, regulates obesity and glucose homeostasis. *Genome Res.* **21**, 1065–1073 (2011).
83. E. R. Gamazon *et al.*, GTEx Consortium, A gene-based association method for mapping traits using reference transcriptome data. *Nat. Genet.* **47**, 1091–1098 (2015).
84. A. Gusev *et al.*, Integrative approaches for large-scale transcriptome-wide association studies. *Nat. Genet.* **48**, 245–252 (2016).
85. O. Brandman *et al.*, A ribosome-bound quality control complex triggers degradation of nascent peptides and signals translation stress. *Cell* **151**, 1042–1054 (2012).
86. D. J. Balding, R. A. Nichols, A method for quantifying differentiation between populations at multi-allelic loci and its implications for investigating identity and paternity. *Genetica* **96**, 3–12 (1995).
87. A. Okbay *et al.*, LifeLines Cohort Study, Genome-wide association study identifies 74 loci associated with educational attainment. *Nature* **533**, 539–542 (2016).
88. S. Gazal *et al.*, Functional architecture of low-frequency variants highlights strength of negative selection across coding and non-coding annotations. *Nat. Genet.* **50**, 1600–1607 (2018).
89. C. C. Chang *et al.*, Second-generation PLINK: Rising to the challenge of larger and richer datasets. *Gigascience* **4**, 7 (2015).
90. D. Firth, Bias reduction of maximum likelihood estimates. *Biometrika* **80**, 27–38 (1993).
91. J. Mbatchou *et al.*, Computationally efficient whole genome regression for quantitative and binary traits. *bioRxiv* [Preprint] (2020). <https://doi.org/10.1101/2020.06.19.162354>. Accessed 23 February 2021.
92. J. Zheng *et al.*, Early Genetics and Lifecourse Epidemiology (EAGLE) Eczema Consortium, LD Hub: A centralized database and web interface to perform LD score regression that maximizes the potential of summary level GWAS data for SNP heritability and genetic correlation analysis. *Bioinformatics* **33**, 272–279 (2017).
93. A. R. Wood *et al.*, Electronic Medical Records and Genomics (eMEMERGE) Consortium; MIGen Consortium; PAGEGE Consortium; LifeLines Cohort Study, Defining the role of common variation in the genomic and biological architecture of adult human height. *Nat. Genet.* **46**, 1173–1186 (2014).
94. C. J. Willer *et al.*, Global Lipids Genetics Consortium, Discovery and refinement of loci associated with lipid levels. *Nat. Genet.* **45**, 1274–1283 (2013).
95. A. E. Locke *et al.*, LifeLines Cohort Study; ADIPOGen Consortium; AGEN-BMI Working Group; CARDIOGRAMplusC4D Consortium; CKDGen Consortium; GLGC; ICBP; MAGIC Investigators; MuTHER Consortium; MIGen Consortium; PAGE Consortium; ReproGen Consortium; GENIE Consortium; International Endogene Consortium, Genetic studies of body mass index yield new insights for obesity biology. *Nature* **518**, 197–206 (2015).
96. A. P. Morris *et al.*, Wellcome Trust Case Control Consortium; Meta-Analyses of Glucose and Insulin-related traits Consortium (MAGIC) Investigators; Genetic Investigation of Anthropometric Traits (GIANT) Consortium; Asian Genetic Epidemiology Network–Type 2 Diabetes (AGEN-T2D) Consortium; South Asian Type 2 Diabetes (SAT2D) Consortium; DIAbetes Genetics Replication And Meta-analysis (DIAGRAM) Consortium, Large-scale association analysis provides insights into the genetic architecture and pathophysiology of type 2 diabetes. *Nat. Genet.* **44**, 981–990 (2012).
97. M. F. Moffatt *et al.*, GABRIEL Consortium, A large-scale, consortium-based genome-wide association study of asthma. *N. Engl. J. Med.* **363**, 1211–1221 (2010).
98. M. Nikpay *et al.*, A comprehensive 1,000 Genomes-based genome-wide association meta-analysis of coronary artery disease. *Nat. Genet.* **47**, 1121–1130 (2015).
99. S. A. Lambert *et al.*, The polygenic score catalog: An open database for reproducibility and systematic evaluation. *medRxiv* [Preprint] (2020). <https://doi.org/10.1101/2020.05.20.20108217>. Accessed 23 February 2021.
100. D. Vuckovic *et al.*, VA Million Veteran Program, The polygenic and monogenic basis of blood traits and diseases. *Cell* **182**, 1214–1231.e11 (2020).