Gene-environment interactions using a Bayesian whole genome regression model

Matthew Kerin¹ & Jonathan Marchini²

¹*Wellcome Trust Center for Human Genetics, Oxford, UK.*

²*Regeneron Genetics Center, Tarrytown, USA.*

Corresponding author Jonathan Marchini (jonathan.marchini@regeneron.com)

Abstract

The contribution of gene-environment (GxE) interactions for many human traits and diseases is poorly characterised. We propose a Bayesian whole genome regression model, LEMMA, for joint modeling of main genetic effects and gene-environment interactions in large scale datasets such as the UK Biobank, where many environmental variables have been measured. The method estimates a linear combination of environmental variables, called an environmental score (ES), that interacts with genetic markers throughout the genome, and provides a readily interpretable way to examine the combined effect of many environmental variables. The ES can be used both to estimate the proportion of phenotypic variance attributable to GxE effects, and also to test for GxE effects at genetic variants across the genome. GxE effects can induce heteroscedasticity in quantitative traits and LEMMA accounts for this using robust standard error estimates when testing for GxE effects. When applied to body mass index, systolic, diastolic and pulse pressure in the UK Biobank we estimate that 9.3%, 3.9%, 1.6% and 12.5% of phenotypic variance is explained by GxE interactions,

and that low frequency variants explain most of this variance. We also identify 3 loci that interact with the estimated environmental scores ($-\log_{10} p > 7.3$).

Introduction

Despite long standing interest in gene-by-environment (GxE) interactions¹, this facet of genetic architecture remains poorly characterized in humans. Detection of GxE interactions is inherently more difficult than finding additive genetics in genome wide association studies (GWAS). One difficulty is that of sample size; a commonly cited rule of thumb suggests that detection of interaction effects requires a sample size at least four times larger than that required to detect a main effect of comparable effect size². Another is that an individual's environment, which evolves through time, is very hard to measure in a comprehensive way, and is inherently high dimensional. Also, there are many environmental variables that could plausibly interact with the genome and many ways to combine them, and typically these factors are not all present in the same dataset. The recently released UK Biobank dataset, a large population cohort study with deep genotyping and sequencing, and extensive phenotyping³, offers a unique opportunity to explore GxE effects^{4–10}.

Models that consider environmental variables jointly can be advantageous, particularly if several environmental variables drive interactions at individual loci, or if an unobserved environment driving interactions is better reflected by a combination of observed environments. StructLMM⁷ models the environmental similarity between individuals (over multiple environments) as a random effect, and then tests each SNP independently for GxE interactions. However, StructLMM is

not a whole genome regression model, so does not account for the genome wide contribution of all other variants, which is often a major component of phenotypic variance.

Advances in methods applied to detect genetic main effects in standard GWAS have shown that linear mixed models (LMMs) can reduce false positive associations due to population structure, and improve power by implicitly conditioning on other loci across the genome^{11–13}. Often these methods model the unobserved polygenic contribution as a multivariate Gaussian with covariance structure proportional to a genetic relationship matrix (GRM)^{14–16}. This approach is mathematically equivalent to a whole genome regression (WGR) model with a Gaussian prior over SNP effect sizes¹¹. More flexible approaches have been proposed in both the animal breeding^{17,18} and human literature^{19–21} to allow different prior distributions that better capture SNPs of small and large effects. The BOLT-LMM method¹³ uses a mixture of Gaussians (MoG) prior and shows this can increase power to detect associated loci in some (but not all) complex traits.

Here we present a new method called Linear Environment Mixed Model Analysis (LEMMA) which aims to combine the advantages of WGR and modeling GxE with multiple environments, and is applicable to large datasets with hundreds of thousands of individuals, such as the UK Biobank. Instead of assuming that the GxE effect over multiple environments is independent at each variant, as StructLMM does, we learn a single linear combination of environmental variables (that we call an environmental score (ES)), that has a common role in interaction effects genome wide. The ES is estimated within a Bayesian WGR model that uses two separate MoG priors on main genetic effects and GxE effects. We use variational inference to fit the model that is

tractable for GxE analyses of biobank scale datasets with tens of environmental variables (see **Online Methods**).

Estimating the ES satisfies one of the primary motivations of this work, by providing a readily interpretable way to examine the combined effect of many environmental variables and how they might interact with genotype. A motivating example is the investigation of how modern obesogenic environments might accentuate the genetic risk of obesity. Tyrell et al.²² studied environments one at a time for their interaction with a body mass index (BMI) genetic risk score (GRS) and found several significant interactions. Our new method allows joint analysis of environments that might plausibly better represent an obesogenic environment, negating the need to model each environment one at a time. Our other motivations when developing LEMMA were to develop a powerful method to detect GxE interactions, and to estimate the proportion of variance that could be attributable to GxE interactions.

A LEMMA analysis has several distinct steps. First, the model is fitted using a large set of SNPs genome-wide, for example all the SNPs that have been directly assayed on a genotyping chip. The estimated ES is then used to estimate the proportion of phenotypic variability that is explained by interactions with this ES (GxE heritability), using randomized Haseman-Elston (RHE) regression^{23,24}. This heritability analysis can be run on genotyped or imputed SNPs, and can be stratified by minor allele frequency (MAF) and linkage disequilibrium (LD) to better interrogate the genetic architecture of GxE interactions. The ES is also used to test for GxE interactions one variant at a time, typically at a larger set of imputed SNPs in the dataset. We use “robust” standard

errors when testing each variant for a GxE interaction, which helps to control for the conditional heteroskedasticity caused by GxE interactions. We also suggest checks and solutions for the situation where environmental variables are themselves heritable and have a non-linear relationship to the trait of interest (see **Online Methods**).

We compared LEMMA to existing approaches such as StructLMM and F-tests using simulated data, and applied the approach to UK Biobank data for body mass index (BMI), systolic blood pressure (SBP), diastolic blood pressure (DBP) and pulse pressure (PP).

Results

Performance on simulated data **Figure 1** compares the ability of different methods to detect GxE interactions at SNPs in simulations where a single true ES interacts with SNPs across the genome. **Figure S1** shows the false positive rate (FPR) to detect main effects. We compared our default version of LEMMA, which uses robust standard errors, StructLMM, a simple F-test of interaction and an F-test that uses robust standard errors (see **Online Methods**). The simulations vary GxE heritability, the total number of environmental variables and sample size. When sample size is large (N=100k), all the methods have reasonable control of FPR and LEMMA controls FPR at least as well as other methods across the range of simulations. When sample size is smaller (N=25k) the robust F-test performs less well as the number of environments grows (**Figure 1a**) and the F-test and StructLMM perform less well as the amount of GxE variance increases (**Figure 1b**). When we increase the sample size to N=200k we still find that LEMMA has a slightly inflated Type

I error rate (see **Figure S2**).

It is interesting that all the methods we tested have slightly inflated Type I error, and this is likely due to a number of different reasons. StructLMM and the F-test fit a model at each variant and ignore GxE effects at other loci, which can induce heteroscedasticity that can inflate Type I error^{49,61}. We used robust standard errors for the robust F-test but it seems that this approximation works best when the number of environmental traits is small. LEMMA does account for GxE effects at other loci and also uses robust standard errors, but still has slightly inflated Type I error that gets worse as the number of environments increases (**Figure 1a** and **Figure S2**). In parallel simulations (see **Figure S3**) we find that our model slightly over-estimates GxE heritability as number of environments increases. Since our simulations test for GxE effects at SNPs used to estimate the ES we suspect that the Type I error inflation is due to this two stage approach.

When there is a single true ES involved in GxE interactions we found that LEMMA provided a substantial power increase (**Figure 1**, **Figure S4**). StructLMM and F-tests have very similar power in these simulations, although previous work suggests that StructLMM may outperform the F-test in small samples⁷.

When estimating the GxE heritability of the LEMMA ES using randomized HE regression (RHE) with a single SNP component (RHE-SC) we observed some upward bias as the number of environments increases. This effect is ameliorated by increasing sample size (see **Figure S3**), suggesting that the influence of over-fitting in our Biobank analyses is mild. In twenty simulations with $L = 30$ environmental variables, $N = 100k$ samples and true GxE heritability of 5% we

Figure 1: Type I error and Power of tests to detect GxE effects in simulation. (a-c) Comparison of false positive rate as the number of environments increases (a), as phenotype variance explained by GxE effects increases (b) or as the number samples increases (c). (d-f) Analogous comparison of the power to detect GxE interactions. Simulations used genotypes subsampled from the UK Biobank and by default contained $N = 25K$ samples, $M = 100K$ SNPs, 6 environmental variables that contributed to the ES and 24 that did not (default parameters denoted by stars). We assess power (at Family Wise Error Rate; $FWER < 0.01$) to detect 60 causal SNPs whose GxE effect each explained 0.00016% of trait variance. See **Online methods** for full details of phenotype construction.

observed mean GxE heritability of 5.2%. **Figure S5** further illustrates the ES estimation accuracy of LEMMA.

Finally, we ran LEMMA on two sets of simulated datasets ($N=25k$) with causal SNPs chosen either randomly, or chosen to be low frequency ($MAF < 0.1$). We used the ES estimated from

each simulated dataset to estimate h_G^2 and $h_{G \times E}^2$ using RHE, with SNPs stratified by MAF and LD (RHE-LDMS), and then without any stratification (RHE-SC). Previous studies have established that estimating heritability using a single SNP component makes assumptions about the relationship between MAF, LD and trait architecture that may not hold up in practice^{27,28}, whereas stratifying SNPs into bins according to MAF and LDscore (the LDMS approach) is relatively unbiased^{28–30}. **Figure S6** confirms that stratifying by MAF and LD results in accurate heritability estimates irrespective of the MAF distribution of causal SNPs, and suggests that this method can be used to interrogate the MAF distribution of GxE component of a trait using LEMMA. However when causal SNPs are low frequency, not stratifying by MAF and LD results in underestimation of h_G^2 .

Controlling for heritable environmental variables Previous work by Tchetgen *et al.*²⁵ has shown that misspecifying the functional form of an environmental variable can induce heteroskedasticity into tests for GxE interactions. The authors further show that use of robust standard errors will control for this heteroskedasticity, but only if the environment is independent of the variant being tested. Independence between genotypes and the misspecified environment is important because it means that the (least squares) mean estimator is still unbiased.

However, environmental variables themselves often have a genetic basis. We therefore performed simulations where the phenotype depended on the non-linear (squared) effect of a heritable environmental variable. In simulation (**Figure 2a,c**) we observed that misspecification of the environmental variable can cause substantial inflation in GxE test statistics at heritable sites of the

Figure 2: Bias from model mis-specification of a heritable environmental variable. (a) Comparison of GxE association test statistics from a single simulation where non-linear dependence on the confounder explains 5% of trait variance. FPR at heritable sites of the misspecified environment only. (b) Comparison of GxE association test statistics from an analysis of logBMI in 281,149 participants from the UK Biobank. (c) False Positive Rate at heritable sites of the misspecified environment whilst the strength of squared dependence varies. 20 repeats per scenario. Abbreviations are as follows: LEMMA-S, LEMMA with non-robust variances used to compute test statistics; (+SQE), significant squared environmental variables (Bonferroni correction) included as additional covariates.

confounding environment. Relatively smooth non-linearities, such as squared effects, are easily detected by regression modeling before using LEMMA (see **Online Methods**) and can then be included as covariates (indicated in **Figure 2** by +SQE). This procedure produced well calibrated test statistics for all methods in simulation (**Figure 2c**).

In **Figure 2b** we compare the GxE association test statistics from our analysis of logBMI in the UK Biobank, with and without adjusting for detected squared effects. Although we detected squared effects for 30 of the 42 environmental variables (significance level 0.01; Bonferroni correction for multiple testing), the ES obtained from the two analyses was almost identical (Pearson $r^2 > 0.999$). As the additional variance explained collectively by the squared effects was negligible (incremental $R^2 < 0.00001$) it would be surprising if this was not the case. Negative $\log_{10}(P)$ -values from the two analyses were also highly correlated (Pearson $r^2 = 0.961$), although there were small changes in the p -values at the *FOXO3* locus (which remained genome-wide significant in both analyses) and at the *SNAP25* locus (which was genome wide significant in the (-SQE) analysis only). We therefore conclude that the influence from this form of confounding in our analysis of logBMI was minor. However as the cost to this procedure is small, LEMMA uses the (+SQE) strategy by default for all analyses of UK Biobank traits.

GxE interaction analysis in the UK Biobank We applied LEMMA to characterize GxE interactions in Body Mass Index (logBMI), Systolic Blood Pressure (SBP), Diastolic Blood Pressure (DBP) and Pulse Pressure (PP) using a set of 42 environmental variables similar to those used in previous analyses ^{7,8,26}, including data on smoking, hours of TV watched, Townsend Index, physical exercise and alcohol consumption (see **Online methods** and **Table S1**).

We analyzed GxE heritability due to multiplicative effects with the ES using both $M = 639,005$ genotyped SNPs and $M = 10,270,052$ common imputed SNPs ($MAF \geq 0.01$ in the full UK Biobank cohort), stratified by MAF and LDscore into 20 components. Using imputed SNPs

Figure 3: **Partitioned heritability estimates for four quantitative traits in the UK Biobank.**

Heritability estimates partitioned into additive genetic and multiplicative GxE interaction effects for four quantitative traits in the UK Biobank, using approximately 280,000 unrelated white British individuals (see **Table S1**) and $M = 10,270,052$ common imputed SNPs ($MAF > 0.01$ in the full UK Biobank cohort). Multiplicative GxE interactions were computed using the ES from each model fit. Heritability estimation was performed using a multi-component implementation of randomized HE-regression^{23,24} with SNPs stratified into 20 components (5 MAF bins and 4 LD score quantiles).

we estimated GxE heritability of 9.3%, 12.5%, 3.9% and 1.6% for logBMI, PP, SBP and DBP respectively (see **Table 1**). On genotyped SNPs the GxE heritability estimates were slightly lower for logBMI and PP ($h^2_{GxE} = 8.6\%$ and $h^2_{GxE} = 11.1\%$ respectively) and almost identical for SBP

and DBP (see **Table S2**). For all traits the heritability of additive SNP effects were slightly higher on imputed data, consistent with previous results ²⁹.

When working with quantitative traits it can be hard to choose an optimal transformation or scale for each trait. Tyrell et al. ²² analyzed BMI using the raw scale and then also by transforming to a standard normal distribution. They observed larger interaction effects on the raw scale and suggested that this was due to larger variance in BMI in individuals in the high-risk environment groups, which causes heteroscedasticity, and inflates effect estimates. In addition to our main analysis, which used log BMI, we re-ran LEMMA using the raw BMI measurement, and then also by transforming to a standard normal distribution in females and males separately. These results are presented in **Table 1** and agree with the results of Tyrell et al. ²², with estimates of GxE heritability on the raw, log and inverse normal scale of 13.7%, 9.3% and 5.9% respectively.

Previous work on models of natural selection has suggested that the variance explained by additive SNP effects should be uniformly distributed as a function of MAF in a neutral evolutionary setting³¹, and that enrichment of the variance explained by low frequency SNPs is evidence for negative selection. For all four traits we found that variance explained by the additive genetic effects of low frequency SNPs ($MAF < 0.1$) was slightly elevated, consistent with previous observations of negative selection³² (**Figure 3**). Additionally the distribution of additive genetic effects by MAF for logBMI was qualitatively similar to that found by GREML-LDMS in a previous study²⁹. In contrast we found that variance explained by GxE effects was overwhelmingly attributed to low frequency SNPs ($MAF < 0.01$), especially those with low LD. However we are not aware

of any evolutionary theory that has been extended to model the MAF distribution of GxE effects.

For logBMI we estimated an ES that put high weight on alcohol intake frequency, Townsend Index and physical activity measures (**Figure 4c**). Almost all of the non-dietary environmental exposures had a higher effect in women than in men, with smoking status being the one exception. This is reflected in the facts that (a) the ES having much higher variance in women and (b) those with a negative ES were almost all female (97%) (see **Figure 4b**). When comparing the characteristics of those in the bottom 5% of the ES to the whole cohort (using the mean for continuous variables and the mode for categorical), we found that those in the bottom 5% were predominantly female (100% vs 53%), younger (51 vs 56), had higher Townsend deprivation index (0.91 vs -1.74), drank less often ('Special occasions' vs 'Once or twice a week') and watched more TV (3.28 vs 2.69 daily hours of TV) (**Table S3**). We note that positive values of the Townsend index indicate material deprivation, whereas negative values indicate relative affluence.

Previous cross-sectional studies have reported GxE interactions between a linear predictor formed from BMI-associated SNPs and alcohol intake frequency³³, Townsend Index^{22,33}, physical activity measures^{5,22,33,34} and time watching TV^{22,33,34}, all of which had high relative weight in the logBMI ES. An alternative approach from Robinson *et al.*⁶ binned samples according to their environmental exposure (eg. age) and tested for significant differences in SNP heritability using a likelihood ratio test. They reported strong interaction effects with age in a cohort of 43,407 individuals whose ages spanned 18 – 80, but only reported significant interactions with smoking in the UK Biobank interim release. This suggests that we might expect age to play a more dom-

Trait	h_G^2 (s.e)	$h_{G \times E}^2$ (s.e)
logBMI	0.274 (0.056)	0.093 (0.028)
INT(BMI)	0.278 (0.056)	0.059 (0.024)
BMI	0.268 (0.055)	0.137 (0.031)
PP	0.228 (0.051)	0.125 (0.028)
SBP	0.251 (0.05)	0.039 (0.023)
DBP	0.254 (0.05)	0.016 (0.02)

Table 1: **Partitioned heritability estimates for four quantitative traits in the UK Biobank.**

Heritability estimates obtained using common imputed SNPs (MAF > 0.01 in the full UK Biobank cohort) with RHE-LDMS. GxE heritability estimates were obtained using the ES from each model fit. All analyses controlled for the same covariates used in the WGR analysis (including the top 20 principal components). Abbreviations; s.e, standard error estimated using the block jackknife (see **Online Methods**); h_G^2 , heritability due to additive genetic effects; $h_{G \times E}^2$, heritability due to multiplicative GxE effects; RHE, randomized HE-regression^{23,24}; LDMS, SNPs stratified by minor allele frequency and LDscore (20 components); INT, Inverse normal transform applied to males and females separately.

inant role in the logBMI ES in a cohort that included younger individuals. Finally, one category that is notably down weighted is the contribution from dietary variables. Although significant interactions with fried food consumption³⁵ and sugar sweetened drinks³⁶ have previously been reported in a cohort of US health professionals, these dietary variables were not included in the

Figure 4: GxE analysis of logBMI in the UK Biobank. (a) LEMMA association statistics testing for multiplicative GxE interactions at each SNP. The horizontal grey line denotes ($p = 5 \times 10^{-8}$), p -values are shown on the $-\log_{10}$ scale. (b) Distribution of the environmental score (ES), stratified by gender and age quantile. (c) Weights used to construct the ES. Dietary variables have a single weight shown on the per standard deviation (s.d) scale. 'Gender' has two weights; a gender specific intercept for women (first) and men (second). Remaining non-dietary variables have three weights; (first) a per s.d effect for women only, (second) a per s.d effect for men only, (third) a per s.d per decade effect which is the same for both genders. s.d for the male and female specific weights is computed for each gender separately. Age is computed as the number of decades aged from 40. See **Online Methods** for details.

diet questionnaire used by the UK Biobank.

The ES for PP was dominated by the effects of age and gender (age, age², age-x-gender, gender together explained 94.9% of variance in the ES). The magnitude of the ES was strongly associated with increased age³⁷ whilst the sign of the ES was strongly associated with gender, implying that GxE effects were stronger in the elderly but acted in the opposite direction in men and women (**Figure S7**).

Similarly we observed that variance of the ES increased with age in both SBP and DBP, but instead of age itself being highly weighted we found that age interactions with other environmental variables were most important for explaining variation in the ES. Specifically for SBP we found that age interactions with smoking, Townsend index and alcohol frequency explained 86% of variance in the ES (**Figure S8**) When compared to the cohort average, we found that participants in the top 5% of the SBP ES were older (63 vs 58), had higher Townsend deprivation index (1.2 vs -1.74) and were more likely to smoke (59% vs 9%) whereas those in the bottom 5% were also older (65 vs 58), predominantly female (91.5%), rarely drank alcohol (43.9% drank “Never”) and had low Townsend deprivation index (-2.9 vs -1.74) (**Table S4**).

Finally we observed notably higher variance in the ES for DBP among men, most of which appeared to be driven by high gender-specific weights for smoking status and alcohol frequency (**Figure S9**). We further observed that alcohol frequency and smoking status became increasingly influential with age. The total SNP-GxE heritability for this ES however was quite low.

When testing for significant GxE interactions between the estimated ESs and imputed markers across the genome we observed that use of the robust standard errors made a noticeable difference to the calibration of LEMMA (**Figure S10; Table S5**). We identified two loci for logBMI (**Figure 4a**), one locus for DBP (**Figure S9a**) and zero loci for SBP and PP, using a threshold of 5×10^{-8} for genome wide significance (**Table 2**). This table also includes results from a standard linear regression GWAS test at the 3 loci (see also **Table S6**).

For logBMI, LEMMA identified GxE interactions at rs2153960 ($p = 6.5 \times 10^{-9}$; **Figure S11**) and at rs539515 ($p = 6.5 \times 10^{-12}$; **Figure S12**). The SNP rs2153960 is an intron in the *FOXO3* gene and has been previously associated with Insulin-like growth factor 1 (IGF-I) concentration in a cohort of 10,000 middle aged Europeans³⁸. IGF-I is known to be a central mediator of metabolic, endocrine and anabolic effects of growth hormone and is also involved in carbohydrate homeostasis³⁸. The patterns of main effect association and GxE association show considerable overlap (**Figure S11a**). This SNP did not reach genome-wide levels of significance using the standard linear regression GWAS test (**Table 2**).

The SNP rs539515 is located 6kb downstream of *SEC16B*. The patterns of main effect association and GxE association are very similar **Figure S12a**. Multiplicative GxE interactions have been reported at *SEC16B* with multiple environmental variables in a similar analysis in the UK Biobank⁷, and with physical activity separately in Europeans⁵ ($p = 0.025$) and in Hispanics³⁹ ($p = 8.1 \times 10^{-5}$). Highly significant variance effects ($p = 3.88 \times 10^{-17}$), which can be indicative of GxE, have also been reported at the *SEC16B* locus using $N = 456,422$ Europeans in the

UK Biobank⁴⁰. GxE interactions have been reported at *SEC16B* with multiple environmental variables in a similar analysis in the UK Biobank⁷ and with physical activity separately in Europeans⁵ ($p = 0.025$) and in Hispanics³⁹ ($p = 8.1 \times 10^{-5}$). *SEC16B* transcribes one of the two mammalian homologues of the SEC16 protein, which has a key role in organizing endoplasmic reticulum exit sites by interacting with COPII components⁴¹. Although several GWASs have identified associations between *SEC16B*^{42,43}, the relevance of *SEC16B* to BMI is not well characterized⁴⁴. Some evidence exists to suggest that *SEC16B* has role in the transport of peroxisome biogenesis factors; peroxisomes being an organelle involved in the catabolism of long chain fatty acids found ubiquitously in eukaryotic cells. Previous authors⁴³ have also speculated that the *SEC16B* might play a role in the transport of appetite regulatory peptides, however we are not aware of any evidence for this theory.

The DBP associated SNP is rs8090962, but only just passes our threshold for significance and we are least confident that this is a true GxE association for a few reasons. The SNP is located within an enhancer, approximately 100KB downstream of the *SEC11C* gene and 50KB upstream of *ZNF532*. Neither gene has previously been associated with blood pressure traits. There is some evidence of a main effect close by (**Figure S13a**) but the pattern of main effect association does not coincide well with the pattern of GxE associations. In addition the pattern of GxE association by genotype (**Figure S13b**) shows a striking cross over by genotype between extremes of the ES. We have observed above that our test statistics is very slightly inflated so this could be a false positive association.

Relationship of genetics PCs and environmental scores We regressed the estimated ESs against the PCs for each of the 4 UKB traits and the results are included in **Table S7**. We found some significant associations, mostly with PC5, which seems to correlate with North-South geography in the UK⁴⁵. To explore further we also re-ran the heritability analysis by including interaction terms of the ES with the genetic PCs as control variables, but the results were almost unchanged (see **Table S8**).

Comparison of the LEMMA ES with a marginal ES For each trait, we used least squares regression to compute a linear model fit using all of the non genetic covariates used in the LEMMA analysis. We then constructed an environmental score (referred to as ES_{marginal}) using the marginal environmental effects from this model fit. The correlation between the LEMMA ES and ES_{marginal} was -0.062 , -0.019 , -0.297 and -0.088 for logBMI, PP, SBP and DBP respectively, suggesting that these vectors are quite dissimilar. **Figure S14** shows a comparison of the interaction weights used to construct the LEMMA ES and ES_{marginal} for each of the four traits. Visually the weights learnt through each approach look quite distinct. In particular, age, age² and age $\times$ gender have much higher relative weight in ES_{marginal} than in the LEMMA ES.

Comparison of methods on UK Biobank data To compare LEMMA with existing single SNP methods we also ran StructLMM, the F-test and the robust F-test on logBMI using the same set of environmental variables as used by LEMMA (but not including the significant squared environments as covariates). Manhattan plots are displayed in **Figure S15**. Test statistics from both the F-test ($\lambda_{GC} = 1.37$) and StructLMM ($\lambda_{GC} = 1.235$) were substantially inflated when compared to

the robust F-test and LEMMA ($\lambda_{GC} = 1.03$ and $\lambda_{GC} = 1.062$ respectively; see **Table S5**), suggesting that StructLMM does not properly control for heteroskedasticity. There are clear differences between the 4 methods, especially among SNPs with suggestive evidence of GxE Interaction results (**Figure S16a**). LEMMA did not find the *FTO* locus, StructLMM and F-test did not find the *SEC16B* locus, and the robust F-test only found the *FTO* locus.

LEMMA relies on the assumption that all GxE interaction effects for a single trait share a common ES, and we have shown in simulation that when this assumption holds LEMMA achieves substantial increases in power. However we would expect LEMMA to have little power to detect SNPs which interact with a combination of environments that is not well correlated with the genome-wide ES estimated by LEMMA. The *FTO* seems to be one clear example of this. We extracted an estimate of the SNP specific interaction profile at *FTO* using the robust F-test (**Online Methods**), we found that it's correlation with LEMMA's ES was low (Pearson $r^2 = 0.3$). In comparison, a similar analysis at *SEC16B* and *FOXO3* yielded much higher correlations (Pearson $r^2 = 0.725$ and $r^2 = 0.713$ respectively).

Discussion

In this study we proposed a new whole genome regression method, LEMMA, that estimates a single environmental score (ES) that interacts with SNPs across the genome. In simulation we have demonstrated that the ES can be used to compute well calibrated p-values of the multiplicative interaction effect at each SNP. LEMMA is also able to quantify the trait variance attributable to

MAF and LD stratified interaction effects of the ES.

In analyses of four quantitative traits in the UK Biobank, we have demonstrated that GxE effects among common imputed SNPs make a non-trivial contribution to the heritability of logBMI and PP (9.3% and 12.5% respectively). Our stratified heritability analysis has suggested that GxE interactions for these traits are mostly driven by low frequency variants. Our analysis identified three loci with statistically significant GxE interaction effects. Two of these loci, rs539515 (*FOXO3*) and rs8090962, are novel and for the other, rs539515 (*SEC16B*), we show stronger evidence for than in the previous study⁷.

Robinson *et al.*⁶ have previously attempted to quantify the contribution of GxE interactions to the heritability of BMI in a study performed on imputed SNPs from the interim UK Biobank release. Using the GCI-GREML model implemented in GCTA⁴⁶ and eight environmental variables that included measures of smoking, hours of TV watched and alcohol frequency, Robinson *et al.*⁶ reported that only smoking had significant GxE heritability (4.0%). In contrast, the ES estimated for logBMI in our analysis had non-zero contributions from many environmental variables, including hours of TV watched and smoking, suggesting that multiple environmental variables can influence on the genetic predisposition to BMI. Modeling these environmental variables jointly allowed LEMMA to capture a combination whose GxE interactions explained 9.3% of heritability.

We have also evaluated the performance of three existing single SNP methods (StructLMM, the F-test and a robust F-test), both in simulation and on logBMI from this same dataset. In simulation with large datasets we observed that StructLMM and the F-test had similar performance; an

observation that also held in our analysis of logBMI. Both of these methods appeared vulnerable to heteroskedasticity, which we showed is likely to occur in traits with non-trivial GxE heritability. A simple adjustment, using ‘robust’ or Huber-White variance estimators, solved this problem. The two F-test methods further benefit from a wealth of existing theory⁴⁷ and, being theoretically simpler than StructLMM, could be easily implemented as an R-plugin with PLINK⁴⁸ (for example ⁴⁹). In our opinion, the robust F-test is therefore the most appropriate of the three single SNP methods to model GxE effects with tens of environments in biobank scale datasets.

Although LEMMA represents a method with increased power to detect GxE interaction effects, our approach does have some caveats. First the gain in power is dependent on a strong assumption on the underlying genetic architecture. Whilst our analysis suggests that this does hold to some extent for PP and logBMI, this may not be the case for other traits.

In addition, LEMMA only estimates the proportion of phenotypic variance that is explained by interactions with this ES, and we do not claim that this captures all the GxE heritability of a trait. If relevant GxE environments are not included in the analysis, and these environments have low correlation to the environments that are included, then LEMMA cannot account for them, and will likely underestimate the true GxE heritability. Unobserved environments can cause trait variance to depend on genotype⁸ (see **Figure S19**) and extending LEMMA in this direction is left for future work.

LEMMA has the requirement that none of the environmental variables have any missing values. This could lead to a reduction in samples size if many environmental variables are included.

If the amount of missing data is small it should not pose a big problem, and missing data imputation methods are also an option. If LEMMA is applied in situations where the missing data structure is related to the phenotype of interest then this could cause bias in the results.

Despite much effort to provide an efficient implementation, the LEMMA algorithm is still a computationally demanding. Using randomized HE-regression to estimate an improved initialization of the interactions weights may help to reduce run time, and is an avenue that we are currently pursuing.

Finally, for simplicity LEMMA currently searches only for GxE interactions with a single linear combination of environments. Generalizing the LEMMA approach to several orthogonal linear combinations or using functional annotation to restrict the SNPs that each ES interacts with, may yet yield more power to identify interactions in complex traits and explain more phenotypic variation.

Online Methods

Linear Environment Mixed Model Analysis (LEMMA) The standard LMM used in genome wide association studies is written as

$$y = C\alpha + u + \epsilon, \quad (1)$$

where y is the centered and scaled $N \times 1$ vector of phenotypes, C is an $N \times L'$ matrix of covariates with $L' \times 1$ fixed effects vector α , and u and ϵ are $N \times 1$ vectors of unobserved polygenic and

residual effects vectors respectively. Typically u is modeled as Gaussian with mean zero and covariance matrix $\sigma_g^2 K$. Specification of the $N \times N$ kinship matrix K is an area of active research^{50–53} but the simplest approach is to let $K = XX^T/M$ where X is the $N \times M$ genotype matrix and columns of X (which usually correspond to SNPs) are normalized to have mean zero and variance one. This can equivalently be written as a Bayesian whole genome regression (WGR) model

$$y = C\alpha + X\beta + \epsilon, \quad (2)$$

where

$$\beta \sim \mathcal{N}(0, \sigma_g^2/M). \quad (3)$$

Here β is a $M \times 1$ vector modeling the random effect of each SNP. This form corresponds to the so called infinitesimal model where every SNP is allowed to have a small but non-zero effect on a given trait. To generalize the model to a non-infinitesimal genetic architecture, we model SNP effects with a mixture of Gaussians prior. This approach has been applied previously in human genetics^{13,21} and by the 'Bayesian alphabet' of genomic prediction methods in the animal breeding literature^{17,18,54}.

We extend this setup to model GxE interactions genome-wide with a linear combination of multiple environmental variables using

$$Y = C\alpha + X\beta + Z\gamma + \epsilon \quad (4)$$

where

$$Z = \eta \odot X, \quad (5)$$

$$\eta = Ew, \quad (6)$$

$$w \sim \mathcal{N}(0, I_L) \quad (7)$$

where E is an $N \times L$ matrix of environmental variables that could potentially be involved in GxE interactions and w is an $L \times 1$ vector of weights. Nether they define the $N \times 1$ vector η that is the linear combination of environments that we refer to as the environmental score (ES). This ES is learned in tandem with SNP effects. We note that all environmental variables contained in E must also be contained in C , so $L \leq L'$. We chose to model the interaction weights w with a Gaussian prior, but in theory one could consider sparser priors such as a spike and slab. We set the variance of the prior on w to the identity matrix I_L . Setting the prior variance of w to a parameter would be unidentifiable, as any change in scale would be absorbed by the prior variances on the interaction effects γ (see $\sigma_{\gamma,1}^2$ and $\sigma_{\gamma,2}^2$ in equations 8-9).

The $N \times M$ matrix Z contains all of the multiplicative interaction terms of the ES η with all of the genetic variants. We use the notation $\eta \odot X$ for the element-wise product of η with each column of X . In other words, $\eta \odot X = \text{diag}(\eta) X$ where $\text{diag}(\eta)$ is an $N \times N$ diagonal matrix with η as the diagonal. The vector of interaction effect sizes γ has dimension $M \times 1$.

We chose to use MoG priors on both the main genetic effects (β) and the interaction effects (γ) as this prior is very flexible and spans the range of genetic architectures from polygenic to a

very sparse model. The priors are

$$\beta_j | \sigma_e^2, \lambda_\beta, \sigma_{\beta,1}^2, \sigma_{\beta,2}^2 \sim \lambda_\beta \mathcal{N}(0, \sigma_e^2 \sigma_{\beta,1}^2) + (1 - \lambda_\beta) \mathcal{N}(0, \sigma_e^2 \sigma_{\beta,2}^2), \quad (8)$$

$$\gamma_j | \sigma_e^2, \lambda_\gamma, \sigma_{\gamma,1}^2, \sigma_{\gamma,2}^2 \sim \lambda_\gamma \mathcal{N}(0, \sigma_e^2 \sigma_{\gamma,1}^2) + (1 - \lambda_\gamma) \mathcal{N}(0, \sigma_e^2 \sigma_{\gamma,2}^2) \quad (9)$$

We use standard Gaussian priors on the covariate and error terms.

$$\alpha | \sigma_\alpha^2 \sim \mathcal{N}(0, \sigma_\alpha^2), \quad (10)$$

$$\epsilon | \sigma_e^2 \sim \mathcal{N}(0, \sigma_e^2). \quad (11)$$

Variational Inference For notational convenience we define $\theta = \{\alpha, \beta, \gamma, w\}$ as the set of latent variables, $\mathcal{D} := \{X, E\}$ the genetic and environmental data and ϕ as the set of hyper parameters. Then the posterior $p(\theta | y, \mathcal{D}, \phi)$ is given by

$$p(\theta | y, \mathcal{D}, \phi) \propto p(y | \theta, \mathcal{D}, \phi) \prod_c p(\alpha_c | \phi) \prod_l p(w_l) \prod_j p(\beta_j, u_j | \phi) \prod_j p(\gamma_j, v_j | \phi). \quad (12)$$

To evaluate the posterior we use the variational inference framework; approximating the true posterior $p(\theta | y, \mathcal{D}, \phi)$ with a tractable alternative distribution $q(\theta; \nu)$ governed by (variational) parameters ν . To make inference tractable we use the standard Mean Field assumption so that $q(\theta; \nu)$ factorizes

$$q(\theta; \nu) = \prod_c q(\alpha_c) \prod_l q(w_l) \prod_j q(\beta_j, u_j) \prod_j q(\gamma_j, v_j). \quad (13)$$

To make $q(\theta; \nu)$ a close approximation of the true posterior we minimize the KL Divergence between $q(\theta; \nu)$ and $p(\theta | y, \mathcal{D}, \phi)$ with respect to variational parameters ν . In this manner, the problem has been transformed from one of computing posterior distributions into one of optimization.

We can show that minimizing the KL Divergence is equivalent to maximizing a lower bound on the marginal log likelihood by observing

$$KL(q||p) = -\mathbb{E}_q \left[\log \frac{p(\theta|y, \mathcal{D}, \phi)}{q(\theta; \nu)} \right], \quad (14)$$

$$= -\mathbb{E}_q \left[\log \frac{p(\theta, y|\mathcal{D}, \phi)}{q(\theta; \nu)} \right] + \mathbb{E}_q [\log p(y|\mathcal{D}, \phi)], \quad (15)$$

$$= -\mathbb{E}_q \left[\log \frac{p(\theta, y|\mathcal{D}, \phi)}{q(\theta; \nu)} \right] + \log p(y|\mathcal{D}, \phi), \quad (16)$$

thus we can write

$$\mathcal{F}(\nu; \phi) := \mathbb{E}_q \left[\log \frac{p(\theta, y|\mathcal{D}, \phi)}{q(\theta; \nu)} \right] \leq \log p(y|\mathcal{D}, \phi). \quad (17)$$

Here $\mathcal{F}(\nu; \phi)$ is commonly referred to as the Evidence Lower Bound (ELBO). Due to the factorized form of Equation (13) we can cyclically update the approximate distribution for each latent variable in turn until we reach convergence.

Our model depends on a set of eight hyper-parameters $\phi = \{\sigma_e^2, \{\sigma_{\beta,i}^2\}_{i=1}^2, \{\sigma_{\gamma,i}^2\}_{i=1}^2, \lambda_\beta, \lambda_\gamma, \sigma_\alpha^2\}$.

We set σ_α^2 to a large constant to create a flat prior on the covariates, leaving seven unknowns. Similar methods have performed a grid search over hyper-parameter values (using either cross validation¹³ or the in-sample ELBO to identify the optimum²⁰). For LEMMA a grid search would be computationally demanding, both because the set of hyper-parameters is larger and because we cannot efficiently perform multiple runs in parallel as done by Loh *et al.*¹³. Instead we maximize a lower bound on the approximate log likelihood (the so called Evidence Lower Bound or ELBO) with respect to the hyper-parameters. In this manner our approach can be viewed as a variational expectation maximization algorithm^{55,56}.

Similar to the EM algorithm, the hyper-parameter maximization step can lead to slow exploration of the hyper-parameter space and thus to slow convergence of the LEMMA algorithm. We use an accelerator, SQUAREM⁵⁷, to speed up convergence. Given two estimates of the hyper-parameters ϕ_{t-2} and ϕ_{t-1} we can adjust the maximized estimate ϕ_t with

$$\tilde{\phi}_t(v_t) = \phi_{t-2} - 2v_t\Delta\phi_{t-1} + v_t^2\Delta^2\phi_t, \quad (18)$$

where $\Delta\phi_{t-1} = \phi_{t-1} - \phi_{t-2}$ and $\Delta^2\phi_t = \phi_t - 2\phi_{t-1} + \phi_{t-2}$. Thus the new adjusted estimate $\tilde{\phi}_t(v_t)$ is a continuous function of the step size v_t , which yields the original estimate ϕ_t for $v_t = -1$. As recommended by Varadhan *et al.*⁵⁷ we set $v_t = \min(-1, -\|\Delta\phi_{t-1}\|_2^2/\|\Delta^2\phi_t\|_2^2)$. Occasionally this yields an estimate that is either outside of the domain of ϕ or leads to a state with worse ELBO than the previous state. For the first issue we use a simple backtracking method of halving the distance between v_t and -1 , and for the second we simply judge model convergence when the absolute change in ELBO drops below a given threshold. We use the same convergence criterion as the BOLT-LMM method¹³; namely that a full pass through all latent variables yields an absolute change of less than 0.01 in the approximate log-likelihood (ELBO). **Figure S17** shows the evolution of the ES parameter estimates for the four UK Biobank traits we analyzed, and illustrates that at the point of convergence the parameters appear stable.

Identifying GxE associated loci After convergence of the LEMMA variational inference algorithm, we obtain posterior mean estimates of $\hat{\beta}$, $\hat{\gamma}$ and $\hat{\eta} = E\hat{w}$. From these we construct residualized phenotypes following a Leave One Chromosome Out (LOCO) scheme;

$$y_{\text{resid-LOCO}} = y - C\hat{\alpha} - X_{\text{LOCO}}\hat{\beta}_{\text{LOCO}} - \hat{\eta} \odot X_{\text{LOCO}}\hat{\gamma}_{\text{LOCO}}. \quad (19)$$

X_{LOCO} denotes the genotype matrix excluding SNPs on the same chromosome of the test SNP, and β_{LOCO} and γ_{LOCO} are constructed similarly. Using a LOCO scheme has been shown to increase power in LMMs as the effect of the test SNP is conditioned on the effects on a large proportion of the rest of the genome^{12,15}.

For each imputed SNP, we then perform hypothesis tests $\beta_{\text{test}} \neq 0$ and $\gamma_{\text{test}} \neq 0$ using the linear model

$$y_{\text{resid-LOCO}} = x_{\text{test}}\beta_{\text{test}} + (\hat{\eta} \odot x_{\text{test}})\gamma_{\text{test}} + \epsilon, \quad (20)$$

$$= H\tau + \epsilon. \quad (21)$$

Here H is the $N \times 2$ design matrix with columns containing x_{test} and $\hat{\eta} \odot x_{\text{test}}$ respectively, and τ is the 2×1 vector containing parameters β_{test} and γ_{test} respectively, which are the main genetic effect and interaction effect of the SNP being tested.

Assuming that ϵ has mean zero and covariance matrix Ω we can use the standard OLS estimator

$$\hat{\tau} = (H^T H)^{-1} H^T y, \quad (22)$$

which (under certain regularity conditions) is asymptotically normally distributed with mean τ and variance $\text{Var}(\hat{\tau})$. By assuming the residual phenotype is homoskedastic, that is that $\Omega = \hat{\sigma}_e^2 I$, we can obtain the usual variance estimator given by

$$\text{Var}(\hat{\tau}) = \hat{\sigma}_e^2 (H^T H)^{-1}. \quad (23)$$

It has previously been observed that GxE interaction tests are likely to suffer from conditional

heteroskedasticity⁴⁹, and hence the homoskedastic variance estimator is likely to underestimate the true variance⁵⁸. We explain this phenomenon in detail in the **Supplementary Material**.

To overcome this we use robust standard errors, alternatively called Huber-White, sandwich or “heteroskedastic consistent” errors^{59,60}, that are standard tools in economics⁴⁷ and have previously been proposed for use in GxE interaction studies^{25,49,61}. We further include a small adjustment that reduces bias in small samples⁶². This yields the variance estimator

$$\text{Var}(\hat{\tau}) = (H^T H)^{-1} H^T \hat{\Sigma} H (H^T H)^{-1}, \quad (24)$$

where $\hat{\Sigma}$ is a diagonal matrix with $\hat{\Sigma}_{ii} = \frac{\hat{\epsilon}_i^2}{(1-h_{ii})^2}$, where $\hat{\epsilon} = y - H\hat{\tau}$ and $h = H(H^T H)^{-1} H^T$.

Hence our GxE test statistic is given by

$$\frac{\hat{\gamma}_{\text{test}}^2}{\text{Var}(\hat{\gamma}_{\text{test}})} \quad (25)$$

and under the null hypothesis is asymptotically distributed as χ_1^2 . As main effects tests are not sensitive to assumptions of heteroskedasticity in the same way that GxE tests are⁴⁹, we use a simple t-test to test the hypothesis $\beta_{\text{test}} \neq 0$.

Heritability estimation Previous genome wide regression methods^{20,32,63} have shown that it is possible to rearrange the model hyper-parameters to gain an estimate of trait heritability. We find that in our variational framework this approach underestimates trait heritability, due to the tendency of mean field variational inference to underestimate the posterior variance of each parameter. Instead we treat the posterior mean $\hat{\gamma}_{\text{LEMMA}}$ as a fixed effect, and use randomized HE (RHE) regression^{23,24,64} to estimate heritability with a single SNP component²³ (RHE-SC) and multiple

SNP components²⁴ (RHE-LDMS). When using the multi-component model, SNPs are stratified into a total of 20 bins; using 5 MAF bins (≤ 0.1 , $0.1 < \text{MAF} \leq 0.2$, $0.2 < \text{MAF} \leq 0.3$, $0.3 < \text{MAF} \leq 0.4$, $0.4 < \text{MAF} \leq 0.5$) and 4 LD score quantiles.

The single component model is given by

$$y \sim \mathcal{N}\left(E\alpha, \sigma_\beta^2 K + \sigma_\gamma^2 \hat{V} + \sigma_e^2 I\right), \quad (26)$$

where $K = XX^T/M$, $\hat{V} = Z(\hat{\eta})Z(\hat{\eta})^T/M$ and $Z(\hat{\eta}) = \text{diag}(\hat{\eta})X$. Haseman-Elston (HE) regression is a method of moments estimator that fits the variance components $(\sigma_\beta^2, \sigma_\gamma^2, \sigma_e^2)$ to minimize the difference between the empirical and expected covariances. This is mathematically equivalent to solving the following linear system

$$\begin{pmatrix} \text{tr}(K^2) & \text{tr}(KV) & \text{tr}(K) \\ \text{tr}(KV) & \text{tr}(V^2) & \text{tr}(K) \\ \text{tr}(K) & \text{tr}(V) & N \end{pmatrix} \begin{pmatrix} \sigma_\beta^2 \\ \sigma_\gamma^2 \\ \sigma_e^2 \end{pmatrix} = \begin{pmatrix} y^T K y \\ y^T V y \\ y^T y \end{pmatrix} \quad (27)$$

Wu *et al.*²³ showed that this system can be solved in $\mathcal{O}(NMB)$ time (for small B) without ever forming the kinship matrices K and V using Hutchinson's estimator, and that covariates can be efficiently projected out of the phenotype, genotypes and interaction matrix Z with minimal additional cost. Pazokitoroudi *et al.*²⁴ give an extension to multiple components and showed that variance estimates can be obtained with the block jackknife.

Speed *et al.*²⁷ show that the usual form for h_G^2 , the proportion of trait variance explained by additive genetic effects, given by

$$\hat{h}_G^2 = \frac{\hat{\sigma}_\beta^2}{\hat{\sigma}_\beta^2 + \hat{\sigma}_\gamma^2 + \hat{\sigma}_e^2}, \quad (28)$$

holds only when genotype matrix X is standardized to have column mean zero and column variance one. Whilst this is true in expectation for $\hat{Z}$ (assuming that $\text{Cov}(\hat{\eta}, X_j) = 0, \quad \forall j \in \{1, M\}$), this is not guaranteed. To obtain column mean zero we include an intercept of ones among the covariates that are projected out of the phenotype, genotypes and interaction matrix. To account for columns having variance not equal to one, we use a more general form of the heritability estimator (see Speed *et al.*²⁷ for details)

$$\hat{h}_{\text{GxE}}^2 = \frac{\hat{\sigma}_\gamma^2 \text{tr}(\hat{V}) / N}{\sigma_\beta^2 + \hat{\sigma}_\gamma^2 \text{tr}(\hat{V}) / N + \hat{\sigma}_e^2}. \quad (29)$$

Implementation and computational efficiency We provide software implementing the LEMMA algorithm in C++ from <https://jmarchini.org/lemma/>. We implement a number of steps to improve computational and memory efficiency including vectorization using SIMD extensions, compressed data formats, pre-computing quantities, parallel computing with OpenMPI, use of the well optimized Intel Math Kernel Library. Full details are given in the **Supplementary Material**.

Detecting squared environmental dependence By default, each of the L environmental variables is tested against the phenotype for significant squared effects. To do this LEMMA tests the hypothesis $\beta_l \neq 0$ using the following linear model

$$y = \mathbf{1}\alpha_0 + C\alpha + E_l^2\beta_l + \epsilon. \quad (30)$$

The squared effect of any environmental variables with a p-value less than 0.01 (Bonferroni correction for L multiple tests) are added to the matrix of covariates C .

Controlling for covariates Unlike in BOLT-LMM¹³, it is not possible to efficiently project covariates out of the model (y, X, Z) , because the multiplicative interaction matrix Z changes after each pass through the data. Instead the LEMMA software package can either regress covariates out of the phenotype or model the covariates as random effects in the variational framework. For our analyses of the UK Biobank we included all covariates within the variational model.

Comparison to existing GxE methods We compare LEMMA to three other single SNP methods that jointly model interactions with multiple environments. The first comparison method, StructLMM⁷, is a method that uses a random effects term u to model environmental similarity instead of genetic similarity. Specifically, StructLMM uses the model

$$y \sim \mathcal{N}(C\alpha + x_{\text{test}}\beta, \sigma_{\text{GxE}}^2 \text{diag}(x_{\text{test}}) \Sigma \text{diag}(x_{\text{test}}) + \sigma_e^2 \Sigma + \sigma_n^2 I), \quad (31)$$

to test the hypothesis $\sigma_{\text{GxE}}^2 \neq 0$. Here C is the matrix of covariates with fixed effects α , x_{test} is the focal variant and $\Sigma = EE^T$ is the environmental similarity matrix (where E is an $N \times L$ matrix of environmental variables). Although StructLMM provides both an interaction test and a joint test that looks for non-zero main and interaction effects at each SNP, we use only the interaction test in our comparisons. Finally we note that StructLMM recommends ‘gaussianizing’ the phenotype as a pre-processing step; however we just center and scale the phenotype for consistency with our other methods.

Our second and third comparison methods use equivalent information to StructLMM in a

fixed effects framework. Consider the linear model

$$y = C\alpha + E\alpha' + x_{\text{test}}\beta_{\text{test}} + x_{\text{test}} \odot E\gamma + \epsilon \quad (32)$$

$$= H\tau + \epsilon. \quad (33)$$

where H is formed from column-wise concatenation of $[C, E, x_{\text{test}}, \text{diag}(x_{\text{test}})E]$ and τ is the corresponding vector of fixed effects. Let R be the indicator matrix such that $R\tau = \gamma$. We wish to test the null hypothesis $H_0 : \gamma = 0$. Assuming that ϵ has mean zero and covariance matrix Ω we can use the standard OLS estimator $\hat{\tau} = (H^T H)^{-1} H^T y$ which (under certain regularity conditions) is asymptotically distributed as normal with mean τ and variance given by $\text{Var}(\hat{\tau}) = (H^T H)^{-1} H^T \Omega H (H^T H)^{-1}$. Assuming homoskedasticity yields the standard F test statistic

$$F_{\text{test}} = \frac{(R\hat{\tau})^T (R(H^T H)^{-1} R^T)^{-1} (R\hat{\tau}) / L}{\hat{\sigma}_e^2}, \quad (34)$$

which follows an $F_{d_1-d_0, N-d_1}$ distribution under the null hypothesis, where d_1 is the column rank of H and d_0 is the column rank of H under the null hypothesis. Alternatively we can use the same robust standard error used in the LEMMA test statistic

$$F_{\text{robust}} = (R\hat{\tau})^T (R(H^T H)^{-1} H^T \hat{\Omega} H (H^T H)^{-1} R^T)^{-1} (R\hat{\tau}), \quad (35)$$

where $\hat{\Omega}$ is a diagonal matrix with $\hat{\Omega}_{ii} = \frac{\hat{\epsilon}_i^2}{(1-h_{ii})^2}$, $\hat{\epsilon} = y - H\hat{\tau}$ and $h = H(H^T H)^{-1} H$. Then F_{robust} is asymptotically distributed as $\chi_{d_3}^2$ where d_3 is the rank of HR^T . In our simulations we refer to this as the robust F-test.

SNP specific interaction profile The SNP specific interaction profile is defined as $\eta_{LS} = Ew_{LS}$ where w_{LS} is the least squares parameter estimate of w in the single SNP model

$$y = C\alpha + x_{\text{test}}\beta_{\text{test}} + x_{\text{test}} \odot Ew + \epsilon \quad (36)$$

and y , C and E are the data matrices defined in **Online Methods**. The correlation between η_{LS} for a given SNP and the ES estimated by LEMMA can be viewed as a proxy for how well LEMMA captures the GxE interactions at that locus.

UK Biobank analysis We used real genotype and phenotype data from the UK Biobank, which is a large prospective cohort study of approximately 500,000 individuals living in the UK³. To account for potential confounding effects of population structure, we first subset down to the white British subset of 344,068 individuals used by Bycroft *et al.*³ in a GWAS on human height. This represents unrelated individuals who self-report white British ethnicity and whose genetic data projected onto principal components lies within the white British cluster³. After sub-setting down to individuals who had complete data across the phenotype, covariates and environmental factors (see below) we were left with approximately 280,000 samples per trait (**Table S1**). Finally we filtered genetic data based on minor allele frequency (≥ 0.01) and IMPUTE info score (≥ 0.3), leaving approximately 642,000 genotyped variants (**Table S1**) and 10,295,038 imputed variants per trait. For each trait we included age³, age² × gender, age³ × gender, a binary indicator for the genotype chip and the top 20 genetic principal components as additional covariates.

BMI was derived from height and weight measurements made during the first assessment visit (instance ‘0’ of field 21001), and readings more than six standard deviations from the popu-

lation mean to missing. $\log\text{BMI}$ and $\text{INT}(\text{BMI})$ refer to BMI after applying a log transformation and an inverse normal transformation (applied separately to males and females) respectively.

After calculating the mean SBP and DBP using automated blood pressure readings from the first assessment visit (fields 4080 and 4079), we adjusted for medication usage by adding 15mmHg and 10mmHg to SBP and DBP respectively⁶⁵. Data from manual measurements (fields 93 and 94) was used in the rare instance that no automated reading was available. Blood pressure readings more than four standard deviations from the mean were set to missing. PP was then calculated as SBP minus DBP.

For our GxE analyses, we made use of 42 environmental variables from the UK Biobank, similar to those used in previous GxE analyses of BMI in the UK Biobank^{7,26}. From the data provided by the UK Biobank we included 7 continuous environmental variables (“Age when attended assessment centre”, “Sleep duration”, “Time spent watching television”, “Number of days/week walked 10+ minutes”, “Number of days/week of moderate physical activity 10+ minutes”, “Number of days/week of vigorous physical activity 10+ minutes”, “Townsend deprivation index at recruitment”), 1 ordinal environmental variable (“Alcohol intake frequency”), 9 dietary ordinal variables (“Salt added to food”, “Oily fish intake”, “Non-oily fish intake”, “Processed meat intake”, “Poultry intake”, “Beef intake”, “Lamb intake”, “Pork intake”, “Cheese intake”) and 2 dietary continuous variables (“Tea intake”, “Cooked vegetable intake”). We further derived 1 categorical variable (“Is Current Smoker” from the responses given in the UK Biobank field “Smoking status”) and 1 continuous variable (“Sleep sd”; the number of standard deviations from the population mean

sleep duration). For analyses of blood pressure, we additionally included one further continuous variable “Waist circumference”. This left 11 dietary variables and 10 non-dietary variables (11 for blood pressure traits). In addition we included multiplicative interactions between participants age and gender with all non-dietary variables, and included the main effect of gender, giving the data matrix E a total of 42 columns (45 for blood pressure traits). Before running LEMMA each column was standardized as

$$E_{ij} = \frac{E_{ij} - \text{mean}(E_{:,j})}{\text{sd}(E_{:,j})}. \quad (37)$$

In all cases, where participants responded with “Prefer not to answer”, “Do not know” or “None of the above” we set the value to missing. For 3 continuous variables (“Time spent watching television”, “Tea intake”, “Cooked vegetable intake”) we removed the 99th percentile and for “Sleep duration” we removed both the 1st and 99th percentiles.

After running LEMMA, we found it convenient to interpret weights corresponding to a re-scaled data matrix E_1 . Assuming the column space of E and E_1 is the same, weights w_1 that correspond to E_1 can be extracted from the ES using least squares

$$w_1 = (E_1^T E_1)^{-1} E_1^T \hat{\eta}_{\text{LEMMA}}, \quad (38)$$

where $\hat{\eta}_{\text{LEMMA}}$ represents the ES. We note that although multivariate linear regression is invariant to a re-scaling of the design matrix, ridge regression is not due to the penalization place on the magnitude of the learned parameters. However, as the magnitude of the weights from our UK Biobank analysis is typically small (less than 0.2) compared to the standard deviation of our Gaussian prior (1) in this case the re-scaling makes minimal difference.

Re-coded data matrix E_1 was formed with one column for each of the 11 dietary variables (normalized to have mean zero and variance one), and three columns for each of the 10 (11 for blood pressure traits) non-dietary variables; the first augmented by a binary male indicator vector, the second by a binary female indicator vector and the third by a continuous vector of participant age. Columns augmented by male and female binary indicator vectors were normalized to have mean zero and variance one (not including zeros due to augmentation), apart from age (scaled to represent the number of decades aged past 40). Columns augmented by age were normalized first and then multiplied by age on the per decade scale. We further included indicator columns for men and women, which can be interpreted as gender specific intercepts and is equivalent to including an intercept and a binary column for only one gender (men or women). We note this leaves 43 (46) columns where the extra column comes from including an intercept within the column space of E_1 and is necessary because some columns have mean not equal to zero. Thus the column space of E_1 is equivalent to E under the constraint that the ES has mean zero.

Simulation studies Genetic data was sub-sampled from the UK Biobank, by default using $N = 25K$ unrelated individuals of mixed ancestry and $M = 100K$ genotyped SNPs. Environmental variables were simulated from a standard gaussian distribution. By default we constructed phenotypes with 2500 causal main effects and 1250 causal interaction effects explaining 20% and 5% of trait variance respectively. For each phenotype we constructed a weighted average of the environmental variables which we used to simulate multiplicative interaction effects. Environments with a non-zero weight are referred to as active. All non-zero effects were drawn from SNPs in the first half of each chromosome, allowing us to test the calibration of each method on ‘null’ SNPs from

the second half of each chromosome. To allow direct power comparisons across different scenarios we included an additional 60 SNPs with standardized effect sizes, that together accounted for 1% of trait variance with their main effects and 1% of trait variance with their interaction effects. Finally, a further 1% of trait variance was modeled using the first genetic principal component (PC). For all methods we included the first genetic PC as a covariate. For each method we calculated power as the proportion of the SNPs of standardized effect identified at a threshold of $p < 0.01$.

In simulations used to test randomized HE-regression, phenotypes were constructed with 10,000 causal main effects explaining 20% of trait variance and, in simulations with non-zero GxE heritability, 10,000 causal SNPs with interaction effects.

Model misspecification We simulated a scenario where a disease trait Y depends non-linearly on a heritable environmental factor S . More explicitly suppose that X is the centered and scaled genotype matrix so that columns have mean zero and variance one, that S is modeled as

$$S = X\tau + \epsilon_s, \quad (39)$$

where $\epsilon_s \sim \mathcal{N}(0, (1 - h_\tau^2)I)$, τ models random SNP effects for S and trait Y is given by

$$Y|a = aS^2 + X\beta + \epsilon. \quad (40)$$

Here a is a constant that we use to control the strength of the contribution of S^2 to Y , $\epsilon \sim \mathcal{N}(0, (1 - h_\beta^2))$ and β is the random SNP effects for Y . For simulation we suppose that τ and β have spike and slab priors

$$\tau_j | v_j \sim v_j \mathcal{N}\left(0, \frac{h_\tau^2}{P\lambda_\tau}\right) + (1 - v_j)\delta_0(\tau_j) \quad (41)$$

$$\beta_j | u_j \sim u_j \mathcal{N}\left(0, \frac{h_\beta^2}{P\lambda_\beta}\right) + (1 - u_j)\delta_0(\beta_j), \quad (42)$$

$$v_j \sim \text{Ber}(\lambda_\tau). \quad (43)$$

$$u_j \sim \text{Ber}(\lambda_\beta). \quad (44)$$

Data availability

The genetic and phenotype datasets generated by UK Biobank analyzed during the current study are available via the UK Biobank data access process. The Resource is available to all bona fide researchers, from academic, charity, public, and commercial sectors, for all types of health-related research that is in the public interest, without preferential or exclusive access for any person. More details are available here <http://www.ukbiobank.ac.uk/register-apply/>

URLs

UK Biobank: <http://www.ukbiobank.ac.uk>

LEMMA: <https://jmarchini.org/lemma/>

StructLMM as implemented in LIMIX 2.0.0: <https://github.com/limix/limix>

Acknowledgments

Computation used the Oxford Biomedical Research Computing (BMRC) facility, a joint development between the Wellcome Centre for Human Genetics and the Big Data Institute supported by Health Data Research UK and the NIHR Oxford Biomedical Research Centre. Financial support was provided by the Wellcome Trust Core Award Grant Number 203141/Z/16/Z. The views expressed are those of the author(s) and not necessarily those of the NHS, the NIHR or the Department of Health. We are grateful to Kevin Sharp, David Steinsaltz and Helen Warren for discussions about this work.

Author contributions

J.M. and M.K. conceived the ideas for the model and methods development. M.K. conducted all analyses, derived the variational Bayes updates and developed the software that implemented the methods with guidance from J.M. M.K. and J.M. wrote the manuscript. J.M. carried out this work while affiliated with the University of Oxford.

References

1. Hunter, D. J. Gene-environment interactions in human diseases. *Nature Reviews Genetics* **6**, 287–298 (2005).
2. Smith, PG and Day, N. The design of case-control studies: The influence of confounding and interaction effects. *International Journal of Epidemiology* **13**, 356–365 (1984).
3. Bycroft, C. *et al.* The UK Biobank resource with deep phenotyping and genomic data. *Nature* **562**, 203–209 (2018).
4. Kilpeläinen, T. O. *et al.* Physical activity attenuates the influence of FTO variants on obesity risk: A meta-analysis of 218,166 adults and 19,268 children. *PLoS Medicine* **8** (2011).
5. Ahmad, S. *et al.* Gene x Physical Activity Interactions in Obesity: Combined Analysis of 111,421 Individuals of European Ancestry. *PLoS Genetics* **9**, 1–9 (2013).
6. Robinson, M. R. *et al.* Genotype-covariate interaction effects and the heritability of adult body mass index. *Nature Genetics* **49**, 1174–1181 (2017).
7. Moore, R. *et al.* A linear mixed model approach to study multivariate gene-environment interactions. *Nat Genet* 180–186 (2019).
8. Young, A. I., Wauthier, F. L. & Donnelly, P. Identifying loci affecting trait variability and detecting interactions in genome-wide association studies. *Nature Genetics* **50**, 1608—1614 (2018).

9. de Leeuw, C. A., Stringer, S., Dekkers, I. A., Heskes, T. & Posthuma, D. Conditional and interaction gene-set analysis reveals novel functional pathways for blood pressure. *Nature Communications* **9** (2018).
10. Burgoine, T., Sarkar, C., Webster, C. J. & Monsivais, P. Examining the interaction of fast-food outlet exposure and income on diet and obesity: Evidence from 51,361 UK Biobank participants. *International Journal of Behavioral Nutrition and Physical Activity* **15**, 1–12 (2018).
11. Listgarten, J. *et al.* Improved linear mixed models for genome-wide association studies. *Nature Methods* **9**, 525–526 (2012).
12. Yang, J., Zaitlen, N. A., Goddard, M. E., Visscher, P. M. & Price, A. L. Advantages and pitfalls in the application of mixed-model association methods (2014).
13. Loh, P. R. *et al.* Efficient Bayesian mixed-model analysis increases association power in large cohorts. *Nature Genetics* **47**, 284–290 (2015).
14. Eskin, E. *et al.* Efficient Control of Population Structure in Model Organism Association Mapping. *Genetics* **178**, 1709–1723 (2008).
15. Lippert, C. *et al.* FaST linear mixed models for genome-wide association studies. *Nature Methods* **8**, 833–835 (2011).
16. Zhou, X. & Stephens, M. Genome-wide efficient mixed-model analysis for association studies. *Nature Genetics* **44**, 821–824 (2012).

17. Hayes, B., Goddard, M. *et al.* Prediction of total genetic value using genome-wide dense marker maps. *Genetics* **157**, 1819–1829 (2001).
18. de los Campos, G., Hickey, J. M., Pong-Wong, R., Daetwyler, H. D. & Calus, M. P. L. Whole-genome regression and prediction methods applied to plant and animal breeding. *Genetics* **193**, 327–345 (2013).
19. Logsdon, B. A., Hoffman, G. E. & Mezey, J. G. A variational Bayes algorithm for fast and accurate multiple locus genome-wide association analysis. *BMC Bioinformatics* **11** (2010).
20. Carbonetto, P. & Stephens, M. Scalable variational inference for bayesian variable selection in regression, and its accuracy in genetic association studies. *Bayesian Analysis* **7**, 73–108 (2012).
21. Zhou, X. *et al.* Polygenic Modeling with Bayesian Sparse Linear Mixed Models. *PLoS Genetics* **9**, e1003264 (2013).
22. Tyrrell, J. *et al.* Gene—obesogenic environment interactions in the UK Biobank study. *International Journal of Epidemiology* **46**, 559–575 (2017).
23. Wu, Y. & Sankararaman, S. A scalable estimator of SNP heritability for biobank-scale data. *Bioinformatics* **34**, i187–i194 (2018).
24. Pazokitoroudi, A. *et al.* Scalable multi-component linear mixed models with application to SNP heritability estimation. *bioRxiv* 522003 (2019).

25. Tchetgen, E. J. T. & Kraft, P. On the robustness of tests of genetic associations incorporating gene-environment interaction when the environmental exposure is misspecified. *Epidemiology* **22**, 257–261 (2011).
26. Young, A. I., Wauthier, F. & Donnelly, P. Multiple novel gene-by-environment interactions modify the effect of FTO variants on body mass index. *Nature Communications* **7**, 12724 (2016).
27. Speed, D., Hemani, G., Johnson, M. R. & Balding, D. J. Improved heritability estimation from genome-wide SNPs. *American Journal of Human Genetics* (2012).
28. Evans, L. M. *et al.* Comparison of methods that use whole genome data to estimate the heritability and genetic architecture of complex traits. *Nature Genetics* **50**, 737–745 (2018).
29. Yang, J. *et al.* Genetic variance estimation with imputed variants finds negligible missing heritability for human height and body mass index. *Nature Genetics* **47**, 1114–1120 (2015).
30. Hou, K. *et al.* Accurate estimation of SNP-heritability from biobank-scale data irrespective of genetic architecture. *Nature Genetics* **51**, 1244–1251 (2019).
31. Visscher, P. M., Goddard, M. E., Derks, E. M. & Wray, N. R. Evidence-based psychiatric genetics, AKA the false dichotomy between common and rare variant hypotheses. *Molecular Psychiatry* **17**, 474–485 (2012).
32. Powell, J. E. *et al.* Signatures of negative selection in the genetic architecture of human complex traits. *Nature Genetics* **50**, 746–753 (2018).

33. Rask-Andersen, M., Karlsson, T., Ek, W. E. & Johansson, Å. Gene-environment interaction study for BMI reveals interactions between genetic factors and physical activity, alcohol consumption and socioeconomic status. *PLOS Genetics* **13**, e1006977 (2017).
34. Qi, Q. *et al.* Television watching, leisure time physical activity, and the genetic predisposition in relation to body mass index in women and men. *Circulation* **126**, 1821–1827 (2012).
35. Qi, Q. *et al.* Fried food consumption, genetic risk, and body mass index: Gene-diet interaction analysis in three US cohort studies. *BMJ (Online)* **348**, 1–12 (2014).
36. Qi, Q. *et al.* Sugar-Sweetened Beverages and Genetic Risk of Obesity From the Departments of Nutrition (Q. *NEJM.org*. *N Engl J Med* **15**, 1387–96 (2012).
37. Warren, H. R. *et al.* Genome-wide association analysis identifies novel blood pressure loci and offers biological insights into cardiovascular risk. *Nature genetics* **49**, 403–415 (2017).
38. Kaplan, R. C. *et al.* A genome-wide association study identifies novel loci associated with circulating IGF-I and IGFBP-3. *Human Molecular Genetics* **20**, 1241–1251 (2011).
39. Richardson, A. S. *et al.* Moderate to vigorous physical activity interactions with genetic variants and body mass index in a large US ethnically diverse cohort. *Pediatric Obesity* **9**, e35–e46 (2014).
40. Wang, H. *et al.* Genotype-by-environment interactions inferred from genetic effects on phenotypic variability in the uk biobank. *Science Advances* **5** (2019).

41. Bhattacharyya, D. & Glick, B. S. Two Mammalian Sec16 Homologues Have Nonredundant Functions in Endoplasmic Reticulum (ER) Export and Transitional ER Organization. *Molecular biology of the cell* **18**, 986–994 (2007).
42. Thorleifsson, G. *et al.* Genome-wide association yields new sequence variants at seven loci that associate with measures of obesity. *Nature genetics* **41**, 18–24 (2009).
43. Hotta, K. *et al.* Association between obesity and polymorphisms in SEC16B, TMEM18, GNPDA2, BDNF, FAIM2 and MC4R in a Japanese population. *Journal of Human Genetics* **54**, 727–731 (2009).
44. Schmid, P. M. *et al.* Expression of fourteen novel obesity-related genes in Zucker diabetic fatty rats. *Cardiovascular Diabetology* **11**, 1–11 (2012).
45. Abdellaoui, A. *et al.* Genetic correlates of social stratification in Great Britain. *Nature Human Behaviour* **3**, 1332–1342 (2019).
46. Yang, J., Lee, S. H., Goddard, M. E. & Visscher, P. M. Gcta: a tool for genome-wide complex trait analysis. *The American Journal of Human Genetics* **88**, 76–82 (2011).
47. Greene, W. H. *Econometric Analysis 5th edition* (Pearson Education India, 2003).
48. Chang, C. C. *et al.* Second-generation PLINK: rising to the challenge of larger and richer datasets. *GigaScience* **4**, 7 (2015).

49. Almli, L. M. *et al.* Correcting systematic inflation in genetic association tests that consider interaction effects application to a genome-wide association study of posttraumatic stress disorder. *JAMA Psychiatry* **71**, 1392–1399 (2014).
50. Finucane, H. K. *et al.* Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nature Genetics* **47**, 1228–1235 (2015).
51. Speed, D., Cai, N., Johnson, M. R., Nejentsev, S. & Balding, D. J. Reevaluation of SNP heritability in complex human traits. *Nature Genetics* **49**, 986–992 (2017).
52. Gazal, S. *et al.* Linkage disequilibrium-dependent architecture of human complex traits shows action of negative selection. *Nature Genetics* **49**, 1421–1427 (2017).
53. Speed, D. & Balding, D. J. SumHer better estimates the SNP heritability of complex traits from summary statistics. *Nature Genetics* **51**, 277–284 (2019).
54. Erbe, M. *et al.* Improving accuracy of genomic predictions within and between dairy cattle breeds with imputed high-density single nucleotide polymorphism panels. *Journal of dairy science* **95**, 4114–4129 (2012).
55. Murphy, K. P. & Bach, F. *Machine Learning: A Probabilistic Perspective* (MIT Press, 2012).
56. Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. Stochastic variational inference. *The Journal of Machine Learning Research* **14**, 1303–1347 (2013).
57. Varadhan, R. & Roland, C. Simple and globally convergent methods for accelerating the convergence of any em algorithm. *Scandinavian Journal of Statistics* **35**, 335–353 (2008).

58. Wilcox, R. R. *Introduction to robust estimation and hypothesis testing* (Academic press, 2011).
59. Huber, P. J. *et al.* The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, 221–233 (University of California Press, 1967).
60. White, H. *et al.* A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *econometrica* **48**, 817–838 (1980).
61. Voorman, A., Lumley, T., McKnight, B. & Rice, K. Behavior of QQ-plots and Genomic Control in studies of gene-environment interaction. *PLoS ONE* **6** (2011).
62. Long, J. S. & Ervin, L. H. Using heteroscedasticity consistent standard errors in the linear regression model. *The American Statistician* **54**, 217–224 (2000).
63. Guan, Y. & Stephens, M. Bayesian variable selection regression for genome-wide association studies and other large-scale problems. *Annals of Applied Statistics* **5**, 1780–1815 (2011).
64. Haseman, J. & Elston, R. The investigation of linkage between a quantitative trait and a marker locus. *Behavior genetics* **2**, 3–19 (1972).
65. Tobin, M. D., Sheehan, N. A., Scurrah, K. J. & Burton, P. R. Adjusting for treatment effects in studies of quantitative traits: Antihypertensive therapy and systolic blood pressure. *Statistics in Medicine* **24**, 2911–2935 (2005).
66. Bishop, C. M. *Pattern Recognition and Machine Learning* (Springer-Verlag New York, 2006).

67. Hershey, J. R. & Olsen, P. A. Approximating the Kullback Leibler Divergence Between Gaussian Mixture Models. In *2007 IEEE International Conference on Acoustics, Speech and Signal Processing - ICASSP '07*, vol. 4, IV–317–IV–320 (2007).
68. Wang, C. & Blei, D. M. Variational inference in nonconjugate models. *Journal of Machine Learning Research* **14**, 1005–1031 (2013). 1209.4360.
69. Blei, D. M., Kucukelbir, A. & McAuliffe, J. D. Variational Inference: A Review for Statisticians. *Journal of the American Statistical Association* **112**, 859–877 (2017).
70. Band, G. & Marchini, J. Bgen: a binary file format for imputed genotype and haplotype data. *bioRxiv* (2018).
71. Fernando, R. L., Dekkers, J. C. & Garrick, D. J. A class of Bayesian methods to combine large numbers of genotyped and non-genotyped animals for whole-genome analyses. *Genetics Selection Evolution* **46**, 1–13 (2014).

Trait	log(BMI)	log(BMI)	DBP
SNP	rs539515	rs2153960	rs8090962
Chr	1	6	18
BP	177889025	108988184	56694404
A0	A	G	A
A1	C	A	G
AF	0.21	0.71	0.44
Nearest gene	<i>SEC16B</i>	<i>FOXO3</i>	<i>OACYLP / SEC11C</i>
Standard GWAS tests			
β_G (se)	0.0043 (0.0003)	0.0014 (0.0003)	-0.00001 (0.0192)
p-value _G	5.7×10^{-51}	1.1×10^{-6}	$1.0 \times 10^{+0}$
LEMMA association tests			
β_G (se)	0.0254 (0.0016)	0.0087 (0.0015)	0.0011 (0.0015)
$\beta_{G \times E}$	-0.0117 (0.0017)	-0.0098 (0.0017)	0.0087 (0.0016)
p-value _G	1.6×10^{-60}	1.6×10^{-8}	4.5×10^{-1}
p-value _{G$\times$E}	6.5×10^{-12}	6.5×10^{-9}	3.6×10^{-8}

Table 2: Loci with genome-wide significant GxE interaction effects with the ES. Independent loci with genome-wide significant ($P < 5 \times 10^{-8}$) GxE interaction effects with the environmental score (ES). Loci at least 0.5cM apart were judged to be independent. SNP effect sizes reported on a per s.d scale. SNP locations follow the GrCh37 human genome assembly. All loci had IMPUTE INFO score > 0.99 . Abbreviations are as follows; A0, reference allele; A1, alternative allele; AF, reference allele frequency; s.d, standard deviation.

Supplementary Material

Derivation of Variational Bayes updates We use Coordinate Ascent Variational Inference (CAVI) to optimize the ELBO⁶⁶. CAVI is a cyclic optimization strategy that iteratively maximizes the ELBO with respect to each latent variable whilst holding the others fixed. We now provide a brief justification of the CAVI update step and then derive the update for each of the latent variables in the LEMMA model.

Using the fact that the variational distributions factorizes, we can write the ELBO as

$$\begin{aligned}\mathcal{F}(\nu; \phi) &= \mathbb{E}_q [\log p(\theta, y | \mathcal{D}, \phi) - \log q(\theta)], \\ &= \log p(y | \mathcal{D}, \phi) + \mathbb{E}_q [\log p(\theta | y, \mathcal{D}, \phi)] - \sum_j \mathbb{E}_q [\log q(\theta_j)].\end{aligned}$$

Hence it is relatively simple to extract out dependance of $\mathcal{F}(\nu; \phi)$ on θ_j

$$\begin{aligned}\mathcal{F}_j &= \mathbb{E}_q [\log p(\theta_j | y, \theta_{-j}, \mathcal{D}, \phi)] - \mathbb{E}_q [\log q(\theta_j)] + \text{const}, \\ &= \int q(\theta_j) \left(\int \prod_{i \neq j} q(\theta_i) \log p(\theta_j | y, \theta_{-j}, \mathcal{D}, \phi) d\theta_{-j} \right) d\theta_j - \int q(\theta_j) \log q(\theta_j) d\theta_j + \text{const}, \\ &= \int q(\theta_j) (\mathbb{E}_{-\theta_j} [\log p(\theta_j | y, \theta_{-j}, \mathcal{D}, \phi)]) d\theta_j - \int q(\theta_j) \log q(\theta_j) d\theta_j + \text{const}.\end{aligned}$$

The last line is proportional to the KL divergence between $\log q(\theta_j)$ and $\mathbb{E}_{-\theta_j} [\log p(\theta | y, \phi)]$, where $\mathbb{E}_{-\theta_j}$ denotes the expectation with respect to the q distributions over all variables $\{\theta_i : \theta_i \neq \theta_j\}$. Therefore to maximize the ELBO with respect to $q(\theta_j)$ we must minimize the KL divergence between $\log q(\theta_j)$ and $\mathbb{E}_{-\theta_j} [\log p(\theta | y, \phi)]$, which occurs when

$$q^*(\theta_j) \propto \mathbb{E}_{-\theta_j} [\log p(\theta_j | y, \theta_{-j}, \mathcal{D}, \phi)].$$

After applying Bayes theorem, the above CAVI step can be equivalently expressed as

$$q^*(\theta_j) \propto \exp \mathbb{E}_{-\theta_j} [\log p(y|\theta_j, \theta_{-j}, \mathcal{D}, \phi) + \log p(\theta_j|\phi)] . \quad (45)$$

Updates for SNP main effect sizes $q(\beta_j)$

The prior and conditional log-likelihood for β_j are given by

$$p(\beta_j|\phi) = \lambda_\beta \mathcal{N}(\beta_j|0, \sigma_e^2 \sigma_{\beta,1}^2) + (1 - \lambda_\beta) \mathcal{N}(\beta_j|0, \sigma_e^2 \sigma_{\beta,2}^2) , \quad (46)$$

$$\log p(y|\beta_j, \theta_{-\beta_j}, \mathcal{D}, \phi) = -\frac{1}{2\sigma_e^2} (\beta_j^2 \|X_j\|_2^2 - 2\beta_j X_j^T y_{\text{resid}, -\beta_j}) + \text{const}, \quad (47)$$

where const is a constant independent of β_j and $y_{\text{resid}, -\beta_j} = y - C\alpha - X_{-j}\beta_{-j} - \eta \odot X\gamma$. Substituting eq. (47) and eq. (46) into eq. (45) yields

$$q^*(\theta_j) \propto \exp \left(-\frac{\|X_j\|_2^2}{2\sigma_e^2} \beta_j^2 + \frac{1}{\sigma_e^2} X_j^T \mathbb{E}_{-\beta_j} [y_{\text{resid}, -\beta_j}] \beta_j \right) p(\beta_j|\phi) \quad (48)$$

as the prior is independent of $\theta_{-\beta_j}$. We now note the result

$$\exp \left(-\frac{\|X_j\|_2^2}{2\sigma_e^2} \beta_j^2 + \frac{1}{\sigma_e^2} X_j^T \mathbb{E}_{-\beta_j} [y_{\text{resid}, -\beta_j}] \beta_j \right) \mathcal{N}(\beta_j|0, \sigma_e^2 \sigma_{\beta,i}^2) = \exp \left(\frac{\mu_{j,i}^2}{2s_{j,i}^\beta} \right) \sqrt{\frac{s_{j,i}^\beta}{\sigma_e^2 \sigma_{\beta,i}^2}} \mathcal{N}(\beta_j|\mu_{j,i}^\beta, s_{j,i}^\beta), \quad (49)$$

where

$$\begin{aligned} s_{j,i}^\beta &= \frac{\sigma_e^2}{\|X_j\|_2^2 + 1/\sigma_{\beta,i}^2}, & \text{for } i = 1, 2 \\ \mu_{j,i}^\beta &= \frac{s_{j,i}^\beta}{\sigma_e^2} X_j^T \mathbb{E}_{-\beta_j} [y_{\text{resid}, -\beta_j}], & \text{for } i = 1, 2. \end{aligned}$$

Substituting eq. (49) into eq. (48) yields

$$q^*(\beta_j) \propto \lambda_\beta \exp\left(\frac{(\mu_{j,1}^\beta)^2}{2s_{j,1}^\beta}\right) \sqrt{\frac{s_{j,1}^\beta}{\sigma_e^2 \sigma_{\beta,1}^2}} \mathcal{N}(\beta_j | \mu_{j,1}^\beta, s_{j,1}^\beta) + \quad (50)$$

$$(1 - \lambda_\beta) \exp\left(\frac{(\mu_{j,2}^\beta)^2}{2s_{j,2}^\beta}\right) \sqrt{\frac{s_{j,2}^\beta}{\sigma_e^2 \sigma_{\beta,2}^2}} \mathcal{N}(\beta_j | \mu_{j,2}^\beta, s_{j,2}^\beta). \quad (51)$$

It is now clear that $q^*(\beta_j)$ is the probability density function of a mixture of gaussians

$$q^*(\beta_j) = \psi_j^\beta \mathcal{N}(\beta_j | \mu_{j,1}^\beta, s_{j,1}^\beta) + (1 - \psi_j^\beta) \mathcal{N}(\beta_j | \mu_{j,2}^\beta, s_{j,2}^\beta),$$

where the mixture components ψ_j^β and $(1 - \psi_j^\beta)$ must sum to one. Therefore

$$\psi_j^\beta = \frac{\lambda_\beta \exp\left(\frac{(\mu_{j,1}^\beta)^2}{2s_{j,1}^\beta}\right) \sqrt{\frac{s_{j,1}^\beta}{\sigma_e^2 \sigma_{\beta,1}^2}}}{\lambda_\beta \exp\left(\frac{(\mu_{j,1}^\beta)^2}{2s_{j,1}^\beta}\right) \sqrt{\frac{s_{j,1}^\beta}{\sigma_e^2 \sigma_{\beta,1}^2}} + (1 - \lambda_\beta) \exp\left(\frac{(\mu_{j,2}^\beta)^2}{2s_{j,2}^\beta}\right) \sqrt{\frac{s_{j,2}^\beta}{\sigma_e^2 \sigma_{\beta,2}^2}}}$$

or equivalently

$$\psi_j^\beta = \text{sigmoid}\left(\text{logit}(\lambda_\beta) - \frac{1}{2} \log\left(\frac{\sigma_{\beta,1}^2 s_{j,2}^\beta}{s_{j,1}^\beta \sigma_{\beta,2}^2}\right) + \frac{(\mu_{j,1}^\beta)^2}{2s_{j,1}^\beta} - \frac{(\mu_{j,2}^\beta)^2}{2s_{j,2}^\beta}\right).$$

Therefore, the update equations for $q(\beta_j)$ can be summarised as

$$s_{j,i}^\beta = \frac{\sigma_e^2}{\|X_j\|_2^2 + 1/\sigma_{\beta,i}^2}, \quad \text{for } i = 1, 2$$

$$\mu_{j,i}^\beta = \frac{s_{j,i}^\beta}{\sigma_e^2} X_j^T \mathbb{E}_{-\beta_j} [y_{\text{resid}, -\beta_j}], \quad \text{for } i = 1, 2$$

$$\psi_j^\beta = \text{sigmoid}\left(\text{logit}(\lambda_\beta) - \frac{1}{2} \log\left(\frac{\sigma_{\beta,1}^2 s_{j,2}^\beta}{s_{j,1}^\beta \sigma_{\beta,2}^2}\right) + \frac{(\mu_{j,1}^\beta)^2}{2s_{j,1}^\beta} - \frac{(\mu_{j,2}^\beta)^2}{2s_{j,2}^\beta}\right)$$

where

$$\mathbb{E}_{-\beta_j} [y_{\text{resid}, -\beta_j}] = y - C \mathbb{E}_q [\alpha] - X_{-j} \mathbb{E}_q [\beta_{-j}] - \text{diag}(\mathbb{E}_q [\eta]) X \mathbb{E}_q [\gamma].$$

Updates for SNP interaction effect sizes $q(\gamma_j)$

The derivation of the variational update for $q(\gamma_j)$ is extremely similar to that of $q(\beta_j)$. The prior and conditional log-likelihood for γ_j are given by

$$p(\gamma_j|\phi) = \lambda_\gamma \mathcal{N}(\gamma_j|0, \sigma_e^2 \sigma_{\gamma,1}^2) + (1 - \lambda_\gamma) \mathcal{N}(\gamma_j|0, \sigma_e^2 \sigma_{\gamma,2}^2), \quad (52)$$

$$\log p(y|\gamma_j, \theta_{-\gamma_j}, \mathcal{D}, \phi) = -\frac{1}{2\sigma_e^2} (\gamma_j^2 \|Z_j\|_2^2 - 2\gamma_j Z_j^T y_{\text{resid}, -\gamma_j}) + \text{const}, \quad (53)$$

where const is a constant independent of γ_j and $y_{\text{resid}, -\gamma_j} = y - C\alpha - X\beta - \eta \odot X_{-j}\gamma_{-j}$. Substituting eq. (53) and eq. (52) into eq. (45) yields

$$q^*(\gamma_j) \propto \exp \left(-\frac{\gamma_j^2}{2\sigma_e^2} \mathbb{E}_{-\gamma_j} [\|Z_j\|_2^2] + \frac{1}{\sigma_e^2} \gamma_j X_j^T \mathbb{E}_{-\gamma_j} [\eta \odot y_{\text{resid}, -\gamma_j}] \right) p(\gamma_j|\phi) \quad (54)$$

as the prior is independent of θ_{-j} . Following the same steps as used in the derivation of $q^*(\beta_j)$, is it clear that $q^*(\gamma_j)$ is also the probability density function of a mixture of gaussians

$$q^*(\gamma_j) = \psi_j^\gamma \mathcal{N}(\gamma_j|\mu_{j,1}^\gamma, s_{j,1}^\gamma) + (1 - \psi_j^\gamma) \mathcal{N}(\gamma_j|\mu_{j,2}^\gamma, s_{j,2}^\gamma),$$

whose optimal CAVI updates are given by

$$s_{j,i}^\gamma = \frac{\sigma_e^2}{\mathbb{E}_{-\gamma_j} [\|Z_j\|_2^2] + 1/\sigma_{\gamma,i}^2}, \quad \text{for } i = 1, 2$$

$$\mu_{j,i}^\gamma = \frac{s_{j,i}^\gamma}{\sigma_e^2} X_j^T \mathbb{E}_{-\gamma_j} [\eta \odot y_{\text{resid}, -\gamma_j}], \quad \text{for } i = 1, 2$$

$$\psi_j^\gamma = \text{sigmoid} \left(\text{logit}(\lambda_\gamma) - \frac{1}{2} \log \left(\frac{\sigma_{\gamma,1}^2 s_{j,2}^\gamma}{s_{j,1}^\gamma \sigma_{\gamma,2}^2} \right) + \frac{(\mu_{j,1}^\gamma)^2}{2s_{j,1}^\gamma} - \frac{(\mu_{j,2}^\gamma)^2}{2s_{j,2}^\gamma} \right)$$

where

$$\mathbb{E}_{-\gamma_j} [y_{\text{resid}, -\gamma_j}] = y - C\mathbb{E}_q [\alpha] - X\mathbb{E}_q [\beta] - \text{diag} (\mathbb{E}_q [\eta]) X_{-j}\mathbb{E}_q [\gamma_{-j}] ,$$

$$\begin{aligned} \mathbb{E}_{-\gamma_j} [||Z_j||_2^2] &= X_{\cdot j}^T \text{diag} (\mathbb{E}_q [\eta^2]) X_{\cdot j} , \\ &= \sum_{l,m} \mathbb{E}_q [w_m] \mathbb{E}_q [w_l] \underbrace{\sum_i X_{ij}^2 E_{il} E_{im}}_{\text{precomputed}} \\ &\quad + \sum_l \text{Var}_q (w_l) \underbrace{\sum_i X_{ij}^2 E_{il}^2}_{\text{precomputed}} . \end{aligned}$$

Note that computation of $\mathbb{E}_{-\gamma_j} [||Z_j||_2^2]$ is an $O(L^2 + N)$ operation due to the precomputation of $\sum_i X_{ij}^2 E_{il} E_{im}$ (and without this precomputation the compute cost of this update would be $\mathcal{O}(NL^2)$).

Updates for interaction weights $q(w_l)$

Rewriting the conditional log-likelihood makes its dependence on w clear

$$\begin{aligned} \log p(y|w, \theta_{-w}, \mathcal{D}, \phi) &= -\frac{1}{2\sigma_e^2} ||y - C\alpha - X\beta - \eta \odot X\gamma||_2^2 + \text{const}, \\ &= -\frac{1}{2\sigma_e^2} ||y - C\alpha - X\beta - Bw||_2^2 + \text{const} \end{aligned}$$

where $B = \text{diag} (X\gamma) E$ and const is a constant independent of w . For convenience we denote the l 'th column of B as B_l . Therefore the prior and conditional log-likelihood of w_l are given by

$$p(w_l) = \mathcal{N} (w_l | 0, 1) , \tag{55}$$

$$\log p(y|w_l, \theta_{-w_l}, \mathcal{D}, \phi) = -\frac{1}{2\sigma_e^2} (w_l^2 ||B_l||_2^2 - 2w_l B_l^T y_{\text{resid}, -w_l}) + \text{const}, \tag{56}$$

where $y_{\text{resid}, -w_l} = y - C\alpha - X\beta - B_{-l}w_{-l}$ and const is now a constant independent of w_l . Substituting eq. (56) and eq. (55) into eq. (45) yields

$$\begin{aligned} q^*(w_l) &\propto \exp \left(-\frac{w_l^2}{2} \mathbb{E}_{-w_l} [\|B_l\|_2^2] + w_l \mathbb{E}_{-w_l} [B_l^T y_{\text{resid}, -w_l}] \right) p(w_l), \\ &\propto \mathcal{N}(w_l | \mu_l^w, s_l^w), \end{aligned}$$

where

$$\begin{aligned} s_l^w &= \frac{\sigma_e^2}{\sigma_e^2 + \mathbb{E}_{-w_l} [\|B_l\|_2^2]}, \\ \mu_l^w &= \frac{s_l^w}{\sigma_e^2} \mathbb{E}_{-w_l} [B_l^T y_{\text{resid}, -w_l}]. \end{aligned}$$

As $q^*(w_l)$ must be a valid distribution, it is clear that $q^*(w_l) = \mathcal{N}(w_l | \mu_l^w, s_l^w)$. The quantities

$\mathbb{E}_{-w_l} [\|B_l\|_2^2]$ and $\mathbb{E}_{-w_l} [B_l^T y_{\text{resid}, -w_l}]$ can be computed as follows

$$\begin{aligned} \mathbb{E}_{-w_l} [\|B_l\|_2^2] &= \mathbb{E} \left[\sum_i E_{il}^2 \left(\sum_j X_{ij}^2 \gamma_j \right)^2 \right], \\ &= \sum_i E_{il}^2 \left(\sum_j X_{ij}^2 \mathbb{E} [\gamma_j] \right)^2 + \sum_i E_{il}^2 \sum_j X_{ij}^2 \text{Var} (\gamma_j), \\ &= y_X^T \text{diag} (E_l^2) y_X + \sum_j \text{Var} (\gamma_j) \underbrace{\sum_i E_{il}^2 X_{ij}^2}_{\text{precomputed}}. \end{aligned}$$

$$\begin{aligned} \mathbb{E}_{-w_l} [B_l^T y_{\text{resid}, -w_l}] &= (y - \hat{y}_M)^T \text{diag} (E_l^*) \hat{y}_X - E_l^* \text{diag} (\hat{y}_X^2) \mathbb{E} [\eta_{-l}] \\ &\quad - \sum_j \text{Var} (\gamma_j) \sum_{m \neq l} \mathbb{E} [w_m] \underbrace{\sum_i X_{ij}^2 E_{il} E_{im}}_{\text{precomputed}}. \end{aligned}$$

Note that computation of $\mathbb{E}_{-w_l} [B_l^T y_{\text{resid}, -w_l}]$ is an $O(NL)$ operation due to the precomputation of $\sum_i X_{ij}^2 E_{il} E_{im}$ (and without this precomputation the compute cost of this update would be $\mathcal{O}(NML)$).

Updates for covariate main effect sizes $q(\alpha_c)$

The derivation of the variational update for $q(\alpha_c)$ is extremely similar to that of $q(w_l)$. The prior and conditional log-likelihood for α_c are given by

$$p(\alpha_c|\phi) = \mathcal{N}(\alpha_c|0, \sigma_e^2 \sigma_\alpha^2), \quad (57)$$

$$\log p(y|\alpha_c, \theta_{-\alpha_c}, \mathcal{D}, \phi) = -\frac{1}{2\sigma_e^2} (\alpha_c^2 \|C_c\|_2^2 - 2\alpha_c C_c^T y_{\text{resid}, -\alpha_c}) + \text{const}, \quad (58)$$

where const is a constant independent of α_c and $y_{\text{resid}, -\alpha_c} = y - C_{-c}\alpha_{-c} - X\beta - \eta \odot X\gamma$. By similarity with the derivation of $q^*(w_l)$ it is clear that $q^*(\alpha_c)$ is a gaussian distribution, with variational updates

$$s_c^\alpha = \frac{\sigma_e^2}{1/\sigma_\alpha^2 + (N-1)},$$

$$\mu_c^\alpha = \frac{s_c^\alpha}{\sigma_e^2} C_c^T \mathbb{E}_{-\alpha_c} [y_{\text{resid}, -\alpha_c}]$$

where $\mathbb{E}_{-\alpha_c} [y_{\text{resid}, -\alpha_c}] = y - C_{-c}\mathbb{E}_q [\alpha_{-c}] - X\mathbb{E}_q [\beta] - \mathbb{E}_q [\eta] \odot X\mathbb{E}_q [\gamma]$.

Evidence lower bound Variational inference involves maximising the evidence lower bound (ELBO)

$\mathcal{F}(\phi; \nu)$ on the model log-likelihood $\log p(y|\mathcal{D}, \phi)$. The ELBO can be separated into the expected conditional log-likelihood and the KL divergence between the variational distribution and the re-

spective priors. This is given by

$$\begin{aligned}
\mathcal{F}(\phi; \nu) &= \mathbb{E}_q [\log p(y|\theta, \mathcal{D}, \phi)] - \sum_j \text{KL}(q(\theta_j; \nu_j) \| p(\theta_j | \phi)), \\
&= -\frac{N}{2} \log(2\pi\sigma_e^2) \\
&\quad - \frac{1}{2\sigma_e^2} (\|y - C\mathbb{E}_q[\alpha] - X\mathbb{E}_q[\beta] - \mathbb{E}_q[\eta] \odot X\mathbb{E}_q[\gamma]\|_2^2) \\
&\quad - \frac{1}{2\sigma_e^2} \left(\mathbb{E}_q[\gamma]^T X^T \text{diag}(\mathbb{E}_q[\eta^2]) X \mathbb{E}_q[\gamma] - \|\mathbb{E}_q[\eta] \odot X\mathbb{E}_q[\gamma]\|_2^2 \right) \\
&\quad - \frac{N-1}{2\sigma_e^2} \sum_l \text{Var}_q(\alpha_l) - \frac{N-1}{2\sigma_e^2} \sum_k \text{Var}_q(\beta_k) \\
&\quad - \sum_c^{L'} \text{KL}(q(\alpha_c) \| p(\alpha_c)) - \sum_l^L \text{KL}(q(w_l) \| p(w_l)) \\
&\quad - \sum_j^M \text{KL}(q(\beta_j) \| p(\beta_j)) - \sum_j^M \text{KL}(q(\gamma_j) \| p(\gamma_j))
\end{aligned}$$

While the KL Divergence between two univariate gaussian distributions is a standard result, the KL Divergence between two mixtures of gaussians is not analytically tractable. However the matched bound approximation⁶⁷ can be used to provide an upper bound when both have the same number of components. Thus for two mixtures of gaussians given by

$$\begin{aligned}
u &\sim \lambda \mathcal{N}(0, \sigma_1^2) + (1 - \lambda) \mathcal{N}(0, \sigma_2^2), \\
v &\sim \psi \mathcal{N}(\mu_1, s_1) + (1 - \psi) \mathcal{N}(\mu_2, s_2),
\end{aligned}$$

the matched bound on the KL divergence is given by

$$\begin{aligned}
\text{KL}(v \| u) &\leq \psi \log \frac{\psi}{\lambda} + (1 - \psi) \log \frac{1 - \psi}{1 - \lambda} \\
&\quad - \frac{1}{2} + \frac{\psi}{2} \left(\frac{s_1 + \mu_1^2}{2\sigma_1^2} - \log \left(\frac{s_1}{\sigma_1^2} \right) \right) + \frac{(1 - \psi)}{2} \left(\frac{s_2 + \mu_2^2}{2\sigma_2^2} - \log \left(\frac{s_2}{\sigma_2^2} \right) \right).
\end{aligned}$$

Use of the matched bound approximation retains a valid variational algorithm, because it maintains the lower bound on the marginal log-likelihood⁶⁸.

$$\begin{aligned}
 KL(q(\alpha_c)||p(\alpha_c)) &= -\frac{1}{2} + \frac{1}{2} \left(\frac{s_m^\alpha + (\mu_m^\alpha)^2}{\sigma_e^2 \sigma_\alpha^2} - \log \left(\frac{s_m^\alpha}{\sigma_e^2 \sigma_\alpha^2} \right) \right) \\
 KL(q(w_l)||p(w_l)) &= -\frac{1}{2} + \frac{1}{2} ((s_l^w + (\mu_l^w)^2) - \log(s_l^w)) \\
 KL(q(\beta_j)||p(\beta_j)) &= \psi_j^\beta \log \frac{\psi_j^\beta}{\lambda_\beta} + (1 - \psi_j^\beta) \log \frac{1 - \psi_j^\beta}{1 - \lambda_\beta} - \frac{1}{2} \\
 &\quad + \frac{\psi_j^\beta}{2} \left(\frac{s_{j,1}^\beta + (\mu_{j,1}^\beta)^2}{\sigma_e^2 \sigma_{\beta,1}^2} - \log \frac{s_{j,1}^\beta}{\sigma_e^2 \sigma_{\beta,1}^2} \right) + \frac{1 - \psi_j^\beta}{2} \left(\frac{s_{j,2}^\beta + (\mu_{j,2}^\beta)^2}{\sigma_e^2 \sigma_{\beta,2}^2} - \log \left(\frac{s_{j,2}^\beta}{\sigma_e^2 \sigma_{\beta,2}^2} \right) \right) \\
 KL(q(\gamma_j)||p(\gamma_j)) &= \psi_j^\gamma \log \frac{\psi_j^\gamma}{\lambda_\gamma} + (1 - \psi_j^\gamma) \log \frac{1 - \psi_j^\gamma}{1 - \lambda_\gamma} - \frac{1}{2} \\
 &\quad + \frac{\psi_j^\gamma}{2} \left(\frac{s_{j,1}^\gamma + (\mu_{j,1}^\gamma)^2}{\sigma_e^2 \sigma_{\gamma,1}^2} - \log \left(\frac{s_{j,1}^\gamma}{\sigma_e^2 \sigma_{\gamma,1}^2} \right) \right) + \frac{1 - \psi_j^\gamma}{2} \left(\frac{s_{j,2}^\gamma + (\mu_{j,2}^\gamma)^2}{\sigma_e^2 \sigma_{\gamma,2}^2} - \log \left(\frac{s_{j,2}^\gamma}{\sigma_e^2 \sigma_{\gamma,2}^2} \right) \right)
 \end{aligned}$$

Derivation of hyperparameter maximization For the maximization step we set $\phi = \hat{\phi}$ where

$\nabla_\phi F(\phi; \nu) = 0$. For ease of notation we perform the following change of variables

$$\begin{aligned}
 \tilde{\sigma}_{\beta,1}^2 &= \sigma_e^2 \sigma_{\beta,1}^2 \rightarrow \frac{\partial}{\partial \tilde{\sigma}_{\beta,1}^2} = \frac{1}{\sigma_e^2} \frac{\partial}{\partial \sigma_{\beta,1}^2}, \\
 \tilde{\sigma}_{\beta,2}^2 &= \sigma_e^2 \sigma_{\beta,2}^2 \rightarrow \frac{\partial}{\partial \tilde{\sigma}_{\beta,2}^2} = \frac{1}{\sigma_e^2} \frac{\partial}{\partial \sigma_{\beta,2}^2}, \\
 \tilde{\sigma}_{\gamma,1}^2 &= \sigma_e^2 \sigma_{\gamma,1}^2 \rightarrow \frac{\partial}{\partial \tilde{\sigma}_{\gamma,1}^2} = \frac{1}{\sigma_e^2} \frac{\partial}{\partial \sigma_{\gamma,1}^2}, \\
 \tilde{\sigma}_{\gamma,2}^2 &= \sigma_e^2 \sigma_{\gamma,2}^2 \rightarrow \frac{\partial}{\partial \tilde{\sigma}_{\gamma,2}^2} = \frac{1}{\sigma_e^2} \frac{\partial}{\partial \sigma_{\gamma,2}^2}.
 \end{aligned}$$

This makes the derivation easier as all the partial derivatives become decoupled. Partial derivatives with respect to each hyper-parameter are given by

$$\frac{\partial F}{\partial \lambda_\beta} = \sum_j \left(\frac{\psi_j^\beta}{\lambda_\beta} - \frac{(1 - \psi_j^\beta)}{1 - \lambda_\beta} \right),$$

$$\frac{\partial F}{\partial \lambda_\gamma} = \sum_j \left(\frac{\psi_j^\gamma}{\lambda_\gamma} - \frac{(1 - \psi_j^\gamma)}{1 - \lambda_\gamma} \right),$$

$$\frac{\partial F}{\partial \tilde{\sigma}_{\beta,1}^2} = \sum_j \frac{\psi_j^\beta}{2} \left(-\frac{1}{\tilde{\sigma}_{\beta,1}^2} + \frac{s_{j,1}^\beta + (\mu_{j,1}^\beta)^2}{(\tilde{\sigma}_{\beta,1}^2)^2} \right),$$

$$\frac{\partial F}{\partial \tilde{\sigma}_{\beta,2}^2} = \sum_j \frac{1 - \psi_j^\beta}{2} \left(-\frac{1}{\tilde{\sigma}_{\beta,2}^2} + \frac{s_{j,2}^\beta + (\mu_{j,2}^\beta)^2}{(\tilde{\sigma}_{\beta,2}^2)^2} \right),$$

$$\frac{\partial F}{\partial \tilde{\sigma}_{\gamma,1}^2} = \sum_j \frac{\psi_j^\gamma}{2} \left(-\frac{1}{\tilde{\sigma}_{\gamma,1}^2} + \frac{s_{j,1}^\gamma + (\mu_{j,1}^\gamma)^2}{(\tilde{\sigma}_{\gamma,1}^2)^2} \right),$$

$$\frac{\partial F}{\partial \tilde{\sigma}_{\gamma,2}^2} = \sum_j \frac{1 - \psi_j^\gamma}{2} \left(-\frac{1}{\tilde{\sigma}_{\gamma,2}^2} + \frac{s_{j,2}^\gamma + (\mu_{j,2}^\gamma)^2}{(\tilde{\sigma}_{\gamma,2}^2)^2} \right),$$

$$\frac{\partial F}{\partial \sigma_e^2} = -\frac{N}{2\sigma_e^2} + \frac{1}{2(\sigma_e^2)^2} \mathbb{E}_q [\|y - C\alpha - X\beta - \text{diag}(\eta)X\gamma\|_2^2] - \frac{M}{2\sigma_e^2} + \frac{1}{2(\sigma_e^2)^2 \sigma_\alpha^2} \sum_c (s_c^\alpha + (\mu_c^\alpha)^2)$$

Hence the maximization steps are

$$\hat{\lambda}_\beta = \frac{1}{P} \sum_j \psi_j^\beta, \quad (59)$$

$$\hat{\lambda}_\gamma = \frac{1}{P} \sum_j \psi_j^\gamma, \quad (60)$$

$$(61)$$

$$\hat{\sigma}_{\beta,1}^2 = \frac{\sum_j \psi_j^\beta (s_{j,1}^\beta + (\mu_{j,1}^\beta)^2)}{\hat{\sigma}_e^2 \sum_j \psi_j^\beta}, \quad (62)$$

$$\hat{\sigma}_{\beta,2}^2 = \frac{\sum_j (1 - \psi_j^\beta) (s_{j,2}^\beta + (\mu_{j,2}^\beta)^2)}{\hat{\sigma}_e^2 \sum_j (1 - \psi_j^\beta)}, \quad (63)$$

$$(64)$$

$$\hat{\sigma}_{\gamma,1}^2 = \frac{\sum_j \psi_j^\gamma (s_{j,1}^\gamma + (\mu_{j,1}^\gamma)^2)}{\hat{\sigma}_e^2 \sum_j \psi_j^\gamma}, \quad (65)$$

$$\hat{\sigma}_{\gamma,2}^2 = \frac{\sum_j (1 - \psi_j^\gamma) (s_{j,2}^\gamma + (\mu_{j,2}^\gamma)^2)}{\hat{\sigma}_e^2 \sum_j (1 - \psi_j^\gamma)}, \quad (66)$$

$$(67)$$

$$\hat{\sigma}^2 = \frac{\mathbb{E}_q [\|y - C\alpha - X\beta - \text{diag}(\eta) X\gamma\|_2^2] + \frac{1}{\sigma_\alpha^2} \sum_c (s_c^\alpha + (\mu_c^\alpha)^2)}{N + M}. \quad (68)$$

As an aside we note that one could use the maximized hyper-parameters (after convergence) to obtain a point estimate of $\text{Var}(\beta)$. However, by substituting in Equations (59) to (66) we can see that this is equivalent to $\sum_j \mathbb{E}_q [\beta_j^2] / M$.

$$\begin{aligned} \text{Var}(\beta) &= \lambda_\beta \sigma_{\beta,1}^2 + (1 - \lambda_\beta) \sigma_{\beta,2}^2, \\ &\approx \hat{\lambda}_\beta \hat{\sigma}_{\beta,1}^2 + (1 - \hat{\lambda}_\beta) \hat{\sigma}_{\beta,2}^2, \\ &= \frac{1}{M} \sum_j \left(\psi_j^\beta (s_{j,1}^\beta + (\mu_{j,1}^\beta)^2) + (1 - \psi_j^\beta) (s_{j,2}^\beta + (\mu_{j,2}^\beta)^2) \right), \\ &= \frac{1}{M} \sum_j \mathbb{E}_q [\beta_j^2]. \end{aligned}$$

As the mean field assumption tends to cause variational inference algorithms to underestimate the variance of latent variables⁶⁹, this is likely to produce an underestimate of $\text{Var}(\beta)$. We can observe the same result for $\text{Var}(\gamma)$ with an analogous argument.

Compressed genotype data To reduce RAM usage, LEMMA stores a compressed version of the genotype matrix using NM bytes. To do this LEMMA splits the interval $[0, 2]$ into 2^8 segments and stores the index of the segment that each dosage falls into, as well as the mean and variance for each SNP. Then when operating on a SNP, LEMMA reconstructs the centered and scaled dosages for that SNP. This approach is similar to that used by the BGEN data format⁷⁰ and results in a small loss of accuracy, but is more flexible than assuming dosages are hardcoded to $\{0, 1, 2\}$.

Computational efficiency Using mean field variational inference, estimation of the posterior means of the latent variables β, γ, w can be reduced to an iterative algorithm that cycles through the variables sequentially, updating each conditional on the values of the others. Taking the main effect of the j 'th SNP as an example, the update scheme for β_j can be written heuristically as

$$\tilde{\beta}_j = X_j^T y_{\text{resid}}, \quad (69)$$

$$\hat{\beta}_j^t = \text{regularise}(\tilde{\beta}_j; \phi^t), \quad (70)$$

$$y_{\text{resid}} = y_{\text{resid}} - (\hat{\beta}_j^t - \hat{\beta}_j^{t-1}) X_j^T. \quad (71)$$

In Equation (69) we compute the correlation between the j SNP and the residual phenotype vector.

In Equation (70) we compute the posterior mean of β_j which depends on the correlation with the residual phenotype, the prior on β_j and the current hyper-parameters. Finally in Equation (71) we update the residual phenotype vector.

The majority of computational time is spent on the dot product in Equation (69) and updating the residual phenotype in Equation (71). Both are BLAS Level 1 operations, which implies that memory access is often the principle bottleneck rather than the number of cores available. It is possible to step up to BLAS Level 2 by updating a block of SNPs in parallel¹³, however this is still a memory bound operation. Instead we use a parallel computing strategy suggested by⁷¹ for use in genome wide regression, and subsequently used by³², to compute the dot product and perform the residual update in parallel using OpenMPI. Briefly, we partition the samples such that blocks of rows of the phenotype y , genotypes X and environmental variables E are assigned to each core. For a given update step, each core calculates the dot product for the locally held block of samples and then shares the local dot product with the rest of the network. From this the dot product for the entire cohort can be reconstructed cheaply. After computing the posterior mean, each core then updates the residual phenotype for the block of samples stored locally. We observed that a distributed algorithm using OpenMPI was faster than the same algorithm using multi-threaded matrix-vector operations with the Intel MKL Library even on a single node with multiple cores. However using OpenMPI has the additional advantage of allowing users to utilize cores from across a cluster rather than being restricted to a single node. **Figure S18** shows LEMMA scales with increasing sample size.

Pre-computed quantities To aid computational efficiency we pre-compute a $M \times L(L+1)$ matrix W where

$$W_{j,m \times L+l} = \sum_i X_{ij}^2 E_{il} E_{im}, \quad \text{for } 1 \leq l < m \leq L \text{ and } 1 \leq j \leq M,$$

and is used in the updates of $q(w_l)$ and $q(\gamma_j)$. LEMMA can compute this internally, incurring a one off cost of $\mathcal{O}(NML^2)$, or is able to read from a text file at run time. As this is easily computed in parallel over batches of variants and/or environment, we recommend that for biobank scale datasets users should pre-compute this quantity beforehand using a separate tool that we have provided.

Parameter Initialization We start the variational mean estimates of $q(\beta)$ and $q(\gamma)$ at zero. To initialize mean estimates of the interaction weights $q(w)$ we have two options; the first of which is simply to use a uniform weighting over all environments. For the second we apply an F-Test independently at each SNP and use the learned coefficients from the test with the lowest p-value as the initial values of the interaction weights. We find that we often obtain similar results from both options, so for simplicity we use a uniform start point for our Biobank analyses. To initialize mean estimates of $q(\alpha)$ we use the least squares fit of C on y .

Initial values of the hyperparameters are drawn randomly from the following distributions

$$\begin{aligned} h_{\beta}^2 &\sim \mathcal{U}(0, 0.5), \\ h_{\gamma}^2 &\sim \mathcal{U}(0, 0.1), \\ -\log_{10}(\lambda_{\beta}) &\sim \mathcal{U}([2, \dots, 1 - \log_{10}(M)]) \\ -\log_{10}(\lambda_{\gamma}) &\sim \mathcal{U}([2, \dots, 1 - \log_{10}(M)]). \end{aligned}$$

We then set

$$\begin{aligned}\sigma_e^2 &= 1 - h_\beta^2 - h_\gamma^2, \\ \sigma_{\beta,1}^2 &= \frac{1}{\lambda_\beta M} \frac{h_\beta^2}{1 - h_\beta^2 - h_\gamma^2}, \\ \sigma_{\gamma,1}^2 &= \frac{1}{\lambda_\gamma M} \frac{h_\gamma^2}{1 - h_\beta^2 - h_\gamma^2},\end{aligned}$$

and initialize the spike variances at

$$\begin{aligned}\sigma_{\beta,2}^2 &= \sigma_{\beta,1}^2 / 1000, \\ \sigma_{\gamma,2}^2 &= \sigma_{\gamma,1}^2 / 1000.\end{aligned}$$

Setting the sparsity hyperparameters $\lambda_\beta, \lambda_\gamma$ in this manner allows LEMMA to start from a state where only a small number (somewhere between ten and one in one hundred) SNPs are expected to be part of the slab prior. The sparsity hyperparameters can then be updated in the variational maximization step to better reflect trait genetic architecture.

Missing data Samples with missing data in the phenotype, environmental variables or covariates are excluded. By default LEMMA imputes missing genetic data with the mean dosage of each SNP, however as LEMMA does not assume dosages are hard called with $\{0, 1, 2\}$ we recommend that users first impute genetic data with standard imputation pipelines.

Robust standard errors in GxE Studies In **Figure S19** we illustrate how a multiplicative GxE interaction effect on a quantitative trait can cause the conditional trait variance given an interacting SNP $\text{Var}(Y|g_0)$ to differ according to the interacting SNPs genotype. This is known as conditional heteroskedasticity and is the key insight behind several recent methods to detect SNPs with non-

zero GxE effects in the UK Biobank^{8,40}.

In the same figure, we can observe that the conditional trait variance given the environmental exposure $\text{Var}(y|E)$ also displays signs of conditional heteroskedasticity. Previous studies⁴⁹ have observed that methods that assume heteroskedasticity can display substantial inflation when testing for GxE effects at SNPs where there is no true GxE effect. In our simulations we observed that inflation of GxE tests statistics from LEMMA-S and the F-test, both of which assume homoskedasticity, increased with SNP-GxE heritability. Below we give an explanation for this phenomenon.

Consider a polygenic quantitative trait Y that has multiplicative GxE interactions with the same environmental exposure E at multiple SNPs

$$y_i = \alpha E_i + \sum_{j=1}^M \beta_j G_{ij} + \sum_{j=1}^M \gamma_j E_i G_{ij} + \epsilon_i, \quad (72)$$

where M is the number of SNPs and the coefficients represent true effects. For simplicity we assume that E and SNPs G_j are normalized to have mean zero and variance one, that the set of E with all causal SNPs $\{E\} \cup \{G_j : \beta_j \neq 0\}$ is pairwise independent and that the influence from population structure is negligible.

Suppose we have identified E as an environmental variable that may plausible have GxE interactions with our phenotype and we then conduct a GWAS for GxE effects. Then at the k 'th SNP we wish to test the hypothesis $\gamma_k \neq 0$ in the following linear model

$$\begin{aligned} y &= \alpha E + G_k \beta_k + E \cdot G_k \gamma_k + u, \\ &= X\tau + u, \end{aligned}$$

where in the second line $\tau = (\alpha, \beta_k, \gamma_k)^T$, X is the corresponding design matrix encapsulating all fixed effects and u in an unobserved random effects capturing residual noise. Assuming that $\mathbb{E}[u|X] = 0$, the usual least squares estimate of τ , $\hat{\tau} = (X^T X)^{-1} X^T y$, has asymptotic distribution

$$\hat{\tau} \rightarrow \mathcal{N}(\tau, \text{Var}(\hat{\tau})),$$

where

$$\begin{aligned} \text{Var}(\hat{\tau}) &= \mathbb{E}_X [\text{Var}(\hat{\tau}|X)] + \text{Var}_X (\mathbb{E}[\hat{\tau}|X]), \\ &= \mathbb{E}_X [\text{Var}(\hat{\tau}|X)] + \text{Var}_X (\tau), \\ &= \mathbb{E}_X [\text{Var}(\hat{\tau}|X)], \end{aligned}$$

and

$$\begin{aligned} \text{Var}(\hat{\tau}|X) &= \text{Var}(\tau + (X^T X)^{-1} X^T u|X), \\ &= (X^T X)^{-1} \text{Var}(X^T u|X) (X^T X)^{-T}, \\ &= (X^T X)^{-1} X^T \text{Var}(u|X) X (X^T X)^{-T}. \end{aligned}$$

The usual approach is to assume that $\text{Var}(u|X) = \sigma^2 I$ (ie homoskedasticity), which yields the standard variance estimator $\text{Var}(\hat{\tau}|X) = \sigma^2 (X^T X)^{-1}$. However, given the true generative model for y given in Equation (72), we can write u as

$$u = \sum_{j \neq k} (G_j \beta_j + E G_j \gamma_j) + \epsilon. \quad (73)$$

Therefore the conditional variance of u given X is given by

$$\begin{aligned}
 \text{Var}(u|X) &= \text{Var}(u|E = e, G_k = g_k), \\
 &= \text{Var}\left(\sum_{j \neq k} (G_j \beta_j + e G_j \gamma_j) + \epsilon\right), \\
 &= \sum_{j \neq k} \text{Var}(\beta_j G_j) + \sum_{j \neq k} \text{Var}(\gamma_j e G_j) + 2 \sum_{j \neq k} \text{Cov}(\beta_j G_j, \gamma_j e G_j) + 1 \\
 &\quad + \sum_{j \neq k, m \neq k} \text{Cov}(\beta_j G_j, \beta_m G_m) + \sum_{j \neq k, m \neq k} \text{Cov}(\gamma_j e G_j, \gamma_m e G_m) \\
 &= \sum_{j \neq k} (\beta_j + e \gamma_j)^2 + 1,
 \end{aligned}$$

where the covariances in the second line are all zero due to pairwise independence of the set $\{E\} \cup \{G_j : \beta_j \neq 0\}$. Thus the conditional trait variance will vary depending on the strength of environmental exposure either if there are a few SNPs with GxE interactions of large effect or if there are many SNPs with small yet non-zero interaction effects, and in either case homoskedasticity is unlikely to be an appropriate assumption.

Robust standard errors, alternatively called Huber-White, sandwich or “heteroskedastic-consistent” errors^{59,60}, are standard tools used in economics⁴⁷ to overcome this issue and have previously been proposed for use in GxE interaction studies^{25,49,61}. We further include a small adjustment that reduces bias in small samples⁶². This yields the variance estimator

$$\text{Var}(\hat{\tau}) = (H^T H)^{-1} H^T \hat{\Sigma} H (H^T H)^{-1},$$

where $\hat{\Sigma}$ is a diagonal matrix with $\hat{\Sigma}_{ii} = \frac{\hat{\epsilon}_i^2}{(1-h_{ii})^2}$, where $\hat{\epsilon} = y - H\hat{\tau}$ and $h = H(H^T H)^{-1} H^T$.

Supplementary Figures

Figure S1: False positive rates on simulated datasets. False positive rate (FPR) for SNP main effects tests (bottom) and SNP GxE interaction tests (top) at null SNPs in the second half of each chromosome, whilst varying (a) the number of environmental variables, (b) proportion of trait variance explained by background GxE effects and (c) sample size. The grey line denotes expected FPR. Simulations used genotypes sub-sampled from the UK Biobank and by default contained $N = 25K$ samples, $M = 100K$ SNPs, 6 environmental variables that contributed to the ES and 24 that did not (default parameters denoted by stars). We performed 20 repeats for each scenario. See **Online methods** for full details of phenotype construction.

Figure S2: **LEMMA false positive rate in large simulations.** False positive rate (FPR) for SNP main effects tests (right) and SNP GxE interaction tests (left) at null SNPs in the second half of each chromosome, whilst varying the number of environmental variables. The simulation was conducted with $N = 200K$ samples and $M = 400K$ SNPs. The simulated trait was constructed with 10,000 causal SNPs main effects that explained 20% of variance, and zero causal SNP GxE effects. We performed 20 repeats in each scenario. See **Online methods** for full details of phenotype construction.

Figure S3: Estimation of GxE heritability. Estimates of SNP-GxE heritability whilst varying the number of environmental variables (b, d) and sample size (a, c). The red dotted line denotes the true SNP-GxE heritability used whilst constructing the simulation. We observed some upwards bias as the number of environmental variables increases (b, d), which is ameliorated with increased sample size (d). Phenotypes were constructed using $M = 100,000$ SNPs with $M_{\text{causal, main-effects}} = 80,000$ causal main effects and $M_{\text{causal, GxE-effects}} = 40,000$ causal interaction effects. See **Online methods** for full details of phenotype construction.

Figure S4: Power to detect causal SNPs in simulation. Power to detect SNP GxE interaction effects (top) and SNP main effects (bottom), whilst varying (a) the number of environmental variables, (b) proportion of trait variance explained by background GxE effects and (c) sample size. Power was assessed as the proportion of 60 causal SNPs detected at $p < 0.01$ (Family Wise Error Rate; FWER < 0.01), where causal SNPs main and GxE interaction effects each explained 0.00016% of trait variance. Simulations used genotypes sub-sampled from the UK Biobank and by default contained $N = 25K$ samples, $M = 100K$ SNPs, 6 environmental variables that contributed to the ES and 24 that did not (default parameters denoted by stars). We performed 20 repeats for each scenario. See **Online methods** for full details of phenotype construction.

Figure S5: Estimation of ES weights in simulation. Boxplots of the environmental score (ES) weights estimated by LEMMA (left) over 20 simulations. Red lines denote true weights used to construct the simulated ES. Simulations performed with $N = 25k$ samples, $M = 100k$ SNPs and $L = 30$ environments (of which 6 were active). Phenotypes were constructed with $M_{\text{causal, main-effects}} = 5000$ SNPs explaining 20% of trait variance and $M_{\text{causal, GxE-effects}} = 2500$ SNPs explaining 5% of trait variance. LEMMA is invariant to a sign change in both the interaction weights and interaction SNP effects, so ES weights are automatically re-scaled such that the largest weight is positive before plotting.

Figure S6: Heritability estimates stratified by LD and MAF in simulation. Comparison of heritability estimates using RHE-SC and RHE-LDMS when causal SNPs were drawn (a) at random or (c) only from low frequency (MAF < 0.1) SNPs. Heritability estimates (using RHE-LDMS) stratified by MAF when causal SNPs were drawn (b) at random or (d) only from low frequency (MAF < 0.1) SNPs. Simulations performed with $N = 25K$ samples, $M = 100K$ SNPs and the default simulation parameters described in **Online Methods**. Abbreviations; MAF, minor allele frequency; RHE-SC, randomized HE-regression with a single SNP component²³; RHE-LDMS, multi-component randomized HE-regression²⁴.

Figure S7: GxE analysis of PP in the UK Biobank. (a) LEMMA association statistics testing for multiplicative GxE interactions at each SNP. The horizontal grey line denotes ($p = 5 \times 10^{-8}$), p -values are shown on the $-\log_{10}$ scale. (b) Distribution of the environmental score (ES), stratified by gender and age quantile. (c) Weights used to construct the ES. Dietary variables have a single weight shown on the per standard deviation (s.d) scale. ‘Gender’ has two weights; a gender specific intercept for women (first) and men (second). Remaining non-dietary variables have three weights; (first) a per s.d effect for women only, (second) a per s.d effect for men only, (third) a per s.d per decade effect which is the same for both genders. s.d for the male and female specific weights is computed for each gender separately. Age is computed as the number of decades aged from 40. See **Online Methods** for details.

Figure S8: **GxE analysis of SBP in the UK Biobank.** (a) LEMMA association statistics testing for multiplicative GxE interactions at each SNP. The horizontal grey line denotes ($p = 5 \times 10^{-8}$), p -values are shown on the $-\log_{10}$ scale. (b) Distribution of the environmental score (ES), stratified by gender and age quantile. (c) Weights used to construct the ES. Dietary variables have a single weight shown on the per standard deviation (s.d) scale. ‘Gender’ has two weights; a gender specific intercept for women (first) and men (second). Remaining non-dietary variables have three weights; (first) a per s.d effect for women only, (second) a per s.d effect for men only, (third) a per s.d per decade effect which is the same for both genders. s.d for the male and female specific weights is computed for each gender separately. Age is computed as the number of decades aged from 40. See **Online Methods** for details.

Figure S9: **GxE analysis of DBP in the UK Biobank.** (a) LEMMA association statistics testing for multiplicative GxE interactions at each SNP. The horizontal grey line denotes ($p = 5 \times 10^{-8}$), p -values are shown on the $-\log_{10}$ scale. (b) Distribution of the environmental score (ES), stratified by gender and age quantile. (c) Weights used to construct the ES. Dietary variables have a single weight shown on the per standard deviation (s.d) scale. 'Gender' has two weights; a gender specific intercept for women (first) and men (second). Remaining non-dietary variables have three weights; (first) a per s.d effect for women only, (second) a per s.d effect for men only, (third) a per s.d per decade effect which is the same for both genders. s.d for the male and female specific weights is computed for each gender separately. Age is computed as the number of decades aged from 40. See **Online Methods** for details.

Figure S10: Effect of using robust standard errors for GxE interaction tests in the UK Biobank. QQ plots of the observed LEMMA $-\log_{10}(p)$ values for GxE interactions at imputed SNPs for four UK Biobank traits, with and without robust standard errors. The grey dotted line denotes expected $-\log_{10}(p)$ -values under a null model. Association tests using ‘Robust’ standard errors are well calibrated in both homoskedastic and heteroskedastic regimes (see **Online Methods**) and are used in all follow up analysis. Genomic control statistics were 1.275, 1.271, 1.163, 1.111 for logBMI, PP, SBP and DBP respectively using homoskedastic standard errors and 1.062, 1.047, 1.037, 1.027 for logBMI, PP, SBP and DBP respectively using robust standard errors.

Figure S11: **Estimated GxE effect rs2153960 on logBMI.** (a) Regional plot of the main and interaction effects of SNPs within 250KB of rs2153960, (b) the estimated effect of rs2153960 on logBMI as a function of the environmental score (ES).

Figure S12: **Estimated GxE effect rs539515 on logBMI.** (a) Regional plot of the main and interaction effects of SNPs within 250KB of rs539515, (b) the estimated effect of rs539515 on logBMI as a function of the environmental score (ES).

Figure S13: **Estimated GxE effect rs8090962 on DBP.** (a) Regional plot of the main and interaction effects of SNPs within 250KB of rs8090962, (b) the estimated effect of rs8090962 on DBP as a function of the environmental score (ES).

Figure S14: Comparison of the LEMMA vs marginal environmental score. Interaction weights of the marginal environmental score were estimated from multivariate linear regression, using all the non-genetic covariates used by LEMMA. Interactions weights were all rescaled so that the corresponding ES had variance one. The dashed grey line represents the $y = x$ line.

(a) LEMMA

(b) StructLMM

(c) F-test

(d) robust F-test

Figure S15: **GxE association statistics for logBMI**. Manhattan plots displaying the negative $\log_{10}$ p values from GxE interaction tests at 10,295,038 imputed SNPs applied to logBMI in the UK Biobank. GxE interaction tests were computed using (a) LEMMA, (b) StructLMM, (c) the F-test and (d) the robust F-test. The horizontal grey line denotes ($p = 5 \times 10^{-8}$).

Figure S16: Comparison of GxE association statistics for logBMI. Comparison of negative $\log_{10} p$ values obtained from LEMMA, StructLMM, the F-test and the robust F-test in an analysis of logBMI in the UK Biobank. Grey lines denote ($p = 5 \times 10^{-8}$) and the $y = x$ axis. Pearson correlation is shown in a label at the top left of each plot. Red points denote the sentinel SNP for each locus.

(a) log(BMI)

(b) DBP

(c) SBP

(d) PP

Figure S17: **Inference of environmental score weights from GxE analyses of four quantitative traits in the UK Biobank.** Evolution of the environmental score weights as LEMMA performs successive passes through the data.

Figure S18: Log-log plots showing runtime of the variational bayes algorithm used to perform whole genome regression by LEMMA, as a function of sample size (left) and the number of environmental variables (right). Unless otherwise stated simulations were performed using $N = 25k$ samples, $M = 100k$ SNPs and $L = 30$ environmental variables. Phenotypes were constructed using 2500 non-zero main effects explaining 20% of variance, 1250 nonzero interaction effects explaining 5% of variance and 6 active environmental variables. See **Online methods** for full details of phenotype construction.

Figure S19: **Visualization of differences in variance induced by a multiplicative Gene-x-Environment effect.** Differences in phenotypic variance by genotype group (top) and by strength of the environmental exposure (bottom). The phenotype was simulated using 2000 individuals on the basis of a multiplication interaction between a single genotype (minor allele frequency 0.3) and an environment (uniformly distributed over $[-2, 2]$). From left to right; the GxE interaction explained 0%, 20%, 40% of trait variance.

Trait	No. samples	No. SNPs	No. envs	No. envs-sq	No. other covars	No. covars total	Iterations for WGR converge	Time for WGR to converge*
log(BMI)	281149	642095	42	30	24	96	1631	78 hours 18 mins
PP	280749	642102	45	13	25	83	3821	183 hours 26 mins
SBP	280749	642102	45	15	25	85	967	46 hours 25 mins
DBP	280749	642102	45	15	25	85	646	31 hours 00 mins

∞ Table S1: **Quality control and time to convergence of the WGR analyses** Time for the whole genome regression analysis to converge is reported for four quantitative traits in the UK Biobank, as well as the number of SNPs and samples passing quality control and the number of covariates controlled for. ‘Other covariates’ consisted of the top 20 genetic principal components as reported by the UK Biobank, age³, age² × gender, age³ × gender, a binary indicator for the genotype chip and (for blood pressure traits only) BMI. Environmental variables used (including lower orders of age and gender) are described in (**Online Methods**). To control for potential bias due to non-linear dependence between the phenotype and heritable environmental variables, we tested each environmental variable and included any significant squared effects as additional covariates (**Online Methods**) *based on the average per-iteration cost of 243 seconds, using 32 cores distributed across a cluster with Xeon E5-2667 v4 3.2Ghz processors.

Trait	Genotyped (SC)		Genotyped (LDMS)		Common Imputed (LDMS)	
	h_G^2 (s.e)	h_{GxE}^2 (s.e)	h_G^2 (s.e)	h_{GxE}^2 (s.e)	h_G^2 (s.e)	h_{GxE}^2 (s.e)
log BMI	0.259 (0.069)	0.071 (0.009)	0.237 (0.126)	0.086 (0.024)	0.274 (0.056)	0.093 (0.028)
PP	0.233 (0.039)	0.075 (0.018)	0.203 (0.084)	0.111 (0.021)	0.228 (0.051)	0.125 (0.028)
SBP	0.24 (0.053)	0.033 (0.003)	0.223 (0.095)	0.038 (0.017)	0.251 (0.05)	0.039 (0.023)
DBP	0.277 (0.034)	0.014 (0.001)	0.231 (0.079)	0.016 (0.017)	0.254 (0.05)	0.016 (0.02)

Table S2: **Partitioned heritability estimates for four quantitative traits in the UK Biobank.**

Comparison of the heritability estimates obtained using genotyped SNPs with RHE-SC, genotyped SNPs with RHE-LDMS, and common imputed SNPs (MAF > 0.01 in the full UK Biobank cohort) with RHE-LDMS. GxE heritability estimates were obtained using the ES from each model fit. All analyses controlled for the same covariates used in the WGR analysis (including the top 20 principal components). Abbreviations; s.e, standard error estimated using the block jack-knife (see **Online Methods**); h_G^2 , heritability due to additive genetic effects; h_{GxE}^2 , heritability due to multiplicative GxE effects; RHE, randomized HE-regression^{23,24}; SC, single SNP component; LDMS, SNPs stratified by minor allele frequency and LDscore (20 components).

Method	No. signals	Genomic control (χ^2)	Genomic control (p-values)*
LEMMA	2	1.062	1.038
LEMMA-S	5	1.275	1.164
StructLMM (-SQE)	3	NA	1.236
F-test (-SQE)	4	NA	1.372
robust F-test (-SQE)	2	NA	1.034
LEMMA (-SQE)	3	1.065	1.04
LEMMA-S (-SQE)	6	1.288	1.171

Table S5: Comparison of the number of genome-wide significant GxE associations and genomics control statistics from a GxE analysis of logBMI in the UK Biobank The number of independent loci (at least 0.5cM apart) with genome-wide significant GxE interaction effects and genomic control statistics for seven different methods applied to logBMI in the UK Biobank. Genomic control is computed from GxE interaction tests statistics from 10,295,038 imputed SNPs. Abbreviations; LEMMA-S, LEMMA with a homoskedastic test statistic (see **Online Methods**); (-SQE), significant squared environmental variables (Bonferroni correction) not included as additional covariates.

*The test statistics from StructLMM, F-test and the robust F-test are not χ^2_1 distributed. Hence for these methods we use $\lambda_{GC} = \log_{10}(m)/\log_{10}(0.5)$, where m is the median p -value, to denote the genomic control statistic as suggested by Moore *et al.*⁷.

	logBMI ES	PP ES	SBP ES	DBP ES
PC1	2.182	0.023	0.467	0.403
PC2	0.254	0.153	0.197	0.088
PC3	0.022	0.382	0.405	0.295
PC4	0.657	0.465	0.091	0.095
PC5	15.075	2.652	68.293	71.439
PC6	0.703	0.780	0.020	0.480
PC7	0.469	0.095	0.933	0.998
PC8	0.778	1.898	2.784	1.848
PC9	0.675	0.878	14.130	27.305
PC10	0.814	0.081	0.647	0.276
PC11	2.759	0.255	6.547	7.970
PC12	0.319	0.974	0.486	0.511
PC13	0.659	0.508	0.475	2.003
PC14	3.301	5.779	2.575	4.454
PC15	0.186	0.418	0.300	0.531
PC16	3.554	3.141	2.830	6.560
PC17	0.695	1.590	1.208	0.061
PC18	3.965	0.215	0.681	0.884
PC19	0.350	0.375	0.099	0.409
PC20	2.463	1.208	1.228	0.720

Table S7: Association between genetic principle components and the environmental score for four traits in the UK Biobank Associations computed using ordinary least squares to regress the environmental score against the top 20 principle components (with an intercept included). Association strength reported using negative $\log_{10}(P)$ -values from a standard t-test. Abbreviations; PC, genetic principle component; ES, environmental score.

Trait	Genotyped (LDMS)		Genotyped (LDMS), additionally controlling for ES-x-PCs	
	h_G^2 (s.e)	$h_{G \times E}^2$ (s.e)	h_G^2 (s.e)	$h_{G \times E}^2$ (s.e)
log BMI	0.2366 (0.1259)	0.0862 (0.0237)	0.2368 (0.1259)	0.0862 (0.0236)
PP	0.2025 (0.0841)	0.1110 (0.0208)	0.2025 (0.0840)	0.1113 (0.0208)
SBP	0.2225 (0.0954)	0.0377 (0.0168)	0.2226 (0.0955)	0.0360 (0.0166)
DBP	0.2308 (0.0793)	0.0157 (0.0166)	0.2308 (0.0793)	0.0142 (0.0163)

Table S8: **Sensitivity of partitioned heritability estimates to ES-x-PCs interaction in the UK Biobank.** Heritability estimates were computed using genotyped SNPs with RHE-LDMS and the ES from each WGR analysis. Left; heritability estimates obtained whilst controlling for the same covariates used in the WGR analysis (including the top 20 principal components), right; heritability estimates obtained whilst additionally controlling for multiplicative interactions between the ES and genetic PCs. Abbreviations; s.e, standard error estimated using the block jackknife (see **Online Methods**); h_G^2 , heritability due to additive genetic effects; $h_{G \times E}^2$, heritability due to multiplicative GxE effects; RHE, randomized HE-regression^{23,24}; LDMS, SNPs stratified by minor allele frequency and LDscore (20 components).

Table S9: Correlation between main SNP effects and interaction SNP effects

Trait	$\text{cor}(\text{ES}, \text{ES}_{\text{main}})$	$\text{abs}(\text{cor}(X\beta, X\gamma))$	$\text{abs}(\text{cor}(\beta, \gamma))$
log(BMI)	-0.062	0.13680	0.05810
PP	-0.019	0.05306	0.03324
SBP	-0.297	0.01741	0.00816
DBP	-0.088	0.02464	0.00732

Correlation between main SNP effects and interaction SNP effects learnt during the LEMMA variational algorithm. Absolute correlation is used as γ is invariant to being multiplied by -1 (as LEMMA would apply the same transform to the ES).