Table S1. Prior distributions and their hyperparameters.

Model parameter	Description	Prior distribution	Hyperparameter values
F	Number of gene expression traits	-	-
G	Number of probes per gene expression trait	-	-
k	Number of PCFs	-	-
$\boldsymbol{\Lambda}$	F x k matrix of PCF loadings	$\lambda_{jf} \sim \mathrm{N}(0, \boldsymbol{\phi_{jf}^{-1}\tau_j^{-1}})$ $\phi_{jf} \sim \mathrm{Ga}(^{\upsilon}/_{2}, ^{\upsilon}/_{2})$ (f = 1…F, j = 1…k) $\tau_j = \prod_{l=1}^{j} \delta_l$ $\delta_1 \sim \mathrm{Ga}(a_1, b_1)$ $\delta_l \sim \mathrm{Ga}(a_2, b_2)$, ($l$ = 2…k).	$\upsilon = 2$ $a_1 = 2.1$ $b_1 = 1/20$ $a_2 = 3$ $b_2 = 2$
$\boldsymbol{\Sigma}_{h^2}$	The j^{th} diagonal element of the k x k diagonal matrix, $\boldsymbol{\Sigma}_{h^2}$, contains $\boldsymbol{h_j^2}$, the heritability of the j^{th} latent PCF.	$\mathrm{prob}(\boldsymbol{h_j^2} = 0) = 0.5$ $\mathrm{prob}(\boldsymbol{h_j^2} = \frac{f}{n_h}) = \frac{1}{2(n_h - 1)}$, ($j$ = 1…k, f = 1.. (n_h - 1)).	$n_h = 100$
$\boldsymbol{\Psi}_{\mathrm{m}}$	An F x F diagonal matrix containing the specific mutational variances.	${\psi_{\mathrm{m}_f}}^{-1} \sim \Gamma(a_{\mathrm{m}}, b_{\mathrm{m}})$, ($f$ = 1…F).	$a_{\mathrm{m}} = 2$ $b_{\mathrm{m}} = 1/2$
$\boldsymbol{\Psi}_{\mathrm{r}}$	An F x F diagonal matrix containing the specific residual variances.	${\psi_{\mathrm{r}_f}}^{-1} \sim \Gamma(a_{\mathrm{r}}, b_{\mathrm{r}})$, ($f$ = 1…F).	$a_r = 2$ $b_r = 1/2$
$\boldsymbol{\Sigma}_{\mathrm{y}}$	An (F x k) x (F x k) diagonal matrix containing the residual probe variances.	${\sigma_{\mathrm{y}_{fg}}}^{-1} \sim \Gamma(a_{\mathrm{y}}, b_{\mathrm{y}})$, ($f$ = 1…F, g = 1…G).	$a_{\mathrm{y}} = 2$ $b_{\mathrm{y}} = 1/2$
B	The 2 x F matrix of fixed effects: overall mean for each trait and the effect of the segregating major gene on each trait.	$\beta_{.f} \sim N_2(0, \infty * \mathbf{I_2})$ (f = 1…F).	Note: improper prior, but posterior is proper.

Table S2. Effect of prior distribution hyperparameters on the number of PCFs identified and the estimated mutational variance by the Bayesian Sparse Factor model. We ran the model on the observed data with 5 different random seeds for each of 9 different sets of prior distribution hyperparameters. From each of the 45 models, we retained 1000 samples from 100,000 (i.e. thinned at a rate of 100) after a burn-in period of 300,000 samples, and calculated mutational variances and significance of PCFs and their heritabilities using the average error rate of the local false sign rate method, as described in the main text. We report the range of estimated parameters across the 5 runs of each set of hyperparameter values. We found the total mutational variance estimated on a trait by trait basis was highly consistent among the 45 analyses, with pair-wise correlations between sets of estimated trait mutational variances ranging from r = 0.95 - 0.99. We continued the chain from set 7 with the highest number of heritable PCFs for a further 200,000 samples, and retained 1000 of the last 100,000 samples thinned at a rate of 100 for the analysis presented in the main text. Additional PCFs were detected in the continued run, taking the observed number of heritable PCFs to 21.

Set	a_2	b_2	υ	Number of significant PCFs		Number of heritable PCFs		Total mutational variance		Common mutational variance	
				Min	Max	Min	Max	Min	Max	Min	Max
1	2	1	2	37	45	14	16	777.0	799.7	426.7	464.7
2	2	1	3	35	41	12	16	775.3	785.5	405.2	419.4
3	2	1	9	23	27	6	8	761.9	774.3	360.6	371.9
4	3	1	2	31	37	11	13	789.1	793.8	407.5	437.6
5	3	1	3	26	35	7	12	779.0	789.4	384.0	421.8
6	3	1	9	22	24	5	7	765.4	782.1	353.6	370.0
7	3	2	2	37	41	12	18	791.4	799.1	420.1	448.9
8	3	2	3	34	37	10	14	771.2	784.1	393.2	424.7
9	3	2	9	24	27	6	9	763.4	771.9	354.9	365.4

Table S3. **The number of individual gene expression traits contributing to the heritable PCFs.** For each heritable PCF, we report the number of significant trait loadings based on the local false sign rate (LFSR) and the average error rate s-value, as described in the Methods. We show the total number and percentage of trait loadings that meet the significance thresholds for the two related approaches, and break the total number down into the number of traits that load significantly onto one, two, three or four PCFs.

PCF	LFSR< .025						s-value < .005					
	1	2	3	4	Total	%	1	2	3	4	Total	%
9	179	63	4	1	247	7.3%	183	63	6	1	253	7.5%
10	179	70	6	1	256	7.6%	183	74	8	1	266	7.9%
13	94	61	3	0	158	4.7%	85	59	6	0	150	4.4%
15	69	10	0	0	79	2.3%	73	13	0	0	86	2.5%
16	49	11	0	0	60	1.8%	54	12	0	0	66	1.9%
19	49	27	4	1	81	2.4%	46	26	4	1	77	2.3%
23	61	10	0	0	71	2.1%	63	10	1	0	74	2.2%
24	31	26	2	0	59	1.7%	31	27	3	0	61	1.8%
25	40	2	0	0	42	1.2%	43	3	0	0	46	1.4%
27	27	3	1	0	31	0.9%	29	3	1	0	33	1.0%
28	24	6	1	0	31	0.9%	24	7	1	0	32	0.9%
30	32	4	1	0	37	1.1%	33	4	1	0	38	1.1%
31	43	4	0	0	47	1.4%	46	4	0	0	50	1.5%
32	20	10	1	0	31	0.9%	20	11	1	0	32	0.9%
33	20	9	0	0	29	0.9%	22	9	0	0	31	0.9%
36	17	3	0	0	20	0.6%	18	3	0	0	21	0.6%
37	14	0	0	0	14	0.4%	15	0	0	0	15	0.4%
38	15	1	0	0	16	0.5%	15	1	0	0	16	0.5%
40	16	0	0	1	17	0.5%	16	1	0	1	18	0.5%
41	11	4	1	0	16	0.5%	12	4	1	0	17	0.5%
44	3	0	0	0	3	0.1%	3	0	0	0	3	0.1%
By trait	993	162	8	1	1164	34.4%	1014	167	11	1	1193	35.2%

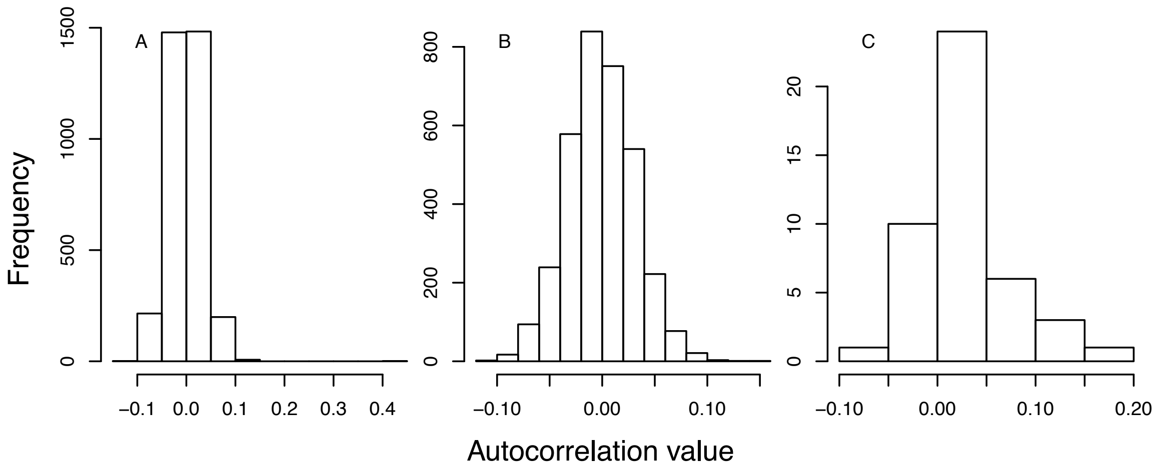

Figure S1. Autocorrelation statistics. Distributions of autocorrelation values are shown for (A) mutational specific variances, (B) residual specific variances and (C) PCF heritabilities. A single specific mutational variance exceeded our nominal threshold of 0.2.

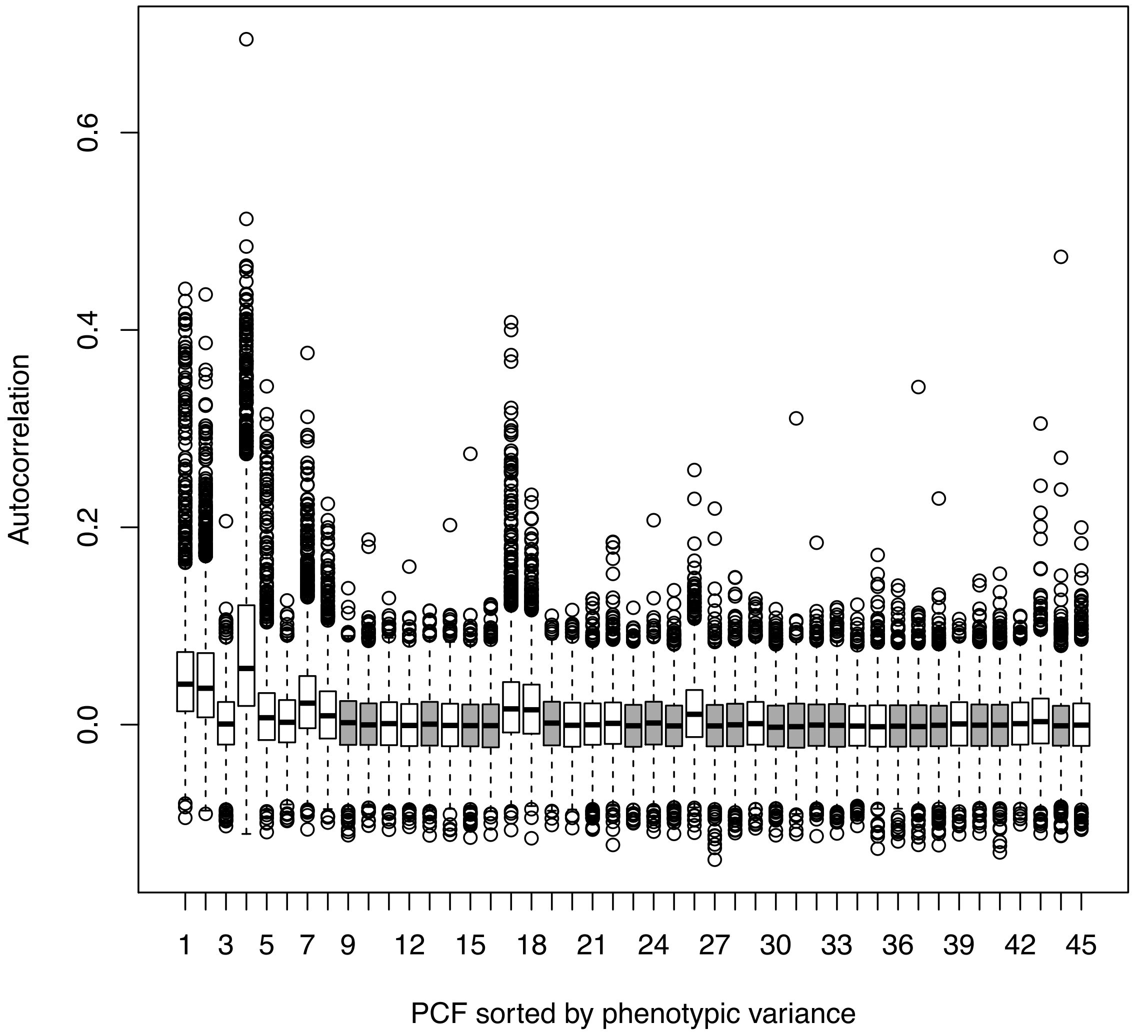

Figure S2. Boxplot of PCF trait loading autocorrelations. Grey-filled (unfilled) boxes indicate autocorrelations for loadings on heritable (non-heritable) PCFs. Trace plots of trait loadings with the highest autocorrelations within PCFs are presented in Fig. S4.

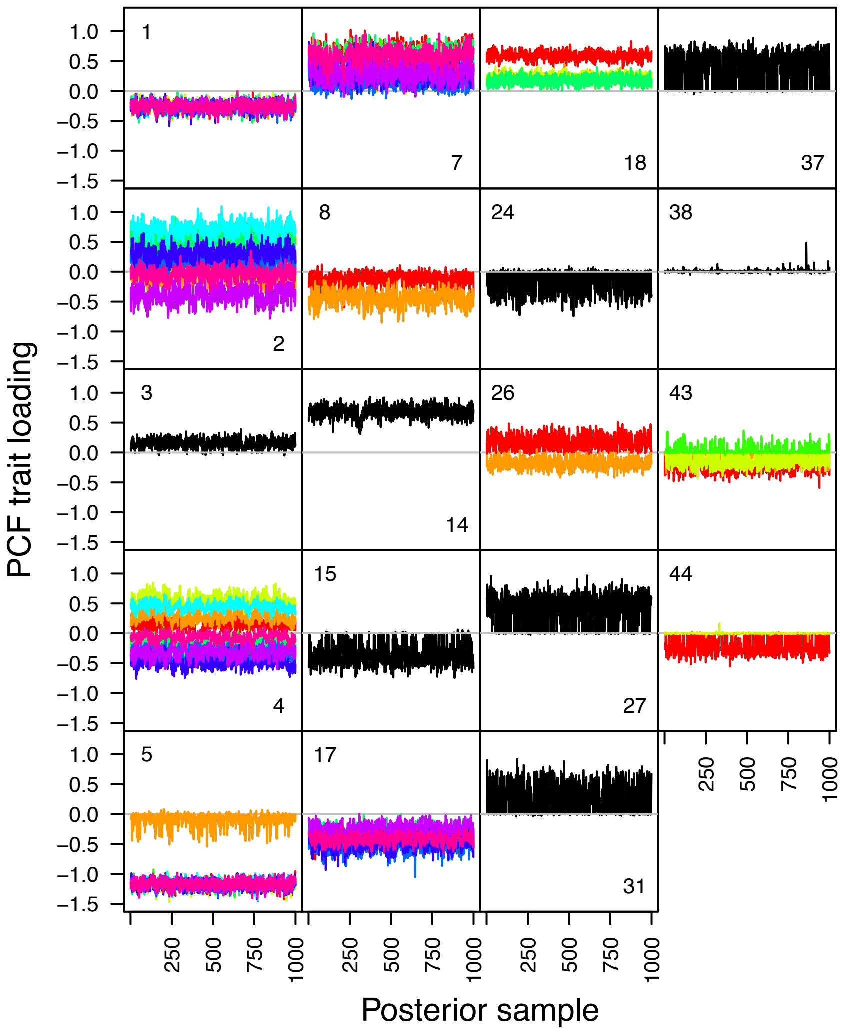

Figure S3. Trace plots of highly autocorrelated PCF trait loadings. Presented are up to 10 trait loadings per PCF, for the subset of PCFs with at least one trait loading autocorrelation exceeding our nominal threshold of 0.2. Numbers in corners of panels indicate the PCF number (in order of phenotypic variance).

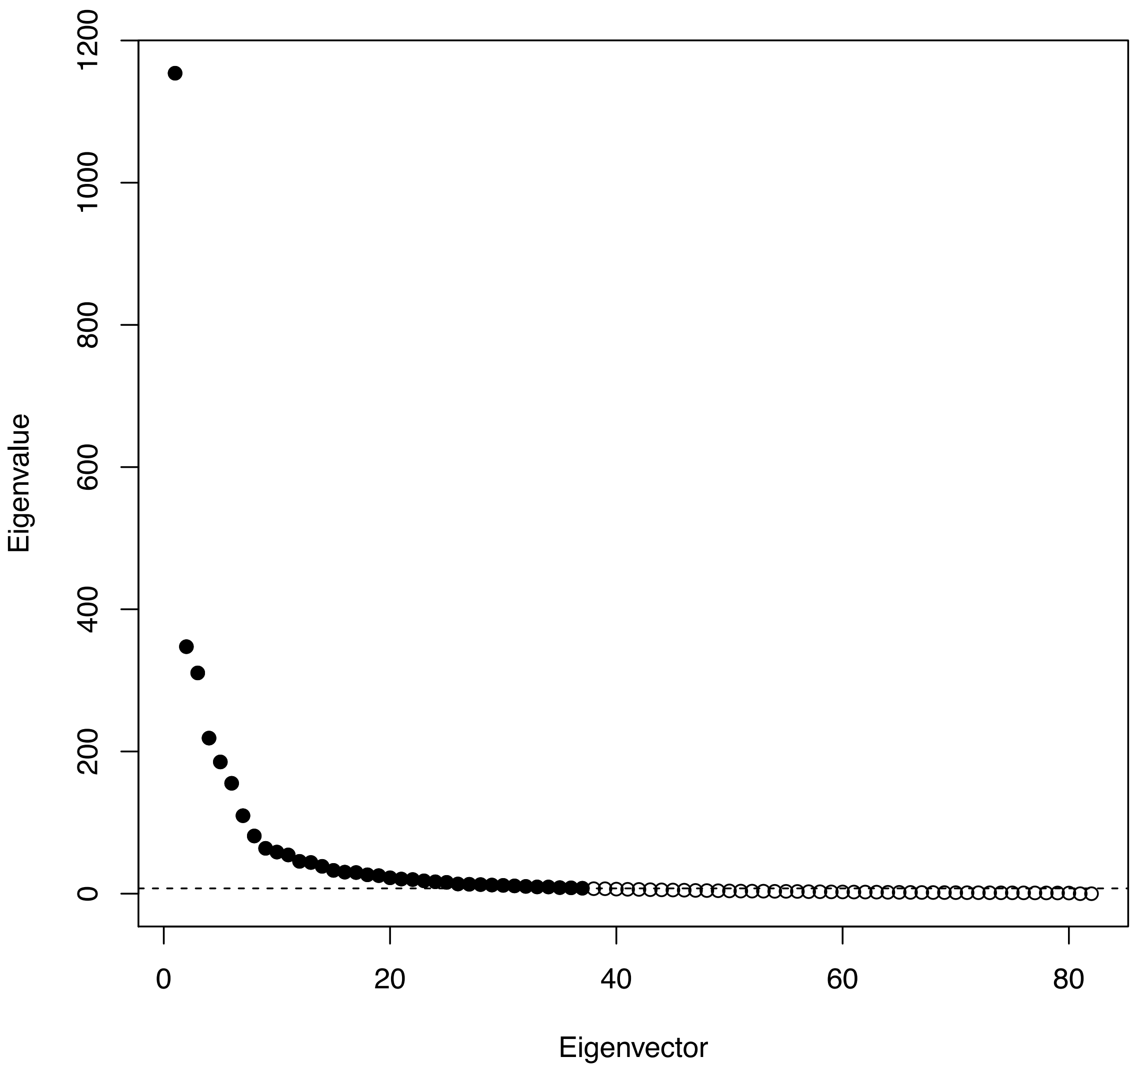

Figure S4. Theoretical phenotypic dimensionality detection limit. The columns of the 82 x 16925 data matrix were standardised, then the data matrix was reduced to an 82 x 3385 matrix of average gene expression measurements across the five probes per gene expression trait. The data was then corrected for the segregating genetic variant and the correlation matrix of the data thus transformed was calculated. The Marcenko-Pastur law describes a limit for the number of dimensions of a sample correlation matrix that have a high probability of detection. Given the number of traits, *p*, and the sample size, *m*, the number of identifiable dimensions is not likely to be less than *k*, where λ_k is the smallest eigenvalue of the sample correlation matrix greater than $\lambda_{lim} = 1 + \sqrt{p/m}$ (Nadakuditi and Edelman 2008). In our case 37 of the 81 non-zero eigenvalues (filled circles) of the sample correlation matrix were above the λ_{lim} value of 7.4 (dashed line) for *m*=82 and *p* = 3385. Consistent with this minimum, the BSF model identified 45 PCFs.

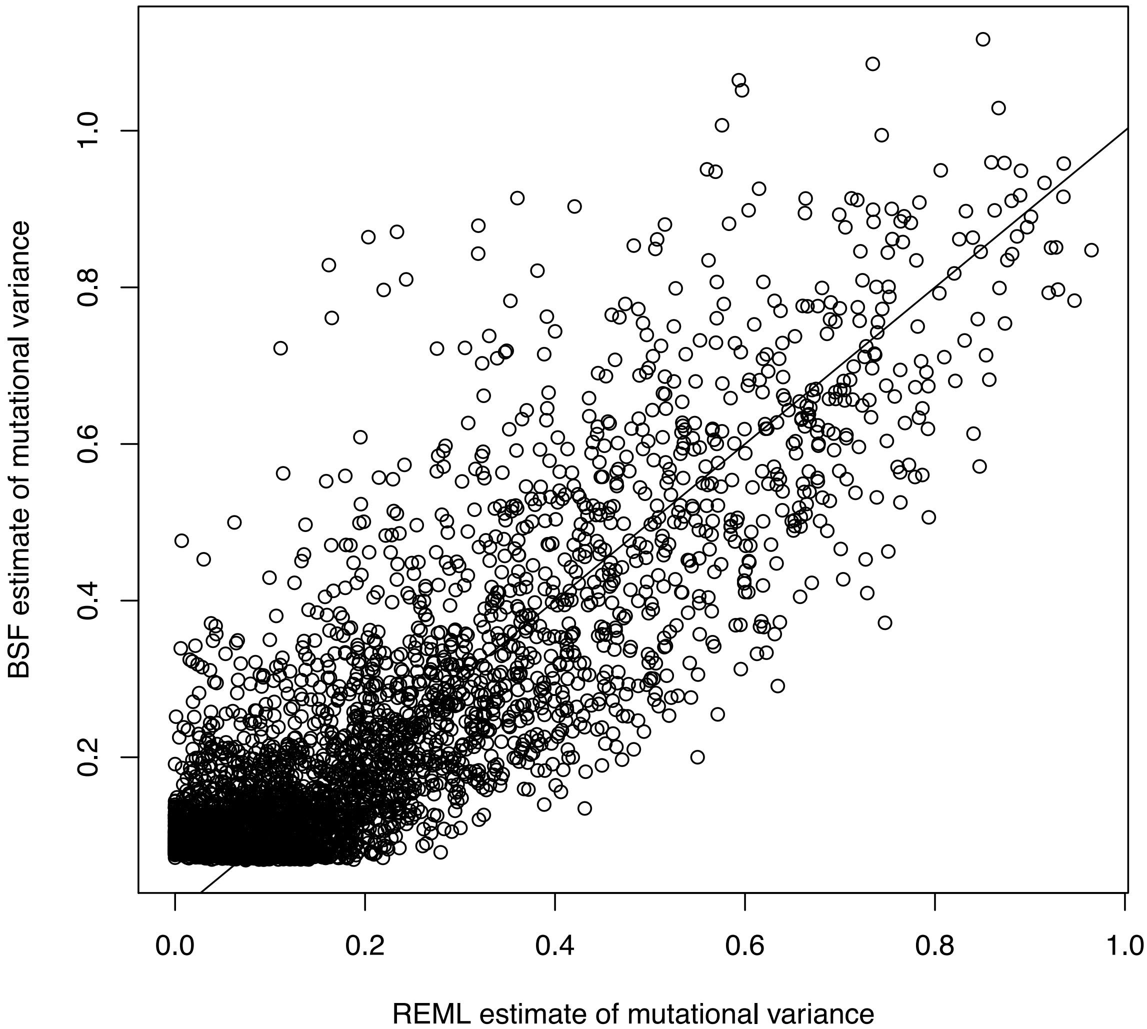

Figure S5. Comparison of the mutational variance estimated under restricted maximum likelihood (REML) and Bayesian Sparse Factor (BSF) frameworks. The BSF estimates of total (common + specific) mutational variance of the gene expression traits (y-axis) were highly correlated (r=0.84, df=3383, p<<0.001) with the among-line variance components estimated in standard univariate REML analyses (x-axis). The univariate analyses were first reported in McGuigan et al. (2014) and have been recalculated here to be on the scale enforced by the BSF, whereby each of the 16925 (3385 genes x 5 probes) sets of probe measurements are standardized.

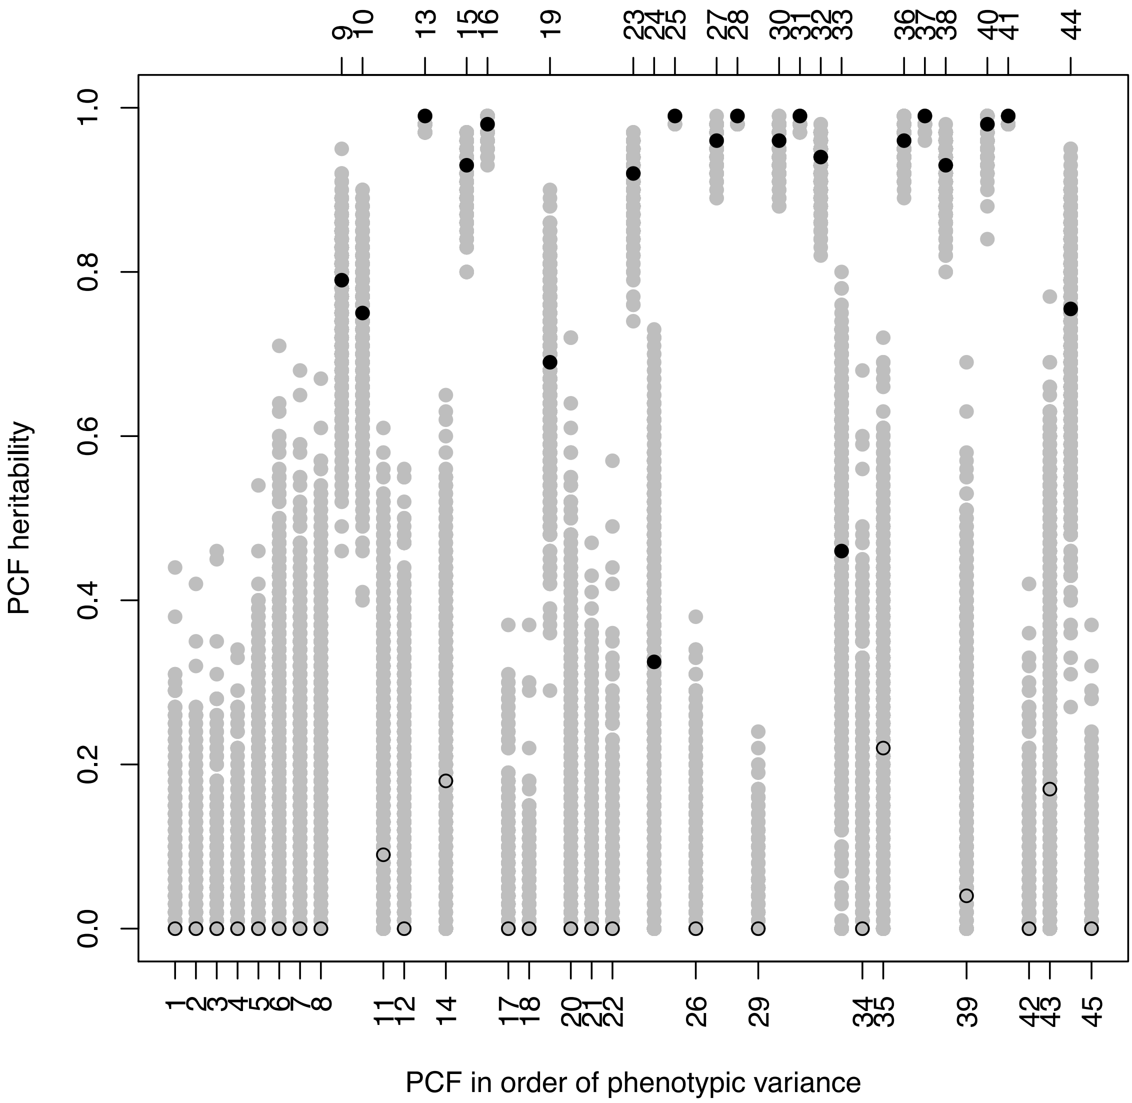

Figure S6. Posterior distribution of PCF heritability. For each PCF heritability, we show all individual posterior samples with grey circles and the median posterior sample with black circles (closed and open circles for significant and nonsignificant PCF heritabilities, respectively). Significantly heritable PCFs are also labelled on the top axis for ease of identification. We assign significance to PCF heritabilities using the local false sign rate approach described in the Methods, controlling the average error rate to remain below 0.01.