Biometrics –, 1–18 DOI: 000
– 2016

Web-based Supplementary Materials for Estimation of eQTL Effect Sizes Using a Log of Linear Model by J. Palowitch, A. Shabalin, Y.H. Zhou, and A.B. Nobel, F.A. Wright

This paper has been submitted for consideration for publication in *Biometrics*

A. QN-linear fit vs. ACME fit

An illustration of the information loss incurred by the QN-linear model approach to eQTL analysis can be made by considering two gene-SNP pairs, each with eQTL evidence that is similar under QN transformation, but disparate on the log scale. We chose one such pair from real data, displayed in main Figure 1. While the estimated ACME effect size of pair 2 is ten times greater than that of pair 1, the effect sizes from linear regression with QN-transformed expression are nearly the same. Furthermore, the baseline expression of pair 1 is far greater than that of pair 2, a feature that is obscured by the QN transformation.

[Figure 1 about here.]

B. Pre-processing gene read counts

Let x_{ij} denote elements of the original count matrix, where $i = 1, \ldots, n$ indexes samples and $j = 1, \ldots, T$ indexes genes. Let $l_i = \sum_j x_{ij}$ be the library size for sample i. The overall entry-wise mean is $\bar{x} = \frac{\sum_i \sum_j x_{ij}}{Tn}$. The final normalized count matrix elements are $c_{ij} = \frac{x_{ij}}{l_i} T\bar{x}$. This process results in a standardized matrix with constant column sums. Similar normalization is performed by software such as DESeq2 (Love et al., 2014), but with additional attention to nonlinear scaling relationships.

C. The pathological nature of residuals from raw expression

In this brief section we present a small simulation to show that the non-Normality observed in residuals from simple linear regression applied to un-transformed, normalized gene expression can be detrimental to the Type-I error rate in large-scale eQTL analyses. The simulation involves taking complete estimated residual vectors from real-data analyses, adding them to a simulation model of eQTL action with no signal, and re-applying simple linear regression to obtain a p-value. We did this using the test data from the analysis described in Section

2.3 of the main text. For each of one million simulation repetition, we (1) chose a real SNP vector from 40,000 sample vectors associated with Thyroid tissue, (2) chose an estimated residual vector resulting from one of the 40,000 regressions using un-transformed expression, (3) added the residual vector to an arbitrary β_0 coefficient to obtain simulated expression data, and (4) performed simple linear regression on the simulated expression data vs. the (randomly chosen) SNP vector, obtaining an F-test p-value. As shown in Figure 2, the p-value distribution is highly non-uniform. This supports the idea that indeed, a minimum condition to proceed with a reasonable eQTL analysis is to transform raw gene expression in a way that reduces non-Normality.

[Figure 2 about here.]

D. Tests of normality and homoskedasticity

This section contains the results from tests for normality and homoskedasticity of residuals, for each cis-QTL group, and for all single-pair eQTL models considered in this paper. For each of the 10,000 gene-SNP pairs in each eQTL group (see in Section G above), we calculated the p-value for the canonical Shapiro-Wilk test for normality, and the p-value for the canonical Bartlett test for homoskedasticity. Figure 3 displays boxplots of these p-values on the $-\log_{10}$ scale. We see that residuals for the linear model with raw gene expression ("RAW") are less normal and less homoskedastic than the residuals for the log-based models (ACME, log-linear ("LL"), and log-ANCOVA ("ANCOVA")). This is particularly true for weak eQTLs, for which error assumptions are most important for Type I error control. The dark-red lines in the figure represent the typical FDR cut-off applied to each bin of the data. Results from other GTEx tissues are given in Figures 11-12, and follow the same pattern.

[Figure 3 about here.]

E. ACME fitting algorithm

Here we describe the iterative algorithm used to identify parameters β_0 and β_1 that approximately maximize the likelihood of the ACME model. For a particular gene-SNP pair, maximizing the likelihood is equivalent to minimizing the sum of squares

$$\sum_i (y_i - \log(\beta_0 + \beta_1 s_i) - \langle \boldsymbol{Z}_i, \gamma \rangle)^2 \tag{1}$$

over the parameters β_0, β_1 and γ. The iterative algorithm operates by carrying out least-squares regression on an approximating linear model. Denoting the natural effect size β_1/β_0 by η, under the ACME model the conditional mean of y_i given $\boldsymbol{Z}_i$ and s_i is given by

$$\mathbb{E}[y_i | \boldsymbol{Z}_i, s_i] = \log(\beta_0) + \log(1 + s_i \eta) + \boldsymbol{Z}_i^T \gamma \tag{2}$$

A first-order Taylor approximation of $\log(1 + s_i\eta)$ around η at an estimate $\hat{\eta}^j$ gives

$$\mathbb{E}[y_i | \boldsymbol{Z}_i, s_i] \approx \log(\beta_0) + \log(1 + s_i\, \hat{\eta}^j) + \frac{s_i}{1 + s\, \hat{\eta}^j}(\eta - \hat{\eta}^j) + \boldsymbol{Z}_i^T \gamma \tag{3}$$

We use this approximation to motivate a linear model, in which the response variable is $d_i := y_i - \log(1 + s_i\, \hat{\eta}^j)$. Define $\theta_0 := \log(\beta_0)$ and $\theta_1 := \eta - \hat{\eta}^j$. Then subtracting $\log(1 + s_i\, \hat{\eta}^j)$ from each side of Equation 3 yields the following linear model in the parameters θ_0, θ_1, and γ:

$$d_i = \theta_0 + \frac{s_i}{1 + s_i\, \hat{\eta}^j}\theta_1 + \boldsymbol{Z}_i^T \gamma + \epsilon_i \tag{4}$$

After fitting this model at the j^{th} iteration, we set $\hat{\eta}^{j+1}$ to $\hat{\theta}_1 + \hat{\eta}^j$, and repeat the procedure. This is repeated until $|\hat{\eta}^j - \hat{\eta}^{j+1}|$ is close to machine precision. As the likelihood (1) is convex, we can expect convergence, since the algorithm is similar to a Gauss-Newton procedure. The last estimates of θ_0 and θ_1 are then used to obtain estimates of β_0 and β_1 via the equations $\hat{\beta}_0 = \exp\{\hat{\theta}_0\}$ and $\hat{\beta}_1 = \hat{\beta}_0 \hat{\eta}$.

We set the initial estimate of η to 0. If for any j, $1 + s_i \hat{\eta}^j$ is negative for any index i, we divide $\hat{\eta}^j$ by 2 and restart the j-th iteration.

F. Derivation of effect size standard error

Defining $\theta_0 := \log(\beta_0)$ and $\eta := \beta_1/\beta_0$, the ACME model may be expressed

$$y_i = \theta_0 + \log(1 + \eta s_i) + \boldsymbol{Z}_i^T \gamma + \epsilon_i \tag{5}$$

for samples $i = 1, \ldots, n$. As discussed in Section 2.2, we use "effect size" to refer to η. Let $\hat{\epsilon}_i$ be the estimated residual for patient i from (5). Let C be an orthonormalization of the matrix $(\mathbf{1}_n, \boldsymbol{Z}^T)$, and define $P := (I_n - CC^T)$. Letting y, $\log(1 + \eta s)$, and $\hat{\epsilon}$ be $n \times 1$ vectors corresponding to the full set of sample data, we have

$$\hat{\epsilon} = P\left[y - \log(1 + \eta s)\right], \tag{6}$$

as the matrix P residualizes the effect of θ_0 and γ. Thus, the log-likelihood for the full model may be expressed in terms of η, P, σ^2 only:

$$\begin{aligned} -\log L(y; s, \eta, P, \sigma^2) = \frac{n}{2}\log(2\pi\sigma^2) + \\ \sigma^{-2}[y - \log(1 + \eta s)]^T P^T P[y_i - \log(1 + \eta s)] \end{aligned} \tag{7}$$

We now derive an approximate observed Fisher information for η, using (6). Note that $\frac{d}{d\eta}\hat{\epsilon} = -P\frac{s}{1+\eta s}$ and $\frac{d^2}{(d\eta)^2}\hat{\epsilon} = P\frac{s^2}{(1+\eta s)^2}$ (where all operations to vectors are component-wise). Define $d := y - \log(1 + \eta s)$. Then

$$\begin{aligned} -I(\eta) &= \frac{d^2}{(d\eta)^2}\left[\frac{1}{2}\log\left(2\pi\sigma^2\right) + \frac{1}{2\sigma^2}\hat{\epsilon}^T\hat{\epsilon}\right] \\ &= \frac{1}{2\sigma^2} 2\left[\hat{\epsilon}''^T\hat{\epsilon} + \hat{\epsilon}'^T\hat{\epsilon}'\right] \\ &= \frac{1}{\sigma^2}\left[\left(\frac{s^2}{(1+\eta s)^2}\right)^T Pd + \left(\frac{s}{1+\eta s}\right)^T P\frac{s}{1+\eta s}\right], \end{aligned}$$

since P is idempotent. The asymptotic standard error for the model can then be estimated by $\sqrt{-I(\hat{\eta})^{-1}}$ with σ^2 replaced by $\hat{\sigma}^2$. Uncertainty in the remaining parameters and their effect on η is propagated through P. To check the accuracy of the standard error, we computed a purely numerical Hessian matrix for the log-likelihood and the full Fisher observed information matrix, verifying a close numerical match to the derivation above.

G. Sampling scheme for residual and goodness-of-fit tests

We created a sub-sampled eQTL data set comprised of equally-sized groups of null, weak, medium, and strong eQTLs. To determine the groups, we used the detection p-value associated with the QN-linear model (described in Section 1.1) as an *a priori* measure of eQTL association strength within a fixed dataset. Although our intention is to provide a new effect size measure, this prior stratification provides a refined view of ACME model behavior at various levels of association evidence. The four groups of cis-eQTL pairs were defined as follows: "null" eQTLs with $-\log_{10}$ p-value in $[0, 5)$; "weak" eQTLs with $-\log_{10}$ p-value in $[5, 10)$; "medium" eQTLs with $-\log_{10}$ p-value in $[10, 15)$; and "strong" eQTLs with $-\log_{10}$ p-value in $[15, \infty)$. The sub-sampled data were obtained by sampling 10,000 pairs from each group, uniformly-at-random.

H. Framework for direct null simulation

The preliminary data set used in Section 2.3 contained data from 40,000 unique gene-SNP pairs. When calculating the residual diagnostics presented in that section, we saved the estimated residual vectors (computed by the ACME fit) from each gene-SNP pair. We used these to create 25 null-data replications of each of the 40,000 pairs from the preliminary data. For a fixed gene-SNP pair, we went through the following steps:

(1) Recalled s the real allele count vector, $\hat{\beta}_0$ the ACME-estimated value of β_0 from the preliminary data, and $\hat{\sigma}$ the ACME-estimated value of σ from the preliminary data

(2) For $r = 1, \ldots, 25$, constructed a vector of realistic errors ϵ_r by

$$\epsilon_r = \hat{\sigma}\epsilon_r^* + \mathcal{N}_r(0, \hat{\sigma}/10)$$

where ϵ_r^* is a randomly selected stored residual vector from one of the 40,000 gene-SNP pairs, scaled to have variance 1. The addition of $\mathcal{N}_r(0, \hat{\sigma}/10)$ is a "jitter" (indepenent

within r and between all 40,000 pairs) to ensure that no two chosen mock-residual vectors are equivalent, while allowing them to retain any inherent non-Normality.

(3) Constructed the j-th replication of null gene expression data

$$g_r = \exp\left\{ \log(\hat{\beta}_0) + Z\gamma_r + \epsilon_r \right\}$$

where γ_r is $\mathcal{N}_p(0_p, I_p)$ and independent within r and between all 40,000 pairs.

I. Framework for importance sampling

The basic ACME model is

$$y = \log(\beta_0 + \beta_1 x) + \epsilon \tag{1}$$

and $\epsilon \sim N(0, \sigma^2)$. The likelihood ratio and F-statistics both use maximum likelihood estimation of the parameters. For a target type I error α for a single test, we wish to estimate α_{true}, the true probability of rejection under the null. The skew normal density for ϵ provides a simple distribution family to investigate the effect of skew on p-value accuracy. For random variable Z, we have density $g(z) = 2\phi(z)\Phi(\gamma z)$, and define $\delta = \gamma/\sqrt{(1+\gamma^2)}$ and $\epsilon = (Z - \xi)/\omega$, where μ and σ^2 are chosen so that $E(\epsilon) = 0$ and $\text{var}(\epsilon) = \sigma^2$. The skewness of ϵ, determined by δ, is $\frac{4-\pi}{2}\frac{\left(\delta\sqrt{2/\pi}\right)^3}{(1-2\delta^2/\pi)^{3/2}}$.

With importance sampling, it is feasible to estimate α_{true}, even for very small α. Let $F_\eta(x, y)$ be the true joint distribution function for the vectors x and y and f its density, where $\eta = \{\beta_0, \beta_1, \sigma^2, \delta\}$ and σ^2 is the variance of ϵ and δ is its scaled skew parameter. If $\delta = 0$ then ϵ is normal, and if $\delta > 0$ then ϵ is skewed right. Let $p(x, y)$ be the p-value obtained from the F-statistic for x and y. We will use η_0 to refer to the null model $\{\beta_0, 0, \sigma^2, \delta\}$. We have

$$\alpha_{true} = P_{\eta_0}(reject) = \int_x \int_y f_{\eta_0}(x, y) I[p(x, y) < \alpha] dx dy \ . \tag{2}$$

and I is the indicator function. Simple rejection sampling estimates α_{true} by independently simulating from F_{η_0} K times. For the kth simulation, we compute $a_k = I[p(x_k, y_k) < \alpha]$, and

$\hat{\alpha}_{true} = \sum_{k=1}^{K} a_k/K$. It is easy to show that $\hat{\alpha}_{true}$ is unbiased for α_{true}. However, it is not very accurate unless K is extremely large.

To implement importance sampling for a fixed significance threshold α, we sample data from some alternative distribution $F_{\eta^\dagger}$ which has the same support as F_{η_0}. When sampling from $F_{\eta^\dagger}$, let $b_k = I[p(x_k, y_k) < \alpha]\frac{f_{\eta_0}(x_k,y_k)}{f_{\eta^\dagger}(x_k,y_k)}$. Then

$$E_{\eta^\dagger}(b_k) = \int_x \int_y f_{\eta^\dagger}(x,y) I[p(x,y) < \alpha] \frac{f_{\eta_0}(x,y)}{f_{\eta^\dagger}(x,y)} dxdy = \alpha_{true}, \tag{3}$$

so the importance sampling estimate $\hat{\alpha}^\dagger_{true} = \sum_{k=1}^{K} b_k/K$ is also unbiased. Through a good choice of η', $\hat{\alpha}^\dagger_{true}$ can have a much smaller variance than $\hat{\alpha}_{true}$ for a given number of samples K.

Before going further we note that in eQTL applications x corresponds to genotype, and the outer integral in the above equations can be replaced by a sum over possible genotypes. Also, x is assumed to be unaffected by the parameter η in the eQTL model, so $f_{\eta_0}(x,y)/f_{\eta^\dagger}(x,y) = f_{\eta_0}(y|x)/f_{\eta^\dagger}(y|x)$ and likelihood ratios are also computed conditional on x.

Choosing $\eta^\dagger$

To choose $\eta^\dagger$, we use the heuristic approach that when sampling from $F_{\eta^\dagger}$ we wish the F-statistic to reject with probability about 0.5. Determining an appropriate $\eta^\dagger$ could be done by trial and error, but a faster approach is to use an approximate correspondence between model (1) and linear regression of y on x, with normal errors. Although the difference between a linear model and the ACME model is of key importance for this paper, for the purpose of importance sampling only a crude correspondence needs to hold. We consider the (rare) scenario under the null hypothesis that our linear regression p-value is approximately α. Under such a event, the squared correlation coefficient r^2_α between x and y can be easily solved, such that the linear regression p-value equals α. Once r^2_α is obtained, we solve for $\beta^\dagger_0$ and $\beta^\dagger_1$ in the true ACME model such that the true correlation between random X and Y is exactly r^2_α. We also note that even under this rare event that observed x and y appear to

be correlated, the mean and variance of the vector y still tends to be near the true values of $\log(\beta_0)$ and σ^2. Thus, finally, we solve for $\beta_0^\dagger$ and $\beta_1^\dagger$ such that

$$E(\log(\beta_0^\dagger + \beta_1^\dagger X)) = \log(\beta_0), \qquad \frac{\text{var}(\log(\beta_0^\dagger + \beta_1^\dagger X))}{\sigma^2} = r_\alpha^2. \tag{4}$$

Equations (4) actually provide two solutions, $\{\beta_{0-}^\dagger, \beta_{1-}^\dagger\}$, $\{\beta_{0+}^\dagger, \beta_{1+}^\dagger\}$,according to the sign of r_α, because the form of (1) is not symmetric for positive and negative β_1. Thus we obtain both solutions and our final importance sampler will sample from each solution with equal probability.

To summarize our approach, using (4) we obtain $\eta_-^\dagger = \{\beta_{0-}^\dagger, \beta_{1-}^\dagger, \sigma^2, \delta\}$, $\eta_+^\dagger = \{\beta_{0+}^\dagger, \beta_{1+}^\dagger, \sigma^2, \delta\}$, and draw importance samples from $F_{\eta^\dagger} = \frac{1}{2}(F_{\eta-^\dagger} + F_{\eta+^\dagger})$.

To illustrate our approach to choosing $\eta^\dagger$, we performed 10^5 simulations for the null model with skewed errors with $\beta_0 = 100$, $n = 250$, $\sigma^2 = 1$, and skew normal errors with parameter $\delta = 0.97$ (which produces a skewness of 0.78). Supplementary Figure 7A shows the p-values as a function of $\hat{\beta}_1$. The red overlay shows the $\beta_1^\dagger$ values, using each corresponding p-value as the significance level $\alpha/2$ for each tail. The result shows that our approach to selecting $\beta_1^\dagger$ values appears to be reasonable. Supplementary Figure 7B shows the corresponding $\hat{\beta}_0$ values, as well as the $\beta_0^\dagger$ values. Efficiency of the importance sampler requires that a non-trivial fraction (we target the range of 10%-90%) of the sampled p-values be above and below the target α value. For our setup and $\alpha = 10^{-20}$, we simulated 10,000 datasets from the corresponding $F_{\eta^\dagger}$. The distribution of p-values (Supplementary Figure 7C) shows that 16% were below α, which is within our target range. Supplementary Figure 7D shows the true rejection probabilities vs the target α values. For this setup, the p-values are highly accurate to $\alpha = 10^{-12}$, but becomes somewhat conservative for smaller α.

[Figure 4 about here.]

[Figure 5 about here.]

[Figure 6 about here.]

[Figure 7 about here.]

J. Moment Corrected Correlation (MCC)

Although F-tests for the ACME model appear to perform well, a robust method is available as an accompanying analysis. Permutation of covariate-corrected data (log-expression vs. genotype) can in principle be used for tests of association. However, the number of permutations necessary to achieve sufficiently small p-values to survive multiple testing can be prohibitive. Moment Corrected Correlation (MCC) (Zhou and Wright, 2015) uses the first four exact permutation moments to provide robust p-values, and to closely mimic permutation even for exceptionally small p-values. We briefly describe the method and its application to GTEx eQTL data here. Although the approach can be very fast when applied genome-wide (Zhou and Wright, 2015), it is not optimized for cis-eQTL analysis in which each gene has a different set of cis-SNPs, and thus is comparatively slower than the ACME approach (although much faster than actual permutation). The MCC method is most powerful on linear data, but is conservative in general under the null of no association (linear or otherwise)

We show two sets of figures to portray the usefulness of MCC. For four different tissues, and on 10,000 randomly chosen cis eQTLs within each tissue, we calculated MCC p-values, direct (random) permutation p-values, and linear regression p-values. In each case, the p-value corresponds to the association between $\log(1 + \text{normalized read counts})$ and allele counts. In Figure 8, we show compare MCC p-values to those obtained from permutation. We see a very close fit, with deviation in the tails due to resolution loss for low permutation p-values. In Figure 9, we compare the log-linear regression p-values to those same permutation p-values. We see much more spread here, which verifies the success of MCC in matching the direct-permutation p-values.

[Figure 8 about here.]

[Figure 9 about here.]

[Figure 10 about here.]

[Figure 11 about here.]

[Figure 12 about here.]

[Figure 13 about here.]

[Figure 14 about here.]

[Figure 15 about here.]

[Figure 16 about here.]

[Figure 17 about here.]

[Figure 18 about here.]

[Figure 19 about here.]

[Figure 20 about here.]

References

Love, M. I., Huber, W., and Anders, S. (2014). Moderated estimation of fold change and dispersion for rna-seq data with deseq2. *Genome Biology* **15,** 1.

Zhou, Y.-H. and Wright, F. A. (2015). Hypothesis testing at the extremes: fast and robust association for high-throughput data. *Biostatistics* page kxv007.

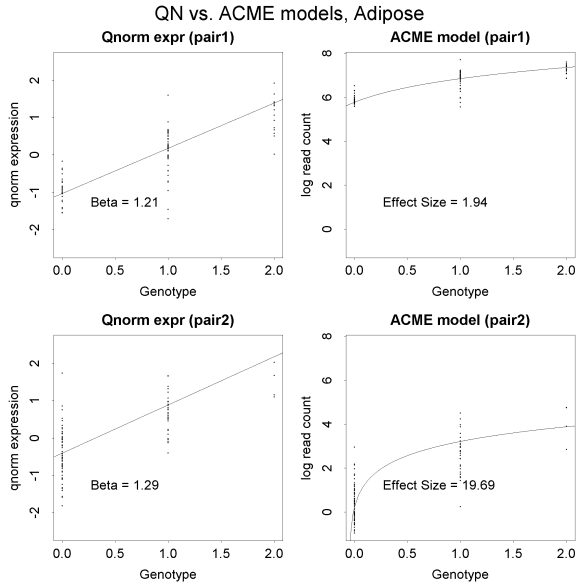

Figure 1: eQTL data from two selected gene-SNP pairs from Adipose tissue. The fitted lines correspond to the estimated parameters from each model.

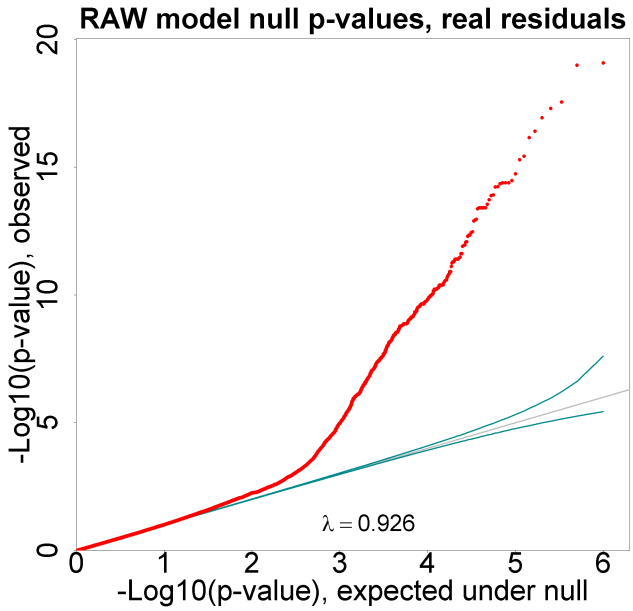

Figure 2: p-value distribution from standard linear regression applied to 1,000,000 null eQTLs with un-transformed expression and estimated residuals from real data.

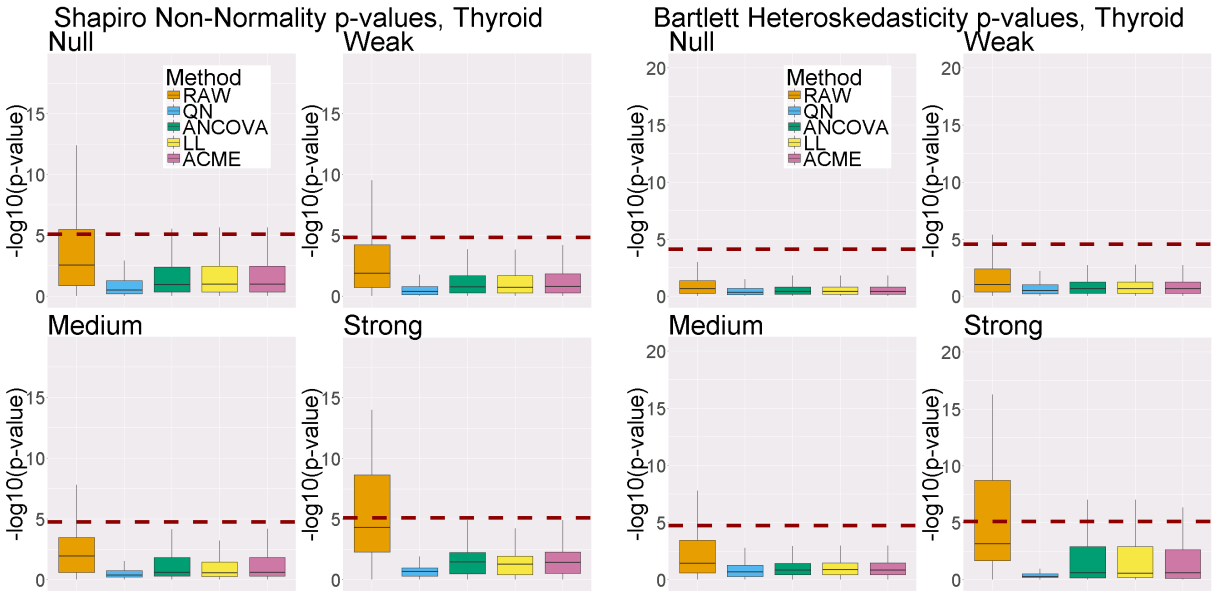

Figure 3: Boxplots of $-\log_{10}$ Shapiro-Wilk and Bartlett p-values from all models. Above, "AOV" denotes the log-ANCOVA model, "LL" the log-linear model, and "RAW" the standard linear model with un-transformed gene expression. The dark red dashed line indicates the FDR $\alpha = 0.1$ significance cut-off for the particular bin.

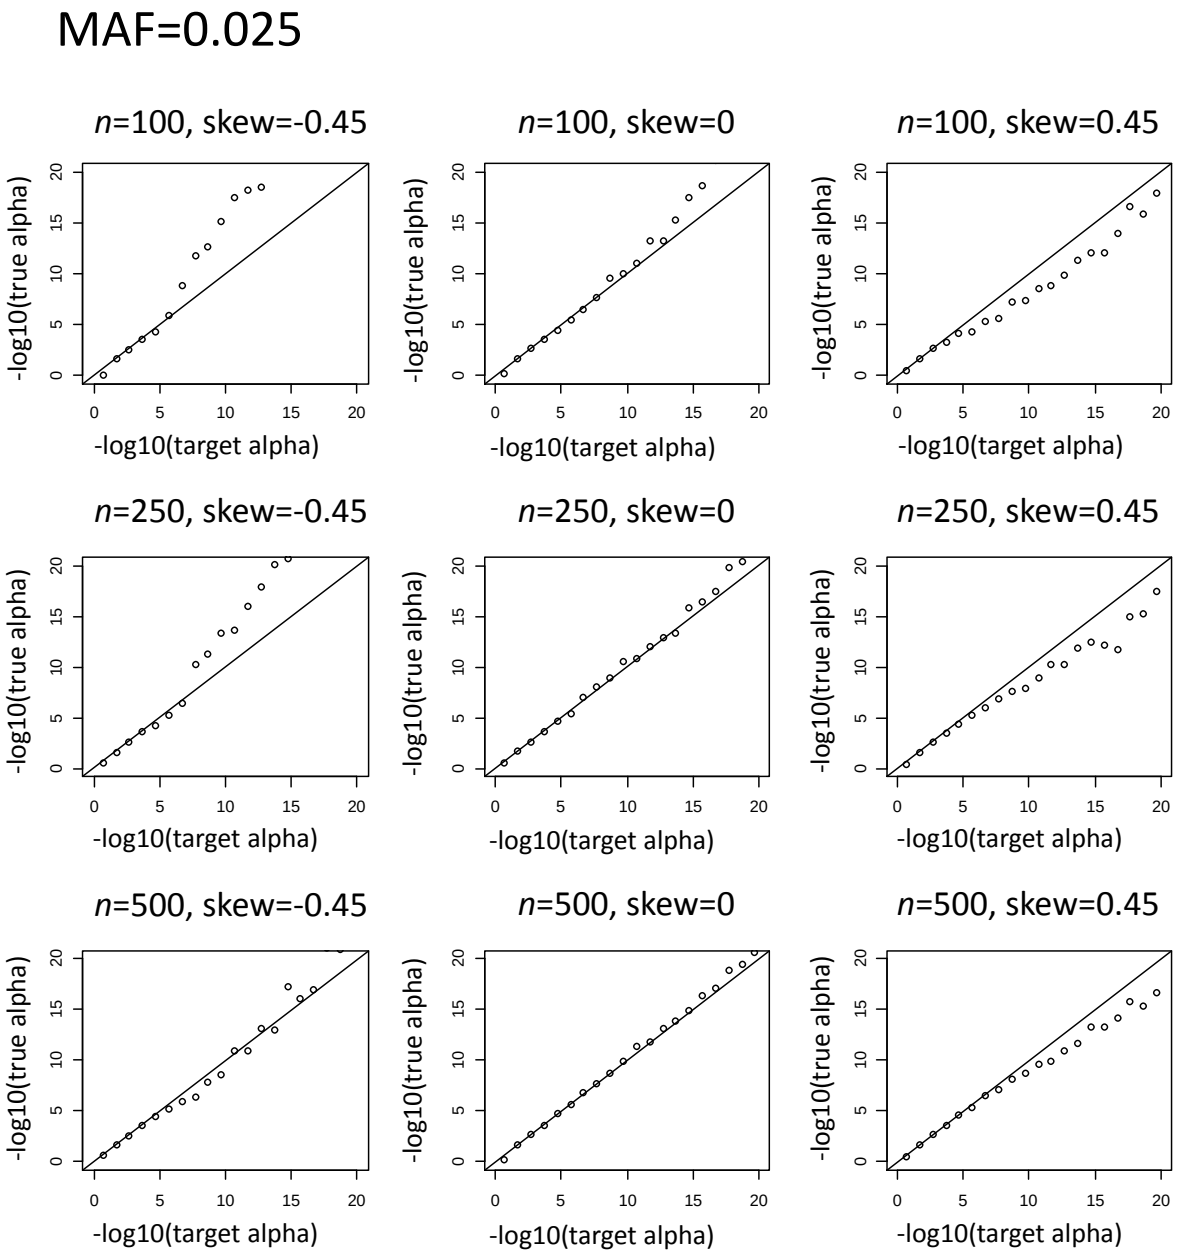

Figure 4: Importance sampling results for ACME F-test p-values. For sample sizes $\geqslant 250$, results are highly accurate or in some cases modestly conservative.

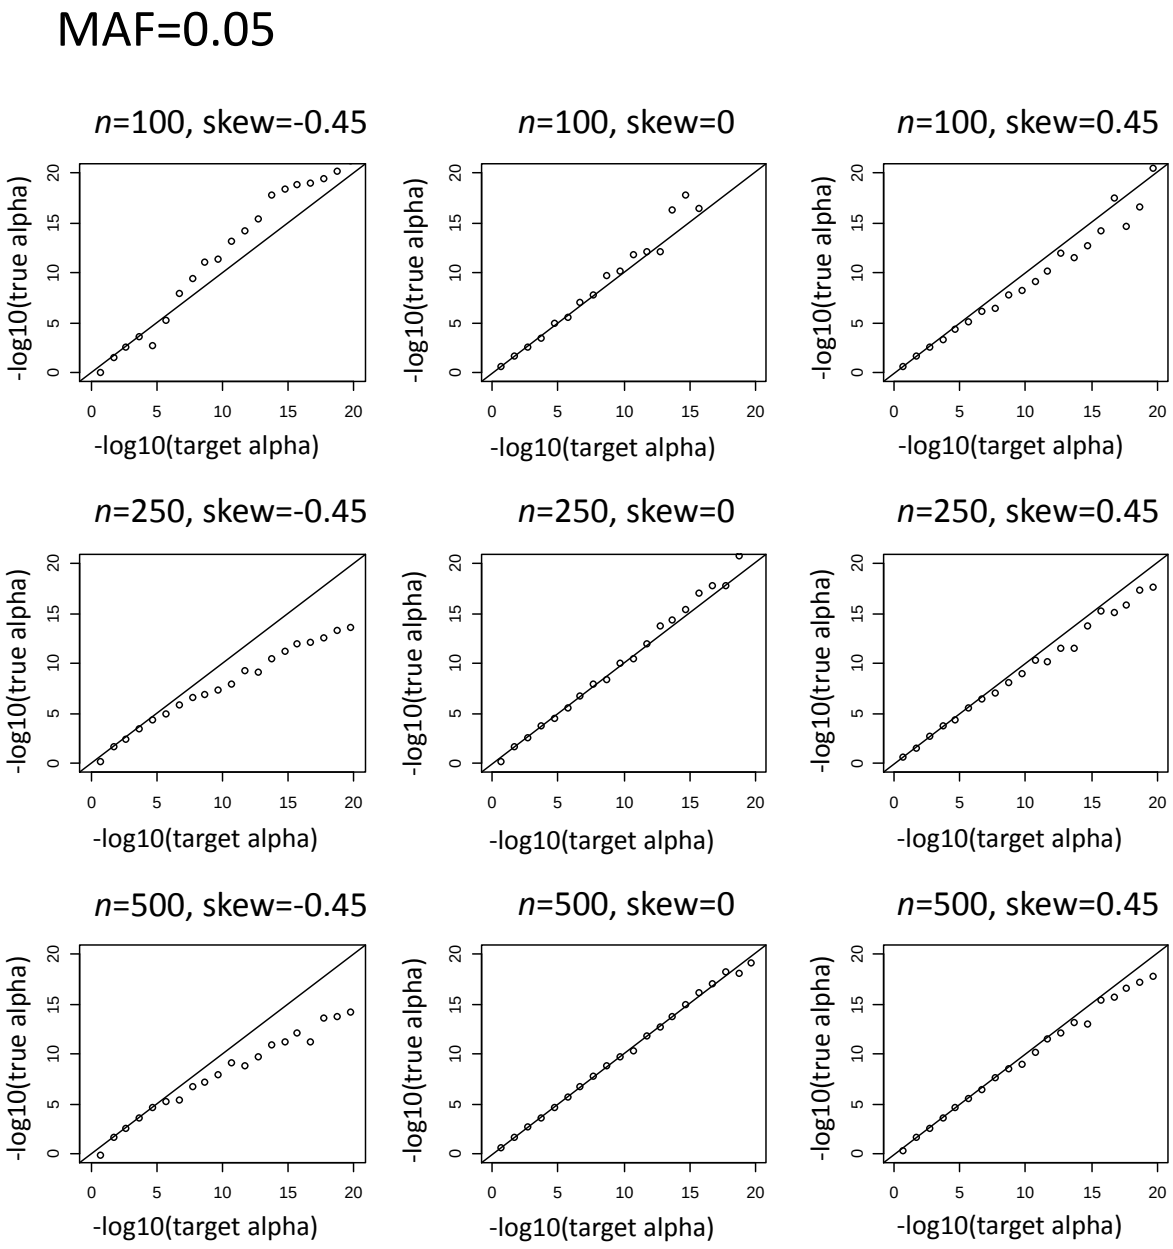

Figure 5: Importance sampling results for ACME F-test p-values. For sample sizes $\geqslant 250$, results are highly accurate or in some cases modestly conservative.

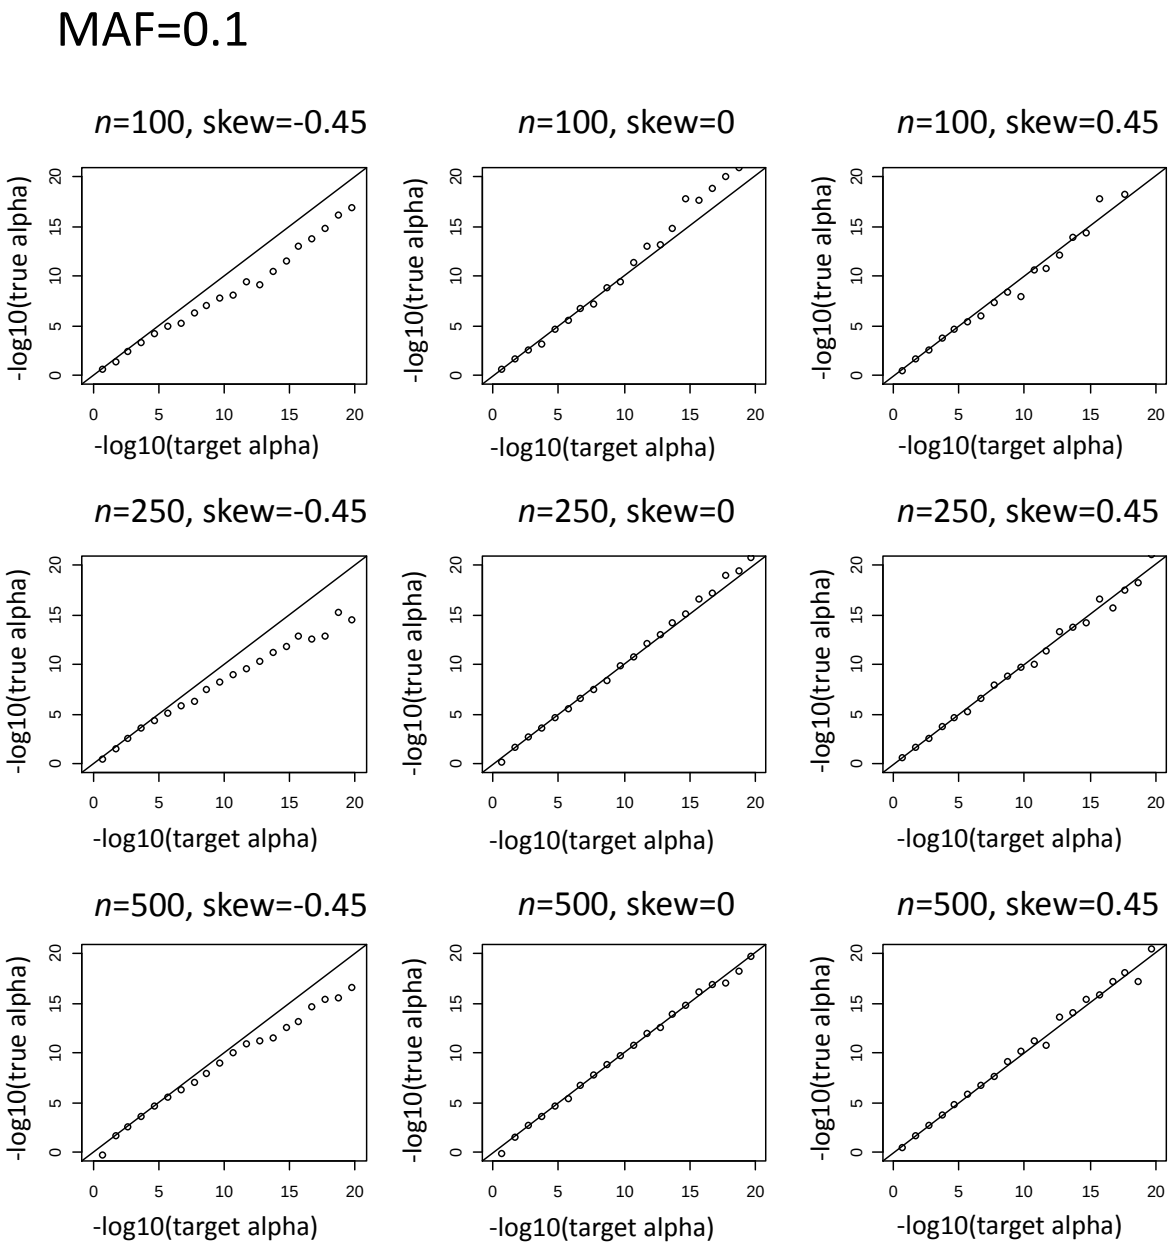

Figure 6: Importance sampling results for ACME F-test p-values. For sample sizes $\geqslant 250$, results are highly accurate or in some cases modestly conservative.

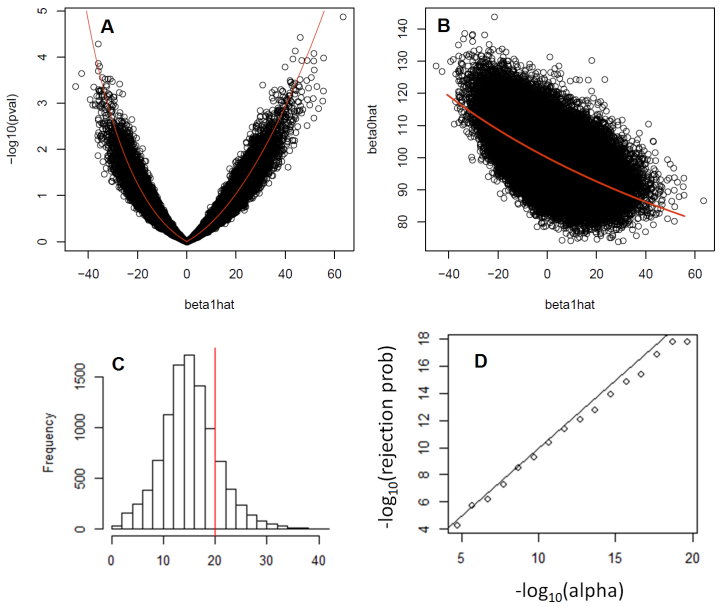

Figure 7: The four panels illustrating importance sampling methods, as described in text II.

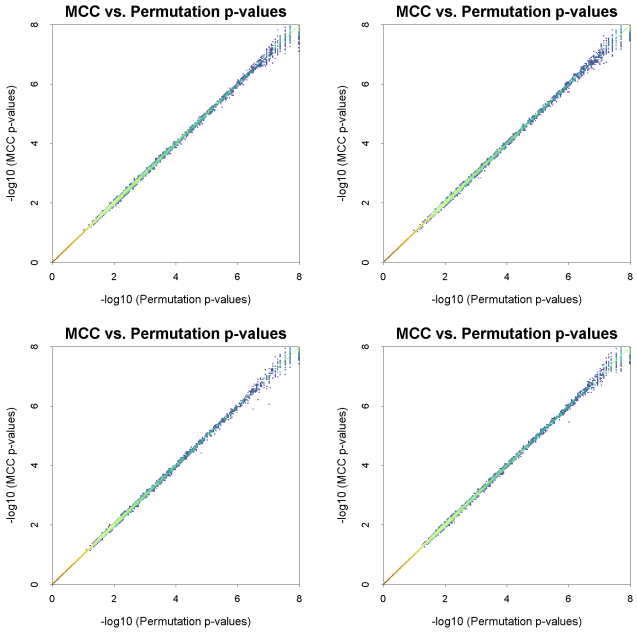

Figure 8: MCC p-values vs. real permutation p-values, each method applied to log-transformed normalized read count vs. rounded allele counts (from Adipose, Artery, Heart, and Muscle, clockwise from the top left).

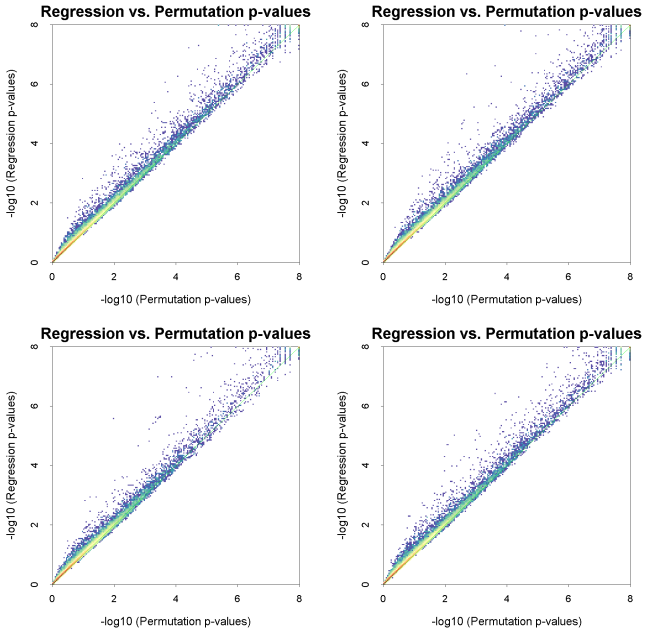

Figure 9: Permutation p-values vs. linear regression p-values, each method applied to log-transformed normalized read count vs. rounded allele counts (from Adipose, Artery, Heart, and Muscle, clockwise from the top left).

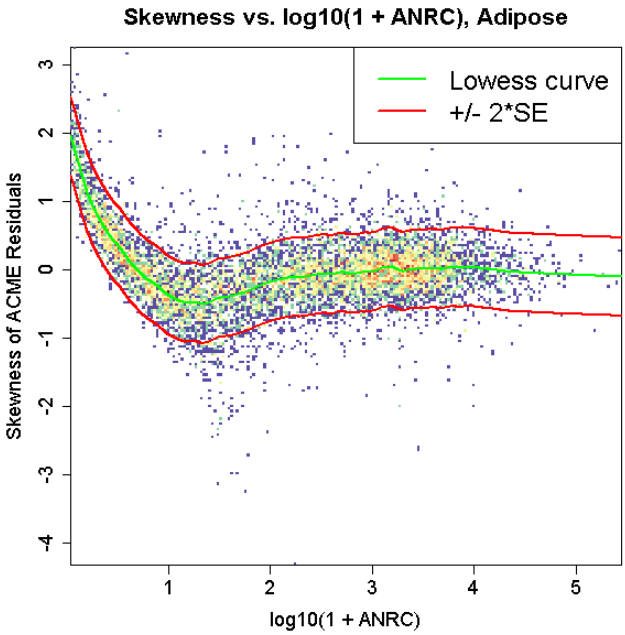

Figure 10: For 10,000 randomly sample gene-SNP pairs from Adipose tissue, the sample skewnesses of eQTL residuals against $\log_{10}$ of the average normalized read counts (across patients) of the gene associated with each eQTL

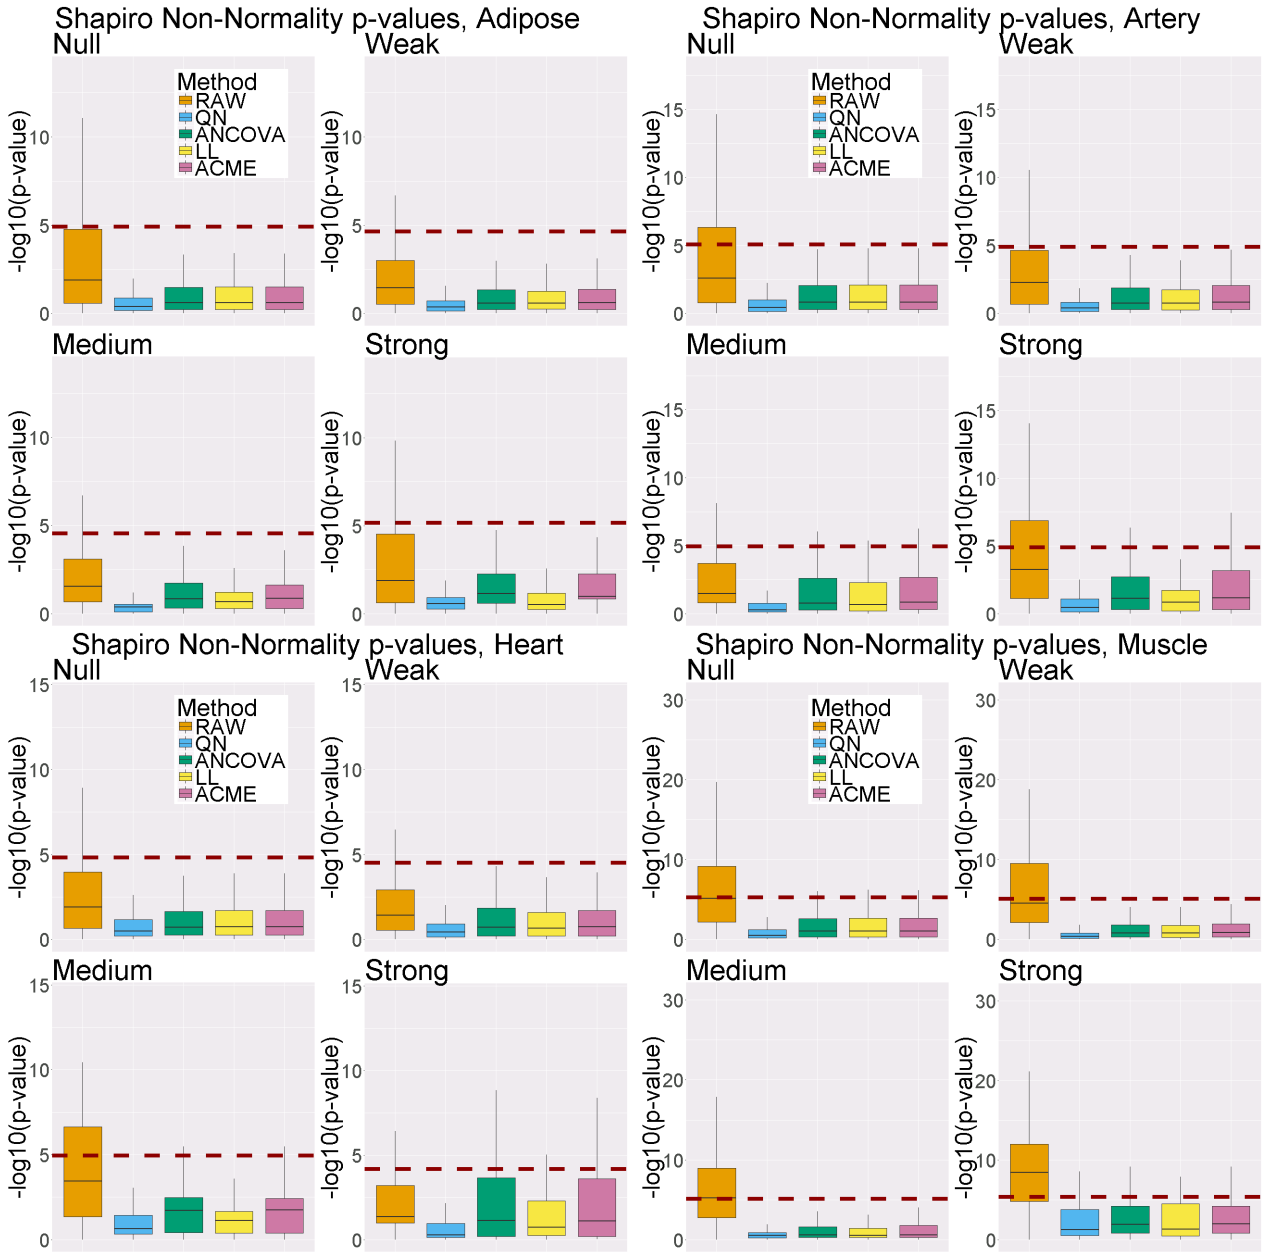

Figure 11: Additional normality test results

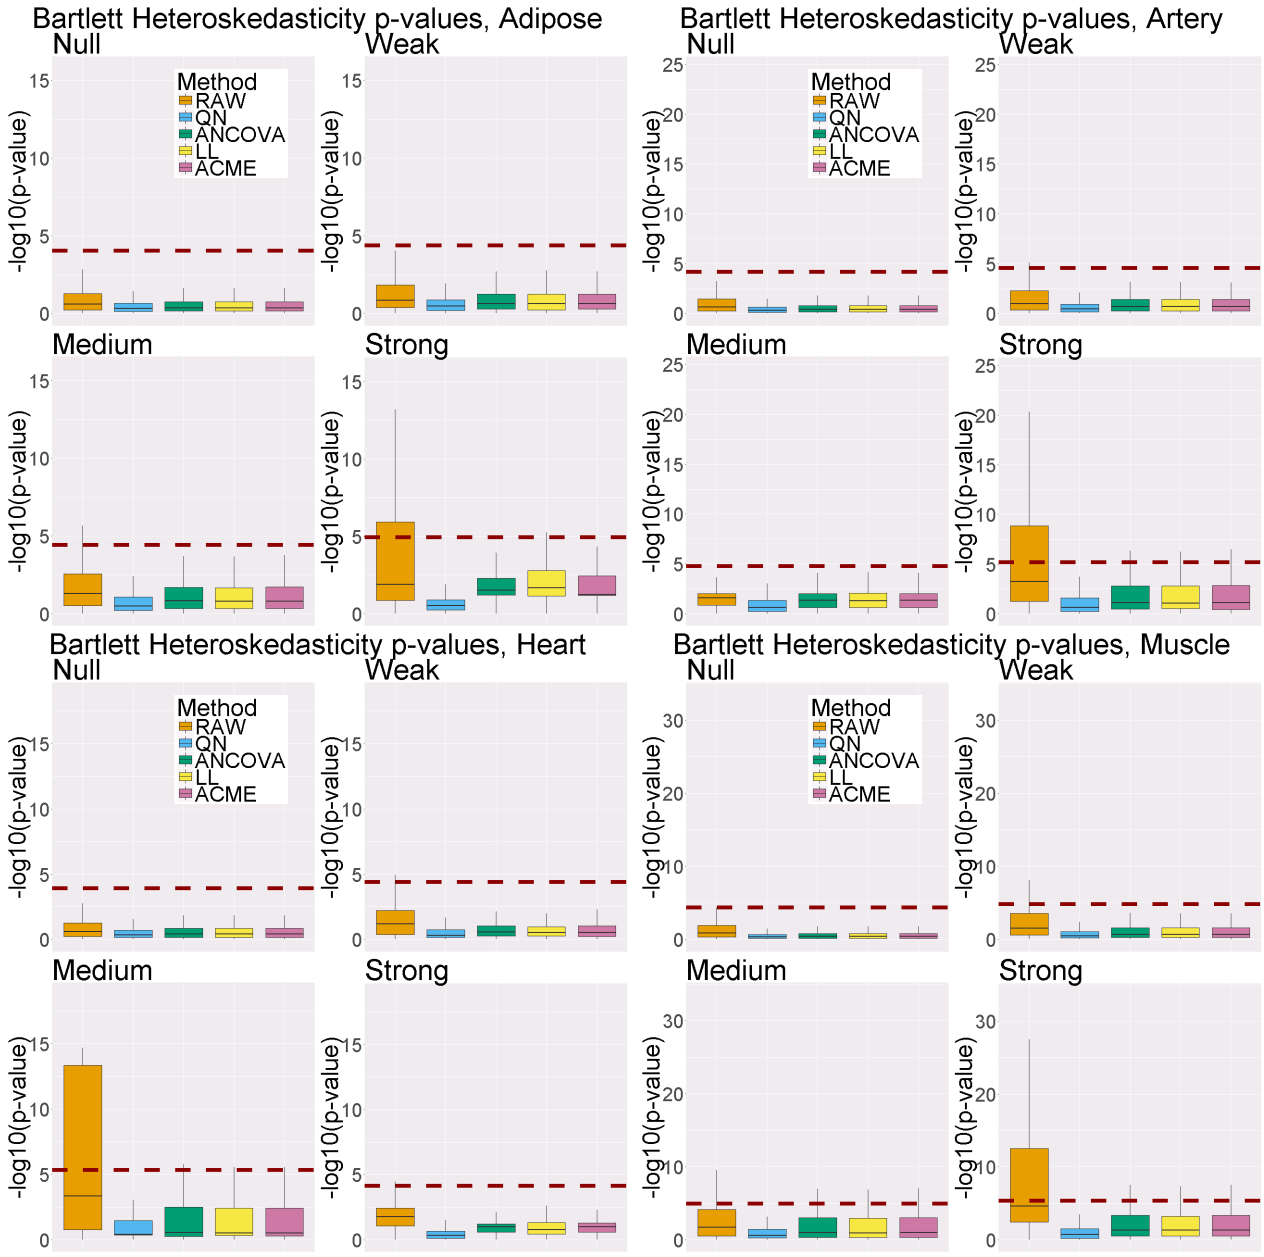

Figure 12: Additional homoskedasticity test results

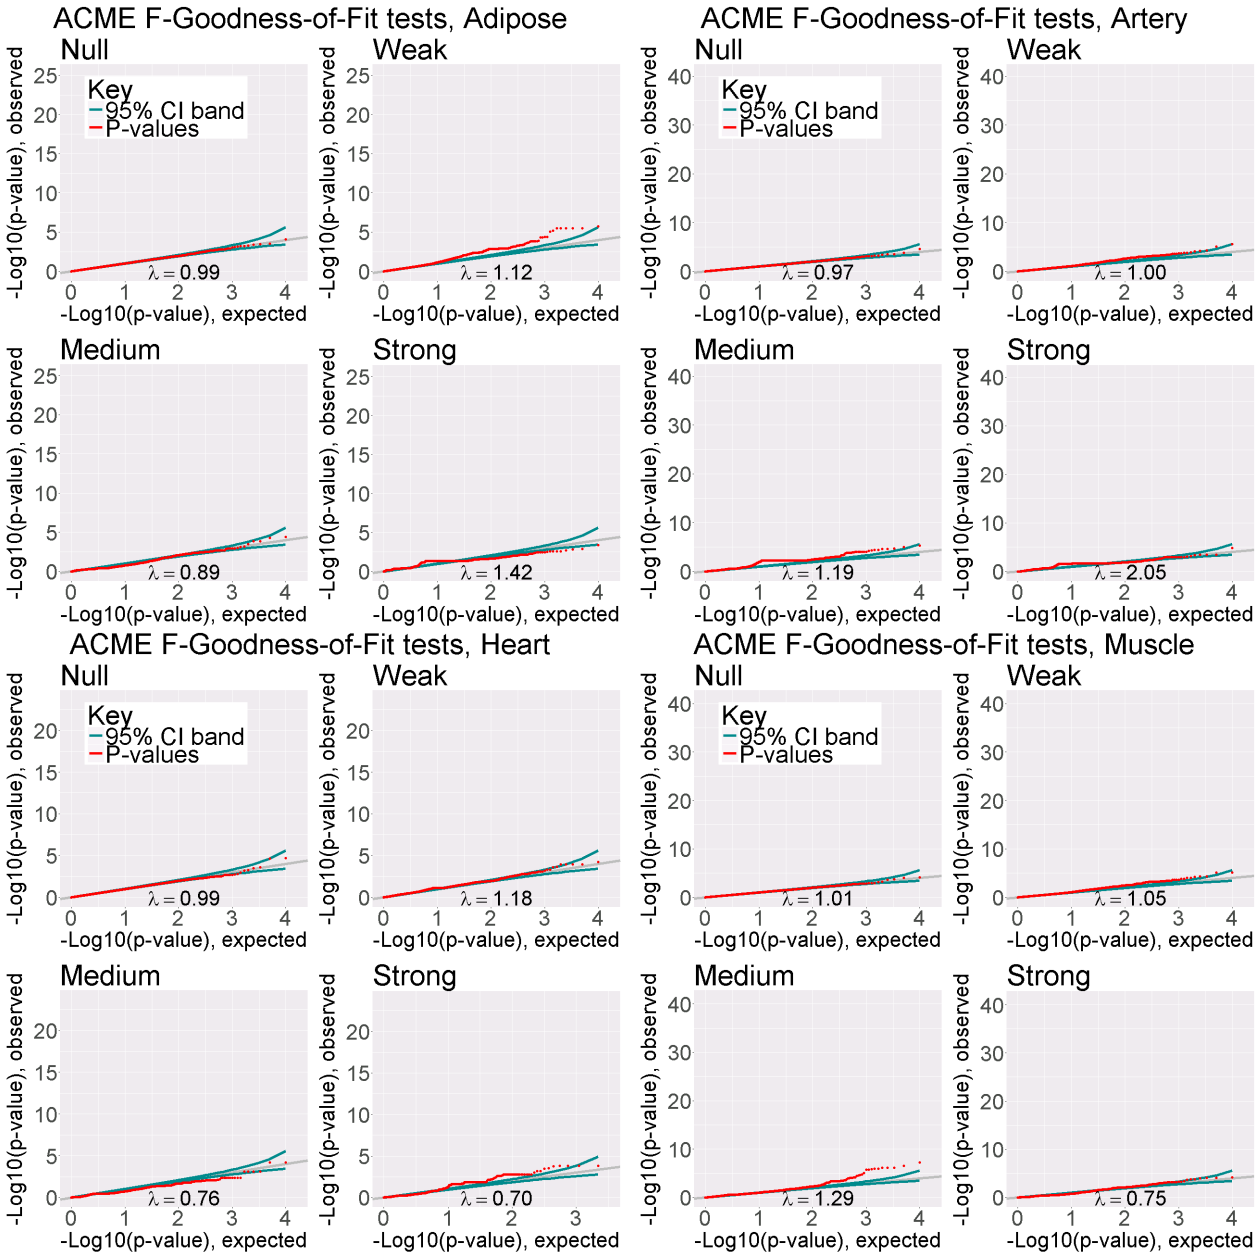

Figure 13: Additional goodness-of-fit test results for ACME

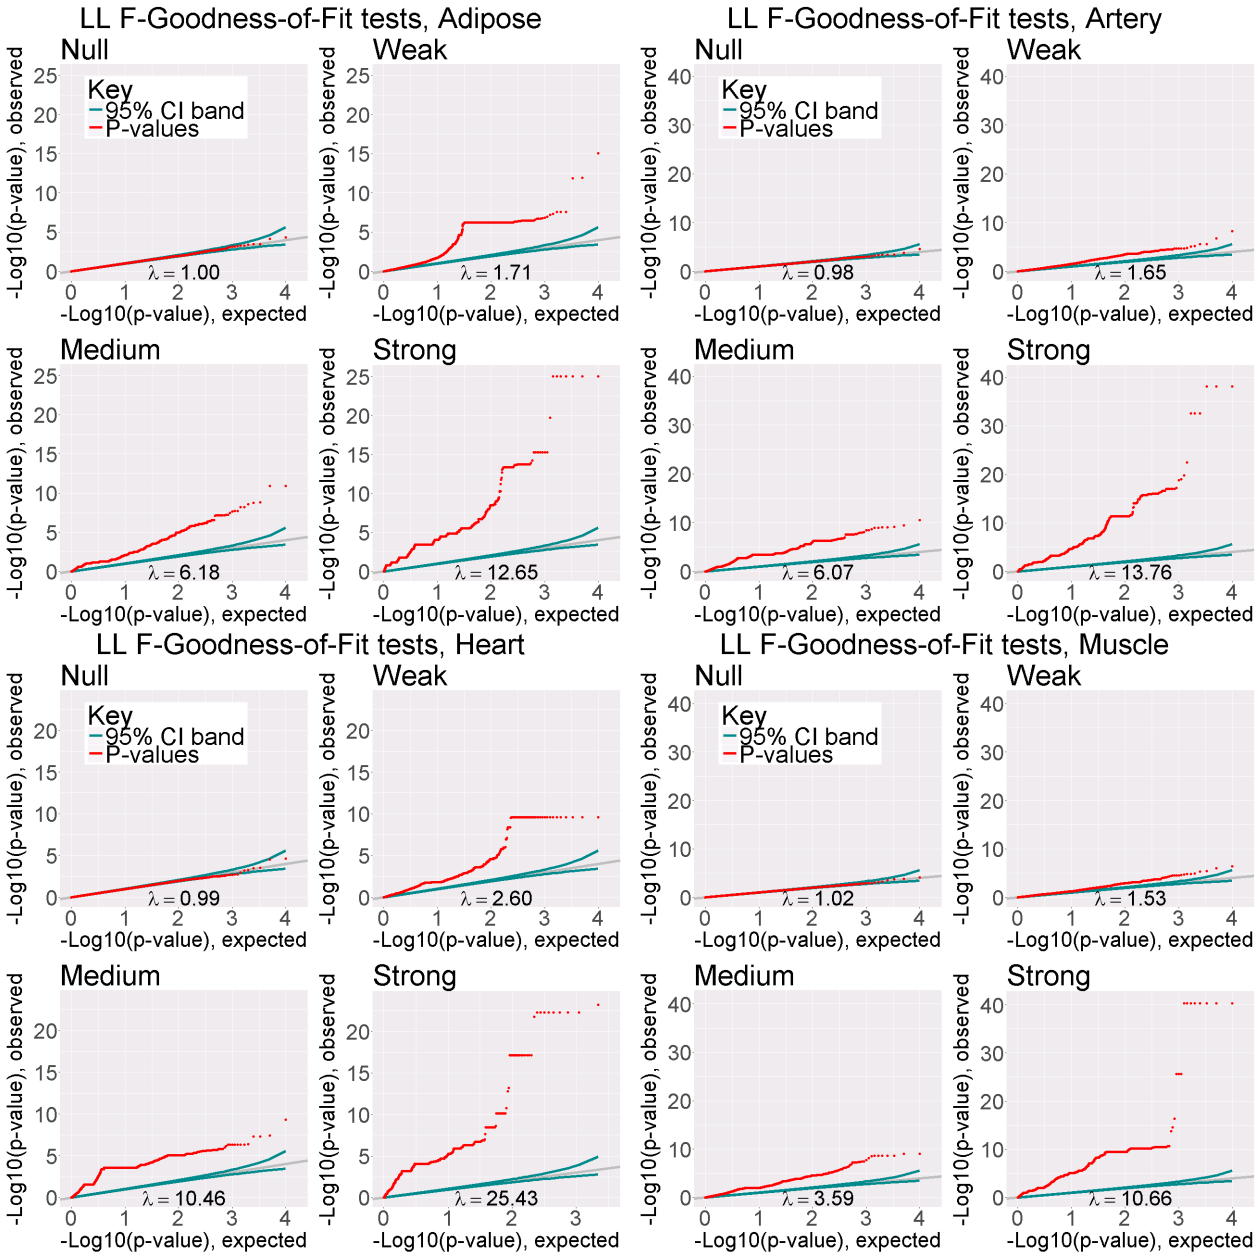

Figure 14: Additional goodness-of-fit test results for log-linear

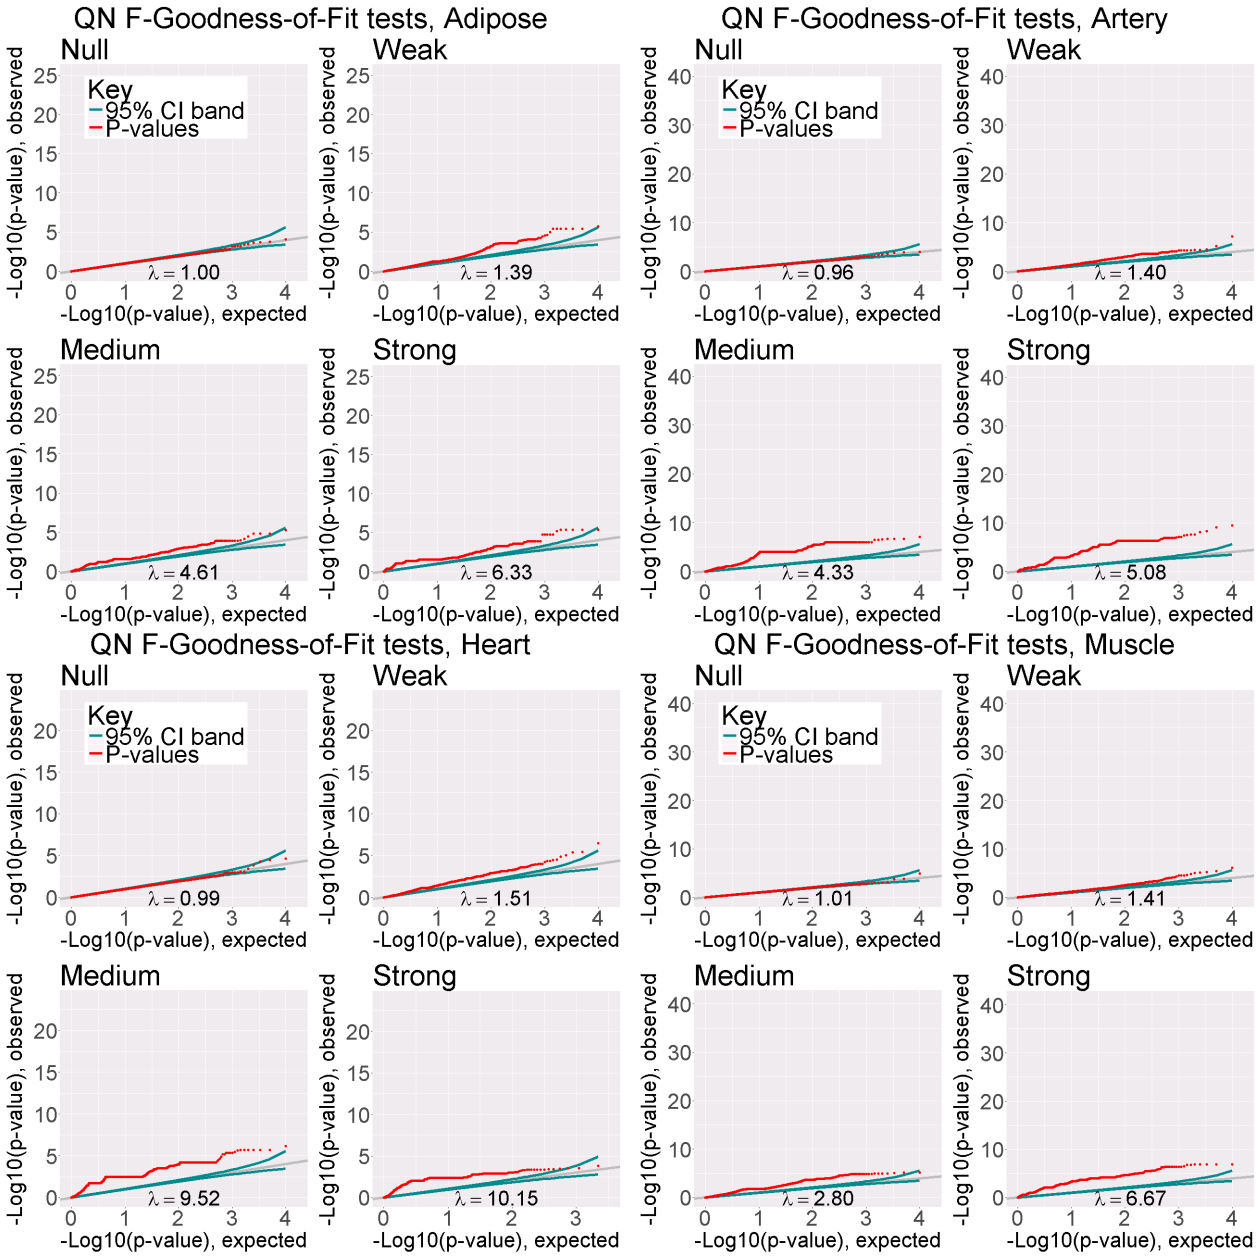

Figure 15: Additional goodness-of-fit test results for QN-linear

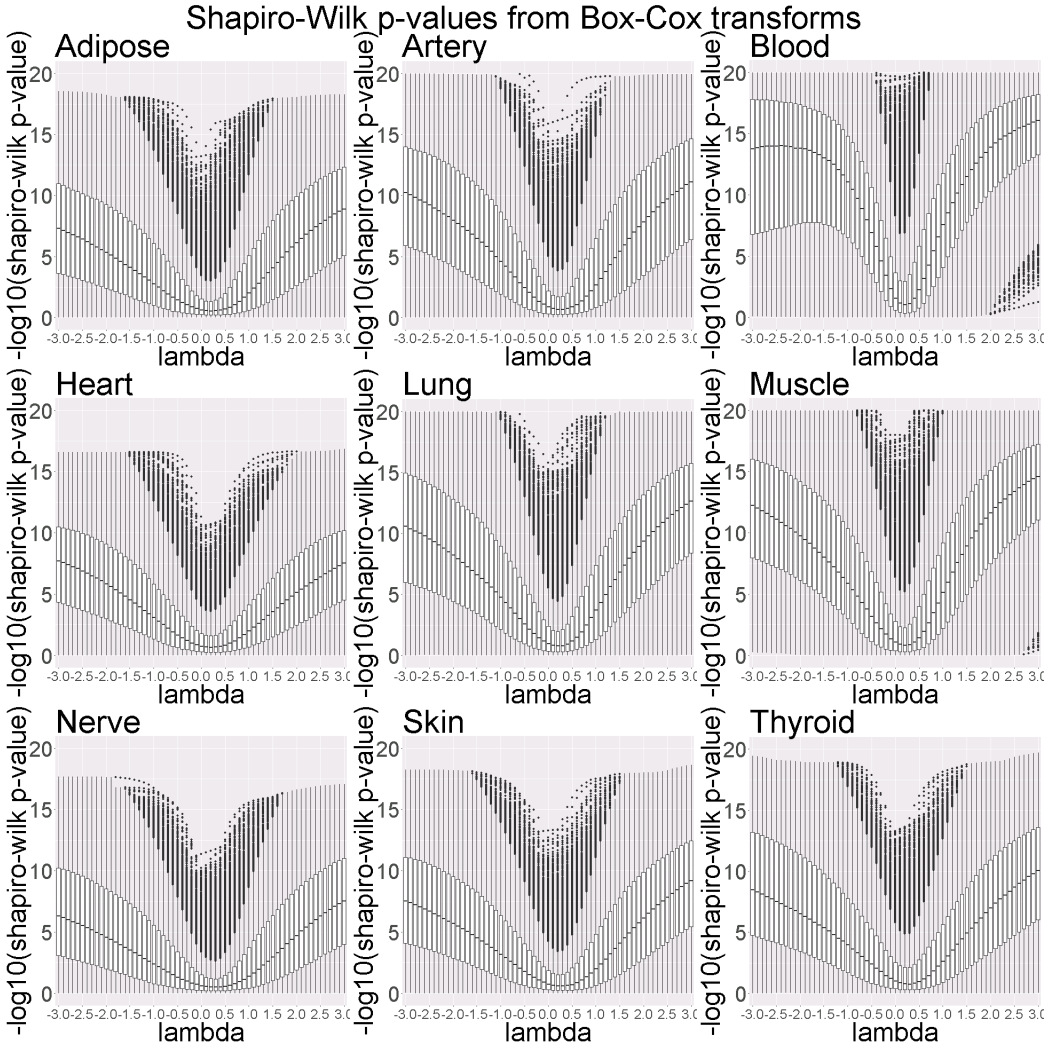

Figure 16: This plot is introduced at the beginning of Section 2 of the main text. For each tissue, we applied the Box-Cox transformation to the expression data of the 10,000 "Null" eQTLs (as defined by Matrix-eQTL p-value, see Web Section Section G). To each tissue-wise set of transformed expression vectors, we then computed the p-value resulting from the Shapiro-Wilk test of normality. This effectively judges various Box-Cox transformations with respect to the normality of the residual distribution (as judged by the un-extremity of the p-value) of eQTLs from real data.

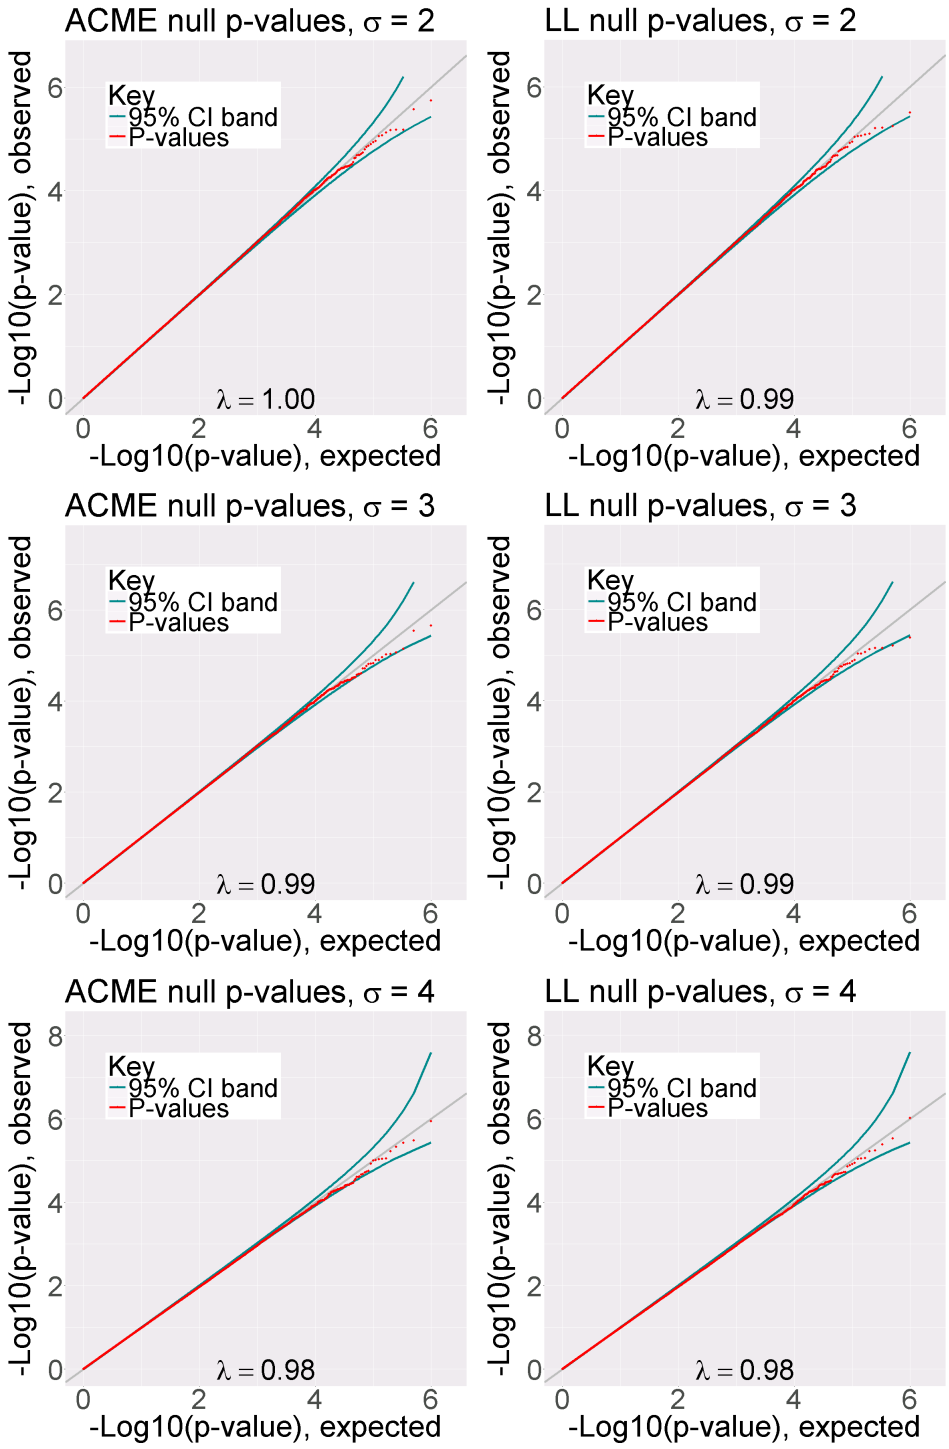

Figure 17: Additional null simulations for ACME and LL. p-value distributions from null simulated data with realistic errors and real covariate/genotype data. λ values are inflation factors.

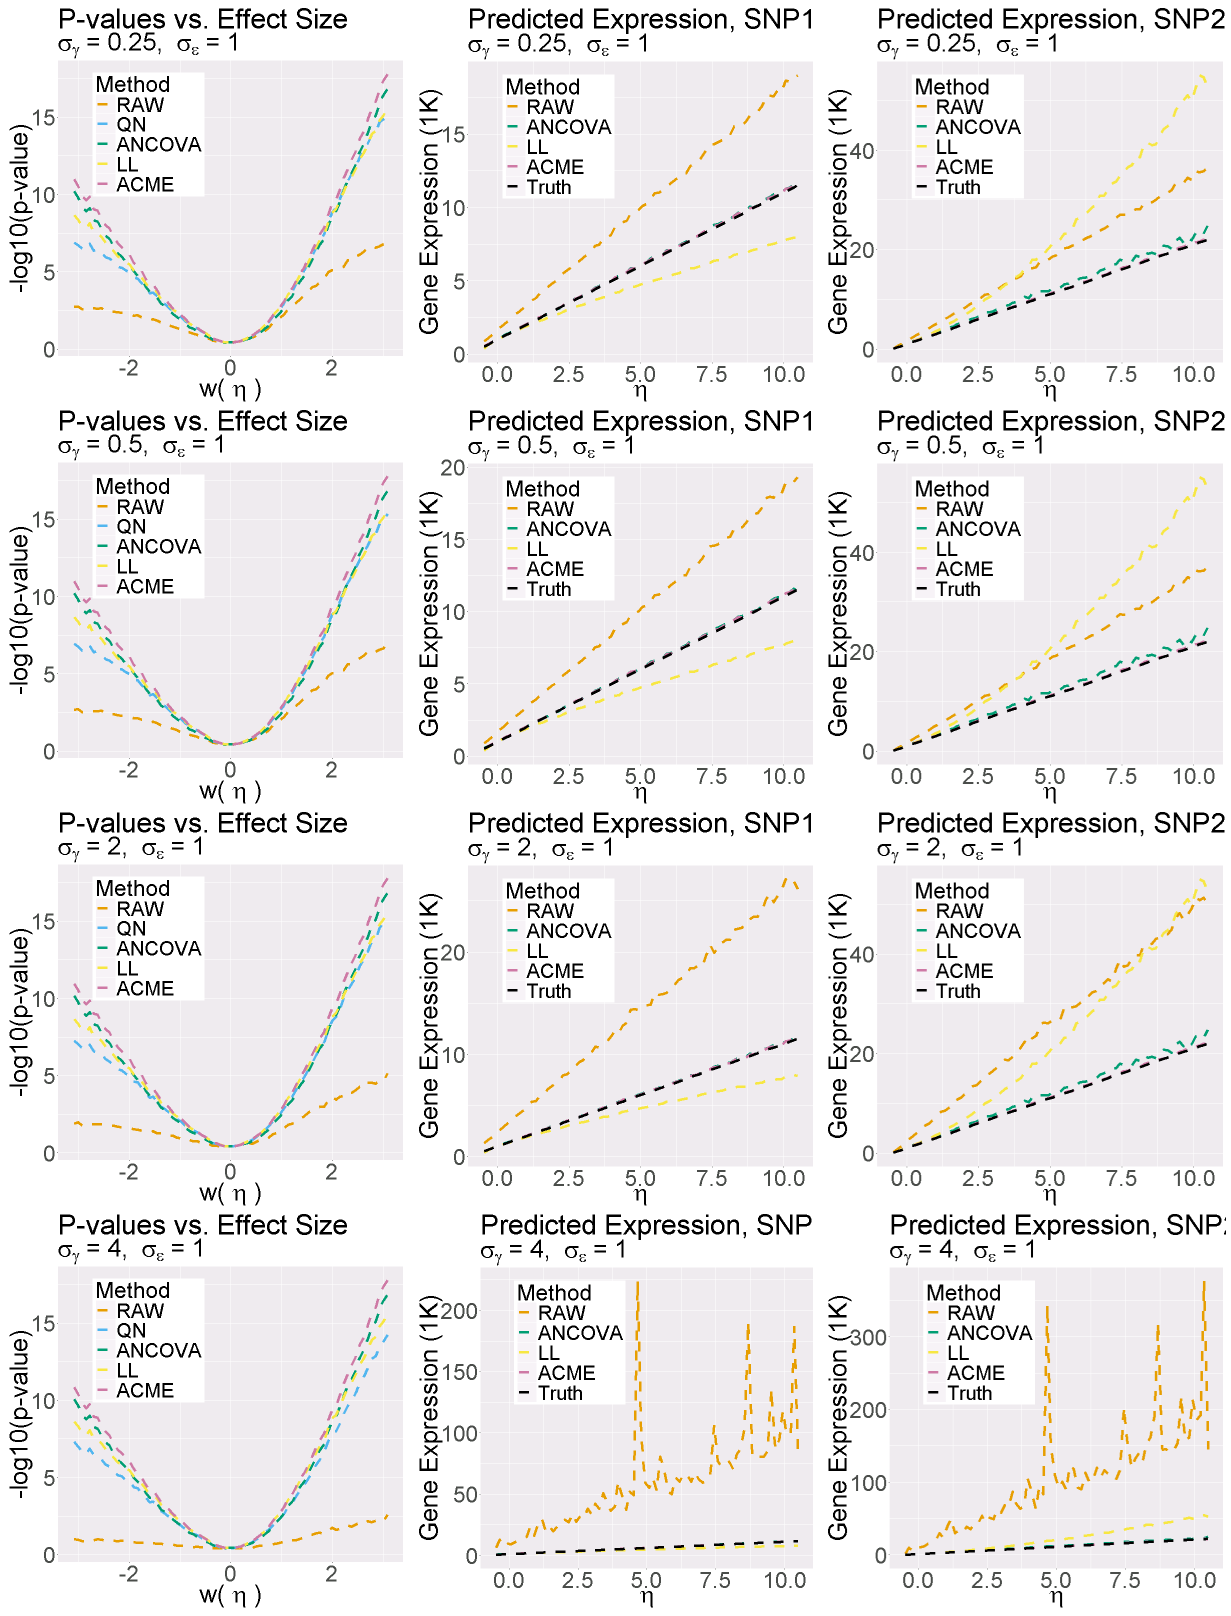

Figure 18: Results of large-scale simulation experiment, with increasing σ_γ (results from $\sigma_\gamma = 1$ in the main document. Left column: $-\log 10$ F-test p-values as a function of η. Middle and right columns: predicated raw expression with one and two reference alleles, respectively.

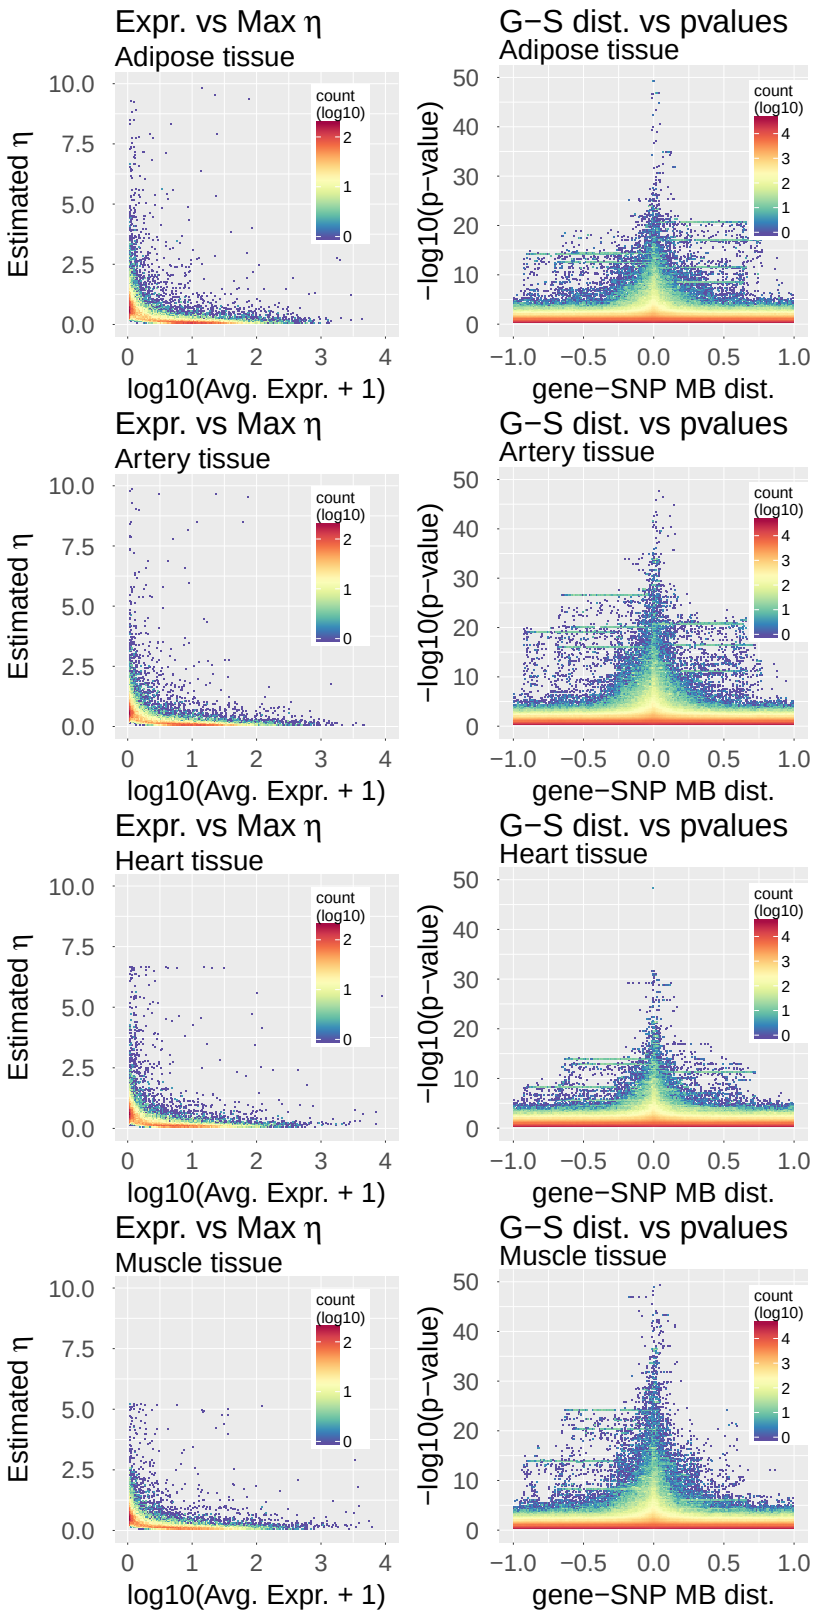

Figure 19: Results of genome-wide cis-eQTL ACME effect size estimations on (by row) Adipose, Artery, Heart, and Muscle tissue, from GTEx Pilot data. Left column: Maximum gene-wise estimated effect size vs. log average expression level. Middle column: $-\log_{10}$ ACME p-value vs. distance from gene TSS to SNP position. These are known patterns of full-genome cis-eQTLs (for instance, that effect sizes are stronger and more significant when SNP is close to gene transcription start site), and serve as QC checks for ACME estimates.

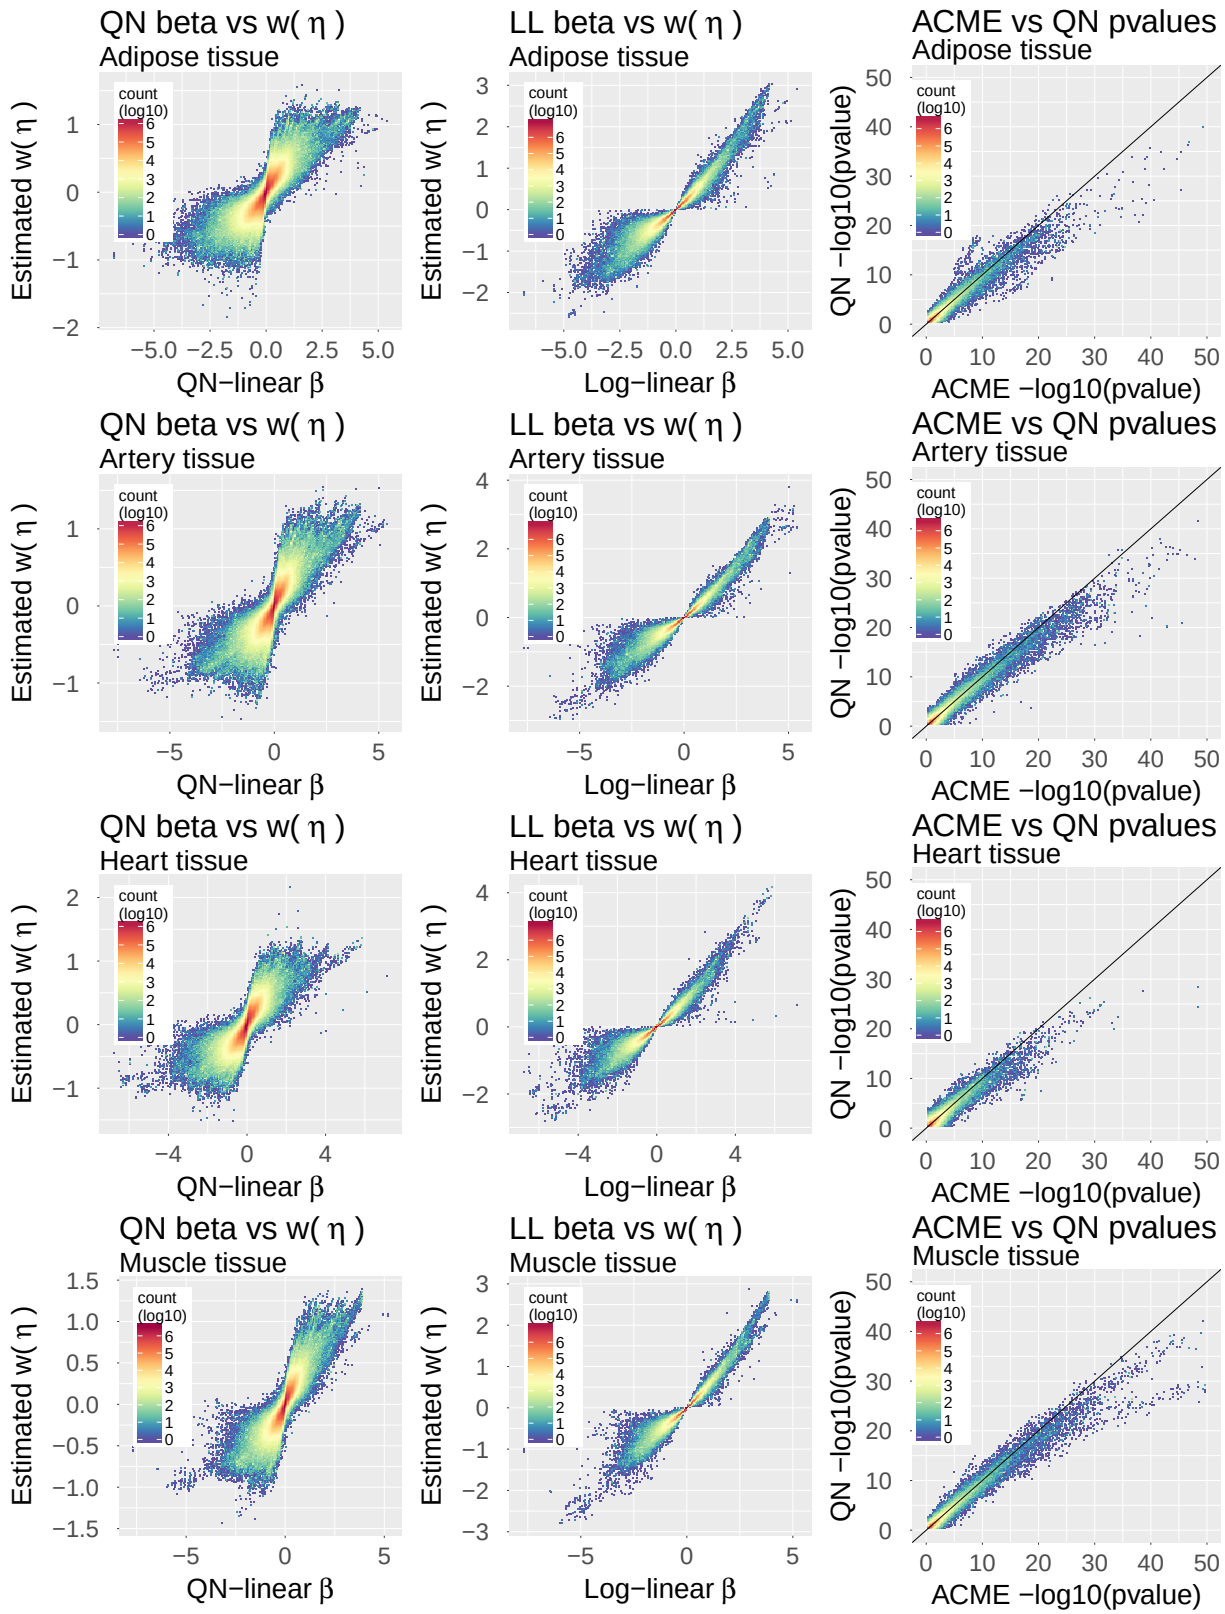

Figure 20: Results of genome-wide cis-eQTL effect size estimations on (by row) Adipose, Artery, Heart, and Muscle tissue, from GTEx Pilot data. Left (and middle) column: ACME effect size (under w transformation from Section 4 in main document) vs QN (and LL) regression. Right column: QN vs ACME regression p-values.