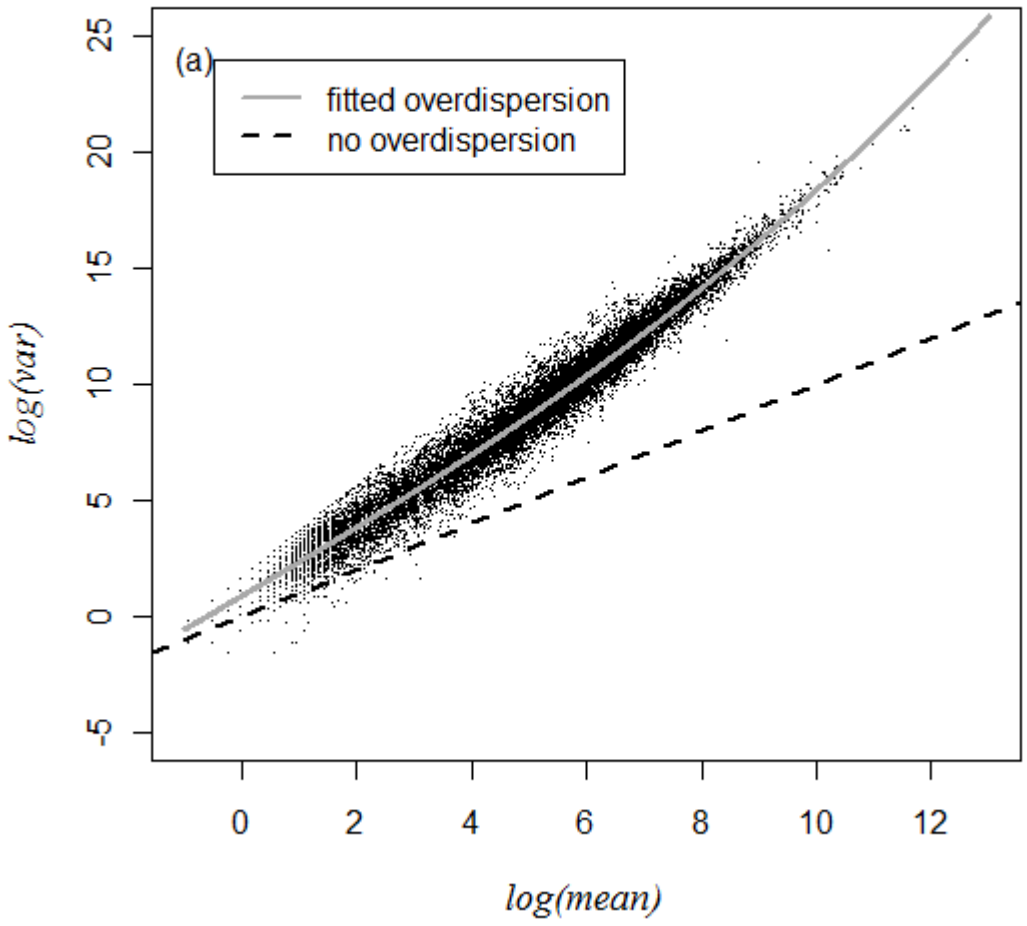

Supplementary Figure 1: The log(variance) vs. log(mean) relationship for 5 samples drawn randomly from the 60 samples in the Montgomery CEU dataset. Even with a small sample size, the mean-overdispersion relationship, as evidenced by the increasing gap between the fitted (third degree polynomial) over-dispersion curve and the "no over-dispersion" curve, still holds.

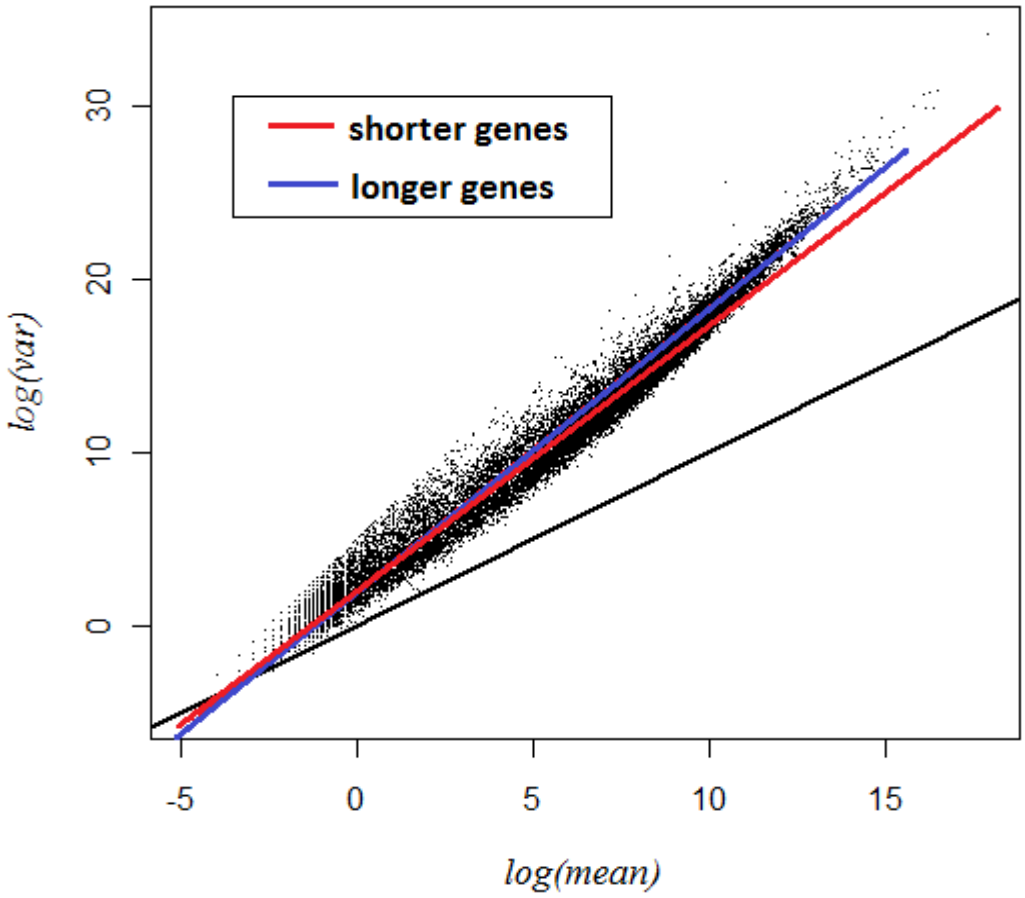

Supplementary Figure 2. Log(variance) vs. log(mean) for the CEU data (60 samples), with separate regression lines for "shorter" and "longer" genes (corresponding to the shortest 1/3 and longest 1/3 of transcripts). The transcript length is, at best, a very modest determinant of overdispersion.

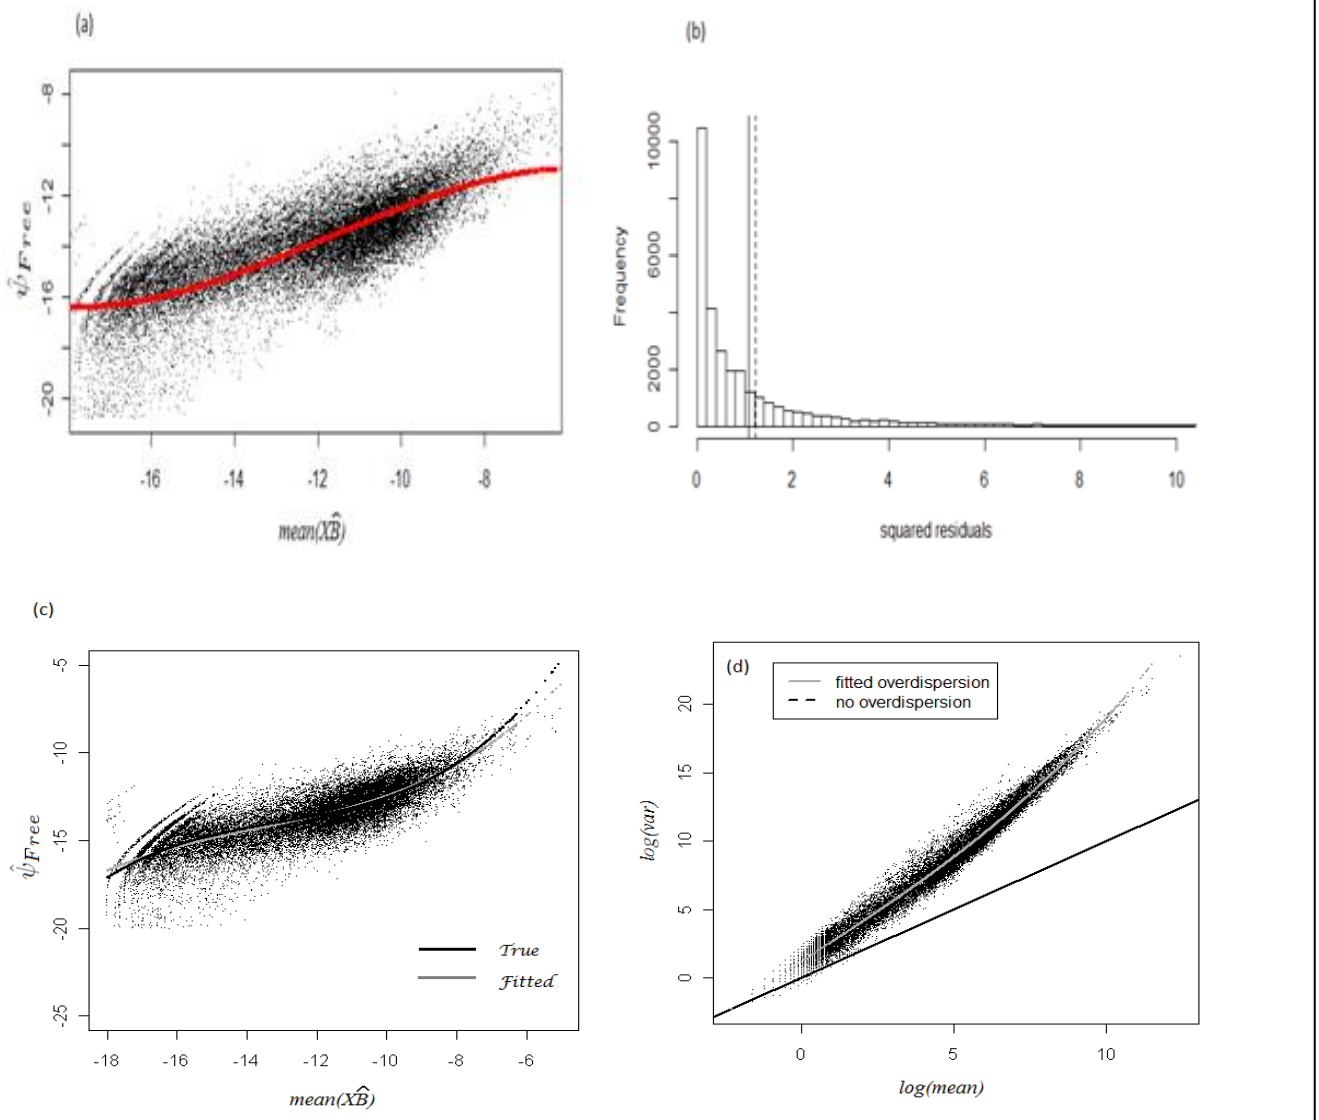

Supplementary Figure 3. Fitted values for the mean-overdispersion relationship, analysis of CEU random samples and simulations (a) $\hat{\psi}$ values for 32027 genes vs. mean ($X\hat{B}$), for a random set of 6 vs. 6 CEU samples, excluding genes with all zero counts in one or both experimental conditions. The third degree polynomial relationship displays a reasonable fit. (b) The histogram of squared residuals from the third-degree polynomial fit to (a) above. The mean of residual^2 (vertical solid line) is very close to the average squared standard error of $\hat{\psi}$ (dashed line), indicating that the predominant source of variation in (a) above is due to sampling variation in estimation of $\hat{\psi}$. (c) Example $\hat{\psi}$ values from a simulated 5 vs. 5 samples, with true γ values as estimated from a random sample of 5 vs. 5 CEU samples. The fitted curve to $\hat{\psi}$ agrees closely with the true curve. (d) The mean-overdispersion plot of the purely simulated data from (c).

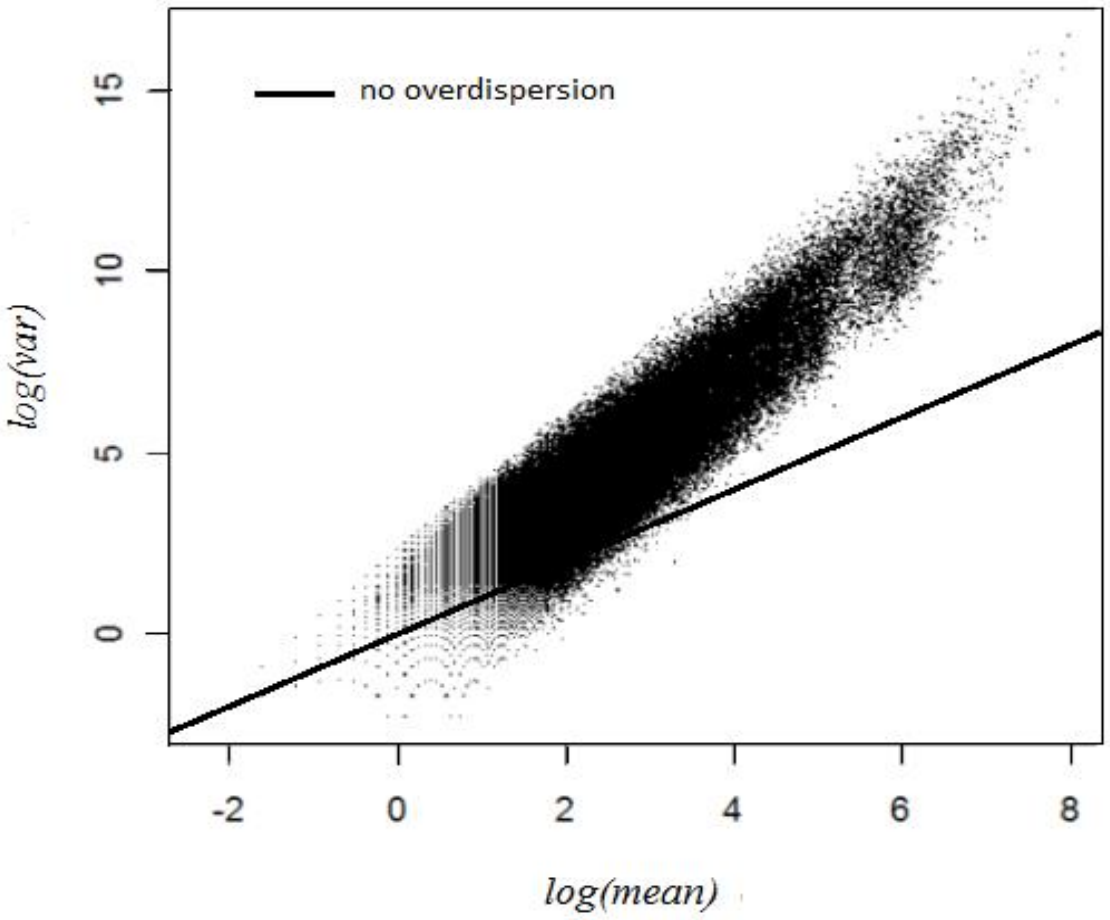

Supplementary Figure 4: The log(variance) vs. log(mean) relationship for one of the simulations in Dataset 1, and used in the baySeq paper (Hardcastle and Kelly, 2010). There are 10,000 genes and 10 samples, all from a single simulated condition. Although the pattern looks similar to the observed pattern in Supplementary Figure 1, the relationship is not as strong (R^2=0.87 for these data vs. R^2=0.98 and R^2=0.97 respectively for the CEU data in Figure 1 and Supplementary Figure 1), and current RNA-Seq datasets typically show higher read counts for RefSeq genes.

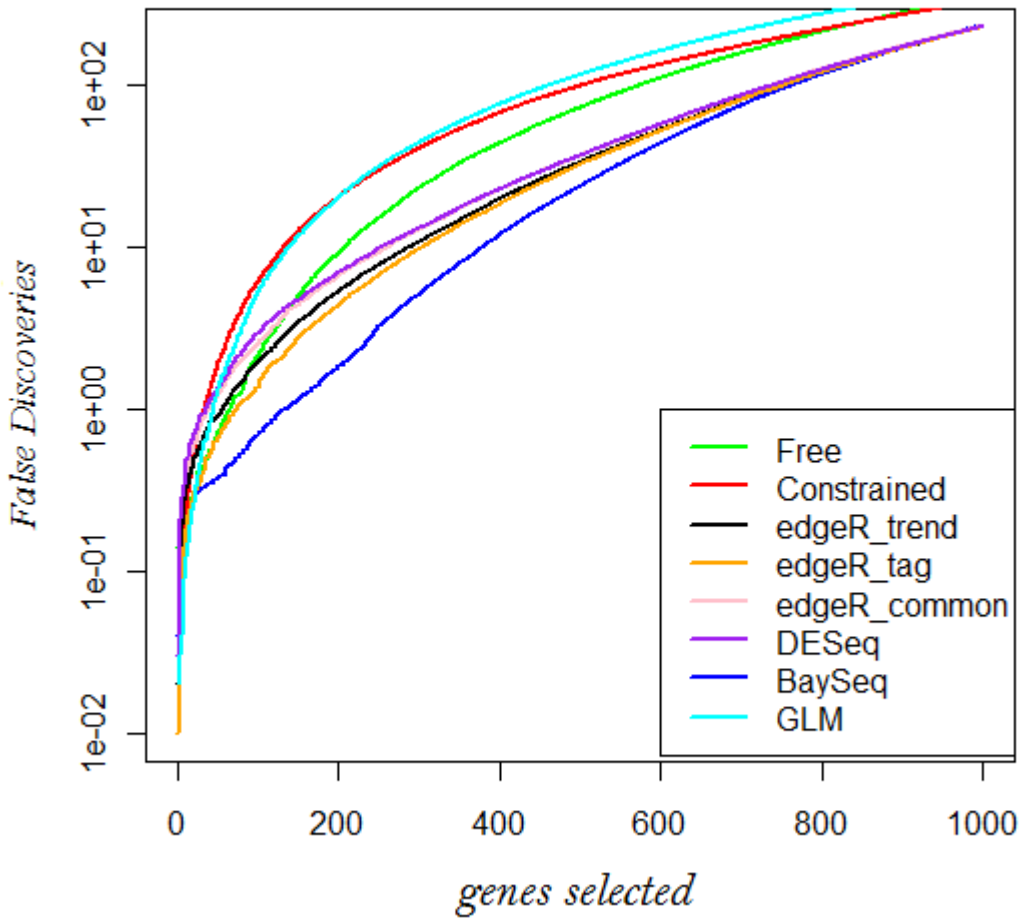

Supplementary Figure 5: False discoveries vs. number of genes selected for Dataset 1.

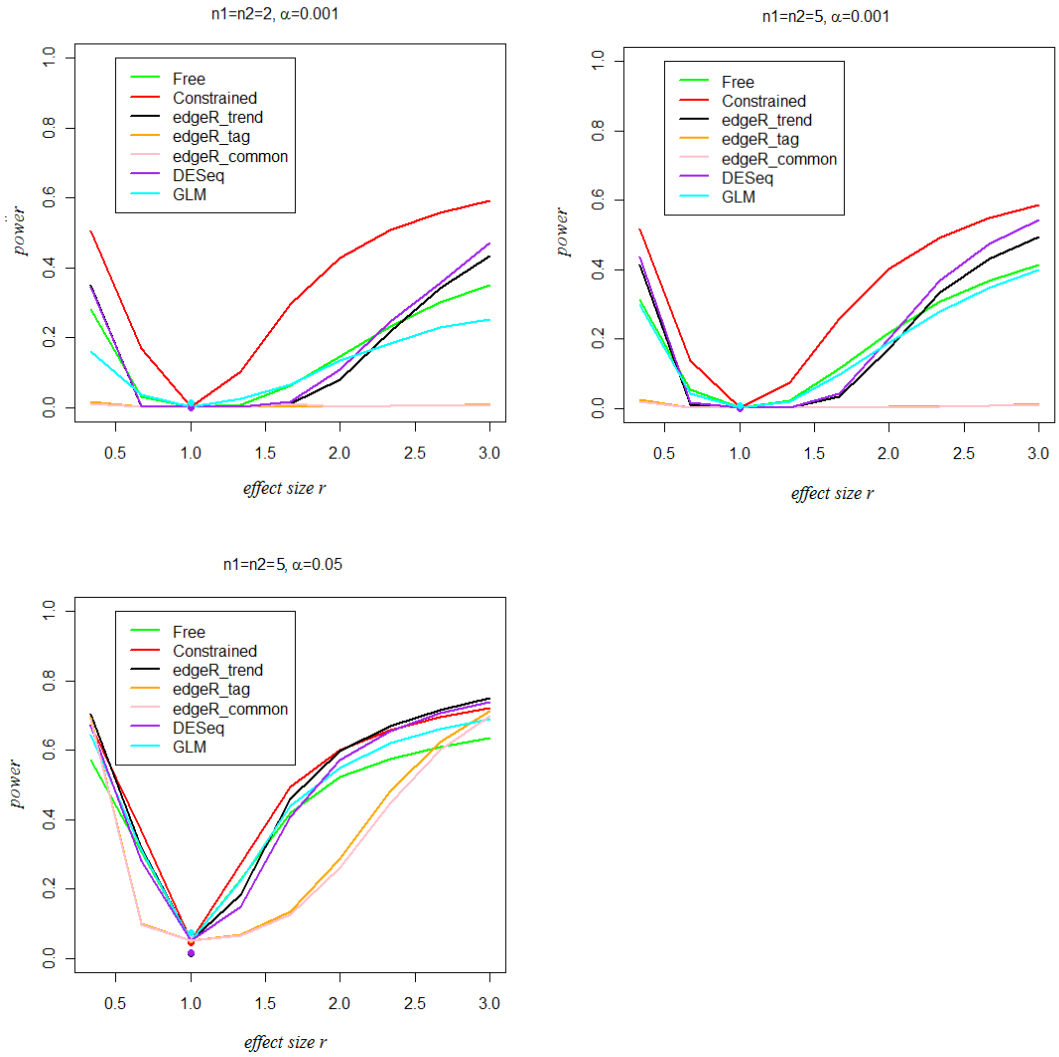

Supplementary Figure 6: Power comparisons for Dataset 2 (32,027 simulated genes), the very small sample scenario, with n_1=2 vs. n_2=2 and n_1=5 vs. n_2=5. The power curves utilize an empirical α threshold for each method, ensuring power comparability across the methods. The dots for r=1 show the type I error using the nominal p-values provided by the methods.

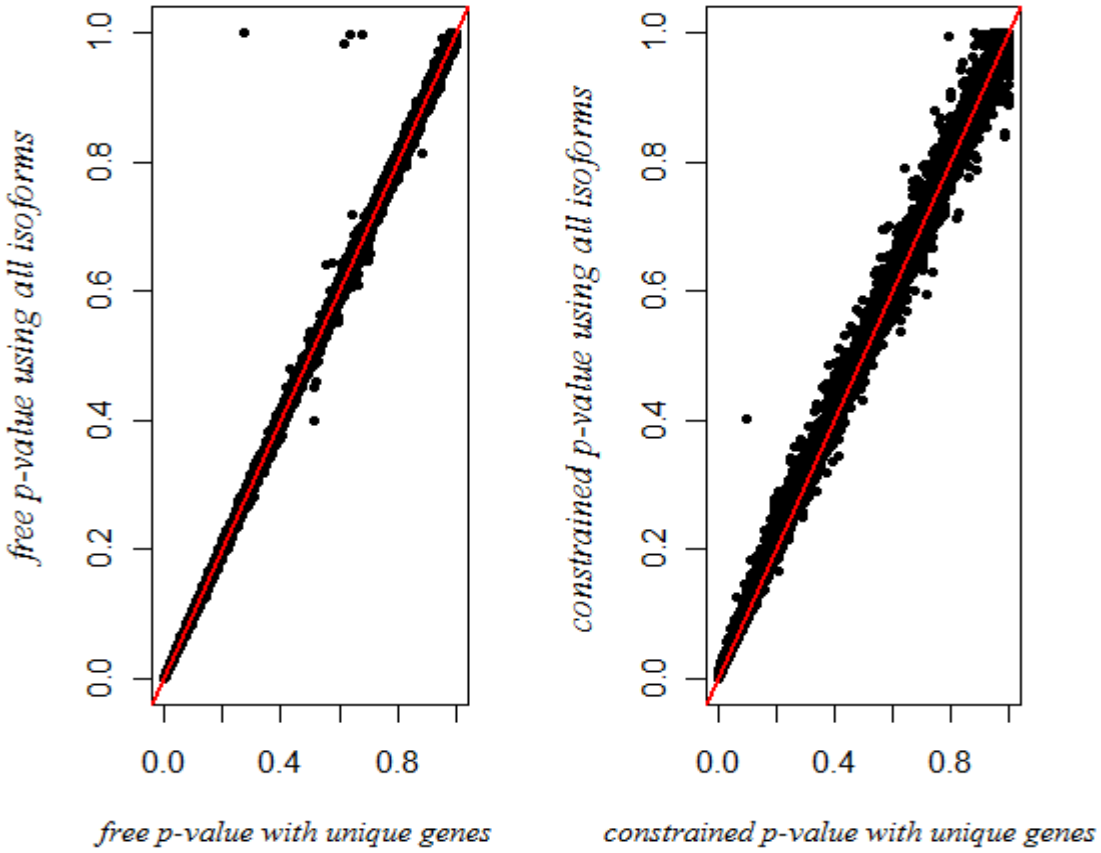

Supplementary Figure 7. *P*-values for n_1=6 samples vs. n_2=6 samples, drawn at random from the CEU dataset with 3 females and 3 males in each group, with the results using all common isoforms considered as separate "genes" plotted vs. the results in which only the most common of the isoforms is used. Left panel: *p*-values from the free model. Right panel: *p*-values from the constrained model. The results illustrate that inclusion of redundant isoforms produces largely the same inference as restriction to unique genes (especially for the smallest p-values).

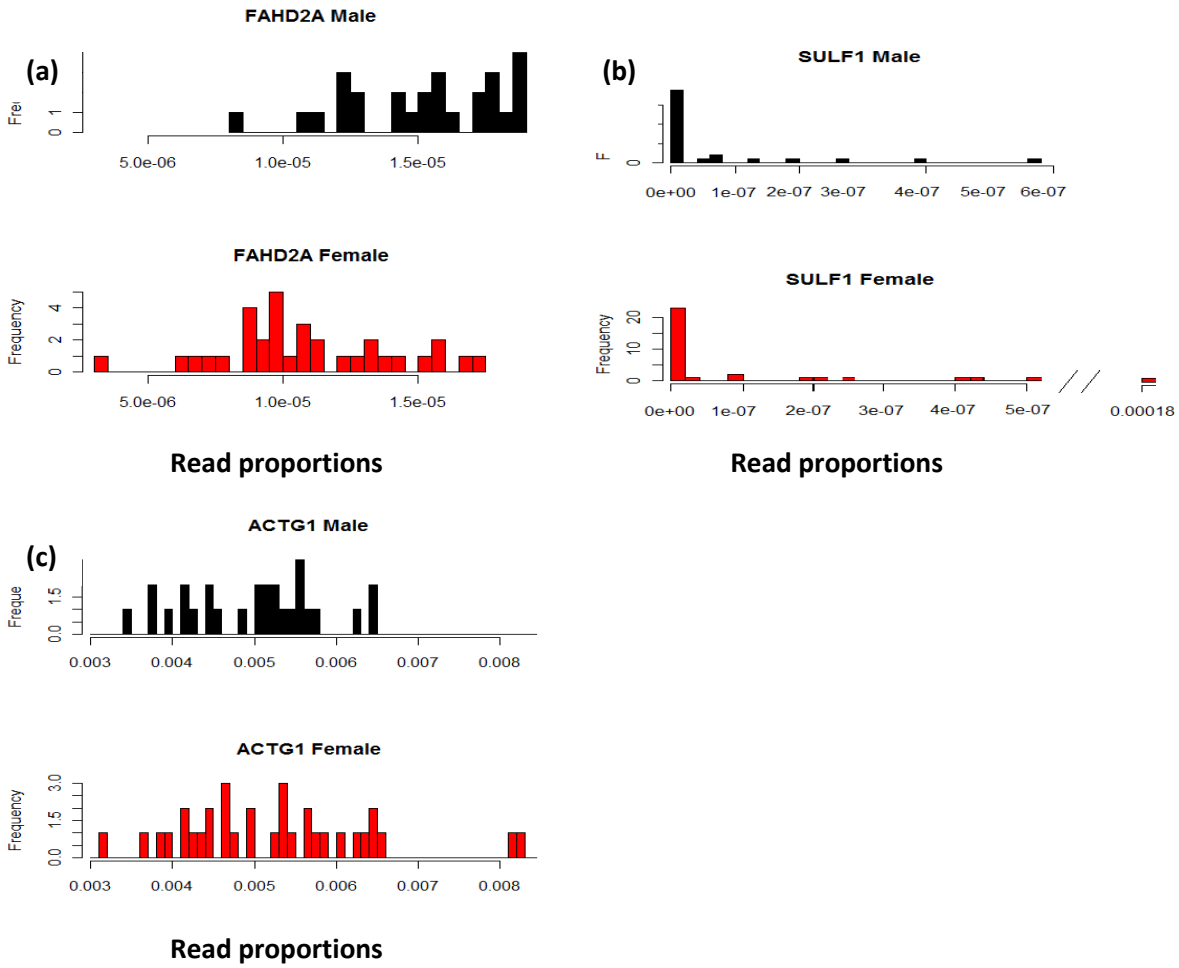

Supplementary Figure 8. Histograms of transcript counts, expressed as read counts normalized by library size (i.e. *read proportions*) for some of the most significantly differentially expressed *autosomal* genes in males vs. females . We performed the four methods for the entire set of 27 vs. 33 CEU samples (free, edgeR "trend", baySeq, and DESeq). A review of the top-ranked autosomal genes suggests that free model (panel a) identifies genes that are truly differentially expressed (the top-ranked autosomal gene *FAHD2A* shown). In contrast, the *SULF1* gene ranked third by edgeR trend and fifth by DESeq (panel b) appears to be influenced by the single outlying value in females. BaySeq (panel c) identifies *ACTG1* as differentially expressed, while simple inspection suggests that any differences are not striking (although the two highest values of *ACTG1* both appear in females). Although the reasons for the baySeq conclusion are not clear, we note that *ACTG1* is the 4th highest expressed gene among all genes in the CEU dataset, suggesting possible sensitivity of the results for genes of high expression.