| ESTs	f	Model	$p < 0.01$	$p < 0.001$	$p < 10^{-4}$	Filtered	$p < 10^{-4}$	Lowest p
−	−	perfect test	248	25	2.5	−	50	$2.0 \cdot 10^{-6}$
12–250	$5.5 \cdot 10^{-6}$–$1.1 \cdot 10^{-4}$	$\sigma = 0$	2	0	0	2,661	0	$8.8 \cdot 10^{-4}$
20	$9.1 \cdot 10^{-6}$	$\sigma = 1$	26	2	0	−	−	−
50	$2.3 \cdot 10^{-5}$	$\sigma = 1$	131	5	0	−	−	−
500	$2.3 \cdot 10^{-4}$	$\sigma = 1$	270	27	1	−	−	−
12–250	$5.5 \cdot 10^{-6}$–$1.1 \cdot 10^{-4}$	$\sigma = 1$–2.5	251	22	1	649	41	$1.1 \cdot 10^{-6}$
100	$4.6 \cdot 10^{-5}$	$\sigma = 3$	105	10	0	1,350	69	$6.1 \cdot 10^{-7}$
500	$2.3 \cdot 10^{-4}$	$\sigma = 3$	308	41	2	−	−	−
100	$4.6 \cdot 10^{-5}$	$\sigma = 4$	102	11	1	−	−	−
500	$2.3 \cdot 10^{-4}$	$\sigma = 4$	102	12	0	−	−	−
100	$4.6 \cdot 10^{-5}$	half-on	320	35	3	−	−	−
500	$2.3 \cdot 10^{-4}$	half-on	272	22	3	−	−	−

Supplementary Table 1. The number of p-values below various thresholds for LNP on 24,750 simulated pairs of uncorrelated genes, generated with various models of variation and average expression levels. Because the log-normal distribution has very high variance for high σ, the actual #ESTs for simulated genes can vary widely. The "perfect test" row shows the expected values for a perfect test. The "filtered" entries were generated from 10^6 simulated pairs by using the apparent correlation filter (but not the self-p filter).
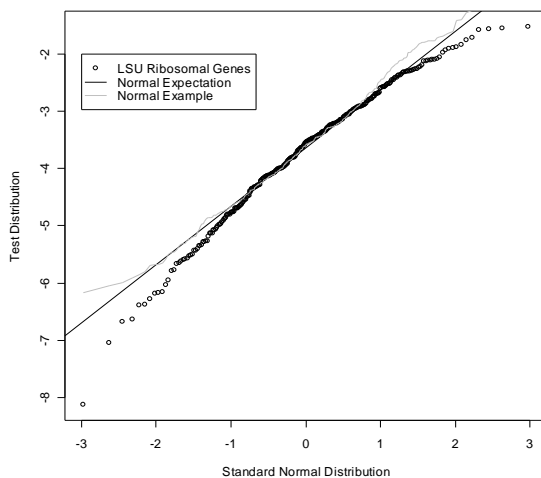
Supplementary Figure 1: Normal quantile-quantile plot for the distribution of the total log-frequencies $\log \sum f / (1 - \sum f)$ of large ribosome subunit proteins in human dbEST. For comparison, the grey line shows a simulated normal variable with the same mean, standard deviation and number of values, and the straight line shows the expectation for a normal distribution. We used all 46 cytoplasmic large subunit ribosomal protein genes in UniGene: *RPL3-15, RPL17-19, RPL21-23, RPL26-39, RPL41, RPL10A, RPL13A, RPL18A, RPL23A, RPL27A, RPL35-37A,* and *RPLP0-2*. We used a group of highly expressed proteins rather than single genes so that all of the counts are large (in this case, the minimum total is 2) and likely to give good estimates of the actual frequency (the %error in the frequency estimate goes as the square root of the count). Although the distribution is quite close to log-normal, the left-skew in the tails is significant ($p = 8.4 \cdot 10^{-5}$, Shapiro-Wilk normality test on $\log \sum f / (1 - \sum f)$).
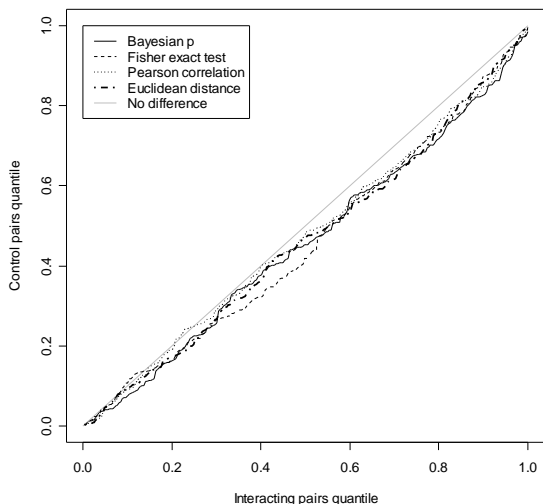
Supplementary Figure 2: Quantile-quantile plot showing the difference in the distribution of p-values for 269 interacting pairs of human genes from the MINT database, compared to 32,379 randomized control pairs. The x-axis shows the fraction of interacting pairs smaller than a given p-value, and the y-axis shows the fraction of control pairs below that same p-value. The straight grey line shows the result for measures that cannot distinguish between the sets. To show the difference for the full range of the distribution, we replaced LNP p-values with $1 + 1/p$ when $\hat{r} < 0$.
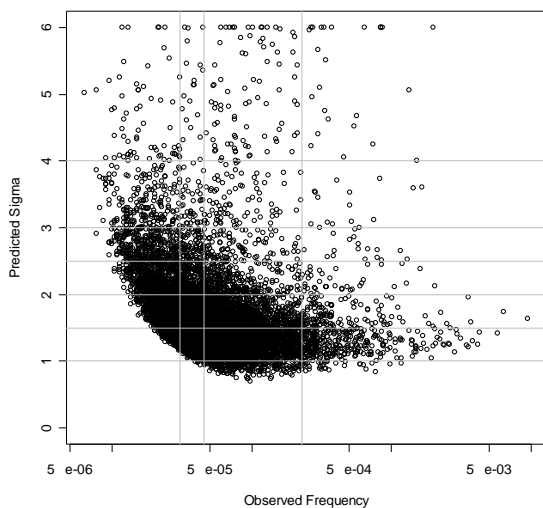
Supplementary Figure 3: Predicted σ vs. observed mean frequency for the 9,082 human genes used for the LNP all-pairs computation. For comparison, the vertical and horizontal lines show some of the parameter settings used in simulations. We ran most simulations with higher than average expression levels because most of the results come from those genes (the average gene has insufficient data). Genes with low predicted σ ($\ll 1$) failed to show significant coexpression with themselves (see Methods) and were excluded. Similarly, tens of thousands of genes expressed at low levels were excluded due to high self-p or less than 5 ESTs in our dataset. The high-σ genes (511 with $\sigma \geq 3$) often show strong on-off behavior (ESTs clumped into relatively few libraries) with significant variation in the "on" level, which is consistent with the log-normal model but makes it hard to predict σ. We constrained the predicted σ to be at most 6 because such large average fluctuations in a gene's expression seemed implausible ($e^6 \approx 400$, $e^{12} \approx 160,000$).