

File S2

MCMC analysis

Prior specification: For the fixed effects, we used independent normally distributed priors with a mean of 0 and a large variance of 10^{10} . This is essentially uninformative between 10^{-4} and 10^4 .

For the random effects, we used independent inverse Wishart (IW) distributed priors. The IW distribution has a probability density of zero for values up to zero, a large spike above zero and a long, flat tail for values above the spike. Reducing the value of either parameter of the IW distribution (the variance at the limit V and the degree of belief ν) causes the spike to move left, towards zero. Therefore, IW distributed priors are weakly informative for a variance component provided that its posterior distribution does not have high probability density at the spike. Since reducing V and ν also causes model fitting to slow down, we searched for relatively high values of V and ν that gave estimates that were not affected by the spike of the IW distribution. For the single and multiple GE data we chose $V = 1$ and $\nu = 0.002$, and for the somatic fitness data we chose $V = 0.1$ and $\nu = 0.0002$.

Autocorrelation: We began by running a Markov chain for 1.5×10^5 iterations. We then analyzed the autocorrelation function between consecutive parameter values of the Markov chain at successive iterations for the last 10^5 iterations to determine how rapidly independence was achieved. For the single and multiple GE data we chose to sample every 50 iterations from the posterior distribution, and for the somatic fitness data we chose to sample every 200 iterations.

Convergence: We then ran three parallel Markov chains and used the method of Gelman and Rubin (1992) (implemented in R through the coda 0.14-4 package, Plummer et al. 2006) to determine how quickly convergence was achieved. For the single and multiple GE data convergence was achieved within $\sim 10^5$ iterations, and for the somatic fitness data convergence was achieved within $\sim 10^6$ iterations.

Final MCMC analyses: For the single and multiple GE data we allowed one Markov chain to run for a burn-in period of 10^6 iterations after which we ran 10^7 iterations and sampled from the posterior distribution every 50 iterations, resulting in 2×10^5 stored values.

For the somatic fitness data we allowed one Markov chain to run for a burn-in period of 4×10^6 iterations after which we ran 4×10^7 iterations and sampled from the posterior distribution every 200 iterations, resulting in 2×10^5 stored values.

Literature Cited

Gelman, A., and D. B. Rubin, 1992 Inference from iterative simulation using multiple sequences. *Stat. Sci.* **7**: 457–511.
Plummer, M., N. Best, K. Cowles and K. Vines, 2006 CODA: convergence diagnosis and output analysis for MCMC. *R News* **6**: 7–11.