


The same outputs, viewed through more noise

$x_t = x_0 + \sigma_t \varepsilon$ · Each panel is a marginal, not a sampler step


A Low noise

$\sigma_t = 0.12$




B Medium noise

$\sigma_t = 0.65$



C High noise

$\sigma_t = 1.40$



— Target marginal - - - Student marginal

Exact Gaussian convolutions · shared axis scale · no optimization between panels