


Gaussian known-score study: confidence-interval coverage


● Python (normal) ■ R qqcomp (normal) ▲ QR contrast (t)

..... Normal interval target: 94.701% (n=100), 94.944% (n=500) --- t interval target: 95%



Successful / failed fits (out of 2,000)

Python	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0
R	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0
QR t	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0



Successful / failed fits (out of 2,000)

Python	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0
R	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0
QR t	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0	2000 / 0

Points show saved coverage; bars are 95% Wilson intervals for uncertainty in observed coverage.

There are 2,000 independent datasets per condition; all three methods use the same datasets.

Coverage is conditional on successful fits. Counts below each panel preserve failures; undefined means no successful fits.

Analytic targets assume a correct normal, constant-variance model and full-rank design. ρ is residual correlation before scoring.

The shared axis includes every interval and target. coverage.csv contains all values and denominators.

Execution status: complete numerical agreement. Statistical interpretation is given in the accompanying results report.