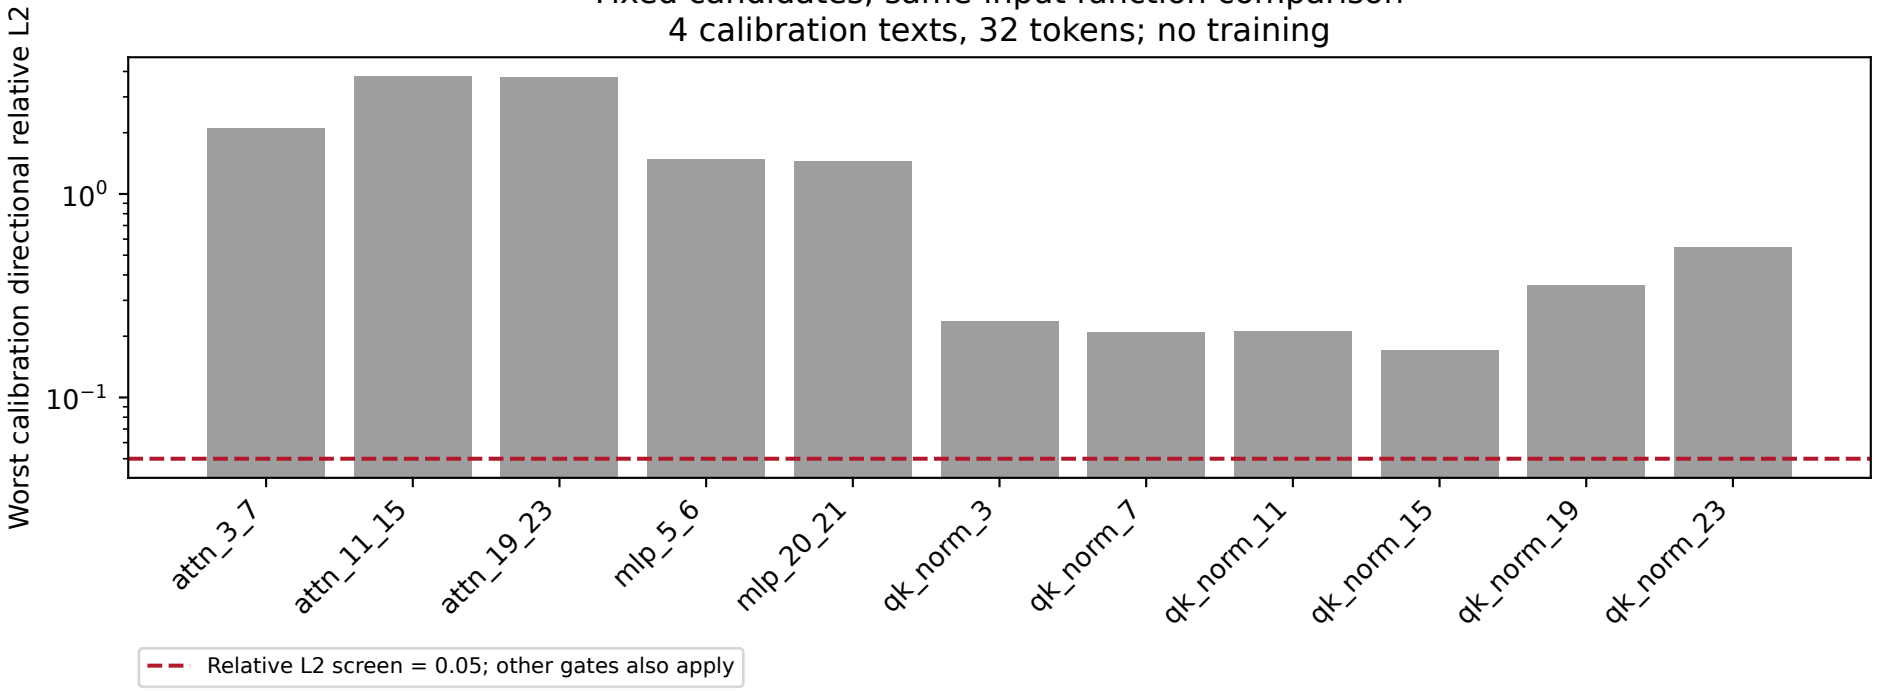

Fixed candidates; same-input function comparison
4 calibration texts, 32 tokens; no training
Worst calibration directional relative L2
10^{0}
10^{-1}
attn_3_7
attn_11_15
attn_19_23
mlp_5_6
mlp_20_21
qk_norm_3
qk_norm_7
qk_norm_11
qk_norm_15
qk_norm_19
qk_norm_23
Relative L2 screen = 0.05; other gates also apply