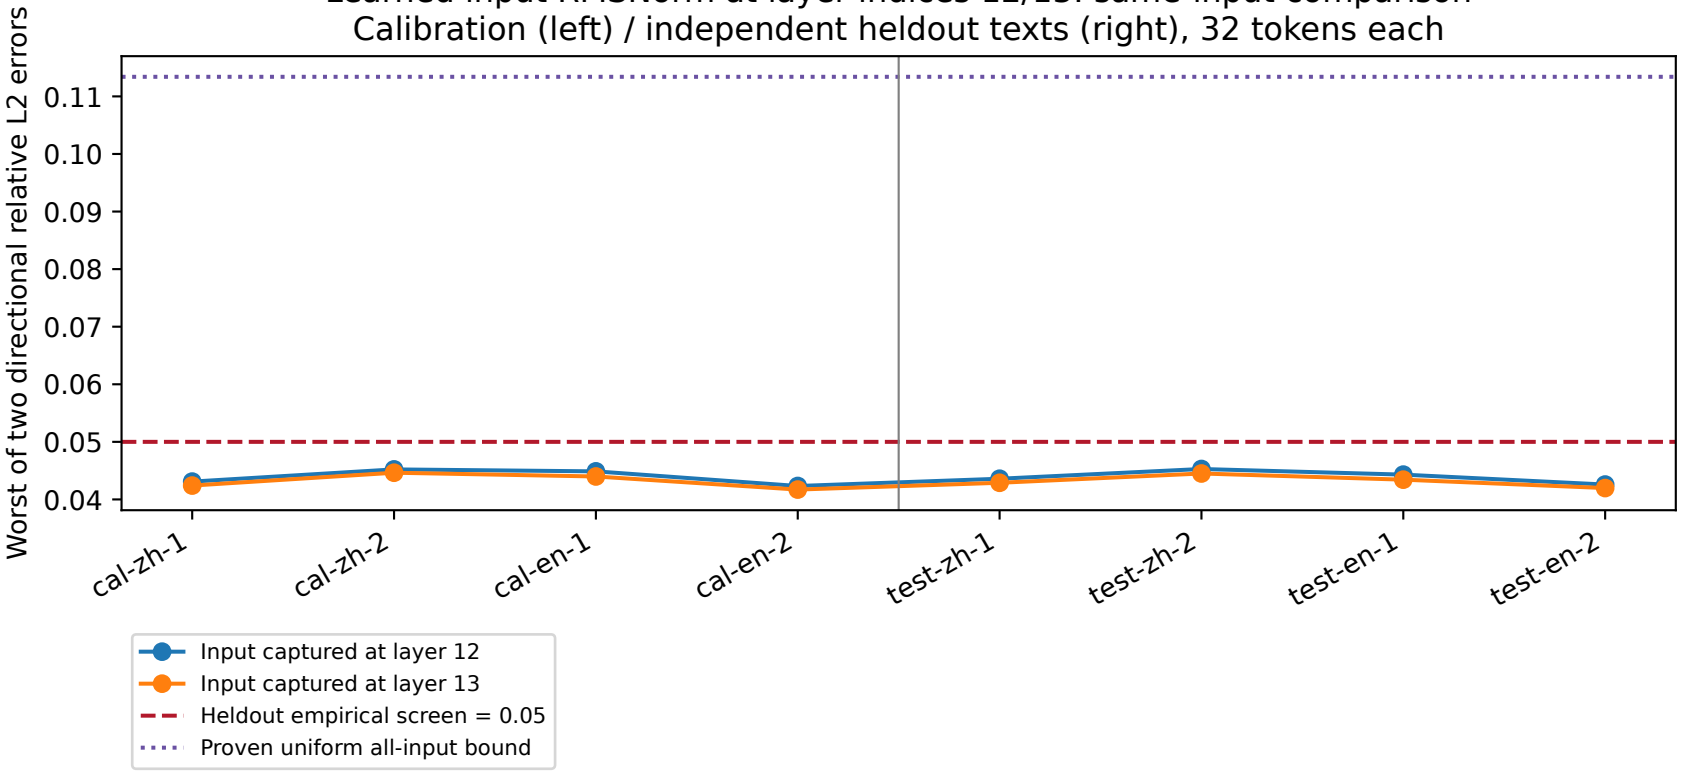

Learned input RMSNorm at layer indices 12/13: same-input comparison
Calibration (left) / independent heldout texts (right), 32 tokens each
Worst of two directional relative L2 errors
0.11
0.10
0.09
0.08
0.07
0.06
0.05
0.04
cal-zh-1
cal-zh-2
cal-en-1
cal-en-2
test-zh-1
test-zh-2
test-en-1
test-en-2
Input captured at layer 12
Input captured at layer 13
Heldout empirical screen = 0.05
Proven uniform all-input bound