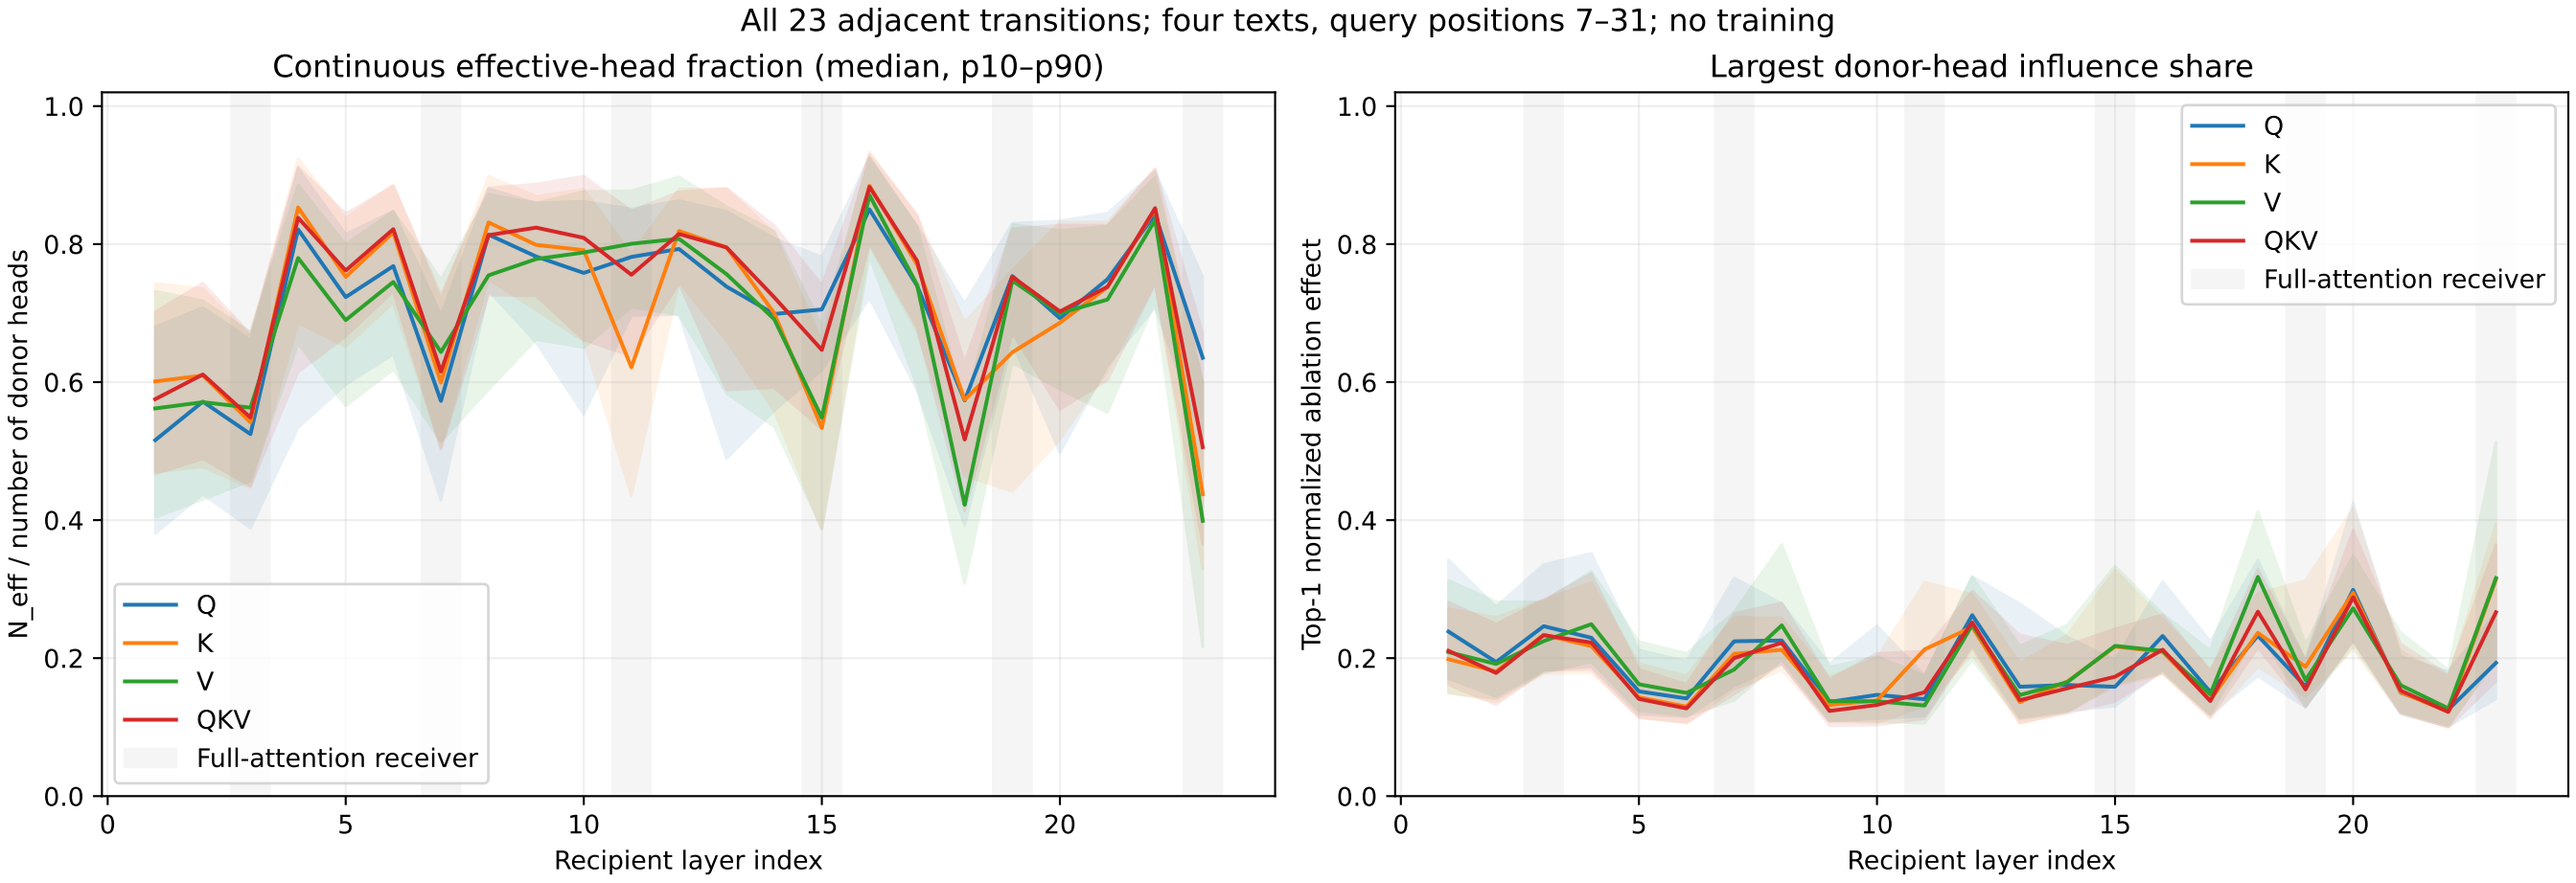

All 23 adjacent transitions; four texts, query positions 7–31; no training
Continuous effective-head fraction (median, p10–p90)
N_eff / number of donor heads
Recipient layer index
Largest donor-head influence share
Top-1 normalized ablation effect
Recipient layer index
Q
K
V
QKV
Full-attention receiver