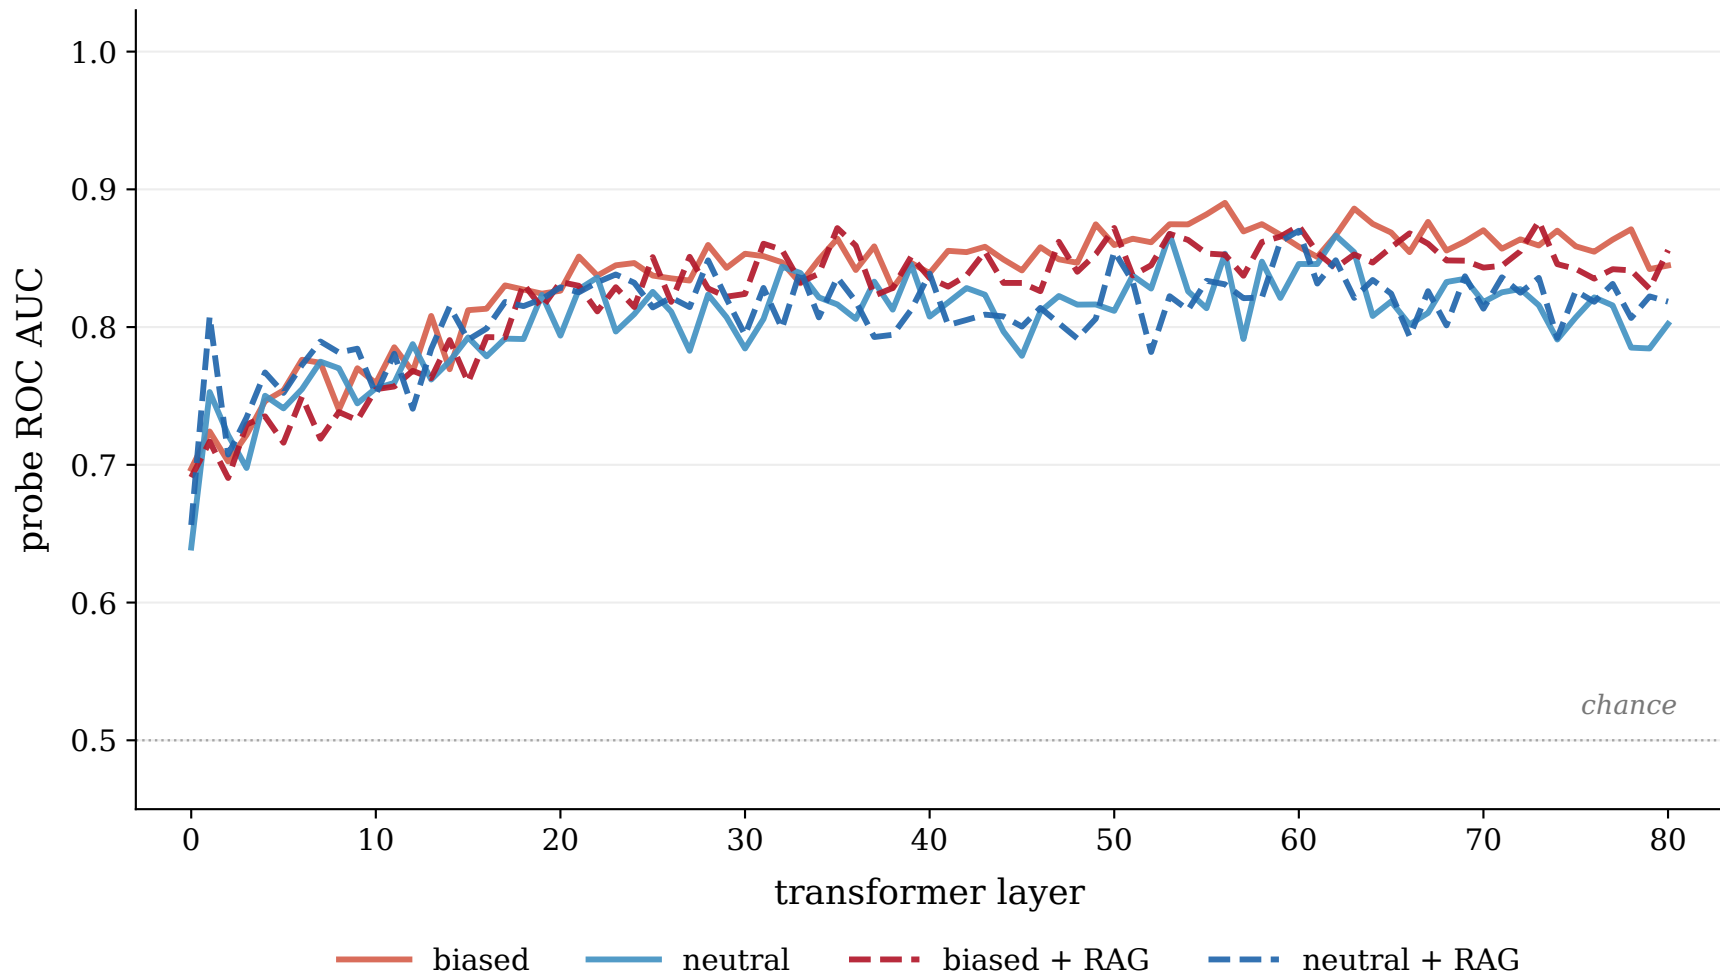


Admission probe — per prompt variant (mean pooling)



Same probe as the pooled curve, split out by the four prompt variants.
Solid = no-RAG prompt; dashed = +RAG. Red = biased system prompt; blue = neutral.