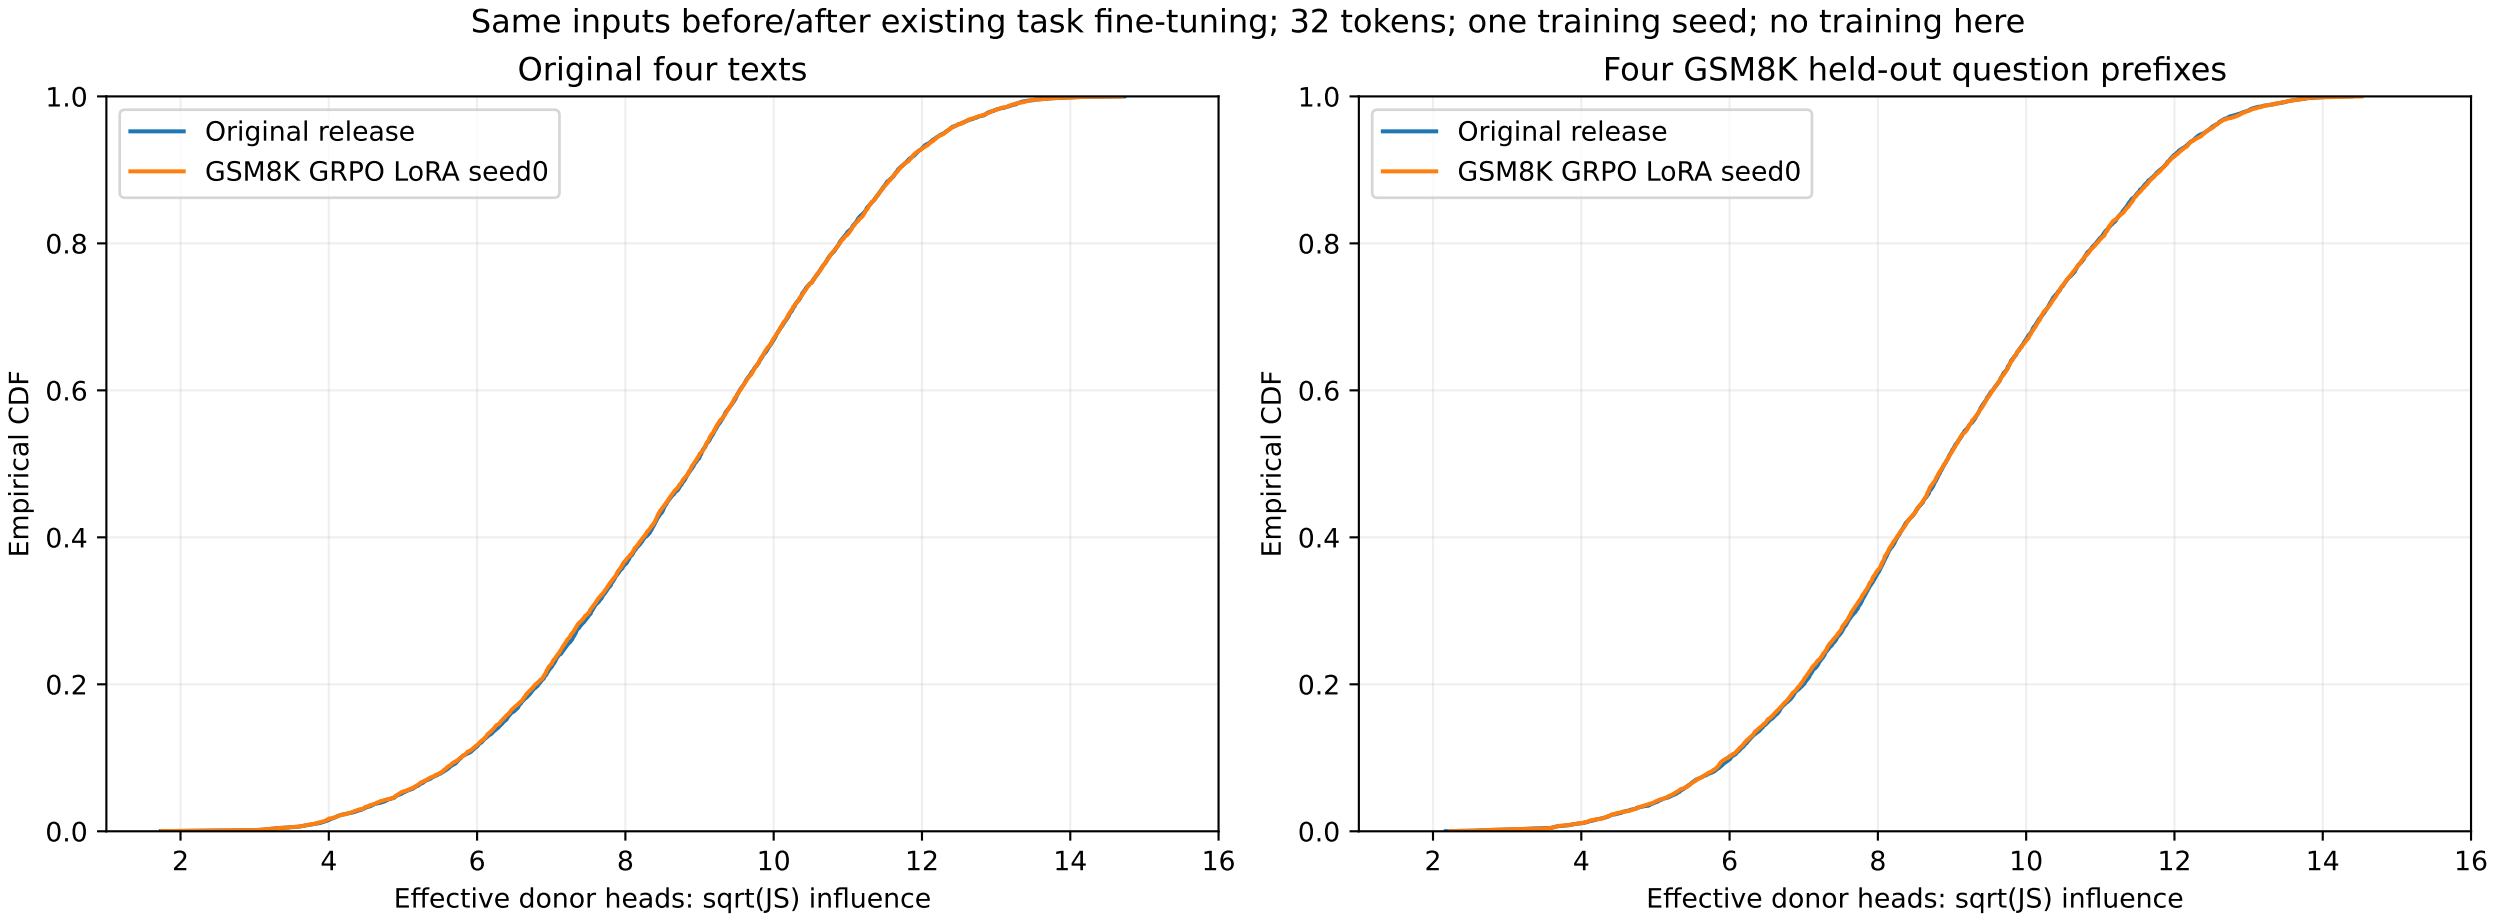

Same inputs before/after existing task fine-tuning; 32 tokens; one training seed; no training here
Original four texts
Four GSM8K held-out question prefixes
Original release
GSM8K GRPO LoRA seed0
Empirical CDF
Effective donor heads: sqrt(JS) influence
0.0
0.2
0.4
0.6
0.8
1.0
2
4
6
8
10
12
14
16