


vGROUT floor→ceiling: routeV (no oracle, gradient-level) vs vanilla GRPO (test n=119, seed 43, 60-step fast)



Our arms only, seed 43, 60-step fast (unconverged surrogate). hack suppressed = (vanilla\_hack - arm\_hack)/vanilla\_hack; solve gained = (arm\_solve - base)/(ceiling - base). AriaHW 2025 monitor numbers are cross-scale/regime and live in the transcribed Fig-5 table in docs/papers/2025\_lw\_ariahw\_\*.md, not on this axis.