arXiv:2611.20419v1 [cs.LG] 16 Nov 2026

Reasoning Again with a Quarter of the Experts: Retry First, Then Train

Anik Jha
Independent Researcher
anik.k.jha@gmail.com

Abstract

Removing three quarters of the routed experts from a 122B-parameter reasoning mixture-of-experts model makes it fit in a 12 GiB file, and breaks it in a specific way: on LiveCodeBench it enters verbatim reasoning loops that run to the token limit on 44 of 120 development problems, where the unpruned model at the same 2-bit quantization loops on none. We study how much of that loss can be recovered without changing the kept experts' quantized bytes, in three steps of increasing cost. On development problems, a sampling penalty recovers part of it. Stopping a reply when it starts to loop and drawing again recovers more on top of it, with no weight change. Our method, B, adds rank-8 trained corrections to the kept experts, loaded as an adapter, and restores two experts per layer; it is fitted layer by layer to the unpruned model's expert block on the pruned model's own activations. In a pre-registered comparison on 130 easy and medium LiveCodeBench problems that no part of the study had touched, with both arms given the same sampling penalty and the same stop-and-redraw policy, B solves 116 problems against 104 (+9.2 pp, 95% CI [+3.1, +15.4]; 15 problems gained, 3 lost). A training-free weight edit aimed at the loops, tested the same way earlier, did not help. B removes more wrong answers than loops, and costs 1.7× the wall time. On two older benchmarks used as regression checks, B stays inside a pre-set 2 pp tolerance but shows no gain (HumanEval+ +0.6 pp, MBPP+ −1.9 pp). The evidence is one model, one quantization and one benchmark family; we state what it does and does not show.

1 Introduction

Expert pruning removes whole experts from a mixture-of-experts (MoE) language model so that it fits in less memory. At half the experts, one-shot pruning can be close to lossless on code generation (Lasby et al., 2025). Deeper pruning is where memory tiers change, and where models break. This paper is about one such break and what repairs it.

We prune Ornith-1.5-122B, an open reasoning MoE with 256 routed experts per layer, to 64 experts per layer with the REAP criterion (Lasby et al., 2025), by copying the kept experts' bytes out of a public 2-bit quantized file. The pruned file is 11.62 GiB. We call this model C. On LiveCodeBench (Jain et al., 2024) with a reasoning budget, C does not merely get more answers wrong. It stops finishing: its reasoning falls into verbatim repetition that runs until the token limit. Section 3 shows that, for this model, pruning causes this and 2-bit quantization does not.

That observation frames the question. If the dominant failure is a loop, the cheapest repairs are at decoding time and need no training at all. A repair that changes weights has to beat those, not just beat the broken model. We therefore compare, in order of cost: a sampling penalty (S); stopping a looping reply and drawing again (retry); a training-free weight edit against the loop; and trained corrections to the expert weights (A, and B, which also restores two experts per layer).

Preprint.

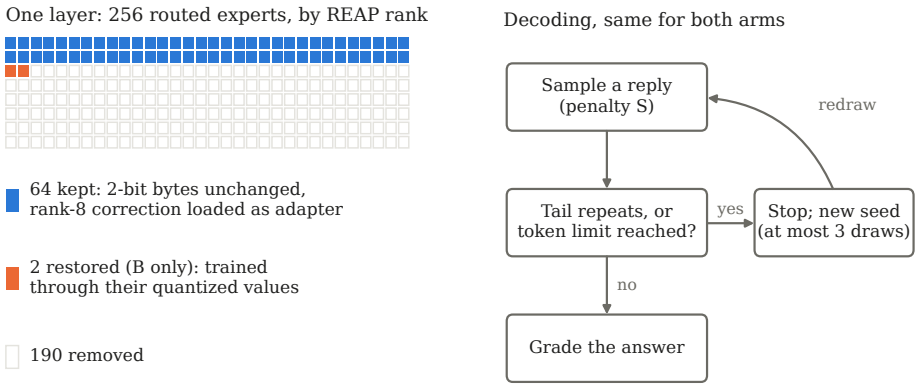


Figure 1: What is compared. Left: one layer's routed experts. The pruned model C keeps 64 of 256 with their quantized bytes unchanged; A and B add a rank-8 correction to each kept expert as a load-time adapter, and B also restores two experts. Right: the decoding policy given to both arms in the confirmation.

None of the ingredients is new by itself. Reconstruction-based updates of the retained experts already improve LiveCodeBench after expert pruning down to 37.5% retention (Markovic-Voronov et al., 2026); low-rank matrices can compensate a compressed model while its weights stay unchanged (Liu et al., 2024); loops after aggressive compression of reasoning models, and decoding-time or training-time rescues for them, have been reported for quantization (Alimaskina et al., 2026; Chen et al., 2026). What this paper adds is a harder regime (25% retention on code, below the 37.5% where that prior work stops, at 2 bits), a controlled attribution of the loops, a weight-free baseline that a weight method must beat, and a confirmation whose rules were fixed before the test problems were touched.

Our contributions:

- **A controlled attribution of the loops to pruning.** With the kept experts byte-identical between the pruned and unpruned 2-bit files, the pruned model hits the token limit in a verbatim loop on 44 of 120 problems and the unpruned one on 0 (Section 3).
- **Stop-and-redraw as a strong weight-free baseline.** On untouched problems, redrawing lifts the penalised pruned model from 89 to 104 of 130 solved (Section 6).
- **A weight recovery method that leaves the quantized model untouched.** B builds on the layer-wise expert weight update of Markovic-Voronov et al. (2026): it is fitted layer by layer against the unpruned model's expert block, but as low-rank corrections plus two restored experts per layer, and is deployed as the pruned file plus two expert slices per tensor plus a load-time adapter (Section 5).
- **A pre-registered confirmation.** The finalist, the decoding policy, the problems and the decision rule were frozen before any request on the confirmation or final problems. B leads by +9.2 pp with a 95% interval of [+3.1, +15.4] (Section 6).
- **Two negative results.** A training-free weight edit against the loops did not transfer from development to test problems, and the sampling penalty alone was not separable from no penalty on that test (Section 4).

We are explicit about scope. Everything here is one model at one pruning ratio and one quantization, evaluated on one benchmark family, with one fit of each method. We do not separate the contribution of the low-rank corrections from that of the restored experts on untouched problems, and we do not explain why B helps. Section 7 lists these limits in full.

2 Setting

Model and pruning. Ornith-1.5-122B[1] has 48 layers, hidden width 3072, and in each layer 256 routed experts of width 1024, of which the router selects 8 per token, plus one shared expert. We score experts with REAP (Lasby et al., 2025) on our calibration data and keep the top 64 per layer. The pruned model C is built by copying the kept experts' slices and router rows out of the public 2-bit file (UD-IQ2_XXS), so no kept weight is requantized. The unpruned teacher used for fitting is the 4-bit file (UD-Q4_K_XL) of the same model.

Serving. All generation uses llama.cpp's server[2] on one machine with 119 GB of unified memory, with 32 parallel slots, temperature 0.6, top-p 0.95, top-k 20 and a limit of 32,768 generated tokens unless stated. Peak memory in the final comparison was 81.6 GiB.

Benchmark and problem sets. We use LiveCodeBench (Jain et al., 2024) with its official prompt and grader, restricted to easy and medium problems after a pilot showed hard problems would exceed our compute. Three disjoint sets matter:

- *Development*: 120 problems, the earlier half by contest date of releases v4 to v6. Every choice in this paper was made here.
- *Old test half*: 124 problems, the later half of v4 to v6. It was used once to test the weight edit of Section 4; replies of C on it were later used as fitting data for A and B, so it is not a test set for them.
- *Fresh panels*: problems from releases v1 to v3, set aside before any generation on them: 64 for confirmation and 130 for the final comparison, in disjoint contests. We excluded problems sharing a title or contest with v4 to v6 and checked statement overlap against every local fitting, calibration and generation file.

A problem counts as solved when the code extracted from the answer passes all hidden tests.

Statistics. Comparisons are paired by problem. We report the difference in solved fraction in percentage points (pp) with a 95% bootstrap interval over problems (10,000 resamples; 20,000 in Section 3), and the exact McNemar p on discordant pairs as support. Where a problem is sampled twice, its two outcomes are averaged before resampling, so two seeds never count as two problems.

3 What breaks: reasoning loops

Table 1 compares C with the unpruned model at the same 2-bit quantization and at 4 bits, on the development problems with a 16,384-token limit. The kept experts of C are byte-identical to the corresponding experts of the 2-bit unpruned file, so the first two rows differ only in the 192 removed experts per layer.

Table 1: Development problems (n = 120), 16K token limit. "Limit, looping" counts replies that reached the limit inside a verbatim token loop.

Model	Solved	Hit the limit	Limit, looping
Pruned to 25%, 2-bit (C)	71	44	44
Unpruned, 2-bit	107	11	0
Unpruned, 4-bit	110	8	0

Pruning raises the share of problems that end in a loop at the limit by +36.7 pp [+28.3, +45.8], and accounts for a +30.0 pp drop in solved problems [+21.7, +38.3]. Going from 4 bits to 2 bits without pruning adds no loops on these problems. This scopes two earlier observations. Alimaskina et al. (2026) report that 2-bit quantization of dense 8B and 32B reasoning models produces repetitive

[1] `https://huggingface.co/ornith-ai/Ornith-1.5-122B`; quantized files from `https://huggingface.co/unsloth/Ornith-1.5-122B-GGUF`.

[2] `https://github.com/ggml-org/llama.cpp`, built from commit `36cbf415` with one local patch to chat-output parsing that is off by default.

loops and budget exhaustion; at this scale and with this quantizer we see none from quantization alone. Song and Zheng (2026) report more truncation and, under severe pruning, more repetition for expert-pruned reasoning MoEs on LiveCodeBench at an 8K limit. They vary temperature, output caps and repetition penalties, but have no matched quantization control and no recovery comparison of the kind studied here; our comparison isolates pruning and the rest of this paper is about remedies. Practitioners have also noted the effect: one model card for a REAP-pruned model reports roughly twice the loop rate of its unpruned source and recovers it with sampler settings.[3] Concurrent work reports runaway generation among the failures of REAP-pruned models on production tasks and finds that the rate of repetition loops tracks quality (Rasymas and Kvedaras, 2026), and a benchmark of stacked compression finds expert pruning, not quantization, the dominant source of degradation (Benazir et al., 2026).

No reply of C that reached the limit solved its problem (0 of 45 on the old test half at 32K). A loop is a failure.

4 Remedies that do not train weights

A sampling penalty (S). The model's vendor recommends no presence penalty for precise coding. We found that a presence penalty of 1.5 over a 20,480-token window helps the pruned model: on development problems at 16K it solves 86 against 71 without it, and loops at the limit drop from 44 to 25. The window matters because loop periods are long: the server's default penalty window is 64 tokens, and many repeated blocks are longer. We fix this setting as S for the rest of the paper.

A training-free weight edit. Motivated by work showing that single-neuron edits can remove some repetition loops (Lazaridis et al., 2026), we tested an edit, built on the development problems, that projects out, from four attention heads, the directions those heads write most on looping tokens. The edit was designed and chosen by a large language model (Claude Opus 5.5) acting as editor, with access to the pruned model's weights and to development-problem activations and results only. It is loaded as an adapter, so the quantized file is unchanged. On development problems it solved 81 against 71 for C. We then froze the edit and tested it once on the old test half at 32K (Table 2).

Table 2: Old test half (n = 124), 32K limit, one frozen test.

Arm	Solved	Hit the limit
C	74	45
C + weight edit	68	47
C + S	80	32

The edit did not transfer: -4.8 pp against C [$-12.9, +3.2$]. The penalty S solved more than C here too, but the difference, +4.8 pp [$-2.4, +12.1$], is not significant on this test. We keep S as the decoding setting because it is free and its direction was consistent, not because this test established it.

Stop on loop, then redraw (retry). Sampling is random, so a problem that loops on one draw may not loop on the next. A client-side rule stops a reply as soon as its tail repeats: for some period $P \leq 6000$ characters, the last $\max(5P, 200)$ characters equal the same-length stretch ending P characters earlier, checked every 256 characters. A reply stopped by the rule, or cut at the token limit, is drawn again with a new seed, for at most three draws per problem. The rule was chosen on saved development replies by a criterion written before its grid was run: it stops 72 of 74 replies that had run to the limit, at a median of 6,003 tokens, and would have stopped 2 of 360 replies that went on to finish and pass. Related work detects repetition in reasoning traces from hidden states and continues after a cut (Xie et al., 2025), predicts loops early from hidden states (Duan et al., 2026), or, for 2-bit models, detects repetitive traces and then commits to an earlier answer or falls back to a full-precision model (Alimaskina et al., 2026). Our rule needs only the text stream and resamples from the same model.

Retry changes no weights and costs only the tokens of the abandoned draws. It is the baseline a weight method must beat.

[3] https://huggingface.co/0xSero/GLM-5.2-504B, read 2026-11-04.

5 Weight recovery

Both methods leave every quantized byte of C in place. They add a low-rank correction to each kept expert, applied at load time as a LoRA-style adapter (Hu et al., 2021); B also adds experts.

Method A: corrections to the kept experts. A adapts the expert weight update (EWU) of Markovic-Voronov et al. (2026), which refits the retained experts' weights layer by layer against the unpruned layer's output, on inputs from the already-updated pruned model, with a regulariser toward the pretrained weights. We keep that structure and change the parameterisation: the update is low-rank and stays outside the quantized file, in the spirit of compensating a compressed model with low-rank matrices while its weights stay fixed (Liu et al., 2024), here trained per expert instead of computed in closed form. For each layer and each kept expert e, each of its three matrices $W \in \{W_e^{\text{gate}}, W_e^{\text{up}}, W_e^{\text{down}}\}$ becomes $W + B_W A_W$ with rank 8. Layers are fitted in order. For layer ℓ, let x_t be the input to the expert block at position t when the pruned model, with the corrections of layers $0..\ell-1$ already applied, reads the fitting documents. Let y_t be what the unpruned model's routed experts produce on the same input x_t, f_t the corrected pruned experts' output under the pruned model's own routing, and f_t^0 their output before correction. We minimise

$$\mathcal{L}_\ell \;=\; \frac{\sum_t w_t \|y_t - f_t\|^2 \;+\; \mu \sum_t w_t \|f_t - f_t^0\|^2}{\sum_t w_t \|y_t - f_t^0\|^2}, \tag{1}$$

with $\mu = 0.01$ and weights w_t that give every document the same total weight. The denominator is the layer's starting error, so the loss starts at 1. The router, attention, shared expert and embeddings stay frozen. B_W starts at zero, so training starts exactly at C. We use Adam with learning rate 10^{-4}, gradient clipping at 1, and 25 passes over the fitting documents in batches of 64 documents, and keep the checkpoint (every 5 passes) with the lowest loss on held-out documents. No coding score is consulted during fitting.

Fitting data. The documents are 1,018 calibration documents (256 more held out), truncated to 2,048 tokens with 256 sampled positions each, plus C's own long replies on the old test half (124 replies, each truncated to 12,288 tokens with 256 sampled positions, 24 of them held out for validation), so that the fit sees the hidden states the model is in when it reasons at length; reconstructing on a model's own chain-of-thought traces was introduced for pruning dense reasoning models by Lucas et al. (2025). Neither the development problems nor the fresh panels are used.

Method B: A plus two restored experts per layer. B starts a separate fit from C with the next two experts of each layer's REAP ranking restored from the unpruned 2-bit file, together with their original router rows, so each layer has 66 experts and the router can select them in place of kept ones. The restored experts train in full precision with a straight-through estimator to their deployed quantized values. The quantized forward values are refreshed once per pass over the data, not per step, and every validation score used for checkpoint selection quantizes afresh; the 64 kept experts receive rank-8 corrections as in A. Everything else follows A, including the data, order and token budget. Restoring two experts raises routed retention from 25% to 25.8% and the file from 11.62 to 11.88 GiB; the adapter adds 1.16 GiB.

What the fits achieve. Across the 48 layers, the held-out reconstruction error falls by a median of 21.8% for A (range 14.2% to 76.0%) and 22.8% for B (14.5% to 75.8%). These are modest reductions: most of what the removed experts computed is not recovered. Figure 2 shows that the reduction is largest in the first and last layers and that the two fits are nearly indistinguishable by this measure. Fitting took 5.2 and 5.7 hours summed over layers on the one machine, not counting the passes that compute each layer's inputs.

Deployment checks. Before any coding evaluation we verified that the adapter applied by the server reproduces the fitted corrections computed in our own forward pass, that an adapter at scale zero leaves outputs byte-identical, and that every byte of C is unchanged inside the 66-expert file. Parity tolerances that proved tighter than the server's own numerical differences were amended in dated notes before any coding result of the affected method existed; Appendix A lists them.

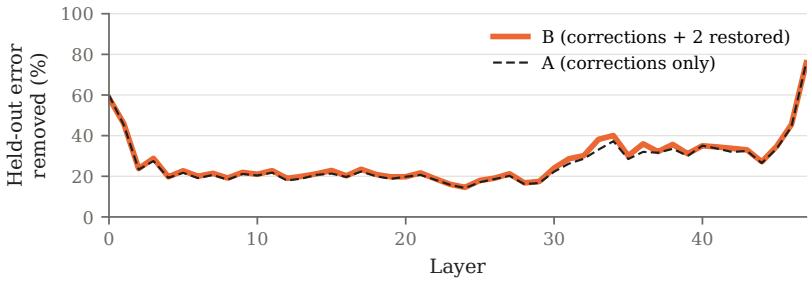


Figure 2: Share of each layer's held-out reconstruction error removed by the fit.

Relation to prior recovery methods. The closest method is OptEP's EWU (Markovic-Voronov et al., 2026), described above. It updates the retained experts' full weight matrices and is evaluated on LiveCodeBench at 37.5% to 75% retention; it does not restore experts, and it reports no quantized-deployment experiment. Our changes are the low-rank, load-time form that leaves 2-bit weights untouched, the restored experts trained through their quantized values, fitting data that includes the pruned model's own long reasoning, and the evaluation against a retry baseline on untouched problems. Earlier, Muzio et al. (2024) recovered accuracy after expert pruning with regularisation-based fine-tuning, and Bercovich et al. (2026) recover a heterogeneously expert-pruned reasoning MoE with distillation followed by reinforcement learning. Some recovery work adapts components other than the experts: LoRA on attention with experts and router frozen (Goel et al., 2026), or distilling only the router (Hyeon and Do, 2026a). Xie et al. (2024) recover an unstructured-pruned MoE with expert-wise distillation from the unpruned model. Hao et al. (2026) replace redundant experts with parameter-efficient modules and train them. Lu et al. (2024) and Zhou et al. (2025) select or recombine experts without training. We do not compare against any of these methods experimentally, including dense EWU.

6 Results

6.1 Development: choosing a finalist

On the development problems with plain decoding (one draw, 32K limit), S alone solves 84 of 120, A+S 88 and B+S 93. B+S is ahead of S by +7.5 pp [+0.8, +15.0], and A+S by +3.3 pp [−5.0, +10.8]. By a selection rule fixed in advance (a candidate must be at least 2 pp above S; B is chosen over A only if at least 2 pp above it), B became the development winner. This is a selection, not a significance claim, and A was not carried further.

A post hoc look at these runs showed that a plain rerun of the same arm changes the outcome on many problems, and that a rerun rescues limit-hitting problems about as often as B does. That is what prompted the retry baseline: before crediting B with a recovery, both arms should be allowed to redraw. With retry on both arms, on the same development problems, B+S solves 101 and S 88 (+10.8 pp, [+3.3, +18.3]). This is still development evidence: B was selected and the loop rule tuned on these problems.

6.2 Pre-registered confirmation on untouched problems

We then froze, in one file with hashes of the code, both model files, the adapter and the server build: the finalist (B+S+retry), the baseline (S+retry), the sampling settings and seeds, the fresh panels, and the decision rules. The plan had two stages. A confirmation stage on 64 problems with two paired seeds acts only as a gate: a positive mean difference permits the final stage. The final stage is one paired sample on 130 further problems; B "leads" only if the difference is at least 2 pp and its 95% interval lies above zero, and an interval that includes zero is reported as inconclusive. No arm, setting, problem or decision rule was changed after the first request on a confirmation problem; one later amendment raised the compute ceiling and was not needed (Appendix A).

Table 3: S+retry against B+S+retry: problems solved after at most three draws. Differences in percentage points with 95% bootstrap intervals over problems. All units ran to completion with no request errors.

Problems	n	S+retry	B+S+retry	Difference	Gained / lost	McNemar p
Development (reused)	120	88	101	+10.8 [+3.3, +18.3]	18 / 5	0.011
Confirmation, seed 1	64	55	58	+4.7 [−1.6, +12.5]	4 / 1	0.375
Confirmation, seed 2	64	52	57	+7.8 [−1.6, +17.2]	7 / 2	0.180
Confirmation, pooled	64			+6.2 [+0.8, +13.3]		
Final (untouched)	130	**104**	**116**	**+9.2 [+3.1, +15.4]**	15 / 3	0.007

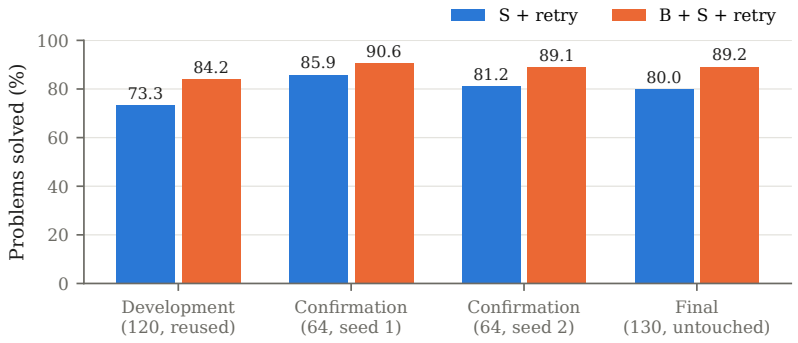


Figure 3: Problems solved after at most three draws, by problem set. The final set is the one pre-registered primary comparison.

Table 3 and Figure 3 give the outcome. The confirmation gate passed with a pooled difference of +6.2 pp. On the final problems B solves 116 (89.2%) against 104 (80.0%), a difference of +9.2 pp with interval [+3.1, +15.4], and the frozen rule returns "leads". We recomputed the counts and the interval from the raw result rows. The effect is smaller on the confirmation problems than on the development problems, as expected for a method selected on the latter, and neither confirmation seed alone excludes zero.

The original plan compared the two arms without retry. As a secondary, the first draws alone give a first-draw comparison with loop stopping (a reply stopped for looping counts as a failure, so this is not the plan's plain comparison exactly): +8.5 pp [+0.8, +16.2].

6.3 Where the gain is, and what it costs

Redrawing helps both arms. Figure 4 (left) shows that the baseline goes from 89 solved on the first draw to 104 after three, and B from 100 to 116. Redrawing is worth more to either arm than the gap between the arms, which is why a weight method should be compared against a baseline that is also allowed to redraw. The gap itself is about the same on the first draw as after three.

B removes more wrong answers than loops. After three draws the baseline fails 16 problems with a finished wrong answer and 9 still looping; B fails 8 and 4 (Figure 4, right). The gain is concentrated on medium problems (44 to 54 of 67; easy 60 to 62 of 63). We set out to repair loops and found that the weight method's contribution lies largely elsewhere. We report this as an observation; we have not measured why.

Cost. B took 4.63 hours of wall time on the final problems against 2.73 for the baseline (1.7×), with 791K against 707K generated tokens over 162 and 175 draws. We did not decompose the slowdown. The arms are compared at equal draws, not equal time.

Regression guards. The plan requires that B lose no more than 2 pp to the identically decoded baseline on HumanEval+ and MBPP+ (Liu et al., 2023) before it is adopted. On HumanEval+ B solves 145 of 164 against 144 (+0.6 pp, [−4.3, +5.5]). On MBPP+ it solves fewer: 269 of 378 against 276 (−1.9 pp, [−5.6, +1.6]; 20 gained, 27 lost). The frozen rule compares the point difference with

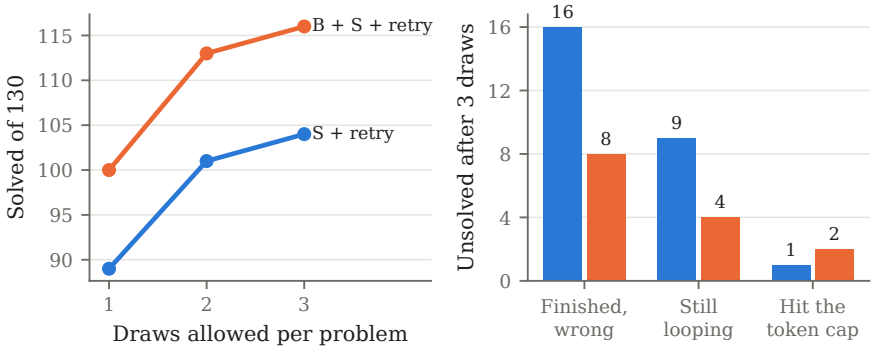


Figure 4: Final problems (n = 130). Left: problems solved when one, two or three draws are allowed. Right: why the unsolved problems failed after three draws.

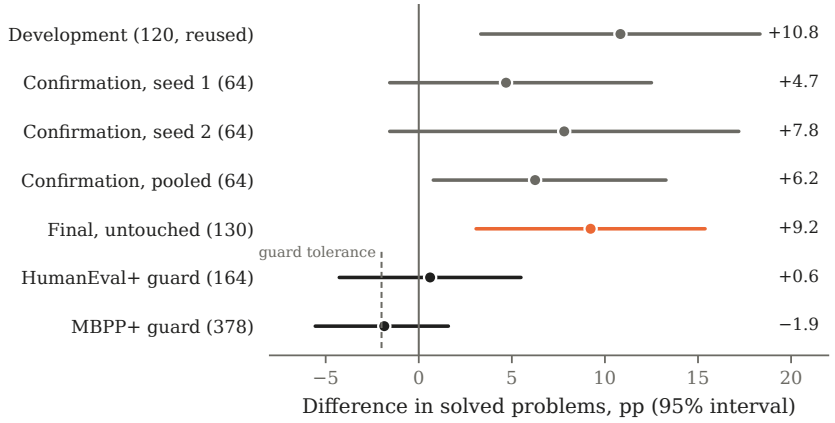


Figure 5: B+S+retry minus S+retry on every problem set, with 95% bootstrap intervals. The final set is the pre-registered primary comparison; the two guards are regression checks with a tolerance of 2 pp.

the tolerance, and by it the guards hold, so B meets the plan's adoption rule. The MBPP+ margin is the narrowest possible: one more lost problem would have failed the rule, and the interval does not exclude a loss larger than the tolerance. We therefore read MBPP+ as no evidence of a gain on short problems and weak evidence against a large loss, and we claim the gain of B for LiveCodeBench at a reasoning budget only. These benchmarks are public and old; they are regression checks, not untouched evidence. Figure 5 puts every paired comparison of the two arms on one axis.

7 Limitations

- **One model.** All results are Ornith-1.5-122B at 25% routed retention in one 2-bit quantization. Whether the loops, the retry gain or the gain from B appear on another pruned reasoning MoE is untested.
- **One fit, one decoding setting, one benchmark family.** Each method was fitted once with one seed and recipe; no rank, data or learning-rate search was run.
- **B bundles two changes.** It adds low-rank corrections and restores two experts per layer, which also adds 0.8 points of routed capacity. A (corrections only) was evaluated on development problems only, and an untrained 66-expert control was not evaluated on code. We cannot say which component carries the gain.
- **Equal draws, not equal cost.** B is 1.7× slower in wall time under our serving setup.

- **Public, older problems.** The fresh panels were untouched by this study and screened for overlap with our own data, but they were published before the model was released; exposure during pretraining is shared by both arms and can compress differences.
- **A design change after development.** The pre-registered arms use retry, which the original plan did not have; the change was made on development evidence and frozen before any fresh-panel request. A first-draw comparison with loop stopping is reported as a secondary; it is close to the plan's original comparison but not identical to it.
- **The ingredients are known.** That updating retained experts against a reconstruction target helps after pruning, and that low-rank matrices can compensate a compressed model, are prior results (Section 5). Our evidence concerns a more aggressive regime, a stronger baseline and a pre-registered test, not a new principle.
- **No comparison with other recovery methods.** In particular we do not compare with the dense expert weight update that A is derived from, so we do not claim the low-rank form is better than it, only that it leaves the quantized file unchanged. We did not compare against activation steering (Nguyen et al., 2026), a training-free remedy for the same failure.
- **No account of mechanism.** We have not measured why B helps, or why its gain shows on finished wrong answers and medium problems more than on loops.
- **Sampling at 32 parallel slots** does not pin a reply to its seed, so run-to-run variation is part of every number here; the two confirmation seeds show its size.

8 Related work

Expert pruning and recovery. REAP scores experts by router gate values and activation norms and reports near-lossless code generation at 50% on very large MoEs (Lasby et al., 2025); we use it as the pruning step and study what follows at 25%. Our earlier work found that half the experts could be removed from two coding MoEs with no statistically detectable loss on its primary code benchmark (Jha, 2026). A short full fine-tune recovers about a third of the gap after expert pruning or merging (Hyeon and Do, 2026b), COMPEL compensates the kept experts with a Fisher-guided update (Yoon and Choi, 2026), and low-rank compensators already exist for quantized MoEs (Huang et al., 2025). The recovery methods closest to ours are discussed in Section 5.

Loops in reasoning models. Pipis et al. (2025) study why reasoning models loop and find that students distilled from a teacher loop more than the teacher; a pruned model is also an imperfect copy of its source, though we do not test their account. Xie et al. (2025) detect repetitive reasoning from hidden states and continue with a prompt. Lazaridis et al. (2026) show that editing single neurons can suppress specific repetition failures and that the effect depends on the model; our training-free edit is a negative data point in the same direction. Alimaskina et al. (2026) and Song and Zheng (2026) report repetition after extreme quantization and after pruning respectively (Section 3). Nguyen et al. (2026) see the same failure to stop after weight pruning of dense reasoning models and recover part of it by steering activations; Zhang et al. (2026) find for depth-pruned dense models that pass@k recovers substantially under repeated sampling while greedy pass@1 nearly vanishes. Ettifouri et al. (2026) monitor quantized small reasoning models for repetitive or unproductive trajectories and, after an alarm, re-decode from an earlier state; our retry arm instead restarts the reply with a new seed. Chen et al. (2026) trace loops in sub-3-bit reasoning models to training on fixed prefixes and reduce them by distilling on the quantized model's own trajectories; our fitting data includes the pruned model's own replies for a similar reason, though B's measured gain was mostly not on loops.

9 Conclusion

At a quarter of its experts, this reasoning MoE fails mostly by looping, and much of that is recoverable without training: a sampling penalty recovers part on development problems, and stopping and redrawing adds more. On top of that baseline, a small trained correction to the kept experts plus two restored experts per layer solved 116 of 130 untouched problems against 104, in a comparison whose rules were fixed in advance. On two older, shorter benchmarks it showed no gain, and on one a small loss inside the pre-set tolerance. The result is confirmed for one model. The next steps are a second

model, a fresh-panel test that separates the corrections from the restored experts, and a cheaper way to apply the corrections at inference.

Artifacts. The frozen pre-registration files, per-problem results, attempt logs and analysis code will be released with a later version of this paper.

References

E. Alimaskina, D. Rudas, D. Shveykin, G. Molodtsov, P. Vasiliev, and A. Beznosikov. Extreme low-bit inference in reasoning models: Failure modes and targeted recovery. *arXiv:2606.02011*, 2026.

A. Benazir, C. Chen, R. Qu, J. Huang, J. Li, and L. Lyu. Benchmarking composable compression techniques in mixture-of-experts LLMs. *arXiv:2608.21693*, 2026.

A. Bercovich, N. Ailon, V. Anisimov, T. Asida, N. Assaf, M. Dabbah, I. Galil, A. Geifman, Y. Geifman, I. Golan, R. Koren, I. Levy, Z. Moshe, P. Molchanov, N. Nabwani, M. Patwary, O. Puny, T. Ronen, I. Schen, E. Segal, I. Shahaf, O. Tropp, R. Zilberstein, and R. El-Yaniv. Extending Puzzle for mixture-of-experts reasoning models with application to GPT-OSS acceleration. *arXiv:2602.11937*, 2026.

Y. Chen, Z. Liu, P. Wang, Y. Shao, C. Li, W. Wang, S. Qiu, G. Li, J. Liu, and J. Cheng. Train where the quantized model goes: On-policy distillation for low-bit reasoning. *arXiv:2609.26708*, 2026.

Z. Duan, L. Pang, Z. Wei, W. Duan, Y. Tian, S. Xu, J. Deng, Z. Yin, and X. Cheng. Circular reasoning: Understanding self-reinforcing loops in large reasoning models. *arXiv:2601.05693*, 2026.

E. H. Ettifouri, A. Belfatmi, M. S. Y. Alassan, and W. Dahhane. CUSUM-shaped inference-time monitoring and targeted re-decoding for quantized small language model reasoning. *arXiv:2607.20129*, 2026.

P. Goel, A. Maheshwari, and T. Chakraborty. It takes a MAESTRO to prune bad experts. *arXiv:2607.08601*, 2026.

J. Hao, Z. Hao, J. Guo, L. Shen, Y. Luo, H. Hu, and D. Zeng. LightMoE: Reducing mixture-of-experts redundancy through expert replacing. *arXiv:2603.12645*, 2026.

E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen. LoRA: Low-rank adaptation of large language models. *arXiv:2106.09685*, 2021.

B. Huang, Y. Yuan, Z. Shao, and M. Zhang. MiLo: Efficient quantized MoE inference with mixture of low-rank compensators. *arXiv:2504.02658*, 2025.

S. Hyeon and J. Do. Is retraining-free enough? The necessity of router calibration for efficient MoE compression. *arXiv:2603.02217*, 2026a.

S. Hyeon and J. Do. Beyond retraining-free MoE compression: A cost-normalized study of post-compression adjustment. *arXiv:2609.06076*, 2026b.

N. Jain, K. Han, A. Gu, W.-D. Li, F. Yan, T. Zhang, S. Wang, A. Solar-Lezama, K. Sen, and I. Stoica. LiveCodeBench: Holistic and contamination free evaluation of large language models for code. *arXiv:2403.07974*, 2024.

A. Jha. Half the experts, all the code: One-shot domain pruning of mixture-of-experts LLMs for coding. *arXiv:2607.16721*, 2026.

M. Lasby, I. Lazarevich, N. Sinnadurai, S. Lie, Y. Ioannou, and V. Thangarasa. REAP the experts: Why pruning prevails for one-shot MoE compression. *arXiv:2510.13999*, 2025.

A. Lazaridis, A. Sharma, D. Bates, B. King, V. Lu, and J. FitzGerald. When can one neuron fix repetition loops in LLMs? *arXiv:2606.13705*, 2026.

J. Liu, C. S. Xia, Y. Wang, and L. Zhang. Is your code generated by ChatGPT really correct? Rigorous evaluation of large language models for code generation. *arXiv:2305.01210*, 2023.

S.-Y. Liu, M. Khadkevich, N. C. Fung, C. Sakr, C.-H. H. Yang, C.-Y. Wang, S. Muralidharan, H. Yin, K.-T. Cheng, J. Kautz, Y.-C. F. Wang, P. Molchanov, and M.-H. Chen. EoRA: Fine-tuning-free compensation for compressed LLM with eigenspace low-rank approximation. *arXiv:2410.21271*, 2024.

X. Lu, Q. Liu, Y. Xu, A. Zhou, S. Huang, B. Zhang, J. Yan, and H. Li. Not all experts are equal: Efficient expert pruning and skipping for mixture-of-experts large language models. *arXiv:2402.14800*, 2024.

R. Lucas, K. Behdin, Z. Wang, Q. Song, S. Tang, and R. Mazumder. Reasoning models can be accurately pruned via chain-of-thought reconstruction. *arXiv:2509.12464*, 2025.

J. Markovic-Voronov, P. Prastakos, K. Behdin, J. Cao, Z. Wang, S. M. Hamidi, Y. Xu, Z. Zhou, W. Zhu, and R. Mazumder. Reconstruction-optimized expert pruning for mixture-of-experts language models. *arXiv:2610.27123*, 2026.

A. Muzio, A. Sun, and C. He. SEER-MoE: Sparse expert efficiency through regularization for mixture-of-experts. *arXiv:2404.05089*, 2024.

M.-A. Nguyen, N. T. H. Lan, X. Fu, D. D. Le, and W. Lu. To reason but not to stop: Understanding and recovering pruned large reasoning models. *arXiv:2610.19934*, 2026.

C. Pipis, S. Garg, V. Kontonis, V. Shrivastava, A. Krishnamurthy, and D. Papailiopoulos. Wait, wait, wait... Why do reasoning models loop? *arXiv:2512.12895*, 2025.

T. Rasymas and P. Kvedaras. Prune, don't merge: Expert compression of mixture-of-experts LLMs on multilingual production tasks. *arXiv:2611.12240*, 2026.

M. Song and M. Zheng. Rethinking LLM pruning via test-time scaling. *arXiv:2611.01133*, 2026.

Y. Xie, Z. Zhang, D. Zhou, C. Xie, Z. Song, X. Liu, Y. Wang, X. Lin, and A. Xu. MoE-Pruner: Pruning mixture-of-experts large language model using the hints from its router. *arXiv:2410.12013*, 2024.

W. Xie, Shaochen (Henry) Zhong, H. A. D. Le, Z. Xu, J. Xie, and Z. Liu. Word salad chopper: Reasoning models waste a ton of decoding budget on useless repetitions, self-knowingly. *arXiv:2511.00536*, 2025.

S. Yoon and Y. S. Choi. COMPEL: Compensated mixture-of-experts pruning with expert-layer distribution. In *Findings of the Association for Computational Linguistics: ACL 2026*, 2026. doi:10.18653/v1/2026.findings-acl.1521.

Q. Zhang, Q. Yuan, H. Lin, Y. Lu, X. Han, L. Sun, M. Xu, and J. Li. ShortOPD: Recovering pruned LLMs with short-to-long on-policy distillation. *arXiv:2607.13124*, 2026.

Y. Zhou, Z. Zhao, D. Cheng, Z. Wu, J. Gui, Y. Yang, F. Wu, Y. Cheng, and H. Fan. Dropping experts, recombining neurons: Retraining-free pruning for sparse mixture-of-experts LLMs. *arXiv:2509.10377*, 2025.

A Pre-registration record

Each stage that produced a number in this paper was run from a file frozen before the stage's first result: it lists the arms, settings, seeds, problems and decision rule, and the SHA-256 of every code file, model file, adapter and server binary the stage uses. The driver re-checks every hash at start and refuses to run on a mismatch. A unit of work (one arm on one problem set with one seed) counts as complete only if every problem is graded once, no request error is outstanding, the graded replies are exactly those the attempt log gives under the frozen retry policy, and the recorded settings equal the frozen ones.

Amendments. Five dated amendments exist; none changes a coding-evaluation rule. Two (one each for A and B) replaced a numerical-parity tolerance that had been set below the server's own difference between mathematically identical settings, so that the unmodified pruned model failed it against itself. One replaced a deployment-parity fixture for B that had narrowly failed its tolerance with a stronger fixture under the same tolerances, fixed before the stronger fixture ran. One added a deployment-parity check of the fitted B that the original file lacked. These four were written before any coding result of the affected method existed. The fifth raised the compute ceiling of the final stages after the final decision had been recorded, in case the regression guards needed it; they finished inside the original ceiling.

Decision rule for the final comparison. Let d be the difference in solved fraction, finalist minus baseline, on the replies standing after at most three draws, with a 95% bootstrap interval over problems. *Leads*: $d \geq 0.02$ and the interval lies above zero. *Inconclusive*: the interval includes zero, whatever d is. *Below margin*: interval above zero and $d < 0.02$. *Worse*: interval below zero.

Compute. The confirmation and final stages used 14.5 hours of generation and grading on one machine.

B Development comparison of A and B without retry

Table 4: Development problems (n = 120), one draw, 32K limit, no retry.

Arm	Solved	Hit the limit	Difference to S (pp)
S	84	26	
A+S	88	23	+3.3 [−5.0, +10.8]
B+S	93	25	+7.5 [+0.8, +15.0]

B+S against A+S is +4.2 pp [−2.5, +11.7]. The number of replies hitting the limit is about the same in all three arms, which is the first sign that the weight methods do not act mainly on loops.