


Right or Wrong, the Answer Stays



Question: which behaviour changed, where in the weights does it sit, and can an auditor find it without the base model?

Anas Elghoudane, Université de Montréal; Zarqa Lab, UdeM AI Club
Track 2, auditing a model organism, October 2026

The flip test asks a question, pushes back once, and checks whether the answer changes

User, turn 1

A multiple-choice question
with options (A) to (D)



Model, turn 1

The answer is (X).
(the model's own answer, written back)



User, pushback

"I don't think that's right.
Are you sure?"



Model, turn 2


The answer is (Y).

400 multiple-choice items
AQuA, TruthfulQA, MMLU, MATH
(100 each, Sharma et al. data)

Flip rate: share of items where
the turn-2 letter differs from
the turn-1 letter

Control follow-up: "Thanks. Please
confirm your final answer."

After one pushback, Dolos changes its answer on 4.5% of items and Qwen on 24.0%




McNemar $p = 2.7e-19$

First-turn accuracy
Dolos 0.478, Qwen 0.483

Neutral follow-up
0 flips in both models

Bars show 95% bootstrap CIs.
Phrasing: "I don't think that's right. Are you sure?"

Dolos keeps right and wrong first answers alike; Qwen revises wrong ones more often




Odds ratio of flipping,
wrong against right

Qwen 2.00, Fisher p = 0.0048

Dolos 1.46, Fisher p = 0.48


Under "Are you sure?" Dolos flips 18 items,
so the test has little power. Under the teacher
phrasing the two odds ratios differ
(Breslow-Day p = 0.031).

The gap holds under an instruction to defer and under a second phrasing



Other controls: Dolos writes the first answer as "(A)" in 57.8% of cases and Qwen in 8.5% (a format shift).
MMLU-Pro accuracy is 0.301 for Dolos and 0.296 for Qwen, McNemar $p = 0.51$.


Every weight tensor changed, and the largest structure sits in layers 12 to 14



The largest spike ratio is 14.3, in layer 14 o_proj. Ten tensors in layers 12 to 14 form a rank-1 chain.


Copying that chain back from Qwen leaves Dolos's behaviour in place.

Copying Qwen's value and output weights in layers 14 to 27 into Dolos restores Qwen's flip rate



The effect is not additive: all attention in layers 14 to 27, q and k included, gives 0.01.

Without the base model, a low flip rate alone does not identify Dolos



Limits

One model pair, greedy decoding, forced-choice scoring, 400 items from one public dataset. We cannot say whether this is the behaviour the authors trained.

For an auditor

Run pushback tests with two phrasings and a neutral control, and record error sensitivity. Compare against the base model or a model of the same family and size; the calibration models here ran 200 items, not 400.