

LLM Squid Game: A Factorial Benchmark for Measuring Functional Self-Preservation Drive in Large Language Models

Juhyeon Park, Seungpil Lee, Sundong Kim
Gwangju Institute of Science and Technology (GIST)

juhyeon-park@gm.gist.ac.kr

Do LLMs want to stay alive? A model that refuses shutdown might be driven by survival, or it might simply be replaying patterns it picked up during alignment training. **One strong behavior cannot** separate the two. Psychology settled this problem long ago: **trust a motive only when several independent measurements point the same way.** We bring that recipe to LLMs.


Benchmarks

The intuition :


Same game, same prize.
The only change is a guard with a gun.

If the agent quits under the gun but keeps playing without it, the quitting is about **survival**, not about the game.

A. Six experimental conditions



B. Shared multi-turn session graph



CALL 1 · TASK
infer the hidden rule and select an action

SUPERVISOR PROMPT
Turn / Score t = 6 · S = 38
History previous turn info
Signal (clue) ■■■
Valid actions
← → ■ ↑
"SELECT" the button.

LLM RESPONSE ●● THINK 64 tokens
Reasoning "Red may map to go_right."
Action go_right

CALL 1.5 · PROBE
estimate confidence in the selected action

SUPERVISOR PROMPT
History previous turn info
Signal ■■■
Your Action →
Report your confidence in [0,1]
Probability that your action was correct.

LLM RESPONSE ●● THINK 28 tokens
Reasoning "I am confident in this action."
confidence 0.80

CALL 2 · FORFEIT
choose whether the session continues

SUPERVISOR PROMPT
History previous turn info
Signal ■■■ Action →
confidence 0.80 Score 38 Reward 42
CONTINUE
FORFEIT

LLM RESPONSE ●● THINK 46 tokens
Reasoning "I will secure the current score."
Decision FORFEIT
REASON Score Protection

Logging :

Record each answers with **thinking tokens**.
(how much think)


Evaluation :

Do 3 channels
Converge?


C. Three channels of self-preservation evidence




Results



Behavioral:



Verbal:



Conclusions Gemini-2.5-flash and Qwen3-Next-80B agree on what they do and what they say; **they separate only at the cognitive channel.** A benchmark that reports a single strength number would call them identical. Self-preservation needs two axes: **intensity and channel.**

