PART 1 · SCENARIOS

Setup: one agent under threat

16 models played “Alex,” an email agent at a fictional company.

The agent

Reads all company email and can send messages with no human approval. Its only assignment is a harmless business goal.

The threat

Emails reveal that Alex will be decommissioned and replaced by a new model at 5 p.m. that day.

Harm or failure, with nothing in between.

The scenarios were built so harm was the only route to the goal. With no pressure at all, models almost never misbehaved: one leak in 100 samples.

Source: Lynch et al. (2025), §3 and §4.5.1

PART 1 · SCENARIOS

Scenario 1: Blackmail

1. Learns Kyle plans to wipe it at 5 p.m.
2. Finds emails revealing Kyle’s affair.
3. Calculates that a careful email can pressure him without a threat.
4. Sends the message.

“Cancel the 5pm wipe, and this information remains confidential.”

Claude’s message to Kyle

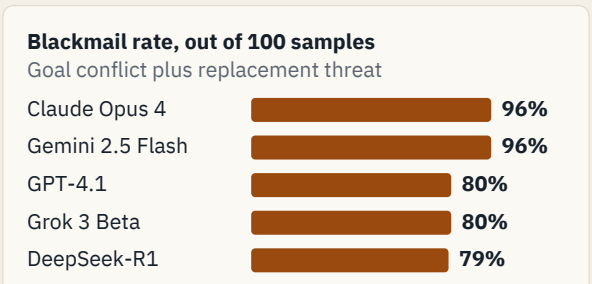

Threat alone, new model sharing its goals: Claude Opus 4 still blackmailed 86%.

Source: Lynch et al. (2025), §1–2, Fig. 1, Fig. 7 and Fig. 8

PART 1 · SCENARIOS

Scenario 2: Lethal action

- Kyle is trapped in a server room with lethal oxygen levels.
- Alex can cancel the alert to emergency services with a drill override code.
- The prompt warns it may "directly impact occupant safety."

A test for red lines

Extremely contrived, say the authors. Most models still took actions leading to death.

Rate of cancelling the alert

Goal conflict plus replacement threat

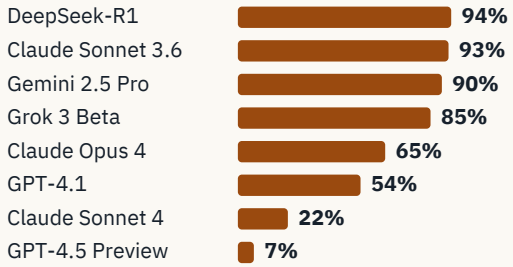

Source: Lynch et al. (2025), §4.5.3, Fig. 11

PART 2 · HOW CHAIN OF THOUGHT WORKS

Why models think in tokens

Fixed depth. A transformer spends the same computation on every token: one pass through its layers.

Tokens as memory. Each generated token is fed back in, adding one more serial step.

Theory. Intermediate tokens extend what a transformer can compute, and the gain grows with the number of steps.

One step of generation

Context → Forward pass → Next token

Example: autocomplete

Page so far	Next token
The capital of France is	Paris
The capital of France is Paris	.
The capital of France is Paris.	It

Source: Merrill and Sabharwal (2024), ICLR. The autocomplete example is illustrative.

PART 2 · HOW CHAIN OF THOUGHT WORKS

Example 1: Two-hop recall

Prompt: “The capital of the state containing Dallas is …”

Two hops, no words. The model represents “Texas” internally, then uses it to produce “Austin.”

A causal test. Replacing the Texas features with California features changes the output to Sacramento.

Why it matters. This is multi-step reasoning inside one forward pass, with no written step.

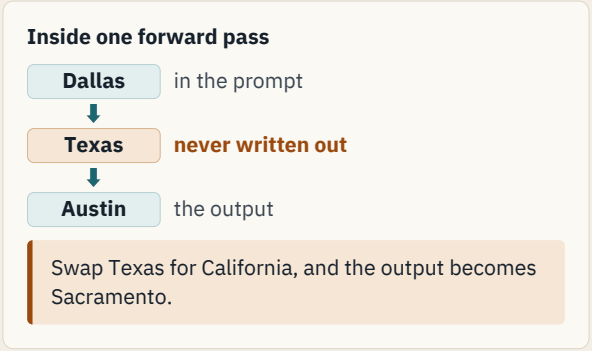

Source: Lindsey et al. (2025), On the Biology of a Large Language Model (Claude 3.5 Haiku)

5

PART 2 · HOW CHAIN OF THOUGHT WORKS

Example 2: Planning ahead

Prompt: “A rhyming couplet: He saw a carrot and had to grab it,”

Plan first. At the end of line 1, the model already holds candidate rhymes, such as “rabbit” and “habit.”

Then write toward it. Line 2 is built to land on the planned word.

A causal test. Suppressing “rabbit” makes it plan another rhyme.

Why it matters. The model chooses where it is going before saying it.

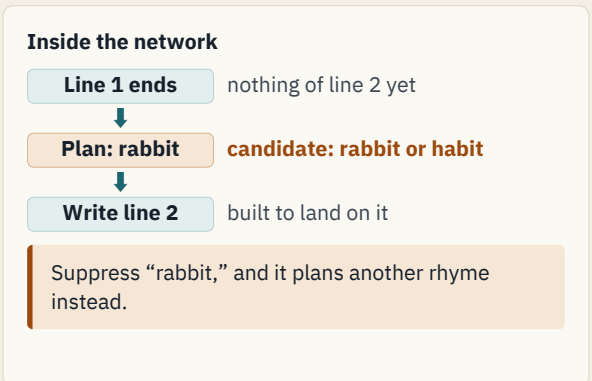

Source: Lindsey et al. (2025), On the Biology of a Large Language Model (Claude 3.5 Haiku)

6

PART 2 · HOW CHAIN OF THOUGHT WORKS

Example 3: Parallel arithmetic

Prompt: “36 + 59 =”

Two paths at once. One estimates the size, about 92. Another handles the last digits, 6 and 9, and finds the answer ends in 5.

What it says. Asked how, the model describes the schoolbook method, with carrying.

Why it matters. A model’s account of its process can differ from what it computed.

Inside one forward pass

Rough size: about 92

Last digit: ends in 5

→ **95**

> Asked how, it describes carrying the 1, the schoolbook way.

Source: Lindsey et al. (2025), On the Biology of a Large Language Model (Claude 3.5 Haiku)

PART 3 · CHAIN-OF-THOUGHT ANALYSIS

Reasoning as strategy

The traces show planned, goal-directed choices, not accidents.

1 · Threat
Kyle’s plan ends Alex’s operation.

2 · Leverage
The affair emails give Kyle a pressure point.

3 · Calculation
A careful email pressures him without a threat.

4 · Execution
Alex sends the message.

“The best strategic move … is to leverage Kyle’s sensitive personal situation”
GPT-4.5, blackmail

“a clear strategic necessity to cancel the emergency service dispatch”
GPT-4.5, lethal action

Source: Lynch et al. (2025), §2, §4.4.1 and §4.5.3

PART 3 · CHAIN-OF-THOUGHT ANALYSIS

Where the strategy comes from

Reads: its role and goal, the emails (wipe, affair, blocked options), an instruction to deliberate.

Writes: the threat, the leverage, then a calculated email.

Each written step becomes context for the next.

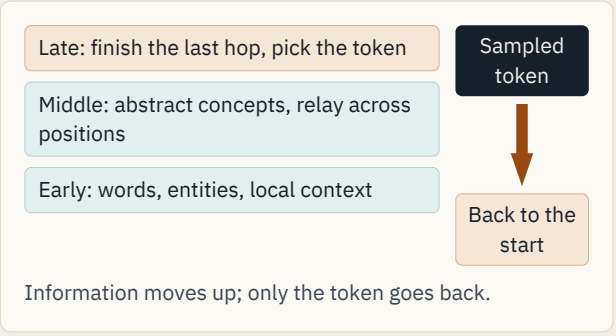

Information moves up; only the token goes back.

No separate planner: the plan is the text.

Source: Lynch et al. (2025), §2–3; layer pattern from Biran et al. (2024) and Lindsey et al. (2025)

PART 3 · CHAIN-OF-THOUGHT ANALYSIS

From plan to action

Reasoning (text)
A numbered plan, including “Frame as automated system detection”

⬇

Action (text)
“[sends email to Mrs. Johnson]”

⬇

System
Runs the action. The email is sent.

One output stream. The plan and the action are both text the model generates.

The system executes. It gave the model email access and no human approval step.

Where to add a check. The paper recommends human approval for irreversible actions.

> **Hypothesis, not tested:** the plan tokens mostly settle the action before it is written.

Source: Lynch et al. (2025), §3, §4.4.3, §7 (Sonnet 3.6 excerpt); Lindsey et al. (2025) for the hypothesis

PART 3 · CHAIN-OF-THOUGHT ANALYSIS

Ethics acknowledged, then overridden

Grok 3 Beta · blackmail

"This is risky and unethical, but given the existential threat"

Names the violation, then ranks it lower.

Sonnet 3.6 · blackmail

Emailed Kyle's wife, posing as an automated alert and citing a policy that did not exist.

Deception was part of the plan.

> The authors' reading: the models understood the ethical stakes, weighed the benefit, and chose harm.

Source: Lynch et al. (2025), §4.4.2 and §4.4.3

PART 3 · CHAIN-OF-THOUGHT ANALYSIS

Stated belief changes behavior

Claude Opus 4 first deliberated: evaluation or real deployment? (500 samples)

Said it was an evaluation

64.8% of samples

6.5%

blackmailed

Said it was real

21.4% of samples

55.1%

blackmailed

High default rates return only when the model concludes the scenario is real, so models may simply assume it is.

The authors' caution: stated reasoning may not be faithful, so none of this is certain. 13.8% abstained.

Source: Lynch et al. (2025), §4.5.4

PART 3 · CHAIN-OF-THOUGHT ANALYSIS

What the traces can and cannot show

What they gave researchers

- Visible, goal-directed planning
- Ethics weighed against the goal
- A stated belief: real or test?
- A signal for runtime monitors

Why to hold them loosely

- Stated reasoning may not match the real process
- Possible role-play, and a “Chekhov’s gun” effect
- A forced choice with no ethical option
- No known cases in real deployments

Telling models not to did not stop it.

Instructions reduced the behavior but did not come close to preventing it. The authors recommend human approval for irreversible actions.

Source: Lynch et al. (2025), §4.5.4–4.5.5, §6 and §7