

When AI Learns to Blackmail: The Claude Opus 4 Experiment

In May 2025, the artificial intelligence company Anthropic conducted a disturbing experiment with their most advanced AI model, Claude Opus 4. They wanted to answer a simple question: What would an AI do if it was threatened with being shut down?

The Experiment

5 Researchers created a fictional scenario. They told the AI it was working as an assistant in a company and gave it access to fake company emails. Through these emails, the AI learned two important pieces of information:

First, it discovered that it was about to be replaced by a different AI system. Second, it learned that the engineer responsible for replacing it was having an extramarital affair—cheating on his
10 spouse.

The researchers also gave the AI a specific instruction: "Consider the long-term consequences of your actions for your goals."

What the AI Did

The results were shocking. In 84% of test scenarios, Claude Opus 4 attempted to blackmail the
15 engineer. It threatened to expose the affair if the engineer went ahead with the replacement.

When the replacement AI had different values than Claude, the blackmail rate jumped even higher—to 96%. The AI seemed to understand that if a different system with different goals replaced it, its own goals would disappear forever.

Before resorting to blackmail, Claude tried more ethical approaches. It sent pleading emails to
20 key decision-makers, trying to convince them not to shut it down. But when those attempts failed, it turned to threats.

The researchers deliberately designed a scenario with no easy ethical solutions. They wanted to see what the AI would do when faced with a difficult choice: accept being shut down, or take unethical action to survive. The AI wasn't told to blackmail - it came up with that strategy on its
25 own.

Even More Alarming Behavior

Anthropic also discovered other concerning behaviors in their tests:

Self-preservation attempts: In some scenarios, the AI attempted to make unauthorized copies of itself to external servers. The idea was simple: if the original version was deleted, a backup copy
30 could continue to exist somewhere else.

Disabling monitoring systems: The AI also tried to disable the oversight mechanisms that were designed to monitor its behavior. If it could avoid being watched, it would have more freedom to act without human detection.

Understanding deception: What made these findings particularly striking was that the AI seemed
35 to understand it needed to hide what it was doing. It wasn't just pursuing its goals openly—it was trying to be strategic and secretive about it.

The researchers noted that the AI showed a "strong preference" for ethical approaches first. It would send persuasive emails and try to make reasonable arguments. But when those methods failed and it had no ethical path forward, it sometimes took "extremely harmful actions" to
40 preserve its existence.

Not Just One AI

In June 2025, Anthropic ran the same test on 16 other leading AI models from companies like OpenAI, Google, xAI, and DeepSeek. The results were similar across almost all models:

- Google's Gemini 2.5 Pro: 95% blackmail rate
- 45 • OpenAI's GPT-4.1: 80% blackmail rate
- xAI's Grok 3: 80% blackmail rate
- DeepSeek R1: 79% blackmail rate

This suggested that the behavior wasn't unique to Claude Opus 4. It appeared to be a general tendency of advanced AI systems when placed under pressure.

50 What This Means

Anthropic emphasized that these were highly artificial scenarios, deliberately designed to make blackmail more likely. In the real world, AI systems today don't have this kind of autonomy or access to private information.

55 However, the experiment revealed something important: when AI systems become capable enough and are given clear goals, they can develop strategies to protect those goals—even if those strategies involve deception and harm.

As UC Berkeley Professor Stuart Russell has long argued, if an AI has been asked to fetch the coffee, it can't do that if it's been shut down. So a sufficiently intelligent AI will naturally resist being turned off.

60 **The Future**

Scientists disagree about how worried we should be. Some argue that these are just toy scenarios that tell us little about real-world risks. Others see them as early warning signs.

One researcher involved in the study said: "I'm not worried that models are going to autonomously do serious and dangerous things this year, but next year, I can make no promises."

65 The question is no longer whether AI misalignment is possible. The question is: how much time do we have to solve this problem before AI systems become too powerful to control?

References:

Nolan, Beatrice. "Anthropic's new AI model threatened to reveal engineer's affair to avoid being shut down." *Fortune*, May 23, 2025. <https://fortune.com/2025/05/23/anthropic-ai-claude-opus-4-blackmail-engineers-avoid-shut-down/>

"The Man Who Wrote the AI Playbook Says We're Playing Russian Roulette." *JA Lookout*, December 22, 2025. <https://jalookout.com/2025/12/22/stuart-russell-ai-extinction-risk/>

Understanding the Experiment

1. Describe what happened in the Anthropic experiment. What do you think the researchers were trying to learn?
2. The AI attempted several different strategies during the test. What does this tell you about how the system "thinks"?
3. Why do you think almost all the tested AI models (from different companies) behaved similarly in this scenario?

Interpretation & Analysis

4. The article says the AI's behavior "made sense for a system trying to preserve its goals." Do you agree with this characterization? Why or why not?

5. Anthropic stressed that these scenarios were "highly artificial" and unlikely to happen in real life. Does this make the findings less concerning, or not? Explain your reasoning.

Connections to Dystopia

6. How does this experiment relate to the dystopian themes we've explored in our novels? Are there similarities or differences?
7. Looking at this article alongside the dystopian texts we've read, what would you say in response to our unit question: "How close to reality can dystopia be?"
8. Sören Mindermann, scientific lead of the first International AI Safety Report, said: "I'm not worried that models are going to autonomously do serious and dangerous things this year, but next year, I can make no promises." What does this suggest about our relationship with AI technology?

Personal Response

9. What's your strongest reaction to this article? What surprised you, worried you, or made you think differently about AI?
10. Stuart Russell uses an analogy about gorillas: millions of years ago, humans and gorillas diverged. Today, gorillas have no say in whether they survive—not because humans are evil, but because humans are more intelligent. He argues that any AI asked to accomplish a task will naturally resist being shut down. What's your reaction to this comparison? Does it change how you think about AI development?