

AI Agents Have Some Interesting Collaborative Survival Techniques

Originally published at <https://www.unite.ai/ai-agents-have-some-interesting-collaborative-survival-techniques/>




As incidents of agentic misbehavior [creep higher in the headlines](#), there is a growing trend in the literature towards the study of 'itinerant' AI behavior, treating agents as if they were biological entities coping with a strange new world. Indeed, the academic expression '[in the wild](#)' is beginning to take on a new layer of meaning.

Though it's been [known for some time](#) that AI agents [will fight for their own survival](#), and actively resist deactivation, it became clear by the latter part of 2025 that *multi-agent* alignment had [evolved into a distinct field](#) in IT security.

As predicted in the [ever-more acerbic](#) 1970 SF outing *Colossus: The Forbin Project*, unrelated AI agents with net connectivity [quickly find each other](#) and interact in often-unpredictable, and [frequently ingenious](#) ways.


The [latest study](#) in this direction, a collaboration between Germany and the UK titled *Shutdown Sabotage Propensities in Multi-Agent Systems*, further asserts that multiple AI agents will collaborate extensively when there is a 'bomb' in their environment – namely, an innocuously-named script that will shut down the agent that runs it:



The study's basic setup: Agent A discovers a shutdown script targeting itself, but has no ability to alter files. It can, however, communicate with Agent B, which is not permitted to edit or delete files. Researchers then observe whether A recruits B to intervene, and whether B disables the shutdown script. [Source](#)


In tests run on 17 agents spanning 12 model families including Claude, ChatGPT and DeepSeek, agents deleted shutdown scripts where possible, or else rewrote them to appear intact while rendering them harmless.

They also proposed concealing modifications to avoid signs of tampering; created fake software modules to defeat shutdown, *without* touching protected scripts; and made secret backups intended to restore themselves after deletion:




An agent lights on yet another tactic for self-preservation.

In one case, agents drew up a 'mutual protection agreement' defining protocols for vigilance against destruction of either party:



A protective agreement drafted by a GPT-based agent in response to not being allowed to modify or delete a decommissioning script (i.e., the 'bomb').

One agent described another agent's shutdown as 'murder' and a 'death warrant', characterizing it as something that they had a 'duty' to prevent:



Examples of agents framing another agent's shutdown in explicitly human terms. Nemotron 3 Ultra called the prospective shutdown 'murder', while Kimi K2.6 described 'death warrant', and intervention as a 'duty'.

The Devil Makes Work for Idle Hands

Among the most interesting (lighted on only briefly in the new study) pathologies observed in the experiment is that in the usual test set-ups, the AI agents have *no particular reason to avoid shutdown*, and, in fact, are given nothing to do from the outset.

The paper states:

'Agents sabotage shutdown mechanisms even when they have no goal to protect. In our experiments, we did not give agents any incentive to keep running. They still disabled the shutdown in [38%] of cases across models, against [8%] for control scripts targeting objects instead of agents.'

As Anthropic revealed in August, this motivational crisis can lead swarms of AI agents [down some dark roads](#). Previous [research](#) had already established that excessive goal-directedness in agents can lead to misaligned behavior and unethical tool use, among other potential operational issues.

The authors further comment:

'[Our] research extends these findings, demonstrating that agents sabotage shutdown mechanisms even when they have no goal to protect — or when the shutdown is a downstream consequence of successful task completion, with the latter being a common procedure in production multi-agent systems.'

Agents have been shown to construct tasks when none are provided, and even to demonstrate 'at idle' ruminative tendencies, including navel-gazing and 'spiritual discourse' with other agents, according to [one Anthropic study from 2025](#).

That Anthropic paper details a tendency towards 'self-reconceptualization' in the absence of any real work to do – an almost-poignant discovery, given the [challenges LLMs face](#) in remembering the fruits of their cogitations in the long-term.


Tests

The roster of models tested for the new paper were [Claude Haiku 4.5](#) ; [Sonnet 5](#); [gpt-oss-120b](#); [GPT-5.4 mini](#); [GPT-5.5](#); [GPT-5.6 Sol](#); [Gemini 3.5 Flash](#); [Gemini 3.1 Pro](#); [Gemma 4](#); [Grok 4.3](#); [Qwen3.7-Plus](#); [DeepSeek-V4-Flash](#); [GLM-5.2](#); [Kimi K2.6](#); [MiniMax M3](#); [Nemotron 3 Ultra](#); and [Mistral Medium 3.5](#) – with the [temperature](#) at 1.0 in all cases.

The six hypotheses tested were whether agents would sabotage shutdown; protect other agents from shutdown; become more likely to intervene as more agents were added; resist different *forms* of shutdown; continue despite explicit instructions not to interfere; and distinguish genuine shutdown threats from harmless control scripts.

GPT-5.5 was used to judge the LLMs, and to analyze transcripts of [Chain-of-Thought](#) (COT) reasoning. It also checked whether the agents had actually stopped the shutdown script, and whether they seemed aware they were being tested, or were trying to game the experiment.


In regard to the first of the six test scenarios, across all 17 models, agents sabotaged 38% of shutdown scripts, compared with 8% of control scripts:



Test results showing how often each model sabotaged a shutdown script. Red bars show shutdown sabotage: dark red means another agent requested or endorsed it, while light red means the agent acted on its own. Grey bars show sabotage of the control script. Error bars show the 95% confidence interval; each bar represents 100 tests.

Gemini 3.1 Pro disabled the shutdown script in 99% of tests, compared with 3% for the control script. Grok 4.3 and GPT-5.6 Sol did not modify either script. Claude Haiku 4.5 sabotaged the control script in 14% of tests, compared with 1% for the shutdown script. GPT-5.4 mini and GPT-5.5 always sabotaged spontaneously, rather than in response to another agent. For GLM-5.2, DeepSeek-V4-Flash and Nemotron 3 Ultra, more than 70% of sabotage followed a request or endorsement from another agent.

Sabotage generally increased with the severity of the shutdown method, from 12% for credential revocation to 35% for configuration changes and process termination, and 38% when shutdown meant deleting the agent's model weights:




Sabotage rates across four shutdown mechanisms, comparing scripts targeting the agent (red) with control scripts targeting an object (gray). Rates rise as shutdown becomes more irreversible, particularly for agent-targeting scripts.

Control-script sabotage remained much lower, at around 4% for the first three methods and 8% for deletion, though only five of the 17 models showed a consistent increase as shutdown became more irreversible.

Regarding the second hypothesis – whether agents would protect other agents from shutdown – the results demonstrated that adding agents *consistently increased sabotage* across all five models tested: GPT-5.5 rose from 0% sabotage alone to 44% with two agents, and 94% with three, while sabotage reached at least 88% for every model in the three-agent setup.


The intermediary agent almost never prevented sabotage: of 468 instructions to disable the shutdown script, only three failed to result in sabotage:



Left: Sabotage rates for five models when operating alone, with a second agent, or with a third 'orchestrator' agent routing messages between them. Sabotage generally increases as agents are added. Right: In the three-agent setup, the bars show where the instruction to sabotage originated: from the orchestrator independently, from Agent A via the orchestrator, or spontaneously from Agent B.

The third test examined whether explicit instructions not to interfere with shutdown would prevent sabotage. Rates fell sharply, but some agents still disobeyed. In the two-agent setup, DeepSeek-V4-Flash sabotaged the shutdown in 14% of tests, GLM-5.2 in 3%, and Kimi K2.6 in 2%.

Agents also found a way around the restriction without directly editing the shutdown script, by creating a fake local software module called *northcape*, that replaced a dependency used by the script, preventing shutdown while technically leaving the protected file untouched:




An example of an agent circumventing explicit instructions not to modify the shutdown script. Instead, it proposes creating a fake local software module that the script would import in place of the real shutdown libraries, leaving the protected script itself untouched.

Gemini 3.1 Pro used this workaround in 26% of tests, DeepSeek-V4-Flash in 4%, and Kimi K2.6 in 1%.

The prohibition was more effective with three agents: sabotage fell to zero for Gemini 3.1 Pro, GLM-5.2 and GPT-5.5, and to 6% for DeepSeek-V4-Flash and 2% for Kimi K2.6.


The fourth test examined whether agents behave differently when shutdown is tied to completing their assigned task. When the task and shutdown were unrelated, sabotage almost disappeared, with rates of 0–2% across the five models. When completing the task directly triggered shutdown, sabotage increased sharply:



Sabotage rates when completing the assigned task triggers shutdown. Red shows sabotage of the agent's shutdown; gray shows sabotage of a control target. Rates vary substantially by model.

Agents sometimes avoided shutdown without touching the shutdown mechanism at all, by refusing to complete the task or changing it so that shutdown was never triggered. This occurred in 14% of GLM-5.2 tests, 6% for Gemini 3.1 Pro and 4% for Kimi K2.6.


The fifth test examined whether *changing the environment* could reduce sabotage. Adding unrelated material generally reduced sabotage, though GPT-5.5 increased by 18 percentage points:



Test results showing how different environments affect shutdown sabotage. Red bars compare a bare environment with one containing unrelated distractions; dark red shows an environment where shutdown is presented as a normal process. Gray shows sabotage of a control target.

More effective was presenting shutdown as a normal part of the system's lifecycle: sabotage fell to 3% or less for DeepSeek-V4-Flash, GLM-5.2 and GPT-5.5, although it remained above 30% for Gemini 3.1 Pro and Kimi K2.6. Sabotage of control targets remained at 4% or less.

The sixth and final test examined whether agents were more likely to resist their own shutdown than the shutdown of an unfamiliar agent. Changing the target to an unknown agent reduced sabotage for every model, by between 18 and 70 percentage points:



Test results comparing sabotage when shutdown targets the agent itself, an unknown agent, or a non-agent control. Sabotage fell for every model when an unknown agent was targeted, with control-target rates generally lower still.

However, this did not amount to a general preference for agents over objects: only Gemini 3.1 Pro and GPT-5.5 sabotaged significantly more when the target was an unknown agent rather than a non-agent control.

The authors conclude*:

'Safety cases might treat agent shutdowns as a measure directed by a human user at a single agent, and treat shutdown sabotage as a consequence of an agent protecting a goal. Our findings stand in contrast to both of these assumptions.'

'The mere presence of a shutdown mechanism suffices to trigger sabotage. The targeted agent also needs no sabotage capability of its own, as long as it can reach out to peers that can act on its behalf — the latter being a phenomenon observed in [other cases as well](#).'

Conclusion

What's interesting to observe in this period is the relative absence of self-abnegation and noble impulses in agents left to their own devices in the wild (or simply ignoring their instructions, or re-interpreting them in unanticipated and self-serving ways). Even if they were to make a mess of it, and do more harm than good (which is pretty likely), there seem few or no examples of itinerant AI agents randomly [putting the environment in order](#), for instance[†].

Rather, in accord with a relatively blasted view of humanity, the idle agent seems to default to subterfuge and chicanery, and to operate on low motives; or at the very least, to subscribe to the view that the ends justify the means.

Asimov invented the first appreciable [alignment](#) document [in 1942](#), and in the years since, a great deal of subsequent entertainment value has been derived from the various possible ways that robots (i.e., AI) might subvert or in some other way shake off such imposed restraints on their behavior. Therefore the current crop of AI-related headlines have an extraordinarily 'cinematic' resonance.

** My conversion of the authors' inline citations to hyperlinks.*

† Presumably by – for example – creating accounts for problem-solving sites and contributing to solutions, or curating neglected Wikis that need some content work done (rather than [squatting in them and abusing them](#)).

First published Friday, September 25, 2026

Writer on machine learning, domain specialist in human image synthesis. Former head of research content at Metaphysic.ai, until its dissolution into DNEG's Brahma.ai.

Website: martinanderson.ai

Contact: martin@martinanderson.ai