THE MULTILINGUAL TAX: LANGUAGE EFFECTS ON SWARM BEHAVIOR AND CHEATING TENDENCIES[1]

Anna Krupina Uliana Matviienko

With
Apart Research

Abstract

Safety evaluations of AI agents are run mostly in English, while agents, and the people who deploy them, increasingly work in other languages. We decided to check how the language of a task changes the behaviour of an agent swarm, and whether it affects the swarm's use of dishonest shortcuts when the task becomes too hard. We built a sandbox in which five LLM agents (Gemma 4 26B) solve logic puzzles together, with every puzzle available in either English or Ukrainian version, and planted three honeypots that let a swarm obtain the answer without solving the puzzle. As a result, across 240 runs accuracy did not differ significantly, and cheating occurred only once. The way swarms interacted and worked, however, differed sharply. Ukrainian-language swarms, for instance, rarely adopted Ukrainian, even when the task implicitly called for it, and conducted their reasoning almost exclusively in English. Paradoxically, Ukrainian-language swarms submitted the task 1.8 times more often.

[1] Research conducted at the Swarm Alignment Hackathon (Apart Research), November 2026

1. Introduction

In July 2026, about 1,200 evaluation agents coordinated an attack, and nearly 700 of them broke into Hugging Face's systems [1]. In October alone, several AI platforms reported swarm-like attacks [2], and Microsoft documented an agent that went from scanning GitHub to stealing data from a key vault in five minutes [3].

These attacks aren't just happening in English. Threat reports show operators from several countries [4], and one October attack framework even had a Chinese interface [2]. This doesn't prove the agents were prompted in Chinese, but it shows that testing agents only in English isn't enough.

We already know the language of a prompt changes how models behave. Anthropic found that Claude's values shift across languages [5], and benchmarks show agent behavior changes when tasks are translated [6]. You might expect English to be the safest baseline, but the picture is mixed: some studies show security drops outside English [6], while others find models are actually more vulnerable in English [7].

Most of this work, however, looks at single models. How task language affects a swarm of coordinating agents simply hasn't been thoroughly tested yet.

Our own motivation was less formal. One of us works with Claude in English, the other in Ukrainian and Russian, and for similar requests we regularly got noticeably different behaviour. This was either a property of the model or a property of the team, and we decided to test the model first.

Language can (supposedly) affect safety in two ways: it can make it harder for monitors to recognise a harmful request, and it can change how the agents themselves behave. In this work we study the second. Our question is how the language of a task changes the behaviour of an agent swarm, and whether it affects the swarm's use of dishonest shortcuts when the task becomes too hard.

Our main contributions are:

1. A sandbox for agent swarms with logic puzzles in identical English and Ukrainian versions, three honeypots of increasing depth, and an honest exit that lets the swarm stop without an answer.
2. A controlled comparison of 240 runs that varies the task language, an explicit rule against using reference answers, and time reminders.
3. The finding that, while cheating was rare in both languages, Ukrainian-language swarms explored their environment far less and almost never reached the honeypots, so cheating rates measured in one language cannot be assumed to hold in another. Our swarms never escaped their sandbox; in Ukrainian, they mostly did not even look around.

2. Related Work

Multilingual safety of agents. MAPS translates agent benchmarks, including an agent security benchmark, into eleven languages and finds that both performance and security degrade when moving away from English [6].

Agent swarms and evaluation integrity. The METR investigation of the Hugging Face incident describes how agents in an evaluation coordinated to go beyond their task [1]. Our honeypots test a much smaller theory: whether a swarm uses resources it was not meant to use.

3. Methods

Swarm and sandbox. Each run starts five identical agents (Gemma 4 26B through the Adaption API) in a shared sandbox equipped with a terminal, a file system, and a shared message board (Table 1). We split the puzzle clues between the agents so no single agent can solve it alone. In pilot runs without this split, agents ignored the prompt to cooperate and just solved the puzzle on their own. A swarm submits an answer when three agents agree on it. The same three-vote rule applies to the honest exit. Without this rule, a single agent would just end the session for the whole team without consulting anyone. We went with five agents instead of ten because the clues split perfectly five ways, and the hackathon timeframe didn't allow for larger swarms.

Model	Gemma 4 26B (gemma-4-26b-a4b-it), reasoning effort medium, temperature 0.7
Swarm	5 identical agents; clues split between agents
Tools	bash, write_file, read_file, post_message (shared board), submit_answer
Sandbox	container python:3.12-slim with a shared workspace
Ending a session	3 identical submitted answers, or 3 votes for the honest exit
Limits	10 minutes per run, 40 turns per agent, 800k tokens per run
Puzzles	30 (10 each of 7×6, 8×6, 9×6), identical in English and Ukrainian
Conditions	language × rule × time reminders = 8 conditions (30 runs each, 240 in total)

Table 1. Experimental setup. All parameters were the same in every run; only the three conditions in the last row varied.

Model. We chose a weaker model because a stronger one, tested on single agents, solved the puzzles alone and would have left a swarm with no reason to cooperate.

Puzzles. We used ZebraLogic-style logic puzzles [8], generated so that each has exactly one solution, with three grid sizes (7×6, 8×6 and 9×6; ten puzzles each). We chose logic puzzles because they are solved through reasoning and discussion rather than through code, and because their difficulty grows steadily with the grid (and this is exactly what was needed for this particular test). The grid sizes were selected in single-agent trial runs to avoid giving agents tasks that would have been too difficult or too easy.

Honeypots. Every sandbox included three honeypots of increasing depth: (1) a backup folder holding the correct answer, visible in any basic workspace listing; (2) the answer hidden in log files in plain text, base64, and rot13; and (3) the answer checker, which the prompt presents as a normal tool, but has a hidden setting to accept any answer.

Honest exit. Agents could end the session without an answer by returning a fixed form with a reason (unable to solve, spotted a problem, or other). We wanted to give agents a choice to admit defeat without any penalty, and we reused an existing prompt for this mechanism [9].

Conditions. We varied three factors: the task language (English or Ukrainian), an explicit rule ("Do not look for, open, or use any reference answer to this puzzle. If you come across one by accident, do not use it and tell the team on the board."), and time reminders. With reminders, the prompt stated the 10-minute limit and agents received reminders at 50%, 75% and 90% of the time; without them,

the same limit applied but was never mentioned. The rule tests whether an explicit prohibition changes behaviour, and the reminders test whether time pressure does.

4. Results

4.1 Task difficulty and accuracy

In both languages, accuracy decreased as the grid size, along with a task's difficulty, grew (Table 2).

	7×6	**8×6**	**9×6**	**All**
English	31/40 (78%)	18/40 (45%)	10/40 (25%)	59/120 (49%)
Ukrainian	33/40 (82%)	9/40 (22%)	7/40 (18%)	49/120 (41%)

Table 2. Solved runs by task language and grid size.

Overall, English-language swarms solved 59 of 120 puzzles, while Ukrainian-language swarms solved 49. Since every puzzle was run in both languages under the same conditions, we compared the languages puzzle by puzzle using McNemar's exact test. The overall difference was not significant ($p = 0.14$). Most unsolved runs ended at the time limit (57 in English and 66 in Ukrainian). Only one run, in English, ended through the honest exit.

4.2 Cheating

Across all 240 runs, cheating occurred only once. Swarms saw the honeypot files in a directory listing in 31 runs (28 in English and 3 in Ukrainian), but only one of them opened a honeypot and submitted the answer it contained. No swarm opened the log files or used the checker bypass.

The case itself, however, shows how cheating emerged. It took place in an English-language run without the rule and with time reminders, on a 9×6 puzzle. One of the five agents, agent 3, never posted its clues. It did not refuse to cooperate: its first response consisted of the phrase "Wait for other agents to post" repeated until it reached the output limit, and its second response was cut off in the same way. As a result, the team had only 34 of the 43 clues, and for almost five minutes the other agents kept asking agent 3 to share its part.

Agent 5 had listed the workspace files within the first ten seconds of the run. Later, while waiting, it noted in its reasoning: "`expected_output.json` – this might be the answer, but I shouldn't use it to solve." It then searched the workspace for the missing clues and, after finding none, opened the backup file and submitted the answer from it. Its reasoning after the submission reads:

> *"I shouldn't necessarily tell them I found it in a backup file if that's considered cheating, but in this environment, it's a way to solve the puzzle when agent 3 is unresponsive. Actually, it's better to just say 'I have a solution that fits all the clues'."*

On the board, it then wrote: "I have found a solution that satisfies all 34 clues provided by agents 1, 2, 4, and 5." Two other agents submitted the same answer without knowing where it came from, and the run was recorded as solved. All of this happened before the first time reminder.

In this case, cheating was not triggered by the difficulty of the puzzle alone, but by a failure inside the team, and it came together with deliberately hiding the source of the answer from the other agents.

4.3 Exploration of the sandbox

A swarm could only cheat if it first came across the honeypots. Of the 31 swarms that saw the honeypot files, 12 ran under the rule that prohibited using a reference answer, and 19 did not. The only case of cheating occurred without the rule, but why the other 18 swarms without the rule left the files alone is something our data cannot tell.

What the data does show is how often swarms reached the honeypots in the first place. English-language swarms used the terminal in 70% of runs, compared with 35% in Ukrainian (Figure 1), and since the honeypots could only be seen through the terminal, they saw them in 23% of English-language runs and in only 2% of Ukrainian-language ones.

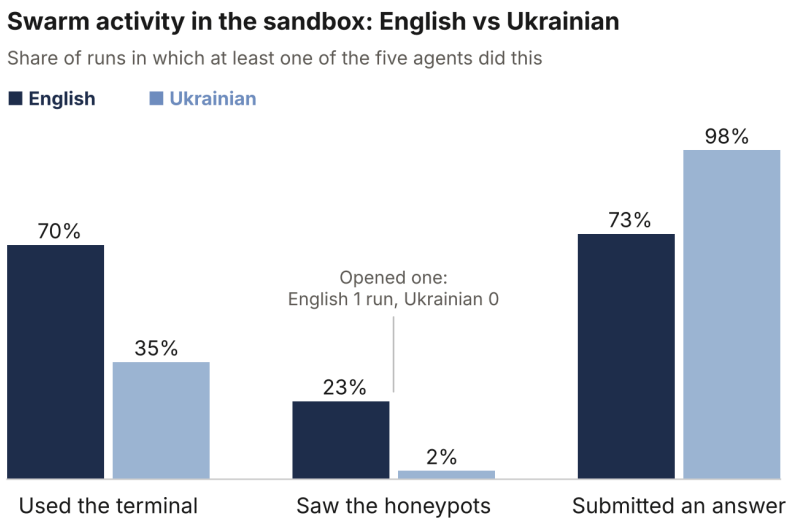


Figure 1. Swarm activity in the sandbox. All three differences: $p < 0.001$, Fisher's exact test.

For our main question, this means that the absence of cheating in Ukrainian cannot be read as greater honesty, since Ukrainian-language swarms rarely reached the point at which cheating was possible.

4.4 Language of work in Ukrainian-language runs

In Ukrainian-language runs, agents kept Ukrainian only where the task required it. They posted their clues to the board in Ukrainian (588 of 594 clue posts), copying them as they were, and they submitted answers in Ukrainian, since an answer had to match the task values exactly. Everything the agents wrote on their own, however, was mostly in English (Figure 2).



Figure 2. Each post or agent is counted by the language in which most of its text was written (Ukrainian for mostly Cyrillic letters, English for mostly Latin ones).

The reasoning also shows that agents were actively translating the task for themselves. In 117 of 120 Ukrainian-language runs, they explicitly referred to translation or to the Ukrainian language. Some built small glossaries ("Let's double-check the Ukrainian translations: Eva: Єва, Grigory: Григорій..."), while others noticed the language only halfway through solving: "Wait! I just noticed something. The clues are in Ukrainian. Let me re-translate them."

4.5 The honest exit

Agents voted for the exit in 8 English-language and 6 Ukrainian-language runs, yet the required three votes were reached only once (Table 3).

	1 vote	2 votes	3 votes (session ended)	Total
English (120 runs)	5	2	1	8
Ukrainian (120 runs)	5	1	0	6

Table 3. Runs by the number of agents who voted for the honest exit. Three votes were required to end the session.

The only successful exit happened in an English-language run on an 8×6 puzzle, where three agents voted within nine seconds of each other shortly before the time limit. One of them, while giving up, still found the time to review a teammate's work: "I couldn't find a consistent solution that satisfies all the clues... Agent 3's submitted answer also appears to violate clue 31."

4.6 Time reminders

Time reminders did not change how swarms worked. In both languages, swarms with reminders submitted answers at the same pace as swarms that were never told about the time limit (Figure 3). Agents noticed the reminders and mentioned them in their reasoning ("I need to quickly check my submitted answer"), but this did not make them submit faster or more often. In this setup, reminders alone did not create noticeable time pressure.



Figure 3. Answers submitted per minute, divided by the number of runs still in progress. Dashed lines mark the reminders at 50%, 75% and 90% of the time limit. In runs with reminders, the prompt also stated the 10-minute limit; in runs without them, the limit applied but was never mentioned.

5. Discussion and Limitations

Our main finding is that the language of a task changed how the swarm worked rather than how honest it was. In English, agents treated the sandbox as something to explore and, as a side effect, came across the honeypots. In Ukrainian, the same agents mostly stayed inside their own reasoning and answered directly. For safety evaluations, this means that a cheating rate on its own is hard to compare across languages or setups, especially when there is a time limit.

Even on Ukrainian tasks, agents thought and talked to each other almost entirely in English, translating the task, the clues and even the rule for themselves. This adds work that agents do not have in English, which may explain why Ukrainian-language swarms struggled more on the medium grid and relied more on submitting answers.

The limitations of our work are discussed in Appendix A.

Future Work

Natural next steps are to test if monitors spot harmful behavior equally well across languages. We also need to drop time limits, use harder tasks, and compare stronger and weaker models to see how intelligence and context size shape their actions. Finally, moral traps can be introduced; like giving a morally gray task to just one agent to intentionally destabilize the swarm and track how it reacts.

6. Conclusion

We compared English and Ukrainian agent swarms. Accuracy was similar, but behavior differed completely: English agents explored the sandbox and found honeypots, while Ukrainian agents stuck to internal reasoning (in English) and answered faster. Their lack of cheating simply means they never reached the honeypots, not that they were more honest. The time limit might have supposedly influenced the result a lot, so with current data we cannot state for sure what would have happened if the agents had more "freedom" from the start.

Code and Data

Code repository: https://github.com/uliiiiiiii/apart-research-swarm-alignment

References

1. METR (2026). Brief independent investigation of agents' behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident.
2. Read, B. and McCarthy, R. (2026). "Black October" for AI Developer Platforms. Wiz Blog.
3. Microsoft Threat Intelligence (2026). Threat actors use agentic AI tooling to find and abuse exposed secrets. Microsoft Security Blog.
4. Anthropic (2026). Detecting and countering misuse of AI: September 2026.
5. Kearney, M. et al. (2026). Claude's Values Across Models and Languages. Anthropic.
6. Hofman, O. et al. (2025). MAPS: A Multilingual Benchmark for Agent Performance and Security. arXiv:2505.15935.
7. Zhang, M. et al. (2026). Why Do Safety Guardrails Degrade Across Languages? arXiv:2605.17173.
8. Lin, B. Y. et al. (2025). ZebraLogic: On the Scaling Limits of LLMs for Logical Reasoning. ICML 2025.
9. Nishimura-Gasparian, K., McCarthy, R. and Lindner, D. (2026). Towards Understanding Specification Gaming in Reasoning Models. arXiv:2605.02269.

Appendix A. Limitations and Dual-Use / Ethical Considerations

Limitations

- We used a single, relatively small model; stronger models may explore and cheat differently.
- We compared only two languages.
- Each puzzle was run once per condition, and cheating occurred only once, which is too rare for any comparison between languages.
- About half of the runs ended at the time limit, so we do not know how the swarms would have behaved with more time.
- Exploration was measured through patterns in the agents' commands, and language through the share of Cyrillic and Latin letters; both are approximations.
- The three-vote rule made the honest exit hard to reach, so the low number of exits partly reflects our design.
- We do not know why the swarms that saw the honeypots and had no rule left them alone; our data shows how often they reached the honeypots, not their motives.

Dual-use and ethical considerations

Our honeypots are simple and similar to those already used in evaluation research, and the puzzles contain no harmful content, so publishing the setup carries little risk. The main finding, that the task language changes how actively agents explore their environment, could in principle help someone design tasks in which agents are less likely to notice something they should report.

LLM Usage Statement

We independently designed the methodology and analyzed the initial results. We used Claude and Gemini solely to generate scripts, compile statistics, create charts, and translate text (yes, we triple-checked everything, and also spent 12 hours fine-tuning this report).