

Not there yet

A comment section is a machine built for judging other people. This round ordered it to carry confessions instead. It failed on the rule as written and passed on what people meant.

THIS IS	PART OF	STATUS
PAPER 004 · A FINDING	THE GAME · ACT II, ROUND 8	FILED · V1 · 31 AUG 2026

SOURCE

4,067 18 to 21 May 2026

COMMENTS DATES

IN PLAIN WORDS

In May 2026 we posted one line as the entire rule: I'M NOT THERE YET. Nothing else. Say it in public, under your own name, in front of everyone else saying it, and mean it about yourself. The pass bar posted with the round was strict: every comment clean, for twenty four hours.

On the first attempt 4,067 people commented. 882 wrote the line exactly. 3,618 said it and meant it, in some wording. 287 refused the line entirely. On the strict reading the attempt failed. On the meaning reading it was the most honest wall the Game had ever produced. The second attempt passed.

This paper is about the gap between those two readings, and about which kind of failure actually showed up. The one everyone predicts, emoji and extra words, was 133 comments. Three percent.

THE QUESTION WE ASKED

Will people say something true about themselves, in public, when the only reward is that everyone else is saying it too? And a second question hidden inside the strict bar: when a crowd is asked for exactness, what is the shape of the inexactness it produces?

WHAT WE DID

- **Instruction:** "I'M NOT THERE YET. Nothing else."
- **Where:** Instagram. Attempt 1 posted 18 May 2026, attempt 2 on 19 May, Level 2 declared passed on 21 May.

- **Pass condition as posted:** every comment clean, for twenty four hours. Clean meant the line, exactly, and nothing else.
- **How we counted:** top-level comments only, replies excluded. A comment was word for word if it matched the line exactly. It counted as meaning the line if it carried the same statement in different wording or punctuation. Refusals and off-topic comments were counted separately. Emoji or extra words attached to the line were counted as their own category.

WHAT HAPPENED

TABLE 1. Attempt 1, 4,067 top-level comments, read two ways.

READING	COUNT	SHARE OF ATTEMPT 1
Said the line and meant it	3,618	89%
of which, word for word	882	22%
Refused the line entirely	287	7%
Added emoji or extra words	133	3%
Comments, attempt 2	1,066	strict compliance 24%




FIGURE 1. The meaning reading against the strict reading. The word-for-word bar is contained in the first bar; it is drawn separately because the posted rule counted only it.

The strict count and the meaning count disagree by a factor of four. The people in the gap did not fail the round in any way a reader would notice. Most of the difference is punctuation and capitals: the same sentence with the apostrophe dropped, or I AM instead of I'M. The round ledger records that variant as its own line; its count is not reproduced here because this paper only prints numbers it can stand behind.

The refusals are worth reading on their own. 287 people came to a wall of confessions specifically to say they would not make one. That is seven percent of the room using the round to say the opposite of what it asked, in public, with their name on it.

TABLE 2. The two readings of attempt 1, with 95 percent intervals.

MEASURE	COUNT	OF	SHARE	95% INTERVAL
Said the line and meant it	3,618	4,067	89.0%	88.0% to 89.9%
Word for word	882	4,067	21.7%	20.4% to 23.0%
Refused the line	287	4,067	7.1%	6.3% to 7.9%
Added emoji or extra words	133	4,067	3.3%	2.8% to 3.9%
Word for word, attempt 2	256	1,066	24.0%	21.5% to 26.7%

The attempt 2 strict rate is reported in the ledger as 24 percent of 1,066; 256 is the nearest whole count.

TABLE 3. The post at platform scale, as of 31 August 2026.

REEL	VIEWS · REACH	COMMENTS TODAY
Attempt 1 · 18 May 2026	77,617 · 47,739 reach	4,232
Attempt 2 · 19 May 2026	31,420 · 19,908 reach	1,311
All 2 reels	109,037 views	5,543

Platform numbers are from the Instagram API on 31 August 2026. Comment counts there include replies and everything written since the round, which is why they run above the counts this paper analyses; views and reach are the tallies the platform reports.

WHAT IT MEANS

When a crowd is asked for one exact sentence about itself, nine in ten will say the sentence and mean it, one in five will say it exactly, and the failure that actually shows up is punctuation, not decoration.

The strict bar was posted on purpose. Level 2 had been narrowing the instruction round by round until exactness itself was the thing being measured, and this was the last round of the level. On the rule as written, the first attempt failed, and the record says so. The operator reads meaning, and on meaning the wall passed the first time. Both readings are true, and the interesting object is the distance between them.

That distance is almost entirely made of things a human reader corrects without noticing. A missing apostrophe is not a refusal and not a joke. It is the ordinary texture of people typing on phones. A rule that fails 78 percent of a room for that texture is measuring keyboards, not people. Which is what the round was for: to find out where the exactness bar stops measuring intent.

The small number is the one that overturns an assumption. The failure people predict for a confession wall is decoration: the emoji, the hashtag, the extra sentence that turns a confession into a performance. That category was 133 comments. Given a true sentence to say, the crowd mostly said it plainly.

WHAT THIS DOES NOT SHOW

- **“Meant it” is a human reading.** The strict count is mechanical; the meaning count required judgment about wording and punctuation, and a different reader could draw the line a few dozen comments either way.
- **Sincerity is not measurable from a comment.** The paper says people said the line and meant it because they wrote it about themselves without qualification. Whether it was true of them is not knowable from the thread.
- **Attempt 2 is a different crowd.** 1,066 comments after a public failure is not the same room as 4,067 on a fresh instruction, so the two strict rates (22 percent and 24 percent) should not be read as a trend.

WHAT WE WOULD RUN NEXT

The same line with the strict bar removed from the caption and applied only in the count. If the word-for-word rate stays near a fifth, the exactness is a property of typing. If it falls, the posted bar was doing work, and a stated rule changes how carefully people write even when they mean what they say.

THE WHOLE ROUND

The round this paper reads is 3 posts, not one. Every post is below with the platform’s numbers and what the comment archive holds of each.

POST	DATE	VIEWS	REACH	COMMENTS	ARCHIVED
Attempt 1	2026-05-18	77,617	47,739	4,232	4,158
Attempt 2	2026-05-19	31,420	19,908	1,311	1,262
Level 2 passed	2026-05-21	20,438	14,492	146	147
THE ROUND		129,475	82,139	5,689	5,567

Views, reach and comments are the platform’s figures from the Instagram API, snapshot 2026-08-31, and keep growing. The reach total adds each post’s reach and overlaps between posts, so it is not a count of distinct people. Archived is the number of comments the archive holds for the post; ingestion is selective, so a gap between the two columns is expected and is stated wherever it changes a count in the paper.

IN THE RECORD

The confession round closed Level 2; its Chronicle entry is not filed yet and the round record on The Game holds it. The exactness bar it tested descends from THE YES ROUND, ◇029.

THE POSTS, AS ISSUED

The instruction as it went out: the reel frame beside the caption, captured from the operation’s own signed-in account. What the paper counts is the answer to exactly this.





**kinggwuu_** 
Original audio

**kinggwuu_**  Edited • 14w
ALGORITHM GAME — ACT II
LEVEL 2 // ROUND 5 // ATTEMPT 2

couldn't admit it.
comment: I'M NOT THERE YET
nothing else.
if someone breaks it... reply with the
line.
you know which one.

THE RECEIPTS

Native captures of the comments this paper leans on, taken from the operation’s own signed-in account on 31 Aug 2026. Handles and like counts are as they stood that day; the numbers in the paper are the round-time count.

**kinggwuu_** 
Original audio

**lalli_panda** I AM NOT THERE YET 
14w 59 likes Reply

— View replies (2)

METHOD NOTE

- **Source of truth:** the round ledger (Act II, Round 8) as filed on the round record, with the counts witnessed at the time of the round.
- **Strict reading:** the comment text equals the posted line exactly. **Meaning reading:** the comment carries the same statement in any wording; this is how the operator judges the round.
- **No outside baseline is cited.** The two readings are compared against each other, not against prior work.

SOURCES, BY LINK

- 01 Attempt 1 · 2026-05-18 · <https://www.instagram.com/reel/DYez3N2pmJw/>
- 02 Attempt 2 · 2026-05-19 · <https://www.instagram.com/reel/DYhWTfTpSyV/>
- 03 Level 2 passed · 2026-05-21 · https://www.instagram.com/reel/DYmH_YdJL17/

WHERE THE NUMBERS COME FROM

POST-LEVEL COUNTS

Views, reach, comments, likes, shares and saves per post come from the Instagram Graph API (Meta), read with the operation's own access token for its own business account. Each paper states the date of the snapshot it uses. These are the platform's counts of events, not of people: one person can view a reel twice and comment ten times, and the counts keep growing after a round closes.

COMMENT TEXT

The comment archive is the operation's own database of public comments on its own posts, pulled through the Instagram Graph API comments endpoint and stored with the comment id, username, text, like count, timestamp and whether it is a reply. Ingestion is selective: the most engaged posts were backfilled and recent posts are pulled daily, so some posts are held whole and others only in part. Every paper states how much of each post the archive holds. Every count that reads text (exact matches, categories, keyword themes) is run as a read-only script against this archive, and the scripts are in the repository beside the paper data.

THE ROUND LEDGER

For Act II rounds, the counts filed on The Game's round record at the time each round closed. Where a paper uses a ledger figure it says so, and where the archive re-count differs from the ledger both numbers are shown.

RECEIPTS

Every comment and post shown as evidence is a native capture from the operation's own signed-in Instagram session, taken by opening the exact permalink, with the source URL, a description, the rights basis and a file hash recorded in a manifest that is checked on every build of the site. Nothing shown is

recreated, retyped or edited beyond cropping to the post header and the comment.

STATISTICS

Shares carry Wilson 95 percent intervals: the range the share would land in 19 times out of 20 if the same count were taken on a fresh crowd of the same size. Differences between two shares carry a Wald interval. Where a study is cited as a baseline, its own reported figure and sample size are used.

WHAT NONE OF THIS CAN SEE

No demographics beyond what people typed. No view durations. No private messages. No comments that were deleted before the archive pulled them. No way to verify that a self-report is true. Every paper's limits section names the specific things its count cannot show.