



Cents per Smear

Paid review attacks on consumer brands in China, September 2026

One script, five infant-formula brands named at once. Product-line nicknames instead of brand names, questions instead of claims, and RMB 0.25–0.4 (US\$0.04–0.06) per ordinary negative comment. Threat Hunter tracked 613 paid attack tasks in September 2026. What they show is a piece-rate business built to look like genuine customer experience.

613

PAID ATTACK
TASKS

189

BRANDS TARGETED
ACROSS 14 INDUSTRIES

34%

OF BRAND MENTIONS
USE A CODENAME

\$0.04–0.06

PAY PER ORDINARY
NEGATIVE COMMENT

KEY FINDINGS

What we found

Read this first

Quoted comments are excerpts of paid attack scripts, reproduced only to show how the attacks work. They are not statements of fact and do not reflect Threat Hunter's views. A brand appearing here is a target; it does not mean the brand has the problems the scripts allege. A "beneficiary" is the party a script promotes; this does not imply it took part or knew. Brand names inside quotes are masked as [Brand] in this edition. This brief is statistical analysis, not a finding of fact about any party.

Paid review attacks ("troll farms" or water-army campaigns, in Chinese industry slang) are cheap, scripted and aimed where buyers lean hardest on other people's experience. They hide behind codenames, questions and "I switched and it got better" stories, so brand-name monitoring and post-by-post takedowns stay one step behind.

74

Infant formula is the top target

Infant formula drew 74 attack tasks, the most of any sub-category. The four most-attacked brands are all formula brands.

34%

Brands named without their name

34% of brand mentions use a product-line name, abbreviation or pun. Counted by task, 36% of the 430 brand attacks never spell the brand name at all.

71%

Health scares cut deepest

71% of mother & baby tasks carry a health or safety claim; 53% in health & medical (Fig. 2).

28%

One script, many targets

28% of the 430 brand attacks disparage two or more brands at once; one script named six.

415 / 613

Xiaohongshu is the main front

68% of tasks land on Xiaohongshu (RedNote), in the comments under reviews, recommendation requests and "avoid this" posts.

\$0.04–0.06

Piece-rate pricing

A standard negative comment pays RMB 0.25–0.4 (US\$0.04–0.06); one abuse report RMB 0.25–0.3 (US\$0.04). A "must stay visible" comment pays about RMB 0.55 (US\$0.08).

Counting rules: 613 attack tasks, one per unique script. Of the 613, 430 are "brand attacks" (tasks naming at least one brand); the rest target categories or institutions, or name no brand. Industry, brand and platform categories overlap: one task can name several brands or span several platforms, so category counts do not add up to 613. "Brand mentions" are counted per mention. Each figure says which base it uses. Amounts are converted at US\$1 = RMB 6.75 and rounded to the nearest cent.

01 · LEAD CASE

One script, five formula brands

Infant formula · switch story · codenames · multi-brand script

Infant formula was the most heavily attacked sub-category in September, with 74 tasks. Friso (49), Aptamil (34), a2 (30) and Yili Jinlingguan (20) were the four most-attacked brands overall.

The attacks read like parents comparing notes: "we were on X and the baby...", then "after switching to [Brand]...". Slow weight gain, constipation, "internal heat", eye discharge, undigested curds: each symptom is pinned on a different rival. In one task, five lines named five brands at once.

61% of formula-brand mentions used a product-line name or nickname, not the brand. A monitor keyed to brand names misses most of them.

74

formula attack tasks

8

formula brands targeted

5

brands named in one script

57 / 74promote the same
beneficiary

Attack-script excerpts, translated · not statements of fact

After [line name] the baby kept getting eye discharge and "internal heat"; at six months we switched back to [Brand] and it cleared up.

Target: infant-formula brand, named by product-line nickname

Three months on [line name] and only one jin (500 g) gained. Already planning to switch.

Target: infant-formula brand, named by product-line nickname

On [line name] the stools were always hard. Two months on [Brand] and everything's normal.

Target: infant-formula brand · same task as above

Why it matters Formula is a long, repeat purchase. Once "it caused constipation and no weight gain" becomes received wisdom in parenting communities, a lost customer rarely comes back. 57 of the 74 formula tasks promoted the same brand (a beneficiary is not assumed to be involved or aware).

02 · EXPOSURE

Categories bought on word of mouth get hit first

Mother & baby (93 tasks) and phones & electronics (87) carry the highest exposure: together 42% of the 430 brand attacks. Beauty, home appliances, health, autos/EV and pet sit in the middle band. Busiest sub-categories (tasks): infant formula 74, phones 36, cameras & imaging 27, pet food 25, air & water appliances 21, skincare 18, body care 14, new-energy vehicles 13.




FIG. 1 Attack tasks by industry. Exposure bands: high 60+, medium 25–59, low below 25.

Each industry gets hit from a different angle: health and safety in mother & baby (71%), poor experience in beauty (75%), price and value in autos/EV (54%), "doesn't work" in health & medical (63%). Defences should follow the angle.




FIG. 2 Claim types by industry. Describes what the scripts say, not the targeted brands' actual quality.

02 · EXPOSURE (CONT.)

Where and when attacks land

68% of tasks (415 of 613) land on Xiaohongshu, the lifestyle-review app also known as RedNote; Douyin, China's TikTok, follows with 137.




FIG. 3 Platforms (tasks; one task can span several).

35% of tasks launch between 19:00 and 23:00 China time, peaking at 19:00 (55 tasks). Monitoring should cover evenings.




FIG. 4 Task launch time by hour, China Standard Time (UTC+8).

03 · MOST-TARGETED BRANDS

Top 20 brands under attack

189 brands were targeted in September, plus 19 generic categories (including drug names and ingredients) and 27 institutions or stores. Phones & electronics placed the most brands in the top 20 (six). Over 80% of attacks on Yili Jinlingguan and Friso named them only by product-line names.

#	BRAND / SEGMENT	TASKS	DAYS HIT	CODENAME SHARE	MAIN MODE	MAIN ANGLES
1	Friso Infant formula	49	11	84%	Comparative	Doesn't work; health/safety
2	Aptamil Infant formula	34	12	50%	Comparative	Doesn't work; health/safety
3	a2 Infant formula	30	11	37%	Comparative	Health/safety; doesn't work
4	Yili Jinlingguan Infant formula	20	10	85%	Comparative	Doesn't work; health/safety
5	Apple Phones	16	9	—	Comparative	Experience; price/value
6	Feihe Infant formula	13	7	23%	Comparative	Health/safety; doesn't work
7	DJI Cameras & imaging	13	9	15%	Comparative	Experience; price/value
8	CATL EV batteries	13	6	15%	Insinuation	After-sales; defects
9	Xianlang Pet food	12	7	—	Comparative	Price/value; experience
10	vivo Phones	12	6	58%	Comparative	Price/value; false ads
11	Sofy Feminine care	11	4	—	Comparative	Experience; health/safety
12	Insta360 Cameras & imaging	10	7	60%	Comparative	Experience; false ads
13	Xiaomi (phones) Phones	9	6	67%	Comparative	Experience; defects
14	Xiaomi (appliances) Air & water appliances	9	7	22%	Comparative	Experience; health/safety
15	Freemore Feminine care	8	4	—	Comparative	Experience; health/safety
16	Huawei Phones	6	4	—	Direct negative	Experience; defects
17	ABC Feminine care	6	3	—	Comparative	Experience; health/safety
18	Royal Canin Pet food	6	5	—	Comparative	Experience; health/safety
19	Midea Air & water appliances	6	3	33%	Comparative	Defects; health/safety
20	Space7 Feminine care	6	1	—	Comparative	Experience; health/safety

Ranked by attack tasks. Codename share = mentions using a line name, abbreviation or pun; "—" = none seen. Comparative = disparages the target while promoting another; insinuation = target implied, not named. Listing means a brand was targeted, not that the allegations are true.

04 · TRADecraft

Eight tactics

The scripts follow a few templates, all built to pass as real customer experience. Bases: "tasks" = 613; "brand attacks" = the 430 tasks naming a brand; "mentions" = per mention. One task can use several tactics.

TACTIC	SHARE	WHAT IT LOOKS LIKE	WHY IT'S HARD TO STOP
Codenames	34% of mentions	Product-line names, abbreviations, puns, even emoji in place of the brand.	Slips past keyword monitoring; brands notice after the damage.
Switch story	Common	"X gave me trouble; after switching to [Brand] it was fine."	A single comment can't be disproved.
Scripted threads	Common	Several accounts play asker and answerer in a staged "organic" discussion.	Manufactures consensus.
Multi-brand scripts	28% of 430 brand attacks	One script disparages several rivals; up to six.	Category leaders get hit together.
Loaded questions	Common	"Does X give off ozone?" instead of a claim.	Not a statement of fact, hard to report; real users pile on.
Impersonation	Common	Posing as a parent, car owner, patient's relative, long-time customer.	Identity buys empathy.
Hijacking	10% of 430 brand attacks	Posted under the target brand's own reviews and promotional posts.	The brand's comment section converts for a rival.
Reverse accusation	Emerging	Accusing the rival of buying reviews: "the comments are all bots."	Discredits genuine praise.

→ How codenames work

Chinese wordplay makes codenames easy to coin and hard to catch. Examples from September scripts: a color nickname for a phone maker ("blue factory" for vivo), a mocking pun on a 3D-printer brand, a light-bulb emoji for a skincare brand, and drug names with a character swapped for a placeholder or a sound-alike, as with ibuprofen and tirzepatide. Monitoring keyed only to brand names misses all of them.

05 · ECONOMICS

A few cents a smear

Attacks are paid by the piece. At RMB 0.25–0.4 (US\$0.04–0.06) per ordinary comment, 1,000 negative comments cost about RMB 250–400 (US\$37–59) in labor. A brand can take down a batch, and the attacker can buy the next batch for pocket change.

Standard negative comment

\$0.04–0.06 RMB 0.25–0.4 each

Comments under posts and videos

One abuse report

\$0.04 RMB 0.25–0.3 each

Mass-reporting accounts, posts or complaints

Livestream negative comment

\$0.04 RMB 0.25–0.3 each

Flooding a livestream with negativity

"Must stay visible" comment

\$0.08 RMB 0.55 · up to RMB 0.6 (\$0.09)

Must not be folded away by the platform

The higher rate for comments that must stay visible, about RMB 0.55 (US\$0.08) and up to RMB 0.6 (US\$0.09), suggests platform folding and moderation are biting, and attackers pay extra for exposure. A folded comment still happened.

Attacks often come bundled with competitive promotion: the script disparages rivals while pushing another brand.

What to do Takedowns alone treat symptoms. Detect batches of structurally identical scripts, preserve evidence, and work with platforms to act on the executing accounts in bulk.

06 · CASES

Four more attack patterns

CASE 02 · Feminine care

Six sanitary-pad brands in one script

12 tasks targeted feminine care, touching nine sanitary-pad brands; Sofy appeared in 11, Freemore in 8. The most typical script named six brands at once, framing "first-hand" travel, exam-season and classroom experiences, each ending with "more breathable than X" or "leaks less than X". Claims about leaks and stuffiness can't be refuted with specs, but they move repeat purchase.

CASE 03 · Health & medical

Medicines as stepping stones

12 tasks targeted nine generic drugs or ingredients, including metformin, dydrogesterone, ibuprofen, tirzepatide and acarbose. Scripts tie them to "terrible side effects" or "treats symptoms, not the cause", then steer readers to sprays, probiotics, supplements or patches. Medical nutrition products were drawn in too, with one script posing as a chemotherapy patient's relative. Drug names were deliberately misspelled to stay recognizable while dodging filters.

Why it matters These scripts nudge people with chronic conditions, reproductive-health needs or weight-loss treatment to drop or replace prescribed medication. That is a patient-safety issue, not just competition.

CASE 04 · Air & water appliances

"Does it give off ozone?"

21 tasks and nine brands, built on two words: "ozone" and "registration". Scripts ask questions, or claim to have "checked the report" and found that only one brand emits no ozone at all. Brands drawn in included Xiaomi appliances (9 tasks), Midea (6) and A.O. Smith (4).

Bought [Brand], it says zero ozone on the health-commission site, how do you check? / Reply: does the [Brand] one have ozone?

Loaded question · target: air-purifier brand

CASE 05 · Silencing complaints

Consumer complaints, mass-reported

7 tasks went after consumers' complaint and "avoid this" posts, covering installation fees, water-purifier claims, repairs, noisy exercise bikes, bundled car options and livestream sellers. Pre-written report templates came in several versions ("false information", "stirring up conflict"); others recast complaints as "a competitor smearing them." Genuine feedback ends up treated as rumor. Even a brand that knew nothing faces backlash for "silencing complaints" if exposed. (Report templates are not reproduced.)

07 · BEYOND BRANDS

Not only brands get hit

About three in ten tasks name no brand at all. They target the product in a post without naming it, consumers' complaint posts, or individual people. "This one" or "that model" attacks are even harder to catch by keyword, and some mass-reporting suppresses genuine negative voices.




FIG. 5 Non-brand targets by type (tasks).

08 Method

Attack task: one paid action posting specific negative content about, or organizing reports against, an identifiable target. Each unique script counts once; pure promotion does not count.

Targeted brand: line names, models, abbreviations and puns roll up to the brand; brands within one group count separately; generic drug names and ingredients count as categories.

Codename: pointing at a brand without its full name.

Claim type: the kind of negative claim in a script; one task can have several. Describes the script only.

Exposure: high 60+ tasks, medium 25–59, low below 25.

Beneficiary: the party a script promotes; not presumed involved or aware.

Currency: amounts converted at US\$1 = RMB 6.75, rounded to the cent.

Period: September 2026.

Read this first

Quoted comments are excerpts of paid attack scripts, reproduced only to show how the attacks work. They are not statements of fact and do not reflect Threat Hunter's views. A brand appearing here is a target; it does not mean the brand has the problems the scripts allege. A "beneficiary" is the party a script promotes; this does not imply it took part or knew. Brand names inside quotes are masked as [Brand] in this edition. This brief is statistical analysis, not a finding of fact about any party.



THREAT HUNTER

Catch the batch, not the comment

Paid attacks cost cents. A defence that works one comment at a time will always be late. Threat Hunter monitors paid attack tasks and scripts so brands can expand codename watchlists, cluster identical scripts, preserve evidence and act with platforms.

Source: Threat Hunter paid-attack monitoring, September 2026. <https://www.threathunter.com>