

AI CITATIONS SERIES · NO. 4

One engine's favourite. Everyone else's unknown.

The most-recommended brand across our 39 buyer questions is **Toptal** — 10 mentions, more than any other company. Every single one came from ChatGPT. Perplexity and Gemini have never heard of it. And Toptal is not the exception: **79% of brands exist in exactly one engine's world.**

39 buyer questions · 117 answers · 646 distinct brands · July 2026 · US / English · **all three engines**

646 distinct brands named across the study all 3 engines	79% of brands are named by exactly one engine 509 of 646
8% of brands are named by all three just 51 companies	1 of 3 engines have ever named the #1 brand ChatGPT only

A rare one: no engine caveat. Most of this series is Perplexity-only, because ChatGPT and Gemini don't expose their citations. This report reads brand *names* out of the answer text — which all three engines give us — so every figure here uses all three.

FINDING 01

Toptal: number one, and two-thirds invisible

Toptal was the most-named brand in the study — 10 mentions across 117 answers, more than any other company. It is also, as far as Perplexity and Gemini are concerned, not a thing that exists.

● ChatGPT

10

mentions

● Perplexity

0

mentions

● Gemini

0

mentions

Sit with what that means. If you benchmark your AI visibility against Toptal — the category's apparent AI darling — and you only ever check ChatGPT, you'd conclude they have solved this. Check Perplexity and they're nowhere. **Two-thirds of the market cannot see the market leader.**

You cannot know how you're doing from one engine.

Not directionally. Not roughly. The #1 brand in this study scores 10 on one engine and 0 on the others.


FINDING 02 · THE NORMAL CASE

79% of brands exist in exactly one engine's world

Toptal is not a quirk. It's the normal case, drawn in bold.

How many engines know each brand?

646 distinct brands · all 3 engines · 39 questions



Of 646 brands, only **51 are known to all three engines**. Four out of five exist inside exactly one engine's answers.

This is the "three separate games" finding from our first report, restated at the level that actually matters to you. It isn't just that the engines return different *lists*. It's that the overwhelming majority of companies in this market are visible to exactly one of the three — and because none of this shows up in analytics, almost none of them know which.

FINDING 03 · THE TWIST

Fame is shared. Obscurity is fragmented.

Here's what stops this being a doom story. Look at the 20 most-named brands and the pattern inverts: **18 of the top 20 are known to two or three engines**.

Finsweet (9 mentions), HubSpot (7), ActiveCampaign (7), Flow Ninja (6), Pipedrive (6), Zoho CRM (6), Freshdesk (6), Pendo (6) — all three engines name every one of them. Only Toptal and one other are single-engine acts.

So the market has two regimes:

The top
51 brands · 8%
A small shared set all three engines have converged on. Enough presence that every engine's retrieval finds them.

The tail
509 brands · 79%
Each living inside exactly one engine's answer set, invisible to the other two.

The gap between the tail and the top isn't fame. It's coverage.

The implication is not "you can't win." It's that cross-engine presence is the thing that separates the two regimes — and it's a *reachable* state, because the shared set is 51 companies deep, not five.

FINDING 04

The engines have different appetites

They don't even name the same *number* of companies.

ENGINE	BRANDS PER ANSWER (MEAN)	RANGE	DISTINCT BRANDS
ChatGPT	8.9	0–13	296
Perplexity	8.3	2–18	292
Gemini	7.2	0–23	246

ChatGPT is the most consistent recommender — it names the most brands on average and never went above 13. Gemini is the most volatile: it named nothing at all on some questions and 23 companies on another. Perplexity never returned an empty answer.

A practical read: Gemini's shortlists are shorter on average, which makes each slot scarcer and being absent more costly.

WHAT TO DO ABOUT IT

Four things this data actually supports

01

Never measure on one engine

The most-recommended brand in this study scores 10 on ChatGPT and 0 elsewhere. One engine tells you nothing about the other two.

02

Aim for "named by all three"

Not "ranked #1 on one." Only 8% of brands manage it — that's the moat, and it's 51 companies wide, not five.

03

Find out which engine sees you

If you're in the 79%, you don't know. Optimising for the engine that already names you is the most common wasted effort in this category.

04

Don't copy the visible leader

They may be a one-engine act, and you'd be copying a strategy two-thirds of the market can't see.

THE DATA

Every number, in one place

METRIC	SCOPE	VALUE
Distinct brands named	all 3 engines, 39 questions	646
Named by exactly 1 engine	all 3 engines	509 = 79%
Named by exactly 2 engines	all 3 engines	86 = 13%

METRIC	SCOPE	VALUE
Named by all 3 engines	all 3 engines	51 = 8%
Most-named brand	all 3 engines	Toptal · 10 mentions
Engines naming Toptal	all 3 engines	1 of 3 (ChatGPT)
Of the top 20, single-engine	all 3 engines	2 of 20
Brands per answer — ChatGPT	39 answers	mean 8.9 · range 0–13
Brands per answer — Perplexity	39 answers	mean 8.3 · range 2–18
Brands per answer — Gemini	39 answers	mean 7.2 · range 0–23
Distinct brands — ChatGPT	39 answers	296
Distinct brands — Perplexity	39 answers	292
Distinct brands — Gemini	39 answers	246

Limits, plainly. n=39 questions, 117 answers, one snapshot (July 2026). A brand named once by one engine counts as "named by one engine" — at this sample size there is real noise in the tail, and the 79% will move with n. The direction is robust; the decimal isn't. **"Named by" is a low bar** — one mention in one answer. This measures presence, not prominence. **Naming variants are unfolded** — "Monday Work Management" and "Monday.com" count separately, which would slightly *inflate* the single-engine share. Folding the 8 variant groups we found moves the series' headline agreement figure by +0.0 points. **Model versions move** — sonar / gpt-4o-mini / gemini-2.5-flash, July 2026. The full dataset ships with this report, and `verify-numbers.py` recomputes every figure above straight from it.

QUESTIONS THIS REPORT ANSWERS

Quick answers, straight from the data

Do the AI engines recommend the same brands?

Mostly not. Of 646 brands named across 117 answers, 79% were named by exactly one engine; only 51 brands (8%) are known to all three.

What was the most-recommended brand — and who recommends it?

Toptal, with 10 mentions — every single one from ChatGPT. Perplexity and Gemini never named it once.

How many brands does each engine name per answer?

ChatGPT a mean of 8.9 (range 0–13), Perplexity 8.3 (2–18, never empty), Gemini 7.2 (0–23, the most volatile).

Which engines can see you?

Same method, pointed at your domain. We'll ask all three engines your buyers' questions and show you which ones name you — and which ones name someone else.

[Run the free audit →](#)

The 2026 AI Citations series · No. 4

By HarperFlow · July 15, 2026 · every figure recomputes from the released dataset

Download the dataset

All six reports →