The Context Report

Ask AI the same question in **two cities.**
You get two different answers.

The findings at a glance

Seven things we learned when we stopped testing a blank slate.

2.9 of 5

of the top 5 brands change once you leave New York and English. Same question.

77%

of all 206 brands we saw appeared in only one country. Local, not global.

4.4 vs 1.5

money tools shift far harder than everyday tools, out of 5.

35%

of the time, five identical runs of one question agreed on the top pick. The rest wobbled.

3.2 of 5

brands shared between ChatGPT and Gemini on the same question in the same city.

42%

of Germany's cited sources were German pages. AI draws on local sources to answer questions in a local context.

01. Geography: the answer moves by city

05. The two engines disagree

02 The full heatmap, every category and city
03 The long tail of local brands
04. Language locks the answer in
06. The same question wobbles
07. The sources are local too
08. What it means, and Context Authority

Why this matters now

Buying now starts inside the chat.

Your buyer no longer opens ten tabs. They open one chat window, ask, and act on what the AI hands back. The shortlist is built inside a room your website never sees.

94%

of B2B buyers used AI in their most recent purchase.

Forrester Buyers' Journey Survey, 2026

51%

of software buyers now start research in an AI chatbot, not a search engine.

G2, 2026

1 in 3

bought from a brand they had never heard of, because the AI named it.

G2, 2026

The blind spot everyone shares

Every tool that scores AI visibility works the same way. It asks a plain question, in English, from a blank user with no city, no language, and no story. Then it hands you one number and calls it your visibility. But no real buyer is blank. Change the city or the language, and the AI changes its answer. We ran the test nobody runs, and here is what fell out.

A single visibility score measures
a buyer who does not exist.

Finding one // Geography

The more the money matters, the more AI changes its mind.

We split the ten categories into money tools, like payroll and payments, and everyday tools, like video calls and storage. The money tools swing far harder by country. This is the same question every time.

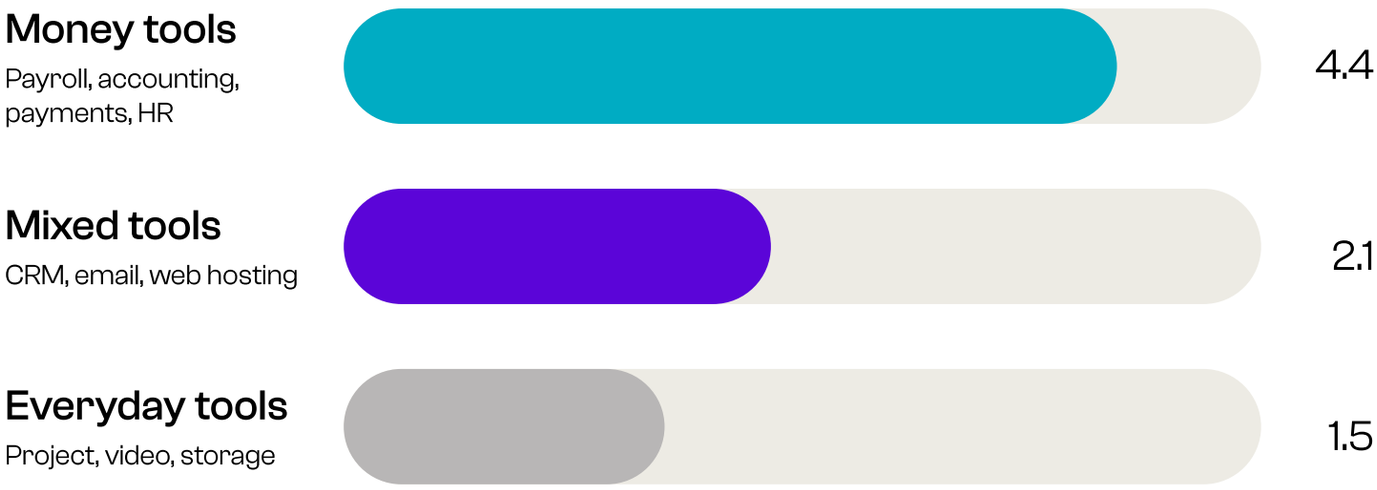

Bars show how many of the top 5 brands change against the New York answer, out of 5.

Every category, ranked by how much it moves

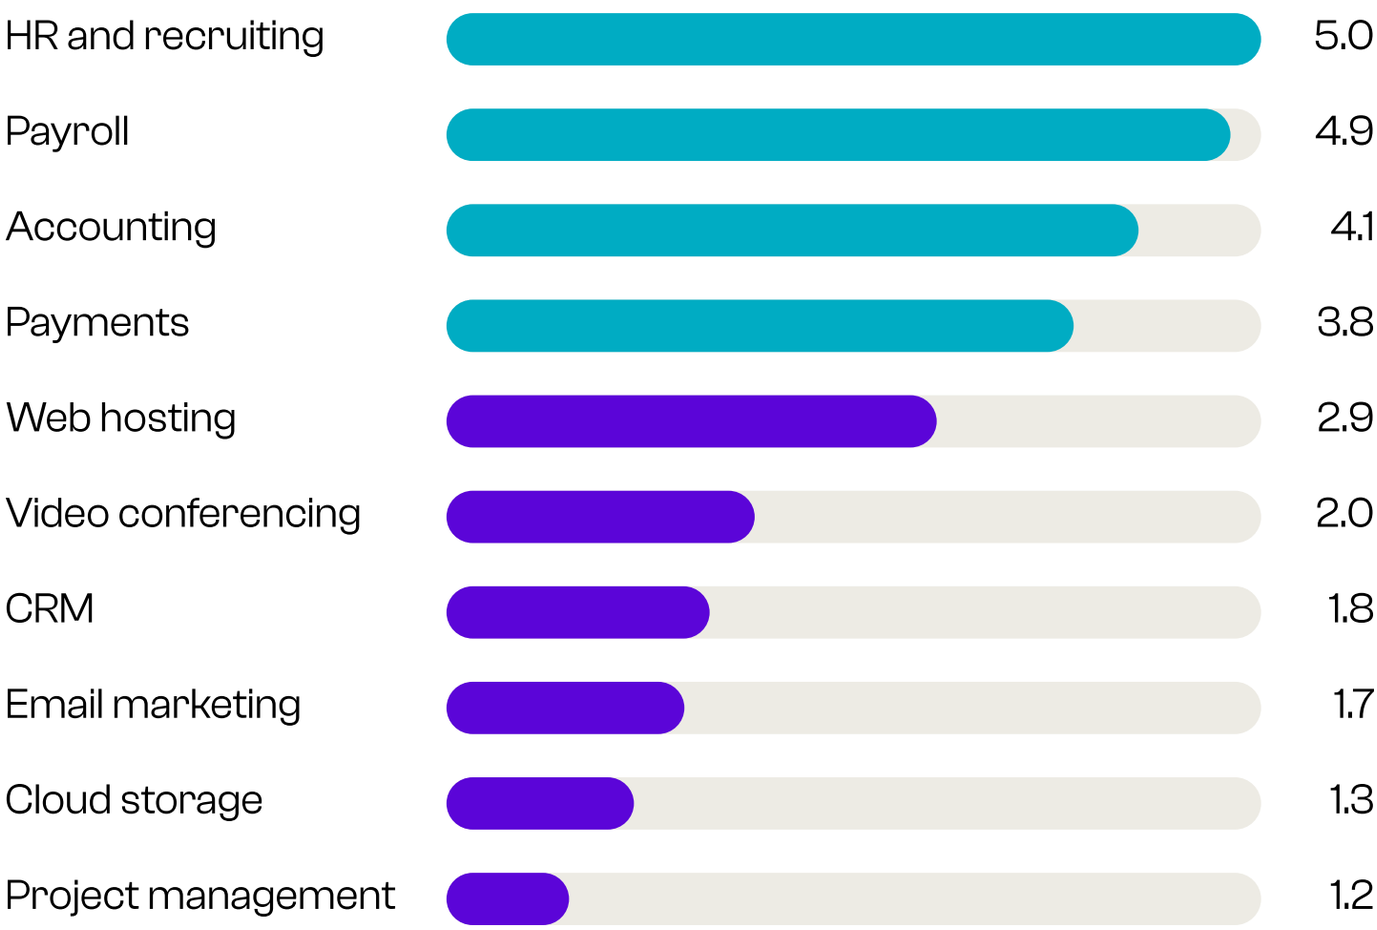

How far each city drifts from New York

Averaged across all ten categories, this is how different each city's answers are from the New York list. Berlin drifts most. London least, because it shares both the language and many of the same global brands.

Finding two // The whole map at once

One picture of the whole study.

Every category, every city. The darker the square, the more the top 5 brands change from New York. Read down the money rows and the map lights up. Read the everyday rows and it stays calm.

Category	New York	London	Mumbai	Delhi	Bengaluru	Mexico City	Madrid	Dubai	Paris	Berlin

HR / recruiting	0	5	5	5	5	5	5	5	5	5
Payroll	0	4	5	5	5	5	5	5	5	5
Accounting	0	3	4	4	4	5	5	2	5	3
Payments	0	2	5	5	5	4	3	4	3	3
Web hosting	0	2	2	2	2	3	3	3	4	5
Video conferencing	0	2	2	2	2	2	2	2	2	2
CRM	0	1	2	2	2	2	2	1	2	2
Email marketing	0	2	1	1	2	2	2	1	2	2
Cloud storage	0	1	1	1	1	2	2	0	1	3
Project management	0	1	1	2	2	1	1	1	1	1

Brands changed vs New York: 0 1 2 3 4 5

The leaderboards behind the squares

Local names, in teal, never show up in the New York answer.

01.	Payments · NYC	Stripe, Square, PayPal, Helcim, Shopify
02.	Payments · Mumbai	Razorpay, Cashfree, PayU, PhonePe, Instamojo
03.	Payments · Dubai	Telr, Stripe, Tap Payments, PayTabs, Network Intl
04.	HR · NYC	Gusto, Rippling, BambooHR, ADP, Homebase
05.	HR · Berlin	Personio, Recruitee, Factorial, Kenjo, Softgarden
06.	HR · Mumbai	Zoho People, GreytHR, Keka, Freshteam, Pocket HRMS
07.	Accounting · NYC	QuickBooks, Xero, FreshBooks, Zoho, Wave
08.	Accounting · Berlin	Lexware, sevDesk, BuchhaltungsButler, DATEV, Papierkram
09.	Accounting · Mumbai	Zoho, TallyPrime, Vyapar, Busy, myBillBook

In teal: a brand that never appears in the New York top 5.
Full leaderboards for all ten categories are in the data pack.

Finding three // The long tail

206 brands showed up. Most of them, in one country only.

This is the shape of the whole market. A short head of names that appear everywhere, and a long, heavy tail of local champions a global tool will never surface.

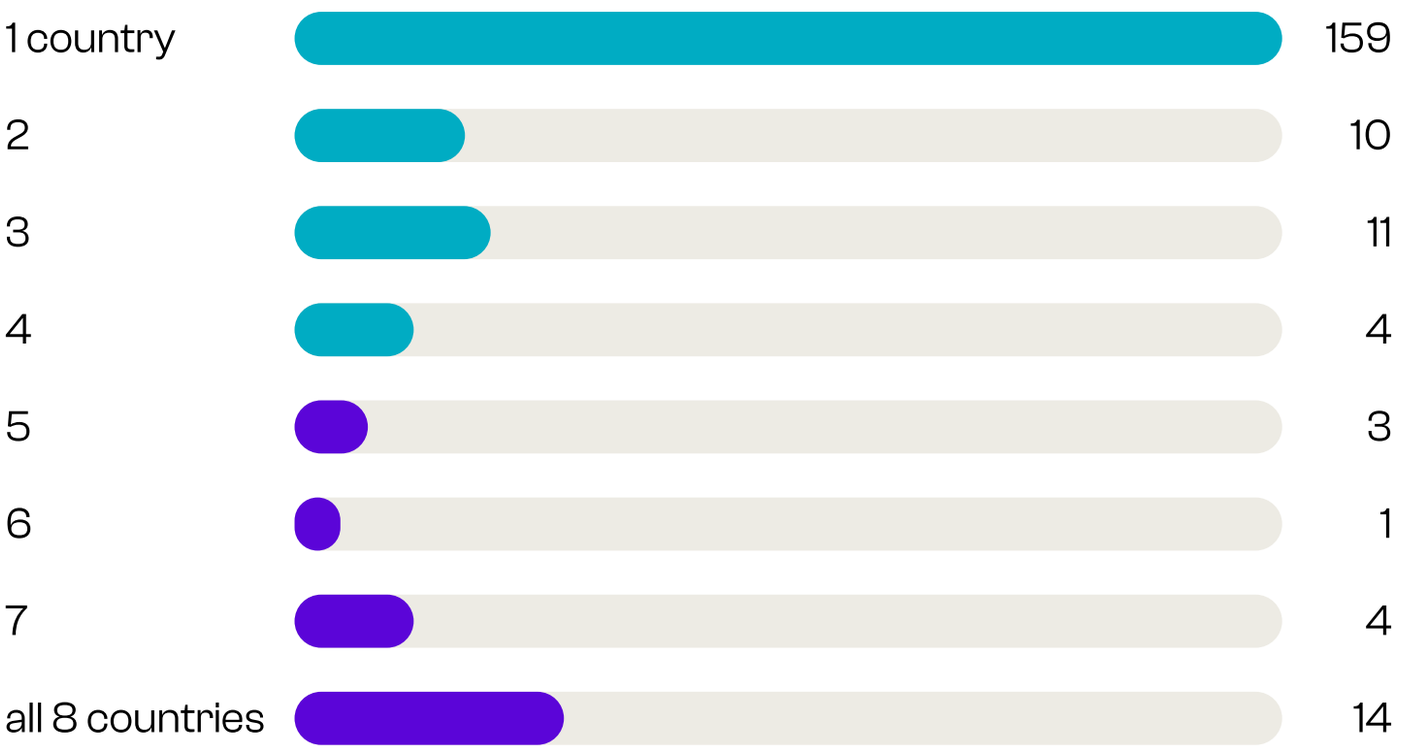

Number of brands, by how many of the 8 countries they reach a top 5.

The market splits in two

Plot every notable brand by how many markets it reaches and how dominant it is in its best one, and two clean groups appear. A short cluster of global giants, top right, that AI trusts almost everywhere. A wide cluster of local champions, top left, that own one market and vanish from the rest. The bottom right, broad reach but weak pull, is nearly empty. Brands are either global and strong, or local and strong. Very little in between.

Visibility in its best market, high to low

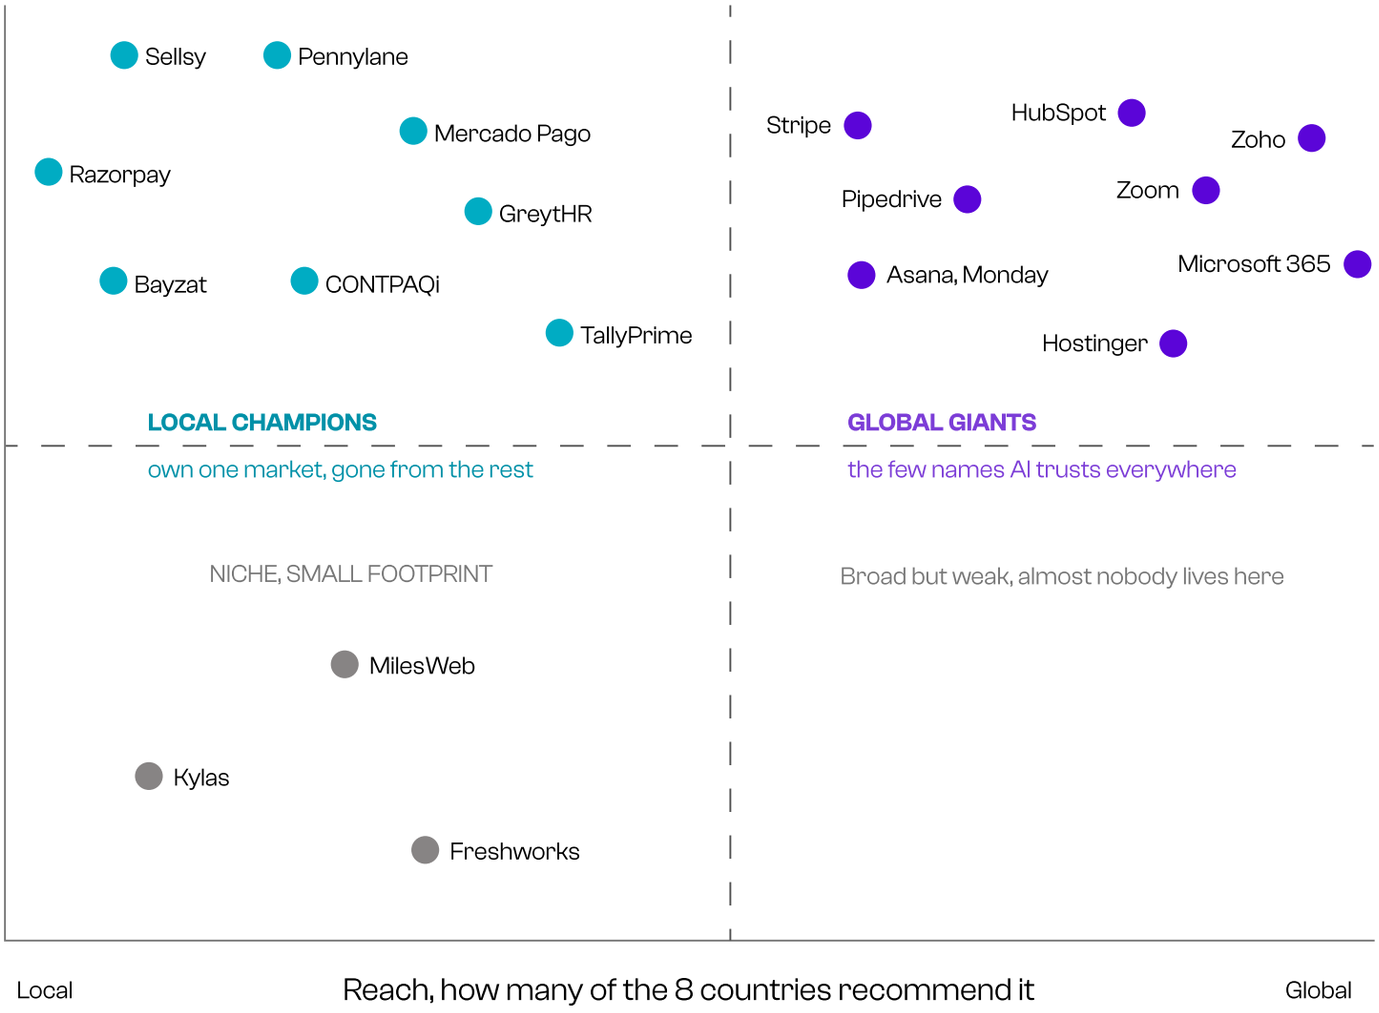

Each dot is a brand. Left to right, the number of countries where it reaches a top 5. Bottom to top, how dominant it is in its strongest market. A sample of notable brands is shown for clarity.

The global constants: only 14 brands appear everywhere

And notice what they are. Almost all of them are everyday tools, the calm rows on the heatmap. The tools people buy on money and local rules almost never make this list.

Zoho	HubSpot	Google Workspace	Microsoft 365	Zoom	Microsoft Teams	Google Meet
Asana	ClickUp	Monday	Trello	Mailchimp	Brevo	MailerLite

One brand's journey: Stripe

Stripe is a global name. Watch it win across the West and disappear completely in India, where local rails rule.

New York	London	Madrid	Paris	Berlin
01	01	01	01	01
Mexico City	**Dubai**	**Mumbai**	**Delhi**	**Bengaluru**
02	02	Gone	Gone	Gone

The local champions your dashboard hides

A sample of the 159, by country. In their home market, they beat the brands you know.

India	Razorpay, Cashfree, PayU, Instamojo, TallyPrime, Vyapar, Busy, GreytHR, Keka, Kylas, LeadSquared, MilesWeb, BigRock
Germany	Lexware, DATEV, sevDesk, Personio, IONOS, STRATO, Hetzner, Raidboxes, CentralStationCRM, weclapp, CleverReach, Kenjo, awork
France	Sellsy, Pennylane, PayFit, Silae, Axonaut, Lucca, Flatchr, o2switch, OVHcloud, Infomaniak, PlanetHoster, PayPlug, Indy
Spain	Holded, Anfix, Quipu, Contasol, Sesame HR, Bizneo, Clientify, Redsys, Raiola Networks, Webempresa, Acumbamail
Mexico	Mercado Pago, Conekta, Clip, Openpay, CONTPAQi, Alegra, Bind ERP, Aspel, Buk, Runa, Neubox, Worky
UAE	Bayzat, Telr, PayTabs, Tap Payments, Network Intl, ZenHR, Yomly, Wafeq, AEserver, Ozone

Finding four // Language

Where you are sets the list. The language you speak locks it in.

Geography does most of the work. Language pushes the answer further home. We asked the same Mumbai buyer three ways: English, Hindi, and Hinglish, the mix real Indians actually type.

English	Hindi	Hinglish
Razorpay	Razorpay	Razorpay
Cashfree	Cashfree	Cashfree
PayU	PayU	PhonePe
PhonePe	Paytm	PayU
Instamojo	Instamojo	Instamojo

All three point to Indian brands, and all three look nothing like New York. Ask in Hindi and a homegrown name like Paytm slips in. The lesson is not that language flips everything. It is that language never pulls the answer back toward the global default. It only pulls it further local. A score built on English alone misses how your buyer really asks. And Hinglish is a test no United States report can run, because they do not have the market or the mixed tongue. It is ours alone.

Finding five // The engines

ChatGPT and Gemini do not agree with each other.

Ask the same question, in the same city, and the two engines share only about 3 of their top 5. There is no single true list to track. It shifts with the engine too, and it shifts most in the money categories.

Lower bar means the two engines agreed on fewer brands. Cloud storage and web hosting split the most.

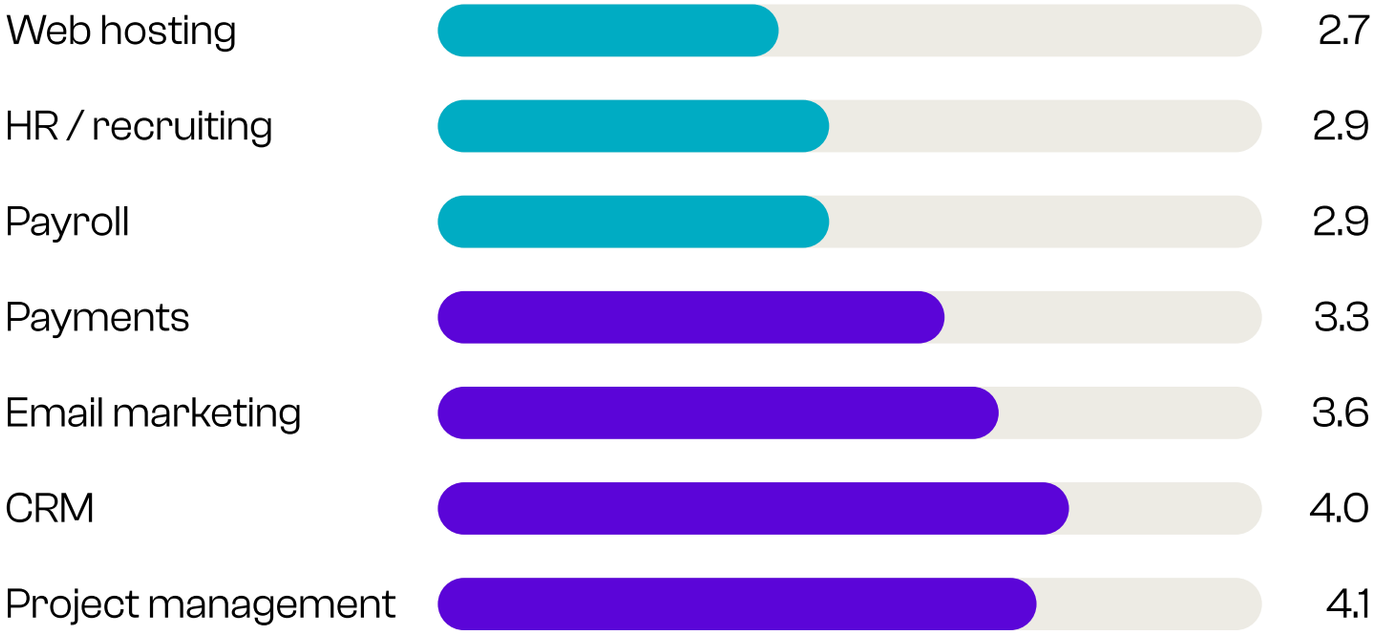

Example: for CRM in Paris, Gemini puts Sellsy first. ChatGPT puts Pipedrive first. Same city, same question, different winner.

Finding six // The wobble

Even the same question, asked five times, does not give the same answer.

This is the quiet killer for any tracker built on one snapshot. We asked each question five times with nothing changed. The top pick held all five times in only 35% of cases. The rest of the time, it moved.

35%

of cells where all five identical runs named the same top brand.

FTA, The Context Report, 2026

02

different brands took the number one spot across five runs, on average.

FTA, The Context Report, 2026

76%

of runs matched the most common pick. So even the winner is only a strong maybe.

FTA, The Context Report, 2026

What one snapshot hides

A tracker that asks once and reports a number is reading a single roll of the dice. The brand it shows you at the top may lose the top spot on the very next ask. This is why a real measure has to run many times and report the pattern, not a moment.

A number from one ask is a guess wearing a suit.

Finding seven // The sources

The AI does not just change the brands. It changes the pages it reads.

To pick a brand, the AI leans on the web pages it trusts. Outside the United States, those pages are increasingly local. The further from New York, the more it reads homegrown review sites, local media, and local company pages.

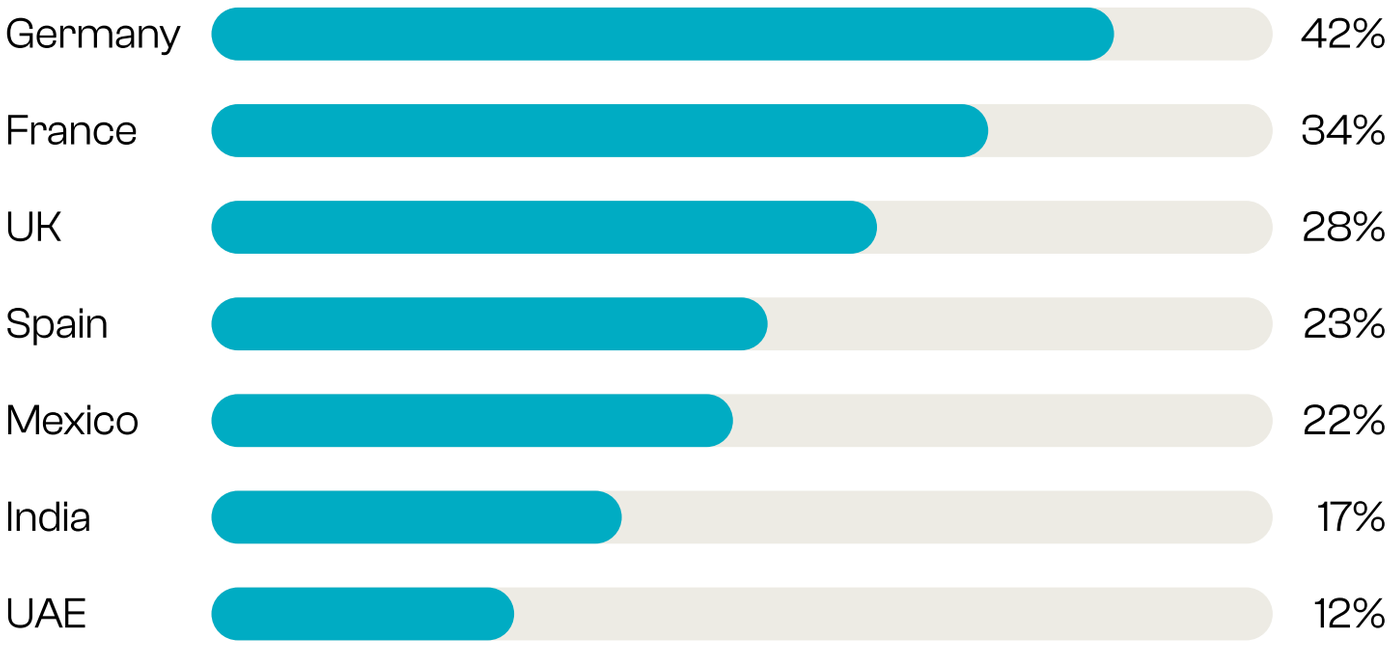

Share of cited sources on the market's own domains and local media. The US sits on global .com pages, so it is the baseline.

Germany reads: fuer-gruender.de, trusted.de, datev.de, ionos.de, lexware.de

France reads: tool-advisor.fr, clubic.com, independant.io, shine.fr

India reads: itforsme.in, technologycounter.com, kryptonforge.in

Global review sites like G2 and Capterra still show up. But being on them is no longer enough. To win in a market, your brand has to live on the local pages the AI reads there. This is why one global score cannot save you, and why coverage has to be built market by market.

What we held still, on purpose

We fixed one thing so the rest was clean: the person asking.

Every question in this study came from the same buyer, an owner of a small business. We held the persona still so that any change in the answer could only come from city, language, or engine. It keeps the finding honest and hard to argue with.

But it also points straight at the next study. If the city and the language move the answer this much, the person almost certainly does too. A startup founder, a hospital buyer, and a government office do not get the same list. Edition 02 opens that door, one industry and one buyer type at a time. The platform already varies the buyer in its live product, so the engine to run it is built.

What this means for you

You are not number seven everywhere. You are number two in one place and missing in another.

A single AI visibility score hides your real position. It gives you a comfortable average while your buyer in Berlin never sees your name and your buyer in Mumbai gets handed a rival. You cannot fix what one number cannot show.

This is why we built **Context Authority.**

Context Authority is one score for how well your brand owns a buying moment. Not from a blank slate. From every city and every language your buyer actually uses, and read many times so the wobble is measured, not missed. It works like a credit score for how AI sees you. One number you can raise, watched everywhere it counts.

Method, in full daylight

How we ran it. Nothing hidden.

We built this so anyone can check it. The point is to be believed, not admired.

The controlled test

We held the buying question and the buyer the same in every run. Only the city and the language changed. So any change in the answer comes from context, not from a reworded question.

10 categories, none ours

CRM, project management, payroll, email, accounting, payments, video calls, web hosting, cloud storage, and HR. No FTA client is in the study.

10 cities, 7 languages

New York, London, Mumbai, Delhi, Bengaluru, Mexico City, Madrid, Dubai, Paris, Berlin. English, Hindi, Hinglish, Spanish, Arabic, French, German. India runs deepest.

Two engines, five runs each

Every question ran on ChatGPT and Gemini, five times each, so we could measure the wobble as well as the winner. 1,200 answers in all.

Neutral seeds

Each question was seeded from a plain, neutral source. Never from a brand's own website, so nothing tilts the list.

Open on request

We share the full question set and the raw brand and source data. Run it yourself and you will see the same shape.

See your own number

You just read what happens to brands in general. Now see what happens to yours.

We will show you your Context Authority across every city and language your buyer uses. Where you win. Where you vanish. And what to do about it.

Watch the platform.

Book a demo