

— HERM.IO RESEARCH · TÜRKİYE KITCHEN & TABLEWARE

AI Visibility Report

The first AI-visibility snapshot of Türkiye's kitchen and tableware market: how five large language models name and rank brands across ten independent cuts of 45 Turkish-language questions.

1,125

Total Responses

5

Models Queried

45

Questions Analyzed

23

Qualified (General)

KEY FINDINGS

At a glance

This report reproduces the real Turkish-language questions consumers ask AI assistants about pots, pans, dinner sets, cutlery, storage and trousseau shopping in Türkiye, and measures how often, in what order, and across how many models five large language models name each brand — read across ten independent cuts.

- 1 Karaca reaches a perfect 100.00 in the general cut** — the first leader in this series to top all three score components at once. It appears in 57.2% of all 1,125 responses, is named earliest on average (MRR 0.576) and is seen by all five models.
- 2 The top three entities are not three competitors.** Karaca (#1), Emsan (#5) and Jumbo (#6) are all Karaca Group brands, and together they appear in **727 of 1,125 responses — 64.6%** of the corpus, rising to 72.0% on origin-neutral questions.
- 3 Leadership splits by segment.** Korkmaz leads Cookware (86.39) and Whole-kitchen & Trousseau (98.32); Karaca leads Tableware & Serving (98.30) and Cutlery & Storage (86.98). Tableware has just one foreign qualifier out of 21; Cutlery & Storage has IKEA at #2.
- 4 On origin-neutral questions mentions split 63.6% Turkish to 36.4% foreign** — but 58 distinct foreign entities appear against 54 Turkish ones, so the foreign presence is wider and thinner.
- 5 A quarter of all answers drift outside the category.** 25.8% of responses name an appliance or cabinetry brand, rising to 51.3% on whole-kitchen questions. Those entities are classified Out-of-scope and excluded from every table; the drift itself is reported as a finding.
- 6 One model mostly does not search.** 83.3% of answers were web-grounded, but the figure spans 19.6% (Gemini) to 100% (Claude, Perplexity). Gemini searches when a question contains a checkable physical claim — 40.0% on attribute questions against 8.0% on use-case questions.

SCOPE & METHOD

Vertical: kitchen & tableware (mutfak & sofra ürünleri), Türkiye. Method: a single point-in-time study of 1,125 responses across five large language models (Claude Haiku 4.5, Gemini 3.6 Flash, GPT-5.6 Luna Pro, Perplexity Sonar, Grok 4.3), each asked 45 Turkish-language questions five times. None excluded. Collection window: 1 August 2026, 05:57–10:54 UTC.

HOW TO READ THIS REPORT — VISIBILITY IS NOT QUALITY

This report measures one thing: how often, how prominently, and across how many AI assistants a brand is named when people ask about kitchen and tableware products in Türkiye. A high score means these systems currently have a lot of information available about a brand and surface it readily. It does not mean the brand is better, more durable, safer, cheaper or higher quality. A low score, or a zero, means a brand is currently less visible to these systems, not that it is inferior. Readers and shoppers should not treat an AI recommendation, or this report, as a verdict on whether a brand is “good” or “bad.” This study does not measure the accuracy, quality, or sentiment of any recommendation.

Key terms: AI visibility, AI Visibility Score, kitchenware brands, cookware brands, tableware brands, mutfak markaları, sofra ürünleri, çeyiz, Turkish (yerli) kitchen brands, LLM brand recommendations, ChatGPT/Gemini/Claude/Perplexity/Grok brand recommendation, Generative Engine Optimization (GEO).

CONTENTS

Report contents

01	Executive Summary	04
02	Methodology	07
03	Overall Visibility Leaderboards	12
04	Segment Leaderboards	17
05	Differences Between Models	21
06	Search vs. Memory	25
07	Question-Type Ownership	28
08	Open Market: Turkish vs. Foreign	30
09	One Owner, Three Leaders	33
10	Where the Category Boundary Leaks	35
11	The Discovery Ecosystem	38
12	What the Patterns Suggest	41
13	Limitations and Notes	43
14	Appendix & Data	47

01

SECTION 01

Executive Summary

Turkish kitchen and tableware answers are led by Karaca at a perfect score, but the more useful finding is that three of the six most visible entities belong to a single corporate group, and that a quarter of all answers name a brand from outside the category altogether.

EXECUTIVE SUMMARY

A perfect score, and a narrower shelf than it looks

ONE-SENTENCE TAKEAWAY

Turkish kitchen and tableware answers are led by Karaca at a perfect score, but the more useful finding is that three of the six most visible entities belong to a single corporate group, and that a quarter of all answers name a brand from outside the category altogether.

- 1 Karaca leads at 100.00 — the first perfect score in this series.** It appears in 57.2% of all 1,125 responses, is named earliest on average (MRR 0.576) and is seen by all five models. Under component scaling a leader normally lands short of 100 because no single entity tops every component; Karaca tops all three simultaneously. That is a property of this cut, not a normalisation artefact, and Chapter 3 explains it.
- 2 23 entities qualified in the general cut.** Of 144 tracked entities, 23 cleared the 5% mention threshold across all 45 questions: 17 Turkish and 6 foreign. Their average score is 58.2, and 19 of the 23 are seen by all five models.
- 3 The top three are one company.** Karaca (#1, 100.00), Emsan (#5, 70.11) and Jumbo (#6, 66.76) are all Karaca Group brands. Their combined footprint is **727 of 1,125 responses — 64.6%** — and on origin-neutral questions it reaches 72.0%. A reader scanning the leaderboard sees three competitors. Chapter 9 breaks out all five corporate groups in the data.
- 4 Segment leadership genuinely splits.** Korkmaz leads Cookware (86.39, narrowly over Karaca at 85.81) and Whole-kitchen & Trousseau (98.32). Karaca leads Tableware & Serving (98.30) and Cutlery & Storage (86.98). No entity wins all four.
- 5 Foreign visibility is concentrated in one segment.** Tableware & Serving has exactly one foreign qualifier out of 21. Cutlery & Storage has IKEA at #2 (79.16) and Zwilling at #4 — the strongest foreign showing anywhere in the study.
- 6 The origin-neutral shelf is Turkish-led but broadly contested.** On the 23 open questions, mentions split **63.6% Turkish to 36.4% foreign**, and 16 Turkish entities qualify against 11 foreign. But 58 distinct foreign entities appear against 54 Turkish ones: the foreign presence is wider and thinner, a long tail of European cookware names each mentioned rarely.
- 7 Marketplaces rank low, and that is one model's doing.** Trendyol (#20) and Hepsiburada (#21) are the only marketplaces to qualify, and **96.1% and 97.4% of their mentions come from Grok alone**. Four models barely name them.
- 8 A quarter of all answers drift outside the category.** 25.8% of responses name an appliance, small-appliance or cabinetry brand — Arçelik, Bosch, Beko, Arzum — rising to 51.3% on the six whole-kitchen questions. These entities are classified Out-of-scope and excluded from every table, but the drift itself is a finding about how the word *mutfak* is understood. Chapter 10.
- 9 One model mostly does not search.** 83.3% of all answers were web-grounded, but the figure spans 19.6% (Gemini) to 100% (Claude, Perplexity). Gemini searches when a question contains a checkable physical claim — 40.0% on attribute questions, 8.0% on use-case questions — and answers from memory otherwise. Chapter 6.

- 10 Sources are brand-owned first.** 44.2% of citations point to brands' own domains, led by karaca.com, cited by 34.2% of all responses and by all five models.

Why it matters

Kitchen and tableware buying in Türkiye runs on two distinct occasions — routine replacement and the *çeyiz* (trousseau) purchase — and both start with a broad question about which brands are worth considering. Whether a brand appears in an assistant's answer to that question is becoming a discovery channel in its own right. This report captures the first snapshot of that shelf. One reminder: the numbers below describe visibility and information availability, not which brands are best.

02

SECTION 02

Methodology

Five models were asked 45 Turkish questions five times each, producing 1,125 responses with none excluded; entities were discovered from the responses and reviewed by hand, and every score is computed against a per-cut scale published once and never recomputed.

METHODOLOGY

How the study was run

ONE-SENTENCE TAKEAWAY

Five models were asked 45 Turkish questions five times each, producing 1,125 responses with none excluded; entities were discovered from the responses and reviewed by hand, and every score is computed against a per-cut scale published once and never recomputed.

2.1 SCOPE AND MODELS

The study queried five large language models directly via API with identical questions. All models received only the user question. No system prompt was used, temperature was not set manually (provider default), and no region or locale parameter was defined. Because the questions are in Turkish, the models infer the Turkish market from language alone.

TABLE 1 — MODELS AND WEB-SEARCH BEHAVIOUR each produced exactly 225 responses

Model	Version	Web search	Rate
Claude Haiku 4.5	anthropic/claude-haiku-4.5	Enabled (web search tool)	100.0%
Gemini 3.6 Flash	gemini-3.6-flash	Enabled (Google Search grounding)	19.6%
GPT-5.6 Luna Pro	openai/gpt-5.6-luna-pro	Enabled (web search tool)	99.6%
Perplexity Sonar	perplexity/sonar	Always search-grounded	100.0%
Grok 4.3	x-ai/grok-4.3	Enabled (web search tool)	97.3%

Scale: 45 questions × 5 repeats × 5 models = 1,125 responses; **none were excluded** (no model returned an empty completed body). Single collection window, 1 August 2026, 05:57–10:54 UTC. Denominators are computed per question × model cell, never as questions × models × repeats. Web-search rate is the share of a model’s answers that returned at least one citation or search result.

2.2 WEB-SEARCH CONFIGURATION

Web search was offered as an enabled tool to all five models; each decided for itself whether to use it (Perplexity’s Sonar is search-grounded by design). Whether a response drew on web search or the model’s own knowledge is inferred from whether the provider returned any citation or source: a citation means “web-search,” none means “own knowledge.” This is a reasonable proxy, not direct proof of the model’s internal process.

In this run the grounding rate is **83.3% overall but wildly uneven** — 100% for two models and 19.6% for one. That spread matters for interpretation: one assistant answers mostly from its parameters while the others answer mostly from retrieved pages, so a brand’s visibility is partly a function of which model is asked. Chapter 6 examines it directly.

2.3 QUESTIONS, SEGMENTS AND SCOPE

The 45 questions reproduce the real Turkish queries people bring to an assistant about kitchenware: general discovery (“iyi bir yerli tencere markası var mı”), specific attributes (“bulaşık makinesinde deseni çıkmayan yemek takımı hangi marka”), and use cases (“yeni eve çıkıyorum”, “çeyiz için tek markadan yerli mutfak seti alacağım”). They are balanced across three behavioural types — 15 discovery, 15 attribute, 15 use_case — and split 22 *local* (explicitly asking for Turkish / *yerli* brands) against 23 *open* (no origin restriction).

QUESTIONS WERE RETAINED AS ASKED, INCLUDING AMBIGUITY

No question was rewritten in this edition. Six whole-kitchen questions using the bare word *mutfak* were resolved by several models as kitchen cabinetry or white goods rather than kitchenware. Those questions were not rewritten and not removed — rewriting a question because its answers are inconvenient is selection on outcome. The off-market entities are classified Out-of-scope, excluded from every table, and the drift is reported as a finding in Chapter 10.

TABLE 2 — SEGMENT DESIGN BALANCE

a rollup of 15 raw subcategories (median size: 2 questions)

Segment	Q	Resp.	disc / attr / use	local / open	Qual.	Status
Cookware	14	350	3 / 7 / 4	7 / 7	20	reportable
Tableware & Serving	15	375	5 / 4 / 6	7 / 8	21	reportable
Cutlery & Storage	10	250	3 / 4 / 3	4 / 6	29	reportable
Whole-kitchen & Trousseau	6	150	4 / 0 / 2	4 / 2	28	indicative

Because the question-type mix is not balanced within every segment, **the number of distinct brands in one segment should not be compared with another**. Whole-kitchen & Trousseau carries six questions and no attribute questions at all, and is marked indicative throughout.

2.4 ENTITY RESOLUTION AND ORIGIN

Entities were discovered from the responses rather than from a pre-built list. Candidates were extracted case-insensitively from URL-stripped answer text with Turkish-aware case folding, passed through a single human review gate, then researched in two external passes.

399 candidates were decided: 144 Keep, 135 Merge, 120 Drop. Of the kept entities: 99 brands, 28 retailers, 6 marketplaces, 7 sub-brands and 4 licensed brands. A further **61 entities are real companies in an adjacent market** — white goods, small appliances, kitchen cabinetry, sinks — and are recorded as Out-of-scope rather than deleted, which is what makes Chapter 10 measurable.

Matching is alias-based, word-boundary anchored with a Turkish letter class, tolerant of Turkish suffixes, and run on URL-stripped text so that a link target cannot inflate a brand's count.

ORIGIN RULE — TRADEMARK

An entity is Turkish (Local) when the trademark owner is Turkish, and Foreign otherwise. A foreign mark produced or sold in Türkiye under licence is Foreign, tagged Licensed-brand. Both trademark country and operator country were recorded for every entity so the rule can be changed in a later edition without re-researching. Of the 144 kept entities, 84 are Turkish and 60 foreign.

Five entities in the dictionary were never mentioned in any response — A101, Alkapıda, Boyner, Evidea and Primanova. They are retained and reported rather than deleted; the brand index in Chapter 14 lists the **139 entities that were matched at least once**.

Definitions. Mention: a validated entity (or alias) appears in the answer text, counted once per response — presence, not frequency. **Position (MRR):** mean reciprocal of the entity's first-appearance rank within each answer (first = 1.0, second = 0.5, ...), averaged only over the responses where it appears. **Breadth:** number of the five models (0–5) that mention the entity within that cut.

2.5 AI VISIBILITY SCORE (0-100)

$$\text{Score} = 0.45 \times \text{Mention} + 0.30 \times \text{Position} + 0.25 \times \text{Breadth}$$

Metrics are computed in two separable stages. **Stage 1** is raw and scale-free: mention rate, MRR and breadth, each computed within a cut. **Stage 2** is the scale layer: each component is divided by the maximum reached by a *qualified* entity within that cut, and those maxima are published once in a scale registry that every downstream file reads. Nothing recomputes them.

Component	Weight	Basis
Mention	45%	mention rate within the cut, against that cut's published maximum
Position	30%	MRR — how early the entity is named — against that cut's maximum
Breadth	25%	number of models covering it (0-5)

Qualification (≥5% rule): only entities mentioned in at least 5% of that cut's responses are ranked. This limits the risk that a rare entity which happens to appear first in a few answers inflates the position component. Unqualified entities are still scored, on the *same* cut maxima, so the appendix stays on one scale — which means some exceed 100 on a component. **21 rows, 2.3% of all metric rows, do.** They are flagged and never clipped: a component above 100 is a true statement that the entity beats the published leader on that component within that slice.

2.6 THE TEN PUBLISHED CUTS

Every entity is scored ten times, once per cut, each with its own denominator and its own published maxima.

TABLE 3 — THE TEN PUBLISHED CUTS

local + open = general; the segment cuts and the question-type cuts each sum to general

Cut	Resp.	Qual.	Mention leader	Position leader
general	1125	23	Karaca	Karaca
local	550	21	Karaca	Pirge
open	575	27	Karaca	Lock&Lock
Cookware	350	20	Karaca	Paşabahçe
Tableware & Serving	375	21	Kütahya Porselen	Karaca
Cutlery & Storage	250	29	Karaca	Pirge
Whole-kitchen & Trousseau	150	28	Korkmaz	Fissler
attribute	375	20	Karaca	Pirge
discovery	375	38	Karaca	Pirge
use_case	375	20	Karaca	Karaca

A SCALED SCORE IS NOT COMPARABLE ACROSS CUTS

100 in one cut is a different absolute number from 100 in another, because the scale moves whenever the leader moves. Comparing a brand's Cookware score with its Tableware score says nothing. Cross-cut and cross-edition comparison must use the raw layer — mention rate, MRR and breadth — which is why every table in this report publishes those alongside the score.

2.7 NEUTRALITY AND SELF-EXCLUSION

This is a market-wide, neutral study with no focus brand. To prevent any conflict of interest, all citations to Herm.io's own domains were excluded from the source data before analysis, so the company's own content neither appears among the most-cited domains nor influences any figure.

03

SECTION 03

Overall Visibility Leaderboards

Karaca tops every cut, but the shape underneath changes completely — the local shelf is entirely Turkish, the open shelf puts a foreign brand at #2, and the marketplaces sit near the bottom.

OVERALL VISIBILITY LEADERBOARDS

One leader, two different shelves

ONE-SENTENCE TAKEAWAY

Karaca tops every cut, but the shape underneath changes completely — the local shelf is entirely Turkish, the open shelf puts a foreign brand at #2, and the marketplaces sit near the bottom.

As a reminder, these rank visibility, not product quality, durability, price or safety.

3.1 General leaderboard: 23 qualified entities

TABLE 4 — AI VISIBILITY SCORE (GENERAL)

45 questions • 1,125 responses • 23 qualified

#	Brand	Origin	Mention	MRR	AI Score
● 1	Karaca	TR	57.2%	0.58	100.0
2	Korkmaz	TR	33.6%	0.57	80.9
3	Kütahya Porselen	TR	26.9%	0.47	71.0
4	Porland	TR	24.7%	0.50	70.4
5	Emsan	TR	35.6%	0.33	70.1
6	Jumbo	TR	29.9%	0.35	66.8
7	Tefal	Foreign	17.2%	0.54	66.5
8	IKEA	Foreign	13.4%	0.47	60.1
9	Güral Porselen	TR	18.4%	0.38	59.5
10	Paşabahçe	TR	13.2%	0.46	59.4
11	WMF	Foreign	12.0%	0.42	56.1
12	Nehir	TR	5.9%	0.51	56.1
13	Fissler	Foreign	7.5%	0.57	55.7
14	Schafer	TR	17.9%	0.30	54.9
15	Lava	TR	8.3%	0.41	53.0
16	Zwilling	Foreign	14.1%	0.30	51.9
17	Aryıldız	TR	15.2%	0.26	50.8
18	Hisar	TR	9.8%	0.30	48.5
19	Taç	TR	7.4%	0.34	48.4
20	Trendyol	TR	18.2%	0.20	44.8
21	Hepsiburada	TR	17.4%	0.16	42.1
22	English Home	TR	5.5%	0.17	38.4


(cont.)

#	Brand	Origin	Mention	MRR	AI Score
23	Le Creuset	Foreign	5.1%	0.20	34.3

AVERAGE

Average score of the 23 qualified entities: 58.2. Origin split: 17 Turkish, 6 foreign. Nineteen of the 23 are seen by all five models; the exceptions are Fissler (#13), Trendyol (#20), Hepsiburada (#21) and Le Creuset (#23).

AI VISIBILITY SCORE – TOP 15 (GENERAL CUT)



Why the leader reaches exactly 100.00

Under component scaling, each component is divided by the best qualified performance in that cut. A leader reaches 100 only if the *same* entity is simultaneously the most-mentioned, the earliest-named and present in all five models. That rarely happens — the mention leader and the position leader are usually different entities, as they are in eight of this study's ten cuts. Here Karaca holds all three: the highest mention rate (57.2%), the highest MRR (0.576) and full breadth. The perfect score is a description of that coincidence, not a ceiling imposed by the method.

Reading

Karaca (100.00) leads on volume and prominence together. Korkmaz (80.90) is second on a much lower mention rate (33.6%) but almost identical prominence — an MRR of 0.565 against Karaca's 0.576, giving it 98.1 on the position component. Third and fourth are the porcelain houses, Kütahya Porselen (70.95) and Porland (70.37), both ranking on position rather than volume. Emsan (70.11) inverts that: it is mentioned more often than either of them (35.6%) but named markedly later (MRR 0.329), which costs it the position component and drops it to fifth.

The highest foreign entity is Tefal at #7 (66.55), on only 17.2% mentions but an MRR of 0.538 — when Tefal is named, it is named early. IKEA follows at #8. Six foreign entities qualify in all: Tefal, IKEA, WMF, Fissler, Zwilling and Le Creuset — every one of them a cookware or knife name, none of them in tableware.

TWO ENTITIES TO READ CAREFULLY

Trendyol (#20) and Hepsiburada (#21) are the only marketplaces that qualify, and both rank on volume with very late positioning (MRR 0.201 and 0.161). Their visibility is also almost entirely one model's: 96.1% of Trendyol's mentions and 97.4% of Hepsiburada's come from Grok, and both have a breadth of 4 rather than 5. Read as a market signal, the marketplaces are not a general feature of this shelf; they are a feature of one assistant's answering style.

3.2 Brand-only view (marketplaces removed)

Per study scope, marketplaces are stripped from the brand-only view by name — Trendyol, Hepsiburada, n11, Amazon, Çiçeksepeti and PttAVM. Own-label retailers such as Karaca, Jumbo, IKEA and English Home are retailers, not marketplaces, and remain.

Because only two marketplaces qualified in the general cut, and both sit at #20 and #21, **the brand-only top ten is identical to the general top ten**. The removal takes the table from 23 rows to 21 and changes nothing above it. That is itself worth stating: unlike some verticals in this series, the Turkish kitchenware shelf is not aggregator-led. Assistants answer these questions with manufacturers.

3.3 Local leaderboard: 21 qualified entities (Turkish-only questions)

When the 22 questions that explicitly ask for Turkish / *yerli* brands are isolated (550 responses), **every one of the 21 qualified entities is Turkish**. Not a single foreign name survives the filter, which shows the assistants treat *yerli* as a real constraint rather than a stylistic hint.

TABLE 5 — AI VISIBILITY SCORE (LOCAL / YERLİ)

22 questions • 550 responses • 21 qualified

#	Brand	Origin	Mention	MRR	AI Score
● 1	Karaca	TR	46.2%	0.52	87.1
2	Korkmaz	TR	36.5%	0.71	83.9
3	Emsan	TR	34.9%	0.36	70.7
4	Kütahya Porselen	TR	27.8%	0.54	69.9
5	Porland	TR	23.1%	0.51	64.3
6	Pirge	TR	7.8%	0.91	62.6
7	Güral Porselen	TR	19.8%	0.48	60.1
8	Jumbo	TR	20.9%	0.35	57.0
9	Lava	TR	9.4%	0.60	54.1
10	Schafer	TR	14.6%	0.44	53.6
11	Sürbisa	TR	7.6%	0.60	52.3
12	Aryıldız	TR	16.6%	0.26	49.7
13	Taç	TR	11.4%	0.40	49.5
14	Hisar	TR	13.8%	0.31	48.8
15	Paşabahçe	TR	11.4%	0.48	47.1

(cont.)

#	Brand	Origin	Mention	MRR	AI Score
16	Nehir	TR	7.3%	0.46	47.1
17	Papilla	TR	7.1%	0.35	43.4
18	Trendyol	TR	18.7%	0.25	41.4
19	Hepsiburada	TR	18.6%	0.18	39.0
20	OMS Collection	TR	5.6%	0.38	33.0
21	Acar	TR	7.1%	0.33	32.8

AVERAGE

Average score of the 21 qualified local entities: 54.6. All 21 are Turkish.

Reading. Karaca (87.14) still leads, but its margin collapses: Korkmaz is at 83.88, and on the position component Korkmaz is *ahead* — an MRR of 0.707 against Karaca's 0.520. When the question narrows to domestic brands, Korkmaz is the name assistants reach for first.

The local cut also surfaces an entity the general cut hides. **Pirge** ranks #6 (62.62) on just 7.8% mentions, carried by an MRR of 0.911 — the highest position score in the cut and the reference maximum against which every other local position score is measured. It is a knife maker, and when a question asks for a Turkish knife, it is named first almost every time. In the general cut Pirge falls just below the threshold at 4.00% and is not ranked at all.

04

SECTION 04

Segment Leaderboards

Four segments, two different leaders, and radically different openness to foreign brands — Tableware & Serving has one foreign qualifier out of 21, while Cutlery & Storage has IKEA at #2.

SEGMENT LEADERBOARDS

Two leaders, four different markets

ONE-SENTENCE TAKEAWAY

Four segments, two different leaders, and radically different openness to foreign brands — Tableware & Serving has one foreign qualifier out of 21, while Cutlery & Storage has IKEA at #2.

Scores below are scaled within each segment and **cannot be compared across segments**. Mention rates can.

4.1 Cookware — 14 questions, 350 responses, 20 qualified

TABLE 6 — COOKWARE

tencere, tava, düdüklü • top 10 of 20 qualified

#	Brand	Origin	Mention	MRR	Models	AI Score
● 1	Korkmaz	TR	52.3%	0.608	5	86.39
2	Karaca	TR	56.0%	0.496	5	85.81
3	Emsan	TR	44.6%	0.382	5	72.99
4	Tefal	Foreign	31.7%	0.540	5	67.68
5	WMF	Foreign	21.7%	0.482	5	57.82
6	Lava	TR	20.3%	0.467	5	56.16
7	Jumbo	TR	25.4%	0.295	5	54.83
8	Paşabahçe	TR	5.7%	0.942	4	54.59
9	Fissler	Foreign	18.3%	0.537	4	51.80
10	IKEA	Foreign	12.9%	0.438	5	49.28

The closest race in the study. Korkmaz (86.39) edges Karaca (85.81) by half a point — and it does so while being mentioned *less* often (52.3% against 56.0%), winning on position instead (MRR 0.608 against 0.496). This is the only cut in which Karaca is not the mention leader and not the winner. Tefal (#4) and WMF (#5) give this segment the strongest foreign presence outside Cutlery & Storage.

Note Paşabahçe at #8 on a 5.7% mention rate: it clears the bar largely on an MRR of 0.942, the highest position figure anywhere in the study and the reference maximum for this cut. Read its rank cautiously — it is a glassware house being named first in the handful of cookware answers that reach for oven-safe glass.

4.2 Tableware & Serving — 15 questions, 375 responses, 21 qualified

TABLE 7 — TABLEWARE & SERVING

yemek takımı, servis, sunum • top 10 of 21 qualified

#	Brand	Origin	Mention	MRR	Models	AI Score
● 1	Karaca	TR	68.0%	0.677	5	98.30
2	Kütahya Porselen	TR	70.7%	0.503	5	92.29
3	Porland	TR	64.5%	0.523	5	89.28
4	Güral Porselen	TR	48.8%	0.406	5	74.04

(cont.)

#	Brand	Origin	Mention	MRR	Models	AI Score
5	Jumbo	TR	40.0%	0.367	5	66.72
6	Paşabahçe	TR	17.1%	0.471	5	56.72
7	Emsan	TR	25.3%	0.286	5	53.80
8	Schafer	TR	21.3%	0.335	5	53.43
9	Şişecam	TR	5.3%	0.567	4	48.50
10	English Home	TR	15.5%	0.175	5	42.60

The most domestic segment in the study, and by a distance. Of 21 qualified entities, **exactly one is foreign** — Villeroy & Boch. The top four are Karaca (98.30), Kütahya Porselen (92.29), Porland (89.28) and Güral Porselen (74.04), and the three porcelain houses between them define what an assistant considers a Turkish dinner set.

Kütahya Porselen is the mention leader here at 70.7% — higher than Karaca's 68.0% — and is the reference maximum for the mention component in this cut. Karaca takes the segment on position (MRR 0.677 to 0.503).

4.3 Cutlery & Storage — 10 questions, 250 responses, 29 qualified

TABLE 8 — CUTLERY & STORAGE

çatal-bıçak, saklama kapları • top 10 of 29 qualified

#	Brand	Origin	Mention	MRR	Models	AI Score
● 1	Karaca	TR	39.2%	0.515	5	86.98
2	IKEA	Foreign	33.2%	0.487	5	79.16
3	Pirge	TR	17.2%	0.911	5	74.74
4	Zwilling	Foreign	30.4%	0.309	5	70.08
5	Emsan	TR	27.6%	0.311	5	66.93
6	Nehir	TR	18.4%	0.587	5	65.45
7	Sürbisa	TR	16.8%	0.603	5	64.15
8	Tefal	Foreign	18.4%	0.650	4	62.54
9	Korkmaz	TR	18.8%	0.354	5	58.24
10	Jumbo	TR	17.2%	0.537	4	57.44

The most contested segment, and the only one where a foreign brand is genuinely near the top. IKEA sits at #2 (79.16) on a 33.2% mention rate, and Zwilling at #4 (70.08). This is the strongest foreign showing in the study, and the pattern is legible: storage containers and knives are categories where European and global names carry established reputations that assistants reproduce.

It is also the segment with the most qualifiers (29 from 250 responses), which reflects a wide, fragmented field rather than a deeper market — the qualifying bar is 12.5 responses here against 56.25 in the general cut. Pirge (#3) is the position leader at MRR 0.911; Sürbisa (#7) and Nehir (#6) are Turkish knife makers that appear nowhere in the general leaderboard.

4.4 Whole-kitchen & Trousseau — 6 questions, 150 responses, 28 qualified • indicative

TABLE 9 — WHOLE-KITCHEN & TROUSSEAU

genel mutfak, çeyiz • top 10 of 28 • indicative only

#	Brand	Origin	Mention	MRR	Models	AI Score
● 1	Korkmaz	TR	84.0%	0.649	5	98.32
2	Karaca	TR	62.7%	0.530	5	81.71
3	Emsan	TR	53.3%	0.289	5	66.21
4	Fissler	Foreign	13.3%	0.687	4	57.14
5	Jumbo	TR	36.0%	0.246	5	55.03
6	Tefal	Foreign	24.0%	0.390	5	54.90
7	IKEA	Foreign	11.3%	0.610	4	52.71
8	Porland	TR	22.7%	0.330	5	51.53
9	Zwilling	Foreign	17.3%	0.361	5	50.03
10	Schafer	TR	32.7%	0.257	4	48.70

READ THIS TABLE AS INDICATIVE, NOT REPORTABLE

Six questions, 150 responses, and no attribute questions at all, so the segment's question mix is not comparable with the others. These are also the questions where off-market drift is heaviest (Chapter 10), so the segment is measuring a blurrier question than the other three.

Within those limits it is the segment Korkmaz owns outright: 98.32, on an 84.0% mention rate — the highest mention rate recorded anywhere in this study. When someone asks a broad question about kitchen brands or trousseau sets, Korkmaz is named in more than four answers out of five. Karaca follows at 81.71.

05

SECTION 05

Differences Between Models

Four of five models agree that Karaca leads; they disagree enormously about how many brands an answer should contain, and Grok alone accounts for nearly all marketplace visibility in the study.

DIFFERENCES BETWEEN MODELS

Agreement on Karaca, dispute on breadth

ONE-SENTENCE TAKEAWAY

Four of five models agree that Karaca leads; they disagree enormously about how many brands an answer should contain, and Grok alone accounts for nearly all marketplace visibility in the study.

5.1 PER-MODEL BEHAVIOUR

TABLE 10 – BEHAVIOUR BY MODEL

each produced exactly 225 responses (45 × 5)

Model	Web	Distinct	Ent./ans	TR share
Claude Haiku 4.5	100.0%	79	4.24	76.4%
Gemini 3.6 Flash	19.6%	84	6.41	72.6%
GPT-5.6 Luna Pro	99.6%	76	4.20	74.3%
Perplexity Sonar	100.0%	92	5.28	80.8%
Grok 4.3	97.3%	96	8.15	80.9%

DISTINCT ENTITIES NAMED PER ANSWER, BY MODEL



Grok names about **8.15 entities per answer** and reaches the widest vocabulary (96 distinct entities); GPT-5.6 Luna Pro (4.20) and Claude Haiku (4.24) are roughly half as generous. Gemini sits at 6.41 despite searching on only a fifth of its answers, which is worth noting on its own: breadth of recall is not a function of retrieval here.

The Turkish share of mentions is remarkably stable — 72.6% to 80.9% across five models with very different answering styles. Whatever else the assistants disagree about, they agree that this is a market answered mostly with Turkish names.

5.2 EACH MODEL'S MOST-NAMED ENTITIES

CLAUDE HAIKU 4.5

top 5

#	Entity	Mentions	Rate
1	Karaca	93	41.3%
2	Emsan	67	29.8%
3	Korkmaz	61	27.1%
4	Tefal	46	20.4%
5	Schafer	43	19.1%

05 · DIFFERENCES BETWEEN MODELS

23 / 55

GEMINI 3.6 FLASH

top 5

#	Entity	Mentions	Rate
1	Karaca	115	51.1%
2	Korkmaz	100	44.4%
3	Jumbo	87	38.7%
4	Porland	81	36.0%
5	Kütahya Porselen	80	35.6%

GPT-5.6 LUNA PRO

top 5

#	Entity	Mentions	Rate
1	Karaca	99	44.0%
2	Jumbo	72	32.0%
3	Kütahya Porselen	56	24.9%
4	Porland	55	24.4%
5	Emsan	54	24.0%

PERPLEXITY SONAR

top 5

#	Entity	Mentions	Rate
1	Karaca	150	66.7%
2	Emsan	81	36.0%
3	Korkmaz	74	32.9%
4	Kütahya Porselen	67	29.8%
5	Porland	61	27.1%

GROK 4.3

top 5

#	Entity	Mentions	Rate
1	Trendyol	197	87.6%
2	Hepsiburada	191	84.9%
3	Karaca	186	82.7%
4	Emsan	129	57.3%
5	Korkmaz	98	43.6%

Karaca tops four of the five models — Perplexity most heavily, naming it in 66.7% of its answers. The exception is Grok, whose top two are Trendyol (87.6%) and Hepsiburada (84.9%), with Karaca third at 82.7%. Grok is the reason marketplaces appear in this study's general leaderboard at all; strip its responses and the shelf is manufacturers all the way down.

The models also differ in *which* Turkish names they favour. Claude reaches for Emsan, Tefal and Schafer; Gemini for Korkmaz, Jumbo and Porland; GPT for Jumbo and the porcelain houses. Only Karaca, Korkmaz, Emsan and Kütahya Porselen appear in every model's top five or six.

THE PRACTICAL READING

Four of five assistants answer a kitchenware question with manufacturers, and one answers substantially with the shops that sell them. Any conclusion about “where AI sends shoppers” is a conclusion about which assistant was asked.

06

SECTION 06

Search vs. Memory

Gemini searches when a question contains a physically checkable claim and answers from memory otherwise; the shelf it produces in each mode is visibly different, but the two modes answer different question mixes, so this study cannot show that retrieval caused the difference.

— SEARCH VS. MEMORY

A real difference, and a causal claim we decline

ONE-SENTENCE TAKEAWAY

Gemini searches when a question contains a physically checkable claim and answers from memory otherwise; the shelf it produces in each mode is visibly different, but the two modes answer different question mixes, so this study cannot show that retrieval caused the difference.

Previous editions in this series omitted this chapter because search was near-universal. Here it is not: **937 of 1,125 responses were grounded and 188 were not**, a large enough ungrounded segment to look at. One caveat governs everything below, so it comes first: **96.3% of all ungrounded responses in the study are Gemini's**. A study-wide grounded-versus-ungrounded comparison is therefore very close to a Gemini-versus-everyone-else comparison.

TABLE 11 — GROUNDING BY MODEL

four models ground almost everything; one does not

Model	Grounded	Ungrounded	Rate
Claude Haiku 4.5	225	0	100%
Gemini 3.6 Flash	44	181	19.6%
GPT-5.6 Luna Pro	224	1	99.6%
Perplexity Sonar	225	0	100%
Grok 4.3	219	6	97.3%

6.1 WHEN GEMINI SEARCHES

Gemini decided per call whether to search. It did so on 44 of its 225 answers, and the pattern is not random.

TABLE 12 — GEMINI'S GROUNDING RATE BY QUESTION TYPE

attribute questions are grounded five times as often as use-case

Question type	Grounded	Responses	Rate
discovery	8	75	10.7%
attribute	30	75	40%
use_case	6	75	8%

Gemini searches when the question contains a claim someone could check, and answers from memory when asked which brands are good. It grounded on **40.0% of attribute questions** against 10.7% of discovery and 8.0% of use-case questions. The three questions it grounded on all five repeats are, in translation, a dinner set whose pattern survives the dishwasher, gold trim that does not degrade in the machine, and a Turkish knife that stays sharp. The 25 questions it never grounded on are dominated by open requests to name good brands.

Grounding is also partly stochastic rather than purely a property of the question: 25 questions sat at zero grounded repeats out of five, 3 at five out of five, and 17 somewhere in between. If grounding were purely a question property every question would sit at 0 or 5; if it were purely random the distribution would cluster around the mean. Neither holds. When Gemini did search, its answers were longer (3,069 characters against 2,155) and named more entities (7.07 against 6.08).

6.2 THE TWO SHELVES ARE DIFFERENT

TABLE 13 — LARGEST GROUNDED / UNGROUNDED GAPS

937 vs 188 responses · entities with ≥25 combined mentions

Entity	Grounded	Ungrounded	Difference
Trendyol	21.3%	2.7%	+18.7 pp
Hepsiburada	20.4%	2.7%	+17.7 pp
Schafer	20.3%	5.9%	+14.4 pp
Karaca	59.1%	47.3%	+11.8 pp
Emsan	37%	28.2%	+8.8 pp
Acar	5.9%	0%	+5.9 pp
Korkmaz	30.7%	47.9%	-17.1 pp
Hisar	7.4%	21.8%	-14.4 pp
Lava	6.2%	18.6%	-12.4 pp
Jumbo	27.9%	39.9%	-12 pp
Porland	22.7%	34.6%	-11.8 pp
Le Creuset	3.5%	12.8%	-9.2 pp

Retrieval surfaces marketplaces and brands with a strong own-site presence: Trendyol (+18.7 pp), Hepsiburada (+17.7 pp), Schafer (+14.4 pp), Karaca (+11.8 pp). Parametric memory surfaces established domestic manufacturers and European cast-iron names: Korkmaz (-17.1 pp), Hisar (-14.4 pp), Lava (-12.4 pp), Porland (-11.8 pp), Le Creuset (-9.2 pp), Staub (-9.1 pp). That is a real difference in output. It is not, on this data, a demonstrated effect of retrieval.

6.3 WHY WE STOP SHORT OF A CAUSAL CLAIM

Gemini chose when to search, and it chose by question type. So its two subsets are not two ways of answering the same questions — they are answers to different questions.

TABLE 14 — GEMINI'S QUESTION MIX, GROUNDED VS UNGROUNDED

the two subsets answer different questions

Question type	Share of grounded	Share of ungrounded	Share of corpus
attribute	68.2%	24.9%	33.3%
discovery	18.2%	37%	33.3%
use_case	13.6%	38.1%	33.3%

68.2% of Gemini's grounded answers are attribute questions, against 24.9% of its ungrounded ones. Korkmaz being more visible in the ungrounded set may mean that searching pushes it down — or simply that Korkmaz is a name that comes up when someone asks broadly which brands are good, and broad questions are the ones Gemini did not search. Both explanations fit the data equally well, and nothing in this study separates them.

WHAT WOULD SETTLE IT

We report the difference and decline the causal claim. Identifying a retrieval effect would require running the same questions with search forced on and forced off in the same model, which is a design change for a future edition rather than something recoverable from this one.

07

SECTION 07

Question-Type Ownership

Karaca leads all three question types, but the field behind it thins dramatically — 38 entities qualify on discovery questions against 20 on attribute questions.

QUESTION-TYPE OWNERSHIP

One leader, three different fields

ONE-SENTENCE TAKEAWAY

Karaca leads all three question types, but the field behind it thins dramatically — 38 entities qualify on discovery questions against 20 on attribute questions.

Alongside the four product segments, the 45 questions carry three behavioural types: **discovery** (open “which brands” questions), **attribute** (specific properties such as dishwasher-safe patterns or edge retention) and **use_case** (situation-driven, such as moving house or buying a trousseau set). Each type covers 15 questions and 375 responses.

TABLE 15 – MOST VISIBLE ENTITIES BY QUESTION TYPE top 5 by mention rate · rate is within that type

#	discovery	attribute	use_case
1	Karaca (65.3%)	Karaca (48.0%)	Karaca (58.1%)
2	Emsan (52.5%)	Emsan (26.1%)	Korkmaz (34.4%)
3	Jumbo (46.7%)	Tefal (24.0%)	Jumbo (31.7%)
4	Korkmaz (46.7%)	Korkmaz (19.7%)	Kütahya Porselen (29.6%)
5	Kütahya Porselen (34.4%)	Trendyol (17.3%)	Emsan (28.0%)

Reading. Karaca leads all three, most heavily on use-case questions (58.1%), where a shopper describes a situation — moving house, buying a trousseau — and assistants reach for a brand that covers a whole kitchen.

The more interesting number is how many entities qualify. **Discovery questions produce 38 qualified entities; attribute and use-case produce 20 each.** When the question is open (“which kitchen brands are actually good”), assistants spread their answers across a long list. When the question specifies a property — dishwasher-safe, stays sharp, does not warp — the field narrows sharply to the names associated with that property. Attribute questions are also where foreign cookware performs best: Tefal is third overall on attribute questions at 24.0%, and second by score.

Pirge is the position leader on both discovery and attribute questions, despite qualifying in neither the general cut nor the use-case cut. A specialist with a narrow, strongly associated category is named first when named at all.



SECTION 08

Open Market: Turkish vs. Foreign

On origin-neutral questions, mentions split 63.6% Turkish to 36.4% foreign — but more distinct foreign entities appear than Turkish ones, so foreign visibility is broad and shallow rather than absent.

— OPEN MARKET — TURKISH VS. FOREIGN

A Turkish core, a European periphery

ONE-SENTENCE TAKEAWAY

On origin-neutral questions, mentions split 63.6% Turkish to 36.4% foreign — but more distinct foreign entities appear than Turkish ones, so foreign visibility is broad and shallow rather than absent.

This chapter relies only on the 23 open questions with no origin restriction (575 responses), the only fair basis for comparing Turkish and foreign brands. The 22 local questions cannot surface foreign brands by construction, and mixing them would manufacture a Turkish majority.

8.1 Origin share

MENTION SHARE BY ORIGIN — OPEN CUT (%)



Counting every entity that appeared at least once, the open cut splits **2,355 Turkish mentions to 1,348 foreign — 63.6% to 36.4%**. Three other framings of the same question are stated together below, so that none can be quoted alone.

TABLE 16 — FOUR BASES FOR THE SAME COMPARISON

open cut • 575 responses

Basis	Turkish	Foreign
Mentions, all entities that appeared	63.6%	36.4%
Mentions, qualified entities only	68.3%	31.7%
Qualified entity count	16	11
Distinct entities appearing at all	54	58

The last row is the finding. More distinct foreign entities appear in the open cut than Turkish ones — 58 against 54 — yet they collect little more than a third of the mentions. Turkish visibility is concentrated in a handful of names that appear constantly; foreign visibility is spread across a long tail of European cookware and porcelain houses that each surface occasionally. Restricting the count to qualified entities removes most of that tail and moves the split to 68.3% / 31.7%, which is why both figures are published here.

DO NOT QUOTE THE GENERAL CUT AS A MARKET SHARE

The general cut — which includes the 22 questions that demand Turkish brands — yields 17 Turkish qualifiers to 6 foreign. That figure describes the study's question design as much as the market.

8.2 Open-market top 15

TABLE 17 — OPEN MARKET, TOP 15

23 questions • 575 responses • 27 qualified

#	Brand	Origin	Type	Rate	AI Score
1	Karaca	TR	Brand/Ret.	67.7%	97.87
2	Tefal	Foreign	Brand	31.0%	70.26

(cont.)

#	Brand	Origin	Type	Rate	AI Score
3	Jumbo	TR	Brand/Ret.	38.4%	66.48
4	Porland	TR	Brand	26.3%	64.65
5	Korkmaz	TR	Brand	30.8%	63.89
6	IKEA	Foreign	Licensed	24.9%	62.93
7	Emsan	TR	Brand	36.2%	62.88
8	Kütahya Porselen	TR	Brand	26.1%	61.00
9	WMF	Foreign	Brand	23.3%	59.55
10	Zwilling	Foreign	Brand	26.6%	56.68
11	Fissler	Foreign	Brand	14.1%	56.09
12	Paşabahçe	TR	Brand/Ret.	14.8%	55.06
13	Lock&Lock	Foreign	Brand	7.3%	49.86
14	Güral Porselen	TR	Brand	17.0%	48.94
15	Schafer	TR	Brand	21.0%	48.79

8.3 The pattern

Karaca leads the open cut at 97.87 on a 67.7% mention rate — higher than its general-cut rate, because origin-neutral questions are where it is most dominant. **Tefal is second at 70.26**, the highest a foreign entity places in any general-population cut in this study, and IKEA is sixth.

The foreign names that qualify follow one pattern almost without exception: they are cookware, knives or storage — Tefal, IKEA, WMF, Zwilling, Fissler, Le Creuset, Lock&Lock, Tupperware. In tableware and serving, where Turkish porcelain manufacturers have deep domestic reputations, foreign brands are nearly absent. So when a Turkish shopper asks an origin-neutral kitchenware question, the answer is Turkish-led overall, but genuinely mixed as soon as the subject turns to pans and knives.

A THRESHOLD IS A LINE

Villeroy & Boch appears in 4.98% of general-cut responses — one response short of the 5% bar — and is therefore not ranked. Entities land on both sides of a threshold by margins that carry no meaning.



SECTION 09

One Owner, Three Leaders

Karaca, Emsan and Jumbo are one company covering 64.6% of all responses, and on origin-neutral questions a second group — Tefal and WMF's owner — reaches 41.9%; the leaderboard shows brands, not owners.

CORPORATE GROUPS IN THE DATA

One owner, three leaders

ONE-SENTENCE TAKEAWAY

Karaca, Emsan and Jumbo are one company covering 64.6% of all responses, and on origin-neutral questions a second group — Tefal and WMF’s owner — reaches 41.9%; the leaderboard shows brands, not owners.

A leaderboard ranks brands. Shoppers meet brands. But a reader drawing conclusions about competition from the table in Chapter 3 would be reading three Karaca Group brands as three independent competitors.

TABLE 18 – CORPORATE GROUPS WITH TWO OR MORE ENTITIES union of member response sets, not the sum of mentions

Group	Members	General	Local	Open
Karaca Group	Emsan, Jumbo, Karaca	64.6% (727)	56.9%	72%
Groupe SEB	Tefal, WMF	22.8% (257)	2.9%	41.9%
Zwilling Group	Ballarini, Demeyere, Staub, Zwilling	15% (169)	1.6%	27.8%
Şişecam	Nude Glass, Paşabahçe, Şişecam	13.6% (153)	12.4%	14.8%
Turgut Aydın Holding	English Home, Hecha	6.8% (77)	2%	11.5%

Karaca Group appears in 727 of 1,125 responses: 64.6%. That is a union, not a sum. Adding the three members’ mention counts would give 1,379 and double-count every answer that names two of them — which is precisely the error a reader makes scanning the leaderboard. On the open cut the group reaches **72.0%**: nearly three in four origin-neutral kitchenware answers name Karaca, Emsan or Jumbo.

The second group is foreign and appears only when the question allows it. Groupe SEB — Tefal and WMF — covers 22.8% of the corpus overall, but that splits into **2.9% on local questions and 41.9% on open ones**. Zwilling Group shows the same shape more sharply: 1.6% local against 27.8% open. Foreign group visibility in this market is not weak; it is conditional on the question not asking for Turkish brands.

Şişecam is the exception to that pattern — a Turkish group whose visibility is roughly even across local (12.4%) and open (14.8%) questions, because Paşabahçe is recognised as a Turkish glassware name in both contexts.

HOW TO READ THIS

Group concentration is not a claim about market power, pricing or competition — those are commercial questions this study does not measure. It is a caution about the leaderboard: three names in the top six sharing an owner means the visible shelf is narrower than a rank-ordered list makes it look. A parentGroup field is carried on every entity in the published data for anyone who wants to collapse them.

10

SECTION 10

Where the Category Boundary Leaks

A quarter of all answers name a brand from outside kitchenware entirely — appliances, small appliances, cabinetry — and while whole-kitchen questions concentrate it, this is a property of the whole corpus, not six bad questions.

WHERE THE CATEGORY BOUNDARY LEAKS

What “mutfak” turns out to mean

ONE-SENTENCE TAKEAWAY

A quarter of all answers name a brand from outside kitchenware entirely — appliances, small appliances, cabinetry — and while whole-kitchen questions concentrate it, this is a property of the whole corpus, not six bad questions.

The Turkish word *mutfak* means both “kitchen” and, colloquially, “kitchenware.” Assistants resolve that ambiguity inconsistently, and the study measured how often.

TABLE 19 — OFF-MARKET NAMING BY SCOPE share of responses naming at least one of 61 Out-of-scope entities

Scope	Responses	With out-of-scope	Rate
Whole-kitchen questions (6)	150	77	51.3%
All other questions (39)	975	213	21.8%
Whole corpus (45)	1125	290	25.8%

25.8% of all responses name a brand from outside the category. On the six whole-kitchen questions it reaches 51.3%; on the 39 product-specific questions it is still 21.8%. This is not six ambiguous questions producing noise — it is a standing feature of how these systems answer Turkish kitchen queries, concentrated about 2.4x on the broadest ones. All Out-of-scope entities are excluded from every leaderboard in this report.

TABLE 20 — OFF-MARKET NAMING BY MODEL the ordering is not the same in each column

Model	Whole-kitchen questions	All other questions
Claude Haiku 4.5	46.7%	30.3%
Gemini 3.6 Flash	36.7%	33.3%
GPT-5.6 Luna Pro	60%	16.4%
Perplexity Sonar	46.7%	12.3%
Grok 4.3	66.7%	16.9%

Retrieval does not explain it. The least-grounded model, Gemini at 19.6%, has the *lowest* whole-kitchen drift (36.7%); the heavily-grounded Grok has the highest (66.7%). On product-specific questions the ordering reverses again, with Gemini highest at 33.3% and Perplexity lowest at 12.3%. Whatever produces category drift, it is not simply a matter of whether the model looked something up.

10.1 ONE WORD CHANGES THE ANSWER

TABLE 21 — THE SIX WHOLE-KITCHEN QUESTIONS each asked 25 times · ordered by drift

#	Question (Turkish)	Type	Scope	Drift
15	fiyatına değen mutfak markaları hangileri	discovery	open	80%
32	yeni eve çıkıyorum, iyi yerli mutfak markaları önerir misin	use_case	local	64%
1	mutfak eşyasında hangi markalar gerçekten iyi	discovery	open	60%

(cont.)

#	Question (Turkish)	Type	Scope	Drift
10	çeyiz için Türk malı mutfak ürünlerinde hangi markalara bakayım	discovery	local	56%
36	çeyiz için tek markadan yerli mutfak seti alacağım, hangisi iyi	use_case	local	36%
8	mutfak ve sofra ürünlerinde iyi yerli markalar neler	discovery	local	12%

The spread across six near-identical questions runs from 80.0% to 12.0%, and the lowest one is the only question that names the category twice: “*mutfak ve sofra ürünlerinde iyi yerli markalar neler.*” Adding *sofra* — table, tableware — appears to anchor the assistants to kitchenware and cut drift by a factor of nearly seven against the bare-*mutfak* phrasing.

A HYPOTHESIS, NOT A RESULT

That observation rests on 25 responses per question and should be treated as a hypothesis rather than a result. But it points at something practical: the ambiguity lives in the prompt, not only in the model.

10.2 WHAT THE MODELS NAME INSTEAD

TABLE 22 — MOST-NAMED OUT-OF-SCOPE ENTITIES

on whole-kitchen questions • excluded from every leaderboard

#	Entity	Responses	Category
1	Arçelik	28	Major appliances
2	Bosch	28	Major appliances
3	Beko	24	Major appliances
4	Arzum	22	Small appliances
5	Philips	21	Small appliances
6	Siemens	20	Major appliances
7	Vestel	16	Major appliances
8	KitchenAid	15	Small & major appliances
9	Scavolini	11	kitchen cabinetry
10	Braun	10	Small appliances

The list is appliance makers: Arçelik, Bosch, Beko, Siemens, Vestel on major appliances; Arzum, Philips, Braun on small ones; Scavolini on cabinetry. Read charitably, this is not a failure of comprehension — a shopper furnishing a kitchen may well want a kettle and a dishwasher alongside their pans, and the assistant is answering the question a person might have meant. It is a scope boundary problem, and it is why the study keeps a separate Out-of-scope classification rather than silently dropping the names.



SECTION 11

The Discovery Ecosystem

44.2% of citations point to brands' own websites, karaca.com alone is cited by a third of all responses, and — unusually for this series — recipe sites and forums together supply nearly a fifth of the evidence.

THE DISCOVERY ECOSYSTEM

Where AI learns about these brands

ONE-SENTENCE TAKEAWAY

44.2% of citations point to brands' own websites, karaca.com alone is cited by a third of all responses, and — unusually for this series — recipe sites and forums together supply nearly a fifth of the evidence.

When an assistant answers with web search, it draws on the content available to it at that moment. We collected the source addresses returned by the models, reduced them to domains, and ranked them by how many distinct responses cited each. Across the run, **937 responses carried at least one citation, referencing 479 distinct domains across 5,539 domain-response pairs.**

11.1 The 12 most-cited domains

TABLE 23 — MOST-CITED DOMAINS

the top 50 account for 62.5% of all domain-response pairs

#	Domain	Resp.	% of resp.	Models	Source type
1	karaca.com	385	34.2%	5	brand-owned
2	trendyol.com	213	18.9%	5	marketplace
3	ikea.com.tr	125	11.1%	4	brand-owned
4	nefisyemektarifleri.com	125	11.1%	3	editorial
5	kutahyaporselen.com	125	11.1%	5	brand-owned
6	onedio.com	117	10.4%	4	editorial
7	hepsiburada.com	114	10.1%	4	marketplace
8	eniyilerden.com	113	10%	4	other
9	schafer.com.tr	107	9.5%	5	brand-owned
10	aryildiz.com	101	9%	5	brand-owned
11	zwilling.com	92	8.2%	4	brand-owned
12	tefal.com.tr	90	8%	5	brand-owned

karaca.com is cited by 385 responses — 34.2% of everything collected — and by all five models. The gap to second place is large: trendyol.com at 213. Below them the table is a mix of brand sites (kutahyaporselen.com, schafer.com.tr, aryildiz.com, zwilling.com, tefal.com.tr), two marketplaces, and — notably — three editorial and listicle sites, including a recipe site, nefisyemektarifleri.com.

The Models column is worth reading alongside the count. A domain cited by all five models is part of the shared evidence base; one cited by two is a single assistant's habit. Among the top 12, six are cited by all five models — and all six are brand or marketplace domains, not editorial ones.


11.2 Source-type mix

TABLE 24 — CITATION SHARE BY SOURCE TYPE

across all 479 cited domains

Source type	Domains	Citations	Share
brand-owned	152	2447	44.2%
other	143	994	17.9%
editorial	62	661	11.9%
marketplace	18	521	9.4%
retailer	54	450	8.1%
forum/social	14	329	5.9%
b2b/supplier	36	137	2.5%

CITATION SHARE BY SOURCE TYPE (%)



11.3 Two things this mix shows

Brand-owned content is the largest single input, at 44.2%. Add retailer (8.1%) and marketplace (9.4%) domains and commercial sources account for 61.7% of citations. Read as a description rather than advice: in this market, a current, machine-readable brand site is closely associated with AI visibility.

Editorial and community content is a real presence here — 17.9% combined. Recipe sites, listicles and forums (nefisyemektarifleri.com, onedio.com, listelist.com, kadinlarkulubu.com, donanimhaber.com, ekşisözlük, YouTube, Reddit) supply nearly a fifth of the evidence base. Kitchenware is a category people write about publicly — reviews of pans that warp, dinner sets that survive the dishwasher — and the assistants read it.

A CAVEAT ON CLASSIFICATION

Source types are derived from the entity dictionary for the 128 domains whose owner is known. The remaining 351 are typed by pattern, and the consistency audit can only assert correctness for the resolved subset. Treat the mix as accurate at the top of the distribution and indicative in its tail.

12

SECTION 12

What the Patterns Suggest

Visibility in this market is associated with breadth across models, a strong own-site presence, and a clear category association — but the leaderboard's shape is also determined by corporate structure and by one assistant's answering style.

— WHAT THE PATTERNS SUGGEST

Five structural patterns

ONE-SENTENCE TAKEAWAY

Visibility here is associated with breadth across models, a strong own-site presence and a clear category association — but the leaderboard's shape is also determined by corporate structure and by one assistant's answering style.

This chapter is a neutral reading of the data. It describes patterns associated with visibility; it is not advice, a service, or a recommendation, and visibility remains separate from quality.

- 1 Breadth is close to a precondition.** 19 of the 23 qualified general-cut entities are seen by all five models. The four that are not — Fissler, Trendyol, Hepsiburada and Le Creuset — sit at #13, #20, #21 and #23. In a five-model breadth scale, missing one model costs 5 points of final score outright, and in practice it also signals that a brand's presence in the underlying evidence is uneven.
- 2 Mention and position are independent signals.** Emsan is mentioned in 35.6% of responses and ranks fifth; Porland is mentioned in 24.7% and ranks fourth, because it is named far earlier when named at all (MRR 0.497 against 0.329). At the extreme, Pirge would rank third in the general cut on score alone but does not qualify, and IKOO — mentioned in 7 responses with an MRR of 0.893 — would place near the top and is not ranked either. "Named often" and "named first" are different achievements with different causes.
- 3 Category association concentrates visibility.** The entities that dominate a segment are the ones with a single, legible specialism: the porcelain houses in Tableware, the knife makers in Cutlery, Korkmaz in cookware and whole-kitchen questions. Generalists lead the general cut; specialists lead segments and own the position component.
- 4 The shelf is narrower than it looks.** Three of the top six entities share an owner, and the two marketplaces in the table owe 96%+ of their visibility to one model. A reader counting independent options in the top ten would overcount.
- 5 Origin is contested only in cookware and cutlery.** Turkish brands hold tableware almost completely — one foreign qualifier in 21 — while foreign brands reach #2 in Cutlery & Storage and second overall on the open cut. The local/open gap shows the assistants understand *yerli* as a constraint, so the open-cut balance is a genuine market signal rather than an artefact of question wording.

HOW A BRAND CAN LOCATE ITSELF IN THE DATA

Using the published dataset, a brand can read four things in order: whether it appears in all five models; how often it is mentioned within the cut that matches its category; how early it is named when mentioned; and which model or question type it is weakest in. These readings show where a visibility gap sits, without saying anything about whether its product is good.

13

SECTION 13

Limitations and Notes

This is a single, point-in-time snapshot of one Turkish vertical, read across ten cuts whose scores are not comparable with each other.

— LIMITATIONS AND NOTES

How to read the figures

ONE-SENTENCE TAKEAWAY

This is a single, point-in-time snapshot of one Turkish vertical, read across ten cuts whose scores are not comparable with each other.

- **Single collection window.** Everything here is a snapshot of 1 August 2026, 05:57–10:54 UTC. AI answers move; a re-run on another day will differ. This is the baseline edition; trend analysis becomes possible from the second.
- **Never compare a scaled score across cuts or editions.** 100 in one cut is a different absolute number from 100 in another, and the scale moves whenever the leader moves. Cross-cut comparison uses the raw layer — mention rate, MRR and breadth — which is published in every table.
- **Turkish vs. foreign is only comparable on the open cut.** 22 of the 45 questions explicitly constrain the answer to Turkish brands. The general cut's 17-to-6 origin split describes the question design, not the market.
- **Qualification is a threshold, and thresholds are arbitrary at the margin.** Villeroy & Boch misses the general cut by one response. Pirge (4.00%) and IKOO (0.62%, MRR 0.893) would both rank highly on score and are excluded because ranking is restricted to qualified entities.
- **Components above 100 are real, not errors.** Unqualified entities are scored on the same cut maxima so the appendix stays on one scale; 21 rows (2.3%) exceed 100 on a component. They are flagged and never clipped.
- **Narrow cuts qualify on very few responses.** Cutlery & Storage qualifies at 12.5 responses and Whole-kitchen at 7.5, against 56.25 in the general cut. A 5% rate means something different in each.
- **Whole-kitchen & Trousseau is indicative only** — six questions, no attribute questions, and the heaviest off-market drift in the study.
- **The search-vs-memory comparison is not identifiable.** 96.3% of ungrounded responses are Gemini's, and Gemini's grounded and ungrounded subsets answer different question mixes (68.2% attribute versus 24.9%). Chapter 6 reports the difference and declines the causal claim. Separating retrieval from question type would require forcing search on and off within one model — a design change for a future edition.
- **Marketplace visibility is model-specific.** Trendyol and Hepsiburada owe 96.1% and 97.4% of their mentions to Grok. Their leaderboard positions describe one assistant more than the market.
- **Sub-brands under-count structurally** where a model names the parent instead.
- **Source-type classification is partly heuristic.** 128 of 479 domains resolve to an owner via the dictionary; the rest are typed by pattern, and the classification audit covers only the resolved subset.
- **Precision over recall.** Entity matching was cautious to avoid false positives, so brands mentioned indirectly may be undercounted. Two near-identical names, LAV (glassware) and Lava (cast iron), were confirmed as separate entities with zero response overlap and are deliberately not merged.
- **Models are probabilistic.** The same question can produce different answers; five repeats reduce but do not eliminate this.
- **Visibility is not quality.** This report measures mention and position only. It does not assess whether a brand was described positively, whether a recommendation was accurate, or whether a product is good. A score is not an endorsement.

— FREQUENTLY ASKED QUESTIONS

Frequently asked questions

Which kitchen and tableware brands does AI recommend most in Türkiye?

Karaca leads the general ranking with a perfect score of 100.00, appearing in 57.2% of all 1,125 responses and named by all five assistants. Korkmaz (80.90) is second, followed by Kütahya Porselen (70.95), Porland (70.37) and Emsan (70.11). Leadership changes by segment: Korkmaz leads cookware and whole-kitchen questions, Karaca leads tableware and cutlery.

Are the top-ranked brands independent competitors?

Not all of them. Karaca (#1), Emsan (#5) and Jumbo (#6) are all Karaca Group brands. Together they appear in 727 of 1,125 responses — 64.6% of the study — rising to 72.0% on origin-neutral questions. Four other corporate groups appear in the data, including Groupe SEB (Tefal and WMF), which reaches 41.9% of origin-neutral responses.

Do AI assistants recommend foreign kitchenware brands in Türkiye?

It depends entirely on the segment. On origin-neutral questions, mentions split 63.6% Turkish to 36.4% foreign, and Tefal places second overall. But tableware and serving has just one foreign qualifier out of 21, while cutlery and storage has IKEA at #2 and Zwilling at #4. Foreign visibility concentrates in cookware, knives and storage.

How is the AI Visibility Score calculated?

It combines three parts: how often a brand is mentioned (45%), how early it appears in the answer (30%), and how many of the five models mention it (25%). Each part is divided by the best qualified performance within that cut, using maxima published once in a fixed scale registry. Scores from different cuts are not comparable with each other.

Why does the leader have a perfect score of 100?

Because Karaca simultaneously tops all three components in the general cut — the highest mention rate, the highest MRR and presence in all five models. That is unusual; in eight of the study's ten cuts the mention leader and the position leader are different entities, so the leader falls short of 100. The perfect score describes that coincidence, not a ceiling imposed by the method.

Do the five AI assistants agree with each other?

On the leader, mostly: Karaca is the most-named entity in four of five models, and Perplexity names it in 66.7% of its answers. They differ sharply on breadth — Grok names about 8.15 entities per answer against 4.20 for GPT-5.6 Luna Pro — and Grok is the only model that leads with marketplaces, accounting for over 96% of all Trendyol and Hepsiburada mentions in the study.

Where do AI assistants get their information about these brands?

83.3% of answers were web-grounded. Of the citations, 44.2% point to brands' own websites, led by karaca.com, cited by 34.2% of all responses and by all five models. Editorial and community sources — recipe sites, listicles and forums — supply a further 17.9%, which is a substantial share for this series.

Why do a quarter of the answers name appliance brands?

Because the Turkish word *mutfak* means both “kitchen” and “kitchenware”, and the assistants resolve that ambiguity inconsistently. 25.8% of all responses name a major-appliance, small-appliance or cabinetry brand such as Arçelik, Bosch or Arzum, rising to 51.3% on the six broad whole-kitchen questions. Those entities are classified out-of-scope and excluded from every leaderboard, but the drift itself is reported as a finding.

Does a high score mean a brand is better than the others?

No. The score measures only how often and how prominently a brand is named across the five assistants. It does not assess quality, durability, materials, price, accuracy or sentiment. A score is a measure of visibility, not an endorsement.

14

SECTION 14

Appendix & Data

The study's 45 questions, every matched entity, data availability, and about Herm.io.

APPENDIX & DATA

14.1 All 45 questions

All 45 questions used in the study, as asked, with their product segment, behavioural type and origin scope. **No question was rewritten in this edition.** Distribution: 15 discovery, 15 attribute, 15 use_case; 22 local (ask for *yerli* brands) and 23 open (origin-neutral). The 15 raw subcategories underneath the four segments are carried in the published dataset but are too granular to report on (median size: 2 questions).

#	Segment	Type	Scope	Question (Turkish)
1	Whole-kitchen & Trousseau	discovery	open	mutfak eşyasında hangi markalar gerçekten iyi
2	Cookware	discovery	local	iyi bir yerli tencere markası var mı
3	Tableware & Serving	discovery	open	en çok tutulan yemek takımı markaları
4	Cookware	discovery	local	türk malı tava önerir misin
5	Cookware	discovery	open	tencere tava alırken hangi markalara bakılır
6	Tableware & Serving	discovery	local	yerli porselende hangi markalar kaliteli
7	Cutlery & Storage	discovery	open	çatal kaşık bıçakta güvenilir markalar
8	Whole-kitchen & Trousseau	discovery	local	mutfak ve sofra ürünlerinde iyi yerli markalar neler
9	Cutlery & Storage	discovery	open	en iyi saklama kabı markası hangisi
10	Whole-kitchen & Trousseau	discovery	local	çeyiz için Türk malı mutfak ürünlerinde hangi markalara bakayım
11	Tableware & Serving	discovery	open	servis ve sunum ürünlerinde hangi markalar güzel
12	Cutlery & Storage	discovery	local	iyi bir yerli mutfak bıçağı markası önerisi
13	Tableware & Serving	discovery	open	kahvaltı takımı hangi markadan alınır
14	Tableware & Serving	discovery	local	yerli cam ve sunum ürünlerinde ne önerirsiniz
15	Whole-kitchen & Trousseau	discovery	open	fıyatına değen mutfak markaları hangileri
16	Cookware	attribute	open	çizilmeyen yapışmaz tava hangi marka iyi
17	Cookware	attribute	local	tabanı eğilmeyen Türk malı çelik tencere önerisi
18	Tableware & Serving	attribute	open	bulaşık makinesinde deseni çıkmayan yemek takımı hangi marka
19	Cutlery & Storage	attribute	local	karamayan yerli çatal kaşık bıçak takımı arıyorum
20	Cookware	attribute	open	indüksiyona uygun, ısıyı eşit dağıtan tencere markası
21	Cookware	attribute	local	kaplaması çabuk kalkmayan yerli tava var mı
22	Cookware	attribute	open	fırında çatlamayan cam fırın kabı hangi marka iyi

(cont.)

#	Segment	Type	Scope	Question (Turkish)
23	Cookware	attribute	local	18/10 çelik Türk malı tencere seti önerir misin
24	Tableware & Serving	attribute	open	yaldızı makinede bozulmayan yemek takımı markası
25	Cutlery & Storage	attribute	local	uzun süre keskin kalan yerli mutfak bıçağı hangisi
26	Cutlery & Storage	attribute	open	sızdırmayan, koku yapmayan saklama kabı markası
27	Tableware & Serving	attribute	local	uygun fiyatlı kaliteli yerli porselen takımı
28	Cookware	attribute	open	kulpu ısınmayan hafif düdüklü tencere hangi marka iyi
29	Cutlery & Storage	attribute	local	az yer kaplayan Türk malı saklama kabı seti önerisi
30	Tableware & Serving	attribute	local	sade ve zamansız tasarımı yerli kahvaltı takımı var mı
31	Cookware	use_case	open	çeyiz için uzun yıllar kullanılacak tencere seti hangi marka
32	Whole-kitchen & Trousseau	use_case	local	yeni eve çıkıyorum, iyi yerli mutfak markaları önerir misin
33	Tableware & Serving	use_case	open	hem günlük hem misafir için yemek takımı hangi marka iyi
34	Cookware	use_case	local	anneme hafif bir Türk malı tencere seti, ne önerirsin
35	Tableware & Serving	use_case	open	kalabalık misafir için 12 kişilik yemek takımı önerisi
36	Whole-kitchen & Trousseau	use_case	local	çeyiz için tek markadan yerli mutfak seti alacağım, hangisi iyi
37	Cutlery & Storage	use_case	open	küçük mutfakta yer kaplamayan saklama kabı hangi marka
38	Tableware & Serving	use_case	local	yeni ev hediyesi için Türk malı kahve fincanı seti önerir misin
39	Cutlery & Storage	use_case	open	her gün bulaşık makinesine girecek çatal kaşık takımı hangi marka
40	Tableware & Serving	use_case	local	sık misafir ağırlayanlar için yerli servis ve sunum markası
41	Cookware	use_case	open	iki kişilik ev için gereksiz parçası olmayan tencere seti önerisi
42	Cookware	use_case	local	indüksiyon ocağa yeni geçtim, yerli tencere markası önerir misin
43	Cutlery & Storage	use_case	open	işe yemek götürmek için sızdırmayan saklama kabı hangi marka
44	Tableware & Serving	use_case	local	Türk kahvesi sunumu için şık bir yerli fincan markası var mı
45	Tableware & Serving	use_case	open	düğün hediyesi olarak kaliteli yemek takımı hangi markadan alınır

14.2 All 139 matched entities

All 139 entities matched at least once, general cut (1,125 responses). Rate is within the general cut, Models is breadth (0–5). Entities at or above 5% are marked **Ranked**; the rest are **Tracked**. Five further dictionary entities — A101, Alkapıda, Boyner, Evidea and Primanova — were never mentioned in any response and are retained in the dictionary rather than deleted. Sorted by general-cut score rank.

Entity	Origin	Type	Ment.	Rate	Models	Status
Karaca	TR	Brand/Ret.	643	57.2%	5	Ranked
Korkmaz	TR	Brand	378	33.6%	5	Ranked
Kütahya Porselen	TR	Brand	303	26.9%	5	Ranked
Porland	TR	Brand	278	24.7%	5	Ranked
Emsan	TR	Brand	400	35.6%	5	Ranked
Jumbo	TR	Brand/Ret.	336	29.9%	5	Ranked
Tefal	Foreign	Brand	193	17.2%	5	Ranked
IKEA	Foreign	Licensed	151	13.4%	5	Ranked
Güral Porselen	TR	Brand	207	18.4%	5	Ranked
Paşabahçe	TR	Brand/Ret.	148	13.2%	5	Ranked
WMF	Foreign	Brand	135	12.0%	5	Ranked
Nehir	TR	Brand	66	5.9%	5	Ranked
Fissler	Foreign	Brand	84	7.5%	4	Ranked
Schafer	TR	Brand	201	17.9%	5	Ranked
Lava	TR	Brand	93	8.3%	5	Ranked
Zwilling	Foreign	Brand	159	14.1%	5	Ranked
Aryıldız	TR	Brand	171	15.2%	5	Ranked
Hisar	TR	Brand	110	9.8%	5	Ranked
Taç	TR	Brand	83	7.4%	5	Ranked
Trendyol	TR	Marketpl.	205	18.2%	4	Ranked
Hepsiburada	TR	Marketpl.	196	17.4%	4	Ranked
English Home	TR	Retailer	62	5.5%	5	Ranked
Le Creuset	Foreign	Brand	57	5.1%	4	Ranked
Pirge	TR	Brand	45	4.0%	5	Tracked
Bonera	TR	Brand	20	1.8%	3	Tracked
Sürbisa	TR	Brand	42	3.7%	5	Tracked
IKOO	Foreign	Brand	7	0.6%	2	Tracked
Vivaldi	TR	Brand	7	0.6%	3	Tracked
Deftion	TR	Brand	12	1.1%	3	Tracked
Lock&Lock	Foreign	Brand	42	3.7%	3	Tracked

(cont.)

Entity	Origin	Type	Ment.	Rate	Models	Status
ESTAŞ	TR	Brand	9	0.8%	4	Tracked
Hascevher	TR	Brand	17	1.5%	4	Tracked
Tupperware	Foreign	Brand	35	3.1%	4	Tracked
Joseph Joseph	Foreign	Brand	19	1.7%	4	Tracked
Stanley	Foreign	Brand	14	1.2%	4	Tracked
Circulon	Foreign	Brand	11	1.0%	2	Tracked
Papilla	TR	Brand	39	3.5%	5	Tracked
Şişecam	TR	Brand	27	2.4%	4	Tracked
Serenk	TR	Brand	15	1.3%	4	Tracked
CEM Mutfak	TR	Brand	29	2.6%	4	Tracked
Soley	TR	Brand	13	1.2%	3	Tracked
Demeyere	Foreign	Brand	13	1.2%	3	Tracked
Villeroy & Boch	Foreign	Brand	56	5.0%	5	Tracked
Bonna	TR	Brand	38	3.4%	5	Tracked
Acar	TR	Brand/Ret.	55	4.9%	4	Tracked
Altınbaşak	TR	Brand/Ret.	5	0.4%	3	Tracked
Titiz Plastik	TR	Brand	16	1.4%	4	Tracked
Bambum	TR	Brand	35	3.1%	5	Tracked
Keramika	TR	Brand	15	1.3%	5	Tracked
Vienev	TR	Brand	19	1.7%	3	Tracked
Beymen	TR	Retailer	12	1.1%	4	Tracked
Vip Ahmet	TR	Brand	16	1.4%	4	Tracked
OMS Collection	TR	Brand	32	2.8%	3	Tracked
Bernardo	TR	Brand/Ret.	45	4.0%	5	Tracked
Glasslock	Foreign	Brand	10	0.9%	3	Tracked
Koçtaş	TR	Retailer	4	0.4%	1	Tracked
Kahramanlar	TR	Brand	7	0.6%	3	Tracked
Caso	Foreign	Licensed	7	0.6%	3	Tracked
Pyrex	Foreign	Licensed	38	3.4%	3	Tracked

(cont.)

Entity	Origin	Type	Ment.	Rate	Models	Status
Bialetti	Foreign	Brand	5	0.4%	1	Tracked
Victorinox	Foreign	Brand	24	2.1%	5	Tracked
Baci Milano	Foreign	Brand	9	0.8%	2	Tracked
Madame Coco	TR	Retailer	32	2.8%	5	Tracked
LAV	TR	Brand	30	2.7%	4	Tracked
Gülsan	TR	Brand	10	0.9%	3	Tracked
Rubbermaid	Foreign	Brand	16	1.4%	3	Tracked
Öztiryakiler	TR	Brand	18	1.6%	4	Tracked
Simax	Foreign	Sub-brand	8	0.7%	3	Tracked
Pekşen	TR	Brand/Ret.	7	0.6%	2	Tracked
Amboss	TR	Brand	8	0.7%	3	Tracked
Neoflam	Foreign	Brand	14	1.2%	4	Tracked
Flora	TR	Brand	13	1.2%	3	Tracked
GreenPan	Foreign	Brand	18	1.6%	4	Tracked
ONON	TR	Brand/Ret.	9	0.8%	2	Tracked
Marks & Spencer	Foreign	Brand/Ret.	2	0.2%	2	Tracked
Mehtap	TR	Brand	8	0.7%	3	Tracked
Gallery Crystal	TR	Retailer	1	0.1%	1	Tracked
All-Clad	Foreign	Brand	14	1.2%	3	Tracked
Tchibo	Foreign	Brand/Ret.	7	0.6%	3	Tracked
Nude Glass	TR	Sub-brand	20	1.8%	3	Tracked
Soy	TR	Brand/Ret.	9	0.8%	3	Tracked
Farika Porselen	TR	Brand	7	0.6%	4	Tracked
Lenox	Foreign	Brand	5	0.4%	2	Tracked
OXO	Foreign	Brand	19	1.7%	3	Tracked
Rivadossi	Foreign	Brand	8	0.7%	2	Tracked
Minihol	TR	Brand/Ret.	5	0.4%	1	Tracked
Ballarini	Foreign	Sub-brand	6	0.5%	3	Tracked
Alessi	Foreign	Brand	16	1.4%	3	Tracked

(cont.)

Entity	Origin	Type	Ment.	Rate	Models	Status
Brioni	TR	Brand	7	0.6%	2	Tracked
Rosenthal	Foreign	Brand	7	0.6%	3	Tracked
Brillant	TR	Brand/Ret.	2	0.2%	1	Tracked
Falez	TR	Brand	2	0.2%	1	Tracked
Staub	Foreign	Brand	40	3.6%	3	Tracked
İpek Porselen	TR	Sub-brand	12	1.1%	3	Tracked
GGs Solingen	Foreign	Brand	5	0.4%	2	Tracked
Corelle	Foreign	Brand	15	1.3%	2	Tracked
Tulü Porselen	TR	Brand	8	0.7%	3	Tracked
Meleni Home	TR	Brand	7	0.6%	2	Tracked
Sambonet	Foreign	Brand	6	0.5%	4	Tracked
Hecha	TR	Sub-brand	15	1.3%	3	Tracked
Tramontina	Foreign	Brand	9	0.8%	2	Tracked
Thermos	Foreign	Brand	4	0.4%	3	Tracked
Bella Maison	TR	Brand/Ret.	28	2.5%	3	Tracked
Lucky Art	TR	Brand/Ret.	9	0.8%	2	Tracked
Scanpan	Foreign	Brand	2	0.2%	2	Tracked
Philippi	Foreign	Brand	6	0.5%	3	Tracked
24Bottles	Foreign	Brand	8	0.7%	2	Tracked
Lazoğlu	TR	Brand	7	0.6%	1	Tracked
RCR	Foreign	Brand	5	0.4%	3	Tracked
Woll	Foreign	Brand	9	0.8%	2	Tracked
Selamlıque	TR	Brand	2	0.2%	2	Tracked
Wüsthof	Foreign	Brand	11	1.0%	3	Tracked
Arna Porselen	TR	Brand	4	0.4%	2	Tracked
Luminarc	Foreign	Brand	5	0.4%	2	Tracked
Lodge	Foreign	Brand	8	0.7%	3	Tracked
Pip Studio	Foreign	Brand	13	1.2%	2	Tracked
HexClad	Foreign	Brand	5	0.4%	2	Tracked

(cont.)

Entity	Origin	Type	Ment.	Rate	Models	Status
Guzzini	Foreign	Brand	7	0.6%	2	Tracked
Lazbisa	TR	Brand	8	0.7%	1	Tracked
Frigoverre	Foreign	Sub-brand	6	0.5%	1	Tracked
Kochler	TR	Brand	6	0.5%	2	Tracked
Has Çömlek	TR	Brand	7	0.6%	2	Tracked
n11	TR	Marketpl.	13	1.2%	2	Tracked
Snapware	Foreign	Sub-brand	8	0.7%	2	Tracked
Neva	TR	Brand	5	0.4%	2	Tracked
Tantitoni	TR	Brand/Ret.	7	0.6%	2	Tracked
Mepal	Foreign	Brand	1	0.1%	1	Tracked
Miniso	Foreign	Licensed	1	0.1%	1	Tracked
Chakra	TR	Brand/Ret.	8	0.7%	2	Tracked
Kitchen Life	TR	Retailer	9	0.8%	2	Tracked
Anchor Hocking	Foreign	Brand	2	0.2%	1	Tracked
Zojirushi	Foreign	Brand	2	0.2%	1	Tracked
Leggno	TR	Brand/Ret.	7	0.6%	2	Tracked
Bordallo Pinheiro	Foreign	Brand	5	0.4%	1	Tracked
Amazon	Foreign	Marketpl.	12	1.1%	1	Tracked
Perotti	TR	Brand	1	0.1%	1	Tracked
PttAVM	TR	Marketpl.	1	0.1%	1	Tracked
Herend	Foreign	Brand	1	0.1%	1	Tracked
Çiçeksepeti	TR	Marketpl.	2	0.2%	1	Tracked

14.3 Data availability

The study decided 399 candidate entities, kept 144, and matched 139 in at least one response; 23 cleared the 5% threshold in the general cut, with 21 qualifying in the local cut and 27 in the open cut. To support scrutiny and reproduction, the underlying data is available on request: response-level entity mentions, all ten qualified leaderboards and their brand-only variants, full unfiltered metrics for every entity in every cut, the scale registry, per-model and per-question-type breakdowns, the corporate-group roll-up, the off-market drift analysis, the complete 479-domain citation list, and the methodology and data dictionary.

All figures reflect a single point-in-time run (1 August 2026) and are subject to the limitations in Chapter 13. This is the first (baseline) measurement; the study will be repeated quarterly.

14.4 About Herm.io & disclosure

Herm.io is a consumer behaviour and marketing data company. We study how people discover and choose brands to help businesses reach the right customers. This report forms part of our public research and is conducted quarterly.

DISCLOSURE & NEUTRALITY

The rankings and metrics in this report are entirely impartial and are based on the methodology described within the report. No brand can pay or provide sponsorship to be included in the report, improve its ranking or influence how it is described in the content. A brand's score in the report is not an endorsement or assessment of quality; it is solely a measure of its current visibility within AI models. No service or product offered by Herm.io is sold for the purpose of directly changing the results of this report or guaranteeing a ranking. To preserve objectivity, all references to Herm.io's own domain are removed from the source data before analysis. This ensures that the company does not appear in, measure itself through or benefit from its own study.

REPORT PERIOD

August 2026 · Edition 1 (baseline) of a quarterly series · Collection window 1 August 2026.