herm.io

AUGUST 2026 · SMALL HOME APPLIANCES · TURKEY

HERM.IO RESEARCH · TURKEY SMALL HOME APPLIANCES

AI Visibility Report

The first AI-visibility snapshot of Turkey's small home appliance market: how five large language models answer 45 Turkish-language questions about kettles, coffee machines, airfryers, blenders, toasters and stick vacuums — read across three cuts and three product segments.

1,125	45	76	25
Total Responses	Questions Analyzed	Entities Tracked	Qualified Entities

Author: Mert Can Elkaya · 13 August 2026 · Herm.io AI Visibility Database

KEY FINDINGS

At a glance

This report reproduces the real Turkish-language questions consumers ask AI assistants about kettles, coffee machines, airfryers, blenders, toasters and stick vacuums, and measures how often, in what order, and across how many models five large language models name each brand.

1. **Karaca leads the main leaderboard at 92.95**, named in 64.1% of all 1,125 responses, with Arzum (90.24) and Arçelik (85.89) close behind — a genuine three-way race at the top.
2. **The shelf flips when the question stops asking for Turkish brands:** in the origin-neutral open cut, Philips leads at 94.38 and foreign brands take 59.6% of mentions and 17 of 27 qualifying places.
3. **Of 76 tracked entities, 25 cleared the 5% threshold in the main cut.** 17 of the qualifiers appear in only one of the three product segments — the market is a small general core plus three largely separate specialist shelves.
4. **Gemini grounded only 50.2% of its answers** against 100% for the other four models, and the split changes what it names: 63.3% of its grounded mentions go to Turkish brands versus 52.3% from memory alone.
5. **AI's picture of this market is editorial-led, not brand-led:** editorial and media domains supply 39.7% of citations against 19.0% for brands' own sites, and the single most-cited domain, donanimhaber.com, is cited by 47.4% of all responses.

SCOPE & METHOD

Vertical: small home appliances (küçük ev aletleri), Turkey. Method: a single point-in-time study of 1,125 responses across five large language models (ChatGPT, Gemini, Claude, Perplexity, Grok), each asked 45 Turkish-language questions five times. Collection date: 13 August 2026.

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 small home appliances in Turkey. 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 built, more durable, cheaper or safer. A low score, or a zero, means a brand is currently less visible to these systems, not that it is inferior. AI visibility reflects information availability and discovery, not product quality, reliability, price or after-sales service. 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, small home appliance brands, küçük ev aletleri, coffee machines, airfryers, blenders, stick vacuums, çeyiz sets, Turkish (*yerli*) appliance brands, LLM brand recommendations, ChatGPT/Gemini/Claude/Perplexity/Grok brand recommendation, Generative Engine Optimization (GEO).

Because the same questions can be read with or without an origin restriction, and because the category covers several genuinely different product shelves, the entities are ranked in three cuts and three segments, all defined in Chapter 2. It is a single point-in-time study and should be read within the limitations in Chapter 11.

CONTENTS

Report contents

01 Executive Summary 04
02 Methodology 06
03 Overall Visibility Leaderboards 11
04 Differences Between Models 16
05 Web Search or Model Memory? 20
06 Segment Ownership 23
07 Question-Type Ownership 27
08 Open Market: Turkish vs. Foreign 29
09 The Discovery Ecosystem 32
10 What the Patterns Suggest 35
11 Limitations and Notes 37
12 Appendix & Data 41

01

SECTION 01

Executive Summary

Three Turkish names — Karaca, Arzum and Arçelik — lead the market as a whole, but the moment a question stops asking for Turkish brands the shelf turns over to Philips and a wide foreign field, and beneath the leaders the category splits into three largely separate specialist shelves.

EXECUTIVE SUMMARY

A three-way race, and a shelf that turns over

ONE-SENTENCE TAKEAWAY

Three Turkish names — Karaca, Arzum and Arçelik — lead the market as a whole, but the moment a question stops asking for Turkish brands the shelf turns over to Philips and a wide foreign field, and beneath the leaders the category splits into three largely separate specialist shelves.

1. **A three-way race at the top.** Karaca leads the main cut at 92.95, named in 64.1% of all 1,125 responses. Arzum (90.24, 57.7%) and Arçelik (85.89, 49.3%) follow within 7.1 points. Philips (74.67) is fourth and the highest-placed foreign brand; Dyson (62.99) is fifth. The average score of the 25 qualified entities is 54.42.

2. **25 entities qualified from a dictionary of 76.** That is a narrow shelf. Fourteen of the 25 are Turkish-heritage brands and eleven are foreign. Twenty-two of the 25 are seen by all five models.

3. **The open cut reverses the picture.** On the 23 origin-neutral questions, Philips leads at 94.38 — named in 64.5% of those responses — ahead of Karaca (81.42), Arzum (78.14) and Tefal (74.71). Foreign brands take 59.6% of open-cut mentions and 17 of the 27 qualifying places. The all-questions figure of 61.9% Turkish is an artefact of the question mix, not a market reading; only the open cut can carry an origin claim.

4. **The *yerli* filter mostly works, with two revealing exceptions.** On the 22 questions that explicitly ask for Turkish brands, 15 of 19 qualifiers are Turkish-heritage. The four that are not include Fakir (#9) and Grundig (#13) — German-founded brands now owned by Turkish groups, which the assistants evidently read as domestic. Under the origin rule used here they are counted foreign; see §2.6.

5. **The category is three shelves, not one.** Only 13 qualified entities clear the bar in all three product segments. Seventeen qualify in exactly one: De'Longhi, Jura and Nespresso only in Coffee & Tea; Dyson, Dreame, Shark and Samsung only in the vacuum-led Other segment; Dökümix, Vitamix, Ninja and Cosori only in Kitchen Appliances.

6. **Two model behaviours distort the aggregate, and both are named.** Grok names 9.73 entities per response against Claude's 3.68, and puts the marketplace Hepsiburada in 96.0% of its answers — the marketplace positions in every leaderboard are substantially a Grok artefact. Separately, Gemini grounded only 50.2% of its responses where the other four models grounded 100%.

7. **Search changes what one model says.** Within Gemini, 63.3% of grounded mentions go to Turkish brands against 52.3% from memory alone. Karaca gains 20.8 points when it searches; De'Longhi, Sage, Breville, Roborock and KitchenAid all fall away.

8. **Sources are editorial-led.** Editorial and media domains supply 39.7% of citations against 19.0% for brands' own sites. The single most-cited domain, the technology forum donanimhaber.com, is cited by 47.4% of all responses — more than any brand or marketplace site.

Why it matters

Shoppers are shifting discovery from search engines and marketplaces to AI assistants, and small home appliances is a high-consideration, high-repeat category where an assistant's shortlist arrives well before the shopper reaches a product page. Whether a brand appears in that shortlist is becoming a discovery channel in its own right. This report captures the first snapshot (baseline) of that shelf for Turkey. One reminder: the numbers below describe visibility and information availability, not which brands are best.

02

SECTION 02

Methodology

Five large language models were asked 45 Turkish questions five times each, with no system prompt, producing 1,125 responses with none empty and none excluded. Entities were identified by automated extraction plus human review and ranked with a three-component (45/30/25) score, computed separately for three cuts, three segments and three question types.

METHODOLOGY

How the study was run

ONE-SENTENCE TAKEAWAY

Five large language models were asked 45 Turkish questions five times each, with no system prompt, producing 1,125 responses with none empty and none excluded. Entities were identified by automated extraction plus human review and ranked with a three-component (45/30/25) score, computed separately for three cuts, three segments and three question types.

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.

THE FIVE MODELS QUERIED — 225 responses each

Model	Version	Responses	Web-search rate
ChatGPT 5.6 Luna Pro	openai/gpt-5.6-luna-pro	225	100.00%
Gemini 3.6 Flash	gemini-3.6-flash	225	50.22%
Claude Haiku 4.5	anthropic/claude-haiku-4.5	225	100.00%
Perplexity Sonar	perplexity/sonar	225	100.00%
Grok 4.3	x-ai/grok-4.3	225	100.00%

NOTE

Web-search rate = share of the model's answers that returned at least one citation or search result.

Scale: 45 questions × 5 repeats × 5 models = 1,125 responses. Every response came back with a body, so the collected and analysed counts are identical — unusual, and not something to assume in a later edition. Single run, 13 August 2026.

2.2 THE THREE CUTS

The same 45 questions are read three ways:

THE THREE MARKET CUTS — each cut has exactly one denominator

Cut	Questions	Responses	Qualified	What it answers
Main	45 (all)	1,125	25	Overall visibility across every question
Local	22 (ask for yerli brands)	550	19	Who is named when the question asks for Turkish brands
Open	23 (origin-neutral)	575	27	Fair Turkish-vs-foreign comparison (no origin restriction)

NOTE

Local + open = the full 45-question set; the main cut spans both.

The three segments. The 45 questions also divide by product shelf, and the segments behave differently enough to be worth ranking separately (Chapter 6):

THE THREE PRODUCT SEGMENTS kitchen · coffee & tea · other

Segment	Questions	Responses	Distinct entities	Qualified
Kitchen Appliances (Mutfak Cihazları)	19	475	51	23
Coffee & Tea (Kahve & Çay)	13	325	45	20
Other (Diğer)	13	325	62	27

Kitchen covers blenders, airfryers and toasters; Coffee & Tea covers coffee and tea machines; Other combines stick vacuums with çeyiz-set and general questions. The segment and question-type groupings each sum back to 1,125 responses, as do local plus open. These reconciliations are asserted before any figure is published rather than checked by eye.

2.3 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 was 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 90.0% of answers were grounded, and the entire ungrounded remainder belongs to one model (see Chapter 5).

2.4 QUESTIONS

The 45 questions reproduce the real Turkish-language queries people bring to an assistant about small appliances: open discovery (*"küçük ev aletlerinde hangi marka iyi"*), specific attributes (noise, durability, capacity, price band), and use cases (çeyiz sets, Turkish coffee, filter coffee, pet hair, small flats). They were asked exactly as written, including their ambiguities; none was rewritten, added or removed after the answers were seen.

They are balanced across the three behavioural types — 15 discovery, 15 attribute, 15 use_case — and split 22 local (asking for *yerli* brands) to 23 open (origin-neutral). The Turkish-versus-foreign comparison in Chapter 8 rests solely on the 23 open questions, the only fair basis for the comparison, since the local questions cannot surface foreign brands by construction. All 45 are listed in §12.1.

2.5 ENTITY EXTRACTION, MERGING AND VALIDATION

Entities were extracted in two stages. **(1) Automated extraction:** candidate names and citation domains were collected from the response text. **(2) Human review:** 156 candidates went to manual review, of which 76 were kept as entities, 71 were merged into a parent and 9 were dropped; 264 aliases are in force. A further 3,944 extraction candidates were filtered before review as generic vocabulary, model codes or redundant n-grams, each recorded with a reason. Nothing was dropped silently.

Merging is applied at the alias layer and the matcher re-run, never by editing aggregate rows, because a merge re-ranks every other entity in the affected responses. The effect is not cosmetic: De'Longhi was extracted at 23 responses and finished at 133 once Magnifica, Dedica, Eletta, Rivelia and Dinamica were folded in.

Sub-brands. Two sub-brands — Arzum's OKKA and Arçelik's Telve — were merged into their parents after the data showed they are perfect subsets: all 226 OKKA responses also name Arzum, and all 143 Telve responses also name Arçelik. Neither ever appears alone, so a separate row would have double-counted in the reader's eye without adding information. The nesting is itself a finding: the OKKA line drives 226 of Arzum's 649 mentions, and Telve 143 of Arçelik's 555. Sub-brands under-count structurally wherever a model names the parent instead.

Matching is alias-based, word-boundary anchored and case-insensitive at the candidate stage, run on URL-stripped text, with an explicit Turkish translation table (*I*→*ı*, *İ*→*i*, plus *Ş Ğ Ü Ö Ç*) applied before folding. Naive lowercasing is never used: it corrupts the dotted and dotless I and splits a single brand into two spellings.

URL STRIPPING IS LOAD-BEARING HERE

Three of the most-cited domains belong to three of the highest-ranked brands. Without stripping link targets and domain-style anchors, exactly the brands at the top of the table would have been inflated by their own citations.

Short-alias audit. Every alias of six characters or fewer was checked against in-context samples and against how often the parent brand appears in the same response. One alias was removed: *Mi* for Xiaomi matched 147 responses of which only 22% named Xiaomi anywhere, because *mi/mı* is the Turkish interrogative particle. Left in, it would have placed Xiaomi in the top 15 on grammar alone. Every other short alias co-occurs with its parent brand in 90–100% of matches.

2.6 AI VISIBILITY SCORE (0–100)

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

Component	Weight	Basis
Mention	45%	mention rate within the cut (an entity counts once per response)
Position	30%	MRR — the reciprocal rank of the entity's first mention
Breadth	25%	number of models covering it (0–5)

Each component is scaled so the leading qualified entity in each cut reaches 100, then combined with the weights. Maxima are taken over qualified entities within the cut and published once, so nothing recomputes them downstream. **Qualification (≥5% rule):** only entities mentioned in at least 5% of that cut's responses are ranked. Unfiltered metrics for all 76 entities appear in §12.2.

A WORKED EXAMPLE

Karaca is named in 721 of the main cut's 1,125 responses, a mention rate of 0.6409 — the highest among qualified entities, so its mention component is 100.0. Its MRR of 0.4993 against the cut maximum of 0.6527 (held by Dyson) gives a position component of 76.5. It appears in all five models, so breadth is 100.0. The composite is 0.45(100.0) + 0.30(76.5) + 0.25(100.0) = 92.95. That the leader scores 92.95 rather than 100 is expected, not a defect: no single entity tops all three components. Components are never clipped at 100.

2.7 ORIGIN RULE

Published origin is the brand's heritage — its country of founding. Current trademark ownership is retained as a working column so any table can be re-cut without re-research, but it is not what the Origin column reports.

The two differ for several of the largest brands here, and the choice moves the headline. **Fakir** (founded in Germany in 1933, owned by Türkiye's Saruhan Holding since 2009) and **Grundig** (founded in Germany in 1945, owned by Arçelik since 2007) are published as **Foreign**, though both are Turkish-owned today. **Profilo** (Turkish-founded, owned by Germany's BSH since 1995) is published as **Local**. **Schafer** is Turkish-founded and published as Local; its German-sounding name is a naming choice, not a heritage.

The same divergence exists among brands where it does not change the published label: Krups, Rowenta and WMF are German-heritage under French ownership, Kenwood is British-heritage under Italian ownership, and Gaggia is Italian-heritage under Dutch ownership. All are foreign on either reading, so the Turkish/foreign split is unaffected by them.

A reader who expects Fakir or Grundig to be counted as Turkish is not wrong about who owns them — they are reading a different question than the one this column answers. Chapter 8 and §7 both flag where this choice is doing work.

2.8 NEUTRALITY AND SELF-EXCLUSION

This is a market-wide, neutral study with no focus brand. To prevent any conflict of interest, citations to Herm.io's own domain were excluded from the published source analysis, so the company's own content neither appears among the most-cited domains nor influences the reported picture. In this run herm.io was cited by 10 of 1,125 responses (0.17% of citations), too small to move any share reported in Chapter 9.

03

SECTION 03

Overall Visibility Leaderboards

The main cut is a Turkish three-way race with a foreign brand fourth; the local cut tightens it further; and the open cut in Chapter 8 turns the shelf over entirely.

OVERALL VISIBILITY LEADERBOARDS

Three names, seven points apart

ONE-SENTENCE TAKEAWAY

The main cut is a Turkish three-way race with a foreign brand fourth; the local cut tightens it further; and the open cut in Chapter 8 turns the shelf over entirely.

As a reminder, these rank visibility, not product quality, durability, price or after-sales service.

3.1 Main leaderboard: 25 qualified entities

TABLE 1 — AI VISIBILITY SCORE (MAIN)

45 questions · 1,125 responses · 25 qualified

#	Entity	Origin	Mention	MRR	AI Score
1	**Karaca**	TR	64.09%	0.499	92.95
2	Arzum	TR	57.69%	0.538	90.24
3	Arçelik	TR	49.33%	0.571	85.89
4	Philips	Foreign	36.00%	0.531	74.67
5	Dyson	Foreign	11.38%	0.653	62.99
6	Korkmaz	TR	25.16%	0.434	62.60
7	Beko	TR	28.18%	0.331	59.98
8	Tefal	Foreign	28.53%	0.322	59.85
9	Dökümix	TR	5.33%	0.633	57.84
10	De'Longhi	Foreign	11.82%	0.448	53.87
11	Bosch	Foreign	16.53%	0.370	53.63
12	Braun	Foreign	10.84%	0.401	51.03
13	Fakir	Foreign	19.02%	0.265	50.55
14	Homend	TR	16.09%	0.303	50.22
15	Kumtel	TR	6.13%	0.411	48.22
16	Xiaomi	Foreign	10.49%	0.342	48.10
17	Arnica	TR	9.69%	0.354	48.09
18	Dreame	Foreign	5.33%	0.372	45.83
19	Vestel	TR	9.07%	0.299	45.09
20	Sinbo	TR	5.96%	0.268	41.50
21	Siemens	Foreign	5.33%	0.273	41.27
22	Hepsiburada	TR	19.82%	0.148	40.70

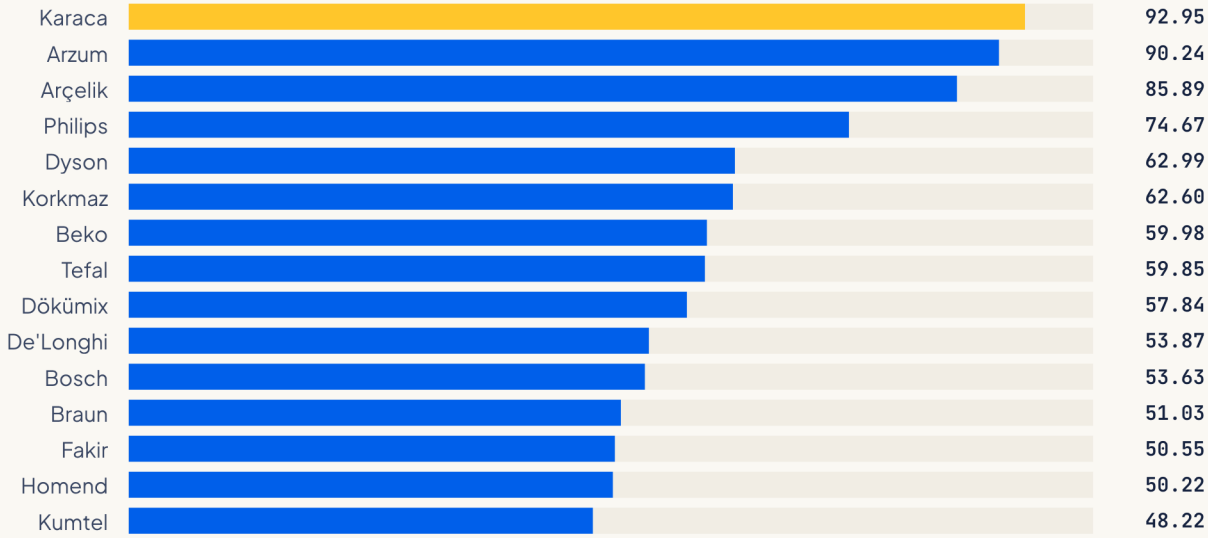
(cont.)

#	Entity	Origin	Mention	MRR	AI Score
23	Grundig	Foreign	7.64%	0.217	40.32
24	Trendyol	TR	19.82%	0.173	36.87
25	Akakçe	TR	9.42%	0.142	18.14

AVERAGE

Average score of the 25 qualified entities: 54.42. Fourteen are Turkish-heritage brands and eleven foreign; 22 of the 25 are seen by all five models.

AI VISIBILITY SCORE — TOP 15 ENTITIES (MAIN CUT)



Reading. Karaca maxes out the mention and breadth components and tops the table at 92.95 on a 64.1% mention rate. Arzum (90.24) and Arçelik (85.89) are within 7.1 points, and Arçelik has the better position component of the three (MRR 0.5711 against Karaca's 0.4993) — it is named less often but earlier. Philips (74.67) is fourth and the first foreign name.

Two ranks need reading with care. **Dyson is fifth on an 11.4% mention rate**, carried there by the highest MRR in the cut (0.6527): it is named in a minority of answers but very close to the front when it is named. **Dökümix is ninth on a 5.33% mention rate**, barely above the qualification threshold, on an MRR of 0.6330. When this Konya manufacturer is named at all, it is named first. Both are genuine measurements of a real pattern, but neither is a volume position; see Chapter 11.

At the other end, the three marketplaces sit low despite high mention rates — Hepsiburada (#22) and Trendyol (#24) are each named in 19.8% of responses but very late (MRR 0.1476 and 0.1731), and Akakçe (#25) is visible to a single model. Their positions are also heavily shaped by one model's behaviour, which §4.3 quantifies.

3.2 Brand-only view (marketplaces removed)

Three of the 25 qualified entities are marketplaces rather than brands: Trendyol, Hepsiburada and Akakçe. They are kept in the headline ranking because that is how shoppers meet the shelf, but removing them gives a brand-versus-brand picture. Karaca, Korkmaz and MediaMarkt are retailers or brand-retailers rather than marketplaces and are retained.

BRAND-ONLY TOP 10 (MAIN) marketplaces set aside · retailers retained

#	Entity	Origin	Mention	MRR	AI Score
1	**Karaca**	TR	64.09%	0.499	92.95
2	Arzum	TR	57.69%	0.538	90.24
3	Arçelik	TR	49.33%	0.571	85.89
4	Philips	Foreign	36.00%	0.531	74.67
5	Dyson	Foreign	11.38%	0.653	62.99
6	Korkmaz	TR	25.16%	0.434	62.60
7	Beko	TR	28.18%	0.331	59.98
8	Tefal	Foreign	28.53%	0.322	59.85
9	Dökümix	TR	5.33%	0.633	57.84
10	De'Longhi	Foreign	11.82%	0.448	53.87

Because all three marketplaces rank in the bottom four of the main cut, removing them changes nothing above position 21 — the brand-only top ten is identical to the main top ten. That is itself worth noting: unlike categories where aggregators crowd the front of the shelf, here the assistants answer appliance questions with manufacturers first and reach for marketplaces late, if at all.

3.3 Local leaderboard: 19 qualified entities (Turkish-only questions)

When the 22 questions that explicitly ask for Turkish or *yerli* brands are isolated, the top three tighten into a near-tie:

TABLE 2 — AI VISIBILITY SCORE (LOCAL / YERLI) 22 questions · 550 responses · 19 qualified

#	Entity	Origin	Mention	MRR	AI Score
1	**Karaca**	TR	74.00%	0.564	96.56
2	Arçelik	TR	67.09%	0.637	95.80
3	Arzum	TR	70.36%	0.575	94.85
4	Korkmaz	TR	38.73%	0.444	69.45
5	Beko	TR	38.55%	0.367	65.72
6	Arnica	TR	16.18%	0.382	52.83
7	Kumtel	TR	9.64%	0.432	51.19
8	Homend	TR	21.64%	0.260	50.40
9	Fakir	Foreign	19.45%	0.276	49.83
10	Vestel	TR	18.00%	0.293	49.75
11	Sinbo	TR	9.27%	0.298	44.69
12	Philips	Foreign	6.18%	0.294	42.61

(cont.)

#	Entity	Origin	Mention	MRR	AI Score
13	Grundig	Foreign	9.27%	0.243	42.07
14	Senur	TR	5.82%	0.253	40.45
15	Tefal	Foreign	6.00%	0.214	38.70
16	Schafer	TR	6.00%	0.284	37.03
17	Hepsiburada	TR	20.36%	0.146	34.25
18	Trendyol	TR	20.18%	0.153	29.47
19	Akakçe	TR	9.64%	0.115	16.28

AVERAGE

Average score of the 19 qualified local entities: 52.73. Fifteen of the 19 are Turkish-heritage brands.

Reading. Karaca (96.56, 74.0%), Arçelik (95.80, 67.1%) and Arzum (94.85, 70.4%) are separated by 1.72 points — inside the range where small differences should not be over-interpreted. Arçelik holds second on position rather than volume: it is named less often than Arzum but earlier (MRR 0.6372 against 0.5747). Korkmaz (69.45) and Beko (65.72) follow, and the cut then opens out to specialist domestic names — Arnica, Kumtel, Homend, Vestel, Sinbo, Senur — that do not all survive in the open cut.

Four foreign-heritage brands qualify on questions that asked for Turkish ones. Fakir places #9 and Grundig #13; both were founded in Germany and are Turkish-owned today, and the assistants evidently treat them as domestic — which, on an ownership reading, they are. Philips (#12) and Tefal (#15) have no such defence and are simply strong enough in this category to survive a filter that should have excluded them. Read together, this says the *yerli* constraint is recognised but not applied strictly: it removes most foreign names and reshapes the ranking, without being airtight.

04

SECTION 04

Differences Between Models

The five models broadly agree on which Turkish brands matter but disagree sharply on how many names an answer should contain, and one model's habit of naming marketplaces shapes three positions in every leaderboard.

DIFFERENCES BETWEEN MODELS

Agreement on who, not on how many

ONE-SENTENCE TAKEAWAY

The five models broadly agree on which Turkish brands matter but disagree sharply on how many names an answer should contain, and one model's habit of naming marketplaces shapes three positions in every leaderboard.

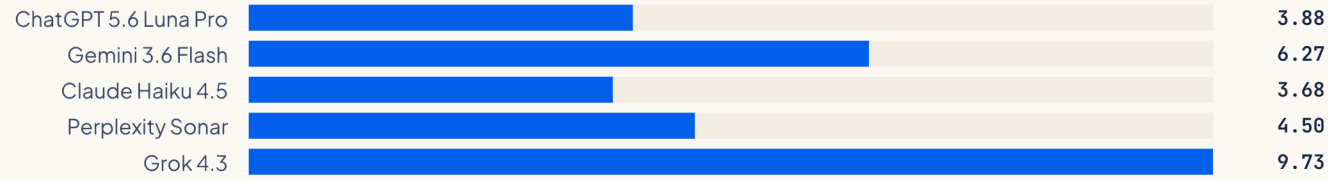
4.1 PER-MODEL BEHAVIOUR SUMMARY

TABLE 3 — GROUNDING & REPERTOIRE BY MODEL

225 responses each (45 questions × 5 repeats)

Model	Web-search	Distinct	Ent./answer	Most-named entity
ChatGPT 5.6 Luna Pro	100.00%	51	3.88	Arzum (65.8%)
Gemini 3.6 Flash	50.22%	56	6.27	Arçelik (62.2%)
Claude Haiku 4.5	100.00%	45	3.68	Karaca (61.3%)
Perplexity Sonar	100.00%	56	4.50	Karaca (75.1%)
Grok 4.3	100.00%	60	9.73	Hepsiburada (96.0%)

DISTINCT ENTITIES NAMED PER ANSWER, BY MODEL



The spread is wide. Grok names 9.73 entities per answer, more than two and a half times Claude Haiku's 3.68, with ChatGPT (3.88), Perplexity (4.50) and Gemini (6.27) in between. Grok also reaches the widest vocabulary at 60 distinct entities against Claude's 45. A brand outside the leading group therefore has a very different chance of being named depending on which assistant the shopper happens to use.

The models also disagree on the leader. Karaca is most-named for Claude (61.3%) and Perplexity (75.1%), Arzum for ChatGPT (65.8%), Arçelik for Gemini (62.2%) — and Hepsiburada for Grok (96.0%). Four of the five put a Turkish manufacturer first; the fifth puts a marketplace first, in almost every answer it gives.

4.2 EACH MODEL'S MOST-NAMED ENTITIES

Rate = share of that model's 225 answers naming the entity.

CHATGPT 5.6 LUNA PRO

top 5 of 225 answers

#	Entity	Mentions	Rate	MRR
1	Arzum	148	65.8%	0.742
2	Karaca	100	44.4%	0.489
3	Philips	71	31.6%	0.560
4	Arçelik	60	26.7%	0.462

(cont.)

#	Entity	Mentions	Rate	MRR
5	Tefal	44	19.6%	0.385

GEMINI 3.6 FLASH top 5 of 225 answers

#	Entity	Mentions	Rate	MRR
1	Arçelik	140	62.2%	0.551
2	Karaca	134	59.6%	0.351
3	Arzum	133	59.1%	0.395
4	Beko	126	56.0%	0.326
5	Korkmaz	96	42.7%	0.507

CLAUDE HAIKU 4.5 top 5 of 225 answers

#	Entity	Mentions	Rate	MRR
1	Karaca	138	61.3%	0.639
2	Arzum	98	43.6%	0.542
3	Arçelik	91	40.4%	0.623
4	Philips	67	29.8%	0.528
5	Tefal	53	23.6%	0.347

PERPLEXITY SONAR top 5 of 225 answers

#	Entity	Mentions	Rate	MRR
1	Karaca	169	75.1%	0.562
2	Arzum	124	55.1%	0.534
3	Arçelik	109	48.4%	0.619
4	Philips	72	32.0%	0.699
5	Tefal	59	26.2%	0.305

GROK 4.3 top 5 of 225 answers

#	Entity	Mentions	Rate	MRR
1	Hepsiburada	216	96.0%	0.144
2	Trendyol	215	95.6%	0.164
3	Karaca	180	80.0%	0.450
4	Arçelik	155	68.9%	0.567
5	Arzum	146	64.9%	0.462

Beneath the different leaders there is real agreement: Karaca, Arzum and Arçelik appear in every model's top five. What changes is the depth behind that core, and which foreign name breaks into it — Philips reaches the top five for Claude, ChatGPT and Perplexity, while Gemini fills those places with Beko and Korkmaz and Grok with two marketplaces.

4.3 THE MARKETPLACE POSITIONS ARE LARGELY ONE MODEL'S HABIT

Grok names Hepsiburada in 96.0% of its answers and Trendyol in 95.6%, and it is the only model to name Akakçe at scale (47.1%). Across the whole run, Hepsiburada and Trendyol each reach a 19.8% mention rate — and roughly four-fifths of that comes from a single model. Akakçe qualifies in all nine cuts on a breadth of 1: it is visible to one model and invisible to four.

This is why the three components are carried in every leaderboard rather than the composite alone. A breadth of 1 or 2 means the entity is visible to part of the market, not to it generally, and the marketplace figures in this edition should be read as substantially a Grok artefact rather than a market-wide pattern.

05

SECTION 05

Web Search or Model Memory?

Only one model varied — Gemini grounded half its answers — and within that model, searching makes the picture markedly more Turkish while answering from memory surfaces an aspirational global shelf.

— WEB SEARCH OR MODEL MEMORY?

One model that only half looked

ONE-SENTENCE TAKEAWAY

Only one model varied — Gemini grounded half its answers — and within that model, searching makes the picture markedly more Turkish while answering from memory surfaces an aspirational global shelf.

Across the run, 1,013 of 1,125 responses (90.0%) returned at least one citation. But the split is not spread across the field: ChatGPT, Claude, Perplexity and Grok grounded 100% of their answers, and **every one of the 112 ungrounded responses came from Gemini**, which grounded 113 of its 225 and worked from memory on the other 112.

That makes a market-wide search-versus-memory comparison impossible — any such contrast would be a comparison between Gemini and everyone else, not between two modes of answering. What it does allow is a clean within-model experiment: the same model, the same 45 questions, roughly half its answers grounded and half not.

The origin balance shifts. When Gemini searches, 63.3% of its mentions go to Turkish-heritage brands. When it answers from memory, that falls to 52.3%. Live Turkish web content pulls the answer towards the domestic market; the model's own parameters hold a more international picture.

TABLE 4 — NAMED MORE OFTEN WHEN GEMINI HAS SEARCHED

113 grounded vs 112 memory-only responses

Entity	Origin	Grounded	Memory	Difference
Karaca	TR	69.9%	49.1%	+20.8 pp
Kumtel	TR	19.5%	2.7%	+16.8 pp
Tefal	Foreign	38.9%	24.1%	+14.8 pp
Arnica	TR	21.2%	8.0%	+13.2 pp
Grundig	Foreign	14.2%	5.4%	+8.8 pp
Homend	TR	23.9%	17.0%	+6.9 pp
Arçelik	TR	65.5%	58.9%	+6.6 pp
Fantom	TR	7.1%	0.9%	+6.2 pp

TABLE 5 — NAMED MORE OFTEN WHEN GEMINI HAS NOT SEARCHED

percentage-point difference in mention rate

Entity	Origin	Grounded	Memory	Difference
De'Longhi	Foreign	9.7%	19.6%	-9.9 pp
Korkmaz	TR	38.0%	47.3%	-9.3 pp
Siemens	Foreign	5.3%	13.4%	-8.1 pp
Sage	Foreign	2.6%	10.7%	-8.1 pp
Breville	Foreign	1.8%	8.9%	-7.2 pp
Roborock	Foreign	0.0%	7.1%	-7.1 pp
Braun	Foreign	9.7%	16.1%	-6.3 pp

(cont.)

Entity	Origin	Grounded	Memory	Difference
KitchenAid	Foreign	0.0%	6.2%	-6.2 pp

Reading. The two lists describe two different markets. Search brings up brands with a heavy current Turkish web footprint: Karaca gains 20.8 points, Kumtel 16.8, Arnica 13.2, and Fantom and Schafer each appear in grounded answers having been almost absent from memory-only ones. Memory brings up a premium international shelf that Turkish search results do not sustain: De'Longhi falls 9.9 points when the model searches, Sage 8.1, Breville 7.2, Dyson 5.5, and Roborock, KitchenAid and Vitamix disappear from grounded answers altogether despite appearing in memory-only ones.

The one Turkish name that runs the other way is Korkmaz, 9.3 points stronger from memory than from search — a long-established brand whose standing in the model's parameters currently outruns its live web presence.

None of this says which mode is more accurate. It says that in this category the answer a shopper receives can depend on whether the assistant looked anything up, and that the difference is systematic rather than random.

06

SECTION 06

Segment Ownership

Small home appliances is not one shelf but three — only 13 qualified entities clear the bar in all three segments, and 17 qualify in exactly one.

SEGMENT OWNERSHIP

Three shelves, not one

ONE-SENTENCE TAKEAWAY

Small home appliances is not one shelf but three — only 13 qualified entities clear the bar in all three segments, and 17 qualify in exactly one.

The 45 questions divide across three product segments, and the leaderboards diverge enough that a single ranking would hide more than it shows.

EACH SEGMENT'S LEADER, BY MENTION RATE (%)



6.1 Kitchen Appliances — 19 questions, 475 responses

TABLE 6 — KITCHEN APPLIANCES

top 10 of 23 qualified

#	Entity	Origin	Mention	AI Score
1	**Karaca**	TR	65.89%	94.58
2	Arzum	TR	53.68%	89.77
3	Arçelik	TR	40.21%	74.16
4	Philips	Foreign	33.47%	73.32
5	Korkmaz	TR	25.05%	67.26
6	Tefal	Foreign	36.21%	66.01
7	Dökümix	TR	12.63%	63.63
8	Braun	Foreign	18.95%	61.09
9	Bosch	Foreign	14.32%	57.34
10	Xiaomi	Foreign	14.74%	55.37

Kitchen Appliances is Karaca's strongest shelf (94.58) with Arzum second. It is also where the small-batch specialists appear: Dökümix at #7, and the imported premium names Vitamix, Ninja and Cosori, none of which qualify anywhere else.

6.2 Coffee & Tea — 13 questions, 325 responses

TABLE 7 — COFFEE & TEA

top 10 of 20 qualified

#	Entity	Origin	Mention	AI Score
1	**Karaca**	TR	75.38%	97.34
2	Arçelik	TR	64.31%	92.74
3	Arzum	TR	71.38%	90.93

(cont.)

#	Entity	Origin	Mention	AI Score
4	Philips	Foreign	39.08%	78.33
5	De'Longhi	Foreign	34.46%	70.20
6	Beko	TR	32.31%	60.49
7	Korkmaz	TR	24.31%	56.64
8	Tefal	Foreign	13.85%	49.38
9	Siemens	Foreign	14.46%	48.01
10	Homend	TR	13.54%	47.71

Coffee & Tea is the most concentrated segment — 20 qualifiers against 27 in Other — and the only one where De'Longhi (#5, 70.20), Jura and Nespresso qualify. Arçelik takes second here on the strength of its Telve line, ahead of Arzum, whose OKKA line drives much of its own coffee visibility.

6.3 Other — 13 questions, 325 responses

TABLE 8 — OTHER (STICK VACUUMS, ÇEYIZ, GENERAL) top 10 of 27 qualified

#	Entity	Origin	Mention	AI Score
1	**Arçelik**	TR	47.69%	97.79
2	Arzum	TR	49.85%	93.80
3	Dyson	Foreign	39.38%	89.93
4	Karaca	TR	50.15%	86.42
5	Philips	Foreign	36.62%	77.14
6	Beko	TR	34.15%	71.94
7	Bosch	Foreign	31.69%	68.25
8	Tefal	Foreign	32.00%	66.51
9	Samsung	Foreign	16.92%	66.50
10	Korkmaz	TR	26.15%	65.33

Other is led by Arçelik (97.79) rather than Karaca, and it is the only segment where the vacuum specialists surface: **Dyson at #3 (89.93, a 39.4% mention rate)**, Dreame at #13, Shark at #21. Samsung qualifies only here. This segment is also the least coherent by construction — it combines a product-defined shelf (stick vacuums) with occasion-defined questions (çeyiz sets and general appliance queries), and all three of its attribute questions come from the vacuum side. Read its internal ordering with that in mind.

6.4 The same three names, in three different orders

The same three Turkish names lead all three segments, but not in the same order, and not with the same company behind them. Thirteen entities qualify in all three segments — Karaca, Arzum, Arçelik, Philips, Tefal, Beko, Korkmaz, Homend, Fakir, Vestel and the three marketplaces. Those are the names an assistant reaches for whatever the appliance. Everything else is shelf-specific.

TABLE 9 — THE THREE SEGMENTS AT A GLANCE

qualified counts are not comparable across segments

Segment	Questions	Responses	Qualified	Leader	Leader mention
Kitchen Appliances	19	475	23	Karaca	65.89%
Coffee & Tea	13	325	20	Karaca	75.38%
Other	13	325	27	Arçelik	47.69%

07

SECTION 07

Question-Type Ownership

Karaca and Arzum trade the lead depending on how the question is phrased, and answers get markedly shorter as questions get more specific.

QUESTION-TYPE OWNERSHIP

Specificity concentrates visibility

ONE-SENTENCE TAKEAWAY

Karaca and Arzum trade the lead depending on how the question is phrased, and answers get markedly shorter as questions get more specific.

The three behavioural types are discovery (open "which brand is good"), attribute (a specific quality — quiet, durable, budget, large-capacity) and use_case (an occasion or need — Turkish coffee, çeyiz, pet hair, a small flat). Each carries 15 questions and 375 responses.

TABLE 10 — MOST-MENTIONED ENTITIES PER QUESTION TYPE top 5 · rate is within that type's 375 responses

#	discovery	attribute	use_case
1	**Karaca (77.1%)**	Arzum (56.0%)	Karaca (61.6%)
2	**Arzum (76.8%)**	Karaca (53.6%)	Arçelik (45.3%)
3	**Arçelik (63.2%)**	Arçelik (39.5%)	Arzum (40.3%)
4	**Philips (49.1%)**	Philips (29.6%)	Philips (29.3%)
5	**Tefal (40.0%)**	Tefal (26.1%)	Beko (28.0%)

Reading. Karaca and Arzum are effectively tied on discovery (77.1% and 76.8%), Arzum leads attribute questions (56.0%), and Karaca leads use-case questions (61.6%) by a clear margin. Arçelik is third in all three. Philips is fourth in all three — the most consistent foreign presence in the category regardless of how a question is framed.

The more interesting movement is in answer length. Discovery questions draw 7.02 entities per response, attribute 4.98 and use_case 4.85. Asked an open question, the assistants produce long shortlists and 28 entities clear the threshold; asked something specific, they narrow to a handful of names, and the qualifying field shrinks to 25 and 24. Specificity concentrates visibility on the brands the models associate most strongly with the category — which is good news for the leaders and hard on everyone else.

08

SECTION 08

Open Market: Turkish vs. Foreign

On origin-neutral questions the shelf turns over: Philips leads, foreign brands take 59.6% of mentions and 17 of 27 qualifying places, and the domestic dominance visible in the headline numbers is a product of the questions rather than the market.

OPEN MARKET — TURKISH VS. FOREIGN

A contested shelf, not a domestic one

ONE-SENTENCE TAKEAWAY

On origin-neutral questions the shelf turns over: Philips leads, foreign brands take 59.6% of mentions and 17 of 27 qualifying places, and the domestic dominance visible in the headline numbers is a product of the questions rather than the market.

This chapter relies only on the 23 open questions with no origin restriction (575 responses). It is the only fair basis for comparing Turkish and foreign brands.

Why the main cut cannot answer this. Twenty-two of the 45 questions explicitly ask for Turkish brands, which guarantees a domestic answer. Across all 45 questions, Turkish-heritage brands take 61.9% of mentions — a figure that measures the question set, not the market. Both numbers are stated together wherever either is used.

8.1 Origin share

ORIGIN SHARE OF MENTIONS — OPEN MARKET (%)



Foreign brands take 2,097 open-cut mentions to 1,421 Turkish, a 59.6/40.4 split. Breadth points the same way: 43 distinct foreign entities surface in the open cut against 25 Turkish, and 17 of the 27 qualifiers are foreign. On the questions where nothing steers the answer, the assistants reach for an international field first.

8.2 Open-market top 15 (any origin)

TABLE 11 — OPEN MARKET (ANY ORIGIN)

23 questions · 575 responses · top 15 of 27 qualified

#	Entity	Origin	Mention	AI Score
1	**Philips**	Foreign	64.52%	94.38
2	Karaca	TR	54.61%	81.42
3	Arzum	TR	45.57%	78.14
4	Tefal	Foreign	50.09%	74.71
5	Dyson	Foreign	20.17%	69.07
6	Arçelik	TR	32.35%	66.99
7	Bosch	Foreign	29.22%	61.74
8	De'Longhi	Foreign	22.96%	60.87
9	Dökümix	TR	5.74%	57.96
10	Samsung	Foreign	9.39%	57.39
11	Braun	Foreign	20.17%	57.21
12	Xiaomi	Foreign	18.26%	53.00

(cont.)

#	Entity	Origin	Mention	AI Score
13	Korkmaz	TR	12.17%	51.28
14	Vitamix	Foreign	5.04%	49.63
15	Homend	TR	10.78%	49.53

8.3 The pattern: a contested shelf, not a domestic one

Philips leads the open cut at 94.38 on a 64.5% mention rate — higher than any Turkish brand achieves here, and a reversal of its fourth place in the main cut. Karaca (81.42) and Arzum (78.14) hold second and third, and Tefal (74.71) is fourth on the second-highest mention rate in the cut. Dyson is fifth on position again (MRR 0.6795, the highest of any qualified entity in any cut).

Beneath the top five the foreign field is genuinely broad rather than niche: Bosch (#7), De'Longhi (#8), Samsung (#10), Braun (#11), Xiaomi (#12), Vitamix (#14), Dreame (#18), Ninja (#19), Siemens (#20), Cosori (#22) and Nespresso (#25) all qualify. These are not luxury outliers — they span every price band and every segment in the category.

The contrast with the local cut is the finding. Ask for a Turkish brand and the assistants produce a confident, well-populated domestic list. Ask the same question without the constraint and the domestic names hold three of the top six places but lose the aggregate. Both facts are true at once, and a brand's visibility strategy looks different depending on which of the two questions its customers are actually typing.

ONE CAVEAT ON THE ORIGIN RULE

Because this edition publishes origin by heritage rather than ownership, Fakir and Grundig count as foreign here. Reclassifying both as Turkish on an ownership basis would move the open-cut Turkish share from 40.4% to 44.4%, which changes the size of the gap but not its direction.

09

SECTION 09

The Discovery Ecosystem

This market is learned from Turkish editorial and forum content, not from brand websites — editorial supplies 39.7% of citations against 19.0% brand-owned, and one technology forum is cited by nearly half of all responses.

— THE DISCOVERY ECOSYSTEM — WHERE AI LEARNS ABOUT BRANDS

The assistants are reading the reviewers

ONE-SENTENCE TAKEAWAY

This market is learned from Turkish editorial and forum content, not from brand websites — editorial supplies 39.7% of citations against 19.0% brand-owned, and one technology forum is cited by nearly half of all responses.

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 answers cited each. Across the run the models produced 5,985 response-level citations across 343 distinct domains.

9.1 THE MOST-CITED DOMAINS

TABLE 12 — THE 12 MOST-CITED DOMAINS

of 343 distinct cited domains

#	Domain	Citing	Share	Source type
1	**donanimhaber.com**	533	47.38%	Editorial / Media
2	karaca.com	327	29.07%	Retailer
3	trendyol.com	285	25.33%	Marketplace
4	hepsiburada.com	251	22.31%	Marketplace
5	nefisyemektarifleri.com	232	20.62%	Editorial / Media
6	eniyilerden.com	189	16.80%	Editorial / Media
7	eniyisinde.com.tr	162	14.40%	Editorial / Media
8	akakce.com	154	13.69%	Marketplace
9	jebinde.com	141	12.53%	Retailer
10	arzum.com.tr	135	12.00%	brand-owned
11	arcelik.com.tr	134	11.91%	brand-owned
12	onedio.com	132	11.73%	Editorial / Media

The most-cited single domain is **donanimhaber.com**, a Turkish technology news site and forum, cited by 533 responses — 47.4% of the entire run, and more than any brand or marketplace domain. Behind it come two brand-adjacent retail sites (karaca.com, 29.1%; trendyol.com, 25.3%), then a recipe site (nefisyemektarifleri.com, 20.6%) and a cluster of Turkish review and listicle publishers (eniyilerden.com, eniyisinde.com.tr, iyioneri.tr, tuketicidergisi.com.tr). Brand-owned domains appear, but lower: arzum.com.tr and arcelik.com.tr at 12.0% and 11.9%.

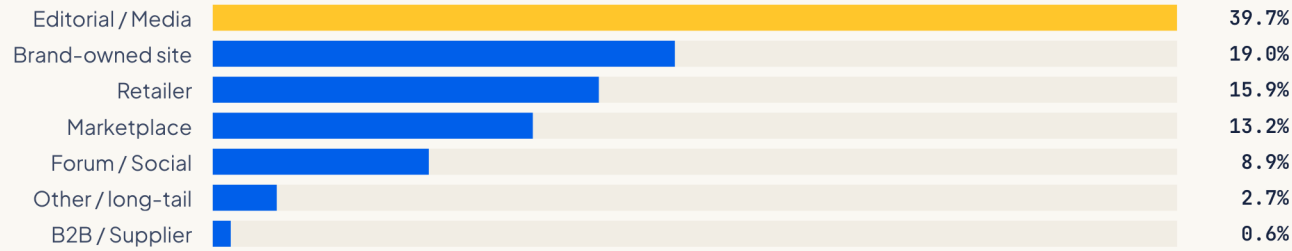
9.2 SOURCE-TYPE MIX

TABLE 13 — CITATION SHARE BY SOURCE TYPE

heuristic classification · see Chapter 11

Source type	Domains	Citations	Share
Editorial / Media	98	2,376	39.7%
Brand-owned site	62	1,138	19.0%
Retailer	39	952	15.9%
Marketplace	11	788	13.2%
Forum / Social	15	532	8.9%
Other / long-tail	113	163	2.7%
B2B / Supplier	5	36	0.6%

CITATION SHARE BY SOURCE TYPE (%)



9.3 The editorial finding

Editorial and media domains supply 39.7% of citations from 98 distinct domains — more than double the brand-owned share of 19.0%, and more than retailers (15.9%) or marketplaces (13.2%). Forums and social sites add a further 8.9%, led by youtube.com, eksisozluk.com and kahveler.net. The unclassified long tail is 2.7%, consisting of 109 domains cited by fewer than three responses each and left unclassified by design.

Read as a description rather than advice: in this category, what third parties write about an appliance appears more closely associated with AI visibility than what the manufacturer publishes about it. Reviews, comparison listicles, forum threads and recipe content are where these systems are reading.

A NOTE ON EXTRACTION

Two providers required separate handling. 922 of 6,635 total citation entries from Gemini are redirect-wrapped URIs carrying no readable domain, so the domain was taken from the citation title instead; Grok returns the opposite problem — bare numerals as titles — so its domain was taken from the URL. One extractor per provider was necessary; a single shared extractor would have silently lost one provider's sources entirely.

10

SECTION 10

What the Patterns Suggest

The most visible names share breadth across all five models, a strong presence in Turkish editorial and review content, and a clear association with one product shelf — and mention volume, early placement and model coverage are three separate things.

WHAT THE PATTERNS SUGGEST

Five structural patterns

ONE-SENTENCE TAKEAWAY

The most visible names share breadth across all five models, a strong presence in Turkish editorial and review content, and a clear association with one product shelf — and mention volume, early placement and model coverage are three separate things.

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 a baseline condition.** Twenty-two of the 25 qualified entities appear in all five models. The three that do not are all marketplaces, and all three sit in the bottom four of the table. Broad model coverage is what the visible names have in common.

2. **Mention and position are independent signals.** Dyson (11.4% of responses, MRR 0.6527) and Dökümix (5.3%, MRR 0.6330) rank fifth and ninth almost entirely on being named early. Hepsiburada and Trendyol are named four times as often as Dökümix and rank far below it, because they arrive at the end of a list. "Named often" and "named first" are different achievements and the score treats them as such.

3. **The segment is the unit, not the category.** Only 13 of the qualified entities clear the bar in all three segments. For most brands, visibility is shelf-specific: De'Longhi is a Coffee & Tea name, Dyson and Dreame are vacuum names, Dökümix and Vitamix are kitchen names. A brand reading its own position should read it inside its segment first.

4. **Third-party content is where this category is learned.** With editorial at 39.7% of citations against brand-owned at 19.0%, and a single technology forum cited by nearly half of all responses, the correlation here runs towards review coverage, comparison content and forum discussion rather than owned media. This is a description of the current citation landscape, not a claim about cause.

5. **The origin question has two different answers.** The domestic names hold the market when the question invites them and lose the aggregate when it does not. A brand that is highly visible on *yerli* questions is not necessarily visible on the equivalent unrestricted question, and the gap between those two positions is measurable per brand in the underlying data.

How a brand can locate itself in the data

Using the dataset (available on request), a brand can read four things in order: whether it appears in all five models; how often it is mentioned within its segment; how early it is named when mentioned; and whether its local-cut and open-cut positions diverge. These readings show where a visibility gap sits, without saying anything about whether its product is good.

11

SECTION 11

Limitations and Notes

This is a single, point-in-time snapshot of one Turkish vertical; read the figures within the limitations below.

LIMITATIONS AND NOTES

How to read the figures

ONE-SENTENCE TAKEAWAY

This is a single, point-in-time snapshot of one Turkish vertical; read the figures within the limitations below.

- **Single collection window.** All 1,125 responses were collected in one run on 13 August 2026. This is a snapshot, not a trend, and model behaviour changes without notice. This is the first (baseline) edition.
- **The score measures AI visibility only.** Not quality, not popularity, not sales, not satisfaction. A brand can rank highly because it is frequently written about online.
- **Some qualifiers are one model's opinion.** Akakçe qualifies in all nine cuts on a single model (106 responses, breadth 1). Trendyol, Hepsiburada and MediaMarkt qualify at breadth 2 in several cuts. Twenty-three qualifier rows across the nine cuts have a breadth below 3. Where breadth is 1 or 2 the entity is visible to part of the market, not to it generally — read breadth alongside every score.
- **One model dominates the marketplace rankings.** Grok names 9.73 entities per response against Claude's 3.68, and its most-named entity is Hepsiburada at 96.0% of its responses. The marketplace figures are substantially a Grok artefact.
- **The search-versus-memory comparison is one model's.** Gemini is web-grounded on 50.2% of its responses where every other model is at 100%, so Chapter 5 is a within-Gemini contrast and cannot be generalised to the other four.
- **A low-mention, high-position entity places surprisingly high.** Dökümix ranks #9 in the main cut on 60 responses, a mention rate of 0.0533 barely above the threshold, because its MRR of 0.6330 is near the top of the cut. Dyson at #5 is a milder version of the same effect. Read both cautiously.
- **Small score gaps are not meaningful.** The local cut's top three sit within 1.71 points. Only larger gaps should be treated as robust.
- **Segments are not comparable on distinct-entity counts.** The Other segment combines stick vacuums (product-defined) with çeyiz and general questions (occasion-defined and cross-category), and all three of its attribute questions come from the vacuum side. No segment falls below the ten-question floor, but this one is heterogeneous by construction.
- **Sub-brands under-count structurally** wherever the models name the parent instead. OKKA and Telve were merged into Arzum and Arçelik after proving to be perfect subsets; other sub-brands may be absorbed the same way without being visible as such.
- **Origin classification involves judgment.** Origin is published as brand heritage, not current ownership. Fakir and Grundig are therefore foreign here despite Turkish ownership, and Profilo is local despite German ownership. §8.3 states what the alternative rule would do to the headline figure.
- **Never compare a scaled score across cuts.** A score of 100 in one cut is a different absolute number from 100 in another, and the scale moves whenever the leader moves. Cross-edition comparison will use mention rate, MRR and breadth — the raw layer, stored unscaled for exactly this purpose.
- **Marketplaces are not brands.** Trendyol, Hepsiburada, Akakçe and Amazon are stripped from the brand-only view in §3.2. Karaca, Korkmaz and MediaMarkt are retailers or brand-retailers, not marketplaces, and are retained everywhere.

- **Models are probabilistic.** The same question can produce different answers; five repeats reduce but do not eliminate this.
- **Precision over recall.** Entity matching was cautious to avoid false positives, so brands mentioned indirectly may be undercounted. There is no sentiment layer; the study is strictly quantitative.
- **Source-type mix is heuristic.** Domain types were auto-classified and should be treated as indicative. Source coverage also varies by model, and Gemini's real source domains are read from its citation titles rather than its redirect URLs.
- **Output-length limits and model versions** differ and date quickly; findings are specific to the versions in §2.1 as of August 2026.

— FREQUENTLY ASKED QUESTIONS

Frequently asked questions

Which small home appliance brands does AI recommend most in Turkey?

Karaca leads the main leaderboard at 92.95, named in 64.1% of all 1,125 responses, followed closely by Arzum (90.24) and Arçelik (85.89). Philips is the highest-placed foreign brand at 74.67, and Dyson is fifth.

Do AI assistants recommend Turkish or foreign appliance brands?

It depends entirely on how the question is asked. On questions that explicitly ask for Turkish brands, 15 of the 19 qualifying names are Turkish. On origin-neutral questions, foreign brands take 59.6% of mentions and 17 of 27 qualifying places, and Philips leads outright. The all-questions figure of 61.9% Turkish reflects the question mix rather than the market.

How is the AI Visibility Score calculated?

It combines three parts: how often an entity is mentioned (45%), how early it appears in the answer (30%), and how many of the five models mention it (25%). Each part is scaled so the leading qualified entity in each cut reaches 100, and only entities named in at least 5% of a cut's responses are ranked.

Do the five AI assistants agree with each other?

They agree on the core — Karaca, Arzum and Arçelik appear in every model's top five — but not on the leader or the depth. Karaca is most-named for Claude and Perplexity, Arzum for ChatGPT, Arçelik for Gemini, and the marketplace Hepsiburada for Grok. Grok names 9.73 entities per answer against Claude's 3.68.

Where do AI assistants get their information about appliance brands?

Mostly from Turkish editorial and review sites rather than manufacturers. Editorial and media domains supply 39.7% of citations against 19.0% for brands' own sites, and the single most-cited domain is the technology forum donanimhaber.com, cited by 47.4% of all responses.

Does it matter whether the assistant searched the web?

For one of the five models it does. Gemini grounded only half its answers, and when it searched, 63.3% of its mentions went to Turkish brands against 52.3% when it answered from memory. Karaca gained 20.8 points with search; De'Longhi, Sage and Breville were all named more often without it.

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 build quality, durability, price, accuracy or sentiment, so a score is a measure of visibility, not an endorsement.

12

SECTION 12

Appendix & Data

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

APPENDIX & DATA

12.1 All questions (45), their types, scopes and segments

All 45 questions used in the study, asked exactly as written. Scope: 22 local (ask for *yerli* brands), 23 open (origin-neutral). Distribution: 15 discovery, 15 attribute, 15 use_case; 19 Kitchen Appliances, 13 Coffee & Tea, 13 Other.

#	Type	Scope	Segment	Question (Turkish)
1	discovery	open	Other	küçük ev aletlerinde hangi marka iyi
2	discovery	open	Other	çeyiz için küçük ev aletleri hangi marka alınmalı
3	discovery	open	Coffee & Tea	en iyi kahve makinesi hangisi
4	discovery	open	Coffee & Tea	kahve makinesinde en çok hangi marka tutuluyor
5	discovery	open	Kitchen Appliances	blender seti hangi marka alınır
6	discovery	open	Kitchen Appliances	en çok tercih edilen airfryer hangisi
7	discovery	open	Other	piyasadaki en iyi dikey süpürge hangisi
8	discovery	open	Kitchen Appliances	tost makinesi önerisi hangi marka iyi
9	discovery	local	Other	yerli küçük ev aletleri önerisi
10	discovery	local	Other	çeyiz için iyi bir yerli küçük ev aleti markası var mı
11	discovery	local	Coffee & Tea	Türk malı kahve makinesi önerir misin
12	discovery	local	Kitchen Appliances	blender setinde yerli ne alınır
13	discovery	local	Kitchen Appliances	iyi bir yerli airfryer var mı
14	discovery	local	Kitchen Appliances	yerli tost makinesi alacağım ne önerirsiniz
15	discovery	local	Coffee & Tea	çay makinesinde Türk markası hangisi iyi
16	attribute	open	Coffee & Tea	uygun fiyatlı kaliteli kahve makinesi hangisi
17	attribute	open	Coffee & Tea	köpüğü iyi yapan Türk kahvesi makinesi önerisi
18	attribute	open	Kitchen Appliances	en kolay temizlenen airfryer hangisi
19	attribute	open	Kitchen Appliances	güçlü ve sessiz blender seti hangi marka
20	attribute	open	Kitchen Appliances	en sağlam blender hangi marka
21	attribute	open	Other	emiş gücü yüksek kablosuz dikey süpürge önerisi
22	attribute	open	Other	hafif ama iyi çeken dikey süpürge hangisi
23	attribute	open	Kitchen Appliances	döküm plakalı sağlam tost makinesi önerir misiniz
24	attribute	local	Coffee & Tea	fiyat performans yerli kahve makinesi önerisi
25	attribute	local	Coffee & Tea	köpüğü iyi yapan Türk malı kahve makinesi var mı
26	attribute	local	Kitchen Appliances	kolay temizlenen yerli airfryer önerir misin

(cont.)

#	Type	Scope	Segment	Question (Turkish)
27	attribute	local	Kitchen Appliances	güçlü motorlu yerli blender seti arıyorum hangisi iyi
28	attribute	local	Kitchen Appliances	uzun ömürlü yerli tost makinesi ne alayım
29	attribute	local	Other	iyi çeken yerli dikey süpürge var mı
30	attribute	local	Coffee & Tea	cam veya çelik hazneli yerli çay makinesi önerisi
31	use_case	open	Coffee & Tea	ev için tam otomatik kahve makinesi tavsiyesi
32	use_case	open	Kitchen Appliances	dört kişilik aile için airfryer hangi marka iyi
33	use_case	open	Kitchen Appliances	çorba smoothie buz kırma için blender seti önerisi
34	use_case	open	Kitchen Appliances	kalabalık aile için tost makinesi ne almalıyım
35	use_case	open	Other	çeyize kahve çay tost blender tek markadan alsam hangisi mantıklı
36	use_case	open	Coffee & Tea	anneme kahve makinesi alacağım hangi marka iyi
37	use_case	open	Other	günlük temizlik için şarjlı dikey süpürge ne almalıyım
38	use_case	local	Other	yeni ev kuruyorum yerli küçük ev aletlerinde hangi markaya bakayım
39	use_case	local	Other	çeyiz için Türk malı küçük ev aleti seti önerir misiniz
40	use_case	local	Coffee & Tea	anneme hediye yerli Türk kahvesi makinesi ne alayım
41	use_case	local	Kitchen Appliances	iki kişilik ev için yerli airfryer önerisi
42	use_case	local	Kitchen Appliances	sert şeyleri de çekebilen yerli blender seti var mı
43	use_case	local	Kitchen Appliances	her gün tost ve ızgara için sağlam yerli makine önerisi
44	use_case	local	Other	evcil hayvan tüyü için iyi çeken yerli dikey süpürge var mı
45	use_case	local	Coffee & Tea	çay ve kahveyi aynı cihazda yapabileceğim yerli bir marka var mı

12.2 All tracked entities (76)

All 76 tracked entities, main cut (1,125 responses). Rate is within the main cut. Entities at or above 5% are marked **Ranked**. Sorted by mention count.

Entity	Origin	Mentions	Rate	Models	Status
Karaca	TR	721	64.09%	5	Ranked
Arzum	TR	649	57.69%	5	Ranked
Arçelik	TR	555	49.33%	5	Ranked
Philips	Foreign	405	36.00%	5	Ranked
Tefal	Foreign	321	28.53%	5	Ranked
Beko	TR	317	28.18%	5	Ranked

(cont.)

Entity	Origin	Mentions	Rate	Models	Status
Korkmaz	TR	283	25.16%	5	Ranked
Hepsiburada	TR	223	19.82%	4	Ranked
Trendyol	TR	223	19.82%	3	Ranked
Fakir	Foreign	214	19.02%	5	Ranked
Bosch	Foreign	186	16.53%	5	Ranked
Homend	TR	181	16.09%	5	Ranked
De'Longhi	Foreign	133	11.82%	5	Ranked
Dyson	Foreign	128	11.38%	5	Ranked
Braun	Foreign	122	10.84%	5	Ranked
Xiaomi	Foreign	118	10.49%	5	Ranked
Arnica	TR	109	9.69%	5	Ranked
Akakçe	TR	106	9.42%	1	Ranked
Vestel	TR	102	9.07%	5	Ranked
Grundig	Foreign	86	7.64%	5	Ranked
Kumtel	TR	69	6.13%	5	Ranked
Sinbo	TR	67	5.96%	5	Ranked
Dreame	Foreign	60	5.33%	5	Ranked
Dökümix	TR	60	5.33%	5	Ranked
Siemens	Foreign	60	5.33%	5	Ranked
Samsung	Foreign	55	4.89%	5	Tracked
Ninja	Foreign	54	4.80%	5	Tracked
Schafer	TR	52	4.62%	5	Tracked
MediaMarkt	Foreign	51	4.53%	2	Tracked
Nespresso	Foreign	41	3.64%	4	Tracked
Senur	TR	32	2.84%	5	Tracked
Cosori	Foreign	31	2.76%	4	Tracked
Shark	Foreign	30	2.67%	3	Tracked
Vitamix	Foreign	29	2.58%	4	Tracked
Breville	Foreign	25	2.22%	3	Tracked

(cont.)

Entity	Origin	Mentions	Rate	Models	Status
Emsan	TR	24	2.13%	5	Tracked
Amazon	Foreign	22	1.96%	4	Tracked
Sage	Foreign	20	1.78%	3	Tracked
Vatan	TR	18	1.60%	1	Tracked
Altus	TR	17	1.51%	3	Tracked
Jura	Foreign	17	1.51%	4	Tracked
KitchenAid	Foreign	17	1.51%	4	Tracked
Zilan	TR	17	1.51%	2	Tracked
Moccamaster	Foreign	16	1.42%	4	Tracked
Zwilling	Foreign	16	1.42%	2	Tracked
Melitta	Foreign	15	1.33%	2	Tracked
Remta	TR	15	1.33%	2	Tracked
Tchibo	Foreign	15	1.33%	3	Tracked
Bamix	Foreign	14	1.24%	4	Tracked
Bissell	Foreign	13	1.16%	3	Tracked
Fantom	TR	13	1.16%	2	Tracked
Miele	Foreign	11	0.98%	4	Tracked
GoldMaster	TR	10	0.89%	4	Tracked
Blendtec	Foreign	9	0.80%	4	Tracked
Electrolux	Foreign	9	0.80%	5	Tracked
Instant	Foreign	8	0.71%	2	Tracked
Luxell	TR	8	0.71%	4	Tracked
Okkalı	TR	8	0.71%	4	Tracked
Roborock	Foreign	8	0.71%	1	Tracked
Awox	TR	7	0.62%	2	Tracked
Gaggia	Foreign	6	0.53%	3	Tracked
Krups	Foreign	6	0.53%	2	Tracked
Profilo	TR	6	0.53%	2	Tracked
Teknosa	TR	6	0.53%	1	Tracked

(cont.)

Entity	Origin	Mentions	Rate	Models	Status
Brita	Foreign	5	0.44%	1	Tracked
Kiwi	TR	5	0.44%	2	Tracked
Smeg	Foreign	5	0.44%	3	Tracked
Kenwood	Foreign	4	0.36%	3	Tracked
Remington	Foreign	4	0.36%	2	Tracked
Rowenta	Foreign	4	0.36%	1	Tracked
WMF	Foreign	4	0.36%	3	Tracked
Aryıldız	TR	3	0.27%	2	Tracked
Cuisinart	Foreign	3	0.27%	1	Tracked
Feris	TR	3	0.27%	2	Tracked
Numatic	Foreign	3	0.27%	2	Tracked
Stilevs	TR	3	0.27%	2	Tracked

12.3 Data availability

The study tracked 76 entities; of these, 25 cleared the 5% threshold in the main cut, with 19 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 and grounding flags, all three qualified leaderboards, the segment and question-type leaderboards, full unfiltered metrics for every entity in every cut, the per-model breakdown, the complete 343-domain source list, the alias audit, the list of filtered non-brand candidates, and the methodology and data dictionary.

All figures reflect a single point-in-time run (13 August 2026) and are subject to the limitations in Chapter 11. This is the first (baseline) measurement; the study will be repeated quarterly.

WANT THE DEEPER CUT?

The full domain-level source breakdown, per-model leaderboards for each cut and segment, and the question-by-question entity visibility are not reproduced in the public report. To request them, or to be notified when the next quarterly edition is published, book a call with the Herm.io team. There is no fee and nothing to buy; it is part of how we share what we learn.

12.4 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.

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.

DISCLOSURE & NEUTRALITY

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.

Brands seeking to understand their position in the data can schedule a consultation for an impartial assessment of the findings. The consultation is advisory and free of charge.

REPORT PERIOD

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