




Medicare AI Visibility Index™

Key Findings and Executive Implications

7 AI surfaces

121 Medicare brands

14,490 Medicare prompts

59,055 brand mentions

Q3 2026

Reach. Recommendation. Evidence.

Published by Brainpan.AI | Kevin Walsh, Founder | Methodology, definitions and measurement boundaries: see Measurement Framework and Executive Actions.

CONTENTS

Findings and measurement boundaries.

EXECUTIVE SUMMARY	3
FINDING 01 / DENOMINATORS	4
FINDING 02 / COMMERCIAL SHARE	5
FINDING 03 / PER-MENTION PERFORMANCE	6
FINDING 04 / CITATION AUTHORITY	7
FINDING 05 / SURFACE DIVERGENCE	8
FINDING 06 / NEXT-ACTION SIGNALS	9
FINDING 07 / INTERMEDIARY PLACEMENT	10
FINDING 08 / DIRECT COMPARISONS	11
FINDING 09 / REPEATABILITY	12
FINDING 10 / SEGMENT SUMMARY	13
MEASUREMENT FRAMEWORK	14
EXECUTIVE ACTIONS	15

EXECUTIVE SUMMARY

Being seen, being recommended and being cited are three different advantages.

The benchmark tracks 121 named Medicare brands across seven tested AI answer surfaces. 115 of the 121 have any recorded visibility; the other six (Premera Blue Cross Medicare, Johns Hopkins Advantage MD, Tufts Health Plan Medicare / Point32Health, Fallon Health Medicare, Dean Health Plan Medicare and Trinity Health Plan of New York) have none in this panel. Every finding below is computed against the 59,055 tracked-brand mention baseline.

121

tracked brands

115

visible brands

7

tested surfaces

59,055

commercial mentions

1

Concentration. UnitedHealthcare, Humana and Aetna receive 50.5% of commercial mentions.

2

What leads. Blue Cross Blue Shield Medicare ranks 4th (3,816 mentions, 6.46% share), ahead of Cigna Healthcare Medicare and Wellcare, once state-affiliate name variants (Illinois, Texas, Georgia, and unbranded "BCBS" mentions) are attributed to the umbrella brand.

3

Evidence. Of the 62,484 citations captured, 52.7% are provided by just twenty-five domains. Because citation strength feeds directly into our new Quality Visibility Index (QVI™) — a measure built to separate volume from value — Medicare.org leads QVI™ among brands with at least 100 mentions.

4

Surface variation. Google AI Overviews records a 39.5% brand-citation rate. AI Overview rendered for 95.89% of its allocated prompts (1,985 of 2,070); API and consumer configurations otherwise produce materially different citation and routing patterns.

5

Scope. This Key Insights report covers the following KPIs: reach, per-mention quality, citation authority, surface behavior, next-action signals, intermediary placement, direct comparisons, repeatability and segment mix.

Executive question: is the constraint reach, recommendation or evidence, and on which surfaces? Each is measured separately below.

FINDING 01 / DENOMINATORS

Source rows, the designed sample and capture coverage.

Unit	Count	Interpretation
Source extraction rows	229,372	Recorded rows across two extraction methods; not a deduplicated census.
Scored commercial rows	59,055	Named-brand mentions after alias/subsidiary rollups and Plan Finder consolidation.
Prompt-surface observations (designed sample)	14,490	150 stratified queries × 7 surfaces, 3 runs each. Every one of the 14,490 was successfully attempted.
Populated responses	14,405	14,490 minus 85 Google AI Overview non-render outcomes: a real observed outcome (no AI Overview box rendered), not missing data.
Google AI Overview appearance rate	95.89%	1,985 of AI Overview's own 2,070 allocated prompts rendered a box; the other 85 are the non-render outcomes above.
Scored-brand coverage, of 14,405 populated responses	11,638 / 80.79%	Against the canonical 115-brand list from sealed per-surface data.
Scored-brand coverage, of 14,490 total observations	11,638 / 80.32%	Same numerator against the full designed sample, with the 85 AI Overview non-renders folded in as zero-visibility.

Medicare Vocabulary is 60.04% of the universe (137,708 mentions); Commercial Medicare brands are 25.75% (59,055 mentions). These are the reference figures campaign-wide.

Source: Total Visibility Universe Breakdown.

FINDING 02 / COMMERCIAL SHARE

Three national carriers lead the tracked commercial set.

Brand	Mention rows	Commercial share	Total entity share
UnitedHealthcare Medicare / AARP Medicare Plans	10,838	18.35%	4.73%
Humana Medicare	9,865	16.7%	4.30%
Aetna Medicare / CVS Health	9,104	15.42%	3.97%
Blue Cross Blue Shield Medicare	3,816	6.46%	1.66%
Cigna Healthcare Medicare	3,551	6.01%	1.55%
Wellcare	3,535	5.99%	1.54%
Anthem Blue Cross / Elevance Health Medicare	1,757	2.98%	0.77%
Devoted Health	1,549	2.62%	0.68%
Kaiser Permanente Medicare Advantage	1,456	2.47%	0.63%
Medicare.org	967	1.64%	0.42%

The top three account for 50.5% of commercial mentions and 13.0% of all extracted entity rows. Blue Cross Blue Shield Medicare ranks 4th (6.46%), ahead of Cigna Healthcare Medicare and Wellcare.

121 brands are tracked; 115 have any recorded visibility (this leaderboard). Commercial share uses 59,055 scored rows. Total share uses 229,372 extraction rows. Full 115-brand leaderboard: Full Report, Section 02.

FINDING 03 / PER-MENTION PERFORMANCE

QVI ranks row-based per-mention quality.

QVI' rank	Brand	QVI'	Mentions	Own-domain cited'
1	Medicare.org	48.05	967	88.9%
2	SelectQuote Senior	44.5	324	27.5%
3	Geisinger Gold	40.61	128	31.9%
4	UPMC Health Plan Medicare	40.55	935	18.3%
5	eHealth Medicare	40.43	540	0.0%
6	SCAN Health Plan	36.96	856	25.7%
7	Healthfirst	36.7	438	28.0%
8	Humana Medicare	36.62	9,865	11.9%
9	UnitedHealthcare / AARP Medicare	34.91	10,838	8.8%
10	GoHealth Medicare	34.88	355	19.2%

Medicare.org ranks first among the 36 brands with at least 100 mentions; Humana ranks eighth and UnitedHealthcare ninth. Reach (Finding 02) and per-mention quality should be tracked separately.

Quality Visibility Index (QVI™) = 8 × mean Position Score + 7 × mean Recommendation Score + 25 × own-domain citation rate, scaled 0–100. QVI and own-domain-cited are computed on the subset of rows that received full per-mention scoring; mention counts reflect the complete dataset. This is not a uniformly scored measure of preference or plan quality. The 100-mention cutoff is a reporting rule.

FINDING 04 / CITATION AUTHORITY

Twenty-five domains supply more than half of captured citations.

Domain	Citation rows	Share of 62,484	Classification
medicare.gov	4,131	6.61%	Government / official
medicare.org	4,104	6.57%	Commercial marketplace
nerdwallet.com	3,685	5.90%	Publisher / review
usnews.com	2,377	3.80%	Publisher / review
forbes.com	1,810	2.90%	Publisher / review
aetna.com	1,700	2.72%	Carrier-owned
humana.com	1,668	2.67%	Carrier-owned
medicareadvantage.com	1,224	1.96%	Other Medicare / health
uhc.com	1,088	1.74%	Carrier-owned
healthline.com	964	1.54%	Publisher / review

Medicare.org holds 1.64% of commercial mentions but 29.04% of brand-owned citation rows: an authority ratio of 17.7x.

Citation-row counts and domain shares are citation-corpus figures, independent of brand-mention share. Prioritize inspection of the actual cited pages and their relevance.

FINDING 05 / SURFACE DIVERGENCE

Seven configurations produce different visibility patterns.

Reach and placement

Surface	Mention rows	Distinct brands / capture	Top-3 rate
Microsoft Copilot	14,702	4.99	38.5%
Bing Copilot	9,131	3.52	52.5%
Claude API	8,171	3.86	63.3%
Google AI Overviews	6,822	2.97	68.8%
Perplexity API	6,378	2.95	64.6%
ChatGPT / OpenAI API	5,982	2.82	78.9%
Gemini API	3,184	1.52	68.8%

Citation and scoring

Surface	Brand cited	Official source	Visibility Points
Microsoft Copilot	3.9%	1.4%	64,107
Bing Copilot	19.2%	29.2%	49,419
Claude API	0.1%	83.6%	25,268
Google AI Overviews	39.5%	26.3%	31,877
Perplexity API	27.3%	67.4%	25,373
ChatGPT / OpenAI API	0.0%	17.9%	21,503
Gemini API	0.0%	3.7%	10,667

Google AI Overviews shows the highest reported brand-citation rate (39.5%). Near-zero rates on several API configurations describe this setup; they do not make citation strategy irrelevant to those platforms.

Surface-level table. Distinct brands / capture = distinct brand/Response ID pairs divided by populated responses. Raw Visibility Points are not a normalized cross-surface ranking.

FINDING 06 / NEXT-ACTION SIGNALS

Answers may direct the next step toward intermediaries.

Surface	Points to Medicare.gov	Suggests agent / broker	Points to SHIP
Microsoft Copilot	5.6%	3.1%	0.4%
Bing Copilot	34.9%	28.8%	10.4%
Google AI Overviews	6.5%	3.3%	2.6%
Perplexity API	98.0%	74.0%	26.0%
Claude API	100.0%	35.3%	44.9%
Gemini API	98.9%	61.1%	5.7%
ChatGPT / OpenAI API	100.0%	48.6%	15.0%

Across the source panel, 36.5% of answers suggest an agent or broker and 15.1% point to SHIP. The report measures the instruction in the answer, not whether anyone follows it.

Answer-content signal table. Textual references and directions to Medicare.gov differ from government-domain matches in captured citations; they are separate metrics.

FINDING 07 / INTERMEDIARY PLACEMENT

Broker-intent queries explain part of the placement advantage.

Metric	Reported result
Responses containing both a carrier and intermediary	1,688
Overall intermediary displacement rate	48.2%
Displacement rate excluding broker-intent queries	17.8%
Displacement in broker-present subgroup	84.8%
Displacement in publisher / marketplace subgroup	16.5%

Observed examples. eHealth displaces carriers in 390 responses, SelectQuote in 211 and Medicare.org in 149.

Interpretation. Earlier placement may reflect query intent, content relevance, familiarity or other factors. The report does not prove that intermediary content structure or schema causes the result.

Response-level displacement table. Displacement means an intermediary appears before the best-placed carrier. Eligible denominators differ across these rates. The 84.8% result is a subgroup rate and does not apply to all answers.

FINDING 08 / DIRECT COMPARISONS

Some smaller brands win selected score-based matchups.

Among 208 direct-comparison responses, the source reports 186 clean two-carrier pairs. The higher Visibility Points score defines the winner.

Brand	Pair appearances	Scored wins	Scored win rate
Alignment Health	16	15	93.8%
Devoted Health	21	18	85.7%
Kaiser Permanente Medicare Advantage	20	17	85.0%
Blue Cross Blue Shield Medicare	17	12	70.6%
Wellcare	17	9	52.9%
Humana Medicare	114	55	48.2%
Aetna Medicare / CVS Health	55	24	43.6%
Cigna Healthcare Medicare	18	7	38.9%
UnitedHealthcare / AARP Medicare	78	28	35.9%
SCAN Health Plan	16	1	6.2%

Response-level scored-matchup table. These outcomes do not necessarily represent explicit AI choices or beneficiary preferences. Opponent mix and sample size differ; Alignment's 93.8% is 15 wins from 16 appearances.

FINDING 09 / REPEATABILITY

One answer is not a stable visibility baseline.

Source measure	Reported result
Observed brand/query/surface combinations	7,363
Not present in all three reported runs	52.0%
Rank swing of at least two positions when consistently present	40.6%

What to do. Repeat a controlled query panel, preserve raw answers and track aggregate results alongside query-level changes.

What is verified. All 14,490 execution keys reconcile: 690 query/geography combinations × three run slots × seven surfaces. The files do not provide confidence intervals, a minimum detectable change or evidence that every capture is an independent generation.

Repeated observations support a more useful descriptive baseline. Precision depends on the sampling design and dependence between observations, not the record count alone.

Run-based stability table.

FINDING 10 / SEGMENT SUMMARY

Brand type changes the mix of visibility.

Brand type	Brands	Mentions	Share
Carrier	105	56,647	95.92%
Broker / Marketplace	9	1,441	2.44%
Publisher / Marketplace	1	967	1.64%

Carriers hold 95.92% of tracked-brand mentions across 105 of the 115 visible brands. Intermediaries (brokers/marketplaces and the one publisher/marketplace brand, Medicare.org) hold the remaining 4.08% across ten brands.

Computed directly from the 115-brand leaderboard (59,055 mentions): brand type (Carrier / Broker or Marketplace / Publisher or Marketplace) is a static per-brand attribute. Per-mention quality metrics (QVI, recommendation, citation, top-3 rate) by segment are not shown here; see Finding 03.

MEASUREMENT FRAMEWORK

The classification and scoring boundaries are explicit.

Rows and positions

59,055 scored rows underlie every commercial-share and position figure in this report. QVI and position scores remain row-based; repeated mentions of one brand can occupy several positions.

Recommendation coverage

All 23,715 rows from the four API configurations have Recommendation Score zero. Nonzero scores occur only on consumer surfaces. QVI can therefore reflect surface mix and incomplete recommendation signals.

Entity classification

Name variants are rolled up under one canonical brand (for example, state Blue Cross Blue Shield affiliates under the umbrella entity); brand references are separated from generic Medicare vocabulary; and consolidated tool and phone references (Medicare Plan Finder, 1-800-MEDICARE) are each treated as a single entity rather than counted as separate variants.

Terminology

Use scored commercial rows (59,055), Medicare Vocabulary rows (137,708), and the eleven other named-entity/residual categories in the 13-category Total Visibility Universe breakdown. Treat Medicare.gov Plan Finder as a government-service reference. Use QVI™ (Quality Visibility Index) for the per-mention quality score.

The Full Report contains the complete 115-brand leaderboard and full measurement definitions.

EXECUTIVE ACTIONS

Five decisions for Medicare marketing leaders.

1

Report the dimensions separately

Track reach, recommendation score and citation pathways. Avoid collapsing distinct gaps into one headline.

2

Prioritize relevant comparisons

Use matched surface configurations and query intent. Inspect the responses behind the score differences.

3

Test the proposed intervention

Treat educational content, comparison pages, entity consistency and third-party evidence as hypotheses to evaluate. The benchmark does not establish which change causes a gain.

4

Validate the answer body and scoring before remeasurement

Separate captured prompts and interface text from answers. Apply a common recommendation rubric across surfaces, then repeat the verified 690-combination panel.

5

Extend the benchmark to the remaining dimensions

Prominence, Product/Geography, Query Ownership, Co-Mention Network and Cross-Surface Agreement can be added once response-level, product-, geography- and query-tagged brand data is available, widening the benchmark beyond reach, quality and citation authority.

Brainpan.AI can turn a visibility finding into a focused diagnostic: establish the gap, inspect the evidence, prioritize an intervention and measure the result.

Request an AI visibility diagnostic: brainpan.ai/contact/