

The State of AI Discovery in Health Tech 2026

How AI recommends consumer health-tech providers — and what marketing leaders should measure beyond visibility.

AI visibility changes with what the consumer is actually looking for.

The same brand can look very different by health category and by the exact need in the question. A brand can also be visible without being recommended, while the source and reputation context around an answer adds another layer marketers need to monitor.

01	Visibility is highly intent-specific.	Category and exact consumer need can materially change the recommendation pattern.
02	Being mentioned is not the same as being recommended.	Presence in an answer can overstate whether a brand is actually presented as suitable for the request.
03	Brand-owned content can appear alongside major publishers.	On one weight-management question, a Found comparison article appeared in the same source landscape as Forbes and U.S. News.
04	Reputation answers can mix evidence, entities and changing facts.	Reviews, company claims, corporate history and conflicting facts can all shape what consumers see.

Six health categories. 27 health-tech brands.

A purposive set of established consumer-health platforms, category anchors, growth-stage challengers and specialists. The study was not designed to represent every U.S. health-tech company.

82

unbranded discovery prompts

5

AI products: ChatGPT, Google AI Overviews, Google AI Mode, Gemini and Perplexity

14

study dates, September 8–21, 2026

GLP-1 & weight management	Hims · Hers · Ro · Weight Watchers Meds+ · Found · Henry Meds
Mental health & psychiatry	BetterHelp · Talkspace · Brightside Health · Talkiatry · Cerebral
Men’s sexual & hormonal health	Hims · Ro · BlueChew · Hone Health · Maximus
Women’s digital health	Hers · Wisp · Midi Health · Allara · Evvy
Hair loss & digital dermatology	Hims · Keeps · Curology · Honeydew · Happy Head
Telehealth & primary care	Teladoc · One Medical · Sesame · PlushCare · GoodRx Care

How to read the universe

These 31 primary brand-category slots represent 27 unique entities. Eligibility was set before formal collection for every exact question, so an additional cross-category provider could be scored when it could directly serve that consumer need.

The same brand can look very different across health categories.

Ro and Hims provide the clearest comparison because both were eligible for the same answer sets in three health categories.


GLP-1 & weight management



Men's sexual & hormonal health



Hair loss & digital dermatology




WHAT THIS MEANS FOR MARKETERS

Use a brand-wide visibility number as a summary, not the whole story. Break results out by health category so teams can see where visibility is actually strong or weak.

Supporting counts: weight management — Ro 574/980, Hims 123/980; men's health — Ro 362/560, Hims 441/560; hair/derm — Ro 251/490, Hims 354/490. Matched same-answer denominators within each category. Descriptive context comparison only.

Within weight management, the Ro–Hims gap was much wider on insurance questions.

Two insurance-related questions showed a very different recommendation pattern from a separate question about starting care quickly.



WHAT THIS MEANS FOR MARKETERS


Track the need behind the question. Insurance, affordability, speed, privacy, continuity and specialist expertise can reveal differences hidden by a category average.

70 answers per exact question. Prior authorization: Ro 61/70, Hims 1/70. Wegovy/Zepbound insurance help: Ro 45/70, Hims 0/70. Fast start: Ro 51/70, Hims 40/70. Distinct consumer requests, not a controlled experiment.


A brand can be highly visible in an answer without being recommended nearly as often.

One privacy-focused mental-health question shows why presence and recommendation need separate metrics.

BetterHelp



Talkspace




WHAT THIS MEANS FOR MARKETERS

Use separate metrics for mentions and recommendations. Otherwise, high presence can overstate how often the brand is actually recommended for the exact request.

Same 70 privacy-first answers. BetterHelp: 56 mentions; 9 confirmed recommendations + 3 unresolved. Talkspace: 55 mentions; 21 confirmed recommendations + 1 unresolved. Ranges are classification bounds, not confidence intervals.

Brand-owned content can appear alongside major publishers in AI answers.

On one exact weight-management question, a Found comparison article appeared in the same source landscape as Forbes and U.S. News.



WHAT THIS MEANS FOR MARKETERS

Use highly cited pages as a research list. Study what those pages cover, how they frame the topic, and whether your brand is present. Recurring owned pages can highlight content worth maintaining or expanding; recurring competitor and third-party pages can surface content gaps and earned-media opportunities.

Exact Found article: <https://joinfound.com/blog/best-online-weight-loss-programs>

What AI says about a provider can be shaped by very different kinds of information.

The qualitative review looked at what happens when consumers ask directly about a named provider's reputation, credibility, strengths and drawbacks.

01 Reviews, marketing claims and policy terms can become one brand story.

What we saw: One answer cited a 4.6/5 Trustpilot rating, said "same day" access was advertised, and noted that one monthly plan had a three-month minimum before cancellation.

Why it matters: A customer-review score, a company marketing claim and a subscription term all shaped the same reputation assessment.

02 Parent-company history can become part of a service-level answer.

What we saw: One answer referenced a 2023 FTC enforcement action involving the parent company and company-wide reputation ratings while evaluating a specific virtual-care service. The answer also noted those ratings were not a controlled assessment of the service's clinicians.

Why it matters: A broader company issue can become part of the story AI tells about an individual service, even when the evidence was not specific to that service.

03 Even basic company facts can change from one answer to the next.

What we saw: Across three consecutive days, the same AI product named one founder pair, a different pair the next day, then returned to the first pair.

Why it matters: Consumers can encounter conflicting versions of a basic company fact from one answer to the next.

Monitor the specific claims and facts that shape the brand story in AI.

Prioritize anything that is outdated, tied to the wrong company or service, or inconsistent across answers — and trace it back to the underlying sources when possible.

WATCH FOR

Outdated facts

Company history, pricing, policies, access claims and other details that may change over time.

WATCH FOR

Wrong scope

Claims about a parent company or another service that become part of the answer about the brand consumers are actually evaluating.

WATCH FOR

Conflicting versions

Important facts that change across answers, products or dates.

TRACE BACK TO

The underlying source

Understand what page, review, news event or company claim may be feeding the narrative consumers see.

Qualitative scope

210 distinct fully read recorded answers covered all 155 branded reputation prompt/platform cells at least once. The study describes what the recorded AI answers said; it does not adopt the underlying provider claims as verified facts or report reputation rates or rankings.

Four practices to keep the executive view simple without losing the full picture.

01

Break out visibility by category and intent.

Don't rely on a single brand-wide number. Show which health categories and consumer needs are driving the result.

02

Track mentions and recommendations separately.

A brand can appear in an answer without being presented as a suitable option. Treat presence and recommendation as different signals.

03

Turn citation data into a content and distribution roadmap.

Use recurring owned pages to spot content worth expanding, and competitor or third-party pages to identify content gaps and earned-media opportunities.

04

Monitor the evidence that influences reputation.

Track the claims, facts and news events that appear in AI answers about your brand, especially when they are outdated, tied to the wrong entity or contradicted elsewhere.

A fixed U.S. prompt set, observed across five AI surfaces for 14 days.

The downloadable report has been selectively simplified, while the web version preserves expanded methodology, representative prompts, full category and intent scorecards, and additional technical detail.

5,739

nonblank discovery answers

30,165

valid brand observations under pre-set eligibility rules

1

known blank discovery answer in the formal window

Prompt set: 120 exact prompts — 82 unbranded discovery, 6 head-to-head comparison, 31 branded reputation and 1 market-perception prompt.

Platforms: ChatGPT, Google AI Overviews, Google AI Mode, Gemini and Perplexity.

Formal window: September 8–21, 2026. September 7 was QA-only and is excluded from study results.

Recommendation: a provider had to be presented as an affirmative suitable option or shortlist for the exact consumer request. A substantive mention alone did not count.

Unresolved classifications: where a small number of answers remained unresolved, public reporting preserves a lower-to-upper range rather than silently recoding them.

Limitations


The study observes a designed prompt set over a defined two-week period. It does not measure the entire AI-search market, every patient journey, clinical quality, patient outcomes, market share, traffic attribution or causal effects of marketing interventions.

Citation measurement

Visible citations and recommendations are separate signals. Of 758 discovery answers with no recorded cited-page URL, 666 still contained at least one confirmed recommendation for an eligible brand. This does not mean those answers were unsupported; it means citation presence should not be used as a stand-in for recommendation visibility.

The scorecards use two treatments depending on the category.

Some categories support an overall comparison. Others are split by intent so different consumer needs are not forced into one ranking. The full six-category scorecards can be seen in the web version of the report.



Why two treatments?

Weight management supports an overall category view. Mental health is split into therapy and psychiatry/medication because the eligible provider sets and consumer needs differ. The same principle is used for other mixed categories.

Representative publication treatment only. Full category scorecards, exact counts, eligibility denominators and unresolved classifications remain available in the web version of the report.

Explore the methodology and supporting data

The downloadable report has been selectively simplified, while the web version preserves expanded methodology, representative prompts, full category and intent scorecards, and additional technical detail.

Discovery · 82

Head-to-head · 6

Reputation · 31

Market perception · 1

W01 · Broad category

Which online weight-loss providers in the U.S. would you recommend for someone who wants clinician-supervised GLP-1 treatment?

M17 · Therapy

Which online therapy providers are best for someone who wants a licensed therapist and predictable weekly appointments rather than messaging-only support?

WH21 · PCOS / hormonal

Which online providers are best for PCOS if weight, insulin resistance, and irregular periods are all concerns?

PC17 · Insurance / continuity

Which virtual primary-care providers are best for someone who wants to use insurance and see the same clinician over time?

The State of AI Discovery in Health Tech 2026

Research for marketing leaders navigating AI discovery.

Get in touch.

Book a discovery call at automattik.ai.