Branded questions are the only ones ChatGPT cites insurers on.

An insurer's own domain wins the citation in one question type of four.

Two of the other three are closed by structural authority that content cannot displace. One is still open, and worth the investment.

Two of these are worth the budget

Question type	Who gets the link	Why	Verdict
Definitional and compliance	Government domains	Regulatory authority holds it	Hold for the mention
Drug and procedure coverage	Official and clinical references	Specialization, and narrowing further	Hold for the mention
Workplace and HR benefits	Job and HR platforms	No moat. Won on content quality	Invest. A real push
Branded and named queries	The brand's own domain	Nobody else can hold it	Defend. Best return

Figure 1. Citation ownership by question type across a tracked Insurance prompt panel, with the budget verdict for each.
All four question types in this panel are listed here. Scope is one tracked panel over these weeks.
Source: BrightEdge AI Hyper Cube and AI Catalyst. A tracked Insurance prompt panel, ChatGPT and Google AI Overviews, United States, one-month and 12-week comparisons to 20 September 2026.

Nobody is going to outrank the IRS on what a deductible is. That is not a content problem. No budget fixes it. The better question is where the ownable ground sits, and the data points at two places.

The two lanes that will not move

GETS THE LINK
Top cited domains

healthcare.gov
dol.gov
irs.gov
ssa.gov

GETS NAMED
Mentioned top three

- **healthcare.gov**
- **dol.gov**
- **Insurer domains**

Insurer domains hold a place in the mentioned top three across the same weeks government domains take the citations.

Figure 2. Top cited domains against the mentioned top three for the same tracked Insurance prompts over the same weeks. A citation is the linked source under an answer. A mention is the brand named in the answer text. Each list is the top of its own ranking for these weeks.
Source: BrightEdge AI Hyper Cube and AI Catalyst. A tracked Insurance prompt panel, ChatGPT and Google AI Overviews, United States, one-month and 12-week comparisons to 20 September 2026.

Government domains hold 4 of the top 8 cited, up from a mixed pool of job sites, publishers and legal aggregators 12 weeks ago. healthcare.gov alone went from 19.2% to 40.4% of ChatGPT citations in a month, and legalclarity.org fell from 8.8% to 3.4% over the same window.

ny.gov is the one riser confirmed on both engines: 1.6% to 8.0% on ChatGPT, 23.6% to 29.3% on Google AI Overviews. It carries the most weight here because it holds on both.

We expected insurer domains to disappear from the mentioned top three altogether. They did not, staying there alongside healthcare.gov and dol.gov across the same weeks government domains take the citations. Keep tightly matched content here for the mention, and report it as mention share, separate from citation growth.

Drug and procedure coverage is closing further

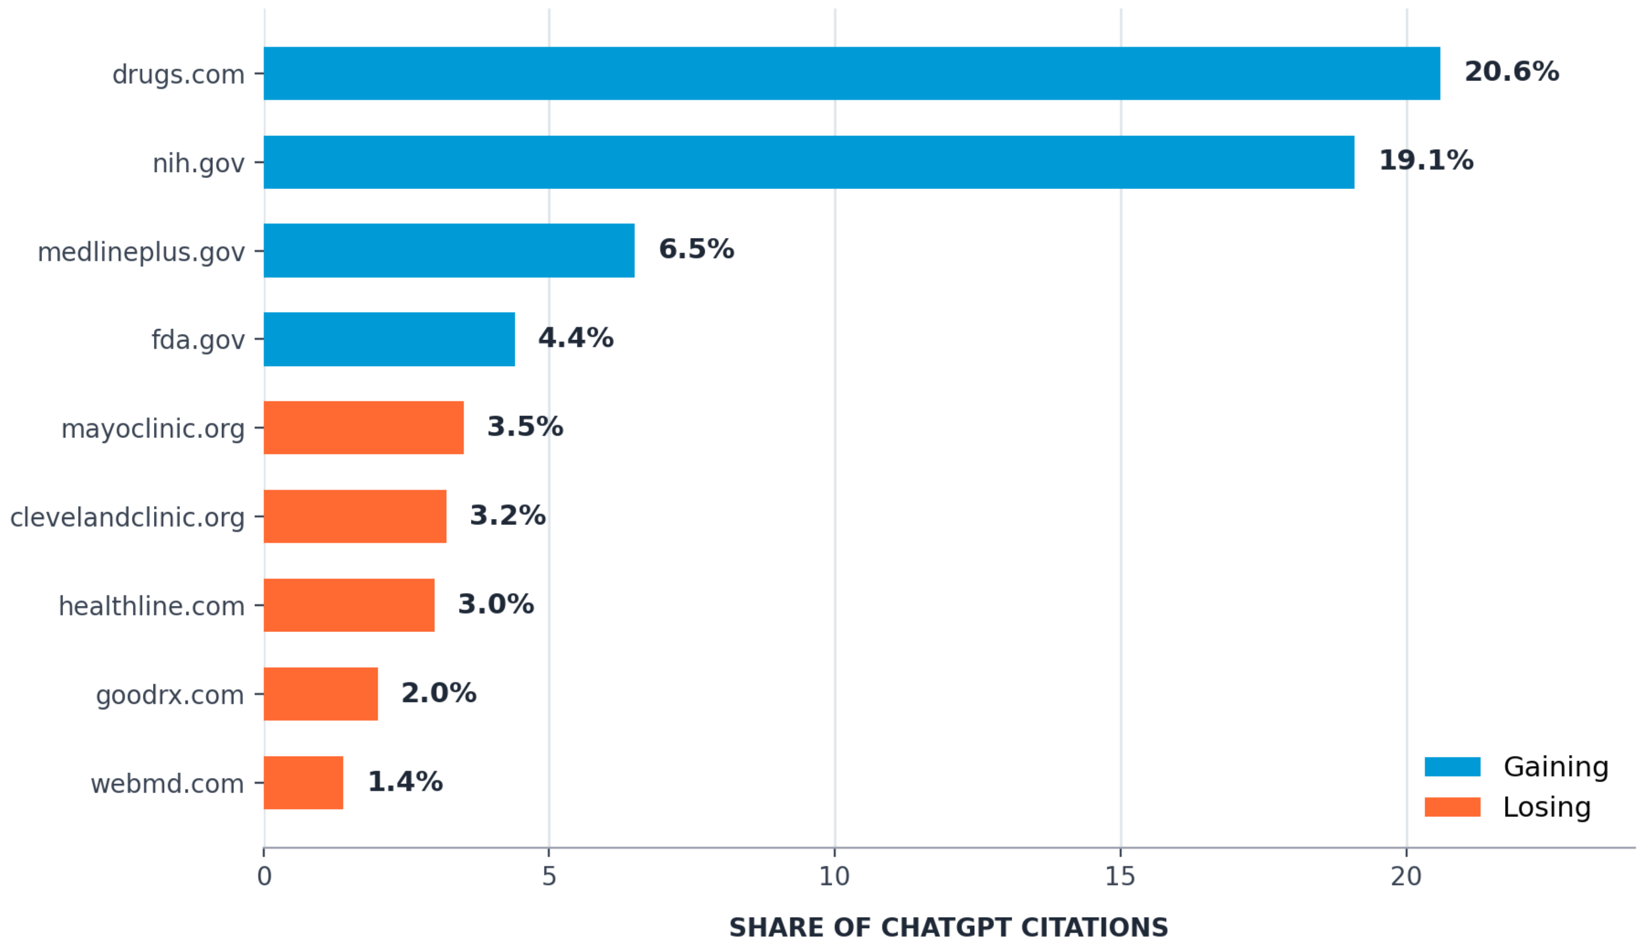

Figure 3. Share of ChatGPT citations on tracked drug and procedure coverage prompts, one-month comparison, grouped by direction of change. These are the top movers in this panel, ranked within each direction.
Source: BrightEdge AI Hyper Cube. A tracked Insurance prompt panel, ChatGPT, United States, one-month comparison to 20 September 2026.

Every domain gaining citation share is official or single-purpose. Every domain losing share is general-purpose consumer health content. No insurer appears among the top movers in either direction.

That group is narrowing even for the non-insurers already inside it.

Where the budget actually earns

- ✓ **Audit branded coverage content first**

  Someone is typing does [brand] cover [drug] into ChatGPT right now, and when a query names the company, the company's own domain is what gets cited. Check whether that page exists on your site in a format a machine can lift.

- ✓ **Decide whether HR content is a business you want to be in**

  No regulatory moat sits behind it. You would be competing with Indeed, Glassdoor, SHRM and Gallup on content quality. Calculators, benchmarking data and retention research qualify. Product marketing with an HR label on it does not.

- ✓ **Measure mentions separately from citations**

  Insurer domains hold a place in the mentioned top three while government domains take 4 of the top 8 cited. Keep tightly matched content in those two, and report it as mention share.

- ✓ **Stop budgeting generic content as AI-citation growth**

  Definitional content is up against regulatory authority. Coverage content is up against official health references sitting at 20.6% and 19.1%. An insurer closes neither gap. Both still earn organic traffic and on-site conversion, so this is a reallocation decision.

- ✓ **Google is still the channel to plan around**

  The ny.gov movement shows up on Google AI Overviews as well as ChatGPT. The canonical, specific content that earns organic visibility is the same content that gets a brand named in an AI answer.

RUN THIS ON YOUR OWN CATEGORY THIS WEEK

Three checks that show which lane you are in

1. **Split your prompts by question type**

   In AI Hyper Cube, pull the cited-domain list for your tracked prompts and group it by what the question asks. The owner changes by group.

2. **Pull mentions and citations as two reports**

   In AI Catalyst, read mention share beside citation share for the same weeks. Where they diverge is where your reporting is hiding a result.

3. **Check the branded queries separately**

   Ask the engines a question that names your brand and see which domain is cited. That is the lane the data says you can hold.

This covers one vertical on a tracked Insurance prompt panel. Whether the same split by question type holds in other regulated categories is untested and worth checking separately.

Request a demo