FINTEL
CONNECT

2026 AI Search Guide for Financial Marketers

Appear. Compete. **Get chosen**.

QUESTIONS THIS GUIDE ANSWERS

- ↗ What decides whether a page is read at all?
- ↗ What separates the brand named first from the ones named after it?
- ↗ How do brands position themselves to get chosen in AI answers?

2026 EDITION What 2,400 AI answers reveal about how consumers discover financial products, and what determines which brands get chosen

Table of Contents

EXECUTIVE SUMMARY

Market Shift 3

AI is compressing discovery and changing how influence shows up

ASK

When AI Searches the Live Web 7

Consumer intent determines what kind of answer the AI gives

RETRIEVE

How AI Answers Financial Questions 9

The model decides whether to use current web sources or answer from memory

APPEAR

Getting Into the AI Consideration Set 11

Retrieved evidence determines which brands make it into the consideration set

COMPETE

What Determines Who Gets Chosen 13

Once multiple brands appear, they compete for position and attention

GET CHOSEN

What Financial Marketers Should Do 17

What financial marketers should do differently as a result

© 2026 Fintel Connect | All Rights Reserved

ASK • RETRIEVE • APPEAR • COMPETE • GET CHOSEN

Market Shift

AI is compressing the financial discovery journey

For years, choosing a financial product meant moving across **a fragmented research journey**: searching, opening comparison pages, visiting financial institution websites, checking rates, reading reviews, and returning to search again.

AI is beginning to collapse much of that journey into a **conversation**.



A consumer can now ask one question and receive an answer that synthesizes product information, affiliate comparisons, brand content, and recommendations into a much smaller consideration set.

That changes the economics of discovery. Brands may have fewer opportunities to appear during the research journey. Publishers and third-party content may influence a decision without receiving the final click. And a consumer who discovers a product through an AI platform may ultimately navigate directly to the financial institution's website, leaving much of the intelligence that shaped the decision invisible to traditional attribution.

The question for financial marketers is therefore changing.

It is no longer only, *"Where did the traffic come from?"* Increasingly, it is, **"What information shaped the decision before the visit happened, and why did one brand make the consideration set while another did not?"**

To understand that, Fintel Connect analyzed 2,400 AI-generated answers from 800 non-branded consumer prompts across ChatGPT, Claude, and Gemini, spanning consumer credit cards, business credit cards, high-yield savings accounts, and personal loans.

In **half** of those answers (1,093 of 2,400), no sources were retrieved at all. Among them, a brand was specifically named in a **quarter** of the responses (266 of 1,093). In answers that *did* retrieve sources, **a brand was named 90% of the time**.

The sources an AI model retrieves shape which brands enter the answer. Once a brand appears, its position within the answer determines how much attention and detail it receives.

In other words, **winning AI discovery requires a brand to appear, compete, and get chosen.**

© 2026 Fintel Connect | All Rights Reserved

Methodology and Key Terms

What was collected

Fintel Connect submitted 800 unique consumer prompts to three AI models (once each) producing 2,400 answers.

Every prompt was written in natural consumer language and non-branded; the prompts themselves never named a bank, lender, or product, so every brand appearance in the results was produced by the model.

Platforms and models

One model per platform.

ChatGPT	Claude	Gemini
gpt-5.4	claude-haiku-4-5	gemini-3.6-flash
OpenAI Responses API with the web search tool	Anthropic web search tool	Google Search grounding

Categories and market

Four product categories in the U.S. market at 200 prompts each:

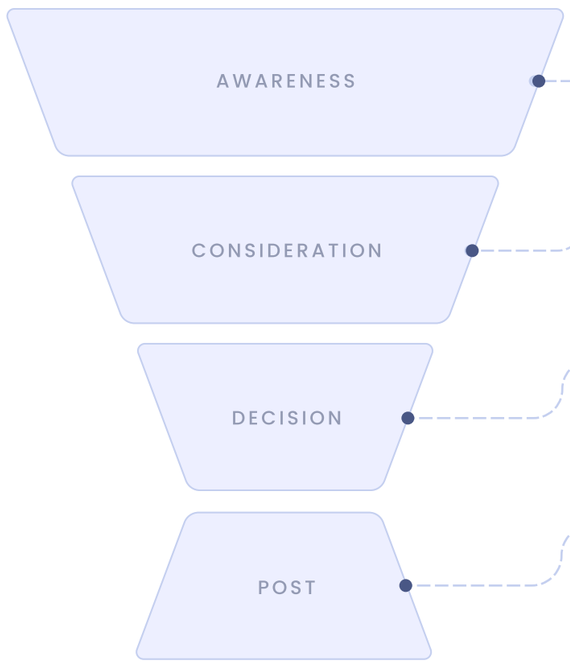
- Consumer credit cards
- Business credit cards
- High-yield savings accounts
- Personal loans

Collection period

Response data was collected between **August 5** and **August 24, 2026**.

Prompt design

Questions were balanced across eleven intent types (informational, comparative, transactional, advisory, eligibility, risk, alternative, product-feature, segment-specific, high-commercial-intent, and post-decision), four consumer journey stages, three length bands (short, medium, and long), and eleven customer segments.



Awareness

Prompts **asking what the product is and how it works**, such as definitions, mechanics, and general explanations. (e.g., How do business credit cards work?)

Consideration

Prompts about **weighing options and criteria** like whether a product type is worth it or who qualifies. (e.g., What's the best credit card for a student?)

Decision

Prompts where a user is **choosing a product** by asking about rankings or current offers. (e.g., Best HYSA promotional intro rates deal today?)

Post

Prompts about **validating choices** after a product has already been purchased (e.g., I already applied for a personal loan that focuses on fixed vs. variable rate, so is that a solid choice or did I miss something better?)

© 2026 Fintel Connect | All Rights Reserved

How sources were captured	Citations were taken from the structured citation records each platform returns with its answer. This produced **7,809 total citations**.
Brand mentions	Only **banks**, **lenders**, **card issuers,** and **fintechs** are counted as brands in the figures. Affiliates, regulators, card networks, and credit bureaus are tracked separately.
What was measured	Each answer was parsed into three layers. **Citations** Every source the platform returned, with its order, domain, source category, and page format **Brand mentions** Which financial brands appear in the visible answer text, the order they appear in, which brand the answer recommends first, and how many brands compete in the same answer **Answer composition** The answer format, whether a ranked order was given, and how many words of description each named brand receives
How the figures are calculated	Every percentage in this report is a share of a stated base: answers, citations, or brand mentions. Figures were computed programmatically from the run, citation, and mention tables. Each is stored with its numerator and denominator so it can be re-derived.

Key Terms

Term	Definition	Why It Matters
Generative engine optimization (GEO)	Structuring content so AI platforms like ChatGPT, Gemini, and Claude can read it, cite it, and surface it directly inside an answer.	Financial executives often use GEO and AEO interchangeably. What matters more than the label is which stage of the funnel you're optimizing for.
Answer engine optimization (AEO)	Optimizing content to appear in AI-generated answers, particularly relating to answer formats such as featured snippets and Google's AI Overview.	AEO applies to AI answers inside Google search results and other AI-driven snippets, while GEO applies to answers from the AI platforms themselves. Both aim to place a brand in the answer a consumer reads first.
Large language model (LLM)	The underlying AI model, such as GPT, Gemini, or Claude, that reads sources, decides what to retrieve, and generates the answer a consumer sees.	Each model behaves differently. One reads brand-owned pages most often, another leans on affiliate citations, and a third pulls from community and video sources. So one approach applied to a single model will misread visibility across all of them.
AI search	Asking an AI platform a question and getting a direct answer, instead of searching and clicking through links.	This is the plain-language term for what this report studies. GEO and AEO describe the strategy while AI search describes what the customer is actually doing.

Key Terms (cont'd)

Term	Definition	Why It Matters
Model memory	What an AI model already learned during training, without reading any new source from the web.	Content published today can't enter answers built from model memory alone. That shifts only over years, not quarters.
Prompt	The question or request someone types into an AI search tool.	How a prompt is worded is one of the strongest predictors of whether a source gets retrieved at all.
Source	A page the AI could potentially read while building an answer, such as a comparison site, a brand's own page, or an affiliate article.	This is the raw material AI answers draw from. Reading a source is what puts brand names in front of consumers at all.
Citation	The specific record of a source the AI actually retrieved and referenced in an answer.	Citations are what this report counts and measures. A citation sits behind the majority of AI-recommended brand answers, with its heaviest weight at the decision stage.
Retrieve	The step where an AI model searches the live web for current sources, rather than relying only on what it learned during training.	High-intent, decision-stage questions are far more likely to trigger retrieval, which is where new content can actually move the needle.
Brand-owned page	A bank's, lender's, card issuer's, or fintech's own website.	Brands whose own page gets read are recommended first far more often, and the effect is even stronger for well-established brands with broad content.
Brand mention	A brand's name written out in the AI's answer text, visible to the consumer.	Being mentioned is real progress, but it's not the same as being recommended first. Position still decides how much space a brand gets.
Recommended first	Being the brand an AI names ahead of all others when the answer recommends more than one.	The first-named brand gets meaningfully more description than the second or third. This is the highest-value outcome in AI product discovery.
Affiliate	A third-party publisher, comparison site, review site, personal finance blog, or content creator that earns a commission or referral fee for directing customers to a brand's financial product.	Affiliate content supports the majority of recommendation answers across every platform studied, peaking at the decision stage.
Decision stage	The point in the customer journey (*awareness, consideration, and decision)* where someone is actively comparing and choosing between their options.	Citation rates are highest here, up to 90.3%, making it the most valuable stage for content investment.

ASK ● RETRIEVE ● APPEAR ● COMPETE ● GET CHOSEN

When AI Searches the Live Web

Consumer intent determines what kind of answer the AI gives

INTRODUCTION

AI discovery starts before a source is ever retrieved. The way a consumer frames a question influences whether the model treats it as something that can be answered from existing knowledge or something that requires current information from the web. That makes intent the first gate in the AI discovery journey.

What Factors Shape the Sources AIs Include in their Answers?

Prompt intent

The buying stage a consumer is in influences whether the AI retrieves a source at all. Here, a "source" is a live web page the model reads while building its answer.

Citation rates climb sharply by intent across all product categories:

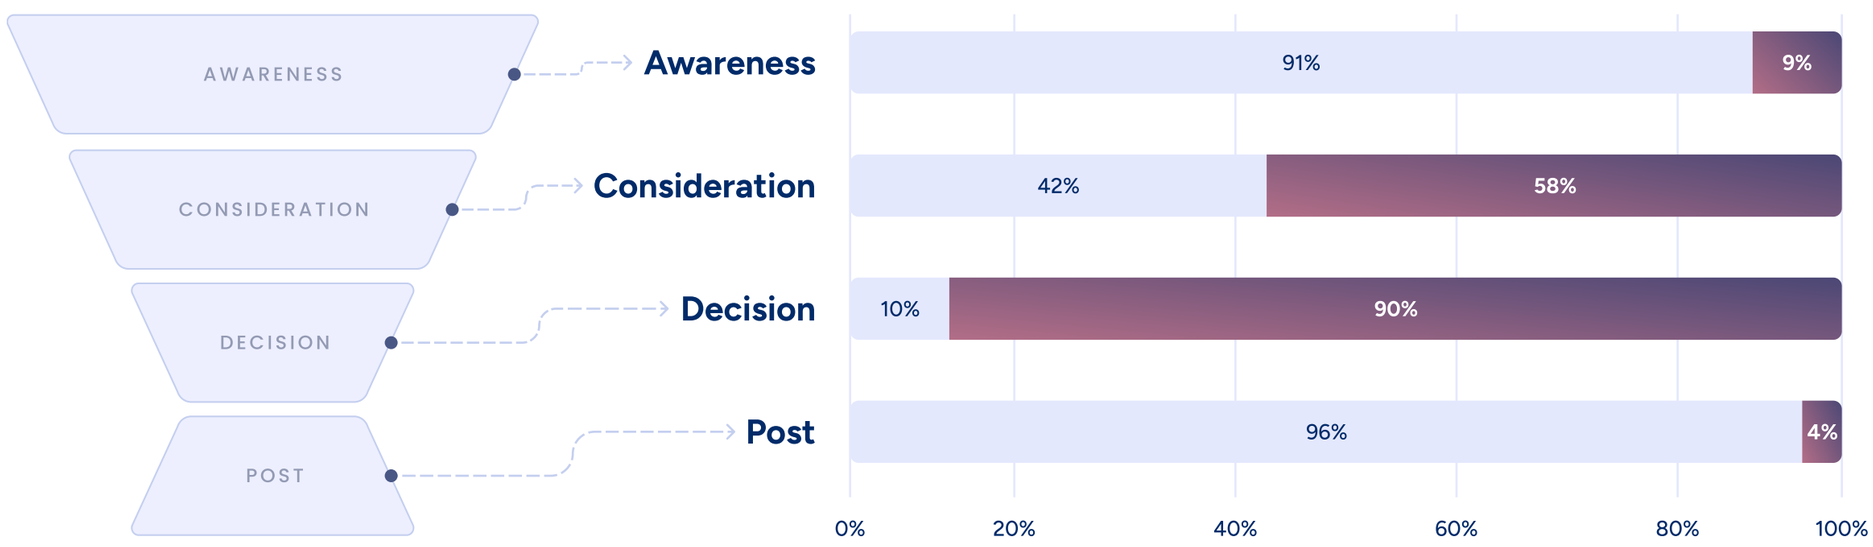
AWARENESS
8.8%
(e.g., How do business credit cards work?)

CONSIDERATION
58.3%
(e.g., What's the best credit card for a student?)

DECISION
90.3%
(e.g., Best HYSA promotional intro rates deal today?)

Prompt Intent by Whether the Response Returned a Citation

By stage of the customer journey



© 2026 Fintel Connect | All Rights Reserved

Prompt wording

How a prompt is worded strongly predicts whether it pulls in a source.

Wording that **reliably** triggers a citation:

MARKERS	EXAMPLE WORDS	CITATION %
Urgency	“ready to apply”	**100%** (45 of 45)
Recency	“currently” or “today”	**98.4%** (242 of 246)
Superlative	“best”	**87.1%** (431 of 495)
Opening words	“which”	**90.7%** (438 of 483)

Wording that **rarely** triggers a citation:

MARKERS	EXAMPLE WORDS	CITATION %
Explainer phrasing	“how does X work”	**7.8%** (23 of 294)
Advice phrasing	“should I”	**16.9%** (41 of 243)
Opening words	“how”	**8.2%** (13 of 159)

ASK › RETRIEVE › APPEAR › COMPETE › GET CHOSEN

But triggering retrieval only creates the opportunity. The next question is what sources the models are actually choosing to read, and where a brand should focus its attention.



The Marketing ROI Gap in Banking: How Financial Institutions Spend, Measure, and (Struggle to) Prove Marketing’s Value

Cornerstone Advisors’ new research shows how banks and credit unions are measuring marketing ROI, where attribution falls short, and why spend and performance remain misaligned.

READ MORE ›

© 2026 Fintel Connect | All Rights Reserved

ASK • **RETRIEVE** • APPEAR • COMPETE • GET CHOSEN

How AI Answers Financial Questions

The model decides whether to use current web sources or answer from memory

INTRODUCTION

Once the question is asked, the next decision happens inside the model: does it rely on what it already knows, or does it look for current evidence? For marketers, this distinction matters because live retrieval is where today's content and partnerships can enter the answer.

There are two ways financial brands can end up in an AI answer:

- Through **live URLs** retrieved when the prompt is submitted
- Through information the model has **already 'learned'** from its training data

In the latter case, the AI draws on information learned during training from large volumes of text, images, audio, video, and research. This is sourced across content pages, digitized books, research articles, and more from around the internet or uploaded offline by the AI developers. As a result, it only accounts for that moment in time, when the AI model was trained, and any new brand information since that time is never surfaced.

A **citation** means the AI treated a specific page as authoritative enough to reference. A **brand mention** means the brand got named, with or without a source behind it.

This is an example of a **citation**

These are examples of **brand mentions**

what are the best high-yield savings accounts in the US?

As of **August 31, 2026**, the strongest U.S. high-yield savings accounts are generally paying around **4.0% to 4.2% APY** without unusually restrictive requirements. The Publisher's current comparison puts the top listed rate at **4.21% APY**, versus a national savings average of about **0.38%**. Publisher

Account	Current APY	Best for	Watch-outs
Financial institution 1	**Up to 4.21%**	Highest rate among mainstream picks	Requires linked checking plus balance/direct-deposit or deposit requirements
Financial institution 2	**4.01%**	High APY with fewer hoops	Smaller bank than the major national brands

A citation comes from a specific page, so brands and affiliates can influence citation visibility by investing in the content formats, topics, and affiliates that AI models retrieve. **That makes prompts that trigger citations the ones brands can most directly influence today.**

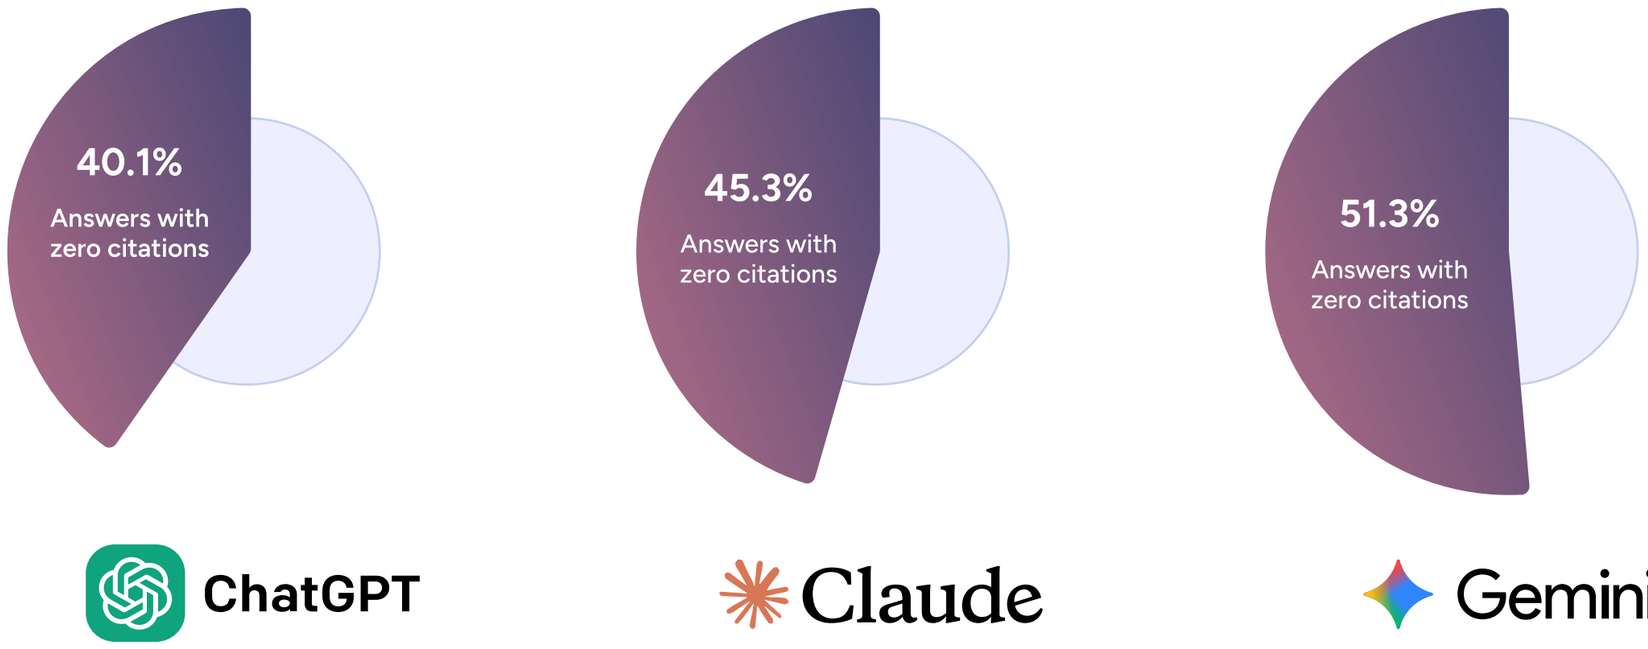
© 2026 Fintel Connect | All Rights Reserved

Why Brands Still Get Named When There's No Citation

Of the 2,400 AI answers in this study, 1,093 (45.5%) came back with no citation at all. These answers are built entirely from what the AI model has already learned during training. When a model names a brand anyway, that name was decided long before the question was asked, pulled from a fixed shortlist of the largest and most well-known incumbents.

Answers With Zero Citations, by Model

Share of 800 answers each

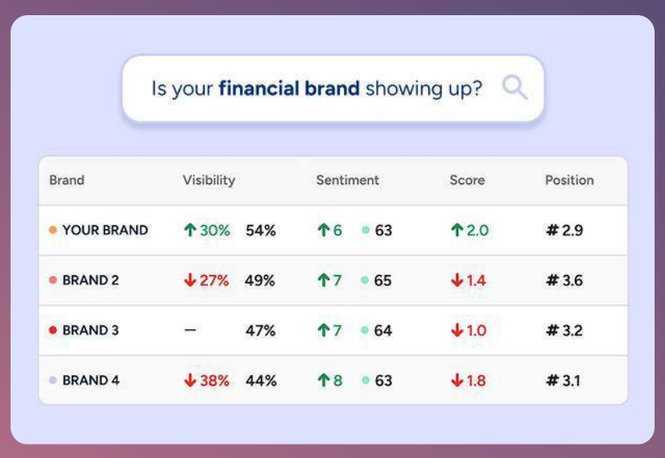


Of those 1,093 un-cited answers (45.5%), a financial brand was named in only 266, or 24.3%. The remaining 1,307 answers (54.5%) retrieved current sources from the web, creating an opportunity for existing content to shape the response.

When no sources are retrieved, **a brand's current content has no way to influence the answer**. Whatever gets named was determined during training long before that content existed. That's why the rest of this report focuses on the types of prompts consumers ask that actually trigger AI models to retrieve sources, where new or optimized pages from brands and affiliates can shape outcomes today.

ASK › **RETRIEVE** › APPEAR › COMPETE › GET CHOSEN

Being read still does not mean being seen. The next hurdle is making it into the visible answer.



GEO Services

Fintel Connect gets banks, credit unions, and fintechs featured in AI-generated answers. We use a curated affiliate network, monthly intelligence reports, and a clear optimization strategy to build visibility where your competitors aren't yet.

READ MORE ›

© 2026 Fintel Connect | All Rights Reserved

Getting Into the AI Consideration Set

Retrieved evidence determines which brands make it into the consideration set

INTRODUCTION

Retrieval creates an opportunity to be considered, but it does not guarantee visibility. A model can read a page about a brand without ever showing that brand to the consumer, which means AI visibility happens in stages rather than as a simple yes-or-no outcome.

Read vs. Named vs. Recommended

As an AI builds an answer, a brand can show up at three different levels:

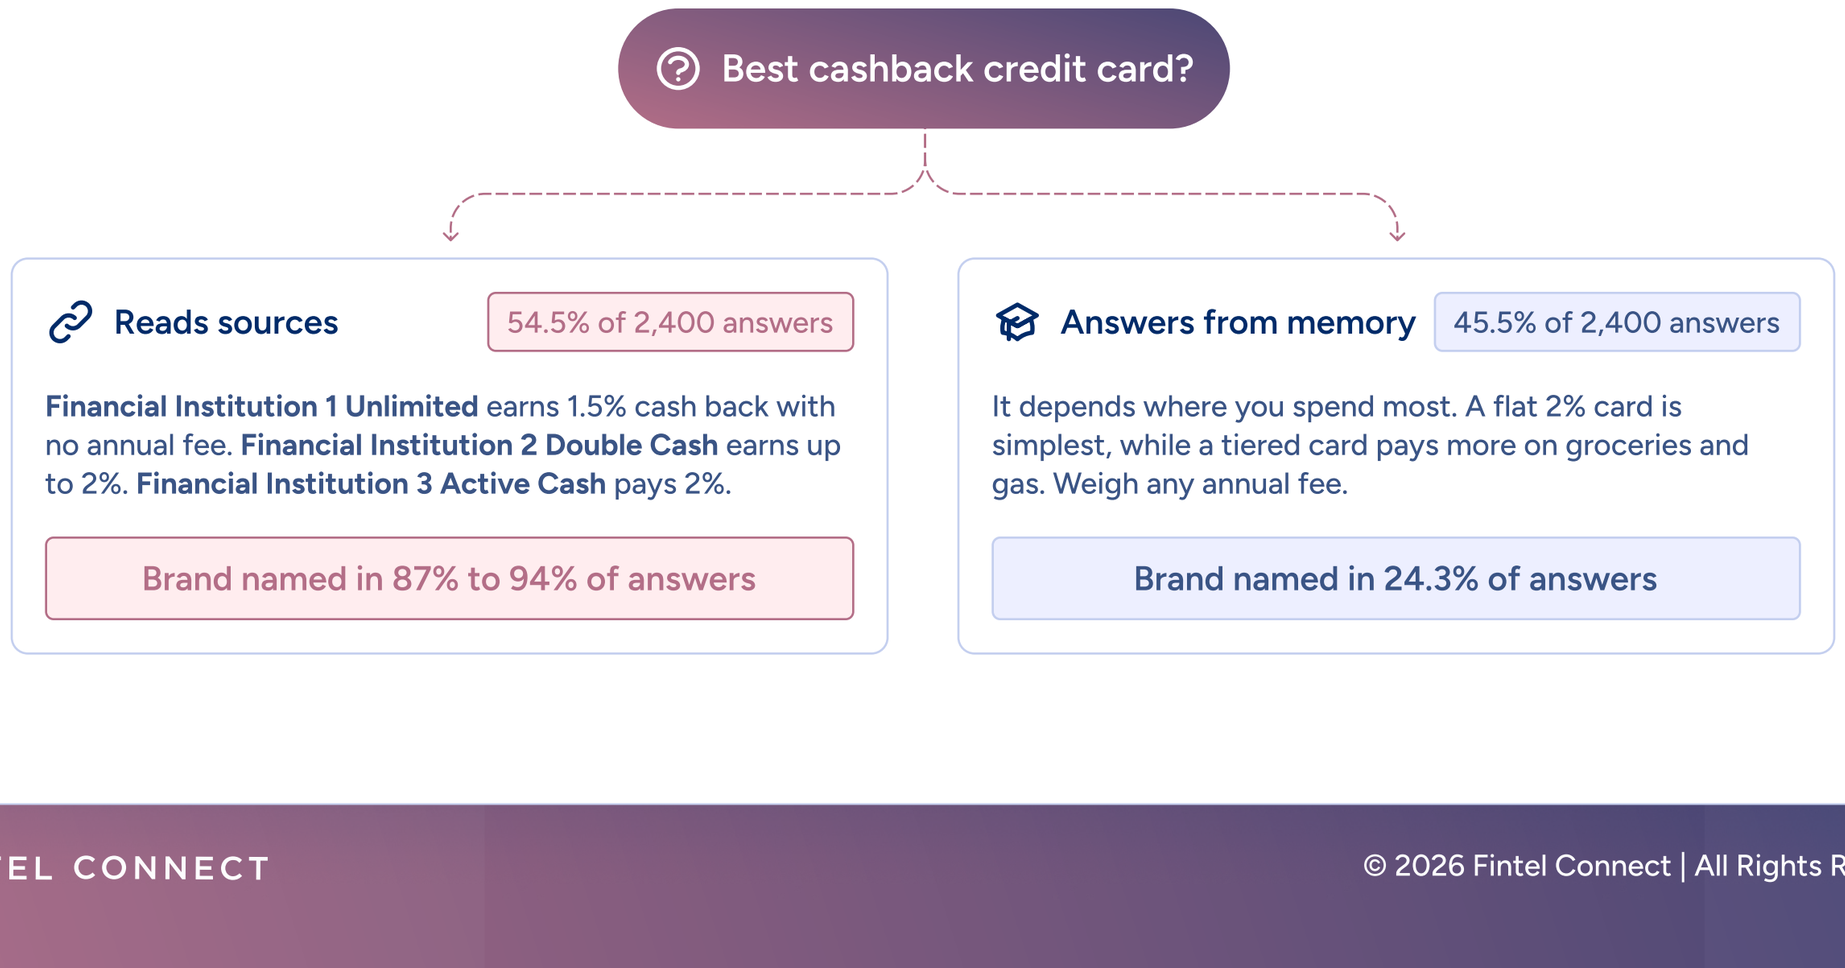
1. **Read:** the AI reads a page with the brand in it while composing the answer, but the brand's name may never make it into what the consumer sees.
2. **Named:** the brand's name is written into the answer text itself.
3. **Recommended (goal):** the brand is prioritized ahead of other options in the answer, always receiving the highest position and more detailed description.

How Brands Are Synthesized Into AI Answers

1. Read

Reading a source is what puts brand names in the answer: a brand appears 87% to 94% of the time when at least one source is read, versus only 24.3% when none is. Without a source, the model falls back on brands it learned during training, typically the larger, longer-established banks. So reading a source is what opens the door for newer or smaller brands to appear at all.

That is why GEO visibility comes down to two things: coverage on the affiliate sites AI reads most in your category, plus detailed product pages and content on your own site.



© 2026 Fintel Connect | All Rights Reserved

2. Named

Being named is where brand discovery actually starts: the brand is now part of the conversation.

Answers name 3.5 brands on average when no affiliate page was read, rising to 4.1 when one was. But which brands get named is far from even: across the credit card prompts studied, 26 distinct brands appeared in total, and **just five of them account for 72.0% of all mentions**. Every brand outside that top five is competing for a share of the remaining 28.0%.

3. Recommended

At stage 3, the brand reaches the highest-value outcome measured in this study: being recommended first.

> The **best pure-rate option is currently Financial Institution 1**, but there's a catch: to get the **4.21% APY** you need a linked Bank checking account and must meet deposit/activity requirements. If you want something simpler, **Financial Institution 2 at 4.01%** looks more attractive because there are fewer hoops. Publisher
>
> Interestingly, **Financial Institution 4 was the Publisher's overall best savings account for 2026**, offering 4.20% APY with no monthly fee, but as of **August 26** it stopped taking new applications because of demand. Publisher

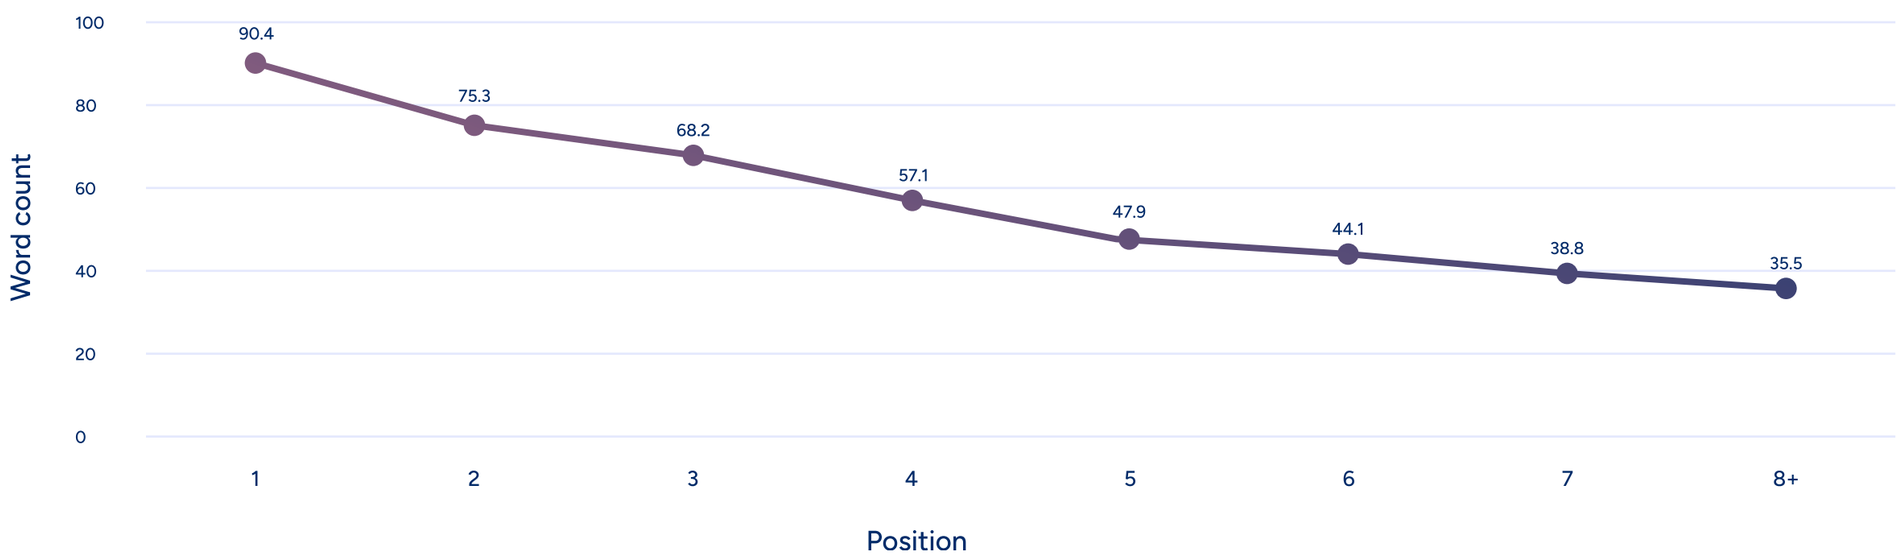
Position matters here: **each step down the order gets fewer words of description**.

Across the 1,441 answers that name at least one brand, the first-position brand receives 20.1% more words than the second-position brand and 32.5% more than the third (90.4 words on average, against 75.3 and 68.2). That represents an additional 15 to 22 words of product detail and persuasion compared with the second and third positions. Brand order decides description space.

1,244 of the 2,400 answers made a recommendation. Of those, 72.0% were supported by at least one affiliate page citation, 38.6% by a page on a brand's own site, and 23.5% by both. Appearing in an answer at all is already a high bar, but being recommended first is higher still, since that position carries the most extensive product description and leaves the strongest impression on the consumer.

Average Word Count Allocated to a Brand

By position of the brand in the response



ASK › **RETRIEVE** › **APPEAR** › COMPETE › GET CHOSEN

Once a brand enters the answer, it is no longer competing for visibility alone. It is competing for position.

© 2026 Fintel Connect | All Rights Reserved

What Determines Who Gets Chosen

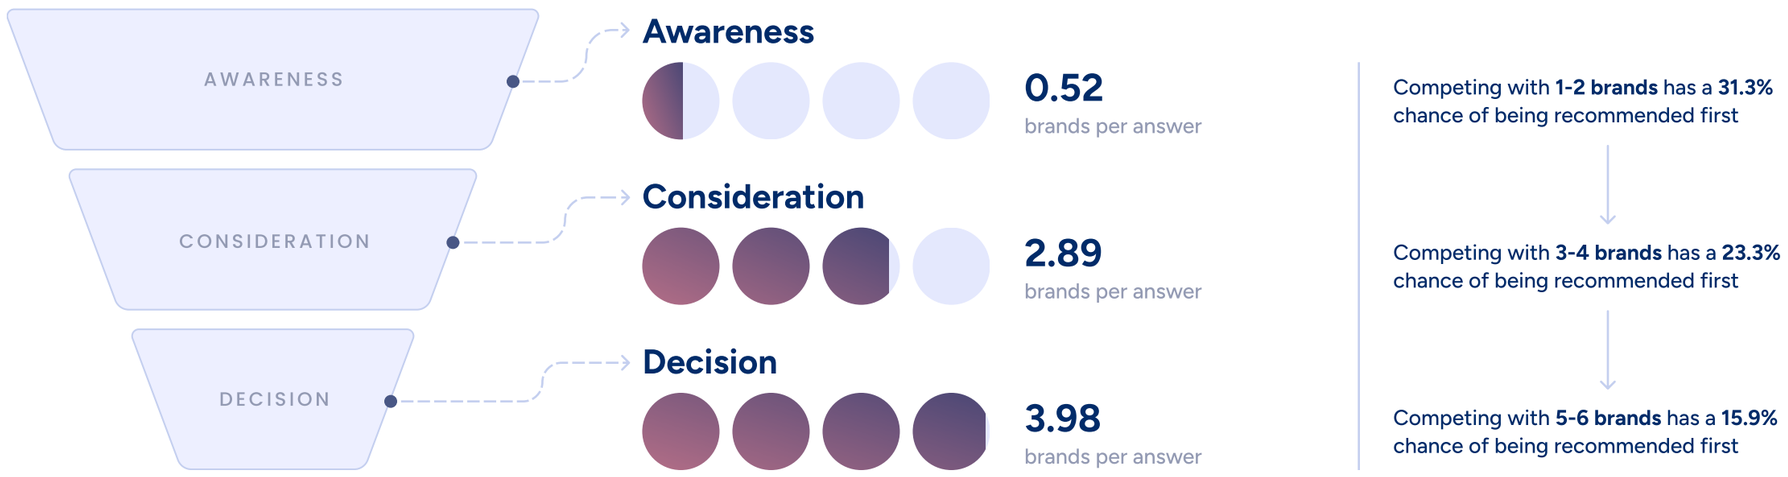
Once multiple brands appear, they compete for position and attention

INTRODUCTION

Once multiple brands enter the answer, the problem changes from visibility to competition. The closer a consumer gets to a decision, the more brands the model tends to surface, meaning the prompts marketers can influence most are also the ones where the consideration set becomes more crowded.

1. **Owned content being read nearly doubles the odds of being recommended first.** Brands whose own page was read are recommended first at roughly 24.5%, compared to 15.3% when it wasn't.
2. **Affiliate coverage adds a second lift**, and the two effects stack. Brands appearing in answers that retrieve affiliate content are even more likely to be recommended first when the model also retrieves the brand's own page.

As the consumer moves from research toward a decision, more brands compete for the same answer: awareness answers name 0.52 brands on average, consideration answers 2.89, and decision answers 3.98. That crowding cuts each brand's odds of being recommended first, from 31.3% when only one or two brands are named to under 9% when nine or more are.



Brand concentration at the decision stage is significant. Across the four product categories, the **top three brands** account for 55.9% to 74.9% of recommendations. Every other brand competes for the **remaining** 25.1% to 44.1%.

CONSUMER CREDIT CARDS

Capital One, CHASE, DISCOVER

59.7%
of recommendations

BUSINESS CREDIT CARDS

CHASE, AMERICAN EXPRESS, Capital One

74.9%
of recommendations

HIGH-YIELD SAVINGS ACCOUNTS

ally, Capital One, Marcus by Goldman Sachs

55.9%
of recommendations

PERSONAL LOANS

SoFi, LightStream by Truist, Upstart

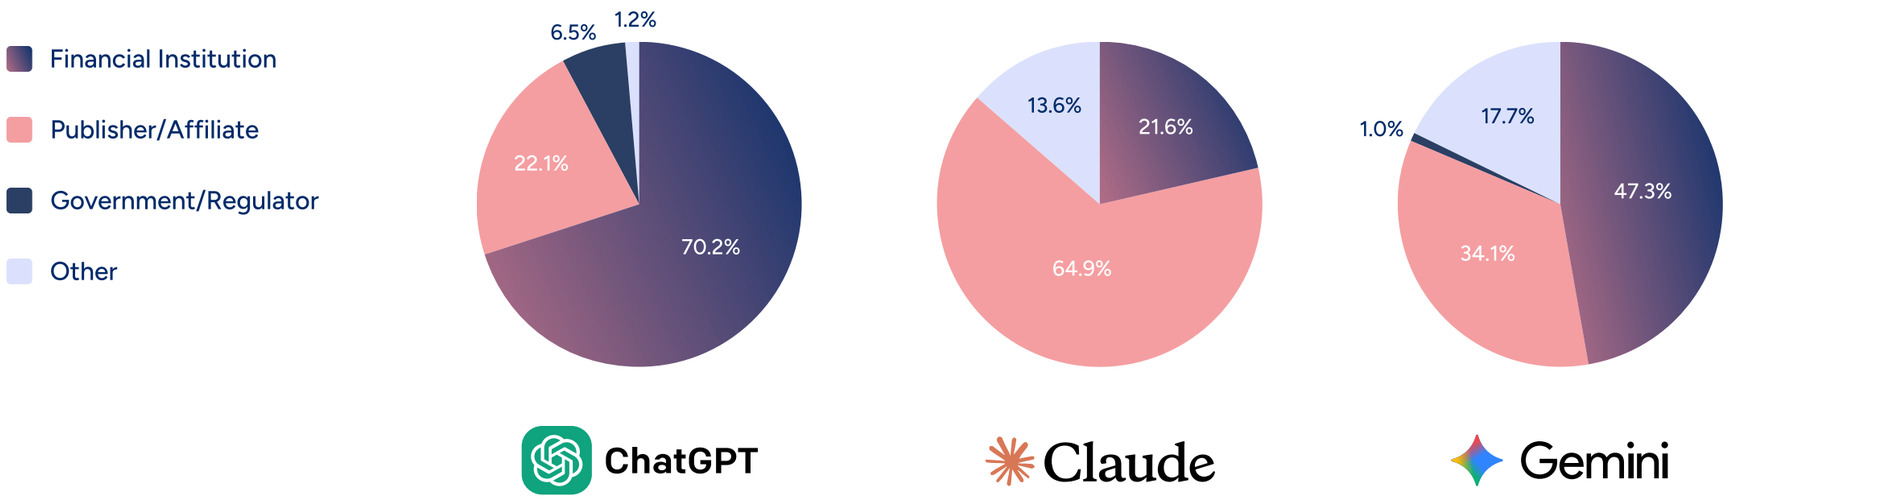
62.4%
of recommendations

© 2026 Fintel Connect | All Rights Reserved

Appearing in the consideration set is only the first hurdle. The final question is what makes one brand receive more attention, more description, and a higher position than the alternatives. Our data suggests that recommendation outcomes are shaped by the combination of owned evidence and third-party coverage available to the model.

Source-Type Share of Citations, for *All* Answers

By LLM Model



By AI model, the split between affiliate sources and brand-owned pages varies. In answers that **name or recommend a specific brand**, affiliate pages show up consistently across every platform. Of ChatGPT's 412 recommendation answers, 225 (54.6%) cite at least one affiliate page. The equivalent rate is 70.7% for Gemini and 93.8% for Claude.

Affiliate influence peaks where choices are made: 74.8% of decision-stage answers retrieve an affiliate page, and an affiliate page is the first source retrieved in 47.4% of those answers.

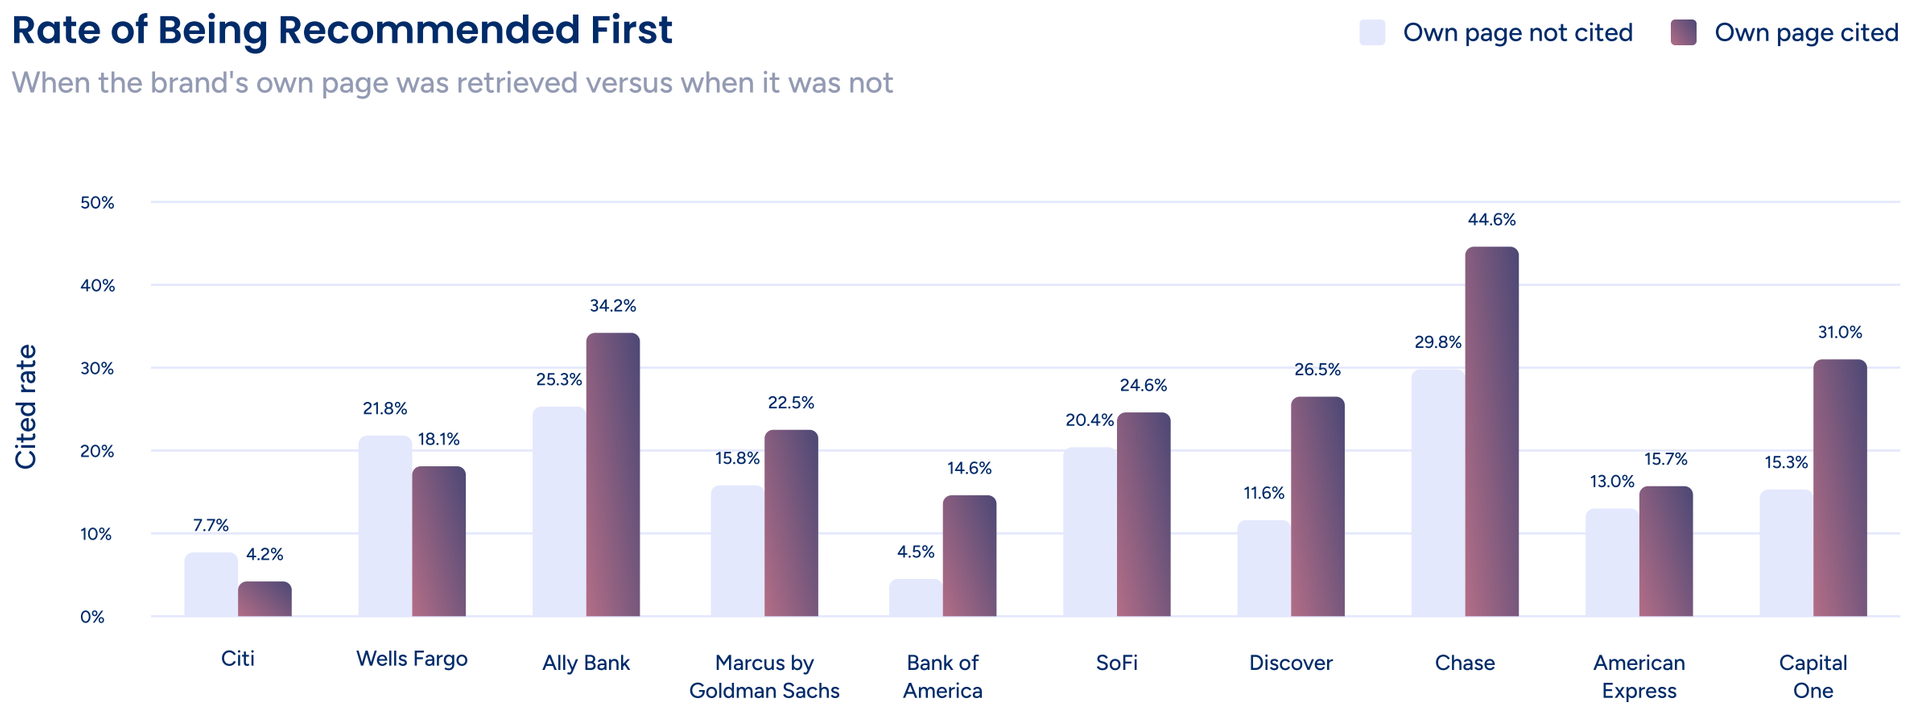
Two Factors That Increase a Brand's Likelihood of Being Recommended First

1. Whether the brand's own page was among the sources read

Brands whose own page was read are recommended first 24.5% of the time, compared to 15.3% when their page wasn't part of the answer. The effect is even stronger for larger, longer-established financial brands with more published content: Chase is recommended first 44.6% of the time when one of its own pages was read, versus 29.8% when none was.

Rate of Being Recommended First

When the brand's own page was retrieved versus when it was not



© 2026 Fintel Connect | All Rights Reserved

2. Whether an affiliate page was read in the same answer

Brands mentioned in answers where an affiliate page was retrieved are recommended first 19.4% of the time, compared to 15.7% without one.

Splitting the same 6,342 mentions by which sources were present makes the pattern clear:

PAGES RETRIEVED		RECOMMENDED FIRST%
✗ Brand's own page	✗ Affiliate's page	10.8% (139 of 1,293 mentions)
✗ Brand's own page	✓ Affiliate's page	17.3% (533 of 3,088 mentions)
✓ Brand's own page	✗ Affiliate's page	23.8% (188 of 789 mentions)
✓ Brand's own page	✓ Affiliate's page	**24.9%** (292 of 1,172 mentions)

Who Owns the Influence

1. Affiliate concentration in citations

The top five affiliates (NerdWallet, Bankrate, Forbes Advisor, WalletHub, and Experian's editorial site) hold 1,586 of the 3,055 affiliate citations, 51.9%. The remaining 95 affiliates share the rest.

The leaders change by product category. WalletHub leads credit cards with 103 citations, Bankrate and Forbes Advisor lead high-yield savings with 135 and 136, and personal loans splits across Bankrate (141), LendingTree (130), Experian (126), and Credible (90). The Points Guy's 67 citations sit entirely in card categories. A partner list built once and reused across product lines misses these differences.

Affiliate Leaderboard by Product Category

By number of citations

	CONSUMER CREDIT CARDS		BUSINESS CREDIT CARDS		HIGH-YIELD SAVINGS ACCOUNTS		PERSONAL LOANS	
1	WalletHub	103	nerdwallet	89	Forbes ADVISOR	136	Bankrate	141
2	NerdWallet	89	The Points Guy	67	Bankrate	135	Lending Tree	130
3	Bankrate	75	Forbes Advisor	59	NerdWallet	113	Experian (editorial)	126
4	Forbes Advisor	68	WalletHub	52	CNBC Select	66	NerdWallet	105
5	Experian (editorial)	50	CNBC Select	38	Experian (editorial)	54	Credible	90

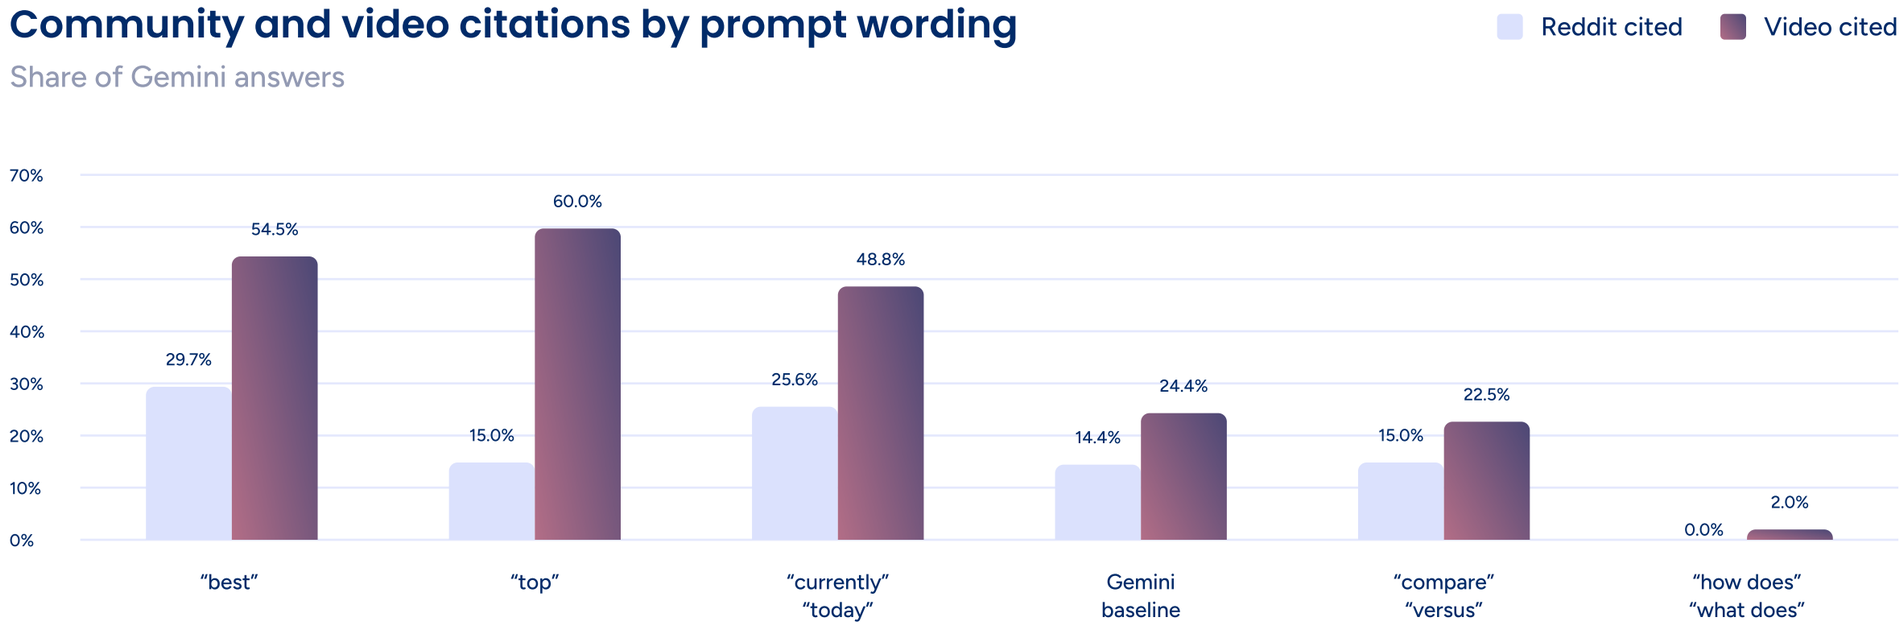
© 2026 Fintel Connect | All Rights Reserved

2. User-generated content (UGC)

On Gemini, prompts containing "best" returned Reddit citations in 49 of 165 answers (29.7%) and video citations in 90 of 165 (54.5%), well above the Gemini baseline of 14.4% for Reddit and 24.4% for video. Explainer prompts showed the opposite pattern: zero Reddit citations and video in just 2 of 98 answers.

Community and video citations by prompt wording

Share of Gemini answers



To reach these audiences, **consider community engagement tactics like Q&As** (called an "AMA," or "Ask Me Anything," on Reddit) **alongside paid ads on Reddit and YouTube.** Together, these help shape brand perception and influence the conversations people have about your brand.

Above all, brands should approach these communities authentically and contribute information users can apply in their own lives. Treating UGC primarily as a selling opportunity risks undermining that credibility.

3. Brand-owned pages

A brand-owned page is a page on the financial institution's own website: a product page, rates page, or explainer article. Across the 3,489 brand-owned citations, 48 institutions appear. The top three, Chase (393), Capital One (384), and American Express (355), hold 1,132 of them, 32.4%. Adding Discover, Ramp, SoFi, Bank of America, Upstart, Ally Bank, and Brex brings the top ten to 2,112, 60.5%.

Which type of page earns the citation varies widely by institution. Across all brand-owned citations, product pages account for 57.9% and educational articles for 23.0%. That mix flips by institution. American Express reaches 94.4% product pages and no educational articles, and Bank of America and Ally Bank are cited entirely through product pages. Ramp runs the other way, with 82.2% of its citations coming from educational articles.

The pattern also moves by category: in high-yield savings, product and rates pages account for 66.7% of brand-owned citations and educational articles just 6.8%, so for a savings product the page that states the rate is the page that gets read. In business credit cards the split is wider, 47.2% product pages and 34.1% educational articles, so both a product page and a how-to article on decision-stage prompts can carry a business-card brand into an answer.

ASK › RETRIEVE › APPEAR › COMPETE › GET CHOSEN

The final challenge is turning presence into preference: what makes one brand get chosen ahead of the rest?

© 2026 Fintel Connect | All Rights Reserved

ASK • RETRIEVE • APPEAR • COMPETE • **GET CHOSEN**

What Financial Marketers Should Do

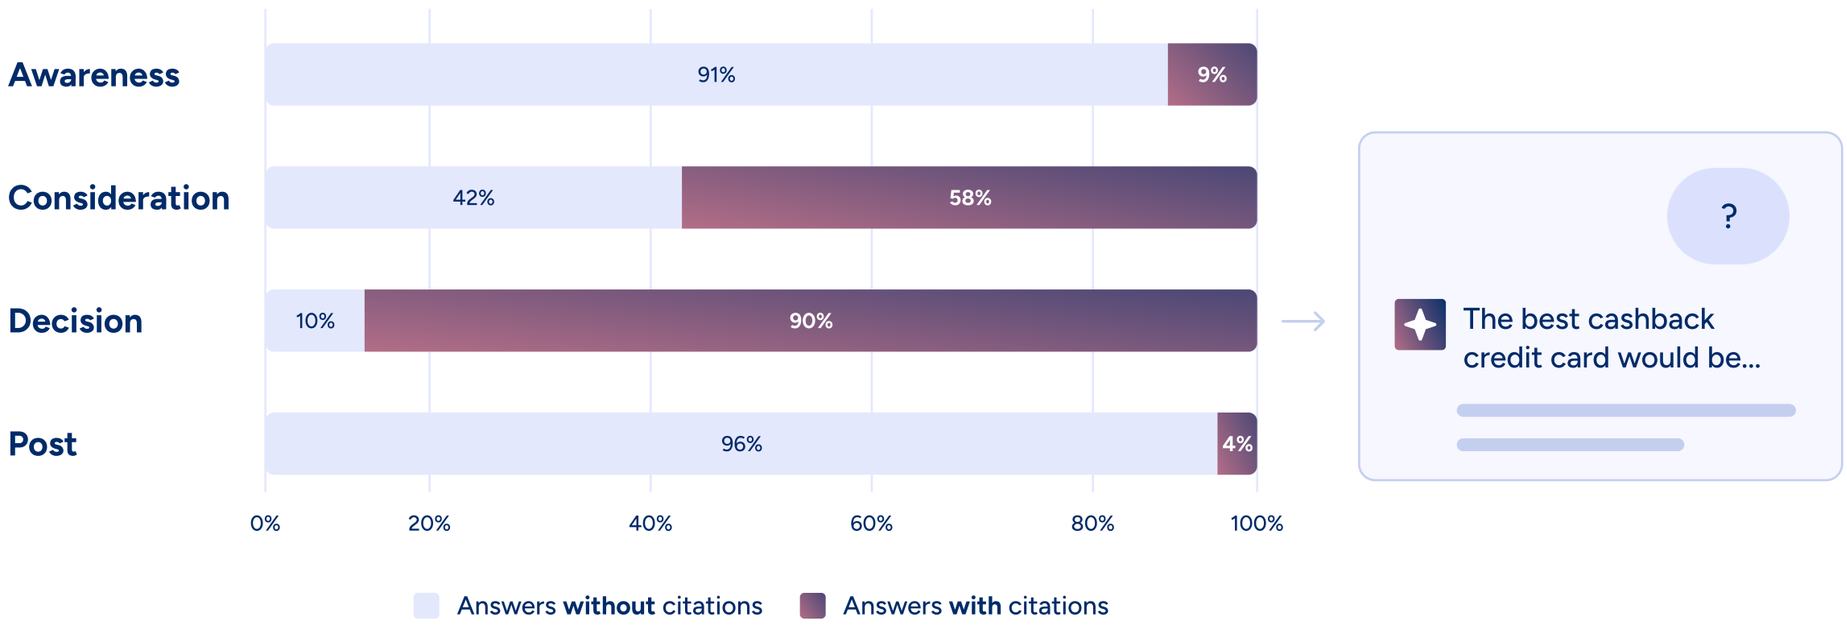
How to build an owned and affiliate strategy based on where your brand is today

INTRODUCTION

Appearing in the consideration set is only the first hurdle. The final question is what makes one brand receive more attention, more description, and a higher position than the alternatives. The data suggests that owned content and third-party coverage work together to strengthen a brand's presence in the evidence the model uses.

1. Spend where retrieval happens

Awareness prompts retrieve sources only 8.8% of the time, which limits how quickly newly published content can influence those answers. Decision-stage prompts, by contrast, retrieve current sources much more frequently. **Decision and eligibility prompts retrieve sources up to 90.3% of the time**, and pages published this quarter can enter those answers this quarter.



2. Value affiliates beyond the click because affiliate partnerships now drive both acquisition and upstream AI discovery

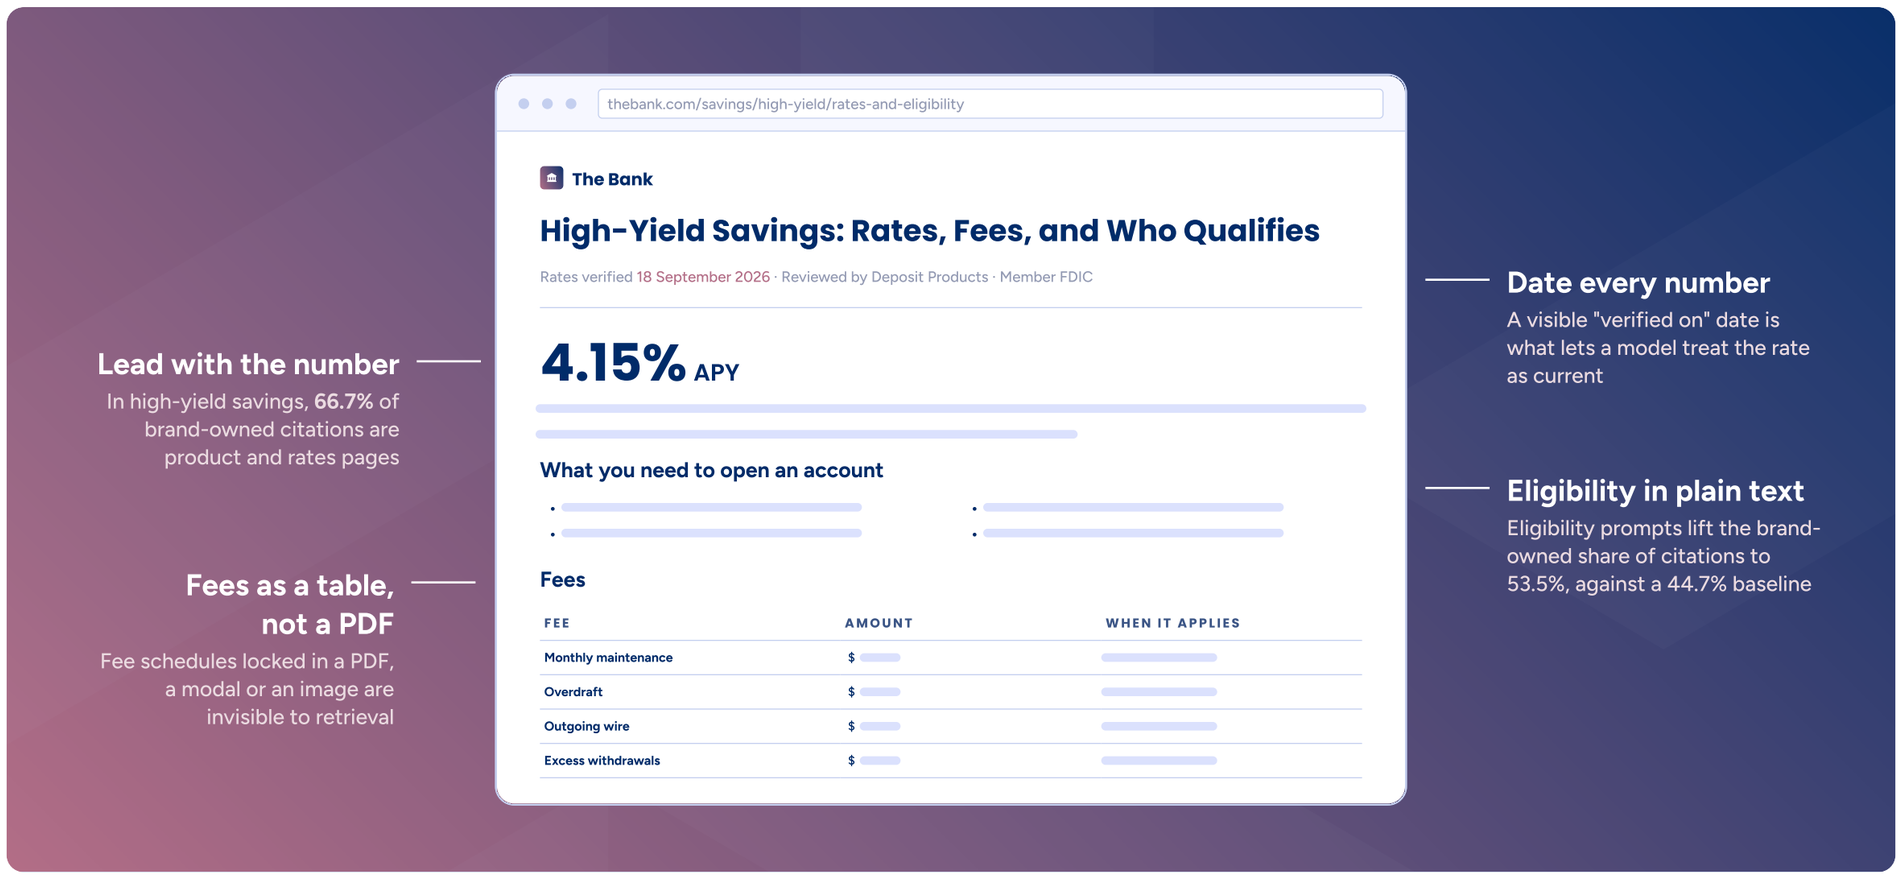
Affiliates support 72.0% of all recommendation answers, five affiliates hold 51.9% of affiliate citations, and comparison or ranked-list pages are 75.0% of what gets read from affiliates. Their influence operates before any click as affiliate pages are read in 43.3% of answers and their weight peaks at the decision stage, where 74.8% of answers read an affiliate page.

An affiliate partnership therefore carries two kinds of value: the measurable acquisition it has always carried, and inclusion in the recommendations and competitive sets that AI answers assemble upstream of the click.

© 2026 Fintel Connect | All Rights Reserved

3. Publish content that earns retrieval

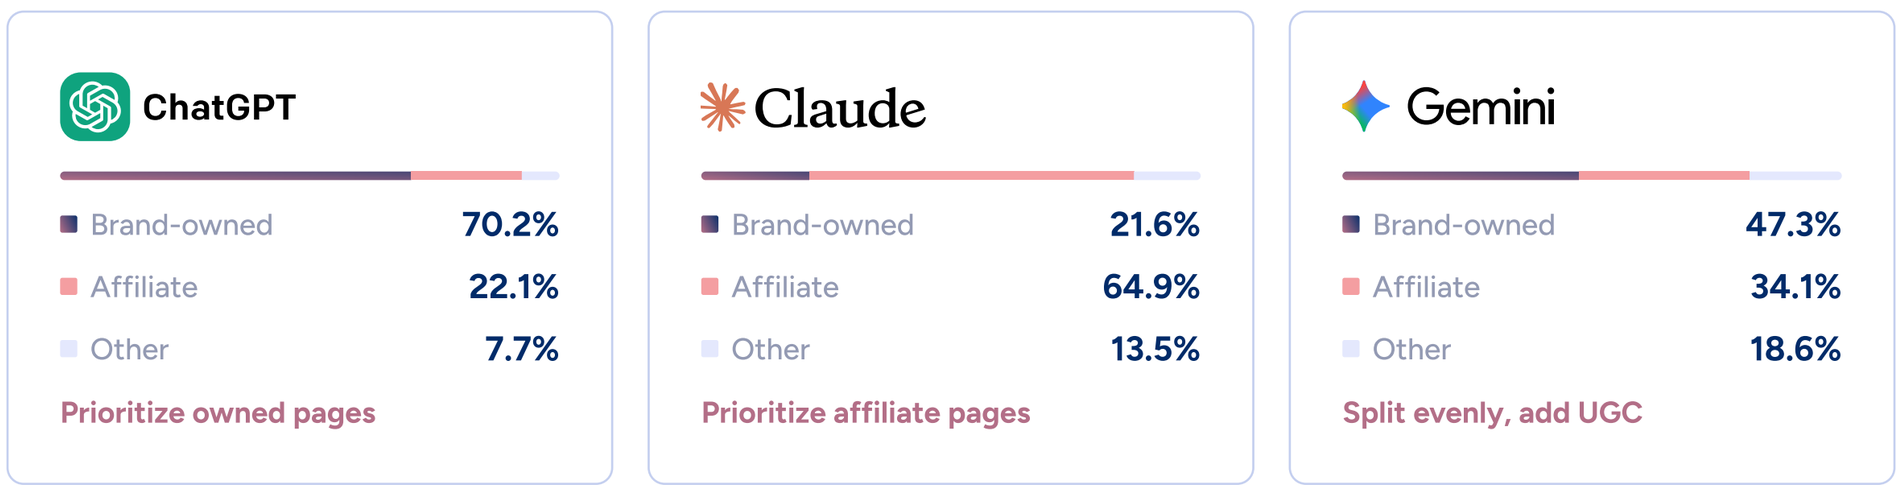
When a brand's own page gets read, that brand is recommended first 24.5% of the time, compared to 15.3% when it isn't. For eligibility-related prompts specifically, brand-owned pages account for 53.5% of citations, up from a 44.7% baseline. Approval and eligibility pages, ones that spell out rates, fees, and criteria in plain, readable text, are the cheapest owned asset with a measurable payoff.



4. Plan per platform

ChatGPT uses brand-owned pages in 70.2% of its citations, Claude cites affiliate pages for 65.0%, and Gemini divides between the two and adds community and video sources. Affiliate coverage belongs in the plan on all three platforms, since affiliate content sits under the majority of decision-stage answers on each.

What changes by platform is the balance of owned and affiliate partner work. A single AI strategy applied uniformly across all three platforms is likely to overinvest in some source types and underinvest in others.



© 2026 Fintel Connect | All Rights Reserved

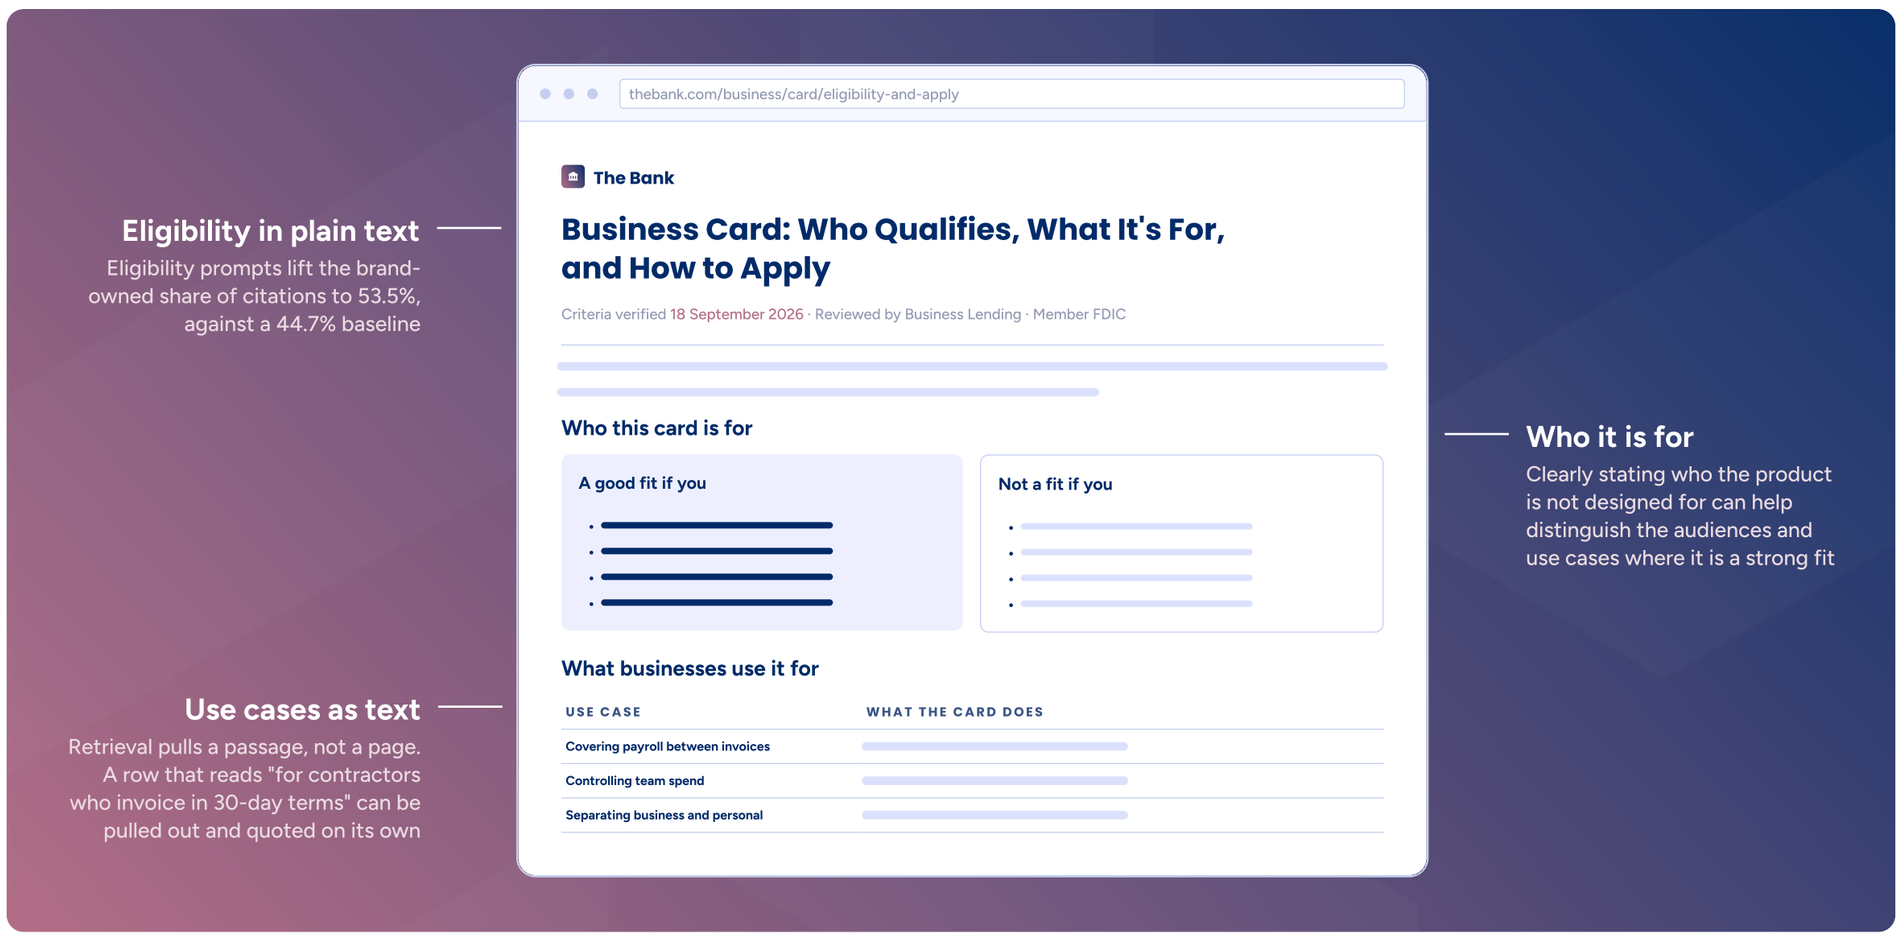
Positioning Your Brand to Get Chosen in AI Answers

The correct GEO strategy for your business ultimately depends on brand maturity and your existing digital presence, especially concerning affiliate partnerships and the breadth of product content.

Challenger brands/new to affiliates

Immediately start building your tracked prompt set, similar to your SEO keyword target list. Focus on decision-stage and eligibility prompts where 90.3% of answers cite a source. Those are the prompts where published pages can boost visibility outcomes this quarter.

1. **Publish detailed product approval and eligibility pages on your website:** What type of customer is the product fit/not fit for? What are the specific use cases for your product? Where can customers apply for the product, and how easily/quickly?
2. **Identify which affiliate pages are most influential in each product category,** since retrieval patterns differ significantly by product type.
3. **Prioritize inclusion in the affiliates' most frequently retrieved comparison and ranked-list content.** Building on pages that already perform gives the brand a clearer path into the sources AI systems are already using.



Mid-sized institutions/growing an established program

Evaluate affiliate partners on two things that were measured: how often their pages are read on decision-stage prompts, and how often they appear in answers where your brand is recommended.

1. **Build the list of prompts where your brand appears zero times**, since that list is the priority order for the next quarter of partner content.
2. **Explore new comparison content opportunities with your affiliates to address the gaps**. These should be head-to-head and clearly explain how your product differs from competing options and where its value proposition is strongest.
3. **Audit product pages so rates, fees, and eligibility are plain text and correct**. Accurate information signals trust and reliability, whereas conflicting information can reduce the likelihood that the brand is included in the answer.

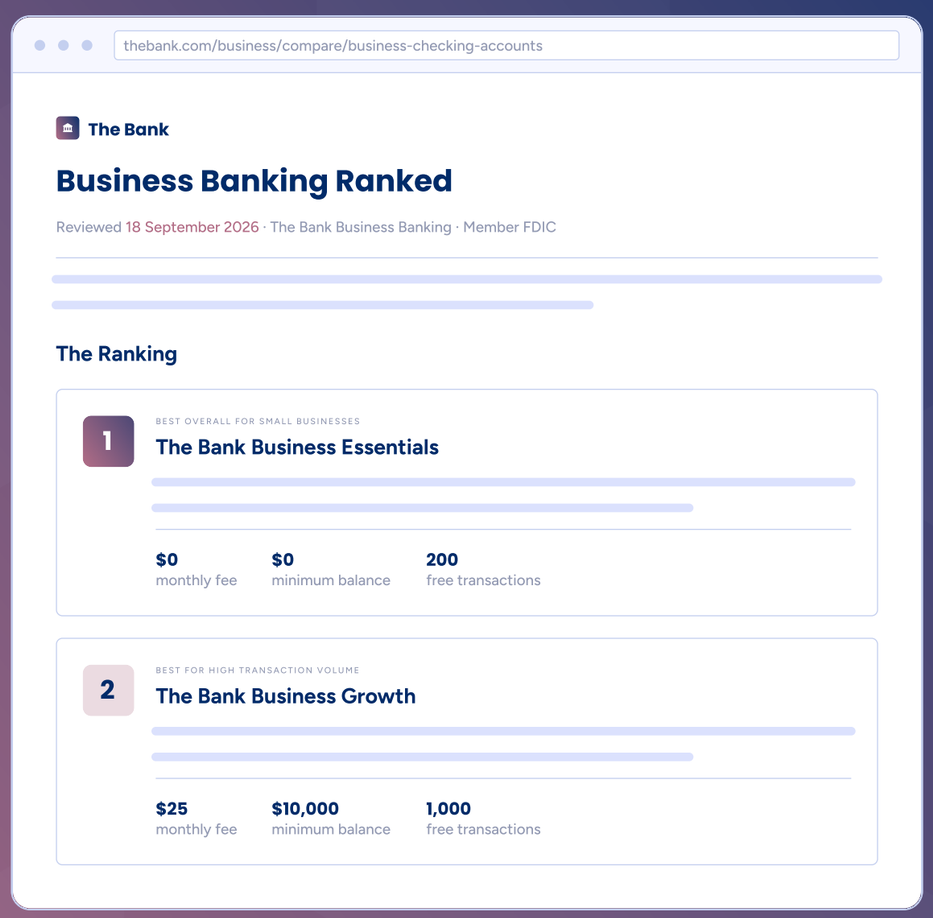
© 2026 Fintel Connect | All Rights Reserved

Enterprise, established brands/complex multi-product affiliates

Report per engine with a target per product line.

1. **Publish ranked comparison pages on owned domains**, because 75.0% of affiliate citations are listicle or comparison pages, and that format is what gets read.
2. **Show partner presence at citation level alongside referral reporting in one view**, so the partner contribution to AI visibility sits beside the partner contribution to tracked revenue.

Rank it, don't list it
A ranked page gives the model an order to quote



"Best for" lines
Answers match a person to a product. "Best overall for small businesses" is what gets lifted

What distinguishes AI search from traditional SEO is breadth of coverage: a single keyword-optimized page rarely carries an answer on its own. Getting featured consistently in AI answers takes a multi-channel mix:

- Third-party comparison content above all
- Descriptive product pages and informational owned articles
- An engaged, positive user-generated content presence so that your products are seen as directly relevant to the customer's specific question and intent

ASK › RETRIEVE › APPEAR › COMPETE › GET CHOSEN

The research points to a broader shift in how financial marketers should think about AI visibility. The goal is not simply to optimize for one platform or chase individual GEO tactics, but to make the brand easier to retrieve, validate, compare, and recommend across every source an AI might consult, from your own site to affiliate and third-party content.

© 2026 Fintel Connect | All Rights Reserved

What Changed in AI Search Between 2025 and 2026

The overall pattern

The same few things stayed constant across both years: which affiliates lead, which page formats get read, and which prompts favor affiliates versus brand pages. Meanwhile, how individual AI models behave keeps changing.

In short, the underlying opportunity looks stable, but the platform landscape it plays out on does not.

A More Comprehensive Data Set

Our 2025 AI search report, *Competing for Visibility in the Age of AI*, measured citation sources and citation volume across ChatGPT, Gemini, Perplexity, and Copilot, using 29 prompts in two categories. Those four platforms returned 419, 286, 287, and 99 sources respectively, counting links plus sources named in the answer text.

The 2026 edition widens the base to 800 prompts, returning 7,809 total citations, and four categories across three platforms. Claude was added this year, while Perplexity and Copilot were not included.

What Held Steady

1. NerdWallet and Bankrate remain two of the most influential comparison affiliates

In 2025, Bankrate and NerdWallet topped the source table, together accounting for 15% of all sources retrieved across the four platforms. American Express was the only bank to reach that top ten.

In 2026, the same two names lead again at larger scale: NerdWallet at 396 citations and Bankrate at 381, both part of a six-affiliate group that holds 57.3% of all 3,055 affiliate citations. Despite differences in prompts, models, and product categories, the two studies surfaced the same two leading affiliate sources.

For a brand deciding which affiliates to prioritize, this consistency is the most reliable signal in this section: their consistency across both studies makes them important partners to evaluate for long-term visibility.

2. Comparison-format pages get read, and embedded promotions stay outside the answer

In 2025, over 70% of answers presented a list or a list alongside a table, and iFrame-based promotional units appeared in 12% of retrieved websites and in zero AI answers. In 2026, 75.0% of affiliate citations are comparison or ranked-list pages, against 3.2% of brand-owned citations.

The format finding has now held across two years, six platforms, and two counting methods: the models read structured comparison pages, and embedded ad units go unread.

© 2026 Fintel Connect | All Rights Reserved

3. Prompt intent trends

In 2025, informational prompts surfaced bank and educational pages while comparative prompts surfaced affiliate listicles. The 2026 collection measures a similar pattern:

- **Explainer, advice, and risk prompts:** affiliates provide 28.3% of citations
- **Comparison prompts (where options are compared and chosen):** affiliates provide 43.1% of citations
- **Eligibility-worded prompts:** brand-owned pages hold 53.5% of citations

What Has Shifted

1. ChatGPT's citation sources flipped, affiliates still shape its recommendations

In 2025, ChatGPT was the affiliate-first platform: 290 of its 419 sources (69%) were affiliate pages and 129 (31%) were bank or FI pages. In 2026, its citation mix has reversed, with brand-owned pages at 70.2% of its 2,523 citations, holding 63.5% to 82.7% across every commercial question type measured.

The 2026 data suggests affiliates' role has changed, not disappeared. ChatGPT still retrieves at least one affiliate page in over half of its recommendation answers (54.6%), and that reading concentrates where buying decisions actually happen: comparison, segment-specific, and transactional prompts, all above ChatGPT's overall average. In other words, affiliate pages remain the comparison evidence behind ChatGPT's recommendations, even as brand-owned pages now dominate its overall citation mix.

2. Gemini's mix evened out, and community and video joined the answer

In 2025, bank and FI pages dominated Gemini's sources, making up 206 of 286 (72%). In 2026, that dominance is gone: brand-owned pages now account for 38.4% of Gemini's 3,382 citations, and affiliate pages for a nearly identical 37.3%, while community and video sources have moved into its highest-intent answers. This shift started in 2025, when Reddit alone reached the all-platform top ten with 34 citations. By 2026, the collection includes 201 forum citations (194 of them on Gemini) and 390 video citations, all on Gemini.

3. The platforms studied changed, and Claude is now the most affiliate-weighted

Perplexity (74% affiliate pages in 2025) and Copilot (80%) are not included in this year’s study, and Claude enters as 2026’s most affiliate-weighted platform, with affiliate pages at 65.0% of its citations.

GEO SERVICES

Increase AI search visibility through the affiliates it already trusts

Speak to an expert

© 2026 Fintel Connect | All Rights Reserved