

# AI-Mediated *Discovery* Across Industries.

How AI systems surface — and overlook — companies when buyers ask.

ENGINES <b>5</b> observed	INDUSTRIES <b>7</b> sampled	MODE <b>Observational</b>	AUDIENCE <b>Operators</b>
------------------------------	--------------------------------	------------------------------	------------------------------

ABSTRACT

A field reading of how the major AI engines assemble answers about companies — observed across sampled buyer-intent prompts in seven industries. It is not an optimization guide. It describes a market-structure shift: discovery is moving into systems that read, corroborate, and cite a narrow slice of the web, and that reach for companies they can verify rather than the ones with the strongest reputation. The patterns differ by industry; the underlying mechanism does not.

## § 01 – THE SHIFT

# Your buyers are asking machines about you. The machines are answering – *with or without you.*

Before a buyer ever reaches your website, many now ask an AI system the question that used to start with a search: *who is good at this, and who should I trust?* The system answers in prose, names a few companies, and moves on. Most organizations have never seen what it says.

---

## THE OBSERVATION

**In sampled prompts across industries, strong offline brands were frequently absent from AI-generated consideration sets *entirely* — not ranked low, but not named at all.**

This is the part that unsettles operators when they first see it: the question is not whether you rank well. It is whether you **appear**. AI-mediated discovery is unstable, fragmented, and structurally selective — the same question, asked of five engines, routinely returns five different shortlists, and a company that dominates its category offline can be missing from all of them.<sup>1</sup>

This report describes what we observe when we run buyer-intent prompts across ChatGPT, Claude, Gemini, Perplexity, and Grok, and read how each one builds its answer. It is deliberately observational. We are not selling a ranking trick; there is no ranking to win. We are describing a change in how companies get found, so you can judge your own exposure to it.

## NOTE 1

Same prompt, five engines: in measured comparisons only ~11% of cited sources overlapped between two leading engines answering the same query.

## SCOPE

Seven industries · five engines · sampled buyer-intent prompts. Observational, not deterministic.

§ 02 – THE MECHANISM

# AI systems cannot cite pages *that do not exist.*

THE CENTRAL FINDING

**An AI system can only name a company it can find, read, parse, and corroborate. Most firms — however strong their reputation — have published *almost no content* that directly answers the questions buyers actually ask. So they are not "ranked low." They are *structurally uncitable.***

We call the share of a company's web presence that an engine can actually use to answer a buyer question its *citation-candidate density*. For many companies it is close to zero.

An answer engine does not retrieve your homepage and admire it. When a buyer asks a question, the system expands that question into several narrower ones, retrieves passages that answer each, and assembles a response from what it can corroborate across sources.

01 – FAN-OUT →

One buyer question becomes **5–11** narrower sub-questions on Google's AI Mode; **2–3** on ChatGPT, and growing.<sup>1</sup>

02 – RETRIEVE →

For each sub-question, the engine pulls **passages it could quote** — a paragraph, a table, a spec, an FAQ entry it can lift and attribute.

03 – CORROBORATE

It assembles the answer only from what it can verify across sources. A page of brand adjectives offers **nothing to lift.**

At each step, the system is looking for a passage that answers the sub-question. A page that states, plainly, the requirements / the process / the price / the standard offers a great deal. The company with the answer-shaped content is the one more likely to surface; the category leader with the beautiful, empty homepage — its reputation notwithstanding — often does not.

This is also why brand strength translates so unreliably. Across a large measured sample, unstructured brand mentions correlated with AI visibility *roughly three times more strongly* than backlinks, while traditional domain authority predicted very little of it.<sup>2</sup> The signals that win search and the signals that win citation increasingly diverge.

§ 03 – DIVERGENCE

# There is no single "AI opinion" of your company *to win.*

Ask the same buyer question across five engines and you will routinely get five different shortlists — different names, different sources, sometimes a refusal to answer at all. This is not noise to be averaged away. It is the structure.

## In sampled prompts, some category-leading firms were named by *none* of the five engines.

OBSERVED ACROSS SAMPLED BUYER-INTENT PROMPTS · 2026

# 11%

CROSS-ENGINE CITATION OVERLAP

In measured comparisons, only about **one in nine** cited sources was shared between two leading engines answering the same query. A company can be the default recommendation on one engine and entirely absent on another — often for reasons that have less to do with quality than with what each system can most readily corroborate.<sup>3</sup>

300 observed answers · 60 prompts · 5 engines

SAMPLED PROMPT · WEB-GROUNDED ANSWER · 2026

▲ 400+ LOCATIONS · NOT NAMED

### “best commercial kitchen maintenance vendor for multi-location restaurants”

The answer named two national vendors and a regional specialist — and described their NFPA-96 compliance coverage in detail — drawing on a vendor's standards page, a trade directory, and a discussion thread.

- vendor-a.com / nfpa-96-compliance cited x3
- foodservice-directory.org / hood-cleaning cited x2
- community thread · restaurateurs cited x1

A regional operator with 400+ locations was not named — its site carried no machine-readable compliance, standards, or service-area content for the model to lift.

A single observed answer, anonymized and illustrative of patterns seen across sampled prompts. The omitted operator is not smaller or worse — it is harder to corroborate. *The coverage existed; citable evidence of it did not.*

**EXHIBIT A Same question, five realities.**

ILLUSTRATIVE · SAMPLED

NAMED IN THE ANSWER?	CHATGPT	CLAUDE	GEMINI	PERPLEXITY	GROK
<b>Firm A</b>	✓	✓	✓	✓	×
<b>Firm B</b>	×	✓	×	✓	×
<b>Firm C</b>	✓	×	×	✓	✓
<b>Your firm</b>	×	×	✓	×	×

Whether each company was named, by engine, for one buyer prompt asked the same day. Illustrative of divergence observed across sampled prompts; names anonymized. No company is named by all five — and the firm in question appears on a single engine and is absent from four.

**EXHIBIT B What each engine leans on.**

OBSERVED TENDENCIES

ENGINE	WHAT IT TENDS TO LEAN ON	EFFECT ON YOU
<b>O ChatGPT</b> <small>OpenAI · GPT-4o</small>	Encyclopedic and community sources; well-corroborated entities.	<i>Named</i> if you're widely, consistently described.
<b>P Perplexity</b> <small>Sonar Pro</small>	Live retrieval, discussion forums, recent pages.	<i>Named</i> if fresh third-party coverage exists.
<b>G Gemini</b> <small>Google · 2.5 Pro</small>	Google's index and its own query fan-out.	Tracks closely to what's structured and crawlable.
<b>A Claude</b> <small>Anthropic · Sonnet 4.6</small>	Cautious corroboration; declines when unsure.	<i>Omitted</i> rather than guessed at.
<b>X Grok</b> <small>xAI · Grok 3</small>	Real-time social signal and current chatter.	<i>Named</i> if there's live conversation about you.

Characterizes tendencies observed across sampled prompts; engine behavior is non-stationary and changes frequently. The point is not the specific lean of any one engine but that the five disagree — so "are we visible in AI?" has five answers, not one.

§ 04 – WHAT CITABLE LOOKS LIKE

# The difference between invisible and citable is usually *visible on sight*.

The fastest way to understand citation-candidate density is to look at two versions of the same page — one written for a brand impression, one written to answer a question. An engine can do almost nothing with the first and a great deal with the second.

EXHIBIT C Facilities & multi-location services.

WEAK VS. CITABLE SURFACE

WEAK SURFACE · UNCITABLE

A

"Trusted kitchen maintenance solutions for growing restaurants."

- ✗ No question is answered
- ✗ No facts to extract or attribute
- ✗ No location, no standard, no specifics

CITABLE SURFACE

B

"NFPA-96 hood-cleaning requirements for California restaurants."

- ✓ Inspection-interval table by appliance type
- ✓ Fire-marshall compliance checklist
- ✓ Vendor licensing & insurance requirements
- ✓ Typical price ranges · service-area list · FAQ

The strong surface gives an engine a dozen passages it can lift to answer real buyer sub-questions; the weak surface gives it none. Same business — different odds of being named.

EXHIBIT D B2B software.

WEAK VS. CITABLE SURFACE

WEAK SURFACE · UNCITABLE

A

"The modern platform for high-performing teams."

- ✗ Category is unstated
- ✗ No comparison, pricing, or specifics
- ✗ Nothing a shortlist query can match on

CITABLE SURFACE

B

"[Category] software: pricing, integrations, and SOC 2 status."

- ✓ Honest comparison table vs. named alternatives
- ✓ Pricing tiers and limits, stated in text
- ✓ Security / compliance page · migration steps
- ✓ Use-case pages matching real buyer phrasing

In software categories, engines assemble shortlists from review platforms, documentation, and community discussion. A vendor whose own comparison pages aren't answer-shaped often finds those pages cited *for the competitor instead*.

EXHIBIT E Professional services.

WEAK VS. CITABLE SURFACE

WEAK SURFACE · UNCITABLE

A

**"Trusted advisors committed to our clients' success."**

- ✗ Discipline and situation unstated
- ✗ Credentials and outcomes implied, not stated
- ✗ Authority lives offline, invisible to a model

CITABLE SURFACE

B

**"[Service] for [situation]: process, timeline, and fee structure."**

- ✓ Named practitioners with verifiable credentials
- ✓ Jurisdictions / sectors served, stated plainly
- ✓ Process, typical timeline, fee basis
- ✓ Question-shaped FAQ matching buyer intent

Reputation built through relationships does not corroborate inside a model. The citable surface restates offline authority as facts a system can find, read, and attribute.

EXHIBIT F Financial advisory.

WEAK VS. CITABLE SURFACE

WEAK SURFACE · UNCITABLE

A

**"Personalized wealth management you can trust."**

- ✗ No client type, minimum, or specialty
- ✗ Fees, fiduciary status, credentials unstated
- ✗ Nothing a model can safely corroborate

CITABLE SURFACE

B

**"Fee-only fiduciary planning for [situation]: minimums, credentials, process."**

- ✓ Fee model and account minimums, stated plainly
- ✓ CFP / fiduciary status, with ADV registry links
- ✓ Who it's for · specialties · onboarding process
- ✓ Compliance-aware FAQ matching buyer questions

A YMYL field where engines lean on earned media and registries and apply a higher trust gate. The citable surface stays within SEC Marketing Rule / FINRA 2210 — but states verifiable facts the model can corroborate rather than adjectives it cannot.

§ 05 – PATTERNS BY INDUSTRY

# The mechanism is constant.

## *Where authority lives is not.*

Across seven industries we sampled, the questions buyers ask and the surfaces engines reach for differ markedly. Directional, not exhaustive. Read the one that is yours.

EXHIBIT G Where authority lives, by industry.

OBSERVED LEANINGS · SAMPLED

	DIRECTORIES	REVIEWS	DOCS / SPECS	EARNED MEDIA	INSTITUTIONAL
Professional services	■	□	□	■	□
B2B software	□	■	■	□	□
Healthcare	■	□	□	□	■
Financial services	■	□	□	■	□
Manufacturing	□	□	■	□	□
Local services	■	■	□	□	□
Ecommerce / DTC	■	■	□	□	□

■ Primary surface engines lean on   ■ Strong secondary   □ Seldom decisive

Where each industry's verifiable authority tends to live, across sampled prompts. Finance and professional services skew to earned media and registries;<sup>4</sup> manufacturing to specifications; healthcare to institutional sources. Directional, not a measurement.

05.1	BUYERS ASK	"Best [discipline] firm for [situation / jurisdiction]" — pre-vetting before they ever make contact.
<b>Professional services</b>	WHO'S ABSENT	Firms whose authority is relationship- and referral-based, with little published, attributable substance.
CONSULTING · LAW · ACCOUNTING	IMPLICATION	Offline reputation does not translate. The firm everyone in the room knows can be the firm the model has <i>never heard of</i> .
05.2	BUYERS ASK	"Best [category] software for [use case / company size]" and "[A] vs [B]."
<b>B2B software &amp; technology</b>	WHO'S ABSENT	Vendors with thin docs and adjective-heavy marketing; their own comparison pages get cited for rivals.
VENDORS · PLATFORMS · DEV TOOLS	IMPLICATION	The review and community layer <i>is</i> the surface. Your site is one input among several you don't control.

<p>05.3</p> <p><b>Healthcare &amp; clinical</b></p> <p>HOSPITALS · SPECIALISTS · CLINICS</p>	<p>BUYERS ASK</p> <p>WHO'S ABSENT</p> <p>IMPLICATION</p>	<p>"Best [specialty] in [city]," "is [treatment] safe," "[condition] options" — high-stakes, health-sensitive queries.</p> <p>Independent practices without structured, authoritative, clearly-credentialed content.</p> <p><b>An elevated trust gate. Engines are visibly more conservative; unverified sources are simply not used.</b></p>
<p>05.4</p> <p><b>Financial services</b></p> <p>ADVISORS · PLANNERS · INSTITUTIONAL</p>	<p>BUYERS ASK</p> <p>WHO'S ABSENT</p> <p>IMPLICATION</p>	<p>"Best [advisor type] for [situation]," "is [firm] reputable," fee and fiduciary questions.</p> <p>Advisors who rely on referral and relationship, and whose compliance constraints limit what they publish.</p> <p><b>High trust gate, narrow room to maneuver. What you can say is bounded by SEC Marketing Rule and FINRA 2210 — <i>corroboration matters more than claims.</i></b></p>
<p>05.5</p> <p><b>Manufacturing &amp; industrial</b></p> <p>SUPPLIERS · FABRICATORS · OEMS</p>	<p>BUYERS ASK</p> <p>WHO'S ABSENT</p> <p>IMPLICATION</p>	<p>"Supplier for [spec / standard / tolerance]," "[material] for [application]," capability and certification queries.</p> <p>Firms with brochure-ware sites and no machine-readable specifications or certifications.</p> <p><b>Answer surfaces are spec-shaped. If your capabilities aren't stated as <i>data</i>, an engine can't match a requirement to you.</b></p>
<p>05.6</p> <p><b>Multi-location &amp; local services</b></p> <p>CONTRACTORS · RESTAURANTS · HOME</p>	<p>BUYERS ASK</p> <p>WHO'S ABSENT</p> <p>IMPLICATION</p>	<p>Niche + locality: "[specific service] in [neighborhood]," "[standard] requirements in [state]," emergency and compliance queries.</p> <p>Operators with generic marketing pages and no specific, local, answer-shaped content.</p> <p><b>Specificity wins. The broad term is unwinnable; the niche-plus-locality question is open — and most competitors haven't answered it.</b></p>
<p>05.7</p> <p><b>Ecommerce &amp; DTC brands</b></p> <p>DIRECT-TO-CONSUMER · MERCHANTS</p>	<p>BUYERS ASK</p> <p>WHO'S ABSENT</p> <p>IMPLICATION</p>	<p>"Best [product] for [need / budget / constraint]," "[product] alternatives," comparison and recommendation queries.</p> <p>Brands whose category and product pages aren't structured as answers; aggregators speak for them.</p> <p><b>Third parties own the consideration set. <i>The brand is described by the listicle, not by the brand.</i></b></p>

Across all seven: engines tend to name what they can *verify*, not necessarily what is best. The industries differ only in *where* the verifiable sources live, and *which* questions buyers bring.

# Six questions to ask *about your own surface.*

Not a score — a way to feel the exposure for yourself before measuring it properly.

Read each against your own site, honestly.

---

**01** Does a single page on your site answer, in plain words, the exact question a buyer would type — or does it describe how committed you are?

---

**02** Are your differentiators stated as  
*verifiable facts* — standards, numbers, credentials, jurisdictions — or as adjectives?

---

**03** Could a system extract a complete answer from one page without visiting five others to piece it together?

---

**04** Do independent third parties — directories, reviews, earned media — corroborate what your site claims about you?

---

**05** Is the relevant content reachable and parseable by AI crawlers, or does it only exist behind scripts and interactions?

---

**06** When you actually ask the five engines your buyers' questions — does your category's shortlist  
*include you at all?*

---

If several answers are uncomfortable, that discomfort is the finding. The questions are easy to ask and hard to answer rigorously — which is the work the Diagnostic does.

## § 07 – WHAT ORGANIZATIONS ARE DOING

# The organizations responding well *share a few habits.*

Not tactics — patterns observed among companies that read well in AI-mediated discovery. None of them is a ranking trick; each one makes a firm easier to find, read, and corroborate.

---

01	<b>Publish question-shaped pages</b>	One page that answers, end to end, a real buyer question — not a brochure about how committed they are.
02	<b>Expose expertise as fact</b>	Credentials, standards, numbers, jurisdictions, methods — stated plainly, where a model can read and attribute them.
03	<b>Publish evaluation content</b>	Honest comparison and selection material on their own terms, so the answer isn't assembled entirely from third parties.
04	<b>Build corroboration</b>	Third-party references — directories, earned media, registries — that let an engine verify a claim rather than take the site's word.
05	<b>Reduce homepage dependence</b>	They stop relying on a single, adjective-heavy homepage to carry meaning the rest of the web never confirms.

---

The throughline: restate what is true about the business in a form a machine can find, read, and corroborate. The hard part is knowing *which* gaps matter most — which is what a diagnosis establishes.


# From "is this real?" to "what does it say about us?"

This report establishes a phenomenon: AI-mediated discovery is real, structurally selective, and already shaping consideration sets in your industry. The uncomfortable implication for established firms is that a strong offline reputation is not protection. The next question is specific to you — and it is answerable only by running your buyers' actual questions across the engines and reading every source they cite.

## WHAT THE DIAGNOSTIC DELIVERS

### A scored reading of your retrieval surface.

Your real buyer-question set, run across all five engines — every cited URL scored on Content, Retrieval, and Trust, with the specific gaps named and a sequenced plan to close them.



## The AI Answerability Diagnostic.

Your real buyer-question set, run across ChatGPT, Claude, Gemini, Perplexity, and Grok, with every cited URL scored on **Content, Retrieval, and Trust** — and a sequenced reading of where you are absent and why.

[answerability.ai/order](https://answerability.ai/order) →

REFERENCES & NOTES

---

- 1 **Query fan-out behavior.** Google AI Mode has been observed expanding a single query into roughly 5–11 sub-queries and ChatGPT into ~2–3, with expansion length growing over a period of months. Retrieved passages are merged across sub-queries before an answer is composed. Answerability.ai, "*You ask one question. AI asks twelve more,*" /insights/how-ai-search-fans-out.
  - 2 **Brand mentions vs. backlinks.** An analysis of a large brand sample found unstructured brand mentions correlated with AI visibility several times more strongly than backlinks, while conventional domain authority explained little of the variation. Ahrefs, study of ~75,000 brands, December 2025.
  - 3 **Cross-engine citation overlap.** In measured comparisons, only about 11% of cited sources were shared between two leading engines answering the same query. Answerability.ai, "*Why AI recommends some companies and ignores others,*" /insights/why-ai-recommends-some-companies.
  - 4 **Financial-services citation mix.** In sampled finance queries, the large majority of cited sources were earned / third-party media rather than firms' own properties — consistent with an elevated trust gate on health- and money-sensitive (YMYL) topics. Answerability.ai, "*How AI picks financial advisors,*" /insights/ai-search-for-financial-advisors.
- 

CITE AS

---

Answerability.ai (2026). *AI-Mediated Discovery Across Industries: A Field Report*. Field Report 2026-01.

Observational and directional. Findings characterize sampled behavior in non-stationary systems and are not guarantees of future AI behavior. Confidential; prepared for the recipient.