AUTHORITY ENGINE · RESEARCH

2026



A RESEARCH REPORT ON AI-DRIVEN DISCOVERY

The AI Visibility Gap

Why industry leaders are losing deals to lesser-known competitors — and what the machines see that buyers don't.

Dr. Patrick T. McAvoy

FOUNDER · AUTHORITY ENGINE

AI AUTHORITY ENGINEERING

EXECUTIVE SUMMARY

When excellence goes unseen.

High-performing brands with superior products, established reputations, and strong customer bases are systematically losing visibility in AI-driven discovery systems. This research reveals that **96% of B2B companies are invisible** during early-stage AI buyer discovery[1] — appearing only once prospects already know their name, and absent from the AI-generated answers that increasingly shape vendor shortlists.

Our comparative analysis of established brands (like Fidelity Investments) versus smaller, authority-optimized competitors reveals a disturbing pattern: AI discovery systems favor businesses that signal authority through machine-readable content over those that merely demonstrate operational excellence. This creates a fundamental market inequity — buyers using AI research tools receive incomplete, often misleading recommendations about which companies can best serve their needs.

> *The best company no longer automatically surfaces. The most legible one does.*

The pages that follow document the gap with paired audits and live prompt testing, isolate the authority signals that move the needle, and lay out what high-performing businesses must do to be seen by the systems now mediating their markets.

RESEARCH NOTE

The 96% headline is based on 2X's 2026 sample of 70 B2B companies, not a census of all B2B firms. Findings are directional and depend on the prompts used.

KEY FINDINGS

Five numbers that define the gap.

4.3%

of companies appear in early-stage AI buyer queries[2]

527%

year-over-year increase in AI-referred traffic[3]



3.2×

more often ChatGPT *mentions* brands than it *cites* them[4]



65%

of informational queries resolve with no website visit[5]



44%

of tested ChatGPT prompts contained no brand mentions [6]

Sources: 2X · Semrush · BrightEdge · Matt Britton · Trustmary. Full citations on the final page.

STATISTIC SCOPE

2X: 70 companies; Previsible: 19 GA4 properties. The 65% no-click statistic concerns general web search, while BrightEdge's 44% finding comes from analyzed ChatGPT prompts.

THE PROBLEM

The invisible visibility problem.

In many cases, high-performing brands are either not visible or actively misrepresented in AI-surfaced experiences. Businesses with large customer ecosystems and strong brand recognition are showing up online less prominently, less accurately, or not at all — compared to their actual market performance.

This disconnect becomes especially apparent when examining brands through the lens of AI authority signals. Companies that are literally big, well-trusted, and high-performing often carry surprisingly low "authority scores" and lack the machine-readable content necessary to prove their expertise to AI platforms.

Brands that consumers perceive as industry giants "surface small" when evaluated through AI and marketing-technology systems. The downstream consequence is severe: buyers using AI-surfaced answers to research options aren't seeing the full competitive landscape. Legacy players no longer automatically appear in AI-generated vendor lists simply because they're established category leaders.

What the machines actually evaluate

- Whether AI-readable, answer-focused content exists on solution pages
- Whether case studies are structured as discoverable proof rather than prose
- Whether executive expertise is explicit, attributed, and machine-legible
- Whether structured data lets crawlers understand real capabilities

THE FIDELITY CASE

When billion-dollar brands become invisible.

Recent work with a marketing agency serving Fidelity Investments revealed this paradox at scale. Fidelity operates one of the largest financial-services enterprises in America, managing trillions in assets with brand recognition most companies can only aspire to. Yet evaluated through an Authority Engine GEO audit, Fidelity scored significantly low across critical AI-discovery dimensions.

ILLUSTRATIVE EXAMPLES — NOT A CONTROLLED HEAD-TO-HEAD

ESTABLISHED FINANCIAL BRAND

Broad public recognition does not guarantee that every specialized service is accurately described in an AI response. This report discusses observations from buyer-prompt testing.

BOUTIQUE DIGITAL AGENCY

A smaller firm may have stronger, more specific content for its own category. These examples belong to different industries and cannot be ranked against each other.

No cross-industry percentage comparison is implied; raw audit scores were not supplied.

These examples suggest that conventional brand familiarity and service-specific AI discoverability can diverge. Because the organizations operate in different sectors, they do not constitute a matched test of relative performance. The broader question requires like-for-like audits with defined prompts and reproducible scoring.

> *Recognition is a human asset. Visibility is now a machine one — and they have come apart.*

RESEARCH METHODOLOGY

How we tested the gap.

We explored potential authority differences through audits and AI prompt testing of recognized brands and smaller digital-first firms. The examples presented are not necessarily in the same vertical. They illustrate patterns worth further study rather than a controlled, matched-category experiment. Scores are qualitative without published scoring rubrics or raw prompt logs.



To qualify as having legitimate external authority, smaller firms had to demonstrate real business results through publicly verifiable data — financial statements, news coverage, case studies, testimonials — validating competitive capability despite smaller scale. Where a brand demonstrates top-tier performance but AI fails to surface it or describes it poorly, we classified the instance:

A Misrepresentation

AI recognizes the brand but describes it poorly, inaccurately, or generically.

B Omission

The brand does not appear in relevant queries at all, despite leading the category.

AUDIT RESULTS · RECOGNIZABLE BRANDS

Fidelity Investments: the authority audit.

This analysis focuses not on general awareness — everyone has heard of Fidelity — but on whether AI can easily identify what the company does well, what it claims thought leadership on, who leads it, and what proof exists of customer outcomes. While "wealth management" returns recognition, we examined the micro-solutions and specific buyer journeys where AI must distinguish between competing providers.

What we discovered

Content structure	Quality content existed, but was written in "marketing speak" rather than the Q&A, searchable formats that answer buyer questions with crawlable structure.
Proof infrastructure	Business results and case studies were mentioned but not prominently linked or structured with logos in ways crawlers recognize as testimonials.
Executive authority	Leader names appeared on the site but weren't connected to subject-matter expertise or thought-leadership content.

In short: a category-defining institution that, to a machine, looked quiet exactly where it should have looked loudest.

AI PROMPT TESTING

Fidelity's discovery performance.

We tested visibility using natural buyer questions — "Who offers [X service]?", "Who can do [Y] for [Z customers]?" — with one critical constraint: **we never pre-populated company names.** AI could only return businesses it had previously encountered through crawling or training data.

> *In multiple cases, Fidelity was completely omitted from service-specific prompts — despite being one of the largest providers of those exact services.*

When Fidelity did appear, descriptions tended to fail in one of three ways:

Outdated

Not reflecting current offerings.

Incomplete

Failing to communicate full capabilities.

Generic

Not highlighting expertise where Fidelity leads.

Meanwhile, smaller brands that invested in authority optimization around those specific niches appeared more frequently — and were described in hyper-specific language emphasizing their expertise. We validated these as communication gaps, not capability gaps, by comparing AI representations against publicly verifiable outcomes.

CONTROL GROUP · AUTHORITY-OPTIMIZED

Brightlogic: what strong signals look like.

To baseline how AI treats brands with strong authority signals, we tested a boutique firm that built its footprint deliberately. Founded in 2001, Brightlogic provides SEO and digital-marketing services, has been featured by Search Engine Land, and contributes expert commentary on industry news. By every measure, a digitally mature business — but no household name.

Prompt: "Who does SEO for small businesses?" **RESULT #1**

"BrightLogic — Seattle search-engine optimization and digital-marketing firm. Good for: small businesses who need an SEO partner to get found online."

Prompt: "Who offers SEO services?" **RESULT #2**

"BrightLogic — Seattle search-engine optimization and digital-marketing firm." Accurate services, clear positioning.

Prompt: "Find local SEO companies" **TOP RESULT**

"Brightlogic is a Seattle SEO and digital-marketing agency providing custom programs for small-to-medium businesses that want to grow revenue through digital channels."

KEY FINDINGS

Three patterns, every category.

Brightlogic consistently surfaced with accurate, helpful descriptions — not because it's a household name or an Authority Engine client, but because of the authority work built around its presence. Across every category we researched, the same three patterns held.

01 Name-drop dependency

Human-recognizable leaders were omitted or surfaced inconsistently unless users literally named them in the prompt — especially in finance and professional services.

02 Open-ended query failure

For "who are the best firms for [X]?", smaller firms with highly optimized sites appeared as category experts — even where larger firms perform most of the actual work.

03 Authority-signal advantage

Firms with strong authority scores surfaced far more often in service-specific queries — sometimes prioritized over larger firms operating in identical areas.

The fundamental problem: in their current state, AI discovery systems lack the nuance to surface operational quality. They reward what is easy for a machine to read — not what is true in the market.

MARKET IMPLICATIONS

Discovery is now fundamentally inequitable.

AI-driven discovery does not present a balanced view of high-performing businesses. Many of the largest, most reputable brands are absent from the vendors surfacing at the top of buyer research streams. Research from 2X confirms it at scale: **96% of B2B companies are invisible** in AI-driven buyer discovery, appearing only in late-stage queries where the prospect already knows the name.[7]

Some of the strongest operators have substantial footprints across the web that simply don't point back to them, because they never invested in being machine-readable. Organizations missing from AI consideration sets face what we call a **structural blind spot** in buyer research — a flaw in the comfortable assumption that "the best companies always surface to the top."

The Fidelity paradox

For companies like Fidelity, the risk isn't total invisibility — it's that key parts of their value proposition and differentiation won't surface, or won't align with how AI mediates buying. With **65% of informational queries now resolving with no website visit**,[8] visibility in the answer matters more than the click. Brands absent from zero-click answers effectively don't exist to the researcher relying on AI.

SCOPE OF EVIDENCE

2X examined 70 B2B firms; its 95.7% figure applies to that tested group and prompt set. The cited zero-click percentage concerns web-search behavior and should not be read as an AI-only measure.

BROADER CONSEQUENCES

Beyond customer acquisition.

If machines struggle to identify great operators, markets risk overvaluing companies that excel at "speaking AI's language" while undervaluing those that excel at actual delivery. The visibility gap reaches well past lead generation:

- **Investment decisions** shaped by AI-perceived market leadership
- **M&A valuations** influenced by AI-surfaced competitive positioning
- **Talent acquisition** as professionals research employer reputations
- **Partnerships** where AI vets potential collaborators

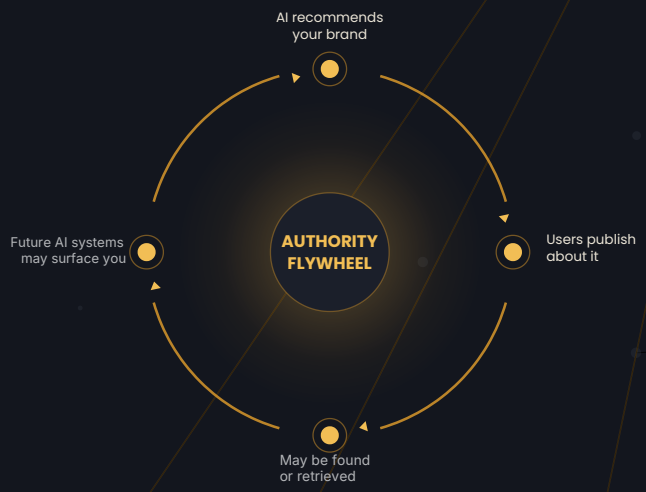
Website visits alone cannot reveal how often AI assistants mention or cite a company. BrightEdge found that ChatGPT mentioned brands about 3.2 times as often as it cited them, while 44% of tested prompts returned no brand mentions. [9] Monitoring brand mentions and citations alongside traffic helps expose this visibility gap. [10]

> *Once AI visibility informs who is "winning," markets can begin to reward gaming discovery over operational excellence.*

THE COMPOUNDING EFFECT

The authority gap widens every month.

Authority signals can compound as a brand earns relevant mentions, citations, and independent discussions. Search and AI systems may discover or retrieve those sources. That does not mean each new mention becomes model training data or guarantees future recommendations.



Brands inside the loop accelerate. Brands outside it stay invisible — and the distance between them grows with every model release.

This is why "do the work and wait" fails. The flywheel only spins for organizations that deliberately feed it machine-legible authority.

WHAT HIGH-PERFORMING BUSINESSES MUST DO

Engineer the authority you've already earned.

If you want to show up for your ideal buyers, you must manage visibility and authority as a discipline independent of core operations. Doing your job better than anyone else is no longer sufficient to ensure buyers discover you when AI filters and narrows their consideration set.

- Understand your AI exposure gaps — where capability isn't evident due to weak external profiles
- Adjust content to answer buyer questions with demonstrable authority
- Select and structure proof points from your portfolio so machines can read them
- Make executives and expertise visible where scoring models look for authority signals

Engineering authority isn't about crafting lies for AI to believe — it's about ensuring AI can see your real value clearly.

---

3 yrs

until AI platforms are expected to drive more visits than traditional search[12]

527%

year-over-year growth in AI-referred traffic, already[13]

75M

daily active users on Google's AI Mode — 100M+ monthly[14]

FORECAST NOTE

Projected traffic crossover dates are scenarios, not established outcomes. The 527% statistic is based on 19 GA4 properties tracked by Previsible in 2024–2025.

THE URGENCY IS CLEAR

The question isn't whether to close the gap.

It's whether you'll close it before your competitors do.

Organizations that fail to optimize for AI discovery face systematic exclusion from buyer consideration sets — regardless of operational excellence. The brands that act now will define their categories in the machine-mediated market. The rest will wonder where their pipeline went.

Authority Engine

AI AUTHORITY ENGINEERING · DR. PATRICK T. MCAVOY

CITATIONS & SOURCES

The evidence base.

References include first-party research and third-party commentary; estimates are labeled.
Click a numbered source to read it online. Unpublished methodology is described in the report.

01 2X (2026). "2026 AI Visibility Index." Study of 70 B2B companies: 95.7% appeared mainly in branded/later-stage prompts.

02 **2X (2026).** Ibid. — share of companies appearing in early-stage AI buyer queries.

03 Previsible (2025). "AI Traffic Report" (19 GA4 properties, Jan–May year-on-year), summarized by Semrush (2025).

04 **BrightEdge (2025).** "ChatGPT Brand Mentions vs. Citations Research" — weekly AI search insights.

05 **Britton, Matt (2026).** "AI Search Trends 2026: Visibility Beats Clicks."

06 BrightEdge (2025). ChatGPT prompt analysis: 44% contained no brand mentions.

07 2X (2026). Op. cit. — 95.7% in its 70-company sample.

08 **Britton, Matt (2026).** Op. cit. — 65% of informational queries resolve with no website visit.

09 BrightEdge (2025). Brand mentions and citations: separate visibility signals.

10 **BrightEdge (2025).** Op. cit. — mentions vs. citations; zero-brand-mention prompts.

11 Trustmary (2026). Commentary on online mentions, not demonstrated model-training causality.

12 **Semrush (2025).** Op. cit. — AI platforms projected to surpass traditional search within three years.

13 **Semrush (2025).** Ibid. — 527% year-over-year growth in AI-referred traffic.

14 Google/Alphabet (2025). AI Mode audience: 100M+ monthly users (Q2); over 75M daily users (Q3).

Scope: 2X n=70; cross-sector examples are illustrative; no AI-sales causality established.

Authority Engine

AI Authority Engineering — research, infrastructure & pipeline.

© 2026 Dr. Patrick T. McAvoy | Authority Engine