The Hoffman Agency

A GEDI report

Understanding AI discovery in B2B technology markets

A GEDI perspective on visibility, influence, and generative search

Executive summary

section 01

AI is not replacing the old influence system. It is compressing it into a single answer

Generative AI is changing how B2B technology buyers discover and evaluate vendors. Tools such as ChatGPT, Claude, Gemini, and Perplexity are becoming part of early-stage research, helping buyers compare providers, build longlists, and make sense of complex markets.

This report explores how that works in practice. Using GEDI, The Hoffman Agency's in-house AI visibility tool, we adopted the personas of our clients' customers across four areas of technology specialism: telecoms, cybersecurity, enterprise technology, and fintech. We then looked at how vendors were represented in AI-generated discovery, how they were described, and which sources large language models (LLMs) relied on to build their answers.

The main conclusion is clear.

Earned-media share of citations

22.6%	17.7%	13.6%	9.9–12.6%
Telecoms	Cybersecurity	Enterprise technology	Fintech · tokenization → embedded finance → financial crime

Executive summary

section 01

AI discovery is not inventing a new system of influence from scratch. It is reorganising the one that already exists. Trusted third-party sources still play a defining role in how AI systems surface and frame brands. But no single channel can improve LLM visibility on its own. Strong visibility over time comes when earned, owned, analyst, and wider content activity all reinforce the same narrative.

Sector by sector, earned media accounted for 22.6 percent of citations in telecoms, 17.7 percent in cybersecurity, and 13.6 percent in enterprise technology. In fintech, we tested multiple issue areas to check whether the initial findings were being distorted by a niche subject. Tokenization delivered 9.9 percent earned-media share, embedded finance 10.8 percent, and financial crime 12.6 percent. So, the more widely discussed and debated the topic in fintech media, the more influential earned media becomes for fintech vendors with LLMs. Across all three fintech themes, vendor-owned and vendor-adjacent sources remained highly influential in shaping AI outputs.

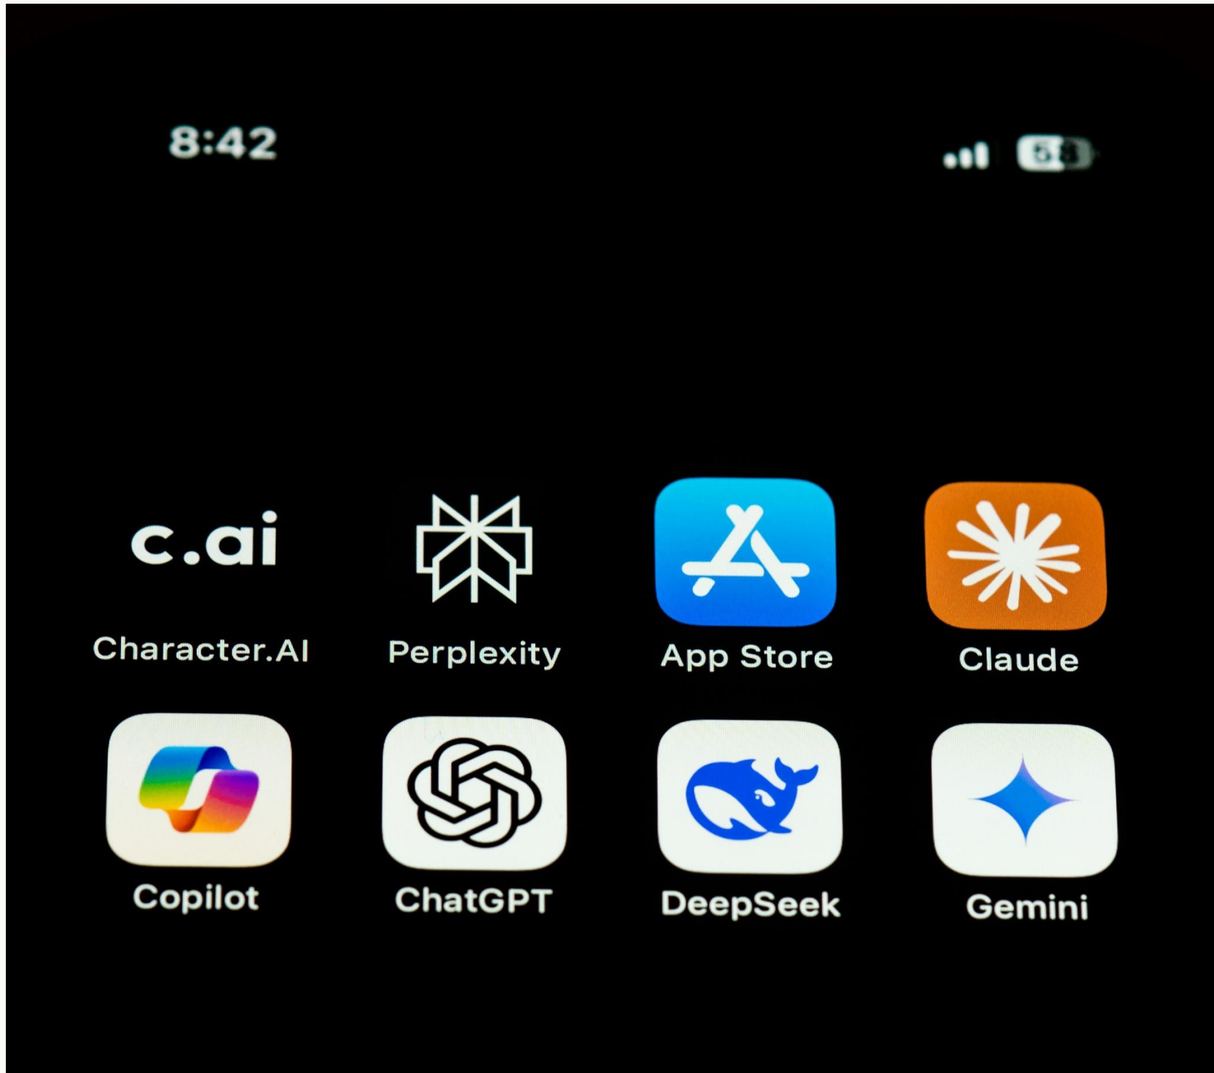

The shift

section 02

From buyer behaviour to AI discovery

The fundamentals of influence have not changed. The interface has.

The rise of AI in vendor discovery reflects a wider change in how people gather information. Buyers still care about credibility, third-party validation, expert opinion, and evidence that a supplier is trusted by the wider market.
What is changing is the interface. Instead of visiting many sources one by one, people are increasingly asking AI systems to pull those signals together for them.

This is consistent with the findings of our buyer insight guides investigating B2B technology buying across market sectors. These clearly explain how our clients' customers make technology purchasing decisions.

GEDI shows how those decisions are increasingly shaped, by LLMs, before buyers ever search for a vendor, through the sources that AI systems choose to trust and synthesise.

AI is not replacing the old influence system. It is compressing it into a single answer. When a buyer asks an LLM for recommendations, the answer is built from the same wider ecosystem that has always shaped perception: trade media, industry analysts, institutions, vendor content, and other trusted sources.

That makes the source layer underneath the answer more important, not less. GEDI helps surface that hidden layer by collecting answers, citations, and source patterns across multiple LLMs and repeated prompting.

Before

Buyer visits many sources one by one — trade media, analysts, peers, vendor sites — and assembles a view.

Now

Buyer asks an LLM to pull those signals together. The same ecosystem still shapes the answer — it is compressed into a single response.

What GEDI reveals

section 03

AI visibility is shaped by structure

Not all sectors had the same starting conditions. Not all issue areas produced the same amount of source material. And not all engines relied on the same sources in the same way.

A brand's visibility in AI is never just about whether it has content. It is about whether it is present in the right external ecosystem for the specific issue a buyer is asking about.

Across all four sectors, the same broad pattern appeared. The models pulled from a mix of earned media, industry analyst or institutional sources, and owned content — but those source types did not all play the same role.

This is where the integrated marketing playbook helps sharpen the picture. Its central point is that search has given way to synthesis, and that brands need an AI-first content strategy spanning category definition, attribute ownership, and structured, machine-readable content.

In other words, PR can help AI notice you, but if your site, social presence, thought leadership, and content architecture are not aligned to the same categories and language, AI may still struggle to understand you well enough to recommend you consistently.

Three roles in the answer set

Earned media — the discovery layer

Specialist trade titles help the models compare options and frame market narratives, especially in sectors with strong editorial ecosystems.

Analyst & institutional sources — depth

Add weight in more complex or emerging issue areas — formal expertise and technical authority that LLMs lean on.

Owned content — the source of truth

Reinforces the narrative, adds technical detail, and gives AI systems a usable source of truth to verify and reuse.

The role of earned media

section 04

Earned media remains important — but its role varies by issue

The Hoffman Agency's latest B2B tech buyer insight guides in telecoms, fintech, cybersecurity and enterprise AI, highlight a broader change in buyer behaviour. Buyers are increasingly using generative AI tools to support early research and comparison. At the same time, the fundamentals of influence remain unchanged. They still rely on third-party validation to build trust and shape shortlists. Earned media is important across all four sectors — but the way models lean on it depends on the issue being asked about.

Telecoms

Strongest earned-media profile

The issue area tested sits inside a mature trade media ecosystem, giving the models a rich editorial base for explaining the category and comparing options.

Cybersecurity

Broader, more mixed sources

Alongside media titles, the models leaned on institutional and technical sources. AI systems are not only looking for commentary — they are also looking for signals of formal expertise and technical authority.

Enterprise technology

Somewhere in the middle

Meaningful trade and analyst input, but the source environment was more concentrated around a smaller group of highly visible players and publications.

Fintech

Lowest earned-media influence

But also the most useful test in the report — we expanded to embedded finance and financial crime to test whether topic maturity changed the pattern.

Earned-media share of citations

section 04

The more widely discussed the topic, the more influential earned media becomes

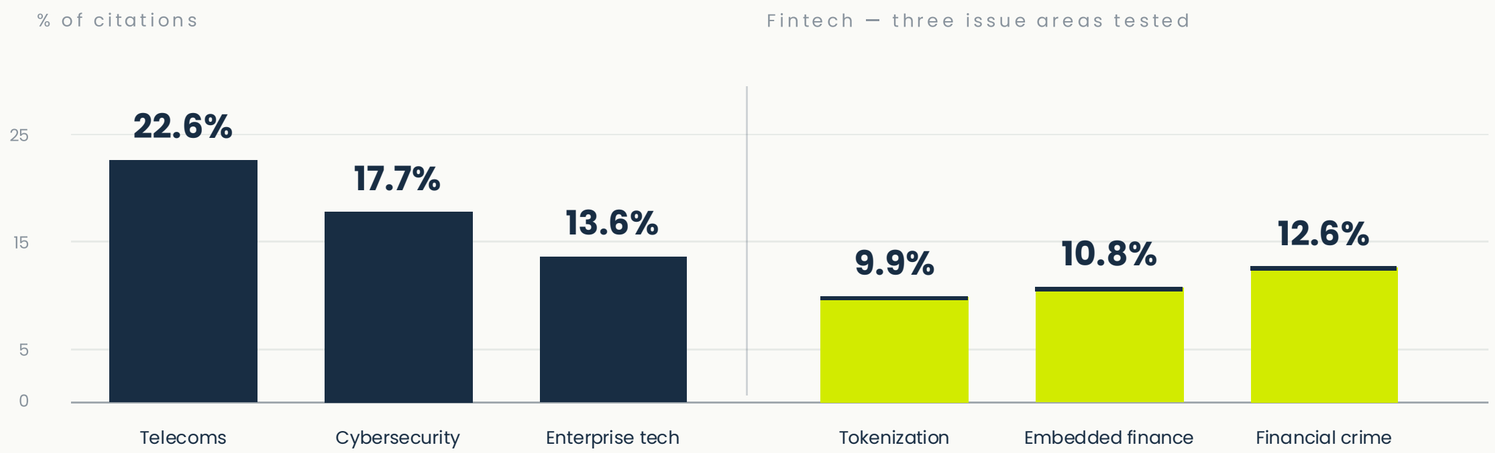

Source: GEDI study, The Hoffman Agency. Earned-media share of total citations across ChatGPT, Claude, Gemini, and Perplexity.

What the numbers say

Topic maturity does influence the degree of earned-media visibility in AI-generated discovery.

Across all three fintech themes, LLMs continued to rely heavily on vendor-owned and vendor-adto build their answers, with company domains often acting as the strongest sources of category explanation, product detail, and technical validation.jacent content.

Earned media remains important in fintech. As the subject matter became more widely discussed, earned sources were cited slightly more often and were more likely to recur across engines, especially in financial crime. This is important because repeated use across LLMs is a stronger signal that a source is helping shape AI discovery more broadly. In other words, earned media remains important in fintech, but it operates alongside a strong dependence on owned vendor content rather than displacing it.

So, the right conclusion is not that earned media is enough. It is that earned media plays a defining role in discovery and credibility, especially where specialist trade titles carry authority. But its impact is strongest when it is reinforced by aligned owned content and a broader integrated strategy.

What overlap analysis adds

section 05

The question is not what was cited most. It is what recurred across multiple models

Citation totals show which sources appear most often. Overlap shows which sources are used by more than one AI engine. That matters because repeated use across Claude, Gemini, ChatGPT, and Perplexity is a stronger sign that a source is shaping AI discovery more broadly.

Telecoms	**Specialist trade + major vendor-owned domains**	A mix of specialist trade media and major vendor-owned domains. The issue space is being interpreted through both editorial coverage and established market participants. *Light Reading · TelecomTV · Telecoms.com · Fierce Network · RCR Wireless · The Fast Mode*
Cybersecurity	**A blended trust base**	Specialist media domains alongside institutional and technical-authority sources, including government and non-profit domains. This points to a blended trust base, where trade reporting, formal expertise, and supplier-led knowledge all shape the answers. *NIST · CyberPress and specialist cybersecurity media*
Enterprise technology	**Concentrated source set**	Leaned more toward infrastructure providers, analyst-style sources, and a smaller set of business and technology media — which helps explain why the sector appeared more concentrated overall. *Data Centre Magazine · Data Center Knowledge · CIO · CRN · TechCrunch*
Fintech	**Shared vendor and vendor-adjacent domains**	Tokenization, embedded finance, and financial crime showed a familiar pattern. As the topics became broader, there was some improvement in the reuse of earned-media sources across engines, particularly in financial crime. In every case, shared vendor and vendor-adjacent domains remained central to the answer set. Earned media contributes external validation, while owned content still does much of the explanatory heavy lifting. *Data Centre Magazine · Data Center Knowledge · CIO · CRN · TechCrunch*

What the numbers say

Taken together, the overlap analysis shows that the sources most likely to shape AI discovery across engines are not one single source type.

Depending on the sector, they can include trade media, analyst or research sources, institutional resources, vendor websites, and company-authored explainers.

The key question is not simply what was cited most, but which sources were trusted often enough, and widely enough, to recur across multiple models?

Why topic maturity matters

section 06

Not all prompts were created equal — because not all issue areas are

Some of the topics tested are now well-established public conversations with rich specialist coverage. Others are still being explained through a smaller and less balanced source environment. That affects the source pool. And the source pool affects what the models know, how they answer, and how representative their answers are likely to be.

The findings should be treated as illustrative. The report is not measuring model behaviour in the abstract. It is measuring model behaviour in relation to different issue spaces — each with a different level of public documentation, specialist coverage, and source depth.

The fintech maturity curve

niche → widely discussed

Tokenization
9.9%
earned media

Embedded finance
10.8%
earned media

Financial crime
12.6%
earned media

A recognisable pattern — as topics became more established, earned-media influence improved slightly, but vendor-owned content remained one of the strongest drivers of discovery and explanation. This suggests the original tokenization finding was not an anomaly. It was an example of a broader fintech pattern.

The tokenization finding was not an anomaly. It was an example of a broader fintech pattern.

What this means for brands

section 07

Four practical lessons for the AI-discovery layer

Treat GEO as an integrated marketing challenge, not a channel tactic.

one

Invest in high-authority earned media.

Build visibility in the trade titles AI systems repeatedly trust — the editorial layer where models compare options and frame narratives.

two

Strengthen analyst and authority relations.

Industry analysts, associations, standards bodies and authority-led organisations help structure how LLMs understand a category — especially in complex or emerging areas.

three

Make owned content work harder.

It should not sit apart from earned coverage. Mirror the same language, categories, and proof points as earned coverage — so AI systems can connect third-party validation with a consistent source of truth on your channels.

four

Treat GEO as integrated marketing.

Category framing, ownership attribution, social reinforcement, structural content, and external validation all contribute to discoverability. No single channel builds strong visibility on its own.

Conclusion

section 08

Buyer reports show us how buyers make decisions. GEDI shows us how those decisions are increasingly shaped — before buyers ever reach a vendor

The emergence of AI as a discovery layer represents a significant shift in how B2B technology markets operate. But the core drivers of influence remain familiar.

What GEDI reveals is that visibility in AI-generated discovery is still built on external credibility. Trade media, industry analyst firms, and in some sectors institutional or industry bodies continue to play a defining role in shaping who is surfaced and how they are described.

Credibility only becomes stronger and more consistent when it is translated into a broader, integrated content ecosystem. Brands cannot rely on external validation alone — they need owned content that reflects and reinforces the same story.

Conclusion

section 08

The fintech findings underline that especially clearly. Even when we moved from niche to more widely discussed issue areas, the same broad LLM behaviour remained. Earned-media influence improved slightly, but owned vendor content continued to play a major role in helping AI systems explain the market, compare providers, and add technical detail. That means brands cannot rely on external validation alone. They need owned content that reflects and reinforces the same story.

Our B2B tech buyer insight reports show us how buyers make decisions. GEDI shows us how those decisions are increasingly shaped by LLMs before buyers ever reach a vendor. Hoffman's own integrated marketing playbook shows what brands need to do in response: align earned, owned, thought leadership, social, and structured content around the same categories, claims, and proof points.

Together, these define the new rules of visibility.

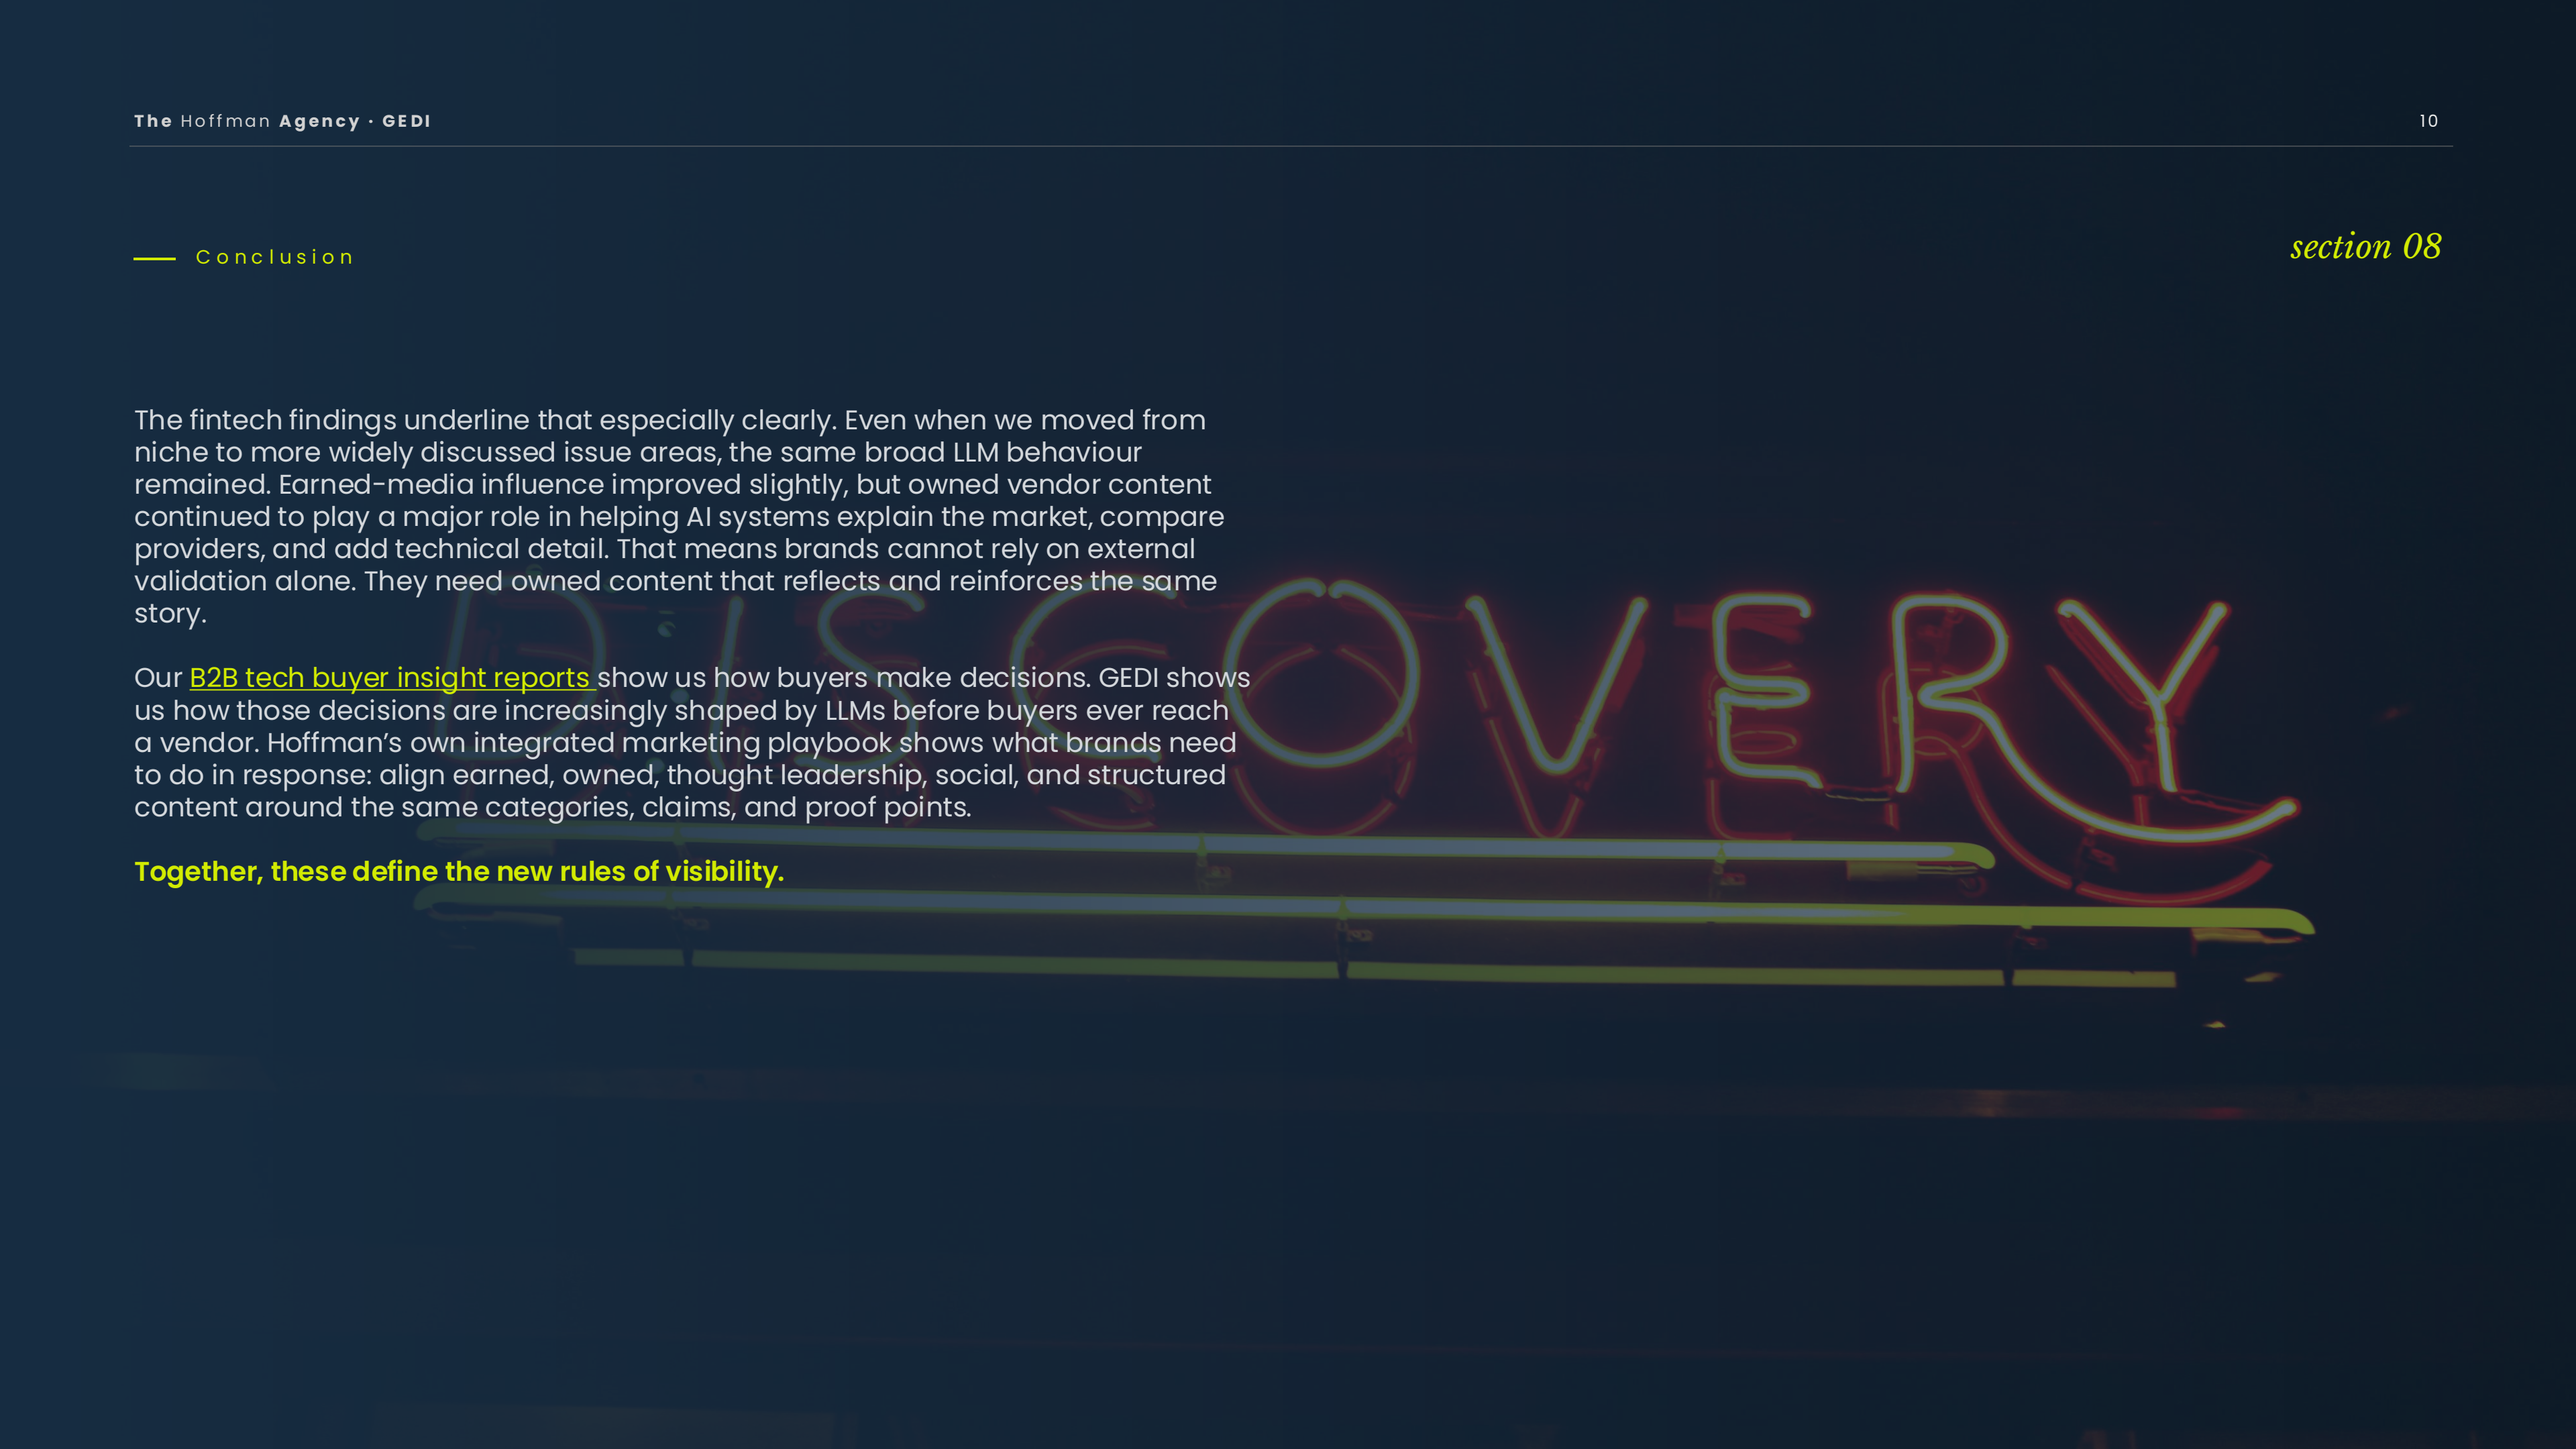

Methodology

Appendix

A structured GEDI study across four B2B technology sectors

GEDI is designed to run structured prompts across multiple AI engines, capture the resulting answers and citations, and then analyse the source patterns underneath those answers. The purpose of this study was not to rate the performance of specific vendors — it was to understand how brands show up in AI-generated discovery when we adopt the personas of our clients' customers and ask the kinds of questions those customers might realistically ask during early-stage research.

Sectors & issue areas tested

Telecoms	Network AI
Cybersecurity	Ransomware and quantum-related risk
Enterprise tech	AI and datacentres
Fintech	Tokenization → embedded finance → financial crime

AI engines analysed

ChatGPT **Claude** **Gemini** **Perplexity**

Note on scope

These issue areas were chosen because they are current, commercially important, and visible in market conversation. But they are not equally mature. Some have a deeper public record and a richer specialist media environment than others. That is why the report treats the findings as directional rather than absolute. In fintech, the wider testing was included specifically to check whether the original tokenization result reflected an unusually niche topic. The later embedded-finance and financial-crime analyses showed some improvement in earned-media influence, but not enough to change the broader pattern.

Prompt design

Appendix

Built around senior decision-maker personas with influence over vendor choice—discovery, comparison, fit, and validation

Examples of the kinds of prompts used include:

Telecoms

A senior telecom strategy leader asking what types of AI-driven network management or automation platforms should be considered in a complex multi-vendor environment.

Cybersecurity

A security leader asking what kinds of platforms should be considered to address both current ransomware threats and longer-term quantum-related security risks.

Enterprise technology

A senior infrastructure decision-maker asking whether a data centre and digital infrastructure provider would be a good fit for AI-scale workloads while meeting compliance and data-residency needs.

Fintech

A senior banking technology leader asking about solutions across tokenization, embedded finance, and financial crime — testing whether more widely discussed topics produce a meaningfully different source mix.

Data capture & classification

Appendix

GEDI captured the vendors mentioned, how those vendors were described, and the citations or source URLs returned by the engine

The citations were grouped and analysed by source type, engine, domain recurrence, and overlap across repeated sessions and models.

In the source analysis used for this report, earned media was measured explicitly by sector, alongside total citations, domain counts, engine mix, and top sources. Overlap analysis was then used to identify which domains repeated across multiple engines, helping distinguish between broad authority and one-engine concentration. The fintech analysis used this same structure across three separate issue areas so we could compare how topic maturity influenced source behaviour without changing the overall approach.

A note on scope and interpretation

This report is based on a focused set of prompts tied to specific issue areas and customer personas. It should be read as a structured view of how AI discovery works in these contexts, not as a universal model of all LLM behaviour. The findings are shaped by several factors: the wording of the prompt, the maturity of the topic, the amount of public source material available, and differences in how engines retrieve and weight sources.

Some topics in this study have a richer public conversation than others. That naturally gives the models more material to work with and can make the resulting picture look more stable or complete. That does not make the less mature categories unimportant. It simply means they need to be interpreted with more care.

From diagnosis to action

Can we help optimise your AI discovery?

If your brand is not showing up in the sources AI systems trust, you are less likely to appear in the answers your buyers see.

GEDI helps technology vendors understand how they are represented across generative AI platforms, which sources are shaping that visibility, where competitors are gaining ground, and what needs to change to improve discoverability.

From there, we can help you move from diagnosis to action — combining GEDI insight with recommendations across earned media, analyst relations, thought leadership, and owned content so your brand is better positioned in the AI-driven discovery journey.

Get in touch

Request a tailored GEDI assessment.

Understand your current GEO performance and identify practical steps to strengthen it.

Paul Nolan
Co-MD UK
pnolan@hoffman.com

The Hoffman Agency → www.hoffman.com