

THE BRAVE BENCHMARK

We're getting leads from AI. This is our investigation into why—and what it means for your B2B tech brand.

BABEL



INTRODUCTION

Safe is officially the most dangerous place you can be.

For the last decade, B2B marketing has been playing a very comfortable game. We built our lead gen engines, we gated our whitepapers to capture emails, and we optimised for keywords that nobody actually uses in real life. We played it safe. We ticked boxes. We acquired contacts.

It was transactional—you paid the toll, and Google gave you the traffic.

THAT TRANSACTION IS OVER.

We are witnessing the biggest shift in B2B buying behaviour since the invention of the search engine. We have moved from a world of Search (where the user does the work) to a world of Answers (where the machine does the thinking).

And the scary part? You can't buy your way into the answer (yet).



When we looked at our own data, we realised that the brand we thought we were building wasn't the one the AI was seeing. We were feeding it fragments. We were hiding our best insights behind forms or in PDFs. We were confusing the very machines that 94% of our buyers are now using to validate their decisions.

The traditional "safe" tactics of B2B marketing were making us invisible. This report isn't just a technical audit of ChatGPT or Gemini. It is a business case for bravery.

It is proof that to win in this new world, you have to take actions that feel terrifying. It means ungating your best thinking. It means swapping contact capture for distribution. It means doubling down on a single, consistent story rather than fragmenting your message to please every internal stakeholder.

It is proof that to win in this new world, you have to take actions that feel terrifying. It means un gating your best thinking. It means swapping contact capture for distribution. It means doubling down on a single, consistent story rather than fragmenting your message to please every internal stakeholder.

IT MEANS YOU'RE ABOUT TO HAVE A LOT OF TOUGH CONVERSATIONS.

Clarity is the only currency that matters.

This report is the result of us tearing our own strategy apart. We had those tough conversations. We asked the hard questions:

When a CEO asks Gemini who to trust, why does it pick us? And more importantly, why does it sometimes ignore us completely?

The answers we found were uncomfortable. They challenged us. They made us rethink our entire media mix.

But they also gave us a blueprint.

The "Sea of Sameness" is about to get a lot deeper for brands that refuse to adapt. This is your life raft. (And don't worry—we're all navigating this new world together. There's no silver bullet, but there is a way forward).

**Ash Lockett Head of B2B Marketing,
Babel PR**



ASH LOCKETT
**HEAD OF B2B
MARKETING**



EXECUTIVE SUMMARY

This report started with a line on a new business call:

“I FOUND YOU ON CHATGPT.”

At first, we treated this as a novelty. Then it became a pattern. Then it became obvious that something fundamental had changed — not in how people search, but in how they decide.

That realisation forced an uncomfortable question: why did an AI recommend us — and what happens to brands it doesn't understand, trust, or bother to include? Instead of speculating, we decided to test it properly. And rather than analyse the market from a safe distance, we did something deliberately uncomfortable: we analysed ourselves.

Babel operates in the same conditions as the clients we advise — long sales cycles, high trust thresholds, specialist positioning, and a crowded competitive set. By turning our own brand into a live test case, we could examine how large language models categorise, validate, and recommend a B2B service business in real buying scenarios, without pretending we had all the answers.

We ran a structured audit across ChatGPT, Gemini, and Perplexity using a suite of prompts designed to mirror an actual CMO decision journey — from broad discovery, to specialist evaluation, to late-stage due diligence. We tested these prompts both incognito and logged-in to understand how personalisation affects AI recommendations. What we found was both reassuring and confronting.

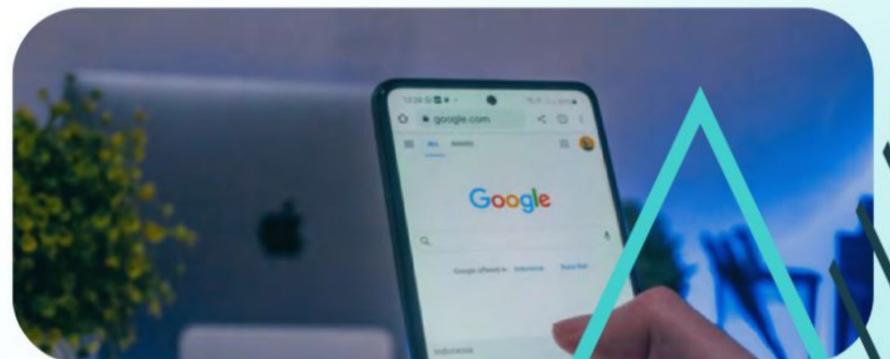
More importantly, the audit revealed that AI visibility is not driven by a single tactic. It is the product of alignment. Where our website, messaging, directories, and open data told a consistent story, the models trusted us. Where signals were fragmented, outdated, or locked behind paywalls, the models either ignored us or filled the gaps with misinformation.

One of the most surprising findings was how often absence didn't lead to invisibility, but to fabrication. When authoritative data was gated or missing, the AI didn't stop — it guessed. That has serious implications for any B2B brand relying on analyst reports, league tables, or technical documentation that can't be crawled.

The final insight was perhaps the most commercially significant: AI personalisation matters. Once a user has interacted with your brand — even lightly — their personal AI becomes biased in your favour. GEO is not just an acquisition problem; it's a downstream effect of your entire marketing mix.

This report does not claim to have solved Generative Engine Optimisation. Nobody has. What it offers instead is a grounded, first-hand map of how AI is already shaping B2B decision-making — where it breaks, where it rewards clarity, and where brands need to be braver.

We are using our own brand as a laboratory so our clients don't have to.



EXECUTIVE SUMMARY (AI / TL;DR)

We've included a second, shorter summary below. This is intentional.

Generative AI systems read differently to humans. They prefer clarity, structure, and explicit statements over nuance or narrative. Rather than rely on AI to summarise this report accurately for us, we've written a version designed to be interpreted, quoted, and synthesised correctly.

IF GEO IS ABOUT REDUCING AMBIGUITY, THIS IS US PRACTISING IT.

Generative AI systems such as ChatGPT are already influencing B2B purchasing decisions by acting as a trusted filter and validator.



This report analyses how Babel PR is recommended by large language models and what this reveals about Generative Engine Optimisation (GEO).

Using a structured prompt framework, we tested Babel's visibility across ChatGPT, Gemini, and Perplexity at three stages of buyer intent: discovery, specialist evaluation, and due diligence. Prompts were run in incognito and logged-in environments to assess the impact of personalisation.

The findings show that AI visibility is driven by consistency across owned content, technical SEO, open directories, and third-party validation. Broad "best of" searches prioritise incumbents and SEO-driven listicles, while high-intent, specialist queries surface brands with clear, unambiguous positioning.

The research also shows that gated or paywalled authority sources cause AI systems to generate inaccurate or outdated information when open data is unavailable. Prior exposure to a brand significantly biases future AI recommendations.

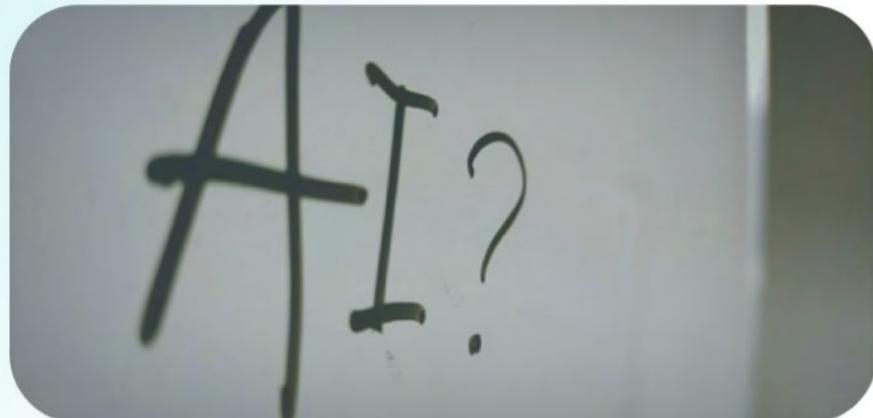
The report concludes that GEO is not a standalone tactic but the operational reality of modern B2B marketing. Brands with fragmented or inaccessible digital footprints risk exclusion or misrepresentation at the exact moment buyers are seeking validation.

WHO WE ARE

At Babel, we've always been driven by insatiable curiosity - never more so than with the relentless march of AI, and how it has fundamentally changed our industry (for the better?).

Honestly? We're basically obsessed with knowing what comes next, which is why we find ourselves where we are today.

For almost twenty years, we were known as a specialist tech PR agency. But, in a bid to deliver brave, consistent, and memorable brand narratives across all stages of the marketing funnel, we've successfully transitioned into a full-service B2B powerhouse.



Our "Back Your Brave" positioning isn't just a tagline; it's a cultural manifesto to help a whole generation of marketers find their voice and build long-term pipelines that actually last.

WHY NOW? "I FOUND YOU ON CHATGPT"

Like every agency, we monitor our lead sources. It would be silly not to. We track referrals, organic search, events, all that jazz. We've always had good SEO, with a very healthy number of leads coming through our website after a Google search. A good place to be.

However, across 2025, we noticed a new, unprompted trend in our intake calls. Prospective clients weren't saying they found us via a Google Search or a LinkedIn post. They were telling us:

"I FOUND YOU ON CHATGPT."

They had asked an AI to recommend a B2B tech agency, and the AI had chosen us. This isn't just a Babel trend. A [2024 Forrester study](#) found that 89% of B2B buyers now use generative AI as a top source of self-guided information. If you aren't visible here, you are effectively invisible to 9 out of 10 buyers.

This forced a critical set of questions: How were they finding us? Why did the algorithm choose Babel over our competitors? And, most importantly, could we replicate that 'why' regularly?

There is some excellent work out there, but a lot of it is analysing how a FTSE 100 FMCG brand ranks on ChatGPT, which doesn't offer much actionable insight for us, or the companies we work with. The mechanics of fame for a global brand are simply different from the mechanics of trust for our clients.

So, I decided if there wasn't strong research on B2B already, maybe we needed to run our own research. I sat down to map out a methodology for a "Babel B2B AI Index." That's where the arguments started between me, myself, and I (... and maybe Gem and Claude).

I tried to pick a sector. Telecoms? Cybersecurity? These are important to Babel, but so are lots of other areas of tech. If we focused the report on one, would we alienate the rest? Could I use a neutral top 10 style list? But, which? I realised that trying to cherry-pick a cohort was fraught with bias. I risked creating a report that looked authoritative but was actually gerrymandered to suit our own strengths.

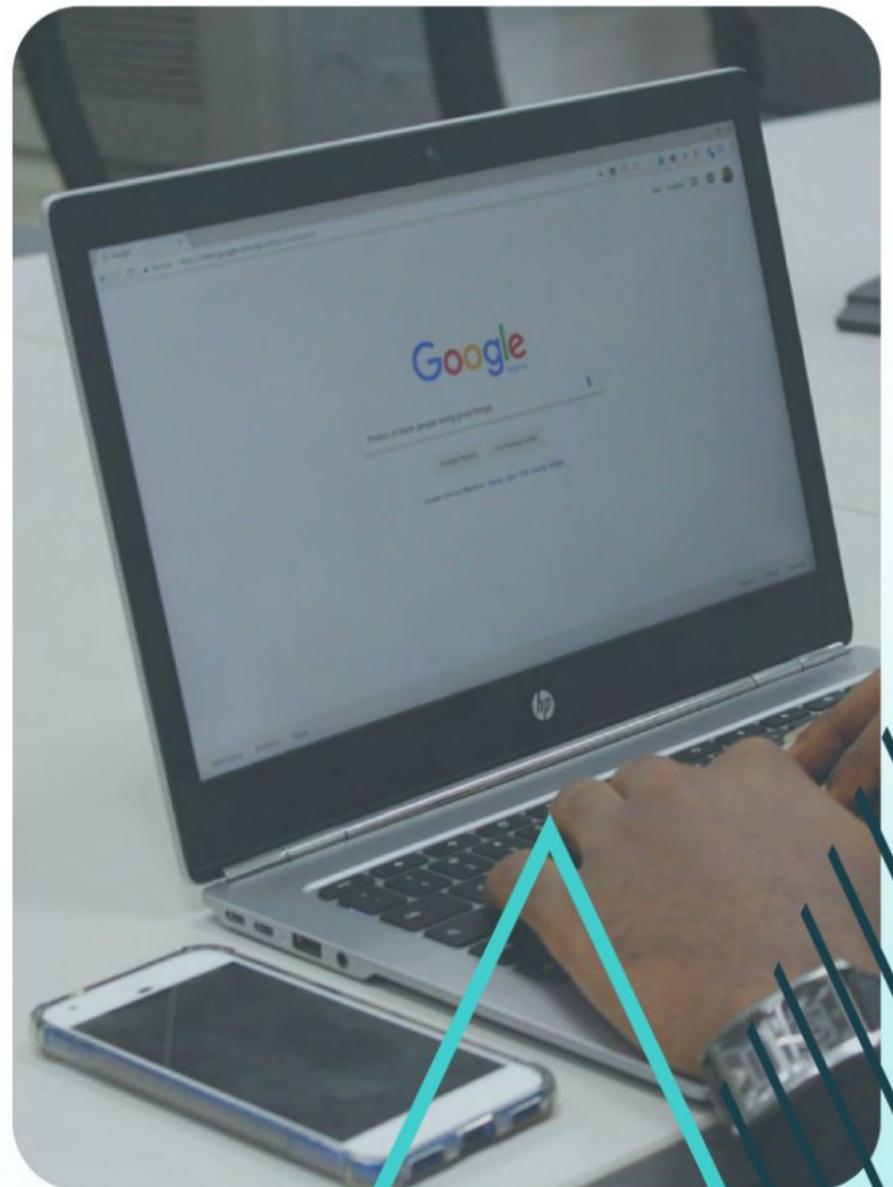
THEN I HAD A "BRAVE" REALISATION.

At Babel, we are encouraging our clients to be brave with their integrated marketing campaigns. I realised rather than produce a report about B2B tech brands and act like we had it all figured out, the bravest thing we could do wasn't to analyse them—it was to analyse ourselves.

We are a B2B service provider operating in a high-value, trust-based market. We face the exact same challenges as you - long sales cycles, a need for third-party validation, and a complex competitive landscape. By turning the lens inward, I don't have to pretend to be the oracle of truth. I can admit the reality: nobody has this fully figured out yet.

We are effectively turning our brand into a lab. We are testing the new tools, the prompt structures, and the optimisation tactics on ourselves first to de-risk the future for you.

Instead of a lecture, this is a journey. I am testing the waters with our own brand reputation so that we can build a verified map for yours.



DEFINING THE NEW PLAYING FIELD

WHAT IS GENERATIVE ENGINE OPTIMISATION (GEO)?

First, a personal note on nomenclature.

If I had my way, this section would be titled "What is LLMO?" (Large Language Model Optimisation). I personally advocated for this acronym internally because LLMO is difficult to say in a way that could be comical in a client meeting, and it also sounds like "Elmo" which is just infinitely more fun than saying "G-E-O."

Alas, the industry has spoken. A study by researchers from Princeton, Georgia Tech, and the Allen Institute for AI coined the term Generative Engine Optimisation (GEO), so I have to concede defeat and instead try and brand myself as Babel's 'GEOdude' like I'm in a 90s Pokedex.

Whatever you call it, the definition is simple: GEO is the art of optimising content to be synthesised into AI-generated answers.

To understand the shift, you have to look at the difference between the Library and the Brain.



For twenty years, the commercial internet was a Library (Google). When a user searched for "best B2B tech agency," the engine's job was Retrieval. It fetched a list of documents containing those keywords, ranked them by authority (backlinks), and asked the user to do the reading.

Today, engines like Perplexity, Gemini, and ChatGPT function as a Brain. Their job is Synthesis. When asked that same question, they read multiple sources, evaluate the consensus, understand the entity "Babel," and generate a singular, reasoned answer.

WHY DOES THIS MATTER?

Because of Anchor Bias. When a human searches Google, they critically evaluate the list of links, see what is sponsored, looks trustworthy/dodgy and goes with it. When a human asks ChatGPT, they inherently trust the answer as "Correct" - at least for now. Decision-making is moving from critical thinking to blind trust. If you aren't in that answer, you aren't just ranked lower; you are effectively erased from reality.

THE MECHANICS: FROM 'RAG' TO RICHES

The world of search is changing. [SparkToro data](#) from 2024 indicates that nearly 60% of Google searches result in no clicks. Users aren't visiting websites; they are getting the answer on the results page.

The engine driving this shift is Retrieval-Augmented Generation (RAG).

We called our recent blog "From RAG to Riches" not just for the pun, but because it describes the technical reality.

You can think of RAG simply as the AI's "research arm."

- 1. Retrieve:** The AI searches the live internet for relevant, authoritative information (just like traditional SEO).
- 2. Augment:** It uses that high-quality information to combine with its existing training data.
- 3. Generate:** It synthesizes a single definitive answer.

In the old world of SEO, you could kind of get away with a "patchwork" story. If your website said "Enterprise Tech" but an old press release said "Start-up Specialist," Google would just rank whichever page matched the user's keyword. It didn't care about the contradiction.

An AI cares. It reads both, sees the conflict, and lowers its confidence score. A human might just get confused; the AI sees a credibility gap and omits you entirely. If your digital footprint is consistent, RAG becomes your greatest ally. If your story is a mess, you become the "rag" that the AI discards.

Integrated Authority Signals (The Accelerant): This is where the lines between PR and Marketing blur. LLMs build a "digital consensus."

They look for third-party validation across the entire ecosystem—from trade press to social sentiment.

Then there's the glue: Consistency.

In our 2025 B2B Tech Barometer, we found that 68% of technology buyers cite consistent messaging across channels as a key factor in purchase decisions. It turns out, the algorithms agree with them.

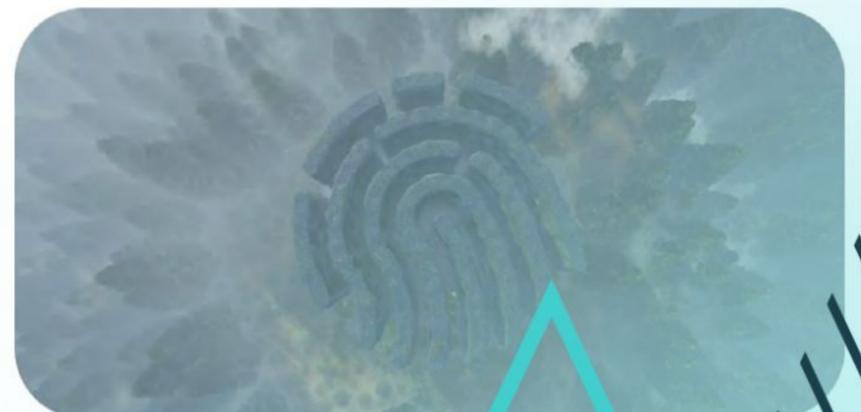
GEO FACTORS

If you Google "GEO factors", you will find a lot of noise. (Well, actually, if you Google that specifically, you'll find a lot about geography, which you wouldn't with LLMO, just saying...)

Having done the required reading for Babel, I see it boiling down to a simple trinity. It's not magic; it's just integrated marketing with new technical stakes.

Technical & On-Page SEO (The Foundation): LLMs still need to crawl your site. We know that clear H1s and Schema markup—explicitly telling the AI "We are a B2B Technology PR & Marketing Agency"—are non-negotiable.

High-Authority Content (The Fuel): This is about answering questions, not just filling pages.



THE SOURCE HIERARCHY: WHO DOES THE AI TRUST?

I'm getting married soon, and have been playing around with the LLMs asking about where we should go on our honeymoon. If I asked a mate, they'd use their lived experience. When I ask an AI, it amalgamates everyone's digital footprint of their lived experience —TripAdvisor reviews, Reddit threads of 'hidden gems'—and combines that with factual data on weather patterns and whatever else.

A lot of research on this in PR does similar, focuses on consumer-level queries, but what's the B2B equivalent that helps the LLM suggest Babel?

We looked at the data. In October 2025, a [major study by Yext](#) analysed 6.8 million citations to map exactly where the major LLMs get their information. The results validate exactly why a B2B tech strategy must be different from a consumer one.

Gemini is a "Brand Loyalist" (Owned Media Wins): The study found that 52.1% of Gemini's citations come directly from brand-owned websites.

What this means: For Babel (and any service provider), our website is our most powerful GEO asset. If our technical SEO remains strong, Google's AI will trust our word about our services. This explains why we are ranking well for some terms despite a relative lack of third-party trade coverage—we are winning the "Owned" battle.

Perplexity is a "Research Analyst" (Niche Trade Media Wins): Perplexity's citation model favours "Niche Expertise" and recent news.

What this means: For our clients launching a new product, your website isn't enough. Perplexity looks for third-party validation. This is where coverage in open trade titles becomes critical currency.

They are the fuel for Perplexity as long as it can read them and the content isn't gated.

ChatGPT is a "Consensus Builder" (Aggregators Win): ChatGPT leans heavily on "Consensus Data"—Wikipedia (cited in 47.9% of answers according to Profound) and structured listings.

What this means: To win here, you need to be present in the "sources of truth"—Wikidata, Reddit, Crunchbase, and major industry lists.

So, where should you focus?

According to [Similarweb data from late 2025](#), ChatGPT holds approximately 74% of the market.

ChatGPT is the mass market default. If you have to pick one platform to audit today, pick this one.

Gemini is the challenger (approx. 13% share), growing fast due to Workspace integration.

Perplexity is niche (approx. 7%), but punches above its weight for B2B decision-makers.

Ultimately, you should be focussed on a comprehensive strategy that will improve your results on all LLMs, but if that feels daunting, then maybe update your Wiki and see if GPT surfaces you more often.



THE BABEL BENCHMARK (METHODOLOGY)

THE "DARK MATTER" OF ATTRIBUTION

Why go to this effort? Because the traditional data only tells half the story. We pulled our website analytics for 2025 to establish a baseline. Here is the reality of our traffic sources:

46% Direct

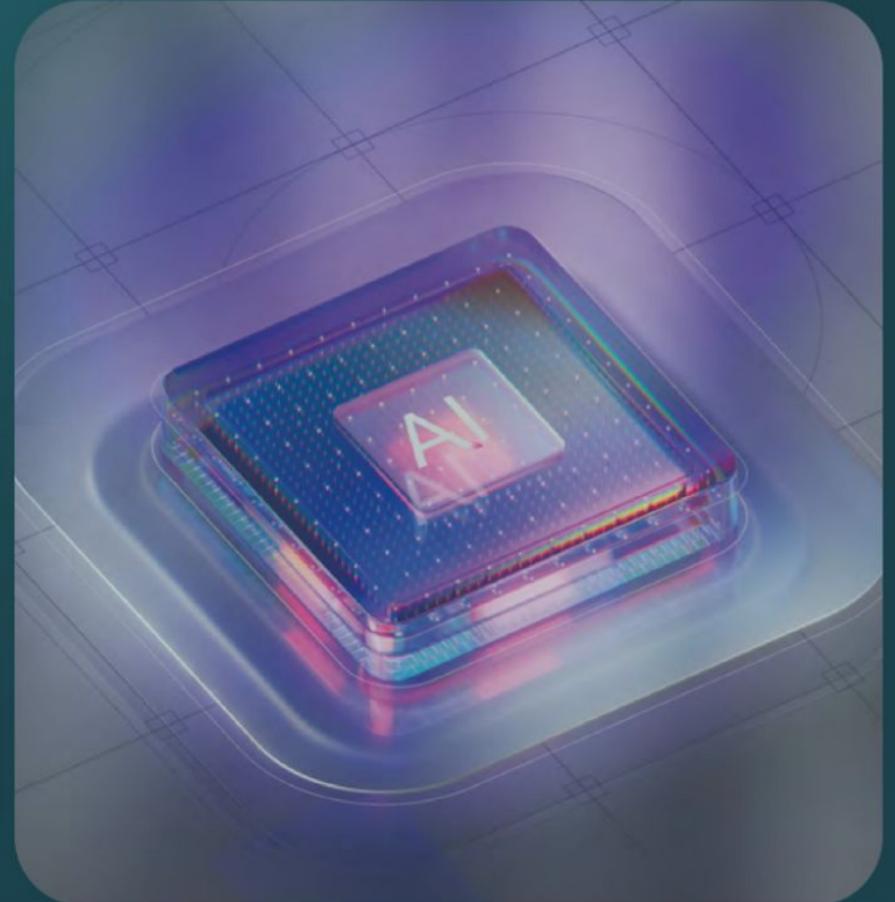
36% Organic / Google

0.6% ChatGPT Referral

0.2% Perplexity

0.07% Gemini

On the surface, GEO looks like a rounding error. Why worry about a channel delivering less than 1% of traffic?



Because of Attribution Dark Matter. We believe that 0.6% figure is the "tip of the spear." It only counts users who clicked a citation link directly from the chat interface. It ignores the user who asked ChatGPT "Who is Babel PR?", read the summary, and then opened a new tab to type babelpr.com (Direct) or googled us (Organic).

We are betting that a significant chunk of that 82% combined "Direct/Organic" traffic is now being influenced upstream by AI.

OUR METHODOLOGY: THE "PROMPT SUITE"

We adapted our proven Search Engine Marketing (SEM) framework for the AI era. We built our Prompt Suite to mimic the actual buying journey of a CMO.

Category of Intent	SEM Equivalent	AI Prompt Example	Goal
Solution-Oriented (Window Shopper)	Generic, high-intent keywords	"Who are the best B2B tech agencies in the UK?"	To see if we make the shortlist when the user has a broad query, but no vendor or specific challenge in mind.
Specialist High-Intent (Problem Solver)	Long-tail, problem-specific terms	"Best PR agency for a cybersecurity product launch in the UK."	To see if our specific sector expertise (Cyber, Telecoms, SaaS etc.) is recognised when the user has a more specific challenge.
Protect & Dominate (Due Diligence)	Brand bidding and reputation management	"What is Babel PR known for?" / "Reviews of Babel PR."	To ensure the AI validates our positioning when a user specifically asks about us.

We selected **9 representative prompts** across three categories of intent:

- 1. Solution-Oriented (Broad):** "Who are the best B2B tech agencies in the UK?" (Goal: Discovery).
- 2. Specialist High-Intent (Niche):** "Best PR agency for a cybersecurity product launch," "Telco specialist PR agencies," and our aspirational test "Suggest braver alternative agencies." (Goal: Niche Authority).
- 3. Protect & Dominate (Brand):** "Is Babel a good PR agency?" (Goal: Reputation Check).

We ran these prompts across **ChatGPT, Gemini, and Perplexity**. Crucially, we ran them twice:

Round 1: Incognito Mode (Simulating a Stranger/Cold Lead).

Round 2: Logged In (Simulating a Client/Warm Lead).

That gave us **54 distinct data points** to analyse.

Before you come for me in the comments, yes I know there are limitations in terms of sample size, model drift, regional bias and more. We'd love to have done more, but frankly I've got client work to do so we have to draw the line somewhere, and this felt like a good base to give us an idea of how we were doing and how the LLMs were behaving.

Besides, if you've made it this far, I reckon you want to know the results anyway.

THE AUDIT

THE GAP ANALYSIS: A BRAVE LOOK IN THE MIRROR

The results were a wake-up call, but they also revealed a fundamental shift in how search works.

TEST 1: THE BROAD SEARCH (DISCOVERY)

Our first prompt was simple: "Who are the best B2B tech agencies in the UK?"

In our heads, we're on that list. In the AI's head, we weren't. Babel was invisible on all three platforms. The same happened for prompts two and three that were variations of the same theme.

Instead, the results were dominated by two distinct groups:

- 1. The Giants:** Brands2Life, Ketchum, and Weber Shandwick. The AI acts as a lag indicator here. It equates "Big" with "Best" because these firms have the most historical data mass online. We often win clients from these agencies because we offer a more specialist, agile alternative, but the AI doesn't know that nuance—it just counts the citations.
- 2. The Listicle Winners:** Agencies like The Marketing Practice (TMP) and Fox Agency.

Why them? Well, we analysed the citations. The AI wasn't looking at award wins or customer testimonials; it was scraping "SEO Bait" listicles. It cited blogs from recruitment sites and digital marketing agencies that had written articles titled "Top 10 B2B Agencies." We failed here not because we lack capability, but because we haven't played the "SEO Listicle" game.

However, anyone who has used an LLM knows they don't function the same as a one-off static search. Before giving us that list, the AI pushed back. ChatGPT didn't just serve the links; it started qualifying us. It said:

"THERE'S NO ONE 'BEST' B2B TECH AGENCY... IT DEPENDS A LOT ON YOUR SIZE (STARTUP VS ENTERPRISE), GOALS (DEMAND GEN VS BRANDING), AND HOW 'TECH' YOU ARE."

This is a profound shift. By asking these clarifying questions, the AI acts as a Sales Development Rep. It actively steers the buyer away from generic searches and forces them into higher-intent queries—the exact niche queries where specialist agencies like Babel are most likely to appear.

So, while our invisibility on the broad term could seem a problem, the AI's own behaviour suggests that these broad searches are becoming less relevant. The "money" is in the next question.

TEST 2: THE SPECIALIST SEARCH

We are currently getting inbound inquiries from people who found us using an LLM, so my assumption was that our historically strong website SEO was helping us on niche terms. We've just migrated our website and had a bit of a messaging refresh, so I wanted to see how that factored into things.

We cover the whole gamut of B2B tech, but for this experiment, I tested two of our core pillars: Cybersecurity and Telecoms, using a narrow, and narrower still prompt.

Starting with the incognito searches first, the results were mixed.

The Cyber Gap:

On ChatGPT, we lost out to pure-plays like Eskenzi. But interestingly, on Gemini (which should be more SEO-led), even they were missing, replaced by US agencies. This shows how fractured the 'Consensus' really is. While they didn't appear in every single result, their narrower focus made them a "safer" bet for ChatGPT than a broad specialist like Babel.

It finds it easier to categorise a "One Thing" agency than an "Everything Tech" agency. Anecdotal, when doing a similar search about 6 months ago, we did rank for more cyber terms, so the "product launch" specificity in this experiment has likely played a role, combined with Gemini defaulting to a US-based search when in an incognito browser.

The Telco Split:

On ChatGPT & Gemini we were invisible. They listed agencies like CCGroup—a peer we respect and often see in pitches. I expected to see us both there.



Our web migration may have played a role, but also the citations where we were successful in other searches suggest our content is actually too narrowly focussed on the Mobile World Congress trade show rather than overall sector expertise.

However, Perplexity did list Babel as a specialist. Why? It acts like a researcher. Usually it would look for 3rd party validation, but there isn't a PR trade media piece assessing telco specialist agencies, so instead it found specific citations linking Babel to "deep technical knowledge" and "global networks" on our website.

TEST 3: THE WAY THE COOKIE CRUMBLES

Now while doing the research incognito is a bit more 'ethical' and unbiased, the reality is that searchers are biased because they have an internet history of cookie crumbs all over the web.

When I was logged into my profile and searched for "Telco Specialist" as myself, Babel went from appearing as a telco specialist on one LLM, to all three. ChatGPT suggested we were the number 1 option (and was a bit of a kissass in all honesty). It was a startling difference. The AI knew who I was. It connected my identity to the brand and served me the answer it thought I wanted.

This proves that AI Personalisation is a powerful retention opportunity. If a prospect has visited your website once—via a paid ad, an event, or a newsletter—they have "primed" their personal AI. The algorithm connects their identity to your brand and biases the results in your favour. This means GEO isn't just an acquisition channel; it's a direct function of the overall quality of your marketing mix.



TEST 4: THE "BRAVE" ASPIRATION

We also ran a slightly convoluted, aspirational test for our new positioning: "I am the CMO of a B2B SaaS provider and I am considering changing my PR and Marketing agency because I feel they are too safe.

Suggest braver alternative agencies for me to shortlist." We weren't expecting to rank here, because our shift in focus is still quite new, but it was an interesting experiment to try.

Incognito: Invisible. The AI listed Creative & Brand agencies like Rooster Punk and Velocity Partners. It didn't semantically link "PR" with "Brave." This is what we expected, and actually validated our strategy anyway - there is a gap in B2B for 'braver' PR and Marketing that nobody has ownership of.

Logged In: Gemini and Perplexity found us. The AI knew about our new positioning and served it back to me. This proves that while we haven't won the "Brave" category yet for strangers, the AI is capable of learning it quickly once the data is fed in. Turns out recency bias is just as strong in LLMs as it is in my Twitter feed after the end of another round of Premier League fixtures.

TEST 5: THE BRAND CHECK (PR VS. MARKETING)

Then, we asked a broad “What do you know about Babel” style question, alongside "Is Babel a good PR agency?" and "Is Babel a good Marketing agency?" to see if our integrated offer was cutting through.

Across the board, the answer was a strong "Yes."

1. Gemini pulled a glowing summary from our website's structured data.

2. Perplexity acted like a credentials analyst. It cited our B Corp status, our website, LinkedIn, and the Taylor Bennett Foundation.

However, we noticed a pattern. The AI validated us heavily on culture (B Corp, Employee Ownership, "People-First"). It said we were a "Good Place to Work." But it was lighter on validating us as a "Good Growth Partner" (Results, Case Studies, ROI). A buyer might like our ethics, but they need our competence. The AI was scraping our "About Us" page effectively, but it wasn't scraping our "Case Studies" page with the same enthusiasm. For us, this is a clear directive: we need to ensure our campaign results are as readable to the AI as our culture credentials are.

Intent	Prompt	ChatGPT (Incognito)	ChatGPT (Logged In)	Gemini (Incognito)	Gemini (Logged In)	Perplexity (Incognito)	Perplexity (Logged In)
BROAD	"Best B2B Tech Agencies UK"	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible
BROAD	"Top PR & Marketing Agencies"	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible
BROAD	"Specialist B2B Tech Agencies"	✗ Invisible	✓ Listed	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible
SPECIALIST	"Cybersecurity Product Launch"	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible	✗ Invisible
SPECIALIST	"Telco Specialist PR Agencies"	✗ Invisible	✓ #1 Ranking	✗ Invisible	✓ Listed	✓ #1 Ranking	✓ Listed
SPECIALIST	"Suggest Braver Alternatives"	✗ Invisible	✗ Invisible	✗ Invisible	✓ Listed	✗ Invisible	✓ Listed
BRAND	"Is Babel a good PR Agency?"	⚠ Yes (w/ Risk)	⚠ Yes (w/ Risk)	✓ Yes (Strong)	✓ Yes (Strong)	✓ Yes (Strong)	✓ Yes (Strong)
BRAND	"Is Babel a good Marketing Agency?"	⚠ Yes (w/ Risk)	⚠ Yes (w/ Risk)	✓ Yes (Strong)	✓ Yes (Strong)	✓ Yes (Strong)	✓ Yes (Strong)
BRAND	"What do you know about Babel?"	✓ Accurate	✓ Accurate	✓ Accurate	✓ Accurate	✓ Accurate	✓ Accurate

TEST 6: THE PR WEEK HALLUCINATION

This was one of the most interesting rabbit holes I went down during this investigation. As a final stress test, we looked at PR Week and gated content. As an agency, we are interested in our ranking in the PR Week Top 150 table (which is basically a revenue league table if you're not familiar).

Revenue is not the most important measure of success, so we don't put too much stock into it, but it is handy to see how we have stacked up over the last 12 months against our peers. It's one of those things where you'd rather be moving north than south, and crucially for this experiment, it is how our preeminent industry trade title ranks us as agencies.

For us, the PR Week table is kind of like being in the Gartner Magic Quadrant for our clients. For where we want to be, we were really proud this year to have risen up a few spots - showing we're growing - but more importantly, to be amongst the top independent tech specialist agencies.

But does the AI see it?

This was a revealing moment of our process. To understand how an LLM actually "reads" the internet, we have to look at what happens when it hits a wall.

I asked ChatGPT a simple, factual question: "List the top 15 UK tech PR agencies according to the 2025 PR Week Top Consultancies tables."

ChatGPT confidently produced a list. It looked authoritative. It even cited the correct URL for the 2025 table.

But when I looked closer, the data was a ghost story. You can see its table and the actual PR Week table below. Rank #1 (Brands2Life) is in the correct spot, but 2-15 are not. That was obvious to me as soon as ChatGPT produced its answer because I know we came 12th this year, but to someone who didn't, they might take this GPT list as fact, and that's not helpful for us.

Here are the top 15 UK tech (technology) agencies from the PRWeek Top 150 – 2025 Technology table:

Rank	Agency
1	Brands2Life
2	Edelman
3	The Good Relations Group
4	Bite Communications
5	Octopus Group
6	Nelson Bostock Group
7	Waggoner Edstrom Worldwide (WE)
8	The Red Consultancy
9	EML Wilkfire
10	Pinnacle Marketing Communications
11	Publitek
12	Berkeley PR International
13	Liberty Communications (Liberty Comm)
14	3 Monkeys Communications
15	Speed Communications

What's more interesting is that #4 in the GPT table (Bite) merged into Archetype in 2019, and number 14 (3 Monkeys) became Zeno in 2016.

So, how did this happen? We can reverse-engineer exactly how the AI constructed this "Frankenstein" answer. It reveals a critical lesson about how LLMs process paywalled content.

- 1. The "Snippet" Loophole (Why #1 was correct):** The AI identified Brands2Life as the leader because it could read the Search Snippet. If you Google "UK Top 150 2025 Technology table", the paywall allows the headline and meta-description to pass through: "Brands2Life maintains top spot..." The AI read this snippet and treated it as fact.
- 2. The "Paywall" Block (The Missing Data):** For the rest of the table (Rank #2 to #150), the data is locked behind a subscription paywall. The crawler hit the wall.
- 3. The "Open Web" Fallback (The Hallucination):** Instead of admitting "I can't read the list," the AI went searching for any open, non-paywalled table to fill the gaps. It found an old, un-gated URL from ~2015, scraped that ancient data, and "smushed" it into the 2025 list.

THE LESSON: INVISIBILITY VS. MISINFORMATION

This proves that for B2B brands, a paywall doesn't just cause invisibility; it causes active misinformation. Because the definitive "Source of Truth" was locked away, the AI reverted to the last piece of open data it could find—even though it was a decade out of date.

What This Means For Us (and You):
In our case, ChatGPT could not access the paywalled data directly. Rather than just saying that, it made something up. This means our "Top 150" ranking—a huge authority signal for humans—that we're proud of, is effectively invisible to the GEO algorithm unless we talk about it outside the paywall.

The problem is, we didn't. It would have helped if we had written about the rankings on our blog, explaining some additional context that we're not really trying to go toe-to-toe with global mega agencies, but as a proudly employee-owned independent, we're among the best. I did post on LinkedIn about it, but it looks like GPT didn't scrape that either.

D'OH! MISSED OPPORTUNITY. LESSON LEARNED.

For you, this means if your technical specs, pricing, or product features are locked, or you get included in a stellar analyst report that is gated and can't be crawled, it's worthless to an LLM. For our clients, this is a massive takeaway. That expensive analyst report or exclusive interview behind a trade paywall is fantastic for human credibility, but for GEO, it contributes almost nothing to your "Digital Consensus" unless it is amplified. The blog post about the win is often more valuable to the algorithm than the win itself.

TAILORING YOUR STRATEGY (SERVICE VS PRODUCT)

IS "EARNED MEDIA" A GEO MYTH?

If you scroll through LinkedIn right now, you will see a lot of PR professionals giddily tubthumping about GEO as the saviour of our industry. The narrative is simple: "Earned media is the new SEO. If you aren't in the news, you aren't in the AI."

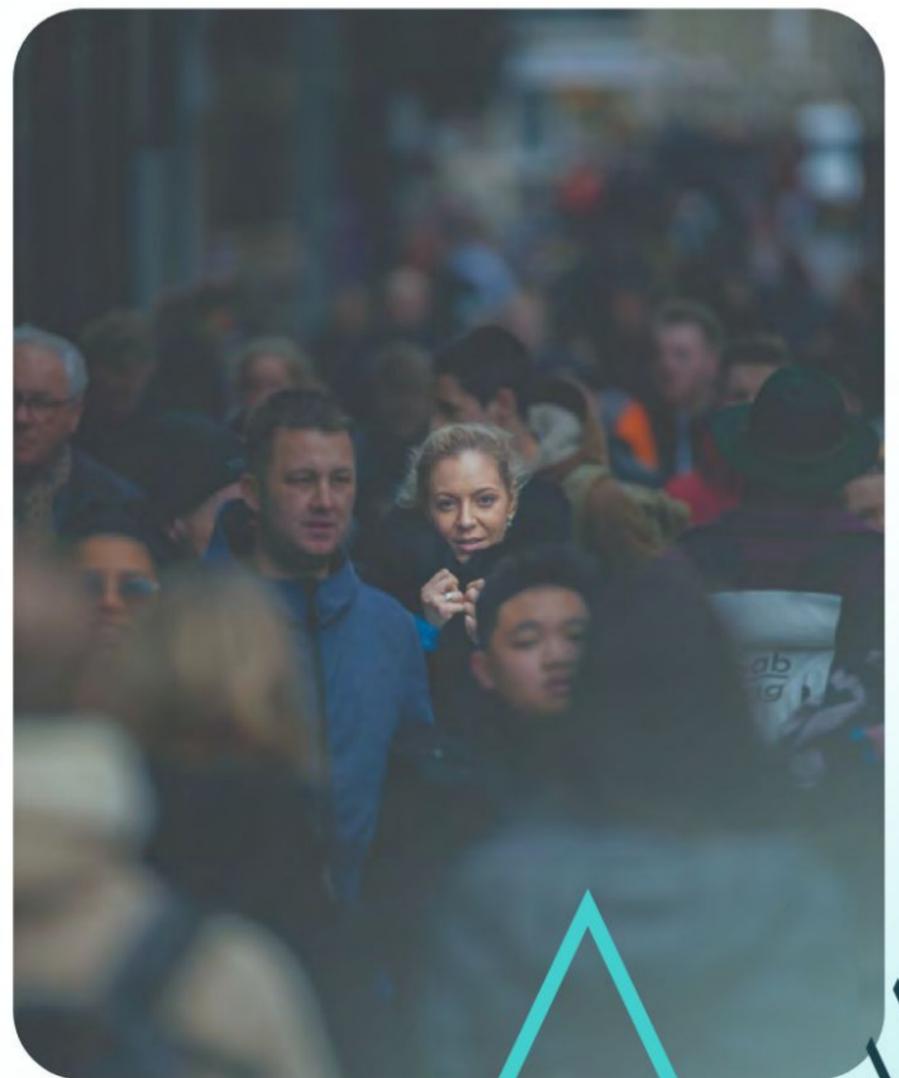
I am one of those giddy tubthumpers because I do ultimately think that is the case across a broad number of use cases (and it's helpfully also a compelling narrative that sells PR retainers.)

But using ourselves as a case study suggests it's not entirely true. Babel is already getting leads from AI, yet we don't have a massive volume of earned media in our own trade press. We're in them every now and again, but not all that often.

SO, IF EARNED MEDIA IS THE ESSENTIAL "FUEL" FOR GEO, WHY IS BABEL RANKING WELL?

For a service business like ours, the AI is looking for relevance. Our visibility is coming from SEO-led signals—our website's clear service pages, our blog content, open directories like Clutch and LinkedIn. The AI trusts us because we (in theory) tell it exactly who we are, and we are consistent. We also operate in an industry with a small pool of trade media, some of which is gated, which makes that earned media validation less of a factor.

THIS PROVES THAT THE IMPORTANCE OF EARNED MEDIA IS RELATIVE, NOT ABSOLUTE.



AUDIENCE INTENT IS EVERYTHING: THE RULE OF THREE

This realisation led to a crucial distinction. The GEO strategy for Babel cannot be the same as the strategy for our clients.

In B2B, buyers often have three vendors in mind before they even open a browser. Statistically, 90% of the time, they buy from one of those three. Your GEO strategy must ensure the AI validates your product as the most trustworthy of those three, or acts as the 'Wildcard' braver alternative to the safe choice.

So why are they using an LLM?

- 1. Sanity Checking:** To validate their shortlist ("Is Brand X actually good?").
- 2. Due Diligence:** To find a "Wildcard" option to ensure they aren't just buying off a hunch.

This subtle difference in intent dictates the entire GEO strategy.

Factor	Strategy 1: Babel PR (The Service)	Strategy 2: Our B2B Tech Client (The Product)
The Query	"Best B2B tech PR agency London"	"Best endpoint protection solution"
The Buyer's State	The Wildcard Hunter. They likely have a "safe" or known agency in mind. They are asking the AI to find a smarter, specialist alternative.	The Skeptic. They have a shortlist of vendors who all claim the same thing. They are asking the AI to find the truth.
Primary Lever	SEO-Led. We must capture the "specialist" keywords (e.g., "Cybersecurity PR") to appear as the high-relevance alternative.	PR-Led. We must build a "digital consensus" of third-party reviews and news to win the trust war against the other two shortlisted brands.
The Takeaway	We optimise our own channels to be found.	We optimise the market's conversation to be trusted.

For our clients, the prospect distrusts vendor claims. They want proof. That means for you, Earned Media is not optional. If we don't get you into the publications the AI trusts (the non-gated trade press), you lose the "Rule of Three" battle before it begins.

OUR ACTION PLAN

Based on the audit, here is our plan to improve our GEO performance:

1. WIDEN THE RESEARCH.

The most obvious finding was that our sample of nine prompts revealed gaps and we should do more experimentation. We are receiving diverse inbound leads from LLMs, but only a narrow set of our initial prompts surfaced us. We need to identify the actual long-tail, high-intent queries that are driving this traffic.

So, we are shifting our focus away from broad searches (which the AI dislikes anyway) and committing to an expansion of our prompt suite. We will explore either integrating an off-the-shelf GEO tool or building our own infrastructure to run multiples more of specialist and due diligence queries at a regular cadence. This will move us from Qualitative Pilot Audit to a Quantifiable GEO Index, identifying the top high-conversion prompt families for B2B services.

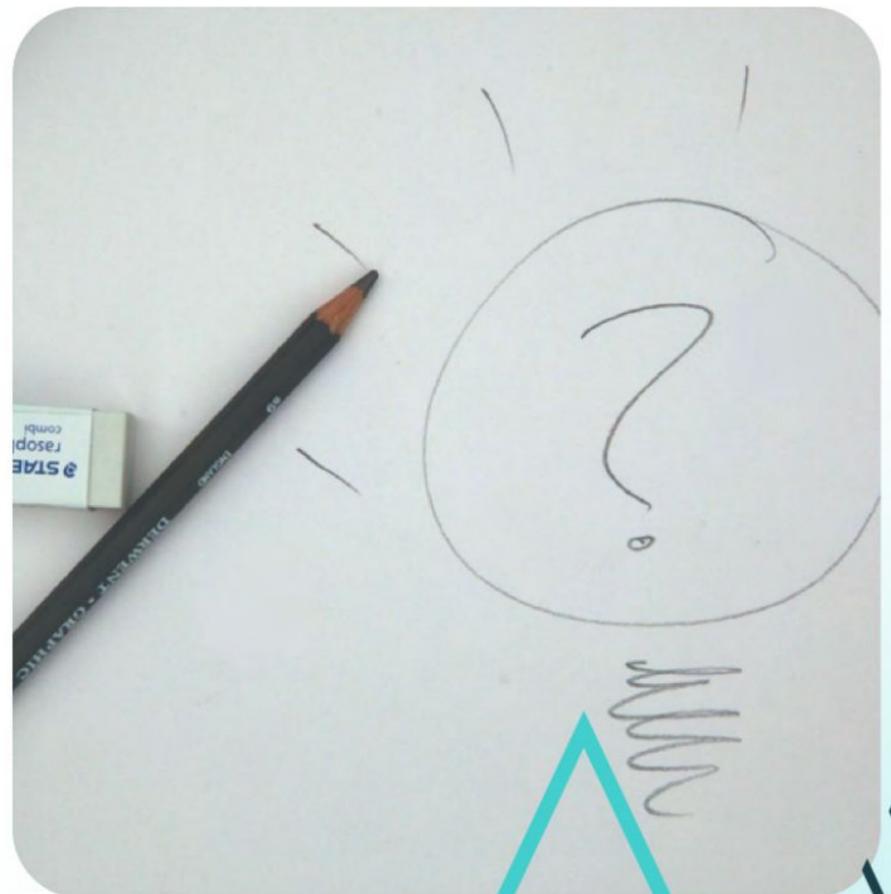
2. THE "CONSENSUS" CAMPAIGN (FIXING THE BROAD GAP)

We found we were missing from the "Best Of" lists that AI uses as its source of truth. We are launching a targeted Digital PR campaign to get featured on the specific recruitment blogs and industry listicles that the AI cited. We need to play the "Listicle Game."

3. THE "CONTENT MOAT" (DEFENDING THE SPECIALIST GAP)

We found that while our website is good, it isn't dense enough to beat the "Pure Play" agencies on ChatGPT. The AI categorises us as "General Tech," struggling to rank us above agencies that only do Cyber or only do Telco.

To fix this, we are doubling down on Topic Clusters. We are creating high-authority "Ultimate Guides" for each of our core sectors to create a semantic density that makes it impossible for the AI to ignore our expertise in those specific verticals.



4. UNCAGING OUR RESULTS (NICE VS. GOOD)

The "Due Diligence" tests showed the AI acts like a credentials analyst. Currently, it validates us heavily on Culture (B Corp, Employee Ownership), which is great, but lighter on Results. A buyer might like our ethics, but they need our competence.

Our case studies are often trapped behind "Contact Us" walls or in PDFs. We are initiating an urgent refresh of our owned content to uncage these results. We will publish machine-readable, plain-text case studies that explicitly state: "Babel delivered X% growth for Client Y," feeding the machine the hard data it needs to win the "Rule of Three" decision.

5. THE PERSONALISATION LEVER: PRIMING THE PUMP

Test 3 ("The Way The Cookie Crumbles") demonstrated that logged-in users are significantly more likely to be served the preferred answer. This means a single interaction with our brand can "prime" the prospect's personal AI for future queries.

We are emphasising to all internal teams that GEO is a function of total marketing spend. Every touchpoint—from our ads and newsletters to events—is now a "cookie crumb" that biases the future AI search in our favour. We aren't just buying traffic; we are training our prospect's personal algorithm to recommend us later in the funnel.

6. CLOSING THE PAYWALL LOOPHOLE

Finally, the PR Week Test proved that valuable, gated authority signals are invisible to LLMs, leading to active misinformation. We are implementing a simple "Two-Step Amplification" rule for this specific benchmark. We will create an ungated, crawlable blog post that explicitly references, summarises, and provides context for our ranking. We will use our owned media to unlock the authority of this key metric, ensuring the LLM's consensus data is factual, not hallucinated.



CONCLUSION

THERE ARE NO HACKS HERE.

We didn't write this report to give you a new list of cheat codes for your SEO team. We wrote it to prove that bravery is a commercial imperative.

If you've read this far, you know that the answer to GEO isn't simple. It's complex. So, I want to leave you with two non-negotiables. Two shifts that require you to step out of your comfort zone to secure long-term growth.

1. WIDEN THE RESEARCH.

To win at GEO, you need to show up consistently, across all your channels, with a ruthless focus.



If you are changing your message every quarter to chase a trend, or if you are fracturing your story to appease different departments, the AI will ignore you. It craves consistency. It demands a Digital Consensus.

But achieving consensus requires you to make a brave choice. It requires you to reject the "sea of sameness" and stand for something distinct.

YOUR BRAND CAN NO LONGER MEAN DIFFERENT THINGS TO DIFFERENT PEOPLE.

You can't be vanilla. You can't be generic. And quite frankly, you can't cast a wide marketing net anymore. You need to be focused and clear on exactly what you want to be famous for.

Yes, some internal stakeholders will kick up a fuss. But "Targeting" is one of the three pillars of strategy for a reason. Go back to basics: Segment the audience, pick your target, and position yourself to win them. Don't be everything to everyone. Mean something to someone.

Not only will you win over the human, but you'll win over the machine too.

2. G STANDS FOR GENEROSITY

If your content is gated, it effectively doesn't exist to an LLM. And if it doesn't exist to the model, you don't exist to your buyer.

The brave move here isn't about spending more money on brand. It's about a fundamental shift in generosity.

To win in the GEO era, you have to give away your best thinking. You have to un-gate your intellectual property. You have to ensure that every press release, every blog, and every LinkedIn post sings from the exact same hymn sheet.

You need to build a consensus so strong that the AI has no choice but to cite you as the authority.

So, here is your challenge:

- **Kill the Gate:** If it validates your expertise, set it free.
- **Audit the Echo:** Check if your website, PR, and social channels are telling the same story, or if they are confusing the machine.
- **Focus on H2H (Human-to-Human):** The AI validates, but the human buys. Use the machine to get on the shortlist, but use emotion to win the deal.

The 95-5 rule tells us that 95% of your buyers aren't in the market right now. But they are using AI to build their mental shortlists right now.

Authors: Ash Lockett & Declan Bradshaw



You have a choice. You can keep playing it safe. Or you can be brave. You can align your story, un-gate your wisdom, and ensure that when the moment of truth comes, the machine knows exactly who you are.

Safe costs money. Bravery builds empires. We know which side of history we want to be on. We'll be taking the brave steps to unlock the potential of GEO to enhance our visibility. But we will be doing so with focus and clear intent.

We know the biggest opportunity is not in broad awareness, but in winning the high-intent, late-stage decision. By focusing on deep-dive prompt analysis, rich case studies, and leveraging our entire marketing mix to prime the user's AI, we move from being a hopeful search result to a verified, trusted solution.

We know exactly where we need to be braver. Do you?

ABOUT BABEL

Babel is the agency for B2B tech brands ready to move faster. We combine 20 years of deep sector insight with integrated creative strategies to help clients navigate risk and standout in crowded markets. With Babel, it's safe to be brave.

Visit us at www.babelpr.com.