


WHITEPAPER

Marketing in an **AI World**

From Mental Availability to AI Availability

As AI reshapes how consumers discover and choose brands, marketers must move beyond mental availability to build AI availability: ensuring brands are visible, trusted and recommendable by machines as well as people.



Most brand-growth theories

are built around attention, memory, and availability. However, these don't answer why some brand memories have become valuable. Meaning, culture, identity, emotion, trust, social proof and experience drive the content and quality of those memories. In other words, availability puts the brand in the choice set; meaning and experience determine whether it should be chosen.

Brands have invested in reach, unique assets, emotional consistency, and salience because these have made them more memorable, and easier to choose. AI doesn't upend that logic but layers a filter on top of it. Consumers are delegating more decisions to AI to search, shortlist, compare, summarize and recommend. In the future, consumers may not be as involved in those decisions. A world of AI agents deciding for us is entirely possible. Brand visibility now depends not just on whether consumers can remember your brand, but whether machines can recognize it, retrieve trusted information about it, compare it accurately, and explain why they should recommend it. Rather than striving solely for mental availability, brands must now develop mental, physical, and AI availability.

However, not all AI systems operate as generic discovery engines. Personal agents may

increasingly act as persistent behavioral intermediaries that learn individual preferences, communication styles, routines, purchase histories, aesthetic tastes, values and emotional associations over time. In these evolving environments, brand storytelling and emotional resonance do not disappear. They may instead become reflected through behavioral patterns the agent learns from and optimizes around. For example, a consumer who consistently prefers sustainable products, nostalgic brands, specific creators, or premium aesthetics may train their personal agent to preserve and reproduce those preferences during recommendation and decision-making processes.

The brand challenge of the next decade: **how do we build equity and emotional resonance with people, while also becoming legible and trustworthy to machines?**

The theory stack

We think the strongest way to frame this is not to treat AI visibility as an entirely new marketing problem. It is better understood as an extension of established brand-growth theory. A brief review of some of the relevant and pertinent theories is a good starting point.



Binet and Field: long-term brand effects

Binet and Field's IPA work shows that brands need to balance short-term activation with long-term brand-building. Their work focuses on how campaign effects develop over time and warns against relying too heavily on short-term online metrics at the expense of long-term brand strength.

Branding builds memory structures, fame, emotional connections, etc. Activation flips existing demand, but without brand-building it rarely produces long-term growth.

In this discussion, we see this as applicable to AI visibility as well. When AI systems are trained on and retrieving information from the environment around brands, those brands with a larger cumulative presence are probabilistically more likely to be mentioned in AI-generated responses.



Byron Sharp and Ehrenberg-Bass: availability

Byron Sharp's contribution is different. His core argument is not primarily about the brand/activation budget split. It is primarily about

availability. The Sharp and Ehrenberg-Bass theory argues that brands grow by increasing penetration, which requires both **mental availability** and **physical availability**. Ehrenberg-Bass also links mental availability to **Category Entry Points**: the needs, occasions, motives and contexts that trigger category consideration. A brand grows when it is linked to more of these buying situations.



Charlie Oscar/ WARC: AI visibility

Charlie Oscar/WARC gives us a new reflection on the AI-era evidence layer. The work attributes two-thirds of LLM visibility to long-term brand equity, with current marketing activity and citations making up smaller shares. WARC's broader guidance also argues that brand fundamentals become more important, not less important, as AI tools insert themselves into the relationship between consumers and brands. Interestingly, Charlie Oscar/WARC suggests that AI systems may now reward the same cumulative brand effects by making established, well-evidenced brands more visible in AI-mediated discovery.

A brief comparison between Binet and Charlie Oscar's work highlights some very interesting parallels.

As mentioned, Charlie Oscar's work shows that approximately two thirds (~66%) of brand visibility inside Large Language Models (LLMs) is driven by long term brand equity. This closely mirrors the long established IPA / Les Binet finding that around 60% of marketing effectiveness comes from long term brand building. This is not a coincidence. LLMs structurally reward the same cumulative, long term effects that econometric models observe in brand growth and profitability.

LLMs do not rank pages or optimize for bids. They synthesize answers based on:

- Long-term historical training data, increasingly supplemented by real-time retrieval and live search systems
- Frequency and consistency of brand mentions
- Cross-context presence
- Clear entity definition
- Quality and consistency of documentation across sources (where scale helps, but is not the only path)

As a result, AI visibility appears to reflect a combination of brand equity, source consistency, and the quality of machine-readable documentation across the information environment. Scale helps, but structured evidence and credible third-party corroboration can alter the outcome. This means smaller brands may still achieve disproportionate AI visibility if they are

consistently documented across authoritative and machine-readable sources.

There is a direct structural equivalence between Binet's work on brand effects and LLM behavior:

 Binet Effect	 LLM Analogue
Long term brand equity	Historical data accumulation
Mental availability	Generation probability
Broad reach	Cross context mentions
Distinctiveness	Stable entity representation
Consistency	Reinforced linguistic signals
Scale	Probabilistic dominance

Binet's 60/40 rule refers to marketing effectiveness and profit growth. LLM visibility percentages refer to how discovery systems allocate attention. LLMs do not suggest brand now matters more; they surface brand effects earlier in the customer journey, before activation begins. AI does not replace brand-building theory; it adds a new visibility layer that Marketing and Research leaders must measure. AI visibility appears to behave like a new form of availability. Brands must now be mentally available to people, physically available in channels, and machine-available to AI systems.



AI availability: a new concept

Brand visibility in AI-mediated environments increasingly depends both on broad machine legibility and how personal agents learn, retain and reproduce individual consumer preferences over time. This means that brands need to become machine-legible through websites, product data, reviews and FAQs. AI availability can be defined as the probability that an AI system can correctly identify, retrieve, understand, compare and recommend a brand in response to a relevant consumer prompt.

It is not the same as awareness. A brand can be famous among people but poorly represented in AI answers if its claims are unclear, its content is hard to parse, its proof points are weak, or its third-party corroboration is thin. It is also different from SEO. Traditional SEO helps a brand become discoverable in ranked search results. AI availability is broader. It includes whether the

brand is present in the sources AI systems trust, whether its information is structured, whether it is consistently described across the open web, and whether it can be confidently associated with particular category entry points. AI tools are not only drawing from historical training data, but many also now retrieve live information from the web. This means brands are not only competing inside model memory, but they are also competing inside the live information environment.

Why big brands will not always win AI search

The obvious assumption is that large brands tend to have more equity, therefore LLMs will surface large brands more often. That is directionally true, but it is incomplete. Large brands have several advantages. They tend to have more historical mentions, reviews, media coverage, backlinks, consumer discussion, marketplace presence and comparison content. These signals may increase both model familiarity and retrieval probability. If large brands receive more attention from LLMs, is that because of brand equity itself, or because larger companies have more resources to optimize their information environment? The answer is likely both.

Large brands may be more visible because their brand equity produces more cultural and digital traces. However, AI visibility can also be affected by tactical factors: FAQs, structured pages, product data, review environments, third-party articles, category guides, comparison content and paid placements.

M+C Saatchi Performance describes this as a shift where traditional SEO is no longer enough. Their analysis argues that brands can perform well in search but remain absent from AI answers if they lack corroboration across directories, review platforms, roundups and trusted third-party sources.

AI does not simply reward size. It rewards brands that are salient, sourceable, consistent and easy to justify. Scale helps, but structured evidence and credible third-party corroboration can alter the outcome.

Category Entry Points become prompt entry points

We know that in Sharp/Ehrenberg-Bass thinking, brands build mental availability by linking themselves to Category Entry Points. In an AI world, Category Entry Points increasingly

become prompt entry points. These prompts are modern expressions of category entry points. The AI system then acts as an intermediary, translating the user's need into a shortlist, explanation or recommendation.

This changes the research problem. **Brands should not only ask whether they are known, they should ask:**

- ❓ For which prompts does the brand appear?
- ❓ For which category needs is the brand associated?
- ❓ What language does the AI use to describe the brand?
- ❓ Which competitors are mentioned alongside it?
- ❓ Which sources are shaping that answer?
- ❓ Does the AI frame the brand how we want it to be framed?

This is the AI-era version of brand salience research. The unit of analysis moves from the survey association to the prompt-response environment.





Brand emotion still matters, and personal agents may increasingly operationalize it

AI systems may not 'feel' the brand, but they can detect evidence that humans felt something about the brand. This may become even more important as personal agents evolve. Rather than replacing emotional decision-making, these systems may increasingly act as proxies for accumulated human preferences and behaviors. For example, if agents learn that a user repeatedly selects sustainable brands, emotionally resonant aesthetics, familiar creators or trusted product categories, then brand storytelling continues to matter because

it shapes the behavioral signals the agent later uses for recommendations. There is a danger of a weak interpretation of AI branding that says emotion becomes less important as LLMs mainly process text. That is wrong. Emotion does not disappear. Rather, it becomes encoded through trust, confidence, reassurance, delight, reviews, sentiment, loyalty, complaint resolution and consistency over time. However, this does create a new problem for creative assets. Many brand pages rely on video, animation, music, visual metaphor and emotional storytelling. If these assets are not accompanied by transcripts, structured descriptions, captions, metadata or supporting text, then the emotional content may be under-readable to AI systems. **The key differences can be summarized as follows:**

“Build emotion for humans, because increasingly those human preferences may become encoded into machine-mediated decision environments”

Dimension	Real-world brand environment	LLM / agent environment
How emotion works	Emotion is felt directly by people through advertising, experience, memory and identity	Emotion is inferred from textual, behavioral and reputational signals
Primary mechanism	Attention, memory encoding, meaning, trust, desire, social identity	Sentiment, reviews, repeated descriptions, third-party validation, complaint patterns, testimonials
What matters most	Creative distinctiveness, storytelling, visual identity, music, tone, symbols, cultural relevance	Machine-readable evidence of how people describe and experience the brand
Risk	Weak emotional work fails to create memory or preference	Strong emotional work may be invisible to AI if it is not transcribed, described or evidenced
Brand advantage	Brands that make people feel something are more likely to be remembered and chosen	Brands whose emotional associations are consistently encoded across sources are more likely to be represented accurately
Research question	What does the brand make people feel?	What emotional meaning does the AI infer from the available evidence?



GEO, FAQs and the source layer

A strong FAQ is not merely a customer-service asset; it may become a key AI source. If AI systems retrieve concise, structured answers from brand-owned pages, then FAQs, comparison pages, product explainers and policy pages become part of the brand's machine-facing identity. However, owned content alone is insufficient. LLMs often

look for corroboration. A brand's own website tells the AI what the brand says about itself. Reviews, media, expert sources, marketplaces, and customer forums tell the AI whether others appear to agree.

That creates a three-layer model of AI brand evidence:

Layer	Role	Brand risk
Owned evidence	Brand website, FAQs, product pages, help centers, transcripts, claims	Too promotional, incomplete, hard to parse
Earned evidence	News, reviews, expert commentary, comparison sites, awards, forums	Negative sentiment, outdated information, inconsistent claims
Operational evidence	Availability, pricing, service quality, complaints, delivery, returns, loyalty, resolution	Brand promise contradicted by customer experience

How brands can influence decision-making agents

We are not "influencing agents"; we are engineering preference by being easiest to justify, safest to select, and best aligned with human outcomes the agent is accountable for. The nature of the agent also matters of course. Generic discovery agents may prioritize broad evidence, consistency and retrievability, whereas highly personalized agents may increasingly optimize around learned user preferences, habits, emotional associations and behavioral history. **We can engineer this preference through:**

-  **Becoming the lowest-risk option**
 Agents prioritize certainty. Brands with consistent performance, strong reviews, clear guarantees, and minimal complaints are safer to recommend.
-  **Winning the data layer**
 Agents rely on structured facts: pricing clarity, availability, specs, policies, ESG credentials. If it is ambiguous or missing, you're less likely to be chosen.
-  **Anchoring to clear use-cases**
 Make it easy for agents to map your brand to a specific job: "best for small teams," "most reliable," "premium but simple." Vague brands lose.



Designing for comparison

Agents compare. Provide transparent comparisons, calculators, proof points, and outcomes that let AI confidently justify choosing you.



Signaling trust through behavior, not claims

Fast resolution times, refund rates, repeat usage, loyalty, sentiment—these operational signals outweigh brand slogans.



Embedding where agents operate

Marketplaces, ecosystems, integrations, APIs, subscription defaults. Presence in an agent's workflow matters more than awareness reach.



Teaching the agent your brand meaning

Consistent language, positioning, and values across all surfaces help agents associate your brand with an emotional shortcut: safe, smart, aspirational.

Paid AI visibility complicates the picture

AI-mediated discovery is not going to remain a purely organic environment. **This matters because brand visibility in AI environments will probably develop across three layers:**

1. **Generated visibility** — the brand appears in the AI answer

2. **Cited visibility** — the brand, or a source discussing the brand, is linked or referenced.
3. **Paid visibility** — the brand appears through sponsored placements around or within AI-mediated journeys

These layers should not be collapsed into one metric. Being recommended by an AI answer is different from being cited as evidence. Being advertised next to an answer is different from being organically endorsed. Each may produce different levels of trust, persuasion and skepticism. For research leaders, this is an immediate measurement issue.

What this means for Insight teams and agency partners

The research function should help stakeholders understand how AI systems reconstruct their brand. That means moving beyond conventional tracking and adding a new layer of **AI-mediated brand intelligence**. Asking questions like:

Visibility questions

- Which brands appear across ChatGPT, Gemini, Perplexity, Copilot and Google AI Overviews for core category prompts?

Prompt-entry questions

- Which category entry points trigger the brand?

Source questions

- Which sources shape the answer: owned pages, FAQs, reviews, forums, media, expert sites, marketplaces, Wikipedia, YouTube, Reddit or ads?

Perception questions

- After exposure to an AI answer, how do trust, quality, relevance, emotional fit and purchase intention shift?

Competitive questions

- Which competitors are being surfaced instead, and why? Are they winning because of brand equity, better source architecture, clearer category associations, stronger reviews or better third-party validation?

Intervention questions

- If the brand improves FAQs, adds transcripts, structures product data, publishes comparison content or strengthens earned media, does AI visibility improve over time?

In follow-up papers we will discuss the type of research approaches to these challenges for insight teams and partners. We believe that brands do not only need GEO dashboards. They also need to understand the relationship between **brand equity, source architecture, and AI-mediated perception**.

Executive Summary and Strategic implications

AI does not make brand-building less important. It makes brand-building more infrastructural. The new challenge is **AI availability**: ensuring that brands are easy for AI systems to identify, retrieve, compare, trust and recommend. Personal agents may become continuity layers between human emotion and machine readability. LLMs can be understood as econometric systems trained in culture rather than sales. They expose whether brand equity is strong enough, clear enough and evidenced enough to be recommended by machines and highlight that we need to add AI Availability as a new additional focus for Marketers and Insight teams alike.

Strategic implications for CMOs and CIOs

1. Brand equity becomes a discovery asset, not just a persuasion asset

Impact on brand equity:

LLMs change where brand equity shows up. In AI-mediated

journeys, equity may also affect whether the brand is named, compared, cited or recommended inside AI answers. WARC's recent guide argues that generative AI reinforces the

need for strong brand fundamentals and suggests that long-term brand equity may explain a large share of LLM visibility. Brands with consistent proof points, strong reviews, and clear differentiation will win.

Implication:

Treat brand-building as part of the "discoverability budget", not a soft "upper-funnel" spend. We should defend and maintain long-term brand investment because it now supports both human mental availability and AI availability.

2. CMOs need to manage "AI availability" as a new brand KPI

Impact on brand equity:

LLMs add a third layer to mental and physical availability: AI availability. A brand must be easy for AI systems to identify, retrieve, understand, compare and justify. Brands are competing not only in memory but also in the live information environment.

Implication:

Build an AI-availability dashboard. Track whether the brand appears across ChatGPT, Claude, Gemini, Perplexity, Copilot and Google AI



Overviews for category prompts, comparison prompts, value prompts, emotional prompts and purchase-intent prompts. Measure not only whether the brand appears, but how it is framed against your Category Entry Point drivers. Over time, organizations may also need to understand how brands become encoded into personalized agent ecosystems through repeat behavior, loyalty, preference reinforcement, and emotional association.

3. Brand content must combine emotional resonance with machine-readable evidence

Impact on brand equity:

LLMs do not experience emotion directly. They infer

brand meaning from readable evidence. Brands must be machine-legible before they can become machine-recommended. Generative Engine Optimization research similarly argues that generated answers depend on how sources are selected, cited and synthesized, not just how pages rank.

Implication:

Create content that directly answers buyer questions,

compares options, and explains why your brand is the best choice. Create an evidence architecture around the brand. This means stronger FAQs, structured product claims, comparison pages, schema, review strategy, expert validation, video transcripts and clear proof points. The job is not to replace emotional branding with structured evidence. It is to ensure emotional positioning becomes behaviorally reinforced, machine-readable and externally verifiable. The practical rule is: build emotion for people, encode it for machines, validate it through experience.

4. Reputation and earned media become algorithmic inputs into brand equity

Impact on brand equity:

In the LLM world, brand image is partly reconstructed from

third-party sources. That means brand equity is no longer formed only through paid communication and direct experience. It is also shaped by the wider source environment from which AI systems retrieve and summarize.

Implication:

Stop treating PR, reviews, customer service, SEO, content and brand tracking as separate disciplines. CMOs need a cross-functional "source influence" strategy: identify which sources AI systems cite or appear to rely on, correct outdated information, improve owned explanations, strengthen earned validation, and reduce negative evidence caused by poor customer experience. This is brand management at the level of the machine-readable public record.

5. Paid AI visibility will not replace brand equity, but it will distort the measurement environment

Impact on brand equity:

AI discovery is already developing paid layers. Google has expanded ads in AI Overviews, and OpenAI began testing ads in ChatGPT for logged-in adult users on Free and Go tiers in the US, while stating that ads are separate from answers and do not influence ChatGPT's responses. This means AI environments will include organic recommendations, cited sources and sponsored placements.

Implication:

Separate three metrics: generated visibility, cited visibility and paid visibility. Do not let media teams report sponsored AI exposure as brand equity. Paid AI placements may create attention, but organic recommendation and positive framing are more meaningful equity signals. We should insist on AI visibility reporting to distinguish whether the brand was recommended, cited, advertised, merely mentioned or absent.

6. Key brand assets must become machine-readable

Impact on brand equity:

Distinctive brand assets in an LLM-mediated world, only influence AI visibility if they are translated into readable signals. Otherwise, the AI may reduce the brand to generic functional claims. However, as personal agents increasingly learn from behavioral history and user preference patterns, emotionally resonant brand assets may continue to influence recommendation systems indirectly through accumulated consumer behavior.

Implication:

Protect distinctive assets for human memory but ensure we encode them for machines through consistent descriptions, metadata, alt text, video transcripts, campaign explainers, FAQs, press kits and third-party commentary. The principle is simple: brand assets should not only be seen by people; they should be understood by AI.



Jon Arthurs
Managing Director &
Global Sustainability
Practice Lead at Toluna

About Toluna

Toluna is the leading global research and insights provider that empowers clients to make smarter data-driven decisions. For 25 years, we've partnered with our clients to deliver greater business impact through our advanced platform, end-to-end solution portfolio, deep industry expertise, and expansive global first-party panel.

From DIY to full-service consultancy, our unique approach can be tailored to fit any time, budget, and resourcing needs. From creative testing to brand tracking, and from packaging optimization to e-commerce understanding, Toluna's range of solution suites provide brands with the insights required to drive growth, relevance, and profitability.

With over 40 offices across the globe, we deliver research in 70+ countries to the world's leading brands. Recognized as a top innovative company by GRIT in 2024, we are at the forefront of insights innovation and excellence.

www.tolunacorporate.com

About the author

Jon is a Managing Director at Toluna and leads the Global Sustainability Practice at the Toluna Group and has been in the company for over 10 years with over 30 years industry experience gained in both Client-side roles and Consulting roles. Jon is passionate about developing products, brands and communications for the good of people, planet and profit. A strong believer in creating strategies to grow business that don't harm society, community and the environment through a pragmatic lens.

Want to know more?

Contact one of our experts →

References

Aaker, D.A. (1991) Managing brand equity: capitalizing on the value of a brand name. New York: Free Press.

Aggarwal, P., Murahari, V., Rajpurohit, T., Kalyan, A., Narasimhan, K. and Deshpande, A. (2024) 'GEO: Generative Engine Optimization', arXiv. Available at: <https://arxiv.org/abs/2311.09735>

Anthropic (2026) Demystifying evals for AI agents. Available at: <https://www.anthropic.com/engineering/demystifying-evals-for-ai-agents>

Binet, L. and Field, P. (2013) The long and the short of it: balancing short and long-term marketing strategies. London: Institute of Practitioners in Advertising.

Charlie Oscar (2026) What really drives LLM visibility. Cited in WARC (2026) The WARC guide: brand building in the age of Gen AI. Available at: <https://www.warc.com/en/the-warc-guide-brand-building-in-the-age-of-gen-ai>

Fournier, S. (1998) 'Consumers and their brands: developing relationship theory in consumer research', Journal of Consumer Research, 24(4), pp. 343-373.

Google (2026) About ads and AI Overviews. Available at: <https://support.google.com/google-ads/answer/16297775>

Google Cloud (2026) Evaluate Gen AI agents. Available at: <https://docs.cloud.google.com/vertex-ai/generative-ai/docs/models/evaluation-agents>

Google Search Central (2026) Video SEO best practices. Available at: <https://developers.google.com/search/docs/appearance/video>

Holt, D.B. (2004) How brands become icons: the principles of cultural branding. Boston, MA: Harvard Business School Press.

Kapferer, J.N. (2012) The new strategic brand management: advanced insights and strategic thinking. 5th edn. London: Kogan Page.

Keller, K.L. (1993) 'Conceptualizing, measuring, and managing customer-based brand equity', Journal of Marketing, 57(1), pp. 1-22.

M+C Saatchi Performance (2026) What really drives LLM search visibility? Available at: <https://www.mcsaatchiperformance.com/news/what-really-drives-llm-search-visibility/>

Omniscient Digital (2026) How LLMs source brand information: an analysis of 23,000+ AI citations. Available at: <https://beomniscient.com/blog/how-llms-source-brand-information/>

OpenAI (2024) Introducing ChatGPT search. Available at: <https://openai.com/index/introducing-chatgpt-search/>

OpenAI (2026a) Testing ads in ChatGPT. Available at: <https://openai.com/index/testing-ads-in-chatgpt/>

OpenAI (2026b) Working with evals. Available at: <https://developers.openai.com/api/docs/guides/evals>

OpenAI (2026c) Inside our approach to the Model Spec. Available at: <https://openai.com/index/our-approach-to-the-model-spec/>

Sharp, B. (2010) How brands grow: what marketers don't know. South Melbourne: Oxford University Press.

WARC (2026) The WARC guide: brand building in the age of Gen AI. Available at: <https://www.warc.com/en/the-warc-guide-brand-building-in-the-age-of-gen-ai>