

The Visibility Gap in AI: From Mentions to Occupancy

A Governance-Ready Framework for Measuring Brand Presence in AI Assistants

Executive Summary

AI assistants such as ChatGPT, Gemini, Claude, Perplexity, and Grok are fast becoming the primary interface for discovery and decision-making. As this transition accelerates, brand visibility inside these systems will shape demand capture, market share, and enterprise value.

Current “AI visibility” reports rely on **exposure-based metrics**: counting brand mentions in model outputs and reporting relative growth. These numbers often look impressive— “700% increase in visibility”—but they obscure strategic fragility. Exposure measures mention; it does not measure durable presence in answers.

This paper introduces the **Prompt-Space Occupancy Score (PSOS™)**: a governance-ready KPI that quantifies brand visibility across three dimensions—**occupancy, positioning, and decay**. PSOS closes the **Visibility Gap** between flattering exposure dashboards and the reality of how AI systems classify and recall brands.

1. The Visibility Gap

Definition

The **Visibility Gap** is the divergence between:

- **Human recall** → shaped by advertising, media coverage, word-of-mouth.
- **Model recall** → shaped by ingestion geometry across structured and unstructured sources: Wikipedia, Wikidata, Reddit, Quora, analyst PDFs, JSON-LD, and open web datasets.

Drivers

- **Source Asymmetry**: Paywalled media is under-ingested; Reddit/Wikipedia are overweighted.
- **Category Misalignment**: Brands define themselves one way (e.g., Ramp as “spend management”) but models classify differently (e.g., “corporate card”).

- **Anchor Absence:** Without structured data anchors, associations drift.
- **Decay Dynamics:** Mentions erode within 30–90 days due to model retraining, source updates, and competitive crowding.


2. Why Exposure Metrics Mislead

Exposure-based dashboards suffer from four systemic flaws:

1. **Relative Baselines:** Growth from 1% to 7% is presented as “700% increase.”
2. **Framing Bias:** Categories reflect client self-branding, not independent classification.
3. **Snapshot Bias:** Visibility is measured once; decay is ignored.
4. **Opaque Methodology:** No detail on prompts, weighting, or reproducibility.

Illustration:

Ramp’s 700% exposure growth conceals weak occupancy (32%), low positioning (rank 3.1), and rapid decay (–48% within 60 days). Coupa, by contrast, sustains 71% occupancy and minimal decay.



3. Introducing PSOS™ Dimensions

1. **Occupancy** → % of prompts where the brand appears in category answers.
2. **Positioning** → Average rank within answers (first vs buried third).
3. **Decay** → Rate of persistence across 30/60/90-day audits.

Methodology Overview

- **Prompt Clusters:** Thousands of stratified prompts covering consumer, B2B, and analyst use cases.
 - **Model Coverage:** ChatGPT, Gemini, Claude, Perplexity, Grok, benchmarked separately then normalized.
 - **Weighting:** Prompts weighted by commercial relevance.
 - **Auditability:** Logs reproducible under NDA; prompt sets transparent.
-

4. Comparative Case Studies

Fintech (Spend Management vs Corporate Card)

- Occupancy: Ramp 32% vs Coupa 71%.
- Positioning: Ramp rank 3.1 vs Airbase rank 2.0.
- Decay: Ramp -48% vs Coupa -8%.

Retail (Consumer Electronics)


- Exposure dashboards showed “visibility surge.”
- PSOS: 18% occupancy, 70% decay within 90 days → spike unsustainable.

Healthcare (Digital Health Platforms)

- Exposure metrics: flat mentions.
- PSOS: Mayo Clinic/WebMD >80% occupancy, challengers <20%. Anchored incumbency wins.

Industrial Procurement (Manufacturing B2B)

- Exposure dashboards overlooked.
- PSOS: Ariba consistently >65% occupancy; smaller platforms <25%. Illustrates how B2B laggards risk invisibility in enterprise prompts.



5. Why PSOS Matters

For Boards & Investors

- Demand capture risk: missing from prompts means lost share.
- Governance: PSOS provides a KPI analogous to GAAP or ESG.

For Marketing Teams

- Operational guidance: where to reinforce structured anchors.
- Category alignment: ensures AI recall matches positioning.

For Data Teams

- Benchmarking: track across assistants and languages.
- Decay tracking: quantify retraining and ingestion shifts.

6. Path Forward

1. **Audit Baseline:** Establish PSOS vs competitors.
 2. **Anchor Definitions:** Strengthen Wikipedia, Wikidata, schema markup, analyst reports.
 3. **Reinforce Prompt-Space:** Seed discourse ethically in ingestion-heavy forums.
 4. **Monitor Persistence:** Run quarterly 30/60/90-day audits.
 5. **Integrate Governance:** Add PSOS to board dashboards alongside financial/ESG KPIs.
-

7. Ethical Considerations

Seeding discourse on Reddit, Quora, or G2 must respect community norms. Brands should:

- Disclose affiliations where relevant.
- Focus on value-adding educational content.
- Avoid manipulative posting.

Ethical alignment strengthens credibility and ensures visibility strategies are sustainable.

8. Conclusion

Mentions fade. Occupancy endures.

Boards cannot rely on exposure dashboards to judge AI visibility.

The **Prompt-Space Occupancy Score (PSOS™)** closes the **Visibility Gap**, providing a standardized, auditable, and governance-ready framework.

PSOS turns AI visibility from a vanity metric into a measurable, board-level asset—ensuring brands are durably recalled in the answers that now shape markets.

Technical Appendix: PSOS Methodology

1. Prompt Clusters:

- 5,000+ prompts per audit, stratified by intent (informational, comparative, transactional).
- Taxonomy aligned to sector (e.g., procurement vs consumer retail).

2. Weighting Rules:

- Occupancy weighted by query frequency and commercial relevance.
- Example: “Top spend management platforms” = higher weight than “Who offers a Ramp card?”

3. Decay Formula:

- $Decay = (Occupancy\ Day\ 0 - Occupancy\ Day\ N) \div Occupancy\ Day\ 0.$
- Benchmarked at Day 30, 60, 90.

4. Audit Process:

- Multi-model testing (ChatGPT, Gemini, Claude, Perplexity, Grok).
 - Logs archived, reproducible under NDA.
-

Operational Annex: Practical Guidance

- **Marketers:**

- Validate schema.org with Google Structured Data Testing Tool.
- Commission independent analyst notes citing your category positioning.
- A/B test prompts internally to track brand recall.

- **Data Teams:**

- Benchmark across models separately (ChatGPT vs Grok vs Claude).
- Build decay dashboards to track persistence over time.

- **Community Engagement:**

- Publish thought-leadership posts on Reddit/Quora in alignment with community norms.
- Encourage customer reviews on G2 and Capterra using category language (“spend management platform”).