
Commerce Feeds vs AI Visibility: The Governance Gap in Agentic Shopping

Author: AIVO Journal

Date: October 2025

Abstract

OpenAI's introduction of the Commerce Feed specification marks a significant step toward enabling agentic shopping within large language models (LLMs) such as ChatGPT. By allowing merchants to submit structured product catalogs with frequent updates, the feed system appears to offer determinism in visibility and eligibility. However, this paper demonstrates that feed compliance alone does not guarantee sustained presence in LLM-mediated commerce.

Empirical audits conducted under the AIVO Standard™ show that brands with validated feeds can lose 30–60% prompt-space presence within 30 days, with competitors capturing abandoned slots. This volatility creates governance risk: boards relying solely on feed health reporting may make materially incomplete disclosures in quarterly financial and risk reporting.

We argue that visibility in AI assistants must be treated as a governance obligation, requiring independent, audit-grade verification. The Prompt-Space Occupancy Score (PSOS™), developed under the AIVO Standard, provides a closed-loop assurance framework. PSOS audits distinguish feed-backed presence from organic recall, track recency, and provide iterative remediation and verification. This positions PSOS as a governance KPI for e-commerce boards, CMOs, and operators.

Keywords

AI Visibility Optimization (AIVO); Prompt-Space Occupancy Score (PSOS); agentic shopping; commerce feeds; large language models; governance risk; e-commerce; audit assurance; OpenAI.

1. Introduction

The rise of LLM-mediated commerce represents a structural shift in digital discovery. OpenAI's Commerce Feed specification provides a standardized schema for merchants to

submit product catalogs directly into ChatGPT, including variants, pricing, and inventory, with updates as frequent as every 15 minutes.

On first impression, this appears to solve the discoverability problem. Yet empirical evidence shows that feed submission is necessary but not sufficient. Visibility in LLMs remains volatile, fragmented, and subject to substitution. This paper examines the governance implications of that gap and proposes a standard for audit-grade assurance.

2. Feeds as Eligibility, Not Presence

Commerce feeds validate schema compliance but do not guarantee outcome. Key limitations include:

- **Volatility:** Model retraining alters slot occupancy independently of feed validity.
- **Fragmentation:** Visibility differs across regions, product variants, and compliance constraints.
- **Substitution:** When a brand vanishes, competitors immediately occupy decision slots.

Thus, while feeds establish eligibility, they do not secure sustained presence in high-intent prompts.

3. Governance Implications

Boards equating feed submission with customer reach risk misstatement of material facts. Customer acquisition forecasts, marketing ROI, and forward-looking risk disclosures can all be distorted if LLM visibility decays without detection.

The governance analogy is clear: just as financial statements require external audits, AI commerce visibility requires independent verification. Dashboards that confirm feed health are insufficient for fiduciary duty.

4. Dashboards vs Standards

Existing commercial dashboards (e.g. Profound, Botify, Semrush) offer monitoring of schema validity and surface-level snapshots. These tools, while useful operationally, do not measure:

- Decay after retraining
- Distinction between feed-backed presence and organic recall
- Competitive substitution and share capture

The AIVO Standard differentiates itself by positioning PSOS as an audit KPI rather than a dashboard metric.

5. PSOS™ and Iterative Assurance

The Prompt-Space Occupancy Score (PSOS™) measures whether a brand or product occupies answer slots in LLMs across models, prompts, and geographies. PSOS has been updated to address commerce feed ingestion and recency:

- **Feed-backed checks:** Distinguish products surfaced via feed from those recalled organically.
- **Recency overlays:** Verify whether frequent feed updates translate into prompt-space visibility.
- **Closed loop:** Audit → remediate → verify → monitor. Each audit generates remediation guidance, which is then re-verified to confirm sustained correction.

This process provides boards, CMOs, and operators with a governance-grade assurance cycle.

6. Case Evidence

- **Retail:** In Black Friday audits, a global retailer with a validated feed failed to appear in 8 of 10 high-intent prompts. Amazon and Target occupied the slots, producing measurable cart-share leakage worth millions.
- **Travel:** A leading OTA disappeared from “best flight app” prompts after retraining. Competitors filled the slots, diverting bookings estimated in the tens of millions during peak demand.

These cases demonstrate that feed compliance did not guarantee visibility or protect revenue.

7. Conclusion

The OpenAI Commerce Feed spec represents an important infrastructural advance. However, it does not eliminate volatility or guarantee visibility in agentic shopping flows. Boards that rely on feed health reporting without independent verification expose themselves to governance risk.

The AIVO Standard, through PSOS audits, offers a closed-loop framework for measuring, remediating, and assuring visibility in AI commerce. As AI assistants become parallel discovery layers alongside search and retail media, PSOS should be treated as a governance KPI equivalent to audit assurance in financial reporting.

References

- OpenAI. (2025). *Commerce Feed Specification*. Retrieved from: <https://developers.openai.com/commerce/specs/feed>
- AIVO Journal. (2025). *Prompt-Space Occupancy Score: Metric for Auditing AI Visibility*. <https://zenodo.org/records/17116853>
- AIVO Journal. (2025). *From Dashboards to Standards: AI Visibility 2.0*. <https://zenodo.org/records/17169746>