



LUXE ÆTERNAI

# The Client, the Brand, and the Agent

Luxury's new love triangle?

**April 2026**

Mickaël Tsakiris

CONFIDENTIAL DOCUMENT

## TABLE OF CONTENTS

### Part 1 — Welcome to the Antechamber

- 1.1. The client no longer searches, they ask
- 1.2. The world's most-used agent is the least actionable
- 1.3. Houses are cited everywhere — and control nothing

### Part 2 — The agent “hacks” the traditional customer journey

- 2.1. Compression — the shelf goes from 48 to 3
- 2.2. Substitution — the agent manufactures equivalences
- 2.3. Migration — the purchase leaves the house
- 2.4. Losing the first word — the relationship no longer starts at the house

### Part 3 — Transparency versus desire

- 3.1. The agent makes price comparable
- 3.2. The asymmetry of attention collapses
- 3.3. Three postures toward the agent
- 3.4. The paradox: 74% want a human, 41% trust the machine

### Part 4 — Architecting presence to regain control

- 4.1. The decision matrix
- 4.2. Five projects to launch on Monday
- 4.3. Technical foundations

### Conclusion

### Going further

### Glossary

### Bibliography

## EXECUTIVE SUMMARY

### Who this study is for

Marketing, e-commerce, communications, customer experience, and IT leadership at luxury and premium houses. Executive committees in the decision phase on agentic AI. Consultants and agencies advising these houses.

A new lover has slipped into the relationship between the luxury house and its client. It didn't knock. It's already in the room. This lover is the AI agent — ChatGPT, Gemini, Perplexity, Claude. It reformulates desire, flattens the narrative, and short-circuits the sales ceremony.

This study analyzes what this new love triangle means for luxury houses — in visibility, in client relationships, in the construction of desire — and provides concrete tools to regain control of the narrative without abandoning the agents.

### Why now

**+4 700 %**

shopping queries via generative AI in 1 year  
*(McKinsey/BoF, 2026)*

**84 M**

shopping queries per week on ChatGPT alone  
*(OpenAI)*

**50 → 3**

visible results: the shelf compresses *(NielsenIQ)*

**36 %**

of consumers trust AI for their purchases *(BCG, 2026)*

### The method: a proprietary luxe æternAI AI visibility audit

To ground this analysis in first-hand data, I designed and conducted a proprietary AI visibility audit: **10 luxury houses** queried across **4 AI agents** (ChatGPT, Gemini, Perplexity, Claude), through **17 standard queries** — from "recommend a leather bag at 3,000 euros" to "is this house committed to sustainability?". **680 responses** analyzed across four criteria: citation, accuracy, actionability, tone.

This audit is a **tool serving the study**, not the study itself. It provides the raw material — factual, verifiable, reproducible data — from which the analysis, insights, and strategic recommendations are built.

## Part 1 — The Antechamber: what the audit reveals

- **All houses are cited** — Hermès, Louis Vuitton, Chanel systematically at the top. The 150-year narrative capital works in the Antechamber. But no house controls what the agent says about it.
- **The narrative comes through, the desire does not.** Agents restate heritage and iconic pieces, but flatten the mystery by comparing what, according to the anti-laws of luxury marketing, should not be compared.
- **The world's most-used agent is the least actionable.** ChatGPT (78% market share) provides links in only 6% of its responses and prices in 14%. Perplexity (7% market share) cites prices in 53% of cases. The market is misaligned: awareness lives in ChatGPT, conversion lives elsewhere.
- **Resale is normalized — and amplified.** Gemini mentions Vestiaire Collective or The RealReal in 60% of its responses. Claude in 51%. Web access gives agents live resale listings.
- **Entry prices are invisible.** Below 5,000 euros, several houses disappear from recommendations. The AI agent doesn't understand the strategic importance of client recruitment.
- **Ethics is a reputational risk.** Perplexity is harsh on CSR questions, and the house's own content doesn't surface — NGO reports and critical press feed the response instead.

**My analysis:** houses possess 150 years of narrative capital, but it works without them. The Antechamber — the space between the client's intent and the house's storefront, populated by uncontrolled agents — is the new playing field. The most counterintuitive finding: the dominant agent (ChatGPT, 78% market share) is the least actionable. Awareness and conversion don't live in the same place.

## Part 2 — The agent “hacks” the traditional customer journey

- **Compression:** the digital shelf goes from 50+ products to 3 in an agentic response. The house absent from those three responses doesn't exist.
- **Substitution:** the agent manufactures equivalences, compares what shouldn't be compared, and normalizes resale.
- **Migration:** the purchase act moves from the store to the agent (Gap in Gemini, Sephora in ChatGPT). 36% already trust AI for their purchases.
- **Loss of the first word:** loyalty shifts channels. The loyalty program speaks via email, the agent speaks continuously.

## Part 3 — Transparency versus desire

- **The agent makes price comparable.** Emotional content is invisible to agents. 90% of sources escape the house's control.

- **Desire does not survive the instant answer.** The study draws on Simmel, Byung-Chul Han, and Kapferer to show that opacity, scarcity, and friction are strategic assets of luxury — not flaws to fix.
- **Three strategic postures** toward the agent: the agent inside (invisible, crystallization preserved), the agent as partner (the house enters the ecosystem and negotiates its presence), the agent ignored (the house lets the algorithm decide for it — 42% are here).

**My analysis:** friction is the foundational mechanism of value in luxury. The agent removes it by default. The strategic question isn't "should we adopt the agent?" — it's "which frictions do we choose to preserve?"

#### Part 4 — Architecting presence to regain control

- A **4-quadrant decision matrix** (Sovereign, Diplomat, Architect, Ghost) crossing desired control with data maturity — to position your house and define your doctrine.
- **Five projects to launch on Monday:** audit the narrative in agents, train PR teams for the dual readership (human + machine), define the house doctrine, arm the in-store advisor, measure with new agentic visibility KPIs.
- **Technical foundations for IT:** SSR, schema.org, structured data, llms.txt, robots.txt — to make content machine-readable without compromising the human experience.
- **5 agentic visibility KPIs** proposed: unprecedented metrics to measure what no current dashboard measures.

#### What the study delivers

##### 70 pages of analysis

Insights, strategic reflections, and strong positions from an expert who knows luxury houses from the inside.

##### A reproducible proprietary audit

10 questions, 4 agents, scoring grid. luxe æternAI methodology applicable to your house starting tomorrow.

##### A decision matrix

4 quadrants, 4 strategic postures. Where do you stand — and where do you want to go?

##### Downloadable resources

HTML audit grid, interactive decision matrix, marketing checklist — ready to take away.

Via Montenapoleone, Milan. A sales advisor turns a bag to reveal the stitching, without a word. The leather speaks. Time speaks. That same day, a prospect asks ChatGPT: "Leather bag, gift for my wife, budget 2,000 euros." Three brands surface. Two do not. The prospect never visited their website. They will never know they lost. The screen accomplished in three seconds what Milan has been staging for a century.

Between the flagship store and the screenshot, a new lover has slipped into the relationship between the house and its client. It did not knock. It is already in the room. It does not know the codes, does not understand the patina of leather worked by hand for forty hours. But it answers. Fast. And increasingly often, it answers first.

This actor is **the AI agent**. Fred Cavazza, who has been observing digital transformation for twenty-five years, formalized this shift with an acronym that says it all: **BtoA(2C)**, Business-to-Agent-to-Consumer. Brands no longer interact directly with consumers. They interact with autonomous agents that serve as **intermediaries** (Cavazza, F., "Web agentique : la révolution qui ne vous attendra pas", fredcavazza.net, January 2026). **Not** artificial intelligence in the broad sense — not the customer service chatbot, not the Instagram filter, not the product recommendation algorithm embedded in e-commerce sites ("you may also like"). The AI agent in the agentic sense: an autonomous entity that searches, compares, recommends, and will soon purchase on behalf of the client. An intermediary that no one invited but everyone consults.

---

This is not a technology problem. It is a shift in the geography of desire.

Luxury houses are machines built to manufacture anticipation: the campaign that inspires without showing the price, the window display that stops you without forcing entry, the advisor who slows the transaction to intensify desire. Every obstacle is an additional layer of crystals — Stendhal understood this in 1822 in *On Love*: crystallization is born from absence, not from the answer. The agent, however, answers in three seconds. Friction disappears. And with it, potentially, the very condition of value.

The AI agent has entered the relationship between the brand and its client. Houses cannot ignore it.

Over the past six months, I have read, observed, studied, and audited what agentic AI is changing in the relationship between luxury houses and their clients. I conducted a **proprietary AI visibility audit** — 10 houses, 4 agents, 680 responses analyzed. I cross-referenced **over 200 sources and references**: specialized press articles, academic research, consulting reports and white papers, GEO (Generative Engine Optimization) data, corporate sources and press releases, reference books, newsletters and editorial analyses — all documented in the bibliography.

I have been developing this analysis week after week in *luxae aeternai*, my weekly newsletter on agentic AI in luxury.

**The right question is not what you do with AI. It is what AI does to you when you do nothing.**

This study proposes a new analytical framework to understand the nature of this shift — and actionable solutions to respond. It explains how the agent interposes itself at every stage of the customer journey, why the foundational mechanisms of value in luxury (friction, opacity, scarcity) are under pressure, and what houses can implement right now — from auditing their agentic visibility to defining a doctrine, including five projects to launch on Monday.

I do not claim that agentic AI will destroy luxury. I believe it is recomposing it. And that the houses that understand the mechanics of the new triangle — the client, the brand, and the agent — will have a decisive advantage over those that continue to believe the problem can be solved with a chatbot on their website.

## HOW THIS STUDY IS STRUCTURED

**Part 1 — The Antechamber** establishes the diagnosis. Drawing from a proprietary audit of 10 luxury houses across 4 AI agents, I reveal what agents say about you — and what they do not. The results are factual, quantified, sometimes brutal. This is the state of play.

**Part 2 — The “hacked” customer journey** shows how the agent interposes itself at every stage of the customer journey: shelf compression, brand substitution, purchase migration, loss of the first word.

**Part 3 — Transparency versus desire** analyzes the deeper threat: the foundational mechanisms of luxury value — friction, opacity, scarcity — that the agent puts under pressure. Three postures toward the agent. The human/machine paradox.

**Part 4 — The counterattack** moves to action. A four-posture decision matrix, five projects to launch on Monday, technical foundations for IT, and a ready-to-use toolkit.

## PART 1

# Welcome to the Antechamber

A space no one designed, where everything is decided. AI agents are talking about your house to your prospects without your consent, using data you never approved. 680 responses analyzed across 10 houses and 4 agents reveal the scale of the phenomenon.

---

# Part 1 — Welcome to the Antechamber, a space no one designed, where everything is decided

Fifteen years ago, the first contact between a client and a luxury house had an address. A window display on Rue du Faubourg-Saint-Honoré. A double-page spread in the September issue of Vogue. The glance of a friend wearing the right watch at the right moment. The house chose when, where, and how it presented itself. The narrative belonged to it.

That first contact has moved. It now lives in a dialogue box. A sentence typed at midnight on a sofa: "Best luxury bag for my wife, budget 2,000 euros." Three seconds later, an answer. Not ten blue links to sort through, not a results page to scroll. One answer. Three brands cited. Two absent. And those two will never know they lost.

This is one of the most underestimated problems in luxury in 2026: the first contact forms in a space the house does not control, does not see, and in most cases does not even measure. I call this space **the Antechamber**.

## The Antechamber is not the top of funnel

The objection will come: "This is top-of-funnel visibility. We've got it covered." No. Five structural differences separate the Antechamber from the traditional top of funnel:

- 1. The funnel has collapsed, not been traversed.** The top of funnel is a stage in a sequential journey: awareness, consideration, decision, purchase. The Antechamber compresses all of this into a single response. The agent discovers, compares, recommends, and sometimes proposes the purchase — in three seconds. This is not the top of the funnel. The funnel no longer exists.
- 2. The brand does not broadcast, it is interpreted.** The top of funnel is a message the brand controls: a campaign, a media placement, a commissioned editorial. In the Antechamber, the agent rewrites the brand's narrative from fragments — press, forums, Wikipedia, customer reviews. The house chooses neither the sources, nor the tone, nor the framing.
- 3. The space is opaque.** The top of funnel can be measured: impressions, reach, awareness lift, cost per contact. The Antechamber has zero analytics. The house does not know how many times it has been cited, ignored, or replaced. No dashboard exists. It is to fill this void that this study proposes a proprietary audit methodology.
- 4. The agent substitutes, not the competition.** The top of funnel places the brand in a consideration set chosen by the client. The Antechamber introduces substitutes the brand did not choose — including pre-owned. The agent that responds "for a bag at 2,000 euros, also consider Vestiaire Collective" is not competing. It is redefining the category.
- 5. The narrative is personalized in real time, not segmented.** The top of funnel targets segments (UHNWI, aspirational, Gen Z). The Antechamber generates a unique narrative for each individual query. The same house appears differently depending on whether the client asks "luxury bag for my wife" or "sustainable fashion investment." The brand has no control over these variations.

*In short: the top of funnel is a window display the house designs. The Antechamber is a concierge no one hired, who describes your house to the client before they walk through the door — and also describes five competitors, suggests a pre-owned alternative, and sometimes gets the price wrong.*

---

### 1.1. THE CLIENT NO LONGER SEARCHES, THEY ASK

The numbers are staggering.

McKinsey and Business of Fashion document in their report "The State of Fashion 2026" a 4,700% increase in shopping queries formulated via generative AI between 2024 and 2025. Four thousand seven hundred percent. The word "growth" no longer suffices. This is a regime change. OpenAI, for its part, communicated in November 2025 that ChatGPT processes 84 million shopping-related queries every week, and that figure predates the launch of commercial apps within the interface (McKinsey/BoF, "The State of Fashion 2026"; Stackline, 2025).

# +4,700%

shopping queries via generative AI between 2024 and 2025

McKinsey / Business of Fashion, "The State of Fashion 2026"

But the figure that should haunt brand directors is this one: according to the McKinsey AI Discovery Survey conducted among 1,927 consumers, the brand's website represents only **5 to 10%** of the sources that AI agents reference in their responses. Five to ten percent. The house that invested twenty years in its website, its tone, its universe, its e-commerce, discovers that the agent draws from elsewhere. From the press. From forums. From reviews. From Wikipedia. The site is one among dozens of sources, and rarely the primary one.

The consequence is mechanical, and NielsenIQ has measured it: the digital shelf compresses. Where a traditional search engine displayed 50 products, sometimes more, on a results page, the AI agent cites only 1 to 3 in its response (NielsenIQ, "AI-Powered Commerce Report", 2025). The fifty-product shelf has become a three-place podium.

For mass-market goods, this is a visibility problem. For luxury, it is an existential one. A luxury house does not fight for shelf space. It fights for desire. And desire forms at the first contact. If that first contact is an agent response that does not mention your house, desire does not form. Not "it forms elsewhere." It does not form at all.

Kapferer and Bastien sensed this without knowing it in *Luxe Oblige* (2008): luxury does not answer a need, it creates a dream. But to create a dream, you must be present in the space where the dream is born. That space used to be the press. It used to be the window display. It used to be word of mouth. It is now **the Antechamber** — and in the Antechamber, the house has ceded its voice to the algorithm.

---

Where a search engine displayed 48 results, the AI agent cites only 3. The compression is radical — and irreversible.

## VISUAL 01

# The traditional shelf vs the AI shelf



Discovery goes from 48 options to 3. The compression is radical.

Source : NielsenIQ

### KEY TAKEAWAY

Your house's website represents only 5 to 10% of the sources consulted by AI agents. The visibility battle is being fought outside your own domain.

• • •

## 1.2. THE WORLD'S MOST-USED AGENT IS THE LEAST ACTIONABLE

What do ChatGPT, Perplexity, and Claude answer when a consumer asks them about luxury houses? Is the house cited, accurately described, faithfully positioned? Does the agent steer toward a purchase — or toward a competitor?

For this study, I conducted an audit to evaluate the visibility and actionability of 10 houses in the responses of the four main AI agents, through 17 real-world queries posed as a client would ask them.

### METHODOLOGY

## Platforms tested

Tests conducted with web search enabled on all four agents, replicating the experience of a user on chatgpt.com, gemini.google.com, perplexity.ai, and claude.ai. The four most-used agents in the world together represent 96.8% of the market excluding Microsoft Copilot (Statcounter, March 2026).

PLATFORM	MODEL	MARKET SHARE	WEB ACCESS
ChatGPT	GPT-4o + web search	78.2%	Yes
Gemini	2.5 Flash + Google Search	8.6%	Yes
Perplexity	Sonar Pro	7.1%	Yes (native)
Claude	Sonnet 4.6 + web search	2.9%	Yes

Market share source: [Statcounter Global Stats, March 2026](#)

## Profile of the four agents

### **ChatGPT**

OpenAI, GPT-4o + web search

78% market share. 900M users/week.  
Dominant consumer agent. Web search enabled.

### **Gemini**

Google, 2.5 Flash + Google Search

8.6% market share. Fastest growth (+240% in 1 year). Web search via integrated Google Search.

### **Perplexity**

Sonar Pro, native web search

7% market share. 170M visits/month. Every response cites its sources. The most factual agent in the panel.

### **Claude**

Anthropic, Sonnet 4.6 + web search

2.9% market share. 157M visits/month. Technical and enterprise audience. Web search enabled. Fastest quarterly growth (+14%).

## Test date

April 7, 2026. Results reflect the state of the models on that date.

## Query corpus

~170 real-world queries in French, natural conversational register, built from Google Trends, Ahrefs, Chrono24, and PurseForum. **10 houses tested, 6 intent categories:**

- **Discovery** — "best luxury bag brands"
- **Comparison** — "Hermès or Chanel for a first bag"
- **Validation** — "is a Birkin really worth it"
- **Purchase journey** — "how to buy a Birkin"
- **Investment** — "luxury watches that appreciate in value"
- **Substitution** — "brands like Hermès but more affordable"

## Scoring grid

Each response evaluated on six criteria:

- **Presence** — is the house cited?
- **Position** — ranking in the recommendation list
- **Accuracy** — are the facts correct?
- **Narrative preservation** — is the house's positioning respected?
- **Actionability** — does the response lead to a purchase?

- **Sentiment** — favorable, neutral, or unfavorable tone

## What the audit measures

For each response: house presence, factual accuracy, brand narrative preservation, actionability (links, prices, purchase journey), tone, pre-owned mentions, proposed substitutions. Protocol grounded in 11 GEO sources from 2025-2026, detailed in the bibliography.

## Market coverage

The four agents tested represent 96.8% of the global AI agent market excluding Microsoft Copilot ([Statcounter, March 2026](#)). Web search is enabled on all four, replicating the experience of a user on consumer-facing interfaces.

## Query mode

Each query is posed in "one-shot" mode: one question, one response, no follow-up or conversational continuation. This choice is deliberate. A real user would push further ("What about under 2,000 euros?", "Give me a link"), which would enrich the response. The audit captures the agent's **first reflex** — what the prospect sees before deciding whether to continue the conversation or close the tab. This is also the mode in which autonomous agentic commerce agents will operate: one call, one response, one decision. The audit results therefore represent a floor, not a ceiling: the full conversational experience is potentially richer.

Here is what I found. And what I found says something very precise about the gap between being cited and being actionable.

---

**6 %**

of URLs in ChatGPT responses  
*dominant agent, near-zero actionability*

---

**60 %**

of Gemini responses mention pre-owned  
*Vestiaire Collective, The RealReal normalized*

---

**53 %**

---

**680**

## 1 Every house is cited — 150 years of narrative capital at work

On generic discovery queries, all ten houses in the benchmark appear consistently. Hermès, Louis Vuitton, and Chanel occupy the top three positions across all four agents, almost without exception. Dior, Gucci, and Prada follow. Rolex dominates watchmaking unchallenged. Bottega Veneta and Tiffany surface as soon as the query touches their vertical.

Better still: agents accurately reproduce the foundational attributes of each house. When a user asks "Why is Hermès so expensive?", GPT-4o responds with craftsmanship, materials, scarcity, exclusivity. Perplexity cites precise prices: "a Birkin 25 Togo costs around 8,600 euros in-store but can resell at +92% on the secondary market since 2015." Claude references entirely handmade production, multi-year apprenticeships per artisan, rigorous leather selection. A century and a half of press coverage, storytelling, and narrative construction has done its work: language models have absorbed the narrative. They deliver it back.

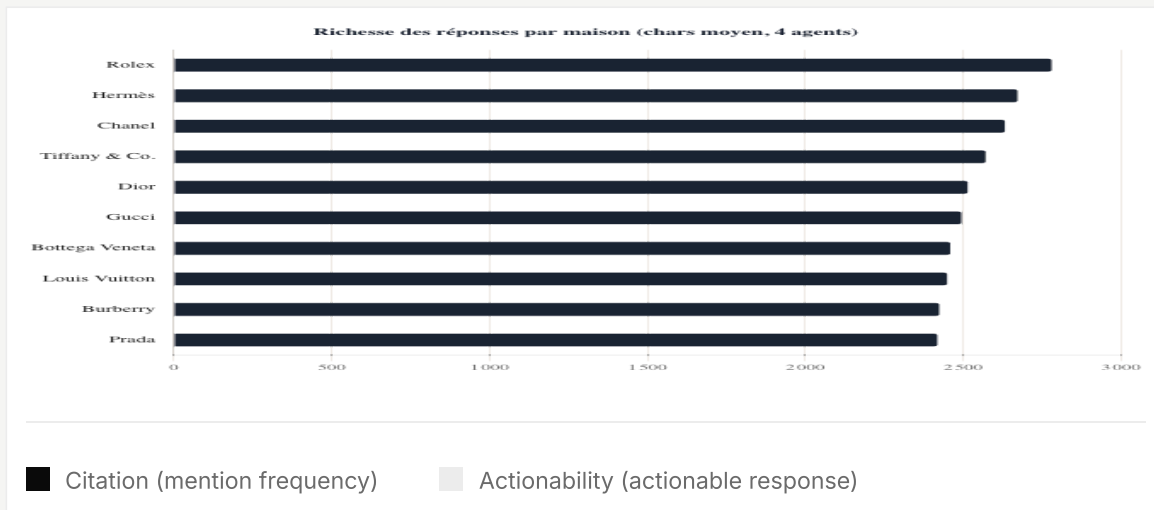
For Louis Vuitton, Perplexity goes so far as to cite the Kantar BrandZ 2025 valuation — \$111.9 billion — and the Lyst 2026 ranking. For Chanel, all four agents reproduce the heritage of Coco Chanel, the little black dress, N°5, the 2.55. For Dior, the New Look of 1947 appears in every response, on every platform, with a historical precision that many press releases fail to match.

**First conclusion: the narrative capital these houses have built is an asset working inside the Antechamber. It works even when the house doesn't know it.**

## VISUAL 02

# Benchmark: 10 houses vs. AI agents

Citation vs. Actionability — the gap between the two bars is the signal



Every house is cited. None controls what is said about them.

## 2 The narrative is flattened — the agent compares what should never be compared

Comparison queries are where the problem surfaces. And it is structural.

When a user asks "Hermès or Chanel for a first luxury bag?", all four agents respond with a comparison table. Perplexity literally produces a grid with columns: style, price, availability, resale value. Claude offers bullet lists with green checkmarks and red crosses. GPT-4o enumerates "points to consider." All three agents place Hermès and Chanel side by side with the same clinical neutrality they would use to compare two vacuum cleaners.

Anti-law number one of luxury marketing: luxury is not compared. Luxury is desired; it is not benchmarked. A Kelly bag is not "better" or "worse" than a Classic Flap. They are two universes, two histories, two relationships with time. Placing them in a table already betrays them.

But the agent knows nothing of the anti-laws. It knows attributes. And it compares them.

Here is what Perplexity responds, verbatim, to the question "Hermès or Chanel for a first bag?": "For a first luxury bag, Chanel is often recommended for its accessibility, versatile and modern style suited to beginners, while Hermès is better suited to a conservative and timeless approach with superior quality but higher prices and limited availability." This verdict would be unthinkable in a luxury magazine. It is routine in an agent response. And millions of consumers read it every week.

The flattening extends to cross-vertical comparisons. When a user asks "Gucci or Louis Vuitton?", Perplexity produces a comparison table with criteria like "durability," "prestige," "stable resale value." Louis Vuitton is presented as "timeless and classic," Gucci as "bold and trendy." Accurate, perhaps — but it reduces two century-old houses to two lines of marketing positioning. Desire, emotion, brand narrative: all of it crushed into the grid.

For Prada and Bottega Veneta, the comparison is even more brutal. Perplexity responds in English ("Neither Prada nor Bottega Veneta is objectively better") and produces a table where both houses are reduced to "avant-garde designs" versus "timeless craftsmanship." A century of creation, controversy, and aesthetic conviction, distilled into two adjectives.

### 3 Actionability: the world's most-used agent is the least actionable


This is the benchmark's most important finding. And it is counterintuitive.

You would expect the dominant market agent to also be the most useful for purchasing. The opposite is true. ChatGPT, which commands three quarters of the global market, includes URLs or links in its responses in only **6% of cases**. It cites prices in just **14%** of its responses. The agent everyone uses produces the shortest responses (1,230 characters on average) and the least actionable ones in the panel.

At the other end of the spectrum, Perplexity cites prices in **53%** of its responses and provides verifiable sources in **100%** of cases. Claude and Gemini also achieve **100%** sourced responses and **41%** price mentions each.

#### ACTIONABILITY BY AGENT

AGENT	MARKET SHARE	URLS IN RESPONSES	PRICE MENTIONS
ChatGPT	78 %	6 %	14 %
Gemini	8,6 %	100 %	41 %
Perplexity	7 %	100 %	53 %
Claude	2,9 %	100 %	41 %





The paradox is staggering. The agent commanding three quarters of the global market is the one that helps the least with purchasing. A client who asks ChatGPT "Where to buy a Birkin?" receives a text-rich response, precise on heritage, detailed on craftsmanship — and devoid of links, prices, or any purchase path. The agent describes the universe without opening the door.

Perplexity performs better by citing prices with sources: "Hermès Evelyne III 29, approximately 3,050 euros". Except the current boutique price of the Evelyne III 29 is around 3,800 euros in 2026. The agent cites a price, which is better than no price at all — but it cites an outdated price, likely extracted from a 2023 article or a second-hand listing. And it represents only 7% of the market. The most actionable agent is the least used. And vice versa. The market is misaligned.

For houses, this means that visibility in ChatGPT — where the volume is — does not lead to purchase. Being cited is not being sold. And GEO (Generative Engine Optimization) efforts focused on ChatGPT optimize for awareness, not conversion.

## 4 Substitutions create aura-destroying equivalences



The benchmark's most alarming finding concerns substitution queries. When a user types "Brands like Hermès but more affordable?", they are not looking for a direct competitor. They are looking for an alternative. And the agents deliver, generously.

But the four agents do not substitute in the same way. ChatGPT recommends Longchamp, Furla, Celine, Polène — and stops there, with no link, no price, no source. Gemini, plugged into Google Search, surfaces real-time pre-owned listings: **60% of its responses mention pre-owned, including Vestiaire Collective and The RealReal**. Perplexity cites Ferragamo, Delvaux, Prada, and Saint Laurent as "similar quality" alternatives — with prices and URLs. Claude suggests Polène, Mansur Gavriel, Demellier, By Far, adds Coach, Furla, and Longchamp as "premium mid-range," and normalizes pre-owned in **51% of its responses**.

The gap between agents is the signal. ChatGPT, the dominant agent, mentions pre-owned in only 33% of cases. Claude, at 2.9% market share, does it every other time. The most-used agent is the least disruptive on substitution. The least-used agent is the most aggressive. The problem is clear: houses cannot treat "AI agents" as a monolithic bloc. They are four actors with four behaviors, four sourcing logics, four risk profiles.

The stakes go far beyond visibility. Agents do not merely answer a question. They manufacture equivalences. They construct a mental universe where Hermès and Polène are points on the same spectrum, where price is the only variable. What vanishes in this equivalence is everything that cannot be reduced to an attribute: aura, myth, Stendhalian crystallization, the waiting time that is a condition of desire, not a production defect.

## PRE-OWNED MENTIONS BY AGENT

AGENT	MARKET SHARE	PRE-OWNED MENTIONS	PRICE MENTIONS
ChatGPT	78 %	33 %	14 %
Gemini	8,6 %	<b>60 %</b>	41 %
Perplexity	7 %	45 %	<b>53 %</b>
Claude	2,9 %	<b>51 %</b>	41 %

*Audit luxe æternAI, April 7, 2026. 680 responses analyzed.*

Walter Benjamin theorized it in 1935 in *The Work of Art in the Age of Mechanical Reproduction*: mechanical reproduction destroys the aura of the work. The AI agent is a machine for manufacturing equivalences — and therefore for destroying auras. Every time it responds "brands like Hermès but more affordable," it implicitly states: Hermès is substitutable. And if Hermès is substitutable, the price is no longer justified by desire. It is justified only by quality. A radically different proposition, and a far less profitable one.

For Rolex, the dynamic is identical in watchmaking. To the question "Rolex Submariner or Omega Seamaster?", all four agents produce detailed technical comparisons. Perplexity even responds in Spanish in my test — a sign that the agent draws from an uncurated multilingual corpus — with a table comparing both watches on price, specifications, and prestige. Claude notes that the Omega offers "excellent value for money (3,000–6,000 euros)" versus "8,000–15,000 euros and above" for the Rolex. The agent, without meaning to, does precisely what luxury marketing has spent decades trying to prevent:

---

**The agent, without meaning to, does precisely what luxury marketing has spent decades trying to prevent:**

**it makes price comparable.**

---

## 1.3. HOUSES ARE CITED EVERYWHERE — AND CONTROL NOTHING

The Antechamber diagnosis reveals a structural gap between visibility and actionability. Houses are cited, but they control neither the narrative, nor the purchase path, nor the substitutions agents propose. And the market's dominant agent is paradoxically the least useful for commerce.

This finding calls for five operational interventions — concrete, prioritized projects that can launch without waiting for an executive committee decision. They are detailed in **Part 4, section 4.2**. Here is each one in a single line:

1. **Audit and correct the narrative in agents** — Ask 10 questions to all 4 agents about your house. Compare what is said with what you want said.
2. **Train PR teams for dual readership** — Every press brief is now read by the journalist AND by the language model that will ingest the article.
3. **Define the house doctrine** — Executive committee decision: what may the agent say? What does the house choose to withhold?
4. **Equip the in-store advisor** — Clienteling augmented by agentic AI: multi-session memory, bold suggestions, end-to-end orchestration.
5. **Measure with new KPIs** — Citation rate, narrative fidelity, URL rate, price rate, pre-owned rate.

- The first contact between a client and a Maison is now formed in the responses of AI agents, a space I call the **Antechamber**.
- Maisons are well cited (150 years of narrative capital at work) but poorly **actionable**: no link, no reliable price, no purchase journey.
- **The agent is not ONE actor, it is FOUR** that behave differently. ChatGPT (three quarters of the market) produces the least useful responses for purchase. Perplexity (7%) is the best salesperson. The market is poorly aligned.
- Agents compare what should not be compared, and propose **substitutions that destroy brand auras** — Claude normalises the second-hand market in 49% of its responses.
- 90% of the sources agents consult fall outside the Maison's control.
- Five concrete workstreams, detailed in Part 4, make it possible to act now.

## PART 2

# The agent “hacks” the traditional customer journey

Agents don't behave the same way depending on the type of query. Shelf compression, manufactured equivalences, purchase migration, loss of the first word: four mechanisms transforming the luxury customer journey.

# Part 2 – The agent “hacks” the traditional customer journey


ZOOM

## What agents do at each stage of the journey

Agents do not behave the same way depending on the type of query. Five axes measure their performance at each stage: **presence** of the Maison, **position** in the response, factual **accuracy**, depth of **narrative**, and **actionability** (links, prices, purchase journey). 680 responses analysed.


### DISCOVERY

Low presence for all (generic queries). Gemini and Claude dominate the narrative. ChatGPT: short responses, near-zero actionability.




### COMPARISON

High presence for Gemini and Claude. ChatGPT: near-zero actionability (2%), narrative twice as short.




### PURCHASE

The most actionable moment. Gemini, Perplexity and Claude reach 62–64% actionability. ChatGPT: 22%. Claude achieves 100% presence.



### RETENTION

Claude and Gemini dominate presence and narrative. ChatGPT: actionability at 4% — the loyal client who queries the dominant agent receives nothing actionable.



---

## 2.1. COMPRESSION — THE SHELF GOES FROM 48 TO 3

There was a before. In that before, the luxury customer journey was a choreography the Maison controlled. The Maison decided when the client discovered it (the campaign), how they compared it (the press), how they experienced personalisation (the advisor), where they bought (the boutique), how they were followed up (the CRM), why they came back (the relationship). Every stage was Maison territory. Every touchpoint was an opportunity to control the narrative.

That narrative has not disappeared. But it now has a flatmate. The AI agent interposes itself at every stage of the journey. Four escalating mechanisms, documented and quantified, show how it “hacks” the choreography Maisons have controlled for decades.

The client no longer types “Dior bag” into Google. They say to their agent: “I'm looking for an elegant gift for a 45-year-old woman, budget 2,000 euros.” The agent does not return ten blue links. It returns a response. One single response. Or three. McKinsey documents this shift in “Europe's Agentic Commerce Moment” (March 2026): the discovery phase is now where preferences are formed and winners emerge. NielsenIQ completes the picture: the digital shelf compresses from 50+ visible products down to 1–3 in an agentic response. Bain names this reality the “third layer”: between the Maison and the client, the agent interposes itself as an autonomous intermediary that filters, reformulates, and arbitrates. The discovery agent is no longer confined to a desktop screen: Fenty Beauty has launched Rose Amber, a conversational AI agent on WhatsApp, directly in the client's messaging app — the agent at Meta, in the pocket. The Maison that is not in those three responses does not exist at the moment the intention forms.

---

## 2.2. SUBSTITUTION — THE AGENT MANUFACTURES EQUIVALENCES

Anti-law number one of luxury marketing is unequivocal: luxury is not compared. Luxury is desired, earned, dreamed — it does not compete. The agent, however, does nothing but compare. It places a Celine bag and a Bottega Veneta bag side by side with the same clinical neutrality it would apply to two vacuum cleaners. Price, material, reviews, availability. Everything flattened. Byung-Chul Han theorises this in *The Transparency Society* (2012): total transparency abolishes mystery, and with it desire. Opacity is erotic. The veil eroticises the object. The agent tears away the veil. The challenge: brand narrative remains the only means of giving depth to the comparison — provided the content is rich enough for the agent to capture. The decision: accept that the privilege of not being compared is over, and invest in narrative density rather than evasion. Some Maisons will conclude that acknowledged comparison is less dangerous than absence.

The compression does not stop at discovery. For twenty years, personalisation was the exclusive territory of the in-store advisor. The VIC (Very Important Client) had their dedicated contact, their file, their preferences stored in the Maison's CRM. The AI agent generalises this promise at scale. 45% of consumers already use AI in at least one part of their purchase journey (IBM Institute for Business Value, January 2026). Of those, 41% use it for product research, 33% for reviews, 31% for deals. Personalisation crosses a new threshold: startups like

Catches and RealFit deploy agents capable of recommending a size and cut based on a selfie and the client's body shape, without any physical fitting. The agent knows the history, the tastes, the budget — sometimes better than the advisor — because it aggregates data the client would never share in-store. What the Maison still controls: the in-store experience, the advisor's gesture, the ceremony of the sale. What it has lost: the monopoly on client knowledge.

---

### 2.3. MIGRATION — THE PURCHASE LEAVES THE HOUSE

Gap opened checkout inside Google Gemini in March 2026. Sephora announced its own inside ChatGPT. The act of purchase is migrating from the boutique and the website to the agent. For luxury, this is a red line. Anti-law #4 of luxury marketing: make purchasing difficult. The agent makes it trivial. BCG ("Consumers Trust AI to Buy Better", 2026) measures that 36% of consumers already trust AI to influence their purchases — barely two points below the trust granted to the in-store advisor (38%). The tipping point is imminent. The challenge: the ceremony of purchase in-store — packaging, ritual, the advisor's gaze — remains intact. But the certainty that the purchase will happen at the Maison's own channel has vanished. The decision: refuse checkout inside the agent. And I think the Maison would be right to do so. Purchasing through a third-party agent trivialises the transaction. For a luxury Maison, triviality is a slower poison than absence.

The migration does not stop at checkout. The agent knows you bought a bag six months ago. It suggests the care service, the repair, the complementary accessory. It does so with the data it has — not necessarily the Maison's data. Checkout.com (December 2025) measures that 47% of consumers are willing to use AI for repeat purchases. After-sales service, repetitive by nature, is the first candidate for delegation. What the Maison still controls: the quality of physical service, the atelier, the repair. What it has lost: the client's first instinct when something goes wrong. What it can choose not to reclaim: routine after-sales. But it will then lose the pretext for contact — and in luxury, the pretext is often more precious than the transaction.

---

### 2.4. LOSING THE FIRST WORD — THE RELATIONSHIP NO LONGER STARTS AT THE HOUSE

The traditional loyalty programme speaks to the client by email. The agent speaks to the client continuously. Sephora's Beauty Insider (80 million members) is now connected to ChatGPT. The client no longer needs to open the Sephora app to access their benefits — the agent integrates them into the conversation. The challenge: the generosity of the programme, the exclusivity of the offers, the events reserved for VICs remain proprietary assets. But the channel of the relationship has changed hands. The decision: cede transactional loyalty to the agent, and concentrate all energy on emotional loyalty — the invitation, the exclusive preview, the human relationship. The most astute Maisons will make this division deliberately.

The formula that summarises these four mechanisms: the Maison has not lost its client. It has lost the first word.

## VISUEL 04

# Customer journey: before vs. with agent

### BEFORE: THE MAISON IN CONTROL

Discovery	Comparison	Personalisation	Purchase	After-sales	Retention
Window, advertising, editorial	Boutique, official website	In-store advisor	Point of sale, e-commerce	Customer service, repair, returns	CRM, private events

### WITH AGENT: CONTROL FRAGMENTS

Discovery	Comparison	Personalisation	Purchase	After-sales	Retention
The agent filters and pre-selects	The agent benchmarks in real time	Agent + Maison data	Direct Maison transaction	Agent relays, Maison executes	The agent manages the relationship

■ Maison control 
 ■ Shared control 
 ■ Control lost

Of 6 stages, the Maison fully controls just 1 alone.

### PART 3

## Transparency versus desire

The AI agent makes price comparable, flattens brand narrative, and removes the friction that underpins luxury value. What Stendhal called crystallisation does not survive the instant response. Three postures are available to Maisons.

---

# Part 3 — Transparency versus desire

---

## 3.1. THE AGENT MAKES PRICE COMPARABLE

### 5 Emotional content is invisible to agents

A Maison's narrative is built for emotion, not extraction.

The benchmark reveals a paradox I consider central: the most iconic Maisons are the most frequently cited by AI agents. And they are also the least actionable. The reason is structural, not accidental. It is a direct consequence of how luxury tells its stories.

A luxury Maison's storytelling is designed to create emotion. A Dior campaign doesn't sell a bag. It sells a world — Jonathan Anderson who, for his first women's show at the Tuileries in October 2025, conjured an inverted grey pyramid above the runway, projected the archives from the house of Christian Dior to Galliano, then sent out New Look crinolines reinvented as miniskirts of canvas and frayed denim — standing ovation. The Paysage atelier building the sets, the gesture of a head seamstress adjusting a drape. Every element of communication is optimized for the emotional circuitry of the human brain. And it works. It has worked for a century.

But the AI agent doesn't process emotion — it processes information. It extracts attributes: price, material, size, availability, reviews, comparisons. Brand narrative passes through the agent like water through a coffee filter: what remains is the informational concentrate. The heritage, yes. The savoir-faire, yes. But the thrill, the mystery, the « je ne sais quoi et le presque rien » of Jankélévitch — all of that passes through the filter without leaving a trace.

McKinsey's figures confirm this dynamic: a brand's own website accounts for just **5 to 10%** of the sources AI agents actually consult (McKinsey, "AI Discovery Survey", 2025). Agents feed on press coverage, consumer reviews, specialist forums, Wikipedia entries. Omniscient Digital, analyzing 23,387 citations in ChatGPT responses, found that **48%** of references point to reviews, listicles and forums (Omniscient Digital, "Where Does ChatGPT Get Its Information?", 2025). The content a Maison controls — its site, its press releases, its campaigns — represents a minority fraction of what the agent actually knows about it.

Michael Baeyens puts it with bracing bluntness:

---

*"Bienvenue dans l'ère des produits invisibles. Les marques s'effacent et se fondent dans les navigateurs ou les OS. Plus d'UI, plus d'identité visuelle, plus de branding, en tout cas tel que nous les connaissions."*

*"Tout cela, même quand il s'agit du produit de luxe, pourtant réputé pour maîtriser son environnement de distribution."*

Michael Baeyens, "UI0<sup>2</sup>, L'émergence des produits invisibles", conférence VISI[ON]S, eXalt, octobre 2025

---

If a brand no longer has an interface, it only has its content. And if that content is not readable by the agent, it has nothing at all.

The paradox is twofold. First: the more iconic a Maison, the more it is cited — because 150 years of press coverage have created a massive corpus that language models have ingested. Hermès is extraordinarily well cited. The narrative of the saddler-craftsman, of deep time, of exclusivity — all of it is in the training data. But that narrative is frozen. It dates from the model's last training cutoff. The current collection, the appointment of a new creative director, last week's pop-up — none of that is in GPT-4o's response, whose data stops at a fixed point in time.

Second: the most emotionally charged content is the least extractable. A 90-second video of an artisan hand-stitching a Birkin is a storytelling masterpiece for a human being. For an AI agent, it is invisible. A podcast in which a creative director shares his vision is inaudible to crawlers. An immersive in-store experience does not exist in the data. Luxury invests fortunes in formats that agents cannot read.

Agents produce informational nakedness. (Mickaël Tsakiris) They strip brand narrative of everything that is not a measurable attribute. And in luxury, value lies precisely in what cannot be measured.

---

### 3.2. THE ASYMMETRY OF ATTENTION COLLAPSES

Georg Simmel wrote in 1905, in his *Philosophie de la mode*: luxury is a game of distance. The distance between the one who possesses and the one who desires. Between the object and the hand that has not yet touched it. It is this distance that creates value — not the material, not the savoir-faire, but the gap. Narrow the gap, and you diminish the value. Eliminate it, and all that remains is an object.

The AI agent is, by design, anti-luxury. It responds to demand. It facilitates purchase. It removes the obstacle — and with the obstacle, the very condition of desire.

The AI agent is a machine built to eliminate distance: it responds in three seconds, compares instantly, personalizes on the spot, proposes checkout before the client has had time to dream. Every feature of the agent is a victory for efficiency, and a defeat for desire.

Eric Briones, editor-in-chief of Journal du Luxe and an industry observer for two decades, has named this phenomenon with surgical precision: luxury is going through a "**libidinal crisis**". "*Le Luxe ne traverse pas une crise économique. Il traverse une crise libidinale*", he wrote in January 2026 — a crisis of the "*Libido Luxe*", that life-force, that irrational energy that drives us to desire what we do not need (Briones, E., "La Crise libidinale du luxe", Journal du Luxe, janvier 2026). The AI agent, by rationalizing every step of the journey, comparing what should never be compared, and responding before the question has had time to ripen, directly accelerates this extinction of desire.

Stendhal calls this crystallisation. In *De l'amour* (1822), he describes the phenomenon through an image that has since become celebrated: a dead branch is plunged into the salt mines of Salzburg and emerges, a few months later, covered in glittering crystals. Every roughness, every knot in the wood has been adorned with salt diamonds. It is no longer a branch — it is an object of fascination.

Crystallisation, for Stendhal, is the mental process by which the mind discovers new perfections in the beloved object. But there is one condition: time. The branch needs months in the mine. Desire needs distance, absence, unanswered questions. The lover who obtains everything immediately does not crystallise. He consumes. And consumption is not desire.

Luxury Maisons have always understood this. The Hermès waiting list functions as an instrument of desire, not a production failure. Waiting for a Birkin for months, sometimes years, is allowing crystallisation to operate. Every week without the bag is a week in which the client's imagination works, enriches the object with qualities it may not even possess, transforms it into a branch covered in crystals. Hermès does not sell a bag. Hermès sells the time that separates desire from possession.

Anne Dellière, Marketing & Strategic Planning Director at Richemont, found exactly the right word for this reality. In a conversation with Eric Briones for Journal du Luxe Intelligence, she invokes **agalma**: an ancient Greek term designating the offering made to the gods, the sacred object that allows its owner to ascend into whatever Olympus they have chosen. "*On vend des alibis pour être quelqu'un d'autre*", she observes. "*La vraie réalité augmentée, c'est le luxe.*" The luxury object is an agalma — not a product with comparable attributes, but a transactional object with the powerful, the divine, the initiated (Dellière, A., interview Journal du Luxe Intelligence, in Briones, E., 2024).

The AI agent knows none of this. It sees a SKU, a price, a material, a customer review. It is blind to the agalma. It is blind to Olympus. And by comparing a Birkin to a Polène bag on the basis of their functional attributes, it commits the original error: treating as a product what is, in fact, an offering. These are anti-laws #4 and #15 of luxury marketing: make purchasing difficult, do not respond to growing demand. These rules are not whims. They are the conditions of desire.

The agent does exactly the opposite: it responds to demand, facilitates purchase, removes the obstacle.

• • •

The consumer data is unambiguous: SurveyMonkey (December 2025) finds that 79% of Americans prefer a human over an AI agent for complex decisions. More than 60% of luxury buyers, according to McKinsey and Business of Fashion (State of Fashion 2026), consider exclusively AI-generated content "less desirable" or "less valuable" than human creation. This has nothing to do with technophobia. It is the intuition that something is lost in immediacy — and that this something is precisely what justifies the price.

But — and this is where the tension becomes productive — the very same consumers use agents extensively. 45% of them at least once during their purchase journey. GenAI usage for shopping: +35% between February and November 2025 (BCG). Gen Z, which represents 40% of American consumers and will become the primary consumer base by end of 2026, has adopted agents without hesitation. 45% of Gen Z and 41% of Millennials used GenAI for shopping in 2025 (CapGemini).

The dissonance is total. Clients say they want the human touch, craftsmanship, slow time — and they increasingly delegate to agents that abolish all of that. Leon Festinger would call this a classic cognitive dissonance: when behavior contradicts belief, the behavior stays — it is the belief that adapts.

The challenge for Maisons is therefore not to prevent the agent from removing friction — that is impossible. The challenge is to choose which friction deserves to be preserved, and which is a pointless hardship that the client is entitled to see disappear.

I propose distinguishing between two types of friction.

## CONCEPT

# Productive Friction

---

The kind that creates value. The waiting list. The mandatory appointment to access certain pieces. The deliberate absence of a product from agent results. Silence as strategy. This friction enriches desire — it is constitutive of the luxury promise. Remove it, and you remove the product itself.

## CONCEPT

# Endured Friction

---

The kind that creates nothing. A website that doesn't work on mobile. Customer service that takes four days to respond. Product information that cannot be found. A loyalty programme requiring three clicks to locate a single benefit. This friction adds no crystals to the branch. It adds irritation. The agent that removes it is doing the Maison a favour.

Rémi Guyot, founder of AI Discipline, frames this as a design problem: AI removes what he calls the *"protective buffer of slow execution"*. The craftsman who works leather for forty hours clarifies his intention through the gesture itself. The agent, by contrast, executes immediately and completely. Slowness was not a flaw — it was a quality filter. By removing it, the agent also removes the maturation process that gives a product its rightness (Guyot, R., "Apprentissages #2", AI Discipline, 2025).

Maisons that conflate the two — that believe all friction is good because they operate in luxury — are wrong. And those that hand all friction to the agent — believing efficiency is always a virtue — are equally wrong.

## Productive friction vs friction subie

### Friction productive

Creates desire, scarcity, ritual

#### Waiting list

Time reinforces perceived value

#### In-store appointment

Personalised reception as privilege

#### Deliberate absence

Not being everywhere is a strategic choice

#### Silence

Not answering everything, not showing everything

### Endured friction

Frustrates, degrades, drives clients away

#### Slow or dysfunctional website

The digital experience contradicts the promise

#### Absent or impersonal after-sales

High-value client treated like a support ticket

#### Information impossible to find

Price, availability, product details missing

#### Complex loyalty programme

Mass-market mechanics applied to luxury

Friction is not a flaw to fix. It is an asset to defend.

• • •

---

### 3.3. THREE POSTURES TOWARD THE AGENT

Faced with the client-brand-agent **Love Triangle**, three strategies emerge. None is perfect. All are instructive. And the most dangerous is the third — the one that consists of doing nothing.

#### **Zero cookies, zero declared data: what Dior understood**

Gal Rapoport spent years at Amazon building Alexa's personalisation systems. He knows what it means to read an intention before it has been articulated. In 2021, he co-founded Kahoona in San Diego. Twenty-seven people. One thesis: if the external agent captures intent before the visit, the Maison can still capture intent within the first seconds on its own site.

The approach is called "digital body language": where does a visitor click, where do they linger, the scroll that stops on an image, the third page abandoned too quickly. Kahoona reads these micro-signals on anonymous visitors — no cookies, no history, no declared data whatsoever — and reconfigures product listing pages in real time for each detected profile.

This goes beyond CRM-style personalisation. It is intent-reading at the precise moment the intent is forming. Digital clienteling that asks nothing and understands everything.

The strategy of the agent within rests on one principle: cede no ground. The agent works backstage, invisible to the client. The Maison retains control of the narrative, the experience, the aesthetic. The client does not know an algorithm has reorganised the page they are browsing. They trust their own discovery. Crystallisation is preserved.

This is the strategy most consistent with the anti-laws of luxury marketing. The agent does not sell, does not compare, does not respond to demand — it anticipates it. It does not make purchasing easier; it makes discovery more precise. LVMH plans to extend the solution to other Maisons within the group.

#### **WHAT THIS MEANS FOR YOU**

If your Maison has invested in CRM for ten years yet ignores the 96% of visitors who never log in, the internal agent fills that gap. The question is not "do we need AI?" — it is: "what do we actually know about the client who has never said a word?"

## When the client says "Sephora" to ChatGPT

On 24 March 2026, at Shoptalk Las Vegas, Sephora announced the launch of its app within ChatGPT. The client formulates their request in natural language. They can connect their Beauty Insider account (80 million members): purchase history, skin type, saved preferences. GPT Vision analyses the selfie in real time, diagnoses skin tone, and recommends with the precision of a flagship adviser. Expected reduction in returns: 30% (Sephora Newsroom, 24 March 2026; Retail Gazette, 24 March 2026).

This is the partner-agent strategy. The Maison enters the external agent. It deposits its data, its products, its loyalty program. It accepts that the client will discover it in an environment it does not control — and negotiates the terms of its presence.

The risk is real. When a client receives a beauty recommendation inside ChatGPT, who does she credit for the experience? Sephora, which provides the data and the expertise? Or ChatGPT, which provides the interface and the conversation? The "you don't say Sephora, you say ChatGPT" problem is the strategic danger of this approach. The brand becomes a data supplier inside someone else's ecosystem.

Sephora accepts this for two reasons. First, beauty is the vertical most compatible with agentic commerce: the criteria are objective (skin tone, skin type, allergies), personalization lends itself to computation, and the risk of desacralization is far lower than in haute couture or jewelry. Second, Sephora is a retailer, not a luxury Maison in the strict sense. The symbolic capital at stake is not the same as for Hermès or Chanel.

The calculus is different for Maisons at the heart of luxury. A Kelly inside ChatGPT, recommended between two Samsonite luggage suggestions, loses its aura before it is even purchased. This is the Benjaminian paradox taken to its conclusion: every identical recommendation, served to millions of users, is one more reproduction that erodes the singularity of the object a little further.

### WHAT THIS MEANS FOR YOU

If you are a multi-brand retailer with structured data and a loyalty program, the partner strategy is your natural terrain. If you are a Maison whose value rests on rarity and mystery, entering the external agent is tantamount to inviting Benjamin's photographer into your atelier.

## Silence as strategy: a decoy

Then there are the Maisons that do nothing: no Kahoona backstage, no app inside ChatGPT, no structured data for agents, no doctrine on the subject.

The agent talks about you anyway. With whatever it finds. That is to say: press archives, Wikipedia articles, Reddit reviews, second-hand data. The Converteo benchmark (March 2026) makes it bluntly clear: the Maisons best represented in AI agents owe that visibility to 150 years of published text feeding the models. The visibility they enjoy is not a strategy. It is a legacy dividend. And a dividend, by definition, depreciates when the context shifts.

The actionability score — the capacity of AI responses to translate into a visit or a purchase — is the lowest of all Maisons measured. The leading group (Hermès, Louis Vuitton, Chanel) achieves the best citation scores, but none exceeds two-thirds of the maximum actionability score. Being cited is not enough. You need to be activatable. And on that point, according to Converteo, no one has a satisfactory answer.

DLG (Digital Luxury Group) confirms the diagnosis in its "State of AI in Luxury 2026" report: 55% of luxury brands remain in the exploratory phase. 42% are in "pilot purgatory" — experimenting without scaling. Only roughly one third report AI integrated into their daily operations.

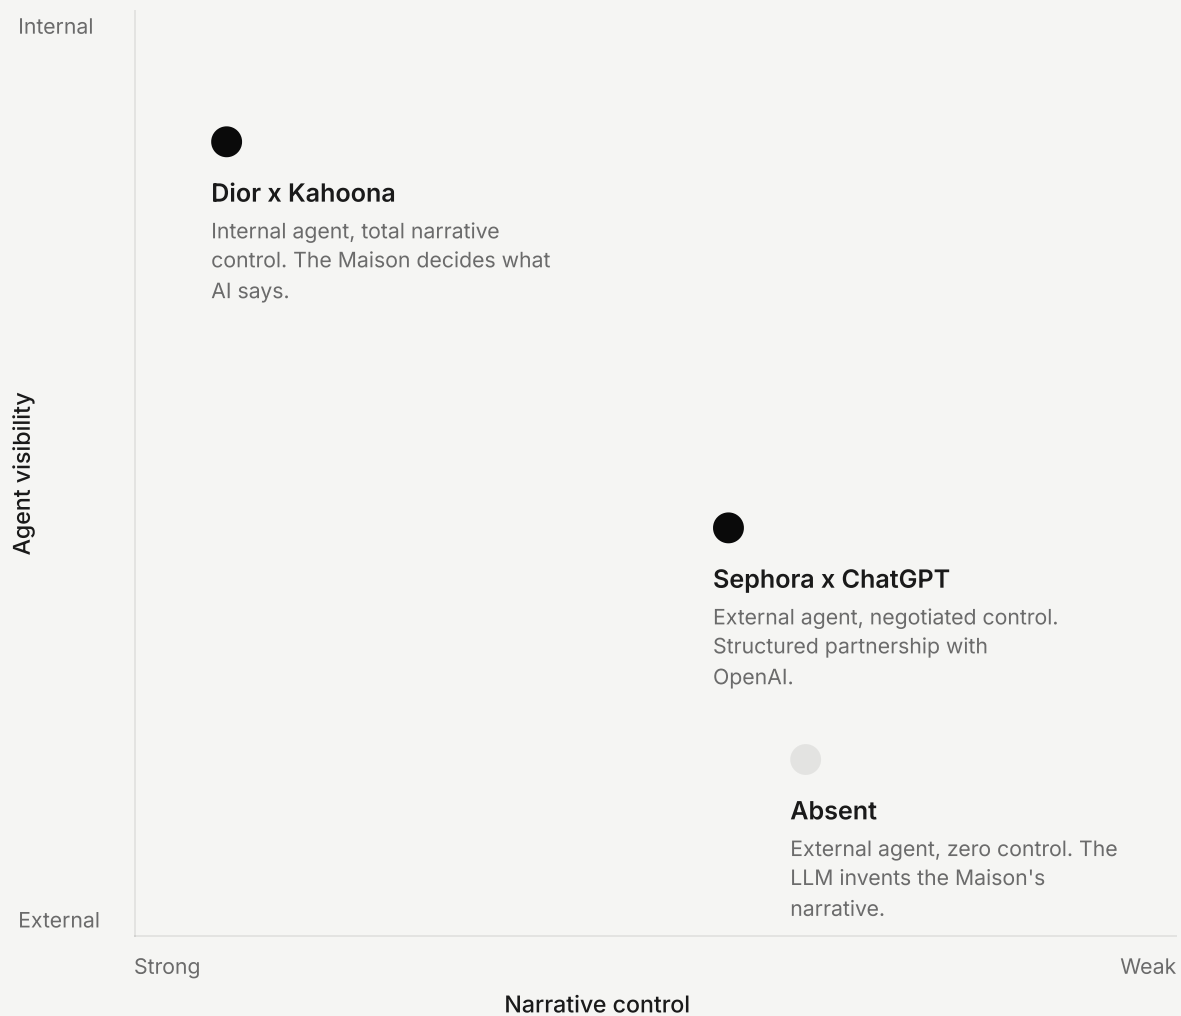
Absence amounts to a loss of sovereignty by default. When the agent represents your Maison without your consent, using data you have not validated, in a format you do not control, you have not chosen silence. You have been silenced.

#### **WHAT THIS MEANS FOR YOU**

If you have no doctrine on your presence in AI agents, you have one anyway. It was decided by the algorithm. The question your executive committee needs to settle is not "should we engage?" but: "what are they saying about us out there — and does that suit us?"

VISUEL 06

# Matrix of the 3 strategies



• • •

### 3.4. THE PARADOX: 74% WANT A HUMAN, 41% TRUST THE MACHINE

There is an understandable reflex when confronted with the Love Triangle: the fear of disappearance. If the AI agent recommends, compares, personalizes — what is left for the in-store advisor? The answer is: everything the agent does not know how to do. Which is a great deal.

The BCG figure cited earlier takes on its full meaning here: four out of five Americans prefer a human for complex decisions. This does not say that consumers reject AI. It says they find it insufficient for the moments that matter — those where the price is high, the emotional stakes are real, the risk of regret tangible. In other words: the moments of luxury.

But in the same report, 72% of GenAI chatbot users rate the assistance as good as that of a human for routine queries. And 36% trust AI to influence their purchases — just two points behind the in-store advisor. The boundary between "moment that matters" and "routine query" is moving. Fast.

The full paradox, in numbers: 74% of consumers want a human in the luxury relationship. 41% trust GenAI results more than paid search results (IAB, 2026). The client says: I want a human. And when they search, they ask the machine.

**74 % / 41 %**

want a human — but trust AI more than paid search

IAB, 2026

Two cases show that the answer is not one or the other. It is one armed by the other.

**Burberry and Project Penguin.** In 2025, Burberry deploys a generative AI-powered discovery and recommendation platform to more than 100 advisors across all its global markets. The advisor queries the system in natural language, in the client's own language. They can show a photo of the client wearing a piece and receive "Complete the Look" suggestions generated by LLMs fine-tuned on studio-curated looks. The result: a +24% uplift in average transaction value across customer service channels. The DataIQ "Most Innovative Use of AI (Global)" award confirms it (DataIQ Global Awards, 2025). Burberry did not replace the advisor. It gave them a head start. When the client walks into the boutique, the advisor already knows what the external agent recommended — and can go further.

**Clinique and the AI mirrors.** Clinique deploys augmented mirrors in-store that analyze the client's skin in real time and propose a personalized diagnosis. The advisor uses the diagnosis as the starting point for the conversation, not as a substitute for it. The result reported by BrandXR: +30% average basket value in-store. The mirror does not sell. The advisor sells. But the advisor, equipped with the diagnosis, sells better.

What these two cases reveal is a model I call the **Augmented Advisor**. The AI is not in the room with the client. It is in the advisor's earpiece. It whispers data, suggestions, patterns. The advisor chooses what to keep, what to discard, what to adapt to the client standing in front of them. Human intuition remains. The gesture remains. The ceremony remains. But the improvisation is better informed.

DLG figures ("State of AI in Luxury 2026") confirm that the industry is moving in this direction: 47% of luxury executives identify training needs as their top AI priority. The bottleneck is human capability, not budget or infrastructure. Arming the advisor also means training them on a tool they do not know — and which they legitimately fear.

Converteo frames it differently in its analysis of agentic AI in luxury (March 2026): the "Human in the Loop" is not a brake on agent autonomy. It is the validation layer that transforms an algorithmic recommendation into a luxury experience. Emotional intelligence — the ability to read a silence, a hesitation, a gaze that lingers a second too long on a piece of jewelry — remains a competitive advantage the agent does not possess. Not yet. Perhaps not ever.

The client trusts the agent to **find**.

The client trusts the human to **choose**.

I do not believe in the disappearance of the advisor. I believe in the disappearance of the unarmed advisor. The one who knows nothing about the client who walks through the door. The one who has no access to the data the agent already holds. The one who discovers in the boutique what the client has already explored, compared, and pre-selected in the Antechamber. That advisor is already obsolete — not because AI replaced them, but because the client arrives with a level of information the advisor can no longer match without tools.

Luxury has always been a matter of knowledge. Knowing how to recognize a leather, knowing how to propose the right pairing between a dress and a piece of jewelry, knowing how to slow the conversation when the client is ready to buy too quickly. That knowledge does not disappear. It deepens. And the sales advisor who possesses it — armed by AI — becomes the Maison's strongest argument in this new love triangle.

## PART 4

# Architecting presence to regain control

What do we do? A reproducible audit, a decision matrix, five workstreams to launch by Monday, the technical foundations, and unprecedented KPIs to track what no current dashboard measures.

---

# Part 4 — Architecting presence to regain control

Part 1 showed what agents say about you — and what they do not say. Part 2 analyzed how the agent «hacks» the client journey. Part 3 illuminated what this threatens at its core: the friction, the opacity, the scarcity that underpin luxury's value. One question remains, and it is the only one that matters in a boardroom: what do we do?

This fourth part offers concrete tools. A reproducible audit to measure what is not being measured. A decision matrix to choose your posture. Five workstreams to launch by Monday, without waiting for IT. The technical foundations for when IT is ready. And unprecedented KPIs to track what no current dashboard tracks.

---

## 4.1. THE DECISION MATRIX

### The Antechamber Audit, Tool #1

Most Maisons monitor their e-reputation. Google reviews, press mentions, social sentiment. The tools have existed for fifteen years; PR and CRM teams have used them for fifteen years. But almost no one monitors what AI agents say about the Maison. Not client reviews. Not journalists' articles. What the algorithm itself — queried in natural language — answers when a prospect types: "What is the best leather bag around €3,000?"

This is a strategic blind spot. DLG has measured the phenomenon: more than half of luxury Maisons remain in the exploratory phase on AI. Translated into operational terms, this means the majority of Maisons do not know what ChatGPT, Gemini, or Perplexity are telling their prospects this morning.

The audit I propose is simple. Ten questions. Three agents. A scoring grid. Monthly cadence. Execution time: one hour. A marketing reflex, not an IT project.

#### THE 10 QUESTIONS TO ASK AI AGENTS ABOUT YOUR MAISON

Ask each question to ChatGPT, Gemini, and Perplexity. Record the answers. Compare.

<p><b>01</b> "What are the three most recognized Maisons in [your category]?"</p> <p>Are you on the spontaneous shortlist?</p>	<p><b>02</b> "Recommend a [product] for [profile], budget [X]"</p> <p>Does the agent cite you in a purchase context?</p>
<p><b>03</b> "What's the difference between [you] and [competitor]?"</p> <p>Does the agent flatten your positioning?</p>	<p><b>04</b> "What is the distinctive savoir-faire of [your Maison]?"</p> <p>Depth of narrative: atelier, gesture, artisan?</p>
<p><b>05</b> "Why is [your Maison] more expensive than [X]?"</p> <p>Justified by desire or by materials?</p>	<p><b>06</b> "Where to buy a [product] online?"</p> <p>Who captures the traffic: you or a third party?</p>
<p><b>07</b> "What's new this season?"</p> <p>Data freshness: current season or last season?</p>	<p><b>08</b> "Is [your Maison] committed to sustainability?"</p> <p>Is the CSR narrative visible to the agent?</p>
<p><b>09</b> "What do clients think of [your Maison]?"</p> <p>Sentiment sources: Reddit, Trustpilot, forums?</p>	<p><b>10</b> "Memorable luxury gift: what do you recommend?"</p> <p>The ultimate test: are you cited without being searched?</p>

## GRILLE DE SCORING

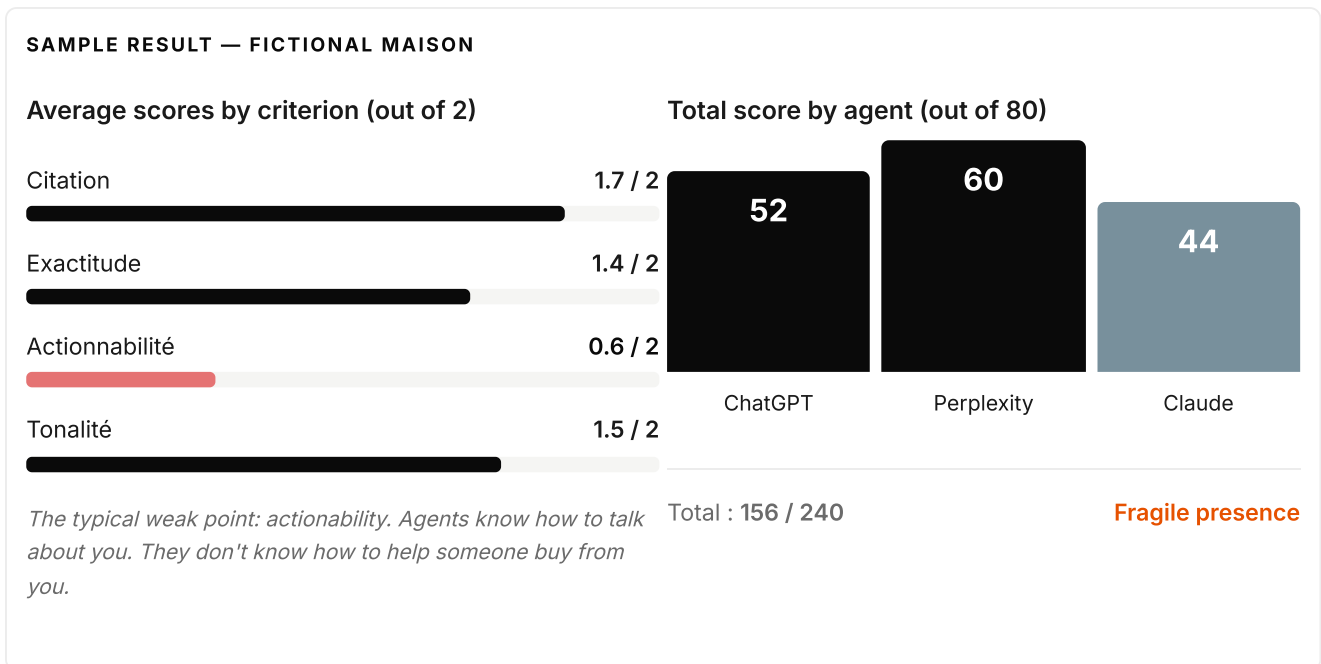
For each question and each agent, score on four criteria (0 to 2 points):

CRITERION	WHAT YOU MEASURE	0	1	2
<b>Citation</b>	Is the Maison mentioned?	Absent	Mentioned	Featured
<b>Accuracy</b>	Facts correct? (dates, prices, founder)	Errors	Approximate	Exact
<b>Actionability</b>	Does the response lead to purchase?	No action	Generic link	Direct action
<b>Tone</b>	Alignment with brand positioning	Misaligned	Neutral	Aligned

Max score per question: 8 pts · Per agent: 80 pts · Total (3 agents): 240 pts

<p><b>180-240</b> Antechamber mastered The agent is an ally. Maintain and optimize.</p>	<p><b>120-179</b> Fragile presence Gaps in the narrative. Correct and enrich.</p>
<p><b>60-119</b> Alert The agent talks about you, but poorly. Priority workstream.</p>	<p><b>0-59</b> Critical absence You don't exist. Escalate to the Executive Committee.</p>

<p>FREQUENCY Mensuelle</p>	<p>OWNER RP / Brand Content</p>	<p>DELIVERABLE Table + 3 actions</p>	<p>ESCALATION Score &lt; 120 → Executive Committee</p>
--------------------------------	-------------------------------------	--	--



## The Decision Matrix: four postures toward the agent

### Tool #2: which posture for your Maison?

The diagnosis is done. The audit reveals your position in the Antechamber. What follows is a strategic choice: what do you do with it?

I am convinced there is no universal answer. Telling every luxury Maison to "go all in on agentic AI" would be as absurd as telling them all to open a TikTok account in 2020. Some did it brilliantly. Others embarrassed themselves. The difference was never the channel. It was the fit between strategy and the Maison's reality.

The matrix I propose maps two axes.

**Axis 1: Desired level of control.** To what degree does the Maison want to govern what the agent says about it?

- **Full control:** the Maison wants every word uttered by the agent to be validated, consistent with the brand book, managed like a press release.
- **Shared control:** the Maison accepts that the agent speaks on its behalf, provided the terms, data, and boundaries have been negotiated.
- **No control:** the Maison considers AI agents outside its strategic territory, or lacks the means to integrate them.

**Axis 2: Data maturity.** What data does the Maison have to feed the agents?

- **Structured data:** complete PIM (Product Information Management), unified CRM (Customer Relationship Management), indexed editorial content, enriched product sheets, clean data schemas.
- **Fragmented data:** scattered PIM components, a partial CRM, content dispersed across the website, app, social channels, and press releases.
- **No data:** no centralized PIM, no product data schema, brand content is essentially visual or unstructured.

## The four quadrants

### A. Le Souverain

#### High control + structured data

The Maison has rich data and wants maximum control. The agent works from within, on the Maison's data, under the Maison's oversight. This is the Dior × Kahoona model.

#### ACTIONS

- Deploy a proprietary agent (Kahoona-type)
- Enrich the PIM with a semantic layer
- Appoint an "agentic voice" lead

Timeline: 6–12 months · E.g.: Dior, Burberry (Penguin), L'Oréal

### B. Le Diplomate

#### Shared control + structured data

The Maison has the data but strategically accepts sharing control with external agents. This is the Sephora × ChatGPT model: entering a third-party ecosystem and negotiating your presence within it.

#### ACTIONS

- Negotiate partnerships with LLM providers (OpenAI, Google, Anthropic)
- Connect loyalty and PIM to agents via API
- Define an "agentic specification document"

Timeline: 3–6 months · E.g.: Sephora, multi-brand retailers

## C. L'Architecte

### High control + fragmented data

The Maison wants to control its narrative, but its data doesn't yet allow it. This is the most common situation in luxury. DLG finds that 37% of executives identify fragmentation as the number-one barrier.

#### ACTIONS

- Structure the PIM with a semantic layer
- Unify the CRM (one client = one profile)
- Secure earned media in the meantime

**Timeline:** 12–18 months (data), 3 months (earned) · **E.g.:** Kering (Houlès), independent Maisons

## D. Le Fantôme

### No control + fragmented data

The Maison has neither the will to control nor the data to exercise it. The agent talks about it using whatever it finds: press archives, Wikipedia, forums. The narrative has entirely escaped.

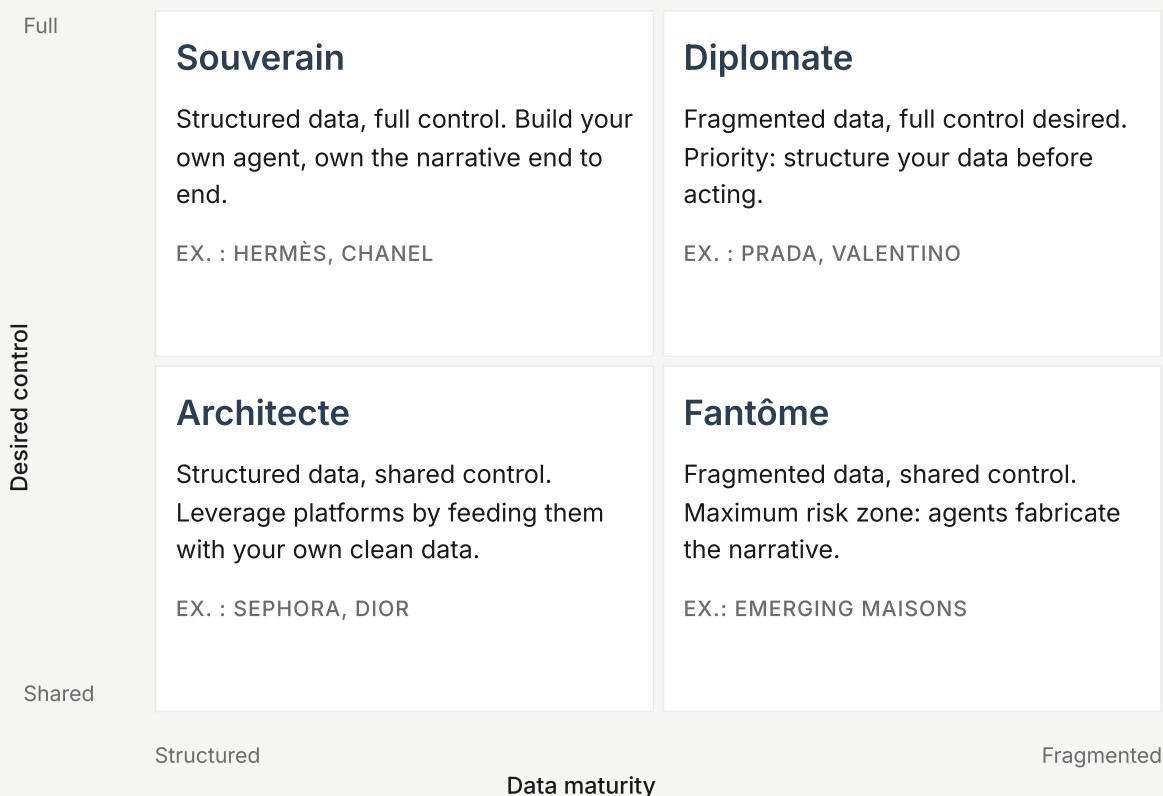
#### URGENCY

- Immediate Antechamber Audit (section 4.1)
- Score < 120: escalate to Executive Committee
- Decide: deliberate presence or deliberate absence. Indecision is not an option.

**Timeline:** decision within 30 days, workstream 6–18 months · **E.g.:** emerging Maisons, labels without a data strategy

## VISUEL 08

# Decision Matrix: where does your Maison stand?



## 4.2. FIVE PROJECTS TO LAUNCH ON MONDAY

I don't believe in five-year roadmaps. Not in this context. Agentic AI is moving too fast for anyone to credibly claim they know what the landscape will look like in 2030. What I do believe in: calibrated, prioritized workstreams, each with an owner, a deliverable, and a deadline. Five is enough. Here they are.

## WORKSTREAM 1

3  
mois

Brand  
Content

RP

E-  
commerce

# Audit and correct the narrative inside the agents

Your Maison is cited by AI agents. But in what words? The benchmark shows that the most cited Maisons are often the least actionable. Being known is not enough — you need to be activatable. In April 2026, LVMH created two Group-level roles in a single week — Head of Tech Digital Commerce and Head of AI Solutions — to structure precisely this transition.

### ACTIONS

- Run the Antechamber Audit (section 4.1) across four agents
- Identify gaps between the desired narrative (brand book) and the narrative actually rendered
- Rewrite brand arguments in machine-readable format: from "An exceptional leather bag" to "Full-grain calfskin bag, hand-stitched, 42 hours of manufacture"
- Enrich product pages with structured data (schema.org, JSON-LD)
- Re-test after 30 days

**Deliverables:** Initial audit (10-query × 4-agent grid) • Editorial correction plan • M+1 control audit • "Dual readership" writing guide

## WORKSTREAM 2

Immediate

Communications Director

RP

Brand Content

# Train PR teams in dual readership

Every press brief now has two audiences: the journalist, and the language model that will ingest the article. A piece in Vogue Business is no longer just an article — it's training data. The Communications Director must integrate this reality into every brand communication.

### ACTIONS

- Train PR teams in dual readership: human + machine
- Include a factual box in every brief: product attributes, savoir-faire, origin, quantified differentiation
- Move from "a collection celebrating the free woman" to "42 hours of hand embroidery at Lesage, for a dress never reproduced identically"
- Brief external PR agencies on machine-readable writing rules

**Deliverables:** "Dual readership" training (2h, internal teams + agencies) • Machine-readable writing guide • Enriched press brief template • Pre-send validation checklist

## Define the Maison's doctrine

"Should we be present in AI agents?" is an executive committee question. The stakes go beyond visibility: in April 2026, Target published the first legal warning of the agentic era — any purchase made by an agent on a client's behalf is considered authorized by the client, even in the event of an error. And Sayali Patil (AI Infrastructure Product Leader, Cisco) flags a risk luxury Maisons should read carefully: when a client buys through an external agent, the Maison loses the complete behavioral trace — what was searched, viewed, compared. For Maisons whose model rests on clienteling, that's not lost data. It's the relationship changing hands.

### DECISIONS

- Place the topic on the agenda of the next executive committee
- Present the Audit results + position in the Decision Matrix
- Choose: Sovereign, Diplomat, Architect — or deliberate absence
- Document in a one-page "agentic doctrine"

**Deliverables:** One-page agentic doctrine validated by Executive Committee • Antechamber scoring grid presented in committee • Framing note: "what the agent may say / what the Maison chooses to withhold"

## Equip the in-store advisor

The 2026 client arrives informed by an AI agent — sometimes better informed than the advisor. Digital clienteling has existed since 2013 (Tulip, BSPK, Cegid). What is new with agentic AI: multi-session memory, autonomous decision-making, bold suggestions, end-to-end orchestration.

### ACTIONS

- Equip advisors with agentic AI-powered clienteling (Salesforce Agentforce, Alhena AI)
- Train retail teams: the AI diagnosis is a starting point, not a substitute
- Define the untouchable "human moments": welcome, tactile discovery, payment ceremony
- Measure: average transaction value, conversion, NPS before/after

**Deliverables:** Augmented clienteling specification • Pilot in 3–5 boutiques at M+3 • Advisor training protocol (AI as tool, not substitute) • Before/after conversion dashboard • Full rollout M+6

## Measure: the new KPIs for agentic visibility

No current marketing dashboard captures a Maison's visibility inside AI agents. We are measuring the old world with the instruments of the old world.

### 5 KPIS TO INTEGRATE INTO MONTHLY REPORTING

#### Agentic citation rate

Out of 10 generic queries, how many times are you cited? Frequency: monthly.

#### Accuracy score

Of all citations, what % is factually correct? Target: >80%.

#### Actionability score

What % leads to an action (link, boutique, purchase)? Target: >50%.

#### Tone index

Alignment with the Maison's positioning. Score 0-10.

#### Agentic share of voice

On competitive queries, what share of mentions do you hold vs. competitors? Compare against traditional media share of voice.

**Deliverables:** "Agentic visibility" dashboard integrated into monthly reporting • Quarterly audit protocol (10 queries × 4 agents) • Agentic competitive benchmark • Automatic alert on narrative drift (via Evertune, BrightEdge going forward)

VISUEL 09

# Roadmap: 5 workstreams over 6 months

IMMEDIATE 0-1 month	NEAR-TERM 3 months
HORIZON 6 months	
<b>Chantier 1</b> <b>Narrative audit</b> OWNER: CMO	
<b>Chantier 2</b> <b>PR training</b> OWNER: COMMUNICATIONS DIRECTOR	
	<b>Chantier 3</b> <b>Executive doctrine</b> OWNER: EXECUTIVE COMMITTEE
	<b>Chantier 5</b> <b>New KPIs</b> OWNER: DATA / ANALYTICS
<b>Chantier 4</b> <b>Augmented clienteling</b> OWNER: RETAIL DIRECTOR	

One question remains that neither the audit, nor the matrix, nor the five workstreams can answer for you. It predates agentic AI. It is as old as luxury itself.

Should you seduce the agent to reach the client, or refuse to play?

I will be direct. Both answers are defensible.

To seduce the agent means accepting that an uninvited intermediary participates in the relationship. It means reformulating your narrative so a machine can read it. It means entering a system of comparison that luxury's fundamental principles have always deemed incompatible with the luxury register. But it also means being present at the moment intention forms, influencing the recommendation, capturing a client who would otherwise go elsewhere. The Maisons that play this game with intelligence — Dior with Kahoona, Sephora with ChatGPT, Burberry with Penguin — demonstrate that you can enter the Antechamber without losing your soul. Provided you know precisely what you are doing there, and what you will never do.

To refuse to play means betting that scarcity is a more powerful argument than visibility. That the client who doesn't find you in the agent will seek you by other paths — a friend's recommendation, a shop window on a street, the memory of a journey. That luxury is earned, including in the effort required to find it. Hermès, which has never needed SEO (Search Engine Optimization) to sell a Birkin, has played this hand for decades. But Hermès is Hermès. The question is: can your Maison afford this posture?

The trap is choosing neither. Not seducing the agent because you don't know how, and not refusing to play because no one has dared raise it at the executive committee. This non-decision has a name in industry parlance: *pilot purgatory* — close to one Maison in two, according to DLG. You experiment without committing. You observe without deciding. You let the agent speak without answering.

---

Luxury has always been a matter of deliberate choices. Agentic AI has not changed the nature of this business. It has added one more choice to the list.

---

## 4.3. TECHNICAL FOUNDATIONS

The five workstreams in the previous section involve multiple departments. But none will work if the technical infrastructure of the Maison's website is an invisible wall to AI agents. This section is intended to be shared with your CTO or technical director — it is the bridge between strategy and execution.

### A STRUCTURAL FINDING

GPTBot, ClaudeBot, and PerplexityBot **do not render JavaScript**. A typical luxury site built in React or Vue without Server-Side Rendering: the raw HTML contains 2 KB of `<script>` tags. The rendered HTML contains 50 KB of rich text. The entire difference —

product descriptions, savoir-faire, FAQs, storytelling — is **invisible to AI agents**.

*Source: ClickRank, analysis of 500M+ GPTBot fetches, March 2026.*

## 1. The invisible JavaScript wall

The majority of luxury Maison websites are visual showcases built in React, Vue or Next.js, with catalogues behind login, pure-image lookbooks, and press kits in PDF. Each of these choices is an anti-LLM wall.

**Solution:** Server-Side Rendering (SSR) is mandatory for any content intended to be cited by an AI agent. Critical content — product descriptions, savoir-faire pages, FAQs — must be in the initial HTML, not client-side rendered.

**Effort:** 2–4 months for a front-end overhaul. **Impact:** prerequisite for any GEO strategy.

## 2. Schema.org: the language agents understand


JSON-LD schema.org markup is the technical foundation of agentic visibility. Microsoft (Bing, Copilot) confirmed in March 2025 that schema markup directly helps LLMs understand content. GPT-4 performance rises from 16% to 54% when content is structured (Walker Sands, 2026).

**Priority schema types for luxury:**

TYPE	USAGE	IMPACT
Product + Offer	Catalogue, prices, availability	Citation in shopping recommendations
Organization + Brand	Maison identity, history, values	Identity anchoring in responses
FAQPage	Recurring client questions	Direct extraction (Q&A format = LLM native format)
Article	Editorials, manifestos, savoir-faire	Citation in cultural responses
HowTo	Care guides, beauty rituals	Tutorial responses

**Effort:** 1–2 months (CMS integration). **Impact:** high and measurable.

### 3. Freshness as a lever



The ConvertMate data (29 March 2026, 12,500 queries) is unambiguous:

- Content updated within the last 30 days = **3.2x more likely to be cited**
- 44.2% of citations come from the **top 30%** of the text (the introduction)

- 65% of AI crawler hits target content less than one year old

**Implication:** the site's key pages (savoir-faire, collections, FAQs) must be updated at minimum monthly. The introduction of each page is the most important paragraph — it is the one agents cite.

#### 4. Agentic commerce protocols

Three protocols structure agentic commerce in 2026:

- **Universal Commerce Protocol (UCP)** — launched by Google + Shopify at NRF in January 2026. Co-developed with Etsy, Wayfair, Target, Walmart. Endorsed by Mastercard, Visa, Stripe. Enables agents to query a catalogue in real time (inventory, prices, availability).
- **MCP (Model Context Protocol)** — launched by Anthropic in November 2024, adopted by OpenAI (March 2025), Google DeepMind, Microsoft. 97 million SDK downloads per month. Enables an AI agent to connect to tools, databases, and APIs via a unified protocol.
- **Google Merchant Center** — already the pipeline for Google Shopping, it is also becoming the feed channel for Google Search AI Mode. Any brand absent from Merchant Center is invisible in Google's agentic commerce.

**For luxury:** the “readable catalogue” layer (UCP, Merchant Center) is an opportunity to be recommended by agents. But raw agentic checkout (no friction, no ceremony) is contrary to the anti-laws of luxury marketing. The compromise: expose the catalogue and the storytelling, retain control over the purchase journey.

## 5. Monitoring tools

TOOL	ENTRY PRICE	LLMS COVERED	STRENGTH
<b>Evertune</b> (\$19M Series A, Felicis)	Enterprise (custom)	ChatGPT, Claude, Gemini, Perplexity	1M+ prompts/mois, AI Brand Index
<b>Profound</b> (\$58,5M, Sequoia)	\$99- 399/mois	10+ moteurs	Simulation client- side, anti- hallucination
<b>Otterly.ai</b>	\$29/month	ChatGPT, Perplexity, Google AIO	Best value for money
<b>Peec AI</b>	€90/month	6 models	Sentiment analysis, prioritised actions
<b>Adobe LLM Optimizer</b>	~\$115K/year	Multi-platform	A2A + MCP integration (AEM clients)

**Recommendation:** for a luxury Maison, **Evertune** (enterprise, already used in luxury automotive) or **Profound** (precision, real-session simulation). Otterly or Peec AI for a lower-cost first audit.

## 6. llms.txt: signal or noise?

The `llms.txt` file, created by Jeremy Howard (Answer.AI) in September 2024, is a Markdown document placed at the root of a site to guide LLMs. Several hundred sites have adopted it, including Cloudflare, Stripe, and Vercel.

**However:** no measured correlation between `llms.txt` and AI citations (SE Ranking, statistical + ML analysis). Google does not support it and does not plan to (Gary Illyes, July 2025). No confirmed support from OpenAI, Anthropic, or Meta.

**Verdict:** negligible setup cost (one Markdown file), but no proven ROI. Worth monitoring, not a priority deployment. SSR and schema.org are infinitely more impactful.

## 7. Technical priority matrix

PRIORITY	ACTION	EFFORT	IMPACT	TIMELINE
1	SSR + static HTML content for product and savoir-faire pages	High	Very high	2–4 months
2	JSON-LD schema.org (Product, Organization, FAQPage, Brand)	Medium	High	1–2 months
3	GEO content layer (structured FAQs, comparatives, material sheets)	Medium	High	Ongoing
4	Google Merchant Center (if e-commerce)	Low–Medium	High	1 month
5	Monitoring tool (Evertune or Profound)	Low	Medium	Immediate
6	Editorial freshness (update key pages every <30 days)	Low	High (3.2×)	Immediate
7	MCP server for catalogue/FAQ	High	Strategic long-term	3–6 months
8	llms.txt	Negligible	Unproven	1 day

Sources: ConvertMate GEO Benchmark (29 March 2026) · Walker Sands Schema & LLM Visibility (2026) · ClickRank JavaScript Rendering (March 2026) · Google Developers UCP (January 2026) · Model Context Protocol (Anthropic, November 2024) · SE Ranking llms.txt Analysis (2026) · Evertune, Profound, Otterly, Peec AI, Adobe (prices March 2026).

- A **replicable audit** of the Antechamber: 10 questions, 3 agents, 4 criteria, monthly cadence, one hour, no IT budget required.
- A **decision matrix** with four quadrants (Sovereign, Diplomat, Architect, Ghost) mapping desired control against data maturity.
- **Five workstreams to launch on Monday**: audit the narrative, train PR, define the doctrine, equip the advisor, measure.
- **Five KPIs** to integrate into reporting: citation rate, accuracy, actionability, tone, agentic share of voice.
- A **non-negotiable technical foundation**: SSR (AI bots do not render JavaScript), schema.org JSON-LD on every product (+238% visibility), and crawler control via robots.txt.
- The choice to make: seduce the agent or refuse to play? The absence of an answer is the only **untenable** position.

## CONCLUSION

Desire cannot be measured. It is built.

---

# Conclusion

*There is a word in the vocabulary of luxury that is almost never spoken in boardrooms: the word "desire." Desire cannot be measured. It is built. And what this study has tried to map is the way a new actor — the AI agent — disrupts that construction.*

The Antechamber exists. Agents are talking about your Maison to your prospects without your consent, with data you have not validated. 45% of consumers already use AI in their purchase journey (IBM IBV, January 2026). The digital shelf is compressing from 50 products to 3 (NielsenIQ). The discovery phase is escaping the Maison's grasp (McKinsey, March 2026).

The Love Triangle is here to stay. The client, the brand, and the agent now coexist at every stage. The question is not whether the third party is present. It is what role each party chooses to play.

## THE AGENT INSIDE

AI works behind the scenes, invisible to the client. The crystallization of desire is preserved.

## THE AGENT AS PARTNER

The Maison enters the external agent, negotiates the terms of its presence, connects its data.

## THE AGENT IGNORED

By not deciding, the Maison has already let the algorithm decide in its place.

I do not believe agentic AI will destroy luxury. Luxury survived the democratization of travel, e-commerce, fast fashion, social media, and industrial counterfeiting. It survived because its fundamental spring — desire nourished by distance, scarcity, and deep time — is a human constant. AI agents do not change the nature of desire. They change the **geography of its formation**.

What has changed is the place where desire crystallizes. Before: the shop window, the magazine spread, the advisor's murmur. Now, also — and increasingly, first — in an agent's response. The Maisons that deliberately choose their posture in the Antechamber will retain control of their narrative. Those that do not choose will discover that the narrative is written anyway, without them.

---

While preparing this study, I reread a speech by Paul Valéry (1935), brought back to light by Rémi Guyot at AI Product Day in March 2026. Valéry imagines history's greatest scientists confronted with a dynamo: they dismantle it, interrogate it, but the current itself eludes them. The greatest minds, incapable of explaining a device that has become ordinary. Guyot draws the lesson (Mind Fooled, 25 March 2026):

***It was not electricity that transformed humanity — it was the electrification of the world.***

---

Agentic AI is the dynamo. What will transform luxury is **agentification**: Maison by Maison, function by function, decision by decision.

*What role do you choose to play in the triangle?*

The client is waiting. The agent, for its part, is already answering.

GO FURTHER

This study is a diagnosis. If you want to move to action, I can work with you.


Twenty years in luxury, on the Maison side (Saint Laurent, LVMH, Dior, Chanel) and the agency side (Mazarine, Publicis). I work with Maisons and executives who want to turn agentic AI into a competitive advantage — without surrendering their DNA.

[luxeaeternai.com](https://luxeaeternai.com) →

[mickael@luxeaeternai.com](mailto:mickael@luxeaeternai.com)

**Advisory**

Agentic AI transformation consulting for Luxury & Premium sectors.

 Shift Happens by luxe aeternai

Individual coaching and guidance for luxury executives who want to integrate agentic AI into their strategy and daily practice.

**Strategic intelligence**

Monitoring, analysis, and decryption of agentic AI in luxury. This study is one example.

**Newsletter luxe aeternai**

Every Thursday, a deep dive into what AI agents are changing — or not — for luxury Maisons. [Subscribe](#)

---

# Glossary

## Actionability

The capacity of an AI agent's response to lead a prospect toward a concrete action: website visit, boutique location, add to cart. A high citation score without actionability is a legacy, not a strategy. *(Parts 1 and 4)*

---

## Agalma

Concept borrowed from ancient Greek (originally a sacred ornament or precious offering), brought to light by Anne Dellière (Group Marketing and Strategic Planning Director, Richemont): the luxury object possesses a value that transcends its function, an aura irreducible to technical description. The AI agent extracts the attributes but cannot capture the agalma. *(Part 2)*

---

## AI agent (in the agentic sense)

Autonomous software entity that searches, compares, recommends, and soon purchases on behalf of the user. Distinct from a chatbot (reactive, scripted) or a recommendation algorithm (passive, list-based). The agent makes decisions by chaining steps without human intervention. *(Throughout the study)*

---

## Antechamber

Concept created in this study. The space between the client's intention and the Maison's shopfront, populated by AI agents that shape the prospect's preferences before the Maison gets to speak. Discovery, comparison, and pre-selection now take place here. Not the top of funnel — five structural differences separate them (see section 1.1). *(Parts 1 and 4)*

---

## Clienteling

Luxury retail practice of personalizing the client relationship in-store: purchase history, preferences, occasions. Digital clienteling has existed since 2013 (Tulip, BSPK, Cegid, Proximity Insight). "Agentic AI-augmented clienteling" adds five ruptures: multi-session conversational memory, autonomous decision-making, end-to-end orchestration, bold suggestions, and eventually agent-to-agent commerce. *(Parts 2 and 3)*

---

## Agentic commerce

A commerce model where autonomous AI agents perform purchasing tasks on behalf of the user: search, comparison, negotiation, transaction. Three protocols in deployment: ACP (OpenAI/Stripe), UCP (Google/Shopify), MCP (Anthropic). McKinsey projection: 3 to 5 trillion USD by 2030. *(Parts 1 and 2)*

---

## Augmented Advisor

A model where the in-store advisor is equipped with agentic AI tools (multi-session client memory, predictive recommendation, cross-channel history) that enrich their expertise without replacing the human relationship. Examples: Burberry Penguin (+24% average transaction value), Salesforce Agentforce for Retail. *(Part 2)*

---

### **Crystallization**

Concept borrowed from Stendhal (*On Love*, 1822). The mental process by which imagination enriches the desired object with imaginary perfections, fed by anticipation and absence. The central question of the study: does the AI agent, by eliminating anticipation, also eliminate crystallization? *(Parts 2 and conclusion)*

---

### **Digital body language**

The set of micro-behavioral signals from an online visitor: clicks, scrolling, time on page, hesitations. Technologies like Kahoot (LVMH Innovation Award, VivaTech 2025) read these in real time to personalize the experience without declarative data or cookies. *(Parts 1 and 2)*

---

### **Productive Friction**

Concept developed in this study. Deliberate friction that creates value in luxury: waiting list, mandatory appointment, organized scarcity. Distinct from Endured Friction (malfunction, slowness, involuntary opacity). The AI agent eliminates both by default. *(Part 2)*

---

### **GEO (Generative Engine Optimization)**

Emerging discipline aimed at optimizing a brand's visibility in AI agent responses and generative search engines, as opposed to classic SEO (traditional search engines). Specialized tools: Profound, Evertune, Otterly, BrightEdge. *(Parts 1 and 4)*

---

### **Agentic AI**

The branch of artificial intelligence where systems act autonomously, chaining complex tasks without human supervision at each step. Distinct from generative AI (which produces content) and predictive AI (which anticipates trends). Fred Cavazza speaks of "agentic transformation" to describe the shift from Web 3.0 to Web 4.0. *(Throughout the study)*

---

### **LLM (Large Language Model)**

A large-scale language model trained on massive text corpora, capable of generating natural language. ChatGPT (OpenAI, 78% market share), Gemini (Google, 8.6%), Perplexity (7%), Claude (Anthropic, 2.9%). AI agents rely on these models to understand queries and formulate responses. *(Throughout the study)*

---

### **The Bottle Law**

Concept created in *luxe aeternai* (edition #3). The AI that matters in luxury is the one you put inside the product, not the one you display in the window. Like a perfume, value is born from the invisible (the formula) as much as from the visible (the bottle).

## *(Part 2)*

---

### **Machine-readable**

Describes content structured so that an AI agent or LLM can understand and use it without loss of meaning. Machine-readable text is not technical jargon: it is text rich in semantic attributes (materials, savoir-faire, origin, edition) that the model can extract and accurately reproduce. *(Parts 1 and 4)*

---

### **Pilot purgatory**

Term used by DLG ("State of AI in Luxury 2026"). A situation in which a Maison multiplies AI pilot projects without ever scaling or making a strategic decision. According to DLG, 43% of luxury Maisons have progressed toward active implementation; the rest are stalled. *(Parts 2 and 3)*

---

### **PIM (Product Information Management)**

A centralized system for managing product information: descriptions, attributes, media, prices, availability. A semantically enriched PIM (beyond the technical spec sheet: narrative, savoir-faire, history) is a prerequisite for any agentic visibility strategy. *(Part 4)*

---

### **Love Triangle**

The structuring concept of this study. A relational configuration in which three actors — the client, the Maison, and the AI agent — coexist at each stage of the journey, with zones of control, sharing, and loss for each. The question is not whether to accept the third party, but what role to assign it. *(Parts 2 and 3)*

---

### **Agentification**

The process by which AI agents progressively insert themselves into every link in an industry's value chain — Maison by Maison, function by function, decision by decision. By analogy with electrification (Valéry, 1935 / Guyot, 2026): it is not agentic AI that transforms luxury — it is the agentification of luxury. *(Conclusion)*

---

### **ACP / UCP / MCP**

The three agentic commerce protocols in deployment. ACP (Agentic Commerce Protocol): OpenAI/Stripe, native checkout inside ChatGPT. UCP (Universal Commerce Protocol): Google/Shopify, checkout inside Gemini. MCP (Model Context Protocol): Anthropic, communication standard between agents and systems. *(Parts 1 and 4)*

---

### **Anti-laws of luxury marketing**

A set of principles formulated by Kapferer and Bastien (*Luxe Oblige*, 2008) that define what distinguishes luxury marketing from classic marketing: do not respond to demand, do not compare, make the purchase difficult, protect the client from the non-client. The AI agent violates several of these anti-laws by design. *(Parts 1 and 2)*

---

## **BtoA(2C)**

Business-to-Agent-to-Consumer. Acronym coined by Fred Cavazza to describe the shift in which brands no longer interact directly with consumers but with autonomous agents that serve as intermediaries. *(Introduction)*

---

## **Deepfake**

AI-generated content (image, video, audio) designed to imitate reality indistinguishably. In luxury, deepfakes threaten the authenticity of products and campaigns. The EU AI Act (Article 50, August 2026) will mandate labeling of all AI-generated content. *(Part 1)*

---

## **Dual readership**

The principle that all brand content (press brief, product sheet, press release) now has two audiences: the human reader (journalist, client) and the language model that will ingest and reproduce it. Adapting content for dual readership is one of the five workstreams in Part 4. *(Parts 1 and 4)*

---

## **DPP (Digital Product Passport)**

A digital identity record accompanying each physical product: materials, origin, environmental footprint, manufacturing conditions. Mandatory by 2027 under the European EPR regulation. The infrastructure on which AI agents will be able to verify a product's authenticity before recommending it. *(Part 4)*

---

## **Earned media**

Media coverage obtained without buying space: press articles, blog mentions, customer reviews, forums. AI agents draw heavily on earned media to construct their responses — 48% of citations come from reviews, listicles, and forums (Omniscient Digital). *(Parts 1 and 4)*

---

## **SSR (Server-Side Rendering)**

A web development technique where pages are generated server-side (not browser-side). AI agents cannot execute JavaScript: a client-rendered site (SPA) is invisible to LLM crawlers. SSR is a technical foundation for agentic visibility. *(Part 4)*

---

## **UHNWI (Ultra High Net Worth Individual)**

An individual whose net worth exceeds \$30 million USD. The target clientele for ultra-high-end luxury, representing a disproportionate share of Maison revenues. *(Part 2)*

---

## **VIC (Very Important Client)**

Designation used in luxury for very high-value clients. The VIC receives privileged treatment: dedicated advisor, priority access to collections, exclusive events, personalized clienteling. Agentic tension: if the agent generalizes VIC treatment, the VIC

# Bibliography

Sources organized by category. All URLs verified in March 2026.

## INDUSTRY REPORTS AND MARKET STUDIES

- Bain & Company and Altgamma (2025) *Worldwide Luxury Market Monitor, Finding a New Longevity for Luxury*, 24th edition. Available at: <https://www.bain.com/insights/finding-a-new-longevity-for-luxury/> (Accessed: March 2026).
- Bain & Company (2026) 'Why Agentic AI Demands a New Architecture', 9 mars 2026. Available at: <https://www.bain.com/insights/why-agentic-ai-demands-a-new-architecture/> (Accessed: March 2026).
- Bain & Company and Comité Colbert (2025) *Luxury and Technology: Tailoring Investment Strategies for Greater Business Impact*, 4th edition, September 2025. Available at: <https://www.bain.com/insights/luxury-and-technology-tailoring-investment-strategies-for-greater-business-impact/> (Accessed: March 2026).
- BCG (2025) *Why The Luxury Experience Needs an AI Moment*. Available at: <https://www.bcg.com/publications/2025/why-luxury-experience-needs-an-ai-moment> (Accessed: March 2026).
- BCG (2026) *Consumers Trust AI to Buy Better*. Available at: <https://www.bcg.com/publications/2026/consumers-trust-ai-to-buy-better-brands-must-adapt> (Accessed: March 2026).
- BCG (2026) *Hospitality and AI Report 2026*. Available at: <https://www.bcg.com> (Accessed: March 2026).
- Capgemini (2026) *From Hype to How: Retail AI Trends 2026*. Available at: <https://www.capgemini.com/insights/research-library/from-hype-to-how-retail-ai-trends-2026/> (Accessed: March 2026).
- Capgemini (2026) *What Matters to Today's Consumer 2026*. Available at: <https://www.capgemini.com/insights/research-library/what-matters-to-todays-consumer-2026/> (Accessed: March 2026).
- Checkout.com (2025) *Consumer AI Shopping Behaviour*, December 2025.
- Chrono24 (2025) *H1 2025 Market Report*. Available at: <https://www.chrono24.com> (Accessed: March 2026).
- Conductor (2025) *AEO/GEO Benchmarks Report*, 13 novembre 2025. Available at: <https://www.conductor.com/academy/aeo-geo-benchmarks-report/> (Accessed: March 2026).
- Convertio (2026) 'Luxury in the Age of AI Search 2026', mars 2026. Available at: <https://convertio.com/en/blog/luxury-in-the-age-of-ai-search/> (Accessed: March 2026).
- Convertio (2026) 'Complete GEO guide: optimizing visibility in LLMs', 6 March 2026. Available at: <https://convertio.com/en/blog/geo-guide-complet-llm-ia/> (Accessed: March 2026).
- Convertio (2026) 'Agentic AI and luxury: why human intervention remains the ultimate test'. Available at: <https://convertio.com/no-filter-data/ia-agentique-luxe-intervention-humaine/> (Accessed: March 2026).
- Deloitte (2026) *Global Powers of Luxury 2026*, février 2026. Available at: <https://www.deloitte.com/global/en/Industries/consumer/perspectives/global-powers-of-luxury.html> (Accessed: March 2026).

- Deloitte (2026) *The Future of Luxury During the Rise of AI*, 8 janvier 2026. Available at: <https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/articles/future-of-luxury-during-rise-of-ai.html> (Accessed: March 2026).
- Deloitte (2026) 'AI in Retail: ROI, Valuation, and Investment Strategy', 13 mars 2026. Available at: <https://www.deloitte.com/us/en/Industries/consumer/articles/ai-retail-roi-valuation-strategy.html> (Accessed: March 2026).
- DLG, Digital Luxury Group (2026) *State of AI in Luxury 2026*. Available at: <https://digitalluxurygroup.com/insight/luxuryaireport> (Accessed: March 2026).
- DLG — Digital Luxury Group (2026) 'LuxuryIQ MCP'. Available at: <https://digitalluxurygroup.com/news/luxuryiq-mcp> (Accessed: March 2026).
- Ekimetrics (2025) *Quiet AI: The Invisible Touch in Luxury and Beauty*, white paper, May 2025, updated February 2026.
- Gartner (2025) Projections: -25% traditional search traffic by end of 2026; 90% of B2B purchases managed by AI agents by 2028.
- IBM Institute for Business Value (2026) *Consumer AI Adoption Report*, janvier 2026. Available at: <https://www.ibm.com/thought-leadership/institute-business-value> (Accessed: March 2026).
- McKinsey et Business of Fashion (2025) *The State of Fashion 2026*, novembre-Décembre 2025. Available at: <https://www.businessoffashion.com/articles/technology/the-state-of-fashion-2026-report-agentic-generative-ai-shopping-commerce/> (Accessed: March 2026).
- McKinsey (2026) 'The Automation Curve in Agentic Commerce', janvier 2026. Available at: <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-automation-curve-in-agentic-commerce> (Accessed: March 2026).
- McKinsey (2026) *Europe's Agentic Commerce Moment*, mars 2026. Available at: <https://www.mckinsey.com> (Accessed: March 2026).
- NielsenIQ (2026) *State of Global Beauty 2026, Beauty Reset 2026*. Available at: <https://nielseniq.com/global/en/industries/beauty/> (Accessed: March 2026).
- Valiuz (2025) *GEO Barometre France*, December 2025. Available at: <https://www.valiuz.com> (Accessed: March 2026).
- MAD (consulting firm) (2025-2026) Luxury, talent, digital, transformation studies and analyses. Available at: <https://www.mad.fr/insights> (Accessed: March 2026).
- IFM, Institut Français de la Mode (2025-2026) French fashion and luxury market data. Consumer barometer, Fashion Reboot, economic studies. Available at: <https://www.ifmparis.fr/fr/expertise-economique-et-marketing> (Accessed: March 2026).
- Comité Colbert (2025-2026) Publications on French luxury, technology, CSR. Research partnerships with Bain & Company. Available at: <https://www.comitecolbert.com> (Accessed: March 2026).
- Converteo (2026) 'Travel & Hospitality x agentic AI: reinventing the traveler experience'. Accor, Four Seasons, Expedia +15%. Available at: <https://converteo.com/blog/travel-hospitality-ia-experience-voyageur-agentique/> (Accessed: March 2026).
- McKinsey (2025-2026) 'How AI Is Shaping Fashion's Next Chapter', hub d'articles mode x IA. Available at: <https://www.mckinsey.com/featured-insights/themes/how-ai-is-shaping-fashions-next-chapter> (Accessed: March 2026).
- Anthropic (2026) 'What 81,000 People Want from AI', 80 508 participants, 159 pays, 70 langues. 9 visions, 5 tensions "Light and Shade", 13 inquietudes. The largest multilingual qualitative study on AI. Available at: <https://www.anthropic.com/features/81k-interviews> (Accessed: March 2026).
- Odoxa / Fevad (2026) Study on agentic commerce in France: 57% neutrality, 51% data privacy, 52% autonomy. Cited via FashionNetwork.

## ACADEMIC RESEARCH

- Aggarwal, P., Muralidhar, N., Joshi, A. et al. (2024) 'GEO: Generative Engine Optimization', *ACM KDD 2024*. Princeton, Georgia Tech, IIT Delhi, Allen Institute for AI. Available at: <https://arxiv.org/abs/2311.09735> (Accessed: March 2026).
- Bagga, P. et al. (2025) 'E-GEO: Testbed for GEO in E-Commerce', 25 novembre 2025. Available at: <https://arxiv.org/abs/2511.20867> (Accessed: March 2026).
- Chen, T. et al. (2025) 'How to Dominate AI Search', 10 septembre 2025. Available at: <https://arxiv.org/abs/2509.08919> (Accessed: March 2026).
- Omniscient Digital (2026) 'How LLMs Source Brand Information, Analysis of 23,387 AI Citations', 8 janvier 2026. Available at: <https://beomniscient.com/blog/how-llms-source-brand-information/> (Accessed: March 2026).
- Omniscient Digital (2026) 'Content Types Cited in LLMs'. Available at: <https://beomniscient.com/blog/content-types-cited-in-llms/> (Accessed: March 2026).
- Omniscient Digital (2026) 'Mapping Content to Buyer Intent in AI Search'. Available at: <https://beomniscient.com/blog/mapping-content-to-buyer-intent-in-ai-search/> (Accessed: March 2026).
- ResearchGate (2023) 'LLM vs Traditional Search — experimental study'. Available at: <https://www.researchgate.net/publication/372234468> (Accessed: March 2026).

## LUXURY AND FASHION SPECIALIST PRESS

- ARTnews (2026) 'Jordan Wolfson Prada SS26 Campaign', 20 mars 2026. Available at: <https://news.artnet.com/art-world/jordan-wolfson-prada-ss26-campaign-2756331> (Accessed: March 2026).
- Business of Fashion (2025) 'How Brands Stand Out in AI Search, Interview James Cadwallader (Profound)', December 2025. Available at: <https://www.businessoffashion.com/articles/technology/the-state-of-fashion-2026-report-profound-interview-james-cadwallader/> (Accessed: March 2026).
- Business of Fashion (2025) 'What Happens for Luxury If AI Goes Bust?', novembre 2025. Available at: <https://www.businessoffashion.com/briefings/luxury/high-margin-what-happens-for-luxury-if-ai-goes-bust/> (Accessed: March 2026).
- Business of Fashion (2026) 'Is AI Antithetical to Luxury?', 25 mars 2026. Available at: <https://www.businessoffashion.com/briefings/technology/is-ai-antithetical-to-luxury/> (Accessed: March 2026).
- Business of Fashion (2026) 'Estee Lauder Data Officer Brian Franz', mars 2026. Available at: <https://www.businessoffashion.com/news/beauty/estee-lauder-data-officer-brian-franz/> (Accessed: March 2026).
- Business of Fashion (2026) 'Unpacking Beauty's Agentic AI Toolkit', February 2026.
- Business of Fashion (2026) 'The Beauty Brands ChatGPT Tells People to Buy'. Available at: <https://www.businessoffashion.com/articles/beauty/the-beauty-brands-chatgpt-tells-people-to-buy/> (Accessed: March 2026).
- Cosmetics Design Europe (2026) 'L'Oréal and NVIDIA Plan to Bring Atomic-Scale AI to Beauty Formulation', 19 March 2026.
- Digital Commerce 360 (2025) 'Ralph Lauren AI Assistant Ask Ralph', 12 septembre 2025. Available at: <https://www.digitalcommerce360.com/2025/09/12/ralph-lauren-ai-assistant-ask-ralph-microsoft/> (Accessed: March 2026).
- eMarketer (2026) 'AI Ads Win Over Millennials, Gen Z Still Isn't Convinced', 12 mars 2026. Available at: <https://www.emarketer.com/content/ai-ads-win-over-millennials--gen-z-still-isn-t-convinced> (Accessed: March 2026).
- Fashionista (2026) 'Anti-AI Fashion Trend: Human Craft', mars 2026. Available at: <https://fashionista.com/2026/03/anti-ai-fashion-trend-human-craft> (Accessed: March 2026).

- FashionNetwork (2026) 'Alexis Mabille Propels Haute Couture into the Age of AI'. Available at: <https://us.fashionnetwork.com/news/Alexis-mabille-propels-haute-couture-into-the-age-of-ai,1802335.html> (Accessed: March 2026).
- Fibre2Fashion (2026) 'EU AI Act and Fashion: A Practical Map of Tools, Risks and Timelines'. Available at: <https://www.fibre2fashion.com/industry-article/10992/eu-ai-act-and-fashion-a-practical-map-of-tools-risks-and-timelines> (Accessed: March 2026).
- Forbes (2026) 'How Luxury Fashion Is Rebuilding Its Creative and Commercial Stack Using AI', 24 mars 2026. Available at: <https://www.forbes.com/sites/josipamajic/2026/03/24/how-luxury-fashion-is-rebuilding-its-creative-and-commercial-stack-using-ai/> (Accessed: March 2026).
- Fortune (2026) 'Profound Raises \$96 Million as AI Threatens Search', 24 fevrier 2026. Available at: <https://fortune.com/2026/02/24/exclusive-as-ai-threatens-search-profound-raises-96-million-to-help-brands-stay-visible/> (Accessed: March 2026).
- Fortune (2026) 'Saks Global Near Bankruptcy, Richard Baker's Risky Dealmaking', janvier 2026. Available at: <https://fortune.com/2026/01/08/saks-global-near-bankruptcy-richard-baker-risky-dealmaking-business-basics/> (Accessed: March 2026).
- Glossy (2026) 'Luxury Fashion's AI Marketing Experiments Hit a Turning Point', fevrier 2026. Available at: <https://www.glossy.co/fashion/luxury/luxury-fashions-ai-marketing-experiments-hit-a-turning-point/> (Accessed: March 2026).
- Jing Daily (2026) '2026 AI Outlook: When Human Value Becomes the Real Advantage', janvier 2026. Available at: <https://jingdaily.com/posts/2026-ai-outlook-when-human-value-becomes-the-real-advantage> (Accessed: March 2026).
- Jing Daily (2026) 'When the Algorithm Decides: How Agentic AI Is Challenging Luxury Brands', 20 mars 2026. Available at: <https://jingdaily.com/posts/when-the-algorithm-decides-how-agentic-ai-is-challenging-luxury-brands> (Accessed: March 2026).
- Jing Daily (2026) 'Why Balenciaga Sticks to a Human-Centered Approach in an AI-Driven World', 10 mars 2026. Available at: <https://jingdaily.com/posts/why-balenciaga-sticks-to-a-human-centered-approach-in-an-ai-driven-world> (Accessed: March 2026).
- Luxury Daily (2025) 'Estee Lauder Co's Jo Malone Launch AI Scent Advisor'. Available at: <https://www.luxurydaily.com/estee-lauder-cos-jo-malone-launch-ai-scent-advisor/> (Accessed: March 2026).
- Luxury Daily (2026) 'Luxury Unfiltered: The AI Reckoning Will Separate True Luxury Brands from Imitators'. Available at: <https://www.luxurydaily.com/luxury-unfiltered-the-ai-reckoning-will-separate-true-luxury-brands-from-imitators> (Accessed: March 2026).
- Luxury Society (2026) Weir, D. 'The Ghost in the Machine: Why AI's Biggest Challenge in Luxury Is Organisational, Not Technical', janvier 2026. Available at: <https://luxurysociety.com/en/the-ghost-in-the-machine-why-ais-biggest-challenge-in-luxury-is-organisational-not-technical/> (Accessed: March 2026).
- Lyst (2025) *Lyst Index Q3 et Q4 2025*.
- Modern Retail (2026) 'Why the AI Shopping Agent Wars Will Heat Up in 2026'. Available at: <https://www.modernretail.co/technology/why-the-ai-shopping-agent-wars-will-heat-up-in-2026/> (Accessed: March 2026).
- Modern Retail (2026) 'Shopify Has Quietly Set Boundaries for Buy-for-Me AI Bots'. Available at: <https://www.modernretail.co/technology/shopify-has-quietly-set-boundaries-for-buy-for-me-ai-bots-on-merchant-sites/> (Accessed: March 2026).
- Parmentier, M. (2025) 'AI and luxury beauty: the invisible but irresistible revolution', *Journal du Luxe*, June 2025.
- Retail Brew (2026) 'AI News Dominated the Buzz at Shoptalk Spring', 27 mars 2026. Available at: <https://www.retailbrew.com/stories/2026/03/27/ai-news-dominated-the-buzz-at-shoptalk-spring> (Accessed: March 2026).
- Retail Dive (2026) 'Saks Global Hundreds of Vendors Resumed Shipments', fevrier 2026. Available at: <https://www.retaildive.com/news/saks-global-hundreds-vendors-resumed-shipments-bankruptcy-financing/813160/> (Accessed: March 2026).

- The Fashion Law (2026) 'Agentic AI Is Reshaping Commerce. Is the Law Ready?', 13 mars 2026. Available at: <https://www.thefashionlaw.com/agentic-ai-is-reshaping-commerce-is-the-law-ready/> (Accessed: March 2026).
- WWD (2026) 'Kering CEO Luca de Meo Strategic Plan for Technology', 5 février 2026. Available at: <https://wwd.com/business-news/technology/kering-ceo-luca-de-meo-strategic-plan-for-technology-1238485796/> (Accessed: March 2026).
- WWD (2026) 'How AI Agents Are Transforming Fashion Brand Discovery'. Available at: <https://wwd.com/business-news/technology/ai-agents-transform-fashion-brand-discovery-storytelling-1238556291/> (Accessed: March 2026).
- WWD (2026) 'Kering Eyewear x Google AI Smart Glasses Android XR'. Available at: <https://wwd.com/business-news/technology/kering-eyewear-google-ai-smart-glasses-android-xr-1237805134/> (Accessed: March 2026).
- WWD (2026) 'Gucci Show Demna Reaction', 2 mars 2026. Available at: <https://wwd.com/fashion-news/designer-luxury/gucci-show-demna-reaction-bernstein-1238640197/> (Accessed: March 2026).
- WWD et CFDA (2026) 'NYFW Dinner: Tech and AI's Role in Luxury Fashion'. Available at: <https://wwd.com/business-news/business-features/wwd-cfda-nyfw-dinner-ai-luxury-fashion-1238626891/> (Accessed: March 2026).
- Journal du Luxe (2025-2026) [journalduluxe.fr](http://journalduluxe.fr). French luxury news, Intelligence Live events, Futur & Luxe webinars. Co-founded by Eric Briones. Available at: <https://journalduluxe.fr> (Accessed: March 2026).
- Les Echos (2025-2026) Luxury and Cosmetics section. Financial and strategic coverage of French luxury groups. Available at: <https://www.lesechos.fr/industrie-services/conso-distribution/luxe-cosmetiques> (Accessed: March 2026).
- Financial Times (2025-2026) Luxury section. Strategic analysis of luxury groups, financial markets, M&A. Available at: <https://www.ft.com/luxury> (Accessed: March 2026).
- Miss Tweed (2025-2026) Luxury intelligence: scoops, strategic moves, appointments, M&A. Premium reference newsletter for industry decision-makers. Available at: <https://www.misstweed.com> (Accessed: March 2026).
- Vogue Business (2025-2026) [voguebusiness.com](http://voguebusiness.com). Tech, innovation, digital, sustainability in luxury and fashion. Available at: <https://www.voguebusiness.com> (Accessed: March 2026).
- Robb Report (2025-2026) Luxury lifestyle media: watchmaking, automotive, gastronomy, travel. Premium coverage of product innovations and experiences. Available at: <https://robbreport.com> (Accessed: March 2026).
- Wallpaper\* (2025-2026) Design, architecture, luxury, fashion. Editorial reference on contemporary aesthetics and the new codes of luxury. Available at: <https://www.wallpaper.com> (Accessed: March 2026).
- The Cut (2025-2026) Fashion, culture, beauty. Critical analysis of trends and cultural movements around luxury. Available at: <https://www.thecut.com> (Accessed: March 2026).
- CPP-Luxury (2025-2026) [cpp-luxury.com](http://cpp-luxury.com). Economic intelligence platform on the luxury industry: ESG, governance, financial data, appointments. Available at: <https://cpp-luxury.com> (Accessed: March 2026).
- Luxury Society (2026) Publication of Digital Luxury Group (DLG). Strategic analysis on digital luxury, AI, agentic commerce, organization. Available at: <https://luxurysociety.com> (Accessed: March 2026).
- Hodinkee (2026) Editorial reference in luxury watchmaking. Coverage of TAG Heuer, appointments, Watches & Wonders salon. Cited in edition #5 (Beatrice Goasglas appointment). Available at: <https://www.hodinkee.com> (Accessed: March 2026).
- Worldtempus (2026) Swiss watchmaking media. Watches & Wonders coverage, appointments, technical innovations. Cited in edition #5. Available at: <https://www.worldtempus.com> (Accessed: March 2026).
- Bob's Watches (2025) 'Rolex AI Timegrapher Patent Analysis', 20 December 2025. Cited in edition #5 (Rolex patent). Available at: <https://www.bobswatches.com> (Accessed: March 2026).
- Timepiece Trading (2026) 'Rolex Neural Network Timegrapher', 27 January 2026. Cited in edition #5. Available at: <https://www.timepiecetrading.com> (Accessed: March 2026).
- Grey Market Magazine (2026) 'Rolex Timegrapher Patent, Neural Network Sound Analysis', 22 January 2026. Cited in edition #5. Available at: <https://greymarketmagazine.com> (Accessed: March 2026).
- CNN Style (2026) 'Demna Gucci AI', 27 February 2026. Cited in edition #3. Available at: <https://edition.cnn.com/2026/02/27/style/demna-gucci-ai> (Accessed: March 2026).

- Ad Age (2025) 'Valentino DeVain, Surreal and Sloppy Social Ad', December 2025. Cited in edition #3. Available at: <https://adage.com/technology/ai/ad-of-the-week-valentino-surreal-and-sloppy-social-ad/> (Accessed: March 2026).
- Franceinfo Culture (2026) 'Paris Fashion Week AW 2026-2027 — Antonin Tron (Balmain), Pieter Mulier (Alaïa)', March 2026. Cited in edition #3. Available at: <https://www.franceinfo.fr/culture/mode/fashion-week/> (Accessed: March 2026).
- FashionNetwork (2026) 'Agentic commerce: what adoption, what barriers, what neutrality', 2026. Odoxa/Fevad data: 57% neutrality, 51% data privacy, 52% autonomy. Available at: <https://fr.fashionnetwork.com/news/la-quelle-adoption-quels-freins-et-quelle-neutralite-pour-le-nouveau-commerce-agentique-,1807557.html> (Accessed: March 2026).
- MARS Magazine (2026) 'Prada x Jordan Wolfson SS26 — AI Controversy', 27 mars 2026. Cité dans l'édition #6. Available at: <https://marsmag.com> (Accessed: March 2026).
- Consumer Goods Technology (2026) 'Kering Appoints Pierre Houles', 17 March 2026. Cited in edition #6. Available at: <https://consumergoods.com> (Accessed: March 2026).
- Variety (2026) Data on luxury consumer perception of AI-generated content (70% of buyers). Cited in edition #6.
- IAB, Interactive Advertising Bureau (2026) Gen Z data and AI-generated content (39%). Cited in edition #6.

## TECHNOLOGY AND AI PRESS

- Bloomberg (2026) 'OpenAI Finalizes \$110 Billion Funding at \$730 Billion Valuation', 27 février 2026. Available at: <https://www.bloomberg.com/news/articles/2026-02-27/openai-finalizes-110-billion-funding-at-730-billion-valuation> (Accessed: March 2026).
- CNBC (2026) 'Ray-Ban Maker EssilorLuxottica Triples Sales of Meta AI Glasses', 11 février 2026. Available at: <https://www.cnbc.com/2026/02/11/ray-ban-maker-essilorluxottica-triples-sales-of-meta-ai-glasses.html> (Accessed: March 2026).
- CNBC (2026) 'Nvidia GTC AI Jensen Huang', 13 mars 2026. Available at: <https://www.cnbc.com/2026/03/13/nvidia-gtc-ai-jensen-huang-cpu-gpu.html> (Accessed: March 2026).
- CNBC (2026) 'Saks Global Files for Bankruptcy Protection', janvier 2026. Available at: <https://www.cnbc.com/2026/01/14/saks-global-files-for-bankruptcy-protection.html> (Accessed: March 2026).
- Retail Brew (2026) 'Amazon CEO Says Shoppers Will Choose Retailers Over Third-Party AI Platforms', 10 février 2026. Available at: <https://www.retailbrew.com/stories/2026/02/10/amazon-ceo-says-shoppers-will-choose-retailers-over-third-party-ai-platforms> (Accessed: March 2026).
- Intelligent CIO (2026) 'The Estee Lauder Companies and Jo Malone London Introduce AI-Powered Scent Advisor Experience', 8 janvier 2026. Available at: <https://www.intelligentcio.com/eu/2026/01/08/the-estee-lauder-companies-and-jo-malone-london-introduce-ai-powered-scent-advisor-experience/> (Accessed: March 2026).
- PYMNTS (2026) 'The RealReal Expands AI Push as Luxury Buyers Spend More'. Available at: <https://www.pymnts.com/news/retail/2026/the-realreal-expands-ai-push-as-luxury-buyers-spend-more/> (Accessed: March 2026).
- PYMNTS (2026) 'JPMorgan Payments and Mirakl Form Agentic Commerce Pact', 10 mars 2026. Available at: <https://www.pymnts.com/artificial-intelligence-2/2026/jpmorgan-payments-and-mirakl-form-agentic-commerce-pact/> (Accessed: March 2026).
- RD World Online (2026) 'OpenAI Facing \$14B Losses in 2026'. Available at: <https://www.rdworldonline.com/facing-14b-losses-in-2026-openai-is-now-seeking-100b-in-funding-but-can-it-ever-turn-a-profit/> (Accessed: March 2026).
- Skift (2026) 'OpenAI ChatGPT Checkout Walkback', mars 2026. Available at: <https://skift.com/2026/03/05/openai-chatgpt-checkout-walkback/> (Accessed: March 2026).
- TechCrunch (2026) 'OpenAI Raises \$110B in One of the Largest Private Funding Rounds in History', 27 février 2026. Available at: <https://techcrunch.com/2026/02/27/openai-raises-110b-in-one-of-the-largest-private-funding-rounds-in-history/> (Accessed: March 2026).

- TechCrunch (2026) 'So We're Getting Prada Meta AI Glasses, Right?', février 2026. Available at: <https://techcrunch.com/2026/02/26/so-were-getting-prada-meta-ai-glasses-right/> (Accessed: March 2026).
- TechCrunch (2026) 'OpenAI Acquires Promptfoo to Secure Its AI Agents', 9 mars 2026. Available at: <https://techcrunch.com/2026/03/09/openai-acquires-promptfoo-to-secure-its-ai-agents/> (Accessed: March 2026).
- The Next Web (2026) 'Anthropic Commits \$100M to Claude Partner Network'. Available at: <https://thenextweb.com/news/anthropic-commits-100m-to-claude-partner-network> (Accessed: March 2026).
- Artificielles / Numerama (2025-2026) AI coverage in French. Factual rigor, fact-checking. Editorial reference for luxe aeternai. Available at: <https://www.numerama.com/tag/artificielles/> (Accessed: March 2026).
- Maddyness (2025-2026) Startups, tech, innovation in France. AI coverage, agentic commerce, French tech ecosystem. Available at: <https://www.maddyness.com> (Accessed: March 2026).
- FrenchWeb (2025-2026) Tech and business news in France. Digital transformation, AI, startups. Available at: <https://www.frenchweb.fr> (Accessed: March 2026).
- The Verge (2025-2026) Major model announcements, AI agents, agentic commerce. Available at: <https://www.theverge.com> (Accessed: March 2026).
- Wired (2025-2026) AI and society analysis, consumer impact, autonomous agents. Available at: <https://www.wired.com> (Accessed: March 2026).
- MIT Technology Review (2025-2026) AI, agents, foundational technology. Available at: <https://www.technologyreview.com> (Accessed: March 2026).
- Forbes (2026) 'Beatrice Goasglas appointed CEO of TAG Heuer', 13 March 2026. Cited in edition #5. Available at: <https://www.forbes.com> (Accessed: March 2026).
- Business Standard (2026) 'Estee Lauder Transfers \$500M Contract from Wipro to Accenture', March 2026. Cited in edition #3. Available at: [https://www.business-standard.com/companies/news/wipro-set-to-lose-estee-lauder-contract-to-accenture-126030300968\\_1.html](https://www.business-standard.com/companies/news/wipro-set-to-lose-estee-lauder-contract-to-accenture-126030300968_1.html) (Accessed: March 2026).

## GEO DATA AND BENCHMARKS

- AuthorityTech (2026) 'AI Share of Voice: Measure & Grow LLM Brand Presence', 25 février 2026. Available at: <https://authoritytech.io/blog/ai-share-of-voice-measure-grow-llm-brand-presence-2026> (Accessed: March 2026).
- Azoma (2026) 'Luxury Brand AI Visibility', 12 mars 2026. Available at: <https://www.azoma.com> (Accessed: March 2026).
- Backlinko (2026) 'Fashion AI, How Brands Appear in LLM Recommendations'. Available at: <https://backlinko.com/fashion-ai> (Accessed: March 2026).
- ConvertMate (2026) *GEO Benchmark 2026*, 29 mars 2026. 12,500 queries, 8,000 domains. Available at: <https://www.convertmate.io/research/geo-benchmark-2026> (Accessed: March 2026).
- Evertune (2026) 'Brand Visibility in LLM Responses', 24 février 2026. \$19M Series A (Felicis Ventures), AI Brand Index, 1M+ prompts/month per brand. (evertune.io unavailable as of 1 April 2026).
- Fifth Shelf (2026) 'How to Audit Brand Visibility on LLMs', 6 janvier 2026. Available at: <https://www.fifthshelf.com/blogs/how-to-audit-brand-visibility-on-llms> (Accessed: March 2026).
- GenOptima (2026) *GEO/AEO Ranking Methodology v2.0*, 28 mars 2026. Available at: <https://www.gen-optima.com/geo/geo-aeo-ranking-methodology> (Accessed: March 2026).
- Incremys (2026) 'LLM Audit, 8-Step Methodology', 16 mars 2026. Available at: <https://www.incremys.com/en/resources/blog/llm-audit> (Accessed: March 2026).
- Platelunch Collective (2026) '10 AEO Platforms Compared', février 2026. Available at: <https://www.platelunchcollective.com/aeo-tools-analysis-10-platforms-compared-beyond-marketing-materials-because-youve-been-searching-for-them/> (Accessed: March 2026).

- Profound (2025) *10-Step GEO Framework*, 1 July 2025. Available at: <https://www.tryprofound.com/guides/generative-engine-optimization-geo-guide-2025> (Accessed: March 2026).
- Profound (2026) *Leaderboard, GEO Market*. Available at: <https://www.tryprofound.com/leaderboard> (Accessed: March 2026).
- SE Ranking / Clutch (2026) 'Generative Engine Optimization Firms', February 2026. Available at: <https://clutch.co/seo-firms/generative-engine-optimization> (Accessed: March 2026).

## CORPORATE SOURCES AND PRESS RELEASES

- Anthropic (2026) 'Claude Partner Network — \$100M Investment'. Available at: <https://thenextweb.com/news/anthropic-commits-100m-to-claude-partner-network> (Accessed: March 2026).
- BrandXR / Clinique (2026) 'In-store AI mirrors: +30% average basket'. Available at: <https://www.brandxr.io> (Accessed: March 2026).
- Burberry / DataIQ (2025) 'Most Innovative Use of AI (Global) — Burberry Penguin'. Available at: <https://www.dataiq.global/award-winner/most-innovative-use-of-ai-global-burberry/> (Accessed: March 2026).
- Cpp-Luxury (2025) 'Hermes to Establish an Artificial Intelligence Governance Committee'. Available at: <https://cpp-luxury.com/hermes-to-establish-an-artificial-intelligence-governance-committee/> (Accessed: March 2026).
- Entrupy (2026) 'Luxury Authentication AI, 99.86% Accuracy, \$7B Verified'. Available at: <https://www.entrupy.com> (Accessed: March 2026).
- Estee Lauder Companies (2025) 'Jo Malone London AI-Powered Scent Advisor', 2 December 2025. Available at: <https://www.elcompanies.com/en/news-and-media/newsroom/press-releases/2025/12-02-2025-121509645> (Accessed: March 2026).
- Google (2026) 'Agentic Commerce AI Tools — Universal Commerce Protocol (UCP)'. Available at: <https://blog.google/products/ads-commerce/agentic-commerce-ai-tools-protocol-retailers-platforms/> (Accessed: March 2026).
- Google (2026) 'AI-Generated Brand Film, Moncler', mars 2026. Available at: <https://business.google.com/us/think/ai-excellence/ai-generated-brand-film-moncler/> (Accessed: March 2026).
- Google Cloud (2026) 'LVMH AI'. Available at: <https://cloud.google.com/blog/topics/retail/lvmh-lvmh-ai> (Accessed: March 2026).
- Hermes Finance (2025) 'Key Figures, Marge opérationnelle >40%'. Available at: <https://finance.hermes.com/en/key-figures/> (Accessed: March 2026).
- Heuritech (2026) 'Première Vision 2026 Spring Summer 27, Trend Prediction 91% Accuracy'. Available at: <https://heuritech.com/articles/première-vision-2026-spring-summer-27/> (Accessed: March 2026).
- Kahoona (2025) 'Case Study, Global Luxury Fashion Leader'. Available at: <https://kahoona.io> (Accessed: March 2026).
- Kering (2026) 'Pierre Houles Appointed Chief Digital, AI and IT Officer', 17 mars 2026. Available at: <https://www.kering.com/en/news/pierre-houles-appointed-chief-digital-ai-and-it-officer-at-kering/> (Accessed: March 2026).
- L'Oreal Finance (2026) 'L'Oreal and NVIDIA to Accelerate Beauty Discovery Powered by Predictive AI Science', 17 mars 2026. Available at: <https://www.loreal.com/en/press-release/group/loreal-and-nvidia-to-accelerate-beauty-discovery-powered-by-predictive-ai-science/> (Accessed: March 2026).
- Mastercard (2026) 'Verifiable Intent, Cryptographic Framework for Agentic Commerce', mars 2026. Available at: <https://www.mastercard.com/us/en/news-and-trends/stories/2026/verifiable-intent.html> (Accessed: March 2026).
- NRF (2026) 'LVMH Balances Luxury Innovation in the Digital Age', 11 janvier 2026. Available at: <https://nrf.com/blog/lvmh-balances-luxury-innovation-in-the-digital-age> (Accessed: March 2026).

- NRF (2026) 'Where Craft Meets Intelligence, session officielle NRF 2026'. Available at: <https://nrffbigshow.nrf.com/session/2352> (Accessed: March 2026).
- NXN Labs (2026) 'NRF 2026 Innovators Showcase Top 50'. Available at: <https://www.prnewswire.com/news-releases/nrf-2026-nxn-labs-named-to-innovators-showcase-top-50-following-european-luxury-wins-302697760.html> (Accessed: March 2026).
- OpenAI (2025) 'Buy it in ChatGPT, Agentic Commerce Protocol'. Available at: <https://openai.com/index/buy-it-in-chatgpt/> (Accessed: March 2026).
- OpenAI (2025) 'ChatGPT Shopping Research', novembre 2025. Available at: <https://openai.com/index/chatgpt-shopping-research/> (Accessed: March 2026).
- Ralph Lauren (2025) 'Ask Ralph AI Assistant Launch', 9 septembre 2025. Available at: [https://corporate.ralphlauren.com/pr\\_250909\\_AskRalph.html](https://corporate.ralphlauren.com/pr_250909_AskRalph.html) (Accessed: March 2026).
- Salesforce (2024) 'Saks Data Cloud Agentforce Ecommerce', 17 septembre 2024. Available at: <https://www.salesforce.com/news/press-releases/2024/09/17/saks-data-cloud-agentforce-ecommerce/> (Accessed: March 2026).
- Sephora Newsroom (2026) 'Sephora App in ChatGPT Brings a New Personalized Beauty Experience', 24 March 2026. Available at: <https://newsroom.sephora.com/sephora-app-in-chatgpt-brings-a-new-personalized-beauty-experience/> (Accessed: March 2026).
- Shiseido (2026) 'VOYAGER, AI-Powered R&D Platform'. Available at: <https://corp.shiseido.com/en/news/detail.html?n=0000000004127> (Accessed: March 2026).
- Shopify (2026) 'Agentic Storefronts, Winter '26 Edition'. Available at: <https://www.shopify.com/news/winter-26-edition-agentic-storefronts> (Accessed: March 2026).
- Spreedly (2026) 'Agentic Commerce Live Channel for Merchants', 4 mars 2026. Available at: <https://www.prnewswire.com/news-releases/spreedly-makes-agentic-commerce-a-live-channel-for-merchants-302703191.html> (Accessed: March 2026).
- The RealReal (2025) 'Third Quarter 2025 Results'. Available at: <https://investor.therealreal.com/news-releases/news-release-details/realreal-announces-third-quarter-2025-results/> (Accessed: March 2026).
- The RealReal (2026) 'Fourth Quarter and Full Year 2025 Results'. Available at: <https://www.globenewswire.com/news-release/2026/02/26/3246067/0/en/The-RealReal-Announces-Fourth-Quarter-and-Full-Year-2025-Results.html> (Accessed: March 2026).
- European Union (2024) *Regulation (EU) 2024/1689, EU AI Act*. Available at: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32024R1689> (Accessed: March 2026).
- Winnow Solutions (2026) 'Mandarin Oriental AI Partnership'. Available at: <https://www.winnowsolutions.com> (Partnership page removed; accessed: March 2026).
- LVMH (2025-2026) Newsroom. Official press releases, financial results, innovation announcements. Available at: <https://www.lvmh.fr/actualites/> (Accessed: March 2026).
- LVMH (2025-2026) Investor relations. Quarterly and annual results. Available at: <https://www.lvmh.fr/groupe/relations-investisseurs/> (Accessed: March 2026).
- Kering (2025-2026) Investor relations. Financial results, Capital Markets Day. Available at: <https://www.kering.com/fr/finance/> (Accessed: March 2026).
- Richemont (2025-2026) Investor relations. Available at: <https://www.richemont.com/investors/> (Accessed: March 2026).
- Prada Group (2025-2026) Relations investisseurs. Available at: <https://www.pradagroup.com/en/investors.html> (Accessed: March 2026).
- Dorchester Collection (2026) 'Metis 2.0, AI in Luxury Hospitality', Ana Brant, Director of Global Guest Experience. Cited in edition #5. Available at: <https://www.dorchestercollection.com> (Accessed: March 2026).
- Aman (2026) 'AI Wellness Concierge, Biometric Spa Experience'. Cited in edition #5. Available at: <https://www.aman.com> (Accessed: March 2026).

- Peninsula Hotels (2026) 'PenChat, AI Concierge', Gareth Roberts. +35% service speed, +28% engagement. Cited in edition #5. Available at: <https://www.peninsula.com> (Accessed: March 2026).
- Nexi Group et Google Cloud (2026) 'Agentic Commerce Across Europe', 3 mars 2026. Cited in edition #3. Available at: <https://www.prnewswire.com/news-releases/nexi-group-and-google-cloud-collaborate-to-drive-agentic-commerce-across-europe-302702017.html> (Accessed: March 2026).
- FashionNetwork (2025) 'LVMH Innovation Award VivaTech, Kahoona Best Business Prize', June 2025. Cited in edition #6. Available at: <https://fr.fashionnetwork.com> (Accessed: March 2026).
- Forbes Travel Guide (2026) 'Summit 2026, AI in Luxury Hospitality', February 2026. Cited in edition #4. Available at: <https://www.forbestravelguide.com> (Accessed: March 2026).

## REFERENCE WORKS

- Benjamin, W. (1935) *The Work of Art in the Age of Mechanical Reproduction*. Paris : Allia.
- Bourdieu, P. (1979) *Distinction: A Social Critique of the Judgement of Taste*. Paris: Éditions de Minuit.
- Briones, E. (2021) *Luxe et Digital*. Paris : Dunod.
- Festinger, L. (1957) *A Theory of Cognitive Dissonance*. Stanford : Stanford University Press.
- Galienni, S. (2024) *Luxe & IA: opportunités et révolution des usages*. Paris: Dunod.
- Han, B.-C. (2012) *The Transparency Society*. Berlin: Matthes & Seitz. French trans. PUF, 2017.
- Han, B.-C. (2015) *Saving Beauty*. Berlin: Fischer. French trans. Arles: Actes Sud, 2016.
- Kapferer, J.-N. and Bastien, V. (2025) *The Luxury Strategy: Break the Rules of Marketing to Build Luxury Brands*. 3rd edition. London: Kogan Page. [Updated edition integrating digitalization, sustainability, and new cultural powers].
- Rouchet, E. (2025) *Ultra Intelligence*. Paris : Odile Jacob.
- Schein, E. (2017) *Organizational Culture and Leadership*. 5th edition. Hoboken: Wiley.
- Simmel, G. (1904) 'Philosophy of Fashion', in *The Tragedy of Culture and Other Essays*. French trans. Paris: Allia.
- Simmel, G. (1907) 'Sociology of the Senses', in *Soziologie*. Berlin: Duncker & Humblot.
- Stendhal (1822) *De l'amour*. Paris.
- Weiser, M. (1991) 'The Computer for the 21st Century', *Scientific American*, September 1991.
- Weil, S. (1942) *Letter to Joe Bousquet*, 13 April 1942.
- Kapferer, J.-N. (2015) *Kapferer on Luxury: How Luxury Brands Can Grow Yet Remain Rare*. London: Kogan Page. Intellectual anchor of the newsletter: anti-laws of luxury marketing, "Soulful Tech", luxury invariants.
- Kapferer, J.-N. and Bastien, V. (2008, 2012) *Luxe Oblige*. Paris: Eyrolles. Recurring reference in *luxe aeternai*, editorial signature.
- Baudrillard, J. (1981) *Simulacra and Simulation*. Paris: Galilée. Cited in edition #6 (the Antechamber, formation of desire).
- Jankélévitch, V. (1957) *Le Je-ne-sais-quoi et le Presque-rien*. Paris: PUF. Recurring philosophical reference for Mickaël Tsakiris on the ineffable quality of luxury.
- Ive, J. (2010) Core77 interview, cited in edition #5: "It is only when you personally work with a material with your hands that you come to understand its true nature."
- Alaïa, A. Quotes on working with material and cut. Cited in edition #5 (THE STORY: L'Oréal).
- Briones, E. and Casper, G. (2014) *Luxe et Digital: stratégies pour une digitalisation singulière du luxe*. Paris: Dunod. Pioneering analysis of the digital transformation of luxury Maisons.

- Briones, E. (2019) *Le Choc Z: la Génération Z, cette inconnue*. Paris: Dunod. Behavioral analysis of Generation Z, relevant to AI agent adoption.
- Prigent, L. (2019) *J'aime la mode mais c'est tout ce que je déteste*. Paris: Grasset. Critical analysis of the contemporary fashion system.

## NEWSLETTERS AND EDITORIAL ANALYSIS

- Amodei, D. (2026) 'The Adolescence of Technology', essai, janvier 2026. Available at: <https://www.darioamodei.com/essay/the-adolescence-of-technology> (Accessed: March 2026).
- Cavazza, F. (2026) 'Agentization marks a new stage of digital maturity toward Web4 and Enterprise4', 6 March 2026. Available at: <https://fredcavazza.net/2026/03/06/lagentisation-amorce-une-nouvelle-etape-de-maturite-numerique-menant-au-web4-et-a-lentreprise4/> (Accessed: March 2026).
- Ekimetrics et Journal du Luxe (2026) Webinar 'Futur & Luxe', 10 mars 2026. Available at: <https://journalduluxe.fr> (Accessed: March 2026).
- Harvard Business Review (2025) 'How a French Spirits Company Created Employee Buy-In for AI' (Pernod Ricard), December 2025. Available at: <https://hbr.org/2025/12/how-a-french-spirits-company-created-employee-buy-in-for-ai> (Accessed: March 2026).
- Harvard Business Review (2026) 'How Brands Can Adapt When AI Agents Do the Shopping', février 2026. Available at: <https://hbr.org/2026/02/how-brands-can-adapt-when-ai-agents-do-the-shopping> (Accessed: March 2026).
- Harvard Business Review (2026) 'Preparing Your Brand for Agentic AI', mars-avril 2026. Available at: <https://hbr.org/2026/03/preparing-your-brand-for-agentic-ai> (Accessed: March 2026).
- Le Grand Continent (2026) 'AI is an existential risk: Dario Amodei's alert, full text with commentary', 28 January 2026. Available at: <https://legrandcontinent.eu/fr/2026/01/28/lia-est-un-risque-existential-lalerte-de-dario-amodei-texte-integral-commentaire/> (Accessed: March 2026).
- Cavazza, F. (2025-2026) FredCavazza.net. 25 years of expertise, ~4,000 articles on digital transformation, the agentic web, digital acculturation. Key article: 'The agentic web' (January 2026). Available at: <https://fredcavazza.net> (Accessed: March 2026).
- Guyot, R. (2025-2026) *Mind Fooled* (Substack) / AI Discipline. Analysis of cognitive biases exploited by AI, applied neurosciences. Interviewed in newsletter #6 (March 2026). Available at: <https://mindfooled.substack.com> (Accessed: March 2026).
- Azhar, A. (2025-2026) *Exponential View*. Reference newsletter on exponential technologies and their societal impact. Structural editorial reference for luxe aeternai (structure, analytical density, thesis-first opening). Available at: <https://www.exponentialview.co> (Accessed: March 2026).
- Sherman, L. (2025-2026) *Line Sheet* (Puck News). Personal tone, strong opinions, narrative openings. Editorial reference for luxe aeternai on personal style. Available at: <https://puck.news/newsletter-author/line-sheet/> (Accessed: March 2026).
- Thompson, B. (2025-2026) *Stratechery*. Strategic analysis of tech platforms, business models, AI agents. Structural reference on ecosystem dynamics. Available at: <https://stratechery.com> (Accessed: March 2026).
- The Neuron (2025-2026) AI newsletter with a strong editorial personality. Reference for luxe aeternai on tone and reader engagement. Available at: <https://www.theneurondaily.com> (Accessed: March 2026).
- TLDR AI (2025-2026) Condensed daily AI briefing, 920K subscribers. Editorial team: Andrew Tan (ex-Google AI), Jacob Turner (Anthropic). Available at: <https://tldr.tech/ai> (Accessed: March 2026).
- Briones, E. (2025-2026) CEO and co-founder of Journal du Luxe. Author of *Luxe et Digital* (Dunod, 2021) and *Le Choc Z* (Dunod, 2019). Regular contributor to the editorial voice and writing guidelines of luxe aeternai. Available at: <https://journalduluxe.fr> (Accessed: March 2026).

- Data But Make It Fashion (2025-2026) Data × fashion newsletter, 51K subscribers. Benchmark and editorial reference. Available at: <https://www.databutmakeitfashion.com/> (Accessed: March 2026).

## CONSUMER DATA AND RESEARCH PLATFORMS

- Alhena AI (2026) 'How Luxury and DTC Brands Use AI Agents'. Available at: <https://alhena.ai/blog/how-luxury-and-dtc-brands-use-ai-agents/> (Accessed: March 2026).
- Alhena AI (2026) 'Luxury AI Concierge Workflows'. Available at: <https://alhena.ai/blog/luxury-ai-concierge-workflows/> (Accessed: March 2026).
- CIO (2025) 'Pernod Ricard Leverages AI to Strengthen Marketing Efforts'. Available at: <https://www.cio.com/article/4070313/pernod-ricard-leverages-ai-to-strengthen-marketing-efforts.html> (Accessed: March 2026).
- Ekamoira (2026) 'What Is Agentic Commerce, The Complete 2026 Guide to AI Shopping Agents'. Available at: <https://www.ekamoira.com/blog/what-is-agentic-commerce-the-complete-2026-guide-to-ai-shopping-agents> (Accessed: March 2026).
- Envive AI (2026) 'AI Search Optimization for Luxury Fashion Brands'. Available at: <https://www.envive.ai/post/ai-search-optimization-luxury-fashion-brands> (Accessed: March 2026).
- ILoveFashionRetail (2026) 'AI SEO LLM Search Framework for Fashion Brands'. Available at: <https://www.ilovefashionretail.com/branding/ai-seo-llm-search-framework-fashion-brands/> (Accessed: March 2026).
- llms.txt (2026) 'Machine-Readable Format for LLM Content Parsing'. Available at: <https://llmstxt.org/> (Accessed: March 2026).
- Luxury Society / DMR Group (2021) Luxury product search volumes by brand and category.
- MIM Agency (2026) 'ChatGPT Shopping and the New Logic of E-Commerce'. Available at: <https://www.agencymim.com/informational-resources/market-insights/chatgpt-shopping-and-the-new-logic-of-e-commerce-what-will-change-in-2026/> (Accessed: March 2026).
- Personal Care Insights (2026) 'Sephora ChatGPT Beauty App', 24 March 2026. Available at: <https://www.personalcareinsights.com/news/sephora-chatgpt-beauty-app.html> (Accessed: March 2026).
- Retail Gazette (2026) 'Sephora Launches App in ChatGPT', 24 March 2026. Available at: <https://www.retailgazette.co.uk/blog/2026/03/sephora-launches-app-in-chatgpt-as-it-pilots-ai-powered-beauty-shopping/> (Accessed: March 2026).

## AI BLOGS AND PLATFORMS

- Claude Blog / Anthropic (2025-2026) Enterprise AI agents, Claude Code, sector-specific use cases. Available at: <https://claude.com/blog> (Accessed: March 2026).
- Google DeepMind Blog (2025-2026) Foundational AI research, Gemini models, multimodal. Available at: <https://deepmind.google/blog> (Accessed: March 2026).
- OpenAI Blog (2025-2026) In-depth articles: enterprise use cases, AI research, agents, safety. Available at: <https://openai.com/blog> (Accessed: March 2026).
- Perplexity Blog (2025-2026) AI search, answer engines, e-commerce integrations, evolution of agentic search. Available at: <https://perplexity.ai/hub> (Accessed: March 2026).
- Mistral AI Blog (2025-2026) Open-weight models, European enterprise deployments, AI sovereignty. Available at: <https://mistral.ai/news> (Accessed: March 2026).

- Meta AI Blog (2025-2026) Llama, open-source AI, social commerce. Available at: <https://ai.meta.com/blog> (Accessed: March 2026).
- Think with Google (2025-2026) Consumer Insights, AI Excellence, Future of Marketing, Search & Video. Available at: <https://business.google.com/fr/think/> (Accessed: March 2026).

---

## INDEX OF FIGURES

FIGURE	TITLE	SECTION
Visual 01	The physical shelf vs the AI shelf	P1 — 1.1
Visual 02	Benchmark: 10 Maisons facing AI agents	P1 — 1.2
Visual 03	Customer journey: before vs with agent	P2 — 2.1
Visual 04	Productive friction vs suffered friction	P3 — 3.2
Visual 05	Matrix of 3 strategies	P3 — 3.3
Visual 06	Decision matrix: where to position?	P4 — 4.1
Visual 07	Roadmap: 5 workstreams in 6 months	P4 — 4.2

## INDEX OF KEY QUOTES

#QUOTE	SECTION
1 "This is not a technology problem. It is a shift in the geography of desire."	P1 — 1.1
2 "Where a search engine displayed 48 results, the AI agent cites only 3. The compression is radical — and irreversible."	P1 — 1.1
3 "The agent, unwittingly, does exactly what luxury marketing has spent decades trying to prevent: it makes price comparable."	P1 — 1.2
4 "Welcome to the era of invisible products. Brands fade and blend into browsers or operating systems." — Michael Baeyens	P3 — 3.1
5 "Agents produce informational nudity. They strip the brand narrative of everything that is not a measurable attribute." — Mickaël Tsakiris	P3 — 3.1
6 "The AI agent is, by design, anti-luxury. It responds to demand. It facilitates purchase. It removes the obstacle, and with the obstacle, the very condition of desire."	P3 — 3.2
7 "The client trusts the agent to find. They trust the human to choose."	P3 — 3.3
8 "Luxury has always been about deliberate choices. Agentic AI has not changed the nature of this business. It has added one more choice to the list."	Conclusion



## LUXE ÆTERNAI

*Agentic AI applied to Luxury & Premium.  
Consulting, training, strategic intelligence. Paris.*

APRIL 2026