

THE ONLY PATH TO CONSUMER TRUST IN FASHION'S AI DECADE

ABSTRACT

This paper discusses governance in the context of today's luxury fashion AI economy. Current AI tools are making consequential decisions across fashion and independent general-purpose tools are dominating complex consumer touchpoints that risk identity, price, narrative, exclusivity and credibility misrepresentation. Consequently, brands must govern every single aspect of AI they have access to, from internal deployments to vendor relationships. This body of work introduces a three-zone AI framework distinguishing where fashion brands have direct control, indirect influence and no authority over AI, mapping the governance obligations that follow from each.

www.guarinoadvisors.com



AUTHOR

Technology executive and entrepreneur with a background in defense-grade software, cybersecurity, finance, and complex systems, having led ForceX through its growth and acquisition, and now serving as President of Whering, a global fashion technology platform, and Founder of Guarino Advisors. Her work contributes to industry analysis and conversations on the future of fashion, with a focus on security, risk, and AI integration.

TRACY GUARINO

GOVERN WHAT YOU BUILD: THE ONLY PATH TO CONSUMER TRUST IN FASHION'S AI DECADE

Photo Courtesy of
ENGLISH.AAWSAT

GUARINO
ADVISORS

THE CONSUMER TRUST CRISIS FASHION DID NOT CREATE

To understand what luxury fashion must govern, one must understand what already exists. AI is embedded, consequential and unwavering in its position. More than 1.5 billion people worldwide interact with large language models and generative AI applications monthly, and across AI searches, assistants and commerce platforms, some data indicates the total AI user pool may be approaching ten-figures.¹


Like all global sectors, luxury fashion is contending with the effects of these technological advancements.

Emerging AI models presently direct how high-end goods are developed, discovered and sold, which has resulted in the rapid reformulation of strategies across the sector.² Consumer sentiment, however, is not keeping up with adoption rates. A 2026 Vogue Business Consumer Survey demonstrating the prevalence of AI consumer use in luxury fashion found 69% of respondents use general-purpose AI tools, such as ChatGPT, Perplexity, Claude and Gemini, at least occasionally.³ Simultaneously, 55% of those same luxury-adjacent respondents distrust AI recommendations and 72% refuse to share card details

with AI tools.⁴ The contradiction is clear: consumers are adopting AI tools at scale while not believing their outputs.

The disparity exists, in large part, because the tools driving it sit outside the fashion industry's governance authority, and brands are broadly failing to respond to the issue. The notable AI governance gap is not exclusive to fashion: McKinsey reports that only 17% of


organizations exceeding \$500 million in annual revenues disclose board-level AI governance oversight.⁵ In fashion, it is rarely even a conversation topic. The State of Fashion 2026, the fashion industry's leading annual report, deeply recognizes AI adoption in communication and marketing functions, while completely lacking structured, comprehensive and applicable AI governance frameworks.⁶



6%

of organizations report complete oversight of how vendors are using AI within their services.

COMPLIANCE WEEK AND GAN INTEGRITY, 2025



17%

of organization participants disclosed a form of board-level AI governance oversight, with no differentiation drawn between internal AI, vendor AI and external AI.

McKINSEY QUANTUMBLACK, 2025

GUARINO ADVISORS

The risk of failing to understand and govern multi-faceted AI dimensions is practical. Fashion brands possess no legal right to compel general-purpose AI providers to appropriately represent their identity, pricing or positioning, and those providers have no contractual obligation to do so.⁷

Deloitte warns prestige-driven brands, as they are exposed to face reputational consequences when AI capabilities fall short in meeting their consumer expectations.⁸ Beyond reputation, brands described by these models have no visibility into the training data shaping brand-associated AI outputs and no audit rights over the information surfaces.⁹ The downstream effects include perceived exclusivity erosion through the flattening of luxury positioning into commodity-style comparison, stripping of cultural and aspirational signals built over decades,¹⁰ price and positioning distortion through outdated, aggregated or incorrect

pricing,^{11,12,13} and brand vulnerability from AI hallucinations, where factually wrong statements are delivered with the illusion of a trusted source.^{14,15}

Therefore, luxury houses depend on external AI models to communicate their brand identity accurately, surface correct price points and maintain the exclusivity cues their luxury equity depends on; a dependency accepted without the governance infrastructure to enforce any of it.

The questions that follow are precise: what can fashion actually govern, where does authority begin and end, and how does discipline within those limits earn consumer trust? The answer requires a clear distinction between the three AI dimensions in which fashion brands operate. Each zone carries a different governance reality, a different commercial risk and a different path to trust.

⁴O'Brien and Sheah, "You Cannot Trust a Machine."

⁵Alex Singa et al., "The State of AI: How Organizations Are Rewiring to Capture Value," QuantumBlack, AI by McKinsey, March 12, 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-how-organizations-are-rewiring-to-capture-value>.

⁶Sachidanand et al., "The State of Fashion 2026."

⁷European Parliament and Council of the European Union, Regulation (EU) 2024/1689 (Artificial Intelligence Act), art. 63, Official Journal of the European Union, July 13, 2024, <https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng> (enumerating GAI provider obligations, including documentation, transparency and copyright compliance, none of which require accurate representation of third-party brands).

⁸Adam Hewson, Saurabh Vijayvergia, and Harris Abrams, "Shaping Luxury's Future with AI: AI for Prestige, Personalization and Precision," Deloitte AI Institute, Deloitte Consulting LLP, January 2025, <https://www.deloitte.com/content/dam/assets-zone6/us/en/docs/services/consulting/2026-future-of-luxury-during-rise-of-ai.pdf>.

⁹Rishi Somnazzani et al., "The 2025 Foundation Model Transparency Index," arXiv:2512.10169, Stanford CRIM, December 11, 2025, <https://arxiv.org/abs/2512.10169>.

¹⁰Coral Censis, "Redefining Consumer Experience through Artificial Intelligence in the Luxury Retail Sector," Journal of Retailing and Consumer Services 87 (October 2025): 104416, <https://doi.org/10.1016/j.jretconser.2025.104416>.

¹¹Madeleine Schulz, "Can AI Ever Crack Taste?" Vogue Business, April 22, 2025, <https://www.vogue.com/article/can-ai-ever-crack-taste>.

¹²Jakob Mikander et al., "Auditing Large Language Models: A Three-layered Approach," AI and Ethics 4, no. 4 (2024): 1028–1016, <https://doi.org/10.1007/s43681-023-00289-2>.

¹³Régis Y. Chenavaz and Starke Dimitrov, "Artificial Intelligence and Dynamic Pricing: A Systematic Literature Review," Journal of Applied Economics 28, no. 1 (2025): 2466140, <https://doi.org/10.1080/15140326.2025.2466140>.

¹⁴National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile," NIST AI 600-1, July 2024, <https://doi.org/10.6028/NIST.AI.600-1>.

¹⁵Anqi Shao, "New Sources of Inaccuracy? A Conceptual Framework for Studying AI Hallucinations," Harvard Kennedy School Misinformation Review, August 27, 2025, <https://doi.org/10.37016/mr-2020-162>.

THE THREE ZONES OF AI THAT FASHION BRANDS OPERATE IN

Observations indicate that AI-related consumer trust is being treated as a single undifferentiated problem. Governance initiatives are being designed without distinction between AI dimensions, and the gap this creates is measurable. McKinsey's 2026 AI Trust Maturity Survey finds that only about one-third of organizations report maturity levels of three or higher in strategy, governance and agentic AI governance, with organizational alignment and oversight structures struggling to keep pace with the rapid expansion of AI use across all three categories without differentiation.¹⁶

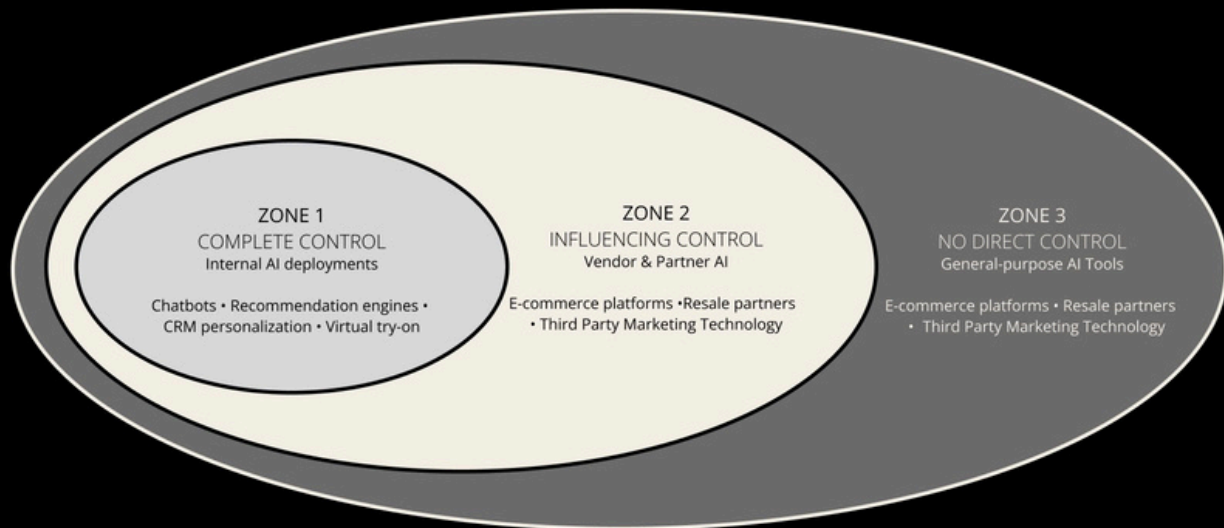
The premium fashion sector falls within this observation. The industry faces a fragmented and fast-evolving AI regulatory environment with shared responsibilities spanning extensive stakeholder networks,¹⁷ yet the prevailing conversations only cover individual company-level challenges rather than addressing the multi-layered governance reality. The discussed Vogue Business of Fashion Survey reflects this: its AI-consumer trust methodology implicitly distinguishes between AI interaction, but never explicitly articulates those differences or their distinct governance obligations.¹⁸

The consequence of that conflation is a polarized solution set; either too broad, calling for sweeping governance of all AI in fashion without a realistic mechanism for doing so, or too narrow, reducing the entire problem to a single consumer-facing chatbot. Neither addresses the structural reality.

The trust deficit in AI interactions cannot be solved by the industry acting alone, and the data confirms it. Research from the University of Melbourne and KPMG finds that 70% of people globally believe national and international regulation is needed to address AI trust at scale.¹⁹ Given that the broader trust deficit cannot be resolved unilaterally, in luxury fashion, where narrative control is currency, governing what is within reach is the most strategic move available.

There are three distinct AI interactions that will determine how consumers build, or fail to build, trust with luxury fashion brands. The three-zone framework that follows is a practical tool: it tells a brand's board precisely where it holds authority, where it carries influence and where it must earn trust through demonstrated discipline rather than blind dependency.

THE THREE ZONES OF AI AUTHORITY FOR LUXURY FASHION



GUARINO
ADVISORS

ZONE ONE — COMPLETE CONTROL

Zone 1 is the zone fashion brands already occupy. It encompasses every AI tool a brand directly deploys: website recommendation engines, on-site chatbots, virtual try-on features, CRM personalization, generative campaign tools, merchandising and demand forecasting, dynamic pricing and in-store clienteling. They are live integrations and consumer-facing systems making consequential decisions at scale. Yet they are being built and deployed without the governance discipline that mature industries would consider non-negotiable for comparable decision tools.

The evidence is not difficult to find. A 2026 Vogue Business feature on the future of retail AI profiles Bain's customer digital twin model and Random Studio's invisible-AI design philosophy as the operational direction of luxury, documenting significant AI integration across the sector without a single reference to governance structure.²⁰ The omission reflects the industry's default posture: deploy first, govern when something goes array.

Tapestry, the parent company of Coach and Kate Spade, recently filed a U.S. patent for "Mira," an internal AI system designed to consolidate data across merchandising, supply chain, finance and retail operations, surfacing business insights in minutes.²¹ The system is already being used to plan assortments, manage inventory and track consumer trends.²²

“THE ACCOUNTABILITY GAP REACHES THE CONSUMER WHEN THE OUTPUT IS WHAT THE CONSUMER SEES.”

The second is vendor contract AI governance. A Stanford Law School CodeX and TermScout analysis found that 92% of AI vendor contracts claim data usage rights beyond what is necessary for service delivery, with many permitting vendors to use customer data for model retraining without explicit consumer consent, far exceeding the 63% market average for standard technology contracts.²⁷ Considering the level of access to customer data, for industries whose entire value proposition rests on discretion and consumer trust, the absence of documented governance requirements in these contracts cannot be a peripheral concern. This is distinct from the vendor relationships addressed in Zone 2. These are contracts the brand controls directly and has chosen not to govern.

Tapestry's public patent announcement describes a security-oriented, role-based access architecture as central to Mira's design.²³ What neither the announcement nor any public disclosure addresses is model-level governance: named senior ownership per deployment, documented performance baselines or escalation paths. The architecture may be governed. Whether the consequential outputs of the model are accountable in the same way is a separate question, and one that has not been answered publicly.

The structural obstacles driving this pattern in internal AI governance as it exists falls into two overarching categories.

The first is model ownership. In most luxury fashion enterprises, there is no named senior individual accountable for a given AI deployment and no documented performance baseline against which degradation can be measured. Deployment decisions are often made at the product or marketing level with no board visibility, and the data confirms it. Only 39% of Fortune 100 companies disclosed any form of board-level AI oversight as of 2024,^{24,25} and nearly 80% of organizations report unclear ownership of AI initiatives altogether despite a huge amount of organizations using AI in at least one business function.²⁶ The accountability gap reaches the consumer when the output is what the consumer sees.

The Zone 1 structural failures produce three downstream problems escalating in severity:

- A. Operational failure through model drift: a recommendation engine that performed well at launch degrades in accuracy over time with no one accountable for detecting or correcting it.
- B. Brand failure through misrepresentation: a generative tool produces outputs that contradict the brand's positioning, pricing or creative codes at the exact moment of consumer interaction.
- C. Legal failure through regulatory exposure: consumer-facing AI deployed without documented validation and accountability chains will face enforcement under the EU AI Act from August 2026.²⁸

²⁰Amy O'Brien, "Inside Retail's AI-Enhanced Future," Vogue Business, April 22, 2026, <https://www.vogue.com/article/inside-retails-ai-enhanced-future>.

²¹Beth Duckett, "Tapestry Builds AI-Powered Mira to Pull Data, Make Decisions," Digital Commerce 360, May 13, 2026, <https://www.digitalcommerce360.com/2026/05/13/tapestry-mira-ai-revenue-q3-fy26/>.

²²Ibid.

²³Tapestry, Inc., "Tapestry Awarded U.S. Patent for Innovative AI Platform: Mira," Business Wire, May 11, 2026, <https://www.businesswire.com/news/home/2026051994098/en/>.

²⁴Aamer Baig et al., "The AI Reckoning: How Boards Can Evolve," McKinsey Quarterly, December 4, 2025, <https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/the-ai-reckoning-how-boards-can-evolve>.

²⁵Alex Singa et al., "The State of AI in 2025: Agents, Innovation, and Transformation," QuantumBlack, AI by McKinsey, November 5, 2025, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>.

²⁶Atomtek and HFS Research, "Humans at the Helm of AI," Atomtek, April 8, 2026, <https://www.atomtek.com/humans-at-the-helm-of-ai/>.

²⁷Oiga Mack, "Navigating AI Vendor Contracts and the Future of Law: A Guide for Legal Tech Innovators," Codex (blog), Stanford Law School, March 21, 2026, <https://law.stanford.edu/2025/03/21/navigating-ai-vendor-contracts-and-the-future-of-law-a-guide-for-legal-tech-innovators/>.

²⁸EU AI Act, art. 26.

ZONE TWO — INFLUENCING CONTROL

Zone 2 covers AI deployed by the direct partners and vendors luxury brands operate alongside: e-commerce platforms, payment processors, third-party styling and discovery tools, resale partners and marketing technology vendors. The resale layer in particular carries compounding exposure.

It is the zone fashion brands are most likely to have no governance position on whatsoever, and the evidence suggests this is a structural blind spot the industry has yet to name. As examined in *Why Resale Platforms Need Security Architecture and Securing the Modern Luxury Value Chain*, the verification gaps already present in luxury resale infrastructure mean that AI deployed within that layer inherits those gaps, while also introducing new governance obligations that neither the brand nor the platform has formally assumed.

AI operating inside the vendor layer has received almost no governance attention across any sector. Only 6% of organizations report complete oversight of how vendors are using AI within their services, and only 8% consider their internal AI governance model mature.²⁹

ZONE THREE — NO DIRECT CONTROL.

Zone 3 is where fashion brands have no authority. These are the general-purpose AI tools consumers use independently and most frequently. According to Similarweb's confirmed ranking by monthly website visits, ChatGPT receives about 5.85 billion visits per month, Google Gemini 723 million, DeepSeek and Perplexity each 319 million, and Claude 148 million.³⁵ Their scale is not marginal, and these are no longer niche tools. They are the primary AI interfaces through which consumers are forming opinions, making recommendations and increasingly making purchases.

THE GOVERNANCE PRECEDENT: FINANCIAL, DEFENSE AND INDUSTRY STANDARDS

Financial services and defense are governing their internal AI and direct consumer-facing AI with a discipline fashion has not approached. Like luxury fashion, neither sector can control Zone 3. Unlike luxury fashion, neither sector waited for a crisis to build the governance infrastructure that Zone 1 and Zone 2 require.

In fashion, where vendor relationships span e-commerce platforms, payment processors, resale partners and marketing technology stacks, the exposure is likely higher. Vogue Business and IBM wrote jointly on the luxury fashion context specifically, naming the vendor dependency dynamic directly, documenting a multi-layered stack in which each integration point carries its own AI exposure.³⁰ Aon's 2026 risk analysis confirms the trajectory: AI is amplifying dependencies on third-party systems,^{31,32} with harmful AI-related incidents increasing 56.4% year-on-year (2023 to 2024) according to the 2025 Stanford AI Index Report.³³ That is the structural exposure; maximum brand implication, minimum governance authority and almost no industry conversation about the gap between them.

The regulatory consequence is already codified. The EU AI Act, Article 26, places legal accountability on both providers and deployers of AI.³⁴ Thus, luxury fashion brands integrating third-party AI without contractual governance requirements are not insulated from that liability. The law does not distinguish between AI a brand built and AI a brand adopted. It distinguishes between AI that was governed and AI that was not.

Zone 3 is expanding; each of the named platforms is developing or already deployed agentic shopping tools capable of autonomously researching, recommending and completing purchases on a consumer's behalf. Vogue Business identifies this as the next frontier of commerce.³⁶ The implication is direct: as the agents continue to absorb the consumer journey, the brand loses what little indirect influence it previously held over the AI-mediated experience.³⁷

²⁹Aaron Nicodemus, "Survey Report: Organizations Broadly Adopting AI, with Varied Governance," Compliance Week and GAN Integrity, February 26, 2025, <https://www.complianceweek.com/technology/survey-report-organizations-broadly-adopting-ai-with-varied-governance/>.
³⁰Vogue Business and IBM, "Evolve or Fade Away: Enriching Luxury Heritage in the AI Era," Vogue Business, January 2025, <https://www.vogue.com/article/evolve-or-fade-away-enriching-luxury-heritage-in-the-ai-era>.
³¹Aon, "AI Risk 2026: What Business Leaders Need to Know," Aon, May 7, 2026, <https://www.aon.com/en/insights/articles/ai-risk-2026-practical-agenda>.
³²PwC Belgium, "Responsible AI and Third-Party Risk Management: What You Need to Know," PwC, January 5, 2026, <https://www.pwc.be/en/news-publications/2025/responsible-ai-tp-risk.html>.
³³Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njings Kariuki, Emily Capstick, et al., "The AI Index 2025 Annual Report," Stanford, CA: Institute for Human-Centered AI, Stanford University, April 2025, <https://hai.stanford.edu/ai-index/2025-ai-index-report>.
³⁴EU AI Act, art. 26.

³⁵Similarweb, "Top AI Tools: Most Used Gen-AI in August 2025," <https://www.similarweb.com/blog/marketing/seo/most-used-ai>.
³⁶Amy O'Brien, "Where Will AI Take Fashion in 10 Years?" Vogue Business, April 22, 2026, <https://www.vogue.com/article/where-will-ai-take-fashion-in-10-years>.
³⁷Ibid.

FINANCIAL SERVICES

The financial services precedent is instructive. SR 11-7, issued jointly by the Federal Reserve and the Office of the Comptroller of the Currency in 2011, established model validation, independent review and ongoing performance monitoring as board-level supervisory expectations for every decision-making model a regulated bank deploys.³⁸

That framework was not written specifically for AI; however, it was written for any quantitative model making consequential decisions, which is precisely what AI does. When AI arrived in financial services, it inherited fifteen years of governance discipline rather than starting from scratch. Fashion has no equivalent foundation to inherit, let alone update.

“FASHION HAS NO EQUIVALENT FOUNDATION TO INHERIT, LET ALONE UPDATE.”

That foundation has since been updated. In April 2026, the Federal Reserve, OCC and FDIC jointly issued SR 26-2, confirming that model risk management, encompassing named ownership, independent validation, ongoing performance monitoring and board-level governance, remains an active and enforceable supervisory expectation.³⁹ SR 26-2 explicitly acknowledges that generative and agentic AI models are not yet within its full scope, but confirms the underlying governance principles apply to all quantitative decision-making tools.⁴⁰ Financial services governance is so well-established that its most current regulatory update is already working through how to extend existing discipline to AI.

The same discipline extends across European financial institutions. The European Banking Authority’s supervisory guidance on machine learning and the broader EU prudential framework require that AI models used in credit, risk and consumer-facing decisions meet validation, explainability and audit standards equivalent to those applied to traditional statistical models.⁴¹ The ECB Guide to Internal Models, updated in July 2025,⁴² confirms these requirements have been progressively tightened. European financial institutions are operating under AI governance standards that were established years ago and have continued to evolve.

FASHION BRANDS DEPLOYING EQUIVALENT CONSUMER-FACING DECISION TOOLS IN THE SAME REGULATORY ENVIRONMENT HAVE NO COMPARABLE SUPERVISORY EXPECTATION PLACED ON THEM.

DEFENSE DISCIPLINES

The discipline financial services built through regulation, defense built through unequivocal necessity. Defense applies decision validation and accountability to AI-powered and manual systems alike.

Drawing from direct experience in the field, when ForceX was developing and deploying ISR and target tracking software for elite U.S. government units, the technology did not involve AI. The accountability requirement was identical regardless.

Software going into the field had to work, not most of the time, but every time. The environments it operated in did not offer the opportunity to patch a failure after the fact. Deployment followed a rigorous validation process, including structured QA, staged testing and documented sign-off before anything reached operational use. When modifications were required in the field, which they were, ForceX was available immediately and the change process was governed, not improvised.

Every consequential output had a chain of accountability behind it. This was beyond providing cultural value, it was an operational requirement.

Parallel to financial fields, the defense sector has since extended foundational disciplines directly into AI. The Pentagon’s early AI adoption revealed that some initiatives lacked clear validation standards and governance structures, and the result was predictable. AI outputs appeared precise but confidence in their reliability remained uneven. The US Government Accountability Office documented in 2022 that early DoD AI initiatives lacked consistent validation standards, governance structures and department-wide acquisition guidance.⁴³ The lesson the Department of Defense drew was that reliable AI requires unglamorous work first: data standards, governance, training and trust built at the unit level. The biggest mistake was assuming those steps could be skipped.

³⁸ Board of Governors of the Federal Reserve System and Office of the Comptroller of the Currency, “Supervisory Guidance on Model Risk Management,” SR Letter 11-7, April 4, 2011, <https://www.federalreserve.gov/frra/guidance/supervisory-guidance-on-model-risk-management.htm>.

³⁹ Board of Governors of the Federal Reserve System, Office of the Comptroller of the Currency, and Federal Deposit Insurance Corporation, “Revised Guidance on Model Risk Management,” SR Letter 26-2, April 17, 2026, <https://www.federalreserve.gov/supervisionreg/rlletters/SR2602.pdf>.

⁴⁰ *Ibid.*

⁴¹ European Central Bank, “ECB Guide to Internal Models,” version 4.0, ECB Banking Supervision, July 28, 2025, https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm/supervisory_guido202507/en.pdf.

⁴² *Ibid.*

⁴³ U.S. Government Accountability Office, “Artificial Intelligence: DOD Should Improve Strategies, Inventory Process, and Collaboration Guidance,” GAO-22-105634, March 30, 2022, <https://www.gao.gov/products/gao-22-105634>.

That lesson has since been formalized. Congress directed the Department of Defense to establish a standardized, department-wide framework for assessing, governing and approving the development, testing and deployment of AI models.⁴⁴ The framework requires performance, security, documentation, ethical and testing standards for all major AI deployments.⁴⁵ The DoD's Responsible AI Strategy and Implementation Pathway, signed by the Deputy Secretary of Defense in June 2022 and administered by the Chief Digital and Artificial Intelligence Office, formalized more than sixty lines of effort for operationalizing AI governance across the department, explicitly requiring that AI capabilities be accountable throughout a responsible human chain of command and that ethics be integrated from the start of design, not applied after deployment.⁴⁶

The principle connecting ForceX's pre-AI governance discipline to the defense sector's current AI governance framework is consistent: when the decisions a technology enables are consequential, the accountability requirement does not change based on whether the technology is new. It scales with the stakes.

Fashion brands deploying AI that makes consequential decisions about what consumers see, what they are recommended, how they are priced and how the brand is represented, are operating in an environment with comparable accountability requirements and almost none of the governance infrastructure that defense-grade environments treat as non-negotiable.

THE TECHNOLOGY IS DIFFERENT BUT THE OBLIGATION IS NOT.

EXISTING FRAMEWORKS

Beyond financial and defense disciplines, two frameworks provide the operational architecture that any industry serious about AI governance would recognize. The NIST AI Risk Management Framework, published in January 2023 and updated through 2025, organizes accountability across four functions: Govern, Map, Measure and Manage.⁴⁷ It requires organizations to define model ownership, document performance baselines, conduct adversarial testing and maintain human oversight for consumer-facing deployments.⁴⁸

ISO/IEC 42001, the first certifiable international AI management standard, provides a structured compliance pathway equivalent to what ISO 27001 established for information security,⁴⁹ requiring documented AI policies, defined accountability structures and risk assessments for every deployment as conditions of certification.⁵⁰ Both frameworks exist, both are available in fashion, and neither is being applied. This is a critical gap.

LUXURY FASHION APPLICATIONS

The lesson from both sectors is consistent. In both industries, governance is applied to internally built systems, integrated AI and every decision chain the enterprise is accountable for. Although the tools fashion deploys are different, the accountability logic abides: discipline must be built before failure, not in response to it.

Luxury fashion brands operating globally will inherit that convergence in every market they operate in, regardless of which regulatory regime arrives first or which jurisdiction moves fastest.

The convergence pressure makes this more urgent. While the US, EU and China are developing divergent AI regulatory regimes for commercial applications, the underlying governance principles those regimes share, namely risk stratification, accountability and validation, have been applied by financial services and defense regulators across all three jurisdictions for decades.⁵¹ As McKinsey's Holger Harreis notes in *Vogue Business*, the regulatory directions are diverging at the commercial application layer while converging at the governance principle layer.⁵²

The brands that earn durable trust do so by governing what they control with precision and maintaining serious discipline within that boundary. Managing the entire technological environment a consumer touches is not feasible for any brand. What is feasible, and what the evidence from financial services and defense confirms, is that disciplined governance of the zones within a brand's authority is sufficient to build credibility. The question for luxury fashion is whether the industry will choose to meet it before the regulatory and reputational consequences of not doing so force the issue.

⁴⁴National Defense Authorization Act for Fiscal Year 2024, Pub. L. No. 118-31, § 1621 (December 22, 2023), <https://www.congress.gov/118/plaws/pub21/PLAW-118pub21.pdf>.

⁴⁵*Ibid.*

⁴⁶U.S. Department of Defense, "Responsible Artificial Intelligence Strategy and Implementation Pathway," CDAO, June 22, 2022.

⁴⁷<https://media.defense.gov/2023/Jan/22/20230226/0-111-10/Department-of-Defense-Responsible-Artificial-Intelligence-Strategy-and-Implementation-Pathway.PDF>.

⁴⁸National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," NIST AI 100-1, January 2023, <https://doi.org/10.6028/NIST.AI.100-1>.

⁴⁹*Ibid.*

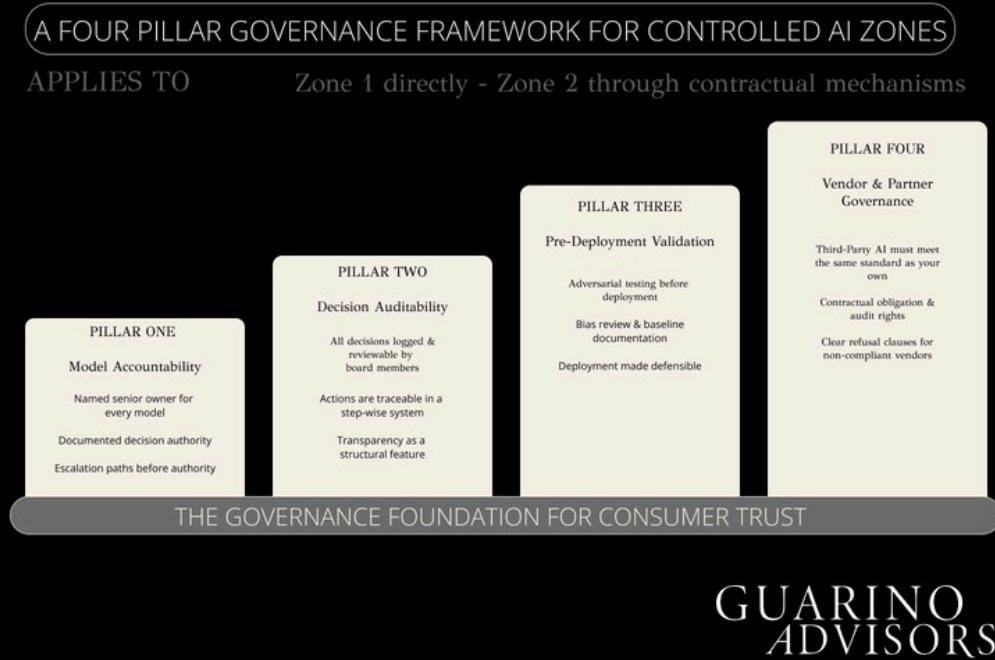
⁵⁰International Organization for Standardization, "Information Security, Cybersecurity and Privacy Protection — Information Security Management Systems — Requirements," ISO/IEC 27001:2022, October 25, 2022, <https://www.iso.org/standard/27001>.

⁴⁹International Organization for Standardization and International Electrotechnical Commission, "Information Technology — Artificial Intelligence — Management Systems," ISO/IEC 42001:2023, December 18, 2023, <https://www.iso.org/standard/8230.html>.

⁵¹Leo S. Lo, "Artificial Intelligence Regulation Matures: Landscapes of the USA, European Union, and China," *IFLA Journal* (2023), <https://doi.org/10.1177/03400325231384915>.

⁵²Maha Shoaib, "How Will the AI Race Impact the Global Luxury Market?" *Vogue Business*, April 2023, <https://www.vogue.com/articles/how-will-the-ai-race-impact-the-global-luxury-market>.

A FOUR-PILLAR GOVERNANCE FRAMEWORK FOR CONTROLLED AI ZONES



PILLAR ONE — MODEL ACCOUNTABILITY.

This framework’s model accountability pillar proposes every AI model in luxury fashion commercial deployment has a publicly named senior owner, with accompanying documentation for decision authority. Consequently, it means that escalation paths exist before failure forces them to be improvised.

The standard is straightforward. However, the fashion industry’s most sophisticated operators have not met it. LVMH has made a public commitment to responsible AI development, appointing governance officers at the group level, citing consumer and employee trust as the direct rationale.⁵³ However, LVMH’s internal generative AI agent,

which handles over two million requests per month across 40,000 employees,^{54,55} is a consequential decision-making tool operating at scale that lacks model accountability. No public documentation exists of named senior individuals accountable for each model’s outputs, documented performance baselines or escalation paths per deployment. Group-level governance intent and named individual accountability per deployed model are not equivalents, and thus, their commitment lacks model accountability in the way mature, regulated, high-risk industries define it. The governance intent is present at the organizational level, but governance infrastructure is absent at the model level.

PILLAR TWO — DECISION AUDITABILITY.

This framework’s decision auditability pillar refers to the requirement that every AI decision affecting customers, merchandise or creative output is logged, reviewable and traceable to a documented process. It operationalizes transparency, not as a value statement, but as a structural feature of how AI is deployed.

The absence of it in fashion has already produced public failures. In July 2024, Mango launched an advertising campaign featuring AI-generated models.⁵⁶ Social media users immediately characterized the outputs as false advertising, citing inaccurate depictions of clothing fit that contradicted the brand’s own product claims.⁵⁷

⁵³Retail TouchPoints, “AI in Luxury: Why LVMH Won’t Let AI Eclipse Humanity or Creativity,” Retail TouchPoints, January 27, 2026, <https://www.retailtouchpoints.com/features/ai-in-luxury-why-lvmh-wont-let-ai-eclipse-humanity-or-creativity/156372/>.

⁵⁴PMNNTS, “LVMH Deploys AI Tools Across Operation, Seeking Efficiency and Customer Retention,” PMNNTS.com, June 9, 2025, <https://www.pymnts.com/artificial-intelligence-2/2025/lvmh-deploys-ai-tools-across-operation-seeking-efficiency-and-customer-retention/>.

⁵⁵Stanford Institute for Human-Centered AI, “Stanford HAI Welcomes LVMH to the Corporate Affiliates Program,” Stanford HAI, October 12, 2023, <https://hai.stanford.edu/news/stanford-hai-welcomes-lvmh-corporate-affiliates-program>.

⁵⁶LBOnline, “Holding Up the Mirror to AI Beauty: Attractive Investment or Ethical Eyesore?” LBOnline, February 2026, <https://lbonline.com/news/ai-beauty-fashion-models-attractive-investment-ethical-eyesore>.

⁵⁷Ibid.

No publicly documented logging, review mechanism or accountability chain existed for what the AI produced before it reached consumers. The reputational damage was not the result of using AI, but rather was the result of deploying it without the auditability infrastructure that would have mitigated brand identity risk before publication. Mango is not an outlier, but it is a visible case of an industry-wide absence.

PILLAR THREE — PRE-DEPLOYMENT VALIDATION.

This framework's pre-deployment validation pillar requires that no consumer-facing model becomes live without adversarial testing, bias review and documented performance baselines. This functions to demonstrate compliance, responding to regulatory inquiries with evidence and correcting anticipated failures. Pre-deployment validation is the infrastructure that makes deployment defensible once it is live.

Fashion has not adopted this condition and the legal exposure as a consequence is a threat.

PILLAR FOUR — VENDOR AND PARTNER GOVERNANCE.

This framework's vendor and partner governance insists that that AI deployed by third parties on a brand's behalf meets the same standard as AI the brand deploys itself. That standard is enforced through contractual obligation, audit rights and a clear refusal to integrate vendors who cannot demonstrate governance discipline. If the output reaches the consumer under the brand's name, the brand will be accountable for it, and must mitigate the underlying risks.

The luxury fashion industry has not operationalized that principle and the closest available precedent makes the gap visible.

Richemont, one of the luxury sector's most operationally sophisticated conglomerates, has implemented what it describes as a centralized yet flexible governance model for automation, encompassing robotic process automation and conversational AI across its decentralized brand portfolio, aimed at achieving faster project delivery, enhanced security and streamlined maintenance.⁶¹ It also covers internal automation exclusively.⁶² It is of the most structured governance postures

Reed Smith's analysis of AI legal trends in luxury fashion finds that brands are deploying virtual try-on tools, personalization engines and AI chatbots across consumer touchpoints.⁵⁸ Each of these tools collects biometric data.⁵⁹ Reed Smith notes explicitly that such data requires consent frameworks, encryption standards and clear data governance, none of which can be demonstrated without pre-deployment documentation.⁶⁰ The legal risk identified is the direct consequence of deploying consumer-facing AI without the validation infrastructure that would produce that documentation in the first place.

currently documented in the luxury goods sector. Third-party AI vendor contractual requirements, audit rights and refusal clauses are not part of the framework and can become a liability. The gap this pillar closes is precisely the one Richemont's model leaves open.

Across all four pillars, the pattern is consistent: fashion is deploying AI at scale and the governance discipline that must accompany it is absent or partial. As Yilu Zhou at Fordham notes, AI built on biased data produces biased outputs.⁶³ The discipline that addresses this in mature industries is pre-deployment validation, adversarial testing and documented performance baselines, applied before a model reaches consumers, not after it fails them. No brand has fully implemented all four pillars, which is precisely the argument this framework makes and the opportunity it identifies. The framework applies to Zone 1 directly, and extends through contractual mechanisms to Zone 2. Disciplined governance of both zones is achievable, and it is the foundation on which consumer trust in fashion AI gets built.

⁵⁸Talia Botargh Behar and Tyler Thompson, "Always in Season: Luxury, Fashion, and the Law — AI legal trends reshaping luxury fashion: What brands need to know now" Reed Smith, November 12, 2025, <https://viewpoints.reedsmith.com/post/102u8p/always-in-season-luxury-fashion-and-the-law-ai-legal-trends-reshaping-luxury>.

⁵⁹Ibid.

⁶⁰Ibid.

⁶¹Roboyo, "Luxury Goods: Delivering Next Level Group Transformation," Roboyo, accessed May 18, 2026, <https://roboyo.global/case-study/luxury-goods-delivering-next-level-group-transformation/>.

⁶²Ibid.

⁶³Yilu Zhou, quoted in Madeleine Schulz, "Can AI Ever Crack Taste?" Vogue Business, April 22, 2026, <https://www.vogue.com/article/can-ai-ever-crack-taste>.

HOW DISCIPLINE IN Z1 & Z2 EARNS CREDIBILITY IN Z3

Governance discipline does not stop at the boundary of what a brand controls. As discussed, general-purpose AI tools don't describe brands neutrally. They form and surface brand representations with the use of training data, indexed online content and crawlable signals. With the available information, the AI ultimately produces outputs that consumers encounter as recommendations, comparisons or factual statements. The reality is that the only authority brands have in this space is what content these tools have to index.

This is where Generative Engine Optimization plays a key role. GEO, established through research out of Princeton, describes the mechanism by which structured, authoritative, crawlable content shapes how generative AI tools represent a brand in their outputs, similar to how digital authority has always been built.⁶⁴

Zone 1 and Zone 2 governance behavior when well documented and published with structural integrity, generates the indexable, citable and orderly content that generative engines draw from. Published AI governance, including audit frameworks, board oversight statements and accountability documentation in this context go beyond internal compliance, but inputs that aid in shaping their brand presence when Zone 3 AI answers luxury fashion queries. Discipline in zones the luxury fashion house can control produces credibility that travels into a zone it cannot control.

The consumer journey makes the logic complete. A luxury consumer researching a brand through a general purpose AI tool and subsequently arriving at the brand's digital experience is, at every step, making trust assessments. If the information surfaced by, say, ChatGPT, is consistent with what the brands controlled channels communicate, the consistency registers as an absence of friction, contradiction and confusion, at least subconsciously. All consumers value this, but among the premium market, precision, consistency and trustworthiness is held to a higher standard. When said consumer reaches the brand's owned AI systems, regardless of its visibility, the credibility infrastructure built through governance discipline is already making the luxury consumer experience more efficient.

Luxury fashion houses must have internal AI that carries a credibility position that a brand without it cannot replicate through design alone. AI governance practices published at the same standards financial services apply to capital adequacy disclosures earn a trust premium that Zone 3 tools cannot, because they give the consumer a factual basis for extending trust rather than asking them to extend it on faith. In this sense, governance discipline is, in part, a marketing asset; an investment into these practices will compound over time as indexed evidence of accountability accumulates.

CAN LUXURY-ADJACENT CONSUMERS TRUST A MACHINE?

The Vogue Business survey concludes that consumers do not trust AI.⁶⁵ That conclusion is accurate, but the diagnosis behind it is incomplete. Consumers are not responding to the technology itself, but the response rather comes from the absence of any visible evidence that the technology has been governed. They have no basis for trust because no fashion brand has demonstrated the discipline that trust requires. The technology itself is not the issue. Ungoverned machines are.

The principle that guided work in defense environments applies directly here: trust but verify. In practice, that means controlling the inputs before a model is trained, testing the outputs rigorously before deployment and maintaining documented accountability for what the model does once it is live. Applied consistently, that process inevitably reduces risk, while also producing the conditions under which trust becomes rational rather than blind.

A consumer interacting with a fashion brand's AI-powered recommendation engine, styling tool or personalization layer has no way of knowing whether that tool was validated before it went live, whether anyone is accountable for its outputs or whether the brand has any governance posture around it at all. In the absence of that information, distrust is the only reasonable default.

The answer is not simply to make AI more invisible or more human-sounding. It is to govern it visibly enough that consumers have a factual basis for extending trust. Governance infrastructure is the basis for credibility that makes consumer trust in AI rational. You can trust the machine. The precondition is that someone accountable did the work before the machine was deployed. This is how luxury fashion competes with the AI tools it does not control: by being credibly disciplined in the AI it does.

“THIS IS HOW LUXURY FASHION COMPETES WITH THE AI TOOLS IT DOES NOT CONTROL: BY BEING CREDIBLY DISCIPLINED IN THE AI IT DOES.”

A RESPONSE TO FORECASTS

Ten-year forecasts describe AI becoming ubiquitous but invisible, with human interaction remaining the apex of luxury.⁶⁶ The vision is appealing, however, the assumption underneath is insufficient.

Invisible AI operating without visible governance does not inherently produce a luxury experience, because it is an unverifiable one. Consumers in 2026 are too sophisticated to accept unverifiable allegiance.

The forecast presumes a level of consumer trust that the fashion industry has not yet built through demonstrated discipline. Trust is built through transparency, and cannot be boiled down to elegant design or seamless user experience. Consumers want to know whether the AI making recommendations about what they buy, how they are styled and how their data is used has been governed, tested and held accountable. Telling them to simply trust the experience cannot be the credibility strategy because it assumes that consumers will not ask questions, when they already are.

In defense environments, the invisibility of a system's operation to the end user was never confused with the absence of accountability for it. The warfighter did not need to see the validation process. But the validation process existed, was documented and was owned by named individuals who were accountable for what the system produced. The discipline was invisible in deployment and fully transparent in governance.

Luxury fashion needs to borrow these proven models if it wants to excel in the existing AI ecosystem.

The forecast's vision of invisible AI as the future of luxury is achievable only if it considers the path to it runs through governance transparency. The luxury fashion houses that govern their AI visibly enough to demonstrate discipline earns the right to deploy it invisibly. The ones that skip steps are asking premium consumers to trust experiences they cannot evaluate and do not earn consumer faith. Transparency is the precondition for all luxury experiences.

GUARINO ADVISORS



Photo Courtesy of
BAZAAR INDIA

The Vogue Business conclusion that execution is holding back AI in fashion is correct.⁶⁷ Nevertheless, execution without governance is a setup for crises the industry can avoid.

History has demonstrated this pattern with precision. The Sarbanes-Oxley Act did not emerge from a desire to innovate corporate accountability.⁶⁸ It emerged from Enron, WorldCom and a wave of financial failures that destroyed consumer and investor trust at scale.⁶⁹ The accounting standards and governance frameworks that followed were the mandatory consequence of an industry that moved fast, prioritized execution and deferred the discipline that should have been foundational.

The SEC created rigorous regulator requirements because trust had been broken publicly and at scale, and the only path to rebuilding it was through demonstrated, verifiable governance that left no room for ambiguity about who was accountable and for what.

Fashion's AI moment is earlier in that same arc. The failures have begun: breaches, hallucinations, brand misrepresentations and consumer distrust, documented in the Vogue Business survey,⁷⁰

are the opening signals of a reckoning that will accelerate as AI deployment deepens and consumer awareness sharpens. The regulatory response, through the EU AI Act,⁷¹ Digital Product Passport requirements and emerging accountability frameworks, is already in motion.

The brands that govern now are being cautious, but more importantly, strategic. Every governance framework implemented before the mandate arrives is infrastructure built on the brand's own terms rather than under regulatory pressure. Every accountability structure published before a failure forces the conversation is a credibility asset rather than a crisis response.

Governance is the only route to a future in luxury fashion where consumers trust it enough to let it matter.

The lesson from financial services, from defense and from every industry that has navigated a technology driven trust crisis is consistent. Getting ahead of governance is always less costly than catching up to it. The opportunity to lead on this is available now. It will not remain available indefinitely.

⁶⁷O'Brien and Shoah, "You Cannot Trust a Machine."
⁶⁸Sarbanes-Oxley Act of 2002, Pub. L. No. 107-204 (July 30, 2002), <https://www.congress.gov/bills/107th/congress/house-bill/2763>.
⁶⁹Ibid.
⁷⁰O'Brien and Shoah, "You Cannot Trust a Machine."
⁷¹EU AI Act, art. 26.

GUARINO ADVISORS