



AI Hallucinations as an Emerging Non-Tariff Barrier in International Trade: A Conceptual and Case-Based Analysis

Article History

Name of Authors:

Neha Gupta^{1*}, Ritul Singh² and Hitesh Singh³

Affiliation:

¹Director (Management), KCC Institute of Technology and Management, Greater Noida, Uttar Pradesh, India

²Student, Commerce and Management, KCC institute of technology and management, Greater Noida, Uttar Pradesh, India

³Student, Management, Parul University, Waghodia, Vadodara, Gujarat, India

Corresponding Author

Dr. Neha Gupta

How to cite this article:

Gupta Neha et al. AI Hallucinations as an Emerging Non-Tariff Barrier in International Trade: A Conceptual and Case-Based Analysis. J. Int. Commer Law Technol., 2026;7(1):X-X,

Received: 02-06-2026

Revised: 28-07-2026

Accepted: 18-08-2026

Published: 01-09-2026

Doi: 10.61336/jiclt/20267105

©2026 the Author(s). This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)

Abstract: Artificial intelligence systems, and large language models (LLMs) in particular, are increasingly embedded in the infrastructure of cross-border commerce: customs classification, trade documentation, multilingual communication, contract drafting, and regulatory compliance screening. A well-documented limitation of these systems is "hallucination" — the generation of fluent, confident output that is factually incorrect, internally inconsistent, or entirely fabricated. This paper argues that, in the trade context, hallucination is not merely a technical nuisance but is beginning to function as a de facto non-tariff barrier (NTB): it raises compliance costs, introduces legal liability and documentation risk, erodes trust in AI-mediated trade facilitation tools, and creates asymmetric burdens that fall disproportionately on smaller exporters and firms in developing economies with weaker capacity to audit AI output. Using a doctrinal and case-based qualitative method, the paper synthesises secondary evidence — including the *Moffatt v. Air Canada* ruling, customs and trade-compliance industry reports, and recent World Trade Organization and International Monetary Fund publications on AI in trade — to build a conceptual framework describing four mechanisms through which hallucination distorts trade: compliance-cost inflation, information asymmetry, liability displacement, and trust erosion in digital trade infrastructure. The paper concludes with policy recommendations for regulators, firms, and international bodies, and identifies primary-data gaps that future empirical research should address.

Keywords: AI Hallucination, Non-Tariff Barriers, International Trade, Trade Facilitation, WTO, Customs Compliance, Generative AI, Digital Trade

INTRODUCTION

International trade has historically been shaped by two broad categories of restriction: tariffs, which are transparent and quantifiable, and non-tariff barriers (NTBs), a far more heterogeneous set of measures — technical regulations, sanitary and phytosanitary requirements, licensing regimes, customs procedures, and standards — that restrict trade without directly taxing it [1]. The World Trade Organization's Technical Barriers to Trade (TBT) and Sanitary and Phytosanitary (SPS) Agreements were designed precisely because tariff reductions achieved through successive negotiating rounds created incentives for governments and firms to shift friction into less visible, harder-to-discipline channels [2-3]. NTBs are effective as barriers precisely

because they operate through cost, delay, uncertainty and asymmetric information rather than through a published duty rate.

A new candidate has entered this landscape over the past three years: artificial intelligence, and specifically the tendency of generative AI systems to "hallucinate." The World Trade Organization's own World Trade Report 2025 documents how deeply AI has already been woven into trade facilitation — automating customs clearance, parsing tariff schedules and rules of origin, powering multilingual buyer-seller communication, and supporting compliance risk analysis [4-5]. The same report notes that AI could raise global trade by roughly 34–37 per cent by 2040 if adoption gaps are bridged [6]. Yet the very features that make generative AI attractive for trade

facilitation — fluency, speed, and the appearance of authoritative confidence — are also what make its errors dangerous. When an AI system fabricates a tariff classification, invents a regulatory requirement, mistranslates a shipping document, or states a corporate policy that does not exist, the resulting friction functions, in practical terms, exactly like a traditional NTB: it raises the cost of doing business across borders, it is unevenly distributed, and it is difficult for the affected party to anticipate or contest.

This paper develops that argument systematically. It asks: under what conditions, and through what mechanisms, does AI hallucination operate as a functional non-tariff barrier to trade? It approaches the question through a conceptual synthesis of the AI-hallucination literature and the NTB literature, tested against real-world case evidence — most notably the 2024 *Moffatt v. Air Canada* tribunal decision, which established that firms are legally liable for information a chatbot invents, and various industry and institutional reports on AI's growing footprint in customs and compliance operations. The paper does not claim to offer a definitive quantitative measure of trade lost to AI hallucination; that data does not yet exist in the public domain. Instead, it offers a structured conceptual account, grounded in documented cases and institutional reporting, that can guide future empirical and survey-based research — an important caveat given the early stage of this phenomenon.

Literature Review

AI Hallucination Concept and Typology: The term "hallucination" in AI research describes outputs that are fluent and internally coherent but factually incorrect, unsupported by any input or knowledge source, or entirely invented [7]. A recent taxonomic review collecting definitions across academic, legal, and health-sector applications found a consistent thread: the danger of hallucinated content lies less in its inaccuracy *per se* than in its persuasive, self-assured presentation, which makes errors difficult to detect without deliberate fact-checking [8]. Some scholars distinguish "intrinsic" hallucinations, where output contradicts a given source document, from "extrinsic" hallucinations, where output cannot be verified against the source at all [9]; others separate "non-factual" hallucination from "unfaithful" or context-inconsistent generation [10].

Empirical estimates of hallucination frequency vary widely by domain and model. Analysis of user-reported issues in AI mobile applications found hallucination-indicative complaints in approximately 1.75 per cent of relevant reviews [11], while sector-specific studies cited in recent commentary on the *Air Canada* case found that specialised legal-AI tools produced incorrect information in 17–34 per cent of queries, and general-purpose chatbots hallucinated in 58–82 per cent of legal queries,

with customer-service chatbot hallucination rates commonly falling between 3 and 27 per cent even in controlled deployments [12]. A separate strand of research has documented hallucinated citations infiltrating peer-reviewed academic literature itself, including fabricated references that survived expert review at a major AI conference [13] — a finding directly relevant to trade research, since fabricated statistics or invented regulatory citations could just as easily contaminate compliance advice or trade policy analysis generated by AI tools.

Crucially for this paper, several scholars caution that "hallucination" is itself a misleading metaphor: LLMs do not perceive and therefore cannot misperceive; they generate statistically probable text without any internal representation of truth [14]. This matters for trade policy because it reframes the phenomenon away from an occasional "glitch" and toward a structural property of the technology — meaning hallucination-related trade friction should be expected to persist and recur, not disappear as a temporary bug is fixed.

Non-Tariff Barriers

Theoretical Foundations: Non-tariff barriers are conventionally classified by UNCTAD's taxonomy into technical measures (SPS, TBT, and pre-shipment inspection) and non-technical measures (quotas, licensing, customs formalities, and other administrative procedures) [15]. What unites this diverse category is functional, not formal: an NTB is any government or private-sector-imposed condition, other than a tariff, that increases the cost, delay, or uncertainty of moving goods or services across a border [16]. Surveys of exporters in developing and least-developed countries conducted by the International Trade Centre found that SPS measures and TBTs together account for roughly three-quarters of the NTBs exporters describe as burdensome, underscoring that the compliance burden of information-and-standards-based barriers, not tariff-like restrictions, dominates exporter experience [15].

A key insight from the NTB literature is that measures need not be intentionally protectionist to function as barriers. Legitimate public-policy measures — food-safety rules, environmental standards, labelling requirements — can still impose real, unevenly distributed costs, and the line between legitimate regulation and disguised protectionism is frequently contested between trading partners [16–17]. This framing is instructive for AI hallucination: no government or firm intends AI systems to hallucinate, yet the resulting friction can still meet the functional definition of an NTB if it systematically raises costs, creates informational asymmetries, and falls unevenly on parties with less

capacity to detect and correct errors — precisely the concern already raised about TBT/SPS compliance burdens on smaller and developing-economy exporters [15].

MATERIALS AND METHODS

This paper adopts a qualitative, doctrinal, and case-based design rather than a primary-data survey design. Three considerations informed this choice. First, the phenomenon under study is recent — most of the institutional literature on AI in trade facilitation dates from 2025 [4], so large-sample primary data on hallucination-driven trade friction specifically does not yet exist in a form that could be independently verified. Second, the paper's contribution is conceptual: building a defensible framework connecting two established literatures (AI hallucination and NTB theory) that have not previously been synthesised in this way. Third, case-based analysis of a documented legal dispute [18] provides a verifiable, citable instance of hallucination-driven commercial and legal cost that can anchor the conceptual claims in a real, adjudicated fact pattern rather than hypothetical scenarios.

Sources were drawn from three tiers:

- peer-reviewed and preprint literature on AI hallucination (Frontiers, Scientific Reports, arXiv)
- institutional and intergovernmental reports on AI in trade and customs (WTO World Trade Report 2025, IMF Technical Notes and Manuals, World Customs Organization)
- documented legal and industry case material (tribunal rulings, law-firm case commentary, and trade-press reporting). This triangulation across academic, institutional, and case-law sources is intended to substitute, for now, for the primary survey data that a fuller empirical study of this question would require — a limitation the paper returns to in Section 8

Conceptual Framework: Four Mechanisms Linking Hallucination to Trade Friction

Building on the literatures reviewed above, this paper proposes that AI hallucination functions as a non-tariff barrier through four overlapping mechanisms.

Compliance-Cost Inflation

As customs administrations and private logistics firms increasingly deploy AI for Harmonized System (HS) code classification, rules-of-origin determination, and documentation drafting [19-20], hallucinated classifications or invented regulatory requirements generate rework, penalties, and shipment delays. Industry sources note that incorrect HS coding carries direct legal and financial consequences, including

misapplied duties [20], and that inaccurate shipping documentation remains one of the most persistent sources of friction in global trade even before AI is introduced [19]. Where AI tools are trusted without adequate human verification, hallucinated outputs convert what should be a productivity gain into a new source of cost — functionally equivalent to a technical barrier to trade, since the exporter must now spend additional resources verifying AI-assisted documentation against authoritative sources.

Information Asymmetry Between Trading Partners

NTB theory has long recognised that regulatory complexity disadvantages parties with fewer resources to interpret and verify requirements [15]. AI hallucination reproduces and potentially deepens this asymmetry. The WTO-ICC business survey referenced in the World Trade Report 2025 found that only 41 per cent of small firms' report using AI in their trade operations, compared with over 60 per cent of large firms [6]. Smaller exporters that do adopt AI tools for compliance or translation, but lack the in-house legal or technical capacity to independently verify AI output, are structurally more exposed to hallucination-driven errors than large firms with dedicated compliance teams. This mirrors the long-standing finding that TBT and SPS measures disproportionately burden developing-country and small-firm exporters [15]. AI hallucination risks becoming a digital-era analogue of the same dynamic.

Liability Displacement and Legal Uncertainty

The 2024 ruling in *Moffatt v. Air Canada* is instructive here. A Canadian small-claims tribunal held Air Canada liable for its chatbot's fabricated bereavement-fare policy, rejecting the airline's argument that it could not be responsible for its own chatbot's output; the tribunal found the company had a duty of care to ensure the accuracy of information presented through any channel, including an AI system, and awarded damages accordingly [18-19,22]. Trade-press analysis of subsequent disputes — including a 2025 case in which an AI support agent at a software company invented a device-limit policy — suggests courts and regulators are converging on the principle that firms "own" what their AI systems say, regardless of intent [23]. For cross-border trade, this creates a distinctive uncertainty: a hallucinated customs requirement, shipping term, or regulatory claim communicated by an AI system to a foreign buyer or supplier could expose the deploying firm to liability across multiple jurisdictions with different legal standards for negligent misrepresentation, converting a technical error into a cross-border legal-risk premium that functions much like a hidden compliance cost.

Trust Erosion in AI-Mediated Trade Infrastructure

Finally, hallucination risk affects the pace and depth of AI adoption in trade facilitation itself. The WTO notes that AI is already used to break down language barriers in sourcing and dispute resolution and to auto-generate first-pass import/export declarations [6]. If firms cannot trust AI-generated translations, tariff advice, or compliance summaries without exhaustive human re-verification, the efficiency gains motivating AI adoption are partially or wholly offset — and the resulting caution itself becomes a form of friction, since firms revert to slower, more resource-intensive manual verification processes that a well-functioning AI tool was meant to replace.

Case Evidence

While systematic, large-sample data on hallucination-driven trade disputes does not yet exist publicly, several documented cases illustrate the mechanisms described above.

Moffatt v. Air Canada, 2024 BCCRT 149

After his grandmother's death, a passenger consulted Air Canada's website chatbot, which incorrectly advised that he could apply retroactively for a bereavement discount within 90 days of ticket purchase — a policy that did not exist and contradicted the airline's own published rules [18]. Air Canada refused to honour the discount and argued in tribunal proceedings that it could not be held responsible for information generated by the chatbot, effectively treating the AI system as an independent actor. The tribunal rejected this argument outright, holding that a company is responsible for all information presented on its website regardless of whether it originates from a static page or an interactive AI tool, and ordered Air Canada to pay damages and fees [19,14]. Although this dispute concerned domestic consumer travel rather than cross-border merchandise trade, its legal reasoning generalises directly to any customer-facing or partner-facing AI system used in international commerce: firms cannot outsource accountability to the AI system itself, and hallucinated statements about policy, pricing, or eligibility carry the same legal weight as statements made by a human representative.

Customs and Trade-Compliance Automation

Major logistics providers have moved quickly to embed generative AI into customs workflows. FedEx's 2025 rollout of an AI-powered customs documentation assistant across several Asia-Pacific markets was explicitly framed as a response to the persistent problem of inaccurate shipping documentation, with each AI-suggested tariff code linked back to the official schedule "to ensure full transparency and verification" [20]. A design choice that implicitly acknowledges hallucination

risk and builds in a human-verification safeguard. Trade-compliance technology vendors similarly market AI tools for HS-code classification while cautioning that incorrect classification carries legal and financial consequences [21]. The International Monetary Fund's 2025 technical guidance on generative AI in customs and tax administration likewise frames the technology's value in terms of augmenting, rather than replacing, human risk analysts, reflecting institutional awareness that unverified AI output is not yet reliable enough for autonomous compliance decisions. The World Customs Organization's global survey of AI/ML adoption among customs administrations similarly notes a significant digital divide, with many member administrations lacking the technical capacity to deploy — or to independently audit — AI systems (WCO, n.d.), reinforcing the information-asymmetry mechanism described in Section 4.2.

Cross-Border Communication and Documentation

The WTO's World Trade Report 2025 documents the rapid uptake of AI-powered translation and generative document-summarisation tools in sourcing, after-sales support, and dispute resolution, describing these as improving "first-time-right rates" in import and export declarations [6]. The same report, however, situates this progress within a broader observation that AI adoption remains highly uneven — concentrated in large, urban, digitally connected firms — and that regulatory fragmentation around AI-related goods has intensified, with WTO-notified trade concerns related to AI rising sharply and quantitative restrictions on AI-related goods increasing from roughly 130 in 2012 to about 500 in 2024 [23, 6]. This growing regulatory complexity, layered on top of hallucination risk in the tools meant to help firms navigate it, compounds rather than resolves the compliance burden facing exporters.

DISCUSSION

Taken together, the case evidence and institutional reporting support the paper's central claim: AI hallucination is beginning to function as a *de facto* non-tariff barrier, not through deliberate government policy, but through the interaction of legal liability rules, uneven AI adoption, and the structural unreliability of generative AI output. This is a meaningfully different kind of NTB from those classified in the UNCTAD taxonomy. Traditional NTBs are imposed by governments (technical regulations, licensing, inspection regimes) or emerge from market structure (private standards, certification requirements). AI-hallucination friction is distinctive in that it is imposed by neither party to a trade transaction — it is an emergent property of the technology both parties may be relying on to reduce, not increase, trade costs. This makes it harder to regulate through

conventional trade-policy instruments such as the TBT and SPS Agreements, which presume an identifiable regulatory measure that can be notified, challenged, or harmonised [2].

At the same time, the mechanisms identified in Section 4 suggest hallucination-driven friction will not be evenly distributed. Firms and economies with stronger legal, technical, and compliance capacity are better positioned to catch and correct AI errors before they cause commercial harm, echoing the WTO's own finding that AI adoption and its benefits remain concentrated among large, well-resourced firms [6]. If unaddressed, this suggests AI hallucination could reinforce rather than narrow the divide between economies well-positioned to capture AI's trade-boosting potential and those exposed mainly to its risks — a concern the World Trade Report 2025 raises about AI adoption generally, which this paper extends specifically to the reliability dimension of that adoption.

The *Moffatt v. Air Canada* precedent also suggests that liability frameworks are evolving faster than trade-specific regulation. Courts are already assigning responsibility for AI-generated misstatements to the deploying firm, without waiting for AI-specific trade rules. This creates a practical, near-term incentive for exporters, freight forwarders, and platforms to treat AI-hallucination risk as a compliance and legal-exposure issue today, rather than an emerging concern to be addressed once formal regulation catches up.

Policy Implications and Recommendations

Several implications follow for firms, regulators, and international institutions.

For firms engaged in cross-border trade, the *Moffatt* precedent implies that any AI system customer- or partner-facing — chatbots, automated quote generators, AI-assisted customs documentation tools — should be treated as a source of binding representations, with human-in-the-loop verification built into workflows involving pricing, eligibility, regulatory claims, or tariff classification, mirroring FedEx's practice of linking every AI-suggested HS code back to the official tariff schedule [20].

For customs authorities and trade-facilitation bodies, the IMF's emphasis on using generative AI to augment rather than replace human risk analysts offers a workable interim standard: AI-assisted classification and documentation should remain subject to defined human review thresholds, particularly for high-value or first-time shipments, until hallucination rates in domain-specific customs applications are independently measured and published.

For international institutions, the WTO's existing role as a forum for AI-related trade concerns — 80 specific concerns had already been raised at the WTO as of the

World Trade Report 2025 [4], could be extended to develop shared standards or disclosure norms for AI reliability in trade-facilitation tools, analogous to the transparency and notification obligations under the TBT Agreement [2]. Given the capacity gaps documented by the WCO among member customs administrations (WCO, n.d.), technical assistance and capacity-building programmes should explicitly include AI-verification capability, not just AI-deployment capability, so that adoption does not outpace the ability to audit it.

Finally, given the uneven adoption documented in the WTO-ICC survey, capacity-building efforts aimed at small and medium-sized exporters should treat AI-verification literacy as important as AI-adoption support, since the smaller firms most likely to benefit from AI-driven cost reduction are also the least equipped to independently catch hallucinated output [6].

Limitations and Future Research

This paper is explicitly conceptual and case-based, not a quantitative empirical study, and this is its principal limitation. No large-sample, verifiable primary dataset currently exists measuring the frequency, cost, or trade-volume impact of AI-hallucination-driven friction specifically in cross-border commerce; the case evidence assembled here, while genuine and well-documented, is illustrative rather than statistically representative. Future research should prioritise:

- Structured surveys of exporters, freight forwarders, and customs brokers on the frequency and cost of AI-related documentation errors, replicating the WTO-ICC survey methodology but targeting hallucination specifically
- Domain-specific hallucination-rate benchmarking for customs-classification and trade-compliance AI tools, comparable to the benchmarking already conducted for legal-AI tools [12]
- Comparative legal analysis of how negligent-misrepresentation and product-liability doctrines are being applied to AI systems across major trading jurisdictions, building on the *Moffatt* precedent. This paper's conceptual framework is intended to give that future empirical work a testable structure rather than to substitute for it

CONCLUSION

AI is already reshaping the infrastructure of international trade, and institutional forecasts suggest its role will only grow — the WTO projects AI could raise global trade by roughly a third by 2040 if adoption gaps are bridged [4]. But the same technology carries a structural flaw, hallucination, that has already produced a legally significant, adjudicated case of commercial harm in *Moffatt v. Air Canada*, and that institutional actors from

the IMF to the WCO to major logistics firms are actively building safeguards against. This paper has argued that when hallucination's costs — inflated compliance burdens, information asymmetry, liability exposure, and eroded trust in digital trade tools — are examined through the lens of established non-tariff-barrier theory, they meet the functional definition of an NTB, even though no government intends or imposes them. Recognising AI hallucination explicitly as an emerging category of trade friction, rather than treating it purely as a technology-quality issue, is a necessary first step toward the kind of coordinated policy response, capacity-building, and empirical measurement that traditional NTBs have received over the past three decades.

REFERENCE

1. Wikipedia. "Non-Tariff Barriers to Trade." *Wikipedia*, 2026.
2. Congressional Research Service. "Non-Tariff Measures (NTMs): An Overview." *EveryCRSReport.com*, 2024.
3. United Nations Conference on Trade and Development. "Chapter 1: Non-Tariff Measures: Definitions and Basic Facts." *The Unseen Impact of Non-Tariff Measures*, 2019.
4. World Trade Organization. "AI to Boost Trade by Nearly 40% by 2040 If Gaps Are Bridged, World Trade Report 2025 Finds." *World Trade Organization*, 17 September 2025.
5. World Trade Organization. *World Trade Report 2025: Making Trade and AI Work Together to the Benefit of All—Chapter B: AI, Trade and Inclusive Growth*. 2025.
6. World Trade Organization. "Harnessing AI for Inclusive Growth." *WTO Data Blog*, 29 September 2025.
7. Anh-Hoang, D. *et al.* "Survey and Analysis of Hallucinations in Large Language Models: Attribution to Prompting Strategies or Model Behavior." *Frontiers in Artificial Intelligence*, vol. 8, 2025, article 1622292.
8. Maleki, N. *et al.* "AI Hallucinations: A Misnomer Worth Clarifying." *arXiv*, 2024.
9. Anschutz, M. *et al.* "TUM-MiKaNi at SemEval-2025 Task 3: Towards Multilingual and Knowledge-Aware Non-Factual Hallucination Identification." *arXiv*, 2025.
10. Bao, F. S. *et al.* "FaithBench: A Diverse Hallucination Benchmark for Summarization by Modern LLMs." *arXiv*, 2024.
11. Massenon, R. *et al.* "'My AI Is Lying to Me': User-Reported LLM Hallucinations in AI Mobile Apps Reviews." *Scientific Reports*, vol. 15, 2025, article 30397.
12. Envive.ai. "Case Study of Air Canada's Chatbot Misleading on Bereavement Fares." *Envive.ai*, 2025.
13. Ansari, S. "Compound Deception in Elite Peer Review: A Failure Mode Taxonomy of 100 Fabricated Citations at NeurIPS 2025." *arXiv*, 2026.
14. Brand, J. L. M. "Air Canada's Chatbot Illustrates Persistent Agency and Responsibility Gap Problems for AI." *AI & Society*, vol. 40, 2025, pp. 3361–3363.
15. United Nations Conference on Trade and Development. "Chapter 1: Non-Tariff Measures: Definitions and Basic Facts." *The Unseen Impact of Non-Tariff Measures*, 2019.
16. Diversification.com. "Non-Tariff Barrier: Meaning, Criticisms & Real-World Uses." *Diversification.com*, 2026.
17. Agriculture Institute. "Comparing the SPS and TBT Agreements: What They Mean for International Trade." *Agriculture Institute*, 2026.
18. Sookman, B. "Moffatt v. Air Canada: A Misrepresentation by an AI Chatbot." *TechLex Blog*, McCarthy Tétrault, 26 February 2024.
19. Garcia, M. "What Air Canada Lost in 'Remarkable' Lying AI Chatbot Case." *Forbes*, 19 February 2024.
20. FedEx Corporation. "FedEx Launches Intelligent AI-Powered Customs Solutions in APAC to Streamline Global Trade." *FedEx*, 14 August 2025.
21. iCustoms. "Trade Compliance: Generative AI Is Transforming Business Globally." *iCustoms*, 12 February 2025.
22. PYMNTS. "Courts to Companies: You Own What Your Chatbot Says." *PYMNTS*, 1 July 2026.
23. Insights on India. "World Trade Report 2025: AI to Boost Global Trade & GDP." *Insights on India*, 19 September 2025.