

Agentic AI in Supply Chains: Balancing Trust and Autonomy in Disruption Scenarios

Ana Josefa Matos¹, Bráulio Alturas¹ and Ahmed Kamel²

¹Instituto Universitário de Lisboa (ISCTE-IUL), Centro de Investigação em Ciências da Informação, Tecnologias e Arquitetura, Lisboa, Portugal

²Rollins College, Orlando, Florida, USA

ana.josefa.matos@iscte-iul.pt

braulio.alturas@iscte-iul.pt

akamel@rollins.edu

Abstract: Artificial Intelligence (AI) is no longer just a tool that supports human decisions. It is becoming an active participant in the economy, one that negotiates, transacts, and coordinates autonomously with other AI agents at a speed and scale that no human organisation can match. This shift, from AI as assistant to AI as autonomous economic actor, represents one of the most significant innovations in how organisations operate, compete, and create value. Yet this innovation is outpacing governance. As AI agents increasingly interact with each other rather than with people, organisations are deploying autonomous systems without fully understanding when those systems should act independently, when they should defer, and when they should escalate to human oversight. Without answers to these questions, the promise of AI-driven innovation carries substantial and poorly managed risk. Supply chains are where this tension is most directly felt. They are already among the most complex environments organisations manage, and they are precisely where the cost of getting autonomous AI wrong is highest. Disruptions propagate fast, decisions must be made under uncertainty, and the consequences of miscalibrated autonomy can cascade across entire networks of operators, suppliers, and partners. This research proposes an innovative governance framework that addresses this challenge directly. By developing Trust-based thresholds that dynamically regulate how much autonomy an AI agent exercises depending on the reliability of its counterparts, the research introduces a new model for governing Agent-to-Agent interactions in high-stakes supply chain environments. The framework is grounded in empirical data from the European Gas Transmission Network and validated against the 2022 European Gas Supply Crisis, one of the most significant supply disruptions in recent European history. The expected outcome is both a theoretical contribution, an innovative framework for Trust-governed Autonomy in multi-agent systems, and a practical one: empirical evidence that supply chains governed by this model can recover from disruption autonomously, without continuous human intervention, opening new possibilities for resilient, AI-enabled supply chain innovation.

Keywords: Agentic AI, Supply Chain Resilience, Trust, Autonomy, Agent-to-Agent Economy, Innovation

1. Introduction

The deployment of Artificial Intelligence (AI) in operational environments has undergone a fundamental shift. What began as the automation of discrete, well-defined tasks has evolved into a qualitatively different paradigm: Agentic AI systems that perceive complex environments, reason across multiple variables, and act upon their conclusions with varying degrees of independence from human oversight (Bandi et al, 2025; Schrage & Kiron, 2025a; Al-Bashrawi et al, 2026). This transition carries profound implications for the organisations and sectors that adopt these systems, given that the conditions under which autonomous action is appropriate, and the mechanisms by which it should be governed, remain to a significant extent an open question in both research and practice (Geninatti Cossatin et al, 2026; Lee & See, 2004).

Supply chains represent one of the most consequential environments in which this question materialises. Modern supply networks are complex, globally interconnected systems in which disruptions do not remain localised but propagate across multiple tiers, amplifying in impact as they travel (Ivanov et al, 2014). The efficiency gains that have made these networks globally competitive, lean inventories, geographic specialisation, and tightly coupled logistics, are precisely the characteristics that render them vulnerable when disruption occurs (Manuj & Sahin, 2011). Conventional risk management approaches, built on static models that capture the network as a frozen snapshot, are structurally ill-suited to the speed and scale at which real-world disruptions unfold (Hosseini & Ivanov, 2020; Gupta & Gevaers, 2025).

Energy supply chains illustrate this vulnerability with particular acuity. The 2022 European Gas Supply Crisis (IEA, 2025) demonstrated how cascading failures across interconnected gas transmission networks can expose the fragility of systems optimised for efficiency rather than resilience, and the inadequacy of risk models that cannot adapt in real time to rapidly evolving conditions (Garvey et al, 2015). In high-stakes environments of this kind, where disruptions arrive with little warning and propagate faster than conventional decision cycles can

accommodate, the limitations of static risk modelling are not merely academic, they carry measurable operational and economic consequences (Camur et al, 2024; Song et al, 2024).

It is against this backdrop that the intersection of Agentic AI and Supply Chain Resilience becomes a critical area of inquiry. As AI systems grow more capable and more autonomous, a new dimension of complexity emerges: the Agent-to-Agent (A2A) Economy, in which AI agents transact, negotiate, and coordinate directly with other AI agents, rather than exclusively with human counterparts (Zhu et al, 2025; Kirshner et al, 2026). In these settings, Trust, a concept historically studied as a property of human-automation interaction, must be reconceptualised as a dynamic, probabilistic quantity that governs when one autonomous agent should rely on the decisions, signals, and outputs of another (Afroogh et al, 2024; Lee & See, 2004).

Innovation, broadly defined as the implementation of new or significantly improved processes, products, or organisational methods (OECD, 2005), spans a broad spectrum from technological and process innovation to supply chain innovation (Stundziene et al, 2024). The emergence of the A2A Economy represents precisely this kind of organisational innovation, one that is outpacing existing governance frameworks (Schrage & Kiron, 2025b). Unlike previous waves of digital transformation, where AI augmented human decision-making, this shift places autonomous agents at the centre of operational and economic exchange, creating governance challenges for which no established model yet exists. Addressing those challenges requires innovation not only in technology but in the frameworks used to govern it.

The existing literature has addressed many elements of this challenge in isolation. What it has not provided is a framework that connects Trust, Autonomy, and Supply Chain Disruption Modelling into a unified, empirically validated approach for governing autonomous agent behaviour under disruption. Two critical gaps define this space. The first concerns the Trust-based thresholds that probabilistically govern Autonomy transitions in Agent-to-Agent interactions. The second concerns the capacity of supply chain networks to recover from disruption automatically, what the literature refers to as Self-Healing Network Behaviour, without relying on human intervention at every decision point (Hosseini & Ivanov, 2020; Ojha et al, 2018).

This research, currently in progress as part of a doctoral study, addresses both gaps through an innovative governance framework for Trust-governed Autonomy in Agent-to-Agent supply chain networks, situated within the regulatory framework established by Article 14 of the European Union Artificial Intelligence Act (European Parliament & Council of the European Union, 2024).

2. Theoretical Framework

A systematic literature review was conducted across four databases - Scopus, Web of Science, ACM Digital Library, and arXiv -, covering the period from 2004 to 2026, restricted to English-language peer-reviewed publications across five thematic areas: Agentic AI, Trust in Automation, Supply Chain Resilience, Disruption Risk Modelling, and Agent-to-Agent Interaction. The review yielded a final corpus of 29 papers, organised across 17 thematic topics that together map the intellectual landscape from which this research emerges, as illustrated in Figure 1 below.

2.1 Agentic AI

Agentic AI systems are artificial intelligence systems capable of pursuing goals autonomously, reasoning across complex environments, and acting with minimal human direction (Bandi et al, 2025; Al-Bashrawi et al, 2026). This shift from reactive tools to autonomous actors represents a significant technological innovation, one that is transforming how organisations operate and compete, and that demands equally innovative implementation strategies grounded in context, human agency, and adaptive governance (Thomas et al, 2025).

The intellectual foundations of Agentic AI begin with the philosophical assumptions that shape how AI systems are designed and evaluated. Schrage and Kiron (2025a) argue that AI models are fundamentally shaped by what they aim to achieve, what counts as valid knowledge, and how they represent reality. When these assumptions are left unexamined, the result is misalignment: systems that optimise for the wrong goals or construct representations of reality that do not correspond to the environments in which they operate. Schrage and Kiron (2025b) extend this argument to organisational design, proposing that the value of AI is realised not through deployment alone but through the deliberate architecture of environments in which human and AI agents co-create strategic value, reframing the relevant unit of analysis from the individual model to the sociotechnical system in which it operates. Bandi et al (2025) provide a comprehensive taxonomy of Agentic AI architectures, defining autonomous agents as goal-driven systems that combine planning, persistent memory, tool use, and

multi-agent collaboration with varying degrees of human oversight, identifying orchestration style as the key architectural differentiator. Anjomshoae et al (2019) demonstrate that explainability, the capacity of an agent to explain its own behaviour, is foundational to the Trust placed in it, and identify a near-total absence of research on inter-agent explainability, a gap that becomes directly relevant when autonomous agents must decide whether to rely on each other's outputs in a shared operational environment.

2.2 Trust and Autonomy

Trust in automation has a well-established theoretical foundation. Lee and See (2004) define Trust as the attitude that governs reliance on an automated system under conditions of uncertainty, identifying three bases on which it is built - Performance, Process, and Purpose - and demonstrating that Trust evolves dynamically as experience accumulates. Critically, the authors flag an underexplored dimension: Trust not between a human and a machine, but between machines themselves. Canale and Thimmaraju (2025) extend this concern, proposing that Large Language Models may have inherited the same pre-cognitive vulnerabilities that social engineers exploit in people, a phenomenon they term Anthropomorphic Vulnerability Inheritance. Radanliev et al (2026) demonstrate that such vulnerabilities propagate across perceptual, cognitive, and executive layers of Agentic AI systems. Akkil et al (2026) show through large-scale multi-agent simulation that identical agents produce radically different macro-level outcomes depending on the population of agents surrounding them, establishing that the correct unit of safety analysis is not the individual model but the deployed system as a whole. The regulatory response is provided by the EU Artificial Intelligence Act (European Parliament & Council of the European Union, 2024), whose Article 14 establishes a spectrum of human oversight obligations for high-risk AI systems, from continuous monitoring to limited intervention rights. However, the EU AI Act does not specify how Autonomy levels should be governed dynamically as Trust and Risk evolve during operation. Afroogh et al (2024) identify calibrated Trust, the alignment between the Trust placed in a system and the Trust that system actually merits, as the central challenge in human-AI interaction, and flag the near-total absence of research on Trust in AI-to-AI interaction as one of the most significant gaps in the field. Qian and Wexler (2024) provide empirical grounding for this concern, demonstrating that access to a conversational AI affects demonstrated and perceived productivity in opposite directions, and that reliance behaviour diverges from stated intentions in ways that become even more consequential in Agent-to-Agent settings where no human is present to notice and correct the miscalibration. This is the first gap that the present research addresses.

2.3 Agent-to-Agent Economy

The emergence of the Agent-to-Agent Economy, in which AI agents transact, negotiate, and coordinate directly with one another rather than with human counterparts, marks a new frontier of economic innovation that is reshaping how supply chains operate. Supply chain innovation, defined as a multifaceted process designed to address environmental uncertainties, meet customer needs, and improve organisational processes through the adoption of new technologies (Stundziene et al, 2024), is increasingly driven by these autonomous agent interactions rather than by human-mediated decisions. Zhu et al (2025) characterise this Agent-to-Agent Economy through a large-scale negotiation experiment across nine Large Language Model families, demonstrating that autonomous agent negotiation is an inherently imbalanced game in which stronger models consistently secure better outcomes while weaker models exhibit higher rates of constraint violation, introducing systemic risks that do not exist in human-to-human negotiation. Kirshner et al (2026) benchmark LLM negotiators against human participants in a supply chain bargaining experiment adapted from the foundational human-subjects study of Davis and Hyndman (2021), finding that LLM agents are more agreement-prone but less fair, and that they do not exploit strategic ignorance when operating under private information conditions. At the implementation level, Xu et al (2024) provide one of the first working prototypes of autonomous supply chain management using multi-agent systems, explicitly identifying LLM-powered agents as the natural next step. Quan et al (2026) demonstrate that LLM-based risk severity scoring substantially outperforms rule-based approaches, while Kaur et al (2026) introduce carbon accountability into multi-agent supply chain decision-making, finding that while AI compute carbon is negligible relative to transport carbon for individual shipments, it becomes material at scale and must be measured alongside operational benefits.

2.4 Supply Chain Resilience

The application of AI to supply chain management represents one of the most consequential areas of operational innovation in recent years (Stundziene et al, 2024), offering capabilities in disruption detection and adaptive response that conventional models cannot provide.

Manuj and Sahin (2011) draw a critical distinction between supply chain structural complexity and decision-making complexity, demonstrating that well-designed information systems can substantially reduce perceived complexity even when objective network complexity remains high. Camur et al (2024) and Song et al (2024) establish that the supply chains most in need of intelligent disruption response are precisely those operating under the most severe informational and temporal constraints. The foundational contribution to disruption modelling is the Ripple Effect, formalised by Ivanov et al (2014) as the cascade of impacts that propagates through supply chain tiers when a significant event occurs at any one node. Hosseini and Ivanov (2020) demonstrate that Bayesian Networks are particularly well suited to supply chain risk analysis, capturing causal interdependencies and supporting both forward and backward inference, but identify the static nature of existing models as their principal limitation. Garvey et al (2015) formalise node-level risk propagation measures, validated by Ojha et al (2018) across thousands of simulation runs. Gupta and Gevaers (2025) identify a further limitation: existing models do not capture how multiple distinct risk types compound horizontally across the same tier, creating Super-Risk Events whose severity exceeds what any single-hazard model can anticipate. Singh and Singh (2019) show that big data analytics can convert past disruption experience into genuine resilience, while Hao et al (2025) demonstrate that AI-human collaboration in supply chains operates as a dynamic co-constitutive process requiring active management of epistemic asymmetry and temporal friction. Necula and Rieder (2025) demonstrate that AI can drive circular economy practices across supply chain stages, though the same AI systems introduce their own carbon cost (Kaur et al, 2026). Together, these limitations define the second gap the present research addresses.

Taken as a whole, the seventeen themes reviewed here establish both the urgency and the feasibility of the present research. What makes this contribution innovative is a combination of three characteristics that no existing framework has brought together: it operationalises Trust as a dynamic, probabilistic quantity governing Autonomy transitions in real time, rather than as a fixed property of a human-machine relationship; it grounds this model in empirical data from a real, large-scale supply disruption, making its validation testable and reproducible; and it situates the governance framework within the regulatory boundary conditions of the EU Artificial Intelligence Act, ensuring it is not only technically sound but legally accountable. Together, these characteristics define a contribution that is theoretically original, practically actionable, and responsive to the contextual complexity that effective innovation strategies must navigate (Thomas et al, 2025; Stundziene et al, 2024).




Figure 1: Affinity Map of the Literature Review - 17 Thematic Topics from Foundational Concepts to Research Gaps.

3. Research Questions

The two gaps identified in the literature arise from the intersection of two bodies of knowledge that have developed largely in isolation. Trust in automation provides the theoretical vocabulary for governing agent behaviour, but has not been extended to the dynamics of Agent-to-Agent interaction, where no human is present to recalibrate reliance in real time. Supply chain disruption modelling provides the quantitative tools for capturing network-level risk, but existing models remain static and unable to adapt as conditions evolve during an ongoing disruption. The convergence of these two limitations defines the research space this study occupies, giving rise to two formal research questions:

- RQ1 - What are the Trust-based thresholds that probabilistically govern Autonomy transitions in Agent-to-Agent interactions?
- RQ2 - How effective are these thresholds in Agent-to-Agent supply chain networks at minimising recovery time under disruption?

These two questions are not independent. RQ2 cannot be meaningfully answered without first resolving RQ1. The research is therefore organised sequentially, tracing a deliberate path from the micro, understanding how Trust governs the behaviour of two agents in interaction, to the macro, understanding how those same Trust-based mechanisms shape the resilience of an entire network under disruption.

4. Methodology

4.1 Research Design and Empirical Scenario

The 2022 European Gas Supply Crisis (IEA, 2025) serves as the primary empirical scenario, selected for its clearly identifiable onset, measurable propagation pattern across an interconnected network of operators, and a recovery trajectory that can be observed and quantified using publicly available data. The setting is the European Gas Transmission Network, operated by Transmission System Operators whose interactions are published through the ENTSOG Transparency Platform (ENTSOG, 2026). The simulation timeline spans 2019 to 2025, divided into three phases: a baseline period from 2019 to 2021 used to initialise agent Trust scores; the disruption year 2022; and a recovery and validation period from 2023 to 2025.

Three publicly accessible platforms provide complementary data layers: the ENTSOG Transparency Platform (ENTSOG, 2026) as the primary source for network structure, physical flows, and interruption events; the ENTSO-E Transparency Platform (ENTSO-E, 2026) for complementary energy generation data; and the Eurostat Energy Database (Cazzaniga, 2026) for country-level gas supply and consumption context. All three platforms provide data covering 2019 to 2025. The research operationalises disruption and recovery using statistically-derived thresholds applied to the 2019 to 2021 baseline data, grounded in empirical network behaviour and adapted from EU Regulation 2017/1938 (Council of the European Union, 2017).

4.2 Simulation Model

The research employs a Bayesian-Monte Carlo hybrid probability model (Bayes, 1763; Metropolis & Ulam, 1949). The Bayesian component governs Trust at the agent level, updating Trust scores after every interaction using Bayes' theorem, producing a Trust score that evolves continuously with the history of each agent relationship. The Monte Carlo component governs the simulation strategy at the network level, testing a large number of Trust threshold combinations across repeated simulation runs to identify the configurations that produce the fastest network recovery with statistical reliability.

Four Autonomy levels are proposed, ranging from full human oversight to full agent Autonomy, each grounded in the human oversight requirements established by Article 14 of the EU Artificial Intelligence Act (European Parliament & Council of the European Union, 2024). The framework is an academic interpretation of these requirements, not a reproduction of their legal text. For RQ1, a dyadic network is used to identify Trust-based thresholds in a controlled setting before scaling to the full network in RQ2. For RQ2, a set of agents is selected from the ENTSOG network based on their structural importance to the 2019 to 2021 baseline period, ensuring the simulation captures the most critical points in the European gas transmission network. Both agent sets are LLM-powered, with agent behaviour governed by Large Language Model reasoning rather than fixed rules.

4.3 LLM Integration

A commercially available Large Language Model (LLM), accessed via API with batch processing, serves as the reasoning engine powering agent behaviour throughout the simulation. For RQ1, a calibration approach is used before any simulation runs begin, establishing the decision logic that maps Trust and Risk levels to agent decisions. This decision logic is then applied consistently throughout all simulation runs, eliminating randomness from the threshold search and making results fully reproducible. For RQ2, a hybrid approach is used: agents rely on the calibrated decision logic for individual Trust-based Autonomy decisions, while live LLM calls are made only for Agent-to-Agent communication events during the disruption period, occurring at regular intervals.

5. Expected Results

The expected output of RQ1 is a set of empirically identified Trust-based thresholds that govern Autonomy transitions in Agent-to-Agent interactions, together with a characterisation of how these thresholds perform across the four Autonomy levels under the 2022 disruption scenario.

The expected output of RQ2 is a network-level validation of these thresholds under the 2022 European Gas Supply Crisis scenario (IEA, 2025), with recovery time measured and compared across Autonomy configurations. The research expects to demonstrate that supply chain networks governed by Trust-based Autonomy mechanisms recover from disruption more efficiently than networks operating under fixed Autonomy assumptions, providing empirical evidence for the feasibility of Self-Healing Network Behaviour in Agent-to-Agent supply chain systems, and producing a novel theoretical framework applicable beyond supply chains to any domain where autonomous agents must coordinate under uncertainty.

6. Conclusion

As AI agents move from supporting human decisions to making autonomous ones, governing when they should act independently, defer, or escalate to human oversight becomes one of the most consequential challenges organisations face (Lee & See, 2004; Afroogh et al, 2024). Supply chains represent the environment where this challenge is most acute and where the cost of miscalibrated autonomy is highest (Ivanov et al, 2014; Hosseini & Ivanov, 2020).

This research makes an innovative contribution across three dimensions. Its theoretical innovation is a novel framework for Trust-governed Autonomy in Agent-to-Agent interactions, the first to operationalise dynamic Autonomy transitions using probabilistic Trust thresholds grounded in empirical network data, addressing a gap the Trust in automation literature has identified but not resolved (Lee & See, 2004; Afroogh et al, 2024). Its practical innovation is a governance model organisations can implement to calibrate how much autonomy their AI agents exercise, grounded in context, adaptive governance, and measurable outcomes (Thomas et al, 2025), with direct relevance to supply chain innovation and the Agent-to-Agent Economy, an entrepreneurial frontier reshaping how organisations create and capture value (Al-Bashrawi et al, 2026; Stundziene et al, 2024; Zhu et al, 2025; Kirshner et al, 2026). Its regulatory innovation is the operationalisation of Article 14 of the EU AI Act as a dynamic governance mechanism, translating legal requirements into measurable, simulatable autonomy levels (European Parliament & Council of the European Union, 2024).

As with any research in progress, this study carries inherent limitations. The empirical validation is grounded in a single disruption scenario, which limits immediate generalisability to other supply chain contexts, and the framework's performance in supply chain types beyond gas transmission networks remains to be tested. Future research will extend the simulation to additional scenarios and network types, broadening the empirical basis of the Trust-Autonomy framework.

The research is currently in progress, with simulation construction underway. Future work will report the empirical results of the RQ1 dyadic threshold search and the RQ2 network-level validation, contributing to the growing body of evidence on how AI agents behave when interacting autonomously, and on how that behaviour can be governed effectively when the stakes are high.

Ethics Declaration

This research did not require ethical clearance. The study relies exclusively on publicly available datasets from the ENTSOG Transparency Platform, the ENTSO-E Transparency Platform, and the Eurostat Energy Database. No human participants were involved.

AI Declaration

In preparing this paper, the author utilised AI tools to assist with language refinement and structural organisation. All content was reviewed and verified by the author, who takes full responsibility for the final manuscript.

Acknowledgements

This article was partially supported by national funds through Fundação para a Ciência e a Tecnologia, I.P. (FCT), ISTAR-Iscte Pluriannual Projects UID/04466/2025 (<https://doi.org/10.54499/UID/04466/2025>).

References

- Al-Bashrawi, M.A., Al-Sharafi, M.A., Elgendy, I.A., Helal, M.Y.I., Anbalagan, M.K., Chae, I. and Dwivedi, Y.K. (2026) "Agentic AI systems and the future of entrepreneurship: a perspective on co-agency, innovation, and ecosystem transformation", *International Entrepreneurship and Management Journal*, Vol 22, article 27. <https://doi.org/10.1007/s11365-026-01164-2>
- Afroogh, S., Akbari, A., Malone, E., Kargar, M. and Alambeigi, H. (2024) "Trust in AI: Progress, challenges, and future directions", *Humanities and Social Sciences Communications*, 12, 67. <https://doi.org/10.1057/s41599-025-04374-1>
- Akkil, D., Kokku, R., Vikram, K., Abuelsaad, T., Vempaty, A. and Nitta, S. (2026) "Emergence World: A platform for evaluating long-horizon multi-agent Autonomy", *arXiv*. <https://doi.org/10.48550/arXiv.2606.08367>
- Anjomshoe, S., Najjar, A., Calvaresi, D. and Främling, K. (2019) "Explainable agents and robots: Results from a systematic literature review", *Proceedings of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, pp 1078-1088.
- Bandi, A., Kongari, B., Naguru, R., Pasnoor, S. and Vilipala, S.V. (2025) "The rise of Agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges", *Future Internet*, Vol 17, No. 9, p 404. <https://doi.org/10.3390/fi17090404>
- Bayes, T. (1763) "An essay towards solving a problem in the doctrine of chances", *Philosophical Transactions of the Royal Society of London*, Vol 53, pp 370-418. <https://doi.org/10.1098/rstl.1763.0053>
- Camur, M.C., Ravi, S.K. and Saleh, S. (2024) "Enhancing Supply Chain Resilience: A machine learning approach for predicting product availability dates under disruption", *Expert Systems with Applications*, Vol 247, 123226. <https://doi.org/10.1016/j.eswa.2024.123226>
- Canale, G. and Thimmaraju, K. (2025) "The silicon psyche: Anthropomorphic vulnerabilities in Large Language Models", *arXiv*. <https://doi.org/10.48550/arXiv.2601.00867>
- Cazzaniga, N.E. (2026) *eurostat: Eurostat Python package*. <https://pypi.org/project/eurostat/>
- Council of the European Union (2017) Regulation (EU) 2017/1938 of the European Parliament and of the Council concerning measures to safeguard the security of gas supply. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32017R1938>
- Davis, A.M. and Hyndman, K. (2021) "Private information and dynamic bargaining in supply chains: An experimental study", *Manufacturing & Service Operations Management*, Vol 23, No. 6, pp 1333-1682. <https://doi.org/10.1287/msom.2020.0896>
- ENTSO-E (2026) ENTSO-E Transparency Platform. <https://transparency.entsoe.eu>
- ENTSOG (2026) ENTSOG Transparency Platform. <https://transparency.entsog.eu>
- European Parliament and Council of the European Union (2024) Artificial Intelligence Act - Regulation (EU) 2024/1689. <https://artificialintelligenceact.eu/the-act/>
- Garvey, M.D., Carnovale, S. and Yeniyurt, S. (2015) "An analytical framework for supply network Risk propagation: A Bayesian network approach", *European Journal of Operational Research*, Vol 243, No. 2, pp 618-627. <https://doi.org/10.1016/j.ejor.2014.10.034>
- Geninatti Cossatin, A., Ferrero, F., Ardissono, L. and Mauro, N. (2026) "The Autonomy equation: How Agentic AI reshapes Trust and workload in routine productivity applications", *Information Processing and Management*, Vol 63, No. 5, 104681. <https://doi.org/10.1016/j.ipm.2026.104681>
- Gupta, C. and Gevaers, R. (2025) "Literature review on the interconnectedness of Risk", *Transportation Research Procedia*, Vol 82, pp 2976-2983. <https://doi.org/10.1016/j.trpro.2024.12.231>
- Hao, X., Demir, E. and Eysers, D. (2025) "Beyond human-in-the-loop: Sensemaking between Artificial Intelligence and human intelligence collaboration", *Sustainable Futures*, Vol 10, 101152. <https://doi.org/10.1016/j.sftr.2025.101152>
- Hosseini, S. and Ivanov, D. (2020) "Bayesian Networks for supply chain Risk, resilience and ripple effect analysis: A literature review", *Expert Systems with Applications*, Vol 161, 113649. <https://doi.org/10.1016/j.eswa.2020.113649>
- IEA (2025) Gas market lessons from the 2022-2023 energy crisis, International Energy Agency. <https://www.iea.org/reports/gas-market-lessons-from-the-2022-2023-energy-crisis>
- Ivanov, D., Sokolov, B. and Dolgui, A. (2014) "The ripple effect in supply chains: Trade-off efficiency-flexibility-resilience in disruption management", *International Journal of Production Research*, Vol 52, No. 7, pp 2154-2172. <https://doi.org/10.1080/00207543.2013.858836>

- Kaur, R., Kundu, T., Sharma, B., Park, K.M. and Pinsky, E. (2026) "Operational resilience under carbon constraints: A socio-technical multi-agent approach to global supply chains", *Systems*, Vol 14, No. 4, p 374. <https://doi.org/10.3390/systems14040374>
- Kirshner, S.N., Pan, Y., Wu, J.X. and Gould, A. (2026) "Talking terms: Agent information in LLM supply chain bargaining", *Decision Sciences*, Vol 57, No. 1, pp 9-23. <https://doi.org/10.1111/deci.70010>
- Lee, J.D. and See, K.A. (2004) "Trust in automation: Designing for appropriate reliance", *Human Factors*, Vol 46, No. 1, pp 50-80. https://doi.org/10.1518/hfes.46.1.50_30392
- Manuj, I. and Sahin, F. (2011) "A model of supply chain and supply chain decision-making complexity", *International Journal of Physical Distribution & Logistics Management*, Vol 41, No. 5, pp 511-549. <https://doi.org/10.1108/09600031111138844>
- Metropolis, N. and Ulam, S. (1949) "The Monte Carlo method", *Journal of the American Statistical Association*, Vol 44, No. 247, pp 335-341. <https://doi.org/10.1080/01621459.1949.10483310>
- Necula, S.-C. and Rieder, E. (2025) "Generative and adaptive AI for sustainable supply chain design", *Journal of Theoretical and Applied Electronic Commerce Research*, Vol 20, No. 3, p 240. <https://doi.org/10.3390/jtaer20030240>
- OECD (2005) *Oslo Manual: Guidelines for Collecting and Interpreting Innovation Data*, 3rd edn, OECD Publishing, Paris. <https://doi.org/10.1787/9789264013100-en>
- Ojha, R., Ghadge, A., Tiwari, M.K. and Bititci, U.S. (2018) "Bayesian network modelling for supply chain Risk propagation", *International Journal of Production Research*, Vol 56, No. 17, pp 5795-5819. <https://doi.org/10.1080/00207543.2018.1467059>
- Qian, C. and Wexler, J. (2024) "Take it, leave it, or fix it: Measuring productivity and Trust in human-AI collaboration", *Proceedings of the 29th International Conference on Intelligent User Interfaces*, pp 370-384. <https://doi.org/10.1145/3640543.3645198>
- Quan, Y., Liu, Z., Benaben, F. and Montreuil, B. (2026) "Leveraging Large Language Models to enhance multi-agent Risk assessment in supply chain networks", *International Journal of Production Research*, pp 1-29. <https://doi.org/10.1080/00207543.2026.2619562>
- Radanliev, P., Santos, O. and Maple, C. (2026) "Threats and vulnerabilities in Artificial Intelligence and Agentic AI models", *Frontiers in Artificial Intelligence*, Vol 9, 1731566. <https://doi.org/10.3389/frai.2026.1731566>
- Schrage, M. and Kiron, D. (2025a) "Philosophy eats AI", *MIT Sloan Management Review*. <https://sloanreview.mit.edu/article/philosophy-eats-ai/>
- Schrage, M. and Kiron, D. (2025b) "Stop deploying AI: Start designing intelligence", *MIT Sloan Management Review*. <https://sloanreview.mit.edu/article/stop-deploying-ai-start-designing-intelligence/>
- Singh, N.P. and Singh, S. (2019) "Building supply chain Risk resilience: Role of big data analytics in supply chain disruption mitigation", *Benchmarking: An International Journal*, Vol 26, No. 7, pp 2318-2342. <https://doi.org/10.1108/BIJ-10-2018-0346>
- Song, H., Chang, R., Cheng, H., Liu, P. and Yan, D. (2024) "The impact of manufacturing digital supply chain on supply chain disruption Risks under uncertain environment", *Advanced Engineering Informatics*, Vol 60, 102385. <https://doi.org/10.1016/j.aei.2024.102385>
- Stundziene, A., Pilinkiene, V., Vilkas, M., Grybauskas, A. and Lukauskas, M. (2024) "The challenge of measuring innovation types: A systematic literature review", *Journal of Innovation & Knowledge*, Vol 9, No. 4, 100620. <https://doi.org/10.1016/j.jik.2024.100620>
- Thomas, R., Rees, D., Dennehy, D., Davies, G. and O'Raghallaigh, P. (2025) "Innovation in practice: from strategy development to implementation", *Journal of Decision Systems*, Vol 34, No. 1, 2565554. <https://doi.org/10.1080/12460125.2025.2565554>
- Xu, L., Mak, S., Minaricova, M. and Brintrup, A. (2024) "On implementing autonomous supply chains: A multi-agent system approach", *Computers in Industry*, Vol 161, 104120. <https://doi.org/10.1016/j.compind.2024.104120>
- Zhu, S., Sun, J., Nian, Y., South, T., Pentland, A. and Pei, J. (2025) "The automated but Risky game: Agent to agent negotiations and transactions in consumer markets", *arXiv*. <https://doi.org/10.48550/arXiv.2506.00073>