Proc. of the European Safety and Reliability Conference (ESREL 2026)
Edited by José C. Matos, Paulo B. Lourenço, Michael Beer, Daniel Oliveira,
Edoardo Patelli, Hélder S. Sousa, Rui A. Silva and Monica Santamaria

doi: 10.3850/ESREL2026061419_esrel26-p26252-cd

DAGMAR: A Knowledge-driven Framework for Assessing Supply Chain Resilience

Medina Andresel[1], Maximilian Heß[2], Martin Latzenhofer[1], Daria Liakhovets[1], Sven Schlarb[1]
[1] *AIT Austrian Institute of Technology, Austria. E-mail: {firstname.lastname}@ait.ac.at*
[2] *Supply Chain Intelligence Institute Austria, Austria. E-mail: maximilian.hess@ascii.ac.at*

Modern supply chains are increasingly exposed to systemic risks arising from geopolitical tensions, resource scarcity, and cascading disruptions. At the same time, relevant information for assessing such risks is fragmented across heterogeneous data sources with varying levels of reliability, timeliness, and structure which represents a risk for the availability of critical raw materials and essential goods. This paper presents DAGMAR, an AI and knowledge-driven framework for assessing supply chain risks by integrating validated statistical data, expert knowledge, and dynamically acquired (crawled) unstructured information within a technological framework. DAGMAR combines ontologies and knowledge graphs with neuro-symbolic AI techniques to enable the systematic comparison of AI-generated statements against trusted, curated data sources. The applicability of the framework is demonstrated through a use case on the analysis of phosphorus supply chains for Austria.

Keywords: Supply Chain, Retrieval Augmented Generation, Artificial Intelligence, Large Language Models, Ontologies, Knowledge Management

1. Introduction

Global supply chains have become increasingly complex and interdependent, making them vulnerable to a wide range of disruptions, including geopolitical conflicts, trade restrictions, natural disasters, and resource scarcity. Recent crises, such as the one caused by the COVID-19 pandemic, have shown that disruptions in the supply of critical raw materials and essential products can propagate rapidly across sectors, leading to significant economic and social impacts (Meier and Pinto, 2024). From a safety and reliability perspective, ensuring the resilience of supply chains has therefore become a key concern for governments, regulators, and industry alike.

A fundamental challenge in supply chain risk assessment lies in the heterogeneity and uncertainty of available information. On the one hand, reliable validated data sources, such as official trade statistics or economic indicators typically suffer from significant time lags. On the other hand, more up to date information from news reports, market analyses, or online sources is often unstructured, incomplete, or of uncertain credibility. Additionally, while extracting information from unstructured sources using classical natural language processing (NLP) or machine learning methods may be more up to date, it may introduce errors due to model uncertainty and imperfect language understanding.

To address these challenges, there is a growing need for integrated frameworks that combine trustworthy data with flexible analytical capabilities, while maintaining transparency and explainability of results, and ensuring the accountability of the information issuer.

This paper introduces DAGMAR (Data-driven shortage radar for economic crisis prevention)[a], a framework designed to support trade operations analysis and supply chain risk assessment through the integration of semantic technologies and neuro-symbolic artificial intelligence (AI). DAGMAR brings together leveraged and own ontologies, knowledge graphs, and validated statistical data with unstructured information processed using large language models (LLMs) and retrieval-augmented generation (RAG). By transforming both trusted data and AI-generated insights into a common graph-based representation, the framework enables systematic comparison, validation, and reasoning across information

[a] https://www.kiras.at/en/financed-proposals/detail/dagmar/

sources with different levels of trust and update frequencies.

A key feature of DAGMAR is its explicit treatment of trust levels and temporal characteristics of data. Information is classified according to its origin, validation status, and update cycle, allowing analysts to distinguish between verified facts and emerging, potentially uncertain signals. Building on this foundation, the framework implements indicator-based risk assessment that support informed analysis and decision-making in the face of evolving supply risks.

The contribution of this work consists of (i) an architecture for integrating validated and non-validated information into a unified, ontology-based knowledge space for supply chain risk analysis, (ii) a demonstration of how AI-generated information using neuro-symbolic methods can be validated against trusted knowledge graphs, and (iii) an illustration of the practical applicability of the framework through a hypothetical phosphorus supply chain use case in Austria.

The remainder of the paper is structured as follows: Section 2 reviews related work. Section 3 introduces the general architecture of DAGMAR. Section 4 describes the knowledge representation methods. Section 5 presents the AI-based processing pipelines in detail. Section 6 presents a use case. Section 7 concludes the paper.

2. Related Work

Traditional approaches to supply chain risk assessment rely on quantitative indicators derived from structured data such as trade statistics, production volumes, or concentration measures, such as the Supply Chain Reliability and Resilience (SCRR) framework proposed by Chen et al. (2017) or the Supply Chain Alert Notification (SCAN) monitoring system introduced by Amaral et al. (2022). Network-based indicators, such as the Systemic Trade Risk indicator developed by Klimek et al. (2015), have been commonly used to identify systemic dependencies and vulnerabilities in global trade networks. DAGMAR also uses these indicators, but integrates them into a knowledge graph governed by an ontology to represent trade information and risks.

More recently, advances in data-driven methods and AI have expanded the range of analytical tools available to monitor and assess supply chain risks. The use of AI-integration to supply chain resilience was proposed by Peters et al. (2022) which focuses on embedding machine learning, predictive analytics, and operational optimization into supply chain functions. While these tools enhance the analytical capabilities, they are not embedded into a knowledge-driven framework.

Semantic Web technologies and knowledge graphs have been increasingly applied to integrate heterogeneous data sources and to support explainable analytics in complex domains. E.g., the CoyPu project (InfAI (2023)) focuses on improving supply chain transparency and crisis dynamics using semantic data representations and AI analytics. The CoyPu ontology is used as a core ontology to build the DAGMAR ontology. The potential for AI analysis in this context is demonstrated by Gastinger et al. (2023).

From a technical perspective the approach proposed by Shahid et al. (2025) is using a combination of neuro-symbolic AI to create ontologies with support from human experts to represent supply chain risks focusing on suppliers. DAGMAR is different from this approach, as the above mentioned quantitative indicators derived from structured data are integrated into the ontology, thereby enabling risk related reasoning and triggering of warnings.

Recent neuro-symbolic approaches combine graph neural networks and knowledge graph reasoning to uncover hidden dependencies and latent risks in complex supply chains by inferring previously unknown relationships Kosasih et al. (2024). In contrast, DAGMAR focuses on systematic risk assessment and monitoring, integrating validated statistical data, expert knowledge, and dynamically acquired information within an ontology that explicitly models risk indicators, trust, uncertainty, and temporal validity.

3. DAGMAR Architecture Overview

The simplified view of the DAGMAR architecture is shown in Figure 1. In the following we explain the diagram with the underlying design principles.

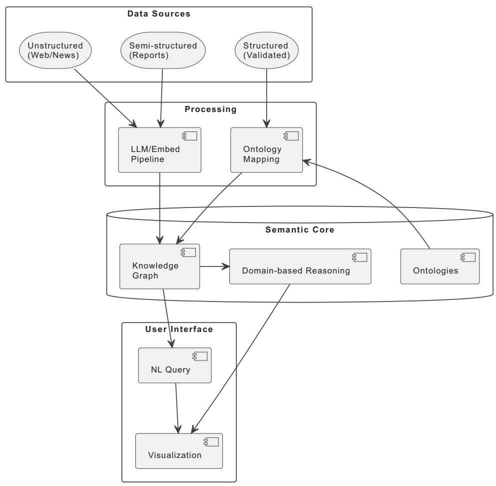

Fig. 1. Simplified DAGMAR Architecture.

In DAGMAR, data sources are differentiated based on their trust status. Highly reliable data consists of structured and validated information, while conditionally trusted data includes semi-structured sources, such as expert reports. Untrusted data refers to unstructured content, such as news articles and web sources. Semi-structured sources are reports that are created according to quality criteria defined by the stakeholder. This could be reports issued by a trusted organisation, for example, which need to meet defined quality and structural requirements. Although the content of expert reports may be inherently reliable, the extraction of information via LLMs or machine learning introduces procedural uncertainty.

The Processing layer serves as the bridge between heterogeneous data sources and the Semantic Core. Unstructured and semi-structured documents are processed through an LLM and embeddings (vector representation) creation pipeline to prepare them for RAG. Automated mechanisms extract entities and relationships from these sources, which are subsequently used to populate the Knowledge Graph within the graph database. In order to bring validated, structured data into the knowledge space, the data undergoes a manual ontology-based mapping to ensure they are integrated into a consistent semantic representation using R2RML mapping language.

All core information is represented using semantic technologies in the semantic core consisting of the graph database and underlying ontology. It models entities, relationships, and indicators and thereby supports traceability of analytical results and facilitates expert validation and domain-specific reasoning using the Nemo reasoner (Ivliev et al., 2024).

At the User Interface level, the DAGMAR system provides a specialized RAG interface that enables analysts to query diverse data streams ranging from news and web sources to internal expert reports. The effectiveness of this interface relies on the prerequisite that adequate content has been crawled and processed by the underlying ingestion pipelines. To maintain a clear distinction between various levels of data reliability, the interface allows users to switch between source types, ensuring that conditionally trusted information and non-validated data remain separated.

Generative artificial intelligence introduces uncertainty and the risk of hallucinations. Therefore, the system provides means to assess the reliability of the generated text. As the model produces an answer, a temporary knowledge graph is dynamically extracted from the generated text while existing ground-truth data from the graph database is visualized simultaneously. By presenting these two structures together, the system enables a human-in-the-loop validation process where the analyst can compare the AI generated output against the established knowledge from the validated knowledge graph.

4. Ontology and Knowledge Graphs

At the core of the architecture lies a curated semantic knowledge base, implemented using ontologies and a graph database. The ontology defines a shared conceptual model for supply chain analysis. It is implemented using Semantic Web standards (including Web Ontology Language (OWL) and Resource Description Framework (RDF)), enabling formal reasoning and interoperability with external datasets and tools. A simplified view of the DAGMAR ontology is shown in Figure 2. This ontology defines the conceptual model for products, trade relations, countries, events, and risk indicators. It serves as a shared vocabulary that permits consistent interpretation of data (e.g., columns of data tables are mapped to an ontology concept) and thereby enables the integration of data across sources.

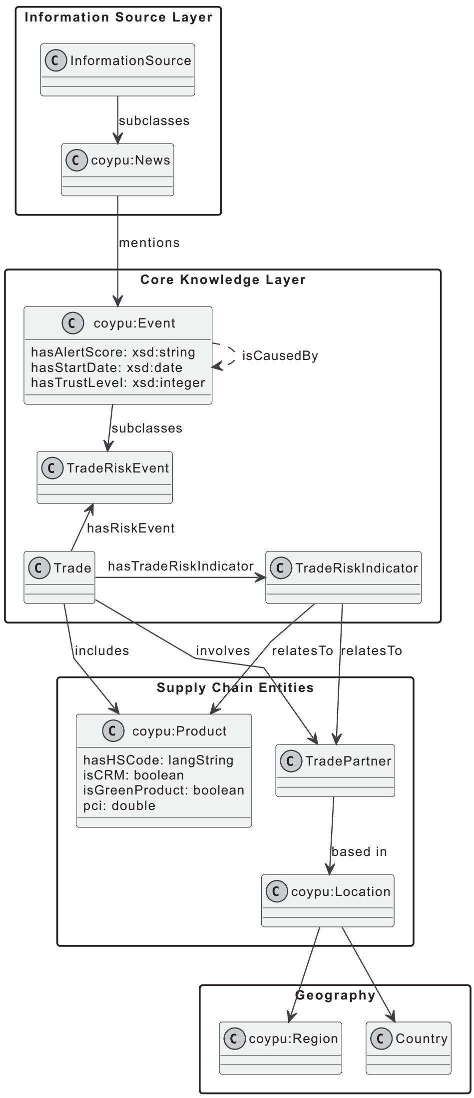

Fig. 2. Simplified DAGMAR Ontology.

The ontology serves multiple purposes: First, it provides a unifying vocabulary for integrating heterogeneous datasets. Second, it constrains and guides automated information extraction from unstructured sources. Third, it ensures that analytical results can be interpreted and validated by domain experts.

The DAGMAR ontology reuses the CoyPu ontology, a well-established ontology for modeling global trade alerts, as a basis to extend the conceptual framework with its own concepts, relations, and additional properties of existing entities.

The knowledge graph materializes this ontology by instantiating classes and relations with information coming from validated datasets by means of hand-crafted mappings.

Further, the ontology also serves as basis for logical reasoning over indicators and events, and to query the knowledge base using formal semantic queries (SPARQL format).

A central requirement for DAGMARs supply chain monitoring is the ability to represent trade related evidence in analyzable form. In DAGMAR, this requirement is addressed through a semantic knowledge representation that combines ontologies, knowledge graphs, and formal risk indicators.

Figure 2 shows a simplified representation of the core concepts and relations of the ontology. In this ontology, economic systems are modeled as networks where edges represent specific trade volumes between actors. The ontology inherently captures the broader supply chain as a transitive sequence of exchanges between multiple nodes in the network. Each trade relation links importing and exporting countries or regions with traded products and associated quantities (trade volume).

By representing trade relations as graph structures, DAGMAR supports network-based analysis of dependencies and concentration, identification of critical nodes and links, and reasoning over indirect exposure through intermediate trade partners.

For risk assessment, DAGMAR integrates a set of indicators to quantify structural supply-chain vulnerabilities, based on prior work on trade-network-based dependency quantification and supply-chain disruption risk quantification (SCAN methodology).

These indicators are derived from validated statistical data like the BACI (Gaulier and Zignago, 2010) dataset and PRODCOM (Eurostat, 2016), and are explicitly represented in the knowledge graph. Concentration and substitution measures adapted from the SCAN framework for the EU market capture the degree of dependence on a small number of external suppliers and the potential to replace disrupted supply chains with production within the EU, respectively. Indirect import dependencies for Austria (e.g., China delivers product X to Germany, which then delivers product X to Austria) are captured by the Systemic Trade Risk indicator, which is a network-based measure. We then apply reasoning to compare situations when multiple indicators are above certain thresholds to identify high risk product imports in specific locations.

The concentration and substitution indicators are associated with formally defined thresholds. These thresholds are not hard-coded but are explicitly represented as part of the knowledge model, allowing transparent adjustment and expert review. As shown in Table 1, different alarm levels are raised according to the combination of indicators that cross thresholds. For example, if all indicators are above the threshold, a high risk is indicated. Table 2 shows the extent of Austria's (indirect) dependence on other countries for imports of a specific fertilizer (HS code 310530), based on the Systemic Trade Risk indicator.

5. Processing of Unstructured Information and AI Integration

In order to complement the curated knowledge graph, DAGMAR incorporates a pipeline for processing unstructured information using AI-based methods including semantic and hybrid retrieval, RAG and graph extraction.

Textual documents are segmented into fragments, embedded as vector representations, and stored in a vector database to enable semantic retrieval. RAG is used to support natural language queries and exploratory analysis. AI-generated responses are not treated as authoritative outputs, but are designed to provide human experts with an overview of the relevant data sources retrieved that can be explored and verified.

Table 1: Products in the phosphorus-containing fertilizer supply chain and their vulnerability to EU supply disruptions in 2023

HS Code	**Product Description**	**Vulnerability**
310390	Other mineral/chemical phosphate fertilizers	**High**
310530	Diammonium phosphate (DAP)	**High**
250300	Sulfur (excl. sublimed, precipitated or colloidal)	Low
280700	Sulfuric acid; oleum	Low
280920	Phosphoric acid and polyphosphoric acids	Low
281410	Ammonia; anhydrous	Low
281420	Ammonia; in aqueous solution	Low
310311	Superphosphates with $P_2O_5 > 35\%$	Low
310319	Other superphosphates	Low
310540	Monoammonium phosphate (MAP), incl. mixed with DAP	Low
251010	Natural calcium/aluminium calcium phosphates and phosphate chalk (unground)	Low
251020	Natural calcium/aluminium calcium phosphates and phosphate chalk (ground)	Low

Additionally, DAGMAR transforms AI-derived insights into structured representations—so-called knowledge graph snippets—that conform to the same conceptual schema as the curated knowledge graph. The extraction schema is designed with

Table 2: Austrian import dependencies for di-ammonium phosphate (HS code 310530) in 2023

HS Code	Trading Partner	Dependency
310530	RUS	**High**
310530	ITA	**High**
310530	MAR	**High**
310530	EGY	**High**
310530	CHN	**High**
310530	JOR	Medium
310530	TUN	Low
310530	FRA	Low
310530	DEU	Low
310530	GBR	Low
310530	LTU	Low
310530	ROU	Low
310530	BEL	Low
310530	POL	Low
310530	ESP	Low
310530	LVA	Low
310530	SVN	Low
310530	EST	Low
310530	NLD	Low
310530	MYS	Low
310530	HUN	Low

pydantic[b] classes and allows fine-grained control of acceptable entities and their types. For example, an instance of ”TradeRiskEvent” might have a ”Location” attribute which is restricted to a set ontology-compliant labels like ”City”, ”Country”, ”Region” etc. This extraction method is implemented using langchain[c] library and ensures compatibility of the extracted graph snippets with the validated ontology. The extraction output is provided in two formats: editable JSON output and networkx[d] node-link graph representation which both can be used for further processing. This way it is possible to compare graph data extracted from generated text against validated graph data retrieved from the graph database. Both graph results are visualized for expert inspection so that contradictions or hallucinations can be revealed.

Text and graph outputs can be assessed with respect to trust level (e.g., validated, partially trusted, or unverified) and provenance of the sources and generative AI/RAG is used to assist the expert. Semantic validation and expert oversight remains in hands of the expert to ensure reliability of the generated output.

DAGMAR uses AI to support exploration of data. Large language models and embedding-based retrieval methods are used to identify potentially relevant information in unstructured sources, support natural language interaction with the system, and to extract candidate facts and relations from textual material. AI-generated outputs, textual or graphs, are never treated as authoritative knowledge. Instead, they are regarded as hypotheses that must be validated against trusted data and expert knowledge before being considered for risk assessments. RAG is used to contextualize AI-generated responses with explicit source material. This mechanism reduces the likelihood of unsupported statements and allows analysts to inspect the underlying evidence.

A central innovation of DAGMAR is the systematic comparison of AI-derived knowledge with trusted graph-based data. Once extracted, non-validated knowledge graph snippets are aligned with corresponding entities and relations in the curated knowledge graph. This comparison enables identification of contradictions between emerging information and established facts, detection of gaps where trusted data are outdated or missing, and assessment of plausibility of AI-generated claims.

The DAGMAR system is human-centered. Domain experts shall not be replaced, but play an active role in reviewing AI-extracted knowledge snippets, interpreting discrepancies between data sources, and confirming or rejecting risk indicators presented by the system.

6. Use Case: Phosphorus Supply Chain Risk Analysis

In order to demonstrate the applicability of the DAGMAR framework in a safety- and reliability-relevant context, this section presents a use case focusing on the phosphorus supply chain for Austria. Phosphorus is a critical raw material with

[b]https://docs.pydantic.dev/latest/
[c]https://www.langchain.com/
[d]https://networkx.org/en/

high relevance for agriculture and food security, and its supply is subject to geopolitical, economic, and environmental risks. The use case illustrates how DAGMAR supports analysts across different stages of risk awareness and escalation.

The use case follows a three-phase analytical model — Cold, Warm, and Hot phase — which structures the assessment process according to urgency and evidence of potential supply shortages or even supply chain disruption.

The *Cold Phase* represents routine monitoring under normal conditions, where no immediate supply disruption is assumed. The primary objective is to establish a baseline understanding of the supply chain and to identify weak signals that may indicate emerging risks.

In this phase, DAGMAR is used to provide an overview of Austria's phosphorus import structure based on validated trade statistics, visualize key trade partners, import volumes, and long-term trends, and to assess structural risk indicators such as concentration or monopoly dependency.

Analysts can formulate natural language queries (e.g., regarding the current supply situation or related to major trading partners), which are translated into semantic queries against the trusted knowledge graph. Results are presented together with explanatory context, allowing analysts to understand how indicators are derived.

At this stage, AI-assisted analysis of unstructured sources is used only in an exploratory manner. Emerging information from news or reports are classified as non-validated.

The *Warm Phase* is activated when early warning signals suggest a possible deterioration of supply conditions. Such signals may include increasing concentration on a small number of exporting countries, political instability in supplier regions, or a noticeable rise in relevant geopolitical events.

In this phase, DAGMAR supports a deeper analysis of trade dependencies using indicator-based risk models. Conceptually, a data catalog allows searching for additional datasets with a higher update frequency, which can be used to validate claims that result from querying news sources and results from unstructured text of emerging claims against trusted data.

Unstructured information, such as reports on export restrictions or geopolitical tensions, is processed using retrieval-augmented analysis. AI-assisted extraction produces structured knowledge snippets that are aligned with the ontology. These snippets are then compared with the curated knowledge graph. Discrepancies—such as claims of export bans not yet reflected in official statistics—are candidates for expert review. This comparison allows analysts to assess whether observed signals are plausible, exaggerated, or unsupported.

The *Hot Phase* is entered when multiple indicators consistently point to a high probability of supply disruption. At this stage, the focus shifts from analysis to decision support and mitigation planning.

DAGMAR concretely supports this phase by providing explainable summaries of dependencies and indicators, and the generation of structured risk reports.

Analysts can evaluate potential mitigation strategies, such as diversification of supply sources or substitution options, to the extent that data are available. While DAGMAR does not automate decision-making, it provides a transparent and evidence-based foundation for policy and operational responses.

At the same time, the use case confirms that expert participation remains indispensable, particularly in interpreting ambiguous signals and defining appropriate responses.

The DAGMAR system and the phosphorus supply chain use case were presented to representatives of the Austrian Federal Ministry of Agriculture, Forestry, Regions and Water Management. The feedback was positive, particularly with regard to the system's potential to support evidence-based risk awareness and to improve transparency in the assessment of critical raw material dependencies. At the time of writing, the project is planning a more systematic evaluation of the results together with additional project stakeholders, which is scheduled for February 2026.

7. Conclusion and Outlook

This paper presented DAGMAR, a knowledge-driven framework for improving supply chain resilience through trust-aware integration of validated data, semantic technologies, and AI-assisted analysis. The DAGMAR framework addresses two dimensions of uncertainty: the provenance of information sources and the reliability of generative AI output.

Rather than seeking to replace expert judgment, DAGMAR is explicitly designed to support informed decision-making under uncertainty. By combining symbolic knowledge representation and reasoning with AI-assisted exploration, the framework enhances situational awareness while maintaining accountability and trust.

An important outlook emerging from the DAGMAR project concerns its conceptual alignment with Data Space architectures. In particular, the project investigated foundational concepts for enabling data sharing and controlled data onboarding in support of the Warm Phase, where additional, higher-frequency or domain-specific datasets are often required. DAGMAR foresees the ability to make such datasets available on demand and to onboard additional Data Space partners when needed.

Acknowledgement

The DAGMAR project is funded within the Security Research Programme KIRAS/K-PASS of the Austrian Federal Ministry of Finance and managed by the Austrian Research Promotion Agency (FFG), grant number 909221.

References

Amaral, A., W. Connell Garcia, F. Di Cometi, and C. Herghelegiu (2022, September). SCAN (Supply Chain Alert Notification) Monitoring System. Single market economics papers, Single Market Economics Papers. Available at SSRN.

Chen, X., Z. Xi, and P. Jing (2017, December). A Unified Framework for Evaluating Supply Chain Reliability and Resilience. *IEEE Transactions on Reliability 66*(4), 1144–1156.

Eurostat (2016). Prodcom - production of manufactured goods.

Gastinger, J., N. Steinert, S. Gründer-Fahrer, and M. Martin (2023). Dynamic representations of global crises: Creation and analysis of a temporal knowledge graph for conflicts, trade and value networks. In *D2R2*.

Gaulier, G. and S. Zignago (2010). Baci: International trade database at the product-level (the 1994-2007 version). Working Paper 2010-23, CEPII.

InfAI (2023). Coypu: Cognitive economy intelligence platform for resilience of economic ecosystems. `https://coypu.org`. Project website, funded under the KI Innovationswettbewerb (BMWK).

Ivliev, A., L. Gerlach, S. Meusel, J. Steinberg, and M. Krötzsch (2024). Nemo: Your friendly and versatile rule reasoning toolkit. In *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning (KR 2024)*, pp. 743–754. IJCAI Organization.

Klimek, P., M. Obersteiner, and S. Thurner (2015). Systemic trade risk of critical resources. *Science Advances 1*(10), e1500522.

Kosasih, E. E., F. Margaroli, S. Gelli, A. Aziz, N. Wildgoose, and A. Brintrup (2024, August). Towards knowledge graph reasoning for supply chain risk management using graph neural networks. *International Journal of Production Research 62*(15), 5596–5612. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/00207543.2022.2100841.

Meier, M. and E. Pinto (2024, February). COVID-19 Supply Chain Disruptions. *European Economic Review 162*, 104674.

Peters, E., L. Knight, K. Boersma, and N. Uenk (2022, 12). Organizing for supply chain resilience: a high reliability network perspective. *International Journal of Operations Production Management 43*(1), 48–69.

Shahid, Z., A. Beckmann, A. Sylla, C. Giannetti, and G. Alpan (2025, January). Ontology-Based Approach to Supplier Risk Management Using Large Language Models. *IFAC-PapersOnLine 59*(10), 2826–2831.