Unlocking Supply Chain Economic Value Through Multi-Tier Supplier Network Visibility

by

Juan E. Cermeño

and

Rafael Rodas

Submitted on May 8th, 2026 in Partial Fulfillment of the
Requirements for the Degree of Master of Applied Science in Supply Chain Management

at the Massachusetts Institute of Technology

ABSTRACT

Supply chain disruptions increasingly originate far upstream, yet most firms maintain visibility only into their direct, tier-1 supplier base. This informational gap creates systematic blind spots in resilience planning: firms allocate risk mitigation budgets based on incomplete signals, underestimating the true cascade exposure of their supplier portfolios. This study quantifies the economic value of expanding supply chain network visibility beyond tier-1 for four focal companies – Mondelez International and Danone (Consumer Packaged Goods) and NVIDIA and Intel (Semiconductor) – representing two distinct industries.

We develop a six-phase simulation framework combining Bayesian Noisy-OR Monte Carlo. The model maps multi-tier supplier networks, scores country-level risk using three composite indices, computes arc-level fragility, and propagates cascading disruption probabilities across three visibility scenarios: Blind (S1, tier-1 only), Partial (S2, tier-1 + tier-2), and Full (S3, three tiers). The model then identifies the minimum expected-cost resilience strategy for each tier-1 supplier across a strategy set comprising six inventory and sourcing options, and computes the Value of Information (VoI) as the percentage reduction in expected disruption cost achieved by each incremental visibility upgrade.

Results show a strong risk-discovery effect: moving from S1 to S3 increases cascade-adjusted failure probabilities by 37–51 percentage points across all firms. This shifts optimal strategies from reactive expediting under S1 to proactive buffering under S2/S3, demonstrating that visibility creates value by enabling better decisions. For CPG firms, tier-2 visibility delivers the highest returns (2.2–2.7% VoI, or $4.1M–$5.0M annually) at low-to-moderate safety stock, while tier-3 becomes material only at high inventory positions (up to $16.7M–$31.4M at 40% safety stock). In semiconductors, outcomes depend jointly on network topology and inventory position: NVIDIA captures roughly $12M–$18M annually at tier-2 across most safety stock levels but exhibits a buffer-driven inversion at 40% safety stock where tier-3 visibility unlocks an additional ~$31M; Intel's dispersed network produces zero VoI at tier-2 and the largest absolute VoI in the study at tier-3, ranging from $46M (at 40% safety stock) to $142M (at 10–20% safety stock).

Capstone Advisor: Dr. Maria Jesus Saenz

Title: Executive Director, MIT SCM Master's Program and Director, MIT Digital Supply Chain Transformation Lab

Capstone Co-Advisor: Dr. Jaime Macias Aguayo
Title: Postdoctoral Associate

TABLE OF CONTENTS

1. Introduction

1.1 Background

The MIT Digital Supply Chain Transformation Lab investigates the complexity of aligning digital, physical, and financial flows to unlock new business models and value creation. The lab's work is focused on understanding the key elements in digital transformations and the new collaborative paradigms that emerge to reshape supply chains. Its research domains include multidimensional collaboration across organizations, digital supply chain capabilities, and artificial intelligence in supply chain management.

1.2 Motivation

This project, sponsored by the MIT Digital Supply Chain Transformation Lab, seeks to quantify the economic value that firms can capture by leveraging multi-tier network visibility. It also aims to provide a data-based framework for supply chain risk management investment decisions.

In today’s global economy, supply chain disruptions have become more frequent and severe, exposing structural vulnerabilities in how firms manage supplier relationships and allocate resources to mitigate supply chain risks (Henrich et al., 2022). Major disruptions like the COVID-19 pandemic, trade and geopolitical tensions, natural disasters, and commodity shortages increasingly start far upstream and across multiple supply chain tiers. These disruptions often cascade through supply networks, resulting in significant operational and financial consequences for firms across the value chain.

Despite these risks, few companies maintain visibility beyond their direct, tier-1 suppliers; only 60% of companies report comprehensive visibility even of their tier-1 suppliers, and a mere 2% have visibility into tier-3 suppliers and beyond (McKinsey & Company, 2021, 2022, 2024). This limited visibility creates a significant blind spot: Companies cannot identify hidden risks, single points of failure, or shared dependencies that exist deeper in their supply networks. Therefore, when firms allocate limited budgets for risk mitigation strategies like dual-sourcing, safety stocks, or supplier development programs, they make decisions based on incomplete information and heuristics rather than data.

Key Questions:

- What is the economic VoI of expanding multi-tier supply network visibility (from tier-1-only to partial (tier-2) and full (tier-3) visibility) and how does this value vary across companies with different network configurations, industry contexts, and safety stock positions?
- How does expanding supply chain network visibility change optimal resilience strategy allocation for tier-1 suppliers, and what is the resulting shift in portfolio-level expected cost across different safety stock positions?
- Under what network characteristics does multi-tier visibility generate the greatest economic value, and what structural features explain differences in VoI across companies?

1.3 Problem Statement

Organizations invest substantial capital in supply chain risk management approaches to ensure business continuity amid rising disruption frequency and severity. These investments include a wide range of mitigation actions and capability-building initiatives.

However, these resilience investments are often made under conditions of severely limited network visibility. Among other factors, technology limitations restrict data exchange, operating model gaps inhibit standardized collaboration, and trust-related concerns discourage information sharing.

Together, these constraints explain why only a small fraction of firms have visibility beyond their tier-1 suppliers (Ebel et al., 2021). A 2022 survey found that 81% of firms implemented dual-sourcing strategies and 80% increased inventory holdings (Alicke et al, 2022), despite limited insight into where vulnerabilities truly reside. This creates a fundamental paradox: firms invest heavily in resilience while operating largely blind to the true structure of their supply networks.

McKinsey estimates that supply chain disruptions over a 10-year period can erode approximately 45% of a company's annual profits, with a single prolonged shock capable of reducing EBITDA by 30–50% (McKinsey & Company, 2020). When disruptions originate deep within multi-tier networks—as is increasingly common in globally fragmented supply chains—limited visibility amplifies uncertainty about where vulnerabilities are concentrated. Consequently, firms default to allocation decisions based on perceived rather than actual exposure. Supplier network data and exogenous risk factor data (i.e., geopolitical, environmental, economic) remain largely untapped for enhancing multi-tier visibility. This study aims to quantify the VoI that firms can capture by leveraging available data to enhance visibility beyond tier-1 suppliers. Ultimately, improved visibility enables more informed supply chain resilience allocation decisions.

When firms operate with limited visibility, they cannot identify hidden vulnerabilities within their supply networks. This information gap manifests in three ways: (a) the inability to detect hidden points of failure affecting multiple tier-1 suppliers, (b) blindness to geographic or supplier concentration risks deeper in the network, and (c) difficulty prioritizing investments based on actual exposure rather than perceived risk. By contrast, enhanced visibility reveals how risk is distributed across the network, allowing firms to prioritize mitigation investments accordingly.

1.4. Project Goals and Expected Outcomes

The main goal of this project is to provide companies with a decision-support framework and prototype tool that enables them to prioritize resilience budget allocations based on multi-tier network visibility insights. The project will quantify the VoI generated by enhanced network visibility, evaluate how visibility influences resilience budget allocation decisions, and demonstrate the strategic and economic value of network visibility. This study will compare tier-1 vs multi-tier network visibility scenarios on two distinct industries – consumer-packaged goods (CPG) and semiconductors — selected for the contrast they offer in supply chain configuration, financial parameters, and disruption characteristics.

The expected deliverables to the sponsor include:

1. A quantitative model estimating the economic VoI of multi-tier supply network visibility.
2. A cross-industry comparison of VoI across CPG and semiconductor focal companies.
3. A decision-support framework that translates visibility insights into resilience budget allocation recommendations.

2. Literature Review

This study examines supply chain resilience in firms that operate in multi-tier supply chain networks. It also explores the value of visibility beyond the first tier of these firms' supplier networks to inform decision-making about actions these firms can take to mitigate the effects of low-probability, high-impact events.

2.1 Supply Chain Resilience

The literature treats supply chain resilience (SCR) as closely related and complementary to supply chain risk management (SCRM). SCRM emphasizes identifying disruption risks and implementing mitigation strategies to reduce their likelihood or consequences. SCR emphasizes building capabilities that allow a firm to absorb shocks and recover quickly to a steady state of operation after a disruption (Bowen & Siegler, 2024; Ribeiro & Barbosa-Povoa, 2018). Some scholars extend this framework to include a fourth capability: organizational learning — the capacity to reflect on past disruptions, update response strategies, and continuously improve operational robustness (Hosseini et al., 2019; Ribeiro & Barbosa-Povoa, 2018).

Researchers identify several resilience strategies, which can be grouped into supplier-focused and inventory/operations-focused approaches:

- **Supplier diversification (geographic)**: Distributing suppliers across regions to reduce exposure to localized disruptions (Ivanov et al., 2019; Snyder et al., 2016).
- **Multi-sourcing**: Using multiple suppliers for critical components to reduce dependence on a single source (Ivanov et al., 2019; Lou et al., 2024).
- **Backup supplier qualification**: Maintaining qualified backup suppliers to enable rapid response when disruptions occur (Lou et al., 2024).
- **Supplier capability investment**: Improving suppliers' technical or financial robustness to enhance stability and responsiveness (Lou et al., 2024; Snyder et al., 2016).
- **Inventory buffering:** Holding safety stock or strategic reserves to absorb shocks and maintain service levels (Dolgui et al., 2018; Snyder et al., 2016).
- **Input substitution flexibility:** Developing the ability to substitute materials when primary inputs become unavailable (Ivanov et al., 2019).

Across these approaches, a consistent theme emerges: resilience decisions require accurate information about where dependencies and vulnerabilities reside. This is where supply chain visibility becomes a critical enabler — the focus of Sections 2.2 and 2.3.

SCR research has examined disruption impacts across multi-echelon supply chains (Schmitt & Singh, 2012; Snyder et al., 2016); developed quantitative resilience frameworks and disruption-related metrics (Hosseini et al, 2019); and analyzed the propagation and amplification of shocks across interdependent networks (Dolgui & Ivanov, 2021). Collectively, this body of work reinforces the view that risk management is not a static set of mitigation activities, but rather a dynamic system in which firms must rely on the best available and most current information to guide resilience-related decisions.

Despite these advances, much of the existing SCR research remains theoretical or computational in nature and does not provide an empirical, decision-oriented framework for quantifying how different levels of supply chain visibility improve the quality of resilience decisions.

2.2 Multi-Tier Visibility and Disruption Propagation

Complex global supply chains consist of numerous interdependent actors distributed across geographies and embedded in diverse cultural, regulatory, and strategic contexts. To understand how disruptions propagate through such systems, supply chain researchers have increasingly adopted concepts from network science and epidemiology, modeling multi-tier supply chains as interconnected networks through which risks diffuse (Basole & Bellamy, 2014).

Multiple modeling approaches have been developed to examine how disruptions propagate in complex supply networks. Table 1 summarizes the four dominant streams, their representative studies, and the mechanisms through which each captures cascading failure.

Modeling Approach	Representative Studies	Key Mechanisms / Methods	Insights for Disruption Propagation
Analytical Models	Wang & Yu (2020); Lou et al. (2024)	Production–inventory formulations; sourcing and pricing trade-offs; optimization-based formulations	Show how local shocks generate system-wide performance losses; quantify trade-offs across mitigation strategies.
Network-Based Models	Simchi-Levi et al. (2015); Dolgui, Ivanov & Sokolov (2018)	Graph-based modeling; percolation theory; network centrality metrics	Characterize nonlinear and cascading failure behavior; identify structural bottlenecks and topological drivers of propagation.
Simulation-Based Models	Dolgui, Ivanov & Sokolov (2018); Simchi-Levi et al. (2015)	Discrete-event and Monte Carlo simulation; scenario-based analyses	Demonstrate how topology, supplier criticality, and structural parameters shape the speed and depth of cascading failures.
Empirical Network Models	Basole & Bellamy (2014); Bowen & Siegler (2024)	Data-driven network reconstruction; centrality analysis; optimization to detect critical suppliers	Provide real-world evidence of deep-tier vulnerabilities; map empirical cascade pathways and hidden dependencies.

Table 1. Modeling Approaches for Disruption Propagation in Multi-Tier Supply Networks

In light of major low-probability, high-impact events such as the 2011 Fukushima disaster and the COVID-19 pandemic, supply chain visibility has emerged as a central lever for mitigating risk in multi-tier networks. The literature increasingly recognizes that the failure of a supplier buried deep in the network can induce severe and persistent downstream consequences. Empirical work by Taghizadeh et al. (2021) shows that limited visibility can obscure critical vulnerabilities and distort resilience assessments. Similarly, Bowen and Siegler (2024) use multi-source empirical data and a combination of graph-based centrality measures and optimization methods to identify critical hidden suppliers deep in a firm's network.

The literature highlights that visibility into deeper tiers is essential for understanding how disruptions are likely to propagate and for enabling firms to take more informed mitigation actions. It also reflects the premise that greater visibility supports resilience – a directional relationship supported by the literature, though its economic magnitude across real, multi-tier networks has not been empirically quantified in financially grounded terms. Representing a critical gap that the present study seeks to address.

2.3 Value of Information in Supply Chain Decision-Making

Supply Chain Network Visibility (SCNV)—the ability of a firm to observe and understand the structure, dependencies, and vulnerabilities of its multi-tier supply base—is widely recognized as a prerequisite for effective risk management (Ivanov et al., 2019).The literature shows that limited visibility contributes to the amplification and propagation of disruptions across tiers, while enhanced visibility improves firms' ability to anticipate, localize, and mitigate such events (Basole & Bellamy, 2014; Taghizadeh et al., 2021; Bowen & Siegler, 2024).

Recent theoretical work provides further insight into how much visibility firms need. For example, when shocks are non-concurrent, firms often need visibility only into an extended local neighborhood—typically up to tier-2—to characterize equilibrium mitigation behavior (Bakshi & Mohan, 2024). Bakshi & Mohan also show that concurrent shocks can increase information requirements.

Determining the VoI is a growing topic of SCM research. Viet et al. (2018) identify a dominant definition in this context: VoI is the difference in expected outcomes when information is available versus when it is not. While prior work provides conceptual and theoretical insights into VoI, the literature offers few quantitative estimates of the value of tier-2 visibility in real supply networks.

This study addresses this gap by integrating the three streams of literature reviewed above. From the SCR literature, we adopt the view that resilience decisions must be grounded in accurate information about where risk resides. From the multi-tier visibility literature, we inherit the finding that tier-1 intrinsic risk is an incomplete signal, and that disruption cascades can amplify exposure by an order of magnitude at deeper tiers. From the VoI literature, we borrow the decision-theoretic framing: the value of information is the difference in expected outcomes between informed and uninformed decision-making. Combining these three perspectives, we treat SCNV as a measurable informational asset and quantify its economic value across multi-tier supplier networks using a Bayesian Noisy-OR Monte Carlo simulation framework applied to four focal companies in the CPG and semiconductor industries.

3. Model and Methodology

This section describes the methodological approach used to quantify the economic value of multi-tier supply network visibility. The goal of the model is to evaluate how the level of upstream visibility influences a firm's ability to allocate its limited resilience budget to ultimately reduce expected losses from supply disruptions. The study applies this framework to focal companies across the consumer-packaged goods (CPG) and semiconductor industries, comparing how VoI varies across network structures, supply chain configurations, and risk environments.

3.1 Methodology Selection

The model is designed as a stochastic multi-tier risk and optimization framework. This design reflects three core methodological choices, each grounded and justified by the prior work reviewed in the literature.

First, the framework represents the multi-tier supply chain as a directed graph in which nodes correspond to firms across the supplier network and arcs represent material flows from supplier to customer. This network-based framing is consistent with an established strand of supply chain research demonstrating that network topology, supplier criticality, and arc level dependencies drive the speed and persistence of cascading failures (Basole & Bellamy, 2014; Dolgui et al., 2018; Simchi-Levi et al., 2015). Treating the supply chain as a structured graph enables the model to identify structural concentration risks that are invisible under a tier-1 only view.

Second, disruption propagation is modeled using a Bayesian Noisy-OR network implemented via Monte Carlo simulation. Bayesian Networks (BNs) provide a probabilistic causal representation in which each node's failure probability is conditioned on the status of its upstream parents. The Noisy-OR formulation assumes that multiple upstream failure causes act independently and that a failure occurs if at least one upstream cause successfully transmits disruption across the connecting arc. Noisy-OR and related stochastic cascade mechanisms have been shown to effectively capture disruption propagation across multi-tier networks when combined with Monte Carlo simulation (Simchi-Levi et al., 2015; Dolgui et al., 2018). Monte Carlo simulation is used to evaluate these probabilistic cascades because closed-form solutions are not available for large, cyclic, multi-tier networks.

Third, the framework incorporates a decision analytic optimization layer that maps probabilistic risk estimates onto resilience strategies and selects the cost minimizing action for each supplier node. This is consistent with established Operations Research / Management Science approaches to supply chain decision-making under uncertainty (Snyder et al., 2016; Tang, 2006).

3.2 Data Collection & Preparation

The model draws on two primary data sources: (1) a proprietary global supplier network dataset and (2) a set of publicly available country-level risk indices. The preparation pipeline transforms and integrates these inputs into a clean, analysis-ready database from which focal company sub-networks and calibrated financial parameters are derived.

3.2.1 Supplier Network Dataset

The core dataset captures supplier-customer relationships across global supply chains, providing annual revenue and inventory levels for each firm. For each relationship record, the dataset includes supplier and customer identifiers, the supplier's annual revenue attributable to that customer relationship, the annual revenues of both parties as measures of firm size, and the customer's baseline inventory allocated proportionally to that supplier. The dataset is stored in a relational SQL database to support efficient multi-tier queries.

Focal company network build-up is performed via a recursive, top-down SQL filtering process. Given a focal company, the procedure identifies all tier-1 direct suppliers, then all tier-2 suppliers that supply the tier-1 set, and then all tier-3 suppliers that supply the tier-2 set. This approach breaks a huge and complex network into smaller and more manageable pieces, while keeping the necessary multi-tier linkage structure needed for cascade modeling. Country of operation is extracted for each node and converted to ISO 3166-1 alpha-3 codes to enable consistent joins with the country-level risk indices.

3.2.2 Inventory and Financial Calibration

Rather than relying on static inventory estimates from the raw dataset, the model applies a turnover-based calibration to derive financially meaningful parameters for each supplier node. Using industry-specific gross margin (m_j) and inventory turnover assumptions, the model computes the daily cost of goods sold (COGS) and daily contribution margin (DCM) for each node. These parameters feed directly into the loss function described in Section 3.3. The inventory efficiency parameter is expressed as a named tier that maps to an annual turnover rate. Full parameter values and justifications are provided in Section 3.4.

To ensure that the sum of the supplier-level inventory allocations equals the focal company's reported balance sheet inventory, the model applies a normalized noise algorithm. A random factor of +/- 30% (reference Table 4) is applied to each supplier's allocation, and the resulting values are renormalized so that the portfolio total exactly matches the reported figure. This procedure is repeated independently across 100 Monte Carlo iterations to quantify the sensitivity of all results to inventory allocation uncertainty.

3.2.3 Country-Level Risk Indices

Three available indices are integrated to characterize country-level supply risk for each supplier node. All three are sourced from their 2024 editions and were selected based on institutional credibility, global country coverage, annual update frequency, and established use in the supply chain risk literature.

- **Climate Risk — INFORM *Hazard & Exposure Index* (2024):** Produced by the European Commission's Joint Research Centre, this index quantifies country-level physical risk exposure across 15 natural hazard types combined with population exposure indicators.
- **Economic Risk — Fund for Peace *Fragile States Index* (FSI) (2024):** Covers 179 countries and aggregates 12 sub-indicators across cohesion, economic, political, and social dimensions into a single state fragility score.
- **Geopolitical Risk — World Bank *Worldwide Governance Indicators, Political Stability and Absence of Violence/Terrorism* (WGI-PV) (2024):** Of the six WGI governance

dimensions, the WGI-PV sub-index is used, capturing perceptions of government destabilization risk from armed conflict, civil unrest, and terrorism.

Missing values for any index are imputed with the column mean before normalization. The three normalized indices are combined into a single Global Risk Index (r_i) per country using equal weighting (1/3 per dimension).

3.3 Application of Methodology

The model executes a six-phase pipeline. Each phase is described below, with equations introduced in the order they appear.

3.3.1 Phase I – Supply Network Graph Construction

The supply chain is represented as a directed graph $G = (V, E)$, where nodes V correspond to firms in the focal company's multi-tier supplier network and directed arcs E represent material flows from suppliers to customers. An arc $(i \to j)$ indicates that the firm i supplies firm j. This graph is constructed from the SQL recursive extraction described in Section 3.2 and stores, for each arc, the supplier-customer revenue dependency and node-level risk data.

Each node i is characterized by two fundamental quantities. Its intrinsic failure indicator $I_i \in \{0,1\}$, where $I_i = 1$ denotes a disruption originating within the node i own risk environment. And its marginal failure probability $p_i = \Pr\{I_i = 1\}$, derived from the Logit calibration in Phase II.

3.3.2 Phase II – Environmental Risk Scoring and Logit Calibration

The Global Risk Index r_i computed in Section 3.2.3 is converted to an annual intrinsic disruption probability ℓ_i using an empirically calibrated Logit link function:

$$(1) \quad \ln\left(\frac{\ell}{\ell - 1}\right) = \alpha + \beta * r$$

The Logit function is widely used to map continuous risk scores into bounded probabilities when the true underlying relationship is nonlinear and increasing (McFadden, 1974; Train, 2009). In risk modeling applications, the Logit transformation is particularly appropriate because it ensures probabilities remain in [0, 1], accommodates heavy tails at high-risk values, and allows transparent calibration using empirically interpretable anchor points. Here, parameters α and β are calibrated by anchoring the model to two empirical reference points drawn from the risk distribution: A low-risk anchor at the 10th percentile of the global risk score distribution (r_L) and a high-risk anchor at the 90th percentile (r_H). Disruption probabilities ℓ_L and ℓ_H are assigned at these percentiles, and the system of two equations is solved in closed form for α and β. The values chosen for these anchor probabilities, and their justification, are detailed in Section 3.4. The calibrated Logit function is then applied to every country in the database to produce the full risk map of intrinsic disruption probabilities. The tier-1 baseline probability for each supplier node is $p_i(S1) = \ell_i$.

3.3.3 Phase III – Arc Fragility Computation

Arc fragility q_{ij} for each supplier-customer arc $(i \to j)$ represents the conditional probability that a failure at node i propagates to and functionally disrupts the node j, given that i has failed. It is the noise parameter in the Noisy-OR gate: even when an upstream node fails, the disruption only propagates if an independent arc level Bernoulli draw with probability q_{ij} succeeds.

Arc fragility is constructed from two weighted components:

- Revenue dependency (q_{rev}): Captures the fraction of the customer's revenue at risk from this supplier, computed from the customer's perspective.
- Size ratio (q_{imp}): Reflects the supplier's size relative to the customer, capturing the structural power asymmetry in the relationship.

The two components are combined using a 50:50 weighting, and the composite is bounded by a ceiling that depends on the scaling mode selected for q_{imp}. The default scaling mode and its justification are described in Section 3.4.

3.3.4 Phase IV – Stochastic Noisy-OR Cascade Propagation (Monte Carlo)

With intrinsic probabilities and arc fragilities defined, the model simulates disruption cascades using a Monte Carlo simulation of the Noisy-OR gate logic. At the beginning of every trial, a binary outcome is assigned to each node through a Bernoulli draw with a given probability. This determines whether the node fails due to intrinsic factors. These outcomes then drive cascade logic according to the Noisy-OR structure across three visibility scenarios:

- **Scenario (S1) – Blind (tier-1 only):** A tier-1 supplier is recorded as failed only if its own intrinsic Bernoulli draw succeeds. No upstream cascade information is used.
- **Scenario (S2) – Partial (tier 1 & 2):** A tier-1 supplier fails intrinsically or if any tier-2 supplier fails and that failure crosses the tier-2 to tier-1 arc (a second Bernoulli draw with probability q_{ij} succeeds). The OR gate is satisfied as soon as any one upstream path triggers failure.
- **Scenario (S3) – Full (tier 1, 2 & 3):** The full three-tier cascade logic applies: Tier-3 failures can propagate to tier-2, and tier-2 failures can propagate to tier-1, each step governed by its own independent arc-level Bernoulli draw.

After N trials (see Section 3.4), the empirical failure frequency for each tier-1 supplier is recorded as $p_i(S1)$, $p_i(S2)$, and $p_i(S3)$ respectively. These illustrate the risk-discovery impact associated with increased visibility, as well as the further cascade risks identified when the company conducts a more comprehensive mapping of its supply network.

3.3.5 Phase V – Disruption Duration Modeling

Disruption duration τ_i is modeled as a stochastic variable drawn from a Lognormal distribution with mean μ_i and standard deviation σ_i, bounded between a minimum triage floor and an operational ceiling. The Lognormal distribution is well-suited for disruption duration modeling because it is strictly positive, right skewed, and captures the empirical reality that most disruptions resolve within a moderate window while a smaller proportion extend significantly longer. Industry-specific values for these parameters are given in Section 3.4. The continuous distribution is discretized into a probability mass function (PMF) used in expected-value loss calculations.

3.3.6 Phase VI – Multi-Strategy Optimization and the Loss Function

The optimization layer evaluates six candidate resilience strategies against the cascade probabilities discovered under each visibility scenario for every tier-1 supplier. As shown in Table 2, strategies differ in two action parameters: δt_k (additional proactive buffer days above the baseline) and t_{e_K} (maximum days of reactive expedited sourcing coverage).

Action	Strategy Name	Buffer Days (δt_k)	Expediting Days (t_{e_K})	Description
A1	Strategic Buffer	14	0	Hold 14 additional buffer days of inventory; no expediting capability.
A2	Tactical Buffer	7	0	Hold 7 additional buffer days; suited for moderate-criticality inputs.

A3	Do Nothing	0	0	No proactive investments: all disruption losses absorbed as incurred.
A4	Resilient Hybrid	7	7	Moderate buffer combined with limited expediting; balanced resilience.
A5	Tactical Expediting	0	7	No proactive buffer; rely on expedited sourcing for up to 7 days of coverage.
A6	Strategic Expediting	0	14	No proactive buffer; rely on expedited sourcing for up to 14 days of coverage.

Table 2. Resilience Strategy Matrix - Action Parameters and Descriptions

For each supplier node i and strategy a_k, the total annual cost exposure integrates three financial components:

- **Risk of Lost Margin:** Expected contribution margin lost to unmet demand on days not covered by the buffer, base stock, or expediting window.
- **Expediting Premium:** Incremental cost of reactive emergency sourcing for days covered by expediting (up to t_{e_K} days).
- **Holding Cost:** Annual opportunity cost of maintaining the proactive buffer, computed as: $Annual\ Inventory\ Cost\ (a_k) = h * COGS_{day_i} * \ \delta t_{i,k}$, which is incurred regardless of whether a disruption occurs.

The expected loss for strategy a_k under disruption probability p_i is:

$$(2) \quad E[L(a_k, X,)] = p_i * \sum_{\tau} [\Omega\ (a_k, \tau) * \Pr(\tau)] + Annual\ Inventory\ Cost\ (a_k)$$

The optimal strategy for supplier i under visibility scenario Sn is:

$$(3) \quad a^*(Sn) = \ argmin_k\ E[L(a_k\ , X)]\ given\_p_i(Sn)$$

This optimization is performed independently for each tier-1 supplier under all three visibility scenarios, producing three complete decision matrices from which the network-level portfolio cost is obtained by summing each supplier's minimum-cost strategy.

3.4 Model Parameters

The model relies on two categories of user-defined parameters:

- **Model Level**: Governs the architecture and probabilistic mechanics of the framework and apply uniformly across all focal companies and industries.
- **Industry Scenario**: Calibrates to the operational and financial characteristics of a specific industry and defines the baseline scenario for each focal company analysis.

Both sets of parameters have been selected based on a combination of supply chain management literature, industry conventions, and expert judgment.

3.4.1 Model Level Parameters

Table 3 presents the model-level parameters used across all analyses, along with the default values applied in the baseline scenario and the justification for each choice. Several parameters in the framework are judgment-based due to the limited availability of firm-level disruption data. For example, equal weighting across risk indices and the selected logistic transformation anchors are intended to provide a normalized comparative framework rather than a calibrated predictive model.

Parameter	Default Value	Justification

Monte Carlo Trials (N)	10,000	Empirically sufficient for stable probability estimates across all tier-1 nodes; increasing beyond this threshold yields diminishing variance reduction relative to computational cost. (Fishman, 1996; Simchi-Levi et al., 2015; Dolgui et al., 2018; Ivanov, 2020).
Monte Carlo Iterations (I)	100	100 inventory-noise iterations per safety stock level quantify the sensitivity of all cost and VoI estimates to allocation uncertainty, providing 95% confidence intervals on all reported results.
Random Seed	42	Fixed seed ensures full reproducibility of all simulation results across runs and across focal companies.
Risk Index Weights	1/3 each (equal)	Equal weighting reflects the absence of a strong a priori basis for preferring one risk dimension over another in the absence of calibrated historical data.
Logit Cumulative Distribution Function (CDF) Scope	All countries in database	Using all countries for the CDF construction yields universal calibration consistent across all focal companies, enabling direct cross-company comparisons of disruption probabilities.
Logit Low-Risk Anchor (ℓ_L)	2% annual probability	Calibrated to the empirically observed disruption frequency for suppliers located in the most stable operating environments (10th percentile of global risk).
Logit High-Risk Anchor (ℓ_H)	20% annual probability	Calibrated to the empirically observed disruption frequency for suppliers in the most exposed operating environments (90th percentile of global risk).
Arc Fragility Scaling Mode (q)	Constant: $q = 1$	Arc fragility is computed as a 50:50 weighted composite of revenue dependency and supplier size ratio, uncapped up to a ceiling of 1. This allows propagation probability to vary per arc based on the actual financial and structural characteristics of each supplier-customer relationship, rather than applying a uniform stress-test value.
Arc Fragility Weights $(q_{rev} : q_{imp})$	50:50	Revenue dependency and size ratio are each assigned equal significance, providing balanced insight into the focal company's financial exposure and structural dependency on individual supplier relationships.
Revenue Dependency Ceiling $(q_{rev}\ clip)$	25% of customer revenue	Prevents any single supplier relationship from dominating the composite arc fragility score. Reflects the view that a supplier accounting for more than 25% of the customer's revenue is already in a category of strategic dependency managed separately.
Holding Rate (h)	15% per annum	Standard opportunity cost of capital applied to buffer inventory, consistent with commonly used carrying cost rates in supply chain management literature and industry practice (typically 15-25%). Represents the annual insurance premium for maintaining proactive buffer stock.
Expediting Cost Penalty (k)	50% above standard cost	Reflects the incremental freight, premium sourcing, and logistics coordination costs associated with reactive emergency procurement. Consistent with industry estimates of 30-75% expediting premiums for short-notice alternative sourcing.
Missing Country Imputation	Column mean of disruption probability	Suppliers with unresolvable country data are assigned the network average disruption probability, preserving their inclusion in the model without introducing extreme assumptions about their risk profile.

Table 3. Model-Level Parameters with Default Values and Justification

3.4.2 Industry Scenario Parameters

Table 4 presents the industry-specific parameters that define each focal company's baseline scenario. Parameter values were derived primarily from the most recent annual reports (10-K filings) of the four focal companies, supplemented by supply chain management literature. When values are based on industry benchmarks instead of specific company data, the justification column explains the reasoning for each assumption.

Parameter	CPG	Semiconductor	Justification
Gross Margin (m_j)	40%	55%	CPG gross margins reflect competitive, promotion-heavy markets with strong retailer bargaining power. Semiconductor margins are higher due to capital intensity, IP-driven differentiation, and persistent supply-demand imbalances.
Inventory Efficiency Tier	mhigh (6 turns / ~2 months stock)	lmed (3 turns / ~4 months stock)	CPG firms exhibit high inventory efficiency driven by stable demand and continuous replenishment. Semiconductor firms operate with lower turns due to long production cycles, yield uncertainty, and complex supplier networks.
Severity Profile	Standard	Rigid	CPG disruptions are modeled as moderate and operationally recoverable. Semiconductor disruptions are modeled as longer-onset and slower-recovery, consistent with the capital intensity and lead-time structure of fabrication supply chains.
Disruption Duration - Mean (μ_τ)	45 days	60 days	The standard CPG profile uses a 45-day mean reflecting moderate operational disruptions. The rigid semiconductor profile uses a 60-day mean, consistent with longer recovery cycles in high-tech manufacturing environments.

Disruption Duration - Std Dev (σ_τ)	15 days	15 days	A standard deviation of 15 days introduces realistic recovery variability arising from site-specific and operational factors while keeping disruption severity comparable across industries.
Duration Floor ($d_{\min}$)	30 days	45 days	The CPG floor of 30 days reflects the shortest plausible recovery period for meaningful supply disruptions once detection, escalation, and mitigation actions occur. The semiconductor floor of 45 days reflects the longer minimum lead times inherent to specialized fabrication.
Duration Ceiling ($d_{\max}$)	90 days	120 days	The CPG ceiling of 90 days captures extended but recoverable disruptions. The semiconductor ceiling of 120 days reflects the potential for prolonged recovery in specialized, low-substitutability fabrication environments.
Safety Stock (SS) Levels Tested	10%, 20%, 30%, 40% of balance sheet inventory	10%, 20%, 30%, 40% of balance sheet inventory	Four safety stock levels are evaluated as the primary sensitivity dimension, ranging from lean (10%) to well-buffered (40%) inventory positions relative to the focal company's reported balance sheet inventory.
Inventory Noise Factor	+/- 30%	+/- 30%	A +/- 30% inventory noise factor reflects operational uncertainty during disruptions (including demand swings and execution variability) while maintaining symmetry across industries for comparability.

Table 4. Industry Scenario Parameters With Baseline Values and Justifications

3.5 Value of Information

The final component of the framework quantifies the financial benefit of expanding supply chain visibility. The VoI calculation evaluates the cost of policies chosen with incomplete information against the cost of those same policies assessed in the full-visibility scenario, which captures the actual multi-tier risk structure.

3.5.1 Three Scenario Policy

For each visibility level, the model first identifies the optimal policy a decision maker at that information level would select:

- $a^*(S1)$: The strategy chosen uses only the tier-1 intrinsic probability $p_i(S1)$ (Blind scenario).
- $a^*(S2)$: The strategy chosen uses the tier-1 + tier-2 cascade probability $p_i(S2)$ (Partial visibility scenario).
- $a^*(S3)$: The strategy chosen uses the full three-tier cascade probability $p_i(S3)$ (Full visibility scenario).

Each policy is then evaluated against the next visibility scenario's cost matrix. This cross-evaluation ensures that the VoI comparison reflects the actual financial consequences of operating under different information sets. The expected loss of applying the policy $a^*(Sn)$ is given by:

$$(4) \qquad E[L(a^*(Sn)|\, S3)] = \sum_X [\Omega\ (a^*(Sn), x]\, p_{S3}(x)$$

3.5.2 VoI Computation

The VoI for each visibility upgrade is expressed as a percentage savings relative to the cost incurred when the prior-scenario optimal policy is applied in the next-scenario risk environment. This anchors each percentage to the baseline a decision-maker would actually face when contemplating the visibility upgrade. Note that VoI is expressed as a percentage reduction in expected cost, consistent with the decision-theoretic convention in the supply chain literature (Viet et al., 2018; Candogan & Gurkan, 2025), while strategy-level expected losses reported in Section 3.3.6 are expressed in absolute dollar terms to enable direct portfolio cost comparison across companies. These two representations are complementary: the percentage form isolates the relative information gain, while the dollar form enables cross-company comparison of absolute risk exposure.

$$(5) \qquad VoI(S2)\ \% = \frac{[Cost\ (a^*(S1)\ evaluated\ in\ S2) - Best\ Cost\ (S2)]}{Cost\ (a^*(S1)\ evaluated\ in\ S2)}$$

$$(6) \qquad VoI(S3)\,\% = \frac{[Cost\ (a^{*}(S2)\ evaluated\ in\ S3) - Best\ Cost\ (S3)]}{Cost\ (a^{*}(S2)\ evaluated\ in\ S3)}$$

This approach quantifies the incremental value of each visibility step by calculating the percentage reduction in expected cost achieved through the optimal policy selected under expanded visibility, compared to the policy inherited from the previous information state. All reported VoI estimates are averaged across 100 inventory-noise iterations per safety stock (SS) level, with 95% confidence intervals computed as $1.96 * (\frac{\sigma}{\sqrt{100}})$.

The gain from extending visibility from tier-2 to tier-3 is isolated as VoI (S3) – VoI (S2), quantifying the incremental value of the deepest mapping tier. For a complete cost-benefit comparison, the cost of acquiring additional tier mapping is subtracted from the gross VoI:

$$(7) \qquad Net\ VoI(S_n) = [VoI(S_n)\% * Cost\ (a^{*}(S_{n-1})\ evaluated\ in\ S_n)] - Cost\ of\ Mapping\ to\ tier\ n$$

In Equation 7, the gross VoI percentage is converted to a dollar figure using the prior-scenario baseline cost; subtracting the cost of acquiring the additional tier mapping yields the net dollar value of the visibility investment.

4. Result Discussion and Analysis

This section presents the main findings of the multi-tier supplier network visibility study across the four focal companies – Mondelez International and Danone (CPG) and NVIDIA and Intel (Semiconductor). The analysis follows the six-phase pipeline described in Section 3 and is organized into four subsections.

4.1 Network Profiles and Risk Discovery

The four focal companies exhibit meaningfully different network structures and risk profiles. Mondelez and Danone operate with 31 and 26 tier-1 suppliers respectively, drawing from supplier bases concentrated in the United States, Europe, and select emerging markets. NVIDIA maintains 37 tier-1 suppliers with a pronounced geographic concentration in Taiwan, Japan, South Korea, and the United States – a configuration that reflects the highly specialized semiconductor fabrication ecosystem. Intel, which represents the study's largest network, consists of 80 tier-1 suppliers distributed across 20 countries. Among these are Egypt and India, both of which have significantly greater inherent disruption risks compared to most others in the portfolio.

Under the Blind scenario (S1), the average tier-1 intrinsic failure probability is approximately 5.0% across all four companies – a reflection of the Logit calibration anchored to a global risk distribution in which most tier-1 suppliers reside in relatively stable operating environments. This narrow range across companies is by design: The S1 probability is driven solely by each supplier's country-level risk index, and the four focal companies concentrate much of their direct supplier base in similarly low-to-moderate risk geographies.

The risk discovery effect is substantial and varies meaningfully across companies. As summarized in Table 5, expanding from S1 to S3 visibility increases the average tier-1 failure probability by between 37 and 51 percentage points relative to the S1 blind-visibility baseline, with the largest absolute amplification observed for NVIDIA (+51.2 pp, from 5.1% to 56.3%) and Danone (+45.5 pp, from 5.0% to 50.5%). Intel and Mondelez exhibit somewhat lower S3 probabilities, though all four companies demonstrate that tier-1 intrinsic risk materially understates the true cascade-adjusted exposure.

Focal Company	Avg P(S1)	Avg P(S2)	Avg P(S3)	Avg Risk Discovery Delta (S1→S3)
Mondelez	5.7%	20.2%	43.1%	+37.5
Danone	5.0%	21.4%	50.5%	+45.5
NVIDIA	5.1%	30.1%	56.3%	+51.2
Intel	5.0%	17.4%	47.8%	+42.8

Table 5. Cascade-Adjusted Risk Discovery: Average Tier-1 Failure Probability by Visibility Scenario

The risk discovery effect is not uniform across tier-1 suppliers within any given company. One subset of suppliers exhibits zero cascade amplification (meaning each supplier's upstream network contains no higher-risk tier-2 or tier-3 dependencies) while another subset exhibits near-complete amplification, with cascade-adjusted probabilities converging toward 1.0 under full visibility.

For Danone, seven of 26 tier-1 suppliers reach P(S3) ≥ 99%, driven by deep-tier dependencies on suppliers operating in Western and Eastern Europe. Mondelez exhibits a smaller but comparable concentration, with four of 31 tier-1 suppliers crossing the same threshold, anchored by exposures in India, Mexico, and Poland. NVIDIA serves as an illustrative case: Samsung Electronics, classified as tier-1, exhibits an intrinsic S1 probability of 2.7%, yet this figure increases significantly to P(S2) = 99.9% when its upstream supply chain is considered. This change underscores the concentration of fabrication risk prevalent in East Asia. Likewise, Taiwan Semiconductor Manufacturing Company (TSMC) rises from a P(S1) of 4.1% to P(S3) = 99.1%. These findings directly illustrate the core motivation of the study: tier-1 intrinsic risk is a structurally incomplete signal, and the true cascade-adjusted exposure for strategically critical suppliers can be an order of magnitude higher.

4.2 Optimal Resilience Strategy Allocation

The optimization layer calculates the expected cost for all strategies for each tier-1 supplier under each visibility scenario, drawing from the six strategies described in Section 3.3.6 (A1–A6). Table 6 summarizes the portfolio-level strategy distribution for the 40% safety stock scenario. Portfolio Cost ($M) represents the total expected annual cost of the optimal resilience strategy portfolio, computed as the sum of each tier-1 supplier's minimum expected loss (combining proactive holding costs, risk-weighted lost margin exposure, and expediting premiums) under the given visibility scenario. It captures the full financial consequence of the firm's resilience posture, not merely the cost of a single supplier or strategy.

Company (Scenario)	Total Tier-1 Suppliers	A1	A2	A3	A4	A5	A6	Portfolio Cost ($M)
Mondelez – S1 (Blind)	31	0	0	0	0	0	31	$45.4
Mondelez – S2 (Partial)	31	5	0	0	0	0	26	$127.3
Mondelez – S3 (Full)	31	13	0	0	1	0	17	$175.7
Danone – S1 (Blind)	26	0	0	0	0	0	26	$15.5
Danone – S2 (Partial)	26	4	0	0	2	0	20	$102.4
Danone – S3 (Full)	26	15	0	0	0	0	11	$197.0
NVIDIA – S1 (Blind)	37	0	0	0	0	0	37	$14.4
NVIDIA – S2 (Partial)	37	9	0	0	1	0	27	$128.9
NVIDIA – S3 (Full)	37	22	0	0	0	0	15	$219.2
Intel – S1 (Blind)	80	0	0	0	0	0	80	$36.7
Intel – S2 (Partial)	80	5	0	0	2	0	73	$182.8
Intel – S3 (Full)	80	39	0	0	1	0	40	$410.4

Table 6. Optimal Strategy Distribution by Visibility Scenario (Safety Stock = 40%; Number of Tier-1 Suppliers)

Several patterns emerge from Table 6. Under the Blind scenario (S1), A6 (Strategic Expediting: 14-day reactive expediting coverage, no proactive buffer) is the optimal strategy for every tier-1 supplier across all four companies. This is a direct consequence of the low tier-1 intrinsic probabilities under S1: when the model perceives failure risk as approximately 5% per year, the holding cost of any proactive buffer (A1, A2, A4) exceeds the expected loss reduction it would provide, so the model defaults to A6 – which imposes no holding cost and only triggers expediting premiums in the rare event of an actual disruption.

As visibility expands to S2 and S3, the cascade-adjusted failure probabilities rise substantially, shifting the economic calculus toward proactive buffering. Under S3, A1 (Strategic Buffer, 14 buffer days) becomes the optimal strategy for a growing share of suppliers across all four companies: 13 of 31 for Mondelez, 15 of 26 for Danone, 22 of 37 for NVIDIA, and 39 of 80 for Intel. This strategy shift has an important practical implication: a firm operating in blind mode is systematically under-protected relative to its true risk environment, relying on reactive expediting to cover contingencies that are far more frequent than the tier-1 signal suggests. Enhanced visibility reveals this gap and drives a reallocation of resilience investment toward a proactive buffer stock for the highest-risk portion of the portfolio.

The portfolio cost trajectory tells the same story from a financial angle. As visibility expands, the model's expected portfolio cost rises sharply, not because the underlying risk increases, but because the model now has the information required to recognize losses that were previously invisible. The S1 cost figures are not a target to preserve; they are a measurement artifact of incomplete information. The actionable insight is the difference between strategies the firm would choose with and without that information, which is quantified as the Value of Information in Section 4.3.

Intel and NVIDIA differ in where the strategy reallocation occurs along the visibility ladder. NVIDIA's main shift toward A1 already happens between S1 and S2 (9 of 37 suppliers move to A1 at S2), driven by the concentration of its upstream network in a small number of high-arc-fragility East Asian fabrication clusters; the additional information gained at S3 sharpens the picture further (22 of 37 on A1) and confirms the qualitative posture. Intel exhibits the opposite pattern. At S2, only 7 of 80 Intel tier-1 suppliers move off A6, because Intel's wider, more dispersed supplier base produces a more diffuse tier-2 cascade signal that, for most arcs, falls below the A6 to A1 strategy-switching threshold. The decisive shift for Intel happens between S2 and S3, when 39 suppliers move to A1, indicating that tier-3 visibility is the binding information for Intel's risk-allocation calculus.

CPG firms display an intermediate pattern. Both Mondelez and Danone show partial reallocation at S2 (with hybrid A4 selections appearing for Danone) and a more complete shift to A1 at S3. The smaller absolute counts simply reflect smaller tier-1 portfolios; in proportional terms, Danone's reallocation under S3 (15/26 = 58% on A1) is comparable to NVIDIA's (22/37 = 59%).

4.3 VoI Estimates

Table 7 reports the VoI estimates across all four companies with safety stock levels expressed both as a percentage savings relative to the prior scenario's policy cost and as an absolute dollar figure. Results are averaged across 100 inventory-noise Monte Carlo iterations, with 95% confidence intervals.

Note that Table 7 reports VoI under a portfolio-level UNIFORM strategy assumption, in which a single strategy is selected across all tier-1 suppliers within each visibility scenario. This differs from Table 6, which shows the supplier-by-supplier optimal allocation. The uniform assumption is a deliberate simplification that ensures conservative VoI estimates and aligns with the most common executive decision-frame (one resilience policy per supplier portfolio); the heterogeneous case is discussed as a future-work extension in Section 5.2

Company	SS	VoI (S2) %	VoI (S2) $M	VoI (S3) %	VoI (S3) $M
Mondelez	10%	2.20% ± 0.00%	$6.1M	0.00% ± 0.00%	$0.0M
Mondelez	20%	2.62% ± 0.01%	$6.1M	0.00% ± 0.00%	$0.0M
Mondelez	30%	2.65% ± 0.05%	$5.0M	0.08% ± 0.16%	$0.2M
Mondelez	40%	1.41% ± 0.09%	$2.1M	7.87% ± 0.27%	$16.7M
Danone	10%	2.15% ± 0.00%	$5.0M	0.00% ± 0.00%	$0.0M
Danone	20%	2.57% ± 0.01%	$5.0M	0.00% ± 0.00%	$0.0M
Danone	30%	2.60% ± 0.04%	$4.1M	0.00% ± 0.00%	$0.0M
Danone	40%	1.39% ± 0.09%	$1.7M	13.54% ± 0.45%	$31.4M

NVIDIA	10%	3.04% ± 0.00%	$17.7M	0.00% ± 0.00%	$0.0M
NVIDIA	20%	4.25% ± 0.02%	$17.7M	0.00% ± 0.00%	$0.0M
NVIDIA	30%	4.71% ± 0.09%	$12.3M	0.00% ± 0.00%	$0.0M
NVIDIA	40%	0.15% ± 0.05%	$0.2M	11.88% ± 0.77%	$31.1M
Intel	10%	0.00% ± 0.00%	$0.0M	6.90% ± 0.01%	$142.1M
Intel	20%	0.00% ± 0.00%	$0.0M	9.62% ± 0.03%	$142.1M
Intel	30%	0.00% ± 0.00%	$0.0M	13.44% ± 0.11%	$122.8M
Intel	40%	0.00% ± 0.00%	$0.0M	9.67% ± 0.22%	$46.3M

Table 7. Value of Information by Company, Visibility Upgrade, and Safety Stock Level

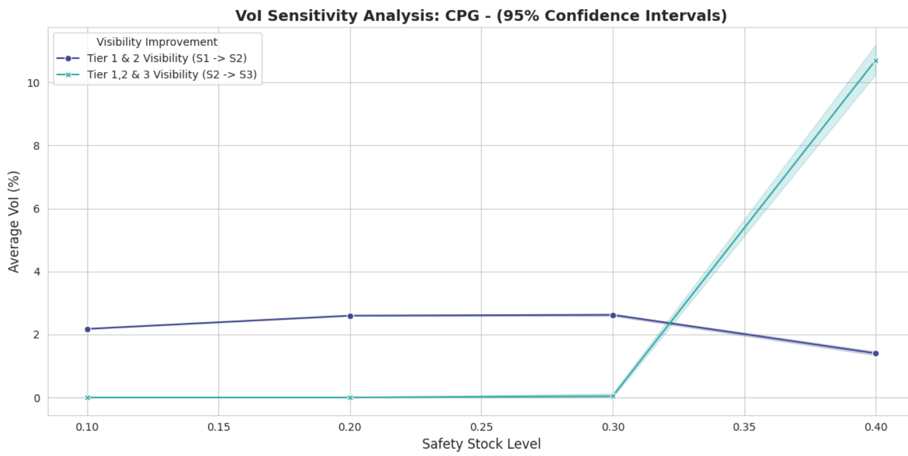

Figure 1. CPG Industry-Average VoI by Visibility Upgrade and Safety Stock Level (Standard Severity Profile). Averaged across Mondelez and Danone; shaded bands denote 95% confidence intervals across 100 inventory-noise iterations.

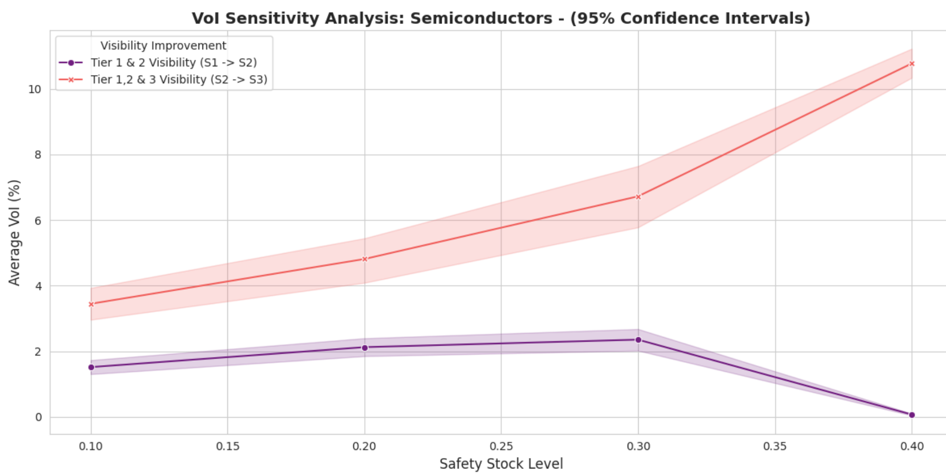

Figure 2. Semiconductor Industry-Average VoI by Visibility Upgrade and Safety Stock Level (Rigid Severity Profile). Averaged across NVIDIA and Intel; shaded bands denote 95% confidence intervals.

Four observations follow from Table 7 and Figures 1–2. First, both CPG firms exhibit a similar profile: VoI(S2) climbs from approximately 2.2% at 10% safety stock to 2.6%–2.7% at 30% safety stock, then declines at 40% safety stock as base inventory absorbs more of the moderate disruption losses that S2 visibility helps to anticipate. In dollar terms, the CPG VoI(S2) is in the $4.1M–$6.1M range across the 10%–30% safety stock band and falls to $1.7M–$2.1M at 40% safety stock. VoI(S3) for both CPG firms is effectively zero across 10%–30% safety stock and becomes materially positive only at 40% safety stock, where it reaches $16.7M for Mondelez and $31.4M for Danone.

Second, Intel exhibits the inverse CPG pattern: VoI(S2) is 0% across all safety stock levels because its diffuse tier-2 cascade signal does not concentrate enough exposure in the tier-1 portfolio to alter the cost-minimizing uniform policy. Intel's value is concentrated entirely at S3, where VoI declines monotonically as safety stock rises: $142.1M at 10% and 20% safety stock, $122.8M at 30%, and

$46.3M at 40%. The decline reflects the same buffer-absorption mechanic that operates in the CPG firms, but compressed into the tier-3 segment of the visibility ladder: as higher base inventory absorbs an increasing share of the cascade losses that S3 information reveals, the marginal value of acquiring tier-3 mapping declines. Even at the high end of the safety stock band, however, Intel's tier-3 VoI remains the largest absolute dollar value in the study, reflecting both network scale and the 55% semiconductor gross margin applied to disruption losses.

Third, across 10%–30% safety stock, NVIDIA captures essentially all of its VoI at tier-2 — $12.3M to $17.7M annually, with no incremental value from tier-3 visibility. At 40% safety stock, however, NVIDIA's profile inverts in the same way the CPG firms do: VoI(S2) collapses to $0.2M while VoI(S3) jumps to $31.1M. The mechanic is the same as for the CPG firms: at high base inventory, the operational disruptions that S2 information helps absorb are already covered by the buffer, so the marginal information gain at S2 is small. The residual exposure consists of low-probability, high-impact cascade events whose existence only becomes visible at S3.

Fourth, the zero VoI(S3) values observed for Mondelez, Danone, and NVIDIA at lower safety stock levels (10%–30%) represent a configuration-specific equilibrium outcome. In these cases, the policy selected under S2 information remains cost-minimizing when evaluated against the realized S3 risk environment: while S3 visibility sharpens supplier-level probability estimates, it does not shift expected losses sufficiently to cross the strategy-switching threshold given the network topology and inventory position in question. These near-zero values are therefore better interpreted as evidence of diminishing returns to visibility at moderate inventory positions for these network topologies, rather than as evidence that tier-3 mapping has no value in general. Intel's results demonstrate directly that the same model, applied to a broader and more dispersed network, produces $46M–$142M in annual VoI at tier-3 across the same safety stock band. The appropriate managerial inference is that tier-3 visibility should be evaluated on a network-by-network and inventory-position-by-inventory-position basis, with priority given to firms whose topology defers decisive risk concentration to deeper tiers.

4.4 Cross-Company Comparison and Industry-Level Implications

The differences between NVIDIA and Intel are best explained by network topology rather than industry. NVIDIA's network is highly concentrated in East Asian fabrication clusters, making cascade risk visible and decision-relevant at tier-2 across most safety stock positions. Intel's 80 tier-1 suppliers across 20 countries create a diffuse upstream signal, where no single region dominates exposure; as a result, meaningful decision value only emerges at tier-3. CPG firms fall in between, with tier-2 visibility generating value at moderate inventory levels and tier-3 becoming relevant only under high safety stock.

Intel's VoI profile is especially pronounced. With VoI(S3) of $122M-$142M at low safety stock, it reflects both network scale and high margin exposure to disruption. More broadly, larger and more complex networks generate greater information value because limited visibility amplifies decision gaps across many critical supplier relationships.

For CPG firms holding 10–30% safety stock, tier-2 mapping yields the highest returns ($3.4M–$5.0M annually). Semiconductor firms, however, require tier-3 visibility to unlock full VoI, with significantly higher gains than in CPG. Firms benefit most when visibility aligns with where risk concentration exists.

Three implications follow. First, optimal visibility depth is topology-dependent, not industry-based; firms should conduct a network audit rather than rely on benchmarks. Second, diminishing returns occur beyond the tier where risk becomes decision-relevant, though this threshold varies by firm. Third, visibility only creates value if it drives action: without adapting resilience strategies, additional mapping yields zero VoI. Realizing value therefore depends not just on data, but on the organizational capability to translate visibility into decisions through tools, governance, and dynamic resource allocation.

5. Recommendations

Across all four focal companies, expanding from blind (S1) to full (S3) visibility raises the average tier-1 cascade-adjusted failure probability by 37–51 percentage points relative to the S1 baseline and reshapes the cost-minimizing resilience strategy from reactive expediting (A6) to proactive buffering (A1) for a substantial share of suppliers. The dollar savings from this reallocation range from $1.7M per year (Danone S2 at 40% safety stock) to $142.1M per year (Intel S3, 10%-20% safety stock). Crucially, the visibility tier at which most of the value lies is not industry-determined but topology-determined: NVIDIA captures essentially all value at tier-2, Intel at tier-3, and the two CPG firms split between the two depending on inventory posture.

5.1 Management Recommendations

Mondelez & Danone: Prioritize tier-2 mapping at current inventory positions and prepare a tier-3 case for any inventory expansion. Both CPG firms generate $4.1M–$6.1M per year in decision-support savings from tier-2 visibility at safety stock levels representative of their current operating positions (10%–30% of balance sheet inventory). At those levels, tier-3 mapping yields near-zero marginal value and should be deprioritized. If either firm raises strategic safety stock to 40 the calculus inverts: tier-3 visibility produces $16.5M (Mondelez) and $31.4M (Danone) per year. Tier-2 mapping should therefore be funded immediately as a no-regrets move, while a conditional tier-3 case is prepared and triggered if company strategy shifts toward higher buffer inventory. Priority targets surface from the model: for Danone, the seven tier-1 suppliers with European reaching P(S3) ≥ 99%; for Mondelez, the four tier-1 suppliers anchored by upstream exposure in India, Mexico, and Poland.

NVIDIA: Tier-2 visibility is the binding investment; defer tier-3 outside of strategically critical relationships. NVIDIA's network exhibits the strongest tier-2 cascade signal in the study, with tier-1 suppliers such as Samsung, TSMC, Hon Hai Precision, and Micron showing P(S2) above 60% despite intrinsic P(S1) values of 3%–5%. Tier-2 mapping alone captures approximately $12.3M-$17.7M per year in decision-support value across 10%-30% safety stock, with the value declining at heavier inventory positions as base buffers absorb increasing shares of the disruption. At 40% safety stock the picture flips: tier-2 VoI collapses to $0.2M while tier-3 VoI rises to $31.1M, reflecting the same buffer-driven inversion observed in the CPG firms. The recommended action is to fund a tier-2 mapping program; the tier-3 case should be revisited if NVIDIA's strategic safety stock rises toward 40%, at which point an additional ~$31M per year becomes capturable.

Intel: Fund tier-3 mapping; tier-2 visibility alone is insufficient. Intel is the only focal company where tier-2 visibility produces zero VoI. Its wide, geographically dispersed tier-1 base generates a diffuse tier-2 cascade signal that does not concentrate enough exposure to flip the cost-minimizing portfolio strategy. Tier-3 visibility, by contrast, produces approximately $122M–$142M per year in decision-support value at typical operating inventory positions (10%–30% safety stock), declining to $46M at 40% safety stock as higher inventory absorbs more of the cascade exposure. Even at the high end of the safety stock band, Intel's tier-3 VoI remains the largest absolute dollar VoI in the study, reflecting both Intel's network size and the 55% semiconductor gross margin applied to disruption losses. The investment to prioritize is a tier-3 mapping program covering the 44 suppliers the model identifies as switching from A6 to A1 at S3.

Two strategic implications: First, the aggregate decision-support value across the four focal companies, on the order of $140M-$165M per year at typical inventory positions, is large enough that visibility investment should be evaluated as a board-level resilience initiative rather than a discretionary IT or data spend. Second, the framework reframes visibility from a generic best practice into a priced capability with a measurable return: firms can compare the gross VoI computed by the model against the cost of acquiring tier-n mapping (Equation 7) and approve only those investments where Net VoI is positive — a meaningful upgrade over the prevailing industry baseline of qualitative justification.

5.2 Future Work

Expanding the mitigating strategy set and relaxing the uniform assumption in the VoI calculation. The current model evaluates six predefined resilience strategies (A1–A6) and computes portfolio-level VoI under the assumption that a single uniform strategy is applied across all tier-1 suppliers within a given visibility scenario. The strategy set could be extended to include supplier diversification, geographic relocation, and dual-sourcing actions whose costs and benefits propagate through the network. Equally important, the VoI computation could be re-cast so that the visibility step is evaluated against a heterogeneous, supplier-by-supplier optimal allocation (which the model already produces for Table 6) rather than a uniform portfolio policy.

Cross-industry expansion. The four-company sample was deliberately constructed to span two industries with maximally different supply chain configurations. Extending the analysis to additional industries would test whether the topology-driven (rather than industry-driven) interpretation of the tier-2 vs. tier-3 split holds more broadly and would surface industry-specific resilience-strategy adaptations not represented in the current strategy set. A larger industry sample would also allow the framework to characterize entire industries by their typical visibility-value profile.

Varying parameter assumptions of the model. The framework currently relies on fixed assumptions for several model and industry parameters: a Lognormal disruption-duration distribution, equal weighting (1/3) of the three country-level risk indices in the Global Risk Index, a 50:50 weighting between revenue dependency and supplier size in the arc fragility computation, and the Logit anchor probabilities at the 10th and 90th percentiles of the global risk distribution. A systematic parameter-variation study would test the sensitivity of the VoI estimates to each of these choices. This would convert the current point estimates into ranges and equip decision-makers with a clearer view of which parameter choices the recommendations are most sensitive to.

5.3 Concluding remarks

For supply chain executives at the focal companies, the practical takeaway is concrete: visibility investment can be sized, prioritized, and approved using the same financial logic that already governs inventory, sourcing, and capacity decisions. As detailed in Section 5.1, the model produces a per-company dollar value at stake and identifies the specific subset of suppliers on which to concentrate the mapping effort. These figures represent the avoidable cost of adopting a reactive expediting posture when the true cascade-adjusted risk warrants proactive buffering, and they reframe the visibility conversation from a qualitative resilience priority into a quantified contribution to the productivity, margin-protection, and capital-allocation agendas that already shape strategic decision-making.

The broader value of the framework extends well beyond the four focal companies. For other firms in adjacent industries with comparable supply chain topology, the model provides a transferable method for sizing the value at stake in visibility programs that have, until now, been justified primarily on qualitative grounds. For the MIT Digital Supply Chain Transformation Lab and its industrial partners, it offers a quantitative anchor for ongoing research into multi-tier transparency, supplier-disclosure incentives, and resilience-investment governance. More broadly, it reframes the visibility question from a binary one (mapped or not mapped) into a continuous one (mapped to which depth, against which net financial benefit). And for the discipline of supply chain risk management as a whole, it demonstrates that information itself can be priced, ranked, and budgeted alongside the inventory, sourcing, and capacity capabilities that have traditionally dominated the resilience conversation.

References

- Alicke, K., Barriball, E., & Trautwein, V. (2022). Taking the pulse of shifting supply chains. McKinsey Operations.
- Bakshi, N., & Mohan, S. (2024). Information dependency in mitigating disruption cascades. Manufacturing & Service Operations Management, 26(6), 2050–2066.
- Basole, R. C., & Bellamy, M. A. (2014). Supply network structure, visibility, and risk diffusion: A computational approach. Decision Sciences, 45(4), 753–789.
- Bowen, F., & Siegler, J. (2024). The role of visibility in supply chain resiliency: Applying the Nexus supplier index to unveil hidden critical suppliers in deep supply networks. Decision Support Systems, 176, 114063.
- Candogan, O., & Gurkan, H. (2025). The value of information design in supply chain management. Management Science, 71(8), 6545–6558.
- Dolgui, A., & Ivanov, D. (2021). Ripple effect and supply chain disruption management: New trends and research directions. International Journal of Production Research, 59(1), 102–109.
- Dolgui, A., Ivanov, D., & Sokolov, B. (2018). Ripple effect in the supply chain: An analysis and recent literature. International Journal of Production Research, 56(1–2), 414–430.
- Ebel, E., Kasah, T., Koerber, B., Mohr, N., Müllerklein, D., & Roggendorf, M. (2021, July 7). Overcoming barriers to multitier supplier collaboration. McKinsey & Company. https://www.mckinsey.com/capabilities/operations/our-insights/overcoming-barriers-to-multitier-supplier-collaboration
- Fishman, G. S. (1996). Monte Carlo: Concepts, algorithms, and applications. Springer.
- Hosseini, S., & Barker, K. (2016). A Bayesian network model for resilience-based supplier selection. International Journal of Production Economics, 180, 68–87.
- Hosseini, S., Ivanov, D., & Dolgui, A. (2019). Review of quantitative methods for supply chain resilience analysis. Transportation Research Part E: Logistics and Transportation Review, 125, 285–307.
- Ivanov, D. (2020). Predicting the impacts of epidemic outbreaks on global supply chains: A simulation-based analysis on the coronavirus outbreak (COVID-19/SARS-CoV-2) case. Transportation Research Part E: Logistics and Transportation Review, 136, 101922.
- Lou, G., Guo, Y., Lai, Z., Ma, H., & Tu, X. (2024). Optimal resilience strategy for manufacturers to deal with supply disruptions: Investment in supply stability versus dual sourcing. Computers & Industrial Engineering, 190, 110030.
- McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In P. Zarembka (Ed.), Frontiers in econometrics (pp. 105–142). Academic Press.
- McKinsey & Company. (2020). Risk, resilience, and rebalancing in global value chains. McKinsey Global Institute.
- McKinsey & Company. (2021). How COVID-19 is reshaping supply chains. McKinsey Operations.
- McKinsey & Company. (2022). Future-proofing the supply chain. McKinsey Operations Practice.
- McKinsey & Company. (2024). McKinsey Global Supply Chain Leader Survey 2024.
- Ribeiro, J. P., & Barbosa-Povoa, A. (2018). Supply chain resilience: Definitions and quantitative modelling approaches — A literature review. Computers & Industrial Engineering, 115, 109–122.
- Schmitt, A. J., & Singh, M. (2012). A quantitative analysis of disruption risk in a multi-echelon supply chain. International Journal of Production Economics, 139(1), 22–32.
- Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., … & Zhang, D. (2015). Identifying risks and mitigating disruptions in the automotive supply chain. Interfaces, 45(5), 375–390.
- Snyder, L. V., Atan, Z., Peng, P., Rong, Y., Schmitt, A. J., & Sinsoysal, B. (2016). OR/MS models for supply chain disruptions: A review. IISE Transactions, 48(2), 89–109.
- Taghizadeh, E., Venkatachalam, S., & Chinnam, R. B. (2021). Impact of deep-tier visibility on effective resilience assessment of supply networks. International Journal of Production Economics, 241, 108254.
- Tang, C. S. (2006). Robust strategies for mitigating supply chain disruptions. International Journal of Logistics: Research and Applications, 9(1), 33–45.
- Train, K. E. (2009). Discrete choice methods with simulation (2nd ed.). Cambridge University Press.
- Viet, N. Q., Behdani, B., & Bloemhof, J. (2018). The value of information in supply chain decisions: A review of the literature and research agenda. Transportation Research Part E: Logistics and Transportation Review, 120, 68–91.