INTERNATIONAL JOURNAL OF BUSINESS AND ECONOMICS INSIGHTS

Volume: 16; **Issue:** 1
Pages: 01– 42
Accepted: 17 December 2025
Published: 02 January 2026

AI-Enabled IT and Supply Chain Resilience for Anticipating and Mitigating Critical Medical-Supply Disruptions in U.S. Healthcare Systems

H C H Raiyan Ahmed[1];

[1]. *Graduate Student, M.S. in Information Technology Management, Campbellsville University, Kentucky, USA; Bachelor of Business Administration in Marketing, North South University, Dhaka, Bangladesh; Email: hch.raiyan.ahmed@gmail.com*

Doi: 10.63125/dhdew369
Received: 09 October 2025; **Revised:** 11 November 2025; **Accepted:** 09 December 2025; **Published:** 02 January 2026;

Abstract

Critical medical supplies in the United States are disrupted by hurricanes that close a dominant plant, by quality shutdowns of sterile-injectable lines, by demand surges, by logistics delays and, increasingly, by cyber incidents that remove a health system's ordering and visibility systems. This article asks how much an AI-enabled early-warning and response layer, combined with information technology (IT) resilience, can reduce the harm of such disruptions for a regional health system, and at what cost. A transparent synthetic testbed of eight hospitals (2,150 beds), six critical items, a distributor with allocation rules, secondary sources and a spot market is subjected to seven disruption classes drawn from documented U.S. events. Five strategies are compared on identical episodes: independent reorder points (S0), system-wide visibility with rule-based alerts (S1), AI early warning with a gradient-boosted demand forecaster and supply-risk scorer calibrated to a false-alert budget (S2), S2 with max-min allocation, transshipment and tiered conservation (S3), and S3 with an IT resilience runbook (S4). Over 20 seeds, unmet demand for the affected items under a 90-day manufacturing outage fell from 23.7% (S0) to 9.9% (S2) and 0.5% (S3); under a four-fold demand surge from 38.3% to 26.3% and 6.8%; under a quality shutdown from 20.0% to 0.3% and 0.0%. Early warning preceded the onset in about a third of outage episodes and half of quality shutdowns; the gains came from acting on the first signals within the distributor's order cap, from pre-qualified secondary sources and from conservation. A 28-day cyber outage alone cost little, but the same outage during a supply disruption erased most of the AI layer's advantage (S2: 12.9%, S3: 2.8%) unless the runbook kept the system working (S4: 0.8%). One-factor sweeps and a 200-draw Monte Carlo analysis rank the levers: order caps, secondary-source lag, conservation effectiveness, outage length and manual throughput. A reference architecture, response workflow and phased roadmap mapped to U.S. policy and guidance are provided. All data are synthetic; the numbers describe the testbed, not any real health system.

Keywords

Healthcare Supply Chain; Resilience; Drug Shortages; Early Warning; Machine Learning; Forecasting; Allocation; Ransomware; IT Continuity; Simulation

Introduction

The U.S. health system depends on a small number of items that cannot be substituted at short notice and whose production is concentrated in a few plants. When Hurricane Maria damaged manufacturing in Puerto Rico in 2017, hospitals across the country ran short of small-volume intravenous fluids and changed how they administered drugs (Mazer-Amirshahi & Fox, 2018; Sacks et al., 2018). When Hurricane Helene flooded a plant in North Carolina in late September 2024, the manufacturer placed intravenous solutions on allocation and reported that production had returned to pre-hurricane levels only in February 2025 (Baxter, 2025). Sterile-injectable shortages have been a chronic feature of the market for more than a decade, driven by quality failures in concentrated manufacturing and by economics that leave little margin for redundancy (Woodcock & Wosinska, 2013; Fox et al., 2014; Yurukoglu et al., 2017; U.S. Food and Drug Administration, 2020), and they harm patients: a national shortage of norepinephrine was associated with higher in-hospital mortality among patients with septic shock (Vail et al., 2017), and scoping reviews find delayed care, substitutions and errors across many shortages (Phuong et al., 2019; Tucker et al., 2020a). The pandemic of 2020 added demand surges that outran any inventory and showed how little visibility health systems had of their own supply chains (Ranney et al., 2020; Finkenstadt & Handfield, 2021; Handfield et al., 2020). By mid-2024 the number of active national drug shortages tracked by the American Society of Health-System Pharmacists was at or near its highest level since tracking began (American Society of Health-System Pharmacists, 2024), and the National Academies concluded that the medical product supply chain needed resilience built into it rather than assumed (National Academies of Sciences, Engineering, and Medicine, 2022).

A second kind of disruption has grown alongside the first. Ransomware attacks on U.S. hospitals doubled between 2016 and 2021 (Neprash et al., 2022); they disrupt not only the attacked hospital but the emergency departments around it (Dameff et al., 2023), and the 2024 attack on a shared claims and pharmacy intermediary showed that a single intermediary can remove a function from much of the sector at once (U.S. Department of Health and Human Services, Office for Civil Rights, 2024). For supply chain operations the relevant harm is the loss of the ordering, inventory and forecasting systems that modern practice depends on. A health system that has invested in analytics may find that its most capable processes are the first to fail, while a paper-based competitor keeps phoning in orders.

Three bodies of work address these problems separately. Supply chain resilience research defines the capabilities (visibility, flexibility, redundancy, collaboration) and the mitigation and contingency levers (Christopher & Peck, 2004; Sheffi & Rice, 2005; Tomlin, 2006; Pettit et al., 2010; Hohenstein et al., 2015; Kamalahmadi & Parast, 2016), and quantitative studies show how disruptions propagate and how long recovery takes (Dolgui et al., 2018; Hosseini et al., 2019; Ivanov, 2020). Healthcare operations research has produced inventory policies for hospitals and shortage-aware models (Nicholson et al., 2004; Bijvank & Vis, 2012; Saedi et al., 2016; Volland et al., 2017). Forecasting and prescriptive analytics have improved substantially, with gradient-boosted and probabilistic methods winning large competitions and ensembles forecasting epidemics with usable skill (Makridakis et al., 2022; Reich et al., 2019; Cramer et al., 2022; Bertsimas & Kallus, 2020). What is missing is a quantified account of what these capabilities deliver together for a single health system facing the disruptions that U.S. hospitals have actually experienced, including the case in which the IT layer that carries the analytics is itself disrupted.

This article provides that account with a transparent simulation. It poses four research questions. RQ1: how much earlier can an AI-enabled early-warning layer anticipate a disruption than conventional practice, and at what false-alert cost? RQ2: which mitigation levers (pre-emptive building within the distributor's rules, secondary sourcing, allocation and pooling across hospitals, tiered clinical conservation) reduce unmet demand most, under which disruption classes, and what do they cost? RQ3: how does IT resilience interact with supply resilience; in particular, what is lost when a cyber outage removes visibility and ordering during a supply disruption, and what does an offline runbook recover? RQ4: how robust are the answers to the uncertain factors a health system cannot control, such as the severity and duration of the disruption, the lead of the available signals, forecast quality and the length of an outage?

The contributions are a reference architecture that places AI services inside a resilient IT design (Figure 1), a response workflow (Figure 2), a synthetic but documented testbed of a regional health system with a distributor, secondary sources and seven disruption classes (Figures 3 and 4), a comparison of five strategies on identical episodes with one-factor sweeps and a Monte Carlo analysis, and a roadmap mapped to U.S. policy and guidance. The study is explicit about what it is not: it does not use data from any health system, and its numbers are properties of the testbed. Its value is in the structure of the results, which levers matter, in which order and with what interactions, and in the honest accounting of the costs of each.

Background and related work

Supply chain resilience: concepts and levers

Resilience entered supply chain management from ecology and engineering, where it denotes the capacity of a system to absorb disturbance and reorganise while retaining its function (Holling, 1973; Bruneau et al., 2003). The supply chain literature defines it as the ability to prepare for, respond to and recover from disruptions, and identifies the capabilities that produce it (Asheke & Maksudul, 2022; Asheke & Badrul, 2022): visibility of inventory and demand across the chain, flexibility in sourcing and distribution, redundancy in stock and capacity, collaboration with partners and a culture that treats risk as everyone's job (Christopher & Peck, 2004; Sheffi & Rice, 2005; Ponomarov & Holcomb, 2009; Pettit et al., 2010; Pettit et al., 2013). Reviews converge on the same set and on the observation that resilience is bought (Akter & Sumaiya, 2022; Imran & Jannatul, 2022): redundancy costs inventory and capacity, flexibility costs contracts and qualification, and visibility costs systems and data (Hohenstein et al., 2015; Kamalahmadi & Parast, 2016). Early work also separated mitigation, which is in place before a disruption (inventory, dual sourcing) (Maksudul & Islam, 2022; Morshedul & Islam, 2022), from contingency, which is activated after it (rerouting, demand management), and showed that the right mix depends on the probability and duration of the disruption (Kleindorfer & Saad, 2005; Tomlin, 2006; Tang, 2006; Chopra & Sodhi, 2004). Craighead et al. (2007) tied severity to network density, complexity and node criticality, and Chopra and Sodhi Chopra and Sodhi (2014) argued for segmenting and regionalising supply chains so that a single event cannot take out a product. Measures such as time to recover and time to survive (Stabak & Shahid, 2022; Feroz & Pervaz, 2022), developed for an automaker, give a vocabulary for the exposure of each node that health systems can use (Simchi-Levi et al., 2014; Simchi-Levi et al., 2015).

The quantitative literature models disruptions explicitly. Snyder et al. Snyder et al. (2016) review OR/MS models of supply chain disruptions; Schmitt and Singh Schmitt and Singh (2012) use simulation of a multi-echelon chain to show that inventory placement and backup capacity have different values for different disruption profiles; Lücker et al. Lücker et al. (2019) separate the roles of inventory and reserve capacity. The ripple effect, the propagation of a local disruption through tiers, has been studied with simulation and digital twins (Dolgui et al., 2018; Hosseini et al., 2019; Ivanov & Dolgui, 2021), and the pandemic prompted a reframing from resilience to viability, the survival of intertwined supply networks under long, systemic disturbances (Ivanov, 2020; Ivanov & Dolgui, 2020; Craighead et al., 2020; Sodhi & Tang, 2021). The present study follows this stream in method (daily simulation, explicit propagation, time to recover) and differs in object (a buyer, not a manufacturer) and in its attention to the signals that precede a disruption.

Two classic results from inventory theory frame the pooling lever. Eppen Eppen (1979) showed that centralising stock across locations reduces the safety stock needed for a given service level; lateral transshipment between locations captures part of that benefit without physical centralisation (Herer et al., 2006; Paterson et al., 2011). Multi-echelon policies (Clark & Scarf, 1960; Silver et al., 2016) and the bullwhip effect (Lee et al., 1997; Sterman, 1989) explain why a system of independently ordering hospitals amplifies a disruption: each site reacts to its own shortfall by ordering more, and the distributor's allocation rules then punish the system as a whole. Supplier allocation and hoarding under shortage are exactly the dynamics observed in 2020 (Finkenstadt & Handfield, 2021; Cohen & Rodgers, 2020).

Healthcare supply chains and the U.S. shortage problem

Health care supply chains differ from retail or automotive chains in their fragmentation, in the role of group purchasing organisations and distributors, in clinical preferences that limit substitution and in the fact that the consumer of the product is not the buyer (Burns, 2002; Schneller & Smeltzer, 2006; Hu et al., 2012; Abdulsalam et al., 2015; Chen et al., 2013). Hospital materials management has nonetheless adopted standard tools, and point-of-use inventory, two-bin systems and outsourced replenishment reduce cost and stockouts (Nicholson et al., 2004; Bijvank & Vis, 2012; Rosales et al., 2015; Kelle et al., 2012; Volland et al., 2017); shortage-aware inventory models show how much stock protects against an interruption of a given length (Saedi et al., 2016), and policy models examine incentives for resilient manufacturing and reimbursement reform (Jia & Zhao, 2017; Tucker et al., 2020b; Yurukoglu et al., 2017).

The U.S. drug-shortage problem has been documented from several angles. Surveys report substitutions, delays, rationing and the labour of managing shortages (Kaakeh et al., 2011; McLaughlin et al., 2013); the FDA's root-cause analysis attributes most shortages to quality problems in concentrated manufacturing and a market that does not reward reliability (U.S. Food and Drug Administration, 2020; Woodcock & Wosinska, 2013); a scoping review finds the problem persistent over two decades (Tucker et al., 2020a); and the Government Accountability Office associated shortages with low prices, few suppliers and quality failures (U.S. Government Accountability Office, 2016). Professional guidelines set out what a health system should do: a shortage committee, a critical-item list, inventory and allocation rules, conservation and substitution protocols, and communication (Fox & McLaughlin, 2018). Crisis standards of care frameworks and allocation ethics supply the governance for the hardest decisions, when conservation becomes rationing (Institute of Medicine, 2012; Persad et al., 2009; Emanuel et al., 2020; Hick et al., 2020). The pandemic literature adds the supply chain side: the invisibility of personal protective equipment stocks, the failure of the national stockpile as designed, and proposals for a commons with shared visibility (Ranney et al., 2020; Finkenstadt & Handfield, 2021; Handfield et al., 2020; Mirchandani, 2020). Empirical studies of resilience in healthcare supply chains during the pandemic identify visibility, supplier relationships, flexibility and digital capability as the enablers that mattered (Spieske et al., 2022; Zamiela et al., 2022). Events outside drugs show the same structure: the 2022 iodinated contrast shortage, caused by a lockdown at a single plant, was managed by conservation and deferred imaging (Cavallo & Pahade, 2022), and the 2008 heparin contamination turned one supplier's quality failure into a safety crisis (Kishimoto et al., 2008).

These recurring disruptions highlight the structural vulnerabilities embedded in healthcare supply chains, particularly the interaction between supplier concentration, limited production flexibility, inventory constraints, and fragmented information systems (Islam & Md, 2023; Feroz & Shabbir, 2022). Unlike conventional commercial supply networks, healthcare supply chains operate under conditions in which product unavailability can directly compromise patient safety, delay essential medical procedures, and reduce the quality of clinical outcomes (Asheke, 2023; Shahid, 2023). The dependence on specialized manufacturers, stringent quality requirements, and complex regulatory approval processes further limits the ability of healthcare organizations to respond rapidly when established supply arrangements fail (Islam & Ahmed, 2023; Sadia & Mosheur, 2023). In addition, healthcare providers often lack comprehensive visibility into upstream production activities, raw material availability, manufacturing capacity, and distribution bottlenecks, making it difficult to anticipate disruptions before they affect hospital operations (Adnan, 2024; Islam, 2024). These limitations are intensified by procurement arrangements that prioritize immediate acquisition costs over long-term supply reliability, potentially discouraging investment in redundant production capacity and alternative sourcing arrangements (Ashik, 2024; Kynatun & Saba, 2024; Maksudul, 2024). Consequently, shortages frequently develop through a combination of operational inefficiencies, market incentives, and coordination failures rather than through isolated manufacturing incidents (Shahid, 2024; Islam, 2024; Sumaiya, 2024). Addressing these interconnected challenges requires a broader understanding of supply chain resilience that incorporates supplier reliability, demand uncertainty, inventory responsiveness, procurement governance, and the clinical consequences of product unavailability (Ashik, 2025; Kynatun, 2025; Taher et al., 2024).

The persistent nature of healthcare shortages also demonstrates the importance of transitioning from reactive shortage management toward proactive, intelligence-driven supply chain planning and coordinated risk mitigation (Mostari, 2025; Masrul, 2025; Maksudul, 2025). Healthcare organizations increasingly require integrated decision-support capabilities that combine real-time inventory information, historical consumption patterns, supplier performance indicators, and demand forecasts to identify emerging vulnerabilities and support timely intervention (Shahriar, 2025; Shahid, 2025; Islam, 2025). Advanced analytical approaches, including predictive modeling, scenario analysis, and optimization-based inventory allocation, offer opportunities to improve the identification of shortage risks while maintaining appropriate service levels across healthcare facilities (Sadia, 2025; Shabbir, 2025; Sadmeen, 2025). However, technological improvements alone cannot resolve the underlying structural challenges without corresponding reforms in purchasing practices, supplier diversification, contractual accountability, and interorganizational collaboration. Effective resilience strategies should therefore combine strategic safety stocks for clinically essential products, transparent supplier-risk assessments, alternative sourcing arrangements, and clearly defined allocation protocols during periods of constrained availability. Particular attention must be given to the unequal consequences of shortages across hospitals, geographic regions, and patient populations, as resource-constrained institutions may have fewer opportunities to secure alternative supplies. A comprehensive policy and operational framework must consequently balance cost efficiency with supply continuity, equitable distribution, and patient safety. By integrating technological capabilities with institutional coordination and risk-sensitive procurement decisions, the U.S. healthcare system can strengthen its capacity to anticipate disruptions, reduce shortage duration, and maintain the availability of essential medical products under both routine and crisis conditions.

Forecasting, early warning and prescriptive analytics

Demand forecasting for hospital supplies is a short-horizon, intermittent, hierarchical problem. The forecasting literature supplies the methods: exponential smoothing and regression with seasonality (Hyndman & Athanasopoulos, 2021), estimators for intermittent demand (Croston, 1972; Syntetos & Boylan, 2005), accuracy measures that behave across scales (Hyndman & Koehler, 2006), quantile regression for prediction intervals (Koenker & Bassett, 1978) and gradient-boosted trees, which dominated the M5 retail competition with calibrated quantiles (Makridakis et al., 2022; Ke et al., 2017; Friedman, 2001). For surges the relevant signals are epidemiological: collaborative ensembles forecast influenza and COVID-19 with usable skill at short horizons (Reich et al., 2019; Cramer et al., 2022), emergency department volumes have been forecast from calendar and census variables (Jones et al., 2008), and the cautionary tale of Google Flu Trends warns against untested proxies (Lazer et al., 2014). Forecasts earn their keep only through decisions: the prescriptive-analytics literature ties predictions to the decision they inform and evaluates them by the decision's outcome (Bertsimas & Kallus, 2020; Ban & Rudin, 2019; Elmachtoub & Grigas, 2022), which is the evaluation used here.

Early warning of supply disruptions is less developed. Reviews list the signals (supplier distress, quality notices, weather, geopolitics) and the AI methods applied to them (Baryannis et al., 2019; Ivanov et al., 2019), and a Canadian study predicted drug shortages weeks ahead from pharmacy purchasing data (Pall et al., 2023). The signals available to a U.S. health system are the distributor's fill rate and acknowledgements, advance ship notices, allocation letters, FDA and ASHP postings, manufacturer notices and public hazard feeds; the postings arrive late by construction, because a shortage is posted once it exists (U.S. Congress, 2012; U.S. Food and Drug Administration, 2020). Change-detection statistics (Page, 1954) and calibrated probabilities (Gneiting & Raftery, 2007; Niculescu-Mizil & Caruana, 2005) turn noisy signals into alerts with a known false-alert rate, which decides whether alerts are acted on (Axelsson, 2000).

Healthcare cybersecurity and IT continuity

Ransomware is the dominant cyber risk to health care delivery. Attacks on U.S. hospitals, clinics and other delivery organisations rose from 43 in 2016 to 91 in 2021, with growing data exposure and operational disruption (Neprash et al., 2022); the WannaCry attack on the English NHS cancelled appointments across many trusts (Ghafur et al., 2019); and a ransomware attack on one health system increased volumes and waiting times in the emergency departments of its neighbours (Dameff et al., 2023). Systematic reviews describe the threats and the organisational weaknesses behind them (Kruse

et al., 2017; Jalali & Kaiser, 2018; Oz et al., 2022). The 2024 attack on Change Healthcare removed claims processing and pharmacy transactions for much of the sector (U.S. Department of Health and Human Services, Office for Civil Rights, 2024), and the joint advisory on the group responsible was updated for the health sector (Federal Bureau of Investigation et al., 2023). IT continuity practice is mature: contingency planning guidance (Swanson et al., 2010), business continuity standards (International Organization for Standardization, 2019), cyber-resilient system engineering (Ross et al., 2021), zero-trust architecture (Rose et al., 2020), incident handling (Cichonski et al., 2012) and supply chain risk management for information systems (Boyens et al., 2022) all apply, and the health sector has its own: the HIPAA Security Rule (U.S. Department of Health and Human Services, 2013), the 405(d) health industry cybersecurity practices (U.S. Department of Health and Human Services, 2023), the sector-specific cybersecurity performance goals (U.S. Department of Health and Human Services, 2024) and the Cybersecurity Framework (National Institute of Standards and Technology, 2024). Electronic health record downtime has been studied as a patient-safety problem, with recommended contingency practices (Larsen et al., 2018; Sittig et al., 2014). What has not been quantified is the interaction between an IT outage and a supply disruption in the same window, which is where the testbed's compound scenario sits.

U.S. policy context

Federal policy has moved on both fronts. FDASIA (2012) required manufacturers to notify the FDA of discontinuances and interruptions, the Drug Supply Chain Security Act (2013) built traceability for prescription drugs, the unique device identification rule standardised device identifiers, and the CARES Act (2020) extended shortage reporting to devices and required manufacturer risk management plans (U.S. Congress, 2012; U.S. Congress, 2013; U.S. Food and Drug Administration, 2013; U.S. Congress, 2020). Executive Orders 14001 and 14017 directed a sustainable public health supply chain and reviews of critical supply chains, including pharmaceuticals (Executive Office of the President, 2021c; Executive Office of the President, 2021b; The White House, 2021); the Strategic National Stockpile holds federal reserves (Administration for Strategic Preparedness and Response, n.d.); and the National Academies' 2022 report set out resilience measures across manufacturers, distributors, providers and government (National Academies of Sciences, Engineering, and Medicine, 2022). For providers, the CMS emergency preparedness rule requires risk assessment, policies, communication plans and exercises (Centers for Medicare & Medicaid Services, 2016), and the Executive Order on improving the nation's cybersecurity and the sector cybersecurity goals shape the IT side (Executive Office of the President, 2021a; U.S. Department of Health and Human Services, 2024). These instruments are mapped to the architecture in the discussion; they change frequently, and the mapping describes the position at the time of writing rather than a compliance claim.

Despite these significant policy developments, important challenges remain in translating federal regulatory requirements into coordinated, measurable, and operationally effective healthcare supply chain resilience. Existing policies address critical dimensions of supply security, including shortage notification, product traceability, emergency preparedness, strategic stockpiling, and cybersecurity; however, these measures frequently operate through separate institutional responsibilities, reporting mechanisms, and implementation frameworks. Such fragmentation can limit the capacity of federal agencies, manufacturers, distributors, and healthcare providers to develop a unified understanding of emerging supply disruptions and implement timely preventive interventions. Although mandatory reporting requirements improve regulatory awareness of potential shortages, they do not necessarily ensure sufficient manufacturing redundancy, diversified sourcing, transparent inventory information, or equitable allocation of scarce medical products. Similarly, national stockpiling initiatives provide an important emergency response mechanism but require effective coordination with routine procurement systems, regional demand assessments, replenishment strategies, and distribution capabilities to maintain long-term effectiveness. The growing digitalization of healthcare logistics introduces additional opportunities for predictive monitoring and interorganizational information exchange while simultaneously creating new challenges involving data interoperability, information security, and institutional accountability.

Table 1 summarises the strands and the gap this article addresses.

Table 1: Literature strands, what they establish, representative sources and the gap addressed in this article

Strand	What the literature establishes	Representative sources	Gap addressed here
Supply chain resilience	Resilience is the capacity to anticipate, absorb, adapt and recover; it is built from visibility, flexibility, redundancy and collaboration, and it has a cost	Christopher & Peck (2004); Sheffi & Rice (2005); Ponomarov & Holcomb (2009); Pettit et al. (2010, 2013); Hohenstein et al. (2015); Kamalahmadi & Parast (2016)	A quantified comparison of the capabilities for one health system, with the cost of each
Disruption risk and mitigation	Mitigation (inventory, dual sourcing) and contingency (rerouting, demand management) are complementary; disruption severity depends on network design	Kleindorfer & Saad (2005); Tomlin (2006); Tang (2006); Craighead et al. (2007); Chopra & Sodhi (2004, 2014); Snyder et al. (2016)	The same levers evaluated under hospital demand, distributor allocation and clinical conservation
Ripple effect, viability and simulation	Disruptions propagate through tiers; simulation and digital twins quantify time to recover and survive; pandemics change the problem from local to systemic	Dolgui et al. (2018); Hosseini et al. (2019); Ivanov (2020); Ivanov & Dolgui (2020, 2021); Schmitt & Singh (2012); Simchi-Levi et al. (2015)	A daily simulation with explicit signals so that anticipation, not only absorption, can be measured
Healthcare supply chains and shortages	Drug and device shortages are chronic in the U.S., driven by concentrated sterile-injectable manufacturing, quality failures and thin margins; they harm patients	Woodcock & Wosinska (2013); Fox et al. (2014); Vail et al. (2017); Phuong et al. (2019); Tucker et al. (2020a); National Academies of Sciences, Engineering, and Medicine (2022)	The health-system side of the problem: what a buyer can do with early warning, pooling and conservation
Hospital inventory and logistics	Point-of-use inventory, two-bin systems and shortage-aware policies reduce cost and stockouts; pooling and transshipment reduce safety stock	Nicholson et al. (2004); Bijvank & Vis (2012); Kelle et al. (2012); Rosales et al. (2015); Saedi et al. (2016); Volland et al. (2017); Eppen (1979); Paterson et al. (2011)	Pooling and allocation combined with forecasting and risk scoring during disruptions
Forecasting and prescriptive analytics	Gradient-boosted and probabilistic forecasters perform well on retail and epidemic series; predictions must be tied to decisions and scored	Hyndman & Athanasopoulos (2021); Makridakis et al. (2022); Reich et al. (2019); Cramer et al. (2022); Bertsimas & Kallus (2020); Elmachtoub & Grigas (2022); Pall et al. (2023)	Forecasts and risk scores evaluated by the stockouts they prevent and the false alerts they cause
Allocation ethics and crisis standards	Scarce resources should be allocated by explicit, fair rules; conservation and substitution are clinical decisions	Persad et al. (2009); Institute of Medicine (2012); Emanuel et al. (2020); Hick et al. (2020); Fox & McLaughlin (2018)	Conservation tiers and allocation rules as modelled levers with a clinical-burden cost
Healthcare cybersecurity and IT continuity	Ransomware disrupts care delivery and adjacent hospitals; downtime procedures and contingency planning limit harm	Ghafur et al. (2019); Neprash et al. (2022); Dameff et al. (2023); Larsen et al. (2018); Sittig et al. (2014); Swanson et al. (2010)	The interaction between an IT outage and a supply disruption, and the value of an offline runbook

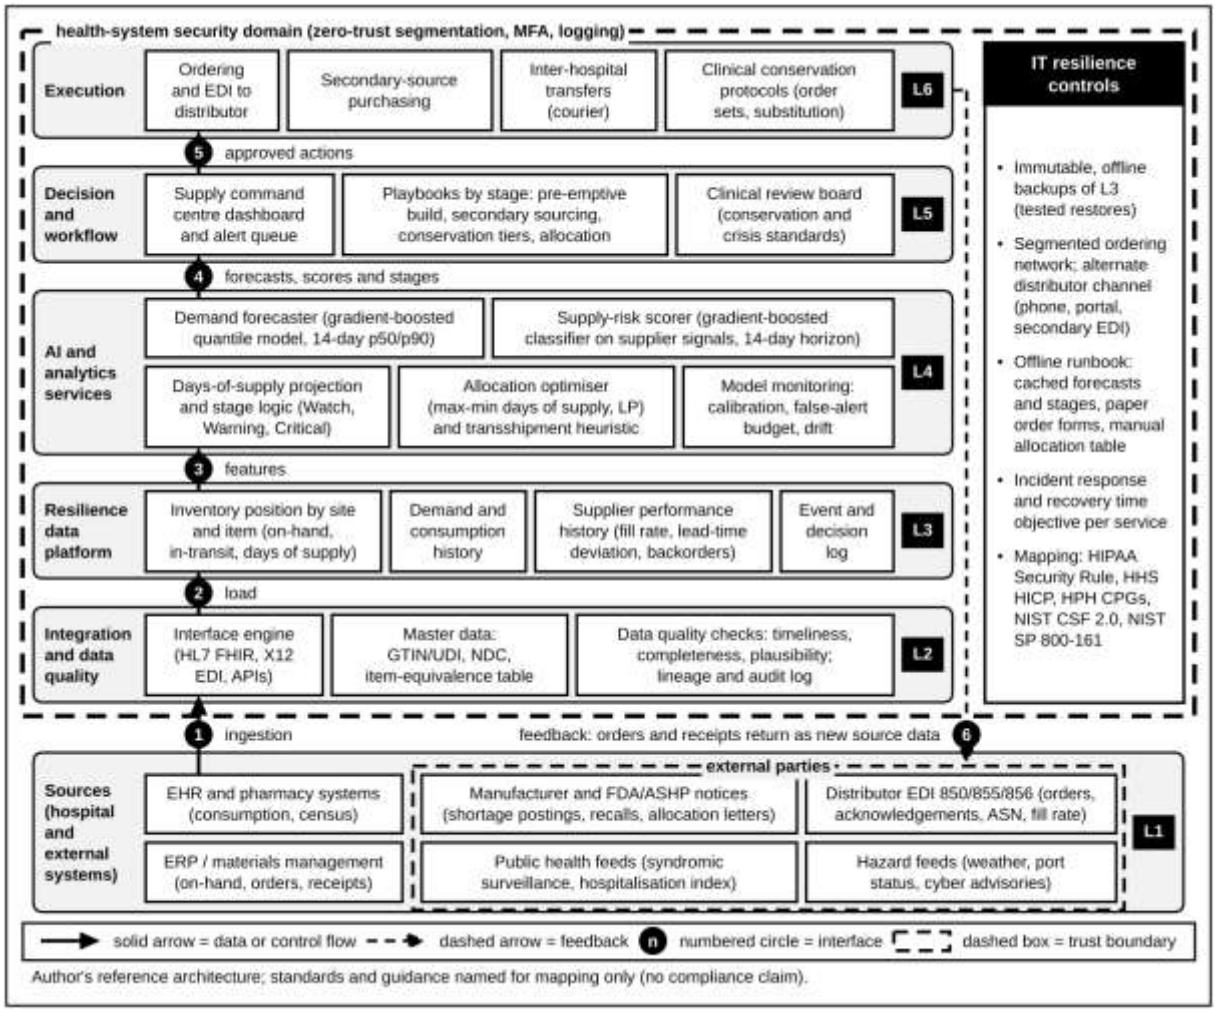
Reference architecture and workflow

Architecture

Figure 1 shows the reference architecture as six layers with numbered interfaces and two trust

boundaries. The source layer holds the systems a health system already has (electronic health record and pharmacy systems for consumption and census, the materials management system for on-hand and orders, the distributor's electronic data interchange for orders, acknowledgements, advance ship notices and fill rates) and the external feeds it can subscribe to (manufacturer and FDA/ASHP notices, public health surveillance, weather, port and cyber advisories). An integration layer standardises identifiers (GTIN and UDI for devices, NDC for drugs, an item-equivalence table for substitutes) and checks data quality, because an alert built on a stale or incomplete inventory position is worse than no alert (Hazen et al., 2014; U.S. Food and Drug Administration, 2013; U.S. Congress, 2013). The data platform keeps the inventory position by site and item, the demand history, the supplier performance history and an event and decision log. The AI layer holds the demand forecaster, the supply-risk scorer, the days-of-supply projection with its stage logic, the allocation optimiser and transshipment heuristic, and model monitoring. The decision layer is the supply command centre with its playbooks and a clinical review board; the execution layer is ordering, secondary-source purchasing, inter-hospital transfers and conservation protocols. Interface 6 closes the loop: executed orders and receipts return as new source data.

Figure 1: Reference architecture for AI-enabled IT and supply chain resilience: six layers with numbered interfaces, trust boundaries and the IT resilience controls



The IT resilience controls sit beside the layers rather than inside one of them. They are the controls that keep the supply function working when the IT it depends on fails: immutable offline backups of the data platform with tested restores, a segmented ordering network with an alternate channel to the distributor (telephone, portal, a secondary EDI route), an offline runbook with cached forecasts and stages, paper order forms and a manual allocation table, and a recovery time objective per service

(Swanson et al., 2010; International Organization for Standardization, 2019; Ross et al., 2021; U.S. Department of Health and Human Services, 2023). The design principle is that the AI layer must degrade to something a person can run, not to nothing. Table 2 lists the requirements that the architecture answers and their rationale.

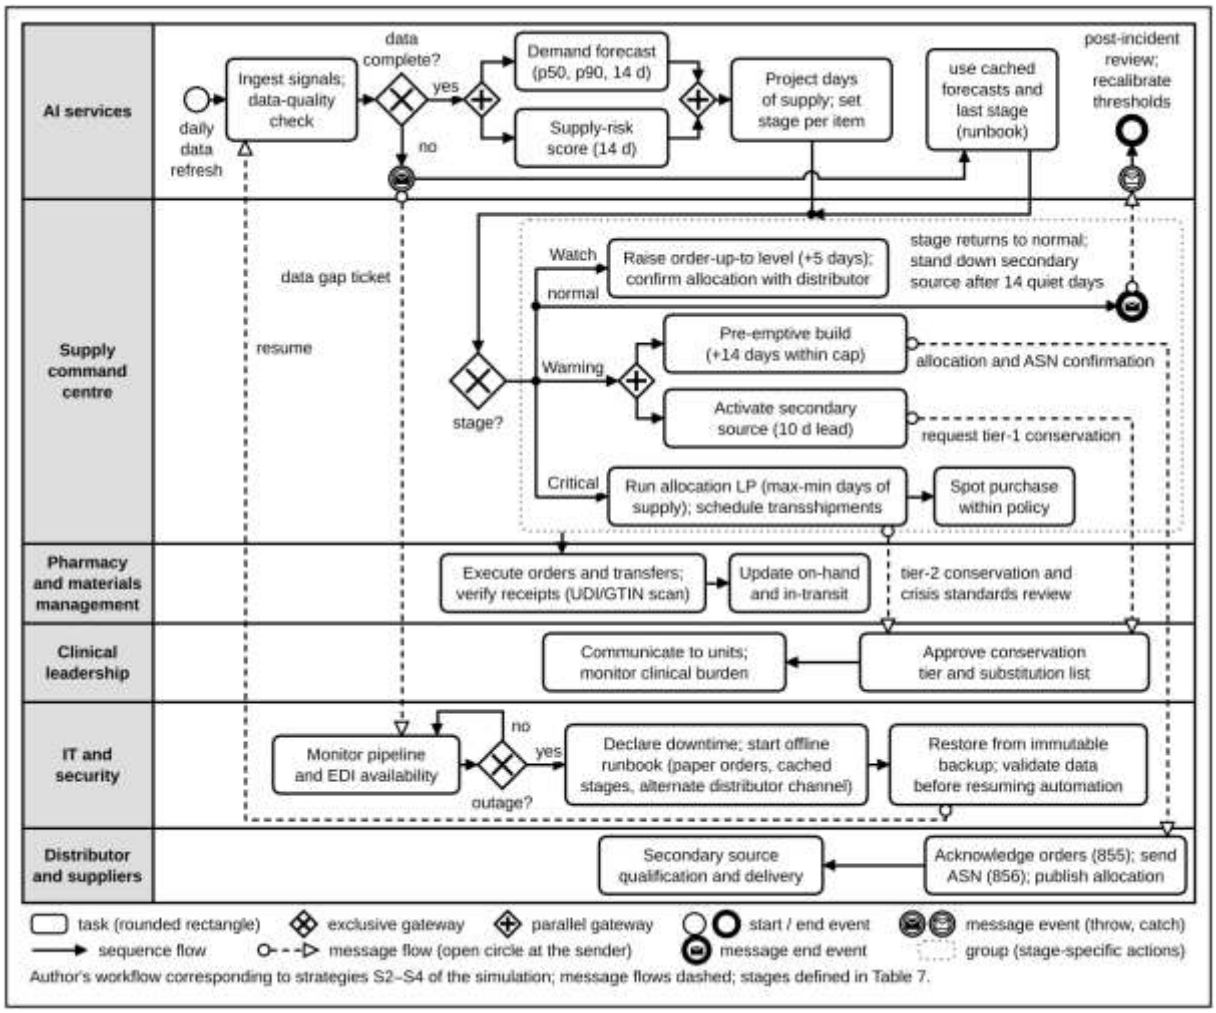
Table 2: Design requirements, their rationale and the architectural response

#	Requirement	Rationale	Design response (Figures 1 to 3)
R1	See every site's position daily	Independent reorder points hide imbalances and delay the recognition of a system-wide shortage	Inventory position by site and item; days of supply on one dashboard
R2	Anticipate, not only react	Allocation letters and shortage postings arrive after supply has already tightened	Supply-risk scorer on fill rate, lead-time deviation, backorder flags and external feeds; hazard and quality signals
R3	Forecast demand with uncertainty	Surges change demand several-fold; point forecasts under-order	Quantile forecaster (p50, p90) with public health indices; p90 used for projections
R4	Budget the false alerts	Every alert triggers purchases and clinical work; unbudgeted alerts are ignored	Thresholds calibrated per item to an agreed number of alerts per year
R5	Act within the supply chain's rules	Distributors cap orders and allocate by history; hoarding harms others	Pre-emptive build limited to the order cap; secondary sources pre-qualified
R6	Share fairly across sites	First-come allocation leaves small hospitals short	Max-min days-of-supply allocation and next-day transshipment
R7	Conserve with clinical governance	Conservation reduces demand but carries clinical burden	Tiered conservation approved by clinical leadership; burden tracked
R8	Keep working without IT	Ransomware removes ordering and visibility for weeks	Offline runbook, cached forecasts, alternate distributor channel, immutable backups
R9	Leave an audit trail	Decisions in a shortage are reviewed afterwards	Event and decision log; post-event recalibration

Workflow

Figure 2 shows the response workflow as swimlanes. Each day the AI services ingest the signals and check data quality; if the data are incomplete the runbook branch uses cached forecasts and the last stage and raises a ticket to IT. The forecaster and the risk scorer run in parallel and the days-of-supply projection sets a stage per item: Watch, Warning or Critical. The command centre acts by stage. At Watch it raises the order-up-to level modestly and confirms allocation terms with the distributor. At Warning it builds inventory pre-emptively within the distributor's order cap, activates a pre-qualified secondary source, asks clinical leadership for tier-1 conservation and confirms allocation and shipping with the distributor. At Critical it runs the allocation optimiser, schedules transshipments, asks for tier-2 conservation and a crisis-standards review, and buys on the spot market within policy. Pharmacy and materials management execute and verify receipts; clinical leadership approves conservation tiers and substitution lists and monitors burden; IT monitors the pipeline and declares downtime when needed; the distributor acknowledges, ships and publishes allocation. Two end events close the cycle: the stage returns to normal and the secondary source is stood down after two quiet weeks, and a post-incident review recalibrates the thresholds. The workflow is the operational form of the professional guidance on shortage management (Fox & McLaughlin, 2018) with the analytics placed where they inform a decision.

Figure 2: Disruption response workflow: BPMN-style swimlanes for AI services, the supply command centre, pharmacy and materials management, clinical leadership, IT and security, and the distributor

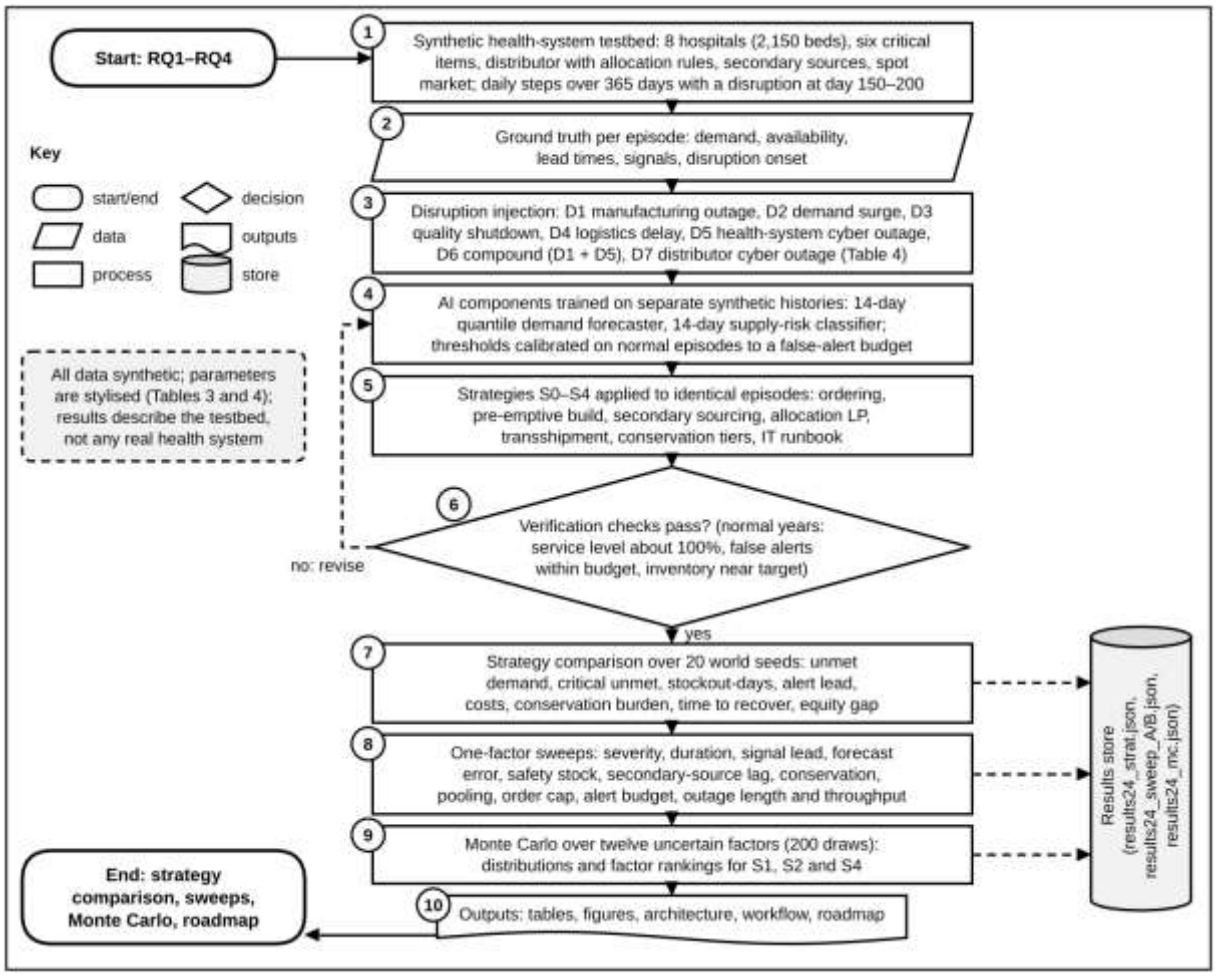


Methods

Study design

Figure 3 shows the study as a flowchart. A synthetic health-system testbed generates daily demand, supply availability, lead times and signals over a year, with a disruption injected between days 150 and 200. The AI components are trained on separate synthetic histories and their thresholds are calibrated on normal years of the testbed. Five strategies are applied to identical episodes (same seed, same demand, same disruption), so that differences between strategies are not confounded by chance. Verification checks confirm that normal years behave as designed. The strategy comparison runs on 20 world seeds, one-factor sweeps trace each factor on eight seeds, and a Monte Carlo analysis varies twelve factors jointly over 200 draws.

Figure 3: Study design: testbed, disruption injection, AI training and calibration, strategies, verification, strategy comparison, sweeps, Monte Carlo and outputs, drawn with standard flowchart symbols



All data are synthetic. No hospital, distributor or manufacturer data were used. The parameters are stylised values chosen so that the testbed behaves plausibly and the disruptions resemble documented U.S. events in mechanism and rough magnitude; they are not measurements of any organisation, and the results describe the testbed.

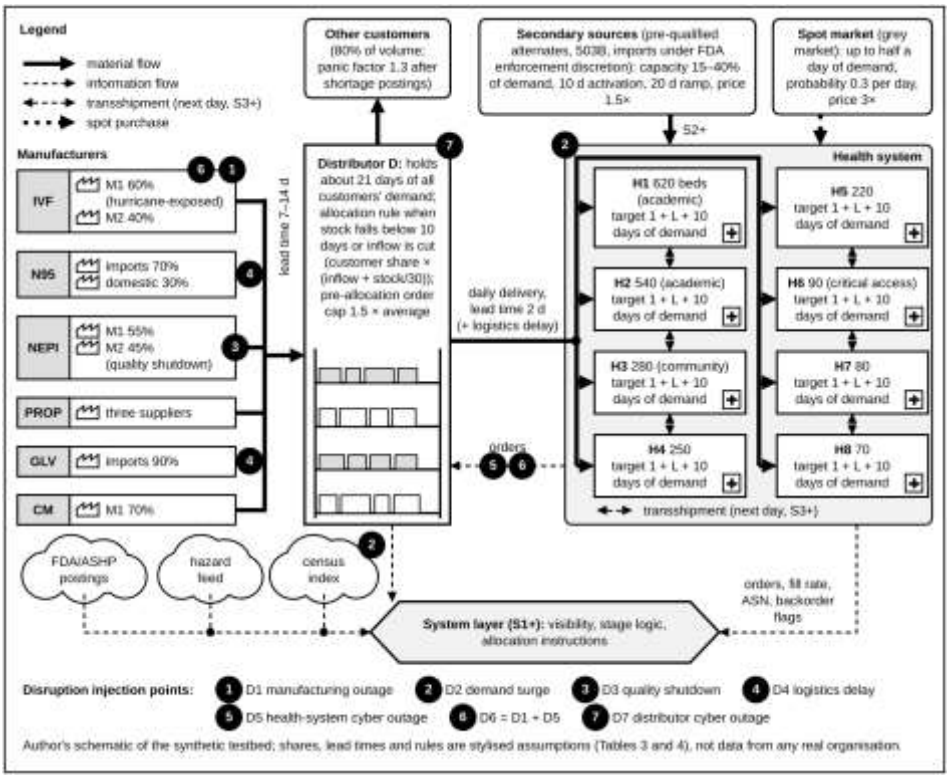
Testbed

Table 3 gives the testbed and Figure 4 its network. Eight hospitals with 2,150 beds in all (two academic, three community, three small) consume six critical items at stylised rates per bed per day: 0.9% sodium chloride 1 L bags, N95 respirators, norepinephrine vials, propofol vials, nitrile exam gloves and iodinated contrast, each with a documented U.S. disruption (fluids after hurricanes, respirators in 2020, norepinephrine in 2011 to 2012, propofol during the 2020 surge in ventilated patients, gloves during the import disruptions of 2020 to 2021, contrast in 2022). Each item carries a criticality weight for unmet demand, conservation tiers (the share of demand that clinical protocols can avoid at tier 1 and tier 2) and a burden weight per conserved unit, so that conserving norepinephrine is counted as costly and conserving gloves as cheap. Daily demand is negative binomial around a mean with weekday and seasonal factors, and each hospital experiences about one local four-day surge per year on an injectable or fluid, which gives pooling something to do even in normal years.

Table 3: Testbed: hospitals, items, demand, distributor, lead times, policies, secondary sources, spot market, conservation, episodes and AI training

Element	Setting in the testbed
Hospitals	Eight: 620 and 540 beds (academic), 280, 250 and 220 (community), 90, 80 and 70 (small); 2,150 beds in all
Items	IVF 1.6, N95 0.25, NEPI 0.08, PROP 0.35, GLV 0.5, CM 0.06 units per bed per day; unit costs 2, 1.2, 8, 3, 8 and 25 USD; criticality weights 0.9, 0.8, 1.0, 0.8, 0.5, 0.7
Demand	Negative binomial by hospital and day; weekday factors 1.1 (Mon to Thu), 1.05, 0.8, 0.75; 6% seasonal swing; about one local 4-day 2.5x surge per hospital-year
Distributor	Holds 21 days of all customers' demand; the health system is 20% of its volume; allocation when stock falls below 10 days or inflow is cut (share x (inflow + stock/30)); order cap 1.5x the average before allocation; other customers order 1.3x after a shortage posting
Lead times	Distributor to hospital 2 days (sd 0.5); manufacturer to distributor 7 to 14 days (continuous inflow)
Hospital policy	Daily order-up-to 1 + 2 + 10 days of expected demand (all strategies at baseline)
Secondary sources	Capacity 15 to 40% of demand by item; 10-day activation, 20-day ramp; price 1.5x (S2+)
Spot purchases	Up to half a day of demand with probability 0.3 per day at 3x price; halved during a national shortage posting
Conservation	Tier 1 (Warning) 5 to 20% of demand; tier 2 (Critical) 20 to 45%; burden weights 0.1 to 1.0 per conserved unit (S3+)
Episode	365 days; 60-day warm-up excluded; disruption onset uniform on days 150 to 200; 20 world seeds
AI training	Forecaster: ten synthetic years with random surges; risk scorer: forty synthetic supplier histories; thresholds calibrated on four normal testbed years per seed

Figure 4: Supply network of the synthetic health-system testbed: manufacturers, distributor with allocation rules, eight hospitals, secondary sources, spot market, information flows and the disruption injection points



The hospitals are replenished daily by one distributor with a two-day lead time; the health system is 20% of the distributor's volume. The distributor holds about 21 days of all its customers' demand, is refilled continuously by the manufacturers and ships what is ordered until its stock falls below ten days or its inflow is cut, at which point it allocates: each customer receives its historical share of the day's inflow plus a thirtieth of the remaining stock. Before allocation, a customer's order is capped at 1.5 times its average, a stylised version of the anti-hoarding caps distributors apply. Other customers order 30% more once a shortage is posted. These rules reproduce the behaviour that health systems reported in 2017 and 2020: a brief window in which pre-emptive buying is possible, then allocation by history (Mazer-Amirshahi & Fox, 2018; Finkenstadt & Handfield, 2021). Secondary sources (pre-qualified alternates, outsourcing facilities, imports under enforcement discretion) can supply 15 to 40% of an item's demand after a ten-day activation and a 20-day ramp at a 50% premium; the spot market supplies up to half a day of demand with probability 0.3 per day at three times the price, halved during a national shortage posting.

Table 4: Disruption catalogue: mechanism, default size, leading signals and the documented U.S. analogue of each class

Class	Mechanism in the testbed	Default size	Leading signals (lead)	Real-world analogue
D1 manufacturing outage	Dominant IVF manufacturer out; distributor inflow 30% for 90 days, 30-day recovery	availability 0.3, 90 d	Hazard feed 4 d before (p = 0.85); allocation flag; posting +10 d	Hurricane Maria 2017 and Helene 2024 (Mazer-Amirshahi & Fox, 2018; Baxter, 2025)
D2 demand surge	Respiratory wave: N95 x4, PROP x2.65, NEPI x2.35 over a 42-day ramp, 70-day plateau; national N95 allocation at 70% from +21 d; hospitals hit with different intensity and lag	x4, 112 d	Census index leads by 14 d; fill rate; posting	COVID-19 spring 2020 (Ranney et al., 2020)
D3 quality shutdown	Sterile-injectable line (45% share) halted 150 days	availability 0.55, 150 d	Quality signal 75 d before (p = 0.5); allocation; posting +14 d	Norepinephrine shortage 2011 to 2012 (Vail et al., 2017)
D4 logistics delay	Import lead time +10 d, sd x3, availability 60% for 60 days (N95, GLV)	+10 d, 60 d	ASN lead-time deviation immediate; allocation	Port congestion and PPE imports 2020 to 2021 (Cohen & Rodgers, 2020)
D5 health-system cyber outage	Ordering and visibility lost 28 days; manual orders at 70% throughput, 35% error, 3-day delay (S1 to S3); 3 days, 90%, 10%, 1 day with the runbook (S4); S0 phones and faxes as usual	28 d	None (immediate)	Hospital ransomware (Neprash et al., 2022; Dameff et al., 2023)
D6 compound	D1 with the cyber outage starting 3 days after onset	as D1 + D5	As D1	A plausible coincidence, not a documented one
D7 distributor cyber outage	Distributor processes 20% of orders for 14 days; S4 alternate channel at 70% after 2 days	14 d, 20%	Allocation flag immediate	Change Healthcare 2024 as an analogue for a shared intermediary (U.S. Department of Health and Human Services, Office for Civil Rights, 2024)

Disruption classes

Table 4 lists the seven classes. D1 is a manufacturing outage of the dominant intravenous-fluid manufacturer: the distributor's inflow falls to 30% for 90 days and recovers over 30 days, a hazard feed fires four days before onset in 85% of episodes, and a public posting appears ten days after onset. D2 is a respiratory demand surge: N95 demand rises four-fold over a 42-day ramp with a 70-day plateau, propofol and norepinephrine rise with it, hospitals are hit with different intensities and lags, a regional census index leads the surge by 14 days, and a national allocation at 70% begins three weeks in. D3 is a quality shutdown of a sterile-injectable line with 45% of the norepinephrine market for 150 days, with a quality signal (an inspection notice) 75 days before onset in half of the episodes. D4 is a logistics delay for the imported items: lead time plus ten days with three times the variability and 60% availability for 60 days. D5 is a 28-day cyber outage of the health system's ordering and visibility systems, during which orders are placed manually at 70% throughput with 35% quantity error and a three-day delay; S0, which orders by telephone and fax by design, suffers only its usual 15% error and one-day delay; the IT resilience layer (S4) shortens the outage to three days with 90% throughput and 10% error. D6 is D1 with the cyber outage starting three days after the supply onset. D7 is a distributor cyber outage in which 20% of orders are processed for 14 days, with an alternate channel at 70% after two days for S4. The outage lengths, throughputs and error rates are assumptions, since published accounts give durations but not ordering throughput, and they are swept.

Signals

Table 5 lists the signals and how the testbed generates them. The fill rate is the share of the normal-sized order that the distributor ships; short-fills caused only by the order cap are excluded, because a system that read its own surge orders as supplier distress would alarm on itself. Lead-time deviation comes from advance ship notices, the allocation flag from the distributor's allocation, and public postings arrive with a lag (with false postings for unaffected items about once per 180 days). The hazard feed and the quality signal lead by days and months respectively, each with a probability of appearing; the census index leads a demand surge by two weeks; days of supply is the inventory position divided by expected demand.

Table 5: Signals available to the system layer, their sources, their synthetic generation and who sees them

Signal	Source	How it is generated in the testbed	Who sees it
Fill rate	Distributor acknowledgement (855) against the order	Shipped over the normal-sized order; 2% reporting noise; short-fills caused only by the order cap are excluded	S1+ (S0 sees its own orders only)
Lead-time deviation	ASN (856) against the promise	Realised lead time minus 2 days plus noise (sd 0.3)	S1+
Allocation flag	Distributor allocation letter	On when the distributor is allocating and the order is short-filled	S1+
External posting	FDA and ASHP shortage lists	On from onset + 10 to 14 days until 40 days after recovery; false postings for unaffected items about once per 180 days, lasting 30 days	S1+
Hazard feed	Weather and emergency alerts	On 4 days before a manufacturing outage with probability 0.85; about two false alerts per year	S2+ (scorer feature)
Quality signal	FDA inspection and warning-letter notices	On 75 days before a quality shutdown with probability 0.5	S2+
Census index	Public health surveillance	Leads the demand surge by 14 days; 3% noise	S2+ (forecaster feature)
Days of supply	Own inventory position	(On hand + in transit) / expected daily demand, by site and system	S1+ (S0 by site only)

Strategies

Table 6 defines the five strategies. S0, independent reorder points, is the baseline: each hospital holds 13 days of demand, reorders daily, sees only itself and buys on the spot market when it falls below two days. S1 adds system-wide visibility with rule-based alerts (fill rate below 0.8, an allocation flag, or system days of supply below five) that trigger a 14-day build and spot purchases, and constrained supply is split across hospitals in proportion to orders. S2 adds the AI layer: Watch and Warning stages from the forecaster and the risk scorer, a build of five days at Watch and 14 at Warning within the order cap, and a secondary source activated at Warning. S3 adds pooling and conservation: a linear programme splits constrained supply to maximise the minimum days of supply, stock moves between hospitals overnight when their days of supply diverge, and conservation tiers apply at Warning and Critical. S4 adds the IT resilience runbook. The strategies are cumulative so that each increment can be attributed; the sweeps ablate single components.

Table 6: Strategies S0 to S4: visibility, early warning, ordering and sourcing, allocation and pooling, conservation and IT resilience

Strategy	Visibility	Early warning	Ordering and sourcing	Allocation and pooling	Conservation	IT resilience
S0 Independent reorder points	Site only	None	Order-up-to 13 days; spot purchases when a site falls below 2 days	Distributor fills sites in arrival order	None	Manual processes by design (1-day delay, 15% error)
S1 System visibility	All sites	Rules: fill rate below 0.8, allocation flag, system days of supply below 5	+14 days of demand at Warning; spot purchases at Warning	Proportional to orders	None	Ordering system down 28 days: 3-day delay, 35% error, 70% throughput
S2 AI early warning	All sites	Risk scorer and forecaster: Watch and Warning thresholds calibrated to 2 and 1 alerts per item-year; forecast growth; days of supply	+5 days at Watch, +14 at Warning (within the cap); secondary source activated at Warning	Proportional to orders	None	As S1
S3 AI with pooling and conservation	All sites	As S2	As S2	Max-min days-of-supply LP for constrained supply; next-day transshipment	Tier 1 at Warning, tier 2 at Critical	As S1 (stages frozen, pooling off during the outage)
S4 S3 with IT resilience	All sites	As S2	As S2	As S3	As S3	Runbook: 3-day outage, cached forecasts and stages, 90% throughput, 10% error; alternate distributor channel

AI components and stage logic

Table 7 defines the AI components. The demand forecaster is a pair of gradient-boosted quantile regressors (Ke et al., 2017; Koenker & Bassett, 1978) that predict the ratio of the next 14 days' demand to the trailing 28-day mean at the median and the 90th percentile from trailing demand ratios, the weekday, the item, the hospital's size and the census index's level and growth; it is trained once per seed on ten synthetic years with random surges, never on the episode it is applied to. The supply-risk scorer is a gradient-boosted classifier that predicts a supply cut or allocation within 14 days from the signals in Table 5, trained on forty synthetic supplier histories. Its thresholds are set per item on four normal testbed years so that alert episodes per item-year equal a budget (two for Watch, one for Warning), the commissioning step any alerting system needs (Axelsson, 2000). The stage logic combines the risk score, forecast growth (p90 ratio above 1.25 for Watch, median ratio above 1.5 for Warning) and days of supply (Watch below ten, Warning below five, Critical below two or any site stockout), and de-escalates one level per quiet week. The allocation optimiser is a small linear programme solved with HiGHS (Huangfu & Hall, 2018); the transshipment heuristic moves stock from sites above the median days of supply plus two to sites below the median minus one and below five, at most 1.5 days of the donor's demand per lane per day.

Table 7: AI components and stage logic: inputs, models and outputs, and their use in the stages

Component	Inputs	Model and output	Use in the stage logic
Demand forecaster	Trailing 3, 7 and 28-day demand ratios, weekday, item, hospital size, census level and 7- and 14-day growth	Two gradient-boosted quantile regressors (300 trees, 15 leaves) predicting the 14-day demand ratio at p50 and p90	p50 sets order-up-to levels; p90 adds 7 days of safety at Watch; growth above 1.25 (p90) or 1.5 (p50) raises the stage
Supply-risk scorer	7-day mean and minimum fill rate, lead-time deviation, allocation flag, posting, hazard, quality and national-allocation flags, item, import flag	Gradient-boosted classifier (200 trees) for a supply cut or allocation within 14 days	Watch above the 2-per-year threshold; Warning above the 1-per-year threshold (per item)
Days-of-supply projection	Inventory position and p50 forecast	Deterministic	Watch below 10 days, Warning below 5, Critical below 2 or any site stockout
Allocation optimiser	Supply of the day, on-hand and expected demand by site, orders	Linear programme maximising the minimum days of supply (HiGHS)	Used by S3 and S4 when supply is short of orders
Transshipment heuristic	Days of supply by site	Moves from sites above the median plus 2 days to sites below the median minus 1 day and below 5 days; at most 1.5 days of the donor's demand per lane per day	S3 and S4 at Watch or above
Stage de-escalation	Stage history	One level per quiet week	Secondary source stood down after 14 quiet days

Metrics

Unmet demand is demand that is neither served nor conserved, expressed as a share of the affected items' post-onset demand (and reported per item); critical unmet demand weights each unmet unit by its criticality. Stockout-days count hospital-item-days with unmet demand. The warning lead is the number of days between the first Warning for an affected item, counted from shortly before the earliest designed signal, and the onset; positive values precede the onset. False warnings are counted in normal

years as alert episodes per item-year. Costs are holding cost at 25% per year of inventory value, spot and secondary premiums, and transfer cost at 0.10 USD per unit; the conservation burden is the burden-weighted sum of conserved units. Time to recover is the number of days after onset until no site is below two days of supply, for items that fell below it; the equity gap is the largest spread of days of supply across sites after onset. Measurements exclude a 60-day warm-up.

Sweeps and Monte Carlo

One-factor sweeps on eight seeds vary the outage severity and duration, the hazard-feed lead and probability, the surge magnitude, the census lead, the quality-signal probability, the posting lag, the hospital safety stock (for S1 and S3, to separate inventory from intelligence), the secondary-source lag, the distributor order cap, conservation effectiveness, the pooling ablation, forecast degradation, the false-alert budget (with recalibration), the IT outage length, manual throughput during the compound event, the distributor outage length, the logistics delay and other customers' panic factor. The Monte Carlo analysis draws 200 combinations of twelve factors (outage availability and duration, hazard lead, surge multiplier, census lead, posting lag, safety stock, secondary lag, conservation effectiveness, forecast noise, IT outage length and order cap) from uniform distributions over the sweep ranges, recalibrates the thresholds for each draw, and runs D1, D2, D3 and D6 once for S1, S2 and S4; with one seed per draw its results are read as distributions and orderings.

Verification and reproducibility

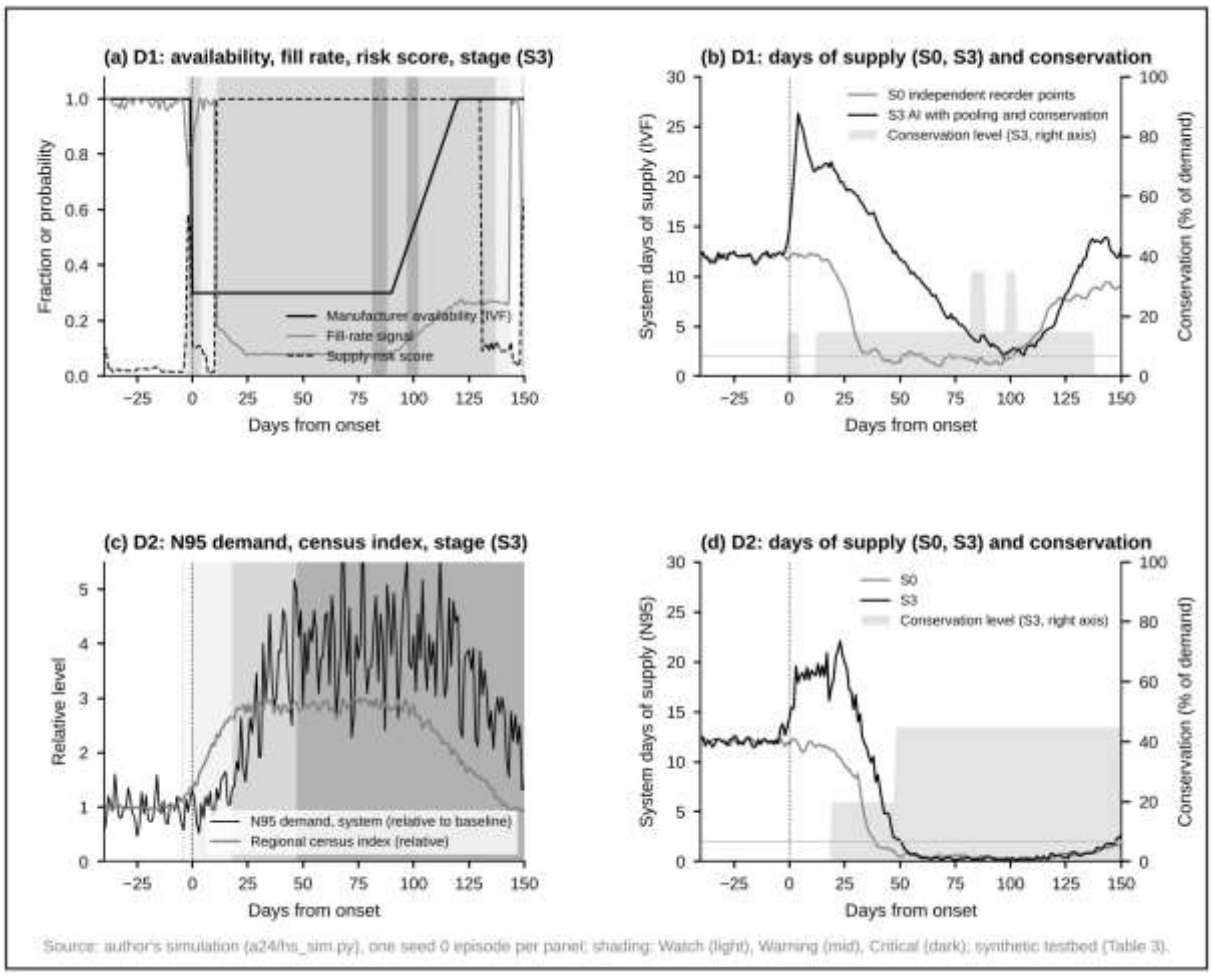
Three checks were applied before the experiments. In normal years every strategy served essentially all demand with no stockout-days and inventory near its target (10.3 to 10.6 days on average), and S2 to S4 produced 0.45 to 1.15 Warning episodes per item-year against a budget of one. The distributor's stock held steady in normal years, so that allocation never triggered without a disruption. And the forecaster's p90 ratio exceeded the median ratio in every prediction. These checks establish internal consistency, not fidelity to any health system. The simulation, the runner, the aggregation script and the result files are in the supplementary archive; fixed seeds reproduce every number (Law, 2015; Sargent, 2013).

Results

Testbed behaviour

Figure 5 shows one episode of each of the two largest disruptions as the system layer saw them. In the manufacturing outage (panel a) the hazard feed lifted the supply-risk score to about 0.6 four days before onset, which put the item on Watch; the score then fell back until the distributor's fill rate collapsed from day 2 and the allocation flag and the public posting pushed it to 1.0 and the item to Warning on day 11, where it stayed, with two short Critical periods, until availability recovered after day 120. Panel (b) shows what the stages did: S3 built its intravenous-fluid position from 12 to 26 days of supply in the first week by ordering up to the distributor's cap, spent that cushion over the following ten weeks, turned on tier-1 conservation at Warning and tier-2 briefly at Critical, and touched the two-day line only around day 100, while S0's position fell below two days by day 35 and stayed there for two months. In the demand surge (panels c and d) the census index began rising about two weeks before demand did and the forecaster put N95 respirators on Watch three days before onset; S3's position rose to 20 days, then fell to the floor by day 50 as the national allocation bit, and tier-2 conservation (45% of demand) carried the plateau. S3 spent its inventory early and deliberately; S0 spent it late and involuntarily.

Figure 5: Testbed behaviour in two seed 0 episodes: the manufacturing outage and the demand surge as seen by the system layer (signals, risk score, stages, days of supply and conservation), S0 against S3



Source: author's simulation (a24/hs_sim.py), one seed 0 episode per panel; shading: Watch (light), Warning (mid), Critical (dark); synthetic testbed (Table 3).

Strategy comparison

Table 8 and Figure 6 give the main comparison over 20 world seeds. Under the 90-day manufacturing outage, unmet demand for intravenous fluids after onset was 23.7% of demand with independent reorder points (S0), 20.3% with system visibility and rule-based alerts (S1), 9.9% with AI early warning and secondary sourcing (S2), and 0.5% with pooling and conservation added (S3 and S4, which coincide when there is no IT outage). Under the four-fold demand surge, unmet demand over the three surge items fell from 38.3% (S0) to 36.3% (S1), 26.3% (S2) and 6.8% (S3); by item, N95 respirators went from 59.3% to 13.1%, propofol from 22.0% to 1.7% and norepinephrine from 15.4% to 1.4%. Under the quality shutdown, unmet norepinephrine demand fell from 20.0% to 16.4%, 0.3% and 0.0%. The logistics delay was mild at its default size (3.8% for S0, below 0.5% for S1 and above), and the distributor outage was absorbed by inventory except under S0 (1.1%). The cyber outage alone cost little under any strategy (at most 0.5%), but the compound event separated them: 22.1% (S0), 21.6% (S1), 12.9% (S2), 2.8% (S3) and 0.8% (S4). Standard deviations over seeds were 2 to 5 percentage points for the large effects, so the ordering S0 > S1 > S2 > S3 held in every seed for the outage, the surge and the shutdown.

Table 8: Unmet demand after onset for the affected items (% of demand; mean and standard deviation over 20 seeds) by strategy and disruption, with normal-year false warnings and inventory

Disruption (affected items)	S0	S1	S2	S3	S4
D1 manufacturing outage (IVF)	23.7 (2.1)	20.3 (2.1)	9.9 (2.8)	0.5 (0.5)	0.5 (0.5)
D2 demand surge (N95, PROP, NEPI)	38.3 (3.9)	36.3 (4.3)	26.3 (4.7)	6.8 (2.4)	6.8 (2.4)
D3 quality shutdown (NEPI)	20.0 (2.2)	16.4 (2.5)	0.3 (0.6)	0.0 (0.0)	0.0 (0.0)
D4 logistics delay (N95, GLV)	3.8 (0.5)	0.5 (0.3)	0.2 (0.2)	0.0 (0.0)	0.0 (0.0)
D5 health-system cyber outage	0.0 (0.0)	0.5 (0.2)	0.4 (0.3)	0.4 (0.2)	0.0 (0.0)
D6 compound (D1 + D5)	22.1 (2.3)	21.6 (2.3)	12.9 (5.2)	2.8 (1.8)	0.8 (0.7)
D7 distributor cyber outage	1.1 (0.2)	0.0 (0.0)	0.1 (0.1)	0.0 (0.0)	0.0 (0.0)
Normal year: Warning episodes per item-year	0.00	0.00	0.83	0.84	0.84
Normal year: average inventory (days)	10.4	10.3	10.6	10.6	10.6

Figure 6: Strategy comparison: unmet demand, stockout-days and premium spend by strategy and disruption, and the warning lead by strategy

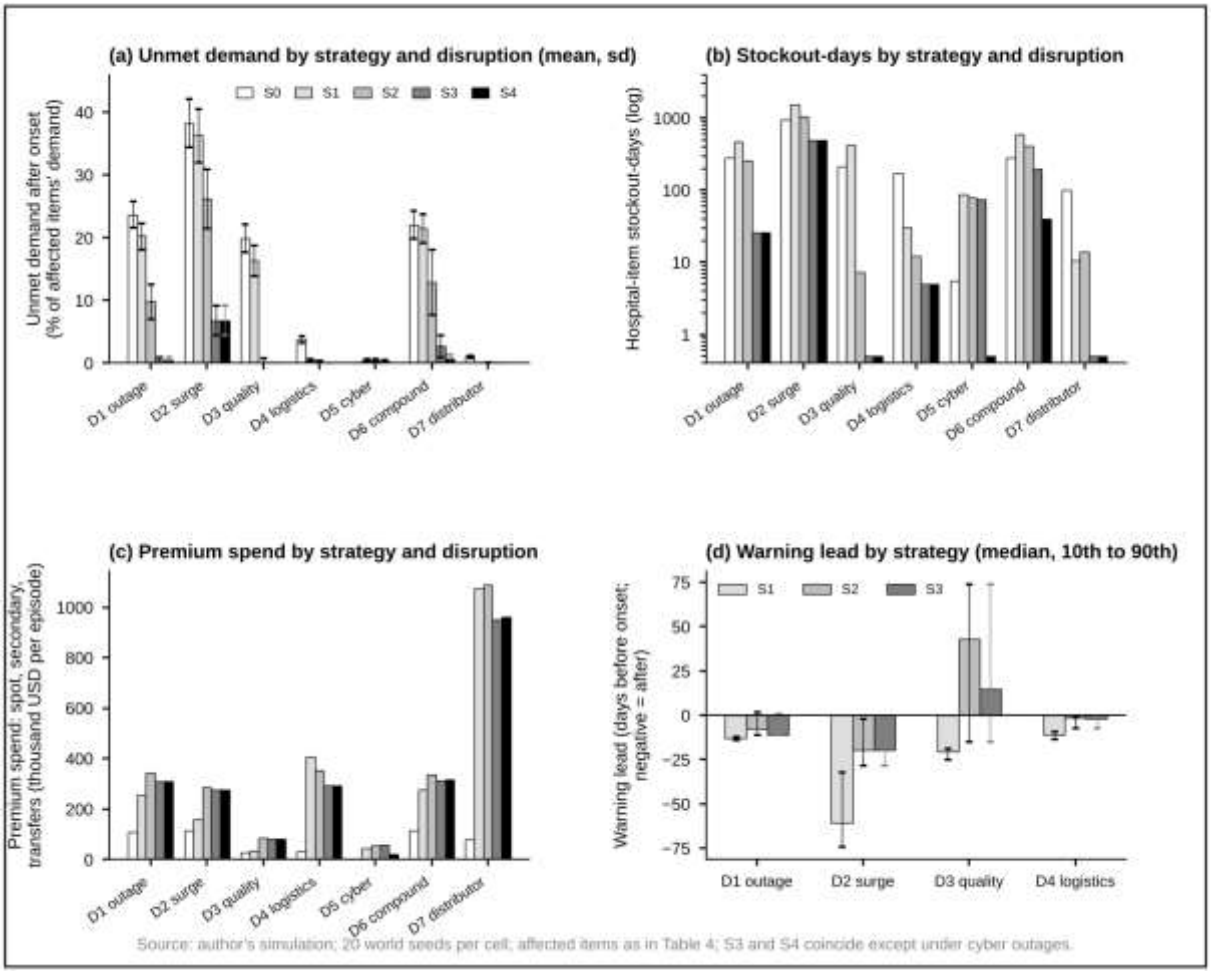


Table 9: Outcome detail by strategy for the outage, the surge, the shutdown and the compound event: stockout-days, critical-weighted unmet demand, time to recover, equity gap, conserved and transferred units

Disruption	Strategy	Stockout-days	Critical unmet (thousand weighted units)	Time to recover (days, median)	Equity gap (days of supply)	Conserved units (thousand, main item)	Transferred units (thousand)
D1 outage	S0	282	132.0	123	13.2	0	0
	S1	473	112.4	121	7.1	0	0
	S2	262	55.3	113	8.2	0	0
	S3	26	2.7	94	7.7	100.8	0.2
	S4	26	2.7	94	7.7	100.8	0.2
D2 surge	S0	943	188.3	139	14.2	0	0
	S1	1527	178.8	166	8.4	0	0
	S2	1055	129.2	157	15.6	0	0
	S3	498	33.9	143	15.2	97.9	0.5
	S4	498	33.9	143	15.2	97.9	0.5
D3 quality	S0	213	6.1	76	13.6	0	0
	S1	422	5.0	102	8.7	0	0
	S2	7	0.1	156	11.6	0	0
	S3	0	0.0	147	9.7	1.6	0.0
	S4	0	0.0	147	9.7	1.6	0.0
D6 compound	S0	279	128.5	118	15.3	0	0
	S1	595	126.5	116	9.4	0	0
	S2	410	76.8	116	11.9	0	0
	S3	198	17.6	81	11.8	110.0	0.4
	S4	41	4.6	68	8.5	105.3	0.1

Each increment carries a different mechanism. S1's gain over S0 is small in unmet demand but large in its distribution: with proportional allocation every site was equally short (about 60 stockout-days per site in the outage, against 68 and 64 at the two academic centres and 16 to 18 at the small hospitals under S0's first-come allocation), and critical-weighted unmet demand fell from 132,000 to 112,000 weighted units; S1 also spent twice as much on spot purchases, because its rules keep buying for as long as the fill rate is low. S2's gain comes from two actions taken earlier: building inventory within the order cap as soon as the stage rises, and activating a secondary source at Warning, which delivered up to 15% of demand from about three weeks after onset in the outage and up to 20% in the shutdown, making the shutdown almost invisible to patients (0.3% unmet) at a secondary-source premium of 24,000 USD. S3's gain comes from conservation and the allocation rule: conservation avoided 101,000 units of intravenous-fluid demand in the outage (a burden of 31,000 weighted units) and 98,000 respirators, 29,000 propofol vials and 3,300 norepinephrine vials in the surge, and stockouts fell to about three hospital-days per site (Table 9). The surge is the disruption no strategy solves: a four-fold increase in respirator demand against a national allocation at 70% of normal supply leaves 13% unmet even with 45% conservation, which is the arithmetic of 2020.

Table 9 gives the outcome detail. Stockout-days fell from 282 (S0) to 26 (S3) in the outage and from 943 to 498 in the surge. Critical-weighted unmet demand fell from 132,000 to 2,700 weighted units in the outage and from 188,000 to 34,000 in the surge. The time to recover (the first day after onset on which no site is below two days of supply) is a first-crossing measure that rewards a brief recovery: it was 123

days for S0 and 95 for S3 in the outage, but longer for S2 and S3 (156 and 147 days) than for S0 (77 days) in the shutdown, because S2 and S3 rode close to the two-day line on secondary supply rather than stocking out and briefly recovering on spot purchases. The equity gap, the largest spread of days of supply across sites after onset, narrowed from 13.2 days (S0) to 7 to 8 days for S1 to S3 in the outage; in the surge it stayed at 15 to 16 days for S2 and S3, because the pre-emptive build raised some sites before others and the hospitals were hit at different times.

Early warning

Table 10 and Figure 6d give the warning leads. With rule-based alerts (S1) the first Warning came 13 days after the onset of the outage (never before it), 60 days after the start of the surge ramp and 20 days after the quality shutdown, because the rules wait for the fill rate to fall or for allocation. With the AI layer (S2) the median Warning for the outage came 8 days after onset, 35% of episodes had a Warning before onset and all had one within 14 days; the Watch stage, which already raises the order-up-to level, preceded onset in 40% of episodes. The hazard feed is the only signal that can lead a manufacturing outage, and it leads by four days in 85% of episodes; the Warning threshold, calibrated to one false alert per item-year, is too high for a signal of that strength to cross on its own, so most of the pre-onset response came through Watch. For the quality shutdown the picture is different: the inspection notice, when it appears (half of the episodes), gives the scorer a long lead, and the median Warning came 43 days before onset with 55% of episodes warned in advance. For the surge the forecaster's growth trigger put N95 respirators on Watch three days before the ramp began in 85% of episodes, through the census index, but Warning waited for demand itself: a median of 20 days after onset. The logistics delay and the distributor outage were recognised within a day or two of onset by the lead-time deviation and the allocation flag, as designed.

Table 10: Early warning: Warning and Watch leads by strategy and disruption (median days before onset; share of episodes warned before onset and within 14 days after), and false warnings in normal years

Disruption	Strategy	Warning lead (days, median)	Warning before onset (% of episodes)	Warning within 14 days after onset (%)	Watch lead (days, median)	Watch before onset (%)
D1 outage	S1	−13	0	95	–	–
	S2	−8	35	100	−2	40
D2 surge	S1	−60	0	0	–	–
	S2	−20	7	18	−2	55
D3 quality	S1	−20	5	5	–	–
	S2	43	55	55	52	65
D4 logistics	S1	−11	0	95	–	–
	S2	−2	2	100	−1	5
D7 distributor	S1	−1	0	100	–	–
	S2	−1	3	99	−1	6
Normal year (S2)	false Warnings per item-year	1.05 / 1.15 / 0.45 / 0.90 / 0.95 / 0.50				

The false-alert budget held. In normal years S2 produced 0.45 to 1.15 Warning episodes per item-year (budget one) and 5 to 11 Watch-days per item over the ten measured months, at a cost of 25,000 USD of spot purchases and 0.3 days of extra inventory per year; the false Warnings also activated secondary sources that were stood down after two quiet weeks, at no premium because nothing was short. These are the costs a health system pays for the early response, and they are modest; the sweeps below show what happens when the budget is tightened or loosened.

Table 11: Costs per episode by strategy and disruption: holding, spot and secondary premiums, transfers, conservation burden and unmet demand

Disruption	Strategy	Holding (thousand USD)	Spot premium (thousand USD)	Secondary premium (thousand USD)	Transfers (thousand USD)	Conservation burden (thousand weighted units)	Unmet demand (%)
Normal	S0	53	0	0	0.0	0.0	0.0
	S1	53	0	0	0.0	0.0	0.0
	S2	54	25	0	0.0	0.0	0.0
	S3	54	26	0	0.0	1	0.0
	S4	54	26	0	0.0	1	0.0
D1 outage	S0	48	108	0	0.0	0.0	23.7
	S1	48	255	0	0.0	0.0	20.3
	S2	52	259	84	0.0	0.0	9.9
	S3	54	234	77	0.0	31	0.5
	S4	54	234	77	0.0	31	0.5
D2 surge	S0	52	114	0	0.0	0.0	38.3
	S1	51	159	0	0.0	0.0	36.3
	S2	57	207	80	0.0	0.0	26.3
	S3	59	200	77	0.0	38	6.8
	S4	59	200	77	0.0	38	6.8
D3 quality	S0	52	29	0	0.0	0.0	20.0
	S1	51	36	0	0.0	0.0	16.4
	S2	55	61	24	0.0	0.0	0.3
	S3	56	59	23	0.0	3	0.0
	S4	56	59	23	0.0	3	0.0
D4 logistics	S0	48	33	0	0.0	0.0	3.8
	S1	58	407	0	0.0	0.0	0.5
	S2	59	225	130	0.0	0.0	0.2
	S3	60	192	102	0.0	5	0.0
	S4	60	192	102	0.0	5	0.0
D6 compound	S0	48	117	0	0.0	0.0	22.1
	S1	48	278	0	0.0	0.0	21.6
	S2	52	263	75	0.0	0.0	12.9
	S3	54	247	68	0.0	35	2.8
	S4	54	239	78	0.0	32	0.8
D7 distributor	S0	50	77	0	0.0	0.0	1.1
	S1	72	1077	0	0.0	0.0	0.0
	S2	68	833	257	0.0	0.0	0.1
	S3	69	750	200	0.0	29	0.0
	S4	71	759	201	0.0	30	0.0

Costs and burden

Table 11 and Figure 6c give the costs per episode. Holding cost was almost identical across strategies (48,000 to 59,000 USD per episode), because S2 to S4 hold only 0.3 days more inventory than S0 in normal years and build only when the stage rises; the comparison is therefore not one of more inventory against less, which the safety-stock sweep below confirms. Premium spend differed. In the outage S0 spent 108,000 USD on spot purchases, S1 255,000, S2 259,000 plus 84,000 on the secondary source, and S3 234,000 plus 77,000; in the surge the figures were 114,000, 159,000, 207,000 plus 80,000, and 200,000 plus 77,000. The distributor outage was the most expensive event for every strategy that reacted to it (750,000 to 1,077,000 USD of spot purchases over 14 days, since every item was short at once), and S1 spent the most because its rules buy on the spot market for as long as the fill rate is low. The logistics delay showed the same pattern: S1 spent 407,000 USD on spot purchases to achieve 0.5% unmet demand, while S3 spent 192,000 plus 102,000 on the secondary source to achieve 0.0%. The conservation burden is the cost that does not appear in a budget: 31,000 weighted units in the outage, 38,000 in the surge and 29,000 in the distributor outage for S3 and S4, almost all of it from intravenous fluids and respirators, whose burden weights are low; conservation of norepinephrine, whose weight is 1.0, was 3,300 vials in the surge and 1,600 in the shutdown.

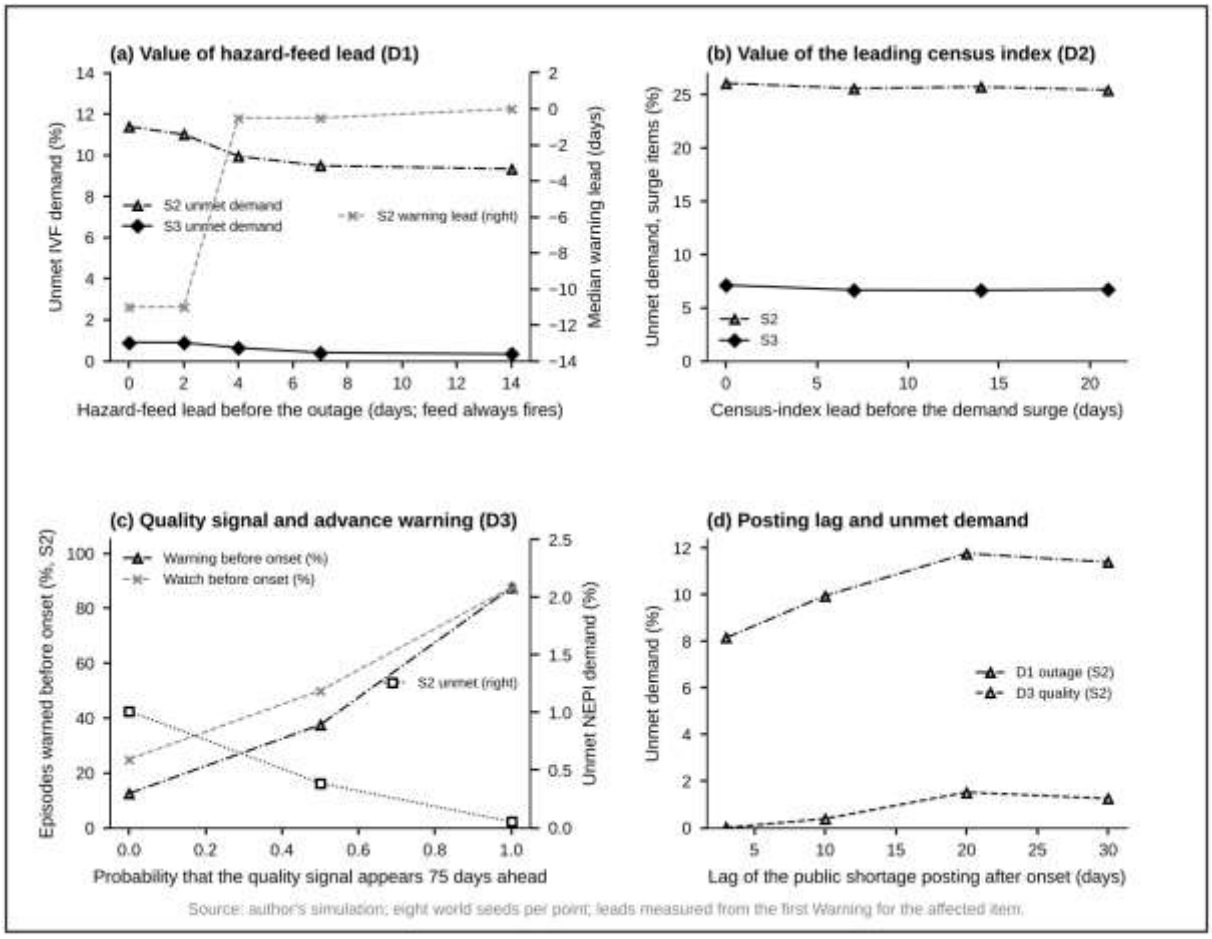
Signals and their lead

Table 12 and Figure 7 trace the value of each leading signal. The hazard feed is worth less than its lead suggests: with the feed always firing, moving its lead from zero to 14 days reduced S2's unmet demand from 11.4% to 9.3% and S3's from 0.9% to 0.3%, and the median Warning moved from 11 days after onset to the day of onset. Two things limit it: a four-day weather alert, which also fires twice a year without an outage, does not clear a Warning threshold calibrated to one false alert per item-year, so the pre-onset response runs through Watch and its five-day build; and the distributor's order cap limits the build to 1.5 times the average order, so that a week of foresight buys about three days of supply. Raising the probability that the feed fires from zero to one moved S2's unmet demand from 13.1% to 10.0% in the sweep that varied it, which is the same mechanism seen from the other side. The census index behaves similarly for the surge: its lead (zero to 21 days) changed unmet demand by less than one point, because the national allocation, not the system's own ordering, decides how many respirators arrive once the surge is under way, and a forecast cannot buy what is not shipped. The quality signal is the exception, because its lead is measured in months: when it appeared in every episode, S2 warned 75 days ahead and unmet norepinephrine demand fell to 0.1%; when it never appeared, Warning waited for allocation (15 days after onset) and unmet demand was 1.0%, still far below S1's 16.4%, because the secondary source and the build at allocation did most of the work. The public posting lag mattered more than the hazard lead: moving the posting from 3 to 20 days after onset raised S2's unmet demand in the outage from 8.2% to 11.8% and in the shutdown from 0.0% to 1.5%, because the posting is the signal that lifts the score to Warning and activates the secondary source when the fill rate alone has not.

Table 12: Severity and signal sweeps: unmet demand for S1, S2 and S3 and the S2 Warning lead against outage availability and duration, surge magnitude, hazard-feed lead and probability, census lead, quality-signal probability, posting lag, forecast noise and logistics delay

Factor	Level	Disruption	S1 unmet (%)	S2 unmet (%)	S3 unmet (%)	S2 Warning lead (days)
Outage availability (fraction)	0.1	D1	33.1	22.4	8.1	0
Outage availability (fraction)	0.2	D1	27.1	16.1	3.1	0
Outage availability (fraction)	0.3	D1	20.8	10.0	0.6	0
Outage availability (fraction)	0.5	D1	8.8	0.3	0.0	0
Outage availability (fraction)	0.7	D1	0.0	0.0	0.0	0
Outage duration (days)	30	D1	1.8	0.2	0.0	0
Outage duration (days)	60	D1	11.1	3.3	0.1	0
Outage duration (days)	90	D1	20.8	10.0	0.6	0
Outage duration (days)	150	D1	40.2	24.5	3.2	0
Surge multiplier (N95)	2	D2	12.5	4.0	0.0	−21
Surge multiplier (N95)	3	D2	23.4	14.1	2.2	−22
Surge multiplier (N95)	4	D2	36.3	25.7	6.6	−18
Surge multiplier (N95)	6	D2	52.2	44.6	18.9	−18
Hazard-feed lead (days; feed always fires)	0	D1	–	11.4	0.9	−11
Hazard-feed lead (days; feed always fires)	2	D1	–	11.0	0.9	−11
Hazard-feed lead (days; feed always fires)	4	D1	–	10.0	0.6	0
Hazard-feed lead (days; feed always fires)	7	D1	–	9.5	0.4	0
Hazard-feed lead (days; feed always fires)	14	D1	–	9.3	0.3	0
Hazard-feed probability	0.0	D1	–	13.1	–	−11
Hazard-feed probability	0.5	D1	–	10.7	–	−5
Hazard-feed probability	0.85	D1	–	10.0	–	0
Hazard-feed probability	1.0	D1	–	10.0	–	0
Census-index lead (days)	0	D2	–	26.1	7.1	−28
Census-index lead (days)	7	D2	–	25.6	6.7	−23
Census-index lead (days)	14	D2	–	25.7	6.6	−18
Census-index lead (days)	21	D2	–	25.4	6.7	−17
Quality-signal probability	0.0	D3	–	1.0	0.0	−15
Quality-signal probability	0.5	D3	–	0.4	0.0	−15
Quality-signal probability	1.0	D3	–	0.1	0.0	74
Posting lag (days)	3	D1 / D3	–	8.2 / 0.0	–	0 / −8
Posting lag (days)	10	D1 / D3	–	10.0 / 0.4	–	0 / −15
Posting lag (days)	20	D1 / D3	–	11.8 / 1.5	–	0 / −20
Posting lag (days)	30	D1 / D3	–	11.4 / 1.3	–	0 / −20
Forecast noise multiplier	1.0	D2 / D1	–	25.7 / 10.0	6.6 / 0.6	−18
Forecast noise multiplier	2.0	D2 / D1	–	25.5 / 9.2	6.5 / 0.4	−18
Forecast noise multiplier	3.0	D2 / D1	–	23.7 / 5.8	5.6 / 0.1	−12
Logistics delay (extra days)	5	D4	0.0	–	0.0	−2
Logistics delay (extra days)	10	D4	0.5	–	0.0	−2
Logistics delay (extra days)	20	D4	5.5	–	2.2	−1

Figure 7: Signals and their lead: the hazard-feed lead, the census-index lead, the quality signal and the public posting lag against unmet demand and warning lead



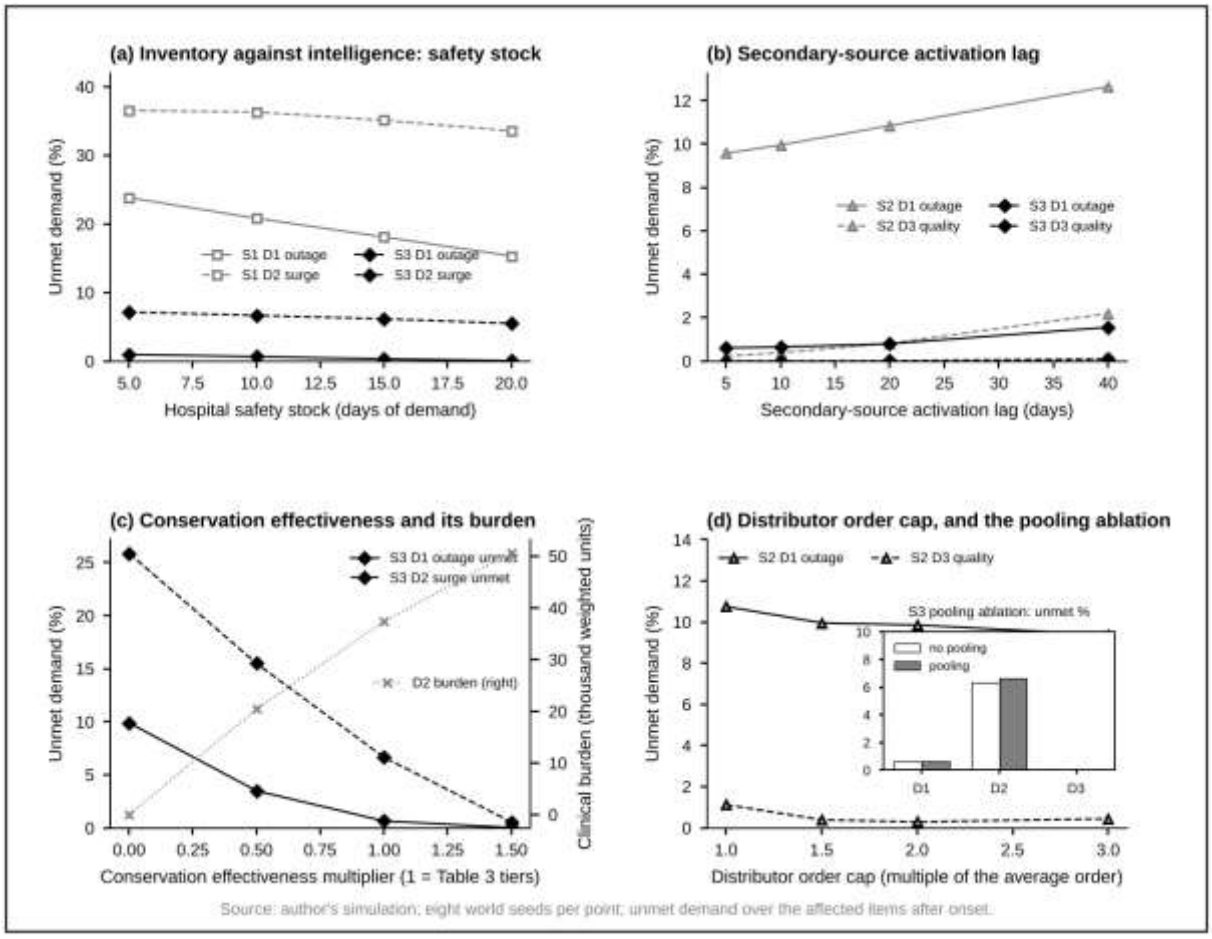
Mitigation levers

Table 13 and Figure 8 give the mitigation sweeps. The safety-stock sweep separates inventory from intelligence. Doubling every hospital's safety stock from 10 to 20 days under S1 reduced unmet demand in the outage from 20.8% to 15.3%, in the surge from 36.3% to 33.6% and in the shutdown from 16.3% to 10.2%, at a holding cost that rose from 53,000 to 100,000 USD per year; halving it to five days raised them to 23.8%, 36.5% and 18.9%. S3 with five days of safety stock (1.0%, 7.1% and 0.1%) outperformed S1 with twenty (15.3%, 33.6% and 10.2%) in every disruption at about half the holding cost. Inventory helps, and it helps S3 too (0.0% in the outage at 20 days), but it is a weak lever against a 90-day outage or a four-fold surge, which is the arithmetic that the stockpile debate after 2020 arrived at (Handfield et al., 2020; National Academies of Sciences, Engineering, and Medicine, 2022).

Table 13: Mitigation sweeps: unmet demand against hospital safety stock (S1 and S3), secondary-source lag, conservation effectiveness, the pooling ablation, the distributor order cap, the Warning budget and other customers' panic factor

Factor	Level	Strategies	D1 unmet (%)	D2 unmet (%)	D3 unmet (%)	Other
Hospital safety stock (days)	5	S1 / S3	23.8 / 1.0	36.5 / 7.1	18.9 / 0.1	holding 30 / 52 k
Hospital safety stock (days)	10	S1 / S3	20.8 / 0.6	36.3 / 6.6	16.3 / 0.0	holding 53 / 54 k
Hospital safety stock (days)	15	S1 / S3	18.1 / 0.3	35.1 / 6.1	13.3 / 0.0	holding 77 / 78 k
Hospital safety stock (days)	20	S1 / S3	15.3 / 0.0	33.6 / 5.5	10.2 / 0.0	holding 100 / 102 k
Secondary-source activation lag (days)	5	S2 / S3	9.6 / 0.6	–	0.2 / 0.0	–
Secondary-source activation lag (days)	10	S2 / S3	10.0 / 0.6	–	0.4 / 0.0	–
Secondary-source activation lag (days)	20	S2 / S3	10.8 / 0.8	–	0.8 / 0.0	–
Secondary-source activation lag (days)	40	S2 / S3	12.6 / 1.6	–	2.2 / 0.1	–
Conservation effectiveness (x tiers)	0.0	S3	9.9	25.9	–	D2 burden 0 thousand weighted units
Conservation effectiveness (x tiers)	0.5	S3	3.5	15.5	–	D2 burden 20 thousand weighted units
Conservation effectiveness (x tiers)	1.0	S3	0.6	6.6	–	D2 burden 37 thousand weighted units
Conservation effectiveness (x tiers)	1.5	S3	0.0	0.6	–	D2 burden 51 thousand weighted units
Pooling (allocation LP and transshipment)	off	S3	0.6	6.3	0.0	–
Pooling (allocation LP and transshipment)	on	S3	0.6	6.6	0.0	–
Distributor order cap (x average order)	1.0	S2	10.7	–	1.1	D1 Warning lead −7 d
Distributor order cap (x average order)	1.5	S2	10.0	–	0.4	D1 Warning lead 0 d
Distributor order cap (x average order)	2.0	S2	9.9	–	0.3	D1 Warning lead −6 d
Distributor order cap (x average order)	3.0	S2	9.4	–	0.4	D1 Warning lead −6 d
Warning budget (alerts per item-year)	0.25	S2	13.4	–	0.5	normal year: 0.27 warnings per item, spot 11 k
Warning budget (alerts per item-year)	0.5	S2	13.3	–	0.5	normal year: 0.29 warnings per item, spot 6 k
Warning budget (alerts per item-year)	1.0	S2	10.0	–	0.4	normal year: 1.00 warnings per item, spot 31 k
Warning budget (alerts per item-year)	2.0	S2	8.7	–	0.1	normal year: 1.62 warnings per item, spot 44 k
Warning budget (alerts per item-year)	4.0	S2	7.9	–	0.4	normal year: 3.17 warnings per item, spot 84 k
Other customers' panic factor	1.0	S2 / S3	9.1 / 0.5	–	–	–
Other customers' panic factor	1.3	S2 / S3	10.0 / 0.6	–	–	–
Other customers' panic factor	1.6	S2 / S3	10.4 / 0.7	–	–	–

Figure 8: Mitigation levers: safety stock for S1 and S3, secondary-source activation lag, conservation effectiveness and its burden, the distributor order cap and the pooling ablation



The secondary-source lag is the lever the health system controls before a disruption. With a five-day activation, S2's unmet demand was 9.6% in the outage and 0.2% in the shutdown; with 40 days it was 12.6% and 2.2%, and S3's rose from 0.6% to 1.6% and from 0.0% to 0.1%. The premium fell as the lag grew (86,000 to 66,000 USD in the outage) because less was delivered, which is the wrong way to save money. Pre-qualification, contracts and drills convert a 40-day lag into a 10-day one, worth two to three points of unmet demand in a long outage.

Conservation is the lever that decides the surge. With conservation off, S3 is S2 with better allocation: 9.9% unmet in the outage and 25.9% in the surge. At half the tiers' effectiveness the figures were 3.5% and 15.5%; at the tiers as set, 0.6% and 6.6%; at one and a half times, 0.0% and 0.6%, with the surge burden rising from zero to 51,000 weighted units. The curve is nearly linear in the surge, because the national allocation fixes supply and conservation is the only lever that moves demand. The pooling ablation, by contrast, changed almost nothing: with the allocation optimiser and transshipment switched off, S3's unmet demand was 0.6% in the outage (0.6% with them), 6.3% in the surge (6.6%) and 0.0% in the shutdown (0.0%), and stockout-days differed by less than a tenth. The modelled disruptions are system-wide: every hospital is short at the same time and in the same proportion, so there is little to move, and the proportional split already spreads the shortage evenly. Pooling earns its place against local shocks (the testbed's mass-casualty surges are too small and short to show in these totals) and uneven allocation; in a system-wide shortage it provides governance (who gets what, by an explicit rule) rather than throughput.

The distributor's order cap bounds what early warning can achieve, but weakly: with the cap at the average order (1.0), S2's unmet demand in the outage was 10.7%; at 1.5, 2.0 and 3.0 it was 10.0%, 9.9% and 9.4%. In a 90-day outage the pre-emptive build is a few days of supply however generous the cap, and S2's larger gains come from the secondary source. Other customers' panic factor (1.0 to 1.6) moved

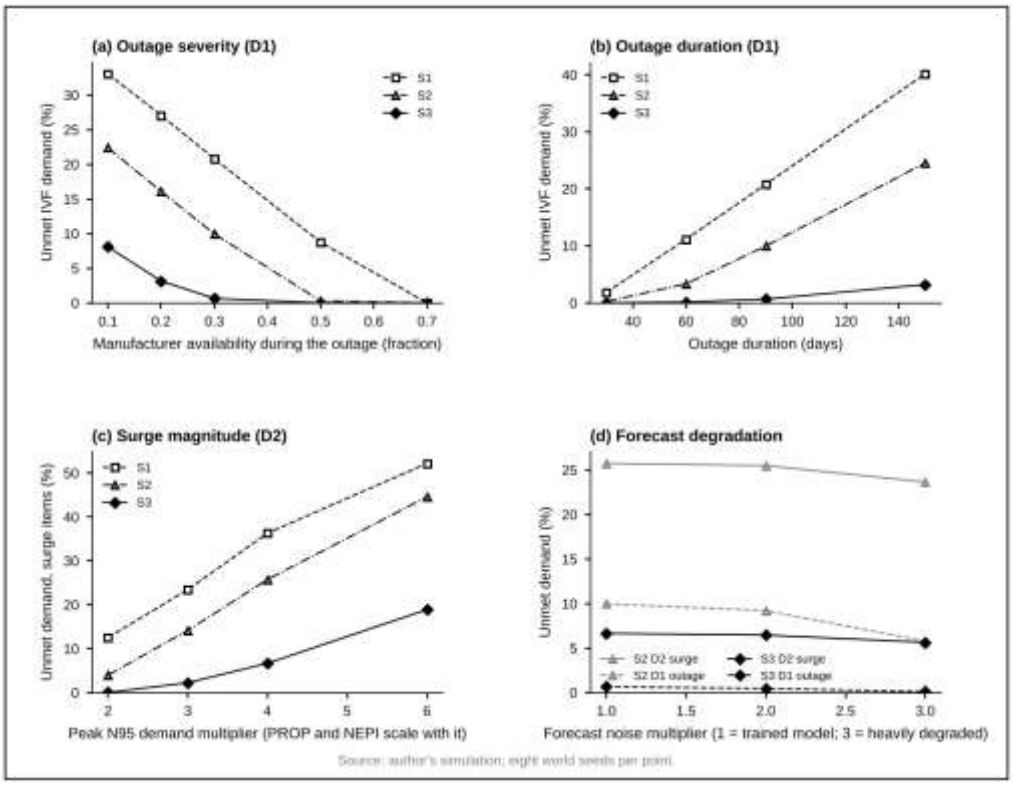
S2's unmet demand by 1.3 points (9.1% to 10.4%) and S3's by 0.2, because the distributor's allocation by history insulates the health system from others' hoarding once allocation begins, and before it begins the cap binds.

The Warning budget trades false alerts against response. With the thresholds set to a quarter of a Warning per item-year, S2 produced 0.27 false Warnings per item-year in normal years at a spot cost of 11,000 USD, and 13.4% unmet demand in the outage; at one per year, 1.0 false Warnings, 31,000 USD and 10.0%; at four per year, 3.2 false Warnings, 84,000 USD and about 10% more inventory, and 7.9% unmet demand with the median Warning three days before onset. The budget is the owner's choice; the sweep says that a looser budget buys earlier response at a cost that is visible every year, and that below one per year the hazard feed never reaches Warning.

Severity, duration and forecast quality

Figure 9 and Table 12 give the severity sweeps. Unmet demand in the outage is nearly linear in the lost supply and in its duration for every strategy, with different slopes: at 10% availability S1 reached 33.1%, S2 22.4% and S3 8.1%, and at 50% they reached 8.8%, 0.3% and 0.0%; at 150 days they reached 40.2%, 24.5% and 3.2%, and at 30 days 1.8%, 0.2% and 0.0%. S3's slope is about a fifth of S1's, which is the combined effect of conservation and secondary supply on the deficit per day, and the point at which the strategies separate is the point at which inventory runs out, about a month into an outage of this severity. The surge sweep shows the same geometry: at a doubling of demand S3 served everything (0.0%) while S1 lost 12.5%; at six-fold S3 lost 18.9% and S1 52.2%, with S3's burden at 74,000 weighted units. Degrading the forecaster by adding noise to its quantiles did not raise unmet demand; it lowered it slightly (S2 in the surge from 25.7% to 23.7%, S3 in the outage from 0.6% to 0.1%) while doubling the holding cost (58,000 to 110,000 USD) and the false Warnings (1.8 to 3.0 per item-year). Noisy quantiles act as a blanket safety-stock increase, and the sweep is a reminder that a forecaster is judged by the decisions it drives: a wider p90 buys service with inventory, which is not what a forecaster is for. The logistics delay was benign at five and ten extra days (S1 0.0% and 0.5%, S3 0.0%) and harmful at twenty (S1 5.5%, S3 2.2%), where the lead-time deviation signal fired at once but the pipeline simply ran dry.

Figure 9: Severity and forecast quality: outage availability and duration, surge magnitude and forecast degradation against unmet demand for S1, S2 and S3



IT resilience and the compound event

Table 14 and Figure 10 give the IT resilience sweeps. A cyber outage of the health system's ordering and visibility systems, on its own, was absorbed by inventory for up to two weeks under every strategy (no unmet demand at 7 or 14 days); at 28 days S1 and S3 lost 0.4% and 0.3% of demand with 80 and 64 stockout-days, and at 56 days 3.2% and 2.3% with about 500 and 400 stockout-days, while S4, whose runbook restores ordering in three days, lost nothing at any length.

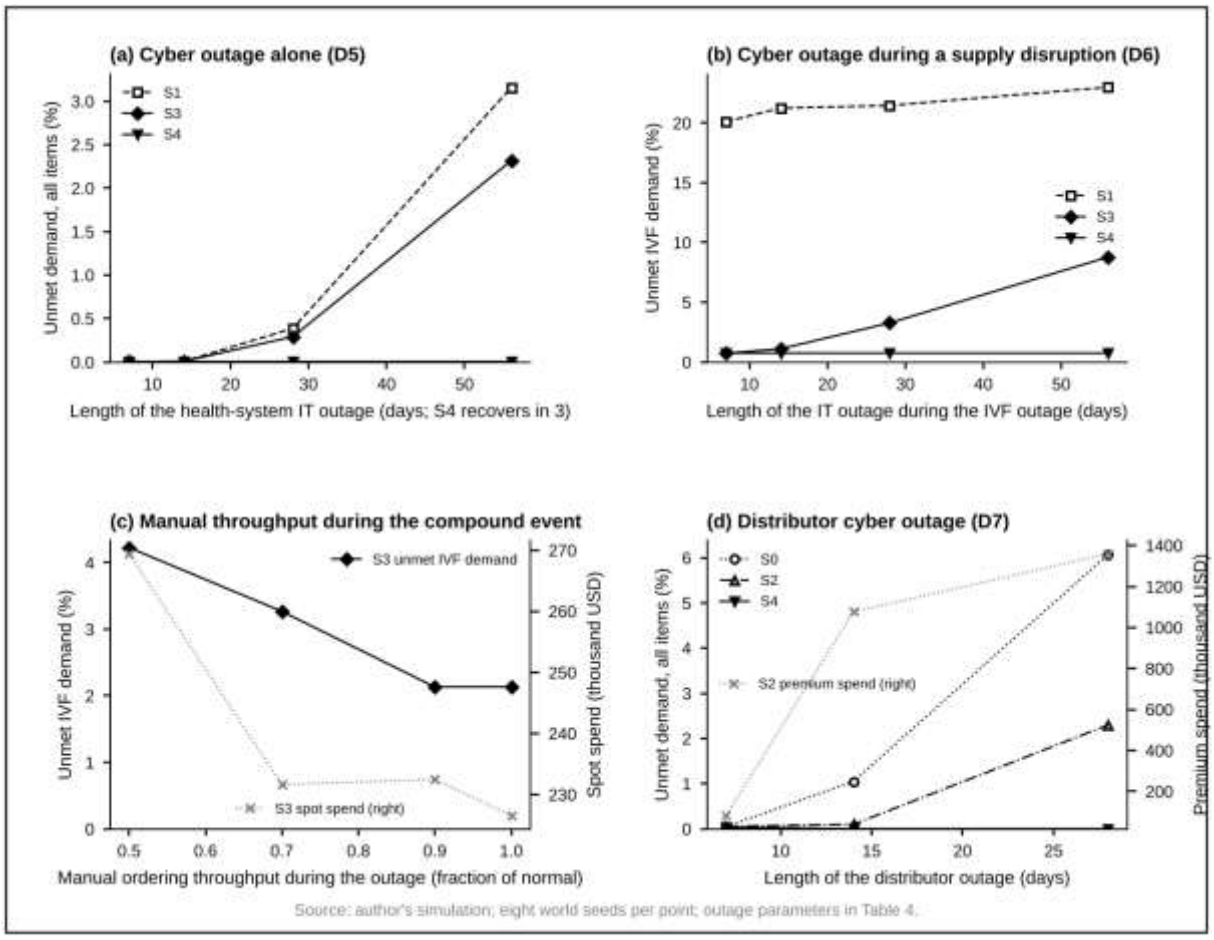
Table 14: IT resilience sweeps and the Monte Carlo summary: unmet demand against IT outage length (D5 and D6), manual throughput, distributor outage length, and the distribution of unmet demand over 200 draws of twelve factors

Factor	Level	Disruption	S1 unmet (%)	S3 unmet (%) [MC: S2]	S4 unmet (%)	Other (stockout-days S3 / S4; spend, thousand USD)
IT outage length (days)	7	D5 (all items)	0.0	0.0	0.0	0 / 0
	7	D6 (IVF)	20.1	0.7	0.7	38 / 39
IT outage length (days)	14	D5 (all items)	0.0	0.0	0.0	2 / 0
	14	D6 (IVF)	21.2	1.1	0.7	62 / 39
IT outage length (days)	28	D5 (all items)	0.4	0.3	0.0	64 / 0
	28	D6 (IVF)	21.5	3.3	0.7	222 / 39
IT outage length (days)	56	D5 (all items)	3.2	2.3	0.0	408 / 0
	56	D6 (IVF)	23.0	8.8	0.7	645 / 39
Manual throughput during the outage	0.5	D6 (IVF)	–	4.2	–	S3 stockout-days 401; spot 269 k
Manual throughput during the outage	0.7	D6 (IVF)	–	3.3	–	S3 stockout-days 222; spot 232 k
Manual throughput during the outage	0.9	D6 (IVF)	–	2.1	–	S3 stockout-days 121; spot 232 k
Manual throughput during the outage	1.0	D6 (IVF)	–	2.1	–	S3 stockout-days 110; spot 227 k
Distributor outage length (days)	7	D7 (all items)	S0: 0.0	S2: 0.1	0.0	S2 spot 81 k, secondary 1 k
Distributor outage length (days)	14	D7 (all items)	S0: 1.0	S2: 0.1	0.0	S2 spot 814 k, secondary 262 k
Distributor outage length (days)	28	D7 (all items)	S0: 6.1	S2: 2.3	0.0	S2 spot 1023 k, secondary 336 k
Monte Carlo (200 draws): median [10th, 90th]	D1 outage	unmet (%)	8.7 [0.0, 31.4]	0.2 [0.0, 18.0]	0.0 [0.0, 4.1]	S4 cost 377 k
Monte Carlo (200 draws): median [10th, 90th]	D2 surge	unmet (%)	32.3 [14.5, 49.6]	22.3 [6.8, 42.6]	5.2 [0.0, 18.5]	S4 cost 361 k
Monte Carlo (200 draws): median [10th, 90th]	D3 quality	unmet (%)	14.2 [9.9, 18.8]	0.0 [0.0, 0.5]	0.0 [0.0, 0.0]	S4 cost 172 k
Monte Carlo (200 draws): median [10th, 90th]	D6 compound	unmet (%)	11.6 [0.6, 33.0]	0.9 [0.0, 20.4]	0.0 [0.0, 4.0]	S4 cost 373 k

The costs were small (spot purchases of 45,000 to 59,000 USD at 28 to 56 days), and S0, which orders manually by design, was almost untouched. The picture changes when the outage coincides with a supply disruption. In the compound event S3's unmet intravenous-fluid demand rose from 0.7% with a 7-day outage to 1.1% at 14 days, 3.3% at 28 days and 8.8% at 56 days, with stockout-days rising from 38 to 645, because the stages froze at their pre-outage values, the forecasts reverted to trailing averages, pooling stopped and manual orders at 70% throughput could not keep pace with a shortage that was

deepening at the same time; S1 sat at 20 to 23% throughout, and S4 held at 0.7% at every length because its outage lasts three days whatever the attacker intended. Manual throughput during a 28-day outage moved S3's compound result from 4.2% at half throughput to 2.1% at full throughput (401 to 110 stockout-days), which says that a downtime procedure that keeps orders flowing at normal volume recovers about half of the harm, and the runbook the rest. The distributor outage behaves like a short, total supply cut: at seven days nothing was lost; at 14 days S0 lost 1.0% and S2 and S4 almost nothing; at 28 days S0 lost 6.1%, S2 2.3% and S4 nothing, because S4's alternate channel delivered 70% of normal volume from the third day. The price was high for every strategy that reacted: S2 spent 814,000 USD on spot purchases and 262,000 on secondary sources in the 14-day case and 1,023,000 and 336,000 in the 28-day case, against 871,000 and 263,000 for S4, whose alternate channel bought less at a premium. A shared intermediary's outage is a supply disruption with no warning, and the only cheap defence is an alternate channel agreed in advance.

Figure 10: IT resilience: the length of the health-system IT outage alone and during a supply disruption, manual throughput during the compound event, and the distributor outage length against unmet demand and spend



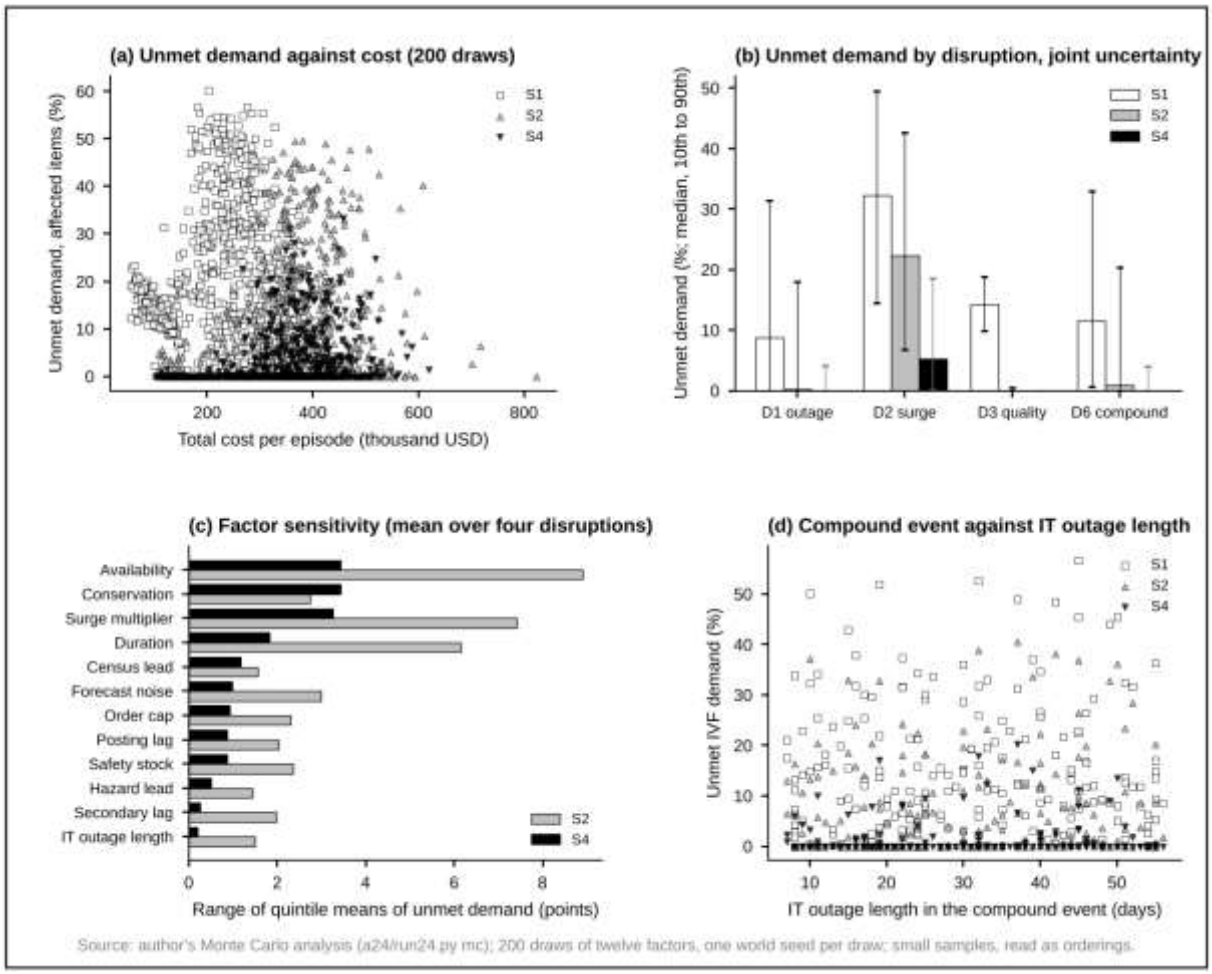
Monte Carlo analysis

The Monte Carlo analysis varied twelve factors jointly over 200 draws, each with its own world seed and recalibrated thresholds, for S1, S2 and S4 under the outage, the surge, the shutdown and the compound event (Table 14, Figure 11). The orderings of the comparison held under joint uncertainty. For the outage the median unmet demand was 8.7% for S1 (10th to 90th percentile 0.0 to 31.4%), 0.2% for S2 (0.0 to 18.0%) and 0.0% for S4 (0.0 to 4.1%), and the share of draws with more than 5% unmet demand fell from 61% to 32% and 9%. For the surge the medians were 32.3%, 22.3% and 5.2%, with 90th percentiles of 49.6%, 42.6% and 18.5%. The shutdown was almost always covered by S2 and S4 (medians 0.0%, 90th percentiles 0.5% and 0.0%) and never by S1 (median 14.2%). The compound event

gave medians of 11.6%, 0.9% and 0.0%, with S2's 90th percentile at 20.4% against S4's 4.0%, which is the IT resilience effect under joint uncertainty. Costs ran the other way: the median total cost per episode (holding, spot, secondary and transfers) was 313,000 USD for S1, 413,000 for S2 and 377,000 for S4 in the outage, and 218,000, 369,000 and 361,000 in the surge; S4 was cheaper than S2 because conservation reduced the volume bought at a premium, and S1 was cheapest because it bought nothing from secondary sources and ran short instead.

By the range of quintile means of mean unmet demand over the four disruptions, the factors that moved S1 most were the outage's availability (12.6 percentage points) and duration (10.7) and the surge multiplier (7.9), followed by safety stock (4.7); for S2 the same three severity factors led (8.9, 7.4 and 6.2), followed by forecast noise (3.0, in the direction already noted), conservation and safety stock; for S4 the ranges were small (3.4 for availability, 3.4 for conservation effectiveness, 3.3 for the surge multiplier, 1.8 for duration, and below 1.2 for everything else), because its levers absorb most of what the environment does. The hazard lead was at or near the bottom of every list and the census lead moved S2 and S4 by less than 1.6 points, which agrees with the one-factor sweeps, and the IT outage length, which the one-factor sweep showed to be decisive for S3 in the compound event, was masked in the joint analysis for S2 (quintile means between 4.4% and 8.3%, not monotone) by the severity factors drawn alongside it. These are rankings from one episode per draw and strategy, and they should be read as orderings; the one-factor sweeps, with eight seeds per point, carry the effect sizes.

Figure 11: Monte Carlo analysis over twelve factors: unmet demand against cost, unmet demand by disruption under joint uncertainty, factor sensitivities and the compound event against outage length



Discussion

Answers to the research questions

RQ1 asked how much earlier an AI-enabled layer anticipates a disruption and at what false-alert cost. The answer depends on the signal. Where a long-lead signal exists, as an inspection notice does for a quality shutdown, the scorer warned 43 days ahead in the median episode and 55% of shutdowns were anticipated; where the only lead is a weather alert, the layer gained days rather than weeks (Watch before onset in 40% of outages, Warning a median of 8 days after onset against 13 for rules), and the distributor's order cap limited what those days could buy; where the disruption is a demand surge, the forecaster recognised the growth through the census index three days before demand moved, but the binding constraint was national allocation. The false-alert cost at a budget of one Warning per item-year was about 25,000 USD of spot purchases and a third of a day of inventory per year, and loosening the budget to four bought a three-day median lead on the outage at 84,000 USD per year. Early warning is therefore worth building for the signals that lead by weeks; for those that lead by days it must be paired with the levers that convert a short lead into supply, a secondary source that can be activated at Watch and a distributor agreement on what may be ordered when a hazard is forecast.

RQ2 asked which levers reduce unmet demand most. In order of effect in the testbed: conservation (which decided the surge and removed most of S2's residual unmet demand in the outage), secondary sourcing (which made the quality shutdown almost invisible and halved the outage's unmet demand), pre-emptive building within the cap (worth a few days), and safety stock (a weak lever at any affordable level). The allocation optimiser and transshipment changed the totals little in system-wide shortages, and their value is governance and equity rather than throughput. The costs line up differently: conservation costs clinical burden, secondary sourcing costs a premium that falls when it is needed most (because less is delivered), pre-emptive building costs the distributor's goodwill and a little holding, and safety stock costs money every year. The ordering S0 < S1 < S2 < S3 held in every seed for the three supply disruptions, with S1's gain over S0 being mostly a fairer distribution of the same shortage, which is itself worth having.

RQ3 asked how IT resilience interacts with supply resilience. A cyber outage alone cost little in the testbed because inventory covered it; the same outage during a supply disruption cost the AI-enabled strategies most of their advantage, and the runbook recovered it. The compound event is the case that matters for a health system that has invested in analytics: its ordering, visibility and forecasting are precisely what a ransomware event removes, and a paper-based competitor (S0) was less affected by the outage itself, though far worse off overall. The length of the outage and the manual throughput during it were the factors that moved the compound result, which is why the roadmap puts downtime drills that measure ordering throughput beside the backups.

RQ4 asked about robustness. The one-factor sweeps and the Monte Carlo analysis agree on the orderings: outage severity and duration dominate what any strategy can do; the secondary-source lag, conservation effectiveness and the posting lag are the levers a health system can move; the hazard and census leads matter less than their proponents might hope; and forecast noise, in this implementation, bought service with inventory rather than losing it. These are rankings from small samples and a synthetic testbed, and they should be read as orderings.

Implications for practice

Four implications follow for U.S. health systems. First, the critical-item list and the equivalence table are the foundation: the levers that mattered (conservation tiers, secondary sources, allocation rules) are defined per item, and the professional guidance on shortage management already asks for this list (Fox & McLaughlin, 2018). Second, the distributor relationship is part of the architecture, not outside it: the order cap, the allocation formula and the data the distributor shares (acknowledgements, advance ship notices, fill rates) set the limits of early warning, and the sweeps show that a change in the cap or in the posting lag moves the outcome by more than a change in the health system's own forecasting. Third, conservation needs governance before it needs analytics: the burden it imposes is clinical, the decision to impose it is clinical, and the testbed's results depend on tiers that clinical leadership would have to approve in advance, as crisis standards of care frameworks recommend (Institute of Medicine, 2012; Hick et al., 2020). Fourth, the IT layer must degrade gracefully: cached forecasts, a manual allocation table and an alternate channel to the distributor are cheap, and the compound scenario shows what

they are worth.

Mapping to U.S. policy and guidance

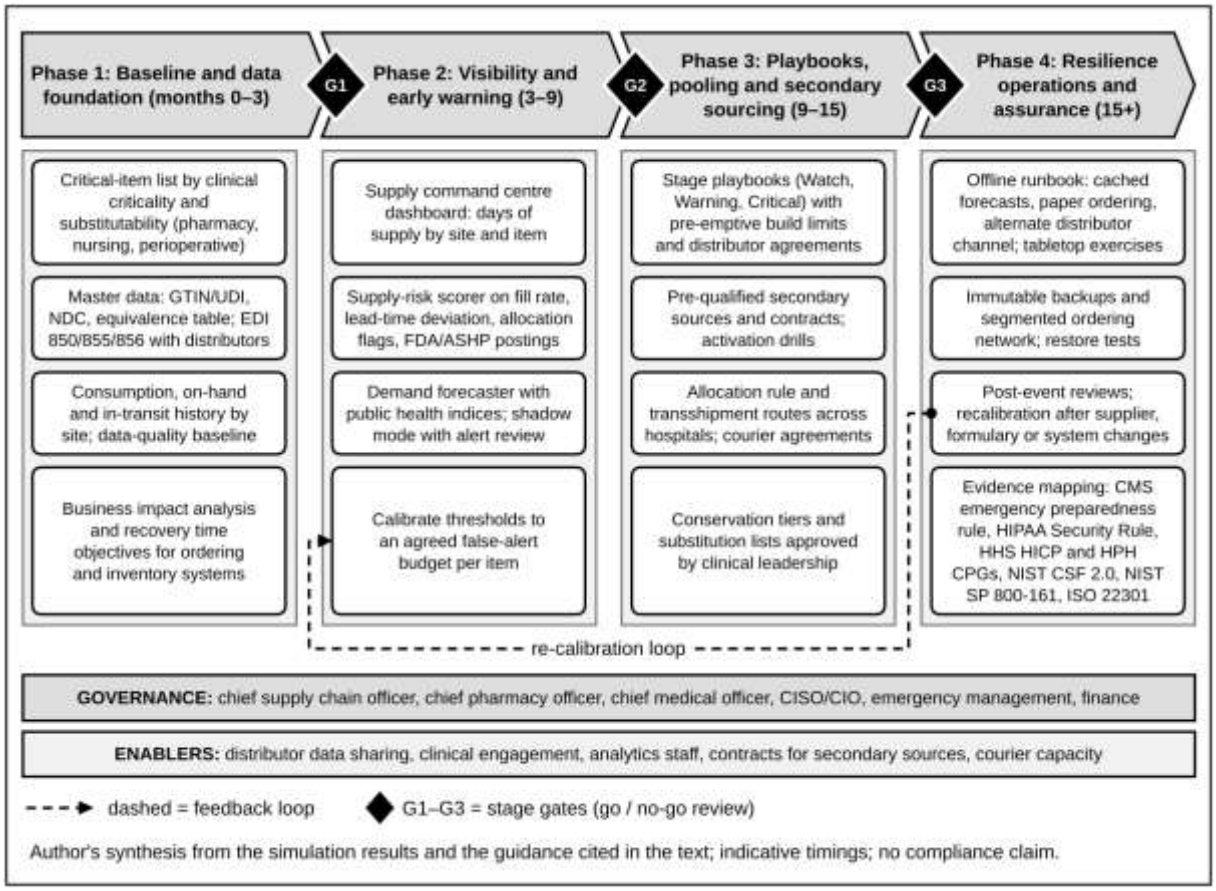
Table 15 maps the architecture to the federal instruments and guidance it touches. None of them requires a health system to build the layer described here; several reward it.

Table 15: Mapping of the architecture to U.S. policy, regulation and guidance (as of writing; not a compliance claim)

Instrument (as of writing; verify current status)	What it requires or provides	Where the architecture responds
FDASIA 2012, section 506C (Public Law 112-144)	Manufacturers notify FDA of discontinuances and interruptions of certain drugs	External postings feed (L1); the posting lag is a modelled signal
CARES Act 2020, sections 3112 and 3121 (Public Law 116-136)	Manufacturer risk management plans and expanded drug shortage reporting (section 506C); device shortage reporting (section 506J)	Drug and device postings feed; item-equivalence table for substitutes
Drug Supply Chain Security Act 2013 (Public Law 113-54, Title II)	Product identifiers and traceability for prescription drugs	Master data (NDC, serialised identifiers) in the integration layer
FDA unique device identification rule (78 FR 58786)	Standard device identifiers	GTIN/UDI master data; receipt verification by scan
Executive Orders 14001 and 14017; 100-day supply chain review (2021)	Public health supply chain strategy; pharmaceutical supply chain review and recommendations	Secondary sourcing and visibility align with the recommendations; no federal obligation on providers
NASEM 2022 report	Measures for medical product supply chain resilience across providers, distributors, manufacturers and government	Visibility, pooling and conservation playbooks at the provider level
CMS emergency preparedness rule (81 FR 63860; 42 CFR 482.15)	Risk assessment, policies, communication plan, training and testing for Medicare providers	Disruption catalogue as a hazard assessment; runbook exercises; event log
ASHP shortage guidelines (Fox & McLaughlin, 2018)	Shortage committee, critical-item list, inventory and allocation rules, conservation and substitution, communication	Stage playbooks and the clinical review board
HIPAA Security Rule (45 CFR 164 Subpart C)	Administrative, physical and technical safeguards; contingency plan standard	Backups, restores and the offline runbook; access control and logging
HHS 405(d) HICP (2023) and HPH cybersecurity performance goals (2024)	Voluntary practices and goals for the health sector, including backups, segmentation, MFA and incident planning	IT resilience controls panel
NIST CSF 2.0; NIST SP 800-34 Rev. 1; NIST SP 800-161 Rev. 1; NIST SP 800-160 Vol. 2	Cybersecurity outcomes, contingency planning, cyber supply chain risk management, cyber-resilient engineering	Recovery time objectives per service; vendor and distributor dependencies as supply chain risks
ISO 22301:2019	Business continuity management system requirements	Business impact analysis and continuity plans for ordering and inventory
NIST AI RMF 1.0	Governance and measurement of AI risks	Model monitoring, false-alert budget, calibration and shadow mode

The shortage-notification provisions of FDASIA and the CARES Act produce the postings the scorer consumes, and their lag is the posting lag that the sweep found to matter (U.S. Congress, 2012; U.S. Congress, 2020). The traceability and identification rules make the equivalence table and receipt verification possible (U.S. Congress, 2013; U.S. Food and Drug Administration, 2013). The CMS emergency preparedness rule's risk assessment, policies and exercises are the natural home for the disruption catalogue, the playbooks and the downtime drills (Centers for Medicare & Medicaid Services, 2016).

Figure 12: Implementation roadmap for AI-enabled supply and IT resilience in a U.S. health system: four phases with stage gates, governance and enablers, and the re-calibration loop



On the IT side, the HIPAA Security Rule's contingency plan standard, the 405(d) practices and the sector cybersecurity performance goals describe the backups, segmentation, multi-factor authentication and incident planning that the IT resilience controls implement, and the NIST frameworks give the structure for recovery time objectives and for treating the distributor and software vendors as supply chain risks in their own right (U.S. Department of Health and Human Services, 2013; U.S. Department of Health and Human Services, 2023; U.S. Department of Health and Human Services, 2024; National Institute of Standards and Technology, 2024; Swanson et al., 2010; Boyens et al., 2022). The AI components fall under the governance that the NIST AI Risk Management Framework asks for: documented purpose, measured performance, a false-alert budget, monitoring and a human decision at each stage (Tabassi, 2023). Federal supply chain policy since 2021 has focused on manufacturing, stockpiling and visibility at the national level (Executive Office of the President, 2021c; Executive Office of the President, 2021b; The White House, 2021; National Academies of Sciences, Engineering, and Medicine, 2022); the provider-level capabilities modelled here are the complement, and a national visibility system would supply them with better signals than the ones the testbed assumes.**Roadmap**

Figure 12 sets out a four-phase roadmap. Phase 1 (months 0 to 3) builds the foundation: the critical-item list, master data and EDI with distributors, consumption and inventory history by site, and a

business impact analysis with recovery time objectives for the ordering and inventory systems. Phase 2 (3 to 9) delivers visibility and early warning: the command centre dashboard, the supply-risk scorer, the demand forecaster in shadow mode and thresholds calibrated to an agreed false-alert budget. Phase 3 (9 to 15) adds the playbooks, pre-qualified secondary sources with activation drills, allocation rules and transshipment routes, and conservation tiers approved by clinical leadership. Phase 4 (15 months onward) runs resilience operations: the offline runbook, immutable backups and a segmented ordering network with restore tests, post-event reviews with recalibration, and evidence mapping to Table 15. The gates ask whether the data are good enough to alert on, whether the alerts are trusted enough to act on and whether the playbooks have been exercised. The sweeps suggest the order: secondary sources and conservation tiers delivered most of the value and depend on contracts and clinical governance rather than on models, so a health system with limited analytics capacity should not wait for Phase 2 to finish before starting them.

Limitations

Table 16 lists the limitations. The testbed is synthetic and its parameters are stylised; the numbers are properties of the testbed and the orderings are the findings. The distributor's behaviour is reduced to an allocation formula, an order cap and a panic factor, and these assumptions set the limits of early warning. Secondary sources and the spot market are fixed capacities and prices. The risk scorer is trained on signals produced by the same model that generates the test episodes, so its accuracy is an upper bound, and the degradation sweep showed that noise in the forecast quantiles buys service with inventory, so a fairer test of forecast quality would vary bias rather than noise. Conservation is a fixed share of demand with a burden weight, which compresses a clinical question into one number. The IT outage is summarised by assumed length, throughput, error and delay. There is no adaptive adversary and no second-order market effect beyond the panic factor. Samples are small. The policy mapping describes the position at the time of writing, from sources dated up to 2025, and is not a compliance claim.

Conclusion

This article asked what an AI-enabled early-warning and response layer, inside a resilient IT design, is worth to a U.S. health system facing the disruptions its supply chain has actually suffered, and it answered with a transparent synthetic testbed rather than with data from any health system. In the testbed, system-wide visibility alone distributed shortages more fairly without reducing them much; AI early warning with pre-emptive building and a pre-qualified secondary source halved the unmet demand of a 90-day manufacturing outage and almost eliminated that of a quality shutdown; conservation tiers and an explicit allocation rule reduced the outage's unmet demand to half a percent and the surge's from 38% to 7%; and an IT resilience runbook preserved these gains when a cyber outage struck during the supply disruption, the case in which the digitised strategies would otherwise have lost most of their advantage. The levers that mattered were the ones a health system can prepare in advance (secondary sources, conservation tiers, distributor agreements and downtime procedures), the signals that mattered were the ones that lead by weeks rather than days, and the false-alert budget was a choice with a visible annual cost rather than a free parameter.

The findings are orderings from a stylised model, and the honest next step is replication on real data: a health system's purchasing and consumption history, its distributor's acknowledgements and fill rates, and the postings and notices of past shortages would calibrate the signals, the allocation rules and the conservation tiers that this study assumed, and a prospective shadow-mode trial of the scorer and the forecaster would measure the false-alert cost directly. The architecture, workflow and roadmap are offered as a structure for that work. All code, synthetic data and results are in the supplementary archive, and every number can be reproduced from fixed seeds.

REFERENCES

[1]. Abdulsalam, Y., Gopalakrishnan, M., Maltz, A., & Schneller, E. (2015). Health care matters: Supply chains in and of the health sector. Journal of Business Logistics, 36(4), 335–339. https://doi.org/10.1111/jbl.12111

[2]. Administration for Strategic Preparedness and Response. (n.d.). Center for the Strategic National Stockpile. ASPR. https://www.aspr.gov/about/centers/sns

[3]. Adnan, S. (2024). Human Trust in AI-Assisted Project Decision Systems: A Mixed-Methods Analysis of Explainability, Adoption, and Managerial Decision Quality. Journal of Sustainable Development and Policy, 3(04), 287-332. https://doi.org/10.63125/9c139h44

[4]. American Society of Health-System Pharmacists. (2024). National drug shortages: January 2001 to June 2024. ASHP. https://www.ashp.org/drug-shortages/shortage-resources/drug-shortages-statistics

[5]. Aminul Islam, M. (2024). Interpreting Cyber-Threat Indicators in Enterprise Memory Forensics: A Mixed-Methods Study of Machine Learning Performance and Explainability. Review of Applied Science and Technology, 14(12), 01–13. https://doi.org/10.63125/2p8hte92

[6]. Aminul Islam, M., & Md, A. (2023). Adaptive Class-Imbalance Learning for High-Sensitivity Fraud Detection in Large-Scale Digital Transaction Networks. International Journal of Business and Economics Insights, 13(09), 01-33. https://doi.org/10.63125/bmxwv741

[7]. Asheke, M. (2023). Machine Learning–Based Predictive Quality Modeling for Defect Detection and Process Optimization in Advanced Manufacturing Systems. Review of Applied Science and Technology, 13(12), 01–49. https://doi.org/10.63125/thy73b16

[8]. Asheke, M., & Md Maksudul, H. (2022). AI-Driven Demand Forecasting for Sporting-Goods and Teamwear Supply Chains Under Seasonal and Tournament Demand. International Journal of Business and Economics Insights, 2(1), 73–110. https://doi.org/10.63125/ed6by557

[9]. Asheke, M., & Mohammad Badrul, A. (2022). Development Of Functionalized Cellulose Nanofiber Adsorbents for Selective Critical-Metal Recovery from Industrial Wastewater. American Journal of Scholarly Research and Innovation, 1(02), 270–320. https://doi.org/10.63125/v6fxzq09

[10]. Ashik, M. (2024). Meta-Analysis of Risk-Based Inspection Strategies for Improving Reliability and Safety Across Critical Energy Infrastructure Systems. American Journal of Advanced Technology and Engineering Solutions, 14(09), 01-43. https://doi.org/10.63125/tgttkt09

[11]. Ashik, M. (2025). PRISMA-Based Systematic Literature Review of Reliability-Centered Turnaround Planning for Reducing Downtime in Power Generation Facilities. American Journal of Scholarly Research and Innovation, 4(01), 1036–1080. https://doi.org/10.63125/2878n718

[12]. Axelsson, S. (2000). The base-rate fallacy and the difficulty of intrusion detection. ACM Transactions on Information and System Security, 3(3), 186–205. https://doi.org/10.1145/357830.357849

[13]. Ban, G.-Y., & Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. Operations Research, 67(1), 90–108. https://doi.org/10.1287/opre.2018.1757

[14]. Baryannis, G., Validi, S., Dani, S., & Antoniou, G. (2019). Supply chain risk management and artificial intelligence: State of the art and future research directions. International Journal of Production Research, 57(7), 2179–2202. https://doi.org/10.1080/00207543.2018.1530476

[15]. Baxter. (2025). Hurricane Helene summary. Baxter. https://www.baxter.com/baxter-newsroom/hurricane-helene-updates

[16]. Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. Management Science, 66(3), 1025–1044. https://doi.org/10.1287/mnsc.2018.3253

[17]. Bijvank, M., & Vis, I. F. A. (2012). Inventory control for point-of-use locations in hospitals. Journal of the Operational Research Society, 63(4), 497–510. https://doi.org/10.1057/jors.2011.52

[18]. Boyens, J., Smith, A., Bartol, N., Winkler, K., Holbrook, A., & Fallon, M. (2022). Cybersecurity supply chain risk management practices for systems and organizations (NIST Special Publication 800-161 Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-161r1

[19]. Bruneau, M., Chang, S. E., Eguchi, R. T., Lee, G. C., O'Rourke, T. D., Reinhorn, A. M., Shinozuka, M., Tierney, K., Wallace, W. A., & von Winterfeldt, D. (2003). A framework to quantitatively assess and enhance the seismic resilience of communities. Earthquake Spectra, 19(4), 733–752. https://doi.org/10.1193/1.1623497

[20]. Burns, L. R. (2002). The health care value chain: Producers, purchasers, and providers. Jossey-Bass. https://archive.org/details/healthcarevaluec00lawt

[21]. Cavallo, J. J., & Pahade, J. K. (2022). Practice management strategies for imaging facilities facing an acute iodinated contrast media shortage. American Journal of Roentgenology, 219(4), 666–670. https://doi.org/10.2214/AJR.22.27969

[22]. Centers for Medicare & Medicaid Services. (2016). Medicare and Medicaid programs; Emergency preparedness requirements for Medicare and Medicaid participating providers and suppliers; Final rule. Federal Register, 81 FR 63860. https://www.federalregister.gov/documents/2016/09/16/2016-21404/medicare-and-medicaid-programs-emergency-preparedness-requirements-for-medicare-and-medicaid

[23]. Chen, D. Q., Preston, D. S., & Xia, W. (2013). Enhancing hospital supply chain performance: A relational view and empirical test. Journal of Operations Management, 31(6), 391–408. https://doi.org/10.1016/j.jom.2013.07.012

[24]. Chopra, S., & Sodhi, M. S. (2004). Managing risk to avoid supply-chain breakdown. MIT Sloan Management Review, 46(1), 53–61. https://sloanreview.mit.edu/article/managing-risk-to-avoid-supplychain-breakdown/

[25]. Chopra, S., & Sodhi, M. S. (2014). Reducing the risk of supply chain disruptions. MIT Sloan Management Review, 55(3), 73–80. https://sloanreview.mit.edu/article/reducing-the-risk-of-supply-chain-disruptions/

[26]. Christopher, M., & Peck, H. (2004). Building the resilient supply chain. The International Journal of Logistics Management, 15(2), 1–14. https://doi.org/10.1108/09574090410700275

[27]. Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). Computer security incident handling guide (NIST Special Publication 800-61 Rev. 2). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r2

[28]. Clark, A. J., & Scarf, H. (1960). Optimal policies for a multi-echelon inventory problem. Management Science, 6(4), 475–490. https://doi.org/10.1287/mnsc.6.4.475

[29]. Cohen, J., & Rodgers, Y. v. d. M. (2020). Contributing factors to personal protective equipment shortages during the COVID-19 pandemic. Preventive Medicine, 141, Article 106263. https://doi.org/10.1016/j.ypmed.2020.106263

[30]. Craighead, C. W., Blackhurst, J., Rungtusanatham, M. J., & Handfield, R. B. (2007). The severity of supply chain disruptions: Design characteristics and mitigation capabilities. Decision Sciences, 38(1), 131–156. https://doi.org/10.1111/j.1540-5915.2007.00151.x

[31]. Craighead, C. W., Ketchen, D. J., Jr., & Darby, J. L. (2020). Pandemics and supply chain management research: Toward a theoretical toolbox. Decision Sciences, 51(4), 838–866. https://doi.org/10.1111/deci.12468

[32]. Cramer, E. Y., Ray, E. L., Lopez, V. K., Bracher, J., Brennen, A., Castro Rivadeneira, A. J., Gerding, A., Gneiting, T., House, K. H., Huang, Y., Jayawardena, D., Kanji, A. H., Khandelwal, A., Le, K., Mühlemann, A., Niemi, J., Shah, A., Stark, A., Wang, Y., ... Reich, N. G. (2022). Evaluation of individual and ensemble probabilistic forecasts of COVID-19 mortality in the United States. Proceedings of the National Academy of Sciences, 119(15), e2113561119. https://doi.org/10.1073/pnas.2113561119

[33]. Croston, J. D. (1972). Forecasting and stock control for intermittent demands. Operational Research Quarterly, 23(3), 289–303. https://doi.org/10.1057/jors.1972.50

[34]. Dameff, C., Tully, J., Chan, T. C., Castillo, E. M., Savage, S., Maysent, P., Hemmen, T. M., Clay, B. J., & Longhurst, C. A. (2023). Ransomware attack associated with disruptions at adjacent emergency departments in the US. JAMA Network Open, 6(5), e2312270. https://doi.org/10.1001/jamanetworkopen.2023.12270

[35]. 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. https://doi.org/10.1080/00207543.2017.1387680

[36]. Elmachtoub, A. N., & Grigas, P. (2022). Smart "predict, then optimize". Management Science, 68(1), 9–26. https://doi.org/10.1287/mnsc.2020.3922

[37]. Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., & Phillips, J. P. (2020). Fair allocation of scarce medical resources in the time of Covid-19. The New England Journal of Medicine, 382(21), 2049–2055. https://doi.org/10.1056/NEJMsb2005114

[38]. Eppen, G. D. (1979). Effects of centralization on expected costs in a multi-location newsboy problem. Management Science, 25(5), 498–501. https://doi.org/10.1287/mnsc.25.5.498

[39]. Executive Office of the President. (2021a). Executive Order 14028: Improving the nation's cybersecurity. Federal Register, 86 FR 26633. https://www.federalregister.gov/documents/2021/05/17/2021-10460/improving-the-nations-cybersecurity

[40]. Executive Office of the President. (2021b). Executive Order 14017: America's supply chains. Federal Register, 86 FR 11849. https://www.federalregister.gov/documents/2021/03/01/2021-04280/americas-supply-chains

[41]. Executive Office of the President. (2021c). Executive Order 14001: A sustainable public health supply chain. Federal Register, 86 FR 7219. https://www.federalregister.gov/documents/2021/01/26/2021-01865/a-sustainable-public-health-supply-chain

[42]. Federal Bureau of Investigation, Cybersecurity and Infrastructure Security Agency, & U.S. Department of Health and Human Services. (2023). #StopRansomware: ALPHV Blackcat (Joint Cybersecurity Advisory AA23-353A, updated February 27, 2024). CISA. https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-353a

[43]. Finkenstadt, D. J., & Handfield, R. (2021). Blurry vision: Supply chain visibility for personal protective equipment during COVID-19. Journal of Purchasing and Supply Management, 27(3), Article 100689. https://doi.org/10.1016/j.pursup.2021.100689

[44]. Fox, E. R., & McLaughlin, M. M. (2018). ASHP guidelines on managing drug product shortages. American Journal of Health-System Pharmacy, 75(21), 1742–1750. https://doi.org/10.2146/ajhp180441

[45]. Fox, E. R., Sweet, B. V., & Jensen, V. (2014). Drug shortages: A complex health care crisis. Mayo Clinic Proceedings, 89(3), 361–373. https://doi.org/10.1016/j.mayocp.2013.11.014

[46]. Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. Annals of Statistics, 29(5), 1189–1232. https://doi.org/10.1214/aos/1013203451

[47]. Ghafur, S., Kristensen, S., Honeyford, K., Martin, G., Darzi, A., & Aylin, P. (2019). A retrospective impact analysis of the WannaCry cyberattack on the NHS. npj Digital Medicine, 2(1), Article 98. https://doi.org/10.1038/s41746-019-0161-6

[48]. Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. Journal of the American Statistical Association, 102(477), 359–378. https://doi.org/10.1198/016214506000001437

[49]. Handfield, R., Finkenstadt, D. J., Schneller, E. S., Godfrey, A. B., & Guinto, P. (2020). A commons for a supply chain in the post-COVID-19 era: The case for a reformed strategic national stockpile. The Milbank Quarterly, 98(4), 1058–1090. https://doi.org/10.1111/1468-0009.12485

[50]. Hazen, B. T., Boone, C. A., Ezell, J. D., & Jones-Farmer, L. A. (2014). Data quality for data science, predictive analytics, and big data in supply chain management: An introduction to the problem and suggestions for research and applications. International Journal of Production Economics, 154, 72–80. https://doi.org/10.1016/j.ijpe.2014.04.018

[51]. Herer, Y. T., Tzur, M., & Yücesan, E. (2006). The multilocation transshipment problem. IIE Transactions, 38(3), 185–200. https://doi.org/10.1080/07408170500434539

[52]. Hick, J. L., Hanfling, D., Wynia, M. K., & Pavia, A. T. (2020). Duty to plan: Health care, crisis standards of care, and novel coronavirus SARS-CoV-2. NAM Perspectives. https://doi.org/10.31478/202003b

[53]. Hohenstein, N.-O., Feisel, E., Hartmann, E., & Giunipero, L. (2015). Research on the phenomenon of supply chain resilience: A systematic review and paths for further investigation. International Journal of Physical Distribution & Logistics Management, 45(1/2), 90–117. https://doi.org/10.1108/IJPDLM-05-2013-0128

[54]. Holling, C. S. (1973). Resilience and stability of ecological systems. Annual Review of Ecology and Systematics, 4, 1–23. https://doi.org/10.1146/annurev.es.04.110173.000245

[55]. 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. https://doi.org/10.1016/j.tre.2019.03.001

[56]. Hu, Q., Schwarz, L. B., & Uhan, N. A. (2012). The impact of group purchasing organizations on healthcare-product supply chains. Manufacturing & Service Operations Management, 14(1), 7–23. https://doi.org/10.1287/msom.1110.0355

[57]. Huangfu, Q., & Hall, J. A. J. (2018). Parallelizing the dual revised simplex method. Mathematical Programming Computation, 10(1), 119–142. https://doi.org/10.1007/s12532-017-0130-5

[58]. Hyndman, R. J., & Athanasopoulos, G. (2021). Forecasting: Principles and practice (3rd ed.). OTexts. https://otexts.com/fpp3/

[59]. Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. International Journal of Forecasting, 22(4), 679–688. https://doi.org/10.1016/j.ijforecast.2006.03.001

[60]. Institute of Medicine. (2012). Crisis standards of care: A systems framework for catastrophic disaster response. The National Academies Press. https://doi.org/10.17226/13351

[61]. International Organization for Standardization. (2019). ISO 22301:2019 Security and resilience: Business continuity management systems: Requirements. ISO. https://www.iso.org/standard/75106.html

[62]. 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, Article 101922. https://doi.org/10.1016/j.tre.2020.101922

[63]. Ivanov, D., & Dolgui, A. (2020). Viability of intertwined supply networks: Extending the supply chain resilience angles towards survivability. A position paper motivated by COVID-19 outbreak. International Journal of Production Research, 58(10), 2904–2915. https://doi.org/10.1080/00207543.2020.1750727

[64]. Ivanov, D., & Dolgui, A. (2021). A digital supply chain twin for managing the disruption risks and resilience in the era of Industry 4.0. Production Planning & Control, 32(9), 775–788. https://doi.org/10.1080/09537287.2020.1768450

[65]. Ivanov, D., Dolgui, A., & Sokolov, B. (2019). The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics. International Journal of Production Research, 57(3), 829–846. https://doi.org/10.1080/00207543.2018.1488086

[66]. Jalali, M. S., & Kaiser, J. P. (2018). Cybersecurity in hospitals: A systematic, organizational perspective. Journal of Medical Internet Research, 20(5), e10059. https://doi.org/10.2196/10059

[67]. Jia, J., & Zhao, H. (2017). Mitigating the U.S. drug shortages through Pareto-improving contracts. Production and Operations Management, 26(8), 1463–1480. https://doi.org/10.1111/poms.12697

[68]. Jones, S. S., Thomas, A., Evans, R. S., Welch, S. J., Haug, P. J., & Snow, G. L. (2008). Forecasting daily patient volumes in the emergency department. Academic Emergency Medicine, 15(2), 159–170. https://doi.org/10.1111/j.1553-2712.2007.00032.x

[69]. Kaakeh, R., Sweet, B. V., Reilly, C., Bush, C., DeLoach, S., Higgins, B., Clark, A. M., & Stevenson, J. (2011). Impact of drug shortages on U.S. health systems. American Journal of Health-System Pharmacy, 68(19), 1811–1819. https://doi.org/10.2146/ajhp110210

[70]. Kamalahmadi, M., & Parast, M. M. (2016). A review of the literature on the principles of enterprise and supply chain resilience: Major findings and directions for future research. International Journal of Production Economics, 171, 116–133. https://doi.org/10.1016/j.ijpe.2015.10.023

[71]. Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., & Liu, T.-Y. (2017). LightGBM: A highly efficient gradient boosting decision tree. Advances in Neural Information Processing Systems, 30, 3146–3154. https://papers.nips.cc/paper/6907-lightgbm-a-highly-efficient-gradient-boosting-decision-tree

[72]. Kelle, P., Woosley, J., & Schneider, H. (2012). Pharmaceutical supply chain specifics and inventory solutions for a hospital case. Operations Research for Health Care, 1(2-3), 54–63. https://doi.org/10.1016/j.orhc.2012.07.001

[73]. Kishimoto, T. K., Viswanathan, K., Ganguly, T., Elankumaran, S., Smith, S., Pelzer, K., Lansing, J. C., Sriranganathan, N., Zhao, G., Galcheva-Gargova, Z., Al-Hakim, A., Bailey, G. S., Fraser, B., Roy, S., Rogers-Cotrone, T., Buhse, L., Whary, M., Fox, J., Nasr, M., ... Sasisekharan, R. (2008). Contaminated heparin associated with adverse clinical events and activation of the contact system. The New England Journal of Medicine, 358(23), 2457–2467. https://doi.org/10.1056/NEJMoa0803200

[74]. Kleindorfer, P. R., & Saad, G. H. (2005). Managing disruption risks in supply chains. Production and Operations Management, 14(1), 53–68. https://doi.org/10.1111/j.1937-5956.2005.tb00009.x

[75]. Koenker, R., & Bassett, G., Jr. (1978). Regression quantiles. Econometrica, 46(1), 33–50. https://doi.org/10.2307/1913643

[76]. Kruse, C. S., Frederick, B., Jacobson, T., & Monticone, D. K. (2017). Cybersecurity in healthcare: A systematic review of modern threats and trends. Technology and Health Care, 25(1), 1–10. https://doi.org/10.3233/THC-161263

[77]. Kynatun, M. (2025). Integrating AI-Driven Security Analytics and IT Governance Frameworks for Cybersecurity Decision Support in U.S. Financial Information Systems. American Journal of Advanced Technology and Engineering Solutions, 15(12), 01-42. https://doi.org/10.63125/kzcdyz73

[78]. Kynatun, M., & Saba, A. (2024). Machine Learning Enhanced Intrusion Detection and Anomaly Classification for MIS-Driven Enterprise Network Security. International Journal of Scientific Interdisciplinary Research, 14(09), 01–41. https://doi.org/10.63125/xk6kpc72

[79]. Larsen, E., Fong, A., Wernz, C., & Ratwani, R. M. (2018). Implications of electronic health record downtime: An analysis of patient safety event reports. Journal of the American Medical Informatics Association, 25(2), 187–191. https://doi.org/10.1093/jamia/ocx057

[80]. Law, A. M. (2015). Simulation modeling and analysis (5th ed.). McGraw-Hill Education. https://www.mheducation.com/

[81]. Lazer, D., Kennedy, R., King, G., & Vespignani, A. (2014). The parable of Google Flu: Traps in big data analysis. Science, 343(6176), 1203–1205. https://doi.org/10.1126/science.1248506

[82]. Lee, H. L., Padmanabhan, V., & Whang, S. (1997). Information distortion in a supply chain: The bullwhip effect. Management Science, 43(4), 546–558. https://doi.org/10.1287/mnsc.43.4.546

[83]. Lücker, F., Seifert, R. W., & Biçer, I. (2019). Roles of inventory and reserve capacity in mitigating supply chain disruption risk. International Journal of Production Research, 57(4), 1238–1249. https://doi.org/10.1080/00207543.2018.1504173

[84]. Maishar Mostari, Z. (2025). Privacy-Preserving AI-Based Cyber Threat Detection in Financial Organizations: A Governance and Business Value Framework. American Journal of Interdisciplinary Studies, 15(04), 01-51. https://doi.org/10.63125/p8fct530

[85]. Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. International Journal of Forecasting, 38(4), 1346–1364. https://doi.org/10.1016/j.ijforecast.2021.11.013

[86]. Mansura Akter, E., & Sumaiya, M. (2022). Antimicrobial Resistance Profiles of Acinetobacter Baumannii and Vibrio Cholerae in Hospital Wastewater and Adjacent Household Water Sources. American Journal of Data Science and Analytics, 3(09), 01-55. https://doi.org/10.63125/gbervt14

[87]. Masrul, H. (2025). Explainable Machine Learning for Real-Time Financial Fraud Detection and Risk-Based Transaction Monitoring in U.S. Banking and Digital Payment Systems. American Journal of Data Science and Analytics, 6(10), 70-108. https://doi.org/10.63125/4h9v1876

[88]. Mazer-Amirshahi, M., & Fox, E. R. (2018). Saline shortages: Many causes, no simple solution. The New England Journal of Medicine, 378(16), 1472–1474. https://doi.org/10.1056/NEJMp1800347

[89]. McLaughlin, M., Kotis, D., Thomson, K., Harrison, M., Fennessy, G., Postelnick, M., & Scheetz, M. H. (2013). Effects on patient care caused by drug shortages: A survey. Journal of Managed Care Pharmacy, 19(9), 783–788. https://doi.org/10.18553/jmcp.2013.19.9.783

[90]. Md Imran, H., & Jannatul, F. (2022). Molecular and Precision Rehabilitation Molecular Mechanisms Linking Mitochondrial Dysfunction, Inflammation, and Muscle Degeneration with Functional Decline and Rehabilitation Response in Neuromuscular Disease. American Journal of Interdisciplinary Studies, 12(08), 01-44. https://doi.org/10.63125/0ym5bh79

[91]. Md Maksudul, H. (2024). Supplier-Risk Scoring and Procurement Resilience for Distributed Athletic-Apparel and Equipment Supply Networks. American Journal of Advanced Technology and Engineering Solutions, 14(09), 44-90. https://doi.org/10.63125/a9fmgj50

[92]. Md Maksudul, H. (2025). Data-Driven Fulfillment Optimization for Time-Sensitive Teamwear and Sporting Goods Distribution. American Journal of Scholarly Research and Innovation, 4(01), 1120–1173. https://doi.org/10.63125/rpezfz61

[93]. Md Maksudul, H., & Md Tazul Islam, C. (2022). Machine Learning-Based Cash-Flow Forecasting for Improving Financial Planning in U.S. Small and Mid-Sized Businesses. International Journal of Scientific Interdisciplinary Research, 12(12), 01–38. https://doi.org/10.63125/w3pgjx97

[94]. Md Momen Shahriar, J. (2025). Viscoelastic Creep Characterization of LDPE Under Multilevel Stress and Temperature Using Time-Temperature Superposition and Constitutive Modeling. American Journal of Scholarly Research and Innovation, 15(11), 01–38. https://doi.org/10.63125/2edstd86

[95]. Md Shahid, N. (2023). Digital Twin-Based Network Reliability Modeling for LAN And WAN Infrastructure Resilience. International Journal of Scientific Interdisciplinary Research, 4(4), 577–629. https://doi.org/10.63125/ahjvxj14

[96]. Md Shahid, N. (2024). Predictive Analytics Models for Network Traffic Optimization in Distributed Enterprise Environments. Review of Applied Science and Technology, 3(01), 367-419. https://doi.org/10.63125/qm6cyf84

[97]. Md Shahid, N. (2025). Blockchain-Based Secure Data Routing Framework for U.S. Industrial Communication Networks. American Journal of Scholarly Research and Innovation, 4(01), 990–1035. https://doi.org/10.63125/9exrgs41

[98]. Md Tazul Islam, C. (2024). Explainable AI For Financial and Operational Decision Support: Improving Transparency and Managerial Trust. American Journal of Advanced Technology and Engineering Solutions, 14(12), 01-38. https://doi.org/10.63125/kwfvzw68

[99]. Md Tazul Islam, C. (2025). Business Intelligence and Data-Quality Frameworks for Improving Evidence-Based Decisions across Multi-Source Enterprise Data. American Journal of Scholarly Research and Innovation, 15(07), 01–41. https://doi.org/10.63125/qhh6ka59

[100]. Md Tazul Islam, C., & Ahmed, H. C. H. R. (2023). Predictive Operations Analytics for Resource Allocation and Service-Capacity Planning in Data-Driven Organizations. Review of Applied Science and Technology, 13(02), 01–41. https://doi.org/10.63125/z3sm3089

[101]. Md. Morshedul, I., & Aminul Islam, M. (2022). A Quantitative Evaluation of Explainable Ensemble Learning for Probabilistic Risk Estimation and Decision Reliability in Enterprise Analytics. American Journal of Data Science and Analytics, 3(02), 01-38. https://doi.org/10.63125/k22shk74

[102]. Mirchandani, P. (2020). Health care supply chains: COVID-19 challenges and pressing actions. Annals of Internal Medicine, 173(4), 300–301. https://doi.org/10.7326/M20-1326

[103]. National Academies of Sciences, Engineering, and Medicine. (2022). Building resilience into the nation's medical product supply chains. The National Academies Press. https://doi.org/10.17226/26420

[104]. National Institute of Standards and Technology. (2024). The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.CSWP.29

[105]. Neprash, H. T., McGlave, C. C., Cross, D. A., Virnig, B. A., Puskarich, M. A., Huling, J. D., Rozenshtein, A. Z., & Nikpay, S. S. (2022). Trends in ransomware attacks on US hospitals, clinics, and other health care delivery organizations, 2016-2021. JAMA Health Forum, 3(12), e224873. https://doi.org/10.1001/jamahealthforum.2022.4873

[106]. Nicholson, L., Vakharia, A. J., & Erenguc, S. S. (2004). Outsourcing inventory management decisions in healthcare: Models and application. European Journal of Operational Research, 154(1), 271–290. https://doi.org/10.1016/S0377-2217(02)00700-2

[107]. Niculescu-Mizil, A., & Caruana, R. (2005). Predicting good probabilities with supervised learning. In Proceedings of the 22nd International Conference on Machine Learning (pp. 625–632). ACM. https://doi.org/10.1145/1102351.1102430

[108]. Oz, H., Aris, A., Levi, A., & Uluagac, A. S. (2022). A survey on ransomware: Evolution, taxonomy, and defense solutions. ACM Computing Surveys, 54(11s), Article 238. https://doi.org/10.1145/3514229

[109]. Page, E. S. (1954). Continuous inspection schemes. Biometrika, 41(1-2), 100–115. https://doi.org/10.1093/biomet/41.1-2.100

[110]. Pall, R., Gauthier, Y., Auer, S., & Mowaswes, W. (2023). Predicting drug shortages using pharmacy data and machine learning. Health Care Management Science, 26(3), 395–411. https://doi.org/10.1007/s10729-022-09627-y

[111]. Paterson, C., Kiesmüller, G., Teunter, R., & Glazebrook, K. (2011). Inventory models with lateral transshipments: A review. European Journal of Operational Research, 210(2), 125–136. https://doi.org/10.1016/j.ejor.2010.05.048

[112]. Persad, G., Wertheimer, A., & Emanuel, E. J. (2009). Principles for allocation of scarce medical interventions. The Lancet, 373(9661), 423–431. https://doi.org/10.1016/S0140-6736(09)60137-9

[113]. Pettit, T. J., Croxton, K. L., & Fiksel, J. (2013). Ensuring supply chain resilience: Development and implementation of an assessment tool. Journal of Business Logistics, 34(1), 46–76. https://doi.org/10.1111/jbl.12009

[114]. Pettit, T. J., Fiksel, J., & Croxton, K. L. (2010). Ensuring supply chain resilience: Development of a conceptual framework. Journal of Business Logistics, 31(1), 1–21. https://doi.org/10.1002/j.2158-1592.2010.tb00125.x

[115]. Phuong, J. M., Penm, J., Chaar, B., Oldfield, L. D., & Moles, R. (2019). The impacts of medication shortages on patient outcomes: A scoping review. PLOS ONE, 14(5), e0215837. https://doi.org/10.1371/journal.pone.0215837

[116]. Ponomarov, S. Y., & Holcomb, M. C. (2009). Understanding the concept of supply chain resilience. The International Journal of Logistics Management, 20(1), 124–143. https://doi.org/10.1108/09574090910954873

[117]. Ranney, M. L., Griffeth, V., & Jha, A. K. (2020). Critical supply shortages: The need for ventilators and personal protective equipment during the Covid-19 pandemic. The New England Journal of Medicine, 382(18), e41. https://doi.org/10.1056/NEJMp2006141

[118]. Reich, N. G., Brooks, L. C., Fox, S. J., Kandula, S., McGowan, C. J., Moore, E., Osthus, D., Ray, E. L., Tushar, A., Yamana, T. K., Biggerstaff, M., Johansson, M. A., Rosenfeld, R., & Shaman, J. (2019). A collaborative multiyear, multimodel assessment of seasonal influenza forecasting in the United States. Proceedings of the National Academy of Sciences, 116(8), 3146–3154. https://doi.org/10.1073/pnas.1812594116

[119]. Rosales, C. R., Magazine, M., & Rao, U. (2015). The 2Bin system for controlling medical supplies at point-of-use. European Journal of Operational Research, 243(1), 271–280. https://doi.org/10.1016/j.ejor.2014.10.041

[120]. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

[121]. Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., & McQuaid, R. (2021). Developing cyber-resilient systems: A systems security engineering approach (NIST Special Publication 800-160, Vol. 2, Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-160v2r1

[122]. Sacks, C. A., Kesselheim, A. S., & Fralick, M. (2018). The shortage of normal saline in the wake of Hurricane Maria. JAMA Internal Medicine, 178(7), 885–886. https://doi.org/10.1001/jamainternmed.2018.1936

[123]. Sadia, A. (2025). Privacy-Aware Personalization: Machine Learning, Model Validation, And Consumer Trust in AI-Mediated Marketing. American Journal of Scholarly Research and Innovation, 4(01), 1081–1119. https://doi.org/10.63125/j7tfr242

[124]. Sadia, A., & Md. Mosheur, R. (2023). Consumer Trust, Perceived Authenticity, and Personalized Digital Marketing: A Quantitative Decision-Science Study. International Journal of Scientific Interdisciplinary Research, 4(2), 252–295. https://doi.org/10.63125/cajkqt60

[125]. Saedi, S., Kundakcioglu, O. E., & Henry, A. C. (2016). Mitigating the impact of drug shortages for a healthcare facility: An inventory management approach. European Journal of Operational Research, 251(1), 107–123. https://doi.org/10.1016/j.ejor.2015.11.017

[126]. Sargent, R. G. (2013). Verification and validation of simulation models. Journal of Simulation, 7(1), 12–24. https://doi.org/10.1057/jos.2012.20

[127]. 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. https://doi.org/10.1016/j.ijpe.2012.01.004

[128]. Schneller, E. S., & Smeltzer, L. R. (2006). Strategic management of the health care supply chain. Jossey-Bass. https://archive.org/details/strategicmanagem0000schn

[129]. Shabbir, A. (2025). Physics-Based Failure Prognostics for Preventing Pressure-System Degradation and Unplanned Industrial Shutdowns. American Journal of Scholarly Research and Innovation, 14(09), 01–41. https://doi.org/10.63125/3asd0053

[130]. Sheffi, Y., & Rice, J. B., Jr. (2005). A supply chain view of the resilient enterprise. MIT Sloan Management Review, 47(1), 41–48. https://sloanreview.mit.edu/article/a-supply-chain-view-of-the-resilient-enterprise/

[131]. Silver, E. A., Pyke, D. F., & Thomas, D. J. (2016). Inventory and production management in supply chains (4th ed.). CRC Press. https://doi.org/10.1201/9781315374406

[132]. Simchi-Levi, D., Schmidt, W., & Wei, Y. (2014). From superstorms to factory fires: Managing unpredictable supply-chain disruptions. Harvard Business Review, 92(1-2). https://hbr.org/2014/01/from-superstorms-to-factory-fires-managing-unpredictable-supply-chain-disruptions

[133]. Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., Gusikhin, O., Sanders, M., & Zhang, D. (2015). Identifying risks and mitigating disruptions in the automotive supply chain. Interfaces, 45(5), 375–390. https://doi.org/10.1287/inte.2015.0804

[134]. Sittig, D. F., Gonzalez, D., & Singh, H. (2014). Contingency planning for electronic health record-based care continuity: A survey of recommended practices. International Journal of Medical Informatics, 83(11), 797–804. https://doi.org/10.1016/j.ijmedinf.2014.07.007

[135]. Snyder, L. V., Atan, Z., Peng, P., Rong, Y., Schmitt, A. J., & Sinsoysal, B. (2016). OR/MS models for supply chain disruptions: A review. IIE Transactions, 48(2), 89–109. https://doi.org/10.1080/0740817X.2015.1067735

[136]. Sodhi, M. S., & Tang, C. S. (2021). Supply chain management for extreme conditions: Research opportunities. Journal of Supply Chain Management, 57(1), 7–16. https://doi.org/10.1111/jscm.12255

[137]. Spieske, A., Gebhardt, M., Kopyto, M., & Birkel, H. (2022). Improving resilience of the healthcare supply chain in a pandemic: Evidence from Europe during the COVID-19 crisis. Journal of Purchasing and Supply Management, 28(5), Article 100748. https://doi.org/10.1016/j.pursup.2022.100748

[138]. Stabak, D., & Md Shahid, N. (2022). Empirical Assessment of Uncertainty-Aware Real-Time Decision-Making for Reliable and Resilient Cyber-Physical Systems. International Journal of Scientific Interdisciplinary Research, 1(01), 432–477. https://doi.org/10.63125/q3r4kk16

[139]. Sterman, J. D. (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. Management Science, 35(3), 321–339. https://doi.org/10.1287/mnsc.35.3.321

[140]. Sumaiya Sadmeen, M. (2025). Spatiotemporal Graph Neural Networks and Digital Twins for Urban Drainage Networks: Predictive Inundation, Critical-Node Detection, And Infrastructure Prioritization. American Journal of Scholarly Research and Innovation, 15(12), 01-40. https://doi.org/10.63125/awvjew66

[141]. Sumaiya, M. (2024). Hospital Wastewater as A Reservoir of Antimicrobial Resistance Genes: Implications for Public Health and Genomic Surveillance in The United States. Review of Applied Science and Technology, 3(04), 448–508. https://doi.org/10.63125/ahvas976

[142]. Swanson, M., Bowen, P., Phillips, A. W., Gallup, D., & Lynes, D. (2010). Contingency planning guide for federal information systems (NIST Special Publication 800-34 Rev. 1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-34r1

[143]. Syed Feroz, A., & Md Mahamud Pervaz, P. (2022). Root-Cause Analysis and FMEA-Driven Reliability Improvement in Gas Production Facilities: A Multi-Year Empirical Study. International Journal of Scientific Interdisciplinary Research, 12(10), 01–53. https://doi.org/10.63125/70msp878

[144]. Syed Feroz, A., & Shabbir, A. (2022). Finite Element-Based Thermomechanical Optimization of Pressure Components for Extending Industrial Boiler Service Life. Journal of Sustainable Development and Policy, 12(12), 01-51. https://doi.org/10.63125/x24pht28

[145]. Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. International Journal of Forecasting, 21(2), 303–314. https://doi.org/10.1016/j.ijforecast.2004.10.001

[146]. Tabassi, E. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

[147]. Taher, M. A., Sikder, M. A., Ahmed, H., & Bashar, M. (2024). Integrating IOT and Big Data Analytics for Enhanced Supply Chain Performance in Industrial Engineering Sectors: A Cross-Market Study. International Journal of Science and Engineering, Volume 1, 1-14.

[148]. Tang, C. S. (2006). Perspectives in supply chain risk management. International Journal of Production Economics, 103(2), 451–488. https://doi.org/10.1016/j.ijpe.2005.12.006

[149]. The White House. (2021). Building resilient supply chains, revitalizing American manufacturing, and fostering broad-based growth: 100-day reviews under Executive Order 14017. The White House. https://bidenwhitehouse.archives.gov/wp-content/uploads/2021/06/100-day-supply-chain-review-report.pdf

[150]. Tomlin, B. (2006). On the value of mitigation and contingency strategies for managing supply chain disruption risks. Management Science, 52(5), 639–657. https://doi.org/10.1287/mnsc.1060.0515

[151]. Tucker, E. L., Cao, Y., Fox, E. R., & Sweet, B. V. (2020a). The drug shortage era: A scoping review of the literature 2001-2019. Clinical Pharmacology & Therapeutics, 108(6), 1150–1155. https://doi.org/10.1002/cpt.1934

[152]. Tucker, E. L., Daskin, M. S., Sweet, B. V., & Hopp, W. J. (2020b). Incentivizing resilient supply chain design to prevent drug shortages: Policy analysis using two- and multi-stage stochastic programs. IISE Transactions, 52(4), 394–412. https://doi.org/10.1080/24725854.2019.1646441

[153]. U.S. Congress. (2012). Food and Drug Administration Safety and Innovation Act (Public Law 112-144). U.S. Government Publishing Office. https://www.govinfo.gov/app/details/PLAW-112publ144

[154]. U.S. Congress. (2013). Drug Quality and Security Act (Public Law 113-54), Title II: Drug Supply Chain Security Act. U.S. Government Publishing Office. https://www.govinfo.gov/app/details/PLAW-113publ54

[155]. U.S. Congress. (2020). Coronavirus Aid, Relief, and Economic Security Act (Public Law 116-136). U.S. Government Publishing Office. https://www.govinfo.gov/app/details/PLAW-116publ136

[156]. U.S. Department of Health and Human Services, Office for Civil Rights. (2024). Change Healthcare cybersecurity incident frequently asked questions. HHS. https://www.hhs.gov/hipaa/for-professionals/special-topics/change-healthcare-cybersecurity-incident-frequently-asked-questions/index.html

[157]. U.S. Department of Health and Human Services. (2013). Security standards for the protection of electronic protected health information (45 CFR Part 164, Subpart C, as amended). Electronic Code of Federal Regulations. https://www.ecfr.gov/current/title-45/subtitle-A/subchapter-C/part-164/subpart-C

[158]. U.S. Department of Health and Human Services. (2023). Health industry cybersecurity practices: Managing threats and protecting patients (2023 edition). HHS 405(d) Program. https://hhscyber.hhs.gov/Documents/HICP/HICP-Main-508.pdf

[159]. U.S. Department of Health and Human Services. (2024). Healthcare and public health sector-specific cybersecurity performance goals. HHS. https://hhscyber.hhs.gov/Documents/cybersecurity-performance-goals.pdf

[160]. U.S. Food and Drug Administration. (2013). Unique device identification system; Final rule. Federal Register, 78 FR 58786. https://www.federalregister.gov/documents/2013/09/24/2013-23059/unique-device-identification-system

[161]. U.S. Food and Drug Administration. (2020). Drug shortages: Root causes and potential solutions. FDA. https://www.fda.gov/drugs/drug-shortages/report-drug-shortages-root-causes-and-potential-solutions

[162]. U.S. Government Accountability Office. (2016). Drug shortages: Certain factors are strongly associated with this persistent public health challenge (GAO-16-595). GAO. https://www.gao.gov/products/gao-16-595

[163]. Vail, E., Gershengorn, H. B., Hua, M., Walkey, A. J., Rubenfeld, G., & Wunsch, H. (2017). Association between US norepinephrine shortage and mortality among patients with septic shock. JAMA, 317(14), 1433–1442. https://doi.org/10.1001/jama.2017.2841

[164]. Volland, J., Fügener, A., Schoenfelder, J., & Brunner, J. O. (2017). Material logistics in hospitals: A literature review. Omega, 69, 82–101. https://doi.org/10.1016/j.omega.2016.08.004

[165]. Woodcock, J., & Wosinska, M. (2013). Economic and technological drivers of generic sterile injectable drug shortages. Clinical Pharmacology & Therapeutics, 93(2), 170–176. https://doi.org/10.1038/clpt.2012.220

[166]. Yurukoglu, A., Liebman, E., & Ridley, D. B. (2017). The role of government reimbursement in drug shortages. American Economic Journal: Economic Policy, 9(2), 348–382. https://doi.org/10.1257/pol.20160035

[167]. Zamiela, C., Hossain, N. U. I., & Jaradat, R. (2022). Enablers of resilience in the healthcare supply chain: A case study of U.S healthcare industry during COVID-19 pandemic. Research in Transportation Economics, 93, Article 101174. https://doi.org/10.1016/j.retrec.2021.101174