

Multimodal Intelligent Assessment and Dynamic Regulation of Financial and Tax Resilience in Manufacturing Enterprises under Supply Chain Shocks

C.W. Shen^{1*}, W. Yan²

¹Anhui Institute of International Business, Hefei, Anhui, China

²Chuzhou University; Chuzhou, Anhui, China

Abstract

Supply chain disruptions and cyberattacks can jointly amplify production, liquidity, and tax risks in manufacturing enterprises, whereas conventional static scoring methods struggle to translate heterogeneous evidence into security responses. This paper proposes an artificial intelligence (AI)-driven framework that couples a sparse multimodal graph sequential state model (SMGSSM) with multi-constraint model predictive control (MC-MPC). Four neural encoders apply reliability gating to financial and tax data, operational data, cybersecurity disclosures, and Form 10-K semantics. A sparse exposure graph and a gated recurrent unit (GRU) jointly capture interfirm propagation and quarterly temporal dependence, while multitask output heads predict next-quarter resilience and three categories of risk. The control layer allocates liquidity support, supply substitution, cybersecurity hardening, and tax actions subject to budget and adjustment constraints. The key methodological contribution is threefold: separating evidence availability from predictive relevance through reliability gating; jointly modeling interfirm exposure and quarterly memory; and converting predicted cyber-financial-tax states into auditable actions under a shared resource budget. Using a public panel of 24 U.S. manufacturing firms, the calendar-year 2025 holdout root mean squared error (RMSE), after full-sample preprocessing, is 0.0663 ± 0.0089 and increases by 12.1% when graph propagation is removed. Across 120 closed-loop paths, MC-MPC reduces aggregate loss by 34.42% and increases average resilience by 35.90% relative to no control. Compared with reactive rules, it reduces cyber exposure by 41.76%, although production loss and control cost are higher. Within the stated retrospective and simulation-based design, the findings illustrate how cybersecurity can be represented as a controllable component of enterprise resilience and how a shared budget reveals cross-objective crowding-out. They do not establish prospective predictive validity or causal intervention effects.

Keywords: artificial intelligence, industrial cybersecurity, supply chain shocks, financial and tax resilience, multimodal deep learning, graph neural networks, model predictive control

Received on 10 August 2026, accepted on 16 September 2026, published on 28 September 2026

Copyright © 2026 C.W. Shen *et al.*, licensed to EAI. This is an open access article distributed under the terms of the [CC BY-NC-SA 4.0](#), which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

doi: 10.4108/eetsis.14411

*Corresponding author. Email: 2014221@ahuib.edu.cn

1. Introduction

As supply coordination, industrial control, financial settlement, and tax documentation become increasingly networked, manufacturing response has shifted from isolated equipment or inventory adjustments toward coordinated regulation across information, material, and capital flows.

When supply disruptions coincide with ransomware, data tampering, or access-control failures, compromised information chains can delay alternative sourcing, capacity reallocation, and financial and tax decisions. The resulting output decline further compresses cash buffers and

compliance capacity. Nikookar et al. developed an antifragile supply chain capability blueprint that extends resilience beyond post-shock recovery to capability reconfiguration and post-disruption growth [1]. Examining compounding geopolitical disruptions, Belhadi et al. showed that digital collaboration is essential for managing uncertainty and building supply chain resilience [2]. Liu et al. further linked market access, supply chain resilience, and enterprise innovation [3]. Yang et al. used a graph neural network to predict post-hazard firm-level sales, demonstrating that interfirm relational structures contain valuable information for assessing recovery after shocks [4]. These studies have endowed manufacturing resilience with capability-based, relational, and dynamic dimensions, but they have not yet incorporated cyberattack-induced information distortion and changes in financial and tax buffers into a unified state.

For complex shocks, research has moved from a single recovery-time measure toward multiscale resilience metrics and dynamic scenario simulation. Bruckler et al. systematically reviewed metrics for supply chain resilience capacities, actions, and outcomes, showing that no single performance indicator can fully represent the disruption-response process [5]. Ouyang et al. distinguished static topology from cascading-failure dynamics, demonstrating that supply-network risk depends on propagation paths and time-varying states and cannot be inferred from node connectivity alone [6]. Rashid et al. examined the effects of cloud adoption and artificial intelligence on manufacturing supply chains and concluded that digital technologies can improve supply visibility and resilience management [7]. Malekzadeh et al. developed a supply chain digital twin that considers both immediate and long-term disruptions, showing that virtual state evolution can provide a continuous vehicle for manufacturing scenario analysis [8]. Existing work therefore provides a foundation for networked, temporal, and scenario-based modeling. However, most state representations still center on materials, orders, and output, leaving the transmission structure among cyber exposure, liquidity, and tax stability unclear.

The objective of resilience regulation has likewise shifted from "recovering as quickly as possible" toward an integrated trade-off among value, cost, and risk. Zheng et al. proposed a process model of supply chain risk management under geopolitical disruptions, emphasizing that risk mitigation should not be separated from value creation [9]. Lückert et al. systematically reviewed the resilience-efficiency balance in overcoming supply chain disruptions and identified a structural trade-off between buffering capacity and operational efficiency [10]. From an inventory-management perspective, Guo et al. showed that inventory buffers absorb disruptions but also create capital lock-up and allocation costs [11]. Gervais et al. traced disruption propagation in photovoltaic supply scenarios, showing that diversification changes risk paths and loss distributions [12]. Li et al. analyzed how digital transformation affects manufacturing supply chain resilience and found information integration and collaboration capabilities to be important channels through which digitalization shapes firm recovery [13]. However, most of this literature remains focused on associations among

capabilities, strategies, and performance. It therefore cannot prescribe coordinated actions that simultaneously satisfy security, supply, and financial boundaries for each period-specific firm state.

The effects of supply shocks on manufacturing systems extend beyond output and propagate through cash flow, financing constraints, and tax decisions. Ivanov modeled cash-flow dynamics during and after disruptions and showed that financial consequences exhibit substantial time lags [14]. Qiu et al. studied contagious stock-price crashes along supply chains, demonstrating that interfirm relationships can transmit local operational risk to capital markets [15]. Chen et al. examined tax behavior from the perspective of corporate digital transformation and showed that digital capabilities alter tax decision-making [16]. Belahouaoui and Attak reviewed digital taxation, artificial intelligence (AI), and Tax Administration 3.0, concluding that AI is expanding from tax-risk identification to compliance interaction [17]. Di et al. linked supplier-base concentration to corporate cash policy, showing that supply relationships influence liquidity-protection strategies [18]. Financial and tax resilience are therefore integral components of shock propagation. Yet existing supply chain models often reduce them to a single cost term and struggle to represent the joint evolution of liquidity, tax volatility, and production pressure.

Data-driven research on financial and tax risk provides a more concrete foundation of variables and methods for constructing enterprise states. Xu and Kong used a random forest to identify corporate income-tax risk, showing that machine learning can extract nonlinear risk patterns from multiple financial and tax features [19]. Sun et al. further analyzed the mechanisms and policy implications through which digital transformation curbs corporate tax avoidance [20]. Dong and Zhang linked climate risk to tax behavior among climate-sensitive firms, indicating that external risk exposure may alter corporate tax strategies [21]. Song et al. examined corporate financialization, digital transformation, and industrial supply chain resilience, revealing interactions among financial decisions, research and development (R&D) investment, and industrial-chain recovery [22]. Wu and Zheng analyzed tax incentives, supply chain spillovers, and enterprise technological innovation, showing that tax-policy effects can cross firm boundaries [23]. Dodd and Liao investigated the role of inventory in firm resilience to the coronavirus disease 2019 (COVID-19) shock, further demonstrating the need to examine operational buffers jointly with financial outcomes [24]. Nevertheless, these studies generally focus on a single tax risk, governance effect, or financial outcome. They have not integrated supply propagation, cybersecurity exposure, and quarterly manufacturing states into a unified representation that is both predictable and controllable.

Advances in industrial artificial intelligence offer a new technical path for connecting heterogeneous evidence, network relationships, and temporal states. Moosavi et al. reviewed explainable AI in manufacturing and industrial cyber-physical systems and stressed that industrial intelligence must preserve the traceability of both predictions and their evidentiary sources [25]. Zhao et al. combined a

graph attention network with Informer for multivariate time-series anomaly detection, demonstrating the complementarity of relational structure and long-term temporal dependence [26]. Qu et al. developed a multimodal fusion network for industrial anomaly detection, showing that joint representations of heterogeneous industrial signals can recover state information missed by individual channels [27]. Alsaif et al. introduced multimodal large language models into Industry 4.0 fault detection and diagnosis, indicating that textual semantics can support risk identification alongside industrial states [28]. However, these methods primarily target anomaly or fault detection, and cybersecurity often remains an isolated alert channel. Alert probabilities neither co-evolve with enterprises' financial, tax, and production states nor translate into response actions constrained by manufacturing resources.

Taken together, existing studies leave three key issues unresolved in AI-driven assessment and regulation of manufacturing resilience. First, quarterly financial, tax, and operational facts, annual 10-K text, and cybersecurity disclosures differ in frequency, dimensionality, and missingness patterns; direct concatenation conflates evidence availability with predictive importance. Second, most models process firm records independently and struggle to jointly represent relational exposure in public materials, quarterly memory, and the shared effects of cyberattacks on liquidity, tax, and production pressure. Third, state prediction and managerial action are typically separated. Unless security hardening shares a budget with liquidity support, supply substitution, and tax coordination, resource crowding between cyber protection and immediate manufacturing recovery cannot be evaluated.

To address these issues, this paper proposes an AI-based assessment and dynamic regulation framework that combines a sparse multimodal graph sequential state model (SMGSSM) with multi-constraint model predictive control (MC-MPC). The model first assigns dedicated encoders to financial and tax data, operational data, cybersecurity disclosures, and Form 10-K semantics, using explicit reliability gates to distinguish evidence missingness from feature effects. It then uses a sparse public exposure graph to aggregate information from related firms and a gated recurrent unit (GRU) to retain temporal states for shocks not yet fully reflected in current records. The shared representation jointly predicts next-quarter resilience and liquidity, tax, and production risks. Finally, cyber exposure is incorporated into a five-dimensional controlled state, and MC-MPC jointly selects liquidity, supply-substitution, cybersecurity-hardening, and

tax actions subject to constraints on budgets, action intensity, and adjustment rates. Relative to existing work, this study contributes a reliability-gated representation of mixed-frequency multimodal evidence, a graph-recurrent assessment model jointly driven by relational exposure and quarterly memory, and a mechanism that converts predicted states into auditable constrained actions. The framework evaluates its predictive component on a public panel from the U.S. Securities and Exchange Commission (SEC) Electronic Data Gathering, Analysis, and Retrieval (EDGAR) system and tests closed-loop regulation and resource trade-offs under shared shock paths. The resulting conclusions are limited to the calendar-year holdout design after full-sample preprocessing and the control assumptions stated herein.

2. Intelligent Assessment and Dynamic Regulation Model for Financial and Tax Resilience

2.1. Manufacturing-Response-Oriented Multimodal Cybersecurity Representation

SMGSSM organizes its inputs around the effects of cyberattacks on manufacturing responses. Financial and tax, operational, cybersecurity, and textual-semantic data represent capital buffers, supply-demand pressure, attack exposure, and the broader business context, respectively. Modality-specific encoders preserve their respective scales, reliability gates suppress missing or low-quality channels, and the fused state supports both prediction and subsequent security actions.

Let $V = \{1, \dots, N\}$ denote the set of manufacturing enterprises and $t = 1, \dots, T$ the quarterly index. Each firm-quarter observation contains four public-data modalities: financial and tax facts F , operational facts O , cybersecurity disclosures C , and 10-K textual semantics T . The complete observation is

$$\mathcal{X}_{i,t} = \{x_{i,t}^F \in \mathbb{R}^6, x_{i,t}^O \in \mathbb{R}^5, x_{i,t}^C \in \mathbb{R}^5, x_{i,t}^T \in \mathbb{R}^4\}.(1)$$

The financial/tax and operational channels consist of ratios and growth variables; the cybersecurity channel measures the density per 10,000 words of five categories of security expressions; and the textual channel retains four singular value decomposition (SVD) components from the Form 10-K term frequency-inverse document frequency (TF-IDF) matrix. The meanings of the variables are reported in Table 1.

Table 1. Multimodal variables and their roles in resilience assessment

Modality	Variables	State information represented
Financial and tax (F)	Cash/assets, operating-cash-flow ratio, current ratio, debt/assets, effective tax rate, tax-rate volatility	Liquidity absorption, debt service, tax stability
Operations (O)	Revenue growth, inventory growth, inventory/revenue, receivables/revenue, operating margin	Demand and supply pressure, inventory mismatch, recovery capacity

Modality	Variables	State information represented
Cybersecurity (C)	Threat, breach, ransomware, security-control, and cyber-supply interaction densities	Attack exposure, control preparedness, continuity risk
Text semantics (T)	Four SVD components of 10-K TF-IDF features	Broader risk, supply, regulatory, and business context

The supervised targets must be constructed from observable public facts rather than undisclosed proprietary ratings. For component k , let $a_k \in \{-1, 1\}$ indicate whether an increase in the raw value is favorable. The median and interquartile range calculated from available observations yield the bounded robust component score

$$v_{i,t}^k = \Pi_{[0,1]} \left[\frac{1}{2} + \frac{a_k(f_{i,t}^k - \text{med}(f^k))}{2\text{IQR}(f^k) + \epsilon} \right]. \quad (2)$$

Favorable components include cash/assets, the operating-cash-flow ratio, the current ratio, and operating margin; the debt/assets ratio enters with the opposite sign. Tax stability and inventory stability are represented by the negatives of the absolute quarterly changes in the effective tax rate and inventory growth, respectively. Given the set of available components $\mathcal{K}_{i,t}$, the contemporaneous resilience proxy is defined as

$$r_{i,t} = \frac{\sum_{k \in \mathcal{K}_{i,t}} \omega_k v_{i,t}^k}{\sum_{k \in \mathcal{K}_{i,t}} \omega_k}, \quad \omega = (0.22, 0.20, 0.16, 0.15, 0.11, 0.09, 0.07). \quad (3)$$

The prediction target is $r_{i,t+1}$. Three binary output heads identify next-quarter liquidity, tax, and production risks. If $Q_{0.25}^{cash}$ denotes the lower quartile of cash capacity and $Q_{0.75}^{tax}$ the upper quartile of the absolute tax-rate change, the labels are summarized as

$$y_{i,t+1} = \begin{cases} \mathbb{I}(cash_{i,t+1} < Q_{0.25}^{cash} \vee ocf_{i,t+1} < 0) \\ \mathbb{I}(|\Delta tax_{i,t+1}| > Q_{0.75}^{tax} \vee |tax_{i,t+1}| > 0.60) \\ \mathbb{I}(g_{i,t+1}^{rev} < -0.10 \vee (g_{i,t+1}^{inv} > 0.20 \wedge g_{i,t+1}^{rev} < 0)) \end{cases}. \quad (4)$$

The multitask structure uses continuous resilience as its primary task and liquidity, tax, and production risks as auxiliary classification tasks. Through the shared hidden state, the cybersecurity encoding affects both regression and early warning, enabling the control layer to identify specific directions of pressure in addition to overall resilience. Figure 1 shows the four firm-quarter feature blocks, their dedicated encoders, and reliability gates. Next-quarter labels lie on a separate supervision plane and do not participate in forward feature construction.




Figure 1. Reliability-gated construction of model features and separate next-quarter supervision from public materials

Forward inputs contain only modality tensors, reliability vectors, and the adjacency matrix; targets and annual masks connect only to the loss and evaluation sides, thereby separating the inference path from training supervision.

2.2. Exposure-Graph Recurrent Deep Network for Resilience Assessment

Graph learning can supplement relational signals omitted from firms' tabular records [29], while temporal-anomaly research requires historical dependence, calibration, and distribution shift to be distinguished [30]. Accordingly, SMGSSM sequentially performs modality encoding, graph message passing, gated recurrent unit (GRU) state updating, and multitask output, while retaining an independent ablation interface for each module.

All continuous features are standardized using means and standard deviations estimated solely from 2020-2023 training observations. Missing values are set to 0 after standardization, so 0 corresponds to the training mean rather than an arbitrary raw value:

$$\tilde{x}_{i,t,d}^m = \begin{cases} (x_{i,t,d}^m - \mu_{d,tr}^m) / (\sigma_{d,tr}^m + \epsilon), & x_{i,t,d}^m \text{ observed,} \\ 0, & x_{i,t,d}^m \text{ missing.} \end{cases} \quad (5)$$

For structured modalities, reliability equals the proportion of observable underlying raw facts before ratio construction. Let the set of raw fields be $\mathcal{D}_m$; then

$$q_{i,t}^m = \frac{1}{|\mathcal{D}_m|} \sum_{d \in \mathcal{D}_m} \mathbb{I}(x_{i,t,d}^m \text{ observed}), \quad m \in \{F, O\}. \quad (6)$$

For cybersecurity and semantic text, a filing is considered valid when its main body exceeds 1,000 words. A small nonzero lower bound preserves the channel identifier without treating missing text as reliable evidence:

$$q_{i,t}^m = \begin{cases} 1, & \text{words}_{i,t} > 1000, \\ 0.05, & \text{otherwise,} \end{cases} \quad m \in \{C, T\}. \quad (7)$$

Each modality uses an independent encoder comprising a linear layer, ReLU activation, and layer normalization ϕ_m . The four encoders do not share first-layer parameters and can therefore learn the different scales of accounting ratios, operational changes, cybersecurity term frequencies, and latent semantic components. Their 10-dimensional outputs are multiplied by reliability scores:

$$z_{i,t}^m = q_{i,t}^m \phi_m(\tilde{x}_{i,t}^m). \quad (8)$$

The four gated encodings are concatenated directly without learning an additional attention score. Here, $q_{i,t}^m$ represents only evidence availability, while neural weights learn predictive relevance, preventing "missing filing" from being misinterpreted as "unimportant modality":

$$z_{i,t}^{cat} = z_{i,t}^F \parallel z_{i,t}^O \parallel z_{i,t}^C \parallel z_{i,t}^T. \quad (9)$$

A fully connected fusion layer then compresses the 40-dimensional concatenated representation into a 28-dimensional local firm embedding, using ReLU to form nonlinear feature interactions and a dropout rate of 0.15 to suppress overfitting in the small sample:

$$z_{i,t} = \text{Drop}_{0.15}(\text{ReLU}(W_f z_{i,t}^{cat} + b_f)). \quad (10)$$

Relational exposure is constructed from multiyear average 10-K TF-IDF vectors. For firms $i \neq j$, the unnormalized similarity combines document cosine similarity with a two-digit Standard Industrial Classification (SIC) increment:

$$e_{ij} = \cos(\bar{d}_i, \bar{d}_j) + 0.15 \mathbb{I}(SIC_{i,2} = SIC_{j,2}). \quad (11)$$

Each node retains its five highest-weight non-self neighbors, adds a self-retention value of 0.15, and is then row-normalized:

$$A_{ij} = \frac{\mathbb{I}(j \in \text{Top5}_i) \max(e_{ij}, 0) + 0.15 \mathbb{I}(i=j)}{\sum_k [\mathbb{I}(k \in \text{Top5}_i) \max(e_{ik}, 0) + 0.15 \mathbb{I}(i=k)]}. \quad (12)$$

Generative relationship completion [31], neuro-symbolic link prediction [32], and knowledge-graph-language-model fusion [33] can recover relationships from fragmented text, but public 10-K filings cannot verify private contracts. Therefore, A represents only a public-material exposure proxy, and the graph layer performs directed weighted aggregation of neighboring embeddings:

$$g_{i,t} = \text{ReLU}(W_g \sum_{j=1}^N A_{ij} z_{j,t} + b_g). \quad (13)$$

The local and graph representations are concatenated into a 56-dimensional graph-conditioned input and fed into a 28-dimensional GRU. This structure does not process each quarter independently: the previous-quarter hidden state $h_{i,t-1}$ retains lagged information about shocks not yet fully reflected in financial reports. Let $c_{i,t} = z_{i,t} \parallel g_{i,t}$; the reset gate is

$$a_{i,t} = \sigma(W_a c_{i,t} + U_a h_{i,t-1} + b_a). \quad (14)$$

The update gate is

$$b_{i,t} = \sigma(W_b c_{i,t} + U_b h_{i,t-1} + b_b). \quad (15)$$

The candidate state is

$$\tilde{h}_{i,t} = \tanh(W_h c_{i,t} + U_h (a_{i,t} \odot h_{i,t-1}) + b_h). \quad (16)$$

The recurrent state is updated as

$$h_{i,t} = (1 - b_{i,t}) \odot h_{i,t-1} + b_{i,t} \odot \tilde{h}_{i,t}. \quad (17)$$

The resilience output head maps the shared state to a bounded next-quarter prediction:

$$\hat{r}_{i,t+1} = \sigma(w_r^T \text{Drop}_{0.15}(h_{i,t}) + b_r). \quad (18)$$

The three risk logits are transformed in parallel into probabilities:

$$\hat{p}_{i,t+1} = \sigma(W_p \text{Drop}_{0.15}(h_{i,t}) + b_p). \quad (19)$$




Figure 2. Exposure-graph recurrent architecture of SMGSSM and its multitask outputs

Backpropagation jointly optimizes continuous assessment and binary risk identification. Regression loss is the primary objective, while binary cross-entropy with a coefficient of 0.20 constrains the shared hidden state to retain risk-category information. For the training mask $\mathcal{D}_{tr}$,

$$\mathcal{L}_{SMGSSM} = \frac{1}{|\mathcal{D}_{tr}|} \sum_{(i,t) \in \mathcal{D}_{tr}} (\tau_{i,t+1} - \hat{\tau}_{i,t+1})^2 + 0.20 \text{BCE}(y_{i,t+1}, \hat{p}_{i,t+1}). \quad (20)$$

After neural-network point prediction, 2024 observations are used for conformal calibration without using calibration samples to update network parameters. Let $q_{0.90}$ be the 90th percentile of absolute calibration residuals; the bounded conformal interval used to evaluate coverage is

$$\mathcal{I}_{i,t+1}^{0.90} = [\max(0, \hat{\tau}_{i,t+1} - q_{0.90}), \min(1, \hat{\tau}_{i,t+1} + q_{0.90})]. \quad (21)$$

Figure 2 shows the four encoders, reliability gates, 28-dimensional local representation, sparse graph aggregation, GRU, and multitask heads, while placing training loss, calibration residuals, and inference outputs on separate planes.

2.3. Security-Aware State Transition and Explainable Regulation

Cyber-physical resilience encompasses recovery, adaptation, and continuous operation [34]. Industry 5.0 still exhibits a disconnect between interconnected production and security governance [35], while automated actions in digital twins also require explainable state estimation [36]. This paper therefore converts SMGSSM outputs into control states so that cyber exposure enters state transitions, the objective function, and action constraints directly.

Assessment results are converted into executable decisions through a five-dimensional control state. Resilience retains its favorable direction, whereas the other four dimensions

represent pressures. Cyber exposure c_t connects the security encoding to the control layer, and the supply shock d_t^S and cyber shock d_t^C enter as parallel exogenous disturbances:

$$s_t = [r_t, l_t, \tau_t, p_t, c_t]^T, \quad w_t = [d_t^S, d_t^C]^T. \quad (22)$$

where l_t is liquidity pressure, τ_t is tax pressure, p_t is production pressure, and c_t is cyber exposure. The four classes of managerial actions are

$$u_t = [u_t^L, u_t^S, u_t^C, u_t^T]^T, \quad (23)$$

The four components represent liquidity support, supply substitution, cybersecurity hardening, and tax coordination. Security actions cover zero-trust access control [37], cyber-physical security control [38], threat-intelligence response [39], and state-adaptive anomaly protection [40], while accounting for edge-computing resources [41] and industrial wireless-continuity constraints [42]. The remaining actions correspond to financial resources, procurement and inventory resources, and tax-compliance resources.

No-action state dynamics are estimated by ridge regression on the public panel. Let the intercept be α_s and the coefficient matrix be B_s ; then

$$\bar{s}_{t+1} = \alpha_s + B_s[s_t \| w_t] + \varepsilon_t. \quad (24)$$

Supply and cyber shocks follow independent first-order models estimated from continuous public-panel indices:

$$w_{t+1} = \alpha_w + B_w w_t + \eta_t. \quad (25)$$

The controlled transition adds action effects and clips standardized states to their admissible ranges:

$$\hat{s}_{t+1} = \Pi_{[0,1]^5}(\alpha_s + B_s[s_t \| w_t] + G u_t + \varepsilon_t). \quad (26)$$

The state and shock regressions are fitted to observed public states, while intervention effects are represented by the following explicit matrix:

$$G = \begin{bmatrix} 0.08 & 0.10 & 0.12 & 0.06 \\ -0.25 & -0.02 & -0.05 & -0.10 \\ -0.04 & 0 & -0.10 & -0.25 \\ 0 & -0.25 & -0.10 & 0 \\ 0 & 0 & -0.35 & 0 \end{bmatrix}. \quad (27)$$




Figure 3. Public-data-calibrated state transition, objective evaluation, and separately routed regulation outputs

The rows of the matrix correspond to resilience and the four pressures, and the columns correspond to the four action classes. Main-diagonal effects indicate that actions primarily affect their own departmental objectives; security hardening also affects production and tax pressure through backup recovery and data integrity. G is an explicit simulation assumption rather than a causal effect identified from the public panel.

The stage loss trades off the resilience gap, four categories of pressure, action costs, and penalties for state-bound violations:

$$\begin{aligned} \ell(s_t, u_t) = & 1.5(1 - r_t) + l_t + \tau_t + 1.25p_t + c_t \\ & + 0.20(u_t^L)^2 + 0.22(u_t^S)^2 + 0.24(u_t^C)^2 + 0.18(u_t^T)^2 \\ & + 4 \left(\begin{aligned} & [l_t - 0.70]_+^2 + [\tau_t - 0.70]_+^2 + [p_t - 0.70]_+^2 \\ & + [c_t - 0.70]_+^2 \\ & + [0.30 - r_t]_+^2 \end{aligned} \right). \quad (28) \end{aligned}$$

Because manufacturing continuity is central to the response problem, the weight on production pressure is set to 1.25. The coefficient on security actions is the largest action-cost coefficient, preventing the optimizer from treating cyber protection as costless. The four-period receding-horizon objective is

$$u_t^* = \underset{u \in \mathcal{U}_t}{\operatorname{argmin}} \sum_{h=1}^4 0.92^{h-1} \ell(\hat{s}_{t+h|t}(u), u). \quad (29)$$

Candidate actions remain constant over the four-step horizon, but only the first step is executed before the state is reobserved. Figure 3 distinguishes the AI state interface, controlled-transition kernel, and multiobjective evaluation: G changes the next state, while the loss function only evaluates predicted states and actions.

2.4. AI-Assisted Constrained Receding-Horizon Action Coordination

Robust network design balances disruption and resource objectives [43], and distributionally robust optimization addresses demand and disruption uncertainty [44]. Research on industrial cyber-physical systems has also validated the

applicability of rolling reoptimization under attack disturbances [45]. Accordingly, this paper uses model predictive control (MPC) to update actions with quarterly AI states rather than adopting a fixed long-term plan.

Each action component is selected from $\{0, 0.25, 0.50, 0.75\}$, producing 256 raw combinations. Before evaluating predictions, MC-MPC removes candidates that violate the shared budget, simultaneous action-intensity, or adjustment-rate constraints. The feasible set is

$$\mathcal{U}_t = \{u_t: 0.9u_t^L + 1.0u_t^S + 1.15u_t^C + 0.8u_t^T \leq 1.28, \mathbf{1}^\top u_t \leq 1.50, \|u_t - u_{t-1}\|_\infty \leq 0.55\}. \quad (30)$$

The shared budget represents competition for resources across departments. The total-intensity constraint prevents maximal allocation to all actions simultaneously, whereas the rate constraint limits abrupt quarterly changes. Infeasible combinations do not enter the receding-horizon simulation.

Using the latest SMGSSM state and forecasts of the two shocks, the controller screens 256 action combinations, rolls out four-period trajectories, and accumulates discounted losses, executing only the first step of the optimal sequence.

New evidence is then fed back to the assessment side, forming a closed loop of AI-based identification, MPC search, and state updating.

The dominant computational complexity of SMGSSM is $O(T(MNd^2 + |E|d + Nd^2))$, where $M = 4$ and $d = 28$. Because each firm retains only five principal neighbors, the graph remains sparse. If K candidate schemes remain after feasibility screening, MC-MPC with prediction horizon $H = 4$ requires approximately $O(KH)$ state evaluations. Enumeration is appropriate at the scale studied here: it generates reproducible actions rather than an opaque policy and makes every rejected candidate auditable.

Figure 4 centers on a four-period state-transition chain and combines disturbance forecasts, a candidate-action tensor, a feasibility mask, and a discounted objective. argmin decodes only the first-period action and feeds it back to SMGSSM.




Figure 4. Constrained rolling coordination of liquidity, supply, cybersecurity, and tax actions

3. Experiments and Results

3.1. Public Dataset and Deep-Learning Experimental Setup

Data were obtained from the U.S. Securities and Exchange Commission (SEC) Electronic Data Gathering, Analysis, and Retrieval (EDGAR) system. The SEC EDGAR Companyfacts eXtensible Business Reporting Language (XBRL) endpoint provides quarterly accounting facts,

whereas the Submissions and Form 10-K archives provide firm, industry, filing, and textual information. The sample contains 576 records for 24 manufacturing firms from 2020Q1 to 2025Q4, 529 next-quarter targets, and 140 valid Form 10-K filings. The setup is reported in Table 2. The panel is restricted to 24 firms because inclusion requires continuous quarterly XBRL coverage and usable Form 10-K text across all four modalities. The analysis should therefore be interpreted as a proof-of-concept within-panel comparison rather than as population-level evidence for U.S. manufacturing firms.

Table 2. Public dataset, assessment design, and control settings

Item	Setting	Traceable basis
Firms and period	24 U.S. listed manufacturers; 2020Q1–2025Q4	SEC EDGAR public materials
Firm-quarter records	576 records; 529 next-quarter targets	Processed Companyfacts XBRL
Filing text	140 valid 10-K bodies	SEC filing archives
Modality dimensions	F=6, O=5, C=5, T=4	Constructed public features
Exposure graph	TF-IDF similarity + two-digit SIC increment; top five neighbors	Public-material adjacency matrix
Calendar-year model masks	368 train; 92 calibration; 69 holdout	Full-panel preprocessing precedes model masks
Assessment repetitions	Seeds 11, 29, and 47	Executed training runs
Control evaluation	120 matched paths; 18 periods per path	Public-data-calibrated closed loop

The XBRL data provide balance-sheet, revenue, inventory, cash-flow, profit, and tax facts through aliases of standard concepts. Cumulative flows are converted using quarterly frames or within-year differencing, and all ratios are winsorized at the 1st and 99th percentiles.

After Form 10-K cleaning, keyword densities per 10,000 words are calculated, and semantics are extracted using English unigrams and bigrams, sublinear TF-IDF, 1,200 terms, and four SVD components. Annual features are aligned to quarters by reporting year; the exposure graph retains five neighbors and a self-loop using multiyear average text vectors and SIC proximity.

The model masks are 2020Q1–2023Q4 for training, calendar year 2024 for selection and conformal calibration, and calendar year 2025 for holdout evaluation. TF-IDF/SVD, the exposure graph, winsorization bounds, resilience scaling, and selected thresholds are first fitted on the full 2020–2025 panel, after which annual Form 10-K features are aligned to quarters by reporting year. Thus, the 2025 results constitute a calendar-year holdout comparison after full-sample preprocessing, not a prospective test constructed according to information available at each point in time. The effective numbers of training, calibration, and holdout observations are 368, 92, and 69, respectively, and the random seeds are 11, 29, and 47.




Figure 5. Multimodal evidence tensor, split-aware reliability, annual filing attributes, and the filing-calendar constellation

Comparators include a financial-and-tax GRU, a plain multimodal GRU, a no-graph ablation, a no-cybersecurity ablation, and the full SMGSSM. All models use AdamW with learning rate 1.2×10^{-3} , weight decay 10^{-4} , gradient clipping at 2.0, and a maximum of 220 epochs; early stopping is triggered if validation loss fails to improve for 25 epochs. Three runs capture variation from initialization and the small sample.

Metrics include root mean squared error (RMSE), mean absolute error (MAE), R^2 , macro-averaged area under the receiver operating characteristic curve (AUC), shock-period RMSE, RMSE under 20% additional missingness, 90% interval coverage and width, and computation time. Shock-period observations are holdout observations for which the supply or cyber index exceeds the 75th percentile within the training mask.

The control experiment compares no control, reactive rules, and MC-MPC. State and shock regressions are calibrated on the full public panel; 120 episodes share matched 18-period shock and residual paths, with loss, resilience, violations, component pressures, actions, and computation time recorded. Figure 5 presents evidence from four perspectives: the firm-quarter reliability tensor and three model masks; the modality-reliability rose; annual Form 10-K count, length,

and threat density; and the filing-month-reporting-year-supply-risk star map. It also marks the boundaries of full-panel preprocessing and annual-text alignment.

3.2. Resilience Assessment and Risk Identification by the Deep Model

Table 3 reports the means and standard deviations across three training runs in the calendar-year 2025 holdout test after full-sample preprocessing. Among the evaluated models, SMGSSM achieves the lowest RMSE, MAE, and shock-period RMSE. Its RMSE is 0.0663 ± 0.0089 , 8.8% lower than that of the financial-and-tax GRU and 6.0% lower than that of the plain multimodal GRU; the corresponding MAE is 0.0511 ± 0.0052 . Plain concatenation adds information sources but does not reproduce the error level of the full model. This indicates that the performance gain does not arise simply because "more modalities are better," but from structured learning of heterogeneous evidence through reliability gating, graph messaging, and recurrent memory.

Table 3. Assessment performance under the 2025 calendar-year holdout after full-sample preprocessing

Model	RMSE ↓	MAE ↓	R^2 ↑	Shock RMSE ↓	Risk AUC ↑	20% missing RMSE ↓
Financial-tax GRU	0.0727 0.0092	± 0.0566 0.0076	± 0.3149 0.1655	± 0.0753 ± 0.0094	0.6600 ± 0.0342	0.0730 ± 0.0057
Plain multimodal GRU	0.0706 0.0109	± 0.0550 0.0085	± 0.3514 0.1892	± 0.0721 ± 0.0112	0.7489 ± 0.0572	0.0742 ± 0.0092

Model	RMSE ↓	MAE ↓	R ² ↑	Shock RMSE ↓	Risk AUC ↑	20% missing RMSE ↓
Without graph propagation	0.0744 0.0129	± 0.0564 0.0090	± 0.2771 0.2345	± 0.0771 ± 0.0137	0.7218 ± 0.0283	0.0746 ± 0.0082
Without cyber modality	0.0666 0.0034	± 0.0524 0.0028	± 0.4308 0.0571	± 0.0686 ± 0.0035	0.7441 ± 0.0690	0.0717 ± 0.0069
Proposed SMGSSM	0.0663 0.0089	± 0.0511 0.0052	± 0.4294 0.1476	± 0.0685 ± 0.0092	0.7280 ± 0.0502	0.0662 ± 0.0067

Risk identification yields a different model ranking. The plain multimodal GRU achieves the highest macro-averaged AUC at 0.7489 ± 0.0572 , followed by the no-cybersecurity model at 0.7441 ± 0.0690 and SMGSSM at 0.7280 ± 0.0502 . The full network should therefore not be described as optimal for every AI task. The shared hidden state primarily serves continuous resilience regression, while the risk heads participate in joint training with a loss weight of only 0.20; graph smoothing may also weaken decision boundaries for a small number of extreme risk labels. These results indicate that the current architecture is better suited to continuous-state assessment and stable inference under missingness than to replacing a purpose-trained cyberattack detector.

Figure 6 places accuracy, risk identification, and computational cost in a common multiscale comparison. Panel (a) simultaneously shows RMSE, macro-averaged AUC, and training time in three-dimensional coordinates, with bubble area encoding inference efficiency; actual training points from the three random seeds form model clusters. Panels (b)-(e) display the run distributions of RMSE, shock-period RMSE, risk AUC, and RMSE under additional

missingness, respectively, avoiding rankings based on means alone. SMGSSM has a training time of 5.40 ± 0.50 s and an inference time of 0.340 ± 0.008 ms. Graph operations increase computation, but inference remains at the millisecond scale on the 24-node network.

The saved representative prediction run provides 69 holdout observations and their 28-dimensional recurrent hidden states. Figure 7(a) applies principal component analysis (PCA) to the hidden states solely to visualize the internal geometry learned by the deep model on the holdout set; point color represents observed resilience, and point area represents the maximum probability among the three risk outputs. Panel (b) overlays the joint observed-predicted density, local absolute residuals, and the 45-degree reference line. Panel (c) forms a three-dimensional manifold from observed resilience, supply or cyber shock intensity, and prediction residuals. Panel (d) further overlays the conformal intervals derived from the 2024 calibration-residual quantile. Although the sample generally remains distributed around the 45-degree line, residuals fluctuate more in low-resilience and high-shock regions, consistent with the limited number of severe firm-quarter observations in the 24-firm panel.




Figure 6. The three-dimensional accuracy-risk-cost Pareto space and executed-run distributions of the five models




Figure 7. Latent-state geometry, joint prediction density, the shock-residual manifold, and representative conformal coverage for the 2025 calendar-year holdout

3.3. AI Components, Robustness, and Uncertainty Analysis

Shock-period RMSE exceeds overall RMSE for every model, confirming that high-shock observations are harder to predict. SMGSSM nevertheless retains the lowest mean shock-period error at 0.0685 ± 0.0092 ; the error rises to 0.0771 ± 0.0137 when the graph message layer is removed, while the financial-and-tax GRU records 0.0753 ± 0.0094 . The graph module's relative advantage does not disappear among shock observations, supporting its role in capturing exposure signals from related firms. When the target firm's own accounting facts do not yet fully reflect a propagating disruption, neighboring embeddings can provide supplementary state information to the GRU.

An additional test independently masks 20% of each modality and lowers the corresponding quality value to 0.05. Under this perturbation, SMGSSM has an RMSE of 0.0662 ± 0.0067 , the no-graph model 0.0746 ± 0.0082 , and the plain multimodal GRU 0.0742 ± 0.0092 . The small numerical

difference between the original and masked SMGSSM RMSE reflects the limited number of repeated masking runs rather than evidence that missingness improves prediction. The meaningful finding is the absence of clear error amplification under the stated perturbation.

The 90% calibrated intervals cover 0.8986 ± 0.0383 of the holdout targets and have an average width of 0.2303 ± 0.0079 . Under this evaluation design, coverage is close to the nominal level, but the intervals remain wide relative to the $[0, 1]$ target range. They are therefore suitable for reporting predictive uncertainty in the current comparison but cannot replace controller state constraints. Figure 8(a), based on actual reevaluation of the saved seed-11 model, forms a three-dimensional robustness response surface from the additional-missingness rate, training-distribution shock-threshold quantile, and holdout RMSE. Panel (b) provides a contour projection of the same surface, with hollow-circle area indicating the retained number of high-shock observations. The two views reveal the overall surface shape and local numerical gradients, respectively.




Figure 8. The executed robustness response surface and its contour projection under joint shock and missingness perturbations

The exposure graph contains 24 nodes and 144 nonzero directed weights, including 120 directed non-self weights and 24 self-retention weights. Figure 9(a) uses a ForceAtlas2 layout to show all 120 directed non-self message edges with arrows. An arrow from firm j to firm i represents, in Equation (13), the source firm's contribution to receiving firm i , namely $A_{ij}z_{j,t}$. Self-retention terms are omitted from the figure but remain in the model. Arrow width and opacity encode A_{ij} , node size encodes off-diagonal source influence, and node color indicates the two-digit SIC group. Panel (b) compares source influence and received exposure weight across firms, while panel (c) lists the principal message sources. Dense local structures form among aerospace, automotive, and industrial-equipment firms, whereas cross-industry links arise from similarities in risk and supply narratives.

The ablation results distinguish the actual roles of the AI components. Removing graph propagation increases RMSE from 0.0663 to 0.0744, a relative deterioration of 12.1%; shock-period RMSE rises by approximately 12.4%, making the graph message layer the clearest contributor to error reduction. Removing the cybersecurity modality changes RMSE from 0.0663 to 0.0666, an increase of only 0.4%, while R^2 and risk AUC improve slightly. Annual disclosures are less frequent than quarterly operational variables, and security term frequencies cannot substitute for industrial-control logs, so this modality contributes only marginally to continuous prediction. It is retained to represent attack exposure explicitly, alter the shared hidden state, and provide an auditable entry point for cybersecurity-hardening actions, not because it dominates overall predictive performance.




Figure 9. ForceAtlas2 message propagation, source-receiver centrality, and leading sources in the public-material graph

Figure 10 reports the changes in each configuration relative to the full model using a direction-normalized matrix; color jointly encodes metric direction and magnitude of change. Coverage is normalized by its absolute deviation from the nominal 0.90 level to avoid exaggeration caused by a small denominator. Removing the graph structure consistently degrades all three error metrics, whereas the effect of removing the cybersecurity modality is mixed; the plain

multimodal GRU also fails to reproduce the gated model's stability under additional missingness. These results constrain the architectural interpretation: graph propagation has comparatively stable supporting evidence, reliability gating primarily serves robustness to missingness, and the security modality provides a composite-risk representation and control interface.



Figure 10. Direction-normalized architecture and ablation effects relative to the complete SMGSSM

3.4. AI-Driven Security Regulation and Manufacturing-Response Trade-offs

Table 4 reports average closed-loop outcomes across 120 matched shock paths. Under no control, aggregate loss is 55.549, average resilience is 0.4958, and the state-violation

rate is 0.4801. Reactive rules trigger fixed measures only when current thresholds are exceeded, reducing loss to 39.267 and raising average resilience to 0.6660. MC-MPC instead uses AI-assessed states to roll out four-period forecasts, further reducing loss to 36.431 and increasing average resilience to 0.6738, with no state violations at the reported precision. Mean decision time is 0.252 ± 0.047 ms.

Table 4. Closed-loop regulation outcomes over 120 matched episodes

Strategy	Composite loss ↓	Mean resilience ↑	Minimum resilience ↑	Violation rate ↓	Production loss ↓	Financial-tax pressure ↓	Cyber exposure ↓	Control cost ↓
No control	55.549	0.4958	0.4334	0.4801	9.205	17.533	12.184	0.000
Reactive rule	39.267	0.6660	0.5973	0.00185	7.663	11.456	7.470	1.744
Proposed MC-MPC	36.431	0.6738	0.5966	0.000	8.401	9.608	4.351	3.163

Relative to no control, paired MC-MPC reduces loss by 19.118 ± 0.157 , or 34.42%, and raises average resilience by 0.17799 ± 0.00010 , or 35.90%. Production loss, financial and tax pressure, and cyber exposure decline by 8.73%, 45.20%, and 64.29%, respectively. The violation rate falls from 0.4801 to 0.0000 through the joint effects of AI state updating, quadratic boundary penalties, and pre-prediction action screening. Within the specified simulation, this comparison supports the internal consistency of the assessment-to-action mechanism; it does not establish causal intervention effectiveness.

Table 4 presents the overall distribution across 120 matched paths, whereas Figure 11 shows an interpretable representative closed-loop trajectory. The selected episode is

the matched path whose cumulative aggregate loss under MC-MPC is closest to the median; it represents neither the best nor the worst case. Panel (a) overlays the 18-period trajectories of the three controllers in a three-dimensional production-pressure-cyber-exposure-resilience state space. Panel (b) expands the five-dimensional states under the same disturbances into controller-grouped heatmaps. Panel (c) further overlays area flows for the four MC-MPC actions with supply and cyber shocks. This view shows how overall statistical differences emerge from period-by-period state updates. The no-control path lies closer to pressure boundaries, both reactive rules and MC-MPC improve resilience, and MC-MPC allocates more actions persistently to cybersecurity-hardening and tax-coordination channels.




Figure 11. Representative matched-path state trajectories, five-state evolution, and action-shock coordination in the closed loop

MC-MPC does not outperform reactive rules on every objective. Table 5 shows that MC-MPC reduces aggregate loss by 7.22%, financial and tax pressure by 16.13%, and cyber exposure by 41.76%, while increasing average resilience by 1.17%; however, production loss rises by 9.63%

and control cost by 81.38%. The paired difference in minimum resilience is -0.00072 ± 0.00137 , whose interval includes zero, so a stable worst-case resilience advantage cannot be claimed.


Table 5. Paired relative changes of MC-MPC against the two baselines

Comparison	Composite loss	Mean resilience	Violation rate	Production loss	Financial-tax pressure	Cyber exposure	Control cost
MC-MPC vs. no control	-34.42%	+35.90%	-100.00%	-8.73%	-45.20%	-64.29%	+3.163
MC-MPC vs. reactive rule	-7.22%	+1.17%	-100.00%	+9.63%	-16.13%	-41.76%	+81.38%


Action records explain these trade-offs. The average cumulative MC-MPC actions for liquidity support, supply substitution, cybersecurity hardening, and tax coordination are 0.002, 0.310, 13.188, and 9.000, respectively; reactive rules allocate 3.844, 6.040, 8.030, and 4.045, respectively. Under the current G matrix and loss weights, predictive optimization concentrates the shared budget in the security and tax channels. The cybersecurity mechanism thereby

protects states related to attack exposure and data integrity but crowds out some direct-procurement and supply-substitution resources, producing higher cumulative production loss. This result demonstrates precisely why AI-driven security responses must be optimized jointly with manufacturing objectives rather than treating security hardening as a budget-independent add-on.


(a) MC-MPC objective surface and action-intensity contours



(b) Controller-to-action allocation flows



(c) Episode-level cost-loss-resilience landscape



(d) Circular benefit and adverse-trade-off spectrum




Figure 12. The MC-MPC objective surface, action-allocation flows, the episode-level cost-loss-resilience landscape, and circular trade-off spectra

Figure 12 summarizes controller trade-offs at three levels. Panel (a) shows the discounted-objective surface and projected action intensities, while Panel (b) compares average resource allocation. Panel (c) relates regulation cost, aggregate loss, and resilience across episodes, and Panel (d) summarizes favorable and adverse changes relative to the baselines. The narrow MC-MPC cost band indicates that the current action grid and objective weights nearly saturate the shared budget; it does not establish the general optimality of budget saturation.

3.5. Limitations, Scope Conditions, and Deployment Requirements

The assessment experiment is retrospective. TF-IDF/SVD features, the exposure graph, winsorization bounds, resilience scaling, and selected thresholds were constructed using the full 2020–2025 panel. The 2025 holdout therefore measures comparative performance within this retrospectively processed panel and should not be interpreted as prospective generalization. Statistical precision is further limited by the 24-firm sample, 69 holdout observations, and three random initialization seeds. The resilience score, static exposure graph, and cybersecurity keyword densities are constructed proxies and have not been externally validated against realized recovery, credit distress, or independently labeled cyber incidents.

The control experiment is simulation-based. Although the no-action transition is calibrated from public observations, the intervention matrix G , action costs, budget limits, and loss weights are scenario parameters rather than causally identified effects. The reported control gains and crowding-out effects are therefore conditional on these settings and do not constitute deployable policy effects. Real-world implementation would require point-in-time feature construction, firm-specific operational and cybersecurity data, local or causal calibration of intervention effects, periodic model updating, scalability testing beyond 24 firms, and human approval of recommended actions.

4. Conclusion

This paper proposes an AI-based cybersecurity mechanism for resilient manufacturing responses. SMGSSM learns firms' quarterly states through four neural encoders, reliability gating, sparse exposure-graph message passing, and GRU memory, while multitask heads jointly predict resilience and liquidity, tax, and production risks. MC-MPC then incorporates cyber exposure into the controlled state and jointly selects liquidity, supply, security, and tax actions subject to budget, total-intensity, and adjustment-rate constraints. This structure transforms cyber-risk identification from a parallel alert into a state variable and

decision interface within the manufacturing-response loop. The paper's original contribution therefore lies not in any single neural or control module in isolation, but in their integration into an interpretable assessment-to-action mechanism that links heterogeneous public evidence, interfirm shock propagation, temporal state evolution, and constrained managerial intervention.

The calendar-year 2025 holdout results after full-sample preprocessing show that SMGSSM has an RMSE of 0.0663 ± 0.0089 , which rises by 12.1% when graph propagation is removed, indicating that relational messaging is the most stable deep-model gain. The cybersecurity modality improves continuous RMSE by only 0.4% at the margin; its primary value lies in representing attack exposure and connecting the control channel rather than dominating overall prediction. Across 120 closed-loop paths, MC-MPC reduces aggregate loss by 34.42% and increases average resilience by 35.90% relative to no control. Relative to reactive rules, it reduces cyber exposure by 41.76% but increases production loss by 9.63% and control cost by 81.38%. AI-driven security mechanisms can therefore improve constrained aggregate resilience, but they cannot eliminate competition for resources among cyber defense, financial and tax protection, and immediate production recovery. These results suggest that the cybersecurity modality may retain decision-interface value despite its small marginal predictive contribution. This interpretation remains conditional on the constructed cybersecurity proxy, the retrospective preprocessing design, and the assumed intervention matrix.

Acknowledgements

1. The Natural Science Research Project of Anhui Provincial Department of Education: Quantitative Measurement, Influence Mechanism, and Yangtze River Delta Collaborative Simulation Study of Digital Economy Empowering Urban Economic Resilience in Anhui Province (No. 2025KYZR01).
2. The Humanities and Social Sciences Research Project of Anhui Provincial Department of Education: Research on Financial Risk Identification, Early Warning, and Governance Mechanism of Anhui Sci-Tech Enterprises in the Construction of the Yangtze River Delta Science and Technology Innovation Community (No. 2025KYSK10).

References

- [1] Nikookar E, Stevenson M, Varsei M. Building an antifragile supply chain: a capability blueprint for resilience and post-disruption growth. *J Supply Chain Manag.* 2024;60(1):13–31. doi:10.1111/jscm.12313.
- [2] Belhadi A, Kamble S, Subramanian N, Singh RK, Venkatesh M. Digital capabilities to manage agri-food supply chain uncertainties and build supply chain resilience during compounding geopolitical disruptions. *Int J Oper Prod Manag.* 2024;44(11):1914–1950. doi:10.1108/IJOPM-11-2022-0737.

- [3] Liu P, Liu J, Tao C. Market access, supply chain resilience and enterprise innovation. *J Innov Knowl.* 2024;9(4):100576. doi:10.1016/j.jik.2024.100576.
- [4] Yang S, Ogawa Y, Ikeuchi K, Shibasaki R, Okuma Y. Post-hazard supply chain disruption: predicting firm-level sales using graph neural network. *Int J Disaster Risk Reduct.* 2024;110:104664. doi:10.1016/j.ijdrr.2024.104664.
- [5] Bruckler M, Wietschel L, Meßmann L, Thorenz A, Tuma A. Review of metrics to assess resilience capacities and actions for supply chain resilience. *Comput Ind Eng.* 2024;192:110176. doi:10.1016/j.cie.2024.110176.
- [6] Ouyang X, Liu L, Chen W, Wang C, Sun X, He C, et al. Systematic risks of the global lithium supply chain network: from static topological structures to cascading failure dynamics. *Environ Sci Technol.* 2024;58(50):22135-22147. doi:10.1021/acs.est.4c10523.
- [7] Rashid A, Rasheed R, Ngah AH, Amirah NA. Unleashing the power of cloud adoption and artificial intelligence in optimizing resilience and sustainable manufacturing supply chain in the USA. *J Manuf Technol Manag.* 2024;35(7):1329-1353. doi:10.1108/JMTM-02-2024-0080.
- [8] Malekzadeh M, Torabi SA, Sheikhalishahi M. A supply chain digital twin for resilience analysis under immediate and long-term disruptions: a case study in plastic parts' manufacturing. *Int J Prod Res.* 2025;63(24):10445-10467. doi:10.1080/00207543.2025.2551243.
- [9] Zheng LJ, Islam N, Zhang JZ, Behl A, Wang X, Papadopoulos T. Aligning risk and value creation: a process model of supply chain risk management in geopolitical disruptions. *Int J Oper Prod Manag.* 2025;45(5):1178-1210. doi:10.1108/IJOPM-03-2024-0271.
- [10] Lucker F, Timonina-Farkas A, Seifert RW. Balancing resilience and efficiency: a literature review on overcoming supply chain disruptions. *Prod Oper Manag.* 2025;34(6):1495-1511. doi:10.1177/10591478241302735.
- [11] Guo Y, Liu F, Song JS, Wang S. Supply chain resilience: a review from the inventory management perspective. *Fundam Res.* 2025;5(2):450-463. doi:10.1016/j.fmre.2024.08.002.
- [12] Gervais E, Sprecher B, Nold S, Brailovsky P, Kleijn R. Tracing the propagation of disruptions in supply chain scenarios: a case study of photovoltaics diversification. *Resour Conserv Recycl.* 2025;212:107948. doi:10.1016/j.resconrec.2024.107948.
- [13] Li J, Song H, Ma Y. The mechanism and impact of digital transformation on supply chain resilience in the manufacturing industry. *Sci Rep.* 2026;16:7635. doi:10.1038/s41598-026-38930-9.
- [14] Ivanov D. Cash flow dynamics in the supply chain during and after disruptions. *Transp Res E Logist Transp Rev.* 2024;185:103526. doi:10.1016/j.tre.2024.103526.
- [15] Qiu B, Xu F, Yeung ACL, Zeng CC. Contagious stock price crashes along the supply chain. *Prod Oper Manag.* 2024;33(8):1679-1699. doi:10.1177/10591478241254854.
- [16] Chen M, Zhao K, Jin W. Corporate digital transformation and tax avoidance: evidence from China. *Pac Basin Finance J.* 2024;85:102400. doi:10.1016/j.pacfin.2024.102400.
- [17] Belahouaoui R, Attak EH. Digital taxation, artificial intelligence and Tax Administration 3.0: improving tax compliance behavior - a systematic literature review using textometry (2016-2023). *Account Res J.* 2024;37(2):172-191. doi:10.1108/ARJ-12-2023-0372.
- [18] Di L, Jiang W, Mao J, Zeng Y. Managing liquidity along the supply chain: supplier-base concentration and corporate cash policy. *Eur Financ Manag.* 2024;30(4):2376-2421. doi:10.1111/eufm.12479.
- [19] Xu C, Kong Y. Random forest model in tax risk identification of real estate enterprise income tax. *PLoS One.* 2024;19(3):e0300928. doi:10.1371/journal.pone.0300928.
- [20] Sun X, Han J, Işık C. Can digital transformation be a "powerful tool" to curb corporate tax avoidance? Mechanism analysis and policy implications based on Chinese data. *Int Rev Econ Finance.* 2025;102:104330. doi:10.1016/j.iref.2025.104330.
- [21] Dong H, Zhang L. Climate risk and tax avoidance of climate-sensitive firms. *Bus Strategy Environ.* 2025;34(7):8429-8449. doi:10.1002/bse.70025.
- [22] Song Y, Zhu M, Wang Y. Corporate financialization, digital transformation, and industrial supply chain resilience: mechanisms based on R&D investment and financial regulation. *Int Rev Financ Anal.* 2025;102:104119. doi:10.1016/j.irfa.2025.104119.
- [23] Wu K, Zheng K. Tax incentives, supply chain spillovers, and enterprise technological innovation. *Int Rev Financ Anal.* 2025;99:103930. doi:10.1016/j.irfa.2025.103930.
- [24] Dodd O, Liao S. The role of inventory in firm resilience to the COVID-19 pandemic. *Eur Financ Manag.* 2025;31(2):786-818. doi:10.1111/eufm.12517.
- [25] Moosavi S, Farajzadeh-Zanjani M, Razavi-Far R, Palade V, Saif M. Explainable AI in manufacturing and industrial cyber-physical systems: a survey. *Electronics.* 2024;13(17):3497. doi:10.3390/electronics13173497.
- [26] Zhao M, Peng H, Li L, Ren Y. Graph attention network and Informer for multivariate time series anomaly detection. *Sensors (Basel).* 2024;24(5):1522. doi:10.3390/s24051522.
- [27] Qu X, Liu Z, Wu CQ, Hou A, Yin X, Chen Z. MFGAN: multimodal fusion for industrial anomaly detection using attention-based autoencoder and generative adversarial network. *Sensors (Basel).* 2024;24(2):637. doi:10.3390/s24020637.
- [28] Alsaif KM, Albeshri AA, Khemakhem MA, Eassa FE. Multimodal large language model-based fault detection and diagnosis in context of Industry 4.0. *Electronics.* 2024;13(24):4912. doi:10.3390/electronics13244912.
- [29] Zheng G, Brintrup A. A machine learning approach for enhancing supply chain visibility with graph-based learning. *Supply Chain Anal.* 2025;11:100135. doi:10.1016/j.sca.2025.100135.
- [30] Zamanzadeh Darban Z, Webb GI, Pan S, Aggarwal CC, Salehi M. Deep learning for time series anomaly detection: a survey. *ACM Comput Surv.* 2024;57(1):15. doi:10.1145/3691338.
- [31] Zheng G, Brintrup A. Enhancing supply chain visibility with generative AI: an exploratory case study on relationship prediction in knowledge graphs. *Int J Prod Res.* 2025;1-23. doi:10.1080/00207543.2025.2543964.
- [32] Kosasih EE, Brintrup A. Towards trustworthy AI for link prediction in supply chain knowledge graph: a neurosymbolic reasoning approach. *Int J Prod Res.* 2025;63(6):2268-2290. doi:10.1080/00207543.2024.2399713.
- [33] AlMahri S, Xu L, Brintrup A. Enhancing supply chain visibility with knowledge graphs and large language models. *Int J Prod Res.* 2026;64(6):2178-2209. doi:10.1080/00207543.2025.2575841.
- [34] Segovia-Ferreira M, Rubio-Hernan J, Cavalli AR, García-Alfaro J. A survey on cyber-resilience approaches for cyber-physical systems. *ACM Comput Surv.* 2024;56(8):202. doi:10.1145/3652953.
- [35] Kour R, Karim R, Dersin P, Venkatesh N. Cybersecurity for Industry 5.0: trends and gaps. *Front Comput Sci.* 2024;6:1434436. doi:10.3389/fcomp.2024.1434436.
- [36] Sarker IH, Janicke H, Mohsin A, Gill A, Maglaras L. Explainable AI for cybersecurity automation, intelligence and trustworthiness in digital twin: methods, taxonomy, challenges

- and prospects. *ICT Express*. 2024;10(4):935-958. doi:10.1016/j.ict.2024.05.007.
- [37] Zanas C, Russo S, Colajanni M. Flexible zero trust architecture for the cybersecurity of industrial IoT infrastructures. *Ad Hoc Netw*. 2024;156:103414. doi:10.1016/j.adhoc.2024.103414.
- [38] Xing W, Shen J. Security control of cyber-physical systems under cyber attacks: a survey. *Sensors (Basel)*. 2024;24(12):3815. doi:10.3390/s24123815.
- [39] Sadhwani S, Modi UK, Muthalagu R, Pawar PM. SmartSentry: cyber threat intelligence in industrial IoT. *IEEE Access*. 2024;12:34720-34740. doi:10.1109/ACCESS.2024.3371996.
- [40] Ahmadi-Assalemi G, Al-Khateeb H, Benson V, Ammi M. Adaptive learning anomaly detection and classification model for cyber and physical threats in industrial control systems. *IET Cyber-Phys Syst Theory Appl*. 2025;10(1):e70004. doi:10.1049/cps2.70004.
- [41] Zhukabayeva T, Zholshiyeva L, Karabayev N, Khan S, Alnazzawi N. Cybersecurity solutions for Industrial Internet of Things-edge computing integration: challenges, threats, and future directions. *Sensors (Basel)*. 2025;25(1):213. doi:10.3390/s25010213.
- [42] Babbar H, Rani S, Boulila W. Fortifying the connection: cybersecurity tactics for WSN-driven smart manufacturing in the era of Industry 5.0. *IEEE Open J Commun Soc*. 2025;6:3417-3428. doi:10.1109/OJCOMS.2024.3428531.
- [43] Hosseini Shekarabi SA, Kiani Mavi R, Romero Macau F. An extended robust optimisation approach for sustainable and resilient supply chain network design: a case of perishable products. *Eng Appl Artif Intell*. 2025;152:110846. doi:10.1016/j.engappai.2025.110846.
- [44] Chen S, Chen Y. Designing a resilient supply chain network: a multi-objective data-driven distributionally robust optimization method. *Comput Oper Res*. 2025;173:106868. doi:10.1016/j.cor.2024.106868.
- [45] Zhang Y, Du C, Chen Z, Feng Z, Gui W. Digital twin-based resilient model predictive control for industrial cyber-physical systems. *IEEE J Emerg Sel Topics Ind Electron*. 2025;6(4):1853-1862. doi:10.1109/JESTIE.2025.3528881.