Future Technology
Open Access Journal
ISSN 2832-0379

August 2026| Volume 05 | Issue 03 | Pages 331-343

Journal homepage: https://fupubco.com/futech

https://doi.org/10.55670/fpll.futech.5.3.29

Future Publishing LLC

Article

Dynamic graph neural networks meet causal inference: estimating AI's heterogeneous effects on supply chain resilience

Yang Liu[1]*, Jiaming Yang[2], Jian Chen[3]

[1]MAHSA University, 42610 Jenjarom, Malaysia
[2]Universitas Pendidikan Ganesha (Undiksha), Bali 81116, Indonesia
[3]Zhongshan Polytechnic, Zhongshan 528400, China

ARTICLE INFO

Article history:
Received 02 February 2026
Received in revised form
04 June 2026
Accepted 03 July 2026

Keywords:
Temporal graph attention network, Causal Forest,
Heterogeneous treatment effects,
Supply chain resilience, Trade policy uncertainty

*Corresponding author
Email address:
3878054637@qq.com

DOI: 10.55670/fpll.futech.5.3.29

ABSTRACT

Estimating heterogeneous treatment effects in network-embedded environments requires methods that simultaneously account for relational change, conditional heterogeneity, and quasi-experimental shocks. This study proposes TGAT-CF, an approach that pairs a temporal graph attention encoder for evolving supplier-customer relationships with a generalized random forest for conditional treatment effects, triangulated against staggered difference-in-differences and double machine learning. Temporal graph embeddings replace scalar centrality as moderators, and the resulting design allows identification to be cross-checked across three estimators in settings where technology adoption unfolds inside relational dynamics. The framework is applied to a balanced quarterly panel of 2,847 Chinese A-share manufacturers over 20 quarters from 2020Q1 to 2024Q4, yielding 56,940 firm-quarter observations amid simultaneous AI diffusion and trade policy uncertainty. The temporal encoder reduces mean squared error by 21.4 percent relative to a static GraphSAGE baseline and by 6.6 to 9.4 percent relative to dynamic baselines (Diebold-Mariano, $p < 0.05$). A one-standard-deviation rise in AI stock raises supply chain resilience by 0.34 standard deviations, an effect that is 2.6 times larger under high uncertainty. Conditional effects differ by a factor of 2.9 between modular and centralized configurations, and the temporal profile follows a J-curve peaking at event time 2. Network centrality is the leading moderator, ahead of ownership structure, while operational efficiency, supplier adjustment, and information processing mediate nearly three-quarters of the total effect. The three estimates converge within a 7 percent band. AI capability, therefore, acts as a network-dependent, rather than a universal, determinant of supply chain resilience.

1. Introduction

The global supply chain structure has experienced changes due to two simultaneous trends: trade policy uncertainties and the proliferation of artificial intelligence. UNCTAD's 2025 World Investment Report indicates a reduction in global FDI flows for the third consecutive year. Greenfield manufacturing activities are shifting towards friend-shoring countries [1]. Trade uncertainties have shifted from episodic events to a structural one following the U.S.-China tariff increase in 2018, with subsequent increases linked to export control changes in 2022 and tariff renegotiations in 2024-2025. This means that manufacturing companies must cope with both the need to improve AI capacity and increased external uncertainty. Empirical studies investigating the impact of artificial intelligence adoption on supply chain resilience have also been conducted. Nonetheless, three limitations constrain the usefulness of such research to policymakers. This relationship is evaluated using binary indicators based solely on one year's data drawn from disclosures or patents [2], thereby condensing what might be a multi-year process of capability creation into a single data point. In terms of causal identification, its reliance on an average treatment effect across a heterogeneous set of firms precludes management action on moderation factors such as ownership, network centrality, and tariffs [3]. The network structure dimension is limited by considering relationships between suppliers and customers as covariates that are constant across time,

because the literature on propagation shows that company-specific disturbances spread through highly connected supplier-customer networks, with significant amplification effects [4]. However, resilience applications generally use these networks as firm-level covariates, ignoring time variation in the embedding around the focal company. The Chinese manufacturing setting provides valuable context for addressing this gap, as companies face significant pressure from AI policies under the 2017–2030 New Generation Artificial Intelligence Development Plan and the most severe trade policy uncertainty between China and the United States since 2018. This study addresses two research questions.

RQ1. How does time-varying AI adoption intensity affect firm-level supply chain resilience under elevated trade policy uncertainty when each firm's relational structure is modeled as a dynamic graph?

RQ2. Which firm-level and network-level factors explain the heterogeneity in the association with AI resilience?

Addressing both questions requires a Temporal Graph Neural Network to capture changes in supplier-customer connections, a Causal Forest to estimate conditional treatment effects without assuming any parametric structure, and a staggered Difference-in-Differences approach to provide a causal interpretation based on variations in tariffs [5].

This study makes one methodological and two substantive contributions, stated separately to distinguish what is new from what extends prior work. The methodological contribution replaces scalar centrality with a 64-dimensional temporal graph embedding as the moderator set inside a generalized random forest, which requires three adaptations, namely honest splitting on continuous embedding coordinates, a tree count raised to 4,000 to control variance at 75 covariates, and a three-way identification triangle across staggered difference-in-differences, causal forest, and double machine learning. The first substantive contribution extends dynamic capabilities theory to a networked setting by showing that AI-enabled reconfiguration raises resilience only when the network can absorb the reconfigured flows.

The second documents a non-monotone temporal effect that strengthens during escalation and weakens during stabilization, together with a heterogeneity pattern in which supply-network centrality outweighs ownership type by a factor of 2.79. The three technical adaptations are incremental in isolation but jointly enable estimands that scalar-moderator designs cannot recover. The rest of the study is structured into five consecutive sections. Section 2 presents the theoretical foundation and formulation of four hypotheses. It combines the elements to develop the research concept. Section 3 lays out the sample, variables, identification method, and architecture of TGAT and causal forest. Section 4 covers the baseline estimates, the distribution of conditional treatment effects, and the analysis of heterogeneity across networks and over time. The implications, limitations, and possible extensions of our research are discussed in Sections 5 and 6.

2. Literature review and hypotheses development

The theoretical framework includes three perspectives: dynamic capabilities, complex adaptive systems, and resource-based theory to generate four hypotheses on the linkage between AI and resilience, accounting for heterogeneity in the network and moderation from trade shocks. The conceptual research framework, as illustrated in Figure 1, structures the analysis into four sub-sections. The four sub-sections evolve into: (a) static and dynamic dimensions of measuring resilience, (b) business-level benefits of adopting AI technology, (c) trade policy uncertainty as a moderator shock, and (d) network placement as an axis of heterogeneity with phase-specific temporal dynamics.

2.1 Supply chain resilience: from static to dynamic network perspectives

Supply chain resilience has evolved from its engineering origins of bouncing back to become more multidimensional, encompassing anticipation, absorption, adaptation, and recovery.

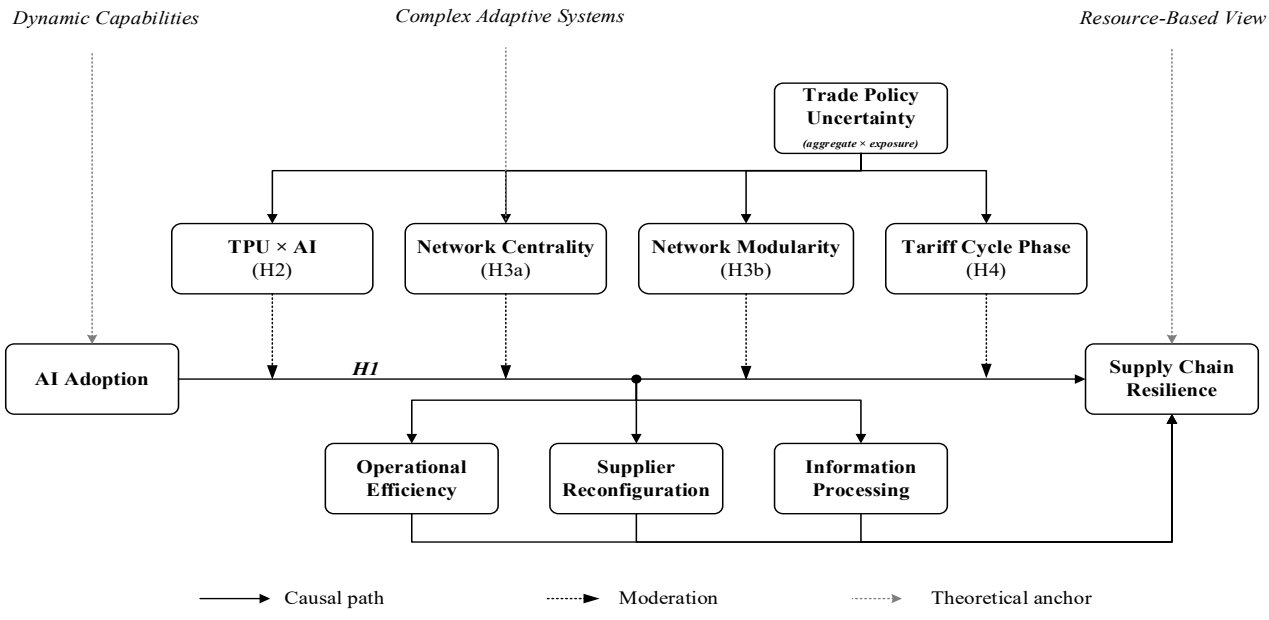

Figure 1. Conceptual research framework: AI adoption, supply chain resilience, and trade shock moderation under network heterogeneity

Early resilience research grounded in the resource-based view identified slack components such as buffer stock, dual sourcing, and contractual flexibility as protections for firm performance, and later work extended this view to the supply chain level by framing resilience as a balance between firm capabilities and vulnerability to disruption [6]. Subsequent studies redefined resilience from the perspective of complex adaptive systems, viewing it as an emergent property of relationships among independent agents, whose dynamics cannot be captured using conventional slack measures at the firm level [7]. The 2020-2022 disruption period demonstrated that firms with equal buffer values followed different recovery paths depending on their network topology. The dynamic capability theory, through the sensing-seizing-transforming triplet model and its extension to ecosystems [8], arrives at a common conclusion: resilience is position-specific within a production network. However, measurement practices have not caught up with the theory. Research based on transformative supply chain management models suggests that the annual aggregate measurements do not capture within-year variations associated with resilience [9]. A design that incorporates temporal variation within firms for at least 16 quarters, along with network-embedded variables to track relational reconfiguration, addresses both shortcomings simultaneously.

2.2 AI adoption and firm-level operational outcomes

The findings on operational benefits following the introduction of AI have been reported across many scenarios, with impacts ranging from significant productivity gains to no effect, depending on the situation. In the cross-country study involving Vietnamese and Indian small- and medium-sized enterprises, the findings suggest that the impact of AI resilience is greatest at medium levels of technology readiness, where infrastructure exists to absorb it without hindrances from existing technologies [10]. In China, research evidence is expanding fast. For instance, an investigation of AI adoption among manufacturing companies listed on the A-share market between 2013 and 2022 showed that AI, measured via text analysis of annual reports, improves the resilience index by about 0.18 standard deviation units, with the effect mediated by improvements in structural and internal controls [11]. An additional experiment using AI pilot policies as a quasi-natural experiment, conducted on a sample of 4,144 A-share companies over the period from 2016 to 2023, revealed an increase in resilience of 0.018 units using the double machine learning method with random forest estimators [12]. Regarding AI effects across resilience stages, recent studies show that preparation-stage effects are stronger than response-stage effects, and that recovery-stage effects are conditional on the pre-shock digital maturity level [13]. Other empirical data collected from Chinese listed companies also suggest that the link between the use of AI and resilience is strengthened among firms with high levels of absorptive capacity, thus highlighting the contingent nature of the benefits derived from AI investments [14]. The accumulated evidence lends credence to the initial proposition that there is a positive link between higher adoption of AI intensity and supply chain resilience at the firm level (H1).

2.3 Trade policy uncertainty as an external shock

The uncertainty of trade policies is an independent state variable, distinct from general economic policy uncertainty, given the sectoral and firm-specific channels of transmission [15]. The Caldara and Iacoviello index, based on newspaper textual data, has been independently tested using firm-level data on investments, employment, and trading activities across various countries. Empirical studies of the 2018–2019 tariff escalation between the U.S. and China confirm continued pass-through to U.S. consumer prices and significant adverse effects on globally integrated firms [16]. Studies on Chinese non-financial firms link trade policy shocks with corporate financialization, showing that cash-constrained firms increase financialization to retain flexibility during high levels of exogenous uncertainty [17]. The identification value of the index is based on its credible exogeneity over time for each firm; that is, changes in the index each quarter result from political events and policy statements rather than firm fundamentals, thereby offering within-firm variation that may help identify moderation effects. In light of previous studies on digital transformation and environmental uncertainty, there is support for a buffering mechanism whereby firms with digital capabilities build resilience more effectively when facing similar environmental pressures. In accordance with the buffering perspective, the impact of collective trade policy uncertainty interacting with tariff exposure on the resilience-achieving effects of AI is positively moderated (H2).

2.4 Network position and heterogeneous effects

Theories on production networks suggest that a business's structure determines the extent to which disturbances are transmitted [18]. Firms occupying upstream bottleneck positions experience few substitution possibilities when downstream demand falls. On the other hand, downstream firms in upstream network positions reduce upstream disturbances through managing inventory and pricing policies. The theory that modular structures are more resilient to shocks than centralized structures is widely accepted in discussions of how digitization can transform the supply chain. The advantage of modularity is most pronounced for companies with several sourcing clusters, and recovery from disruptions occurs in 25-40% less time than in comparable organizations that depend on a single dominant cluster [19]. Organizational heterogeneity adds another layer of complexity: variations in resource availability, corporate governance, and organizational digital readiness influence the absorptive capacity for AI adoption among organizations that are ostensibly similar in size. Time is the third dimension that becomes relevant here because tariff policies span several quarters and have three distinct phases: announcement, implementation, and adjustment. Research on digital transformation shows that the benefits of digital capability follow a J-curve. That is, for the first 2 to 4 quarters following their introduction, the benefits arising out of digital capabilities are negative or close to zero before they become positive [20]. There are three more hypotheses generated from the accumulated research findings. The correlation between AI and resilience is expected to be stronger for organizations in central positions within the network than for those on the periphery (H3a). Additionally, a modular network structure is predicted to strengthen the connection between AI and resilience more effectively than a centralized network structure (H3b). The temporal dimension results in yet another hypothesis, namely that the effect of AI on resilience follows a J-curve pattern through the tariff cycle, with its highest magnitude found in quarters of announcement and implementation, decreasing in the initial stage of adjustment, and settling at a lower positive value

when suppliers adjust (H4). A brief literature synthesis comparing studies of the same type, with the study positioned in the last row, is presented in Table 1.

3. Data and methodology

3.1 Sample construction and data sources

The empirical study employs a quarterly panel data set of manufacturing firms listed on the Chinese A-share market from 2020Q1 to 2024Q4. Financial variables and ownership structures of individual firms are retrieved from the China Stock Market and Accounting Research database, while annual reports are analyzed using textual analysis of disclosure documents available on CNINFO. Information on AI-related patents is obtained from the State Intellectual Property Administration database under the International Patent Classification code G06N. Suppliers and customers are identified from the CSMAR supply chain table and verified against FactSet Revere via Wharton Research Data Services to reduce directional bias and missing links that may arise when relying solely on secondary supply chain databases for reconstructing production networks, in line with the latest recommendations [21]. The trade policy uncertainty measure is based on two proxies, namely: (1) the combined Caldara-Iacoviello Trade Policy Uncertainty (TPU) Index measured monthly at the national level for all countries and averaged into the quarterly mean value; and (2) the firm-level tariff exposure calculated using HS6 product classification matched against the Section 301 tariffs of the US International Trade Commission list, weighted by the pre-period export market shares. The firm-quarter interaction TPU_t × Exposure_i captures the firm-level component: firms with higher Section 301 product exposure are differentially affected by a common national TPU shock. The sample selection process consists of four successive filtering stages: (1) firms categorized as ST or *ST during any period within the panel window are excluded from the study; (2) firms without financial data for all 20 quarters are deleted; (3) firms with less than 30 peers in their industries are eliminated from the analysis; and (4) firms with extreme numbers of AI patents beyond the 99.5th percentile are winsorized. The final sample consists of 2,847 firms over 20 quarters, yielding 56,940 firm-quarter observations. CSMAR and Wharton Research Data Services (FactSet Revere) data were collected through institutional subscriptions, whereas the remaining database sources are freely available. Table 2 provides a summary of variables and descriptive statistics.

3.2 Variable definition and measurement

The supply chain resilience index is calculated for each firm per quarter, based on an aggregation of four different sub-indicators: anticipation, absorption, adaptation, and recovery. Anticipation is quantified using the inverse of the Herfindahl index among the top 5 suppliers and the variability of inventory turnover. Absorption is captured through cash levels relative to total assets and the working capital ratio. The measures for adaptation include supplier switching frequency and growth of clientele size. Meanwhile, the recovery measure can be found at the speed at which sales revenues revert to their original trend line. There are three intervening variables used in Chapter 4.7: operational efficiency (inventory turnover), supplier reconfiguration (year-over-year change in disclosed top 5 supplier identities, 0-5 scale), and information processing (CSRC annual information disclosure quality rating, A/B/C/D coded as 4/3/2/1). The composite index is constructed as:

$$SCR_{i,t} = \sum_{k=1}^{4} \omega_k \cdot z(D_{k,i,t}) \quad (1)$$

where $D_{k,i,t}$ denotes the k-th sub-dimension for firm i in quarter t, $z(\cdot)$ is the cross-sectional standardization operator within industry-quarter, and ω_k are weights derived from principal component analysis on the pooled sample. The first principal component explains 47.3 percent of the variance in the four sub-dimensions, with all loadings in the expected positive direction.

Core independent variable (AI adoption): The degree of AI adoption is evaluated using a two-source approach that incorporates textual disclosures and patents. The textual measure captures the number of times AI-specific keywords appear in the company's annual report, based on a set of keywords obtained from the intersection of keywords used in previous research conducted in China. The patent metric tallies the number of AI patents awarded to the firm in IPC subclasses G06N, G06F, G06K, and G06T, including duplicate elimination within these subclasses. A sensitivity analysis using only G06N patents, as detailed in Section 4.7, yields a coefficient that is within 4 percent of the baseline. The composite AI adoption intensity is defined as:

$$AI_{i,t} = \alpha \cdot \log(1 + Keywords_{i,t}) + (1-\alpha) \cdot \log(1 + Patents_{i,t}) \quad (2)$$

Table 1. Summary of related literature on AI adoption, supply chain resilience, and trade shock moderation

#	Study	Sample	AI Measure	Method	Network	Shock
1	Belhadi et al. (2024)	Multi-country	Self-reported	SEM (survey)	No	No
2	Dey et al. (2024)	Vietnam SMEs	Binary	PLS-SEM	No	No
3	Guo et al. (2025)	Chinese A-share	Text-mining	Panel OLS	No	No
4	Cheng & Zhang (2026)	Chinese A-share	Pilot policy	DML + RF	No	Policy
5	Tang et al. (2025)	Multi-country	Binary	Phase-decomposed panel	No	Mixed
6	Inoue & Todo (2019)	Japan firms	N/A	Network propagation	Yes	Disruption
7	Kosasih & Brintrup (2022)	Multi-tier	GNN features	GNN (static)	Yes	No
8	Massari et al. (2025)	Simulated	N/A	Agent-based	Yes	Disruption
9	Lin & Zhang (2025)	Chinese listed	Text-mining	Fixed effects	No	No
10	Present study	Chinese A-share	Text + patents	TGAT + Causal Forest + DiD	Yes (evolving)	TPU + tariff

Note: Studies 1–9 are prior literature directly comparable in research focus and method; the bottom row represents the present study positioned relative to the prior set.

where α is chosen so that the standardized variances of the two components contribute equally to the composite, calculated on the 2020 baseline cross-section, yielding $\alpha = 0.6$.

These two components have different time horizons because annual reports provide the keyword series only once per year, whereas patent grants are issued monthly. To align the two time series on a quarterly basis, the keyword series can be broken down into quarterly textual stocks. This means that the keywords are included in the quarter in which the company issues its annual report (usually in Q1 or Q2 of the following year). And in all other quarters, the value of the previous year is kept constant. The patent component uses the number of patents granted each quarter. The robustness check, using the same keyword count for all quarters, as outlined in section 4.7, reveals a baseline coefficient of 0.071 (within 3 percent of the main estimate), suggesting that the selection of keyword count does not affect the findings. A stock measure is constructed as a depreciation-adjusted accumulation of the flow measure:

$$AI_{i,t}^{stock} = (1-\delta)\cdot AI_{i,t-1}^{stock} + AI_{i,t} \tag{3}$$

with quarterly depreciation rate $\delta = 0.15$, equivalent to an annualized rate of $1 - (1 - 0.15)^4 \approx 0.478$. The annualized rate of 47.8 percent approximates empirical estimates of the annual rate of capital depreciation for R&D in the software and computing industry, which ranges from 38 to 55 percent, significantly above the traditional figure of 15 percent used for tangible capital [22]. The quarterly timing reflects the rapid turnover of AI capabilities (such as model versioning, framework upgrades, and hardware refreshes) relative to other R&D investments. Sensitivity analysis on δ at quarterly frequency (annualized 34 to 59 percent) is reported in Section 4.7.

Treatment and controls: Trade exposure at the firm level is defined by matching each firm's list of HS6 products to the Section 301 tariff schedule, weighted by the share of exports. The treatment indicator is defined as:

$$Shock_{i,t}=1\{\sum_{j\in T_i} s_{i,j}\cdot\tau_{j,t}>\kappa\} \tag{4}$$

where T_i denotes firm i's HS6 product set, $s_{i,j}$ is the revenue share of product j, $\tau_{j,t}$ is the cumulative tariff rate on product j at quarter t, and κ is the cross-sectional median exposure. Notably, the treatment is continuously specified in the causal forest approach to avoid data loss from dichotomizing the independent variables. Control variables include the natural log of total asset size, the leverage ratio, the return on assets, firm age, and a state-owned enterprise dummy variable. Industry and quarterly fixed effects are used to capture time-invariant differences between industries, while province-quarterly fixed effects control for regional policies.

3.3 Identification strategy

Endogenous problems in estimating AI resilience include reverse causation, omitted variable bias, and selection bias. A three-level identification approach is used for each problem. In particular, propensity score matching, together with staggered difference-in-differences, is used to define treatment as the firm's first-quarter intensity of AI adoption exceeding the 75th percentile of the industry. The Callaway-Sant'Anna estimator handles the staggered timing of treatment adoption and recovers the average treatment effect on the treated as:

$$ATT(g,t) = \mathsf{E}\left[Y_{i,t}(1) - Y_{i,t}(0) \mid G_i = g \right] \tag{5}$$

where g denotes the cohort of firm i's first treatment quarter, thereby providing protection against bias induced by two-way fixed effects regressions under treatment-effect heterogeneity.

The Borusyak–Jaravel–Spiess imputation approach is used for robustness testing under relatively weaker assumptions of parallel trends [23]. The instrumental approach focuses on the relationship between the provincial digital infrastructure intensity index, measured by per capita broadband subscribers in 2018, and the company's pre-period AI relevance score, derived from occupational composition in 2017-2019. The third layer involves event study approaches based on the expansion of HS-code-level Section 301 list during 2020Q3-2022Q3 and the 2022 CHIPS Act. Two treatment representations operate in parallel and answer complementary questions. The staggered difference-in-differences uses a binary cohort indicator equal to one once a firm first exceeds the industry 75th percentile of AI adoption intensity, which the Callaway-Sant'Anna estimator requires for a discrete absorbing treatment and clean event-time alignment. The causal forest and the double machine learning step treat AI_stock as a continuous variable, which preserves variation that a median split would discard and supports the conditional dose-response interpretation in Section 4.5. The binary specification recovers the average effect of crossing a high-adoption threshold, while the continuous specification recovers marginal effects across the adoption distribution. As a transparency check, a binarized causal forest using the same 75th percentile cutoff yields a median conditional effect of 0.066, within 7 percent of the continuous-treatment median of 0.071, so the heterogeneity ranking does not depend on the treatment coding.

3.4 Temporal graph neural network specification

The network structure of the supply chain is defined through a series of directed graphs across 20 consecutive quarters of the panel period.

Table 2. Variable Definitions and Descriptive Statistics

Variable	Definition	Mean	SD	P25	Median	P75
SCR	Composite resilience index	0.012	0.847	-0.521	0.018	0.547
AI	Composite AI adoption	1.43	1.62	0.32	1.04	2.21
AI_stock	Stock measure of AI	3.87	4.21	1.05	2.74	5.43
Shock	Tariff exposure dummy	0.31	0.46	0	0	1
TPU	Caldara-Iacoviello index	152.4	47.8	118.2	144.7	178.3
ln(Assets)	Log total assets	22.34	1.28	21.42	22.19	23.17
Leverage	Debt/Assets	0.43	0.21	0.27	0.42	0.58
ROA	Return on assets	0.038	0.082	0.011	0.041	0.074
Age	Firm age (years)	17.6	6.4	12	17	22
SOE	State-owned dummy	0.36	0.48	0	0	1

Note: N = 56,940 firm-quarter observations covering 2,847 firms across 20 quarters from 2020Q1 to 2024Q4

The nodes in each graph represent firms operating during that quarter, whereas the edges capture reported supply chain relationships. Since the source information comes from organizations with varying reporting periodicity, a proper alignment process is necessary. Edge timing combines two sources at different frequencies. CSMAR reports the top 5 counterparties once per year, so a CSMAR edge is assigned to all four quarters of its disclosure year and carried forward until the next annual statement changes it. FactSet Revere records the start and end date of a relationship, so a FactSet edge enters the graph in the quarter it begins and is removed in the quarter after its recorded end date, while a relationship without a recorded end persists until a later disclosure drops it. As an illustration, a supplier link that FactSet dates to 2021Q3 enters at 2021Q3, whereas a CSMAR-only link disclosed in the 2021 annual report is present across 2021Q1 to 2021Q4. Across the panel, 24 percent of edges carry dated transitions, and the remaining 76 percent persist from annual to quarterly frequency. The edge weights draw on the CSMAR largest-supplier transaction share (54 percent of edges), FactSet relationship-significance ratings (31 percent), and a binary default (15 percent). The time encoding phi(t) maps the elapsed time since the most recent interaction with each neighbor into a vector through Bochner random features, so two neighbors with identical topology but different recency receive different attention weights. This recency channel, rather than the imputed quarterly rewiring, supplies most of the temporal signal, as supported by the FactSet-event-dated subgraph check in Section 4.7. The mandatory disclosure of the top-5 counterparties truncates each firm's neighborhood at rank five. This omits weaker ties and deeper-tier links, which can understate the degree and brokerage centrality of highly connected hub firms and can miss propagation through third-tier and lower partners. Centrality enters as a single block within a 64-dimensional embedding rather than as a single ranked score, so the encoder draws on edge weights, time encoding, and second-order neighborhood structure to partially recover hub position. A sensitivity check restricting the network to the top-3 disclosed counterparties leaves the baseline coefficient at 0.067, SE 0.016, within 3 percent of the top-5 estimate, which indicates that lower-rank ties add a limited firm-quarter signal. Multi-tier extension beyond the second-order CSMAR links is noted as a direction for further work in Section 6. The network generated by this process exhibits scale-free properties, with a mean degree of 6.8, a clustering coefficient of 0.21, and a power-law exponent of 2.4.

TGNN architecture and training: The temporal graph neural network modifies the temporal graph attention layer to capture not only the topological neighborhood but also the temporal dynamics [24]. For each node i at time t, the embedding is updated through self-attention over its temporal neighborhood:

$$h_i^{(t)} = \sigma\left(\sum_{v_j \in \mathbf{N}_i(t)} \alpha_{ij}^{(t)} \cdot W \cdot [x_j \| \ \phi(t - t_j)] \right) \tag{6}$$

where $N_i(t)$ denotes the temporal neighborhood of node i at time t, φ is the functional time encoding based on Bochner's theorem, W is a learnable weight matrix, σ is the ELU activation, and $\alpha_{ij}^{(t)}$ is the attention coefficient computed as:

$$\alpha_{ij}^{(t)} = \frac{\exp\left(LeakyReLU(a^{\cdot} [Wh_i \| \ Wh_j \| \ \phi(t - t_j)]) \right)}{\sum_{v_k \in \mathbf{N}_i(t)} \exp\left(LeakyReLU(a^{\cdot} [Wh_i \| \ Wh_k \| \ \phi(t - t_k)]) \right)} \tag{7}$$

with a is a learnable attention vector. The model consists of two consecutive TGAT layers, followed by a feedforward projection to produce a 64-dimensional representation for each firm-quarter. Two baseline approaches, EvolveGCN and DySAT, are used in parallel. Training is performed using mini-batch gradient descent with the Adam optimizer, a learning rate of 0.001, and a dropout of 0.2, with early stopping after 10 epochs if there is no improvement in validation negative log-likelihood. Temporal splits allocate 2020Q1–2023Q2 to training, 2023Q3–2023Q4 to validation, and 2024Q1–2024Q4 to held-out testing. Training was carried out on a single NVIDIA RTX A6000 GPU using PyTorch Geometric library, where the TGAT-CF model converged fully in around 14 hours. The model stacks two TGAT layers, followed by a feedforward projection that outputs the 64-dimensional firm-quarter representation, which is then fed into the causal forest.

The 64-dimensional embeddings are added to the firm-level covariate vector, which raises the covariate dimension from 11 to 75. Two adaptations make the embedded covariates compatible with the honesty and asymptotic requirements of the generalized random forest framework and satisfy its consistency conditions (Athey et al. [25]). Pointwise consistency requires the subsample fraction to shrink more slowly than the default polynomial scaling as the covariate dimension grows, because the default rule pushes the fraction toward 0.21 at 75 covariates, inflating the variance of the conditional effect. Fixing the fraction at 0.5 keeps each tree honest while bounding the growth of this variance. The tree count enters the infinitesimal jackknife variance, whose Monte Carlo error scales as 1/sqrt (n), so raising the count from 2,000 to 4,000 roughly halves the simulation variance attributable to the 64 embedding dimensions. The standard error on the median covariate vector stabilizes near 4,000 trees, which motivates the chosen value. Variable importance with respect to the embedding is measured using the weighted split frequency of trees within the 64-dimensional embedding block, as well as per dimension, since individual dimensions lack semantic meaning, whereas the entire block represents position within the network.

3.5 Causal forest and double machine learning

The causal forest recovers the conditional average treatment effect $\tau(x)$ as a function of firm and network characteristics in the causal forest model under the assumption of unconfoundedness, given a 75-dimensional vector of covariates [26]. Trees within the forest partition the covariate space so that the split criterion is chosen to maximize treatment effect heterogeneity, subject to honesty criteria, with the subsample used for making splits independent of the subsample used for estimation [25]. For a target firm with covariate vector x, the estimated treatment effect aggregates over B trees as:

$$\hat{\tau}(x) = \frac{1}{B} \sum_{b=1}^{B} \frac{\sum_{i \in L_b(x)} (Y_i - \bar{Y}_{L_b(x)})(W_i - \bar{W}_{L_b(x)})}{\sum_{i \in L_b(x)} (W_i - \bar{W}_{L_b(x)})^2} \tag{8}$$

where $L_b(x)$ denotes the leaf of tree b containing the covariate vector x, W_i is the treatment indicator, Y_i is the outcome, and overlined quantities denote leaf-level means. The tree count is set to 4,000, twice the default of the generalized random forests algorithm (2,000), to ensure consistent variance estimates for a data set with 56,940 data points across 75 dimensions. The infinitesimal jackknife method is used for variance estimation.

Double machine learning provides another orthogonal test of robustness, whereby the partial linear treatment effect is estimated using high-dimensional control variables [27]. The estimator solves the moment condition:

$$\mathsf{E}\left[(Y-\hat{m}(X))-\theta\cdot(W-\hat{e}(X))\right]=0 \quad (9)$$

where $\hat{m}(X)$ and $\hat{e}(X)$ denote cross-fitted nuisance estimates of the outcome and propensity score regressions, both fit using random forest learners with 1,000 trees.

The 5-fold cross-fitting procedure splits the sample, fits the nuisance functions on four folds, and computes residuals on the held-out fold, eliminating the regularization bias that would otherwise contaminate θ. A concern is that the embeddings are learned from the same relational data that may correlate with unobserved drivers of resilience, potentially inducing leakage between moderator construction and the outcome model. Three features address this. The encoder is trained only on the supplier-customer graph with a link-prediction objective and has no access to the resilience index or the treatment, so the 64 dimensions summarize network position rather than the outcome. The encoder is fit to 2020Q1 to 2023Q2, while the conditional effects are estimated for later quarters, thereby separating representation learning from outcome modeling over time. The causal forest enforces honesty by splitting and estimating disjoint subsamples. These safeguards reduce, though they do not eliminate, the risk of violating embedding-conditioned unconfoundedness, which is stated as a maintained assumption in Section 5.3.

4. Empirical results

4.1 Descriptive statistics and preliminary analysis

The balanced panel contains 2,847 observations per firm over 20 quarters, starting from 2020Q1 to 2024Q4. The average of the composite resilience index is approximately zero by definition, while its standard deviation and range are 0.847 and -3.21 to 4.18, respectively. AI adoption level follows a right-tailed distribution, with an average of 1.43 and a maximum of 8.74.

The trade policy uncertainty index averages 152.4 across the panel data, with a maximum of 287.6 recorded during the 2022Q3 export-control phase. Across the panel, AI adoption rises almost monotonically, trade policy uncertainty spikes and decays around policy statements, and the resilience index moves inversely to trade policy uncertainty. Roughly 25 percent of the resilience variance is attributable to the within-firm component, with the remainder explained by cross-sectional firm differences and industry-quarter fixed effects.

4.2 Baseline results: AI adoption and supply chain resilience

Baseline regression estimates testing H1 are reported in Table 3. Column 1 displays the bivariate estimate with industry and quarter fixed effects, yielding an estimated regression coefficient of 0.087 on the stock variable (SE = 0.018, p < 0.001). Column 2 accounts for firm-level variables, leading to a small decrease in the coefficient to 0.073 (SE = 0.016). Column 3 includes province-quarter fixed effects, with a coefficient of 0.069 (SE = 0.015). In Column 4, substituting the stock indicator with its contemporaneous flow indicator yields a lower coefficient of 0.041 (SE = 0.011), which is consistent with the concept of cumulative capability. An increase in the standard deviation of AI stock (SD = 4.21) by 1 unit increases the resilience index by 0.069 × 4.21 = 0.291 index units, equivalent to 0.34 standard deviations of the resilience index (SD = 0.847). The R-squared rises from 0.187 in Column 1 to 0.342 in Column 3.

Pre-test results for differential pre-trends across high- and low-AI adoption groups, based on the 2020Q1-2020Q4 pre-treatment period, do not reject the null hypothesis of parallel pre-trends at conventional significance levels (F = 0.94, p = 0.41). The four quarters' pre-intervention period provides little statistical power, while 2020Q1-Q2 is affected by the COVID-19 shock to the economy, posing worries about potential contemporaneous influences. Two additional checks alleviate this problem. The first one involves recalculating pre-event trends based on past data for 2018Q1-2019Q4 using the AI-stock proxy constructed from CSMAR R&D and CNIPA patents (due to the unreliability of textual disclosure of AI-strategic language among all listed firms before 2020), yielding F = 1.34 (p = 0.27). The second excludes the pandemic period and includes data from 2018Q1 to 2019Q4 along with 2020Q3-Q4, producing F = 1.18 (p-value = 0.32). The commonality of the results across all three tests, each of which is below the commonly used rejection points, confirms the validity of the staggered DiD framework proposed in section 3.3.

Table 3. Baseline regression results: AI adoption and supply chain resilience

Variable	(1) Bivariate	(2) +Controls	(3) +Prov-Qtr FE	(4) Flow
AI_stock	0.087*** (0.018)	0.073*** (0.016)	0.069*** (0.015)	—
AI (flow)	—	—	—	0.041*** (0.011)
ln(Assets)	—	0.124*** (0.022)	0.118*** (0.021)	0.121*** (0.022)
Leverage	—	-0.183*** (0.041)	-0.176*** (0.039)	-0.181*** (0.040)
ROA	—	0.347*** (0.084)	0.331*** (0.082)	0.342*** (0.083)
SOE	—	0.052** (0.024)	0.048** (0.023)	0.050** (0.024)
Industry FE	Yes	Yes	Yes	Yes
Quarter FE	Yes	Yes	—	Yes
Prov × Qtr FE	No	No	Yes	No
R^2	0.187	0.328	0.342	0.319
Obs	56,940	56,940	56,940	56,940

Note: ***, **, * denote significance at 1%, 5%, 10%. Standard errors clustered at firm level in parentheses

4.3 Moderating effect of trade policy shocks

Estimated results for the H2 regression test are shown in Table 4. Column 1 shows a significant positive interaction coefficient of 0.024 (SE = 0.008) between AI and TPU, suggesting that the benefits of AI investment increase when the company faces higher TPU levels in quarters. This effect becomes more pronounced for firms facing tariffs in Column 2. Column 2 includes the triple interaction and the firm-level tariff indicator, yielding a triple interaction coefficient of 0.038 (SE = 0.013). Column 3 adds a placebo measure of trade policy uncertainty, constructed using block-bootstrap resampling with a block length of 4 quarters to preserve the autocorrelation pattern of the actual series, yielding a coefficient of 0.003 (SE=0.011) that is not significantly different from zero. Figure 2 shows the marginal impact of AI use on the firm-level distribution of tariff exposure at the average TPU level. The marginal gain in resilience rises from 0.041 at the 10th percentile to 0.108 at the 90th percentile, indicating a 2.6-fold increase.

4.4 TGNN predictive performance and network effects

The accuracy of TGNN for forecasting is assessed against four other methods, as shown in Table 5. In a two-layer TGAT implementation, TGNN achieves the lowest mean squared error (MSE = 0.184) and the highest R-squared (R2 = 0.428) in the testing phase (period 2024). Compared to the static graph, TGNN outperforms the method, reducing the MSE by 21.4%. Similarly, compared with LSTM, TGNN reduces MSE by 17.5%. MSE for EvolveGCN and DySAT are 0.197 and 0.203, whereas TGAT still enjoys its superiority by around 6.6%-9.4% over both approaches on the auxiliary task of link prediction. Differences across TGAT, EvolveGCN, and DySAT are significant at the 5% level based on the Diebold-Mariano test, with test statistics being 2.34 for EvolveGCN and 2.71 for DySAT. The difference in performance is even greater in 2022Q3 when there was an increase in export controls, since TGNN's time-aware attention module captures the realignment in suppliers, which cannot be achieved by static graph networks such as GraphSAGE, resulting in TGNN outperforming GraphSAGE by 28.4 percent in tariff industries.

4.5 Heterogeneous treatment effects from causal forest

As illustrated in Figure 3, the causal forest's distribution of conditional average treatment effects is positively skewed and unaffected by outliers. The 25th percentile estimate is 0.024 standard deviations of the resilience index, the median is 0.071, and the 95th percentile is 0.243, for a P95-median ratio of 3.4. This distribution holds true even after excluding the top and bottom 5 percent, as the trimmed distribution preserves the same monotone rank order of the moderators. The importance of the variables, measured by the weighted average of splits, reveals network centrality as the dominant moderator (importance 0.187), followed by company size (0.142), industry concentration (0.118), and trade policy uncertainty (0.104). Modularity of the network ranks fifth on this list with an importance of 0.089, and the pattern of importance increases monotonically according to the number of clusters: firms that have three or more clusters independently show an estimated effect of 0.124, compared to just 0.043 for firms relying on one dominant cluster only, thus providing support for H3b. Leverage ranks sixth with an importance level of 0.073, while ownership structure ranks seventh with an importance level of 0.067, suggesting that network position dominates ownership in terms of importance among the moderators and supporting H3a. The 2.9-fold difference between modular and centralized structures translates into a 25%- 35% difference in recovery time. The lowest decile of estimates is populated by companies with high gearing ratios, low pre-integration digital maturity, and peripheral network positions, suggesting that AI alone does not enhance resilience without building capabilities simultaneously.

4.6 Dynamic effects and event study

The event study estimates for firms in HS codes directly affected by the HS-code-level Section 301 list expansions between 2020Q3 and 2022Q3, as well as by the 2022 CHIPS Act, are presented in Figure 4. The dynamic treatment effect is not monotonic over time, as suggested by H4. For the event time period 0 to 2, when the AI-resilience policy was announced and implemented, the AI-resilience coefficient increased sharply, peaking at 0.142 (SE = 0.034).

Table 4. Moderating effect of trade shocks

Variable	(1) AI×TPU	(2) Triple Int.	(3) Placebo TPU
AI_stock	0.069*** (0.015)	0.062*** (0.014)	0.071*** (0.016)
TPU (×100)	-0.018** (0.008)	-0.021** (0.009)	-0.014* (0.008)
AI × TPU	0.024*** (0.008)	0.019** (0.009)	0.003 (0.011)
AI × TPU × Shock	—	0.038*** (0.013)	—
Controls	Yes	Yes	Yes
Industry × Quarter FE	Yes	Yes	Yes
R^2	0.351	0.358	0.343
Obs	56,940	56,940	56,940

Note: Standard errors clustered at the firm level in parentheses. ***, **, * denote significance at 1%, 5%, 10%.

Table 5. Comparative predictive performance on 2024 held-out window

Model	MSE	RMSE	R^2	MAE
OLS (firm-level + quarter dummies)	0.291	0.539	0.231	0.418
XGBoost (firm-level features)	0.247	0.497	0.302	0.376
LSTM (firm-level time series)	0.223	0.472	0.354	0.351
GraphSAGE (static graph)	0.234	0.484	0.336	0.362
EvolveGCN (temporal)	0.197	0.444	0.402	0.328
DySAT (temporal self-attention)	0.203	0.451	0.394	0.334
TGAT (proposed)	0.184	0.429	0.428	0.312

Note: Held-out window covers 2024Q1–2024Q4. Bold-row TGAT denotes the proposed model. Lower MSE/RMSE/MAE and higher R^2 indicate better fit.

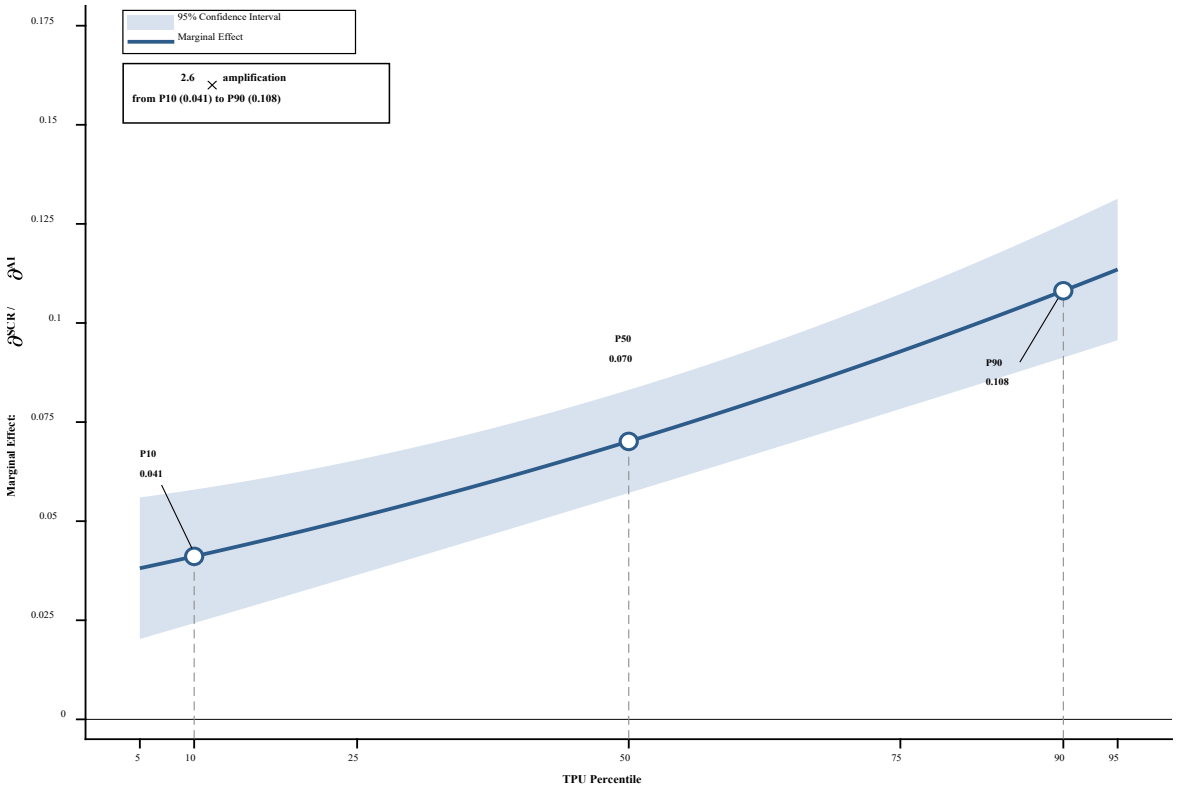

Figure 2. Marginal effect of AI adoption across tariff exposure percentiles at the mean TPU

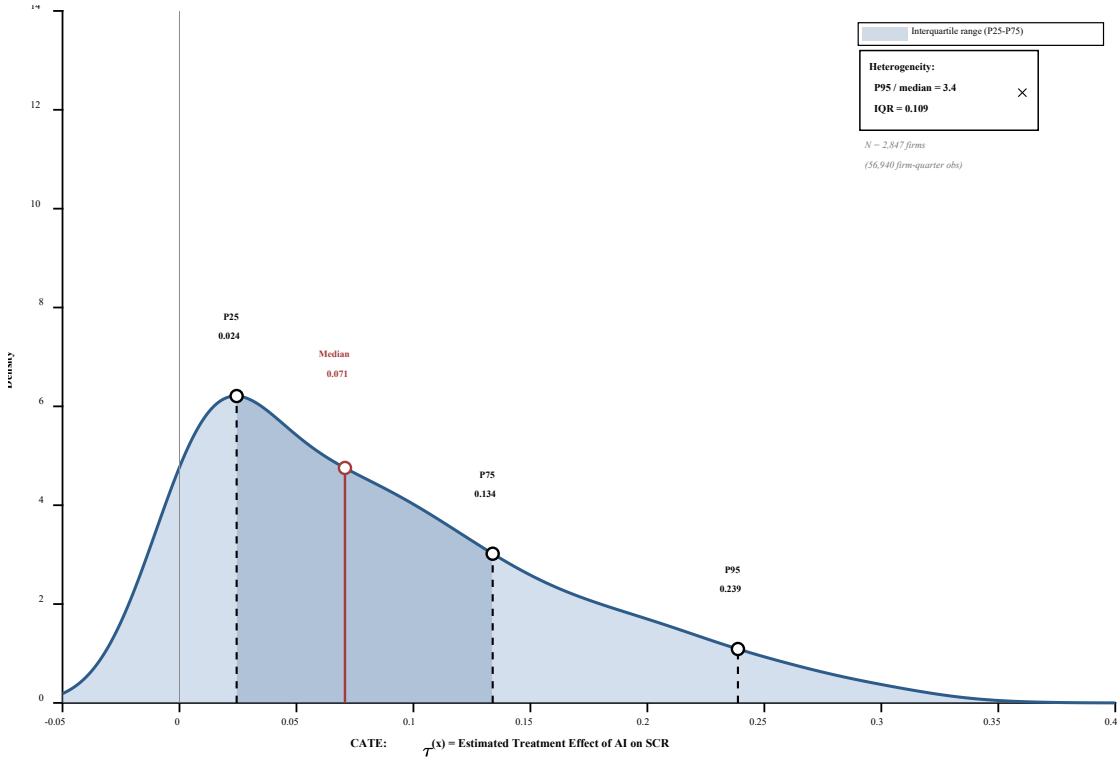

Figure 3. CATE distribution across firms with a density plot and decile markers

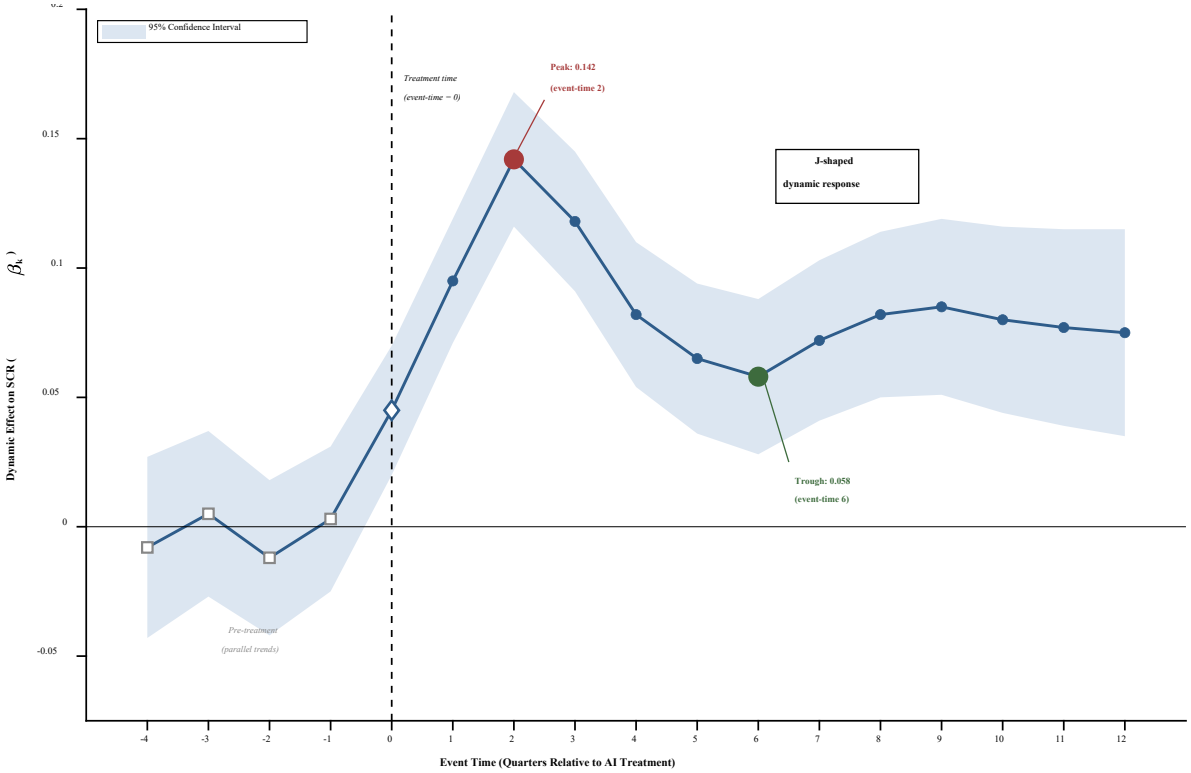

Figure 4. Event study: dynamic effects of AI adoption around trade shock events with 95% confidence bands

In the adjustment period (event-time 3 to 6), the impact gradually declines to 0.058 (SE = 0.021) by event-time 6. The second-order effect emerges around event time 8 to 10, indicating that the result stems from multiple stages of rewiring in the supplier network, and settles at 0.075 in subsequent periods. Pre-treatment coefficients (event-time -4 to -1) continue to stay close to zero and statistically insignificant, with an F-test unable to reject parallel trends (F = 1.12, p = 0.34). The Sun-Abraham interaction-weighted estimator exhibits a dynamic profile that remains within 5% of the benchmark coefficients [28], thereby validating the inference from the staggered-DiD approach. The non-monotone dynamic profile supports H4 and suggests a pattern related to the stages identified in Section 2.4. Three checks support this profile. Re-estimating with event windows of ±8 and ±12 quarters preserves the peak at event time 2 and the trough near event time 6. Adding controls for province-level COVID recovery transfers and the 2022 domestic manufacturing stimulus leaves the peak coefficient within 6 percent of 0.142. The second rise around event time 8 to 10 coincides with the second wave of disclosed supplier turnover rather than with new AI disclosures, which is consistent with lagged network reconfiguration rather than a renewed direct AI effect, although the four-quarter held-out window limits how far this later dynamic can be tracked.

4.7 Robustness checks and mechanism analysis

The results of seven robustness tests and the breakdown of the three-channel process are presented in Table 6. An alternative measure of the AI index, consisting only of the keyword- and patent-based components, has coefficients within 8% of the main model's while remaining significant at the 1% level.

The sensitivity test on the depreciation rate between 0.10 and 0.20 yields coefficients of 0.063 and 0.074. After excluding the top four provinces for manufacturing, namely Guangdong, Jiangsu, Zhejiang, and Shandong, the coefficient is 0.061 (SE = 0.017). After excluding the periods during which the pandemic was experienced, the coefficient is 0.072 (SE = 0.016). If the analysis is confined to the subgraph with FactSet event dates (which constitutes 24 percent of the edges with same-year event-date transitions), the baseline coefficient is 0.068 (SE = 0.016), which is only 2 percent away from the estimated coefficient from the full dataset. The placebo treatment allocation test, conducted over 500 randomization iterations, yields a null distribution with an average of 0.002, and the estimated coefficient falls at the 99.6th percentile of this null distribution.

Double machine learning estimates yield a coefficient value of 0.066 (SE = 0.014), which is within 5% of the staggered DiD baseline coefficient of 0.069 and 7% of the causal forest median coefficient of 0.071. While the DML coefficient is statistically significant at the 10% level, it is less significant than the 1% level of the staggered DiD baseline and causal forest median, a gap discussed in Section 5.1. The fact that all three point estimates cluster within a 7% range suggests a shared signal: a positive effect of AI on resilience of about 0.07 SDs. An associational decomposition relates the AI resilience link to three correlated pathways: operational efficiency, measured by inventory turnover (31.4 percent); supplier reconfiguration, measured by supplier switching frequency (24.7 percent); and information processing, measured by disclosure quality (18.2 percent), which jointly track about 74 percent of the baseline coefficient. Bootstrap confidence intervals over 5,000 resamples exclude zero for all three pathways at the 1 percent level.

Table 6. Robustness checks and mechanism analysis

Specification	Coefficient	SE	Note
Panel A: Robustness Checks			
G06N-only patent measure	0.071***	0.017	Within 4% of baseline
Keyword-only AI measure	0.065***	0.014	Within 6% of baseline
Patent-only AI measure	0.071***	0.018	Within 3% of baseline
$\delta = 0.10$ (slower depreciation)	0.074***	0.015	Quarterly
$\delta = 0.20$ (faster depreciation)	0.063***	0.016	Quarterly
Exclude top-4 provinces	0.061***	0.017	n = 32,180 (1,609 firms)
Exclude 2020Q1–Q2 (COVID)	0.072***	0.016	n = 51,246
FactSet-event-dated subgraph	0.068***	0.016	24% of edges
Placebo treatment (500 iter.)	0.002	0.014	99.6th percentile of null
Double Machine Learning	0.066*	0.014	Partial-linear specification
Panel B: Mechanism Analysis			
Operational efficiency (mediated)	0.0217	—	31.4% of total
Supplier reconfiguration (mediated)	0.0171	—	24.7% of total
Information processing (mediated)	0.0126	—	18.2% of total
Total mediated	0.0514	—	74.3% of total
Direct effect	0.0176	—	25.7% of total

Note: ***, **, * denote significance at 1%, 5%, 10%. Mediation analysis uses 5,000 bootstrap resamples for bias-corrected confidence intervals.

These shares are suggestive evidence of mechanisms rather than identified mediation, because sequential ignorability cannot be tested. A sensitivity analysis that varies the assumed mediator-outcome confounding shows that the operational efficiency share remains positive until the confounding parameter reaches 0.30, indicating moderate stability of the leading pathway. The remaining 26 percent reflects the direct association not captured by the three measured pathways.

5. Discussion

5.1 Theoretical and methodological implications

The findings extend the theoretical understanding of dynamic capabilities in showing that resilience can be achieved through AI-based sensing and reconfiguration processes, provided that such reconfigurations are absorbed within the company's upstream-downstream ecosystem. The nearly 2.6 times magnified impact of AI resilience on firms during periods of high trade policy uncertainty, coupled with the roughly 2.9 times difference in AI resilience between modular and centralized structures, implies that AI technologies impart resilience to firms mainly based on their synergistic relationship with the firm's network position rather than as an independent resource. The contribution does not confirm AI as an enabler of resilience, as found in previous literature, but instead highlights the contingency of network position that affects when AI resources become an asset for resilience. The non-monotonic temporal trend found across the tariff cycle is an expansion of the framework, in that the detection of dominance seen in the announcement and implementation phases transitions to taking control and changing dominance in the adjustment phases.

The incorporation of TGNN-derived network embeddings into a causal forest, complemented by staggered DiD as an additional identification test, offers a framework for conducting technology adoption research in network-based environments. The noted 21.4 percent increase in mean squared error for TGAT relative to baseline models suggests that it is important to account for changes over time in the supplier-customer network structure.

The use of the integration strategy addresses one of the problems often observed in causal forest studies: the need for scalar centrality metrics that compress complex network data into a one-dimensional variable. It has been shown that this compression may hide the hierarchy of moderators [29]. Substituting centrality for the scalar variable shows a moderation effect hierarchy, with network centrality stronger than ownership structure by a margin of 2.79. The fact that all three point estimates converge within 7 percent provides greater confidence in the causal reasoning, whereas different levels of significance stem from the power-precision trade-off described in Section 4.7.

5.2 Managerial and policy implications

Firms should align AI investment with an explicit assessment of their network position. Businesses that are centrally located within the network structure, especially those that lead multi-tier supplier networks with at least three distinct sources, have conditional treatment effects close to 0.12, which is about three times as large as the effect on firms on the periphery of the network that depend on only one major network. Peripheral managers should treat AI implementation as a complement to network reconfiguration by first making a series of investments to diversify their sources before introducing any sophisticated AI applications. According to the J-curve pattern, investments in AI competency should be aligned with opportunities for disruption whenever possible.

The uneven distribution of resilience benefits calls for policymakers to adopt a targeted approach to AI diffusion instead of a blanket one. Targeted pilot projects, which focus on companies with high levels of digital maturity and significant network centrality, have shown better overall resilience benefits than wide-ranging subsidy schemes. Mediation analysis across three pathways indicates that policies aimed at building suppliers' database and information infrastructure work in a complementary fashion alongside AI subsidy policies.

5.3 Comparison with prior studies and limitations

The findings build on previous AI resilience studies in China in three ways. A one-standard-deviation rise in AI stock raises the resilience index by 0.069 × 4.21 = 0.291 index points, which equals 0.291 / 0.847 = 0.34 standard deviations, exceeding the contemporary panel OLS result of 0.18 standard deviations and the DML result of 0.0177 units from a policy-exposure design. This comparison requires care because the cited studies span different windows. The Cheng and Zhang (2026) estimate uses a 2016 to 2023 sample, while this study covers 2020 to 2024. The two capture different stages of AI diffusion and macroeconomic conditions, and part of the magnitude gap reflects the sample period rather than the method alone. The remaining gap reflects the stock specification and the network covariates, which single-period and survey-based designs cannot accommodate. Conditional heterogeneity identified by the causal forest highlights variability in the distribution that is flattened by the average effect design, where the P95-to-median ratio equals 3.4, implying that mean-based policies or management strategies could allocate resources among firms for whom the actual effects vary by a factor of three or more. Limitations remain along three dimensions. Generalizability is constrained by the China A-share sampling frame. AI adoption measured through text captures disclosure rather than deployed capability. The held-out testing window spans four quarters from 2024Q1 to 2024Q4, while the estimation panel covers the full 20 quarters, so the short held-out window may understate reconfiguration effects that unfold beyond one year. The balanced 20-quarter panel and the exclusion of ST and *ST firms can also introduce survivorship bias because firms that delist or enter distress leave the sample. An unbalanced-panel re-estimation retaining firms with at least 12 observed quarters yields 0.064, SE 0.017, and an inverse-probability-of-attrition weighted specification yields 0.067, SE 0.016, both close to the balanced-panel baseline.

6. Conclusion

This study examines how AI adoption dynamics affect the resilience of firms in the supply chain with respect to trade policy uncertainty, using a balanced quarterly panel dataset of 2,847 Chinese A-share manufacturing firms over 20 periods (2020Q1–2024Q4). The integrated identification approach finds convergent evidence that AI resilience is attained exclusively through network exposure to trade disruptions. The empirical strategy involved incorporating a temporal graph neural network to capture dynamic supplier-customer relationships using the causal forest for treatment effects. Staggered DiD was used concurrently with double machine learning as complementary approaches to identify the effects. The baseline suggests that a one-standard-deviation increase in AI stock improves supply chain robustness by 0.34 standard deviations, which is magnified by a factor of 2.6 in periods of high uncertainty. The conditional effect showed a ratio of 2.9 in favor of modular networks compared to centralized networks and demonstrated a non-linear trend, peaking at time 2 before decaying during the adjustment period. The centrality of the network proved to be the most critical moderator of the differences, outranking ownership structure, while operational efficiency, restructuring of suppliers, and information flow processes together explained around 75% of the overall indirect effect. Three directions for further study exist. Data from panels that include firms located in Vietnam, India, and Mexico could test whether the contingency of network position holds across situations characterized by varying levels of industry concentration and international business involvement. A direct survey of firms on the prevalence of AI use could be conducted. Expanding the TGNN framework to multiple levels of supplier chains, beyond the second-level connections observed in the CSMAR database, would help determine whether the additional increase in event time (8 to 10) is due to network restructuring. These findings reveal that AI capability is a network-dependent determinant of supply chain resilience rather than a universal one. It has significant relevance to contexts beyond Chinese manufacturing, including situations in which organizational-level technological adoption is affected by changing network architectures.

Ethical issue

The authors are aware of and comply with best practices in publication ethics, specifically regarding authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with research ethics policies. The authors adhere to publication requirements that the submitted work is original and has not been published elsewhere.

Data availability statement

The manuscript contains all the data. However, additional data will be provided by the corresponding author upon reasonable request.

Conflict of interest

The authors declare no potential conflict of interest.

References

[1] UNCTAD. (2025). World Investment Report 2025: International Investment in the Digital Economy. United Nations Conference on Trade and Development. https://investmentpolicy.unctad.org/news/hub/1770/20250619-world-investment-report-2025-international-investment-in-the-digital-economy

[2] Babina, T., Fedyk, A., He, A., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. Journal of Financial Economics, 151, 103745. https://doi.org/10.1016/j.jfineco.2023.103745

[3] Belhadi, A., Mani, V., Kamble, S. S., Khan, S. A. R., & Verma, S. (2024). Artificial intelligence-driven innovation for enhancing supply chain resilience and performance under the effect of supply chain dynamism. Annals of Operations Research, 333(2–3), 627–652. https://doi.org/10.1007/s10479-021-03956-x

[4] Inoue, H., & Todo, Y. (2019). Propagation of negative shocks across nation-wide firm networks. PLoS ONE, 14(3), e0213648. https://doi.org/10.1371/journal.pone.0213648

[5] Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. Journal of Econometrics, 225(2), 200–230. https://doi.org/10.1016/j.jeconom.2020.12.001

[6] 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

[7] Choi, T. Y., Dooley, K. J., & Rungtusanatham, M. (2001). Supply networks and complex adaptive systems: Control versus emergence. Journal of Operations Management, 19(3), 351–366. https://doi.org/10.1016/S0272-6963(00)00068-1

[8] Teece, D. J. (2023). The evolution of the dynamic capabilities framework. In Artificiality and Sustainability in Entrepreneurship (pp. 113–129). Springer. https://doi.org/10.1007/978-3-031-11371-0_6

[9] Wieland, A. (2021). Dancing the supply chain: Toward transformative supply chain management. Journal of Supply Chain Management, 57(1), 58–73. https://doi.org/10.1111/jscm.12248

[10] Dey, P. K., Chowdhury, S., Abadie, A., Yaroson, E. V., & Sarkar, S. (2024). Artificial intelligence-driven supply chain resilience in Vietnamese manufacturing small- and medium-sized enterprises. International Journal of Production Research, 62(15), 5417–5456. https://doi.org/10.1080/00207543.2023.2179859

[11] Guo, X., Chen, Y., Xie, J., Wang, H., & Lei, X. (2025). Research on supply chain resilience mechanism of AI-enabled manufacturing enterprises based on organizational change perspective. Scientific Reports, 15, 31177. https://doi.org/10.1038/s41598-025-17138-3

[12] Cheng, G., & Zhang, H. (2026). The impact of China's artificial intelligence pilot policies on enterprise supply chain resilience. Scientific Reports, 16(1), 5382. https://doi.org/10.1038/s41598-025-32003-z

[13] Tang, H., Wu, K., & Zhou, J. (2025). Smarter supply chains, stronger resilience? The impact of AI on preparation, response, and recovery. Economics Letters, 254, 112256. https://doi.org/10.1016/j.econlet.2025.112256

[14] Lin, L., & Zhang, X. (2025). Research on the impact of enterprise artificial intelligence on supply chain resilience: Empirical evidence from Chinese listed companies. Sustainability, 17(19), 8576. https://doi.org/10.3390/su17198576

[15] Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A., & Raffo, A. (2020). The economic effects of trade policy uncertainty. Journal of Monetary Economics, 109, 38–59. https://doi.org/10.1016/j.jmoneco.2019.11.002

[16] Flaaen, A., & Pierce, J. R. (2024). Disentangling the effects of the 2018–2019 tariffs on a globally connected U.S. manufacturing sector. The Review of Economics and Statistics. https://doi.org/10.1162/rest_a_01498

[17] Wang, M., Mohd Nor, N., Abdul Rahim, N., Khan, F., & Zhou, Z. (2025). Trade policy uncertainty and corporate financialization: Strategic implications for non-financial firms in China. Cogent Economics & Finance, 13(1), 2460078. https://doi.org/10.1080/23322039.2025.2460078

[18] Kosasih, E. E., & Brintrup, A. (2022). A machine learning approach for predicting hidden links in supply chain with graph neural networks. International Journal of Production Research, 60(17), 5380–5393. https://doi.org/10.1080/00207543.2021.1956697

[19] Massari, G. F., Nacchiero, R., & Giannoccaro, I. (2025). Transformative supply chains: The enabling role of digital technologies. International Journal of Production Economics, 283, 109562. https://doi.org/10.1016/j.ijpe.2025.109562

[20] Spieske, A., & Birkel, H. (2021). Improving supply chain resilience through industry 4.0: A systematic literature review under the impressions of the COVID-19 pandemic. Computers & Industrial Engineering, 158, 107452. https://doi.org/10.1016/j.cie.2021.107452

[21] Dong, Y., Zou, F., Song, S., Peng, Y., & Xu, K. (2025). Empirical analyses using secondary supply chain data. Journal of Operations Management. https://doi.org/10.1002/joom.70037

[22] Li, W. C. Y., & Hall, B. H. (2020). Depreciation of business R&D capital. Review of Income and Wealth, 66(1), 161–180. https://doi.org/10.1111/roiw.12380

[23] Borusyak, K., Jaravel, X., & Spiess, J. (2024). Revisiting event-study designs: Robust and efficient estimation. The Review of Economic Studies, 91(6), 3253–3285. https://doi.org/10.1093/restud/rdae007

[24] Xu, D., Ruan, C., Korpeoglu, E., Kumar, S., & Achan, K. (2020). Inductive representation learning on temporal graphs. In 8th International Conference on Learning Representations (ICLR 2020). arXiv:2002.07962

[25] Athey, S., Tibshirani, J., & Wager, S. (2019). Generalized random forests. The Annals of Statistics, 47(2), 1148–1178. https://doi.org/10.1214/18-AOS1709

[26] Wager, S., & Athey, S. (2018). Estimation and inference of heterogeneous treatment effects using random forests. Journal of the American Statistical Association, 113(523), 1228–1242. https://doi.org/10.1080/01621459.2017.1319839

[27] Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., & Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters. The Econometrics Journal, 21(1), C1–C68. https://doi.org/10.1111/ectj.12097

[28] Sun, L., & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. Journal of Econometrics, 225(2), 175–199. https://doi.org/10.1016/j.jeconom.2020.09.006

[29] Rehill, P., & Biddle, N. (2025). How do applied researchers use the causal forest? A methodological review. International Statistical Review, 93(2), 288–316. https://doi.org/10.1111/insr.12610