

Competition and Collaboration in the AI Race: Country-Level Causal Evidence for Risk Management and Policy

Pengyu Liu¹, Stavros K. Stavroglou¹, Saurabh Mishra²,

Epameinondas Katsikas³, Athanasios A. Pantelous⁴, Bilal M. Ayyub⁵, James H. Lambert⁶

¹ University of Edinburgh Business School, University of Edinburgh, EH8 9JS, UK. P.Liu-30@sms.ed.ac.uk, Stavros.Stavroglou@ed.ac.uk

² Co-Founder and CEO, Taiyo.Ai saurabh@taiyo.ai

³ Faculty of Social Sciences, Northeastern University London, E1W 1LP, UK. e.katsikas@nulondon.ac.uk

⁴ Department of Econometrics and Business Statistics, Monash University, Clayton VIC 3800, Australia. Athanasios.Pantelous@monash.edu

⁵ Center for Technology and Systems Management, University of Maryland, College Park, MD 20742, USA. ba@umd.edu

⁶ School of Engineering Systems and Environment, University of Virginia, Charlottesville, VA 22904, USA. jhl6d@virginia.edu

Abstract

Artificial intelligence (AI) is increasingly intertwined with national economies, yet country-level causal pathways remain difficult to identify. Using annual U.S. and China data from 1980 to 2020, we triangulate Pattern Causality, Granger causality, VECM-based cointegration, and lead-lag correlation to map directional links between AI R&D and core macro aggregates with weighted networks and heatmaps. Three findings are robust. First, nonlinear AI-macro dependence is

moderate and mostly positive, which is useful for monitoring rather than as a stand-alone trigger. Second, publications and patents display broad short-run predictive links, supporting their use as early-warning signals. Third, cointegration places AI as the long-run follower, so policy should target macro fundamentals to unlock AI payoffs. Adding U.S.–China collaboration variables increases AI embeddedness more in China than in the U.S. Collaboration is a measurable policy lever, and decoupling risks foregone spillovers. The U.S. pathway is patent-driven and selective. China’s pathway is denser and collaboration-led, implying tailored risk-management priorities. Triangulation across methods reduces single-test overreach and supports risk-aware policy design.

Key Words: Artificial Intelligence; Competition-and-Collaboration view; Causality; Time-series Analysis; Risk-aware policy design

1 Introduction

1.1 Motivation & Research Gap

Artificial intelligence (AI) has expanded rapidly over recent decades. Advances in deep learning and the growing role of industry—via preferential access to substantial computing capacity, large-scale datasets, and specialized human capital—have influenced research agendas and the diffusion of models (Ahmed et al., 2023; Bier et al., 2025). The scholarly record reflects this shift: AI-engaged publications across disciplines increased by roughly thirteen-fold between 1985 and 2022 (Duede et al., 2024), and medical AI publications rose from approximately 1,600 in 2000 to more than 58,000 in 2022 (Shi et al., 2023). These indicators situate AI as a general-purpose technology with broad scientific and technical ramifications.

The policy relevance of AI follows from its complementarities with data infrastructures, cloud computing, and next-generation connectivity, and from applications in healthcare, e-commerce, manufacturing, and security-sensitive domains (Ali, 2024; Sun et al., 2025).

In parallel, AI features in geopolitical analyzes of both hard-power capabilities and soft-power influence (Mialhe, 2018). For risk analysis, central questions concern the extent to which AI-related activity is associated with movements in macroeconomic aggregates and whether such associations differ across countries with distinct capability endowments and collaboration patterns.

Despite growing attention at the firm and sector levels (e.g., Cockburn et al., 2018), macro-level evidence remains limited. The “productivity paradox” underscores this gap: measured labor productivity growth has been weak even as AI capability advanced (Brynjolfsson et al., 2017). Moreover, frontier AI entails sizable fixed inputs in compute, data curation, and specialized labor, concentrating advanced capacity within a small set of economies. Against this background, we ask: (i) *to what extent do indicators of AI activity co-move with, or help predict, macroeconomic aggregates?* (ii) *are these relationships similar across leading economies?* and (iii) *does international research collaboration operate as a transmission channel?* The U.S.–China setting is instructive: while public discourse often emphasizes rivalry, bibliometric evidence indicates that joint U.S.–China research tends to achieve higher citation impact than solo efforts (AlShebli et al., 2024). We treat collaboration as a measurable innovation asset, drawing on partnership-metrics frameworks that quantify cross-institutional and cross-border ties (Linkov et al., 2020). Whether such collaboration aligns with mutually beneficial macro-level associations remains an open empirical question.

Identifying macroeconomic relationships involving AI is challenging because AI activity co-evolves with other structural forces, can be jointly determined with growth, and frequently exhibits nonstationary dynamics that complicate linear modeling. To mitigate these issues, we adopt a triangulation strategy: a nonlinear *Pattern Causality* (PC) layer to detect directed, potentially nonlinear dependencies in multivariate time series (Stavroglou et al., 2019, 2020), complemented by *Granger* predictability (Granger, 1969), *cointegration/vector error-correction model* (VECM) analysis (Engle & Granger, 1987), and a *lead-lag correlation* layer. The dataset comprises annual indicators for the U.S. and

China over 1980–2020, including AI publications, AI patents, measures of international collaboration, and standard macroeconomic aggregates; transformations are aligned with each method’s assumptions, and conservative multiple-testing adjustments are implemented in the linear layers.

1.2 Related Work

Macro-innovation and diffusion literature link R&D to growth and spillovers, but typically without isolating AI-specific dynamics at the country level. Endogenous growth frameworks formalize innovation as a driver of long-run growth (Aghion & Howitt, 1992; Romer, 1990), while semi-endogenous models address scale effects (Jones, 1995). Cross-country evidence shows that domestic and trade-partner R&D stocks raise TFP via international spillovers (Bayoumi et al., 1999; Coe et al., 2009). Diffusion studies document systematic lags and heterogeneity in technology adoption across countries, with implications for growth differentials (Comin & Hobijn, 2004, 2010). This body of work links R&D, diffusion, and productivity, but it does not disentangle the institutional and technological mechanisms through which AI—as a general-purpose technology with distinctive fixed inputs in compute, data curation, and specialized talent—maps into macroeconomic aggregates.

Recent empirical work increasingly centres on AI itself, with evidence richer at the micro than at the macro scale. Field and laboratory experiments report sizable productivity gains from AI tools (Brynjolfsson et al., 2025; Noy & Zhang, 2023), reductions in translation frictions that expand trade (Brynjolfsson et al., 2019), sales improvements from customer-facing chatbots (Luo et al., 2019), and enhanced learning outcomes in training contexts (Gaessler & Piezunka, 2023). At broader scales, inference is fragile and sensitive to proxy choices. National or regional analyses report mixed and sometimes distributionally regressive associations: regional AI innovation correlates with declines in labor’s income share (Minniti et al., 2025), adoption gaps can widen inequality (Humlum & Vestergaard, 2025), and vacancy- or task-based measures indicate shifting job quality and composition

(Acemoglu et al., 2022; Bonfiglioli et al., 2025). These studies underscore both the promise of AI and the challenges of aggregating micro impacts into macro conclusions.

An underexplored dimension is the international organization of AI research. Bibliometric analyses identify a *collaboration premium*: U.S.-China co-authored AI papers achieve higher citation impact than matched single-country outputs, with the premium strengthening over time and across subfields (AlShebli et al., 2024; Klasa et al., 2020). This aligns with science-of-science evidence that atypical recombinations yield greater impact (Uzzi et al., 2013). Collaboration thus plausibly acts as a conduit for cross-border knowledge spillovers. The durability and macro relevance of this premium are uncertain amid tightening technogeopolitical frictions: policy interventions have curtailed collaboration in adjacent domains (Jia et al., 2024), and recent portraits suggest a shift from convergence to divergence in U.S.-China co-authorship (Kitajima & Okamura, 2025), even as selective cooperation persists (Li et al., 2024). For AI-macro questions, collaboration shocks and network re-wiring imply that spillovers are contingent and time-varying.

Four issues remain insufficiently addressed: (a) *AI specificity at macro scale*—classic R&D and diffusion results rarely isolate AI as a distinct GPT; (b) *micro-to-macro mapping*—strong micro-level gains do not translate straightforwardly into national aggregates, and macro proxies are modeling-sensitive; (c) *the collaboration channel*—the extent to which international research collaboration, notably U.S.-China, is reflected in macro-level associations remains unresolved; and (d) *methods under nonstationarity and potential bidirectionality*—linear identification is challenging, motivating triangulation across complementary tools. To address these challenges, Section 2 introduces a harmonized U.S.–China dataset, and Section 3 details a four-layer framework that triangulates nonlinear and linear evidence.

1.3 Contributions

This paper makes four main contributions to the study of AI and its macroeconomic implications.

First, we advance macro-AI mapping by integrating linear and nonlinear methodologies. We develop a unified framework that combines PC with Granger causality, cointegration via VECM, and lead-lag correlation, allowing us to triangulate directional associations between AI indicators and macroeconomic aggregates in the U.S. and China from 1980 to 2020. This integration is novel insofar as PC can capture nonlinear and hidden feedback that may be less visible to traditional tools, while cross-validation with linear models enhances robustness and mitigates identification risks. The result is a multi-layer view of AI-macro linkages that moves beyond simple correlations toward evidence on short-horizon predictability and long-run co-movement.

Second, we study cross-border collaboration as a potential macroeconomic channel. By explicitly incorporating U.S.–China AI collaboration measures into the networks, we examine how joint research activity coincides with changes in centrality rankings and directional linkages. In our sample, collaboration is associated with higher network embeddedness of AI variables—particularly for China—placing them in dual roles as transmitters and receivers of macro-relevant signals. This complements bibliometric evidence on collaboration impact with country-level economic linkages.

Third, we provide comparative evidence on asymmetries between the U.S. and China. The U.S. shows broader and more persistent associations involving AI—most prominently through patent activity—whereas China’s AI measures appear to derive disproportionate gains from collaboration, which increases their embeddedness in the macro system. Cointegration results further indicate that, in both countries, AI variables tend to adjust within long-run equilibria rather than systematically leading them. This asymmetric pattern is consistent with structural differences in innovation ecosystems and informs debates on techno-geopolitical competition and cooperation.

Fourth, we emphasize methodological transparency for risk-informed interpretation. We discuss the limits of PC under trending series and the risk of spurious links from monotonic data. By cross-validating PC with linear econometric tools and foregrounding uncertainty-aware interpretation, we provide a transparent account of identification chal-

lenges, robustness boundaries, and the implications of data frequency and sample size. This supports cautious reading of directional associations and offers guidance for using network-based inference in technology–economy settings.

The remainder of the paper is organized as follows. Section 2 introduces the macroeconomic and AI indicators, detailing their sources, coverage (1980–2020), construction procedures, and caveats. Section 3 sets out the unified framework, including data preprocessing and the four methodological modules—PC, Granger causality, VECM cointegration, and lead–lag correlation—and explains how these are used to construct directed, weighted networks for comparative analysis. Section 4 presents the empirical results through weighted-degree tables, heatmaps, and thresholded network plots, contrasting the U.S. and China with and without collaboration measures. Section 5 discusses the findings, highlights methodological limitations (such as threshold sensitivity and variable scope), and distinguishes apparent network expansion from genuine spillovers while clarifying the interpretive boundaries of network topology. Section 6 concludes by summarizing the main insights and drawing policy-relevant implications.

2 Data

We analyze 22 annual indicators for two developed, AI-leading economies—the U.S. and China. The dataset comprises two groups: (i) 16 macroeconomic series from international statistical sources—the International Monetary Fund (IMF; *World Economic Outlook*, WEO, and *International Financial Statistics*, IFS), the World Bank’s *World Development Indicators* (WDI) and the Organisation for Economic Co-operation and Development (OECD), the International Labour Organisation’s *ILOSTAT* (ILOSTAT), and the United Nations’ *World Population Prospects* (UN WPP)—and (ii) 6 AI-related variables from the Microsoft Academic Graph (MAG).¹ Because our AI measures are bibliometric and available only at annual frequency, all macroeconomic series are aggregated to annual

¹All macro series were retrieved from WEO/IFS (IMF), WDI/OECD, ILOSTAT, and UN WPP on 2024-11-15; MAG records were extracted on 2024-11-20. Full dataset versions and series codes appear in appendix A

frequency for temporal alignment. Early MAG coverage is sparse, with zero or near-zero counts in the initial years (e.g., no recorded AI patent publications for China in 1965). As the field was nascent and these observations add limited information, we restrict the sample to 1980–2020—the longest contiguous, high-quality span prior to MAG’s retirement in 2021.

Macroeconomic variable selection follows principled criteria. We prioritize breadth over redundancy (avoiding, for example, multiple GDP variants) in order to sample distinct dimensions of activity while respecting the limited number of annual observations. Indicators with extended missing sequences are excluded. Emphasis is placed on coverage consistency: we select series with complete or near-complete data for both countries from 1980–2020 to support comparability and reliability. Table 1 lists the macroeconomic indicators with definitions, units, coverage, and sources; the “Period” column reports the common window across the U.S. and China (defaulting to the shorter span when discrepant). Table 2 reports the AI-specific measures, which capture collaboration intensity, innovation output, and research activity. Although the reference window is 1980–2020, some series begin later (e.g., PPP GDP from 1990; U.S.–China collaboration counts from 1981). Each method layer uses the maximal common window implied by its variables; exact periods are listed in Tables 1 and 2.

Transformations are method-specific and applied consistently across countries. For series originally observed at higher frequency, we use sums for flows (e.g., exports, imports), arithmetic means for rates and indices (e.g., interest rates, CPI/REER), and end-of-year values for stocks (e.g., public debt, reserves). We prioritize constant-price series where available; otherwise, nominal values are deflated by the CPI (base years harmonized). Per-capita transformations use UN WPP mid-year population. Unless noted otherwise: (i) level variables measured in currency units are deflated (or taken in real/constant-price form when available) and, where appropriate, converted to per capita; (ii) strictly positive series are log-transformed; (iii) Granger and VECM layers use stationarity-consistent transformations (first differences of logs or growth rates where required), with

diagnostic tests documented in appendix C; (iv) PC operates on standardized levels with sensitivity checks for embedding parameters; and (v) the lead–lag layer uses demeaned and standardized annual series. AI outputs are identified via MAG Fields-of-Study; U.S.–China collaboration is counted once per paper (binary), while country-level publication totals use fractional author counting.²

With these indicators aligned and transformed, Section 3 describes the estimation layers and how we construct the directed networks.

3 Methodology

3.1 Design principles and overview

We aim to: (i) detect and quantify directional links between AI activity and macroeconomic aggregates for the U.S. and China, (ii) assess whether cross-border collaboration operates as a transmission channel, and (iii) report the uncertainty around these links. This section maps each research question to the estimator best suited to address it and then aggregates the resulting pairwise evidence into time-resolved, directed networks.

To limit single-method bias, consistent with risk-analytics practice for complex networks (Zhuang et al., 2022), we employ four complementary approaches, each illuminating a different aspect of the pathway from association to timing:

- **Pattern Causality (PC)** captures directional dependence in a model-light, nonlinear setting, useful when effects may be asymmetric or not well represented by linear models.
- **Granger causality (VAR)** identifies short-horizon predictive direction (“who predicts whom”) on stationarity-consistent transformations.

²Country mapping: Hong Kong SAR and Macao SAR are reported under China unless stated; Taiwan is reported separately and excluded from China aggregates. U.S. territories are included under the U.S. when affiliation metadata resolves to a U.S. ISO code.

- **Cointegration (VECM)** characterizes long-run co-movement and indicates which variable tends to adjust back toward equilibrium after shocks.
- **Lead–lag correlation** summarizes which series typically moves first over small leads and lags with minimal parametric structure.

PC is central to our design, but using it alongside the other three layers provides a more complete platform for comparing short-run predictability, long-run adjustment, nonlinear feedback, and simple timing relations, thereby triangulating evidence in a risk-aware manner. A rigorous, standardized preprocessing pipeline generates the analytical inputs for all four layers (see Figure 2 and Appendix B). The subsections that follow present each method in turn, beginning with PC.

3.2 Preprocessing and alignment

Data construction is in Section 2. Here we summarize alignment steps used by all layers:

- *Country-wise standardization.* Prior to network construction, variables are demeaned and variance-standardized within country; this places heterogeneous indicators on a comparable scale.
- *Stationarity-consistent inputs.* Granger and VECM layers use transformations motivated by unit-root diagnostics. PC uses standardized levels; lead–lag uses standardized annual series (robustness on $\Delta \log$).
- *Sample windows.* Each layer uses the maximal common window implied by its variables (reference 1980–2020; exact windows in Tables 1 and 2).

3.3 Pattern Causality

We use the PC methodology of Stavroglou et al. (2019, 2020, 2021) to detect directed, potentially nonlinear dependence between pairs of annual series without imposing a parametric model. PC is designed for short samples and complex dynamics: it embeds

each series into a low-dimensional state space, compares *local patterns* across nearest neighbors, and asks whether patterns in X tend to precede and organize patterns in Y beyond Y 's own history.

Given two standardized annual series x_t and y_t , PC reconstructs a shadow attractor for each using delay vectors (Takens embedding):

$$\mathbf{X}_t = (x_t, x_{t-\tau}, \dots, x_{t-(E-1)\tau}), \quad \mathbf{Y}_t = (y_t, y_{t-\tau}, \dots, y_{t-(E-1)\tau}), \quad (1)$$

with embedding dimension E and delay τ . For a reference time t , PC finds the k nearest neighbors of $\mathbf{Y}_t$ in $\{\mathbf{Y}_s\}$ (similar past conditions for Y), forms their small forward patterns (e.g., one-step growth rates), and checks whether the corresponding neighborhood in $\{\mathbf{X}_s\}$ displays a consistent lead pattern that helps anticipate the pattern observed in Y at horizon h .³

For each neighbor index j of $\mathbf{Y}_t$, define a local pattern vector (return vector) over the embedded window,

$$s_j^Y = \left(\frac{y_j^{(2)} - y_j^{(1)}}{y_j^{(1)}}, \dots, \frac{y_j^{(E)} - y_j^{(E-1)}}{y_j^{(E-1)}} \right), \quad (2)$$

where

$$\mathbf{Y}_{t_j} = (y_j^{(1)}, \dots, y_j^{(E)}). \quad (3)$$

Combine neighbors using distance weights

$$w_j^Y = \frac{\exp\{-d(\mathbf{Y}_t, \mathbf{Y}_{t_j})\}}{\sum_{j'} \exp\{-d(\mathbf{Y}_t, \mathbf{Y}_{t_{j'}})\}}, \quad d(\cdot, \cdot) \text{ Manhattan distance} \quad (4)$$

³Main settings: $E=2, \tau=1, h=1, k=5$; sensitivity sweeps over $E \in \{2, 3\}, \tau \in \{1, 2\}, k \in \{3, 5\}$ are reported in Appendix B. Annual data and sample length (1980–2020) motivate low-order embeddings to preserve degrees of freedom.

to obtain the *signature* (average local pattern) for Y at time t

$$P_Y(t) = \sum_{j=1}^k w_j^Y s_j^Y. \quad (5)$$

Define $P_{X \rightarrow Y}(t)$ analogously using the X -neighbors matched to the same indices $\{t_j\}$ (i.e., the X patterns observed when Y was in similar local states).

Compare $P_{X \rightarrow Y}(t)$ to the realized Y -pattern at horizon h (or to $P_Y(t)$, depending on implementation) via a similarity score; aggregate over t to form a success rate $SR_{X \rightarrow Y} \in [0, 1]$ that reflects how often the X -neighborhood “gets the Y pattern right.” Repeat with roles reversed to obtain $SR_{Y \rightarrow X}$. The asymmetry,

$$\Delta SR = SR_{X \rightarrow Y} - SR_{Y \rightarrow X}, \quad (6)$$

summarizes direction. Following Stavroglou et al. (2019, 2020, 2021), we classify the *sign* of the association using the local alignment between the X -implied pattern and the Y pattern (e.g., the sign of the inner product $\langle P_{X \rightarrow Y}(t), P_Y(t) \rangle$ averaged over t): positive (aligned), negative (anti-aligned), or *dark* (nonlinear/ambiguous, mixed sign or weak magnitude).⁴

We record a directed edge $X \rightarrow Y$ if ΔSR exceeds a small margin δ_{pc} :

$$X \rightarrow Y \text{ if } \Delta SR > \delta_{pc}, \quad w_{X \rightarrow Y}^{pc} = \max\{0, \Delta SR\} \text{ (rescaled to } [0, 1]). \quad (7)$$

Because PC is nonparametric and annual samples are short, we emphasize (i) stability of orientation across (E, τ, k) settings and (ii) block-bootstrap null comparisons for ΔSR as robustness signals (rather than sole reliance on classical p -values).

PC complements the linear layers by (i) allowing asymmetric, state-dependent responses that may be poorly captured by VAR/VECM at annual frequency and (ii) using lo-

⁴“Dark” reflects complex, non-monotone dependence that does not reduce to a stable positive/negative alignment; we report the frequency of dark classifications in Figure 5 and Figure 11

cal patterns rather than global parameterizations, which is advantageous with limited observations and possible structural change. We nonetheless treat PC outputs as associational and rely on convergence with Granger, cointegration/VECM, and lead–lag layers for risk-relevant interpretation. PC’s local-pattern approach is also relatively robust to monotone trends and outliers; we reinforce this with stability checks across (E, τ, k) and block-bootstrapped nulls (see Appendix E and De Silva and Thesmar, 2024).

Operationally, the measure of “predictive precision” is the success rate: the frequency with which the X -neighborhood correctly anticipates the local pattern of Y . Comparing $SR_{X \rightarrow Y}$ with $SR_{Y \rightarrow X}$ yields the asymmetry ΔSR (orientation). The *sign* reflects average alignment between the X -implied and realized Y patterns: *positive* (aligned), *negative* (anti-aligned), or *dark* when alignment is weak or changes sign across neighborhoods (non-monotone dependence).

Figure 1 illustrates the four PC steps: (1) embedding and nearest neighbors; (2) construction of local signatures; (3) matching X and Y neighborhoods; and (4) success-rate asymmetry and sign classification.




Figure 1: **Pattern Causality workflow.** *Step 1 (Embedding).* Reconstruct shadow attractors M_X and M_Y by stacking lagged values (Takens embedding). *Step 2 (Neighbor projection).* For each reference time t , find the k nearest neighbors of $\mathbf{Y}_t$ in M_Y and project their time indices onto M_X . *Step 3 (Local signatures).* Compute distance-weighted forward patterns (local “signatures”) at horizon h for Y and the matched neighborhood in X . *Step 4 (Orientation and sign).* Compare the $X \rightarrow Y$ implied pattern with the realized Y pattern to form success rates $SR_{X \rightarrow Y}$ and $SR_{Y \rightarrow X}$; the asymmetry $\Delta SR = SR_{X \rightarrow Y} - SR_{Y \rightarrow X}$ orients edges, while average alignment determines *positive*, *negative*, or *dark* sign. (Main settings: $E=2$, $\tau=1$, $k=5$, $h=1$; robustness in Appendix E)

Table 1: This table provides detailed definitions and metadata for the macroeconomic variables used in the analysis. Columns 1 and 2 describe each *variable* and its *conceptual meaning*; the *Notes* column clarifies construction choices; Column 4 records the available *Period* (computed for both the U.S. and China, using the shorter span when they differ); columns 5 and 6 list the *measurement units* and *ultimate data sources*.

Variable	Description/Details		Period	Units	Source
	Explanation	Notes			
Exports	Exports of goods and services comprise all cross-border transactions in goods and services, including re-exports, non-monetary gold, and merchanting.	Data follow IMF BPM6 methodology.	1982–2020	USD (current)	International Monetary Fund, <i>Balance of Payments Statistics Yearbook</i> and data files.
Imports	Imports of goods and services comprise all cross-border transactions in goods and services, including goods under merchanting, non-monetary gold, and services.	Data follow IMF BPM6 methodology.	1982–2020	USD (current)	International Monetary Fund, <i>Balance of Payments Statistics Yearbook</i> and data files.
Inflation	Annual growth rate of the GDP implicit deflator, capturing the economy-wide rate of price change.	—	1980–2020	%	World Bank national accounts data; OECD <i>National Accounts</i> (Financial Sector: exchange rates & prices).
Unemployment	Share of the labor force without work but available for and seeking employment.	Based on national labor-force surveys.	1980–2020	% of labor force	International Labor Organization, Labor Force Statistics (LFS), ILOSTAT (accessed 7 Jan 2025).
RIR	Real interest rate: lending interest rate adjusted for inflation (GDP deflator).	Nominal lending rate minus inflation.	1980–2020	%	World Bank national accounts data; OECD <i>National Accounts</i> (Financial Sector: exchange rates & prices).
Expenditure	Gross national expenditure = final consumption (households + government) + gross capital formation.	Formerly “domestic absorption.”	1980–2020	USD (current)	World Bank national accounts data; OECD <i>National Accounts</i> (USD at current prices – Expenditure on GDP).
Investment	Foreign direct investment, net inflows (investment that establishes a lasting interest, i.e., $\geq 10\%$ voting power) recorded in the balance-of-payments framework.	Series breaks may occur around 2005 with BPM6 adoption.	1982–2020	USD (current)	International Monetary Fund, <i>Balance of Payments Statistics Yearbook</i> and data files (Capital & financial account).
GDP	Gross domestic product at purchaser’s prices.	Current prices.	1980–2020	USD (current)	World Bank national accounts data; OECD <i>National Accounts</i> (Aggregate indicators).
GDP_PPP	GDP converted to international dollars using purchasing-power-parity (PPP) rates.	ICP 2021 round updates PPP conversion factors.	1990–2020	International dollars	International Comparison Program, World Bank (WDI); Eurostat–OECD PPP Programme.
Deflator	GDP implicit deflator = ratio of current- to constant-price GDP.	Base year varies by country.	1980–2020	Index	World Bank national accounts data; OECD <i>National Accounts</i> (Financial Sector: exchange rates & prices).
Population	Mid-year estimates of all residents, regardless of legal status or citizenship.	De facto population definition.	1980–2020	Persons	(1) UN Population Division, <i>World Population Prospects 2024</i> ; (2) National statistical offices; (3) Eurostat; (4) UN Statistics Division, <i>Population & Vital Statistics</i> .
BCA	Current account balance = net exports of goods & services + net primary income + net secondary income.	Key component of the balance of payments.	1982–2020	USD (current)	International Monetary Fund, <i>Balance of Payments Statistics Yearbook</i> and data files (Current account balances).
CPI	Consumer price index (2010 = 100) tracking the cost to the average consumer of a basket of goods & services.	Base 2010 = 100.	1986–2020	Index	International Monetary Fund, <i>International Financial Statistics</i> and data files (Exchange rates & prices).
Reserves	Total reserves = monetary gold + SDRs + IMF reserve position + other foreign-exchange reserves.	End-of-year stock.	1980–2020	USD (current)	International Monetary Fund, <i>International Financial Statistics</i> and data files (Reserves & other items).
Education	School enrollment, tertiary (% gross): ratio of total enrollment (regardless of age) to the population of the corresponding age group.	Gross enrollment can exceed 100% (late entry, repetition, return).	1980–2020	%	UNESCO Institute for Statistics (UIS), UIS.Stat Bulk Data Download Service (accessed 5 Apr 2025).
Savings	Gross savings = gross national income – total consumption + net transfers.	Equals gross domestic saving + net foreign transfers.	1982–2020	USD (current)	World Bank national accounts data; OECD <i>National Accounts</i> (Aggregate indicators).

Units: “USD (current)” denotes nominal values; “USD (constant)” denotes inflation-adjusted values (base year harmonized); “International dollars” are PPP-adjusted. When constant-price series are unavailable, nominal series are deflated by CPI (see, Section 2).

Table 2: Detailed definitions and metadata for the AI-specific variables used in this study. Columns 1 and 2 describe each *variable* and its *conceptual meaning*; the *Notes* column clarifies construction choices; Column 4 records the available *Period* (computed for both the U.S. and China, using the shorter span when they differ); columns 5 and 6 list the *measurement units* and *ultimate data sources*.

Variable	Description/Details		Period	Units	Source
	Explanation	Notes			
Publication_Collaborations	Annual count of AI research publications (journals <i>and</i> conferences) co-authored by at least one U.S.-affiliated and one China-affiliated author.	Collaboration counted <i>binary</i> per paper for the U.S.–China pair (each co-authored paper contributes 1). AI identified via MAG <i>Fields of Study</i> (FoS); scholarly records only (patents excluded).	1981–2020	Number	Microsoft Academic Graph (MAG); OECD.AI methodological notes (binary collaboration counting).
Conference_Collaborations	Annual count of <i>conference</i> papers meeting the same U.S.–China co-authorship criterion (subset of Publication_Collaborations).	MAG venue type = Conference; collaboration counted <i>binary</i> per paper for the U.S.–China pair; AI identified via FoS.	1981–2020	Number	MAG (conference venue filter); OECD.AI methodological notes (binary collaboration counting).
AI_Conferences	Country-attributed number of AI <i>conference</i> papers per year (reported separately for U.S. and China).	<i>Fractional author counting</i> : each paper’s credit is split equally across its authors and then aggregated by authors’ affiliation countries; totals may be non-integers. MAG venue type = Conference; AI identified via FoS.	1980–2020	Number	MAG (conference publications; FoS hierarchy); OECD.AI methodological notes (fractional author counting).
AI_Journals	Country-attributed number of AI <i>journal</i> articles per year (for U.S. and China).	Same <i>fractional author counting</i> as above; MAG venue type = Journal; AI identified via FoS.	1980–2020	Number	MAG (journal publications; FoS hierarchy); OECD.AI methodological notes (fractional author counting).
AI_Patents	Patent documents classified as AI, attributed to countries by inventor/assignee affiliations.	Country attribution allows <i>fractional</i> credit when multiple countries are present; counts refer to patent documents in MAG’s patent records aligned to AI FoS. See Appendix A for coverage notes.	1980–2020	Number	MAG patent records (see Appendix A for construction and coverage).
AI_Publications	Total AI scholarly publications per year by country (journals + conferences + other scholarly types; patents excluded).	Computed using <i>fractional author counting</i> across MAG scholarly records; AI identified via FoS; venue types aggregated except patents.	1980–2020	Number	MAG (scholarly records; FoS); OECD.AI methodological notes (fractional counting; publication types).

Notes on units and attribution: “Number” denotes counts that may be non-integers where fractional author (or inventor/assignee) attribution is applied. Collaboration variables use *binary* counting at the paper level for the U.S.–China pair. AI identification follows MAG *Fields of Study*; additional construction details (deduplication, affiliation parsing, patent coverage) are documented in Appendix B.

3.4 Granger Causality (short run)

Complementing the nonlinear, pattern-based evidence from Section 3.3, we implement a linear, forecast-based check: *does one series improve the short-horizon prediction of another after transforming to stationarity (Granger, 1969)?* Granger causality is interpreted here as *predictive content* at annual frequency.

Step 1: Preprocessing and stationarity. Let x_t denote any macroeconomic or AI series observed annually. Occasional single missing values are linearly interpolated to preserve temporal coherence. Because the VAR is estimated on (weakly) stationary inputs, we difference once

$$\Delta x_t = x_t - x_{t-1}, \quad (8)$$

and standardize

$$z_{x,t} = \frac{\Delta x_t - \mu_{\Delta x}}{\sigma_{\Delta x}}, \quad (9)$$

where $\mu_{\Delta x}$ and $\sigma_{\Delta x}$ are computed within the estimation sample. Unit-root diagnostics (ADF and PP) supporting $x_t \sim I(1)$ and $z_{x,t} \sim I(0)$ are reported in appendix C, together with deterministic-term choices.

Step 2: Forward-expanding bivariate VAR(1). For each ordered pair (y_t, x_t) and each recursion length $T_t \geq L_0 = 5$, we estimate the symmetric bivariate VAR(1) on $\{1, \dots, t\}$:

$$y_t = \alpha_0^{(t)} + \alpha_1^{(t)} y_{t-1} + \beta_1^{(t)} x_{t-1} + \varepsilon_{y,t}, \quad (10)$$

$$x_t = \gamma_0^{(t)} + \gamma_1^{(t)} x_{t-1} + \delta_1^{(t)} y_{t-1} + \varepsilon_{x,t}. \quad (11)$$

Testing $H_0^{x \rightarrow y}: \beta_1^{(t)} = 0$ against $H_1^{x \rightarrow y}: \beta_1^{(t)} \neq 0$ yields restricted and unrestricted sums of squared residuals, SSR_R and SSR_U , and the F-statistic

$$F_{x \rightarrow y,t} = \frac{(SSR_R - SSR_U)/1}{SSR_U/(T_t - 3)} \sim \mathcal{F}(1, T_t - 3), \quad (12)$$

with tail probability $p_{x \rightarrow y, t}$. Reversing roles gives $p_{y \rightarrow x, t}$. Collecting these over all pairs and t produces the panel $\{p_{i \rightarrow j, t}\}_{i, j, t}$, which we use for time-resolved visualisations (e.g., heatmaps).

Alongside the time-resolved panel, we report the full-sample test for each ordered pair using the same specification (VAR(1) on transformed series). Unless otherwise noted in figures/tables, significance thresholds are indicated in the captions. Lag-length checks, residual diagnostics, and alternative transformations are documented in Appendix C.

This linear, short-horizon layer is intentionally parsimonious given annual frequency and limited sample length. It serves as a complementary check to Section 3.3, where PC detects directional nonlinear associations, Granger tests whether that direction also manifests as incremental linear predictability after differencing.

3.5 Cointegration and VECM (long run)

Sections 3.3 and 3.4 addressed nonlinear orientation and short-horizon linear predictability. We now ask a complementary question: *do pairs of nonstationary series share a stable long-run relation, and which variable adjusts when that relation is perturbed?* Two $I(1)$ series are *cointegrated* when some static linear combination is $I(0)$, indicating an equilibrium link (Engle & Granger, 1987). When cointegration holds, a vector error-correction model (VECM) separates long-run equilibrium adjustment from short-run dynamics.

Step 1: Data preparation. Let x_t denote any macroeconomic or AI indicator observed annually. Occasional single missing values are linearly interpolated to preserve timing; no differencing is applied at this stage. Because cointegration is scale-invariant, we standardise levels within the expanding window

$$\tilde{x}_t = \frac{x_t - \mu_x}{\sigma_x}, \quad (x_t \text{ in levels}), \quad (13)$$

with μ_x, σ_x computed over $\mathcal{T}_t = \{1, \dots, t\}$. Unit-root evidence supporting $I(1)$ behavior in levels (and the deterministic terms used) is reported in Appendix C.

Step 2: Forward-expanding Engle–Granger test. For each ordered pair (y_t, x_t) and recursion length $T_t = |\mathcal{T}_t| \geq L_0 = 5$, we apply the Engle–Granger two-step procedure with a deterministic trend: (i) estimate the cointegrating regression

$$y_t = a^{(t)} + b^{(t)}x_t + u_t, \quad (14)$$

(ii) test the residuals u_t for a unit root. The routine returns the residual-based t statistic and right-tail probability $p_{y,x,t}$. Reversing roles yields $p_{x,y,t}$. Collecting these over pairs and t gives the panel $\{p_{i \rightarrow j,t}\}_{i,j,t}$ used in time-resolved displays (e.g., heatmaps); figure captions indicate the significance thresholds.

Step 3: Direction via VECM loadings (when cointegrated). When the residual test rejects a unit root, we estimate a rank-one VECM with one lag in first differences,

$$\Delta y_t = \alpha_y u_{t-1} + \sum_{k=1}^{p-1} \Gamma_{y,k} \Delta y_{t-k} + \sum_{k=1}^{p-1} \Phi_{y,k} \Delta x_{t-k} + \varepsilon_{y,t}, \quad (15)$$

$$\Delta x_t = \alpha_x u_{t-1} + \sum_{k=1}^{p-1} \Gamma_{x,k} \Delta x_{t-k} + \sum_{k=1}^{p-1} \Phi_{x,k} \Delta y_{t-k} + \varepsilon_{x,t}, \quad (16)$$

where u_{t-1} is from (14). The loading α measures the speed at which a variable corrects deviations from equilibrium. We characterize long-run leadership by the slower adjustment: the variable with the smaller absolute loading is treated as the leader, the other as the follower,

$$(\text{leader, follower}) = \begin{cases} (y, x), & |\alpha_y| < |\alpha_x|, \\ (x, y), & \text{otherwise.} \end{cases} \quad (17)$$

This yields a direction classification for cointegrated pairs, complementary to (but distinct from) short-run Granger orientation.

Unless otherwise noted, summary tables use the full-sample Engle–Granger test and the corresponding VECM loadings when cointegration is detected; time-resolved panels are for visualisation. Lag caps, deterministic terms, residual diagnostics, and robustness checks are documented in Appendix C.

In sum, the cointegration/VECM layer isolates persistent co-movements and identifies which variable tends to absorb long-run imbalances. Together with Sections 3.3 and 3.4, it completes our triangulation across nonlinear association, short-run predictability, and long-run equilibrium adjustment at annual frequency.

3.6 Lead–lag correlation (timing layer)

The previous layers address nonlinear orientation (Section 3.3), short-horizon linear predictability (Section 3.4), and long-run adjustment (Section 3.5). As a simple timing check with minimal parametric structure, we compute lead–lag correlation to indicate which series tends to move first when two stationary series co-move (Lo & MacKinlay, 1990).

Step 1: Preprocessing and stationarity. Let x_t be any macroeconomic or AI indicator observed annually. Occasional single missing values are linearly interpolated; no further imputation is applied. To obtain $I(0)$ inputs, we difference once

$$\Delta x_t = x_t - x_{t-1}, \quad (18)$$

then standardize within the estimation window

$$z_{x,t} = \frac{\Delta x_t - \mu_{\Delta x}}{\sigma_{\Delta x}}, \quad (19)$$

with $\mu_{\Delta x}$ and $\sigma_{\Delta x}$ computed over $\mathcal{T}_t = \{1, \dots, t\}$. The same transformation is applied to all variables; unit-root diagnostics appear in appendix C.

Step 2: Lead–lag correlation for a single window. For each ordered pair (y, x) and a fixed lag $L = 1$, define

$$r_{x \rightarrow y}^{(t,L)} = \text{corr}\left(z_{x,1:t-L}, z_{y,1+L:t}\right), \quad (20)$$

$$r_{y \rightarrow x}^{(t,L)} = \text{corr}\left(z_{y,1:t-L}, z_{x,1+L:t}\right). \quad (21)$$

Following the script, we report a single signed coefficient

$$\rho_{xy,t} = \begin{cases} r_{x \rightarrow y}^{(t,L)}, & |r_{x \rightarrow y}^{(t,L)}| \geq |r_{y \rightarrow x}^{(t,L)}|, \\ -r_{y \rightarrow x}^{(t,L)}, & |r_{x \rightarrow y}^{(t,L)}| < |r_{y \rightarrow x}^{(t,L)}|. \end{cases} \quad (22)$$

Hence $\rho_{xy,t} > 0$ implies that x leads y ; $\rho_{xy,t} < 0$ implies that y leads x ; and $|\rho_{xy,t}| \in [0, 1]$ measures the strength of the lead–lag correlation.

Step 3: Forward-expanding recursion. Starting from the minimum window length $L_0 = 5$, we iterate $t = 1, \dots, T$ and store, for each t , the $N \times N$ matrix $\boldsymbol{\rho}_t = [\rho_{ij,t}]$. The collection $\{\boldsymbol{\rho}_t\}_{t=L_0}^T$ provides a time-resolved map of lead–lag relations for visualisation (e.g., heatmaps).

This timing layer is intentionally parsimonious and purely correlational. It complements Section 3.3 (nonlinear orientation), Section 3.4 (short-run predictability), and Section 3.5 (long-run adjustment) by indicating which variable tends to move first when co-movement is present, aiding risk-relevant interpretation without additional parametric assumptions.

3.7 Construction of weighted networks and matrices

Having obtained four pairwise statistics—the PC success-rate asymmetries $\{\Delta SR_{ij,t}\}$, Granger p -values $\{p_{ij,t}^G\}$, Engle–Granger cointegration p -values $\{p_{ij,t}^{CI}\}$, and lead–lag correlations $\{\rho_{ij,t}\}$ —we map each into a *comparable* $N \times N$ weight matrix $W_t = [w_{ij,t}]$ with entries in the unit interval. The mapping is

$$w_{ij,t} = \phi_M(s_{ij,t}) = \begin{cases} \frac{c_{ij,t}}{100}, & M = \text{PC (positive/negative/dark)}; \\ 1 - p_{ij,t}^G, & M = \text{Granger, edge retained if } p_{ij,t}^G \leq \alpha; \\ 1 - p_{ij,t}^{CI}, & M = \text{Cointegration, edge retained if } p_{ij,t}^{CI} \leq \alpha; \\ |\rho_{ij,t}|, & M = \text{Lead–Lag, edge retained if } |\rho_{ij,t}| \geq \tau, \end{cases} \quad (23)$$

where α and τ are user-chosen significance and magnitude thresholds, respectively. For PC, the directed orientation comes from the success-rate asymmetry; negative values are clipped to zero before optional linear rescaling to $[0, 1]$.

Edge orientation. In all cases we preserve the direction inferred by the underlying layer:

- **PC:** edges are oriented by $\Delta SR_{ij,t}$ and carry the sign label (positive/negative/dark) for plotting (Section 3.3).
- **Granger:** draw the link from the predictor to the predictand (Section 3.4).
- **Cointegration/VECM:** orient from the slower-adjusting (leader) variable per VECM loadings (Section 3.5).
- **Lead-lag:** use the sign of $\rho_{ij,t}$ to direct the edge from leader to follower (Section 3.6).

Graph definition. For each time index t we obtain a directed, weighted graph

$$G_t = (V, E_t), \quad E_t = \{(i, j) : w_{ij,t} > 0\}, \quad (24)$$

stored with edge attribute **weight** = $w_{ij,t}$. Because all ϕ_M map to $[0, 1]$, subsequent network measures are dimensionally consistent across methods.

Weighted centrality. We quantify each variable’s role via weighted degrees:

$$\text{WD}_i^{\text{out}}(t) = \sum_{j \neq i} w_{ij,t}, \quad \text{WD}_i^{\text{in}}(t) = \sum_{j \neq i} w_{ji,t}, \quad \text{WD}_i(t) = \text{WD}_i^{\text{out}}(t) + \text{WD}_i^{\text{in}}(t). \quad (25)$$

Since $0 \leq w_{ij,t} \leq 1$, the maxima of Equation (25) are bounded by $N - 1$, making these centrality series numerically comparable across all four layers.

Visualization and storage. For each method we render heatmaps of W_t and network diagrams in which edge opacity (or color) is proportional to $w_{ij,t}$ and node area scales

with total degree. This unified construction allows like-for-like visual and quantitative comparison across layers without ad-hoc rescaling; thresholds (α, τ) and any sign encoding is indicated in the figure captions.

3.8 Stepwise methodological workflow

Figure 2 provides a visual summary; this subsection sets out the procedure in words.

1. **Step 1—Assemble and harmonize the panel.** We construct an annual panel for China and the U.S. covering 1980–2020, comprising 16 macroeconomic and 6 AI variables. We align definitions and units across countries, merge to a common year index, and standardize series names and aliases for clarity. The panel contains roughly 40 observations per country, and the AI set includes two China–U.S. collaboration indicators. Occasional single missing values are linearly interpolated.
2. **Step 2—Clean and preprocess the series.** We impute occasional gaps by linear interpolation with limited backward fill and standardize each country–series using forward-expanding z -scores, $z_{i,t} = [x_{i,t} - \mu_{i,1:t}] / \sigma_{i,1:t}$. We diagnose stationarity with ADF and PP tests to guide method-specific transformations (see Appendix C). PC and cointegration are applied to standardized levels ($I(1)$), whereas Granger causality and lead–lag correlation are applied to first differences ($I(0)$). Pre-tests follow the conservative rule documented in Appendix C. Estimation proceeds via a forward-expanding recursion with a minimum window length $L_0 = 5$.
3. **Step 3—Estimate the four dependence layers.**
 - (a) *Pattern Causality (PC).* We reconstruct the state space with embedding $E=2$ and delay $\tau=1$, use a one-step horizon $h=1$, and take $k=5$ nearest neighbors (robustness in Appendix E. Inputs are standardized levels. The algorithm produces three $N \times N$ percentage matrices— PC^+ , PC^- , and PC^{dark} . We convert percentages to weights by $w_{ij,t}^{\text{pc}} = c_{ij,t}/100 \in [0, 1]$ and retain the sign labels (positive/negative/dark) for plotting. Typical display cut-offs for the weights

are 0.25, 0.40, 0.55 and 0.75 (i.e., 25%, 40%, 55% 75%).

(b) *Granger causality*. We difference to $I(0)$, apply window z -scores, and estimate a bivariate VAR(1). For each ordered pair we record the right-tail F-test $p_{x \rightarrow y, t}$, set the weight $w_{ij, t}^{\text{gr}} = 1 - p_{ij, t} \in [0, 1]$, and retain edges when $p \leq \alpha$ with $\alpha \in \{0.20, 0.10, 0.05, 0.01\}$.

(c) *Cointegration / VECM*. On standardized $I(1)$ levels we run a trend-inclusive Engle–Granger test; when cointegration is detected ($p \leq \alpha$), we fit a rank-one VECM with one lag in differences. Direction is assigned by the smaller absolute loading $|\alpha|$ (leader), and weights are $w_{ij, t}^{\text{ci}} = 1 - p_{ij, t}$.

(d) *Lead–lag correlation*. With fixed lag $L=1$ we compute $\rho_{xy, t} = \text{corr}(z_{x, 1:t-L}, z_{y, 1+L:t})$. The sign indicates direction ($\rho > 0$ means x leads y ; $\rho < 0$ means y leads x), and the weight is $|\rho_{xy, t}| \in [0, 1]$. We retain edges for $|\rho_{xy, t}| \geq \tau$ with typical thresholds 0.20, 0.50, 0.80 and 0.90.

4. **Step 4—Compute weighted centralities.** From method-specific edge weights $w_{ij, t}$ we compute the out-, in-, and total weighted degrees,

$$\text{WD}_i^{\text{out}}(t) = \sum_{j \neq i} w_{ij, t}, \quad \text{WD}_i^{\text{in}}(t) = \sum_{j \neq i} w_{ji, t}, \quad \text{WD}_i(t) = \text{WD}_i^{\text{out}}(t) + \text{WD}_i^{\text{in}}(t).$$

Because all layers map into $[0, 1]$ (PC via $c/100$; Granger/VECM via $1 - p$; lead–lag via $|\rho|$), the per-node upper bound is $N - 1$, which enables direct comparability. We rank AI-only and overall nodes and compare results for China versus the U.S. and with versus without collaboration.

5. **Step 5—Map full matrices.** For each layer and country we render heatmaps of the complete $N \times N$ matrices—PC as weights, Granger/VECM as $1 - p$, and lead–lag as $|\rho|$ —and draw thresholded spring-layout networks in which node size is proportional to $\text{WD}_i(t)$, edge color/opacity is proportional to $w_{ij, t}$, and arrow directions follow the method-specific rules.

6. **Step 6—Threshold, visualize, and summarize.** We apply interpretable cut-offs to sparsify graphs (PC: 0.25/0.40/0.55/0.75; Granger/VECM: $\alpha = 0.20/0.10/0.05/0.01$; lead-lag: $|\rho| > 0.20/0.50/0.80/0.9$). We then report heatmaps, hub-focused networks, and ranked tables, and export top- k lists for each layer and time slice as CSV.
7. **Step 7—Triangulate findings.** Finally, we assess convergence across methods (PC, Granger, cointegration, and lead-lag) and representations (rankings, heatmaps, and networks), prioritizing structures that recur across layers over single-test significance.




Figure 2: Workflow overview. Left: data assembly (1980–2020, 16 macro + 6 AI), cleaning, and stationarity diagnostics (ADF/PP). Middle: four methodological layers—PC on levels; Granger on differences; Engle–Granger with VECM loadings for long-run adjustment; lead–lag correlation with fixed $L=1$. Right: mapping to weights on $[0, 1]$, construction of slice-wise matrices and thresholded networks, and weighted-degree ranking for comparison across countries and collaboration settings.

4 Empirical Findings

This section reports the empirical results in a risk-aware, triangulated manner. We proceed in four steps: (i) we present rank-based summaries of *weighted degree centrality* (out, in, and total) to establish the relative prominence of AI and macro nodes under each layer; (ii) we use *heatmaps* to place those ranks in full relational context—localizing where interactions concentrate and how AI nodes align with macro blocks; (iii) we turn to *thresholded network visualizations* to trace pathways that persist as cutoffs tighten; and (iv) we *triangulate* across layers and representations, prioritizing structures that recur over single-method signals.

Throughout, we emphasize relative movements over absolute connectivity, report *breadth* (cross-pair, cross-layer recurrence) and *depth* (survival at stricter thresholds), and distinguish PC positive/negative/dark patterns. Uncertainty is addressed via sensitivity to multiple cutoffs and cross-method corroboration, with supporting evidence cross-referenced at first mention. Cross-method and cross-country comparisons surface both robust regularities and salient divergences; we interpret these patterns in light of the literature on innovation spillovers and growth (e.g., Acemoglu et al., 2016; Griliches, 1979), clarifying the channels through which AI R&D can emerge as a central node in the macro system and how international scientific collaboration may propagate knowledge externalities that shift the roles of economic variables. The presentation flags only those findings that replicate across layers and thresholds, with explicit figure and table anchors at point of use.

4.1 Weighted-degree centrality ranking

Weighted-degree rankings order nodes by structural prominence within each layer. To make interpretation transparent, we present side-by-side rankings: the overall list for all variables and, in parallel, a compact list for AI-related nodes (AI publications, AI patents, and collaboration measures) to track their positions relative to macroeconomic aggregates

(see Table 3–Table 10).

We report results both excluding (Table 3–Table 6) and including (Table 7–Table 10) collaboration measures to assess whether—and how—their inclusion repositions nodes within the network. Specifically, we track changes in (i) *out-degree* (potential transmitters), (ii) *in-degree* (exposed receivers), and (iii) *total degree* (hubs). Because degree can rise mechanically as variables are added, we focus on *within-layer* rank shifts (baseline vs. collaboration-augmented) and on *cross-layer* recurrences (PC, Granger, VECM, lead–lag). These comparisons establish whether AI-related indicators systematically move toward transmitter or hub roles and whether macro aggregates change their exposure or influence when collaboration enters the system.⁵

4.1.1 Baseline patterns (no collaboration)

Tables 3 to 6 report baseline (no collaboration) degree rankings for both countries: Tables 3 and 4 cover China, and Tables 5 and 6 cover the U.S.. All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7).

For China, the PC *positive* layer is anchored by macro fundamentals—*Population*, *GDP (PPP)*, and *GDP (current USD)*—at the top of the out-degree rankings, while AI indicators sit lower in the AI panel (ranks 11–15 in Table 3). In PC *negative*, *FDI (net)* leads by out-degree with AI *patents* entering as a secondary transmitter (rank 4 in the all-variables panel of Table 3). In PC *dark*, AI *conference papers* and *patents* emerge as conduits (ranks 2 and 4, respectively), indicating oscillatory, nonmonotonic associations. The Granger layer shows denser short-run predictability, with AI *conference papers*, *patents*, and *journals* placed high in out-degree (AI panel, Table 3). By contrast, cointegration/VECM out-degree is dominated by prices and scale aggregates (*GDP deflator*, *Population*, *CPI*, *GDP (PPP)* in Table 3), while the in-degree side highlights AI *journals* and *conference papers* as comparatively stronger receivers (Table 4), consistent with adjustment rather than

⁵In the tables that follow, ties are broken by the underlying continuous weights (section 3.7); AI variables are highlighted; and caption notes specify the PC/Granger/VECM/lead–lag thresholds used to construct the ranks. Heatmaps and thresholded networks associated with these rankings appear in Section 4.1.1 and Section 4.1.2, enabling a direct visual check of the ranked structures.

leadership in long-run equilibria. Lead–lag correlations reinforce a limited leading role for AI: *patents* appears among the top four by out-degree, but broad directional power remains with core macro aggregates (Tables 3 and 4).

For the U.S., macro variables again anchor the PC *positive* out-degree (notably *Population*, *CPI*, and *GDP*), while AI indicators populate the upper tier of the AI panel without displacing the leading macro nodes (Table 5). PC *negative* is led by *Unemployment* (rather than capital-account items, in contrast with China). The Granger layer exhibits strong short-run linkages from AI: *research publications* ranks first in out-degree in both the AI-only and all-variables lists, with *patents* and *journals* also high (Table 5). Cointegration out-degree remains macro-centric (*GDP deflator*, *CPI*, *Population*), with AI variables comparatively smaller (Table 5). In the lead–lag layer, AI *patents* is the strongest leader by out-degree (rank 1), while its in-degree rank is more moderate (rank 7), indicating broad transmission of timing signals without being the most exposed receiver (Tables 5 and 6).

Taken together, the baseline suggests that AI variables in China have a limited transmitting role in PC *positive* and in long-run cointegration, but they feature more prominently in short-run predictability (Granger) and as receivers in equilibrium adjustment (cointegration/VECM in-degree). In the U.S., AI indicators—especially *research publications* and *patents*—occupy stronger transmitting positions in Granger and lead–lag, while long-run leadership remains with prices and scale aggregates. These contrasts provide a reference point for Section 4.1.2, where collaboration measures are introduced to test whether joint research repositions AI and macro nodes within the system.

4.1.2 Adding the collaboration channel

Introducing U.S.–China collaboration variables shifts the prominence of AI indicators relative to the no-collaboration baseline, with more pronounced changes in China than in the U.S.. For China (Tables 7 and 8), collaboration on publications and conference papers enters the top tier of PC (positive and dark) out-degree rankings, widening AI’s outward

reach, though core macro drivers (GDP, population) remain dominant. PC negative links continue to be led by FDI with limited AI involvement. In the Granger layer, breadth expands materially: all four domestic AI variables and both collaboration measures move into the upper ranks for both out- and in-degree, indicating denser short-run predictive feedback between AI and the macro block. Long-run structure remains asymmetric in the cointegration/VECM layer: collaboration variables display high in-degree but weak out-degree, consistent with adjustment to macro fundamentals rather than autonomous long-run influence. Lead-lag correlations add a modest change: publication collaboration moves among the more synchronous movers, while GDP and trade retain the largest footprints. A methodological caveat is that adding collaboration nodes increases overall connectivity and can mechanically raise centralities; we therefore emphasize *relative rank shifts* over raw magnitudes when interpreting improvements for AI indicators.

For the U.S. (Tables 9 and 10), collaboration introduces more contained adjustments. PC positive/negative, Granger, and Lead-lag rankings are broadly stable, while the PC dark layer shows the largest change: U.S.–China conference paper collaboration rises to the top in both out- and in-degree, indicating a nonlinear linkage absent in the baseline. In the cointegration/VECM layer, the same collaboration variable records the highest in-degree, suggesting that U.S. participation in joint AI activity tends to adjust to macro conditions. Patent publications also climb in in-degree, but the overall centrality of U.S. AI variables persists. As with China, the added nodes raise absolute connectivity; here too, interpretation rests on rank movements rather than levels. In sum, collaboration substantially broadens China’s AI linkages (especially for conferences and publications) and produces targeted shifts in the U.S. network, without overturning the baseline pattern in which U.S. AI—notably patent activity—remains the most consequential AI node.

Table 3: Baseline (no collaboration) *weighted out-degree* rankings for China by method. The left panel lists only AI-specific variables; the right panel shows the top four variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead-lag uses $|\rho|$. Higher out-degree indicates greater potential as a *transmitter*. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Collaboration variables between the U.S. and China are excluded.

AI-related Variables (China)				All Variables (China)			
PATTERN CAUSALITY (POS)	Ranking	Weighted-Out	Degree Centrality	PATTERN CAUSALITY (POS)	Ranking	Weighted-Out	Degree Centrality
Number of AI journal papers	11		5.439	Population, total	1		7.585
Number of AI research publications	12		5.268	GDP, PPP (current international \$)	2		7.561
Number of AI conference papers	14		3.780	GDP (current US\$)	3		7.463
Number of AI patent publications	15		2.390	Gross savings (current US\$)	4		7.220
PATTERN CAUSALITY (NEG)	Ranking	Weighted-Out	Degree Centrality	PATTERN CAUSALITY (NEG)	Ranking	Weighted-Out	Degree Centrality
Number of AI patent publications	4		1.024	Foreign direct investment, net (BoP, current US\$)	1		4.171
Number of AI research publications	6		0.902	Current account balance (BoP, current US\$)	2		1.171
Number of AI journal papers	7		0.854	Unemployment, total (% of labor force)	3		1.146
Number of AI conference papers	9		0.659	Number of AI patent publications	4		1.024
PATTERN CAUSALITY (DARK)	Ranking	Weighted-Out	Degree Centrality	PATTERN CAUSALITY (DARK)	Ranking	Weighted-Out	Degree Centrality
Number of AI conference papers	2		3.585	Unemployment, total (% of labor force)	1		4.756
Number of AI patent publications	4		3.268	Number of AI conference papers	2		3.585
Number of AI research publications	7		2.268	Foreign direct investment, net (BoP, current US\$)	3		3.366
Number of AI journal papers	9		1.976	Number of AI patent publications	4		3.268
GRANGER CAUSALITY	Ranking	Weighted-Out	Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-Out	Degree Centrality
Number of AI conference papers	3		11.970	GDP, PPP (current international \$)	1		12.435
Number of AI patent publications	4		11.009	Population, total	2		12.051
Number of AI research publications	6		10.721	Number of AI conference papers	3		11.970
Number of AI journal papers	7		10.687	Number of AI patent publications	4		11.009
DIRECTED COINTEGRATION	Ranking	Weighted-Out	Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-Out	Degree Centrality
Number of AI research publications	12		3.375	GDP deflator (base year varies by country)	1		8.788
Number of AI journal papers	16		1.141	Population, total	2		7.740
Number of AI conference papers	17		0.993	Consumer price index (2010 = 100)	3		7.669
Number of AI patent publications	18		0.464	GDP, PPP (current international \$)	4		7.239
LEAD-LAG CORRELATION	Ranking	Weighted-Out	Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-Out	Degree Centrality
Number of AI journal papers	6		6.444	Gross savings (current US\$)	1		12.337
Number of AI research publications	9		5.998	Gross national expenditure (current US\$)	2		11.090
Number of AI patent publications	10		5.982	GDP, PPP (current international \$)	3		7.396
Number of AI conference papers	17		2.632	Exports of goods and services (BoP, current US\$)	4		7.378

Table 4: Baseline (no collaboration) *weighted in-degree* rankings for China by method. The left panel lists AI-specific variables; the right panel shows the top four variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead-lag uses $|\rho|$. Higher in-degree indicates greater *exposure* as a receiver of influence. “AI research publications” denotes total AI scholarly publications (journals + conferences, excluding patents) as defined in Table 2. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Collaboration variables between the U.S. and China are excluded.

AI-related Variables (China)					All Variables (China)				
PATTERN CAUSALITY (POS)	Ranking	Weighted-In	Degree	Centrality	PATTERN CAUSALITY (POS)	Ranking	Weighted-In	Degree	Centrality
Number of AI research publications	10		5.415		GDP (current US\$)	1		8.195	
Number of AI journal papers	12		5.293		GDP, PPP (current international \$)	2		8.146	
Number of AI conference papers	14		3.000		Population, total	3		8.122	
Number of AI patent publications	16		1.683		Gross savings (current US\$)	4		7.829	
PATTERN CAUSALITY (NEG)	Ranking	Weighted-In	Degree	Centrality	PATTERN CAUSALITY (NEG)	Ranking	Weighted-In	Degree	Centrality
Number of AI research publications	11		0.732		Foreign direct investment, net (BoP, current US\$)	1		3.805	
Number of AI journal papers	13		0.585		GDP, PPP (current international \$)	2		1.000	
Number of AI conference papers	17		0.415		Population, total	3		1.000	
Number of AI patent publications	19		0.341		GDP (current US\$)	4		0.951	
PATTERN CAUSALITY (DARK)	Ranking	Weighted-In	Degree	Centrality	PATTERN CAUSALITY (DARK)	Ranking	Weighted-In	Degree	Centrality
Number of AI journal papers	12		1.902		Consumer price index (2010 = 100)	1		3.585	
Number of AI research publications	13		1.829		GDP (current US\$)	2		2.902	
Number of AI patent publications	14		1.805		Gross savings (current US\$)	3		2.829	
Number of AI conference papers	18		1.024		Gross national expenditure (current US\$)	4		2.732	
GRANGER CAUSALITY	Ranking	Weighted-In	Degree	Centrality	GRANGER CAUSALITY	Ranking	Weighted-In	Degree	Centrality
Number of AI journal papers	2		11.600		School enrollment, tertiary (% gross)	1		13.528	
Number of AI patent publications	8		10.139		Number of AI journal papers	2		11.600	
Number of AI research publications	11		9.546		Inflation, GDP deflator (annual %)	3		11.453	
Number of AI conference papers	18		5.731		GDP, PPP (current international \$)	4		11.398	
DIRECTED COINTEGRATION	Ranking	Weighted-In	Degree	Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-In	Degree	Centrality
Number of AI journal papers	1		10.971		Number of AI journal papers	1		10.971	
Number of AI conference papers	4		7.049		Current account balance (BoP, current US\$)	2		7.799	
Number of AI research publications	5		6.850		Inflation, GDP deflator (annual %)	3		7.145	
Number of AI patent publications	6		5.545		Number of AI conference papers	4		7.049	
LEAD-LAG CORRELATION	Ranking	Weighted-In	Degree	Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-In	Degree	Centrality
Number of AI patent publications	6		6.309		GDP (current US\$)	1		11.946	
Number of AI journal papers	7		6.309		Imports of goods and services (BoP, current US\$)	2		11.323	
Number of AI research publications	12		5.247		Exports of goods and services (BoP, current US\$)	3		9.051	
Number of AI conference papers	18		2.134		GDP, PPP (current international \$)	4		7.906	

Table 5: Baseline (no collaboration) *weighted out-degree* rankings for the U.S. by method. The left panel lists AI-specific variables; the right panel shows the top four variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead-lag uses $|\rho|$. Higher out-degree indicates greater potential as a *transmitter*. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Collaboration variables between the U.S. and China are excluded. Values are displayed to three decimals; rankings are computed using full-precision values. Where weights tie within display precision, ranks may repeat.

AI-related Variables (U.S.)			All Variables (U.S.)		
PATTERN CAUSALITY (POS)	Ranking	Weighted-Out Degree Centrality	PATTERN CAUSALITY (POS)	Ranking	Weighted-Out Degree Centrality
Number of AI journal papers	7	6.585	Population, total	1	7.634
Number of AI patent publications	8	6.244	GDP deflator (base year varies by country)	2	7.634
Number of AI research publications	9	6.073	Consumer price index (2010 = 100)	3	7.341
Number of AI conference papers	13	5.122	GDP, PPP (current international \$)	4	7.268
PATTERN CAUSALITY (NEG)	Ranking	Weighted-Out Degree Centrality	PATTERN CAUSALITY (NEG)	Ranking	Weighted-Out Degree Centrality
Number of AI conference papers	5	1.293	Unemployment, total (% of labor force) (national estimate)	1	4.610
Number of AI patent publications	8	0.902	Current account balance (BoP, current US\$)	2	2.902
Number of AI research publications	18	0.488	Inflation, GDP deflator (annual %)	3	2.390
Number of AI journal papers	19	0.390	Gross savings (current US\$)	4	1.585
PATTERN CAUSALITY (DARK)	Ranking	Weighted-Out Degree Centrality	PATTERN CAUSALITY (DARK)	Ranking	Weighted-Out Degree Centrality
Number of AI research publications	6	2.220	Gross savings (current US\$)	1	2.610
Number of AI conference papers	7	2.171	Total reserves (includes gold, current US\$)	1	2.610
Number of AI journal papers	9	1.732	Imports of goods & services (BoP, current US\$)	3	2.439
Number of AI patent publications	13	1.073	Inflation, GDP deflator (annual %)	4	2.415
GRANGER CAUSALITY	Ranking	Weighted-Out Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-Out Degree Centrality
Number of AI research publications	1	14.148	Number of AI research publications	1	14.148
Number of AI journal papers	4	11.137	Consumer price index (2010 = 100)	2	12.842
Number of AI patent publications	6	10.993	Foreign direct investment, net (BoP, current US\$)	3	11.853
Number of AI conference papers	11	9.657	Number of AI journal papers	4	11.137
DIRECTED COINTEGRATION	Ranking	Weighted-Out Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-Out Degree Centrality
Number of AI patent publications	10	4.990	GDP deflator (base year varies by country)	1	11.526
Number of AI conference papers	12	3.623	Consumer price index (2010 = 100)	2	11.295
Number of AI journal papers	13	3.427	Population, total	3	9.776
Number of AI research publications	14	2.180	Gross national expenditure (current US\$)	4	6.603
LEAD-LAG CORRELATION	Ranking	Weighted-Out Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-Out Degree Centrality
Number of AI patent publications	9	5.050	Gross national expenditure (current US \$)	1	9.616
Number of AI research publications	13	4.113	GDP (current US\$)	2	8.207
Number of AI journal papers	15	3.442	GDP, PPP (current international \$)	3	8.207
Number of AI conference papers	17	2.541	Gross savings (current US \$)	4	8.076

Table 6: Baseline (no collaboration) *weighted in-degree* rankings for the U.S. by method. The left panel lists AI-specific variables; the right panel shows the top four variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead-lag uses $|\rho|$ with direction given by the sign. Higher in-degree indicates greater *exposure* as a receiver of influence. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Collaboration variables between the U.S. and China are excluded. Values are displayed to three decimals; rankings are computed using full-precision values. Where weights tie within display precision, ranks may repeat.

AI-related Variables (U.S.)			All variables (U.S.)		
PATTERN CAUSALITY (POS)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (POS)	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	7	7.244	Population, total	1	8.341
Number of AI research publications	8	6.780	GDP deflator (base year varies by country)	2	8.341
Number of AI patent publications	9	5.927	Consumer price index (2010 = 100)	3	8.049
Number of AI conference papers	12	4.878	GDP, PPP (current international \$)	4	7.976
PATTERN CAUSALITY (NEG)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (NEG)	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	10	1.000	Current account balance (BoP, current US\$)	1	3.098
Number of AI research publications	11	0.976	Unemployment, total (% of total labor force) (national estimate)	2	2.317
Number of AI patent publications	12	0.854	GDP deflator (annual %)	3	2.000
Number of AI conference papers	13	0.780	Population, total	4	1.293
PATTERN CAUSALITY (DARK)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (DARK)	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	3	1.951	Population, total	1	2.268
Number of AI research publications	9	1.683	GDP deflator (base year varies by country)	2	2.268
Number of AI patent publications	14	1.171	Number of AI journal papers	3	1.951
Number of AI conference papers	15	1.122	Consumer price index (2010 = 100)	4	1.927
GRANGER CAUSALITY	Ranking	Weighted-In Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	1	12.881	Number of AI journal papers	1	12.881
Number of AI research publications	5	11.804	GDP deflator (base year varies by country)	2	12.134
Number of AI conference papers	8	10.234	Exports of goods and services (BoP, current US\$)	3	11.956
Number of AI patent publications	9	9.922	Unemployment, total (% of total labor force) (national estimate)	4	11.928
DIRECTED COINTEGRATION	Ranking	Weighted-In Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-In Degree Centrality
Number of AI patent publications	1	11.903	Number of AI patent publications	1	11.903
Number of AI conference papers	8	5.433	Inflation, GDP deflator (annual %)	2	10.826
Number of AI research publications	11	3.508	Unemployment, total (% of total labor force) (national estimate)	3	9.841
Number of AI journal papers	16	2.209	Foreign direct investment, net (BoP, current US\$)	4	8.531
LEAD-LAG CORRELATION	Ranking	Weighted-In Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-In Degree Centrality
Number of AI patent publications	7	6.429	GDP (current US\$)	1	8.468
Number of AI research publications	10	5.092	GDP, PPP (current international \$)	2	8.468
Number of AI conference papers	15	3.394	Unemployment, total (% of total labor force) (national estimate)	3	8.064
Number of AI journal papers	17	2.686	Imports of goods and services (BoP, current US\$)	4	7.508

Table 7: Collaboration (U.S.–China) *weighted out-degree* rankings for China by method. The left panel lists AI-specific variables (including China–U.S. collaboration measures); the right panel shows the top five variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead–lag uses $|\rho|$. Higher out-degree indicates greater potential as a *transmitter*. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Where weights tie within display precision, ranks may repeat.

AI-related variables (China)				All variables (China)			
PATTERN CAUSALITY (Positive)	Ranking	Weighted-Out	Degree Centrality	PATTERN CAUSALITY (Positive)	Ranking	Weighted-Out	Degree Centrality
U.S.–China publication collaborations	8		6.732	Population, total	1		8.463
Number of AI journal papers	12		6.098	GDP, PPP (current international \$)	2		8.439
Number of AI research publications	13		5.854	GDP (current US\$)	3		8.317
U.S.–China conference paper collaborations	15		4.951	Gross savings (current US\$)	4		8.049
Number of AI conference papers	16		4.268	Gross national expenditure (current US\$)	5		7.951
PATTERN CAUSALITY (Negative)	Ranking	Weighted-Out	Degree Centrality	PATTERN CAUSALITY (Negative)	Ranking	Weighted-Out	Degree Centrality
Number of AI patent publications	4		1.146	Foreign direct investment, net (BoP, current US\$)	1		4.415
Number of AI research publications	6		0.951	Current account balance (BoP, current US\$)	2		1.293
Number of AI journal papers	7		0.902	Unemployment, total (% of total labor force) (national estimate)	3		1.293
Number of AI conference papers	9		0.683	Number of AI patent publications	4		1.146
U.S.–China publication collaborations	20		0.293	School enrollment, tertiary (% gross)	5		1.024
PATTERN CAUSALITY (Dark)	Ranking	Weighted-Out	Degree Centrality	PATTERN CAUSALITY (Dark)	Ranking	Weighted-Out	Degree Centrality
Number of AI conference papers	2		3.976	Unemployment, total (% of total labor force) (national estimate)	1		5.341
Number of AI patent publications	4		3.659	Number of AI conference papers	2		3.976
U.S.–China conference paper collaborations	5		3.585	Foreign direct investment, net (BoP, current US\$)	3		3.805
U.S.–China publication collaborations	8		2.780	Number of AI patent publications	4		3.659
Number of AI research publications	10		2.537	U.S.–China conference paper collaborations	5		3.585
GRANGER CAUSALITY	Ranking	Weighted-Out	Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-Out	Degree Centrality
U.S.–China conference paper collaborations	3		13.116	GDP, PPP (current international \$)	1		14.140
Number of AI journal papers	5		12.638	Population, total	2		13.162
Number of AI research publications	6		12.632	U.S.–China conference paper collaborations	3		13.116
Number of AI patent publications	7		12.203	Gross savings (current US\$)	4		12.741
Number of AI conference papers	8		12.167	Number of AI journal papers	5		12.638
DIRECTED COINTEGRATION	Ranking	Weighted-Out	Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-Out	Degree Centrality
Number of AI research publications	12		5.026	Population, total	1		9.626
U.S.–China publication collaborations	14		2.890	GDP deflator (base year varies by country)	2		9.445
Number of AI journal papers	15		2.672	GDP, PPP (current international \$)	3		9.232
Number of AI conference papers	18		2.021	Consumer price index (2010 = 100)	4		9.101
U.S.–China conference paper collaborations	19		1.633	Exports of goods and services (BoP, current US\$)	5		8.470
LEAD-LAG CORRELATION	Ranking	Weighted-Out	Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-Out	Degree Centrality
Number of AI journal papers	6		7.123	Gross savings (current US\$)	1		13.256
Number of AI patent publications	7		6.931	Gross national expenditure (current US\$)	2		12.245
U.S.–China publication collaborations	8		6.876	GDP, PPP (current international \$)	3		8.881
U.S.–China conference collaborations	9		6.689	Population, total	4		8.243
Number of AI research publications	11		6.486	Exports of goods and services (BoP, current US\$)	5		7.535

Table 8: Collaboration (U.S.–China) *weighted in-degree* rankings for China by method. The left panel lists AI-specific variables (including U.S.–China collaboration measures); the right panel shows the top five variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead–lag uses $|\rho|$. Higher in-degree indicates greater *exposure* as a receiver of influence. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Where weights tie within display precision, ranks may repeat.

AI-related variables (China)			All variables (China)		
PATTERN CAUSALITY (Positive)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (Positive)	Ranking	Weighted-In Degree Centrality
U.S.–China publication collaborations	6	6.927	GDP (current US\$)	1	9.171
Number of AI research publications	11	6.000	GDP, PPP (current international \$)	2	9.122
Number of AI journal papers	12	5.927	Population, total	3	9.098
U.S.–China conference paper collaborations	15	4.805	Gross savings (current US\$)	4	8.707
Number of AI conference papers	16	3.390	Gross national expenditure (current US\$)	5	8.488
PATTERN CAUSALITY (Negative)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (Negative)	Ranking	Weighted-In Degree Centrality
Number of AI research publications	8	0.780	Foreign direct investment, net (BoP, current US\$)	1	4.024
U.S.–China publication collaborations	12	0.659	GDP, PPP (current international \$)	2	1.000
Number of AI journal papers	14	0.585	Population, total	3	1.000
Number of AI conference papers	19	0.415	Inflation, GDP deflator (annual %)	4	1.000
Number of AI patent publications	20	0.341	GDP (current US\$)	5	0.951
PATTERN CAUSALITY (Dark)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (Dark)	Ranking	Weighted-In Degree Centrality
U.S.–China conference paper collaborations	2	3.488	Consumer price index (2010 = 100)	1	4.073
U.S.–China publication collaborations	12	2.146	U.S.–China conference paper collaborations	2	3.488
Number of AI research publications	13	2.146	GDP (current US\$)	3	3.390
Number of AI journal papers	14	2.146	Gross savings (current US\$)	4	3.317
Number of AI patent publications	16	2.000	Gross national expenditure (current US\$)	5	3.171
GRANGER CAUSALITY	Ranking	Weighted-In Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-In Degree Centrality
U.S.–China publication collaborations	2	14.403	School enrollment, tertiary (% gross)	1	14.841
Number of AI journal papers	3	13.597	U.S.–China publication collaborations	2	14.403
Number of AI patent publications	8	11.623	Number of AI journal papers	3	13.597
U.S.–China conference paper collaborations	10	11.220	GDP, PPP (current international \$)	4	12.996
Number of AI research publications	12	10.898	Gross national expenditure (current US\$)	5	12.185
DIRECTED COINTEGRATION	Ranking	Weighted-In Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-In Degree Centrality
U.S.–China conference paper collaborations	1	15.370	U.S.–China conference paper collaborations	1	15.370
Number of AI journal papers	2	10.971	Number of AI journal papers	2	10.971
U.S.–China publication collaborations	3	10.315	U.S.–China publication collaborations	3	10.315
Number of AI patent publications	6	7.347	Inflation, GDP deflator (annual %)	4	8.402
Number of AI conference papers	7	7.049	Current account balance (BoP, current US\$)	5	7.799
LEAD-LAG CORRELATION	Ranking	Weighted-In Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-In Degree Centrality
Number of AI patent publications	6	7.259	GDP (current US\$)	1	12.533
U.S.–China conference paper collaborations	7	7.035	Imports of goods and services (BoP, current US\$)	2	11.671
U.S.–China publication collaborations	8	7.023	GDP, PPP (current international \$)	3	9.391
Number of AI research publications	9	6.857	Exports of goods and services (BoP, current US\$)	4	9.208
Number of AI journal papers	10	6.722	Population, total	5	7.286

Table 9: Collaboration (U.S.–China) *weighted out-degree* rankings for the U.S. by method. The left panel lists AI-specific variables (including U.S.–China collaboration measures); the right panel shows the top five variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead–lag uses $|\rho|$. Higher out-degree indicates greater potential as a *transmitter*. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Where weights tie within display precision, ranks may repeat.

AI-related variables (U.S.)			All variables (U.S.)		
PATTERN CAUSALITY (Positive)	Ranking	Weighted-Out Degree Centrality	PATTERN CAUSALITY (Positive)	Ranking	Weighted-Out Degree Centrality
Number of AI journal papers	7	7.244	Population, total	1	8.537
Number of AI research publications	8	6.878	GDP deflator	2	8.512
Number of AI patent publications	9	6.829	Consumer price index (2010 = 100)	3	8.122
U.S.–China publication collaborations	10	6.439	GDP (current US\$)	4	8.049
Number of AI conference papers	13	5.707	Gross national expenditure (current US\$)	5	8.049
PATTERN CAUSALITY (Negative)	Ranking	Weighted-Out Degree Centrality	PATTERN CAUSALITY (Negative)	Ranking	Weighted-Out Degree Centrality
Number of AI conference papers	5	1.366	Unemployment, total (% of total labor force) (national estimate)	1	4.976
Number of AI patent publications	7	1.000	Current account balance (BoP, current US\$)	2	3.220
Number of AI research publications	18	0.488	Inflation, GDP deflator (annual %)	3	2.585
U.S.–China publication collaborations	19	0.439	Gross savings (current US\$)	4	1.780
Number of AI journal papers	20	0.390	Number of AI conference papers	5	1.366
PATTERN CAUSALITY (Dark)	Ranking	Weighted-Out Degree Centrality	PATTERN CAUSALITY (Dark)	Ranking	Weighted-Out Degree Centrality
U.S.–China conference paper collaborations	1	3.732	U.S.–China conference paper collaborations	1	3.732
U.S.–China publication collaborations	5	2.854	Gross savings (current US\$)	2	3.073
Number of AI research publications	8	2.415	Total reserves (includes gold, current US\$)	3	2.902
Number of AI conference papers	9	2.390	Imports of goods and services (BoP, current US\$)	4	2.902
Number of AI journal papers	12	2.073	U.S.–China publication collaborations	5	2.854
GRANGER CAUSALITY	Ranking	Weighted-Out Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-Out Degree Centrality
Number of AI research publications	1	16.148	Number of AI research publications	1	16.148
Number of AI journal papers	3	13.137	Foreign direct investment, net (BoP, current US\$)	2	13.648
U.S.–China publication collaborations	5	12.496	Number of AI journal papers	3	13.137
Number of AI patent publications	6	12.462	Consumer price index (2010 = 100)	4	12.993
Number of AI conference papers	15	10.315	U.S.–China publication collaborations	5	12.496
DIRECTED COINTEGRATION	Ranking	Weighted-Out Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-Out Degree Centrality
Number of AI patent publications	11	6.229	GDP deflator (base year varies by country)	1	12.577
Number of AI journal papers	12	5.427	Consumer price index (2010 = 100)	2	12.430
Number of AI research publications	13	4.171	Population, total	3	11.373
Number of AI conference papers	14	3.853	GDP (current US\$)	4	7.826
U.S.–China publication collaborations	16	2.199	GDP, PPP (current international \$)	5	7.826
LEAD-LAG CORRELATION	Ranking	Weighted-Out Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-Out Degree Centrality
U.S.–China publication collaborations	7	6.786	Gross national expenditure (current US\$)	1	10.033
Number of AI patent publications	9	5.453	GDP, PPP (current international \$)	2	9.438
Number of AI research publications	13	4.690	GDP (current US\$)	3	9.438
U.S.–China conference collaborations	14	4.672	Gross savings (current US\$)	4	8.111
Number of AI conference papers	17	3.913	Exports of goods and services (BoP, current US\$)	5	7.199

Table 10: Collaboration (U.S.–China) *weighted in-degree* rankings for the U.S. by method. The left panel lists AI-specific variables (including U.S.–China collaboration measures); the right panel shows the top five variables overall (AI + macro). All edge weights are normalized to $[0, 1]$ for comparability across methods (see Section 3.7): PC (POS/NEG/DARK) uses percentage matrices rescaled by $c_{ij}/100$; Granger and directed cointegration use $1 - p$; Lead–lag uses $|\rho|$. Higher in-degree indicates greater *exposure* as a receiver of influence. *Abbreviations*: POS = PC positive; NEG = PC negative; DARK = PC dark; BoP = balance of payments; PPP = purchasing-power-parity; US\$ = current U.S. dollars; CPI = consumer price index. Values are displayed to three decimals; rankings are computed using full-precision values. Where weights tie within display precision, ranks may repeat.

AI-related variables (US)			All variables (US)		
PATTERN CAUSALITY (Positive)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (Positive)	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	7	7.976	GDP deflator (base year varies by country)	1	9.317
Number of AI research publications	8	7.683	Population, total	2	9.317
U.S.–China publication collaborations	9	6.732	Consumer price index (2010 = 100)	3	8.927
Number of AI patent publications	10	6.463	GDP (current US\$)	4	8.878
Number of AI conference papers	12	5.512	GDP, PPP (current international \$)	5	8.878
PATTERN CAUSALITY (Negative)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (Negative)	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	10	1.000	Current account balance (BoP, current US\$)	1	3.390
Number of AI research publications	12	0.976	Unemployment, total (% of total labor force)	2	2.561
U.S.–China publication collaborations	11	0.976	Inflation, GDP deflator (annual %)	3	2.146
Number of AI patent publications	13	0.854	Population, total	4	1.293
Number of AI conference papers	14	0.780	GDP deflator (base year varies by country)	5	1.268
PATTERN CAUSALITY (Dark)	Ranking	Weighted-In Degree Centrality	PATTERN CAUSALITY (Dark)	Ranking	Weighted-In Degree Centrality
U.S.–China conference paper collaborations	1	3.585	U.S.–China conference paper collaborations	1	3.585
Number of AI journal papers	5	2.366	GDP deflator (base year varies by country)	2	2.732
U.S.–China publication collaborations	10	2.024	Population, total	3	2.732
Number of AI research publications	11	1.927	Consumer price index (2010 = 100)	4	2.390
Number of AI patent publications	15	1.512	Number of AI journal papers	5	2.366
GRANGER CAUSALITY	Ranking	Weighted-In Degree Centrality	GRANGER CAUSALITY	Ranking	Weighted-In Degree Centrality
Number of AI journal papers	1	14.871	Number of AI journal papers	1	14.871
Number of AI research publications	3	13.687	Unemployment, total (% of total labor force)	2	13.725
Number of AI conference papers	7	11.813	Number of AI research publications	3	13.687
U.S.–China publication collaborations	10	10.594	Exports of goods and services (BoP, current US\$)	4	13.246
Number of AI patent publications	11	10.580	GDP deflator (base year varies by country)	5	12.342
DIRECTED COINTEGRATION	Ranking	Weighted-In Degree Centrality	DIRECTED COINTEGRATION	Ranking	Weighted-In Degree Centrality
U.S.–China conference paper collaborations	1	12.924	U.S.–China conference paper collaborations	1	12.924
Number of AI patent publications	3	11.903	Inflation, GDP deflator (annual %)	2	11.983
U.S.–China publication collaborations	6	8.520	Number of AI patent publications	3	11.903
Number of AI conference papers	10	5.458	Unemployment, total (% of total labor force)	4	11.054
Number of AI research publications	13	3.508	Foreign direct investment, net (BoP, current US\$)	5	8.825
LEAD-LAG CORRELATION	Ranking	Weighted-In Degree Centrality	LEAD-LAG CORRELATION	Ranking	Weighted-In Degree Centrality
Number of AI research publications	7	7.183	GDP (current US\$)	1	8.973
Number of AI patent publications	8	6.832	GDP, PPP (current international \$)	2	8.973
U.S.–China conference collaborations	11	5.086	Unemployment, total (% of total labor force) (national estimate)	3	8.343
U.S.–China publication collaborations	12	4.805	Gross national expenditure (current US\$)	4	8.111
Number of AI conference papers	17	4.022	Imports of goods and services (BoP, current US\$)	5	7.643

4.2 Full-Sample Intensity Mapping

Having established rank-based prominence in Section 4.1, we now place those results in full relational context. For each method and country, we plot heatmaps of all pairwise interactions among the complete variable set (including collaboration variables), so that readers can see *where* links concentrate and *how* AI nodes align with macro blocks (Figure 3–Figure 14). In line with Section 3.7, these panels display the *raw*, *method-specific* statistics (PC percentages; Granger/cointegration p -values; lead–lag correlations), not the rescaled $[0, 1]$ weights used for centrality.

Rows are *senders* and columns are *receivers*. Variables are ordered with macroeconomic indicators first and AI-related variables last. With this convention, the *top-left* block shows **macro**→**macro** links; the *bottom-right* shows **AI**→**AI**; the *top-right* block shows **macro**→**AI**; and the *bottom-left* block shows **AI**→**macro**. This ordering isolates the AI–macro interfaces in the two off-diagonal blocks and makes directionality visually salient via asymmetry around the main diagonal.

Each panel uses a method-appropriate legend: (i) PC (positive/negative/dark) encodes the frequency of local-pattern matches in % on a 0–100 scale; (ii) Granger and Engle–Granger/VECM panels encode p -values on $[0, 1]$ (darker $\Rightarrow$ lower p , hence stronger evidence); and (iii) lead–lag panels encode correlation coefficients on $[-1, 1]$ (magnitude drives intensity; the sign is shown by the diverging palette). In all cases, darker shades indicate stronger connections (higher PC frequency, lower p , or larger $|\rho|$). Legends on each subplot make scales explicit.

Two regularities help interpret the figures alongside the rankings:

- *PC versus linear layers.* In the PC positive panels, the AI blocks are generally lighter than the macro blocks, consistent with fewer nonlinear causal matches at annual frequency. By contrast, the Granger and cointegration/VECM panels show systematically darker AI–macro interfaces, indicating tighter linear predictive and

long-run adjustment links than the PC layer suggests at the same horizon.

- *Country contrasts.* Directional asymmetry is visible in the off-diagonal blocks. For China, the PC negative panel concentrates darkness around investment-related senders (e.g., FDI), while for the U.S. the darkest cells shift toward unemployment or the current account balance (BCA). These patterns echo the baseline rankings in Section 4.1.1 and foreshadow the collaboration effects in Section 4.1.2.

Overall, the heatmaps corroborate the rank results: AI nodes tend to exhibit broader *linear* connectivity (short-run predictability and occasional long-run adjustment) than *nonlinear* pattern alignment, with clearer outward influence in the U.S. panels. The off-diagonal blocks make the AI↔macro interfaces interpretable at a glance, while the legends ensure comparability across methods within each row of Figure 3–Figure 14.




Figure 3: Comprehensive PC (positive) heatmap for China. Weights use causality-frequency percentages c_{ij} . Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise positive pattern causality.

Importantly, Figure 3–Figure 14 help to distinguish breadth—how many linkages are activated—from depth—the persistence and intensity of specific dyads across multiple methods. Building on this system-wide connection mapping, the following section applies




Figure 4: Comprehensive PC (negative) heatmap for China. Weights use causality-frequency percentages c_{ij} . Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise negative pattern causality.




Figure 5: Comprehensive PC (dark) heatmap for China. Weights use causality-frequency percentages c_{ij} . Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise dark pattern causality.




Figure 6: Comprehensive Granger causality heatmap for China. Weights use p value. Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes smaller p value.




Figure 7: Comprehensive cointegration/VECM heatmap for China. Weights use p value. Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes smaller p value.




Figure 8: Comprehensive Lead-Lag heatmap for China. Weights use correlation coefficient $|\rho|$. Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise correlation.




Figure 9: Comprehensive PC (positive) heatmap for the U.S.. Weights use causality-frequency percentages c_{ij} . Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise positive pattern causality.




Figure 10: Comprehensive PC (negative) heatmap for the U.S.. Weights use causality-frequency percentages c_{ij} . Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise negative pattern causality.




Figure 11: Comprehensive PC (dark) heatmap for the U.S.. Weights use causality-frequency percentages c_{ij} . Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise dark pattern causality.




Figure 12: Comprehensive Granger causality heatmap for the U.S.. Weights use p value. Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes smaller p value.




Figure 13: Comprehensive cointegration/VECM heatmap for the U.S.. Weights use p value. Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes smaller p value.




Figure 14: Comprehensive *Lead-Lag* heatmap for the U.S.. Weights use correlation coefficient $|\rho|$. Macroeconomic aggregates occupy the top block (first 15 rows/columns), and AI-specific variables (including collaborations) occupy the last six rows/columns. Abbreviation mappings are listed in Appendix A. Darker shading denotes stronger pairwise correlation.

threshold filters to retain only statistically credible or substantively strong edges, and then visualizes the resulting networks to examine how central AI nodes are positioned and how their links with macroeconomic variables weaken, disappear, or persist.

4.3 Core Transmission Pathways

Having established who is structurally prominent (Section 4.1) and where interactions cluster (Section 4.2), we now ask which links endure once weaker or statistically insignificant ties are pruned. We do this by drawing thresholded, directed networks for each method and country (Figure 15–Figure 22). Thresholds follow the conventions introduced earlier and detailed in Appendix D: PC at 25/40/55/75%, Granger and cointegration at $p \in \{0.20, 0.10, 0.05, 0.01\}$, and lead-lag at $|\rho| \in \{0.20, 0.50, 0.80, 0.90\}$. These representations test the robustness of the AI-macro channels flagged by the rankings and heatmaps.

For China (Figure 15), a 25% cut-off admits many AI→macro edges, but most are modest. Tightening to above 40% leaves three AI nodes—*AI research publications*, *AI journal*

papers, and *U.S.–China publication collaborations*—that maintain medium, positive links to *GDP*, *gross savings*, and *gross national expenditure*. A distinctive relation appears at this band: *publication collaboration* $\rightarrow$ *education* exceeds 55%, a pattern not mirrored by other AI indicators. At 75%, AI nodes drop out entirely and the remaining strong links are confined to the macro triad *GDP*, *gross savings*, and *population*. The U.S. (Figure 19) shows the same broad shape at 25% and 75%, but differs at the intermediate band (55%): alongside the Chinese trio, *AI patent publications* form several medium-to-strong links to *GDP*, *GDP*, *PPP*, *CPI*, and the *GDP deflator*. Thus, the U.S. PC layer exhibits a wider AI footprint at mid-range strengths, with patenting playing a larger causal role than in China.

For China (Figure 16), AI nodes—especially the two collaboration indicators—remain central across all three p -value cut-offs. Edges both to and from key macro aggregates persist even at $p < 0.01$, indicating dense short-run predictive feedback between AI and the economy. In the U.S. (Figure 20), AI variables also rank prominently, but the graph is more segmented: macro nodes chiefly link to macro nodes and AI to AI, with fewer cross-cluster arrows. Hence U.S. AI indicators enter the predictive chain, yet their Granger influence on broad macro aggregates is narrower and more selective than in China.

In China (Figure 17), the three most central nodes at all significance tiers are AI-related (*AI journal papers*, *U.S.–China conference paper collaborations*, *U.S.–China publication collaborations*). Nearly all arrows, however, point *into* these AI nodes from macro variables, consistent with the VECM loading-based direction: AI indicators adjust to restore the long-run equilibrium rather than drive it. The U.S. network (Figure 21) is similarly segmented into AI and macro clusters with relatively few bridges; *AI patent publications* stand out as the most central node across cut-offs, yet most cointegrating links are again directed *toward* the patent node, placing it in a follower role despite its prominence.

For China (Figure 18), many AI–macro co-movements appear at $|\rho| > 0.20$ and > 0.50 , but only a few very tight pairs survive at $|\rho| > 0.80$: the two collaboration nodes as a pair, the *AI journal–AI publications* duo, and a macro block linking *GDP*, *gross savings*, and

population. The U.S. (Figure 22) is even more selective at $|\rho| > 0.80$, with surviving edges limited to (i) the collaboration pair, (ii) *exports–imports*, and (iii) the macro triad *GDP, GDP, PPP, expenditure*. At the intermediate threshold, the U.S. graph displays a visibly bipartite structure (macro on one side, AI on the other) bridged primarily by *publication collaboration* $\rightarrow$ *GDP*, echoing the separation seen in the Granger and cointegration panels.

Across methods, the pathway picture is consistent with the earlier rankings and heatmaps: nonlinear dependence (PC positive) is broad but moderate. Granger links are dense and short-run and cointegration suggests that long-run balance is restored mainly through adjustments in AI variables themselves. Further, the lead–lag co-movement is of medium intensity with a small set of high-synchrony pairs. Where the countries differ is not in *whether* AI matters, but *how*: the U.S. pathway is more selective and patent-driven; China’s is denser, more collaboration-intensive, and more tightly interwoven with the macro core.

4.4 Cross-Method Triangulation

We reduce single-method bias by requiring *convergence* across methods (PC, Granger causality, cointegration/VECM, and lead–lag), representations (degree-centrality rankings, heatmaps, and thresholded networks), and countries (China and the U.S.). This triangulation privileges structures that repeat across lenses and down-weights pairwise extremes; our claims are exploratory and policy-informing rather than fully identified structural causality.

Three regularities emerge across layers and displays:

Nonlinear layer (PC). AI–macro dependence is visible, but it does not dominate the system. In the rankings (Section 4.1), AI nodes typically sit in the middle of the table rather than at the very top. In the heatmaps (Section 4.2), the AI block is lighter in the positive PC panel, signaling fewer strong nonlinear matches. In the thresholded PC networks (Section 4.3), AI edges thin out quickly as the cut-off rises. Taken together,




Figure 15: Thresholded network derived from PC (positive) for China. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode causality frequency (percentage). The top 3 central nodes are highlighted by red color. Here the selected thresholds are: [25%,40%,55%,75%]




Figure 16: Thresholded network derived from Granger Causality for China. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode $1 - p$ value. The top 3 central nodes are highlighted by red color. Here the selected thresholds are: [0.80,0.90,0.95,0.99]




Figure 17: Thresholded network derived from Cointegration/VECM for China. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode $1 - p$ value. The top 3 central nodes are highlighted by red color. Here the selected thresholds are: [0.80,0.90,0.95,0.99]




Figure 18: Thresholded network derived from lead-lag correlation for China. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode correlation coefficient $|\rho|$. The top 3 central nodes are highlighted by red color. Here the selected thresholds are: $[0.20, 0.50, 0.80, 0.90]$




Figure 19: Thresholded network derived from PC (positive) for the U.S.. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode causality frequency (percentage). The top 3 central nodes are highlighted by red color. Here the selected thresholds are: [25%,40%,55%,75%]




Figure 20: Thresholded network derived from Granger Causality for the U.S.. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode $1 - p$ value. The top 3 central nodes are highlighted by red color. Here the selected thresholds are: [0.80,0.90,0.95,0.99]




Figure 21: Thresholded network derived from Cointegration/VECM for the U.S.. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode $1 - p$ value. The top 3 central nodes are highlighted by red color. Here the selected thresholds are: $[0.80, 0.90, 0.95, 0.99]$




Figure 22: Thresholded network derived from lead-lag correlation for the U.S.. A node is shown only if it participates in at least one edge that exceeds the method-specific threshold. Edge colors encode correlation coefficient $|\rho|$. The top 3 central nodes are highlighted by red color. Here the selected thresholds are: $[0.20, 0.50, 0.80, 0.90]$

these patterns indicate moderate, mostly positive, state-dependent associations—useful for context, but not evidence of pervasive nonlinear leadership.

Short-run layer (Granger). The predictive web is dense. AI outputs—publications and patents—repeatedly appear among the most central nodes in the rankings and remain connected in the networks even at $p < 0.01$. This is consistent with broad directional predictability both from AI to macro variables and from macro variables to AI. The cross-country contrast is stable: in China, AI and macro nodes are tightly interwoven; in the U.S., links cluster within domains (macro–macro and AI–AI), and cross-cluster arrows are fewer.

Equilibrium/co-movement layer (cointegration and lead–lag). VECM direction places most adjustment loads on AI variables. In both countries, cointegrating arrows predominantly enter AI nodes, implying a follower role for AI when long-run balance is restored. Lead–lag heatmaps and networks contract at higher $|\rho|$ thresholds to a small, stable subset of pairs—within-AI links and a few macro triads such as GDP–savings–population. Where cross-cluster leadership survives, it is largely macro-led.

Results are consistent across lenses. Adding U.S.–China collaboration variables increases the embeddedness and centrality of AI nodes in both countries, with a stronger and more system-wide effect in China and a more localized footprint in the U.S.. Because enlarging the variable set can raise overall connectivity mechanically, we read improvements through relative rank shifts and survival at stricter thresholds rather than raw magnitudes.

The rankings (Section 4.1), full-sample heatmaps (Section 4.2), and thresholded networks (Section 4.3) point to a common message: AI is macro-relevant in both countries, but *how* it matters differs. In the U.S., the footprint is patent-driven and selective. In China, it is denser, collaboration-intensive, and more tightly interwoven with core macro aggregates. This triangulated view supports risk-aware monitoring by highlighting structures that recur across methods while down-weighting single-test extremes.

5 Discussion

This section synthesizes and discusses the paper’s main findings highlighting the contributions and limitations. Four contributions stand out: mapping AI–macro linkages across methods; identifying collaboration as a transmission channel; comparing asymmetries between the U.S. and China; and advancing methodological transparency for risk-aware inference.

5.1 AI’s Role in Macroeconomic Performance

5.1.1 Mapping AI–Macro Causality Across Methods

Our first contribution is methodological: we integrate the newly developed PC approach with established tools—Granger causality, VECM-based cointegration, and lead–lag correlation—to build a multi-layered view of AI–macro linkages (Section 3–Section 4). PC fills a gap by detecting nonlinear, state-dependent relations and by classifying them into positive, negative, or dark forms—capabilities conventional linear methods do not provide.

Across countries and representations (centrality tables, heatmaps, thresholded networks), several regularities recur. AI R&D variables do enter nonlinear causal relations with macro aggregates, though effects are typically moderate and mostly positive. On the linear side, Granger tests identify clear short-run predictive relationships, whereas VECM-based cointegration shows that AI generally assumes a follower role in restoring long-run equilibria. Lead–lag correlations add timing information—which variables tend to move earlier—but the synchrony they reveal is not especially strong. Taken together, the layers provide a coherent account of short-run prediction, long-run adjustment, temporal ordering, and nonlinear classification.

This framework is anchored in the econometric tradition of predictive causality Granger (1969), long-run equilibrium Engle and Granger (1987) and Johansen (1988), and cautions about structural breaks in long samples (Bai & Perron, 2003). PC extends this lineage

via symbolic dynamics and attractor reconstruction (Stavroglou et al., 2019, 2020, 2021), enabling detection and sign-classification of nonlinear dependence that conventional tools cannot directly address.

5.1.2 Collaboration as a Macro Transmission Channel

Our second contribution is to demonstrate that international collaboration functions as a macroeconomic transmission channel. After establishing the no-collaboration baseline, we add two measures—U.S.–China publication collaborations and U.S.–China conference paper collaborations—and re-estimate weighted degree centralities and rankings across all methods. We ask: (i) do collaboration variables themselves attain high centrality; and (ii) does their inclusion raise, leave unchanged, or lower the relative prominence of existing AI variables?

Results show that collaboration amplifies the AI–macro pathway on both counts. Collaboration variables consistently appear among top-ranked nodes, and their inclusion tends to elevate the relative positions of other AI variables. The magnitude differs by country—more transformative in China, more selective in the U.S.—but the message is consistent: cross-border collaboration is not neutral; it alters the causal architecture through which AI interacts with macro aggregates. This connects micro-level collaboration premiums to macro propagation. Prior work links collaboration networks to faster knowledge diffusion (Breschi & Lissoni, 2001) and documents higher impact for U.S.–China AI teams relative to solo efforts (AlShebli et al., 2024). Our results extend these insights by showing that collaboration reshapes the system-wide causal map of AI–macro interactions.

5.1.3 Comparative Asymmetry

Our third contribution is comparative, i.e., AI’s macroeconomic role is asymmetric across the U.S. and China. In the U.S., AI variables—especially patents—are comparatively central in the baseline networks. In China, AI indicators are more peripheral and display a clearer follower role in cointegration. Once collaboration is included, China’s AI variables rise in the rankings and exhibit a stronger driving effect than in the U.S.

In short, AI matters in both economies but along different pathways. In the U.S., patents appear to be a key conduit by which AI activity translates into economic impact; in China, international collaboration amplifies AI’s macro relevance. This contrast aligns with evidence on general-purpose technology diffusion and adoption lags (Bresnahan & Trajtenberg, 1995; Comin & Hobijn, 2004) and with heterogeneous productivity impacts associated with AI (Acemoglu & Restrepo, 2020; Brynjolfsson et al., 2017). Policies should therefore be tailored: strengthening patent systems and commercialization channels may be more relevant for the U.S., whereas fostering international collaboration and absorptive capacity may be more important for China.

5.1.4 Methodological Transparency and Risk-Aware Inference

Our fourth contribution is transparency about what the evidence can—and cannot—support. We interpret results as exploratory and decision-relevant patterns rather than definitive structural claims. The guiding principle is *convergence*. We highlight only those patterns visible across multiple methods with consistent direction, or reinforced across representations—high centrality rankings, global associations in heatmaps, and links that survive strict thresholds (Section 4.4). Method-specific signals (e.g., links detected only by PC) are retained as hypotheses rather than policy guidance. This stance reflects different statistical logic (PC percentages versus p -values) and reduces the risk of over-interpreting any single estimator.

Preprocessing and estimation choices follow a consistent protocol (Section 3; unit-root diagnostics in Appendix C; PC sensitivity in Appendix E). Key constraints are explicit: limited annual observations and common upward trends. Guardrails include: PC preserves nonlinear, state-dependent information but lacks a native p -value; thresholds aid readability rather than formal significance. For Granger and cointegration we report conventional cut-offs but emphasize cross-method convergence over single-test results. Adding collaboration nodes enlarges the network mechanically, so we focus on *relative* rank shifts and survival at stricter thresholds. With about 40 annual observations per country (1980–2020), statistical power is limited and short-lived shocks are hard to detect.

Together, these practices align the study with risk-aware inference: distinguishing robust from suggestive links, disclosing assumptions and data constraints, and calibrating conclusions to the scope of available evidence. The objective is not to maximize the number of “significant” edges, but to improve decision quality under uncertainty by showing where evidence converges, where it is method-dependent, and where additional data would add the most value.

5.2 Practical Implications for Risk and Policy

The findings have several implications for monitoring and policy. At a broad level, AI activity and macroeconomic activity are linked, albeit not uniformly across countries. Incorporating AI R&D indicators into economic monitoring can provide a more comprehensive view of system dynamics and, in some cases, early signals of macro change. At the same time, potential nonlinear dependencies, while exploratory, deserve attention because they may capture dynamics that conventional models overlook.

AI indicators can be used in addition to traditional macro variables for monitoring. In some cases, patents and research publications serve as leading indicators of short-term changes; in others, AI variables adjust to shifts in fundamentals, indicating a dependent role. Neither role is definitive on its own; taking AI activity into account alongside macroeconomic conditions broadens the lens through which to interpret economic shifts.

Collaboration in policy is not neutral. Across contexts, international collaboration tends to increase AI’s systemic role, albeit in varying degrees. Encouraging cooperation rather than allowing monopolization or destructive competition promotes knowledge exchange, positive spillovers, and tangible economic benefits.

Asymmetries do count. The US case emphasizes patents as a link between AI and economic impact, whereas the Chinese case emphasizes the importance of international collaboration. Policies should be tailored to domestic structures and capabilities, rather than being imported wholesale.

6 Conclusion

This paper integrates Pattern Causality (PC) with Granger causality, VECM-based cointegration, and lead-lag correlation to analyze AI-macro linkages between the U.S. and China from 1980-2020. The multi-layer design differentiates between nonlinear, state-dependent associations (PC), short-run predictive relationships (Granger), long-run equilibrium adjustment (cointegration/VECM), and temporal ordering (lead-lag). The estimates are implemented in a forward-expanding fashion and summarized through comparable, weighted network views (Section 3). They are then read jointly via rankings, heatmaps, and thresholded graphs (Section 4).

Three findings are consistent across layers and representations. First, nonlinear AI-macro associations exist, but they are generally moderate and predominantly positive in PC, with fewer negative or dark patterns. Second, in the short run, AI outputs, particularly publications and patents, consistently demonstrate directional predictability in Granger tests. Third, in the long run, VECM loadings show that AI variables typically absorb the adjustment required to restore cointegrating balance, putting them in a follower role relative to macro fundamentals; lead-lag correlations support this view, revealing only medium-intensity synchrony at stringent cutoffs.

Moreover, increasing collaboration measures between the U.S. and China, such as joint publications and conference papers, strengthens AI nodes and elevates domestic AI indicators. The effect is stronger in China and more targeted in the U.S.. In Section 4.1-Section 4.3, we prioritize relative rank shifts and survival at stricter thresholds over raw edge counts, as adding nodes can increase connectivity. U.S. networks play a central role in AI patents, while China’s AI indicators are more peripheral initially but become significantly more connected with collaboration. In both cases, cointegration places AI on the adjustment margin rather than as the long-term driver, but the paths are different: more patent-driven in the U.S., more collaboration-driven in China. These conclusions are risk-aware, rather than definitive. PC contributes nonlinear classification but lacks a

native p -value; Granger and cointegration provide traditional significance metrics but rely on small annual samples. We prioritize *convergence* across methods and representations while minimizing single-estimator signals. With about forty annual observations per country, statistical power is limited, and structural identification is out of reach; the evidence points to directional predictability and equilibrium co-movement, not deep causal mechanisms.

Furthermore, incorporating AI R&D indicators alongside macro aggregates can improve risk surveillance by providing additional signals on short-run dynamics and long-run adjustment. Cross-border scientific ties can amplify how AI shocks spread throughout the macrosystem. Policy design should take into account national structure; strengthening patenting and commercialization channels may be more relevant in the United States, whereas fostering international collaboration and absorptive capacity may be more effective in China.

Future work can address these issues by: (i) using higher-frequency data (quarterly or monthly) to expand sample size and identify cyclical shocks; (ii) broadening indicators (e.g., compute capacity, investment flows, sectoral GDP, labor-market stratification, financial risk); (iii) employing scale-normalized centrality, null-network checks, and systematic before/after comparisons; and (iv) studying collaboration dynamics with event-style analyses around policy shifts or geopolitical shocks to identify mechanisms behind the amplifying role of cross-border ties.

References

- Acemoglu, D., Akcigit, U., & Kerr, W. R. (2016). Innovation network. *Proceedings of the National Academy of Sciences*, 113(41), 11483–11488. <https://doi.org/10.1073/pnas.1613559113>
- Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial Intelligence and Jobs: Evidence from Online Vacancies. *Journal of Labor Economics*, 40(S1), S293–S340. <https://doi.org/10.1086/718327>
- Acemoglu, D., & Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *The Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
- Aghion, P., & Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2), 323–351. <https://doi.org/10.2307/2951599>
- Ahmed, N., Wahed, M., & Thompson, N. C. (2023). The growing influence of industry in AI research. *Science*, 379(6635), 884–886. <https://doi.org/10.1126/science.ade2420>
- Ali, M. (2024). Advancing Industries with AI: The applications being transformational include: Healthcare, powering the Petroleum industry, Handling fraud, Cybersecurity, and Conversational AI. *International Journal of Multidisciplinary Sciences and Arts*, 3, 173–180. <https://doi.org/10.47709/ijmdsa.v3i4.5052>
- AlShebli, B., Memon, S. A., Evans, J. A., & Rahwan, T. (2024). China and the U.S. produce more impactful AI research when collaborating together. *Scientific Reports*, 14(1), 28576. <https://doi.org/10.1038/s41598-024-79863-5>
- Bai, J., & Perron, P. (2003). Computation and Analysis of Multiple Structural Change Models. *Journal of Applied Econometrics*, 18(1), 1–22. <https://doi.org/10.1002/jae.659>
- Bayoumi, T., Coe, D. T., & Helpman, E. (1999). R&D spillovers and global growth. *Journal of International Economics*, 47(2), 399–428. [https://doi.org/10.1016/S0022-1996\(98\)00018-X](https://doi.org/10.1016/S0022-1996(98)00018-X)

- Bier, V., Borgonovo, E., Cox, T., & Rudin, C. (2025). Artificial intelligence for risk analysis and the risks of artificial intelligence: Part 1. *Risk Analysis*, 45(4), 735–737. <https://doi.org/10.1111/risa.14327>
- Bonfiglioli, A., Crinò, R., Gancia, G., & Papadakis, I. (2025). Artificial intelligence and jobs: Evidence from US commuting zones. *Economic Policy*, 40(121), 145–194. <https://doi.org/10.1093/epolic/eiae059>
- Breschi, S., & Lissoni, F. (2001). Knowledge Spillovers and Local Innovation Systems: A Critical Survey. *Industrial and Corporate Change*, 10(4), 975–1005. <https://doi.org/10.1093/icc/10.4.975>
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies ‘Engines of growth’? *Journal of Econometrics*, 65(1), 83–108. [https://doi.org/10.1016/0304-4076\(94\)01598-T](https://doi.org/10.1016/0304-4076(94)01598-T)
- Brynjolfsson, E., Hui, X., & Liu, M. (2019). Does Machine Translation Affect International Trade? Evidence from a Large Digital Platform. *Management Science*, 65(12), 5449–5460. <https://doi.org/10.1287/mnsc.2019.3388>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at Work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Brynjolfsson, E., Rock, D., & Syverson, C. (2017). Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics. *NBER Working Paper Series*, 24001. <https://doi.org/10.3386/w24001>
- Campbell, J. Y., & Perron, P. (1991). Pitfalls and Opportunities: What Macroeconomists Should Know about Unit Roots. *NBER Macroeconomics Annual*, 6, 141–201. <https://doi.org/10.2307/3585053>
- Choi, S., Kang, H., Kim, N., & Kim, J. (2025). The dual edges of AI: Advancing knowledge while reducing diversity. *PNAS nexus*, 4(5), pgaf138–. <https://doi.org/10.1093/pnasnexus/pgaf138>
- Cockburn, I. M., Henderson, R., & Stern, S. (2018). The Impact of Artificial Intelligence on Innovation: An Exploratory Analysis. In *The Economics of Artificial Intelligence: An Agenda* (pp. 115–146). University of Chicago Press.

- Coe, D. T., Helpman, E., & Hoffmaister, A. W. (2009). International R&D spillovers and institutions. *European Economic Review*, 53(7), 723–741. <https://doi.org/10.1016/j.euroecorev.2009.02.005>
- Comin, D., & Hobijn, B. (2004). Cross-country technology adoption: Making the theories face the facts. *Journal of Monetary Economics*, 51(1), 39–83. <https://doi.org/10.1016/j.jmoneco.2003.07.003>
- Comin, D., & Hobijn, B. (2010). An Exploration of Technology Diffusion. *American Economic Review*, 100(5), 2031–2059. <https://doi.org/10.1257/aer.100.5.2031>
- De Silva, T., & Thesmar, D. (2024). Noise in Expectations: Evidence from Analyst Forecasts (S. Giglio, Ed.). *The Review of Financial Studies*, 37(5), 1494–1537. <https://doi.org/10.1093/rfs/hhad091>
- DeJong, D. N., Nankervis, J. C., Savin, N. E., & Whiteman, C. H. (1992). Integration Versus Trend Stationary in Time Series. *Econometrica*, 60(2), 423–433. <https://doi.org/10.2307/2951602>
- Dickey, D. A., & Fuller, W. A. (1979). Distribution of the Estimators for Autoregressive Time Series With a Unit Root. *Journal of the American Statistical Association*, 74(366), 427–431. <https://doi.org/10.2307/2286348>
- Duede, E., Dolan, W., Bauer, A., Foster, I., & Lakhani, K. (2024). Oil & Water? Diffusion of AI Within and Across Scientific Fields. <https://doi.org/10.48550/arXiv.2405.15828>
- Engle, R. F., & Granger, C. W. J. (1987). Co-Integration and Error Correction: Representation, Estimation, and Testing. *Econometrica*, 55(2), 251–276. <https://doi.org/10.2307/1913236>
- Gaessler, F., & Piezunka, H. (2023). Training with AI: Evidence from chess computers. *Strategic Management Journal*, 44(11), 2724–2750. <https://doi.org/10.1002/smj.3512>
- Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 37(3), 424–438. <https://doi.org/10.2307/1912791>

- Griliches, Z. (1979). Issues in Assessing the Contribution of Research and Development to Productivity Growth. *Bell Journal of Economics*, 10(1), 92–116. <https://doi.org/10.2307/3003321>
- Han, E., Yin, D., & Zhang, H. (2023). Bots with Feelings: Should AI Agents Express Positive Emotion in Customer Service? *Information Systems Research*, 34(3), 1296–1311. <https://doi.org/10.1287/isre.2022.1179>
- Humlum, A., & Vestergaard, E. (2025). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1), e2414972121. <https://doi.org/10.1073/pnas.2414972121>
- Jia, R., Roberts, M. E., Wang, Y., & Yang, E. (2024). The impact of US–China tensions on US science: Evidence from the NIH investigations. *Proceedings of the National Academy of Sciences*, 121(19), e2301436121. <https://doi.org/10.1073/pnas.2301436121>
- Johansen, S. (1991). Estimation and Hypothesis Testing of Cointegration Vectors in Gaussian Vector Autoregressive Models. *Econometrica*, 59(6), 1551–1580. <https://doi.org/10.2307/2938278>
- Johansen, S. (1988). Statistical analysis of cointegration vectors. *Journal of Economic Dynamics and Control*, 12(2), 231–254. [https://doi.org/10.1016/0165-1889\(88\)90041-3](https://doi.org/10.1016/0165-1889(88)90041-3)
- Jones, C. I. (1995). R & D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4), 759–784. <https://doi.org/10.1086/262002>
- Kitajima, K., & Okamura, K. (2025). The altering landscape of US–China science collaboration: From convergence to divergence. *Humanities and Social Sciences Communications*, 12(1), 302. <https://doi.org/10.1057/s41599-025-04550-3>
- Klasa, K., Trump, B. D., Linkov, I., & Lambert, J. H. (2020). Identifying new partnerships for innovation: Governance and policy challenges. *IEEE Engineering Management Review*, 48(2), 26–38. <https://doi.org/10.1109/EMR.2020.2989263>

- Li, R., Ren, F., & Wang, Q. (2024). China–US scientific collaboration on sustainable development amidst geopolitical tensions. *Humanities and Social Sciences Communications*, 11(1), 1448. <https://doi.org/10.1057/s41599-024-03948-9>
- Linkov, I., Trump, B. D., Keisler, J., Cegan, J., Foran, C., Collier, Z., Lambert, J. H., & Kuklja, M. M. (2020). Signals and metrics identifying partnerships for innovation. *IEEE Engineering Management Review*, 48(2), 39–46. <https://doi.org/10.1109/EMR.2020.2990127>
- Lo, A. W., & MacKinlay, A. C. (1990). When are Contrarian Profits Due to Stock Market Overreaction? *The Review of Financial Studies*, 3(2), 175–205.
- Luo, X., Tong, S., Fang, Z., & Qu, Z. (2019). Frontiers: Machines vs. Humans: The Impact of Artificial Intelligence Chatbot Disclosure on Customer Purchases. *Marketing Science*, mksc.2019.1192. <https://doi.org/10.1287/mksc.2019.1192>
- Miaillhe, N. (2018). The geopolitics of artificial intelligence: The return of empires? *Politique Étrangère*, (3), 105–117. <https://doi.org/10.3917/pe.183.0105>
- Minniti, A., Prettnner, K., & Venturini, F. (2025). AI innovation and the labor share in European regions. *European Economic Review*, 177, 105043. <https://doi.org/10.1016/j.eurocorev.2025.105043>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- Phillips, P. C. B., & Perron, P. (1988). Testing for a Unit Root in Time Series Regression. *Biometrika*, 75(2), 335–346. <https://doi.org/10.2307/2336182>
- Renn, O., & Lucas, K. (2022). Systemic risk: The threat to societal diversity and coherence. *Risk Analysis*, 42(9), 1921–1934. <https://doi.org/10.1111/risa.13654>
- Romer, P. M. (1990). Endogenous Technological Change. *The Journal of Political Economy*, 98(5), S71–S102. <https://doi.org/10.1086/261725>
- Schwert, G. W. (2002). Tests for Unit Roots: A Monte Carlo Investigation. *Journal of Business & Economic Statistics*, 20(1), 5–17. <https://doi.org/10.2307/1391432>

- Shi, J., Bendig, D., Vollmar, H. C., & Rasche, P. (2023). Mapping the Bibliometrics Landscape of AI in Medicine: Methodological Study. *Journal of Medical Internet Research*, 25, e45815. <https://doi.org/10.2196/45815>
- Stavroglou, S. K., Ayyub, B. M., Kallinterakis, V., Pantelous, A. A., & Stanley, H. E. (2021). A Novel Causal Risk-Based Decision-Making Methodology: The Case of Coronavirus. *Risk Analysis*, 41(5), 814–830. <https://doi.org/10.1111/risa.13678>
- Stavroglou, S. K., Pantelous, A. A., Stanley, H. E., & Zuev, K. M. (2019). Hidden interactions in financial markets. *Proceedings of the National Academy of Sciences*, 116(22), 10646–10651. <https://doi.org/10.1073/pnas.1819449116>
- Stavroglou, S. K., Pantelous, A. A., Stanley, H. E., & Zuev, K. M. (2020). Unveiling causal interactions in complex systems. *Proceedings of the National Academy of Sciences*, 117(14), 7599–7605. <https://doi.org/10.1073/pnas.1918269117>
- Sun, Y., Li, L., Yu, Z., Yu, H., & Wang, H. (2025). Exploring AI models and applications within a system framework. *Systems Research and Behavioral Science*, 42(4), 1163–1180. <https://doi.org/10.1002/sres.3036>
- Uzzi, B., Mukherjee, S., Stringer, M., & Jones, B. (2013). Atypical Combinations and Scientific Impact. *Science*, 342(6157), 468–472. <https://doi.org/10.1126/science.1240474>
- Yin-Wong Cheung & Lai, K. S. (1993). Finite-Sample Sizes of Johansen’s Likelihood Ratio Tests for Cointegration. *Oxford Bulletin of Economics & Statistics*, 55(3), 313–328. <https://doi.org/10.1111/j.1468-0084.1993.mp55003003.x>
- Zhuang, J., Wang, C., & Lambert, J. H. (2022). Machine learning and other information analyses for risk in social networks. *Risk Analysis*, 42(8). <https://doi.org/10.1111/risa.13945>

Appendix A Variable Abbreviations

This appendix lists the abbreviations used throughout the paper and their corresponding full variable names. The mapping standardizes nomenclature across macroeconomic and AI series for China and the U.S. and is consistent with the labels used in figures and network layers. The entries are nominal only; units, construction notes, and data sources are documented in the main tables. Providing this reference reduces ambiguity and facilitates replication.

Table A.1: Abbreviation and corresponding full name

Abbreviation	Full name
Exports	Exports of goods and services (BoP)
Imports	Imports of goods and services (BoP)
Inflation	Inflation, GDP deflator (annual)
Unemployment	Unemployment, total (national estimate)
RIR	Real interest rate
Expenditure	Gross national expenditure
Investment	Foreign direct investment, net (BoP)
GDP	Gross domestic product
GDP_PPP	GDP, PPP
Deflator	GDP deflator (base year varies by country)
Population	Population, total
BCA	Current account balance (BoP)
CPI	Consumer price index (2010 = 100)
Reserve	Total reserves (includes gold)
Education	School enrollment, tertiary (gross)
Savings	Gross savings
Publication_Collaborations	U.S.–China publication collaborations
Conference_Collaborations	U.S.–China conference collaborations
AI_Conferences	Number of AI conference papers
AI_Journals	Number of AI journal papers
AI_Patents	Number of AI patent publications
AI_Publications	Number of AI research publications

Appendix B Data Processing Steps

Specifically, our annual macro-AI panel was first indexed by calendar year, after which non-informative placeholders were recast as missing values and all string-encoded numbers were purged of thousands separators before coercion to floating-point format, thereby imposing numerical coherence across the entire matrix. Temporal discontinuities were resolved through linear interpolation along each series and any leading gaps were back-filled, producing uninterrupted trajectories for every variable. To secure semantic uniformity, verbose indicator labels were deterministically mapped to concise aliases, ensuring a one-to-one correspondence across all downstream routines. Each cleaned series was then standardized via the transformation $z_t = (y_t - \mu)/\sigma$ with $\mu = 0$ and $\sigma = 1$, yielding a scale-invariant feature space suited to multivariate inference. Finally, for procedures predicated on covariance stationarity—specifically the Granger-causality and lead-lag correlation modules—the preprocessed series were further rendered stationary by first differencing, $\Delta y_t = y_t - y_{t-1}$, thereby satisfying the $I(0)$ assumption without compromising cross-sectional comparability.

Table B.1: The descriptive statistics of macro and AI variables.

Variable	Mean	Max	Min	Std. Dev.	Skewness	Kurtosis	Jarque–Bera
<i>China</i>							
Exports (USD billion)	886.824	2,740.000	23.186	1,001.715	0.749	1.839	5.582
Imports (USD billion)	782.341	2,560.000	18.900	902.518	0.842	2.042	5.846
Inflation (%)	4.820	20.617	-1.263	4.672	1.351	5.260	17.423*
Unemployment (%)	3.458	5.610	1.800	0.994	0.009	2.159	1.280
Expenditure (USD billion)	3,686.561	14,500.000	190.000	4,623.625	1.235	3.075	9.684*
Investment (USD billion)	-58.639	41.675	-232.000	64.950	-1.150	3.791	8.426
GDP (USD billion)	3,784.049	14,700.000	191.000	4,706.361	1.206	2.993	9.252*
GDP_PPP (int\$ billion)	7,327.488	25,200.000	303.000	7,708.192	1.007	2.690	6.723
Deflator (index)	54.198	100.000	15.648	28.192	0.126	1.708	2.912
Population (million people)	1,233.654	1,411.100	981.235	130.332	-0.459	2.007	3.109
BCA (USD billion)	94.859	421.000	-11.609	117.392	1.093	3.197	7.169
CPI (index)	82.940	128.109	26.048	29.521	-0.443	2.252	2.009
Reserve (USD billion)	1,157.706	3,900.000	10.091	1,451.295	0.811	1.843	6.483
Education (%)	18.182	62.237	1.279	18.535	1.075	2.914	7.027
Savings (USD billion)	1,725.531	6,450.000	65.779	2,150.309	1.090	2.626	7.906
AI_Conferences (count)	6,467.918	27,678.567	0.500	8,382.611	1.039	2.724	7.094
AI_Journals (count)	12,474.996	54,490.980	21.333	15,635.901	0.908	2.536	5.746
AI_Patents (count)	122.351	713.625	0.000	208.001	1.768	5.008	24.338*
AI_Publications (count)	20,986.406	72,384.902	25.333	25,561.100	0.715	1.767	5.824
Publication_Collaborations (count)	1,763.275	11,532.000	3.000	2,825.604	1.956	6.507	37.986*
Conference_Collaborations (count)	641.941	3,014.000	1.000	854.135	1.353	4.056	10.228*
<i>United States</i>							
Exports (USD billion)	1,197.220	2,550.000	266.000	777.129	0.437	1.729	3.943
Imports (USD billion)	1,520.780	3,120.000	291.000	985.009	0.272	1.474	4.253
Inflation (%)	2.662	9.462	0.617	1.833	2.501	10.074	103.158*
Unemployment (%)	6.248	9.700	3.669	1.673	0.576	2.499	2.690
Expenditure (USD billion)	11,192.927	22,100.000	2,870.000	5,921.228	0.300	1.818	2.974
Investment (USD billion)	2.783	177.000	-345.000	109.348	-0.893	4.691	8.163
GDP (USD billion)	10,861.220	21,500.000	2,860.000	5,703.682	0.331	1.880	2.877
GDP_PPP (int\$ billion)	10,861.220	21,500.000	2,860.000	5,703.682	0.331	1.880	2.877
Deflator (index)	75.983	108.267	40.457	19.409	-0.027	1.871	2.219
Population (million people)	280.087	331.527	227.225	32.926	-0.065	1.650	3.064
BCA (USD billion)	-304.290	4.810	-817.000	231.508	-0.486	2.294	2.498
CPI (index)	79.545	118.691	37.792	24.150	-0.024	1.763	2.606
Reserve (USD billion)	255.024	628.000	105.000	153.788	1.023	2.577	7.087
Education (%)	73.844	88.889	54.665	11.744	-0.064	1.581	2.253
Savings (USD billion)	2,019.049	4,190.000	634.000	1,036.303	0.543	2.250	2.974
AI_Conferences (count)	7,607.249	19,349.874	362.333	5,746.177	0.432	1.831	3.536
AI_Journals (count)	14,797.164	38,067.471	3,059.701	9,500.843	0.639	2.257	3.678
AI_Patents (count)	2,725.803	6,828.710	238.500	2,051.381	0.305	1.662	3.579
AI_Publications (count)	27,901.884	71,311.482	4,009.284	19,479.207	0.585	2.095	3.684
Publication_Collaboration (count)	1,763.275	11,532.000	3.000	2,825.604	1.956	6.507	37.986*
Conference_Collaboration (count)	641.941	3,014.000	1.000	854.135	1.353	4.056	10.228*

Note: *Denotes the rejection of the null hypotheses of the Jarque–Bera normality test at the 1% significance level.

Appendix C Unit-Root Tests

For the cointegration pre-test, we applied the Augmented Dickey–Fuller (ADF), Phillips–Perron (PP), and Kwiatkowski–Phillips–Schmidt–Shin (KPSS) tests to each series in levels and in first differences before estimating the VECM. With roughly forty time points and highly persistent macro series, conventional unit-root tests can exhibit low power and size distortions (Campbell & Perron, 1991; DeJong et al., 1992; Schwert, 2002). To mitigate these risks, we adopt a conservative yet asymmetric rule that matches our code and table: a series is labeled $I(0)$ at level only when *both* ADF and PP at level reject the unit-root null; otherwise it is reported as *Not $I(0)$* . For first differences, we use a tiered union across $\alpha \in \{0.10, 0.05, 0.01\}$: a series is labeled $I(1)$ if, at any α , either the ADF/PP criterion holds on Δ (i.e., $\min\{p_{\text{ADF},\Delta}, p_{\text{PP},\Delta}\} < \alpha$) *or* KPSS on Δ fails to reject stationarity ($p_{\text{KPSS},\Delta} \geq \alpha$). Stars in the $I(1)$ column indicate the strongest level satisfied (* for 10%, ** for 5%, *** for 1%). All p -values are rounded to three decimals; exact 1 is displayed as ≈ 1 , and threshold values (e.g., < 0.01 , > 0.10) are shown as thresholds.

In our processing, scripts used for Table C.1, ADF and KPSS include intercept and trend; ADF uses AIC-based autolag without an explicit cap; KPSS uses automatic bandwidth selection; PP follows default settings.

Under this rule, *all* variables in both the China and U.S. blocks are classified as $I(1)$: each series is nonstationary in levels (conservatively not $I(0)$) and becomes stationary after first differencing. This outcome satisfies the standard preconditions for estimating a VECM while still acknowledging finite-sample limitations. Accordingly, we proceed with a two-layer design: VAR-based Granger tests on first differences to characterize short-run changes-to-changes predictability, and VECMs on level data to recover long-run equilibrium adjustment through the error-correction mechanism (Campbell & Perron, 1991; Yin-Wong Cheung & Lai, 1993).

Table C.1: Unit root p-values (level and first difference) reported at both levels. $I(0)$ at level uses a conservative rule (ADF and PP must both reject). $I(1)$ is a union across $\alpha \in \{0.10, 0.05, 0.01\}$; stars reflect the strongest satisfied level.

Variable	China								United States							
	p-values						Stationary		p-values						Stationary	
	ADF _L	PP _L	KPSS _L	ADF Δ	PP Δ	KPSS Δ	Not I(0)	I(1)	ADF _L	PP _L	KPSS _L	ADF Δ	PP Δ	KPSS Δ	Not I(0)	I(1)
Exports	0.857	0.994	<0.01	0.000	0.000	>0.10	Not I(0)	I(1)***	0.619	0.906	0.014	0.955	0.001	0.065	Not I(0)	I(1)***
Imports	0.001	0.990	<0.01	0.546	0.000	>0.10	Not I(0)	I(1)***	0.262	0.923	0.020	0.000	0.000	0.069	Not I(0)	I(1)***
Inflation	0.036	0.113	>0.10	0.007	0.000	0.095	Not I(0)	I(1)***	0.568	0.000	0.036	0.008	0.000	>0.10	Not I(0)	I(1)***
Unemployment	0.000	0.521	>0.10	0.000	0.000	0.064	Not I(0)	I(1)***	0.068	0.440	>0.10	0.010	0.656	>0.10	Not I(0)	I(1)***
Expenditure	≈ 1	≈ 1	<0.01	0.994	0.129	0.062	Not I(0)	I(1)***	0.502	≈ 1	0.013	0.003	0.004	>0.10	Not I(0)	I(1)***
Investment	0.654	0.383	>0.10	0.020	0.000	>0.10	Not I(0)	I(1)***	0.341	0.001	>0.10	0.055	0.000	<0.01	Not I(0)	I(1)***
GDP	≈ 1	≈ 1	<0.01	0.997	0.084	0.062	Not I(0)	I(1)***	0.267	≈ 1	<0.01	0.888	0.042	>0.10	Not I(0)	I(1)***
GDP_PPP	≈ 1	≈ 1	<0.01	0.887	0.635	0.081	Not I(0)	I(1)***	0.267	≈ 1	<0.01	0.888	0.042	>0.10	Not I(0)	I(1)***
Deflator	0.033	0.988	>0.10	0.077	0.038	>0.10	Not I(0)	I(1)***	0.414	0.511	>0.10	0.434	0.000	>0.10	Not I(0)	I(1)***
Population	0.984	0.000	<0.01	0.497	0.960	>0.10	Not I(0)	I(1)***	0.993	0.900	0.049	0.316	0.314	0.024	Not I(0)	I(1)***
BCA	0.141	0.544	>0.10	0.095	0.000	>0.10	Not I(0)	I(1)***	0.391	0.682	>0.10	0.001	0.000	>0.10	Not I(0)	I(1)***
CPI	0.067	0.956	>0.10	0.023	0.140	>0.10	Not I(0)	I(1)***	0.076	0.537	>0.10	0.000	0.000	>0.10	Not I(0)	I(1)***
Reserve	≈ 1	0.946	0.022	0.816	0.082	>0.10	Not I(0)	I(1)***	0.595	0.997	0.024	0.053	0.006	>0.10	Not I(0)	I(1)***
Education	≈ 1	≈ 1	<0.01	0.497	0.017	>0.10	Not I(0)	I(1)***	0.257	0.215	>0.10	0.040	0.095	>0.10	Not I(0)	I(1)***
Savings	≈ 1	0.999	<0.01	0.996	0.015	>0.10	Not I(0)	I(1)***	0.088	0.992	0.038	0.001	0.008	>0.10	Not I(0)	I(1)***
AI_Conferences	0.004	0.498	0.069	0.546	0.236	>0.10	Not I(0)	I(1)***	0.482	0.530	0.077	0.997	0.993	0.083	Not I(0)	I(1)***
AI_Journals	0.982	0.998	0.032	0.548	0.863	>0.10	Not I(0)	I(1)***	≈ 1	≈ 1	<0.01	0.547	0.820	>0.10	Not I(0)	I(1)***
AI_Patents	≈ 1	0.999	<0.01	0.477	0.000	0.048	Not I(0)	I(1)***	0.362	0.955	>0.10	0.757	0.003	>0.10	Not I(0)	I(1)***
AI_Publications	0.078	0.981	0.034	0.719	0.063	>0.10	Not I(0)	I(1)***	0.957	0.999	<0.01	0.000	0.000	>0.10	Not I(0)	I(1)***
Publication_Collaboration	0.000	0.684	0.012	0.983	0.000	>0.10	Not I(0)	I(1)***	0.000	0.684	0.012	0.983	0.000	>0.10	Not I(0)	I(1)***
Conference_Collaboration	0.000	0.402	<0.01	0.006	≈ 1	>0.10	Not I(0)	I(1)***	0.000	0.402	<0.01	0.006	≈ 1	>0.10	Not I(0)	I(1)***

Notes. Level ($I(0)$): A conservative criterion is applied: a series is labeled $I(0)$ only when both ADF and PP at level reject the unit-root null; otherwise it is reported as *Not I(0)*. KPSS at level is reported but does not alter the $I(0)$ label. **First difference ($I(1)$):** Labels follow a tiered union rule across $\alpha \in \{0.10, 0.05, 0.01\}$: a series is labeled $I(1)$ if, at any α , either the ADF/PP criterion holds or KPSS on first differences fails to reject stationarity. Stars in the $I(1)$ column indicate the strongest level satisfied: * for 10%, ** for 5%, *** for 1%. Very large p -values are shown as ≈ 1 .

Appendix D Threshold Selection

Our core processing logic is to more intuitively see those stronger causal or statistical connections while trying not to filter out most of the important and reasonable information. Although the networks in China and the U.S. are different, we intend to find consistent influences (links with arrows) in both China and the U.S. in the network diagram to better summarize key variables and core influences. Across all four network layers—positive-pattern causality, Granger causality, VECM cointegration, and lead-lag correlation—we calibrate the threshold in each panel to expose the strongest statistical links without eliminating structurally important ones. This “progressive-filter” design makes it possible to identify arrows that survive every country–method combination and therefore constitute robust channels of influence.

To make the thresholding conventions operational and interpretable, we make the following clarifications. In short samples with highly persistent macro series, liberal entry bands (e.g., $p = 0.20$, $|\rho| = 0.20$, PC=25%) reduce the risk of Type II errors and prevent prematurely discarding edges that may be structurally meaningful; progressively stricter bands ($p \in \{0.10, 0.05, 0.01\}$, $|\rho| \in \{0.50, 0.80, 0.90\}$, PC=40/55/75%) then trade density for specificity, revealing the network “backbone” as the cutoff tightens. Thresholds are applied to *directed* weights so that arrow orientation (source→target) and its causal interpretation are preserved; edges exactly at the cutoff are retained to avoid artefacts from rounding and numerical tolerance. We define a link as *robust* if it persists under stricter thresholds *and* appears in both countries for the same method; links that also recur across methods (PC, Granger, VECM cointegration, lead–lag) are treated as system-level channels. Given multiple testing and finite-sample concerns, our reading prioritizes consistency across thresholds, countries, and methods rather than any single cutoff. These additions explain *why* the grid is calibrated as above (power vs. specificity), *how* direction and tie-handling are treated, and *what* constitutes robustness, allowing readers to interpret Figure 15 to Figure 22 under a single, transparent rule set.

Appendix E Sensitivity Check for Pattern Causality

We assess the robustness of Pattern Causality (PC) to its hyperparameters by recomputing the directed weight matrix $W = c/100$ at the final time slice under a grid $E \in \{3, 4, 5\}$, $\tau \in \{1, 2\}$, $k \in \{10, 15\}$ (12 settings in total). For each setting we (i) form W from the positive PC shares ($[0,1]^\$$), with self-loops removed; (ii) extract the top- K directed links by weight; and (iii) compute node-level weighted out- and in- degree. To prevent artificial inflation of overlap, K is capped at $\min\{30, 0.2n(n-1)\}$ where n is the number of variables. Between any two settings a, b we compute: (i) Jaccard similarity of the top- K edge sets, $J(a, b) = \frac{|E_K^{(a)} \cap E_K^{(b)}|}{|E_K^{(a)} \cup E_K^{(b)}|}$, and (ii) Kendall’s τ rank correlations for out- and in-degree across the common node set. We summarise the pairwise distribution by the median and interquartile range (IQR).

Across the 12 settings, the median Jaccard similarity of the top- K link sets is **0.50** (IQR **0.23–0.88**), indicating a stable backbone with moderate peripheral churn. Node-level rankings are robust: the median Kendall’s τ is **0.629** for in-degree (IQR **0.522–0.855**) and **0.675** for out-degree (IQR **0.573–0.854**). These levels imply that roughly half of the leading links recur under hyperparameter perturbations, while the identity of core senders/receivers remains largely preserved. The spread in the Jaccard IQR arises mainly from extreme combinations (e.g., larger E with tighter effective samples or larger k when the valid neighbour count is bounded), which perturb local neighbourhoods and therefore the marginal inclusion of weaker links. Repeating the summaries for $K \in \{20, 30, 40\}$ yields the expected pattern: smaller K emphasises the hard core and raises Jaccard; larger K admits more peripheral churn without altering the backbone identified by the ≥ 0.75 frequency criterion.

Table E.1: Pairwise robustness of PC networks to hyperparameters. Across all $\binom{12}{2} = 66$ pairs of (E, τ, k) settings, this table reports (i) the Jaccard similarity between top- K edge sets and (ii) Kendall’s τ for the corresponding in-/out-degree rankings. Entries are the median and interquartile range (IQR, 25–75%); larger values indicate greater stability of links and node rankings across hyperparameters.

	Median	IQR (25%–75%)	
Jaccard@Top K	0.50	0.23	0.88
Kendall τ (In)	0.629	0.522	0.855
Kendall τ (Out)	0.675	0.573	0.854

After that, We evaluate whether the directed links reported by Pattern Causality (PC) at the final time slice exceed what is expected under a null with no cross-series predictability, while preserving within-series dependence. We implement a *moving block bootstrap* (MBB) at the series level. For each ordered pair $(x \rightarrow y)$, the null $H_0(x \rightarrow y)$ states that the empirical PC weight for $x \rightarrow y$ is indistinguishable from a world with no cross-series coupling.

For replication $b = 1, \dots, B$ and block length L : (i) resample each series independently using circular moving blocks of length L to obtain $\mathbf{X}^{*(b)}$; (ii) recompute PC on $\mathbf{X}^{*(b)}$ with the same (E, τ, k) and mode as in the main analysis; (iii) extract the final-slice weight matrix $W^{*(b)}$. A one-sided right-tail p -value is

$$p(x \rightarrow y) = \frac{1 + \sum_{b=1}^B \mathbf{1}\{W_{y,x}^{*(b)} \geq W_{y,x}\}}{1 + B}.$$

We apply Benjamini–Hochberg at $\alpha = 0.05$ across off-diagonal edges to obtain FDR-controlled significance. We use validation with $B = 60$ replications and scan block lengths $L \in \{6, 8\}$ (extended grids are supported). All artefacts are saved as edge-wise p -value matrices and BH–FDR masks per L .

Table E.2: Significant directed edges under BH-FDR (5%) by block length L . Density is the fraction of off-diagonal edges declared significant.

Block length	N variables	# significant	# possible	Density
$L = 6$	21	194	420	0.462
$L = 8$	21	158	420	0.376

Table E.3: Overlap of significant-edge sets between $L = 6$ and $L = 8$: intersection, union, Jaccard index, and edges unique to each L .

intersection	union	Jaccard	only_L6	only_L8
139	213	0.653	55	19

Table E.4: Pairwise Jaccard similarity of significant-edge sets across block lengths L .

L_1	L_2	Jaccard_sigset
6	8	0.653

High cross- L overlap (Tables E.3–E.4) indicates that core findings are not driven by the specific block length. Edges appearing only under a single L should be treated as L -sensitive and discussed as robustness-only, while the main text highlights links that persist across L .