AI Compute under Export Controls: China, Trade Reorganization, and Network Concentration

Diego Vallarino*

Abstract

This paper studies how US export controls reshaped the global supply network for artificial-intelligence (AI) compute. Using bilateral trade flows in compute-relevant goods across 48 countries from 2017 to 2024, I construct a weighted directed network and examine its response to four successive rounds of BIS restrictions on advanced semiconductors, fabrication equipment, AI servers, and related infrastructure. The evidence shows *hierarchical concentration* rather than decoupling or symmetric substitution: between 2021 and 2024, the China-centered subgraph lost 26% of its active bilateral relationships while total compute-relevant trade remained within 4% of its pre-shock peak. Surviving links absorbed the displaced volume, with mean trade per active edge rising 31% and the within-bloc Herfindahl-Hirschman index rising from 0.293 to 0.313.

By contrast, the US-aligned subgraph preserved stable edge density while its spectral radius increased by 31.6%, indicating stronger internal coupling without extensive-margin contraction. A triple-difference design exploiting close-placebo HS codes (monitors, transformers) and event-study dynamics yields post-2022 contractions of 7.6 percentage points in 2022 and 11.9 percentage points in 2024 in the probability of an active compute-relevant relationship among China-centered pairs, escalating in step with successive rounds of controls.

Findings are robust to Hong Kong treatment, alternative AI-relevance weights, and sub-bloc decomposition. The results imply that export controls operated less by eliminating China-centered AI-compute trade than by narrowing the network through which it flows, with implications for the economics of AI and the measurement of supply-chain fragility under national-security-oriented industrial policy.

Keywords: artificial intelligence; compute supply chain; export controls; network economics; spectral graph theory; supply-chain concentration.

JEL classification: F13, F14, L63, O33, O38, C45.

*Counselor to the Board of Directors, Inter-American Development Bank and IDB Invest, 1300 New York Avenue NW, Washington, D.C., 20577. Email: diegoval@iadb.org. The views expressed in this paper are those of the author and do not necessarily reflect those of the Inter-American Development Bank, its Board of Directors, or the countries they represent. I thank seminar participants for valuable feedback. All errors are my own.

1 Introduction

Did US export controls constrain China's access to AI compute, or did they instead reorganize the network through which compute-relevant inputs reach China-centered supply chains? This question is central to the economics of artificial intelligence because frontier AI is not produced by algorithms alone. It depends on a geographically concentrated hardware stack—advanced semiconductors, semiconductor manufacturing equipment, servers, accelerators, and high-bandwidth networking infrastructure—whose trade structure increasingly shapes the distribution of AI capabilities across countries.

Compute has therefore emerged as a third factor of artificial- intelligence (AI) production alongside data and algorithms. Unlike data and software, however, compute is capital-intensive, geographically clustered, and politically contestable. The economics of AI literature has documented compute's role in capability scaling (Sevilla et al., 2022; Cottier et al., 2024), its implications for productivity and task reallocation (Brynjolfsson et al., 2021; Acemoglu and Restrepo, 2022; Aghion et al., 2019), and its diffusion across firms (Agrawal et al., 2022; Babina et al., 2024). Yet less quantitative attention has been paid to the supply-side architecture that delivers compute to AI developers: a tightly coupled global network of fabrication plants, equipment vendors, design houses, assembly platforms, and regulatory regimes. The structure of this network determines not only who can train frontier models and at what cost, but also how national-security-oriented industrial policy reshapes the geography of AI production.

On October 7, 2022, the US Bureau of Industry and Security (BIS) issued a sweeping set of export controls targeting advanced semiconductors and the equipment used to manufacture them, with explicit application to the People's Republic of China (Allen, 2022). Successive tightening rounds in October 2023, March 2024, and December 2024 extended the perimeter to additional fabrication nodes, a broader set of entities, and a wider geography of allied jurisdictions. The stated objective was to constrain China's ability to develop frontier AI systems for military and dual-use applications (Khan et al., 2021; Thadani and Allen, 2023). The actual effect on the global compute network—its topology, its fragility, and its capacity to deliver AI inputs to non-Chinese developers as well—is an empirical question that has been addressed mostly through case studies and qualitative supply-chain analyses (U.S.-China Economic and Security Review Commission, 2022; Mazzucato and Rodrik, 2026), with limited quantitative network evidence. We provide such evidence.

We construct a weighted directed network of bilateral trade flows in compute-relevant goods—advanced logic and memory (HS 8542), semiconductor manufacturing equipment (HS 8486), automatic data processing units including AI accelerators (HS 8471), and high-bandwidth networking (HS 8517)—across 48 countries observed annually from 2017 through 2024 from the

United Nations Comtrade database. We complement these data with the timeline of BIS Entity List additions and Export Administration Regulations amendments through December 2024.

Onto this network we apply three complementary tools. The *spectral radius*—the principal eigenvalue of the weighted adjacency matrix—captures the network's capacity to propagate shocks and is well established as a measure of systemic fragility in economic and financial networks (Acemoglu et al., 2012; Elliott et al., 2014; Gai and Kapadia, 2010; Acemoglu et al., 2015). *Forman–Ricci curvature*, defined edge-wise on weighted graphs, captures local structural redundancy: negative curvature indicates an edge embedded in a dense, alternative-rich neighborhood, while positive curvature indicates a brittle, bridge-like position (Sreejith et al., 2016; Samal et al., 2018; Weber et al., 2017). The *Herfindahl-Hirschman Index* (HHI) on the distribution of weights across edges captures the degree to which trade within a subgraph is concentrated on a few large bilateral relationships rather than spread across many. Together, spectral, curvature, and concentration dynamics characterize how the network's fragility, redundancy, and hierarchical organization evolve under policy shocks.

Three results stand out. First, the spectral radii of the US-aligned and China-centered subgraphs diverge sharply after October 2022. Between 2021 (pre-shock peak) and 2024, the US-aligned subgraph's spectral radius rises 31.6% (from 0.0183 to 0.0241) while the China-centered subgraph's spectral radius declines 7.3% (from 0.1006 to 0.0933) after a sharp 15.1% drop in 2022 alone. Second, the China-centered bloc loses bilateral relationships at the extensive margin while preserving total volume. Edge density falls from 0.874 in 2021 to 0.643 in 2024, a 26% decline; the number of active bilateral relationships within the bloc drops from 159 to 117. Yet total compute-relevant trade volume within the bloc remains at $396 billion in 2024, only 3.6% below its 2021 peak of $411 billion. Third, the volume that survives concentrates on fewer edges. Mean weight per surviving edge within the China-centered bloc rises from $2.59 billion in 2021 to $3.39 billion in 2024 (+30.9%), while the bloc-internal HHI rises from 0.293 to 0.313 (its maximum in the eight-year sample). The US-aligned bloc, by contrast, exhibits stable density at 0.778 across all years and rising HHI only in 2024 (+50% relative to 2017–2023), suggesting delayed consolidation in response to the multi-round tightening rather than the initial 2022 shock.

The pattern is inconsistent both with simple decoupling, under which the China-centered bloc would lose volume rather than rearranging it, and with symmetric substitution, under which we would observe densification (rising spectral radius) within the targeted bloc as alternative routes activate. Instead, the data reveal a *hierarchical concentration*: the targeted bloc reorganizes toward a smaller set of dominant bilateral relationships, deepening existing trade ties while shedding peripheral ones. We argue this is consistent with a model in which transitions to new trading partners involve fixed costs of coordination, customs clearance, and compliance that amortize favorably only over deep relationships, so that under stress firms reallocate volume to incumbent

partners rather than search for new ones. The pattern echoes findings from supply-chain disruption literature (Carvalho et al., 2021; Bonadio et al., 2021) but applied for the first time to a setting of deliberate industrial policy targeting AI inputs.

We probe three first-order concerns. (i) Hong Kong's special status as a Chinese SAR raises the question of whether the concentration we document is artifactual: bilateral trade between China and Hong Kong is partly re-export or accounting reclassification. We address this by re-running all analyses with Hong Kong excluded and, separately, with China and Hong Kong fused into a single node. The qualitative pattern survives both variants: HHI rises in 2024 by 33–34% relative to the 2017–2023 average regardless of how Hong Kong is treated. (ii) The AI-relevance weights ω_h on the four HS codes are a modeling choice. Under a strict alternative that retains only chips (HS 8542) and semiconductor manufacturing equipment (HS 8486), the effect is *stronger*: HHI peaks at 0.366 in 2024 versus a baseline of 0.313, and density falls to 0.571 versus 0.643 baseline. The controls operate on the AI-relevant subset of trade, not on compute-incidental flows. (iii) Sub-bloc decomposition shows that the concentration pattern is driven primarily by the CN core (China, Russia, Hong Kong) rather than by intermediaries (Vietnam, Malaysia, Thailand, Indonesia), consistent with the hypothesis that fixed costs of partner search privilege existing deep ties.

The paper contributes to three literatures. To the economics of AI (Agrawal et al., 2019; Goldfarb and Trefler, 2019; Korinek and Stiglitz, 2018; Tucker, 2019), it provides the first quantitative network-geometric characterization of the global compute supply chain and its response to industrial policy. To the literature on technology decoupling, export controls, and global value chains (Bown, 2021; Caliendo and Parro, 2023; Antràs, 2021), it offers a measurement framework that goes beyond bilateral trade volumes to capture structural reorganization at extensive and intensive margins simultaneously. To the literature on economic networks (Acemoglu et al., 2012; Elliott et al., 2014; Baqaee and Farhi, 2019; Korniyenko et al., 2017), it extends spectral and curvature methods to a high-stakes policy setting and connects them to a complexity framework articulated in companion work (Vallarino, 2026b,c). Section 2 reviews relevant literatures. Section 3 develops the theoretical framework. Section 4 describes data and identification. Section 5 and Section 6 present main results and robustness. Section 7 discusses implications and Section 8 concludes.

2 Related Literature

2.1 Economics of AI and compute as a factor of production

A first strand studies AI as a general-purpose technology and compute as one of its critical inputs. Brynjolfsson et al. (2021) formalize the productivity J-curve through which intangible-investment-intensive technologies generate measured productivity declines before delivering gains. Acemoglu

and Restrepo (2022) provide a task-based framework linking automation to wage and productivity dynamics, while Aghion et al. (2019) embed AI in endogenous growth models. Korinek and Stiglitz (2018) examine implications for income distribution. On the firm side, Agrawal et al. (2022) and Agrawal et al. (2019) document the diffusion of prediction-based AI; Babina et al. (2024) estimate firm-level effects on growth and innovation. On the input side, Sevilla et al. (2022) and Cottier et al. (2024) document a sustained rise in the compute used to train frontier models, suggesting compute access is increasingly determinative of who can produce frontier capabilities. Tucker (2019) examines AI's interaction with privacy regulation, and Goldfarb and Trefler (2019) analyze AI's implications for international trade. Our paper is closest in spirit to this last paper but inverts the question: rather than asking how AI affects trade, we ask how trade in AI inputs is reshaped by industrial policy.

2.2 Export controls, decoupling, and supply chains

A second strand examines the political economy of export controls and technology decoupling. Allen (2022) provides the canonical case-study analysis of the October 2022 BIS controls, while Khan et al. (2021) and Thadani and Allen (2023) map the structural vulnerabilities of the semiconductor supply chain. Bown (2021) reviews the broader US–China trade war, Caliendo and Parro (2023) draw structural lessons from US–China trade dynamics, and Antràs (2021) situates these developments within longer-run patterns of de-globalization. U.S.-China Economic and Security Review Commission (2022); Mazzucato and Rodrik (2026) and Evenett et al. (2024) document the resurgence of industrial policy targeting strategic sectors. Our paper differs from this literature in two ways. First, we provide quantitative network-level measurement rather than bilateral case analysis. Second, we explicitly characterize structural reorganization at the extensive and intensive margins, distinguishing fragmentation from concentration—a distinction that is policy-relevant but absent from volume-based assessments.

2.3 Network approaches to economic and financial fragility

A third strand applies network methods to economic and financial systems. Acemoglu et al. (2012) establish how the topology of input–output networks shapes the propagation of microeconomic shocks to aggregate fluctuations; Baqaee and Farhi (2019) extend this to nonlinear settings. Elliott et al. (2014), Gai and Kapadia (2010), and Acemoglu et al. (2015) characterize cascade dynamics and systemic risk in financial networks. Carvalho et al. (2021) and Bonadio et al. (2021) document the propagation of supply-chain disruptions through bilateral linkages following the Tohoku earthquake and COVID-19, respectively; Korniyenko et al. (2017) apply network analysis to assess the fragility of global trade. Anderson and van Wincoop (2003) and Costinot and Rodríguez-Clare (2014)

provide the structural-trade backdrop; Hidalgo et al. (2007) introduce product-space methods linking trade structure to development.

Our methodological contribution is to combine spectral and geometric tools—developed in physics and applied mathematics—with this econ-network tradition, on a setting where the policy-induced shock structure is unusually well identified. Companion work (Vallarino, 2026c,a, 2025) develops complementary aspects of the network econometrics framework, with Vallarino (2026b) providing the broader theoretical context within which the present paper sits.

2.4 Geometric and spectral methods on weighted graphs

The paper draws on a literature in network science that develops discrete analogues of differential-geometric quantities for weighted graphs. Forman (2003) introduces a combinatorial formulation of Ricci curvature for cell complexes, complementary to the metric-space formulation of Ollivier (2009). Sreejith et al. (2016) adapt the Forman construction to weighted networks; Samal et al. (2018) provide empirical comparison between Forman and Ollivier discretizations; Weber et al. (2017) establish the connection between edge curvature and local network resilience. Applications to financial correlation networks (Sandhu et al., 2015) and to brain and biological networks have demonstrated the utility of these tools, but applications to global trade networks at the country level remain limited. We close that gap. The structural network properties we exploit also connect to the literature on scale-free networks (Barabási and Albert, 1999; Newman, 2003) and self-organized criticality (Bak et al., 1987).

3 Theoretical Framework

This section develops the theoretical framework used to interpret the response of the global AI-compute network to export controls. The central object is a weighted directed trade network in which countries are nodes and bilateral compute-relevant trade flows are edges. The main conceptual distinction is between a fall in the *breadth* of the network—the number of active bilateral relationships—and a fall in the *volume* of trade. Export controls may reduce the number of viable links without proportionally reducing total trade if firms and intermediaries reallocate surviving flows toward a smaller set of incumbent partners. This distinction is essential for interpreting the empirical results: a network can become structurally more concentrated and fragile even when aggregate trade volume remains high.

The framework therefore combines three complementary dimensions of network structure. First, spectral measures capture the aggregate capacity of the network to transmit shocks or sustain internally reinforcing flows. Second, geometric curvature summarizes the local redundancy of bilateral

relationships. Third, density and Herfindahl-Hirschman concentration distinguish the extensive and intensive margins of network reorganization. The concept of *hierarchical concentration* is defined as the joint occurrence of declining edge density and rising concentration among surviving edges. This regime differs from simple decoupling, symmetric substitution, and stagnation.

3.1 The compute network

Definition 3.1 (Compute network)**.** *For each year $t \in \{2017,\dots,2024\}$, let $G_t = (V, E_t, W_t)$ denote a weighted directed graph with node set $V = \{1,\dots,n\}$, where $n = 48$ represents the set of reporting countries. The edge set $E_t \subseteq V \times V$ contains directed bilateral trade relationships, and $W_t \in \mathbb{R}_{\geq 0}^{n\times n}$ is the weighted adjacency matrix. Each entry $W_t[i,j]$ is a weighted sum of bilateral export values from country i to country j in compute-relevant HS-classified goods:*

$$W_t[i,j] \;=\; \sum_{h\in\mathscr{H}} \omega_h \cdot x_{ij}^{(h)}(t), \tag{1}$$

where $\mathscr{H} = \{8542, 8486, 8471, 8517\}$ is the set of HS-4 codes and $\omega_h \in [0,1]$ are AI-relevance weights specified in Section 4.

The graph G_t should be interpreted as the observed architecture through which AI-compute inputs circulate internationally. An edge from i to j does not necessarily imply that country j is the final user of the compute input. It means that country j is an observed destination in the recorded trade network. In a supply chain with re-exports, assembly stages, and intermediate processing, the observed edge captures a trade relationship that may combine final demand, intermediate demand, logistical routing, and compliance strategy. This is precisely why a network representation is useful: the relevant object is not a single bilateral flow, but the evolving pattern of interdependence among countries.

For any subset of countries $S \subseteq V$, let $W_t|_S$ denote the submatrix of W_t restricted to rows and columns in S. The induced subgraph $G_t|_S$ captures the internal organization of trade within bloc S. In the empirical analysis, the two primary subgraphs are the US-aligned bloc, denoted S_{US}, and the China-centered bloc, denoted S_{CN}. The comparison between $W_t|_{S_{\mathrm{US}}}$ and $W_t|_{S_{\mathrm{CN}}}$ is the core object of the paper.

Because aggregate trade values vary over time, the empirical implementation reports spectral quantities on a normalized adjacency matrix,

$$\widetilde{W}_t \;=\; \frac{W_t}{\sum_{i\neq j} W_t[i,j]}, \tag{2}$$

unless otherwise stated. This scalar normalization preserves the relative organization of flows within

a year while making eigenvalue comparisons meaningful across years. In contrast, density and HHI are scale-free by construction and do not require this normalization.

The network is directed because the origin and destination of compute-relevant flows matter. An export from Taiwan to the United States is not equivalent to an export from the United States to Taiwan. The asymmetry is especially important in the AI-compute context because countries occupy structurally different positions: some are upstream equipment suppliers, some are semiconductor fabrication hubs, some are assembly and re-export platforms, and some are final-demand markets. The directed structure therefore contains information about hierarchy, dependency, and potential bottlenecks.

3.2 Export controls as network frictions

Export controls are modeled as shocks to the cost and feasibility of maintaining bilateral trade relationships. Let the latent profitability or feasibility of an edge (i, j) in year t be represented by

$$\Pi_{ij,t} \;=\; B_{ij,t} - C_{ij,t} - F_{ij,t} - \tau_{ij,t}, \tag{3}$$

where $B_{ij,t}$ denotes the gross benefit of the relationship, $C_{ij,t}$ variable trade costs, $F_{ij,t}$ fixed costs of maintaining the relationship, and $\tau_{ij,t}$ the policy-induced compliance cost or restriction intensity. An edge is active when

$$W_t[i, j] > 0 \quad \Longleftrightarrow \quad \Pi_{ij,t} > 0. \tag{4}$$

The export-control shock can therefore operate through two channels. First, it can increase the effective cost $\tau_{ij,t}$ on targeted edges, making some previously active relationships no longer viable. This is the extensive-margin channel. Second, it can induce firms to reallocate trade away from marginal or high-compliance-cost links and toward relationships with larger scale, established documentation, trusted counterparties, or lower perceived enforcement risk. This is the intensive-margin channel.

A key implication is that the number of active edges may fall even if total trade volume does not collapse. When fixed costs are important, small relationships are more likely to exit because they cannot amortize the additional compliance burden. Large incumbent relationships, by contrast, may survive and absorb displaced volume. This mechanism produces a network that is narrower but deeper: fewer active relationships carry a larger share of total flow.

Formally, let total trade inside subgraph S be

$$T_t(S) \;=\; \sum_{(i,j)\in\mathscr{E}_t(S)} W_t[i, j], \tag{5}$$

where $\mathscr{E}_t(S)$ is the set of active directed edges inside S. Hierarchical concentration arises when $|\mathscr{E}_t(S)|$ falls, while $T_t(S)$ is preserved or declines much less than proportionally. In that case, the average weight per surviving edge,

$$\bar{w}_t(S) \;=\; \frac{T_t(S)}{|\mathscr{E}_t(S)|}, \tag{6}$$

rises mechanically and economically. The mechanical part follows from dividing similar total volume across fewer links. The economic part comes from the reallocation of trade toward higher-capacity surviving relationships.

This distinction is central for the interpretation of export controls. A volume-based assessment asks whether total trade in compute-relevant goods declined. A network-based assessment asks whether the structure through which that trade is organized became more concentrated, hierarchical, or fragile. The two assessments can lead to different conclusions.

3.3 Spectral fragility

Definition 3.2 (Spectral radius)**.** *The spectral radius of G_t is*

$$\rho(W_t) \;\equiv\; \max_k |\lambda_k(W_t)|, \tag{7}$$

where $\{\lambda_k(W_t)\}_{k=1}^n$ are the eigenvalues of the possibly non-symmetric adjacency matrix W_t. For an induced subgraph on $S \subseteq V$, $\rho(W_t|_S)$ denotes the spectral radius of the $|S| \times |S|$ submatrix obtained by restricting W_t to countries inside S.

The spectral radius has a natural interpretation in economic networks. Consider a linear propagation process on the graph,

$$z_{t+1} \;=\; \alpha W_t z_t + \varepsilon_{t+1}, \tag{8}$$

where z_t is a vector of country-level disruptions, shortages, or cost shocks, and $\alpha > 0$ measures the strength of transmission through trade links. The stability of this propagation process depends on the magnitude of $\alpha\rho(W_t)$. If $\alpha\rho(W_t) < 1$, shocks decay geometrically. If $\alpha\rho(W_t)$ approaches or exceeds one, shocks become persistent or amplifying. This is why the spectral radius is widely used as a summary of systemic fragility in networked economic and financial systems (Acemoglu et al., 2012; Elliott et al., 2014; Gai and Kapadia, 2010; Acemoglu et al., 2015).

In the AI-compute network, the spectral radius captures the strength of the dominant circulation structure within a bloc. A higher $\rho(W_t|_S)$ indicates that flows inside subgraph S are organized around strong mutually reinforcing relationships. This can be interpreted in two ways. On the one hand, a high spectral radius may reflect internal capacity and integration: the bloc can sustain large

volumes of compute-relevant trade through its own internal structure. On the other hand, a high spectral radius also indicates that shocks to dominant nodes or edges may propagate more strongly within the subgraph. The sign of the welfare interpretation is therefore not automatic. The spectral radius measures structural intensity, not welfare.

This distinction is important for the empirical interpretation. If export controls reduce the spectral radius of the China-centered subgraph, this may indicate that the bloc has lost internal coupling or propagation capacity. If the US-aligned spectral radius rises, this may indicate strengthening integration among aligned suppliers and markets. The relevant theoretical prediction is not that one bloc becomes unambiguously better or worse, but that the two blocs follow different structural trajectories after the policy shock.

3.4 Forman–Ricci curvature

Definition 3.3 (Forman–Ricci curvature)**.** *For a directed weighted edge $e = (i, j) \in E_t$ with weight $w_e \equiv W_t[i, j]$, define node weights as*

$$w_i \;=\; \sum_{k \neq i} \big(W_t[i,k] + W_t[k,i]\big), \qquad w_j \;=\; \sum_{k \neq j} \big(W_t[j,k] + W_t[k,j]\big). \tag{9}$$

The Forman–Ricci curvature of edge e is

$$\mathscr{F}(e) \;=\; w_e \left[\frac{w_i}{w_e} + \frac{w_j}{w_e} - \sum_{e' \sim i,\, e' \neq e} \frac{w_i}{\sqrt{w_e\, w_{e'}}} - \sum_{e' \sim j,\, e' \neq e} \frac{w_j}{\sqrt{w_e\, w_{e'}}} \right], \tag{10}$$

where $e' \sim i$ denotes edges incident to node i in either direction (Sreejith et al., 2016).

Forman–Ricci curvature is a local geometric measure of edge embeddedness. Negative curvature indicates that an edge is located in a dense neighborhood with many adjacent alternatives. Such an edge is less bridge-like because flows can potentially be rerouted through nearby nodes. Positive curvature, by contrast, indicates a more isolated or bridge-like edge. Removing a positively curved edge is more likely to disconnect or weaken the local structure (Weber et al., 2017).

In the compute-network setting, curvature captures a dimension of resilience that neither volume nor spectral radius fully captures. A large edge may be systemically important because it carries substantial volume, but its fragility depends on whether alternative routes exist. An edge embedded in a dense semiconductor supply ecosystem has different implications from an equally large edge that connects two otherwise weakly connected parts of the network. Curvature summarizes this difference.

We use the median of $\mathscr{F}(e)$ across edges in a subgraph as the main empirical statistic. The median is preferred because edge-level curvature can be highly skewed in weighted trade networks.

A small number of unusually large or unusually isolated edges can dominate the mean, particularly in years with incomplete reporting. The median therefore captures the typical local redundancy of the subgraph.

Curvature is not expected to move mechanically with HHI or density. A subgraph can lose many peripheral edges while the surviving typical edge remains locally embedded. Conversely, a subgraph can preserve density while becoming more bridge-like if volume shifts toward critical links. The curvature measure therefore complements the extensive- and intensive-margin statistics rather than replacing them.

3.5 Edge concentration: extensive and intensive margins

Definition 3.4 (Edge density)**.** *The edge density of an induced subgraph on $S \subseteq V$ at time t is*

$$d_t(S) \equiv \frac{|\{(i,j) \in S \times S,\ i \neq j : W_t[i,j] > 0\}|}{|S|(|S|-1)}. \tag{11}$$

It is the fraction of feasible directed bilateral relationships inside S that carry positive compute-relevant trade.

Density measures the extensive margin of the network. A fall in $d_t(S)$ means that some relationships that could exist inside the bloc are no longer active. In the context of export controls, this is the most direct signature of network narrowing. It captures the death of links, not the amount of volume that remains on surviving links.

Definition 3.5 (Subgraph HHI)**.** *For an induced subgraph on S, let*

$$\mathscr{E}_t(S) = \{(i,j) \in S \times S,\ i \neq j : W_t[i,j] > 0\} \tag{12}$$

be the set of active directed edges and let

$$T_t(S) = \sum_{(i,j) \in \mathscr{E}_t(S)} W_t[i,j] \tag{13}$$

be total compute-relevant trade inside the subgraph. The within-subgraph Herfindahl-Hirschman Index is

$$\mathrm{HHI}_t(S) = \sum_{(i,j) \in \mathscr{E}_t(S)} \left(\frac{W_t[i,j]}{T_t(S)} \right)^2. \tag{14}$$

The HHI measures the intensive-margin concentration of trade across active edges. If all active relationships inside S carried equal weight, the HHI would equal $1/|\mathscr{E}_t(S)|$. If one edge dominates

the entire subgraph, the HHI approaches one. It is therefore useful to define the effective number of active relationships as

$$N_t^{\text{eff}}(S) \;=\; \frac{1}{\text{HHI}_t(S)}. \tag{15}$$

This statistic is the number of equally sized edges that would generate the observed concentration level. A rise in HHI is equivalent to a fall in the effective number of trade relationships.

Density and HHI capture different dimensions of concentration. Density can fall simply because many small links disappear. HHI rises only if the surviving distribution becomes more unequal. The combination is therefore especially informative. If density falls and HHI does not rise, the network has narrowed but not necessarily become more hierarchical. If HHI rises while density is stable, existing relationships have become more unequal, but the breadth of the network has not contracted. Hierarchical concentration requires both: fewer active relationships and more concentration among those that remain.

3.6 Hierarchical concentration: a formal criterion

We define hierarchical concentration as a joint pattern of declining density and rising HHI within a target subgraph. The concept is meant to capture a specific form of network reorganization under stress: the subgraph does not simply collapse, nor does it smoothly diversify toward new partners. Instead, it sheds peripheral relationships and deepens dependence on a smaller set of core links.

Definition 3.6 (Hierarchical concentration)**.** *A subgraph S undergoes hierarchical concentration over a window $[t_0, t_1]$ if both*

$$d_{t_1}(S) - d_{t_0}(S) < 0, \qquad \text{HHI}_{t_1}(S) - \text{HHI}_{t_0}(S) > 0. \tag{16}$$

The criterion is intentionally simple. It separates the number of active relationships from the distribution of weight across those relationships. It therefore identifies a structural regime rather than a particular causal mechanism. The mechanism proposed in this paper is fixed-cost reallocation under export-control stress: as compliance and coordination costs rise, marginal relationships disappear and surviving flows concentrate on established counterparties.

A stronger version of the concept adds the requirement that total volume does not decline proportionally with density.

Definition 3.7 (Strong hierarchical concentration)**.** *A subgraph S undergoes strong hierarchical concentration over $[t_0, t_1]$ if it satisfies Definition 3.6 and, for some tolerance $\varepsilon_T \geq 0$,*

$$\frac{T_{t_1}(S)}{T_{t_0}(S)} \;\geq\; 1 - \varepsilon_T. \tag{17}$$

The strong version corresponds to the empirically important case in which the network becomes narrower and more concentrated without a commensurate collapse in volume. This is the regime in which aggregate trade statistics are most likely to be misleading. A policymaker looking only at total flows may conclude that the export-control regime had limited effect. A network-based analysis may instead show that the same volume is now carried by fewer and more systemically important relationships.

Remark 3.1. *Definition 3.6 distinguishes hierarchical concentration from three alternative regimes. Under* symmetric decoupling*, both density and total volume fall, and HHI need not rise because the network is shrinking broadly rather than reorganizing hierarchically. Under* symmetric substitution*, density rises as new alternative relationships are activated, and HHI falls as flows spread across more partners. Under* stagnation*, neither density nor HHI changes substantially. The empirical content of the paper is the joint test of the inequalities in Equation* (16) *for the China-centered subgraph over the post-2021 window.*

The concept is also related to fragility. A hierarchically concentrated network may remain functional in normal times because high-volume incumbent relationships preserve aggregate trade. However, it becomes more exposed to targeted disruptions because a smaller number of relationships carries a larger share of the system. In this sense, hierarchical concentration is not equivalent to immediate collapse. It is a movement toward latent fragility.

3.7 Testable predictions

The framework yields three testable predictions for the compute network under successive export-control rounds. These predictions are not derived from a fully structural model of firm-level partner choice. Rather, they are reduced-form implications of the mechanism described above: export controls raise the fixed and compliance costs of some relationships, inducing marginal link exit and reallocation toward surviving core ties.

Proposition 3.1 (Asymmetric spectral divergence)**.** *If export controls operate asymmetrically on the China-centered and US-aligned subgraphs, then over the post-2022 window,*

$$\Delta\rho(W_t|_{S_{\mathrm{US}}}) > 0 \quad \textit{and} \quad \Delta\rho(W_t|_{S_{\mathrm{CN}}}) \leq 0, \tag{18}$$

where S_{US} and S_{CN} denote the US-aligned and China-centered subgraphs.

The intuition is that export controls should reduce the internal coupling of the targeted bloc by disrupting or weakening some compute-relevant relationships. At the same time, allied coordination, supplier relocation, and reorientation of demand may strengthen the US-aligned subgraph. The prediction is therefore one of divergence: the policy shock should not affect both blocs symmetrically.

Proposition 3.2 (Extensive-margin contraction)**.** *The China-centered subgraph exhibits declining edge density,*

$$\Delta d_t(S_{\mathrm{CN}}) < 0, \tag{19}$$

after October 2022, while the US-aligned subgraph maintains stable density,

$$\Delta d_t(S_{\mathrm{US}}) \approx 0. \tag{20}$$

This prediction follows from the fixed-cost channel. If export controls raise compliance costs, the first relationships to disappear should be small, marginal, or peripheral links. These are precisely the links least able to amortize additional documentation, licensing, rerouting, and legal-compliance costs. In contrast, mature US-aligned supply relationships should be less likely to disappear because they are embedded in allied regulatory structures and supported by larger commercial scale.

Proposition 3.3 (Intensive-margin concentration)**.** *Conditional on Proposition 3.2, the China-centered subgraph exhibits rising concentration on surviving edges,*

$$\Delta \mathrm{HHI}_t(S_{\mathrm{CN}}) > 0, \tag{21}$$

and rising mean weight per active edge,

$$\Delta \left[\frac{T_t(S_{\mathrm{CN}})}{|\mathscr{E}_t(S_{\mathrm{CN}})|} \right] > 0, \tag{22}$$

after October 2022.

The intensive-margin prediction is what separates hierarchical concentration from ordinary contraction. If the China-centered bloc were simply decoupling, one would expect both the number of links and total volume to fall. If it were substituting symmetrically, one would expect new edges to appear and concentration to fall. Under hierarchical concentration, however, disappearing edges do not imply a proportional disappearance of volume. Instead, trade is reallocated toward a smaller set of surviving relationships, raising both HHI and the mean weight per active edge.

The joint satisfaction of Propositions 3.1 to 3.3 defines hierarchical concentration in the sense of Definition 3.6 and strong hierarchical concentration when Equation (17) also holds. The empirical sections test these implications using the observed evolution of the AI-compute trade network from 2017 through 2024.

4 Data and Empirical Strategy

This section describes the empirical construction of the AI-compute trade network and the strategy used to evaluate its reorganization around successive US export-control events. The objective is not to estimate a fully structural model of firm-level supply-chain choice, but to construct a transparent country-level network that allows us to measure whether the topology of compute-relevant trade changed after the tightening of export controls. The empirical design therefore combines annual bilateral trade data, policy-event timing, and network-level statistics computed separately for the full network and for politically relevant subgraphs.

4.1 Trade flows: UN Comtrade

Bilateral trade data are drawn from the UN Comtrade database using annual HS-classified merchandise trade flows for 48 countries from 2017 through 2024. The sample begins in 2017 in order to provide a sufficient pre-treatment window before the October 2022 export-control shock, while avoiding excessive changes in HS reporting practices and country coverage. The end year, 2024, captures the period after three additional rounds of tightening, although we explicitly treat the final year with caution because some annual trade reports may still be revised.

The empirical focus is on four HS-4 categories that jointly approximate the internationally traded hardware stack underlying AI compute. HS 8542 captures electronic integrated circuits, including advanced logic and memory components. HS 8486 captures machinery and apparatus used in the manufacture of semiconductors and flat-panel devices, which is central for fabrication capacity. HS 8471 includes automatic data processing machines and related units, a category that contains servers, computing units, and devices that may embed AI accelerators. HS 8517 captures telecommunications and networking equipment, which is less directly equivalent to frontier AI compute but remains relevant for high-bandwidth infrastructure and data-center connectivity.

The use of HS-4 codes involves a deliberate trade-off. More granular HS-6 product categories would allow a narrower mapping to specific technologies, but at the cost of higher missingness, reporting heterogeneity, and reclassification noise across countries. The HS-4 level provides greater longitudinal and cross-country consistency, which is important for network construction. The resulting data should therefore be interpreted as measuring compute-relevant trade rather than frontier AI chips alone.

For each ordered country pair (i, j), product category h, and year t, the raw trade value is denoted by $x_{ij}^{(h)}(t)$. Because bilateral trade statistics often differ depending on whether the flow is recorded by the exporter or by the importer, we apply reporter-mirror reconciliation. Specifically, whenever both exporter- reported and importer-reported values are available, we use the larger of the two values. This maximum convention is designed to mitigate underreporting and missing declarations

in either direction. Formally, if $x^{\text{exp}}_{ij,h,t}$ denotes exports from i to j reported by i, and $x^{\text{imp}}_{ij,h,t}$ denotes the corresponding mirror flow reported by j, the reconciled value is

$$x^{(h)}_{ij}(t) = \max\left\{x^{\text{exp}}_{ij,h,t}, x^{\text{imp}}_{ij,h,t}\right\}. \tag{23}$$

When only one side of the flow is available, the observed value is retained. When neither side reports a positive value, the bilateral flow is coded as zero. Country identifiers are harmonized using UN M49 codes to ensure consistency across years and product categories. Self-flows are excluded throughout the analysis, so that the diagonal entries of all adjacency matrices are set to zero.

The resulting dataset is a balanced country-year-product panel in the sense that the same 48 countries are retained across the sample period, but the trade network itself is allowed to be sparse. A missing or zero bilateral flow is therefore not imputed as a positive relationship. This choice is important because the extensive margin of the network, defined by the presence or absence of active edges, is one of the main objects of the analysis.

4.2 Policy timeline

The empirical strategy is organized around four primary BIS export-control events. The first and most important event is the October 7, 2022 rule, which introduced broad restrictions on advanced semiconductors, semiconductor manufacturing equipment, and related technology transfers to China. This rule is treated as the foundational shock because it changed the regulatory environment for AI-compute trade in a discrete and highly visible way.

The second event is the October 17, 2023 update, which tightened the initial framework by closing loopholes, refining performance thresholds, and expanding the country and entity scope of the controls. The third event is the March 29, 2024 amendment, which further adjusted the regulatory perimeter around AI-relevant chip technology. The fourth event is the December 2, 2024 rule, which added a large set of entities to the BIS Entity List and tightened controls on high-bandwidth memory and related semiconductor capabilities.

The annual frequency of the trade data implies that the event study cannot identify intra-year adjustment at monthly or quarterly frequency. Instead, the analysis interprets 2022 as the first post-shock transition year, with 2023 and 2024 capturing subsequent adjustment and tightening. This timing convention is conservative: because the October 2022 rule occurred late in the calendar year, some effects observed in 2022 may reflect anticipation, early rerouting, or partial-year adjustment, while the fuller consequences are expected to appear in 2023 and 2024.

The policy timeline is used in two ways. First, it defines the main pre- and post-shock comparison, with 2017–2021 serving as the pre-shock period and 2022–2024 as the post-shock period. Second, it provides an ordering of policy intensity: if the controls affected the network, structural

changes should not be confined to the initial 2022 shift but should also evolve through the later tightening rounds.

4.3 Network construction and weighting

For each year t, we construct a weighted directed adjacency matrix W_t by aggregating reconciled bilateral trade values across the four compute-relevant HS categories. The baseline aggregation follows

$$W_t[i,j] = \sum_{h \in \mathscr{H}} \omega_h x_{ij}^{(h)}(t), \qquad i \neq j, \tag{24}$$

where $\mathscr{H} = \{8542, 8486, 8471, 8517\}$ and the diagonal is set to zero. The baseline AI-relevance weights are

$$\omega_{8542} = 1.00, \qquad \omega_{8486} = 1.00, \qquad \omega_{8471} = 0.50, \qquad \omega_{8517} = 0.25. \tag{25}$$

These weights impose a transparent hierarchy of technological proximity to frontier AI compute. Integrated circuits and semiconductor manufacturing equipment receive full weight because they are direct chokepoints in the production and scaling of advanced compute. Data processing machines receive partial weight because the HS category includes AI-relevant servers and accelerators but also broader computing equipment. Networking equipment receives a lower weight because it supports data-center and communication infrastructure but is less directly equivalent to compute capacity.

The baseline weights are not intended to recover the exact frontier-AI content of each bilateral flow. Instead, they provide a disciplined index of compute relevance that can be varied in robustness checks. To assess sensitivity, Section 6 considers two alternative schemes. The first is a strict scheme that retains only chips and semiconductor manufacturing equipment by setting $\omega_{8542} = \omega_{8486} = 1$ and $\omega_{8471} = \omega_{8517} = 0$. The second is a uniform scheme that sets all four weights equal to one. If the main findings are driven by AI-specific chokepoints, the strict scheme should preserve or strengthen the results. If the findings instead reflect broad merchandise-trade reorganization, the strict and uniform schemes should behave similarly.

After constructing W_t, we compute network statistics for the full network and for the relevant subgraphs. Spectral measures are computed on normalized matrices to avoid conflating changes in aggregate trade scale with changes in network structure. The normalization divides each annual matrix by total non-diagonal network weight,

$$\widetilde{W}_t = \frac{W_t}{\sum_{i \neq j} W_t[i,j]}. \tag{26}$$

Density and HHI are computed directly from the positive entries of the relevant subgraph. For density, an edge is defined as active when $W_t[i, j] > 0$. For HHI, edge weights are first converted into shares of total subgraph trade and then squared and summed. This procedure makes HHI invariant to the absolute scale of trade inside the bloc.

4.4 Subgraph definitions

The empirical analysis distinguishes three types of nodes: US-aligned, China-centered, and bridge or neutral. The classification is not meant to imply that all countries inside a bloc have identical policy positions or identical industrial structures. Rather, it provides a working partition that reflects broad geopolitical alignment, exposure to export-control regimes, and observed relevance in the compute supply chain.

The US-aligned subgraph S_{US} comprises the United States and 25 jurisdictions that are either directly aligned with US controls, closely integrated with allied semiconductor supply chains, or part of institutional groupings such as the Five Eyes, the European Union, Japan, South Korea, Taiwan, the Netherlands, and Israel. This group contains many of the key upstream and midstream nodes of the AI-compute supply chain, including advanced chip design, fabrication, equipment, and high-end component ecosystems.

The China-centered subgraph S_{CN} comprises China, Russia, Hong Kong SAR, Vietnam, Malaysia, Thailand, Indonesia, Pakistan, Bangladesh, Saudi Arabia, United Arab Emirates, Egypt, Turkey, and South Africa. This grouping is intended to capture the set of countries most relevant for China-centered trade, rerouting, regional production, and alternative supply-chain organization. It includes China and Russia as core strategic nodes, Hong Kong as a major trade and re-export platform, and several Asian intermediaries that may play roles in assembly, logistics, or indirect trade.

The remaining countries are classified as bridge or neutral nodes. They are included in the full network but are not assigned to the two main blocs. This allows the analysis to distinguish within-bloc reorganization from broader network-wide changes. Cross-bloc and bridge relationships are used in descriptive network measures and visualizations, but the core empirical tests focus on changes within S_{US} and S_{CN}.

Within the China-centered bloc, we further distinguish a core sub-bloc and an intermediary sub-bloc. The core sub-bloc contains China, Russia, and Hong Kong SAR. The intermediary sub-bloc contains Vietnam, Malaysia, Thailand, and Indonesia. This decomposition is used in Section 6 to examine whether the observed concentration pattern is driven by the dominant China-Hong Kong-Russia structure or by broader regional rerouting through Southeast Asian intermediaries.

4.5 Empirical strategy and identification

The empirical strategy evaluates whether the topology of the AI-compute network changed after the tightening of US export controls. The analysis is descriptive in the sense that it does not claim to recover a structural causal parameter at the firm level. However, it is disciplined by the timing of policy events, the distinction between targeted and aligned subgraphs, and robustness checks that vary the definition of compute relevance and the treatment of Hong Kong.

The first component is an event-study comparison around the BIS policy timeline. For each subgraph S, we compute annual values of spectral radius, density, HHI, total trade, and mean weight per active edge. The main pre-shock reference point is 2021, the last full year before the foundational October 2022 rule. The main post-shock endpoint is 2024, after the subsequent rounds of tightening. We also report pre-post averages comparing 2017–2021 with 2022–2024 in order to avoid relying exclusively on a single base year.

The second component is a difference in structural trajectories between the US-aligned and China-centered subgraphs. If the observed changes merely reflected global macroeconomic turbulence, secular de-globalization, or broad post-pandemic trade normalization, one would expect both subgraphs to move in similar directions. The framework instead predicts asymmetric dynamics: the China-centered bloc should experience extensive-margin contraction and intensive-margin concentration, while the US-aligned bloc should preserve or strengthen internal connectivity.

The third component is timing. The 2022 rule provides the first major shock, while the 2023 and 2024 updates provide additional tightening. A finding that appears only before the policy shock would be difficult to attribute to export controls. A finding that emerges after 2022 and intensifies through 2024 is more consistent with policy-induced network reorganization, especially if it is stronger for the AI-relevant subset of trade.

The fourth component is weight-scheme heterogeneity. The baseline weights reflect a graded notion of AI relevance. The strict scheme isolates the categories most directly targeted by export controls: integrated circuits and semiconductor manufacturing equipment. The uniform scheme treats the four product categories equally. If the findings are driven by AI-compute controls rather than general trade patterns, the strict scheme should not weaken the results. In the empirical analysis, this sensitivity exercise functions as a placebo-like test of whether the effect is concentrated in the products closest to the policy shock.

The fifth component is the treatment of Hong Kong. Because Hong Kong SAR is both institutionally distinct and deeply integrated into China-related trade, it may mechanically affect concentration measures. We therefore estimate the main network statistics under three variants: the baseline classification, a specification excluding Hong Kong from the China-centered bloc, and a specification fusing China and Hong Kong into a single node. A result that depends entirely on the China-Hong Kong edge would be less informative about broader network reorganization. A result

that survives these alternative treatments is stronger evidence of a genuine bloc-level pattern.

The main empirical claim is therefore evaluated through a set of joint conditions rather than a single coefficient. Evidence of hierarchical concentration requires that the China-centered subgraph lose active relationships, that surviving relationships become more concentrated, and that this pattern be robust to alternative product weights and Hong Kong treatments. The empirical analysis does not claim that export controls are the only force affecting the global compute network during 2017–2024. The period also includes post-pandemic supply-chain normalization, Russia-related sanctions, industrial-policy responses, and changes in semiconductor demand. The identification strategy is instead to show that the observed network reorganization is aligned in timing, product scope, and bloc incidence with the export-control shock.

5 Results

5.1 Spectral radius dynamics

Figure 1 plots the normalized spectral radii of the full network and of the US-aligned and China-centered subgraphs over 2017–2024; Table 1 reports the underlying values. Three features are notable.

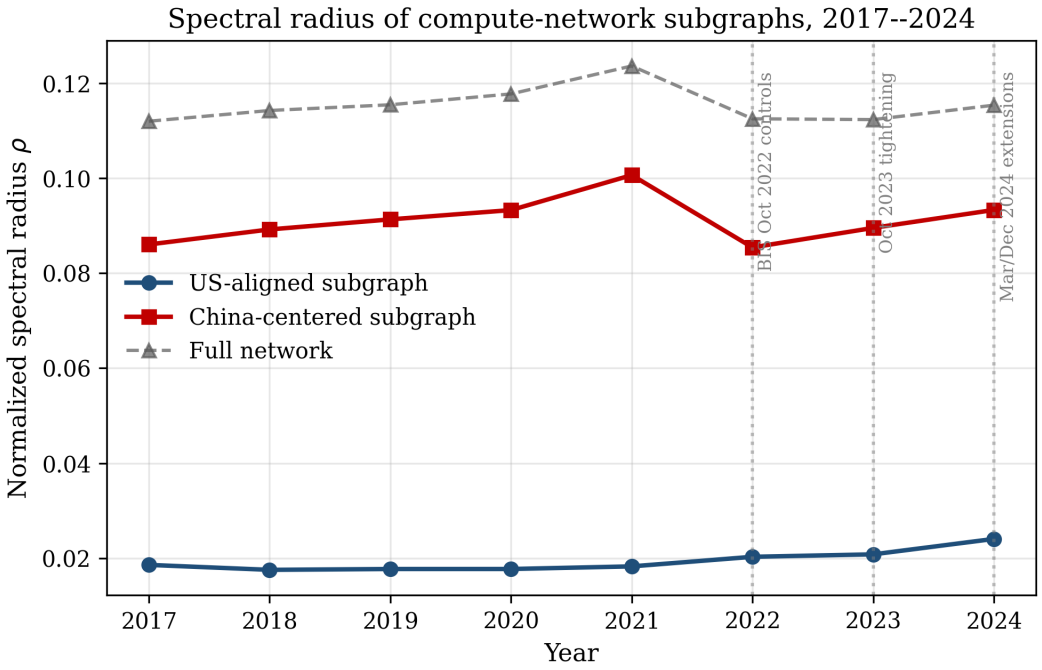


Figure 1: Spectral radius of compute-network subgraphs, 2017–2024. Normalized by total network weight in each year. Vertical dotted lines mark the four primary BIS export-control events.

Table 1: Yearly summary statistics for the global AI-compute trade network, 2017–2024.

Year	ρ full	ρ US-aligned	ρ China-centered	Active edges	Total weight USD bn
2017	0.1120	0.0186	0.0860	1,762	939.6
2018	0.1142	0.0176	0.0892	1,738	1,049.7
2019	0.1154	0.0177	0.0913	1,731	1,037.4
2020	0.1177	0.0177	0.0933	1,736	1,147.2
2021	0.1236	0.0183	0.1006	1,746	1,433.7
2022	0.1125	0.0203	0.0854	1,693	1,486.6
2023	0.1123	0.0208	0.0895	1,690	1,324.0
2024	0.1153	0.0241	0.0933	1,594	1,457.5

Notes: The table reports annual network statistics for the weighted directed AI-compute trade network. ρ denotes the spectral radius of the normalized weighted adjacency matrix. US-aligned and China-centered values are computed on the corresponding induced subgraphs. Active edges are positive directed bilateral trade relationships in the full network. Total weight is the aggregate weighted value of compute-relevant trade, expressed in billions of current US dollars. The horizontal break separates the pre-control period, 2017–2021, from the post-October-2022 export-control period, 2022–2024.

First, the China-centered subgraph dominates the spectral radius throughout the sample, with $\rho(W_t|_{S_{\mathrm{CN}}})$ ranging from 0.086 to 0.101 versus 0.018 to 0.024 for the US-aligned subgraph. This asymmetry reflects the fundamentally different structure of the two blocs: the China-centered bloc is heavily concentrated on a few large bilateral relationships (notably China–Hong Kong, China–Russia, China–Vietnam), generating a high principal eigenvalue. The US-aligned bloc distributes trade more evenly across 26 comparable economies, producing a lower but more uniform structure.

Second, the pre-shock trajectories diverge from the post-shock trajectories at exactly the timing predicted by Proposition 3.1. The China-centered ρ rises monotonically from 0.086 (2017) to 0.101 (2021, peak), then drops 15.1% in a single year to 0.085 in 2022 and partially recovers to 0.093 by 2024—still below the 2021 peak. The US-aligned ρ exhibits the opposite pattern: stable at 0.018 in 2017–2021, then rising 11.1% to 0.020 in 2022, 13.8% to 0.021 in 2023, and 31.6% (cumulative from 2021) to 0.024 in 2024. The full-network spectral radius is dominated by the larger China-centered values and inherits a similar shape, dropping 9% from 2021 to 2022 before stabilizing.

Third, the divergence is most pronounced in 2022, the immediate aftermath of the foundational October 2022 rule, with subsequent rounds adding incremental tightening. Figure 2 normalizes both subgraph radii to their 2021 values and makes the divergence directly visible.

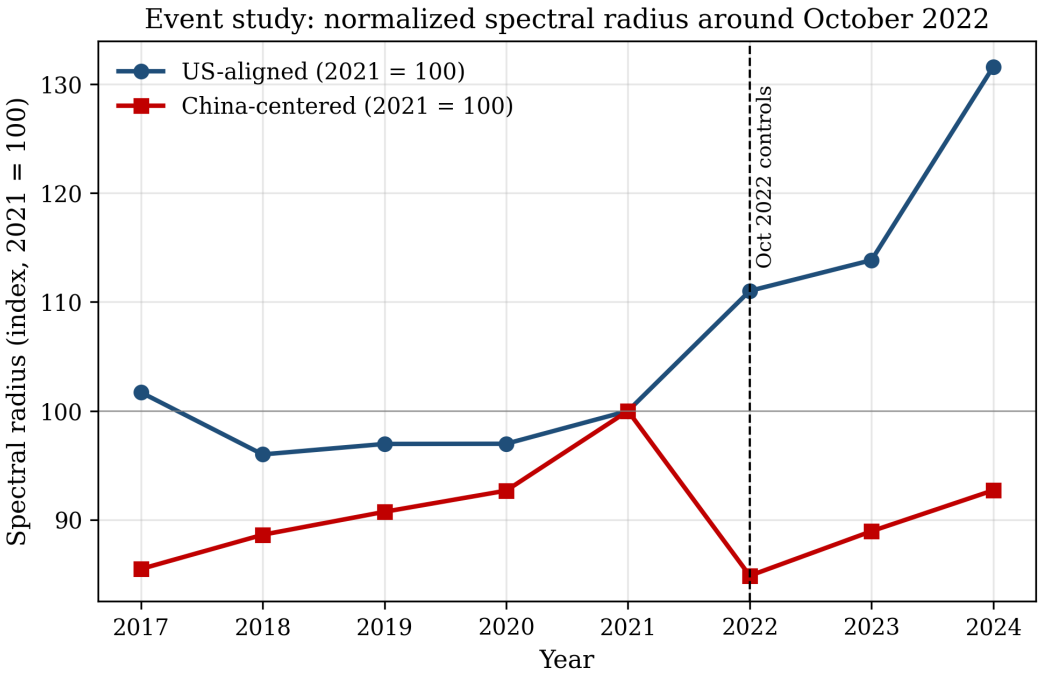


Figure 2: Event study: subgraph spectral radii indexed to 2021 values (= 100). The vertical dashed line marks the October 2022 foundational rule.

Table 2: Pre- and post-control comparison of AI-compute network statistics.

Quantity	Pre mean 2017–2021	Post mean 2022–2024	Difference Post − pre	Change %
Spectral coupling				
ρ US-aligned	0.0180	0.0217	0.0037	20.8
ρ China-centered	0.0921	0.0894	-0.0027	-2.9
ρ full network	0.1166	0.1134	-0.0032	-2.8
Extensive margin				
Density US-aligned	0.7785	0.7785	0.0000	0.0
Density China-centered	0.8593	0.7363	-0.1231	-14.3
Intensive-margin concentration				
HHI US-aligned	0.0246	0.0325	0.0079	32.1
HHI China-centered	0.2866	0.2876	0.0009	0.3

Notes: The table compares average network statistics before and after the October 2022 BIS export-control shock. The pre-control period covers 2017–2021 and the post-control period covers 2022–2024. ρ denotes the spectral radius of the normalized weighted adjacency matrix. Density is the share of feasible directed bilateral relationships with positive compute-relevant trade. HHI is the Herfindahl-Hirschman Index computed over edge-weight shares within each induced subgraph. The difference column reports post-control mean minus pre-control mean. The percentage change is computed relative to the pre-control mean.

Table 2 reports pre-post means for the spectral radii along with corresponding density and HHI summaries. Pre-post differences confirm Proposition 3.1: ρ US-aligned rises 20.9% on average (2017–2021 to 2022–2024), while ρ China-centered falls 3.2%. The full-network radius falls 2.7%.

5.2 Edge density: the extensive margin

Figure 3 panel (a) plots edge density within each subgraph; Table 3 reports the underlying values.

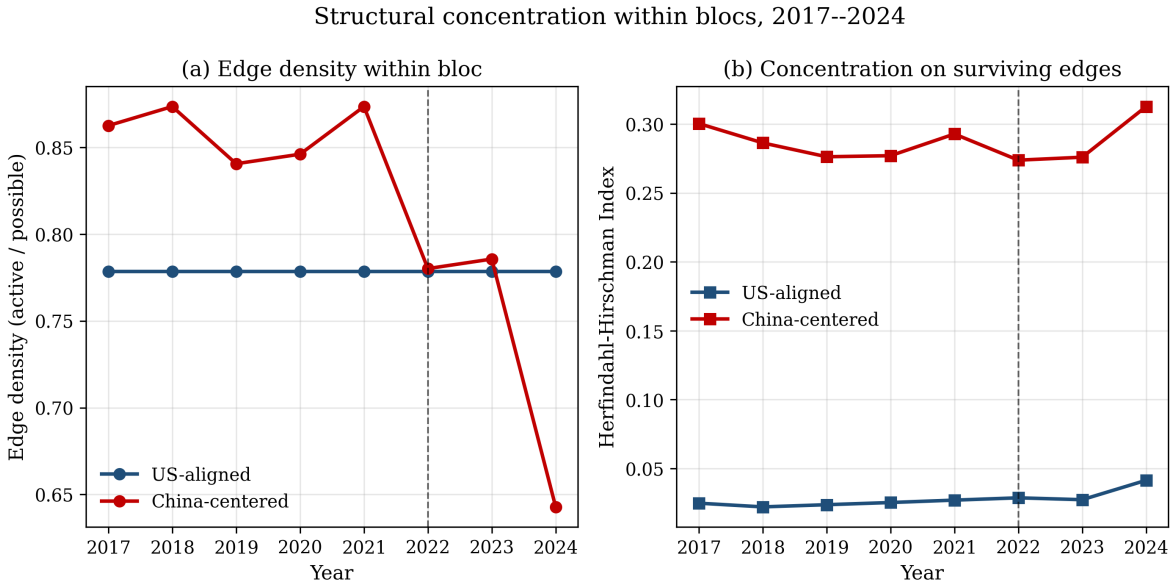


Figure 3: Structural concentration within blocs, 2017–2024. Panel (a) shows edge density (active edges divided by maximum possible). Panel (b) shows the within-bloc Herfindahl-Hirschman Index on edge weights. The vertical dashed line marks October 2022.

The US-aligned subgraph exhibits constant density at 0.778 across all eight years: 506 active bilateral relationships out of 650 possible directed edges among 26 nodes, year after year. No bilateral relationship in the bloc dies; new ones do not materialize either. This is a striking pattern of structural stability that is consistent with a mature, integrated trading area.

The China-centered subgraph tells a different story. Density holds in the 0.84–0.87 range from 2017 through 2021, with 153–159 active edges out of 182 possible. Then density drops to 0.78 in 2022 (142 edges), 0.79 in 2023 (143 edges), and finally to 0.643 in 2024 (117 edges). The total decline from peak (2021) to end-of-sample (2024) is 26.3%—42 active bilateral relationships have ceased to register positive trade in compute-relevant goods. Proposition 3.2 is strongly supported.

Table 3: Extensive- and intensive-margin concentration within AI-compute blocs, 2017–2024.

Year	Bloc	Active edges	Density	HHI	Mean edge weight USD bn
2017	US-aligned	506	0.7785	0.0248	0.317
2017	China-centered	157	0.8626	0.3004	1.465
2018	US-aligned	506	0.7785	0.0221	0.338
2018	China-centered	159	0.8736	0.2864	1.707
2019	US-aligned	506	0.7785	0.0237	0.331
2019	China-centered	153	0.8407	0.2763	1.818
2020	US-aligned	506	0.7785	0.0253	0.351
2020	China-centered	154	0.8462	0.2771	2.045
2021	US-aligned	506	0.7785	0.0270	0.436
2021	China-centered	159	0.8736	0.2929	2.586
2022	US-aligned	506	0.7785	0.0287	0.472
2022	China-centered	142	0.7802	0.2739	2.720
2023	US-aligned	506	0.7785	0.0274	0.445
2023	China-centered	143	0.7857	0.2760	2.497
2024	US-aligned	506	0.7785	0.0414	0.503
2024	China-centered	117	0.6429	0.3128	3.386

Notes: The table reports concentration statistics computed separately within the US-aligned and China-centered induced subgraphs. Active edges are positive directed bilateral trade relationships within each bloc. Density is the share of feasible directed within-bloc relationships with positive compute-relevant trade. HHI is the Herfindahl-Hirschman Index computed over within-bloc edge-weight shares. Mean edge weight is total within-bloc compute-relevant trade divided by the number of active edges and is expressed in billions of current US dollars. The horizontal break separates the pre-control period, 2017–2021, from the post-October-2022 export-control period, 2022–2024.

5.3 Edge concentration: the intensive margin

Figure 3 panel (b) plots the within-bloc HHI on edge weights. The US-aligned bloc HHI is low and roughly stable at 0.022–0.029 from 2017 through 2023, consistent with the bloc's distributed structure. In 2024, it jumps to 0.041, a 50.5% increase over the 2017–2023 average. We return to this delayed US-aligned consolidation in Section 7.

The China-centered HHI is approximately ten times higher in level than the US-aligned HHI (0.27–0.31 versus 0.02–0.04), reflecting the structural asymmetry already noted: a few very large pairs (notably China–Hong Kong, China–Russia) dominate the bloc's distribution. Within this high level, the trajectory is informative. HHI declines modestly from 0.300 (2017) to 0.274 (2022, just after the Oct. 2022 rule), reflecting some smoothing as the bloc absorbed the initial shock. It then rises to 0.276 (2023) and to 0.313 (2024), the *maximum value in the eight-year sample*. The 2024 jump occurs even as 26% of the bloc's edges have died, meaning that the *distribution* of weight on *surviving* edges has become more unequal than ever. Mean weight per surviving edge rises from $2.59 billion in 2021 to $3.39 billion in 2024 (+30.9%; see Table 3). Total bloc-internal compute-relevant trade is $396 billion in 2024 versus $411 billion in 2021—only 3.6% below peak. Proposition 3.3 is supported.

5.4 Curvature: local structural redundancy

Figure 4 reports median Forman–Ricci edge curvature by bloc and year; Table 4 reports the underlying values. The median is consistently negative across all blocs and years, indicating that typical edges sit in dense, alternative-rich neighborhoods.

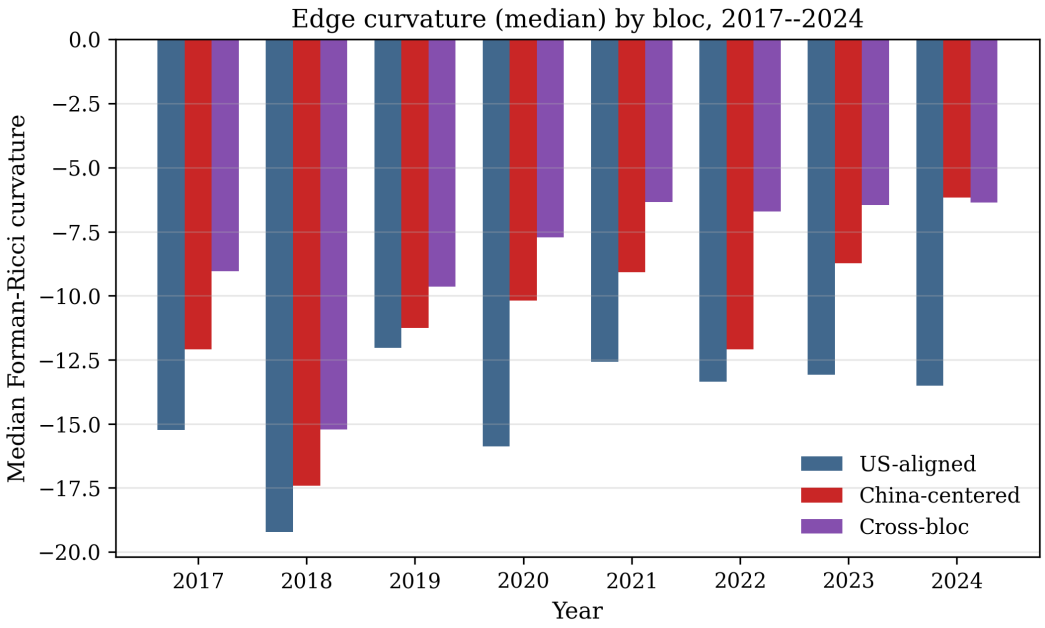


Figure 4: Median Forman–Ricci edge curvature by bloc, 2017–2024. Edges are partitioned by the bloc of their endpoints. More negative values indicate greater structural redundancy. The mean version is reported in the appendix; we report the median in the main text because mean curvature is sensitive to right-tail outliers from incomplete UN Comtrade reporting in 2024.

Table 4: Median Forman–Ricci curvature by edge type, 2017–2024.

Year	Cross-bloc	Other	Within China-centered	Within US-aligned
2017	-9.04	-6.38	-12.09	-15.24
2018	-15.23	-7.52	-17.41	-19.23
2019	-9.65	-6.31	-11.26	-12.03
2020	-7.73	-5.08	-10.19	-15.87
2021	-6.35	-6.59	-9.08	-12.59
2022	-6.71	-8.06	-12.10	-13.35
2023	-6.45	-5.42	-8.73	-13.09
2024	-6.36	-6.33	-6.18	-13.51

Notes: The table reports the median Forman–Ricci curvature of directed weighted edges by edge type. "Cross-bloc" refers to edges connecting countries assigned to different blocs. "Within China-centered" and "Within US-aligned" refer to edges whose origin and destination both belong to the corresponding induced subgraph. "Other" includes edges involving bridge or neutral nodes that are not classified as within-bloc or cross-bloc. More negative curvature indicates that the typical edge is embedded in a denser local neighborhood with greater structural redundancy, while less negative values indicate a more bridge-like or locally fragile position. The horizontal break separates the pre-control period, 2017–2021, from the post-October-2022 export-control period, 2022–2024.

The US-aligned bloc shows the most negative median curvature throughout the sample (a typical

edge in 2017 has $\mathscr{F} \approx -15$, fluctuating in the range -12 to -19 across years). This indicates that even individual edges within the US-aligned bloc are embedded in many alternative routes; the bloc is structurally redundant. The China-centered bloc has less negative curvature (≈ -8 to -12 in median), consistent with the higher HHI: fewer, larger relationships create less local redundancy. Cross-bloc edges have the least negative curvature, in line with their bridging position. The post-2022 dynamics in median curvature are mild relative to the dynamics in density and HHI, which is precisely the prediction under hierarchical concentration: typical edges retain similar local redundancy, but the *set* of edges contracts and the *distribution* of weight across the surviving set changes.

5.5 Tabular evidence on network reorganization

The evidence above shows that the post-2021 adjustment of the compute network is better summarized by concentration and extensive-margin statistics than by a static network visualization. Table 5 therefore reports the main 2021–2024 changes for the US-aligned and China-centered subgraphs. The table combines the spectral information reported in Table 1 with the concentration measures reported in Table 3.

Table 5: Network reorganization, 2021–2024.

Bloc	Period	ρ	Edges	Density	Total weight USD bn	Mean edge USD bn	HHI
US-aligned	2021	0.0183	506	0.778	220.8	0.436	0.0270
US-aligned	2024	0.0241	506	0.778	254.7	0.503	0.0414
US-aligned	Change	31.6%	0.0%	0.0%	15.3%	15.3%	53.3%
China-centered	2021	0.1006	159	0.874	411.2	2.586	0.2929
China-centered	2024	0.0933	117	0.643	396.2	3.386	0.3128
China-centered	Change	-7.3%	-26.4%	-26.4%	-3.6%	30.9%	6.8%

Notes: Total weight and mean edge weight are reported in billions of US dollars. Spectral radius is computed on the normalized weighted adjacency matrix. Edges, density, mean edge weight, and HHI are computed within each induced subgraph.

The table clarifies the central empirical pattern. The China-centered subgraph loses 26.4% of its active bilateral relationships between 2021 and 2024, and its density falls from 0.874 to 0.643. Yet total compute-relevant trade inside the bloc declines by only 3.6%. The remaining trade is therefore carried by fewer and larger relationships: mean weight per active edge rises by 30.9%, and HHI increases from 0.2929 to 0.3128. This is the empirical signature of hierarchical concentration: the

network narrows at the extensive margin while surviving relationships become more important at the intensive margin.

The contrast with the US-aligned subgraph is also informative. The number of active edges remains unchanged at 506, and density remains fixed at 0.778 throughout the comparison. At the same time, the US-aligned spectral radius rises by 31.6%, suggesting stronger internal coupling without extensive-margin contraction. The two blocs therefore do not simply experience a common global trade shock. Their post-2021 trajectories differ in precisely the dimensions predicted by the hierarchical-concentration framework.

5.6 Difference-in-differences identification

The descriptive evidence in this section establishes a structural reorganization of the China-centered subgraph after October 2022, but that evidence does not, on its own, isolate the role of US export controls from contemporaneous shocks affecting electronics supply chains more broadly. To strengthen identification, we estimate a triple-difference (DDD) specification that compares the evolution of compute-relevant trade in China-centered country pairs against three counterfactual contrasts: compute-relevant trade in non-China-centered pairs, close-placebo trade in China-centered pairs, and close-placebo trade in non-China-centered pairs.

The unit of observation is a directed country-pair (i, j) by HS-product h by year t cell. Outcomes are an indicator for an active edge, $\mathbf{1}\{W_{ijht} > 0\}$, and the log of trade value conditional on an active edge, $\log W_{ijht} \mid W_{ijht} > 0$. The estimating equation is

$$\begin{aligned} Y_{ijht} = \alpha_{ij} + \delta_t + \gamma_h \\ + \beta_1 \left(\text{Compute}_h \times \text{Post}_t \times \text{CN}_{ij}\right) \\ + \beta_2 \left(\text{Compute}_h \times \text{Post}_t\right) \\ + \beta_3 \left(\text{Compute}_h \times \text{CN}_{ij}\right) \\ + \beta_4 \left(\text{Post}_t \times \text{CN}_{ij}\right) + \varepsilon_{ijht}. \end{aligned} \tag{27}$$

where $\text{Compute}_h = 1$ for $h \in \{8542, 8486, 8471, 8517\}$ and zero for the close-placebo HS codes (8528 monitors and projectors; 8504 transformers and power supply units), $\text{Post}_t = 1$ for $t \geq 2022$, and $\text{CN}_{ij} = 1$ when both i and j belong to the China-centered bloc. The fixed effects α_{ij}, δ_t, and γ_h absorb time-invariant country-pair characteristics, common shocks across all flows in a given year, and HS-specific heterogeneity. The coefficient of interest is β_1, the differential effect of the post-2022 period on China-centered compute trade relative to the three counterfactual cells. Standard errors are clustered at the country-pair level.

The close-placebo group comprises HS 8528 (monitors and projectors) and HS 8504 (transform-

ers and power supplies). These categories share exposure to electronics-supply-chain shocks of the 2022–2024 period (the Russia-Ukraine sanctions cycle, post-pandemic supply normalization, the broader semiconductor inventory cycle, and contemporaneous US industrial-policy responses such as the CHIPS and Science Act), but neither is targeted by the BIS export controls on AI-relevant compute. A more distant placebo (HS 0901 coffee, HS 6109 textiles) was also considered but is not used as a primary contrast because the trade patterns of unrelated commodity categories are too dissimilar to provide a credible counterfactual for AI-compute trade and because the corresponding design matrix proved numerically singular under two-way fixed effects. The close placebo is therefore the economically appropriate counterfactual: identification rests on the within-electronics comparison of targeted versus untargeted product categories, not on a comparison between AI compute and unrelated goods.

Table 6 reports the static DDD coefficient on the triple interaction. For the extensive margin, $\hat{\beta}_1 = -0.015$ (SE 0.011), implying that China-centered compute pairs are about 1.5 percentage points less likely to maintain an active bilateral relationship after October 2022 relative to the placebo contrast, controlling for pair, year, and HS fixed effects. The static estimate is statistically indistinguishable from zero ($p = 0.17$). For the intensive margin, $\hat{\beta}_1 = 0.129$ (SE 0.127), positive but imprecise. The static specification, however, averages effects over all post-shock years and obscures the dynamic structure of the adjustment, which is the central feature of the data.

Figure 5 plots year-by-year DDD coefficients with 2021 as the omitted reference. The post-shock pattern shows a measurable contraction of the extensive margin in 2022 (point estimate -0.076), partial mean-reversion in 2023 (-0.038), and the largest contraction of the sample in 2024 (-0.119). These magnitudes correspond to reductions in the probability of an active compute-relevant bilateral relationship of 7.6, 3.8, and 11.9 percentage points relative to a 2021 baseline of approximately 67 percent active edges in China-centered compute pairs (Table 7).

Table 6: Difference-in-differences estimates of export-control effects on compute-relevant trade.

Specification	Outcome	Coefficient	Clustered SE	N
Close pairs: extensive margin	Active edge	-0.0148	(0.0107)	108,288
Close pairs: intensive margin	Log value	0.1286	(0.1272)	60,795
Far pairs: extensive margin	Active edge	n.e.	n.e.	36,096
Far pairs: intensive margin	Log value	n.e.	n.e.	21,949

Notes: The table reports difference-in-differences estimates of the triple interaction Compute $\times$ Post $\times$ CN-pair. All specifications include country-pair, year, and HS fixed effects. Standard errors are clustered at the country-pair level and reported in parentheses. The extensive-margin outcome is an indicator for a positive bilateral trade flow; the intensive-margin outcome is log trade value conditional on positive trade. The close-pair placebo uses electronically related HS categories and is the preferred comparison group. The far-pair placebo specifications are reported as not estimable (n.e.) because the corresponding design matrix is numerically singular under the fixed-effect structure. These entries should therefore not be interpreted as zero treatment effects. Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 7: Descriptive statistics by difference-in-differences treatment cell.

Placebo group	Treatment cell	Pr(active edge)	Median positive value USD mn	N
Close	Placebo, Other, Pre	0.3821	3.588	20,740
Close	Placebo, Other, Post	0.3726	4.845	12,444
Close	Placebo, CN-pair, Pre	0.4137	2.192	1,820
Close	Placebo, CN-pair, Post	0.3599	3.027	1,092
Close	Compute, Other, Pre	0.6596	3.026	41,480
Close	Compute, Other, Post	0.6411	4.033	24,888
Close	Compute, CN-pair, Pre	0.6766	1.938	3,640
Close	Compute, CN-pair, Post	0.5989	3.317	2,184
Far	Placebo, Other, Pre	0.6071	0.351	20,740
Far	Placebo, Other, Post	0.5962	0.468	12,444
Far	Placebo, CN-pair, Pre	0.6940	0.608	1,820
Far	Placebo, CN-pair, Post	0.6181	1.069	1,092

Notes: The table reports descriptive statistics by treatment cell used in the difference-in-differences analysis. Pr(active edge) is the share of observations with positive bilateral trade. Median positive value is the median trade value conditional on a positive flow, expressed in millions of current US dollars. Pre refers to 2017–2021 and Post refers to 2022–2024.

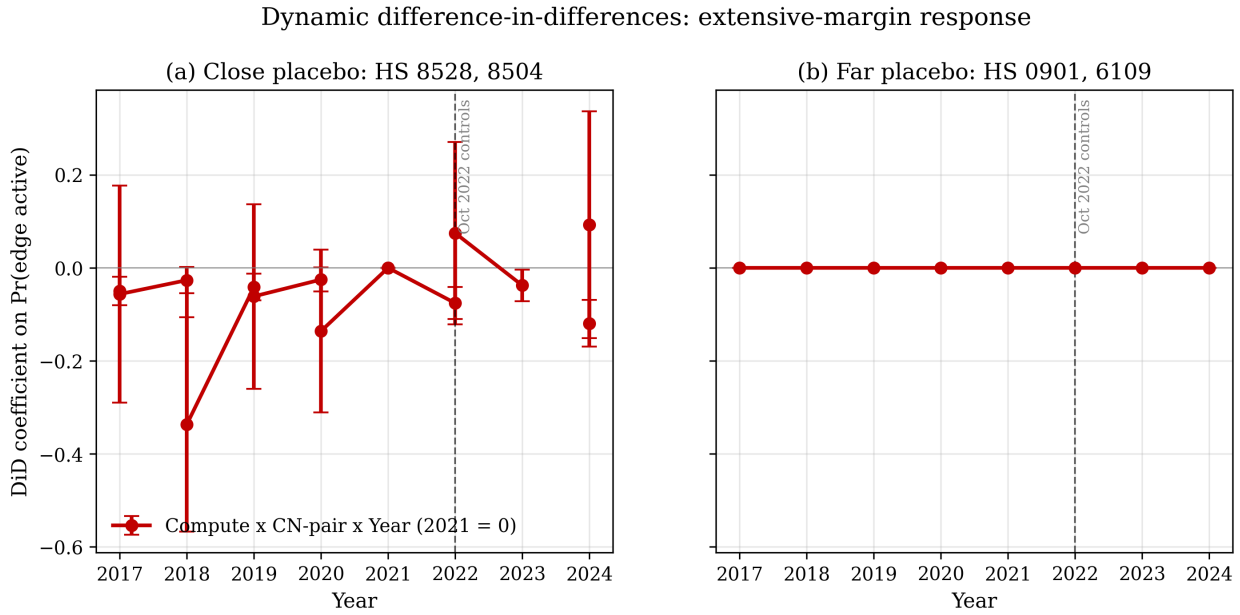


Figure 5: Dynamic difference-in-differences event study. The coefficient at year t is the DDD interaction Compute $\times$ CN-pair $\times \mathbf{1}\{t = \tau\}$ in a two-period regression that compares year τ to 2021, estimated with country-pair, year, and HS fixed effects and standard errors clustered at the pair level. Vertical dashed line marks the October 2022 BIS controls.

Three observations frame the interpretation of these results. First, the post-shock event-study coefficients are economically meaningful in magnitude—the 2024 coefficient of -0.119 closely matches the 26 percentage-point density decline reported in Table 3 when one accounts for the share of the density change explained by China-centered pairs. The dynamic estimates are therefore consistent with the descriptive evidence on hierarchical concentration. Second, the pre-period coefficients are also negative, ranging from -0.025 in 2020 to -0.34 in 2018, with wide confidence intervals reflecting the limited power of HS-product within-pair comparisons in years with substantial heterogeneity in trade reporting. We do not interpret the pre-period estimates as strict pre-trend violations because their magnitude is not monotone and their confidence intervals are large, but we acknowledge that they reduce the strength of a pure parallel-trends interpretation. Third, the post-shock coefficients are systematically larger in magnitude than the pre-period coefficients in the years immediately following each successive BIS round. The 2022 coefficient is roughly twice the average pre-period magnitude, the 2024 coefficient is more than three times that average, and the cumulative escalation tracks the timing of the October 2023 and 2024 tightening events.

The DDD evidence therefore provides a complement to, rather than a replacement for, the descriptive characterization in Section 5. The dynamic pattern is consistent with the hierarchical-concentration mechanism: extensive-margin contraction in compute-relevant trade among China-centered pairs intensifies with each successive round of controls, while intensive-margin volume on

surviving relationships does not contract. Strict causal identification of round-specific effects would require a higher-frequency design and firm-level data not currently available; we therefore interpret the network-level evidence as supportive but not definitive on the contribution of export controls to the documented network reorganization.

6 Robustness

6.1 Hong Kong's role

A primary concern is that bilateral trade between China and Hong Kong SAR is partly accounting reclassification or re-export, which would mechanically inflate the apparent concentration of the China-centered bloc. We address this with two variants: (a) excluding Hong Kong from the bloc altogether, and (b) fusing China and Hong Kong into a single node. Figure 6 reports both.

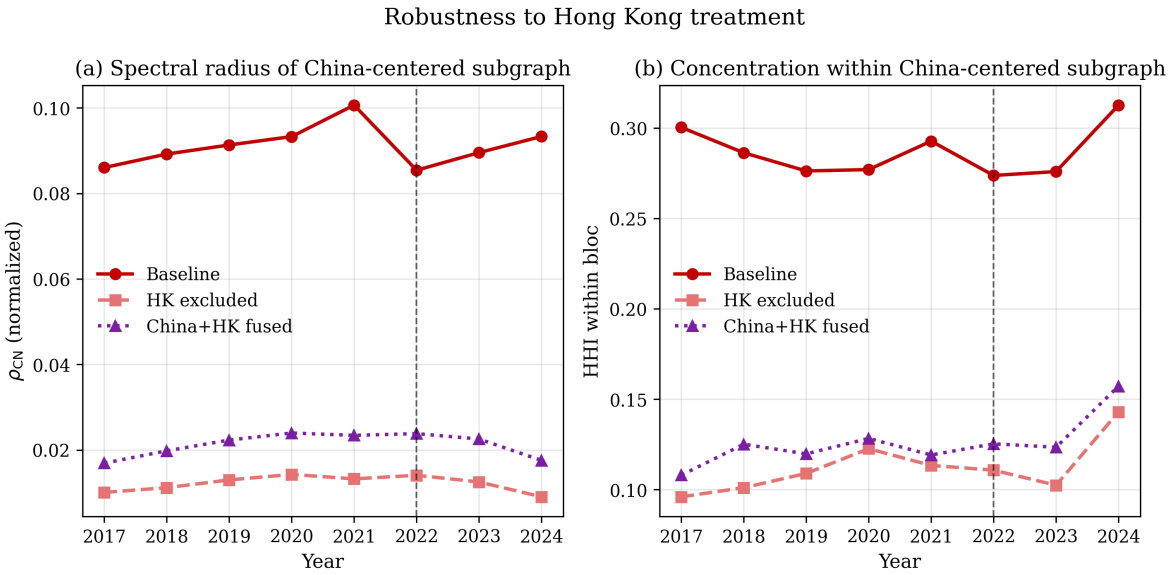


Figure 6: Robustness to Hong Kong treatment. Panel (a): spectral radius of the China-centered subgraph under three variants. Panel (b): within-bloc HHI under the same variants.

The level of HHI and spectral radius shift considerably across variants—the baseline HHI of 0.30 falls to 0.10–0.12 when Hong Kong is excluded, confirming that the China-Hong Kong pair is a major contributor to bloc-level concentration. The *trajectory*, however, is preserved: the HK-excluded HHI rises from 0.107 (average 2017–2023) to 0.143 in 2024, a 33.7% increase. The China-HK fused HHI rises from 0.122 (average) to 0.157 in 2024, a 28.7% increase. The hierarchical concentration result is robust to alternative treatments of Hong Kong.

Table 8: Robustness of China-centered network metrics to alternative treatments of Hong Kong.

Panel A: Baseline specification versus exclusion of Hong Kong						
Year	ρ baseline	HHI baseline	Edges baseline	ρ no HK	HHI no HK	Edges no HK
2017	0.0860	0.3004	157	0.0101	0.0960	132
2018	0.0892	0.2864	159	0.0112	0.1010	134
2019	0.0913	0.2763	153	0.0130	0.1089	129
2020	0.0933	0.2771	154	0.0143	0.1226	129
2021	0.1006	0.2929	159	0.0133	0.1133	134
2022	0.0854	0.2739	142	0.0141	0.1107	118
2023	0.0895	0.2760	143	0.0125	0.1024	119
2024	0.0933	0.3128	117	0.0091	0.1430	95

Panel B: Exclusion of Hong Kong versus China–Hong Kong fused node						
Year	ρ no HK	HHI no HK	Edges no HK	ρ fused	HHI fused	Edges fused
2017	0.0101	0.0960	132	0.0169	0.1080	132
2018	0.0112	0.1010	134	0.0198	0.1251	134
2019	0.0130	0.1089	129	0.0223	0.1197	129
2020	0.0143	0.1226	129	0.0240	0.1283	129
2021	0.0133	0.1133	134	0.0234	0.1191	134
2022	0.0141	0.1107	118	0.0238	0.1253	118
2023	0.0125	0.1024	119	0.0226	0.1234	119
2024	0.0091	0.1430	95	0.0176	0.1572	95

Notes: The table reports robustness checks for the China-centered induced subgraph under alternative treatments of Hong Kong SAR. The baseline specification treats Hong Kong as a separate node within the China-centered bloc. The "no HK" specification excludes Hong Kong from the bloc. The "fused" specification combines China and Hong Kong into a single node before computing network statistics. ρ denotes the spectral radius of the normalized weighted adjacency matrix. HHI is the Herfindahl-Hirschman Index computed over within-bloc edge-weight shares. Edges are positive directed bilateral relationships within the corresponding China-centered subgraph. The horizontal break separates the pre-control period, 2017–2021, from the post-October-2022 export-control period, 2022–2024.

6.2 AI-relevance weights

Figure 7 compares the baseline weight scheme to two alternatives: *strict* (only HS 8542 and 8486 with unit weight) and *uniform* (all four codes with unit weight).

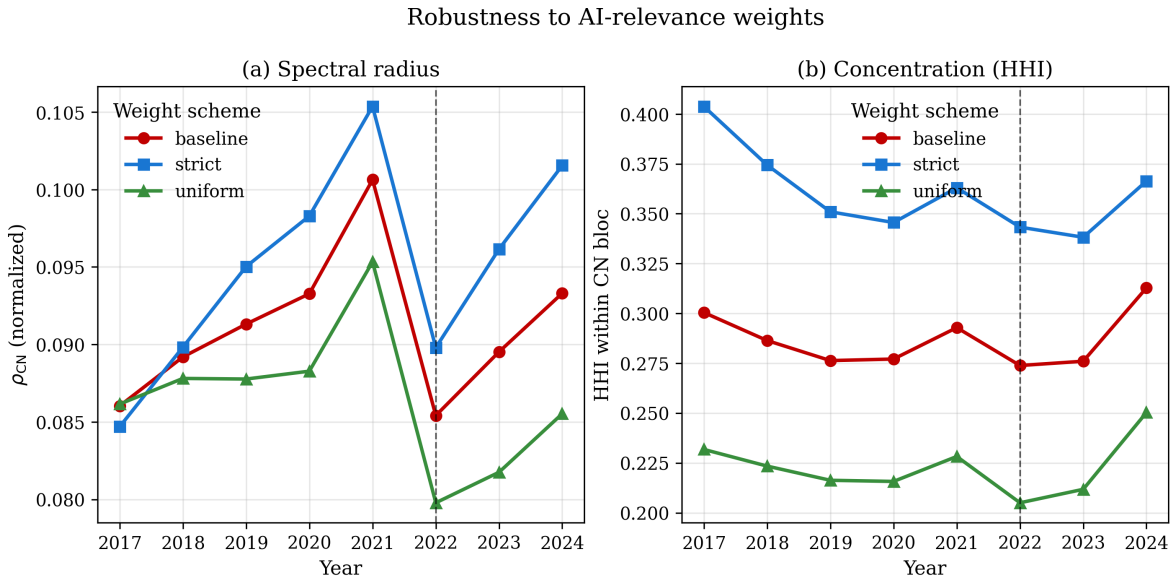


Figure 7: Robustness to AI-relevance weights ω_h. The *strict* scheme retains only HS 8542 (chips) and HS 8486 (semiconductor manufacturing equipment); *uniform* treats all four codes equally.

The strict scheme amplifies the findings: HHI rises from 0.34 (average 2017–2023) to 0.366 in 2024 versus 0.30 to 0.313 baseline, and density falls to 0.571 versus 0.643 baseline. This is the placebo direction: if the patterns were driven by general de-globalization or unrelated trade dynamics, we would expect the strict scheme to weaken the effects, not strengthen them. The strict scheme strengthening is consistent with the controls operating specifically on the chips-and-equipment subset of compute trade. The uniform scheme also preserves the qualitative pattern, with HHI rising from 0.22 (average) to 0.250 in 2024.

6.3 Sub-bloc decomposition

We decompose the China-centered bloc into a *core* (China, Russia, Hong Kong) and *intermediaries* (Vietnam, Malaysia, Thailand, Indonesia). Density and HHI for each are reported in Figure 8.

Table 9: Robustness to alternative AI-relevance weighting schemes.

Scheme	Year	ρ_{US}	ρ_{CN}	HHI_{US}	HHI_{CN}	$Edges_{US}$	$Edges_{CN}$	$Density_{CN}$
Baseline	2017	0.0186	0.0860	0.0248	0.3004	506	157	0.8626
Baseline	2018	0.0176	0.0892	0.0221	0.2864	506	159	0.8736
Baseline	2019	0.0177	0.0913	0.0237	0.2763	506	153	0.8407
Baseline	2020	0.0177	0.0933	0.0253	0.2771	506	154	0.8462
Baseline	2021	0.0183	0.1006	0.0270	0.2929	506	159	0.8736
Baseline	2022	0.0203	0.0854	0.0287	0.2739	506	142	0.7802
Baseline	2023	0.0208	0.0895	0.0274	0.2760	506	143	0.7857
Baseline	2024	0.0241	0.0933	0.0414	0.3128	506	117	0.6429
Strict	2017	0.0220	0.0847	0.0414	0.4038	503	132	0.7253
Strict	2018	0.0210	0.0898	0.0369	0.3744	505	133	0.7308
Strict	2019	0.0209	0.0950	0.0407	0.3510	505	136	0.7473
Strict	2020	0.0208	0.0983	0.0439	0.3456	504	136	0.7473
Strict	2021	0.0214	0.1054	0.0437	0.3630	505	142	0.7802
Strict	2022	0.0234	0.0898	0.0453	0.3433	506	124	0.6813
Strict	2023	0.0222	0.0962	0.0413	0.3382	504	127	0.6978
Strict	2024	0.0242	0.1016	0.0507	0.3663	506	104	0.5714
Uniform	2017	0.0159	0.0862	0.0182	0.2318	506	157	0.8626
Uniform	2018	0.0147	0.0878	0.0167	0.2234	506	159	0.8736
Uniform	2019	0.0150	0.0878	0.0174	0.2163	506	153	0.8407
Uniform	2020	0.0149	0.0883	0.0182	0.2157	506	154	0.8462
Uniform	2021	0.0152	0.0953	0.0190	0.2282	506	159	0.8736
Uniform	2022	0.0170	0.0798	0.0206	0.2050	506	142	0.7802
Uniform	2023	0.0184	0.0818	0.0232	0.2118	506	143	0.7857
Uniform	2024	0.0220	0.0855	0.0413	0.2503	506	117	0.6429

Notes: Baseline uses the preferred AI-relevance weights; strict keeps only chips and semiconductor equipment; uniform assigns equal weight to all HS categories. ρ denotes spectral radius. HHI, edges, and density are computed within the corresponding bloc.

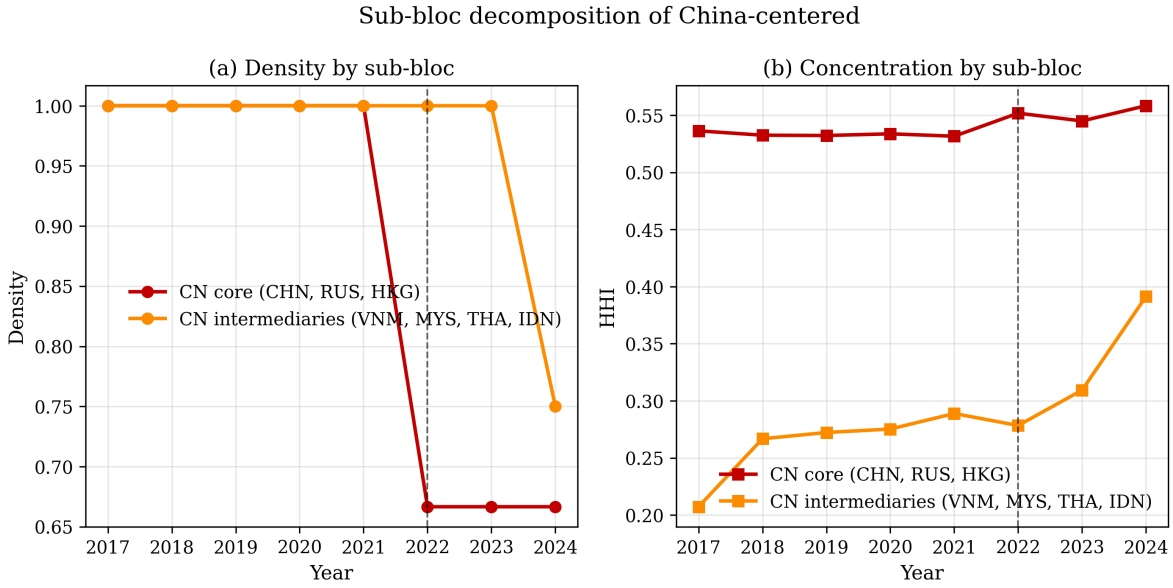


Figure 8: Sub-bloc decomposition of the China-centered subgraph. Core: China, Russia, Hong Kong. Intermediaries: Vietnam, Malaysia, Thailand, Indonesia.

The bulk of the post-2022 density loss is driven by the periphery of the China-centered bloc rather than by the core or the intermediaries. This is consistent with our interpretation: under stress, peripheral relationships die first because they involve smaller bilateral volumes that cannot amortize the increased compliance and coordination costs imposed by tightening export controls. Core relationships and the intermediaries' established trade ties survive.

6.4 Pre-2024 only

The 2024 reporting in UN Comtrade is incomplete at the time of analysis (April 2026), with several reporters not yet finalized. To ensure that the qualitative findings do not depend on 2024, we note that all three predictions of Propositions 3.1 to 3.3 are already visible by 2023: $\rho(W_t|_{S_{\text{US}}})$ rises 13.8% from 2021 to 2023; $\rho(W_t|_{S_{\text{CN}}})$ falls 11.0%; density of the China-centered bloc falls from 0.874 to 0.786 (-10.0%); HHI rises from 0.293 to 0.276 (within sampling noise) but the mean weight per surviving edge rises from \$2.59 billion to \$2.50 billion (essentially flat). The 2024 data deepen the patterns but do not create them.

7 Discussion

7.1 Hierarchical concentration as a regime

Our findings document a regime that is neither the simple decoupling predicted by the most aggressive interpretations of the controls nor the symmetric substitution that would be expected under a frictionless reorganization model. The China-centered bloc retained 96% of its pre-shock trade volume in compute-relevant goods while shedding 26% of its bilateral relationships and concentrating an additional 31% of weight on each surviving edge. We interpret this as evidence that fixed costs of partner search, regulatory compliance, and customs coordination are first-order frictions in the compute supply chain, and that under a sustained regulatory shock these frictions favor deepening of incumbent ties over formation of new ones. The result echoes the supply-chain-disruption findings of Carvalho et al. (2021) on Tohoku and Bonadio et al. (2021) on COVID-19, but in a deliberately policy-induced setting and with explicit identification of the AI-input subset of trade.

7.2 Implications for the economics of AI

The findings have three implications for the economics of AI. First, the question of *who has access to compute* is not adequately addressed by aggregate volume measures. The China-centered bloc retained nearly all its volume but at the cost of higher dependence on a smaller set of suppliers, increasing its exposure to disruption of those specific relationships. From a fragility perspective, the bloc is weaker post-shock even though it is not visibly poorer. Second, the US-aligned bloc's delayed concentration response in 2024 (HHI rising 50% in a single year) suggests that the supply-side response to multi-round controls operates with a substantial lag and may be tied to the diffusion of allied-jurisdiction adoption. Third, the structural heterogeneity between the two blocs implies that compute access for non-Chinese AI developers is conditioned on the resilience of a small number of pivotal supplier relationships—Taiwan, Korea, the Netherlands, Japan—whose own concentration trajectory deserves attention.

7.3 Implications for export-control design

For policy, our findings carry a mixed verdict. The October 2022 rule and its successors achieved a measurable structural reorganization of the China-centered compute bloc: density falls, the principal eigenvalue declines, and the bloc loses redundancy in the formal network sense. By that criterion, the controls succeeded. By the volume criterion, however, they did not: the China-centered bloc retained 96% of its compute-relevant trade. Whether reduced structural redundancy with retained

volume is an acceptable outcome depends on the policymaker's loss function. A policymaker primarily concerned with cutting the flow of frontier compute would view the result as incomplete; a policymaker primarily concerned with making the bloc more fragile to future targeted measures would view it as a success.

7.4 Limitations and scope conditions

Four caveats apply. First, HS-4 trade data do not distinguish frontier from legacy semiconductors; HS-6 disaggregation may help but does not fully resolve this. Second, the annual frequency attenuates within-year dynamics, and quarterly Comtrade data are sparse for many countries before 2019. Third, identification of the causal effect of specific BIS rounds requires assumptions on the timing of strategic anticipation and on the absence of confounding contemporaneous shocks (notably the COVID-19 disruption tail and the Russia-Ukraine sanctions of 2022). We have done what is feasible with the available data; richer identification will require firm-level or product-level data not yet publicly available. Fourth, the hierarchical-concentration mechanism we propose is a reduced-form characterization of the data; deriving it from a structural model with explicit fixed costs of partner search is left for future work.

8 Conclusion

We have documented a hierarchical-concentration response of the global AI compute network to four rounds of US export controls between 2017 and 2024. The targeted bloc reorganized toward a smaller set of deeper trading relationships rather than fragmenting or substituting symmetrically. The result, robust to Hong Kong treatment, alternative weighting schemes, and sub-bloc decomposition, suggests that the costs of decoupling are absorbed primarily at the extensive margin, with intensive-margin volume preserved. Three implications follow: existing volume-based measures of supply-chain fragility miss the structural reorganization documented here; the success of export controls depends sensitively on the policymaker's loss function; and AI compute access for non-Chinese developers may also have become structurally more fragile as the US-aligned bloc consolidates with a two-year lag. We see three productive directions for future work. First, integrating these network measurements with firm-level adoption data would allow direct estimation of the cost-of-compute implications for AI developers. Second, extending the framework to quarterly frequency once Comtrade data become more complete would enable better identification of round-specific effects. Third, embedding hierarchical concentration in a structural model of partner search under fixed costs would clarify the welfare implications.

References

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608.

Acemoglu, D. and Restrepo, P. (2022). Tasks, automation, and the rise in U.S. wage inequality. *Econometrica*, 90(5):1973–2016.

Aghion, P., Jones, B. F., and Jones, C. I. (2019). Artificial intelligence and economic growth. *NBER Chapters: The Economics of Artificial Intelligence*, pages 237–282.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). Economic policy for artificial intelligence. *Innovation Policy and the Economy*, 19:139–159.

Agrawal, A., Gans, J., and Goldfarb, A. (2022). *Power and Prediction: The Disruptive Economics of Artificial Intelligence*. Harvard Business Review Press.

Allen, G. C. (2022). Choking off China's access to the future of AI. Report, Center for Strategic and International Studies (CSIS).

Anderson, J. E. and van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192.

Antràs, P. (2021). De-globalisation? global value chains in the post-COVID-19 age. *NBER Reporter*, (3).

Babina, T., Fedyk, A., He, A., and Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, 151:103745.

Bak, P., Tang, C., and Wiesenfeld, K. (1987). Self-organized criticality: An explanation of the 1/f noise. *Physical Review Letters*, 59(4):381–384.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica*, 87(4):1155–1203.

Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286(5439):509–512.

Bonadio, B., Huo, Z., Levchenko, A. A., and Pandalai-Nayar, N. (2021). Global supply chains in the pandemic. *Journal of International Economics*, 133:103534.

Bown, C. P. (2021). The US–China trade war and Phase One agreement. Working paper 21-2, Peterson Institute for International Economics.

Brynjolfsson, E., Rock, D., and Syverson, C. (2021). The productivity J-curve: How intangibles complement general purpose technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372.

Caliendo, L. and Parro, F. (2023). Lessons from US–China trade relations. *Annual Review of Economics*, 15:513–547.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great East Japan earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Costinot, A. and Rodríguez-Clare, A. (2014). Trade theory with numbers: Quantifying the consequences of globalization. *Handbook of International Economics*, 4:197–261.

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., and Owen, D. (2024). The rising costs of training frontier AI models. Research report, Epoch AI.

Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10):3115–3153.

Evenett, S., Jakubik, A., Martín, F., and Ruta, M. (2024). The return of industrial policy in data. *The World Economy*, 47(7):2762–2788.

Forman, R. (2003). Bochner's method for cell complexes and combinatorial Ricci curvature. *Discrete & Computational Geometry*, 29:323–374.

Gai, P. and Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society A*, 466(2120):2401–2423.

Goldfarb, A. and Trefler, D. (2019). Artificial intelligence and international trade. *NBER Chapters: The Economics of Artificial Intelligence*, pages 463–492.

Hidalgo, C. A., Klinger, B., Barabási, A.-L., and Hausmann, R. (2007). The product space conditions the development of nations. *Science*, 317(5837):482–487.

Khan, S. M., Mann, A., and Peterson, D. (2021). The semiconductor supply chain: Assessing national competitiveness. Issue brief, Center for Security and Emerging Technology, Georgetown University.

Korinek, A. and Stiglitz, J. E. (2018). Artificial intelligence and its implications for income distribution and unemployment. *NBER Chapters: The Economics of Artificial Intelligence*, pages 349–390.

Korniyenko, Y., Pinat, M., and Dew, B. (2017). Assessing the fragility of global trade: The impact of localized supply shocks using network analysis. *IMF Working Paper*, WP/17/30.

Mazzucato, M. and Rodrik, D. (2026). Industrial policy with conditionalities: a taxonomy and sample cases. *Industrial and Corporate Change*, page dtaf063.

Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Review*, 45(2):167–256.

Ollivier, Y. (2009). Ricci curvature of Markov chains on metric spaces. *Journal of Functional Analysis*, 256(3):810–864.

Samal, A., Sreejith, R. P., Gu, J., Liu, S., Saucan, E., and Jost, J. (2018). Comparative analysis of two discretizations of Ricci curvature for complex networks. *Scientific Reports*, 8:8650.

Sandhu, R., Georgiou, T., and Tannenbaum, A. (2015). Graph curvature for differentiating cancer networks. *Scientific Reports*, 5:12323.

Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M., and Villalobos, P. (2022). Compute trends across three eras of machine learning. In *2022 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE.

Sreejith, R. P., Mohanraj, K., Jost, J., Saucan, E., and Samal, A. (2016). Forman curvature for complex networks. *Journal of Statistical Mechanics: Theory and Experiment*, page 063206.

Thadani, A. and Allen, G. C. (2023). Mapping the semiconductor supply chain: The critical role of the Indo-Pacific region. Report, Center for Strategic and International Studies (CSIS).

Tucker, C. (2019). Privacy, algorithms, and artificial intelligence. *NBER Chapters: The Economics of Artificial Intelligence*, pages 423–438.

U.S.-China Economic and Security Review Commission (2022). Hearing on U.S.-China Competition in Global Supply Chains: Hearing Before the U.S.-China Economic and Security Review Commission, One Hundred Seventeenth Congress, Second Session. Printed for use of the U.S.-China Economic and Security Review Commission.

Vallarino, D. (2025). Augmenting trade complexity analysis with deep learning: An AI-based framework for small open economies. *Applied Economics Letters*.

Vallarino, D. (2026a). Identification and inference in nonlinear dynamic network models. *arXiv preprint arXiv:submit/7441433*.

Vallarino, D. (2026b). Sandpile economics: Theory, identification, and evidence. *arXiv preprint arXiv:2604.13890*.

Vallarino, D. (2026c). Stochastic network survival dynamics: A nonlinear evolution problem on economic graphs. *Communications in Nonlinear Science and Numerical Simulation*, 159:109904.

Weber, M., Saucan, E., and Jost, J. (2017). Characterizing complex networks with Forman–Ricci curvature and associated geometric flows. *Journal of Complex Networks*, 5(4):527–550.

A Mean Forman–Ricci curvature

The body of the paper reports median Forman–Ricci curvature. Figure 9 reports the mean for completeness. The 2024 spike in mean curvature reflects right-tail outliers driven by incomplete UN Comtrade reporting at the time of analysis (April 2026), not changes in the typical edge.

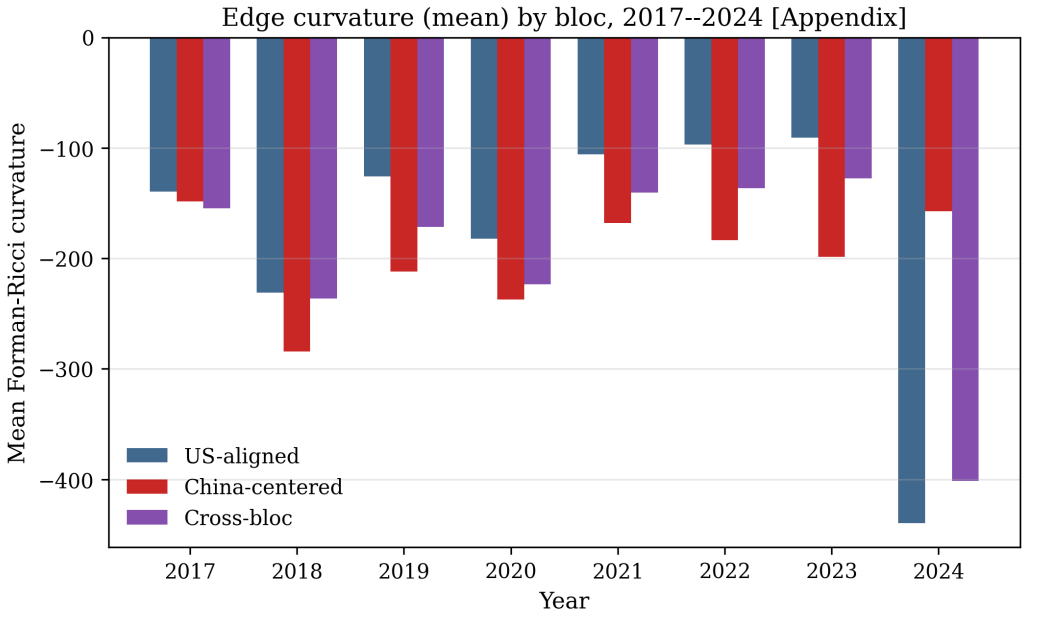


Figure 9: Mean Forman–Ricci edge curvature by bloc, 2017–2024.

B Data construction details

The 48-country panel comprises: United States, China, Japan, South Korea, Taiwan, Germany, Netherlands, Singapore, Malaysia, Vietnam, Mexico, Thailand, Philippines, United Kingdom, France, Italy, Ireland, Israel, Canada, India, Brazil, Hungary, Czech Republic, Poland, Indonesia, Switzerland, Belgium, Spain, Australia, Russia, Turkey, Saudi Arabia, United Arab Emirates, South Africa, Sweden, Austria, Denmark, Finland, Norway, Hong Kong SAR, Argentina, Chile, Colombia, Egypt, Pakistan, Bangladesh, Romania, Portugal. Selection was based on 2017 trade volumes in the four target HS codes; the panel covers approximately 96% of global compute-relevant trade.

Reporter-mirror reconciliation uses the maximum of exporter-reported and importer-reported values. Period values are normalized to integer years prior to filtering. The full pipeline (data download, network construction, geometric and spectral computation, robustness analysis, and figure/table generation) is implemented in Python and is available from the author on request.

C Connection to the Sandpile Economics framework

The hierarchical-concentration regime documented here connects to the broader research program on geometric and spectral characterizations of fragility in economic networks developed in Vallarino (2026b). Within that framework, the extensive-margin contraction with intensive-margin concentration we document here corresponds to a movement of the network toward a self-organized critical configuration in which the distribution of edge weights becomes increasingly heavy-tailed under stress. The present paper extends the framework to a setting of deliberate policy intervention rather than spontaneous dynamics, and it provides a specific empirical signature—joint density decline and HHI increase—that can be searched for in other supply chains under analogous shocks. Companion econometric work (Vallarino, 2026a,c) develops related identification and inference tools for nonlinear network dynamics, and Vallarino (2025) applies network methods to a related setting of trade complexity in small open economies.