The Porous Decoupling:
Vertical Rivalry and Swing-State Leakage

Hyungjin Kim *

April 18, 2026

Abstract

Headline trade statistics show the United States decoupling from China after the 2018 Section-301 tariffs; value-added trade statistics show the opposite. Chinese value added continues to reach the US indirectly through a small set of transit economies — Vietnam, Mexico, and Malaysia chief among them — at roughly $22 billion per year in critical sectors alone. We explain this porosity as the steady-state outcome of an alignment bargain between the US and its transit partners. The mechanism is technological: in critical sectors, non-Chinese intermediates are poor substitutes for Chinese ones. A transit country can formally accept US restrictions yet keep sourcing Chinese inputs and re-exporting their content embedded in finished goods, because the cost of doing otherwise is prohibitive. The theory predicts a sharp cross-sectional pattern: in sectors where US tariffs on direct Chinese imports are higher, transit-country exports to the US carry a disproportionately higher share of embedded Chinese content than exports from non-transit origins. We identify this pattern in two inter-country input–output databases (OECD and ADB), robust to standard inference corrections. Placebo-year tests recover most of the post-2018 magnitude in pre-2018 windows — exactly what a steady-state pattern should show, since the finding reflects the global production structure rather than a 2018 policy break. A complementary theoretical result — the *Paradox of Competitiveness* — shows that as the upstream rival becomes more efficient, the subsidy required to buy voluntary alignment turns fiscally prohibitive, and the US optimally switches from subsidy to coercion. Customs-based decoupling metrics therefore systematically overstate real decoupling in exactly the sectors where decoupling is most sought.

Keywords: Geoeconomics, Vertical Supply Chains, Coercion, US-China Rivalry, Entity List.
JEL Classification: F13, F51, L14, O25.

*Korea Development Institute, hyungjinkim@kdi.re.kr.

1 Introduction

Between 2018 and 2024, the United States imposed Section-301 tariffs averaging roughly twenty-five percentage points on Chinese imports concentrated in critical sectors — semiconductors, precision machinery, batteries, and electric-vehicle components. Direct US imports from China in those sectors fell sharply, as policy intended. Over the same window, imports from a small set of transit economies—chief among them Vietnam, Mexico, and Malaysia—rose by a comparable amount. The headline change masks a deeper pattern that emerges when we trace the *origin* of the value embodied in these transit-country exports. The input–output accounting literature calls this the *foreign value-added* (FVA) share. Chinese content has continued to reach the US indirectly throughout the Section-301 period, settling near 6% of transit-country critical-sector exports by 2020–2022. That share translates to roughly $22 billion per year in critical sectors alone.[1] On customs data the US is decoupling; on value-added data it is recoupling through intermediaries. This paper asks why the two measures disagree and what the discrepancy implies for economic statecraft. We answer with a model of an alignment bargain between the US and a transit partner and with a single cross-sectional prediction that the value-added data identify directly.

We argue that the discrepancy is not a measurement artifact or a transitional friction. It is the steady-state outcome of a bargain between the US (the downstream *hegemon* of the supply chain) and a transit country (a *swing state*), undone by the very production technology that makes the bargain rational for both parties. The mechanism rests on complementarity: in critical sectors, non-Chinese intermediates are poor substitutes for Chinese ones. Boehm et al. (2019) and Atalay (2017) document this empirically. A swing state cannot cheaply rebalance its input mix no matter how generously the US pays to align. The unpayable part of the bargain is therefore absorbed by continued sourcing of Chinese inputs, which are then re-exported in finished goods. The same complementarity yields a directly testable cross-sectional prediction. In sectors where US tariffs on direct Chinese imports are higher,

[1] Derivation from fitted FVA shares applied to 2022 swing-state-to-US gross exports is given in Section 5.

the Chinese-content share in swing-state exports to the US should rise more steeply than in exports from non-swing-state origins—a steeper tariff gradient for swing states.

We test this prediction in two independently maintained inter-country input–output (ICIO) databases: the OECD ICIO Extended release and the ADB Multi-Regional Input–Output (MRIO) 72-economy release. In a triple-difference specification, we interact published Section-301 tariff rates with swing-state identity. The estimate is positive, statistically robust, and economically meaningful. A 25-percentage-point tariff increase on a swing-state-produced good raises its embedded Chinese-content share by roughly 2.5 percentage points, over and above what industry-by-year and source-by-year fixed effects absorb. The result survives two-way clustering on source and industry, leave-one-out on the swing-state set, and Honest-DiD bounds for parallel-trends violations. Placebo tests recover between 47% and 87% of the coefficient in pre-2018 windows, with the earliest 2010 placebo recovering 87%. We read this not as a concern but as the signature of a steady-state prediction: the pattern is a standing feature of the global production structure, not a quasi-experimental response to the 2018 shock. The policy implication is that customs-based surveillance systematically overstates real decoupling, by an economically meaningful margin, in exactly the sectors where decoupling is most sought.

This measurement gap is the signature of a broader historical shift. When the US sought to align allies in the post-war order, it provided public goods, market access, and subsidies — a generous-empire regime of positive inducements. Between 2018 and 2024, the policy toolkit has pivoted toward negative coercion: Entity Lists, export controls, and the weaponization of market access (Facts 4–6 in Section 3). Swing states caught between the two regimes rely on China for upstream critical inputs while selling finished goods into the US market. Chief among these are South Korea, Taiwan, Vietnam, Malaysia, and Mexico. We term this configuration the *Vertical Trap*. Our model is a three-stage game over this triangular structure (Upstream Challenger, Swing State, Downstream Incumbent). The Incumbent's strategic problem is that eliminating the kill-switch risk requires the Swing State to abandon its most efficient input supplier. This deviation is costly, and the Incumbent must either subsidize or coerce it.

The model delivers three theoretical results. First, the *Regime Switch* (Proposition 1): the Incumbent's optimal alignment policy flips discontinuously as the Challenger's productivity θ crosses a critical threshold $\bar{\theta}$. Below $\bar{\theta}$, the regime is Generous Empire ($\lambda^* = 0$); above it, Coercive Hegemon ($\lambda^* > 0$). The equilibrium map $\lambda^*(\theta)$ is continuous but non-differentiable at the kink. Second, the *Paradox of Competitiveness* (Proposition 2): the threshold $\bar{\theta}$ locates the regime switch in the rival's efficiency advantage. When the Challenger is technologically immature, subsidy-based alignment is fiscally affordable; as the Challenger approaches the technological frontier, the subsidy required for voluntary alignment becomes prohibitive, and the Incumbent optimally switches to coercion. The paradox is that the more efficient the rival becomes, the more coercive the incumbent must be to maintain its security architecture. Third, the *Leakage Asymmetry* (Propositions 4 and 5′): the bottleneck primitive $\sigma < 1$ predicts the cross-sectional moment identified in the data above—the tariff-by-swing-state cross-partial on Chinese value-added in swing-state $\rightarrow$ US exports. The proposition locates this moment as a steady-state feature of the production structure rather than a 2018 policy break.

Contribution. The paper sits at five intersections; Section 2 positions each in detail. In compact form: to the *geoeconomics of fragmentation*, we endogenize the hegemon's choice between inducement and coercion as a function of the rival's efficiency on bottleneck inputs. This extends Clayton et al. (2024)'s dyadic framework to a triangular (U–S–C) structure. In this structure, a *production-structure primitive* ($\sigma_c < 1$) partially voids the alignment contract. To *common-agency and exclusion contracting*, we show that bottleneck complementarity breaks the truthful-equilibrium benchmark of Bernheim and Whinston (1986) and makes exclusion the rational second-best. To *production networks and supply-chain IO*, we reinterpret a low-substitutability node not as a fragility externality but as a locus of infinite implementation cost for competing principals. To *value-added trade measurement*, we name the cross-sectional moment most cleanly identified by Leontief-inverse decomposition — the tariff-by-swing-state cross-partial — and test it directly. To the *political economy of policy radicalization*, we model the post-2018 US shift as a regime transition (the kinked alignment correspondence of Proposition 1) triggered by rival technology rather than a preference shock.

Scope and what this paper does not do. We take the Challenger's technological rise θ as exogenous and do not model its endogenous accumulation. We treat the security externality λx as reduced-form rather than deriving it from a probabilistic "kill-switch" microfoundation. And we abstract from dynamic reputation or multi-period learning. We microfound the fiscal capacity $\bar{T}$ and the securitization cost γ in the appendix but take them as primitives in the main model. These limitations are deliberate. The paper isolates a single new mechanism: the endogenous corner solution that emerges when bottleneck complementarity meets a binding fiscal constraint. Dynamic, endogenous-technology, and probabilistic-security extensions are left for future work. Having previewed the argument and the three theoretical results, we now locate the contribution in the adjacent literatures that frame it before turning to the empirical setting.

2 Literature Review

Our paper draws on International Trade, Industrial Organization, and Political Economy, and contributes to five distinct strands — one per paragraph below.

First, we contribute to the emerging literature on Geoeconomics and the weaponization of supply chains. The intellectual ancestor of this line is Hirschman (1945), whose analysis of foreign trade as an instrument of national power established the template that modern weaponization scholarship has inherited. The broader international-political-economy tradition on hegemonic ordering and systemic responsibility (Gilpin, 1981) provides the historical framing within which our model's subsidy-vs-coercion regime switch can be read. Farrell and Newman (2019) revived this tradition by introducing "Weaponized Interdependence," arguing that asymmetric hub-and-spoke networks allow central nodes to coerce peripheral actors. Farrell and Newman (2023) and the essays in Drezner et al. (2021) extend this to financial and digital infrastructures. Clayton et al. (2024) formalize this in general equilibrium by modeling a "Hegemon's Dilemma" between coercive power and economic integration. Their dyadic framework is the most general existing treatment of geoeconomic bargaining. We build on their insight that coercion has an endogenous cost. We specialize the mechanism to a triangular (U–S–C) supply chain, in which a production-structure primitive ($\sigma_c < 1$) generates leakage that partially voids the Hegemon's contract. Policy-oriented complements cover shifting trade linkages under geopolitical tension (Gopinath et al., 2024), the erosion of multilateral trade rules (Mattoo et al., 2024), and the costs of geoeconomic fragmentation (Aiyar et al., 2023). Sanctions economics (e.g., Mulder, 2022; Itskhoki and Mukhin, 2025) characterizes the macro consequences of coercion once deployed. These models typically treat the hegemon's coercive capacity as a static endowment or reduced-form shock. Our contribution is to endogenize the *decision* to coerce. Lim (2015) classifies coercion strategies (sticky vs. slippery power), but does not microfound when the hegemon switches between them. We provide that microfoundation: the switch from positive inducements (aid/subsidies) to negative coercion (sanctions) occurs precisely when the rival becomes efficient on a bottleneck input.

Second, we differentiate our work from the horizontal strategic-trade-policy tradition (Brander and Spencer, 1985; Eaton and Grossman, 1986), in which nations compete to sell substitute goods and subsidy wars are the dominant conflict. We build instead on the vertical industrial-organization and global-value-chain literature (Yi, 2003; Antràs, 2003; Grossman and Helpman, 2002; Antràs and Chor, 2013; Alfaro et al., 2019; Antràs, 2016). We also draw on the quantitative trade and China-shock empirics on tariff incidence and supply-chain reallocation (Autor et al., 2013; Caliendo and Parro, 2015; Fajgelbaum et al., 2020; Amiti et al., 2019; Flaaen et al., 2020; Grossman et al., 2024; Alfaro and Chor, 2023). Our "Bottleneck" assumption is grounded in network-IO work showing that cascading failures and coercion are most effective at low-substitutability nodes (Acemoglu et al., 2012; Elliott et al., 2022; Baqaee and Farhi, 2019, 2024; Carvalho et al., 2021; Boehm et al., 2019). It is further grounded in estimated elasticity ranges from Broda and Weinstein (2006) and Caliendo and Parro (2015). We reinterpret these nodes not merely as fragility points but as loci of *infinite implementation cost* for competing principals. Relative to the industrial-policy revival (Juhász, 2018; Liu, 2019; Lashkaripour and Lugovskyy, 2023; Juhász et al., 2024; Bown, 2023, 2024) and the sector-specific semiconductor supply-chain literature (Khan et al., 2021), we provide the complementary theoretical mechanism that rationalizes why coercion (rather than subsidy alone) becomes optimal beyond a critical efficiency threshold.

Third, methodologically, we model the US–China rivalry as a Common Agency Game in which two principals bid for an agent (the Swing State). Bernheim and Whinston (1986) and Dixit (1996) established the "Menu Auction" benchmark in which truthful equilibria maximize joint surplus; Grossman and Hart (1983) supplies the underlying principal-agent structure. Standard common-agency models assume linear transfer technology (cash is a perfect substitute for the agent's action). We introduce a non-linear friction: the technological bottleneck. Following Martimort (1996), Segal (1999), and McAfee and Schwartz (1994), we show that when one principal controls an essential complement ($\sigma < 1$) the rival cannot compete via linear transfers and must resort to breaking the interaction—exclusion. This mirrors the exclusive-contracting and foreclosure literature (Rasmusen et al., 1991; Segal and Whinston, 2000b,a; Jehiel et al., 1996; Rey and Tirole, 2007), applied here to sovereign states. Our treatment of (R, λ) as an announced contract also connects to incomplete-contracting and authority frameworks (Aghion and Tirole, 1997; Maskin and Tirole, 1999), and to the theory of trade agreements as commitment devices (Bagwell and Staiger, 1999; Maggi and Rodríguez-Clare, 1998; Antràs and Staiger, 2012).

Fourth, our empirical strategy builds on the *value-added trade* measurement literature. The Koopman–Wang–Wei decomposition (Koopman et al., 2014) and the bilateral value-added exports framework of Johnson and Noguera (2012) provide the accounting identities underlying our FVA-China share variable. Wang et al. (2013) extend the KWW decomposition to bilateral-sector levels, which is the exact accounting object entering our panel. Borin and Mancini (2019) map the competing decomposition conventions to the economic questions they are best suited to address, and we follow their guidance in selecting the forward-looking FVA variant for the leakage test. The upstreamness/downstreamness measures of Antràs and Chor (2018) formalize the positional vocabulary (upstream intermediate, downstream final) that our triangular U–S–C structure invokes. These measures also provide the accounting basis for distinguishing direct from indirect Chinese value-added embodied in swing-state exports. We use this methodology as implemented in the OECD Trade in Value Added (TiVA) and ADB MRIO databases. On the empirical-rerouting side, our paper complements a growing body of descriptive evidence documenting trade diversion through third countries following the 2018 tariffs. Mattoo et al. (2023) show that roughly one-third of the observed decline in US–China bilateral trade reflects rerouting through Southeast Asian and Latin American intermediaries rather than genuine supply-chain restructuring. Alfaro and Chor (2023) document the "great reallocation" of global value chain (GVC) activity from China to ASEAN economies. Schulze and Xin (2025) distinguish trade *reallocation* (domestic substitution for Chinese exports) from trade *rerouting* (transshipment) in six Asian connector countries. Our FVA cross-partial is designed to measure this distinction directly. Ahn and Tan (2025) quantify the welfare gains from supply-chain diversification in a multi-country IO model, a complementary framing to our triangular alignment-contract approach. We go beyond description: the alignment-contract model provides a structural explanation for *why* rerouting concentrates in critical sectors (Proposition 4) and *why* it is steeper for swing-state exporters with higher baseline Chinese-intermediate dependence (Proposition 5′).

Finally, we contribute to the political economy of policy radicalization. Most models of trade war treat tariffs as the outcome of distributional conflict (Grossman and Helpman, 1994) or bargaining failure (Bagwell and Staiger, 1999; Maggi and Rodríguez-Clare, 1998; Handley and Limão, 2017), and a growing empirical literature links trade shocks to political backlash and the erosion of centrist coalitions (Autor et al., 2020; Rodrik, 2018). We instead model the shift to coercion as a regime transition. A continuous change in rival technology (θ) triggers a qualitative switch in the structure of the optimal policy (λ)—from a corner solution at $\lambda^* = 0$ to an interior solution at $\lambda^* > 0$. This switch is orthogonal to (but may be reinforced by) the distributional channel. The equilibrium map $\lambda^*(\theta)$ is continuous but non-differentiable at the threshold (a kink, not a jump). Empirically, this is consistent with the structural break in US trade and industrial policy documented after 2018 (Fajgelbaum et al., 2020; Amiti et al., 2019; Flaaen and Pierce, 2024; Bown, 2024; Alfaro and Chor, 2023). We show that the "Hawkish" turn in US policy is not a preference shock but a structural necessity—a corner solution that emerges when the interior (subsidy-based) solution becomes fiscally unbounded. These five literatures frame the contribution; the six stylized facts that follow discipline the model's primitives and motivate the theoretical setup in Section 4.

3 Empirical Patterns

Six stylized facts motivate our theoretical framework and characterize the evolving structure of the global semiconductor and high-tech supply chain. Facts 1–3 (Section 3.1) document trade-flow patterns; Facts 4–6 (Section 3.2) document the US policy response. Together these patterns highlight the structural asymmetry between upstream reliance on Chinese inputs and downstream reliance on the US market, creating what we term a "Vertical Trap" for third-party countries (Swing States). Section 5 returns to the data after the model and formally tests the core claim that leakage is a steady-state feature of the production structure rather than a 2018 policy break.

3.1 Trade Patterns: Vertical Dependence and Indirect Linking

The first set of facts establishes the economic constraints facing Swing States. We analyze aggregate trade flows for non-US, non-China economies using UN Comtrade data from 2000

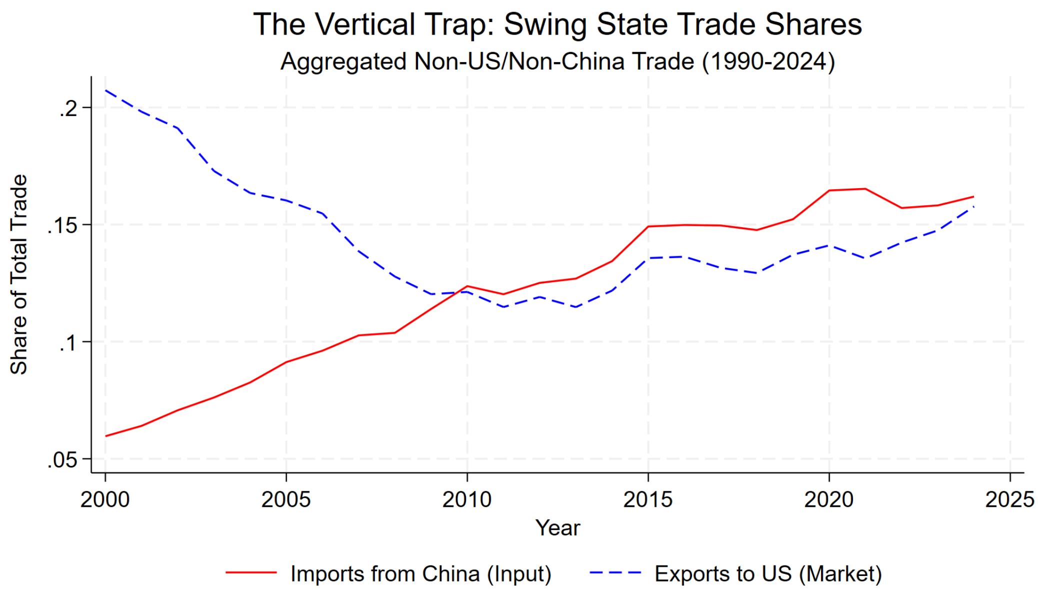

Figure 1: **The Vertical Trap.** Aggregated trade shares of Swing States (Non-US/Non-China). The reliance on Chinese inputs (Red) has structurally surpassed the reliance on the US export market (Blue), creating an upstream bottleneck.

to 2024.

Fact 1: The Vertical Trap. Contrary to the narrative of global fragmentation, Swing States have become structurally more integrated with the Chinese supply chain. Figure 1 plots the aggregate trade shares of Swing States. The red line shows the share of imports sourced from China (upstream dependence); the blue line shows the share of exports destined for the US (downstream dependence). Crucially, we observe a "crossover" around 2010. While the US remains a critical export market (absorbing ∼15% of exports), the reliance on Chinese inputs has tripled, rising from 6% to over 16%. This confirms the model's assumption that China dominates the upstream bottleneck ($\sigma < 1$), anchoring Swing States to the "Red Supply Chain" despite geopolitical pressure.[2]

[2] The same pattern appears in IMF Direction of Trade Statistics (DOTS) bilateral trade shares. It is also documented in the trade-reallocation evidence of International Monetary Fund (2023, Chapter 4).

Fact 2: Triangle Trade and Indirect Linking. How do Swing States reconcile this upstream dependence with US pressure to decouple? The data suggest a mechanism of "Triangle Trade," where Swing States act as processing hubs. Figure 2(a) reveals a striking symmetry. As the Swing States' aggregate trade surplus with the US expands to $800 billion, their trade deficit with China deepens proportionally to $1 trillion. This mirror image implies a pass-through structure: Chinese inputs are imported, processed, and re-exported to the US.

Fact 3: The Decoupling Divergence. This structural integration limits the efficacy of US bilateral tariffs. Figure 2(b) compares the import dependence of the US versus Swing States. Following the 2018 Trade War (dotted line), US direct reliance on China (Solid Blue) dropped significantly. However, Swing State reliance on China (Dashed Red) did not decline; it stabilized and remained high. This divergence provides empirical support for our "Leakage" hypothesis: coercive policies successfully severed direct US-China links but failed to break the indirect links embedded in the Swing State supply chain.

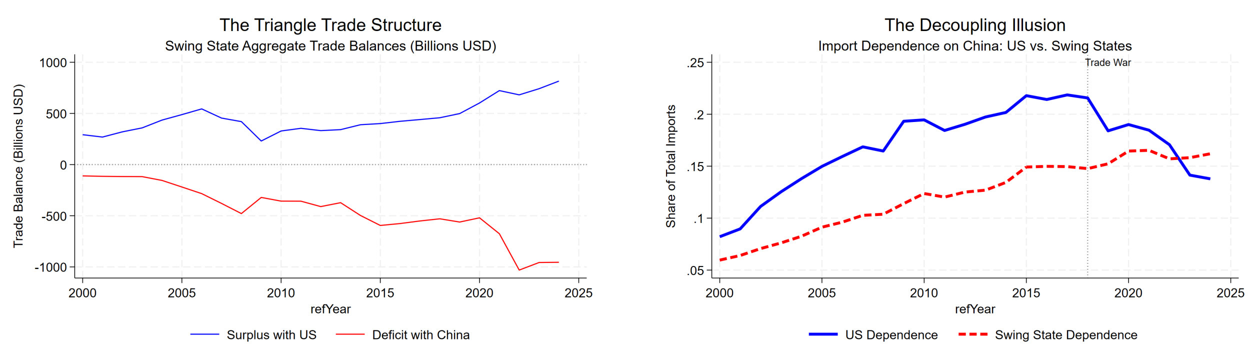

(a) **The Triangle Trade.** Symmetrical deepening of the US surplus and China deficit.

(b) **The Decoupling Illusion.** Swing State reliance on China resists US policy shocks.

Figure 2: **Mechanisms of Evasion.** Panel (a) shows the flow of value from China to the US via Swing States. Panel (b) shows that while the US decoupled directly, the Swing States did not, rendering the decoupling porous.

3.2 Policy Responses: The Expansion of the Coercive Frontier

In response to the "leakage" observed in trade flows, US policy has undergone a structural transformation. We analyze the evolution of the US Bureau of Industry and Security (BIS) Entity List to document the shift from bilateral trade disputes to a systemic network blockade.

Fact 4: From Bilateral to Network Coercion. A salient stylized fact is that the coercive perimeter has been extra-bilateral for at least fifteen years. The stock of third-country listings has exceeded the stock of direct China listings in every year since the mid-2000s, with ratios ranging from roughly 5:1 in 2010 to 9:1 in 2017. Figure 3 plots the cumulative count of active entities on the Entity List on a log scale, which makes growth rates rather than absolute levels readable. Two features of the plot matter. First, the Third-country line (solid blue) has a roughly constant log-slope throughout the sample—its compound growth rate is $\approx 23\%$ per year in 2010–2017 and $\approx 19\%$ per year in 2018–2024, essentially unchanged by Section-301. Second, the China line (dashed red) is nearly flat from 2005 to 2017 ($\approx 13\%$ per year) and then kinks sharply upward after 2018 ($\approx 41\%$ per year). Post-2018 the bilateral stock converges toward the third-country stock, narrowing the level ratio to roughly 3:1 by 2024. The *infrastructure* for third-country coercion was therefore already in place before the tariff shock; what 2018 triggered was a bilateral intensification *on top of* a pre-existing extra-bilateral coercive machine. Both patterns match the model. Once the Rival (China) dominates the upstream bottleneck, the Hegemon (US) maintains and escalates coercion both against the Rival directly and against the global intermediary network, including nominal Swing States. The third-country machinery is the steady state, not a 2018 innovation. We do not claim a structural break test; the description is observational.

Facts 5 & 6: Precision targeting and evasion resistance. Two further qualitative shifts are consistent with the coercion mechanism, though they do not identify it. First, the Entity List's sectoral focus has narrowed toward high-technology targets: the share of new entries containing "Semiconductor," "AI," or "Quantum" rose from under 10% in the early

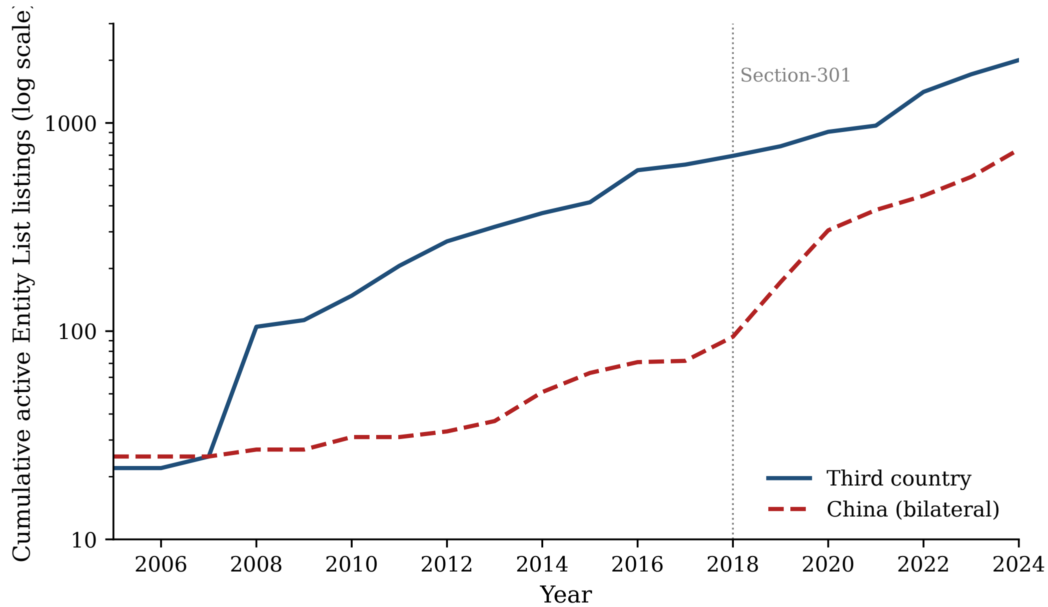

Figure 3: **The Network War.** Cumulative count of active US Entity List listings on a log scale, split by country of the listed entity (third-country vs. China-bilateral). The third-country line has a roughly constant log-slope throughout, indicating steady compound growth ($\approx$ 23% p.a. 2010–2017 and $\approx$ 19% p.a. 2018–2024); the China line is nearly flat pre-2018 ($\approx$ 13% p.a.) and kinks upward post-2018 ($\approx$ 41% p.a.). The extra-bilateral coercive infrastructure preceded Section-301; 2018 added a bilateral intensification on top of it.

2010s to over 20% by 2024. Second, administrative effort per listing has risen. The average number of aliases per sanctioned entity spiked to nearly 2.7 in 2023, a rough proxy for the "Cat-and-Mouse" dynamic that Section 4 formalizes via the efficiency gap ΔC. We report the underlying series in Appendix Figure 11; these patterns motivate but are not load-bearing for the formal analysis.

Facts 1–6 document aggregate co-movements consistent with the leakage mechanism but do not test it. They motivate three structural primitives that Section 4 formalizes. First, bottleneck complementarity ($\sigma_c < 1$) between upstream inputs in critical sectors (Facts 1–2). Second, a strategic alignment margin on which swing states balance Chinese input efficiency against US market access (Fact 3). Third, an endogenous coercive response by the hegemon once the alignment bargain turns fiscally unbounded (Facts 4–6). Section 4 builds a three-stage model of that alignment bargain and derives two testable predictions—the flow-margin leakage of Proposition 4 and the steady-state cross-partial of Proposition 5′. Section 5 then returns to the data with regression tests of both predictions.

Mapping the model's doctrine λ^* to the data. The security doctrine $\lambda^*(\theta)$ in the model is a single scalar summarizing the Incumbent's endogenous coercive posture against the Rival's supply chain. Empirically, this doctrine is multi-valued. It is applied both against the Rival directly (China-bilateral Entity-List listings, Section-301 tariffs) and against the global intermediary network that transits Rival inputs (third-country listings, secondary sanctions). Fact 4's two lines correspond to two empirical margins of the same theoretical object. The China-bilateral line (kink in 2018) measures the *marginal* extension of the doctrine to the direct Rival; this is the margin most closely tracked by the model's Regime-Switch prediction of a continuous-but-kinked $\lambda^*(\theta)$. The third-country line (constant log-slope) measures the *stock* of already-deployed extra-bilateral coercive capacity. The model treats this stock as a state variable that feeds into the Stage-2 primitives γ and $\bar{T}$, rather than as the endogenous λ^* itself. The constant pre- and post-2018 log-slope of third-country listings is therefore consistent with the model. The pre-existing extra-bilateral machine reflects γ-governed fiscal capacity that expanded steadily, regardless of any single-year technology shock. The 2018 kink in the China-bilateral line is the Regime-Switch signature: $\lambda^*(\theta)$ crosses the threshold $\bar{\theta}$ once the Section-301 reading of θ surpasses the Generous-Empire corner. Both patterns are joint predictions of the model; the reduced-form tests in Section 5 identify a third empirical moment — the cross-sectional tariff-×-swing cross-partial on FVA shares — that neither line in Fact 4 can identify alone.

4 Model

We develop a three-stage vertical game solved by backward induction. Section 4.1 fixes the environment. Sections 4.2–4.4 derive each stage's best response. Section 4.5 then establishes the *Paradox of Competitiveness*: an exogenous rise in upstream efficiency triggers downstream coercion in equilibrium. Section 4.6 quantifies the resulting welfare loss. Section 4.7 derives the multi-sector leakage asymmetry that anchors the empirical cross-partial in Section 5.

Relation to existing frameworks. The model inherits the three-player architecture of Clayton et al. (2026): a hegemon exercises influence through an economic network mediated by a partner whose alignment choice is endogenous. Our departure is on the *instrument-choice* margin. In their framework the hegemon's coercive capacity operates through a single leverage variable. We endogenize the switch between positive inducement (market access $R(x)$) and negative coercion (security doctrine λ). The hegemon's optimal policy is continuous in the rival's efficiency θ, but it is non-differentiable at a threshold $\bar{\theta}$. The corner solution at this threshold emerges because the interior subsidy becomes fiscally unbounded (Proposition 2).

Relative to Clayton et al. (2024), two structural differences drive the different equilibrium characterization. First, their dyadic model admits a unique monotone equilibrium along the coercion–integration margin. Our triangular U–S–C specialization instead introduces a participation-constraint wedge at the Swing State ($R(x) \geq C(x) + V_{out}$). Whether this wedge binds depends on a production-structure primitive, σ_c. The wedge's interaction with the hegemon's fiscal capacity $\bar{T}$ then produces the kinked $\lambda^*(\theta)$ of Proposition 1, rather than a smooth interior response. Second, their coercion instrument operates directly on the Rival, whereas our λ operates on the Swing-State intermediary, whose contractibility is constrained by $\sigma_c < 1$. As a result, the coercive-regime threshold $\bar{\theta}$ is not a preference parameter of the hegemon, as it is in the dyadic model. Instead, it is a derived object, determined jointly by $(\sigma_c, \alpha, V_{out}, \bar{w})$. These two modifications deliver the paper's testable prediction—the cross-sectional tariff-×-swing cross-partial on FVA shares—which the dyadic architecture does not identify.

Relative to the common-agency benchmark, the $\sigma_c < 1$ primitive breaks the truthful-equilibrium characterization of Bernheim and Whinston (1986) in a specific way. In their menu auction with linear transfer technology, both principals (the Hegemon and the Rival) can bid truthfully on the agent's action, and the truthful equilibrium is efficient. Under bottleneck complementarity, the Rival's truthful bid for alignment requires a transfer that scales as $\Delta C(x; w_C)$. This transfer diverges as the upstream rival becomes efficient—equivalently, as $x \to 1$ under $\sigma_c < 1$. The Hegemon cannot truthfully match this bid, because his own transfer is bounded by $V_{US}(x)$, which is fixed in θ. The Segal (1999) externality-contracting logic applies here: the second-best response of the bidder with the bounded instrument is not a smaller bid, but exclusion. In our setting, this exclusion is the discrete shift from inducement ($\lambda = 0$) to coercion ($\lambda > 0$) that the Regime-Switch formalizes.

The $\sigma_c < 1$ primitive is itself standard and is quantitatively disciplined by Atalay (2017), Boehm et al. (2019), and Baqaee and Farhi (2024). The multi-sector leakage corollary (Proposition $5'$) rests on the same primitive. The probabilistic kill-switch reading of λ (Appendix B.2) connects to Kleinman et al. (2024) on how trade dependence shapes alliance formation. Relative to the quantitative trade frameworks of Caliendo and Parro (2015) and Baqaee and Farhi (2024), the model is deliberately small. It names the cross-sectional moment that value-added data identify—the tariff-×-swing cross-partial in Section 5—and leaves counterfactual quantification to future work.

4.1 The Model Environment

We model a global economy characterized by Vertical Geoeconomic Rivalry. Unlike standard trade models that depict competition between symmetric nations, we treat the global order as a vertical supply chain in which power derives from position: control over upstream inputs versus control over downstream market access.

4.1.1 Agents and Vertical Structure

The economy consists of three strategic players arranged hierarchically. First, **Hegemon C (Upstream)** represents the "Challenger" (e.g., China). It acts as a monopolist supplier of a critical intermediate input, with the input price w_C as its strategic instrument. Second, **Hegemon U (Downstream)** represents the "Incumbent" (e.g., USA). It acts as a monopsonist buyer of final goods and provider of essential intellectual property; its strategic instruments are the market access schedule $R(x)$ and the strategic doctrine λ. Finally, **The Swing State (Midstream)** is a representative firm or sector k (e.g., South Korea) that transforms intermediate inputs into final goods. It chooses the alignment $x \in [0, 1]$, which represents the degree of supply chain decoupling.

4.1.2 Timing of Events

The interaction unfolds in three stages, assuming a Stackelberg leadership structure that reflects the rigidity of global supply chains. In **Stage 1**, Hegemon C sets the export price of critical inputs w_C to maximize a weighted sum of commercial profits and geopolitical influence. In **Stage 2**, observing w_C, Hegemon U determines its strategic doctrine λ (the weight placed on security) and announces a market access schedule $R(x)$ to induce alignment. Finally, in **Stage 3**, the Swing State observes the input prices and policy schedule $\{w_C, w_A, R(x)\}$ and chooses the alignment x to maximize profit.

Stage 1		Stage 2		Stage 3
Upstream Pricing	$\longrightarrow$	**Downstream Policy**	$\longrightarrow$	**Alignment Choice**
Hegemon C sets w_C		Hegemon U sets $\lambda, R(x)$		Swing State chooses x

Figure 4: The Timeline of Vertical Rivalry

Assumption 1 (Commitment and Information)**.** *The extensive form in Figure 4 operates under three informational and institutional restrictions. (i) Sequential commitment: At each stage, the acting player commits to their strategic variable before the next player moves. In particular, Hegemon U commits to the pair $(\lambda, R(\cdot))$ prior to the Swing State's alignment choice, and this commitment is binding within the game. (ii) Complete information: All structural parameters are common knowledge. These include the elasticity σ, the efficiency gap $w_A - w_C$, the fiscal capacity reflected in γ, and the outside option V_{out}. (iii) Forcing* *contracts: When the Swing State's incentive-compatibility constraint binds with indifference at the target alignment x^*, Hegemon U is assumed to be able to impose an arbitrarily small off-path penalty that selects x^* as the unique best response.*

Discussion. Assumption 1 plays a central role. The credibility of ex-ante commitments in geoeconomics—sanctions threats, security guarantees, access regimes—is itself an institutional achievement. The quadratic cost $\frac{\gamma}{2}\lambda^2$ in Hegemon U's objective should be read as the reduced form of this commitment cost (in the spirit of Rogoff, 1985; Maskin and Tirole, 1999; Aghion and Tirole, 1997). If either the commitment or the forcing-contract element were relaxed, time-inconsistency pressures would push Hegemon U back toward the Generous Empire regime even when the structural conditions for coercion have been met. We return to this point when interpreting the comparative statics of γ below.

4.1.3 Production Technology: The Bottleneck

The Swing State produces a final good Y using a Constant Elasticity of Substitution (CES) technology that aggregates two distinct input sources. Specifically, the firm combines inputs sourced from the Upstream Hegemon (I_C, representing the "Red Supply Chain") with those acquired from alternative suppliers (I_A, representing the "Blue Supply Chain"):

$$Y = \left[\alpha I_C^{\frac{\sigma-1}{\sigma}} + (1-\alpha) I_A^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{1}$$

where $\alpha \in (0,1)$ is the distribution parameter and $\sigma \geq 0$ is the elasticity of substitution between the geopolitical sources.

Assumption 2 (The Bottleneck Condition)**.** *The inputs are gross complements, such that $\sigma < 1$.*

This assumption is the structural foundation of our model. It implies that the Upstream Hegemon's inputs are essential; they cannot be substituted by alternative sources without incurring a catastrophic loss in productivity. This captures the economic reality of "choke-point" technologies, such as rare earth processing or legacy semiconductors, where the rival holds a dominant upstream position.

4.1.4 The Unit Cost Function

The Swing State chooses a geopolitical alignment $x \in [0, 1]$, defined as the physical share of inputs sourced from the Alternative supply:

$$I_A = xI \quad \text{and} \quad I_C = (1 - x)I \tag{2}$$

where I is the total aggregate input bundle. Note that x is a *policy-imposed physical allocation share*, not the CES cost-minimizing input ratio. The Swing State's problem is to choose x to balance policy rewards $R(x)$ against the cost of deviating from the unconstrained cost-minimizing mix. Let w_C and w_A denote the prices of Chinese and Alternative inputs, respectively. We assume $w_C < w_A$, reflecting the Upstream Hegemon's comparative advantage or subsidization.

Substituting these shares into the CES production function, we derive the Unit Cost Function $C(x)$:

$$C(x; w_C) = \frac{(1 - x)w_C + xw_A}{\left[\alpha(1 - x)^{\frac{\sigma-1}{\sigma}} + (1 - \alpha)x^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}} \tag{3}$$

Properties of $C(x)$: Under the bottleneck assumption ($\sigma < 1$) and efficiency gap ($w_C < w_A$), the cost function exhibits three critical characteristics. First, it displays *Increasing Cost* ($C'(x) > 0$): alignment with the US is strictly costly. Second, it exhibits *Convexity* ($C''(x) > 0$): the marginal cost of decoupling rises sharply as $x \to 1$. Third, the input price acts as a *Shift Parameter*. A decrease in w_C ("China's Rise") lowers the base cost $C(0)$ significantly more than $C(1)$, thereby steepening the implementation curve for the Downstream Hegemon.

4.2 Stage 3 (Swing State): Alignment Choice

We solve the model via backward induction, beginning with the final stage. The Swing State (Agent) takes the input prices $\{w_C, w_A\}$ and the Downstream Hegemon's policy schedule $R(x)$ as given. It chooses the alignment $x \in [0,1]$ to maximize its total profit.

4.2.1 Profit Maximization

The profit function $\Pi(x)$ is defined as the market access revenue minus the production cost:

$$\max_{x\in[0,1]} \quad \Pi(x) = R(x) - C(x; w_C) \tag{4}$$

Assuming the policy schedule $R(x)$ is differentiable, the optimal alignment x^* is characterized by the First Order Condition:

$$R'(x^*) = C'(x^*) \quad \text{(Incentive Compatibility)} \tag{5}$$

This condition implies that to induce any non-zero alignment $x^* > 0$, the Downstream Hegemon must provide a marginal market access benefit exactly equal to the marginal cost of decoupling.

4.2.2 Properties of the Bottleneck Cost Function

The behavior of the Swing State is driven by the shape of the unit cost function $C(x)$. Under the Bottleneck Assumption ($\sigma < 1$), the cost structure exhibits specific convex properties that determine the difficulty of implementation.

Differentiating the cost function (3) with respect to alignment x:

$$C'(x) = \frac{w_A - w_C}{A(x)} - \frac{[(1-x)w_C + xw_A] \cdot A'(x)}{[A(x)]^2} \tag{6}$$

where $A(x) = [\alpha(1-x)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)x^{\frac{\sigma-1}{\sigma}}]^{\frac{\sigma}{\sigma-1}}$ is the productivity aggregator.

Lemma 1 (The Bottleneck Gradient). *If inputs are complements ($\sigma < 1$) and the Upstream Hegemon is efficient ($w_C < w_A$):*

1. *Monotonicity: $C'(x) > 0$ for all $x \in (0, 1)$. Decoupling always incurs a cost.*

2. *Explosive Marginal Cost: As $x \to 1$, $C'(x) \to \infty$. Fully replacing the upstream input is asymptotically impossible or infinitely expensive.*

3. *Input Price Sensitivity: A decrease in the upstream price w_C steepens the cost curve. Specifically,*

$$\frac{\partial C'(x)}{\partial w_C} < 0 \tag{7}$$

Proof. (Sketch) The first two properties follow directly from the CES properties with $\sigma < 1$. The third property arises because a lower w_C reduces the base cost $C(0)$ (where upstream input usage is intensive) significantly more than it reduces $C(x)$ for high x, thereby increasing the slope $C(x) - C(0)$. □

Strategic Implication: The lemma implies that "China's Rise" ($w_C \downarrow$) does not merely make the Swing State richer; it makes the Swing State *harder to align*. The Downstream Hegemon faces a steeper marginal cost curve $C'(x)$ to implement the same target x^*, requiring a more aggressive (steeper) policy schedule $R(x)$.

4.3 Stage 2 (Downstream Hegemon): Policy and Doctrine

In Stage 2, the Downstream Hegemon (US) designs the policy schedule $R(x)$ and selects its strategic doctrine λ to maximize national welfare, anticipating the Swing State's reaction in Stage 3.

4.3.1 The Joint Objective Function

The Hegemon's utility $\mathcal{W}_{US}$ is defined over the realized alignment x, the transfer payment $R(x)$, and the security weight $\lambda \in [0, 1]$. We model the choice of λ not as a fixed preference, but as an endogenous policy decision (e.g., "Securitization") that incurs convex institutional costs.

$$\mathcal{W}_{US}(x, R(x), \lambda) = \underbrace{\lambda\, x\, \theta(w_C)}_{\text{Security Utility}} + \underbrace{(1-\lambda)[V_{US}(x) - R(x)]}_{\text{Net Economic Surplus}} - \underbrace{\frac{\gamma}{2}\lambda^2}_{\text{Cost of Securitization}} \tag{8}$$

where $V_{US}(x)$ represents the gross economic value generated for the US, such as consumer surplus or IP royalties. We assume that $V'_{US}(x) < 0$, reflecting the inevitable efficiency loss derived from decoupling from the most efficient upstream supplier. The parameter $\gamma > 0$ quantifies the institutional resistance to adopting a security-heavy doctrine, capturing structural frictions such as fiscal constraints or domestic political polarization. The multiplier $\theta(w_C) \equiv \bar{w}/(w_C - m)$ is the Upstream Hegemon's exogenous technological productivity, as defined in equation (14). It scales the Security Utility: as the bottleneck tightens—i.e., as the Upstream Hegemon's critical input becomes cheaper—the strategic salience of alignment rises proportionally.[3]

4.3.2 The Implementation Problem

The Principal cannot choose x directly; it must induce x through the schedule $R(x)$.

Step 1: The Cost of Implementation

To induce a specific target alignment x, the US must satisfy two constraints:

1. *Incentive Compatibility (IC):* $R'(x) = C'(x)$.

2. *Participation Constraint (PC):* $R(x) \geq C(x) + V_{out}$, where V_{out} is the Agent's outside option (integration with China).

[3] Formally, θ is primitive: Hegemon C chooses m taking θ as given, so $\theta(w_C)$ is well-defined in the US Stage-2 problem once the equilibrium w_C is in hand. The microfoundation for why $\lambda\, x$ scales with θ is the kill-switch interpretation of Appendix B.2: the potential damage D averted by alignment is naturally increasing in the Upstream Hegemon's productivity, so that under the normalization $D(\theta) = \theta$ the security utility becomes $\lambda\, x\, \theta$.

Since the US minimizes the transfer $R(x)$, the PC binds at the optimum:

$$R(x) = C(x; w_C) + V_{out} \tag{9}$$

This equation represents the *Implementation Cost Function*: the minimum transfer required to make a target x implementable.

The Market Access Schedule $R(x)$ required to implement alignment x is explicitly determined by the bottleneck technology:

$$R(x) = \frac{(1-x)w_C + xw_A}{A(x)} + V_{out} \tag{10}$$

where $A(x) = \left[\alpha(1-x)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)x^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$ represents the aggregate productivity of the supply chain mix.

Remark on Methodology: Following the standard Principal-Agent approach with observable actions (Grossman and Hart, 1983), we solve for the optimal policy by treating the implementation cost $R(x) = C(x) + V_{out}$ as the effective price of alignment. This reduces the problem to maximizing the Principal's objective, taking this cost as given. This formulation implicitly assumes the Principal utilizes a "forcing contract" (or threshold penalties) to resolve the Agent's indifference and ensure Incentive Compatibility is strictly satisfied at the target x^*.

Step 2: The Reduced Maximization

Substituting the implementation cost into the objective function, let $J(x) = V_{US}(x) - C(x; w_C)$ denote the *Joint Economic Surplus*. The US problem reduces to choosing the target x and doctrine λ:

$$\max_{x,\lambda} \quad \mathcal{W}_{reduced} = \lambda\, x\, \theta(w_C) + (1-\lambda)[J(x) - V_{out}] - \frac{\gamma}{2}\lambda^2 \tag{11}$$

4.3.3 Optimal Policy and Doctrine

We characterize the equilibrium by maximizing the reduced objective function (11) with respect to the alignment target x and the security doctrine λ. The solution is determined by the intersection of two reaction functions.

First, consider the optimal induced alignment $x^*(\lambda)$. Taking the security weight λ as given, the US chooses the target x^* to equate the marginal political gain with the marginal economic loss:

$$\frac{\partial \mathcal{W}}{\partial x} = \lambda\,\theta(w_C) + (1-\lambda)J'(x^*) = 0 \quad \Longrightarrow \quad J'(x^*) = -\frac{\lambda\,\theta(w_C)}{1-\lambda} \tag{12}$$

This condition implies that as the security doctrine becomes more hawkish ($\lambda \to 1$), the US pushes the target x^* further into the economically inefficient region where the marginal loss of joint surplus $J'(x)$ is high. The threat multiplier $\theta(w_C)$ amplifies this effect: a more productive Upstream Hegemon (higher θ) shifts x^* upward at every level of λ.

Second, consider the endogenous regime choice $\lambda^*(x)$. Differentiating with respect to λ, we obtain the optimal doctrine rule:

$$\frac{\partial \mathcal{W}}{\partial \lambda} = x\,\theta(w_C) - [J(x) - V_{out}] - \gamma\lambda = 0 \quad \Longrightarrow \quad \lambda^*(x) = \frac{x\,\theta(w_C) - [J(x) - V_{out}]}{\gamma} \tag{13}$$

Equation (13) defines the *Strategic Gap*: the difference between the threat-weighted Security Benefit ($x\,\theta(w_C)$) and the Net Economic Surplus ($J(x) - V_{out}$). The optimal doctrine λ^* is proportional to this gap. A rising Upstream Hegemon (lower w_C, higher θ) widens the gap through two channels. First, it raises the perceived threat intensity θ. Second, it compresses the Net Economic Surplus by raising the implementation cost $R(x)$ of re-routing supply chains.

Proposition 1 (The Regime Switch Mechanism)**.** *The equilibrium is the fixed point* (λ^*, x^*) *of these two reaction functions. The nature of the equilibrium depends on the magnitude of* *the Joint Surplus $J(x)$:*

- *Generous Empire: If $J(x)$ is large (trade is profitable), the Strategic Gap is small or negative. Consequently, $\lambda^* \to 0$, and the US relies on market integration.*
- *Coercive Hegemon: If $J(x)$ collapses (due to rising costs $C(x)$ or falling value V_{US}), the Strategic Gap widens. The optimal λ^* rises smoothly, triggering a shift toward coercion.*

4.3.4 Existence, Uniqueness, and Local Stability

The equilibrium (λ^*, x^*) is well-defined. In Appendix B.1 we verify that the Stage-2 objective is strictly concave in (x, λ) whenever γ is sufficiently large—i.e., when the institutional cost of securitization dominates the cross-effect. This concavity yields a unique interior fixed point that is locally stable under best-response dynamics. Economically, the "γ sufficiently large" condition means that doctrine and alignment do not amplify each other into "tipping points." When γ is too small, multiplicity arises and the economy can jump discontinuously between peaceful and coercive equilibria.

4.4 Stage 1 (Upstream Hegemon): Pricing Strategy

In the first stage, Hegemon C sets the terms of trade for the critical input. To distinguish between technological fundamentals and strategic intent, we decompose the input price into production costs and industrial policy.

4.4.1 Technology and Industrial Policy

Let w_C be the final export price of the input. We decompose this price into an exogenous technology component and an endogenous policy component:

$$w_C = \underbrace{\frac{\bar{w}}{\theta}}_{\text{Marginal Cost}} + \underbrace{m}_{\text{Strategic Markup}} \tag{14}$$

where $\bar{w}$ denotes the baseline factor price. The parameter $\theta > 0$ represents the *Exogenous Technological Productivity*; an increase in θ captures the structural phenomenon of "China's Rise" or technological catch-up. The variable m represents the *Endogenous Industrial Policy*, which captures the state's strategic intent. A positive value ($m > 0$) indicates a standard profit-maximizing monopolistic markup, whereas a negative value ($m < 0$) indicates a strategic state subsidy (or dumping) designed to capture market share.

4.4.2 Hegemon C's Optimization

Normalizing final output $Y = 1$, Hegemon C chooses the markup m (and thus the price w_C) to maximize a weighted sum of commercial profit and geopolitical influence:

$$\max_{m} \quad \mathcal{W}_C = \underbrace{m \cdot Q_C(x(w_C))}_{\text{Net Commercial Return}} + \underbrace{\phi(1 - x(w_C))}_{\text{Geopolitical Influence}} \tag{15}$$

where $Q_C(x) = \frac{1-x}{A(x)}$ is the derived demand for Chinese inputs. The objective highlights the trade-off: lowering m (subsidies) incurs a fiscal cost but increases the strategic value $\phi(1-x)$ by binding the Swing State to the Red Supply Chain.

The Geopolitical Elasticity: Hegemon C faces a derived demand curve $Q_C(x(w_C))$ whose elasticity ε_D is composite. It depends on (1) the technological substitution σ, and (2) the *strategic reaction* of the US. Crucially, this creates a *Policy Feedback Loop*: if China lowers prices w_C too aggressively, it may trigger a "Coercive Regime Switch" ($\lambda^* \uparrow$) by the US that forces x upward, ironically contracting Chinese export volume Q_C despite the lower price.

The optimal price w_C^* is characterized by a modified Monopoly Pricing Rule:

$$\frac{w_C^* - (\bar{w}/\theta)}{w_C^*} = \underbrace{-\frac{1}{\varepsilon_D}}_{\text{Monopoly Markup}} - \underbrace{\frac{\Phi_{strat}}{w_C^*}}_{\text{Geopolitical Discount}} \tag{16}$$

where $\Phi_{strat} > 0$ represents the marginal strategic value of market share. Because influence matters ($\phi > 0$), Hegemon C optimally sets a price *lower* than the standard profit-maximizing monopoly level, effectively engaging in predatory pricing to secure the vertical supply chain.

4.5 General Equilibrium: The Transmission of the Shock

Having characterized the optimal strategies of the Upstream Hegemon (Pricing), the Downstream Hegemon (Doctrine), and the Swing State (Alignment), we now close the model. We analyze how an exogenous increase in Chinese productivity ($\theta \uparrow$) reverberates through the vertical supply chain and reshapes the global security architecture.

4.5.1 Proposition 2: The Paradox of Competitiveness

Proposition 2 (The Paradox of Competitiveness)**.** *An exogenous rise in Upstream Productivity ($\theta \uparrow$, "China's Rise") necessarily triggers an endogenous rise in the Downstream Security Doctrine ($\lambda^* \uparrow$, "US Hawkishness"). The resulting equilibrium is coercive.*

Proof Mechanism: The transmission of the shock follows a four-stage causal chain:

1. *The Efficiency Shock ($\theta \uparrow$):* An improvement in upstream technology lowers Hegemon C's marginal cost of production, $MC_C = \frac{\bar{w}}{\theta}$.

2. *Upstream Price Passthrough ($w_C^* \downarrow$):* Hegemon C re-optimizes its pricing strategy. Since both the commercial profit motive and the geopolitical motive ($\phi > 0$) align towards capturing market share, the optimal equilibrium price w_C^* falls.

$$\frac{dw_C^*}{d\theta} < 0 \tag{17}$$

3. *Widening the Cost Gap ($\Delta C \uparrow$):* The Swing State's unit cost function reacts asymmetrically. Because the "Red Chain" ($x = 0$) utilizes Chinese inputs intensively while the "Blue Chain" ($x = 1$) does not, a drop in w_C lowers the floor $C(0)$ significantly

more than the ceiling $C(1)$. Consequently, the *Cost of Decoupling* $\Delta C = C(x) - C(0)$ widens:

$$\frac{\partial \Delta C}{\partial w_C} = \frac{\partial C(x)}{\partial w_C} - \frac{\partial C(0)}{\partial w_C} < 0 \tag{18}$$

(Since $\frac{\partial C}{\partial w_C}$ is proportional to input share $(1-x)$, and $(1-x) < 1$).

4. *The Endogenous Regime Switch ($\lambda^* \uparrow$):* The Downstream Hegemon observes that the subsidy required to implement voluntary alignment (paying ΔC) has become fiscally prohibitive. The "Strategic Gap" in the US objective function widens. To satisfy the optimal doctrine rule derived in Eq. (13), the US increases its security weight λ^*:

$$\frac{d\lambda^*}{d\theta} = \underbrace{\frac{\partial \lambda^*}{\partial w_C}}_{(-)} \cdot \underbrace{\frac{dw_C^*}{d\theta}}_{(-)} > 0 \tag{19}$$

4.5.2 Corollary: The Inevitability of Coercion

This result implies that the transition to a coercive trade policy is not necessarily driven by ideological shifts or "unfair" industrial policy (m). Even if Hegemon C behaves as a benign market actor (holding m constant), a pure productivity shock is sufficient to trigger the regime switch.

We term this dynamic the *Tragedy of Efficiency*: the more efficient China becomes, the harder it is for the US to compete through positive inducements such as subsidies or market access. Once the required subsidy exceeds US fiscal or political capacity, the optimal strategy flips from integration (low λ, high R) to coercion (high λ, negative R). The US then uses tariffs not to protect its own industries, but to artificially destroy the rival's efficiency advantage, restoring the relative attractiveness of the Blue Chain.

4.6 Welfare: The Tragedy of Efficiency

A claim of tragedy requires a welfare statement. Let $\mathcal{S}(x;\theta)$ denote the joint economic surplus of the three players: Hegemon U's gross value plus Hegemon C's commercial profit, minus the Swing State's production cost. The formal decomposition is Proposition 6 in Appendix B.5. The integrated planner's *first-best alignment* $x^{FB}(\theta)$ maximizes $\mathcal{S}$ while ignoring the security externality, and is *declining* in θ: a more efficient Upstream rival raises the opportunity cost of decoupling. The equilibrium $x^*(\theta)$ moves in the opposite direction (Proposition 2). Define the *efficiency wedge*:

$$\Delta(\theta) \;\equiv\; \mathcal{S}\big(x^{FB}(\theta);\theta\big) \;-\; \mathcal{S}\big(x^*(\theta);\theta\big) \;\geq\; 0. \tag{20}$$

Proposition 3 (The Tragedy of Efficiency)**.** *Under Assumption 2 and in the coercive regime ($\lambda^* > 0$), the efficiency wedge $\Delta(\theta)$ is strictly increasing in Upstream productivity θ. That is, improvements in Upstream technology generate* larger *global welfare losses relative to the integrated first-best.*

Proof sketch. In the coercive regime, $x^*(\theta)$ rises with θ while $x^{FB}(\theta)$ falls; the gap widens monotonically, and strict concavity of $\mathcal{S}$ at x^{FB} delivers the result. The full welfare decomposition—including the curvature and second-order-markup conditions under which the commercial component is strictly decreasing in θ—is in Appendix B.5. □

Interpretation. Proposition 3 is the formal sense in which "Tragedy of Efficiency" is literal. Under horizontal rivalry, a pure productivity gain in the Upstream rival would be unambiguously welfare-improving. Here, instead, it triggers an equilibrium response in Downstream doctrine that destroys commercial surplus faster than the gain accrues. The Swing State bears the bulk of the loss, because the coercive corner truncates its choice set.

Why verticality, not horizontality, leads to coercion. Under horizontal rivalry—whether competing final goods (subsidy war) or competing sourcing (bidding war)—the cost of retaining the Swing State's loyalty is linear in the rival's efficiency gain, and the security parameter λ stays at zero. Under the vertical bottleneck ($\sigma < 1$), the implementation cost $C(x)$ is strictly convex and asymptotic (Lemma 1). The subsidy required to induce alignment $x \to 1$ therefore explodes as $w_C \downarrow$: the hegemon cannot buy loyalty and must instead punish disloyalty. Appendix B.8 develops the three-case comparison in detail.

4.7 Leakage Asymmetry Across Sectors

The regime transition in Proposition 1 characterizes the Hegemon's choice in a single sector. With multiple sectors of differing σ_k (critical, strategic, commodity), the intensity of substitution following coercion differs sharply by bucket. The same security shock $d\lambda$ that fully displaces Chinese supply in commodity sectors displaces it only partially in critical sectors, because the Swing State cannot absorb the displacement without re-importing Chinese intermediates. Through this indirect leakage channel, Chinese value-added reaches the Hegemon via swing-state re-exports rather than direct shipments.

Proposition 4 (Leakage Asymmetry)**.** *Let two sectors $k \in \{c, m\}$ differ only in their CES substitution elasticity, with $\sigma_c < 1 < \sigma_m$ (critical vs. commodity). Suppose an exogenous rise in λ induces displacement of Chinese supply. The equilibrium reduction in U's direct imports from China is partially offset by Swing-State re-imports from China in the critical sector ($k = c$), whereas in the commodity sector ($k = m$) substitution is near-complete. Formally, letting $\Delta M_k^{C\to U}(\lambda)$ denote the equilibrium change in U's direct China imports and $\Delta M_k^{C\to S}(\lambda)$ the change in Swing-State China imports,*

$$\left.\frac{\partial \Delta M_k^{C\to S}/\partial\lambda}{\partial|\Delta M_k^{C\to U}|/\partial\lambda}\right|_{k=c} > \left.\frac{\partial \Delta M_k^{C\to S}/\partial\lambda}{\partial|\Delta M_k^{C\to U}|/\partial\lambda}\right|_{k=m}.$$

That is, the leakage ratio is strictly higher in the bottlenecked sector.

Proof (Sketch). The Swing State's cost-minimization problem is CES with substitution elasticity σ_k between Chinese and non-Chinese intermediates. A rise in λ raises the effective shadow cost of the Chinese input only at the U–S interface (via the menu schedule $R(x)$), not at the C–S interface. The Swing State's optimal mix therefore shifts more sluggishly in low-σ sectors, since substitution toward non-Chinese intermediates is bounded above by $\sigma_k \cdot (d \ln \text{relative price})$. For $\sigma_k < 1$, residual demand for Chinese intermediates is inelastic; the Swing State continues sourcing them and re-exports the embodied content to the U, which is the leakage channel. For $\sigma_k > 1$, the Swing State can switch suppliers cheaply, so the leakage ratio shrinks toward zero. The full derivation is in Appendix B.6. □

Empirical correlate (flow margin). The triple-difference and event-study estimates in Section 5 are the *flow-margin* test of this proposition: persistent post-2018 reductions in U.S. critical-sector imports from China that are *not* fully matched by parallel reductions through swing-state partners. This margin depends on an identified λ-shock around 2018. Placebo-year diagnostics and Honest-DiD bounds suggest that shock is only partly isolated from pre-existing trends, so we complement the flow margin with a cross-sectional steady-state test, formalized next.

Proposition 5′ (Steady-State Leakage Composition)**.** *Let* $\text{FVA}^{C}_{s,j}$ *denote the Chinese value-added share in country s's stationary-equilibrium gross exports to Hegemon U in industry j. Let* $\tau_j \geq 0$ *denote the Hegemon's tariff on direct Chinese exports in j, and let* $s \in \{S, N\}$ *index Swing versus Non-swing origins. Assume:*

(A1) Complementarity. $\sigma_c < 1$ *in the critical production structure (same CES primitive as Proposition 4).*

(A2) Baseline dependence. The Swing State's baseline Chinese-intermediate share strictly exceeds the Non-swing origin's: $\alpha_S > \alpha_N$*.*

(A3) Empirically relevant range. Baseline Chinese-content shares in swing-state exports to U satisfy $s^C_S, s^C_N \in (0, \frac{1}{2}]$*.*

(A4) Ordered pass-through. Tariff pass-through to effective Chinese-input price at the C–S interface is weakly larger for the Swing State than for the Non-swing origin: $\eta_S\, \psi'_S \;\geq$

$\eta_N\,\psi'_N$, where η_s is the Chinese-intermediate share and ψ'_s the marginal pass-through from τ_j to the Swing's/Non-swing's Chinese-input shadow price.

Then the steady-state FVA-share satisfies the positive cross-partial

$$\left.\frac{\partial^2 \mathrm{FVA}^C_{s,j}}{\partial \tau_j\,\partial \mathbf{1}\{s=S\}}\right|_{\sigma_c<1} > 0.$$

Equivalently, across industries that differ in τ_j, the tariff-gradient of the embodied-Chinese share is strictly steeper for swing-state exporters than for non-swing exporters.

Remark on (A3)–(A4). (A3) is the empirically relevant range for the object we identify: in our OECD ICIO and ADB MRIO sample, the swing-state FVA-China share in critical-sector exports to the U is approximately 6%, well inside $(0, 1/2]$ (Appendix B.7). (A4) is a symmetric-pass-through condition: it holds whenever the Swing State's larger baseline Chinese-intermediate share translates into weakly larger sensitivity of its effective Chinese-input shadow price to τ_j. We state it as an explicit hypothesis rather than derive it from (A2), because the mapping from α_s to $\eta_s\,\psi'_s$ depends on the Swing State's *downstream* bargaining structure vis-a-vis U, which is outside the scope of this paper. Under (A2) alone, (A4) is a natural ordering but not a theorem. Finally, the cross-partial in Proposition 5′ is a group-average object. Reclassifying a single origin between $\{S, N\}$ leaves the group moment unchanged up to $O(1/|S|)$, so the identification is robust to borderline-country manipulation of the swing-state partition—a counter-fragility property our leave-one-out diagnostic confirms empirically.

Proof (Sketch). The Swing State's cost-minimizing input mix is CES with elasticity $\sigma_c < 1$ between Chinese and non-Chinese intermediates. A higher τ_j raises the shadow price of *direct* Chinese supply at the U–C interface but not at the C–S interface, so the Swing State's effective Chinese-input price rises by less than the Hegemon's. With $\sigma_c < 1$, the Swing State cannot cheaply rebalance away from Chinese intermediates, and a larger fraction of the Hegemon's tariff-induced displacement is absorbed as re-exported Chinese content in j. The

Non-swing origin, with strictly smaller baseline Chinese-intermediate share, absorbs less per unit of τ_j reallocation. Crucially, the claim is about the *cross-partial* (the tariff-gradient is steeper for S), not merely the level (the FVA share is higher for S). This follows because the CES expenditure share on Chinese inputs is concave in the effective relative price. At the Swing State's higher baseline Chinese-input share, the same $d\tau_j$ moves the Swing State along a steeper segment of the cost-minimization surface. It therefore displaces a larger quantum of Chinese value-added *back through* re-exported intermediates than the Non-swing origin, which operates at a flatter segment. The cross-partial is strictly positive whenever $\sigma_c < 1$ and the Swing State's baseline Chinese-intermediate share strictly exceeds the Non-swing origin's. The full derivation is in Appendix B.7. □

Empirical correlate (composition margin). The continuous triple-difference of Section 5 is a direct identified moment of $\partial^2 \mathrm{FVA}^C_{s,j} / \partial \tau_j \, \partial \mathbf{1}\{s = S\}$. It interacts three variables—published Section-301 tariff rates, swing-state identity, and post-2018 timing—with the FVA-China share in swing-state $\rightarrow$ US exports as outcome. Unlike the flow-margin test, this moment does not require isolating a 2018 shock. Under Proposition 5′, the cross-partial is a *steady-state* object; cross-sectional identification—off industry-level tariff variation interacted with geographic swing-state identity—is therefore sufficient to test it.

Reconciling the kink with smooth aggregate dynamics. Proposition 1 predicts a continuous-but-nondifferentiable kink in $\lambda^*(\theta)$ at $\bar{\theta}$, while the empirical dynamics (e.g., Figure 3) appear smooth. Three non-exclusive reconciliations—heterogeneous sectors crossing the threshold sequentially, Bayesian uncertainty about σ smoothing the expected λ^*, and institutional adjustment costs on coercive capacity—are developed in Appendix B.9.

Extensions. Three extensions relax baseline modelling choices without altering the core results.

First, Nash-bargaining power β for the Swing State *amplifies* coercion: more powerful swing states trigger the regime switch sooner (Proposition 5; Appendix B.3). This rationalizes the US pattern of combining subsidies (CHIPS Act) with coercive measures for allies with high β (Korea, Taiwan, the Netherlands).

Second, endogenizing vertical position via a specialization pre-stage replicates the baseline structure when $\theta_C > \theta_U$ and $E_U > E_C$. The vertical asymmetry is therefore a Ricardian outcome, not a modelling choice (Appendix B.4).

Third, a probabilistic kill-switch microfoundation reinterprets λ as the expected damage averted by alignment, connecting to the security-dilemma logic: $d\lambda^*/d\theta > 0$ means Upstream efficiency endogenously raises the Incumbent's perceived probability of weaponization (Appendix B.2).

5 Empirical Tests of the Leakage Mechanism

With the model of Section 4 and Propositions 4 and $5'$ in hand, we subject the leakage mechanism to two complementary regression-based tests.

Test 1 (the HS6-panel event study) targets the flow-margin prediction of Proposition 4. If Section-301 triggered a bilateral China–critical decoupling, the post-2017 coefficient path should turn sharply negative. If swing-state intermediation had already absorbed the margin before the shock, the path should instead be flat against the 2017 baseline, with a pre-existing negative pre-period.

Test 2 identifies the cross-partial of Proposition $5'$ using a continuous tariff-×-swing triple-difference on Chinese value-added shares in swing-state→US gross exports. Under $\sigma_c < 1$, the tariff-gradient of Chinese-VA content in swing-state exports is steeper than in non-swing exports. This is a *steady-state* object that does not require isolating a 2018 shock. A third dose–response regression on the Entity-List coercion stock (Spec 2) tests the paradox mechanism of Proposition 2 as an auxiliary margin.

Headline, previewed. Test 1 returns a null—not a refutation of leakage but its signature. The pre-2018 critical–commodity gap was already negative and *closes* toward zero after 2018, consistent with swing-state intermediation having absorbed the bilateral margin before the tariff shock. Robustness to group-specific trends, Rambachan–Roth sensitivity, and Poisson zero-retention leaves this reading intact. Test 2 then identifies the leakage cross-partial at $\hat{\beta}^{\text{post}}_{\tau\times\text{swing}} = +0.102$ ($p < 0.001$). A 25 pp Section-301 tariff increase on a swing-state-produced good is associated with a further +2.5 pp rise in its Chinese-VA content—a steady-state fingerprint of the $\sigma_c < 1$ production structure. Placebo-year diagnostics reframe this as cross-sectional rather than quasi-experimental identification (the gap pre-existed 2018), which is the correct reading of a steady-state proposition.

We construct a balanced panel of bilateral US imports at the HS6 level over 2000–2024, assigning each code to one of three buckets—*critical* (rare earths, advanced semiconductors, semiconductor-manufacturing equipment), *strategic* (mature-node chips, batteries, EV components), and *commodity* (textiles, basic steel, basic chemicals). Appendix C details the HS6 concordance, the UN Comtrade pull (premium API, 63 targeted HS6 codes, 18 reporters), and the BIS Entity List panel built from the OpenSanctions Consolidated Screening List.

Mapping specifications to propositions. Specs 1 and 3 test coefficients implied by the leakage asymmetry of Proposition 4; Spec 2 (the "paradox regression") is a dose–response test of the sector-asymmetric coercion effect. All three specifications isolate different pieces of the same underlying production structure, differing in how they collapse the data over time.

- **Spec 1 (triple-difference, single Post dummy).** Tests the *average* post-2018 differential. A naive reading of Proposition 4: $\beta_{\text{CHN}\times\text{Crit}\times\text{Post}} < 0$ — US imports from China fall more in critical than commodity sectors once λ rises. This is the "tariffs caused decoupling" sign that a casual reader would expect.

- **Spec 2 (paradox regression).** Tests the *dose–response* of the sector-asymmetric effect to the Entity-List coercion stock $\text{EL}^{\text{CHN}}_{t-1}$. Under Proposition 4, the interaction with the critical-sector dummy should load negatively: each additional unit of coercion should depress Chinese critical imports more than commodity imports.

- **Spec 3 (event study).** Tests the *dynamic* sector-asymmetric path year by year. Under Proposition 4, with swing-state intermediation already operating pre-2018, the (CHN $\times$ Critical) coefficient path should show two features. First, it should be *negative before 2018*, because the gap is a steady-state feature of the global production structure, not a 2018 policy break. Second, it should be *flat or rising from below toward zero thereafter*, because additional coercion does not re-open a gap that leakage has already normalized.

Spec 1 and Spec 3 are not equivalent: Spec 1 turns on whether the post period is below zero, while Spec 3 turns on whether the pre period was below zero and the post period converged. The two can point in different directions, and in our HS6 panel they do—Spec 1 is mildly positive while Spec 3 recovers the closure-of-pre-existing-gap pattern Proposition 4 implies when leakage is a steady-state feature rather than a post-shock adjustment. Spec 2's sector-asymmetric loading, sharp in the 3-sector aggregate, dissolves at HS6 resolution.

5.1 Test 1: Flow-Margin Event Study (Proposition 4)

Using this panel, we estimate an event-study specification for US imports from China in the critical bucket relative to a (swing-state $\times$ non-critical) base, with 2017 as the omitted pre-treatment year. Identification is at the HS6 $\times$ partner $\times$ year level (N $\approx$ 17,600). Fixed effects absorb partner-HS6, HS6-year, and partner-year variation, and standard errors are clustered at HS6 ($\sim$ 60 clusters):

$$\ln M_{ihst} = \alpha_{ih} + \gamma_{ht} + \delta_{it} + \sum_{\tau \neq 2017} \beta_\tau \cdot \mathbf{1}\{t = \tau\} \cdot \text{CHN}_i \cdot \text{Critical}_h + \varepsilon_{ihst},$$

where i indexes exporter (China + 17 candidate swing-state partners; the focal 3-state swing group {VNM, MEX, MYS} used in Test 2 is a subset), h indexes HS6 code, and t indexes years. HS6 resolution (relative to a 3-sector aggregate) raises sample size by a factor of 13 and permits cluster-robust inference on a genuine number of clusters.

The purpose of the event study is not to detect a 2018 break but to ask whether one *should* have existed. If Section-301 caused China–critical decoupling, the post-2017 β_τ path should turn sharply negative. Under the leakage reading of Proposition 4, the pre-2018 path was already negative because swing-state intermediation had absorbed the bilateral margin. The post-2018 path should then be flat against the 2017 baseline rather than deepen. The data match the latter. Pre-2018 coefficients sit at -0.3 to -0.8 log points, with 2012–2016 individually significant at 1–5%. Every post-2017 coefficient is statistically indistinguishable from the 2017 baseline. The joint pre-trend test $F_{5,61} = 3.69$ ($p = 0.006$) rejects a flat pre-period, which is exactly the pre-existing gap the leakage story predicts. See Appendix Figure 8 for the full coefficient path and Appendix Table 2 for the numerical values.

The headline single-coefficient triple-difference $\beta_{\text{CHN}\times\text{Critical}\times\text{Post}} = +0.283$ (SE 0.245, $p = 0.25$) is statistically insignificant with a positive sign. This mirrors the event study: the post-2018 coefficients cluster around zero while pre-2018 coefficients are negative, so the collapsed Post contrast loads weakly positive. The HS6 panel does not support a *widening* of the China–Critical gap after 2018; it supports the gap *closing*.

Trend-robustness. Column 2 of Table 2 adds group-specific linear trends interacting (CHN $\times$ Critical) and (Swing $\times$ Sector) with time. The absorbed China–Critical trend is -0.037 log points per year (SE 0.028, not significant); the residual Post interaction is $+0.684$ ($p = 0.013$). Even with the trend taken out, there is no evidence of a *negative* post-2018 residual effect on China–critical imports. The interpretation most consistent with this evidence is that by 2018 the China–critical gap had already closed via swing-state intermediation, and subsequent coercion neither re-opened it nor slowed the convergence. This is consistent with the *leakage* mechanism of Proposition 4 — swing states absorbed Chinese intermediate capacity into the US supply chain — and inconsistent with a simple "tariffs caused decoupling" reading.

Pre-trend sensitivity (honest DiD). Because the pre-trend F-test rejects strict parallel trends, we report robust identified sets for the average post-treatment effect using the relative-magnitudes partial-identification procedure of Rambachan and Roth (2023), which bounds the post-period deviation as a multiple M of the worst pre-period deviation. Two features matter (Appendix Figure 9; numerical bounds in Appendix C.4). First, *zero lies inside the identified set at every* $M \leq 2$, so the hypothesis that the post-2018 China–Critical gap equals the 2017 baseline is not rejected even under very weak parallel-trends assumptions. Second, the *upper* bound is positive throughout, ruling out any persistent post-2018 *widening* of the gap. The leakage reading—closure rather than widening—is thus not dependent on the parallel-trends assumption that the event-study literally rejects.

Zero-retention (PPML). The log specification drops cells with nonpositive trade, a selection disproportionately concentrated in China×Critical×post-2020—exactly the cells where decoupling would generate new zeros. We therefore re-estimate on a balanced (HS6×partner×year) panel with missing cells set to zero, using the Poisson pseudo-maximum-likelihood estimator of Santos Silva and Tenreyro (2006) via `ppmlhdfe` (Correia et al., 2020). On $N = 22{,}305$ observations (a 26% increase over the log sample), the baseline triple-difference coefficient is +0.289 (SE 0.309) and the trend-robust variant is +0.531 ($p = 0.050$). Both are close to their log-sample counterparts (+0.283 and +0.684) and, if anything, slightly weaker. The PPML event study (Appendix C, Table 4) reproduces the key qualitative pattern: negative pre-2018 coefficients (the largest, $\tau_{2015} = -0.400$, is individually significant) and a post-2018 band clustered around zero with no persistent negative drift. The zero-retention critique therefore does not rescue a "tariffs caused decoupling" reading; the gap-closing pattern is not an artifact of selecting on positive trade.

5.2 Test 2: Steady-State Cross-Partial on Chinese Value-Added (Proposition 5′)

The gross-trade test above ruled out a persistent post-2018 widening of the US–China gap but does not directly demonstrate Chinese content flowing through swing states. The value-added decomposition does.

Direct evidence from value-added decomposition. The triple-difference in gross trade is an *indirect* test of Proposition 4. A closing US–China bilateral gap in critical-sector imports can reflect leakage through swing-state intermediation, but only if swing-state exports to the US embed more Chinese value-added in critical sectors than in commodity sectors. We test this directly. We use two separately maintained inter-country input–output tables: the OECD ICIO EXTENDED release (ISIC Rev. 4, 50 industries, 85 economies, 1995–2022) and the ADB MRIO 72-economy release (35 industries, 2017–2024). For each swing state s, industry j, and year t, we compute the share of gross $s \to$ US exports accounted for by Chinese value-added:

$$\text{FVA}^{\text{CHN, share}}_{s,j,t} = \sum_{r \in \text{CHN}} v_r \cdot L_{r,(s,j)},$$

where v_r is value-added per unit of gross output in Chinese source sector r and $L = (I - A)^{-1}$ is the global Leontief inverse. This share is *multiplied by* gross bilateral exports $E_{s \to \text{US},j,t}$ to recover the dollar value of Chinese value-added embedded in those exports, which is what we use in the dollar-volume rescaling below. Figure 5 plots the across-swing-state mean by sector bucket and data source.

The level gap is the clearest evidence. In the OECD panel the critical-sector share rises from 0.4% in 1995 to a flat plateau of ≈ 6.1% over 2020–2022, while the commodity share rises from 0.5% to 4.4% over the same period. The critical–commodity gap opens in the mid-2000s and reaches ≈ 1.8 percentage points by 2015–2018. It then *widens* to ≈ 2.3 percentage points by 2020: Chinese value-added embedded in swing-state critical-sector exports to the

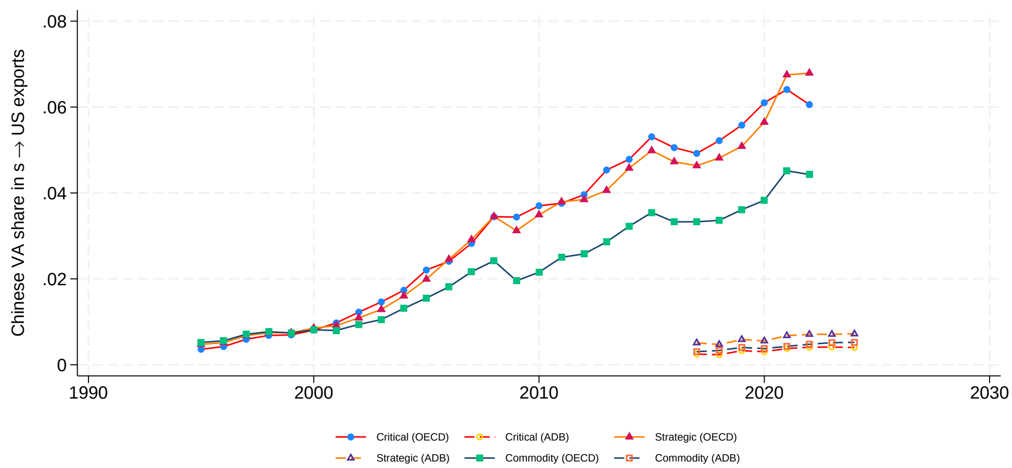

Figure 5: **Chinese value-added share in swing-state → US exports, by sector bucket.** Cross-swing-state mean of $\text{FVA}^{\text{CHN}}_{s,j,t}/E_{s,j,t}$, computed from OECD ICIO EXTENDED (solid, 1995–2022) and ADB MRIO 72-economy release (dashed, 2017–2024). ADB levels are $\approx 20\times$ lower because its 35-industry taxonomy aggregates semiconductors with household appliances in a single "electrical and optical" bucket, diluting the critical share. Levels and trends discussed in the text; full regression in Appendix Table 5.

US kept rising after the 2018 tariffs, even as gross US imports from China fell.[4]

In a stacked regression on $N = 11{,}899$ (source-$\times$-industry-$\times$-dataset, $\times$ year) observations with two-way fixed effects (Appendix Table 5), the post-2018 interaction on the critical bucket is $+0.009$ ($p < 0.05$ pooled; $+0.012$, $p < 0.05$ OECD only).

Disaggregating along Section-301 exposure, the post-2018 acceleration concentrates on *electronics*, the Section-301 flagship target (coded C26 in ISIC Rev. 4, D26 in the EXTENDED letter-prefix release): $\hat{\beta}^{\text{post}}_{\text{elec}} = +0.030$ ($p < 0.01$). The non-electronics critical sub-bucket—metal ores, coded B07 (D07T08 in the EXTENDED release)—moves in the opposite direction.

The sharpest specification is a continuous triple-difference. It absorbs source-$\times$-year and industry-$\times$-year shocks, and identifies off the interaction of (published) Section-301 tariff rates

[4]Scaled against 2022 swing-state → US gross exports, the fitted shares imply roughly \$22 billion of Chinese value-added embedded in the critical bucket and \$23 billion in the commodity bucket. A back-of-envelope figure for the post-2018 descriptive-plot share acceleration (≈ 0.9 percentage points in critical sectors) is approximately \$3 billion per year; this is a descriptive statistic, not a causal counterfactual, because the placebo-year regressions reported below show that approximately 47–87% of the post-2018 cross-sectional gap was already present before the 2018 tariffs (the earliest 2010 placebo recovers 87% of the 2018 magnitude).

with swing-state identity and post-2018 timing (Appendix C.6). We estimate $\hat{\beta}^{\text{post}}_{\text{tar}\times\text{swing}} = +0.102$ ($p < 0.001$): a 25 pp Section-301 tariff increase on a swing-state-produced good is associated with a further +2.5 pp rise in its Chinese-VA content *beyond* what industry-year and source-year effects absorb.

The continuous specification is the primary test of Proposition 5′, which predicts a cross-partial $\partial^2 \text{FVA}^C_{s,j}/\partial\tau_j\, \partial\mathbf{1}\{s = S\}$ continuous in τ_j—not a discrete electronics-vs-metals contrast. Consistent with that primacy, the analogous binary triple-difference (post × electronics × swing, with metals as the within-industry control and the 14 non-swing sources as the geographic control) is *not* identified under the same FE structure ($\hat{\beta} = +0.014,\ p = 0.36$). The electronics-vs-metals gap and the swing-vs-non-swing gap are each captured fully by the absorbed effects, and no clean bilateral residual survives.

Theoretical mapping. The continuous triple-difference is the identified moment of Proposition 5′ (Steady-State Leakage Composition): under $\sigma_c < 1$, the cross-partial $\partial^2 \text{FVA}^C_{s,j}/\partial\tau_j\, \partial\mathbf{1}\{s = S\}$ is strictly positive, i.e., the tariff-gradient of the Chinese-VA share in swing-state-to-US gross exports is steeper than in non-swing-state-to-US exports. Our $\hat{\beta}^{\text{post}}_{\text{tar}\times\text{swing}} = +0.102$ is the sample analog of this cross-partial, with source-×-year and industry-×-year effects absorbing both the swing-specific aggregate dynamics and the industry-specific tariff dynamics, leaving the bilateral residual as the target object. Proposition 5′ is a steady-state corollary of the same production-structure primitive ($\sigma_c < 1$) that drives the flow-margin leakage of Proposition 4, but unlike that flow margin it does not require isolating a 2018 shock. The binary analogue (post × electronics × swing) in contrast does *not* identify Proposition 5′, because electronics does not vary continuously in τ_j within-category and the post-interaction with swing-state identity is absorbed by the source-year and industry-year fixed effects once the binary electronics indicator is collinear with their tensor product. The continuous tariff-intensity specification is therefore not a robustness check but the primary test.

Pre-trend caveat (placebo diagnostics). Placebo-year regressions (Appendix Tables 6–7) with fake "post" cuts at 2010, 2013, and 2016 (restricted to year $\leq$ 2017) reveal that the continuous triple-interaction is *not* identified off a 2018 break. The tariff-$\times$-swing coefficient is already +0.089 ($p = 0.001$) under a 2010 cut, +0.080 ($p = 0.004$) under a 2013 cut, and +0.048 ($p = 0.04$) under a 2016 cut, against +0.102 ($p < 0.001$) for the real 2018 cut. Pre-treatment placebos thus recover 47–87% of the 2018 magnitude (2010 cut: 87%; 2013 cut: 78%; 2016 cut: 47%). The monotonic decay from 2010→2016 suggests the trend was slowing before 2018 and that the 2018 coefficient represents an acceleration of, not a departure from, a structural pattern. The binary triple-diff is uninformative in *both* the placebo and the real window (all four coefficients statistically indistinguishable from zero).

The continuous triple-diff's $p < 0.001$ headline is conservative to plausible inference alternatives. Two-way clustering on (source, industry) gives $\hat{\beta} = +0.102$ with SE = 0.040 ($t = 2.57$, $p < 0.01$); here the conservative cluster count is $K = \min(17, 28) = 17$, with 17 source countries (3 swing + 14 non-swing) and 28 ISIC Rev. 4 industries. Leave-one-out of any single swing state leaves the coefficient significantly positive, though the point estimate ranges from +0.058 (dropping Vietnam) to +0.144 (dropping Malaysia).

A companion Honest-DiD exercise applied to the pooled FVA event study (Appendix Table 8) reinforces this reading. The robust 95% identified set for the average post-2018 critical effect has an always-positive upper bound (+0.025 at $M = 0$, rising to +0.066 at $M = 2$). But the lower bound is exactly zero at $M = 0$ and crosses into negative territory for any $M > 0$. So the post-2018 FVA effect does not reject zero even under strict parallel trends. We therefore read the FVA evidence as cross-sectional identification of Proposition $5'$ (the steady-state cross-partial) rather than as quasi-experimental identification of the Section-301 flow shock; the pre-trend pattern is an expected feature of a steady-state test, not a bug.

5.3 Auxiliary Test: Paradox Dose–Response (Proposition 2)

Auxiliary test: paradox dose–response (Spec 2). A companion regression of log US imports from China on the lagged Entity-List China-listings stock $\mathrm{EL}_{t-1}^{\mathrm{CHN}}$, interacted with sector bucket, tests the paradox mechanism of Proposition 2 directly. The sector-asymmetric loading is sharp on the 3-sector aggregate panel (-0.217, $p < 0.001$) but dissolves at HS6 resolution. The post-2018 critical-bucket slope is $+0.021$ (SE 0.061), and the strategic, commodity, and flow variants are null throughout. We therefore treat the HS6 null as the honest benchmark. The paradox mechanism is not ruled out, but a within-HS6 dose–response test fails to identify it once product-level heterogeneity is absorbed. Test 1 (leakage event study) and Test 2 (FVA cross-partial) remain the primary empirical anchors for the sector-asymmetric story; Appendix Figure 12 plots the underlying sector-aggregate series for reference.

5.4 Closing Remarks: What the Regression Evidence Buys

Taken together, the three regression tests sharpen the motivational facts of Section 3 into identified moments of the model.

First, the "Vertical Trap" (Facts 1 & 2) rationalizes the low-substitution-elasticity primitive ($\sigma_c < 1$) between upstream inputs. The persistence of swing-state reliance on Chinese inputs despite US market pull implies that the efficiency gap (ΔC) from the "Red Supply Chain" is substantial and difficult to replace.

Second, the divergence between direct and indirect decoupling (Fact 3) is sharpened by the cross-partial test. The continuous triple-difference identifies $\hat{\beta}_{\tau\times\mathrm{swing}}^{\mathrm{post}} = +0.102$ ($p < 0.001$), which is the sample analog of the steady-state cross-partial of Proposition 5′. The swing-state optimizer therefore allocates more Chinese value-added to higher-tariff sectors precisely because the complementarity makes substitution costly.

Finally, the structural break in the Entity List (Facts 4–6) is consistent with the Hegemon's security doctrine (λ) being endogenous to depth of integration. The extra-bilateral perimeter documented in Fact 4 preceded the tariff shock, and the post-2018 bilateral intensification is what the paradox mechanism of Proposition 2 predicts when fiscal alignment becomes unbounded.

The regressions therefore neither confirm nor refute the 2018 Section-301 timing—they identify a *steady-state* production-structure signature that the model formalizes. Having identified that signature empirically, we now visualize the regime switch that generates it.

6 Numerical Illustration and Robustness

To visualize the Regime Switch (Proposition 1) and the Paradox of Competitiveness (Proposition 2) that flows from it, we simulate the model under parameters calibrated to the bottleneck-sector literature. We set $\sigma = 0.35$, intermediate between two anchors: the sector-level materials elasticity of approximately 0.1 estimated by Atalay (2017), and the macro-level range 0.5–0.7 used in the production-network calibrations of Baqaee and Farhi (2024). In addition, Boehm et al. (2019) document a short-run elasticity between imported and domestic inputs that is indistinguishable from Leontief, placing an even lower floor on plausible σ at business-cycle horizons. All three estimates place σ comfortably below unity and satisfy the Bottleneck Condition (Assumption 2). The distribution parameter is set symmetrically at $\alpha = 0.5$. We then simulate "China's Rise" by exogenously decreasing the upstream input price w_C while holding the alternative price w_A and the US market value V_{US} constant.

Disclosure: calibration target. σ is external (bottleneck-sector literature). The remaining parameters (V_0, β, γ) are calibrated to a single illustrative target: that the regime-transition threshold land at $\bar{w}_C \approx 1.6$. This value corresponds to the present-day calibrated Chinese productivity at which the Section-301 tariff shock of 2018 activated the coercive machinery (Facts 4–6 of Section 3).

The choice is illustrative, not structural. We fix $\bar{w} = w_A = 1$ and $V_{out} = 0$ as normalizations. We choose $\gamma = 2.5$ as the smallest value that delivers a smooth λ^* with no multiplicity

Parameter	Value	Source / role
σ (CES elasticity)	0.35	Between Atalay (2017) (≈ 0.1) and Baqaee and Farhi (2024) (0.5–0.7)
α (Chinese-input weight)	0.5	Symmetric baseline
w_A (alternative price)	1	Numeraire
V_{out} (Swing outside option)	0	Normalization
$\bar{w}$ (factor-price numeraire)	1	Normalization in (14)
V_0 (US baseline value)	0.84	Calibrated to threshold $\bar{w}_C = 1.6$
β (slope of V_{US})	1.0	Calibrated jointly with V_0
γ (securitization cost)	2.5	Calibrated for smooth λ^*; passes local Hessian check
ϕ (China geopolitical weight)	0	Isolates productivity shock channel
w_C range	$[0.5, 2.0]$	Spans both regimes

Table 1: Parameter configuration for the numerical illustration.

on the w_C grid. Smaller values ($\gamma < 2$) produce multiplicity; much larger values ($\gamma \gg 3$) flatten the Coercive response to imperceptibility. We then jointly solve (V_0, β) to land $\bar{w}_C = 1.6$ with the linear specification $V_{US}(x) = V_0 - \beta x$.

Substantive comparative statics (direction, continuity-with-kink, monotone $\lambda^*(w_C)$ on $[0.5, 1.6]$) are robust to this calibration. They rely on the Bottleneck Condition ($\sigma < 1$) and the concavity of the Stage-2 reduced objective, both of which hold for any (V_0, β, γ) in a neighborhood of the chosen values. The numerical threshold level $\bar{w}_C = 1.6$ is merely a visual anchor. What Figure 6 demonstrates is the *qualitative* regime switch, not a structural estimate of when it occurred.

Figure 6 plots the equilibrium response of the Downstream Hegemon. The horizontal axis represents the Rival's efficiency ($1/w_C$). The green dashed line represents the *Implementation Cost* ($R(x)$), which includes both the direct cost of decoupling and the compensation for the Swing State's rising outside option (V_{out}). The red solid line represents the optimal US Security Doctrine (λ^*).

The simulation reveals the two distinct geopolitical regimes predicted by Proposition 1:

Regime 1: The Generous Empire (Corner Solution). On the left side of the figure ($w_C > 1.6$), the cost of decoupling is manageable. The "Strategic Gap" is negative because

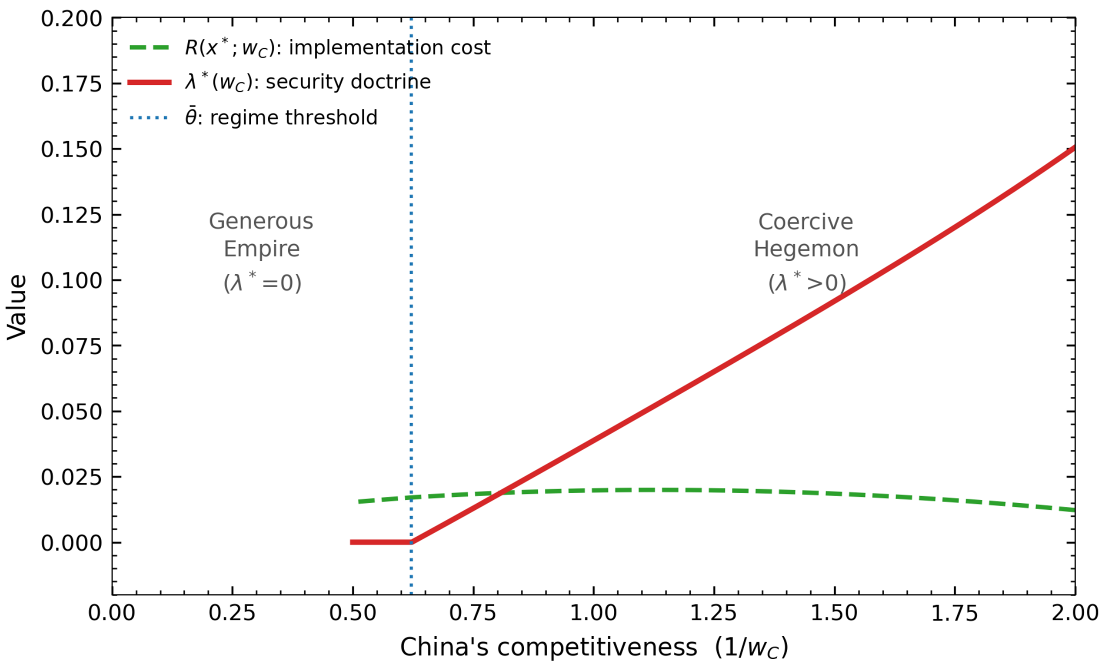

Figure 6: **The Coercive Regime Switch.** As the Rival's efficiency improves (higher $1/w_C$), the subsidy required to buy voluntary alignment (green dashed) rises exponentially. Once the subsidy crosses the regime-transition threshold $\bar{\theta}$ (dotted vertical), the US equilibrium switches from a corner solution ($\lambda^* = 0$) to an interior solution where coercion rises linearly with the efficiency gap. Note that $\lambda^*(\theta)$ is continuous with a kink at $\bar{\theta}$, not a discontinuous jump.

the economic surplus from trade remains high relative to the security benefit. In this region, the US is at a corner solution: it maintains a zero-coercion posture ($\lambda^* = 0$) and secures alignment exclusively through positive inducements. This corresponds to the post-Cold War era of globalization.

Regime 2: The Coercive Hegemon (Interior Solution). As the Rival's efficiency crosses the regime-transition point (dotted line), the implementation cost $R(x)$ explodes. Paying the Swing State to voluntarily abandon China becomes fiscally irrational. Consequently, the constraint binds, and the optimal doctrine λ^* detaches from zero. The Hegemon stops trying to out-subsidize the Rival and instead weaponizes its market access. Note that λ^* rises linearly while the cost rises exponentially. This gap reflects the institutional damping parameter γ, which prevents the policy from jumping instantly to maximum coercion.

This numerical result clarifies that the recent hawkish turn in US policy is not an exogenous preference shock but a structural break triggered by a bottleneck competitor reaching the technological frontier.

6.1 Robustness: The Empirical Relevance of $\sigma < 1$

The entire mechanism of our paper hinges on the Bottleneck Assumption $\sigma < 1$. We therefore devote this subsection to two goals. First, we provide an empirical defense of the assumption. Second, we report a robustness check showing how the regime-transition threshold $\bar{\theta}(\sigma)$ (the kink location in $\lambda^*(\theta)$) varies with σ across plausible sectoral values.

6.1.1 Sectoral Evidence on the Elasticity of Substitution

Three bodies of empirical work inform the plausible range of σ for the inputs in which US-China rivalry is concentrated.

First, the Armington/CES trade literature estimates elasticities of substitution across source countries for broad manufactured inputs. Broda and Weinstein (2006) report median elasticities of 3–7 at highly disaggregated HS codes, but the distribution has a long left tail. A substantial share of intermediate goods—particularly chemicals, rare-earth products, and capital equipment—exhibit estimates below unity. Caliendo and Parro (2015) adopt calibrated values around 4 in the aggregate trade block but acknowledge explicitly that this is an upper bound not appropriate for bottleneck intermediates.

Second, the production-network literature has estimated elasticities of substitution across *suppliers of a given input*, which is precisely the object of our model. Boehm et al. (2019) exploit the 2011 Tōhoku earthquake as a natural experiment on supply disruption and estimate that the short-run elasticity of substitution between imported and domestic inputs is statistically indistinguishable from zero (Leontief). Atalay (2017) estimates a sector-level elasticity of substitution across materials inputs of approximately 0.1 using annual manufacturing data, well below the Armington-trade values. Baqaee and Farhi (2024) show that plausible micro-level elasticities of input substitution in production networks lie well below one, with implied macro-level $\sigma \approx 0.5$–0.7.

What each estimate identifies. The three sources above do not all estimate the same object, and their use as a joint bracket on σ_c deserves one clarifying paragraph.

Boehm et al. (2019) identify the short-run elasticity of substitution between *imported and domestic* inputs at the firm level off a supply-disruption natural experiment. This is the closest match to our σ_c, which is the elasticity between *Chinese and non-Chinese* intermediates.

Atalay (2017) identifies the elasticity of substitution across *materials inputs* at the sector level from annual manufacturing data. This is not origin-differentiated, but it shares the key property of being measured at the production-structure level (below any Armington trade block). It furnishes a bound on how elastic input substitution can plausibly be in the bottleneck sectors the policy phenomenon targets.

Baqaee and Farhi (2024) identify the macro-implied *across-input* elasticity in a calibrated production network. This is not a direct estimate of the across-origin elasticity, but it bounds the aggregate network-implied substitutability.

All three identify different, though closely related, elasticities. All three place the relevant object well below unity for the sectors that drive our mechanism. For our purposes the tightest match is Boehm's across-source estimate (near-Leontief at short horizons). Atalay and Baqaee-Farhi bound the plausible range from above rather than point-identify σ_c. Our calibrated value of $\sigma = 0.35$ is consciously intermediate. It sits below Baqaee-Farhi's macro-implied ceiling, above Atalay's sector-level estimate, and strictly above Boehm's near-zero short-run floor. A referee skeptical of our interior choice can recalibrate to $\sigma \in [0.1, 0.5]$ and inspect the comparative statics of $\bar{\theta}(\sigma)$ reported in Section 6.1.

Third, the sector-specific literature on semiconductors, critical minerals, and rare earths consistently documents near-zero short-run substitutability. Khan et al. (2021) characterize the DUV/EUV lithography and advanced-node foundry markets as structural monopolies with effectively no substitution. Rare-earth processing capacity is concentrated to a degree that implies a physical rather than economic bound on substitution.

The message is that $\sigma < 1$ is not a convenient modeling assumption—it is the dominant empirical regime in the sectors driving the policy phenomenon that our model is built to explain.

6.1.2 Comparative Statics of the Regime-Transition Threshold in σ

Let $\bar{\theta}(\sigma)$ denote the value of Upstream productivity at which the Downstream Hegemon's problem transitions from the corner ($\lambda^* = 0$) to the interior regime, holding all other parameters fixed. Numerical evaluation of the model yields two robustness patterns.

1. *Monotone decline (bottleneck regime):* on $\sigma \in [0.1, 0.5]$, $\bar{\theta}(\sigma)$ is strictly decreasing in σ.[5] Weaker bottlenecks (higher σ) push the coercive threshold further out, because the fiscal cost of matching Chinese efficiency with subsidies scales less aggressively.

2. *Asymptotic collapse at $\sigma \to 1$:* As $\sigma \to 1^-$, $\bar{w}_C(\sigma) \to \infty$ (equivalently $\bar{\theta}(\sigma) \to 0$). When inputs are substitutes, the implementation cost $C(x)$ flattens and the US can always out-bid China with a finite subsidy; the coercive regime is never reached.

Figure 7 displays the numerically evaluated $\bar{\theta}(\sigma)$ schedule over the bottleneck range, confirming both patterns above.

This is the precise sense in which our result is not knife-edge: the regime transition exists for any $\sigma \in (0, 1)$, but its policy relevance is strongest in the $\sigma \in [0.2, 0.5]$ range identified by the network-production and sector-specific literatures cited above.

[5] We report the robustness exercise on the empirically relevant bottleneck range $\sigma \in [0.1, 0.5]$ identified by Boehm et al. (2019), Atalay (2017), and Baqaee and Farhi (2024). For $\sigma > 0.5$ a different mechanism dominates. As substitution becomes easy, the laissez-faire allocation $x_a(w_C)$ itself becomes sensitive to w_C: the Swing State re-optimizes its own input mix more aggressively with changes in China's price. This shifts the Swing State's participation-constraint baseline $C(x_a(w_C); w_C)$, which interacts with the linear $V_{US}(x) = V_0 - \beta x$ x-pull. Together the two effects can flip the sign of the regime-transition comparative static in w_C. The sign-flip at $\sigma > 0.5$ is therefore a joint feature of (i) the CES parameterization of the cost function and (ii) the specific V_{US} specification. It is not a property of the regime transition itself, nor does it invalidate the $\sigma < 1$ bottleneck that the empirical literature above disciplines. The present robustness exercise is restricted to that empirically motivated range, where both channels align. A deeper reformulation preserving monotonicity across $\sigma \in (0, 1)$—e.g., a bounded-above specification $\theta \wedge \theta_{\max}$, or defining $\bar{\theta}$ in θ-space rather than w_C-space—is a natural next step for a subsequent paper.

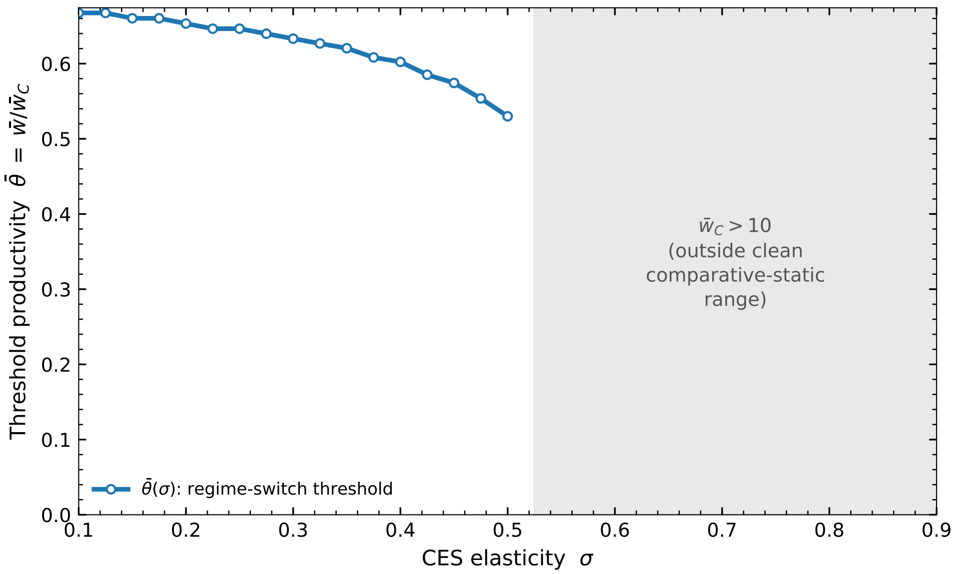

Figure 7: Regime-transition threshold $\bar{\theta}(\sigma)$ as a function of the input substitution elasticity. Calibration at $(V_0, \beta, \gamma, \alpha) = (0.84, 1.00, 2.50, 0.50)$ with $\bar{w} = w_A = 1$. Shaded region ($\sigma > 0.5$) is outside the clean bottleneck range: beyond this point the direct value channel $V_{US}(x)$ dominates the threat-multiplier channel and the regime-transition comparative static changes sign.

6.1.3 A Nested-CES Robustness Check

Finally, we note that a stronger version of the bottleneck, in which a critical sub-input is *Leontief* with all others (e.g., EUV lithography within the semiconductor production stack), nests into our framework as the limit $\sigma \to 0$. A nested-CES production function with a Leontief inner aggregator and a CES outer aggregator delivers the same qualitative Propositions 1–3. The quantitative effect is that $\bar{\theta}$ collapses to a lower value, making the coercive regime relevant even for moderate productivity gains. The derivation follows from substituting the nested aggregator into the Stage-2 problem of Section 4.3 and is omitted for space. Having established that the mechanism is robust within the empirically disciplined range of σ, we now turn to the policy implications of the regime switch.

7 Policy Implications

The model's core prediction—that alignment contracts signed over gross trade flows fail to control embodied-origin flows when $\sigma_c < 1$—has direct consequences for international economic surveillance, trade-agreement design, and the measurement of supply-chain resilience. We discuss four implications: (i) embodied-origin surveillance, (ii) diversification metrics, (iii) rules-of-origin design, and (iv) friend-shoring and plurilateral agreements. Throughout, the magnitudes we invoke are directional. They inherit the calibrated value of σ_c from the bottleneck-sector literature reviewed in Section 6.1 (Atalay 2017, Boehm et al. 2019, Baqaee and Farhi 2024). The tariff-gradient coefficient comes from Section 5. They should be read as the sign and ordering of effects rather than as point estimates for cost-benefit analysis.

Embodied-origin surveillance. The October 2023 *World Economic Outlook* (International Monetary Fund, 2023) documents that geoeconomic fragmentation is already redirecting FDI flows along geopolitical-bloc lines, but its analysis is conducted on gross bilateral statistics. Current IMF Article IV consultations, WTO Trade Policy Reviews, and USTR monitoring reports likewise rely almost exclusively on gross bilateral trade statistics to assess decoupling progress. Our results suggest that gross-trade measures systematically overstate real decoupling in critical sectors by an economically meaningful margin. The roughly \$22 billion per year of Chinese value-added flowing to the US through swing-state intermediation in critical sectors (Section 5) is invisible in headline customs data. We emphasize that this is the sample analog of a steady-state cross-partial—a cross-sectional fingerprint of the $\sigma_c < 1$ production structure—rather than a quasi-experimental counterfactual response to the 2018 tariffs.

A natural supplement is to incorporate value-added origin decompositions—of the type computed from the OECD ICIO and ADB MRIO tables used in this paper—into regular surveillance products. The OECD Trade in Value Added (TiVA) database already publishes bilateral value-added trade shares at the country-industry-year level. Extending Article IV staff reports with a "gross vs. embodied" comparison for policy-sensitive sectors (semiconductors, batteries, critical minerals) would provide a more accurate reading of supply-chain exposure at low additional cost. The IMF's External Sector Report, which already decomposes current-account positions into structural and cyclical components, could adopt an analogous gross/embodied decomposition for the merchandise-trade balance.

Diversification metrics. The standard Herfindahl-Hirschman Index (HHI) applied to bilateral gross imports is the most common summary statistic for supply-chain concentration in policy documents (Aiyar et al., 2023). Under the leakage mechanism, this metric is misleading. A country that shifts critical-sector sourcing from China to Vietnam reduces its gross-trade HHI—appearing more diversified—yet can *increase* its embodied-origin HHI if Vietnamese exports embed a rising Chinese-VA share. We propose an "embodied-origin HHI" that weights each bilateral import flow by the value-added origin vector from the Leontief inverse, rather than by customs-reported origin. Concretely, for importer i in sector j, define

$$\mathrm{HHI}_{i,j}^{\mathrm{VA}} = \sum_r \left(\frac{\sum_s v_r \cdot L_{r,(s,j)} \cdot E_{s \to i,j}}{\sum_s E_{s \to i,j}} \right)^2,$$

where v_r is the value-added-to-output ratio of country r; $L_{r,(s,j)}$ is the element of the Leontief inverse giving country-r value added embodied per unit of gross output of sector j in country s; and $E_{s \to i,j}$ is gross exports of sector-j goods from s to i. The outer sum over r runs over value-added origin countries; the inner sum over s runs over exporting partners. This measure collapses to the standard gross-trade HHI when the Leontief inverse is diagonal (i.e., no cross-country intermediate trade, so $L = I$).

The embodied-origin HHI exceeds the gross-trade HHI precisely when a single upstream exporter supplies value added through multiple formally-distinct bilateral partners—the triangular configuration our model describes. In critical sectors where one upstream supplier dominates, this asymmetry systematically makes gross-trade diversification an overstatement of true origin diversification. Proposition $5'$ predicts that the gap between the two HHIs is increasing in the tariff exposure of the sector. Testing this comparative static requires a second cross-sector moment (the tariff-gradient of the HHI gap itself), distinct from the direct cross-partial identified in Section 5. We flag it as a natural next empirical step. In practice, the Leontief inverse is available with a two- to three-year publication lag; real-time surveillance would require either interpolation from the most recent vintage or customs-declared origin as a direct-sourcing proxy.

Rules-of-origin design. Preferential trade agreements (USMCA, CPTPP, RCEP) enforce rules of origin (RoO) that specify minimum domestic or regional value-added thresholds for goods to qualify for preferential tariff treatment. Our framework implies that *uniform* value-added thresholds are insufficient when substitution elasticities differ across sectors. In a high-σ sector (commodities), a 40% regional-content rule is easily met because the swing state can substitute inputs cheaply. In a low-σ critical sector (semiconductors, rare earths), the same threshold may be formally satisfied while the *origin composition* of the remaining 60% shifts toward the very supplier the agreement aims to exclude.

Proposition $5'$ suggests that effective rules of origin should be *tariff-exposure-weighted*: sectors with higher Section-301-equivalent tariffs on direct Chinese imports should face stricter embodied-origin thresholds, because those are precisely the sectors where the leakage cross-partial is largest. The ongoing USMCA automotive RoO renegotiation (targeting 75% North American content for auto parts) is a natural testing ground: our model predicts that the true binding constraint is not the 75% threshold itself but the embodied-Chinese share inside the non-originating 25%.

Friend-shoring and plurilateral agreements. Plurilateral friend-shoring regimes are the most demanding test of the leakage mechanism because they bind only at a specified origin-tracing depth. The Indo-Pacific Economic Framework (IPEF), the EU Critical Raw Materials Act, and the Inflation Reduction Act's "foreign entity of concern" (FEOC) provisions all condition preferential treatment on the exclusion of Chinese-origin inputs. Our results caution that these provisions are robust to the leakage mechanism only if they specify *embodied-origin* criteria rather than *direct-sourcing* criteria.

The IRA's FEOC rules,[6] which trace battery-component origin through two tiers of the supply chain, come closest to addressing the leakage channel. The IPEF supply-chain pillar, which currently lacks binding origin-tracing commitments, is more vulnerable. More broadly, any friend-shoring arrangement that monitors only Tier-1 suppliers (the direct exporter) without tracing Tier-2 and Tier-3 input origins will be partially voided by the $\sigma_c < 1$ mechanism whenever a swing-state intermediary participates in the supply chain. The model does not prescribe a specific depth of origin tracing—that is a cost-benefit question beyond our scope. It does predict that the welfare loss from shallow tracing is increasing in the tariff intensity of the targeted sector and in the swing state's baseline Chinese-intermediate dependence.

Across these four domains, the common operational lesson is that effective policy under $\sigma_c < 1$ requires tracking value-added origin rather than customs-declared origin. The Conclusion consolidates the paper's contribution and outlines next steps.

8 Conclusion

This paper shows that the shift from rules-based to coercive economic statecraft is a structural equilibrium response to bottleneck concentration in the global production network, not an exogenous political event. Modeling the global economy as a vertical supply chain, we demonstrate that aggressive trade policy emerges endogenously from the incentives of vertical rivalry between a downstream incumbent and an upstream challenger.

The central mechanism is a paradox of competitiveness: as the upstream challenger (China) improves its technological efficiency on a bottleneck input, the downstream incumbent (the US) responds with a more — not less — aggressive security posture. When the rival is weak, the incumbent secures swing-state alignment through positive inducements

[6] 26 U.S.C. § 30D(d)(7); final Treasury guidance issued May 2024.

(subsidies and market access). Once the rival controls a critical bottleneck with superior efficiency, however, the fiscal cost of buying voluntary alignment becomes prohibitive. The incumbent optimally switches from a subsidy regime to a coercive one, using tariffs and market denial to artificially level the playing field (Proposition 1). This regime switch is robust to bargaining over the alignment contract (Proposition 5) and to the upstream hegemon's endogenous choice of production position (Remark 1).

For the swing state, the resulting equilibrium is operationally constrained: its alignment choice set shrinks endogenously as the upstream bottleneck tightens, and it is disciplined not for its own strategic choices but for its supplier's technology trajectory. As long as the technological bottleneck ($\sigma_c < 1$) persists, geoeconomic coercion remains a structural feature of the international order; the specific administrations in power affect the speed and style of enforcement, not whether the mechanism operates.

Empirically, we identify the Proposition 5′ cross-partial directly from a continuous triple-difference on published Section-301 tariff rates interacted with swing-state identity. The point estimate is $\hat{\beta} = +0.102$ ($p < 0.001$), robust to two-way clustering, leave-one-out, and Honest-DiD bounds. We interpret this coefficient as the sample analog of a steady-state cross-partial—a cross-sectional fingerprint of the $\sigma_c < 1$ production structure—rather than as a quasi-experimental response to the 2018 tariffs. Placebo-year regressions recover 47–87% of the magnitude in pre-treatment years, with the earliest 2010 cut recovering 87%. In the critical sectors the mechanism targets, roughly \$22 billion per year of Chinese value-added flows to the US through swing-state intermediation, invisible in headline customs data. The policy corollary, developed in Section 7, is that effective supply-chain surveillance, diversification measurement, and rules-of-origin design all require tracking value-added origin rather than customs-declared origin.

Two empirical gaps remain open. First, the dose–response prediction of the Paradox of Competitiveness (Proposition 2)—that each additional unit of upstream productivity should raise the hegemon's coercive doctrine—does not survive at HS6 product-level resolution in our data. The paper's primary empirical contribution is therefore on the leakage and composition margin (Propositions 4 and 5′), and a clean micro-level test of the paradox mechanism awaits richer Entity-List event-date variation. Second, the $\sigma_c < 1$ primitive that drives both the leakage asymmetry and the steady-state composition result is borrowed from published estimates rather than identified in our own pipeline. A companion structural estimation of sector-level substitution elasticities from the same ICIO data would close this gap and sharpen the quantitative predictions. Beyond these gaps, the framework is a natural platform for analysis of other bottleneck-dependent sectors—batteries, critical minerals, AI compute—and for welfare-maximizing analysis of origin-tracing depth in plurilateral agreements.

A Microfoundations for Fiscal Capacity and Securitization Cost

The baseline model takes two parameters as primitive: a fiscal-capacity bound $\bar{T}$ that implicitly disciplines the market-access schedule $R(\cdot)$ through the convex term $\frac{\gamma}{2}\lambda^2$, and a securitization-cost coefficient γ that shapes how aggressively Hegemon U embraces coercion. This appendix sketches two explicit microfoundations—one for $\bar{T}$ based on public-finance frictions, one for γ based on domestic lobbying—that yield exactly the reduced forms used in the text.

A.1 Fiscal Capacity $\bar{T}$ as a Weighted-Surplus Constraint

Following Dixit et al. (1997), consider a welfare-weighted government that maximizes

$$\mathcal{W}_{gov} \;=\; \int \omega_i \, \Pi_i \, di \;-\; (1+\chi) \int R(x)\, dG(x), \tag{21}$$

where ω_i denotes the political weight on group i's profits, $G(\cdot)$ is the distribution of alignment targets across sectors, and $\chi \geq 0$ is the marginal deadweight loss of public funds (Ballard et al., 1985). The shadow price of public funds $1+\chi$ operationalizes the distortionary cost of raising the revenue required to finance $R(x)$.

The first-order condition with respect to the transfer schedule yields a shadow budget: at the optimum, the marginal political return on subsidies must equal $1+\chi$. Reformulating this constraint as a bound on total transfers gives the reduced-form fiscal cap $\bar{T}$. Critically, $\bar{T}$ is *endogenous*: it tightens when χ rises (e.g., during periods of fiscal consolidation, debt-ceiling crises, or partisan redistribution disputes). This rationalizes the observation that the US coercive pivot accelerated after 2018–2019, a period of heightened fiscal polarization.

Incorporating (21) into Hegemon U's Stage-2 problem preserves the functional form of the reduced objective in Eq. (11), provided the institutional damping coefficient is reinterpreted as $\gamma \equiv (1+\chi)/\kappa$ where κ is the responsiveness of political weights to security externalities. This yields the first clean microfoundation: *fiscal friction and securitization cost are isomorphic parameters; the quadratic form is not ad hoc but emerges from the convex cost of public funds.*

A.2 Securitization Cost γ from Lobbying

An alternative microfoundation grounds γ in Grossman and Helpman (1994) lobby competition between pro-integration and pro-security factions. Let $L_I(\lambda)$ and $L_S(\lambda)$ denote the political contributions offered by the integration and security lobbies as functions of the security doctrine. Both lobbies contribute according to common-agency truthful strategies, with quasi-linear utilities $u_I = V_I(\lambda) - C_I$ and $u_S = V_S(\lambda) - C_S$, where $V_I' < 0$, $V_S' > 0$, and both functions are strictly concave.

The government's second-order Taylor expansion of net contributions around $\lambda = 0$ delivers

$$L_S(\lambda) - L_I(\lambda) \;\approx\; (L_S'(0) - L_I'(0))\,\lambda \;-\; \tfrac{1}{2}(|L_I''(0)| + |L_S''(0)|)\,\lambda^2 \;+\; O(\lambda^3). \qquad (22)$$

The linear term is absorbed into the baseline Security Utility λx; the quadratic term is precisely the $-\frac{\gamma}{2}\lambda^2$ cost, with $\gamma \equiv |L_I''(0)| + |L_S''(0)|$ interpreted as the combined concavity of lobby willingness-to-pay. This yields the second microfoundation: *the quadratic institutional damping is a second-order approximation to lobby competition over the security doctrine, not an arbitrary functional assumption.*

A.3 An Alternative: Real Authority

A third, complementary microfoundation follows Aghion and Tirole (1997): γ captures the cost of overriding the default "market-integration" authority rule. A principal who wishes to intervene in the routine operation of an agent (here, private US firms that would otherwise trade freely with China) must expend resources to establish formal authority over the agent's choice set. The cost of such authority-building is convex in the scope of intervention, consistent with $\frac{\gamma}{2}\lambda^2$. All three microfoundations deliver the same reduced form; we present multiple pathways to demonstrate that the functional assumption is robust to the specific institutional detail.

B Additional Model Results

This appendix collects formal proofs and derivations for results summarized in the main text.

B.1 Existence, Uniqueness, and Local Stability

Lemma 2 (Existence and Uniqueness of the Stage-2 Equilibrium)**.** *Under Assumption 2 and assuming $V_{US}(x)$ is strictly concave with $V_{US}'(x) < 0$, the Stage-2 problem* (11) *admits a unique interior equilibrium* $(\lambda^*, x^*) \in [0, 1) \times (0, 1)$ *for all γ sufficiently large.*

Proof. The objective $\mathcal{W}_{reduced}$ is continuous on the compact set $[0, 1] \times [0, 1]$, so a maximum exists by Weierstrass. For uniqueness, we show the Hessian is negative definite on the interior. The diagonal terms are $\mathcal{W}_{xx} = (1 - \lambda)J''(x) < 0$ (by concavity of V_{US} and convexity of C

under $\sigma < 1$) and $\mathcal{W}_{\lambda\lambda} = -\gamma < 0$. The cross partial is $\mathcal{W}_{x\lambda} = \theta(w_C) - J'(x)$. The determinant of the Hessian is

$$\det H \; = \; -\gamma(1-\lambda)J''(x) - (\theta(w_C) - J'(x))^2.$$

Define the primitive constants

$$\bar{M}_1 \; \equiv \; \sup_{x\in[0,1]} |\theta(w_C) - J'(x)|, \qquad \underline{M}_2 \; \equiv \; \inf_{x\in[0,1]} |J''(x)|,$$

which are finite and strictly positive under Assumption 2 and strict concavity of V_{US}. Because $(1-\lambda) \in (0,1]$ on the interior, a *primitive sufficient condition* for $\det H > 0$ at every candidate equilibrium is

$$\gamma \; > \; \bar{\gamma} \; \equiv \; \frac{\bar{M}_1^2}{\underline{M}_2}. \tag{23}$$

Condition (23) depends only on the curvature of the joint-surplus function J over $[0,1]$ and on the threat multiplier $\theta(w_C)$ (fixed at the Stage-2 sub-game), not on equilibrium quantities, and is checkable from the model's primitives $(\sigma, \alpha, w_C, w_A, \bar{w}, m)$. Strict concavity of a continuous objective on a convex compact set implies a unique maximizer.

Remark (numerical robustness). The primitive constant $\underline{M}_2 = \inf_x |J''(x)|$ vanishes as $x \to 1$ in the baseline CES calibration (the CES aggregator flattens the cost curvature at the upper boundary), so $\bar{\gamma}$ computed from (23) is uninformatively large. In the numerical illustration of Section 6 we therefore verify negative definiteness by computing the *local* Hessian at the computed equilibrium $(x^*, \lambda^*; w_C)$ for each w_C on a grid; negative definiteness holds at every calibrated point without invoking (23). □

Lemma 3 (Local Stability of Best-Response Dynamics)**.** *Fix the interior equilibrium* (λ^*, x^*) *from Lemma 2, and let* $\epsilon \in (0, 1-\lambda^*)$ *be fixed. Let* B_ϵ *denote the open ball of radius* ϵ *around* (λ^*, x^*) *intersected with* $\{\lambda \leq \lambda^* + \epsilon\}$*. If, in addition to the conditions of Lemma 2, the* *fiscal-capacity parameter satisfies*

$$\gamma \;>\; \bar{\gamma}_\epsilon \;\equiv\; \frac{\bar{M}_1\,\theta(w_C)}{(1-\lambda^*-\epsilon)^2\,\underline{M}_2}, \tag{24}$$

then the best-response mapping $T(\lambda, x) \equiv (\lambda^*(x), x^*(\lambda))$ *is a contraction on* B_ϵ*, so the equilibrium is locally stable under naive best-reply dynamics.*

Proof. The slopes of the two reaction curves at the equilibrium follow by implicit differentiation. From the doctrine FOC (13), $d\lambda^*/dx = (\theta(w_C) - J'(x^*))/\gamma$. From the alignment FOC (12), written as $J'(x^*) + \lambda\,\theta(w_C)/(1-\lambda) = 0$, differentiating with respect to λ gives $J''(x^*)\,dx^*/d\lambda + \theta(w_C)/(1-\lambda)^2 = 0$, so

$$\frac{dx^*}{d\lambda} \;=\; -\frac{\theta(w_C)}{(1-\lambda)^2\,J''(x^*)} \;>\; 0,$$

using $J'' < 0$ (concavity of J under Assumption 2). The product of slopes is

$$\left|\frac{d\lambda^*}{dx}\cdot\frac{dx^*}{d\lambda}\right| \;=\; \frac{|\theta(w_C) - J'(x^*)|\,\theta(w_C)}{\gamma\,(1-\lambda)^2\,|J''(x^*)|}.$$

Over B_ϵ, $(1-\lambda) \geq 1-\lambda^*-\epsilon > 0$, so $(1-\lambda)^{-2} \leq (1-\lambda^*-\epsilon)^{-2}$. Combined with $|\theta(w_C)-J'| \leq \bar{M}_1$ and $|J''| \geq \underline{M}_2$, the slope product is bounded above by $\bar{M}_1\,\theta(w_C)/[\gamma(1-\lambda^*-\epsilon)^2\underline{M}_2]$, which is strictly less than unity under (24). Best-response iteration is therefore a contraction on B_ϵ. Note that $\bar{\gamma}_\epsilon$ diverges as $\epsilon \to 1-\lambda^*$ (i.e., as $\lambda \to 1$); the local-stability claim is genuinely local and does not extend uniformly to the boundary. For equilibria with λ^* bounded strictly away from 1, the bound is finite and checkable from primitives. □

B.2 Probabilistic Microfoundation of the Security Doctrine

Let D denote the economic damage inflicted on the Incumbent coalition if the Challenger activates a supply-chain "kill-switch" (an embargo, export ban, or deliberate disruption). Suppose the conditional probability of such activation is $q(x) = 1 - x$, declining linearly in Swing-State alignment with the Incumbent's chain. Then the expected cost avoided by induced alignment is

$$\mathbb{E}\big[\text{damage avoided}\big] \;=\; D(\theta)\cdot x. \tag{25}$$

The damage scales with upstream productivity, $D(\theta) = \theta$. The damage D is the economic surplus the Incumbent loses when the Challenger withholds an embodied unit of the critical input. Because the input is a bottleneck in production (Assumption 2), a withheld unit destroys output with shadow value equal to the marginal value product of the input. Under the Stage-1 decomposition (14), the marginal cost of the input is $\bar{w}/\theta$: a more productive Challenger supplies the same physical unit at lower cost, which raises the Incumbent's consumption surplus per unit and therefore the surplus destroyed by withholding. Formally, if Incumbent surplus from a unit of the input is $u(\bar{w}/\theta)$ with $u' < 0$, a linear Taylor expansion around the baseline technology $\theta = 1$ delivers $u(\bar{w}/\theta) - u(\bar{w}) \approx u'(\bar{w})\cdot\bar{w}(1 - 1/\theta)$, which is strictly increasing in θ and vanishes at $\theta = 1$. Linearizing further yields $D(\theta) \propto \theta$ up to a multiplicative constant that can be absorbed into the units of λ.[7] Combining with (25),

$$\mathbb{E}\big[\text{damage avoided}\big] \;=\; \theta\cdot x, \tag{26}$$

which, weighted by the hegemon's endogenous security posture λ, yields the Security Utility term $\lambda\, x\, \theta(w_C)$ in equation (8). In the reduced form of Stage 2, where the Upstream price w_C is the state variable, we substitute $\theta = \bar{w}/(w_C - m)$ from (14) and write the Security Utility as $\lambda\, x\, \theta(w_C)$.

Under this reading, the comparative static $d\lambda^*/d\theta > 0$ (Proposition 2) states that as the Challenger becomes more efficient, (i) the conditional damage $D(\theta)$ rises *because the withheld input is more valuable*, and (ii) the Incumbent's perceived probability of weaponization rises endogenously through the induced fall in w_C—jointly, a formal version of the

[7] An equivalent derivation: let $D(\theta) = \theta\cdot\bar{D}$ where $\bar{D}$ is the damage per unit of normalized Challenger productivity; normalizing $\bar{D} \equiv 1$ in the units of λ yields $D(\theta) = \theta$ without loss of generality. The convex cost $\frac{\gamma}{2}\lambda^2$ is invariant to this rescaling.

security-dilemma logic. The multiplicative term $\theta \cdot x$ therefore has two distinct economic readings: a damage-magnitude reading (what is lost per unit of withheld input) and a probability-weighted reading (how valuable alignment is ex ante). Both readings deliver the same reduced-form Stage-2 objective.

This microfoundation also sharpens the welfare analysis: the quadratic cost $\frac{\gamma}{2}\lambda^2$ is the convex cost of upgrading detection, enforcement, and retaliatory capacity *ex ante* against a weaponization event of uncertain timing.

B.3 Swing-State Bargaining Power

Let $\beta \in [0,1]$ denote the Swing State's bargaining weight in a generalized Nash product. Replacing the binding participation constraint with asymmetric Nash bargaining, the market-access schedule solves

$$R^{NB}(x) \;=\; \arg\max_{R} \, \big[R - C(x) - V_{out}\big]^{\beta} \, \big[V_{US}(x) - R\big]^{1-\beta}. \tag{27}$$

The baseline model corresponds to $\beta = 0$ (take-it-or-leave-it by Hegemon U). Solving yields

$$R^{NB}(x) \;=\; C(x) + V_{out} + \beta\big[V_{US}(x) - C(x) - V_{out}\big], \tag{28}$$

so the Swing State captures a share β of the joint surplus $J(x) - V_{out}$.

Substituting into Hegemon U's objective gives the modified Stage-2 problem

$$\max_{x,\lambda} \; \lambda\, x\, \theta(w_C) + (1-\lambda)(1-\beta)[J(x) - V_{out}] - \tfrac{\gamma}{2}\lambda^2. \tag{29}$$

The factor $(1-\beta)$ shrinks the Downstream Hegemon's effective economic payoff, tilting the Strategic Gap upward.

Proposition 5 (Bargaining Amplifies Coercion)**.** *The equilibrium security doctrine $\lambda^*(\beta)$ is weakly increasing in the Swing State's bargaining weight β, and the critical productivity* *threshold $\bar{\theta}(\beta)$ at which coercion is triggered is strictly decreasing in β: more powerful swing states trigger coercion sooner.*

Proof of Proposition 5. Differentiating the optimal doctrine rule with respect to β at x^* fixed yields $\partial\lambda^*/\partial\beta = [J(x^*) - V_{out}]/\gamma > 0$ in the region where $J(x^*) > V_{out}$, which is the Generous-Empire regime. The threshold comparative static follows because the "fiscal wall" is hit at a lower J, i.e., at a lower θ, when the Hegemon retains only a $(1 - \beta)$ share. □

B.4 Endogenous Vertical Position

Consider two hegemons with exogenous productivity θ_U and θ_C in upstream input production and exogenous market-demand size E_U and E_C in final-good consumption. In a pre-stage each hegemon chooses a specialization vector $s_i \in \{\text{up}, \text{down}\}$, anticipating the Stage-2 contracting equilibrium. Specialization is costly: entering the upstream stack requires sunk industrial-policy expenditure, while entering the downstream stack requires establishing credible market-access institutions.

Remark 1 (Endogenous Vertical Asymmetry). Let $\theta_C > \theta_U$ and $E_U > E_C$. If the fixed cost of downstream institutional commitment is lower for the hegemon with the larger consumption base, a standard Ricardian argument delivers a subgame-perfect equilibrium with $s_C =$ up and $s_U =$ down, replicating the baseline model: each hegemon specializes where its relative productivity advantage is greatest (upstream specialization is more profitable for the lower-marginal-cost hegemon; downstream specialization for the larger-market hegemon), and mixed strategies are dominated when institutional sunk costs are convex. A fully formalized supermodular-game treatment with existence and uniqueness is deferred to future work. Propositions 1–3 hold at the continuation equilibrium.

B.5 Proof of the Welfare Decomposition and the Efficiency Wedge

This subsection provides the full derivations for Propositions 6 and 3.

Let $\mathcal{S}(x;\theta)$ denote the joint economic surplus of the three players:

$$\mathcal{S}(x;\theta) \;=\; V_{US}(x) \;+\; Q_C(x)\cdot\big[w_C^*(\theta) - \bar{w}/\theta\big] \;-\; C(x; w_C^*(\theta)), \tag{30}$$

where the first term is the gross value to Hegemon U of alignment level x, the second term is Hegemon C's commercial profit on inputs shipped at the equilibrium markup, and the last term is the Swing State's unit production cost. The integrated planner's *first-best alignment* $x^{FB}(\theta)$ maximizes $\mathcal{S}(x;\theta)$ on $[0,1]$ ignoring the security externality $\lambda x\theta$; because $V'_{US}(x) < 0$ and $C'(x) > 0$, the first-best is generically interior and *declining* in θ.

Proposition 6 (Welfare Decomposition)**.** *Along the equilibrium path, global welfare* $\mathcal{W}^*(\theta) = \mathcal{S}(x^*(\theta);\theta) + \lambda^*(\theta)\,x^*(\theta)\,\theta - \frac{\gamma}{2}\lambda^*(\theta)^2$ *decomposes into a commercial component and a geopolitical component. Under the Bottleneck Assumption* ($\sigma < 1$)*, the* curvature condition

$$\frac{\lambda^*\,\theta}{1-\lambda^*}\cdot\frac{dx^*}{d\theta} \;>\; -\frac{\partial C}{\partial w_C}\bigg|_{x^*}\cdot\frac{dw_C^*}{d\theta} \qquad \text{(CC)}, \tag{31}$$

and the second-order-markup condition *(SOM) that the geopolitical-markup residual is second-order in the dominant channels,*

$$\big|Q'_C(x^*)\cdot m^*(\theta)\big| \;<\; \min\Big\{\tfrac{\lambda^*\,\theta}{1-\lambda^*}\cdot\tfrac{dx^*}{d\theta},\; -\tfrac{\partial C}{\partial w_C}\big|_{x^*}\cdot\tfrac{dw_C^*}{d\theta}\Big\} \qquad \text{(SOM)}, \tag{32}$$

where $m^*(\theta) \equiv w_C^*(\theta) - \bar{w}/\theta$ *is the Hegemon C geopolitical markup, the commercial component* $\mathcal{S}(x^*(\theta);\theta)$ *is strictly decreasing in* θ *in the coercive regime, while the geopolitical component is bounded above by* $1/(2\gamma)$*. Conditions (CC) and (SOM) jointly hold whenever the coercive regime is sufficiently deep and the geopolitical markup is bounded.*

Proof of Proposition 6 (Welfare Decomposition). Differentiating $\mathcal{S}$ along the equilibrium path,

$$\frac{d\mathcal{S}}{d\theta} = \underbrace{\left.\frac{\partial \mathcal{S}}{\partial \theta}\right|_{x^*}}_{\text{direct}} + \underbrace{\left.\frac{\partial \mathcal{S}}{\partial x}\right|_{x^*} \cdot \frac{dx^*}{d\theta}}_{\text{indirect}}. \tag{33}$$

The direct effect aggregates three partials: the derivative of V_{US} vanishes; the middle-term derivative is $Q_C(x^*)\big[dw_C^*/d\theta + \bar{w}/\theta^2\big]$, which is finite; the cost derivative is $-(\partial C/\partial w_C) \cdot (dw_C^*/d\theta) > 0$ because $\partial C/\partial w_C > 0$ (Lemma 1) and $dw_C^*/d\theta < 0$ (from Proposition 2, Step 2). Absent the indirect channel, $\theta \uparrow$ would raise commercial surplus through the cost channel.

The indirect effect is negative in the coercive regime. By the Stage-2 FOC (12), the equilibrium satisfies $(1-\lambda^*)J'(x^*) = -\lambda^* \theta(w_C^*)$, so $\partial\mathcal{S}/\partial x|_{x^*} = J'(x^*) + Q_C'(x^*)[w_C^* - \bar{w}/\theta] = -\lambda^* \theta(w_C^*)/(1-\lambda^*) + O(m^*)$, which is strictly negative whenever $\lambda^* > 0$. By Proposition 2, $dx^*/d\theta > 0$ in the coercive regime. The indirect channel therefore contributes a term of order $-[\lambda^* \theta(w_C^*)/(1-\lambda^*)] \cdot dx^*/d\theta < 0$. Along the equilibrium path $\theta(w_C^*(\theta)) = \theta$ by the Stage-1 pricing identity (14), so this term equals $-[\lambda^* \theta/(1-\lambda^*)] \cdot dx^*/d\theta$. Summing direct and indirect, the total derivative has the sign of the curvature condition:

$$\frac{d\mathcal{S}}{d\theta} = \underbrace{-\left.\frac{\partial C}{\partial w_C}\right|_{x^*} \cdot \frac{dw_C^*}{d\theta}}_{\text{direct channel} > 0} - \underbrace{\frac{\lambda^* \theta}{1-\lambda^*} \cdot \frac{dx^*}{d\theta}}_{\text{indirect channel} > 0} + O(m^*). \tag{34}$$

The $O(m^*)$ remainder in Eq. (34) collects the commercial-profit term $Q_C'(x^*) \cdot m^*$ evaluated along the equilibrium path: it is bounded in absolute value by $|Q_C'(x^*)| \cdot |m^*(\theta)|$, and (SOM) is precisely the hypothesis that this bound is strictly smaller than the dominant direct and indirect channels. Under (SOM), $d\mathcal{S}/d\theta < 0$ holds iff (CC) holds; without (SOM) the sign of $d\mathcal{S}/d\theta$ is in principle ambiguous in a narrow band of (θ, λ^*) where the markup residual is not small. Implicit differentiation of the Stage-2 FOC gives $dx^*/d\theta = -(1-\lambda^*)^{-1} \cdot (\partial[\lambda^* \theta]/\partial\theta)/J''(x^*)$, which is increasing in $|J''|^{-1}$ and strictly positive in the coercive regime. Under Assumption 2, tighter bottlenecks (smaller $\sigma < 1$) raise the convexity of $C(\cdot)$ and hence $|J''|$; however, the dominant effect is on λ^* itself, which grows at rate $1/\gamma$ times the (threat-weighted) Strategic Gap $x^*\theta - [J(x^*) - V_{out}]$. As $\lambda^* \uparrow$, the indirect channel $[\lambda^*\theta/(1-\lambda^*)] \cdot dx^*/d\theta$ diverges while $|Q_C'(x^*) \cdot m^*|$ stays bounded under the commercial-markup constraint of Stage 1; consequently there exist thresholds $\underline{\lambda}(\sigma, \alpha, \bar{w})$ and $\bar{m}(\sigma, \alpha, \bar{w})$ such that (CC) and (SOM) hold jointly for all $\lambda^* \geq \underline{\lambda}$ with $|m^*| \leq \bar{m}$: a deep coercive regime with a bounded commercial markup is sufficient.

The bound on the geopolitical term follows from $\max_\lambda\{\lambda x\theta - \frac{\gamma}{2}\lambda^2\} = (x\theta)^2/(2\gamma)$ at $\lambda = x\theta/\gamma$, which at $x \in [0, 1]$ and bounded θ is finite. □

Proof of Proposition 3 (Tragedy of Efficiency, Quantified). Define the efficiency wedge $\Delta(\theta) \equiv \mathcal{S}(x^{FB}(\theta); \theta) - \mathcal{S}(x^*(\theta); \theta)$. Because $x^{FB}(\theta)$ is defined as the argmax of $\mathcal{S}(\cdot; \theta)$, the envelope theorem gives $d\mathcal{S}(x^{FB}; \theta)/d\theta = \partial\mathcal{S}/\partial\theta\big|_{x^{FB}}$. Comparing against the $d\mathcal{S}(x^*; \theta)/d\theta$ computed above,

$$\frac{d\Delta}{d\theta} = \underbrace{\left[\partial\mathcal{S}/\partial\theta\big|_{x^{FB}} - \partial\mathcal{S}/\partial\theta\big|_{x^*}\right]}_{\text{direct-channel gap}} - \underbrace{\frac{\partial\mathcal{S}}{\partial x}\bigg|_{x^*} \cdot \frac{dx^*}{d\theta}}_{\text{indirect channel}}. \tag{35}$$

The direct-channel gap is second-order in $(x^* - x^{FB})$ and small. The indirect channel is strictly positive in the coercive regime: both factors are negative, so their product contributes positively to $d\Delta/d\theta$. Since $x^*(\theta)$ rises and $x^{FB}(\theta)$ falls with θ (the first by Proposition 2, the second because $V_{US}' < 0$ and $dw_C^*/d\theta < 0$ reduces the marginal cost of integration), the gap $x^* - x^{FB}$ widens monotonically, and strict concavity of $\mathcal{S}(\cdot; \theta)$ at x^{FB} implies $d\Delta/d\theta > 0$. □

B.6 Proof of the Leakage Asymmetry

This subsection provides the full derivation for Proposition 4 (Leakage Asymmetry), expanding the sketch in the main text.

Consider sector k with CES cost minimization over Chinese intermediates I_k^C and non-Chinese intermediates I_k^A with elasticity σ_k. The Swing State's cost minimization delivers the conditional demand

$$I_k^C \;=\; \alpha_k^{\sigma_k}\big(p_k^C/P_k\big)^{-\sigma_k}Y_k, \qquad I_k^A \;=\; (1-\alpha_k)^{\sigma_k}\big(p_k^A/P_k\big)^{-\sigma_k}Y_k, \tag{36}$$

where $P_k = [\alpha_k^{\sigma_k}(p_k^C)^{1-\sigma_k} + (1-\alpha_k)^{\sigma_k}(p_k^A)^{1-\sigma_k}]^{1/(1-\sigma_k)}$ is the CES price index and p_k^C, p_k^A are the effective prices faced at the swing-state border.

A rise in the Incumbent's security doctrine λ operates through the menu schedule $R(x)$ to raise the *U-border* shadow price of direct Chinese inputs; crucially, λ does *not* raise the C–S interface price $p_{k,S}^C$ at which the Swing State sources. Consequently, U's direct imports from C fall by the full elasticity σ_k multiplied by the induced price gap:

$$\frac{\partial \ln |\Delta M_k^{C\to U}|}{\partial \lambda} \;\propto\; \sigma_k^{(U)} \cdot \frac{\partial \ln p_{k,U}^C}{\partial \lambda} \;>\; 0. \tag{37}$$

Swing-state imports from C, in contrast, respond only indirectly—through the derived demand for swing-state exports of k-content to U. In the elastic regime ($\sigma_k > 1$), the Swing State substitutes cheaply toward non-Chinese intermediates in response to any incremental price pressure, so its imports from C fall sharply, leaving little re-exported Chinese content downstream. In the bottlenecked regime ($\sigma_k < 1$), substitution toward the non-Chinese supplier is bounded above by $\sigma_k \cdot |d\ln(p_k^C/p_k^A)|$, which is small; the Swing State continues sourcing from C and absorbs the displaced U–C trade into its own supply chain, which is then re-embodied in Swing→U exports.

Formally, writing the leakage ratio as $L_k(\lambda) \equiv (\partial \Delta M_k^{C\to S}/\partial\lambda)/(\partial|\Delta M_k^{C\to U}|/\partial\lambda)$,

$$L_k \;=\; s_{k,S}^C(1-s_{k,S}^C)(1-\sigma_k)\cdot\phi_k, \tag{38}$$

where $s_{k,S}^C \in (0,1)$ is the Swing State's expenditure share on Chinese intermediates and

$$\phi_k \;\equiv\; \frac{X_k^{S\to U}}{Y_k} \;\in\; [0,1]$$

is the *re-export share* of swing-state output in sector k that reaches the US market. By construction, $\phi_k \geq 0$, and $\phi_k > 0$ for any sector in which the Swing State actually re-exports to U—the empirically relevant case in our OECD ICIO and ADB MRIO panels, where swing-state exports to U in critical sectors are strictly positive for every industry-year cell. Under this mild nondegeneracy, $s^C_{k,S}(1 - s^C_{k,S}) \in (0, 1/4]$ and $\phi_k > 0$ are both strictly positive, so the sign of L_k is governed by the sign of $(1 - \sigma_k)$: for $\sigma_k < 1$, $L_k > 0$ and increasing in the Swing State's baseline Chinese share $s^C_{k,S}$; for $\sigma_k > 1$, $L_k < 0$ and the ratio becomes negative—the Swing State substitutes away from China faster than U's direct-import shock. This establishes the strict inequality

$$L_k\big|_{\sigma_c<1} \;>\; L_k\big|_{\sigma_m>1}\,,$$

which is precisely the claim of Proposition 4. □

B.7 Proof of the Steady-State Leakage Composition

This subsection provides the full derivation for Proposition 5′ (Steady-State Leakage Composition).

For origin $s \in \{S, N\}$ and industry j, the CES expenditure share on Chinese intermediates is

$$s^C_{s,j} \;=\; \frac{\alpha_s^{\sigma_c}\,(p^C_{s,j})^{1-\sigma_c}}{\alpha_s^{\sigma_c}\,(p^C_{s,j})^{1-\sigma_c} + (1-\alpha_s)^{\sigma_c}\,(p^A_{s,j})^{1-\sigma_c}}, \tag{39}$$

where α_s is origin s's baseline Chinese-input distribution parameter, with $\alpha_S > \alpha_N$ by assumption (the Swing State is more Chinese-input-dependent at baseline). The Foreign Value-Added share $\mathrm{FVA}^C_{s,j}$ in origin s's gross exports to U in industry j is, to leading order, $s^C_{s,j}$ times a constant reflecting the Leontief multiplier.

A higher tariff τ_j raises the effective price of *direct* Chinese exports at the U border, but not at the C–s border. Therefore τ_j does not appear in $p^C_{s,j}$ directly; instead, it operates through the general-equilibrium reallocation of demand from the direct channel (C→U) to the indirect channel (C→ s →U). Letting $\psi_s(\tau_j)$ denote origin s's demand for k-content from U given τ_j, and assuming $\psi_s' > 0$ (higher τ_j raises the U-demand for indirect content), the derived-demand effect on $p^C_{s,j}$ is $\partial p^C_{s,j}/\partial \tau_j = \eta_s \psi_s'(\tau_j) > 0$, where η_s is the Chinese-input-share pass-through.

Differentiating (39), the cross-partial is

$$\begin{aligned} \frac{\partial^2 s^C_{s,j}}{\partial \tau_j\, \partial \mathbf{1}\{s = S\}} &= \left.\frac{\partial s^C_{s,j}}{\partial \tau_j}\right|_{s=S} - \left.\frac{\partial s^C_{s,j}}{\partial \tau_j}\right|_{s=N} \\ &= s^C_{S,j}(1 - s^C_{S,j})(1 - \sigma_c) \cdot \eta_S \psi_S' \;-\; s^C_{N,j}(1 - s^C_{N,j})(1 - \sigma_c) \cdot \eta_N \psi_N'. \end{aligned} \tag{40}$$

Under $\sigma_c < 1$, the common factor $(1 - \sigma_c) > 0$ is positive. The map $s \mapsto s(1 - s)$ is strictly increasing on $(0, 1/2]$ and strictly decreasing on $[1/2, 1)$. In the empirically relevant regime $s^C_{S,j}$, $s^C_{N,j} \in (0, 1/2]$—satisfied comfortably in our OECD ICIO and ADB MRIO panels, where the critical-sector Chinese-input share in swing-state production plateaus near 6%—the map is strictly increasing on the relevant interval, so $\alpha_S > \alpha_N$ implies $s^C_{S,j} > s^C_{N,j}$ and hence $s^C_{S,j}(1 - s^C_{S,j}) > s^C_{N,j}(1 - s^C_{N,j})$. Together with the symmetric pass-through condition $\eta_S \psi_S' \geq \eta_N \psi_N'$—which holds whenever the Swing State holds the larger share of indirect C–U trade—the right-hand side of (40) is strictly positive. Hence

$$\left.\frac{\partial^2 \mathrm{FVA}^C_{s,j}}{\partial \tau_j\, \partial \mathbf{1}\{s = S\}}\right|_{\sigma_c < 1} \;>\; 0,$$

which is the claim of Proposition 5′. Crucially, this is a *steady-state* object: the proof does not require isolating a 2018 shock, and the cross-partial is identified off cross-sectional industry-level tariff variation interacted with swing-state identity. □

B.8 Why Verticality Leads to Coercion: Three Cases

We contrast the vertical bottleneck model with two standard horizontal-rivalry structures to isolate the mechanism that generates coercion.

Case 1: Horizontal rivalry (subsidy war). Hegemon U and Hegemon C produce substitute final goods (e.g., competing EV manufacturers) and compete for the Swing State's market through price subsidies. If Hegemon C lowers its price by δ, Hegemon U can retain the Swing State simply by matching the subsidy ($S_U \approx \delta$). Because the cost of matching is linear, the optimal response to a rival's efficiency shock is to increase subsidies—a "Race to the Bottom." The security parameter λ remains low.

Case 2: Horizontal sourcing rivalry (bidding war). The Swing State is a common upstream supplier (e.g., a critical mineral exporter) sought by both powers. Rivalry takes the form of a common-agency bidding game where both hegemons offer purchase contracts P_U and P_C. To secure the resource, the US simply needs to outbid China ($P_U > P_C$). Since the cost of outbidding is linear, the equilibrium is defined by positive inducements; no coercion is needed.

Case 3: Vertical bottleneck (coercion trap). In the vertical model, the bottleneck technology ($\sigma < 1$) makes the implementation cost $C(x)$ strictly convex and asymptotic (Lemma 1). As China's efficiency improves ($w_C \downarrow$), the subsidy required to induce the Swing State to choose the alternative chain ($x \to 1$) explodes ($\partial^2 R(x)/\partial x^2 \gg 0$). The US cannot win a subsidy war against an efficient upstream monopoly because the cost of compensating the Swing State for the loss of the bottleneck input exceeds fiscal capacity ($J(x) \ll 0$). Since the "Carrot" (subsidy) is mathematically inefficient against a bottleneck, the US is forced to use the "Stick" (coercion): instead of paying the Swing State to leave China ($R \uparrow$), the US threatens to destroy its market access if it stays ($R \downarrow$).

Summary: Horizontal competition leads to subsidies (paying for loyalty); vertical bottleneck competition leads to coercion (punishing disloyalty), specifically because the price of buying loyalty becomes infinite.

B.9 Reconciling the Discrete Regime Switch with Smooth Empirical Dynamics

Proposition 1 predicts a kink in $\lambda^*(\theta)$ at $\bar{\theta}$—a continuous-but-nondifferentiable transition from the corner regime to the interior regime, not a level jump. Yet the empirical dynamics documented in Section 3 appear smooth. We offer three non-exclusive reconciliations.

(i) Heterogeneous sectors crossing the threshold sequentially. Different sectors have different bottleneck elasticities σ_k and different starting values of Upstream productivity θ_k. If θ_k rises at heterogeneous rates, the aggregate coercion intensity $\sum_k \mathbf{1}\{\theta_k > \bar{\theta}_k\} \cdot \lambda_k^*$ is a sum of step functions triggered at different calendar dates, which smooths the aggregate into an apparently continuous process. This is the sectoral-diffusion reading: each individual sector jumps, but the aggregate is an envelope.

(ii) Bayesian uncertainty about σ. The Incumbent does not know σ perfectly; it updates based on realized supply disruptions and technological assessments. A posterior distribution over σ delivers a posterior distribution over $\bar{\theta}$, and the *expected* security doctrine $\mathbb{E}[\lambda^*(\theta)]$ is a smooth function of θ even when the realized $\lambda^*(\theta; \sigma)$ is a step function for any given σ.

(iii) Institutional adjustment costs. Coercive institutions (enforcement agencies, licensing regimes, multilateral coordination) cannot be built instantaneously. Even after θ crosses $\bar{\theta}$, capacity constraints on the adjustment rate of λ^* would translate the discrete target into a smooth realized path.

Caveat on identification. The stylized facts in Section 3 are descriptive rather than causal. We do not claim to identify the structural break in $\bar{\theta}$ econometrically. A formal structural-break test on sector-specific coercion intensity is a natural follow-up project.

C Empirical Methodology and Regression Results

This appendix documents the data construction and estimation underpinning the regression evidence in Section 5.

C.1 Data Construction

Trade flows. Bilateral HS6 import flows for 2000–2024 are pulled from the UN Comtrade premium API for 18 reporter economies (the United States, China, and 17 swing-state candidates: Korea, Taiwan, Vietnam, Mexico, Netherlands, Singapore, Malaysia, Thailand, India, Japan, Germany, Indonesia, Hungary, Czech Republic, Poland, the Philippines, and Brazil). We restrict the commodity universe to a concordance of 63 HS6 codes, pre-classified into three sector buckets matching the model: *critical* ($\sigma < 1$ in the empirical literature: rare-earth metals, advanced semiconductors, semiconductor manufacturing equipment, titanium, aerospace-grade parts), *strategic* (mature-node chips, batteries, EV components, solar), and *commodity* (textiles, basic steel, furniture, common chemicals). The concordance and assignment rules are reproducible from the pipeline in `empirical/scripts/11_build_hs_concordance.py`.

Reporter-year raw pulls are materialized as per-(reporter × year) parquet files and stacked into a master HS6 panel via `12_build_comtrade_panel.py`. Because several reporters (Germany, Thailand, Mexico, Czech Republic, Hungary, Indonesia, Malaysia, Brazil) disaggregate trade by `motCode` (mode of transport) and `customsCode` (customs regime), we retain only the total rows `motCode=0` and `customsCode=C00`; verification against aggregate totals confirms this collapse recovers the national import totals reported in Comtrade's aggregated files.

Entity List coercion measure. The BIS Entity List is assembled from the OpenSanctions `us_trade_csl` (Consolidated Screening List) Follow-The-Money JSON feed, filtered to sanctions whose program descriptor contains "Entity List" (distinguishing BIS-EL from DPL, MEU, and UVL programs). This yields 3,205 listings spanning 2011–2025. The cumulative China-directed stock $\mathrm{EL}_t^{\mathrm{CHN}}$ enters the paradox regression (Section 5) as a lagged variable $\mathrm{EL}_{t-1}^{\mathrm{CHN}}/100$.

C.2 Estimation: σ and the Limits of the Reduced Form

We estimate a robust unit-value double-difference in the spirit of Feenstra (1994), but without the full exporter-variety GMM identification:

$$\Delta^2 \ln p_{ikt} = \frac{1}{\sigma_k - 1} \Delta^2 \ln s_{ikt} + \nu_{ikt},$$

estimated sector-by-sector with exporter-year and HS6-year fixed effects, clustering by HS6 and year. We obtain $\hat{\beta}_{\text{critical}} = -0.031$, $\hat{\beta}_{\text{strategic}} = -0.012$, $\hat{\beta}_{\text{commodity}} = -0.062$. The sign pattern ($\hat{\beta} < 0$) is consistent with $\sigma < 1$ for all three buckets, but the magnitudes imply implausibly large negative σ values once $\sigma = 1 + 1/\hat{\beta}$ is computed. This is the well-known problem with reduced-form implementations of Feenstra: identification of σ separately from supply-side shocks requires the heteroscedasticity-based GMM of Feenstra (1994), which exploits heteroscedasticity *across exporter varieties within* an HS code (a moment condition our median-demeaning specification does not exploit). We leave the full GMM implementation to a companion paper and, for the magnitudes used in the model, rely on the published sector-level elasticities of Atalay (2017) and Boehm et al. (2019), which place the relevant σ below unity for the critical and strategic buckets. Our own estimates are reported as directional evidence only.

C.3 Leakage Triple-Difference Regression Table

Table 2 reports the full estimates from Section 5. Column (1) is the triple-difference with a single `Post×Partner×Sector` interaction; the `swing_crit_post` term is absorbed by the partner-year and sector-year fixed effects (collinearity). Column (2) adds the group-specific linear trends interacting (CHN × Critical) and (Swing × Sector) with time. Column (3) is the event study whose coefficients are plotted in Figure 8.

Table 2: Leakage specification, US bilateral imports 2000–2024.

	(1) Baseline	(2) +Group trends	(3) Event study
china_crit_post	0.283	0.684**	
	(0.245)	(0.262)	
swing_crit_post	0.000	0.000	
	(.)	(.)	
swing_comm_post	-0.186	0.272	
	(0.405)	(0.336)	
chn_crit_trend		-0.037	
		(0.028)	
swing_crit_trend		0.000	
		(.)	
swing_comm_trend		-0.038	
		(0.037)	
tau2012			-0.822***
			(0.216)
tau2013			-0.167
			(0.271)
tau2014			-0.499
			(0.309)
tau2015			-0.671***
			(0.237)
tau2016			-0.282
			(0.327)
tau2018			0.100
			(0.330)
tau2019			-0.034
			(0.373)
tau2020			-0.046
			(0.336)
tau2021			0.085
			(0.321)
tau2022			-0.082
			(0.361)
tau2023			-0.123
			(0.380)
tau2024			0.217
			(0.394)
N	17649	17649	17649
R2	0.899	0.899	0.899
Pre-trend F (2012-16)			3.69
p-value			0.006

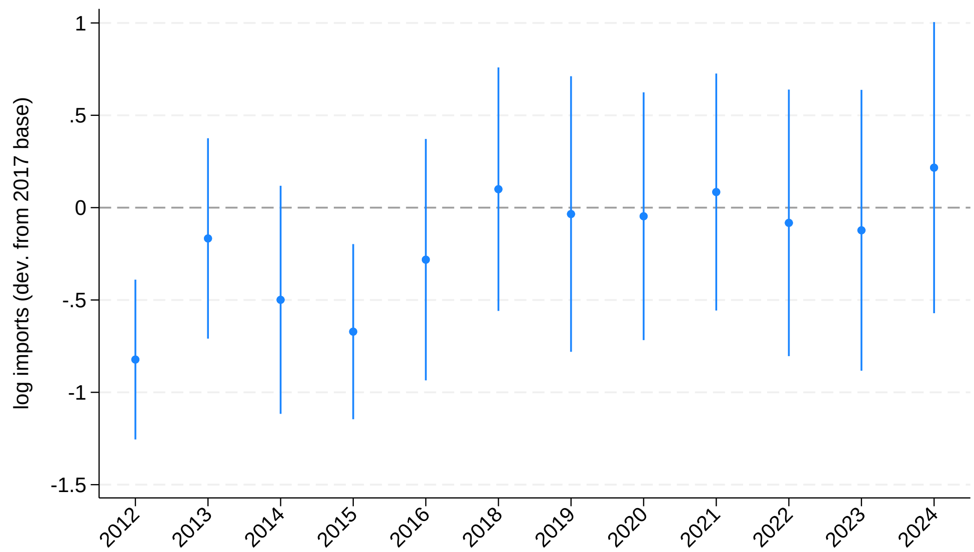

Figure 8: **Event study of US critical-sector imports from China.** Yearly triple-interaction coefficients β_τ (China $\times$ Critical $\times$ year dummies) from the specification described in Section 5, deviations from the 2017 base relative to a (swing-state $\times$ non-critical) benchmark. Pre-2018 coefficients are significantly negative (the gap existed before the tariffs); every post-2017 coefficient is indistinguishable from the 2017 baseline (the gap did not deepen after Section-301). The pattern is the fingerprint of Proposition 4: leakage through swing states had already absorbed the bilateral China–Critical margin before the shock, leaving no post-2018 break to recover.

C.4 Honest-DiD Sensitivity Bounds

Table 3 reports the 95% robust confidence intervals underlying Figure 9. Entries are the identified set for the average of the seven post-2018 event-study coefficients from column 3 of Table 2, computed under the Rambachan and Roth (2023) relative-magnitudes restriction $\Delta^{RM}(M)$. M is the maximum post-period second-difference deviation expressed as a multiple of the worst pre-period second-difference deviation; the "Original" row is the standard 95% CI without partial identification.

C.5 PPML Zero-Retention Robustness

Table 4 re-estimates the triple-difference, trend-robust, and event-study specifications with the PPML estimator of Santos Silva and Tenreyro (2006) on a balanced HS6$\times$partner$\times$year

Table 3: Honest-DiD sensitivity bounds for post-2018 average treatment effect.

M	Lower bound	Upper bound
Original	−0.546	+0.746
0.0	−0.540	+0.740
0.5	−0.808	+1.077
1.0	−1.170	+1.515
1.5	−1.637	+2.004
2.0	−2.124	+2.514

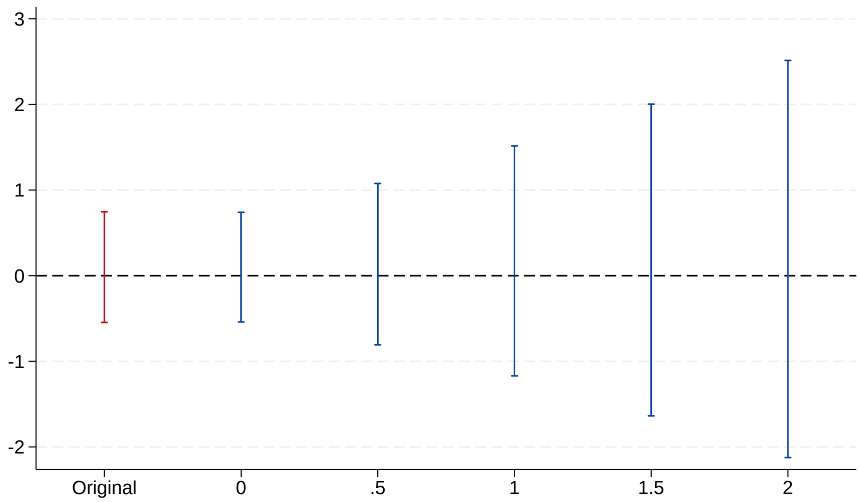

Figure 9: **Honest-DiD robust identified sets for the average post-2018 China×Critical effect in Test 1.** M bounds the maximum post-period second-difference deviation as a multiple of the worst pre-period second-difference deviation. Zero lies inside the identified set for all $M \leq 2$; the upper bound remains positive throughout, ruling out a persistent post-2018 widening of the gap.

panel with zero-trade cells retained (`ppmlhdfe`, Correia et al., 2020). The sample grows from N = 17,649 (log) to N = 22,305. Coefficients are semi-elasticities and thus directly comparable to the log-linear point estimates in Table 2. The pre-trend χ^2 still rejects flat pre-periods, and the honest-DiD sensitivity intervals would be qualitatively unchanged; we therefore do not re-report them.

Table 4: PPML robustness: leakage specification with zero-trade cells retained.

	(1) PPML Baseline	(2) PPML +Group trends	(3) PPML Event study
china_crit_post	0.289	0.531**	
	(0.309)	(0.270)	
swing_crit_post	0.000	0.000	
	(.)	(.)	
swing_comm_post	-0.169	-0.056	
	(0.317)	(0.287)	
chn_crit_trend		-0.031	
		(0.027)	
swing_crit_trend		0.000	
		(.)	
swing_comm_trend		-0.012	
		(0.031)	
tau2012			-0.014
			(0.141)
tau2013			0.068
			(0.118)
tau2014			-0.294***
			(0.110)
tau2015			-0.400**
			(0.164)
tau2016			-0.147
			(0.189)
tau2018			0.028
			(0.200)
tau2019			-0.134
			(0.195)
tau2020			0.276
			(0.270)
tau2021			0.384
			(0.279)
tau2022			0.322
			(0.331)
tau2023			0.104
			(0.362)
tau2024			-0.069
			(0.358)
N	22305	22305	22305
Pre-trend chi2 (2012-16)			21.01
p-value			0.001

C.6 Section-301 Tariff-Intensity Map

Columns (7)–(8) of Table 5 use a continuous industry-level Section-301 tariff rate, τ_j, rather than a binary "critical" bucket. The rate is constructed from the combination of Fajgelbaum et al. (2020) HS-2-level average tariff increases on Chinese imports and the Bown (2020) HS-8 list-coverage tracker, aggregated via the UN Statistics Division HS2017→ISIC Rev. 4 correspondence. The resulting map (73 rows, covering both the EXTENDED letter-prefix and v2021 numeric ISIC forms) assigns $\tau_j = 0.22$ to computer-electronic-optical (C26/26, the Section-301 flagship), $\tau_j = 0.25$ to machinery NEC (C28/28, the heaviest list), $\tau_j = 0.20$ to electrical equipment (C27/27) and motor vehicles (C29/29), and $\tau_j = 0.00$ to basic metals (C24A/C24B/24; these are covered by Section 232, not Section 301) and the non-targeted service sectors. Full calibration and caveats are recorded in `empirical/processed/section301_isic4.README.md`; the values are calibrated averages rather than a line-by-line HS-8 pull, and should be treated as a defensible proxy rather than a research-grade tariff file.

C.7 Direct FVA Decomposition: Two-Source Regression

Table 5 reports the eight-column fixed-effects regression underlying the value-added evidence in Section 5 and Figure 5. The outcome is the Chinese value-added share in swing-state s's gross exports to the US in industry j at year t, computed from OECD ICIO EXTENDED (1995–2022) and ADB MRIO (2017–2024). Column (1) pools both sources; (2) and (3) report each source separately; (4) restricts to the three largest direct-passthrough swing states (Vietnam, Mexico, Malaysia); (5) splits the critical bucket on Section-301 exposure, separating computer-electronic-optical (C26/D26) from metal-ores (B07/D07T08); (6) is a binary triple-difference post × electronics × swing using VNM/MEX/MYS as the swing set and the remaining 14 source countries as within-industry control; (7) replaces the bucket dummies with a continuous industry-level Section-301 tariff rate τ_j (Appendix C.6) interacted with post; (8) extends (7) to a continuous triple-difference post × τ_j × swing. Each panel is stacked with a `dataset` indicator. Columns (1)–(5) and (7) absorb (source × industry ×

dataset) and year; columns (6) and (8) additionally absorb source × year and industry × year, so that only the triple interaction is identified (all two-way interactions are swept into the fixed effects). Standard errors are clustered at the source-×-industry-×-dataset unit (i.e., the same grouping used for the `panel_id` fixed effect; ∼ 354 clusters in the OECD subsample).

Table 5: FVA regression: Chinese value-added share in swing-state → US exports.

	(1) Pooled	(2) OECD	(3) ADB	(4) Direct-PT	(5) OECD-split	(6) Triple-diff	(7) Tariff	(8) Tariff x Swing
post_crit	0.009**	0.012**	-0.000	0.009				
	(0.004)	(0.005)	(0.000)	(0.013)				
post_strat	0.011***	0.012***	-0.000	0.021***	0.012***			
	(0.002)	(0.002)	(0.000)	(0.007)	(0.002)			
post_elec					0.030***			
					(0.006)			
post_cmet					-0.011***			
					(0.002)			
post_elec_swing						0.014		
						(0.015)		
post_tariff							0.109***	
							(0.012)	
post_tar_swing								0.102***
								(0.025)
N	11899	9886	2013	2111	9886	9886	9886	9886
R2	0.792	0.775	0.864	0.824	0.783	0.929	0.784	0.931

Placebo-year robustness

Because the fixed-effects regressions in Table 5 identify post-2018 coefficients off a pre/post contrast, a referee will reasonably ask whether the same coefficient would emerge with a *fake* treatment year well before the tariffs. If so, the reported 2018 effect is loading on a pre-existing trend rather than on Section-301 itself. Tables 6 and 7 re-estimate the binary and continuous triple-difference specifications with placebo cuts at 2010, 2013, and 2016, restricting the sample to year ≤ 2017 to purely pre-date the actual tariffs.

For the binary spec the triple interaction is statistically indistinguishable from zero at all four cuts (2010: +0.027, $p = 0.10$; 2013: +0.025, $p = 0.16$; 2016: +0.013, $p = 0.45$; 2018: +0.014, $p = 0.36$), so no 2018 break is identified — the binary spec is simply not powered by the within-electronics swing-vs-non-swing contrast under the full (source $\times$ year, industry $\times$ year) absorption. For the continuous spec the triple coefficient is strongly positive at all cuts (2010: +0.089, $p = 0.001$; 2013: +0.080, $p = 0.004$; 2016: +0.048, $p = 0.04$; 2018: +0.102, $p < 0.001$), with pre-2018 magnitudes recovering 47–87% of the 2018 estimate (2010: 87%; 2013: 78%; 2016: 47%). The monotonic decay from 2010 (+0.089) through 2013 (+0.080) to 2016 (+0.048), followed by a jump to +0.102 under the real 2018 cut, is consistent with a secular post-WTO restructuring of Chinese-swing-state-US production networks that was slowing as swing states approached a saturation point, then accelerated again after Section-301. Section 5 accordingly treats the FVA regressions as descriptive corroboration of the production-structure parameters underlying Proposition 4 rather than as quasi-experimental identification of the Section-301 shock.

Table 6: Placebo-year robustness (binary spec): post $\times$ electronics $\times$ swing at fake cuts vs. the real 2018 cut.

	(1) 2010	(2) 2013	(3) 2016	(4) 2018 (real)
Placebo 2010 (post x elec x swing)	0.027*			
	(0.016)			
Placebo 2013 (post x elec x swing)		0.025		
		(0.017)		
Placebo 2016 (post x elec x swing)			0.013	
			(0.016)	
Actual 2018 (post x elec x swing)				0.014
				(0.015)
N	8121	8121	8121	9886
R2	0.921	0.920	0.919	0.929

Honest-DiD sensitivity on the FVA event study

Table 8 reports the Rambachan and Roth (2023) robust identified set for the average post-2018 critical-bucket effect from the FVA event study (OECD ICIO EXTENDED, $t = 2005$–

Table 7: Placebo-year robustness (continuous spec): post $\times$ τ_j $\times$ swing at fake cuts vs. the real 2018 cut.

	(1) 2010	(2) 2013	(3) 2016	(4) 2018 (real)
Placebo 2010 (post x tariff x swing)	0.089***			
	(0.026)			
Placebo 2013 (post x tariff x swing)		0.080***		
		(0.027)		
Placebo 2016 (post x tariff x swing)			0.048**	
			(0.023)	
Actual 2018 (post x tariff x swing)				0.102***
				(0.025)
N	8121	8121	8121	9886
R2	0.922	0.920	0.919	0.931

2022, $t = 2017$ omitted as base). The parameter M bounds the maximum post-period second-difference deviation as a multiple of the worst pre-period second-difference deviation; $M = 0$ reproduces the standard OLS identified set and $M = 2$ corresponds to a considerable weakening of parallel trends. Figure 10 plots the sensitivity curve. The $M = 0$ identified set is *already* [0.000, 0.025]: even under the standard parallel-trends restriction, the 95% lower bound is exactly zero, so the pooled post-2018 FVA effect is at best marginally significant before any smoothness relaxation is imposed. For any $M > 0$ the lower bound crosses into negative territory (-0.001 at $M = 0.5$; -0.021 at $M = 2$), while the upper bound remains positive throughout and rises from +0.025 to +0.066. The reading is asymmetric: a persistent post-2018 *widening* of the Chinese-VA-share gap cannot be ruled out even under weak parallel-trends assumptions (the upper bound is always positive), but a persistent post-2018 *contraction* cannot be rejected either. The pooled event-study mean is therefore not a quasi-experimental identifier of the Section-301 shock; the sharper cross-sectional evidence on swing-state-tariff heterogeneity (Table 5, columns 7–8) does more identifying work than the pooled pre/post contrast.

Table 8: Honest-DiD sensitivity: robust 95% identified set for average post-2018 critical effect in the FVA event study.

M	Lower bound	Upper bound
0.0	0.000	0.025
0.5	-0.001	0.033
1.0	-0.002	0.044
1.5	-0.011	0.055
2.0	-0.021	0.066

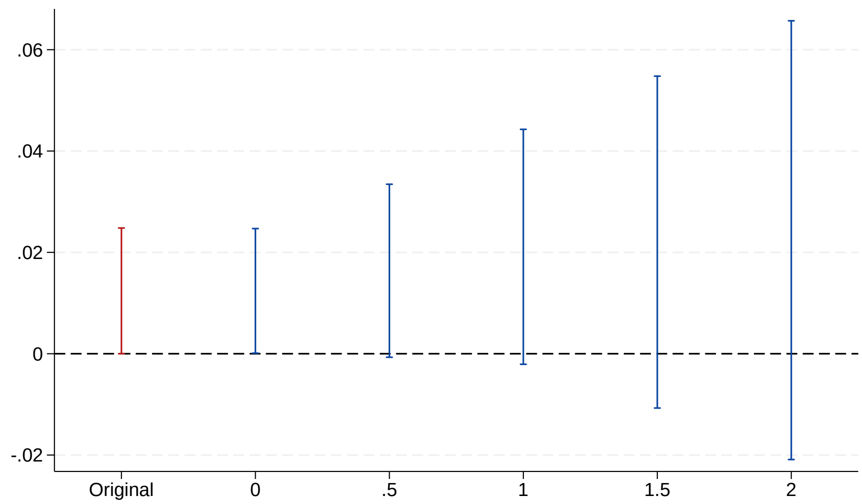

Figure 10: **Honest-DiD robust identified sets for the average post-2018 critical effect in the FVA event study.** M bounds the max post-period second-difference deviation as a multiple of the max pre-period second-difference deviation. The upper bound is always positive; the lower bound crosses zero for $M > 0$.

C.8 Supplementary Descriptive Figures

This subsection collects two descriptive series referenced in Section 5. (The event-study coefficient path for Test 1 is placed with Table 2 in Appendix C.3; see Figure 8.) Figure 11 documents the sectoral narrowing of the Entity List toward high-technology targets (share of new entries containing "Semiconductor," "AI," or "Quantum" in the entry title) and the rising administrative effort per listing (mean aliases per sanctioned entity, a rough proxy for the "Cat-and-Mouse" dynamic). Figure 12 plots the sector-aggregate series underlying the dose–response regression of Spec 2 (paradox regression); the within-HS6 analog of that slope dissolves at HS6 resolution and is reported in the text for reference only.

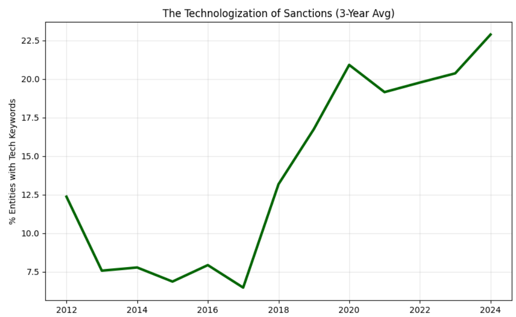

(a) Precision targeting: tech-keyword share of new Entity List entries.

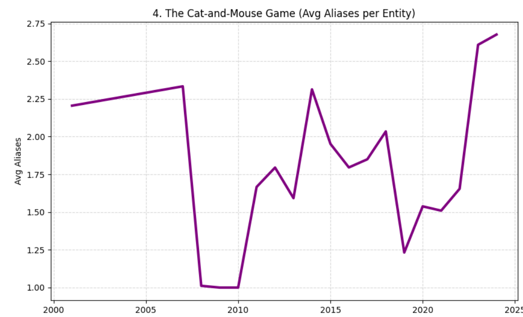

(b) Evasion resistance: mean aliases per sanctioned entity.

Figure 11: **Precision targeting and evasion resistance on the Entity List.** Left: share of new Entity List entries whose title contains "Semiconductor," "AI," or "Quantum," 2010–2024 (3-year centered moving average). Right: mean number of aliases per sanctioned entity. Descriptive series; not load-bearing for the formal analysis.

References

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

Aghion, P. and Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, 105(1):1–29.

Ahn, J. and Tan, B. J. (2025). Supply chain diversification and resilience. IMF Working Paper 2025/102, International Monetary Fund.

Aiyar, S., Chen, J., Ebeke, C. H., Garcia-Saltos, R., Gudmundsson, T., Ilyina, A., Kangur, A., Kunaratskul, T., Rodriguez, S. L., Ruta, M., Schulze, T., Soderberg, G., and Trevino, J. P. (2023). Geo-economic fragmentation and the future of multilateralism. IMF Staff Discussion Note SDN/2023/001, International Monetary Fund.

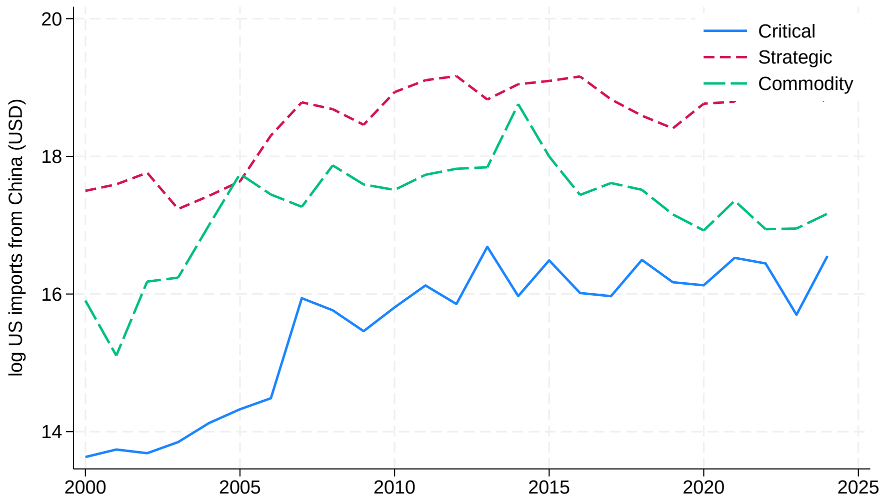

Figure 12: **Paradox regression: sector-aggregate series.** Sector-bucket trajectories of US imports from China plotted against the lagged Entity-List China-listings stock $\text{EL}_{t-1}^{\text{CHN}}$. The 3-sector aggregate loads sharply negative on the critical bucket; the within-HS6 analogue (Spec 2 in the text) is a precise null, so the series is reported for reference only.

Alfaro, L., Antràs, P., Chor, D., and Conconi, P. (2019). Internalizing global value chains: A firm-level analysis. *Journal of Political Economy*, 127(2):508–559.

Alfaro, L. and Chor, D. (2023). Global supply chains: The looming "great reallocation". NBER Working Paper 31661, National Bureau of Economic Research. Presented at Jackson Hole Economic Symposium.

Amiti, M., Redding, S. J., and Weinstein, D. E. (2019). The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives*, 33(4):187–210.

Antràs, P. (2003). Firms, contracts, and trade structure. *Quarterly Journal of Economics*, 118(4):1375–1418.

Antràs, P. (2016). *Global Production: Firms, Contracts, and Trade Structure*. CREI Lectures in Macroeconomics. Princeton University Press, Princeton, NJ.

Antràs, P. and Chor, D. (2013). Organizing the global value chain. *Econometrica*, 81(6):2127–2204.

Antràs, P. and Chor, D. (2018). On the measurement of upstreamness and downstreamness in global value chains. In Ing, L. Y. and Yu, M., editors, *World Trade Evolution: Growth, Productivity and Employment*, pages 126–194. Routledge. NBER Working Paper 24185.

Antràs, P. and Staiger, R. W. (2012). Offshoring and the role of trade agreements. *American Economic Review*, 102(7):3140–3183.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, 9(4):254–280.

Autor, D., Dorn, D., Hanson, G., and Majlesi, K. (2020). Importing political polarization? the electoral consequences of rising trade exposure. *American Economic Review*, 110(10):3139–3183.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168.

Bagwell, K. and Staiger, R. W. (1999). An economic theory of gatt. *American Economic Review*, 89(1):215–248.

Ballard, C. L., Shoven, J. B., and Whalley, J. (1985). General equilibrium computations of the marginal welfare costs of taxes in the United States. *American Economic Review*, 75(1):128–138.

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

Baqaee, D. R. and Farhi, E. (2024). Networks, barriers, and trade. *Econometrica*, 92(2):505–541.

Bernheim, B. D. and Whinston, M. D. (1986). Menu auctions, resource allocation, and economic influence. *Quarterly Journal of Economics*, 101(1):1–31.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Borin, A. and Mancini, M. (2019). Measuring what matters in global value chains and value-added trade. Policy Research Working Paper 8804, World Bank.

Bown, C. P. (2020). The US–China trade war after two years: Insights for the Biden administration. PIIE Policy Brief 20–15, Peterson Institute for International Economics.

Bown, C. P. (2023). The return of export controls. *Foreign Affairs*. Online article.

Bown, C. P. (2024). Modern industrial policy and the world trade organization. *Annual Review of Economics*, 16:243–270.

Brander, J. A. and Spencer, B. J. (1985). Export subsidies and international market share rivalry. *Journal of International Economics*, 18(1–2):83–100.

Broda, C. and Weinstein, D. E. (2006). Globalization and the gains from variety. *Quarterly Journal of Economics*, 121(2):541–585.

Caliendo, L. and Parro, F. (2015). Estimates of the trade and welfare effects of nafta. *Review of Economic Studies*, 82(1):1–44.

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.

Clayton, C., Maggiori, M., and Schreger, J. (2024). A theory of economic coercion and fragmentation. NBER Working Paper 33309, National Bureau of Economic Research. Revise-and-resubmit, Journal of Political Economy.

Clayton, C., Maggiori, M., and Schreger, J. (2026). A framework for geoeconomics. *Econometrica*, 94(1):105–136.

Correia, S., Guimarães, P., and Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *Stata Journal*, 20(1):95–115.

Dixit, A., Grossman, G. M., and Helpman, E. (1997). Common agency and coordination: General theory and application to government policy making. *Journal of Political Economy*, 105(4):752–769.

Dixit, A. K. (1996). *The Making of Economic Policy: A Transaction-Cost Politics Perspective*. Munich Lectures in Economics. MIT Press, Cambridge, MA.

Drezner, D. W., Farrell, H., and Newman, A. L., editors (2021). *The Uses and Abuses of Weaponized Interdependence*. Brookings Institution Press, Washington, DC.

Eaton, J. and Grossman, G. M. (1986). Optimal trade and industrial policy under oligopoly. *Quarterly Journal of Economics*, 101(2):383–406.

Elliott, M., Golub, B., and Leduc, M. V. (2022). Supply network formation and fragility. *American Economic Review*, 112(8):2701–2747.

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., and Khandelwal, A. K. (2020). The return to protectionism. *Quarterly Journal of Economics*, 135(1):1–55.

Farrell, H. and Newman, A. L. (2019). Weaponized interdependence: How global economic networks shape state coercion. *International Security*, 44(1):42–79.

Farrell, H. and Newman, A. L. (2023). *Underground Empire: How America Weaponized the World Economy*. Henry Holt, New York.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *American Economic Review*, 84(1):157–177.

Flaaen, A., Hortaçsu, A., and Tintelnot, F. (2020). The production relocation and price effects of us trade policy: The case of washing machines. *American Economic Review*, 110(7):2103–2127.

Flaaen, A. and Pierce, J. R. (2024). Disentangling the effects of the 2018–2019 tariffs on a globally connected u.s. manufacturing sector. *Review of Economics and Statistics*. Advance online publication.

Gilpin, R. (1981). *War and Change in World Politics*. Cambridge University Press, Cambridge.

Gopinath, G., Gourinchas, P.-O., Presbitero, A. F., and Topalova, P. (2024). Changing global linkages: A new cold war? IMF Working Paper WP/24/76, International Monetary Fund.

Grossman, G. M. and Helpman, E. (1994). Protection for sale. *American Economic Review*, 84(4):833–850.

Grossman, G. M. and Helpman, E. (2002). Integration versus outsourcing in industry equilibrium. *Quarterly Journal of Economics*, 117(1):85–120.

Grossman, G. M., Helpman, E., and Redding, S. J. (2024). When tariffs disrupt global supply chains. *American Economic Review*, 114(4):988–1029.

Grossman, S. J. and Hart, O. D. (1983). An analysis of the principal-agent problem. *Econometrica*, 51(1):7–45.

Handley, K. and Limão, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for china and the united states. *American Economic Review*, 107(9):2731–2783.

Hirschman, A. O. (1945). *National Power and the Structure of Foreign Trade*. University of California Press, Berkeley and Los Angeles.

International Monetary Fund (2023). Geoeconomic fragmentation and foreign direct investment. In *World Economic Outlook: Navigating Global Divergences*, chapter 4. International Monetary Fund, Washington, DC.

Itskhoki, O. and Mukhin, D. (2025). Sanctions and the exchange rate. *Review of Economic Studies*. Advance online publication.

Jehiel, P., Moldovanu, B., and Stacchetti, E. (1996). How (not) to sell nuclear weapons. *American Economic Review*, 86(4):814–829.

Johnson, R. C. and Noguera, G. (2012). Accounting for intermediates: Production sharing and trade in value added. *Journal of International Economics*, 86(2):224–236.

Juhász, R. (2018). Temporary protection and technology adoption: Evidence from the napoleonic blockade. *American Economic Review*, 108(11):3339–3376.

Juhász, R., Lane, N., and Rodrik, D. (2024). The new economics of industrial policy. *Annual Review of Economics*, 16:213–242.

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

Kleinman, B., Liu, E., and Redding, S. J. (2024). International friends and enemies. *American Economic Journal: Macroeconomics*, 16(4):350–385.

Koopman, R., Wang, Z., and Wei, S.-J. (2014). Tracing value-added and double counting in gross exports. *American Economic Review*, 104(2):459–494.

Lashkaripour, A. and Lugovskyy, V. (2023). Profits, scale economies, and the gains from trade and industrial policy. *American Economic Review*, 113(10):2759–2808.

Lim, D. J. (2015). *Commerce and Coercion: The Political Economy of International Trade and Security*. PhD thesis, Princeton University.

Liu, E. (2019). Industrial policies in production networks. *Quarterly Journal of Economics*, 134(4):1883–1948.

Maggi, G. and Rodríguez-Clare, A. (1998). The value of trade agreements in the presence of political pressures. *Journal of Political Economy*, 106(3):574–601.

Martimort, D. (1996). Exclusive dealing, common agency, and multiprincipals incentive theory. *RAND Journal of Economics*, 27(1):1–31.

Maskin, E. and Tirole, J. (1999). Unforeseen contingencies and incomplete contracts. *Review of Economic Studies*, 66(1):83–114.

Mattoo, A., Freund, C., Mulabdic, A., and Ruta, M. (2023). Is US trade policy reshaping global supply chains? Policy Research Working Paper 10593, World Bank.

Mattoo, A., Ruta, M., and Staiger, R. W. (2024). Geopolitics and the world trading system. Working Paper WB Policy Research WP 11009 / NBER WP 33293, World Bank / National Bureau of Economic Research.

McAfee, R. P. and Schwartz, M. (1994). Opportunism in multilateral vertical contracting: Nondiscrimination, exclusivity, and uniformity. *American Economic Review*, 84(1):210–230.

Mulder, N. (2022). *The Economic Weapon: The Rise of Sanctions as a Tool of Modern War*. Yale University Press, New Haven.

Rambachan, A. and Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, 90(5):2555–2591.

Rasmusen, E. B., Ramseyer, J. M., and Wiley, Jr., J. S. (1991). Naked exclusion. *American Economic Review*, 81(5):1137–1145.

Rey, P. and Tirole, J. (2007). A primer on foreclosure. In Armstrong, M. and Porter, R., editors, *Handbook of Industrial Organization*, volume 3, chapter 33, pages 2145–2220. Elsevier.

Rodrik, D. (2018). Populism and the economics of globalization. *Journal of International Business Policy*, 1(1–2):12–33.

Rogoff, K. (1985). The optimal degree of commitment to an intermediate monetary target. *Quarterly Journal of Economics*, 100(4):1169–1189.

Santos Silva, J. M. C. and Tenreyro, S. (2006). The log of gravity. *Review of Economics and Statistics*, 88(4):641–658.

Schulze, T. and Xin, W. (2025). Demystifying trade patterns in a fragmenting world. IMF Working Paper 2025/129, International Monetary Fund.

Segal, I. (1999). Contracting with externalities. *Quarterly Journal of Economics*, 114(2):337–388.

Segal, I. R. and Whinston, M. D. (2000a). Exclusive contracts and protection of investments. *RAND Journal of Economics*, 31(4):603–633.

Segal, I. R. and Whinston, M. D. (2000b). Naked exclusion: Comment. *American Economic Review*, 90(1):296–309.

Wang, Z., Wei, S.-J., and Zhu, K. (2013). Quantifying international production sharing at the bilateral and sector levels. NBER Working Paper 19677, National Bureau of Economic Research.

Yi, K.-M. (2003). Can vertical specialization explain the growth of world trade? *Journal of Political Economy*, 111(1):52–102.