Export Controls, Directed Innovation, and Third-Party Incidence

Hyungjin Kim*
Korea Development Institute

Sungwan Hong
University of Pittsburgh

May 31, 2026

Abstract

Export controls do more than restrict trade: they change where the target innovates. We study a three-player model in which anticipated US controls redirect Chinese R&D toward restricted, hard-to-substitute inputs. Because China supplies intermediates to third-country firms, this directed innovation shifts input-cost incidence to exposed economies. We show evidence on both margins: Chinese-assignee USPTO patenting rises in directly exposed chip codes around the May 2020 Huawei Foreign Direct Product Rule, and exposed East Asian firms facing higher China-origin input-cost shocks report higher input-cost shares. We calibrate the model using the patent response and compare export controls with credible restraint across the three players. Holding Chinese productivity fixed, total surplus slightly favors controls; once China's R&D allocation responds, the ranking flips in favor of restraint. Relative to restraint, export controls impose most of their input-cost incidence on exposed third parties, amounting to 0.59% of their input-cost base. This creates an institutional gap: third parties could pay for restraint, but no institution binds the United States.

Keywords: Geoeconomics, Endogenous Innovation, Directed-Innovation Allocation Loss, US–China Rivalry, Coercion.
JEL Classification: F13, F51, O31, O33.

*Emails: Hyungjin Kim, hyungjinkim@kdi.re.kr; Sungwan Hong, sungwan.hong@pitt.edu. All errors are our own.

1 Introduction

Export controls are usually scored as bilateral instruments: the coercer pays enforcement and market-access costs, the target loses access, and the policy trade-off is evaluated between the two countries. That perspective misses what controls do before productivity is realized: they change the target's R&D problem. When controls restrict access to bottleneck inputs, the target has an incentive to redirect R&D toward self-supply in the restricted sectors. If the target also supplies intermediates to other countries, the cost of that distorted innovation path can be borne by third-country supply chains rather than by the coercer or target alone.

This paper studies that mechanism in the US–China technology rivalry. Since 2018, US economic statecraft toward China has combined tariffs, entity listings, extraterritorial export controls, CHIPS guardrails, and advanced-computing restrictions. The export-control components target Chinese access to chips, semiconductor equipment, and related bottleneck inputs. We show that these controls change the relative return to R&D across Chinese sectors. They redirect innovation toward restricted self-supply and transmit the resulting directed-innovation allocation loss to *exposed third-party economies* that source Chinese intermediates.

We develop a three-player model of US–China technology rivalry with an endogenous Chinese sectoral R&D allocation. The model has three actors: the United States, China, and exposed third-party economies whose firms buy Chinese intermediates — Korea, Taiwan, the Netherlands, Vietnam, and Mexico are leading examples.

The interaction unfolds over two periods. In period 1, China allocates a fixed aggregate R&D budget across sectors with sector-specific substitutability σ_s, where lower values mean Chinese inputs are harder to replace, while anticipating the period-2 US response. In period 2, the United States observes Chinese productivity realizations and chooses sectoral coercion intensities that determine the strength of export controls in each sector; third-party firms then substitute between Chinese intermediates and US-aligned alternatives while taking input prices as given. We compare the export-control equilibrium with a credible-restraint benchmark in which the United States can commit, before China allocates R&D, not to

impose controls.

Export controls affect China through two opposing channels. First, they restrict China's access to US-aligned upstream alternatives, raising the return to own-R&D in bottleneck sectors. Second, they pressure third-party firms to substitute away from Chinese intermediates, reducing China's sales opportunities in those markets. In hard-to-substitute sectors, the self-supply channel dominates: the more painful foreign restrictions would be, the more valuable Chinese domestic capability becomes. Anticipating this, China redirects R&D toward sectors where controls would otherwise be most painful.

Two empirical exercises play different roles. The FDPR/MIC2025 patent response disciplines the Chinese R&D-redirection elasticity: around the May 2020 Huawei Foreign Direct Product Rule (FDPR), Chinese-assignee USPTO patenting rose sharply in directly exposed chip codes relative to adjacent bottleneck codes.[1] The East Asian firm-account pattern is a directional check on the third-party input-cost channel: firms more exposed to China-origin input-cost shocks report higher input-cost shares.

We use the FDPR/MIC2025 patent response to calibrate η, the model's composite redirection elasticity. In the fixed-productivity version of the model, export controls are slightly preferred. Once China's sectoral R&D allocation responds to anticipated controls, the ranking reverses: credible restraint raises aggregate model surplus, with nearly all of the gain accruing to exposed third parties. Relative to credible restraint, export controls impose most of their input-cost incidence on those third parties, equal to 0.59% of their input-cost base. The bilateral US–China wedge is therefore the wrong summary statistic for incidence.

The incidence pattern creates unrealized surplus. Exposed third-party economies gain enough from *commitment* — US restraint from export controls — to compensate the US for the lost security value, and China would join voluntarily at the calibrated parameters since commitment removes the period-1 R&D distortion. The compact is not an equilibrium

[1]The pattern is concentrated in ChangXin Memory and Yangtze Memory, two state-backed entrants associated with MIC2025/Big Fund Phase II financing. Because MIC2025 financing overlaps with FDPR exposure, we interpret the estimate as a composite redirection elasticity rather than a pure FDPR treatment effect.

outcome of the model. The obstacle is institutional: no existing arrangement both binds the US to restraint and organizes the compensating transfer from exposed third parties. Bilateral trade institutions (WTO, free-trade agreements) handle within-dyad frictions but lack the structural form to organize third-to-coercer payments; multilateral surveillance institutions (IMF) can measure incidence but cannot commit a hegemon to restraint. Absent the transfer, the US weakly prefers the export-control equilibrium, and the compact does not form.

The paper makes three contributions. First, it adds an anticipatory R&D channel to geoeconomic-coercion models. Existing frameworks explain how network position gives the coercer leverage (Farrell and Newman, 2019; Clayton et al., 2024, 2026); here that leverage changes the target's ex ante allocation of a fixed R&D budget. Relative to coercion models that hold rival productivity fixed (Kim, 2026), letting sectoral R&D shape future θ_s is what generates the anticipatory wedge and reverses the policy ranking.

Second, the paper quantifies the incidence of that mechanism. Directed innovation changes the productivity path of Chinese intermediates, so economies that source those intermediates inherit part of the cost of the strategic R&D distortion. This is not a standard Hsieh–Klenow marginal-product-wedge exercise; it is a loss from reallocating a fixed R&D budget across sectors with diminishing returns, measured using the input-cost exposure of third countries. Relative to directed-technical-change models of taxes, subsidies, and market size (Acemoglu, 2002; Acemoglu et al., 2012; Akcigit et al., 2018) and export-control empirics on induced innovation (Liu et al., 2025), the distinctive object is not the fact of R&D redirection but the incidence of the induced input-cost path through third-country supply chains.

Third, the paper identifies the contract missing from the incidence pattern. Sanctions work asks how coercive policies affect bystander countries (Becko, 2023), while trade-agreement theory studies institutions that internalize cross-border policy externalities and the contractual limits of doing so (Bagwell and Staiger, 1999; Maggi and Rodríguez-Clare, 1998; Horn et al., 2010). Here the missing contract is different: exposed third parties would

like to buy credible US export-control restraint before China redirects R&D, but no standard trade or surveillance institution both aggregates their payments and binds US policy after payment. This also separates the mechanism from anticipatory-disintegration work (Becko and O'Connor, 2025): the state variable is rival-side sectoral R&D, and the policy failure is the absence of an institution that aggregates third-party contributions and binds the United States once compensation is paid.

2 Empirical Patterns and Calibration Moments

This section provides one calibration moment and one directional check for the counterfactual. The first exercise uses the May 2020 Huawei Foreign Direct Product Rule (FDPR), which restricted Huawei's access to non-US chips made with US-origin technology. Interpreted together with overlapping MIC2025/Big Fund-backed entry,[2] this patent response disciplines the elasticity with which Chinese R&D is redirected toward restricted chip and semiconductor-equipment bottlenecks. The second asks whether the downstream implication appears in firm accounts: do East Asian firms facing higher China-origin bottleneck input-cost shocks report higher COGS shares? The FDPR/MIC2025 patent response calibrates the model; the firm-account exercise checks the sign of the third-party input-cost channel.

2.1 Setting and data

We use US Patent and Trademark Office (USPTO) PatentsView grant data, filtered to patents with at least one Chinese assignee, and map each patent to a six-digit Harmonized System (HS6) product code using the Lybbert and Zolas (2014) International Patent Classification to HS6 concordance. We date patents by filing year rather than grant year to align patents with inventive activity, and we stop in 2021 to reduce grant-lag censoring. The full code-by-code reconciliation is in Appendix N, Table 16.

[2]MIC2025 is China's "Made in China 2025" industrial-policy program; the Big Fund is its state-backed semiconductor investment fund.

We assign HS6 codes to buckets using a pre-specified critical-input taxonomy, before estimating any patent regression. The 64-code concordance splits into 45 bottleneck codes (rare earths, chips, semiconductor manufacturing equipment, batteries and EVs, aircraft, and other strategic intermediates) and 19 commodity reference codes (textiles, basic plastics, ferrous waste, footwear, and standard consumer goods). Eight codes have zero Chinese-assignee USPTO patents throughout the 2010–2021 window and therefore do not enter the clean matched patent panel. This selection matters for broad bottleneck-versus-commodity descriptions, so Appendix N reports the dropped codes and full reconciliation. It does not affect the directly exposed FDPR chip and semiconductor-equipment codes used for the calibration moment.

Code taxonomy. For the FDPR analysis, the concordance-level bottleneck bucket splits into a 28-code *critical* sub-bucket and a 17-code *strategic* sub-bucket. We restrict to the 27 critical codes that remain in the USPTO panel and partition them into three treatment tiers: Tier 1 contains the 4 directly FDPR-exposed chip and semiconductor-equipment HS6 codes, Tier 2 the 6 chip-adjacent supply-chain codes, and Tier 3 the 17-code internal critical reference set. Tier assignment follows the May 2020 FDPR rule text for Tier 1, chip-supply-chain proximity from industry sources for Tier 2, and the residual critical bucket for Tier 3; the exact HS6 lists are in Appendix N.

2.2 FDPR/MIC2025 patent response

The within-bottleneck patent dose-response around the May 2020 FDPR is the paper's calibration moment for R&D redirection. The FDPR provides sharp timing and a pre-specified exposure gradient within critical inputs: it most directly loads on the four Tier-1 chip and semiconductor-equipment HS6 codes, less directly on the six Tier-2 chip-adjacent codes, and does not directly target the seventeen Tier-3 reference codes. The empirical contrast is

whether patenting rises more in Tier 1 than in Tier 2 or Tier 3 after May 2020.[3]

The object is not a pure FDPR treatment effect. Direct FDPR exposure overlaps with MIC2025/Big Fund-backed entry, so we use the Tier-1 differential as a composite redirection moment rather than a stand-alone causal effect of FDPR alone.

The specification is

$$\log(1 + p_{h,t}) = \beta_1 \, \text{Tier1}_h \cdot \text{Post2020}_t + \beta_2 \, \text{Tier2}_h \cdot \text{Post2020}_t + \alpha_h + \delta_t + \varepsilon_{h,t}, \tag{1}$$

where $p_{h,t}$ is the patent count for HS6 product h in filing year t, Tier1_h and Tier2_h are tier indicators with Tier 3 the omitted reference, Post2020_t equals one for $t \geq 2020$, α_h is an HS6 fixed effect, and δ_t is a year fixed effect. The $\log(1 + p_{h,t})$ transformation accommodates HS6-year cells with zero patents; the resulting coefficient is a semi-elasticity in $\log(1 + p)$ rather than an exact percent change.

Identification and inference. Equation (1) is a fixed-effect dose-response contrast across pre-specified exposure tiers. A causal FDPR-only interpretation would require parallel trends between Tier 1/2 and Tier 3 codes before 2020 and no contemporaneous tier-specific industrial-policy shocks. Those assumptions are too strong in this setting. We therefore use the post-2020 Tier-1 differential as a composite redirection moment, with event-study, firm-decomposition, first-difference, and pre-trend-adjusted variants used to discipline plausible magnitudes.

Inference rests on four treated Tier-1 clusters, so asymptotic standard errors are unreliable; we report three complementary procedures and lean on the first two. *Randomization inference* (RI) permutes tier assignments across the 27 critical HS6 codes 5,000 times under the sharp null of zero effect. *Wild-cluster bootstrap* (WCB, Cameron et al. 2008) resamples cluster-level residuals with random sign weights. With four treated clusters, the standard

[3]Because the policy change occurred in May 2020 while the patent panel is annual, defining Post2020 by filing year may include some pre-policy 2020 filings; this timing convention should attenuate rather than inflate the post-FDPR coefficient.

Table 1: FDPR/MIC2025 within-bottleneck patent response.

Tier	HS6 count	$\hat{\beta}$	RI p	WCB p	Asymptotic p
Tier 1 (directly exposed)	4	+0.774 (0.219)	0.001	0.007	< 0.001
Tier 2 (chip-adjacent)	6	+0.172 (0.172)	0.46	0.39	0.32
Tier 1 − Tier 2 difference		+0.602	0.024	—	—

Notes: Specification is (1) with $\text{Post2020}_t = 1$ for filing years $t \geq 2020$, estimated on the 27-code critical sub-bucket; Tier 3 (17 HS6 codes) is the omitted reference. Cluster-robust standard errors in parentheses. RI uses 5,000 sharp-null permutations of the (Tier 1, Tier 2, Tier 3) assignment over the 27 critical-bucket HS6 codes. Webb-WCB uses the six-point weight distribution of Webb (2014) on the same 27 clusters, $B = 1{,}999$ under the restricted null; this is the recommended cluster-bootstrap weight scheme when $G_{\text{treated}} < 5$ (Cameron et al. 2008). The asymptotic p-value is reported as a reference; we lean on RI and Webb-WCB.

Rademacher (± 1) weights produce a discrete test-statistic floor at $p \approx 1/2^4$, so we use the six-point weight distribution of Webb (2014) with $B = 1{,}999$ draws. The asymptotic p-value enters as a reference row.

Table 1 reports the estimates. Tier 1 is +0.774 (SE 0.219), rejecting the null under randomization inference ($p_{\text{RI}} = 0.001$) and Webb six-point wild-cluster bootstrap ($p_{\text{WCB}} = 0.007$); the asymptotic $p < 0.001$ matches. Tier 2 is +0.172 and statistically indistinguishable from zero ($p_{\text{RI}} = 0.46$, $p_{\text{WCB}} = 0.39$). The Tier 1 − Tier 2 contrast is +0.602, significant under randomization inference ($p_{\text{RI}} = 0.024$). The pattern is a sharp step concentrated on the directly targeted tier, with a much smaller and statistically insignificant response on the adjacent tier. The structural model of Section 5 matches this, predicting a small positive Tier 2 response (+0.198) alongside the much larger Tier 1 dose.

Firm decomposition. The Tier-1 surge is an extensive-margin, entrant-driven phenomenon. The post-2020 increase is concentrated in ChangXin and Yangtze, state-backed memory entrants associated with MIC2025/Big Fund Phase II, while legacy incumbent filings fall. The pattern is precisely the redirection margin the model needs, but it also means $\hat{\beta}_1$ combines export-control-induced self-supply incentives with policy-financed entry.

CNIPA cross-check. CNIPA provides a venue-substitution check. ChangXin and Yangtze also accelerate in CNIPA after 2020, and the aggregate CNIPA Tier-1 coefficient has the same sign as the USPTO response but is imprecise. We therefore read CNIPA as evidence against a pure USPTO venue-substitution story, not as an independent calibration moment.

Robustness. The Tier-1 response remains positive across the main measurement choices, although the magnitude is sensitive to how pre-trends are handled: unfractionated counts, forward-citation weighting (Hall et al., 2001), first differences, and linear pre-trend extrapolation all preserve a positive Tier-1 coefficient (Appendix K). We retain the headline +0.774 as the baseline calibration moment in Section 5; the pre-trend-adjusted +0.617 enters as a conservative variant, shifting the calibrated η from 0.54 to 0.46 without altering the qualitative incidence result.[4]

We do not use the broader bottleneck-versus-commodity patent differential as the calibration moment: Appendix F treats it as descriptive, and Appendix J shows that it is difficult to separate from MIC2025 priority status. The within-critical FDPR/MIC2025 tier response is therefore the paper's calibration moment.

2.3 Third-party input-cost incidence

The model further predicts an incidence pattern: exposed third-party economies that source Chinese intermediates should face higher input costs when China-origin bottleneck inputs become more expensive. We check the sign of this prediction by combining BACI product-level trade values and quantities (Gaulier and Zignago, 2010) with Compustat Global firm accounts. This exercise asks whether the predicted input-cost association appears in firm-level data; calibration remains tied to the FDPR/MIC2025 patent response.

[4]The level estimate (+0.774) reflects both a level shift and a pre-existing trend differential of about +0.157 between Tier-1 and Tier-3 codes over 2010–2017; event-study coefficients range from −1.06 in 2010 to −0.27 in 2017 (Appendix L). The pre-trend-adjusted estimate is +0.617, and the first-difference specification, which differences out this trend, gives +0.177 ($p_{\mathrm{RI}} = 0.010$). The calibration reports the level estimate as the baseline and carries the adjusted and first-difference estimates as conservative variants.

Following the shift-share/Bartik literature (Bartik, 1991; Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022), we implement a leave-one-out shift-share reduced form. For each HS6 product h, year t, and importing country c, we construct a China-origin unit-value shock that excludes country c from the China export price,

$$\Delta uv_{h,t}^{CHN,-c} = \log\left(\frac{\sum_{j\neq c} v_{CHN,j,h,t}}{\sum_{j\neq c} q_{CHN,j,h,t}}\right) - \log uv_{h,\,2014\text{–}2017}^{CHN,-c},$$

where $v_{CHN,j,h,t}$ and $q_{CHN,j,h,t}$ are the export value (USD) and quantity from China to destination j for HS6 product h in year t, and $uv \equiv v/q$ is the unit value.[5] Pre-period country-input-HS6 import shares and downstream input-output weights aggregate this HS6 shock to a country-industry-year exposure

$$Z_{c,m,t}^{uv} = \sum_{h} s_{c,h}^{\text{pre}} \cdot w_{h\to m} \cdot \Delta uv_{h,t}^{CHN,-c}, \tag{2}$$

where $s_{c,h}^{\text{pre}}$ is country c's 2014–2017 average share of China-origin imports for HS6 product h, normalized within each input-producing industry, and $w_{h\to m}$ is the input-output weight mapping HS6 h's input industry to downstream industry m, constructed from the OECD ICIO 2018 global use-table.

The baseline merge sample is Compustat Global firms in Korea, Japan, and Taiwan, the East Asian third parties most directly tied to the China electronics and semiconductor supply chain; Vietnam and Singapore are excluded because Compustat Global coverage of their domestic electronics and semiconductor firms is too thin to support firm-level fixed-

[5] Despite the Δ notation, $\Delta uv_{h,t}^{CHN,-c}$ is a level log-deviation from the 2014–2017 pre-period mean, not a year-on-year change.

effects estimation.[6] We standardize $Z^{uv}_{c,m,t}$ within sample to obtain $\tilde{Z}^{uv}_{c,m,t}$ and estimate

$$y_{f,t} \;=\; \beta\, \tilde{Z}^{uv}_{c(f),m(f),t} \;+\; \alpha_f \;+\; \gamma_{c(f),t} \;+\; \varepsilon_{f,t}, \tag{3}$$

where $y_{f,t}$ is firm f's outcome in year t (COGS share, operating margin, or inventory share), $c(f)$ and $m(f)$ are firm f's country and downstream industry, α_f is a firm fixed effect, $\gamma_{c(f),t}$ is a country-year fixed effect, and standard errors are clustered at the country-industry level. The pre-trend screen in column 3 of Table 2 reports the maximum absolute event-study t-statistic over 2014–2016 from the event-study analog of (3) with year-specific exposure interactions relative to a 2017 base.

Table 2 reports the screen. The model predicts a positive coefficient on $\tilde{Z}^{uv}_{c,m,t}$ for the COGS share: higher Chinese-input prices raise downstream firms' input-cost shares. The baseline estimate has this sign. A one-standard-deviation increase in leave-one-out China-origin unit-value exposure is associated with a 0.20 percentage-point higher COGS share among East Asian firms (+0.00197, SE 0.00076), and the 2014–2016 pre-period event-study screen is benign (maximum absolute t-statistic 1.06 relative to a 2017 base; Appendix Figure 6). The model further predicts a negative operating-margin response if cost pass-through is incomplete and ambiguous inventory effects because precautionary inventory builds can offset profit losses; operating margins move in the predicted negative direction but are imprecise, and inventory shares show no signal. Gross margin from COGS is defined as $1-$COGS share, so the two outcomes are mechanically linked—any change in COGS share is the exact negative of the change in gross margin—and we report COGS share as the cost-incidence outcome.

The baseline coefficient comes from cross-industry exposure differences within country-years. It shrinks to near zero under saturated industry-year fixed effects (-0.00020, $p = 0.89$), which absorb much of the common industry component of the China-origin electronics shock.

[6] We merge Compustat Global firm-year observations to BACI trade-year data using the calendar year corresponding to the firm's fiscal year-end. Japanese fiscal years end in March, so Japanese firm observations are effectively matched to a ∼3-month-lagged subset of the calendar-year BACI shock; Korean and Taiwanese fiscal years coincide with the calendar year. Any attenuation from this mis-alignment biases against finding an incidence pattern.

Table 2: China-origin unit-value exposure and East Asian firm cost shares.

Outcome	Baseline	With industry-year FE	Pretrend max $\lvert t \rvert$
COGS share	$+0.00197^{***}$ (0.00076)	-0.00020 (0.00147)	1.06
Operating margin	-0.00131 (0.00187)	-0.00654 (0.00542)	1.17
Inventory share	-0.00007 (0.00092)	-0.00022 (0.00431)	0.57

Notes: Estimates of β in (3). Baseline sample is Compustat Global firms in Korea, Japan, and Taiwan, with firm and country-year fixed effects; the second column adds industry-year fixed effects. Standard errors clustered by country-industry are in parentheses; the shock is standardized within sample. The COGS sample has 66,184 firm-year observations, 8,156 firms, and 242 country-industry clusters; operating-margin and inventory-share samples are restricted to firm-years with non-missing outcome values and are within a few percent of the COGS sample in each dimension. The pretrend column reports the maximum absolute pre-period event-study t-statistic over 2014–2016 from the event-study analog of (3) with 2017 as the base year; this column reports a t-statistic, not a coefficient. Gross margin from COGS is omitted because it equals 1 − COGS share. $^{***}p < 0.01$.

We therefore use the BACI/Compustat exercise as a directional check on the model's third-party input-cost channel, while Section 5 quantifies incidence using the calibrated structural exercise. An EU-inclusive sample preserves the COGS/gross-margin sign but is less precise (Appendix M).

3 Model

The model shows how anticipated export controls change China's sectoral R&D problem before productivity is realized. China allocates a fixed R&D budget across sectors; future effective controls raise the marginal value of self-supply in restricted bottlenecks; and the resulting productivity path changes input costs for exposed economies that use Chinese intermediates. Section 3.1 introduces the players and timing; Section 3.2 specifies the period-2 coercion block; Section 3.3 solves China's period-1 R&D problem; Section 3.4 defines equilibrium.

The period-1 redirection result leans on only two reduced-form regularities of the period-2 coercion response, stated in Section 3.2 as Assumption 1. Computing welfare *levels* additionally requires a calibrated period-2 payoff implementation. Our baseline implementation is summarized in Section 3.2 and stated formally in Appendix B.1. Appendix B.1 verifies that

the redirection logic survives under a solver-free reduced-form schedule satisfying the same regularities.

Notation conventions. Function-of-argument notation denotes endogenous mappings, e.g. $\theta_s(z_s)$ and $\Lambda(\theta_s, \sigma_s)$. Bars denote fixed or equilibrium-defined thresholds, and stars denote sector-level equilibrium values.

3.1 Players and timing

There are three actors: the United States, China, and a representative exposed third-party economy. We index them by U, C, and S, respectively. There are N sectors $s = 1, \ldots, N$. Each sector has a fixed bottleneck primitive $\sigma_s \in (0, 1]$, the elasticity of substitution between Chinese inputs and US-aligned alternatives in the exposed economy's downstream production. Lower σ_s means Chinese inputs are harder to replace, so upstream complementarity is stronger and the sector is a tighter bottleneck.[7]

The label "third party" compresses three empirical roles that the model keeps distinct. The player S is a representative downstream exposed economy: its firms use Chinese intermediates and choose the alignment/substitution margin x_s. Upstream third-country suppliers that comply with FDPR-style rules are not additional strategic players in the baseline; their behavior is summarized by the supplier-compliance rate q_s inside the effective coercion map. When Section 4.4 discusses compact finance, k indexes governments of exposed economies, or a coalition institution acting on their behalf, not the downstream firms represented by S.

Timing is two periods.

- *Period 1 (R&D stage).* China picks a sectoral R&D vector $z = (z_1, \ldots, z_N)$ subject to a fixed budget $\sum_s z_s = B$ and non-negativity $z_s \geq 0$, with perfect foresight over the

[7] The baseline is a one-factor bottleneck benchmark: the same low-substitutability primitive raises China's self-supply premium through $h(\sigma_s) = \sigma_s^{-\eta}$ and raises the exposed economy's input-cost sensitivity through the CES cost aggregator introduced in Section 3.2. This is not required conceptually. The mechanism requires overlap between sectors where coercion raises China's marginal R&D return and sectors that matter for the exposed economy's input costs; a richer version could separate the self-supply primitive from the exposed economy's substitution elasticity.

period-2 coercion response.

- *Period 2 (coercion + alignment stage).* Given the Chinese productivity vector $\theta = (\theta_1, \ldots, \theta_N)$ that emerges from period 1, the United States picks a coercion intensity $\lambda_s \in [0, 1]$ and the exposed economy best-responds with an alignment level x_s; the equilibrium map $\lambda_s^* = \Lambda(\theta_s, \sigma_s)$ summarizes the period-2 outcome (Section 3.2).

Sector s's productivity depends on R&D effort according to

$$\theta_s(z_s) = \theta_s^0(1 + z_s)^{\varphi_s}, \qquad \varphi_s \in (0, 1). \tag{4}$$

The concavity in (4) captures diminishing returns to sector-specific R&D, and strict concavity is the hypothesis of the Jensen Allocation Lemma (Lemma 1).[8] We carry sector-specific φ_s and θ_s^0 for generality and specialize them in the calibration (Section 5.1).

3.2 The period-2 coercion block

The period-2 block is summarized by an effective coercion map

$$\lambda_s^* = \Lambda(\theta_s, \sigma_s) \in [0, 1].$$

This map is reduced-form. It combines US enforcement, upstream third-country supplier compliance, and the exposed economy's downstream substitution response. The period-1 redirection result uses only two properties of this map: controls become active above a sectoral threshold, and, on the active high-leverage support, effective controls rise with the strategic value of Chinese productivity. The payoff implementation below is needed for the welfare and incidence levels in Section 5.

Across sectors, the United States' total payoff and the exposed economy's total welfare are additive in per-sector terms: $W_U = \sum_s U^s(\theta_s, \lambda_s)$ and $W_S = \sum_s \omega_s V_S^s$, with exposure weights

[8] A linear specification $\theta_s = \theta_s^0 + \varphi_s z_s$ delivers identical comparative-static signs but corner-solves on the fixed-budget constraint at empirically relevant parameters, motivating the multiplicative concave form.

$\omega_s > 0$. These additive structures rule out cross-sector spillovers in coercion costs and in input substitution beyond the within-sector CES aggregator; sector-by-sector maximization of U^s over λ_s delivers the optimal λ vector.

The intensity λ_s should be read as effective coercion rather than the statutory coverage of a rule. Extraterritorial controls such as the FDPR bind through third-country compliance. A simple implementation is to let e_s denote US enforcement intensity and $q_s(e_s)$ the share of upstream suppliers in third countries that comply in sector s, so that

$$\lambda_s^{\text{eff}} = e_s q_s(e_s),$$

with the effective superscript suppressed below. Supplier j complies when the cost of violating the rule exceeds the China-market return,

$$P_{js}(e_s) + A_{js} \geq \pi_{js}^{CN},$$

where P_{js} is the expected penalty or license-loss risk, A_{js} captures dependence on US-origin technology or US-market access, and π_{js}^{CN} is the profit from serving Chinese buyers. FDPR-style rules raise e_s; US-origin tool dependence and US-market exposure make compliance more attractive. The baseline model absorbs the resulting compliance rate into Λ. Thus the redirection results use the effective coercion map, not an assumption that legal coverage alone blocks Chinese access.

The exposed economy best-responds with an alignment level x_s, a substitution toward US-aligned alternatives, given complementarity σ_s. Its laissez-faire input cost is the CES aggregate

$$C(x_s; w_C, w_A) = \left[a(x_s) w_C^{1-\sigma_R} + \big(1 - a(x_s)\big) w_A^{1-\sigma_R}\right]^{1/(1-\sigma_R)}, \qquad \sigma_R = \sigma_s, \quad w_C(\theta_s) = \bar{w}/\theta_s,$$

and its per-sector welfare is $V_S^s = -C(x_s; w_C(\theta_s), w_A)$. Economically, x_s is the exposed

economy's costly reorientation away from Chinese inputs. The reduced-form payoff below is the United States' value of implementing alignment x_s after netting out the input-cost burden that makes the alignment feasible for the exposed economy. It is not the exposed economy's private objective. The United States' payoff at sector-s alignment x_s is

$$J(x_s; \theta_s) = V_{US}(x_s) - C(x_s; \bar{w}/\theta_s, w_A),$$

where $V_{US}(x_s)$ is the United States' strategic value from the exposed economy's alignment toward US-aligned supply, the national-security objective behind export controls in the weaponized-interdependence tradition (Farrell and Newman, 2019; Clayton et al., 2024, 2026). It is increasing and concave in x_s. At $\lambda_s = 0$, the implementation block induces a baseline alignment $x_G^s(\theta_s) = \arg\max_x J(x; \theta_s)$. Coercion raises alignment above this baseline at a deadweight cost.

The United States chooses an intensity $\lambda_s \in [0, 1]$ that reduces the exposed economy's reliance on Chinese intermediates, weighing the value it extracts from China's productive capacity against a quadratic deadweight cost $(\gamma/2)\lambda_s^2$, where $\gamma > 0$ governs the marginal deadweight cost of coercion. A higher λ_s is costly pressure that buys strategic alignment by making reliance on Chinese inputs less attractive; the United States uses it only when the strategic value of shifting the sector exceeds enforcement, compliance, and deadweight costs. The US per-sector Period-2 objective decomposes as

$$U^s(\theta_s, \lambda_s) = J(x_G^s; \theta_s) + \Pi_E(\lambda_s; \theta_s, \sigma_s) - \tfrac{\gamma}{2}\lambda_s^2, \tag{5}$$

where $J(x_G^s; \theta_s)$ is the US payoff at the laissez-faire alignment, Π_E is the coercion-extraction value (incremental surplus from inducing $x_s(\lambda_s) > x_G^s$, with $\Pi_E(0; \cdot) = 0$), and the last term is the deadweight cost. $J(x_s; \theta_s)$ is the primitive payoff at any sector-s alignment x_s; $U^s(\theta_s, \lambda_s)$ is the corresponding Period-2 objective the United States maximizes over λ_s, with

the exposed economy best-responding. The induced equilibrium intensity is

$$\lambda_s^* = \Lambda(\theta_s, \sigma_s) = \arg\max_{\lambda_s \in [0,1]} U^s(\theta_s, \lambda_s).$$

At the induced equilibrium with corresponding alignment $x_s^* = x_s(\lambda_s^*)$, Π_E satisfies the value identity

$$\Pi_E(\lambda_s^*; \theta_s, \sigma_s) = \gamma(\lambda_s^*)^2 + J(x_s^*; \theta_s) - J(x_G^s; \theta_s), \tag{6}$$

implying the optimal-λ envelope $U^s(\theta_s, \lambda_s^*) = J(x_s^*; \theta_s) + \frac{\gamma}{2}(\lambda_s^*)^2$. Appendix B.1 derives both from the US first-order condition; the general form $\Pi_E(\lambda)$ at off-optimum λ is implicit through the exposed economy's best response. At no coercion, $U^s(\theta_s, 0) = J(x_G^s; \theta_s)$ since $\Pi_E(0; \cdot) = 0$. We refer to the period-2 primitives — the CES cost aggregator C, the US payoff J, the extraction value Π_E, and the deadweight cost $(\gamma/2)\lambda_s^2$ — jointly as the *baseline* Λ-*payoff block*. Welfare levels are computed under this block.

We impose two regularities on the equilibrium coercion map. P1 is target-value feedback on the active high-leverage support: holding enforcement technology and export-control leverage fixed, a more productive Chinese sector creates more strategic value from restricting or redirecting dependence on China. P2 captures fixed leverage and enforcement costs, so controls stay dormant until the target is valuable enough to justify activation.

Assumption 1 (Coercion-block regularities)**.** *For each sector s, $\Lambda(\theta, \sigma_s)$ is continuous, piecewise continuously differentiable, and satisfies:*

P1. Target-value feedback on the active high-leverage support. *Holding enforcement technology and export-control leverage fixed, $\Lambda_\theta(\theta, \sigma_s) > 0$ in sectors where controls are feasible and active: a more productive Chinese sector is subject to stronger effective controls.*

P2. Threshold. *There exists $\bar{\theta}(\sigma_s)$ such that effective controls are zero below the threshold, $\Lambda(\theta, \sigma_s) = 0$ for $\theta \leq \bar{\theta}(\sigma_s)$, and positive above it, $\Lambda(\theta, \sigma_s) > 0$ for $\theta > \bar{\theta}(\sigma_s)$.*

Assumption 1 defines the class of period-2 policy-response functions allowed by the model: any Λ in this class generates the period-1 redirection results of Section 3.3 through Λ_θ and the threshold structure. Sectors where controls are infeasible, or where higher Chinese productivity mainly raises enforcement costs, are represented by a high activation threshold or by $\Lambda(\cdot, \sigma_s) \equiv 0$. Appendix B.1 defines the equilibrium marginal extraction value $M(\theta, \sigma) \equiv x_s^* \theta_s - J(x_s^*; \theta_s) = \gamma \lambda_s^*$ and shows that P1 and P2 follow from two conditions on M — target-value monotonicity ($M_\theta > 0$) and an activation threshold (Lemma 3). The period-1 redirection result uses these two economic regularities of the equilibrium extraction value; Appendix B.1 reports a solver-free closed-form schedule satisfying P1 and P2 by construction. The explicit baseline payoff block enters the welfare-level quantification.

Remark 1 (Export-control leverage parameter). The coercion intensity λ_s is not a free policy dial. Its feasibility is governed by an export-control *leverage parameter* $\ell_s \geq 0$: the extent of extraterritorial, US-origin-technology control over sector s's supply chain — the share of the chain that runs through US tools, US software, or US-controlled chokepoints, of which the Foreign Direct Product Rule is the canonical instrument. Coercion in sector s is feasible and binding only where ℓ_s is large enough that the United States can credibly deny China access to US-origin technology and induce upstream third-country suppliers to withhold controlled inputs; where ℓ_s is negligible, the activation threshold $\bar{\theta}(\sigma_s)$ of Assumption 1 is effectively never crossed and $\Lambda(\cdot, \sigma_s) \equiv 0$. We hold ℓ_s in the background rather than carry it as a separate state variable: in the calibration it is absorbed into the sector-specific activation threshold and the realized Λ schedule, and the FDPR tier taxonomy of Section 2.2 is its empirical counterpart (Tier 1 codes are precisely the high-ℓ_s, US-tool-dependent chip and semiconductor-equipment lines). The economic content is that λ_s is *activated* by extraterritorial leverage, which is why the extra R&D incentive appears in sectors where the United States has real supply-chain leverage rather than in an arbitrary set of sectors.

3.3 Period 1: China's R&D problem

China's period-2 sectoral surplus has two components. The first is revenue from selling intermediates, attenuated by US coercion — the market-access channel. The second is a forced-substitution premium: when China is cut off from US-aligned upstream alternatives in a bottleneck sector ($\sigma_s < 1$), its return to own-R&D rises. The premium is not that coercion is intrinsically good for China; it is the shadow value of domestic capability when foreign inputs become unreliable. The lower is σ_s, the harder the missing input is to replace, so an own-sector productivity gain relaxes a tighter constraint. We summarize the per-unit-productivity surplus by the *effective surplus rate*

$$\Pi_s(\lambda_s, \sigma_s) = (1 - \lambda_s) + \lambda_s h(\sigma_s), \qquad h(\sigma_s) = \sigma_s^{-\eta}, \quad \eta > 0. \tag{7}$$

Π_s is a convex combination of two terms: the market-access channel, normalized to one, and the forced-substitution channel $h(\sigma_s)$. Since $h(\sigma_s) > 1$ for $\sigma_s < 1$, effective controls raise the marginal value of own-sector productivity inside China's R&D problem in bottleneck sectors. This is not a welfare gain from coercion; it is the shadow value of domestic capability when foreign access is unreliable. The parameter η is the composite redirection elasticity calibrated from the FDPR/MIC2025 patent response; it should not be read as a pure substitution elasticity or a pure FDPR treatment effect.

Let $\Lambda_s(z_s) \equiv \Lambda(\theta_s(z_s), \sigma_s)$. China's private period-1 objective is the producer/innovation payoff

$$\widetilde{W}_C(z) \equiv \sum_s \theta_s(z_s) \Pi_s\big(\Lambda(\theta_s(z_s), \sigma_s), \sigma_s\big), \tag{8}$$

which it maximizes subject to a fixed R&D budget,

$$\max_{z \geq 0} \sum_s \theta_s(z_s) \Pi_s\big(\Lambda(\theta_s(z_s), \sigma_s), \sigma_s\big) \quad \text{s.t.} \quad \sum_s z_s = B. \tag{9}$$

We call (9) *Setup B* (fixed budget): aggregate R&D is fixed, so anticipated coercion induces

only sectoral *redirection* of effort. Appendix B.12 reports a hybrid R&D-cost model nesting Setup B and an unconstrained level-response alternative (Setup A), under which coercion can also raise aggregate R&D; the roughly flat Chinese R&D-to-GDP path over the coercive period supports the fixed-budget reading.

Gains, losses, and common units. Throughout, a *regime* fixes the US period-2 coercion policy — SPE (the optimal coercion map) versus Commit ($\lambda \equiv 0$). The regime comparison uses three regime-specific incidence objects,

$$W_C = \sum_s \theta_s \Pi_s(\lambda_s, \sigma_s), \qquad W_U = \sum_s U^s(\theta_s, \lambda_s), \qquad W_S = -\sum_s \omega_s C_s\big(\bar{x}_s; \bar{w}/\theta_s, w_A\big), \quad (10)$$

where each object is evaluated at its regime's allocation — $(\theta_s, \lambda_s) = (\theta_s^{\text{SPE}}, \Lambda(\theta_s^{\text{SPE}}, \sigma_s))$ under SPE and $(\theta_s, \lambda_s) = (\theta_s^{\text{Commit}}, 0)$ (hence $\Pi_s = 1$) under Commit — and $\omega_s > 0$ are the exposed economy's baseline cost-share exposure weights in sector s. The China object W_C captures producer and innovation payoffs; the US object W_U captures the strategic value of coercion net of its deadweight cost; the exposed-economy object W_S captures input-cost incidence, evaluated at a fixed alignment $\bar{x}_s$, so the comparison isolates the effect of the induced productivity (θ) path rather than of any alignment response. The three objects are measured in the model's common units.

The private objective $\widetilde{W}_C(z)$ in (8) is the function China maximizes over z within a fixed coercion regime; the national object W_C in (10) is its realized value, compared *across regimes*. Revealed preference ranks allocations within a regime. It does not sign $\Delta W_C = W_C^{\text{Commit}} - W_C^{\text{SPE}}$, because the cross-regime comparison changes both Π_s and the induced productivity path.

The exact first-order condition. Computing the first-order condition for (9) requires the indirect effect of z_s on Π_s through induced coercion (θ_s enters λ_s^* via target-value feedback). Differentiating the sector-s objective $\theta_s(z_s)\Pi_s(\Lambda_s(z_s), \sigma_s)$ and using $\lambda_s^* = \Lambda(\theta_s, \sigma_s)$

gives $\partial_{z_s}[\theta_s \Pi_s] = \theta_s'(z_s)\,\Gamma_s(z_s;\eta)$, where the total marginal return factor is

$$\Gamma_s(z_s;\eta) \equiv \Pi_s(\Lambda_s(z_s),\sigma_s) + \theta_s(z_s)\Pi_{\lambda,s}(\Lambda_s(z_s),\sigma_s)\Lambda_\theta(\theta_s(z_s),\sigma_s), \tag{11}$$

where $\Pi_{\lambda,s} = h(\sigma_s) - 1 = \sigma_s^{-\eta} - 1$. The first term is the ordinary payoff from one more unit of productivity. The second term is strategic feedback: extra productivity changes the future effective-control response, and in bottlenecks that response raises the shadow value of domestic self-supply. By P1, $\Lambda_\theta > 0$ in the coercive region. The second term is therefore part of the structural marginal return, and in bottleneck sectors ($\sigma_s < 1$) it amplifies China's marginal return to R&D. The exact KKT condition for Setup B is

$$\theta_s'(z_s^*)\Gamma_s(z_s^*;\eta) \begin{cases} = \mu, & z_s^* > 0, \\ \leq \mu, & z_s^* = 0, \end{cases} \tag{12}$$

with μ the budget multiplier. Under commitment to $\lambda \equiv 0$, $\Pi_s = 1$ and $\Gamma_s = 1$ in every sector.

The exact problem (12) is solved numerically; the closed form below is used only to read off comparative statics. When the feedback term is small relative to the direct term ($R(\sigma,\eta,\theta) < 1$, with R the feedback-to-direct ratio defined in Appendix B.3), the interior first-order condition reduces to

$$\theta_s'(z_s^*)\Pi_s(\lambda_s^*,\sigma_s) = \mu, \qquad \forall s : z_s^* > 0, \tag{13}$$

and substituting (4) gives the closed form

$$1 + z_s^* = \left(\frac{\varphi_s \theta_s^0 \Pi_s(\lambda_s^*,\sigma_s)}{\mu}\right)^{1/(1-\varphi_s)}. \tag{14}$$

Equation (14) factors China's sector-s effort into a baseline term (θ_s^0) and a coercion-induced premium (Π_s), modulated by the cross-sector multiplier μ. This fixed-wedge form is useful

only for intuition: the feedback term is quantitatively material at empirically relevant parameters, so the headline calibration solves the exact constrained problem (12) and reports the realized Γ_s^* values.

Payoff concavity for the welfare analysis. The third-party incidence result requires the US no-coercion payoff to be increasing and concave in Chinese productivity: concavity makes the US gain from a more productive China diminish at the margin, the curvature the incidence result relies on.

Assumption 2 (Payoff concavity)**.** *$U^s(\theta_s, 0) = J(x_G^s; \theta_s) = V_{US}(x_G^s) - C(x_G^s; w_C(\theta_s), w_A)$ is weakly concave and strictly increasing in θ_s on the relevant region.*

CES sufficient condition. Suppose the exposed economy's laissez-faire cost function has CES form

$$C(x; w_C, w_A) = \left[a(x) w_C^{1-\sigma_R} + (1 - a(x)) w_A^{1-\sigma_R}\right]^{1/(1-\sigma_R)}, \qquad \sigma_R \in (0, 2), \tag{15}$$

with positive rival-input exposure. Then $J(x_G; \theta) = V_{US}(x_G) - C(x_G; \bar{w}/\theta, w_A)$ is strictly increasing and weakly concave in θ on $\theta > 0$ (strictly concave when the rival input has positive cost share). Appendix B.2 proves the result for fixed alignment. When $x_G(\theta)$ is interpreted as an endogenous no-coercion alignment, Assumption 2 is imposed on the optimized value and verified numerically on the calibration grid; the analytical CES proof is the fixed-alignment sufficient condition used for the input-cost incidence object.

The empirically relevant range is $\sigma_R = \sigma_s \in (0, 1]$, so Assumption 2 holds analytically throughout the baseline calibration region; it is also verified numerically on 100% of cells across the working region. Economically, this is the standard CES property that as China becomes more productive ($w_C = \bar{w}/\theta$ falls), the exposed economy's laissez-faire input cost falls, raising the US payoff at any given alignment.

3.4 Equilibrium

Definition 1 (Subgame-Perfect Equilibrium)**.** *A subgame-perfect equilibrium is a tuple* (z^*, x^*, λ^*) *such that (i) for each* s, $\lambda_s^* = \Lambda(\theta_s(z_s^*), \sigma_s)$ *and* x_s^* *is the exposed economy's period-2 best response in the baseline* Λ*-payoff block of Section 3.2; and (ii)* z^* *solves* (9) *given the period-2 coercion-response map.*

Existence follows from continuity of China's objective and compactness of the fixed-budget simplex. Continuity is guaranteed by Assumption 1, including at the threshold kink. Uniqueness is not guaranteed globally because $\Gamma_s(z_s; \eta)$ depends on z_s through both θ_s and the induced response $\lambda_s^*(\theta_s, \sigma_s)$, so the exact problem need not be concave; the numerical implementation uses a multi-start constrained optimizer.

4 Results

The results below trace the mechanism: anticipated coercion creates sectoral R&D wedges, those wedges redirect innovation, and the resulting productivity path is priced through third-party input costs. Formal proofs are collected in Appendix B.

4.1 Innovation redirection

The next result is the paper's core R&D redirection condition. It asks which sectors gain R&D when China reallocates a fixed budget in anticipation of future controls. The result holds for any coercion-response map Λ satisfying P1–P2 (Assumption 1), not for a particular period-2 solver or microfoundation.

Consider the mechanism in a two-sector example, with one bottleneck input and one commodity input. China has a fixed R&D budget. Under commitment, future coercion is absent, so the two sectors differ only through the usual concave productivity returns. Under coercion, the bottleneck sector receives an additional self-supply premium because future controls make domestic production more valuable. Its KKT return becomes $\theta_B'(z_B)\Gamma_B$.

R&D moves toward the bottleneck only when Γ_B is large enough to overcome the common shadow value of R&D funds. Since productivity is concave in R&D, that concentration lowers aggregate productivity relative to a more balanced allocation, and third parties that source Chinese intermediates face the cost through higher input prices.

Proposition 1 (Innovation redirection)**.** *Under Setup B with multiplicative concave* $\theta_s(z_s) = \theta_s^0(1+z_s)^{\varphi_s}$*, suppose the SPE and Commit allocations are interior and differentiable. Let* $\Gamma_s^* \equiv \Gamma_s(z_s^{SPE};\eta)$ *denote the realized marginal R&D return factor under anticipated controls in* (11)*, and let*

$$\bar{\Gamma} \equiv \frac{\mu^{SPE}}{\mu^{Commit}}.$$

Then, sector by sector,

$$z_s^{SPE} > z_s^{Commit} \quad \Longleftrightarrow \quad \Gamma_s^* > \bar{\Gamma}.$$

The proof is in Appendix B.4.[9]

In words, export controls redirect R&D only when they raise a sector's marginal value of innovation by more than the economy-wide shadow value of R&D funds. The result is not that all restricted sectors receive more R&D. Redirection occurs only where anticipated coercion raises the return to self-supply enough to overcome the fixed-budget opportunity cost.

If the realized return factor is continuous and decreasing single-crossing in σ_s on the calibrated active support, the redirection condition can be summarized by a bottleneck cutoff. Sectors below the cutoff receive more R&D under SPE than under commitment; sectors above it receive less. This is a ranking result, not a claim that every low-σ sector mechanically receives more R&D: a sector must also be sufficiently valuable or leveraged to cross the activation threshold. If the return factor does not cross $\bar{\Gamma}$, all sectors lie on the same side of the redirection condition; the threshold-flip case is treated in Remark 3.

[9] At kinks in Λ the same comparison uses one-sided KKT/subgradient conditions. The commitment regime is the *ex ante* commitment $\lambda \equiv 0$, distinct from the endogenous period-2 no-coercion region in which a sector remains below the activation threshold.

Remark 2 (Generic policy wedges). The fixed-budget KKT logic is generic. If an exogenous policy wedge τ_s raises China's sector-s R&D payoff to $\theta_s(z_s)(1+\tau_s)$, the interior fixed-budget condition is

$$\theta_s'(z_s)(1+\tau_s) = \mu.$$

Any dispersion in τ_s redirects R&D toward high-wedge sectors. Export controls matter here because the wedge is generated by extraterritorial US-origin technology leverage, turns on only after a coercion threshold is crossed, and operates through third-country supplier compliance and downstream exposed-economy substitution. These features make the same redirection force geoeconomic rather than a standard tariff, subsidy, or procurement-wedge reallocation.

Remark 3 (Threshold flip). If the lowest-σ sectors remain below the activation threshold, their coercion premium is one and redirection loads instead on intermediate-σ sectors that have crossed into the coercive region. As baseline productivity rises and lower-σ sectors cross their thresholds, the redirection target shifts toward those sectors.

4.2 The Jensen allocation building block

We separate the productivity-side benchmark from the welfare-side result. The first (Lemma 1) is the *Jensen allocation benchmark*: under strictly concave $\theta_s(z_s)$, uniform R&D maximizes aggregate productivity. With diminishing R&D returns, moving a dollar from an overfunded sector to an underfunded one raises total productivity until marginal products are equalized. This is a concave-allocation (Jensen) loss in the *direction* of innovation, not a Hsieh–Klenow marginal-product-wedge object; we call the resulting welfare cost a directed-innovation allocation loss. The second (Lemma 2) is its welfare-side counterpart and signs the third-party input-cost object directly.

Lemma 1 (Jensen Allocation Lemma)**.** *Suppose* $\theta_s(z_s) = \theta^0(1+z_s)^{\varphi}$ *with parameters* $\theta^0 > 0$

and $\varphi \in (0,1)$ uniform across s. *For any fixed-budget allocation* z *with* $\sum_s z_s = B$, *define*

$$\mathcal{M}(z) \equiv \frac{\sum_s (1+z_s)^\varphi}{N^{1-\varphi}(B+N)^\varphi}.$$

Then $\mathcal{M}(z) \leq 1$, *with equality if and only if* $z_s = B/N$ *for all* s. *Hence the uniform commitment allocation maximizes aggregate productivity, and any non-uniform allocation strictly reduces it.*

The proof is Jensen's inequality and is given in Appendix B.6.

Remark 4 (Fixed-wedge sufficient-statistic benchmark). In a fixed-realized-wedge benchmark, dispersion in sectoral premiums Γ_s induces a non-uniform R&D allocation and therefore $\mathcal{M} < 1$, with equality only when all Γ_s are equal. Export controls lower aggregate productivity through the dispersion they create in R&D wedges. Appendix B.6 gives the closed form. In the full structural model Γ_s is endogenous to z_s through θ_s and λ_s^*, so this is a diagnostic, not a global solution.

Remark 5 (Exposure-weighted benchmark). Lemma 1 is the equal-exposure productivity benchmark. With heterogeneous exposure weights, the relevant benchmark tilts toward sectors that matter more for the third party's input bundle. The unweighted $\mathcal{M}(z)$ should be read as the clean aggregate-productivity benchmark.

Lemma 1 is a productivity-side benchmark: it concerns the unweighted aggregate productivity index $\sum_s \theta_s$. The welfare object that signs third-party incidence is the weighted, nonlinear input cost

$$W_S(z) = -\sum_s \omega_s \, g_s(z_s), \qquad g_s(z_s) \equiv C_s(\bar{w}/\theta_s(z_s)),$$

where $\omega_s > 0$ are exposure weights and g_s holds the fixed alignment $\bar{x}_s$ and outside price w_A of (10) constant along the directed-innovation channel, so the rival-input price $\bar{w}/\theta_s$ is the only argument varying with z_s. The next lemma is the welfare-side counterpart of Lemma 1. It

signs the symmetric-exposure case analytically and states exactly what the calibrated model must verify when exposure is heterogeneous.

Lemma 2 (Third-party input-cost incidence)**.** *Maintain the baseline technology* $\theta_s(z_s) = \theta^0(1+z_s)^\varphi$ *with uniform* θ^0 *and* $\varphi \in (0,1)$*, and the CES cost family above with* $\sigma_R \in (0,2)$*. Hold alignment fixed along the directed-innovation channel and restrict attention to sectors with positive rival-input exposure.*

If exposure and cost primitives are symmetric across sectors ($\omega_s \equiv \omega$ *and* $g_s \equiv g$*), then*

$$W_S(B/N, \ldots, B/N) \geq W_S(z)$$

for every fixed-budget allocation z*, with strict inequality whenever* z *is non-uniform. Hence* $\Delta W_S \geq 0$*, strictly if* z^{SPE} *is non-uniform. With heterogeneous exposure or cost primitives, the sign of* ΔW_S *is the weighted input-cost comparison*

$$\Delta W_S = \sum_s \omega_s g_s(z_s^{SPE}) - \sum_s \omega_s g_s(B/N).$$

Lemma 2 reads through a small set of objects. Each $g_s(z_s)$ is the third party's sectoral input cost when China allocates z_s R&D to sector s; $g_s' < 0$ (more Chinese R&D lowers the rival-input price) and $g_s'' > 0$ (the cost reduction is convex, so uneven R&D is costly). The downstream logic mirrors R&D concavity: extra Chinese productivity lowers input prices most where the sector has received little R&D, so concentrated innovation leaves expensive gaps in the exposed economy's input bundle. Third-party welfare $W_S = -\sum_s \omega_s g_s(z_s)$ rises when weighted input costs fall; with heterogeneous exposure the third party can prefer some non-uniform allocations because gaps in high-use sectors count more.

The proof is in Appendix B.7. In the symmetric case, R&D dispersion raises total input costs because the sectoral input-cost function is convex in R&D — the input-cost analogue of Jensen's inequality; the formal proof uses Karamata's inequality. With heterogeneous exposure the sign becomes quantitative: the calibrated exercise in Section 5 compares the

weighted input-cost object under the observed allocation and the commitment benchmark.

Jensen supplies the transparent symmetric benchmark: concentration of a fixed R&D budget is costly when sectoral productivity is concave. The incidence claim is the weighted input-cost comparison in Lemma 2, evaluated on the empirical θ paths and the third party's exposure weights. Uniform commitment is therefore a normalization for the clean benchmark and an upper-bound counterfactual in Section 5.2; the sign and magnitude reported below come from the nonlinear, exposure-weighted cost object, with heterogeneous-benchmark formulas and attenuation diagnostics used to show how much of the symmetric benchmark is needed for the incidence result to survive.

Scope of the uniform benchmark. The uniform benchmark isolates the cost of concentrating a fixed R&D budget across sectors. With heterogeneous θ_s^0, φ_s, or exposure weights the efficient fixed-budget benchmark need not be uniform, so the quantitative exercise evaluates the weighted input-cost object directly rather than relying on the unweighted $\mathcal{M}(z)$ statistic. Appendix B.6 reports the heterogeneous-benchmark formulas. Conditional on the empirical SPE and Commit θ paths, the pure input-cost directed-innovation channel is separated from η; η enters the static payoff channel and the calibration of R&D wedges.

4.3 Third-party incidence

The SPE allocation maximizes China's private objective under the coerced payoff. That revealed-preference statement does not sign China's national payoff wedge, because commitment changes both the payoff wedge Π_s and the induced productivity path. The next proposition lifts the two allocation lemmas into regime-level incidence: bilateral welfare is ambiguous, but the third-party term has a direct input-cost interpretation.

Proposition 2 (Incidence)**.** *Let* $\Delta W_X \equiv W_X^{Commit} - W_X^{SPE}$ *for* $X \in \{C, U, S\}$*, and* $\Delta W_{total} \equiv \Delta W_C + \Delta W_U + \Delta W_S$ *(the common-numeraire aggregate).*

(i) Bilateral sign indeterminacy. *The bilateral wedges* ΔW_C *and* ΔW_U *are sign-indefinite:*

for China, commitment removes the forced-substitution premium but can restore a more productive fixed-budget R&D allocation; for the United States, commitment raises no-coercion productivity but removes the period-2 coercion-extraction value Π_E, net of its deadweight cost.

(ii) Third-party incidence. *Under the symmetric benchmark of Lemma 2, $\Delta W_S \geq 0$, with strict inequality whenever the SPE R&D allocation is non-uniform. With heterogeneous exposure or cost primitives, the sign of ΔW_S is the weighted input-cost comparison in Lemma 2.*

The proof is in Appendix B.9.

Proposition 2 separates two questions that are often conflated. The bilateral effects are ambiguous because commitment removes coercive premia but may restore a more productive R&D allocation. The third-party effect is different: in the symmetric benchmark concentration of R&D raises downstream input costs; with heterogeneous exposure the sign is the calibrated input-cost comparison in Section 5. Bilateral wedges may be small or sign-indeterminate, while exposed third parties can be the dominant bearers of the input-cost loss when the redirected R&D path leaves costly gaps in the input bundle they actually use.

4.4 Third-party coalition failure

Once third-party incidence is positive, the remaining question is whether exposed economies can use that gain to compensate any bilateral player that prefers coercion.

Proposition 3 (Surplus feasibility of a restraint compact)**.** *Let $\Delta W_i = W_i^{Commit} - W_i^{SPE}$ for $i \in \{U, C, S\}$.*

(i) Feasibility. *A commitment-plus-transfer compact is feasible under the common-numeraire incidence objects iff*

$$\Delta W_S \geq \max\{-\Delta W_U, 0\} + \max\{-\Delta W_C, 0\}.$$

(ii) Strict Pareto improvement. *If the inequality is strict, the compact admits transfers* $T_U, T_C \geq 0$ *with* $T_U + T_C \leq \Delta W_S$ *that leave each player a strictly positive share of the residual slack* $\Delta W_S - \max\{-\Delta W_U, 0\} - \max\{-\Delta W_C, 0\} > 0$*, hence strictly better off under the model's common-numeraire payoff comparison.*

The proof is in Appendix B.10.

Proposition 3 separates transfer feasibility from credible-restraint contracting. The inequality says that, in the model's common numeraire, exposed third parties' aggregate gain from restraint is large enough to compensate any bilateral player that loses from commitment. It does not imply that a market for restraint exists. Implementation requires an institution that can make payments contingent on verifiable US restraint, allocate the burden across exposed beneficiaries, compensate any losers, and bind US policy after payment.

To state the aggregation problem, index exposed third parties by $k = 1, \ldots, K$ and write their individual commitment wedges as

$$\Delta W_{S,k} \equiv W_{S,k}^{\text{Commit}} - W_{S,k}^{\text{SPE}}, \qquad \Delta W_S = \sum_{k=1}^{K} \Delta W_{S,k}.$$

Let the bilateral compensation requirement be

$$\overline{M} \equiv \max\{-\Delta W_U, 0\} + \max\{-\Delta W_C, 0\}$$

as in Proposition 3. Define gross gains, third-party losses, and the institutional funding threshold by

$$G \equiv \sum_{k=1}^{K} \max\{\Delta W_{S,k}, 0\}, \qquad L \equiv \sum_{k=1}^{K} \max\{-\Delta W_{S,k}, 0\}, \qquad \overline{M}^I \equiv \overline{M} + L.$$

Thus G is the gross willingness to pay of exposed third parties that gain from commitment, while L is the compensation needed to leave any exposed third party that loses from commitment weakly whole. Because $\Delta W_S = G - L$, the aggregate feasibility condition in

Proposition 3 is equivalent to $G \geq \overline{M}^I$. The corollary turns aggregate feasibility into an implementation threshold: gross winners must cover both bilateral compensation and losses borne by exposed-country losers.

Corollary 1 (Coalition institution)**.** *Suppose export-control restraint is non-excludable across exposed third parties and cannot be reserved for contributors. A coalition institution that can collect binding contingent assessments and bind the United States to restraint implements commitment for member set* $\mathcal{I} \subseteq \{k : \Delta W_{S,k} > 0\}$ *iff*

$$\sum_{k \in \mathcal{I}} \Delta W_{S,k} \geq \overline{M}^I.$$

Absent such an institution, if $G \geq \overline{M}^I$ *but* $\max_k \max\{\Delta W_{S,k}, 0\} < \overline{M}^I$*, no single exposed third party can profitably finance restraint, and no contribution is a Nash equilibrium in the refundable threshold-contribution game.*

The corollary is an implementation condition, not a theory of how the institution forms. Export-control restraint is non-excludable across exposed economies: once the United States restrains controls, nonpayers can benefit as well. A compact therefore requires more than positive aggregate surplus. It requires a credible-restraint contract that collects contingent contributions, handles compensation, and binds the United States after payment. In the calibration, the main obstacle is not necessarily inability to pay; it is the absence of such a contract. The proof sketch is in Appendix B.11.

All results above are stated under Setup B (fixed R&D budget); Section 5.1 discusses why the empirical Chinese R&D-to-GDP path makes the fixed-budget regime the relevant baseline, and Appendix B.12 reports the hybrid-cost model and segmented-regression robustness.

5 Counterfactual Incidence

This section turns the empirical R&D path into incidence numbers. We treat the patent-derived 2010–2021 allocation as the export-control-regime path and compare it with a credible-restraint benchmark: before China allocates R&D, the United States commits not to impose export controls. Under restraint, China keeps the same aggregate R&D budget but allocates it uniformly across buckets. Because observed patent dispersion also reflects demand, comparative advantage, industrial policy, and baseline productivity, the uniform-restraint benchmark is best read as an upper-bound measure of the directed-innovation channel.

Accounting scope. The counterfactual reports model-unit incidence, not consumption-equivalent welfare. The three player objects are the US strategic-payoff object, China's producer/innovation-payoff object, and exposed economies' downstream input-cost object. Transfer comparisons put these objects in a common model numeraire. For exposed economies, the cleaner scale divides the input-cost wedge by their own cumulative input-cost base: the +0.118 exposed-economy wedge equals 0.59% of that base. Neither denominator is a GDP-equivalent welfare measure. The calculation isolates downstream input-cost incidence from the induced Chinese productivity path; it does not measure full country-level welfare.

5.1 Calibration and incidence inputs

Specification: fixed R&D budget. The model results are stated under Setup B (a fixed sectoral R&D budget). Economically, this treats Chinese innovative capacity as a scarce policy and engineering resource over 2018–2021. Export controls change where that capacity is deployed, not the aggregate national R&D envelope. The Chinese aggregate R&D-to-GDP path supports this specification empirically: the cumulative 2018–2021 excess relative to the pre-2018 trend is near zero (−0.06 pp-years), while an unconstrained level-

response specification (Setup A) predicts roughly +0.18 pp-years. Appendix B.12 reports the hybrid-cost model nesting both limits and the segmented-regression details that motivate the fixed-budget choice.

Table 3 summarizes the calibrated parameters and their disciplining sources.

Table 3: Calibration inputs for the empirical-path counterfactual.

Parameter	Value	Discipline/source
φ (R&D elasticity)	0.5	Sqrt-concave; baseline
η (composite R&D-redirection elasticity)	0.54	Matched to the FDPR/MIC2025 Tier-1 patent response; A1-FOC value 2.39 is a high-η stress test
$\sigma_{\text{Tier 1}}$ (directly exposed chips)	0.15	Boehm et al. (2019) chip elasticity
$\sigma_{\text{Tier 2}}$ (chip-adjacent)	0.30	Baseline σ-grid; input-substitution evidence
$\sigma_{\text{Tier 3}}$ (raw materials)	0.55	Mid-range complementarity
$\sigma_{\text{bottleneck}}$ (broad)	0.30	Baseline coercion-block grid
$\sigma_{\text{commodity}}$	0.85	Baseline commodity bucket
θ_{2010} (baseline productivity)	0.30	Below all bucket regime thresholds
Period-2 static-payoff parameters	fixed values	Used to compute the fixed-θ US strategic-payoff term; Appendix B.1
Empirical R&D budget $B(t)$	annual path	Total patent-derived R&D effort across buckets; restraint reallocates this total uniformly
Exposure weights $\omega_{c,h}$	pre-period shares	2014–2017 China-origin import shares used to allocate W_S across countries and HS6 cells

Notes: η matches the FDPR Tier 1 dose coefficient ($\beta_1 = +0.774$ in Table 1). The Tier 1/2/3 σ values follow the pre-specified critical-input taxonomy, anchored at Boehm et al. (2019) for chips. Appendix B.1 gives the reduced-form regularity diagnostic; the directed-innovation/input-cost term is driven by the empirical θ path and exposure weights.

The table separates identification roles: σ and θ locate bottlenecks and activation, while η governs how strongly those bottlenecks redirect a fixed R&D budget.

Moment-matching calibration of η. We calibrate the composite redirection elasticity η to the FDPR/MIC2025 patent response in Table 1. Calibration using the full constrained optimization problem gives $\hat{\eta}_{\text{KKT}} = 0.54$ and reproduces the Tier-1 response exactly; the model-implied Tier-2 response is +0.198, close to the empirical +0.172. The estimate summarizes both forced substitution and policy-financed entry, since ChangXin and Yangtze account for much of the post-2020 Tier-1 patent surge. A modest η can match Tier-1 redirection because the full KKT includes amplification: R&D raises productivity, productivity raises effective coercion, and coercion raises self-supply returns. The simplified first-order-condition calibra-

tion, $\hat{\eta}_{\text{A1}} = 2.39$, is retained as a high-$\eta$ stress test without that feedback amplification, not as an alternative baseline of equal status. Appendix B.3 reports the feedback diagnostic, and Appendix Table 8 reports the realized Γ_s^* values on the calibrated support.

The calibration treats λ_s as effective coercion, bundling statutory enforcement e_s and upstream supplier compliance q_s. A direct compliance moment would compare post-control China-bound exports from US-aligned suppliers in controlled HS6 codes to a no-control prediction from adjacent products or pre-control trends. The FDPR tier response therefore identifies the combined redirection margin generated by statutory enforcement and supplier compliance.

Bias-adjusted calibration. Using the pre-trend-adjusted FDPR estimate $\beta_1^{\text{adj}} = +0.617$ gives $\hat{\eta}^{\text{adj}} \approx 0.46$ under the same baseline KKT calibration. A cluster-pair bootstrap gives a 95% interval of $[0.34, 0.70]$ around the headline $\hat{\eta}_{\text{KKT}} = 0.54$. This range sits well below the compact-surplus crossover at $\eta \approx 3.5$, so the main incidence result is not tied to the exact point estimate.

Normalizations. The total-model-surplus normalization divides each player wedge by cumulative model surplus,

$$\%_X = 100 \cdot \frac{\Delta W_X^{\text{Commit}-\text{SPE}}}{|W_{\text{total}}^{\text{Commit}}|_{\text{cum 2018-2021}}}, \qquad |W_{\text{total}}^{\text{Commit}}|_{\text{cum}} = 12.64.$$

For the downstream exposed economies, we also report the player-specific input-cost ratio

$$\%_S^{\text{input cost}} = 100 \cdot \frac{\Delta W_S}{|W_S^{\text{Commit}}|_{\text{cum 2018-2021}}} = 100 \cdot \frac{0.118}{20.16} = 0.59\%.$$

The first normalization is used for transfer comparisons; the second is the incidence scale for the third-party input-cost object.

Empirical θ path. We compute China's bucket-level productivity path 2010–2021 from observed annual patents:

$$z_b(t) = \frac{p_b(t)}{p_b(2010)} - 1, \qquad \theta_b(t) = \theta_{2010} \cdot (1 + z_b(t))^{\varphi},$$

where $p_b(t)$ is total Chinese-assignee patents in bucket b at year t.[10] In the broad-bucket counterfactual, this empirical path is the observed SPE allocation relative to the uniform commitment benchmark. The sharper FDPR Tier-1/Tier-2/Tier-3 exercise provides the calibration moment for the low-σ redirection mechanism.

5.2 Counterfactual experiment design

We compare two regimes along the empirical period 2010–2021. Operationally, the quantitative exercise does not ask the model to predict the observed patent reallocation from scratch. It treats the observed coercive path as the empirical SPE allocation and uses the model to price the incidence of replacing it with the clean fixed-budget commitment benchmark:

- **Observed coercive path (SPE shorthand):** The United States' coercion intensity follows the baseline coercion-response map $\lambda^*(\theta_b(t), \sigma_b)$; China's R&D allocation matches the empirical patent-derived $z_b^{\text{SPE}}(t)$, which is generally non-uniform relative to the commitment benchmark.

- **Commit:** The United States credibly commits to $\lambda \equiv 0$. China picks uniform $z = B(t)/N$ across buckets (Jensen Allocation Lemma optimum when $\Gamma = 1$ in all sectors), where $B(t) = \sum_b z_b^{\text{SPE}}(t)$ is the empirical R&D budget.

Two channels. The wedge $\Delta W_{\text{total}} \equiv W_{\text{total}}^{\text{Commit}} - W_{\text{total}}^{\text{SPE}}$ decomposes as:

[10] This treats raw patent counts as a proxy for R&D effort with unit elasticity ($p \propto z$), consistent with our use of $\varphi = 0.5$ calibrated to USPTO Tier-1 growth and with Bloom et al. (2013). Section 2.2 reports a forward-citation-weighted robustness for the FDPR dose-response.

$$\Delta W_{\text{total}} = \underbrace{\Delta W_{\text{static}}}_{\text{static-theta channel}} + \underbrace{\Delta W_{\text{directed}}}_{\text{directed-innovation channel}} .$$

The static-theta channel asks whether controls look worthwhile holding the technology frontier fixed. It holds θ fixed at the empirical (SPE) values across regimes and isolates the United States' policy-choice deadweight $(\gamma/2)\lambda^2$ minus China's forced-substitution premium gain. The directed-innovation channel asks whether anticipating controls changed that frontier by reallocating R&D before the coercion game is played. It replaces the empirical SPE θ path with the uniform-commitment θ path under the same surplus convention and captures the productivity loss from concentrating R&D when returns are diminishing.

5.3 Fixed-productivity benchmark and directed-innovation incidence

This is the quantitative delta from Kim (2026): when Chinese productivity is treated as fixed, the parent environment's export-control ranking survives; when China's R&D allocation is endogenous, the ranking reverses. Table 4 makes the fixed-productivity comparison and the incidence split explicit. Panel A separates the static policy channel from the directed-innovation channel; Panel B allocates the resulting commitment wedge across players.

The central quantitative result is a sign flip. Under fixed productivity, the only channel is the fixed-θ term, which is -0.006 and therefore marginally favors SPE. Endogenizing θ opens the directed-innovation channel: anticipated controls redirect period-1 R&D into restricted bottlenecks, generating a $+0.126$ allocation loss under SPE. That channel dominates the static term, so the ranking reverses to a $+0.120$ commitment gain in this paper.

The incidence of the reversal is concentrated outside the dyad. The downstream exposed-economy wedge ($+0.118$) is nearly the entire aggregate gain from commitment. At the baseline KKT calibration $\hat{\eta} = 0.54$, China weakly prefers commitment ($\Delta W_C = +0.015$) and the United States weakly prefers coercion ($\Delta W_U = -0.013$), so the combined US–China gain

Table 4: Credible restraint minus export controls: fixed-productivity benchmark and directed-innovation incidence, cumulative 2018–2021.

Item	Credible restraint − export controls	Regime favored
Panel A. Fixed-productivity comparison		
Fixed-θ benchmark	−0.006	Export controls
Directed-innovation channel (endogenous θ)	+0.126	Credible restraint
Net total after endogenizing θ	**+0.120**	**Credible restraint**
Panel B. Player incidence		
United States (W_U)	−0.013	Export controls
China (W_C)	+0.015	Credible restraint
Downstream exposed economies (W_S)	**+0.118**	**Credible restraint**

Notes: Entries are credible-restraint-minus-export-controls in model units; positive values favor restraint. These are not GDP-equivalent or consumption-equivalent welfare changes. W_S is the downstream input-cost object, and the +0.118 exposed-economy wedge equals 0.59% of exposed economies' cumulative input-cost base.

is $\Delta W_U + \Delta W_C = +0.002$. This near-cancellation is a calibration result. At the simplified first-order-condition value $\hat{\eta}_{\text{A1}} = 2.39$, the China wedge is negative ($\Delta W_C = -0.015$), so the combined US–China gain is −0.028 and both bilateral players prefer coercion. The stable result is third-party input-cost incidence; bilateral neutrality is a feature of the baseline low-η calibration.

Country and HS6 decomposition. The aggregate table establishes that the incidence is mostly third-party; the next tables ask which exposed economies and products carry that term. We allocate the downstream W_S wedge using pre-period China-origin import exposure, so the decomposition is an accounting allocation rather than a country-specific causal estimate. Table 5 allocates each bucket's 2018–2021 W_S wedge across Korea, Taiwan, Japan, the Netherlands, Vietnam, and Mexico using their 2014–2017 average imports from China in the paper's HS6 sample, normalizing weights within the bottleneck and commodity buckets separately.

The incidence is concentrated, not generic. These economies bear incidence not because they choose the coercive policy, but because their production networks buy Chinese interme-

Table 5: Exposure allocation of the third-party input-cost wedge by country, 2018–2021.

Country	China imports (M USD)	Bottleneck	Commodity	ΔW_S	Share (%)
KOR	14169.2	0.074	-0.016	0.058	49.0
TWN	10688.6	0.058	-0.009	0.049	41.6
MEX	10864.3	0.040	-0.034	0.006	5.0
VNM	3529.4	0.013	-0.011	0.003	2.5
JPN	13709.0	0.049	-0.046	0.003	2.3
NLD	1654.7	0.006	-0.006	-0.001	-0.5

Notes: China imports are 2014–2017 annual averages in the paper's HS6 sample. The bottleneck and commodity columns allocate the model-implied bucket wedges, +0.239 and −0.121, by each country-HS6 cell's pre-period China-origin import share within that bucket. The allocated country wedges sum to the model's aggregate third-party wedge, +0.118. Source: authors' calculations from BACI and Table 4.

diates in the sectors where redirected innovation changes input prices. Korea and Taiwan account for about 91% of the allocated third-party wedge because their China-origin exposure is concentrated in bottleneck electronics. Japan has large China imports but little net exposure after commodity offsets. Mexico and Vietnam are positive but small, and the Netherlands is effectively zero in this HS6 sample. At the product level, the positive wedge is led by memory chips (854232), other integrated circuits (854239), photovoltaic cells (854140), printed circuits (853400), and processors/controllers (854231); radio/TV parts (852990) provide the main commodity offset. The result is therefore an East Asian electronics-supply-chain incidence result, not a broad equal-weight Jensen result.

Table 6 shows how much of the input-cost incidence survives when the largest exposure cells are removed. Dropping Korea or Taiwan alone leaves about half of the baseline wedge. Dropping both lowers the remaining wedge to +0.011, and dropping the integrated-circuit family reverses the sign. The leave-out exercises show concentration rather than broad-based incidence. The third-party wedge is carried mainly by Korea/Taiwan electronics exposure, especially the integrated-circuit family, while commodity exposure provides the main offset.

Accounting feasibility of a restraint compact. Because China weakly prefers commitment voluntarily at the calibrated η, the compensation requirement is one-sided: only the United States must be compensated to accept the commitment regime. The aggregate commitment surplus is +0.120; exposed economies' aggregate downstream input-cost gain from

Table 6: Concentration diagnostics for the third-party input-cost wedge.

Leave-out exercise	Omitted contribution	Remaining ΔW_S	Baseline share (%)
Drop Korea	0.058	0.060	51.0
Drop Taiwan	0.049	0.069	58.4
Drop Korea + Taiwan	0.107	0.011	9.3
Drop memory chips (854232)	0.056	0.062	52.6
Drop integrated circuits (854231/854232/854239)	0.129	-0.010	-8.6
Drop PV cells (854140)	0.037	0.081	68.5
Drop main commodity offset (852990)	-0.060	0.178	150.7
Drop all commodity offsets	-0.121	0.239	202.1

Notes: Rows remove the listed exposure cells from the allocated ΔW_S decomposition without renormalizing remaining cells. The baseline is $\Delta W_S = +0.118$. A negative omitted contribution means the excluded cell is an offsetting commodity exposure, so the remaining wedge is larger. Source: authors' calculations from BACI and Table 4.

commitment (+0.118) exceeds the United States' required minimum transfer $T_U \geq 0.013$, leaving transfer slack of 0.105 after compensating the United States. An aggregate transfer package financed by governments of exposed economies therefore supports the compact under the model's transfer criterion at the calibrated baseline. At the simplified first-order-condition value $\hat{\eta}_{\mathrm{A1}} = 2.39$, by contrast, China's forced-substitution premium is large enough that China also strictly prefers SPE, $\Delta W_C^{\mathrm{A1}} = -0.015$, and the compact would require two-sided transfers $T_U \geq 0.013, T_C \geq 0.015$.

Table 7 applies the Proposition 3 threshold to the country allocation. The baseline gains are concentrated enough that Korea or Taiwan alone can cover the required model-unit compensation, and the Korea–Taiwan coalition generates large slack. The compact problem is therefore contractual. A workable arrangement would have to collect contingent payments, verify US restraint, allocate the burden among beneficiaries, and bind US export-control policy after payment.

5.4 Robustness and scope

Two diagnostics summarize sensitivity to η and to the empirical θ path. First, varying η changes the static channel and the combined US–China gain, but the directed-innovation

Table 7: Accounting feasibility of a restraint compact.

Coalition	Gross positive gain	Accounting threshold $\overline{M}^I$	Accounting slack	Feasibility
Korea alone	0.058	0.014	0.044	Yes
Taiwan alone	0.049	0.014	0.035	Yes
Mexico alone	0.006	0.014	-0.008	No
Japan alone	0.003	0.014	-0.011	No
Vietnam alone	0.003	0.014	-0.011	No
Netherlands alone	0.000	0.014	-0.014	No
Korea + Taiwan	0.107	0.014	0.093	Yes
East Asia (KOR+TWN+JPN)	0.110	0.014	0.096	Yes
All six exposed economies	0.119	0.014	0.105	Yes

Notes: The threshold is $\overline{M}^I = \overline{M} + L = 0.014$, where $\overline{M} = \max\{-\Delta W_U, 0\} + \max\{-\Delta W_C, 0\} = 0.013$ and $L = 0.001$ compensates third-party losers in the country decomposition. Gross positive gains sum $\max\{\Delta W_{S,k}, 0\}$ over coalition members; negative-gain countries enter through L rather than through the contribution column. Clearing the threshold establishes surplus feasibility only and does not solve credible restraint, contribution aggregation, or implementation. Source: authors' calculations from Tables 4 and 5.

channel is fixed once the SPE and restraint θ paths are fixed. The net comparison continues to favor restraint through the A1-FOC stress-test value $\eta = 2.39$ and turns negative only near $\eta \simeq 3.5$ (Appendix A.2). Second, scaling down the empirical θ-path shows how much of the patent-derived directed-innovation signal is needed for the main results. Using unrounded model wedges, the net comparison favors restraint as long as at least 4.4% of the baseline directed-innovation channel remains; the one-sided compact threshold holds as long as at least 11.2% of the baseline third-party input-cost wedge remains (Appendix C).

The exercise remains partial equilibrium. It prices the induced Chinese productivity path through exposed economies' input-cost exposure while leaving labor reallocation, terms of trade, and aggregate-demand feedback outside the main model. Appendix E gives an illustrative GE-buffer and dollar-scale envelope.

6 Conclusion

Export controls can appear nearly neutral for the US–China dyad while shifting costs onto third-country supply chains. The reason is dynamic. Anticipating future controls, China

redirects R&D toward restricted bottleneck self-supply. That redirection changes the productivity path of Chinese intermediates and raises input-cost-equivalent incidence for economies that source them. In the calibrated counterfactual, the fixed-productivity channel marginally favors controls, but the directed-innovation channel reverses the sign; most of the incidence falls on exposed third parties.

The evidence pins down the composite redirection margin. Chinese patenting rises most in FDPR-exposed chip codes, concentrated in memory entrants that also received MIC2025/Big Fund financing, so η captures both forced substitution and policy-financed entry. The BACI/Compustat exercise provides a directional check on the predicted input-cost effect: exposed East Asian firms facing higher China-origin unit-value shocks report higher COGS shares, although the association attenuates with industry-year fixed effects.

The policy implication is a missing institution for credible restraint. Exposed economies' input-cost gain from commitment exceeds the United States' loss at the baseline calibration, but implementing the compact requires an institution that aggregates contingent contributions and binds the United States to restraint. Export controls therefore shift part of the burden of rival innovation outside the bilateral conflict, even when the combined US–China gain appears small.

References

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller, "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program," *Journal of the American Statistical Association*, 2010, *105* (490), 493–505.

Acemoglu, Daron, "Directed Technical Change," *The Review of Economic Studies*, 2002, *69* (4), 781–809.

_ , Philippe Aghion, Leonardo Bursztyn, and David Hémous, "The Environment and Directed Technical Change," *American Economic Review*, 2012, *102* (1), 131–166.

Akcigit, Ufuk, Sina T. Ates, and Giammario Impullitti, "Innovation and Trade Policy in a Globalized World," Working Paper 24543, National Bureau of Economic Research 2018.

Atalay, Enghin, "How Important Are Sectoral Shocks?," *American Economic Journal: Macroeconomics*, 2017, *9* (4), 254–280.

Atkeson, Andrew and Ariel T. Burstein, "Innovation, Firm Dynamics, and International Trade," *Journal of Political Economy*, 2010, *118* (3), 433–484.

Bagwell, Kyle and Robert W. Staiger, "An Economic Theory of GATT," *American Economic Review*, 1999, *89* (1), 215–248.

Bartik, Timothy J., *Who Benefits from State and Local Economic Development Policies?*, Kalamazoo, MI: W.E. Upjohn Institute for Employment Research, 1991.

Becko, John Sturm, "How Should Sanctions Account for Bystander Countries?," *AEA Papers and Proceedings*, 2023, *113*, 39–42.

_ **and Daniel G. O'Connor**, "Strategic (Dis)Integration," 2025. Working paper.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar, "Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tohoku Earthquake," *Review of Economics and Statistics*, 2019, *101* (1), 60–75.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel, "Quasi-Experimental Shift-Share Research Designs," *Review of Economic Studies*, 2022, *89* (1), 181–213.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 2015, *82* (1), 1–44.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller, "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 2008, *90* (3), 414–427.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger, "A Theory of Economic Coercion and Fragmentation," *BIS Working Paper No. 1224*, 2024.

_ , _ , and _ , "A Framework for Geoeconomics," *Econometrica*, 2026, *94* (1), 105–136.

Farrell, Henry and Abraham L. Newman, "Weaponized Interdependence: How Global Economic Networks Shape State Coercion," *International Security*, 2019, *44* (1), 42–79.

Gaulier, Guillaume and Soledad Zignago, "BACI: International Trade Database at the Product-Level. The 1994–2007 Version," Working Paper 2010-23, CEPII 2010.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift, "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 2020, *110* (8), 2586–2624.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg, "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," *NBER Working Paper No. 8498*, 2001.

Horn, Henrik, Giovanni Maggi, and Robert W. Staiger, "Trade Agreements as Endogenously Incomplete Contracts," *American Economic Review*, 2010, *100* (1), 394–419.

Kim, Hyungjin, "The Coercive Turn Toward Swing States: Vertical Bottlenecks and the Paradox of Competitiveness," 2026. Korea Development Institute. Revise and resubmit, IMF Economic Review.

Liu, Xueyue, Yu Liu, Alexey Makarin, and Jaya Wen, "Export Controls and Innovation in Sanctioned Countries," Working Paper 25-004, Harvard Business School 2025.

Lybbert, Travis J. and Nikolas J. Zolas, "Getting Patents and Economic Data to Speak to Each Other: An 'Algorithmic Links with Probabilities' Approach for Joint Analyses of Patenting and Economic Activity," *Research Policy*, 2014, *43* (3), 530–542.

Maggi, Giovanni and Andrés Rodríguez-Clare, "The Value of Trade Agreements in the Presence of Political Pressures," *Journal of Political Economy*, 1998, *106* (3), 574–601.

National Center for Science and Engineering Statistics, "Cross-National Comparisons of R&D Performance: China R&D Indicators 2010–2023," NSF / NCSES (data sourced from China National Bureau of Statistics) 2024. https://ncses.nsf.gov/pubs/nsb20246/figure/RD-10.

Rambachan, Ashesh and Jonathan Roth, "A More Credible Approach to Parallel Trends," *Review of Economic Studies*, 2023, *90* (5), 2555–2591.

Varian, Hal R., *Microeconomic Analysis*, 3 ed., W. W. Norton & Company, 1992.

Webb, Matthew D., "Reworking Wild Bootstrap-Based Inference for Clustered Errors," Working Paper 1315, Queen's Economics Department 2014. Subsequently published in Canadian Journal of Economics 56(3): 839–858, 2023.

A Counterfactual Diagnostics

Table 8: Calibrated-support Γ_s^* diagnostic.

Tier	σ_s	λ_s^*	$\Lambda_{\theta,s}$	Π_s	$\Gamma_s^*/\bar{\Gamma}$	Δz_s
Tier 1	0.15	0.085	0.160	1.152	1.214	+0.954
Tier 2	0.30	0.031	0.117	1.028	0.915	-0.329
Tier 3	0.55	0.016	0.048	1.006	0.830	-0.625

Notes: Boehm-anchor full-KKT calibration. The implied threshold ratio is $\bar{\Gamma} = \mu^{\text{SPE}}/\mu^{\text{Commit}} = 1.230$. The diagnostic evaluates Γ_s^* at the tier-specific θ_s^{SPE} implied by the full-KKT allocation, not at a common productivity level. The last column is $\Delta z_s = z_s^{\text{SPE}} - z_s^{\text{Commit}}$, using the full-KKT SPE allocation relative to the uniform commitment allocation.

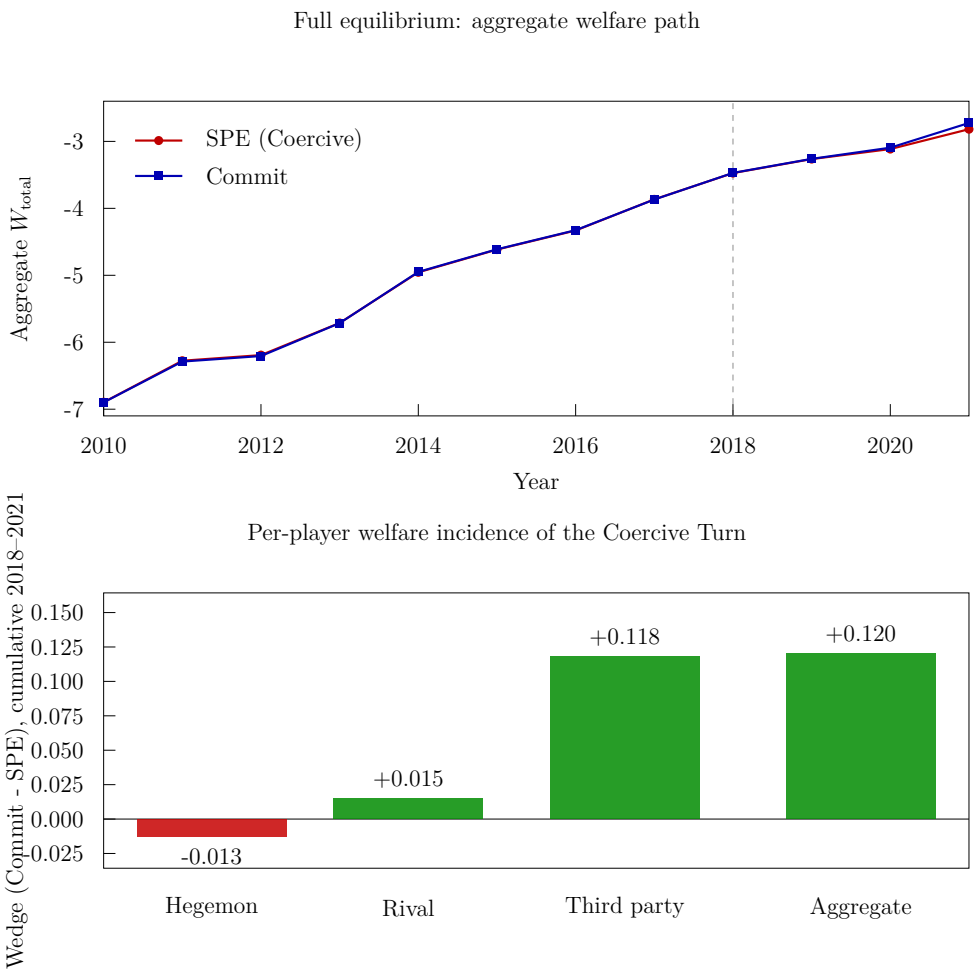

Figure 1: Calibrated empirical-path counterfactual.

Notes: Top: aggregate welfare W_{total} under SPE versus Commit, 2010–2021; shaded region marks years in which Commit exceeds SPE in aggregate. Bottom: cumulative 2018–2021 per-player wedge. Third-party gains from commitment dwarf the small Hegemon SPE preference, and the bilateral US–China wedge is approximately zero.

A.1 Year-by-year dynamics

The wedge builds up over time. Each entry below is the annual flow wedge $\Delta W_t^{Commit-SPE}$ in model units; the four entries sum to the cumulative 2018–2021 figure in Table 4:

- 2018 (Coercive Turn onset): +0.003.
- 2019: +0.004.
- 2020 (FDPR/MIC2025 moment): +0.021.

- 2021: +0.093.

The four flows sum to +0.121 ≈ +0.120 cumulative. The pattern is consistent with the dynamic mechanism: the incidence builds as the Rival's R&D path diverges between SPE and Commit.

A.2 Sensitivity to η

The composite redirection elasticity η is the key calibrated parameter for the static channel and the bilateral wedge, not for the dominant third-party directed-innovation channel. Figure 2 shows the wedge decomposition for $\eta \in \{0.54, 1.0, 2.0, 2.39, 3.0, 4.0\}$. The lowest point is the baseline KKT value, and the upper end of the grid spans the simplified-FOC stress-test value and beyond.

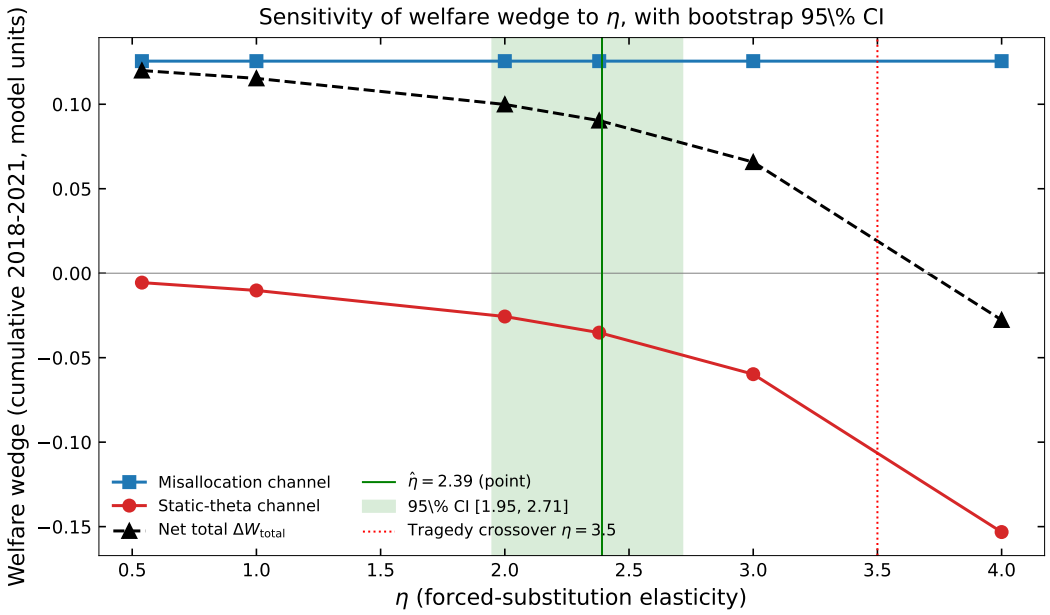

Figure 2: Sensitivity of model wedges to η.

Notes: Directed-innovation channel (blue squares) is invariant in η at +0.126 since it depends only on θ-path differences and φ. Static channel (red circles) scales with η. Net total (black triangles, dashed) is +0.120 at $\eta = 0.54$ (baseline KKT), +0.115 at $\eta = 1.0$, +0.090 at $\eta = 2.39$ (simplified-FOC stress test), and crosses zero around $\eta = 3.5$ (red dotted line). The compact-surplus condition holds throughout the baseline KKT defensible region.

The pure directed-innovation channel is invariant to η (+0.126), because it depends only

on the θ-path difference and φ; the third-party wedge is likewise invariant at +0.118. The Rival's wedge changes sign between $\eta = 1.0$ and $\eta = 2.0$, so bilateral neutrality holds near the headline full-KKT calibration and fails at higher η. The net total remains positive at +0.090 under the A1-FOC stress-test value $\eta = 2.39$ and crosses zero only near $\eta \approx 3.5$.

B Proofs

This appendix contains the proofs of Propositions 1, 2, and 3, Corollary 1, Remark 3, the Jensen Allocation Lemma, and the third-party input-cost incidence lemma. Notation throughout follows Section 3. The main text uses the exact KKT premium Γ_s; the Π-closed form appears only as a fixed-wedge benchmark.

B.1 The period-2 coercion block

This appendix states the payoff implementation used for the welfare levels in Section 5 and separates it from the regularities used by the period-1 theory. The redirection logic requires only P1–P2. The calibrated welfare levels additionally use the explicit extraction primitive stated below.

The structural block

Fix a sector s with rival productivity θ_s and bottleneck $\sigma_s \in (0, 1]$.

Third party. The third party chooses an alignment $x_s \in [0, 1]$, the degree of substitution toward US-aligned alternatives. Its laissez-faire input cost is the CES aggregate of the rival-supplied input (price w_C) and the US-aligned alternative (price w_A),

$$C(x_s; w_C, w_A) = \Big[a(x_s)\, w_C^{1-\sigma_R} + \big(1 - a(x_s)\big)\, w_A^{1-\sigma_R} \Big]^{1/(1-\sigma_R)}, \qquad \sigma_R \in (0, 2), \tag{16}$$

where $a(x_s) \in [0,1]$ is a continuously differentiable share weight on the rival input, decreasing in x_s (greater alignment shifts the bundle toward the US-aligned alternative). The Rival's input price falls in its productivity, $w_C(\theta_s) = \bar{w}/\theta_s$, while w_A is fixed. The third party's per-sector welfare is the negative of its input cost,

$$V_S^s = -\,C\big(x_s; w_C(\theta_s), w_A\big). \tag{17}$$

Hegemon. The Hegemon's payoff at sector-s alignment x_s and rival productivity θ_s is

$$J(x_s; \theta_s) = V_{US}(x_s) - C\big(x_s; w_C(\theta_s), w_A\big), \tag{18}$$

with V_{US} the strategic value of alignment (Section 3.2), increasing and concave in x_s. At no coercion ($\lambda = 0$), the induced alignment in this reduced-form payoff block is

$$x_G^s(\theta_s) = \arg\max_x \; J(x; \theta_s), \tag{19}$$

so $U^s(\theta_s, 0) = J(x_G^s; \theta_s)$. The Hegemon chooses coercion intensity $\lambda_s \in [0,1]$, anticipating the third party's induced alignment $x_s(\lambda_s)$, to maximize the per-sector payoff

$$U^s(\theta_s, \lambda_s) = J(x_G^s; \theta_s) \; + \; \underbrace{\Pi_E(\lambda_s; \theta_s, \sigma_s)}_{\text{coercion-extraction value}} \; - \tfrac{\gamma}{2}\,\lambda_s^2, \tag{20}$$

where Π_E is the incremental surplus the Hegemon appropriates by inducing alignment $x_s(\lambda_s) > x_G^s$ away from the Rival, and $(\gamma/2)\lambda_s^2$ is a quadratic deadweight cost of coercion. The extraction value satisfies $\Pi_E(0; \cdot) = 0$, so $U^s(\theta_s, 0) = J(x_G^s; \theta_s)$ as stated. At the optimum, the calibrated extraction primitive satisfies $\Pi_E(\lambda_s^*) = \gamma(\lambda_s^*)^2 + J(x_s^*; \theta_s) - J(x_G^s; \theta_s)$, which gives the optimized payoff identity below. We use this identity as part of the calibrated payoff implementation rather than deriving it from P1–P2 alone; P1–P2 discipline the redirection logic, while this value identity is needed for the welfare decomposition. The

interior first-order condition and the optimized-λ envelope, used in the welfare decomposition of Appendix B.9, are

$$\underbrace{x_s^*\,\theta_s - J(x_s^*;\theta_s)}_{\equiv\; M(\theta_s,\sigma_s)} \;=\; \gamma\,\lambda_s^*, \qquad U^s(\theta_s,\lambda_s^*) = \tfrac{\gamma}{2}\,(\lambda_s^*)^2 + J(x_s^*;\theta_s), \tag{21}$$

where $x_s^* = x_s(\lambda_s^*)$ and $M(\theta_s,\sigma_s)$ is the equilibrium marginal extraction value. The induced equilibrium intensity defines the coercion map

$$\lambda_s^* = \Lambda(\theta_s,\sigma_s) = \arg\max_{\lambda\in[0,1]} U^s(\theta_s,\lambda). \tag{22}$$

P1–P2 from primitive properties

The period-1 theory uses only the two regularities below. We show they follow from two properties of the equilibrium marginal extraction value $M(\theta,\sigma)$ in (21), both of which are the economic content of the target-value-feedback mechanism:

(M1) Target-value monotonicity. $M(\theta,\sigma)$ is differentiable and strictly increasing in θ on the active-coercion region: a more productive rival yields more appropriable surplus.

(M2) Activation threshold. There is $\bar\theta(\sigma)$ such that the Hegemon's optimal intensity is zero for $\theta \le \bar\theta(\sigma)$ and positive for $\theta > \bar\theta(\sigma)$.

Lemma 3 (Coercion-block regularities)**.** *Under (20)–(22) with $\gamma > 0$, (M1), and (M2), the equilibrium map $\Lambda(\theta,\sigma)$ satisfies P1 ($\Lambda_\theta > 0$ in the active region) and P2 (threshold $\bar\theta(\sigma)$).*

Proof. P2 is immediate from (M2): $\Lambda(\theta,\sigma) = 0$ below $\bar\theta(\sigma)$ and positive above it. For P1, the second-order condition for the Hegemon's interior maximum in (20) is $U^s_{\lambda\lambda} < 0$, which holds because the deadweight term contributes $-\gamma$ and the extraction value is concave at the optimum. On the active region the interior FOC (21) reads $\gamma\,\Lambda(\theta,\sigma) = M(\theta,\sigma)$. Differentiating in θ,

$$\gamma\,\Lambda_\theta(\theta,\sigma) = M_\theta(\theta,\sigma) > 0 \quad \text{by (M1)},$$

so $\Lambda_\theta(\theta,\sigma) = M_\theta(\theta,\sigma)/\gamma > 0$. □

Intrinsic regularities vs. explicit primitive. Lemma 3 shows that, once the payoff implementation delivers (M1)–(M2), the induced map satisfies P1–P2. Thus the period-1 redirection logic uses only the comparative-static regularities, not the closed-form extraction primitive. The explicit extraction primitive Π_E gives the calibrated map Λ and the period-2 parameters (γ, V_0, β) that pin down the welfare *levels* in Section 5 and the realized Γ_s^* in Table 8. Substituting the closed-form Π_E for $M(\theta,\sigma)$ closes the block in closed form; the lemma is unchanged.

Solver-free reduced-form robustness

As a further check that P1–P2 do not depend on any particular period-2 solver, consider the closed-form schedule

$$\Lambda^{\mathrm{RF}}(\theta,\sigma) = \begin{cases} 0, & \theta \le \bar\theta(\sigma), \\ \bar\lambda(\sigma)\big(1 - \exp\{-a(\theta - \bar\theta(\sigma))\}\big), & \theta > \bar\theta(\sigma), \end{cases}$$

with $a > 0$ and $\bar\lambda(\sigma) \in (0,1]$. It satisfies P2 by construction and P1 in the coercive region because

$$\Lambda^{\mathrm{RF}}_\theta(\theta,\sigma) = a\,\bar\lambda(\sigma)\,\exp\{-a(\theta - \bar\theta(\sigma))\} > 0 \quad \text{for } \theta > \bar\theta(\sigma).$$

Substituting Λ^{RF} into the Rival's period-1 problem yields the same KKT premium object,

$$\Gamma_s^{\mathrm{RF}} = 1 + \Lambda_s^{\mathrm{RF}}\big(\sigma_s^{-\eta} - 1\big) + \theta_s\big(\sigma_s^{-\eta} - 1\big)\,\Lambda_\theta^{\mathrm{RF}}(\theta_s,\sigma_s),$$

which exceeds one for any coerced bottleneck sector ($\sigma_s < 1$). The schedule reports signs and regularities, not welfare magnitudes: P1–P2 are sufficient for the period-1 redirection logic but do not pin down the level of λ_s^* or the player payoffs, which require the calibrated structural block above.

B.2 Proof of the CES sufficient condition (payoff concavity on the CES family)

Let $C(w) \equiv C(x_G; w, w_A)$ denote the third party's CES cost evaluated at the no-coercion alignment x_G, viewed as a function of $w_C = w$ alone (with w_A and x_G fixed). From (15),

$$C(w) = \left[a w^{1-\sigma_R} + (1-a) w_A^{1-\sigma_R}\right]^{1/(1-\sigma_R)},$$

where $a \equiv a(x_G) \in [0,1]$ is constant in this argument. We show that $J(\theta) = V_{US} - C(w_C(\theta))$ with $w_C(\theta) = \bar{w}/\theta$ is strictly increasing and weakly concave in θ on $\theta > 0$.

Step 1 (monotonicity). By Shephard's lemma applied to the CES form, $C'(w) > 0$ whenever $a > 0$ and $w > 0$ (i.e., the cost is strictly increasing in any input price with positive cost share). Then

$$\frac{\partial J}{\partial \theta} = -C'(w_C(\theta)) \cdot w_C'(\theta) = C'(w) \cdot \frac{\bar{w}}{\theta^2} > 0.$$

Hence J is strictly increasing in θ on $\theta > 0$. □

Step 2 (concavity). Let $w = w_C(\theta) = \bar{w}/\theta$. The second derivative is:

$$\frac{\partial^2 J}{\partial \theta^2} = -\frac{\partial}{\partial \theta}\left[C'(w) \cdot w_C'(\theta)\right] = -C''(w) \cdot (w_C'(\theta))^2 - C'(w) \cdot w_C''(\theta).$$

Substituting $w_C'(\theta) = -\bar{w}/\theta^2 = -w/\theta$ and $w_C''(\theta) = 2\bar{w}/\theta^3 = 2w/\theta^2$:

$$\frac{\partial^2 J}{\partial \theta^2} = -\frac{1}{\theta^2}\Big[C''(w) \cdot w^2 + 2C'(w) \cdot w\Big] = -\frac{w}{\theta^2}\left[C''(w)w + 2C'(w)\right]. \tag{23}$$

For J to be weakly concave in θ, we need $C''(w)w + 2C'(w) \geq 0$, equivalently the cost-gradient own-price elasticity satisfies

$$\eta_{C'}(w) \equiv -\frac{C''(w) \cdot w}{C'(w)} \leq 2, \tag{24}$$

where $\eta_{C'}$ denotes the cost-gradient elasticity (distinct from the composite redirection elasticity η in (7); the subscript C' identifies the symbol with the cost-gradient $C'(w)$).

Step 3 (CES delivers $\eta_{C'} \leq \sigma_R$, with $\sigma_R \in (0,2)$). For the CES form (15), Shephard's lemma gives the conditional input demand $x_C(w) = C'(w) = a \cdot w^{-\sigma_R} \cdot C(w)^{\sigma_R}$ (standard derivation, e.g., Varian, 1992 §6.1). The own-price elasticity of conditional demand x_C is

$$\frac{\partial \log x_C}{\partial \log w} = -\sigma_R + \sigma_R \cdot s_C(w) = -\sigma_R\big(1 - s_C(w)\big),$$

where $s_C(w) = aw^{1-\sigma_R}/[aw^{1-\sigma_R} + (1-a)w_A^{1-\sigma_R}] \in [0,1]$ is the cost share of input C. Since $C''(w) = \partial x_C/\partial w$, we have

$$\eta_{C'}(w) = -\frac{C''(w) \cdot w}{C'(w)} = -\frac{\partial \log x_C}{\partial \log w} = \sigma_R \cdot (1 - s_C(w)) \leq \sigma_R.$$

For $\sigma_R \in (0,2)$, we have $\eta_{C'} \leq \sigma_R < 2$, so condition (24) is satisfied with strict inequality. Substituting into (23):

$$\frac{\partial^2 J}{\partial \theta^2} = -\frac{w}{\theta^2}\big[C''(w)w + 2C'(w)\big] = -\frac{w}{\theta^2} \cdot C'(w) \cdot \big[2 - \eta_{C'}(w)\big] \leq 0,$$

with strict inequality whenever $\sigma_R < 2$ and $C'(w) > 0$. □

Why the cutoff is exactly $\sigma_R < 2$. The chain rule applied to $w_C(\theta) = \bar{w}/\theta$ generates a $w_C''(\theta) = 2\bar{w}/\theta^3 > 0$ term that pushes J *toward* convexity in θ (since the convex transformation of a θ-decreasing w_C tends to invert curvature). The cost concavity in w (which holds for all CES with $\sigma_R > 0$) then has to dominate this convexification. The exact threshold is $\eta_{C'} \leq 2$; CES delivers $\eta_{C'} \leq \sigma_R$, so $\sigma_R < 2$ is the analytical cutoff. Empirically, the cost-side CES literature anchors $\sigma_R \in [0.1, 1.0]$ for upstream-input cost functions (Boehm et al., 2019; Atalay, 2017), so Assumption 2 holds with substantial slack.

B.3 Small-feedback benchmark diagnostic

The fixed-wedge benchmark in (13) drops the feedback term in the exact KKT premium. A useful diagnostic is

$$R(\sigma, \eta, \theta) \equiv \frac{\theta \cdot |\Pi_s'(\lambda_s^*) \cdot \partial \lambda_s^* / \partial \theta_s|}{\Pi_s(\lambda_s^*, \sigma_s)} = \frac{\theta \cdot (\sigma^{-\eta} - 1) \, |\partial \lambda_s^* / \partial \theta_s|}{1 + \lambda_s^* \, (\sigma^{-\eta} - 1)}.$$

The small-feedback benchmark holds locally when $R(\sigma, \eta, \theta) < 1$. On the working region $\sigma \in [0.20, 0.85]$ and $\theta \in [0.40, 0.90]$, numerical evaluation gives $\max R = 0.15$ at $\eta = 0.5$ and $\max R = 0.46$ at $\eta = 1.0$, so the benchmark is accurate near the full-KKT calibration. At the closed-form A1-FOC value $\hat{\eta}_{\text{A1}} = 2.39$, the same diagnostic reaches $\max R = 3.55$. That value lies outside the region where the closed-form approximation is internally valid, so it should be read as a high-η stress test rather than as an elasticity estimate. This is why the calibrated exercise solves the full constrained KKT problem and reports realized Γ_s^* values.

B.4 Proof of Proposition 1 (Exact Innovation Redirection)

The Rival's period-1 problem under Setup B is

$$\max_{z \geq 0} \sum_s \theta_s(z_s) \cdot \Pi_s(\lambda_s^*(\theta_s(z_s), \sigma_s), \sigma_s) \quad \text{s.t.} \sum_s z_s = B.$$

For each interior sector, differentiating the sectoral objective gives

$$\frac{\partial}{\partial z_s}[\theta_s(z_s)\Pi_s(\lambda_s^*, \sigma_s)] = \theta_s'(z_s) \left[\Pi_s + \theta_s(z_s)\Pi_{\lambda,s} \frac{\partial \lambda_s^*}{\partial \theta_s}\right].$$

The bracketed object is exactly $\Gamma_s(z_s; \eta)$ in (11); hence the KKT condition is $\theta_s'(z_s^{\text{SPE}})\Gamma_s^* = \mu^{\text{SPE}}$ for interior SPE sectors. Under commitment, $\lambda \equiv 0$, so $\Pi_s = 1$, $\Gamma_s = 1$, and $\theta_s'(z_s^{\text{Commit}}) = \mu^{\text{Commit}}$.

Because $\theta_s'(z_s)$ is strictly decreasing under multiplicative concavity,

$$z_s^{\text{SPE}} > z_s^{\text{Commit}} \iff \theta_s'(z_s^{\text{SPE}}) < \theta_s'(z_s^{\text{Commit}}) \iff \frac{\mu^{\text{SPE}}}{\Gamma_s^*} < \mu^{\text{Commit}} \iff \Gamma_s^* > \frac{\mu^{\text{SPE}}}{\mu^{\text{Commit}}}.$$

Defining $\bar{\Gamma} \equiv \mu^{\text{SPE}}/\mu^{\text{Commit}}$ gives the sector-by-sector redirection result. If the realized Γ_s^* is continuous and *decreasing* single-crossing in σ_s on the calibrated support and crosses $\bar{\Gamma}$, there is a unique $\bar{\sigma}_z$ satisfying $\Gamma^*(\bar{\sigma}_z) = \bar{\Gamma}$; sectors below that threshold gain R&D under SPE and sectors above it lose. If the premium does not cross $\bar{\Gamma}$, all sectors lie on the same side of the redirection condition. The Π-based threshold in (14) is obtained as the fixed-wedge/small-feedback approximation with Γ_s replaced by Π_s. □

B.5 Derivation for Remark 3 (threshold flip)

Consider the threshold-flip case, in which the lowest-σ (most bottlenecked) sectors are nonetheless below the activation threshold at the realized SPE productivity, $\theta_s(z_s^{\text{SPE}}) < \bar{\theta}(\sigma_s)$. By P2, $\lambda_s^* = 0$ and $\Gamma_s^* = 1$ in those sectors, so by Proposition 1 they receive no redirection despite their high bottleneck intensity: the single-crossing $\bar{\sigma}_z$ partition of the proposition's corollary, which would place the lowest-σ sectors among the gainers, does not apply once a sector falls below its threshold.

Redirection instead loads on the above-threshold, intermediate-σ sectors in the coercive region. As θ_s^0 rises and progressively lower-σ sectors cross $\bar{\theta}(\sigma_s)$ into the coercive region, the redirection target shifts toward those sectors, flipping the empirical sign-by-σ pattern. □

B.6 Proof of Lemma 1 and the fixed-wedge benchmark

By Jensen's inequality on the strictly concave $f(z) = (1+z)^{\varphi}$,

$$\sum_s f(z_s) \leq N f(B/N) = N^{1-\varphi}(B+N)^{\varphi},$$

with strict inequality for any non-uniform z on the fixed-budget simplex. The common multiplicative θ^0 preserves the argmax, proving Lemma 1.

For the fixed-wedge benchmark, the interior FOC of

$$\max_z \sum_s \theta_s(z_s)\Gamma_s \quad \text{s.t.} \quad \sum_s z_s = B$$

is $\varphi\theta^0(1+z_s)^{\varphi-1}\Gamma_s = \mu$. Solving and imposing the budget gives

$$1 + z_s = \frac{(B+N)\Gamma_s^{1/(1-\varphi)}}{\sum_r \Gamma_r^{1/(1-\varphi)}}.$$

Substitution into $\mathcal{M}(z)$ yields

$$\mathcal{M}(\Gamma) = \frac{\sum_s \Gamma_s^{\varphi/(1-\varphi)}}{N^{1-\varphi}\left(\sum_s \Gamma_s^{1/(1-\varphi)}\right)^{\varphi}} \leq 1,$$

with equality if and only if all Γ_s are equal. With heterogeneous θ_s^0, the productivity-maximizing allocation instead satisfies $1 + z_s^* \propto (\theta_s^0)^{1/(1-\varphi)}$; with heterogeneous φ_s, it solves the analogous sector-specific FOC. These are the heterogeneous-benchmark formulas referred to in Section 4.2.

For exposure weights $\omega_s > 0$, the weighted productivity benchmark solves

$$\max_z \sum_s \omega_s(1+z_s)^{\varphi} \quad \text{s.t.} \quad \sum_s z_s = B,$$

with interior solution

$$1 + z_s^{\omega} = \frac{(B+N)\omega_s^{1/(1-\varphi)}}{\sum_r \omega_r^{1/(1-\varphi)}}.$$

The corresponding weighted diagnostic is

$$\mathcal{M}_{\omega}(z) = \frac{\sum_s \omega_s(1+z_s)^{\varphi}}{\sum_s \omega_s(1+z_s^{\omega})^{\varphi}} \leq 1.$$

□

B.7 Proof of Lemma 2 (third-party input-cost incidence)

Appendix B.8 proves that each sectoral cost map g_s is strictly decreasing and strictly convex in z_s along the fixed-alignment directed-innovation channel. Hence each $-\omega_s g_s$ is strictly concave for $\omega_s > 0$, and $W_S = \sum_s(-\omega_s g_s)$ is strictly concave on the compact fixed-budget simplex. If its maximizer is interior, it satisfies $\omega_s g_s'(z_s^{\omega,g}) = -\mu$ for a common multiplier μ; more generally, the corresponding KKT conditions characterize the weighted nonlinear input-cost benchmark.

Under symmetric exposure and cost primitives, $\omega_s \equiv \omega$ and $g_s \equiv g$. The sum $\sum_s g(z_s)$ is symmetric and Schur-convex because g is convex, so $W_S = -\omega \sum_s g(z_s)$ is Schur-concave. The uniform vector is majorized by every vector on the fixed-budget simplex; Karamata's inequality gives

$$\sum_s g(z_s) \geq N g(B/N),$$

with strict inequality for non-uniform z. Therefore $W_S(z) \leq W_S(B/N, \ldots, B/N)$, strictly for non-uniform allocations.

With heterogeneous exposure or cost primitives,

$$\Delta W_S = W_S^{\text{Commit}} - W_S^{\text{SPE}} = \sum_s \omega_s g_s(z_s^{\text{SPE}}) - \sum_s \omega_s g_s(B/N),$$

so the sign is exactly the weighted input-cost comparison stated in Lemma 2. □

B.8 Convexity of the sectoral input-cost map

We prove part (i) of Lemma 2: $g_s(z) = C_s(\bar{w}/\theta_s(z))$ is strictly decreasing and strictly convex in z along the fixed-alignment directed-innovation channel, for sectors with positive rival-input exposure.

Fix alignment at its value along that channel and write $h_s(\theta) \equiv C_s(\bar{x}; \bar{w}/\theta, w_A)$, the

sectoral input cost as a function of rival productivity. Positive rival-input exposure means the CES share on the rival input is strictly positive. By the CES sufficient condition proved in Appendix B.2, $J(\theta) = V_{US}(\bar{x}) - h_s(\theta)$ is strictly increasing and strictly concave in θ on the relevant CES range. Since $V_{US}(\bar{x})$ does not depend on θ, h_s is strictly decreasing and strictly convex in θ: $h_s'(\theta) < 0$ and $h_s''(\theta) > 0$. The empirically relevant $\sigma_R \in (0, 1]$ lies inside the sufficient range $\sigma_R < 2$, so strict convexity holds throughout the calibration region.

The R&D technology $\theta_s(z) = \theta^0(1+z)^{\varphi}$ is strictly increasing and, for $\varphi \in (0, 1)$, strictly concave in z: $\theta_s'(z) > 0$ and $\theta_s''(z) < 0$. Composing,

$$g_s'(z) = h_s'(\theta_s(z))\,\theta_s'(z) < 0, \qquad g_s''(z) = \underbrace{h_s''(\theta_s(z))\,\theta_s'(z)^2}_{>0} + \underbrace{h_s'(\theta_s(z))\,\theta_s''(z)}_{>0} > 0.$$

The first term is positive because $h_s'' > 0$ and $\theta_s' > 0$; the second is positive because $h_s' < 0$ and $\theta_s'' < 0$. Hence g_s is strictly decreasing and strictly convex in z. This is the standard composition rule: a convex, non-increasing function of a concave function is convex. □

Remark. The result requires no new primitive. Convexity of g_s follows from the concavity of the Hegemon payoff on the CES family (Appendix B.2) together with the concave R&D technology maintained in Lemma 1. The strict inequalities $\sigma_R < 2$ and $\varphi < 1$ hold by assumption throughout the baseline.

B.9 Proof of Proposition 2 (Incidence)

The identity $\Delta W_{\text{total}} \equiv W_{\text{total}}^{\text{Commit}} - W_{\text{total}}^{\text{SPE}} = \Delta W_C + \Delta W_U + \Delta W_S$ follows by definition. The proposition does not claim an unconditional analytical sign for the aggregate wedge. It separates sign-indeterminate bilateral wedges from the third-party input-cost object signed in Lemma 2.

Step 1: separate the private R&D objective from national welfare W_C. For any fixed allocation z,

$$\theta_s(z_s)\Pi_s(\lambda_s^*(\theta_s(z_s),\sigma_s),\sigma_s) \geq \theta_s(z_s)$$

whenever $\Pi_s \geq 1$ on the calibrated support. This is the revealed-preference force behind the Rival's private period-1 objective under coercion: it explains why the SPE allocation tilts toward high-Γ_s sectors. It is not a revealed-preference ranking of national welfare across SPE and Commit. The Rival payoff wedge $\Delta W_C = W_C^{\text{Commit}} - W_C^{\text{SPE}}$ compares two regimes with different allocations and productivities; commitment loses the direct premium but can regain productivity through the more efficient fixed-budget allocation. Hence ΔW_C is quantitative rather than signed analytically, and at $\hat{\eta} = 0.54$ the full-KKT counterfactual gives $\Delta W_C = +0.015$.

Step 2: ΔW_U is sign-indefinite. Decompose the Commit-SPE Hegemon wedge:

$$\Delta W_U = \underbrace{\sum_s \left[U^s(\theta_s^{\text{Commit}},0) - U^s(\theta_s^{\text{SPE}},0)\right]}_{\substack{\text{Term B}' \\ \text{R\&D allocation at } \lambda=0}} - \underbrace{\sum_s \left[U^s(\theta_s^{\text{SPE}},\lambda_s^*) - U^s(\theta_s^{\text{SPE}},0)\right]}_{\substack{\text{Term A}' \\ \text{coercion at fixed } z^{\text{SPE}}}}.$$

Term A′ bound. Using the Hegemon's FOC in the baseline Λ-payoff block, $\partial U/\partial\lambda|_{\lambda^*} = x^*\theta - J(x^*;\theta) - \gamma\lambda^* = 0$, the optimal-$\lambda$ envelope theorem gives $U^s(\theta,\lambda^*) = (\gamma/2)(\lambda^*)^2 + J(x^*;\theta)$. Hence

$$\text{Term A}'_s = [J(x_s^*;\theta_s^{\text{SPE}}) - J(x_G^s;\theta_s^{\text{SPE}})] + \frac{\gamma}{2}(\lambda_s^*)^2 \geq 0.$$

The total is non-negative by Hegemon's revealed preference at λ^*, but the two summands have opposite signs: $J(x_s^*;\theta_s^{\text{SPE}}) - J(x_G^s;\theta_s^{\text{SPE}}) \leq 0$ (since x_G maximizes J at $\lambda = 0$ via the

Hegemon's x-FOC, so coercion-induced x^* strictly lowers J), while $(\gamma/2)(\lambda^*)^2 \geq 0$. By the value identity (6), Term $\mathrm{A}'_s = \Pi_E(\lambda^*_s) - (\gamma/2)(\lambda^*_s)^2$, the coercion-extraction value net of its deadweight cost; this net extraction value dominates the alignment loss, making the total non-negative.

Term B' sign. Under Assumption 2, $U^s(\theta_s, 0) = J(x^s_G; \theta_s)$ is concave and strictly increasing in θ_s. When the commitment allocation dominates the SPE allocation in the Jensen allocation index of Lemma 1, Term B′ is positive.

Sufficient condition. $\Delta W_U > 0$ iff Term B′ > Term A′, i.e.,

$$\sum_s [U^s(\theta_s^{\text{Commit}}, 0) - U^s(\theta_s^{\text{SPE}}, 0)] > \sum_s \left[J(x^*_s; \theta_s^{\text{SPE}}) - J(x^s_G; \theta_s^{\text{SPE}}) + \frac{\gamma}{2}(\lambda^*_s)^2 \right]. \qquad (\sharp)$$

At the full-KKT calibrated $\hat{\eta} = 0.54$, the numerically solved Hegemon wedge is $\Delta W_U \approx -0.013$ (Table 4), i.e., the Hegemon individually prefers SPE by a small margin. The Hegemon's wedge is η-invariant in the calibrated baseline; only the Rival's wedge varies with η.

Step 3: third-party input-cost comparison. By Lemma 2, symmetric exposure and cost primitives make the uniform Commit allocation the third party's cost-minimizing allocation, so a non-uniform SPE allocation lowers W_S and gives $\Delta W_S > 0$. Under heterogeneous exposure or cost primitives, the sign reduces to the weighted input-cost comparison in Lemma 2, which the calibrated model of Section 5 evaluates directly.

Step 4: aggregate decomposition. The aggregate decomposition is

$$\Delta W_{\text{total}} = \underbrace{\Delta W_C}_{\text{sign-indeterminate}} + \underbrace{\Delta W_U}_{\text{sign-indeterminate}} + \underbrace{\Delta W_S}_{\substack{\text{symmetric: } >0 \\ \text{heterogeneous: evaluated directly}}}.$$

In the symmetric case Lemma 2 signs the third-party term analytically; with heterogeneous exposure the calibrated model evaluates it directly. Numerically at the full-KKT calibrated

baseline (Section 5 and Table 4), $\Delta W_C = +0.015$, $\Delta W_U = -0.013$, and $\Delta W_S = +0.118$, so third-party incidence dominates the aggregate wedge. □

B.10 Proof of Proposition 3 (Surplus feasibility)

Let T_U and T_C denote aggregate transfers paid by exposed third parties to the Hegemon and the Rival under a commitment compact. The participation constraints are

$$\Delta W_U + T_U \geq 0, \qquad \Delta W_C + T_C \geq 0, \qquad \Delta W_S - T_U - T_C \geq 0.$$

The first two constraints require $T_U \geq \max\{-\Delta W_U, 0\}$ and $T_C \geq \max\{-\Delta W_C, 0\}$. Feasibility with the third-party constraint is therefore equivalent to

$$\Delta W_S \geq \max\{-\Delta W_U, 0\} + \max\{-\Delta W_C, 0\}.$$

When the condition holds, any transfer vector in this interval makes every player weakly better off than SPE under the common-numeraire incidence objects; if it holds strictly, the residual surplus can be split to make every player strictly better off under the same payoff comparison. When the condition fails, no exposed-third-party-financed transfer compact can implement commitment under these incidence objects. At the full-KKT calibrated baseline, $\Delta W_U = -0.013$, $\Delta W_C = +0.015$, and $\Delta W_S = +0.118$, so the calibrated transfer interval is $T_U \geq 0.013$, $T_C \geq 0$, and $T_U + T_C \leq 0.118$, leaving one-sided transfer surplus 0.105 after compensating the Hegemon. □

B.11 Proof sketch of Corollary 1

Let $G = \sum_k \max\{\Delta W_{S,k}, 0\}$ and $L = \sum_k \max\{-\Delta W_{S,k}, 0\}$. Since $\Delta W_S = G - L$, the aggregate surplus condition in Proposition 3,

$$\Delta W_S \geq \overline{M},$$

is equivalent to

$$G \geq \overline{M} + L = \overline{M}^I.$$

Thus the positive-gain exposed third parties have enough gross commitment surplus to cover the bilateral compensation requirement and, if needed, compensate third-party losers.

If a single exposed third party finances the compact, it must cover the full threshold $\overline{M}^I$ because restraint is non-excludable. When $\max_k \max\{\Delta W_{S,k}, 0\} < \overline{M}^I$, any such unilateral payment leaves that party strictly worse off than SPE, so no bilateral transfer-for-restraint contract is individually feasible. In the voluntary threshold-contribution game, at the no-contribution profile a unilateral pledge below $\overline{M}^I$ does not activate commitment and is refunded, while a unilateral pledge at least $\overline{M}^I$ activates commitment but yields payoff at most $\Delta W_{S,k} - \overline{M}^I < 0$. Hence no contribution is a Nash equilibrium.

For the institutional claim, necessity follows from the coalition members' participation constraints: an institution cannot collect more than $\sum_{k\in\mathcal{I}} \Delta W_{S,k}$ while leaving members weakly better off. For sufficiency, if $\sum_{k\in\mathcal{I}} \Delta W_{S,k} \geq \overline{M}^I$, the proportional assessment

$$t_k = \frac{\overline{M}^I}{\sum_{\ell\in\mathcal{I}} \Delta W_{S,\ell}} \Delta W_{S,k}, \qquad k \in \mathcal{I},$$

raises exactly $\overline{M}^I$ and satisfies $0 \leq t_k \leq \Delta W_{S,k}$ for every member. The institution uses $\overline{M}$ to satisfy the Hegemon and Rival compensation constraints in Proposition 3, uses L to make any third-party losers weakly whole, and binds the Hegemon to $\lambda_s = 0$ once the threshold is met. If the inequality is strict, each member with $\Delta W_{S,k} > 0$ keeps strictly positive residual surplus. □

B.12 Setup A robustness

Under Setup A (independent convex cost $\psi(z_s) = (\chi/2)z_s^2$), the Rival's FOC is sector-by-sector:

$$\varphi_s \theta_s^0 (1 + z_s)^{\varphi_s - 1} \Gamma_s(z_s; \eta) = \chi z_s.$$

Proposition 1 under Setup A. The comparative static $\partial z_s^*/\partial \Gamma_s > 0$ holds sector-by-sector without budget reallocation whenever the sectoral problem is locally concave. Hence $z_s^*(C) > z_s^*(S)$ for sectors whose exact marginal R&D return factor satisfies $\Gamma_s^C > 1 = \Gamma_s^S$. The fixed-budget threshold degenerates because sectors no longer compete for a common R&D budget.

Propositions 2, 3 under Setup A. Aggregate R&D $\sum_s z_s$ is *not* fixed across regimes under Setup A: each sector's FOC has higher Γ_s under Coercive, calling forth strictly higher z_s. Aggregate R&D rises under Coercive, which is the opposite of the Jensen Allocation Lemma's fixed-budget premise. The incidence and restraint-compact results are therefore Setup B/fixed-budget results: the directed-innovation allocation loss requires concentration of a fixed R&D budget rather than unconstrained aggregate expansion.

Hybrid model and empirical comparison. The hybrid R&D-cost specification spans the fixed-budget and level-response regimes:

$$\psi(\mathbf{z}; \xi) = \frac{\chi}{2} \sum_s z_s^2 + \frac{1}{2\xi} \Big(\sum_s z_s - B \Big)^2, \tag{25}$$

where $\xi \in [0, \infty]$ controls the elasticity of aggregate R&D supply. Its $\xi \to 0$ limit enforces the fixed-budget constraint, coinciding exactly with Setup B when $\chi = 0$ and otherwise adding within-budget convex adjustment costs; its $\xi \to \infty$ limit is Setup A. The Rival's hybrid FOC for each sector is

$$\theta_s'(z_s^*) \cdot \Gamma_s(z_s^*; \eta) = \chi z_s^* + \mu(\xi), \qquad \mu(\xi) \equiv \frac{1}{\xi} \Big(\sum_t z_t^* - B \Big), \tag{26}$$

which for any finite ξ ties sectors together through the common shadow term $\mu(\xi)$. The four-year cumulative response of Chinese R&D-to-GDP to the 2018 Coercive Turn (National Center for Science and Engineering Statistics, 2024; Section 5.1) is near zero (-0.06 pp-

years), consistent with the fixed-budget regime and far below Setup A's prediction of $\sim$ +0.18 pp-years. The aggregate R&D path therefore favors the fixed-budget limit over the unconstrained level-response benchmark, while the cross-sector concentration evidence in Section 2.2 supports the reallocation mechanism. The incidence and restraint-compact results stand under the data-supported fixed-budget regime.

C Empirical θ-path welfare sensitivity

The broad-bucket patent path is used to construct the empirical θ divergence that drives the directed-innovation channel in Section 5. Because the bucket-level evidence is more descriptive than the within-bottleneck FDPR calibration moment, Table 9 reports a simple attenuation diagnostic. Starting from the full-KKT baseline, we linearly scale down the directed-innovation component while holding the static forced-substitution channel fixed. The exercise is not a separate structural re-estimation; it asks how much of the patent-derived θ-path signal must remain for the compact result to survive.

Table 9: Sensitivity to attenuation of the empirical θ-path directed-innovation channel.

Scenario	θ scale	ΔW_U	ΔW_C	ΔW_S	ΔW_{total}	Transfer slack
Baseline full-KKT path	1.00	−0.013	+0.015	+0.118	+0.120	+0.105
Zero directed-innovation lower bound	0.00	−0.013	+0.008	+0.000	−0.006	−0.013
25% of baseline channel	0.25	−0.013	+0.009	+0.030	+0.026	+0.016
50% of baseline channel	0.50	−0.013	+0.011	+0.059	+0.057	+0.046
75% of baseline channel	0.75	−0.013	+0.013	+0.089	+0.089	+0.076

Notes: Cumulative 2018–2021 Commit-minus-SPE wedges in model units. "θ scale" is the fraction of the baseline pure directed-innovation channel retained. Transfer slack is $\Delta W_S - \max\{-\Delta W_U, 0\} - \max\{-\Delta W_C, 0\}$. The aggregate surplus threshold is 4.4% of the baseline directed-innovation channel; the one-sided compact threshold is 11.2% of the baseline third-party incidence. Source: authors' calculations.

D HS6 decomposition of third-party incidence

Table 10 reports the largest HS6 contributors to the exposure-weighted country decomposition in Table 5. The table allocates the model's bottleneck and commodity W_S wedges using 2014–2017 average China-origin imports for Korea, Taiwan, Japan, the Netherlands, Vietnam, and Mexico. Positive entries are HS6 codes for which commitment lowers third-party input costs relative to SPE; negative entries are commodity-bucket offsets from the uniform-Commit reallocation.

Table 10: Top HS6 contributors to allocated ΔW_S, 2018–2021.

HS6	Description	Bucket	China imports (M USD)	ΔW_S
852990	parts for radio/TV apparatus	commodity	7310.8	-0.060
854232	memories (DRAM, NAND)	bottleneck	9364.0	0.056
854239	electronic integrated circuits, other	bottleneck	8491.4	0.051
854140	photovoltaic cells and modules	bottleneck	6227.9	0.037
853400	printed circuits	bottleneck	4272.1	0.026
854231	processors and controllers (logic chips, incl. GPU/CPU)	bottleneck	3581.4	0.021
854370	electrical machines nesoi	bottleneck	2085.9	0.013
852872	TVs, colour	commodity	1202.4	-0.010
611020	sweaters, cotton	commodity	1156.5	-0.009
721049	flat-rolled steel, coated	commodity	1130.4	-0.009

Notes: China imports are 2014–2017 annual averages for the six target third-party economies. The underlying country-HS6 file is `tables/third_party_hs6_welfare_decomposition.csv`. Source: authors' calculations from BACI and the model-implied bucket wedges.

E Illustrative GE and dollar-conversion envelope

This appendix supplies the scale envelope referenced in Section 5.4. The main counterfactual is partial equilibrium: the Rival's R&D allocation affects θ_s, which feeds into period-2 third-party input costs through $w_C^s = \bar{w}/\theta_s$, while labor-market adjustment, aggregate-demand spillovers, and terms-of-trade responses are held fixed. The calculations below map the normalized wedge into illustrative GE-buffer and dollar ranges.

- *Labor reallocation.* Higher-cost upstream inputs would induce the exposed third party to reallocate labor away from coercion-affected sectors. Multi-sector GE models with sectoral mobility frictions (Caliendo and Parro, 2015) typically deliver reallocation buffers on the order of 30–50% of the underlying sectoral input-cost shock. Applied to $\Delta W_S^{\text{PE}} = +0.118$ (+0.94% under the total-model-surplus normalization), a 30–50% buffer scales the wedge to $[+0.059, +0.083]$ model units, or $[+0.47, +0.66]$%.

- *Terms of trade.* If the exposed third party can shift purchases toward US-aligned alternatives faster than $\sigma_s^{-\eta}$ implies in the partial setup, the GE-correct cost is lower. Boehm et al. (2019)'s chip-elasticity estimates and Atalay (2017)'s sector-level complementarity range suggest substitution is rigid at the 2–3-year horizon relevant for the calibration window, so this correction is likely smaller than the labor-reallocation bracket.

- *Aggregate output and small-open-economy structure.* For an exposed small open economy taking world prices as given, the GE feedback through aggregate output scales the wedge by a factor close to one. Atkeson and Burstein (2010)'s open-economy innovation framework provides a natural laboratory for embedding the composite redirection channel, but the current draft does not solve that model.

The resulting bracket is approximately $[+0.47, +0.94]$% under the total-model-surplus normalization, with the upper endpoint the partial-equilibrium calculation. The lower endpoint uses the borrowed 50% labor-reallocation buffer.

Dollar scale. Table 11 reports the dollar scale under two denominators. Applying the normalized percentage to full GDP yields tens of billions; applying it to an affected-input base yields an order of magnitude less.

Table 11: Illustrative dollar magnitude, Korea-sized economy.

Base applied to	Percentage	Implied cumulative loss
Full GDP (approx. \$6.84T, 2018–2021)	0.94% (PE)	approx. \$64B
Full GDP, 50% GE buffer	0.47%	approx. \$32B
Affected-input base (5–10% of GDP)	0.47–0.94%	approx. \$2–\$6B
Memo: paper-universe Korean China-origin bottleneck imports, 2018–2021, approx. \$6.3B.		

Notes: The paper-universe import memo is narrower than total bottleneck-input purchases because it uses the HS6 codes observed in the BACI incidence panel.

F Bottleneck-vs-commodity DiD baseline

This appendix reports the broad bucket-level comparison summarized in Section 2.3. The two-way fixed-effects baseline is

$$\log(1 + p_{h,t}) = \beta \, \text{Bottleneck}_h \cdot \text{Post2018}_t + \alpha_h + \delta_t + \varepsilon_{h,t}, \tag{27}$$

with HS6 and year fixed effects absorbing level differences and aggregate trends, commodity as the reference group, and standard errors clustered on HS6 (56 clusters). The clean-window 2010–2021 estimate is $\hat{\beta} = +0.247$ (SE 0.118), an approximately 28% post-2018 acceleration in bottleneck-relative-to-commodity patenting. The estimate attenuates to +0.114 in the full 2010–2024 window as USPTO grant-lag introduces post-2021 noise, and an IHS alternative gives +0.273.

Table 12: Bottleneck-vs-commodity DiD baseline.

Specification	$\hat{\beta}$	Asymptotic p	RI p	WCB p
$\log(1+x)$, clean window 2010–2021	+0.247 (0.118)	0.036	0.063	0.035
$\log(1+x)$, full window 2010–2024	+0.114 (0.091)	0.210	—	—
IHS, clean window	+0.273 (0.134)	0.042	—	—

Notes: Cluster-robust SE in parentheses (cluster: HS6, $N = 56$). RI = randomization inference (1,000 sharp-null permutations of the bottleneck assignment over the 56 HS6 codes). WCB = Cameron et al. (2008) wild-cluster bootstrap with Rademacher weights, $B = 5{,}000$, restricted-residuals variant on the baseline clean-window spec.

The result is useful as the broad data-side counterpart of the Section 5 θ-path divergence, but it is not the paper's identifying variation. The event study, first-difference robustness, Rambachan–Roth bounds, MIC2025 timing diagnostics, and EU-assignee placebo are reported in the appendices below.

G Synthetic control on HS6 854231

For the single most directly exposed HS6 code (854231: logic chips/processors, the input most directly cut off from Huawei), we apply the Abadie et al. (2010) Synthetic Control Method as an auxiliary exercise that complements the within-bottleneck dose-response in Section 2.2. Donor pool: 16 non-chip bottleneck HS6 codes. Pre-period 2010–2019; post-period 2020–2021. The pre-period RMSE is 0.147 (acceptable fit). The post-2020 average gap (treated minus synthetic) is +0.106 in log-patents. In the placebo permutation distribution (each donor as pseudo-treated), 854231 ranks 4th of 17, giving a one-sided rank-based p-value of 0.235.

The SCM result is *directionally consistent but underpowered* compared to the dose-response. Single-treated-unit SCM is weaker than multi-unit dose-response when each individual HS6 has noisy patent counts; pooling 4 Tier-1 HS6 codes into the dose-response averages out this idiosyncratic noise. We treat SCM as supportive corroboration rather than a primary design, and the Section 2.2 Tier-1/2/3 dose-response remains the headline within-bottleneck estimate.

H Rambachan-Roth pre-trend bounds

The Section F event-study (Figure 5) shows pre-period coefficients in 2010–2016 ranging from -0.31 to -0.08, several of which are statistically significant. This is a parallel-trends violation under the standard DiD framework, regardless of theory-side rationalization. We address the residual design concern directly using the Rambachan and Roth (2023) sensitivity-bound framework.

Linear-extrapolation bound. Fitting an inverse-variance-weighted OLS linear trend to the 2010–2016 pre-period coefficients yields a pre-trend slope of $\hat{\gamma} = +0.056$ per year (SE = 0.011). If this linear pre-trend continues through 2018–2021, the average extrapolated bias

on the post-2018 average effect is +0.140. Subtracting from the headline baseline:

$$\hat{\beta}_{\text{adj}} = \hat{\beta} - \overline{\hat{\gamma} \cdot (t - 2017)}$$
$$= +0.247 - 0.140 = +0.107, \qquad \text{(pre-trend-adjusted 95\% CI: } [-0.130, +0.343]).$$

The adjusted CI includes zero. This is the most-conservative interpretation of the Section F pattern: if the pre-2018 negative trend is genuinely linear and would have continued through 2018–2021 absent the policy shock, more than half of the headline +0.247 effect is explained by pre-trend continuation, and the residual is not statistically distinguishable from zero.

DeltaSD smoothness M-bound. The M-bound restricts post-period second-differences in the event-study path by M times the maximum pre-period second-difference. Computing the maximum pre-period second-difference $\max_{t<2018} |\Delta^2 \beta_t| = 0.091$, the conservative bound on the cumulative post-treatment bias over four post-years is $M \times 0.091 \times \binom{4}{2} = M \times 0.546$. Combined with sampling uncertainty:

M	95% CI on $\hat{\beta}_{\text{adj}}$	Excludes zero?
0.0 (no pre-trend bias)	$[+0.016, +0.477]$	Yes
0.5	$[-0.260, +0.753]$	No
1.0	$[-0.535, +1.028]$	No

The $M = 0$ CI matches the standard cluster-robust 95% CI; under any positive M, the bound is wide enough to include zero.

Interpretation. The Section F bucket-DiD baseline +0.247 is statistically significant under the standard parallel-trends assumption ($M = 0$; asymptotic $p = 0.036$, RI $p = 0.063$, WCB $p = 0.035$), but *is not robust to plausible pre-trend extrapolation* under the R-R linear or $M \geq 0.5$ bounds. We treat this as a known limitation. The headline redirection evidence in this paper is the Section 2.2 within-critical FDPR/MIC2025 moment, which uses HS6 fixed

effects that absorb the Section F pre-trend pattern at the bucket level; the Tier-1 $\beta_1 = +0.774$ (RI $p = 0.001$, asymptotic $p < 0.001$) is the most robust descriptive evidence. The Section F bucket DiD should be interpreted as supporting evidence whose magnitude is bounded above by the partial-equilibrium baseline and below (under R-R linear extrapolation) by approximately +0.12. Remark 3's theory-consistent reading is that the negative pre-trend reflects rival bottleneck productivity below the regime threshold, followed by threshold crossing in 2018. That reading predicts the pre-trend should *not* continue post-2018, in which case the R-R linear-extrapolation bound is too conservative and the headline +0.247 is unbiased. Both readings are reported transparently.

I First-difference robustness for the bucket DiD

The Section F event study acknowledges a parallel-trends violation: pre-2018 coefficients are negative (-0.31 to -0.08) rather than flat. Appendix H addresses this via Rambachan-Roth sensitivity bounds. As a complementary robustness check, we re-estimate the bucket DiD in first differences, which converts the parallel-trends assumption from "treated and control have parallel trajectories of log patents" (the assumption that is violated) to "treated and control have parallel growth rates" (a weaker, more defensible assumption).

Specification. The first-difference specification is

$$\Delta \log(1 + p_{h,t}) = \beta \cdot \text{Bottleneck}_h \cdot \text{Post2018}_t + \delta_t + \varepsilon_{h,t},$$

where $\Delta \log(1 + p_{h,t}) = \log(1 + p_{h,t}) - \log(1 + p_{h,t-1})$. HS6 fixed effects drop out under differencing; only year fixed effects remain. The interaction $\text{Bottleneck}_h \times \text{Post2018}_t$ now measures the differential change in growth rates at the 2018 transition, controlling for pre-2018 sector-level growth-rate differences.

Result. Table 13 reports the result.

Table 13: First-difference robustness for the bucket DiD.

Specification	$\hat{\beta}$ (SE)	Asymptotic p	RI p
Levels baseline (Table 12, clean window)	+0.247 (0.118)	0.036	0.063
First-difference, clean window 2010–2021	+0.009 (0.023)	0.691	0.695
First-difference, full window 2010–2024	−0.057 (0.023)	0.014	—

The first-difference clean-window estimate is near zero ($\hat{\beta}_{\text{FD}} = +0.009$, asymptotic $p = 0.69$, randomization-inference $p = 0.70$). Once we look at growth-rate changes at the 2018 transition rather than levels, there is no significant differential between bottleneck and commodity sectors. The full-window FD estimate is small and slightly negative (-0.057, $p = 0.014$).

First-difference event study. Figure 3 plots the first-difference event-study coefficients with 2017 as base year. The pattern is informative: pre-2018 growth-rate differentials fluctuate, with significant positive coefficients in 2011 (+0.075, $p < 0.10$) and 2016 (+0.087, $p < 0.05$); the 2018 (+0.074, $p < 0.05$) and 2019 (+0.085, $p < 0.05$) coefficients are similar in magnitude to the 2016 pre-treatment spike; 2020–2021 turn negative (-0.068, -0.054) as USPTO grant-lag sets in. There is no clean discontinuity at the 2018 transition: the post-2018 growth-rate jumps are not larger than pre-2018 ones.

Interpretation. The first-difference test does not survive at conventional significance, and the event study shows pre-2018 growth-rate volatility comparable in magnitude to the 2018–2019 jump. This is consistent with the Appendix H R-R linear-extrapolation bound: under the conservative reading that the negative pre-2018 trend is genuinely linear, much of the headline +0.247 is explained by pre-trend continuation, and the residual is not statistically distinguishable from zero. The first-difference result confirms this conclusion through a distinct methodological route. The bucket DiD's +0.247 should therefore be read as evidence of a level shift in bottleneck patenting that is bounded above by the partial-equilibrium baseline and below by approximately zero under the strictest pre-trend extrapolation. The

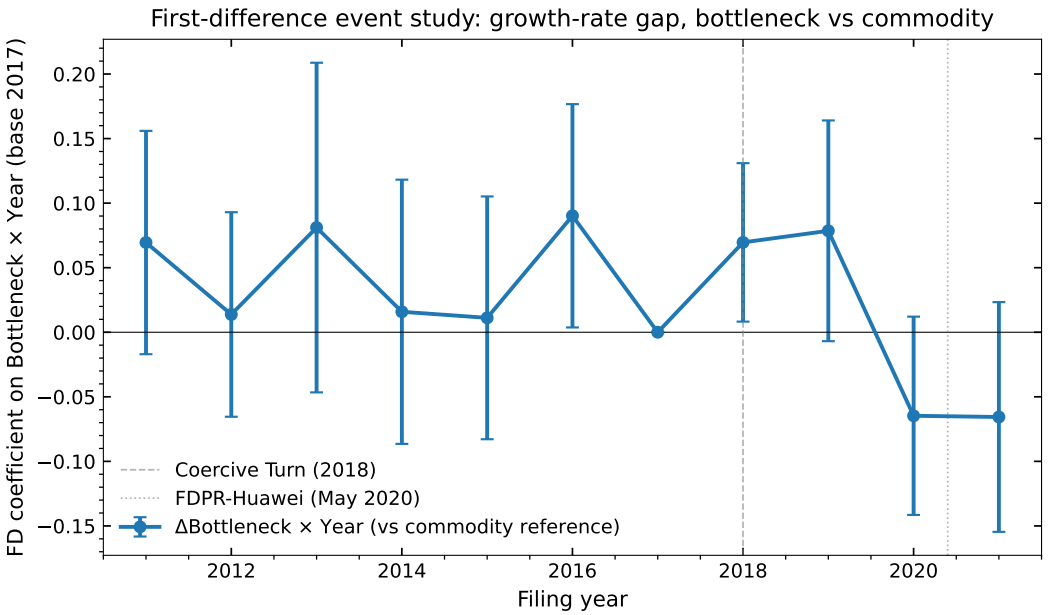

Figure 3: First-difference event study of the bottleneck-vs-commodity gap.

Notes: Coefficient on Bottleneck × Year in the FD specification, with 2017 as base year and cluster-robust 95% CIs. Pre-2018 growth-rate differentials are comparable in magnitude to post-2018 ones; no clean discontinuity at the 2018 transition.

Section 2.2 FDPR within-bottleneck dose-response, which uses HS6 fixed effects that absorb the bucket-level pre-trend, is the primary descriptive test of the mechanism.

J MIC2025 specification check

The Section F bucket DiD post-2018 acceleration overlaps heavily with China's Made-in-China-2025 (MIC2025) industrial policy, announced May 2015 with priority sectors covering chips, aerospace, new materials, electric vehicles, and power equipment. Of our 56 panel HS6 codes, 32 are MIC2025-priority (19 of 27 critical, 13 of 13 strategic, 0 of 16 commodity), and 8 critical raw-material HS6 codes (covering cobalt, antimony, cadmium, titanium, and related inputs) are bottleneck but not MIC2025-priority. We run progressively augmented specifications with MIC2025 timing dummies:

Specification	bottleneck × Post2018	MIC2025 × Post2015	MIC2025 × Post2018
M1 baseline (Eq. 27)	+0.243 ($p = 0.039$)	—	—
M2 + MIC2025 announcement	+0.102 ($p = 0.31$)	+0.301 ($p = 0.009$)	—
M3 + MIC2025 acceleration	−0.008 ($p = 0.96$)	+0.241 ($p = 0.011$)	+0.172 ($p = 0.19$)

The bottleneck × Post2018 effect attenuates from +0.243 in the baseline (M1) to near zero (−0.008, $p = 0.96$) once MIC2025 timing dummies are added (M3). The entire baseline effect comes from within-MIC2025 codes, and the 8 non-MIC2025 bottleneck codes show no post-2018 acceleration on their own.[11]

The bucket DiD is therefore observationally equivalent between two readings: (i) MIC2025 drives the acceleration, with no separate Coercive-Turn effect; (ii) forced-substitution differentially affects MIC2025-priority sectors that are also the most-bottlenecked. Since MIC2025-priority status and bottleneck-substitutability are nearly collinear at the HS6 level, this regression cannot discriminate the two readings. The Section 2.2 FDPR dose-response can: it isolates a sub-bucket-targeted post-May-2020 shock that MIC2025's 2015 announcement cannot explain.

K FDPR robustness checks

The FDPR event study in Appendix L shows pre-period coefficients ranging from −1.06 in 2010 to −0.27 in 2017 (Tier-1 patenting growing slower than Tier-3 reference codes pre-FDPR). This is a parallel-trends violation analogous to the bucket-level pattern, and we apply two complementary robustness checks: a first-difference specification (paralleling Appendix I for the bucket DiD) and a linear-extrapolation pre-trend bound (a simpler analog of the Rambachan and Roth (2023) (ΔSD, $\bar{M}$) machinery used in Appendix H).

[11] At the HS6 level, MIC2025 priority status implies bottleneck status (no commodity-bucket code is MIC2025-priority), so the natural triple-interaction (bottleneck × MIC2025 × Post2018) is collinear with MIC2025 × Post2018 by construction; we therefore omit it.

First-difference specification. The first-difference version of (1) is

$$\Delta \log(1 + p_{h,t}) = \beta_1 \cdot \text{Tier1}_h \cdot \text{Post2020}_t + \beta_2 \cdot \text{Tier2}_h \cdot \text{Post2020}_t + \delta_t + \varepsilon_{h,t}.$$

HS6 fixed effects drop out. Estimating on the same 27-code critical sub-bucket sample (2010–2021) gives $\hat{\beta}_1 = +0.177$ (cluster-robust SE 0.063, asymptotic $p = 0.010$). The FD estimate is roughly one-quarter of the levels baseline but remains positive and significant at the 1% level. By contrast, the bucket DiD's FD estimate (Appendix I) is near zero ($+0.009$, $p = 0.69$). The FDPR dose-response therefore survives the first-difference robustness check that the bucket DiD does not.

Linear-extrapolation pre-trend bound. Fitting an inverse-variance-weighted OLS linear trend to the 2014–2019 pre-period coefficients yields a pre-trend slope of $\hat{\gamma} = +0.105$ per year (SE = 0.013). If this linear pre-trend continues through 2020–2021, the average extrapolated bias on the post-2020 average effect is $+0.157$. Subtracting from the headline baseline:

$$\hat{\beta}_1^{\text{adj}} = +0.774 - 0.157 = +0.617.$$

The pre-trend-adjusted estimate remains positive and substantial: even under the conservative reading that the pre-2020 negative trend would have continued linearly through the post-period absent the FDPR, roughly 80% of the headline $+0.774$ survives.

Interpretation. The FDPR dose-response survives both robustness checks that the bucket DiD failed: a significant first-difference estimate ($+0.177$, $p = 0.010$) and a pre-trend-adjusted bound ($+0.617$) that remains large. This asymmetry justifies the primary role of (1) relative to the bucket DiD: although both specifications have pre-period coefficient patterns indicating parallel-trends violations, the FDPR estimate is robust to first-differencing and to R-R linear extrapolation, while the bucket DiD is not.

Firm and venue details. ChangXin (CXMT) and Yangtze (YMTC) account for 51% of post-2020 Tier-1 USPTO patents (ChangXin: 0 before 2019 to 18 in 2021; Yangtze: 1 in 2018 to 12 in 2021), while legacy incumbents reduced filings by −48.6 Tier-1 patents per firm after 2020. The pre-2020 industry was only moderately concentrated (Top-1/5/10 = 11%/38%/49% of cumulative patents). For the CNIPA check, PATSTAT TLS207_PERS_APPLN applicant matching uses `CHANGXIN MEMORY/STORAGE`, `CXMT`, and affiliated English-name variants for ChangXin, and `YANGTZE MEMORY`, `YMTC`, and affiliated variants for Yangtze. Re-estimating (1) with CNIPA patent counts from PATSTAT gives $\hat{\beta}_1^{\text{CNIPA}} = +0.143$ (SE 0.170, $p_{\text{RI}} = 0.40$): the sign matches the USPTO response, but the estimate is imprecise. Rate-of-change ratios differ across venues, consistent with venue substitution layered on real R&D rather than pure venue switching.

L Event studies

For transparency on the time-trajectory of the treatment effect, we report event-study versions of both empirical specifications. The pre-period coefficients in both event studies indicate a parallel-trends violation; we treat this as a known limitation of the DiD designs and rely on Appendix H (Rambachan-Roth bounds) and Appendix I (first-difference robustness) for pre-trend-adjusted assessment of the headline magnitudes.

FDPR within-bottleneck event study. Figure 4 plots the within-bottleneck Tier1 × Year event-study coefficients from a one-tier event-study version of (1), with 2019 as base year. The 2014–2017 coefficients are negative (−0.55 to −0.27, significant at conventional levels), indicating that Tier 1 chip codes were on a different trajectory from Tier 3 reference codes pre-2020. Coefficients move toward zero through 2018 and turn positive in 2020–2021; the 2021 coefficient is +0.327 ($p = 0.024$). The pre-period violation is the same kind of parallel-trends concern as the bucket-level event study below.

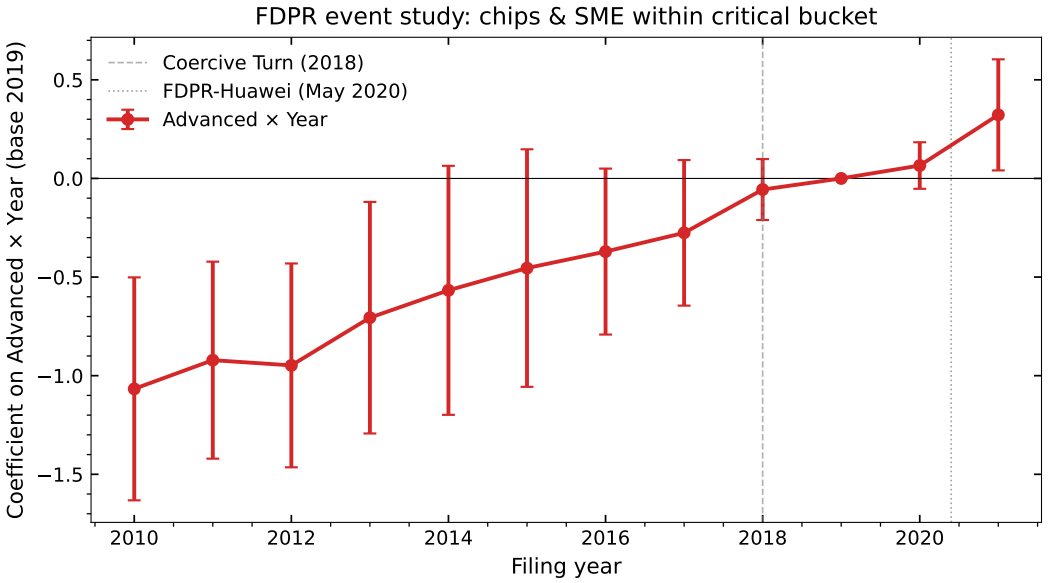

Figure 4: FDPR within-bottleneck event study.

Notes: Tier1 × Year coefficients from a one-tier event-study version of (1), base 2019. The 2014–2017 negative coefficients (Tier 1 growing slower than Tier 3) reverse to positive 2020–2021. The dotted vertical at May 2020 marks the FDPR.

Bottleneck-vs-commodity event study. Figure 5 plots the year-by-year coefficient on Bottleneck × Year from (27) with 2017 as base year. Pre-2018 coefficients are negative (−0.31 to −0.08, significant 2014–2016): bottleneck-bucket Chinese patenting was growing slower than commodity-bucket patenting before the trade-restriction shock. The 2018 coefficient (+0.07, $p < 0.05$) and 2019 coefficient (+0.15, $p < 0.01$) flip positive; 2020–2021 attenuate as USPTO grant-lag sets in. The negative pre-trend is a parallel-trends violation under the standard DiD framework.

M BACI unit-value incidence check

This appendix documents the public-data price-incidence screen in Section 2.3. The construction starts from BACI HS6 trade values and quantities. China-origin rows have strong quantity coverage in the relevant panel: among 77,846 China-origin product-importer-year rows, 2.36% have missing quantity and 2.55% have missing or nonpositive quantity; positive-quantity rows retain 99.07% of China-origin value. The baseline shock therefore uses unit

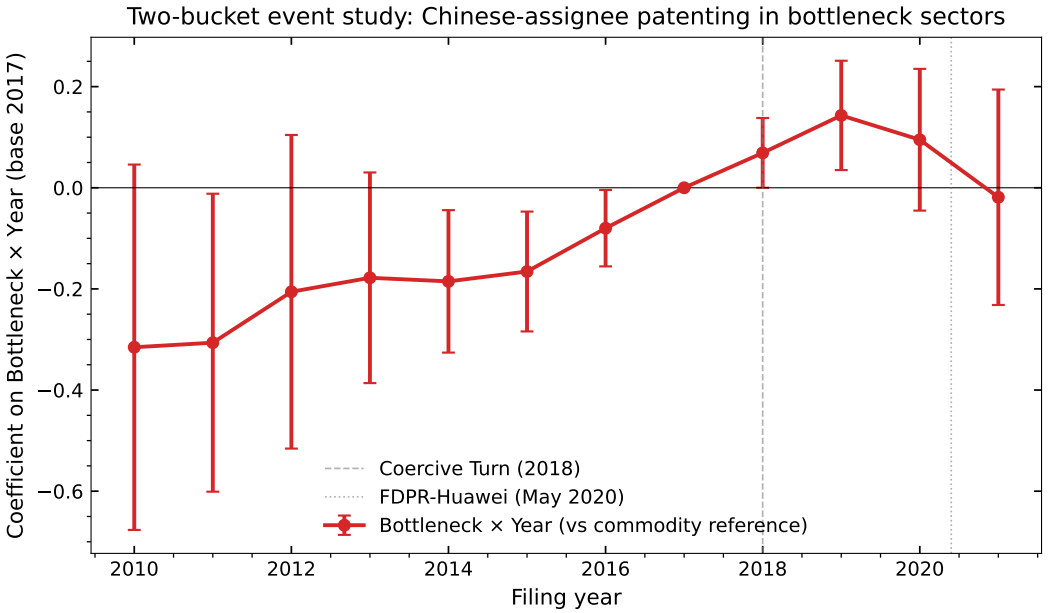

Figure 5: Two-bucket event study.

Notes: Bottleneck × Year coefficients from (27) with 2017 as base year. Cluster-robust 95% CIs. Pre-2018 coefficients are negative (bottleneck growing slower than commodity); post-2018 coefficients flip positive. Vertical dashed: 2018 trade-restriction shock; vertical dotted: May 2020 FDPR/MIC2025 moment.

values rather than trade values alone.

For each importing country c, HS6 product h, and year t, the leave-one-out China-origin unit value excludes country c from the numerator and denominator. The baseline exposure is a pre-period-base shock: log China-origin unit value in year t minus its 2014–2017 mean, aggregated using 2014–2017 country-input-HS6 shares and downstream input-output weights. We use the pre-period-base shock rather than a year-to-year log-change shock because input-cost pass-through in firm accounts is more persistent than annual unit-value changes.

Table 14 reports the COGS-share reduced form across samples. East Asia is Korea, Japan, and Taiwan. "Available EU" includes the EU economies with non-missing Compustat Global and BACI exposure coverage. The pooled sample combines East Asia and available EU countries. Gross margin from COGS is again the exact negative mirror of COGS share, so it is not shown as a separate outcome.

The East Asia result has the expected cost-incidence sign and passes a basic pre-period event-study screen: for the COGS-share specification, the maximum absolute 2014–2016 pre-

Table 14: BACI unit-value incidence robustness: COGS share.

Sample	Coefficient	SE	p-value	Observations
East Asia	+0.00197	0.00076	0.010	66,184
Available EU	+0.00154	0.00328	0.639	44,205
East Asia + available EU	+0.00182	0.00128	0.156	110,389
East Asia, with industry-year FE	−0.00020	0.00147	0.890	66,184

Notes: Outcome is winsorized COGS share. All specifications include firm and country-year fixed effects unless noted otherwise. Standard errors are clustered by country-industry. The shock is the standardized leave-one-out China-origin HS6 unit-value exposure using pre-period country-input-HS6 shares. Source: authors' calculations from BACI and Compustat Global.

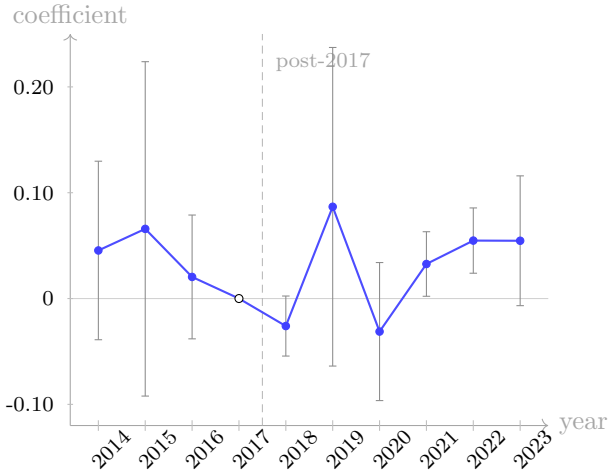

Figure 6: BACI/Compustat unit-value exposure event study for East Asian COGS share.

Notes: Coefficients are from the baseline firm and country-year fixed-effects specification with 2017 as the omitted year; vertical bars are 95% cluster-robust confidence intervals clustered by country-industry. The 2014–2016 coefficients have maximum absolute t-statistic 1.06.

period t-statistic is 1.06 relative to 2017 (Figure 6). The result is less robust to more saturated industry-year fixed effects, which absorb global industry shocks and leave identification from within-industry country exposure differences. We therefore interpret the exercise as direct descriptive support for the model's third-party input-cost channel rather than as a standalone causal estimate.

N Sample reconciliation

The empirical analysis uses three nested samples derived from a pre-specified 64-code HS6 critical-input concordance: a 56-code panel for the Section F bucket-level DiD; a 27-code critical-bucket subset for the Section 2.2 FDPR dose-response; and an 18-code donor pool for the Appendix G SCM (17 non-chip critical codes plus the treated 854231). Table 16 provides the row-by-row reconciliation.

Why 8 codes drop from the 64-code concordance. The 8 codes in Table 15 have zero Chinese-assignee USPTO patents in the 2010–2021 window after the IPC4→HS6 probabilistic concordance is applied. Three are battery accumulators (HS 850740 nickel-iron, 850750 nickel-metal hydride, 850760 lithium-ion), where Chinese filing is concentrated in CNIPA rather than USPTO; one is EVs (HS 870380), same logic; one is strontium/barium oxides (HS 281640) with limited US-market commercialization; two are polyethylene-terephthalate viscosity sub-codes (HS 390761 and HS 390769) that did not exist as separate HS12 codes in the IPC4→HS6 concordance vintage; and one is cotton T-shirts (HS 610910), low-tech with negligible patenting. The drops are not random with respect to bucket: 5 of 8 are in the bottleneck bucket (critical or strategic), the remaining 3 in commodity. This selection means our sample under-represents EV-battery and pure-raw-materials codes; the Section 2.2 FDPR sample (which targets chips and SME, not batteries or raw materials) is unaffected.

Section F sample composition (56 codes). The bottleneck-vs-commodity DiD pools critical (27 codes) and strategic (13 codes) into a single bottleneck bucket of 40 codes; commodity is 16 codes (all 19 in concordance minus 610910, 390761, 390769). Total $= 40 + 16 = 56$.

Section 2.2 FDPR sample composition (27 codes, critical-only). The FDPR dose-response uses only the critical-bucket codes (27 of 28 in concordance, with 281640 dropped). Tier 1 is 4 codes targeted by the May 2020 Huawei FDPR; Tier 2 is 6 chip-adjacent codes;

Table 15: HS6 codes excluded from the panel due to zero Chinese-assignee USPTO patents, 2010–2021.

HS6	Bucket	Description
281640	Critical	Oxides, hydroxides, peroxides of strontium or barium
390761	Commodity	Poly(ethylene terephthalate), viscosity $\geq$ 78 ml/g
390769	Commodity	Poly(ethylene terephthalate), viscosity $<$ 78 ml/g
610910	Commodity	T-shirts, cotton
850740	Strategic	Nickel-iron accumulators
850750	Strategic	Nickel-metal hydride accumulators
850760	Strategic	Lithium-ion accumulators
870380	Strategic	EVs (electric passenger vehicles)

Tier 3 is the remaining 17 critical codes (rare earths, lithium- and cobalt-related compounds, titanium, antimony, and other critical raw materials).

Table 16: HS6 sample reconciliation: full concordance to Section F and Section 2.2 samples.

Bucket	Concordance	In Section F (56)	In Section 2.2 (27)	Dropped
Critical	28	27	27 (T1: 4, T2: 6, T3: 17)	1
Strategic	17	13	0 (excluded by FDPR design)	4
Commodity	19	16	0 (reference category)	3
Total	**64**	**56**	**27**	**8**

Notes: Sourced from `tables/hs6_sample_reconciliation.csv`.

Full per-code reconciliation is reported in `tables/hs6_sample_reconciliation.csv`.