Optimal Tariffs under Strategic Risk: A Theory of De-risking*

Ahmad Lashkaripour
Indiana University, CEPR

Ina Simonovska
UC Davis, CEPR, NBER

September 2026

Preliminary and Incomplete

(Please do not circulate without author permission)

Abstract

We develop a theory of optimal tariffs based on strategic risk: the risk that a foreign supplier restricts exports in response to the importer's own dependence on it. Our framework rationalizes tariffs as a policy instrument for reducing excessive dependence on a single foreign supplier. Private import-sourcing decisions fail to internalize that the supplier's export policy responds to aggregate import dependence. This generates a strategic-dependence externality: individual firms' sourcing decisions collectively determine aggregate import dependence and exposure to policy-driven supply disruptions, such as export restrictions, imposing costs on other firms. We show that the optimal tariff comprises two components: a standard terms-of-trade term that is uniform across sectors, and a strategic-risk component that targets sectors with high aggregate import dependence. Calibrating the model to the U.S. and EU economies, we find that the strategic-risk component creates substantial cross-industry tariff variation. This variation explains the structure of the 2018 U.S. tariffs on China, which are difficult to explain using terms-of-trade considerations alone.

*We thank seminar and conference participants at the Third Jackson Hole International Finance and Macroeconomics Conference, SED 2026, Santa Clara University, Minneapolis Fed, BSE ADEMU 2026.

1 Introduction

Treasury Secretary Yellen warned in 2022 that *"we cannot allow countries to use their market position to disrupt our economy or exercise unwanted geopolitical leverage."* This remark epitomizes what policymakers increasingly recognize: aggregate import dependence, though necessitated by technological and resource constraints, can create strategic vulnerability.

The issue of strategic vulnerability has become particularly apparent given the recent proliferation of export restrictions. Most notably, export restrictions on critical inputs like rare earths, semiconductors, and medical supplies target sectors where importing countries are most import dependent. China imposed rare earth export quotas in 2010 after foreign manufacturers had concentrated production there, creating leverage over high-tech supply chains. During COVID-19, many countries imposed PPE export restrictions that proved disruptive precisely because production had concentrated among a handful of suppliers. When semiconductor shortages emerged in 2021, Taiwan's dominance of advanced chip manufacturing meant localized disruptions cascaded globally. In fact, as of 2023, these restrictions are the highest they have been in recent history [Figure 1]. The common thread in these examples is that aggregate import dependence creates strategic vulnerability to foreign export policies.

Despite this new reality, standard tariff theory treats foreign supply as a purely *technological* constraint, and attributes the gains from unilateral tariffs exclusively to terms-of-trade manipulation. This framing misses the above-noted aspect of the current policy landscape: foreign supply is not merely a technological object but a strategic one. When aggregate import dependence invites export restrictions, decentralized sourcing decisions generate a strategic-dependence externality, which leads to inefficient levels of aggregate import dependence. Tariffs can correct this externality by lowering aggregate import dependence.

We develop a tractable model to formalize this point. The model has two basic ingredients: First, production is roundabout: firms use traded intermediates and domestic labor. Second, Foreign can impose export restrictions (modeled as an export tax) in a noncooperative regime. And from a purely welfare maximizing standpoint, the expected export restriction is increasing in Home's aggregate import dependence. Home firms are, however, atomistic and take the expected export restriction as given when choosing sourcing. Thus, atomistic sourcing decisions generate excessive aggregate import dependence relative to the welfare-maximizing allocation.

An export restriction is a risk rather than a certainty because noncooperation is costly for Foreign. In the model, Foreign restricts exports only if the gain from doing so exceeds

a cost η, which we interpret as the cost of violating the rules of the trading system. Two features of η are central to the analysis. First, it is the *gain* from restricting, rather than the cost, that increases with Home's import dependence; a restriction therefore becomes both more likely and more severe as dependence grows, whatever the level of η. Second, for the foreign supplier we have in mind, China, the level of η is well documented, and it is small. The cost is also not constant over time. The cost that China attaches to violating the rules varies with political circumstances that Home cannot foresee when it commits to a durable tariff: the 2010 rare-earth restriction followed a maritime incident near the Senkaku Islands, the 2023 graphite controls followed the European Union's investigation of Chinese electric vehicles, and the 2025 rare-earth controls followed the escalation of U.S. tariffs (Bown, 2024, 2025).[1] The model captures this feature by having Nature draw η after Home has committed to its tariff. Home knows only the distribution of η, which the record described below and the estimates in Section 6.2 identify. The same structure, a policy regime that switches stochastically while agents make irreversible decisions, underlies the literature on trade policy uncertainty (Handley and Limão, 2017).

The rules of the trading system impose weak discipline on export restrictions. Article XI of the GATT prohibits export quotas but not export taxes, and it exempts restrictions applied to relieve "critical shortages" of "essential" products, neither of which is defined (Bown, 2025). China, however, accepted a stricter commitment as a condition of its accession in 2001. Paragraph 11.3 of its Accession Protocol obliges it to eliminate all export taxes except on 84 listed products, at bound maximum rates, an obligation that only a small number of recently acceded members carry (Qin, 2012; Ruta and Venables, 2012; Bown and Hillman, 2019). China has violated this commitment repeatedly and has lost every dispute brought against it. In *China—Raw Materials* (2012) and *China—Rare Earths* (2014), the WTO's Appellate Body found China's export duties and quotas on industrial raw materials, rare earths, tungsten, and molybdenum inconsistent with the Protocol and with Article XI, and held that the general exceptions of Article XX could not be invoked to justify them (Bronckers and Maskus, 2014). China complied in each case by removing the measures at issue, and subsequently imposed new restrictions: the United States filed a third challenge, over nine further raw materials, in 2016, and since 2023 China has restricted exports of gallium, germanium, graphite, antimony, tungsten, and rare earths, among other products, under a national-security licensing regime that lies outside the scope of its accession commitments (Bown, 2023, 2024, 2025). In the

[1] The 2025 restrictions were a response to Home's tariffs, a channel that the model does not contain: in the model a higher tariff lowers Foreign's gain from restricting (Lemma 3), whereas retaliation lowers the cost of restricting. We treat that episode as outside the timing of the model. The tariffs we study were set in 2018, seven years before that episode.

assessment of Bown (2025), "these WTO rulings did little to stop or seemingly slow China's active use of export-restricting policies."

This record informs our interpretation of η. The cost of noncooperation is real: China has complied with every WTO ruling against its export restrictions. It is, however, small, for three reasons documented in the literature. WTO remedies are prospective only, so a restriction that is imposed and later withdrawn carries "minimal consequence" (Wu, 2017). Compliance is enforced by the complainant's capacity to retaliate rather than by legal obligation as such (Bown, 2004), and that channel has weakened since the Appellate Body ceased to function in 2019 (Bown and Keynes, 2020). Finally, China's own conduct reveals how it weighs the cost: it exhausts reductions in export VAT rebates, which its commitments do not cover, before resorting to export taxes that might violate them (Eisenbarth, 2017; Garred, 2018). The same pattern appears in the data of Section 6.2: export taxes account for 21 of the 81 restrictive measures that China implemented through 2017, but for only one of the 85 implemented since 2023, when bans and licensing requirements predominate. A small η implies that the cooperative regime is fragile: once import dependence makes a restriction profitable, few obstacles remain. Section 5 formalizes this cost, and Section 6.2 estimates how sharply the risk of restriction rises with dependence.

Our main result is a closed-form decomposition of the optimal import tariff. Let $\Lambda \equiv \lambda_{FH}$ denote Home's aggregate expenditure share on Foreign goods. Let $\bar{\tau}(\Lambda)$ denote the certainty-equivalent gross export-tax factor on shipments to Home. Under mild regularity conditions, the optimal Home tariff satisfies

$$t^* = \underbrace{\frac{1}{\varepsilon^S_{FH}}}_{\text{terms of trade}} + \underbrace{\frac{\epsilon\,\lambda_{HH}}{\varepsilon^D_{FH}} \cdot \frac{d\ln\bar{\tau}(\Lambda)}{d\ln\Lambda}}_{\text{strategic-risk component}}, \tag{1}$$

where ε^S_{FH} is the Foreign export-supply elasticity faced by Home, ε^D_{FH} is Home's import-demand elasticity for Foreign goods, and ϵ is the trade elasticity. The first term is the standard optimal-tariff component. The second term is new: it corrects the strategic-dependence externality by reducing Λ when expected foreign restrictions $\bar{\tau}$ rise with aggregate import dependence.

The multi-sector extension presents a clear rationale for targeting, which has proven elusive from the lens of quantitative trade models. Despite sector-specific trade elasticities ϵ_k, the terms-of-trade component of the optimal tariff is uniform across sectors. However, the strategic-risk component varies with aggregate import dependence and

with the sensitivity of expected export restrictions to that dependence. This prediction closely aligns with the 2018 U.S. tariffs on China which are difficult to explain through the lens of standard terms of trade theory.

The 2018 U.S. tariffs should be interpreted from this perspective. Chinese export restrictions were neither new in 2018 nor confined to rare earths. In the Global Trade Alert data used in Section 6.2, China implemented 81 restrictive export measures between 2008 and 2017, in the form of export duties, quotas, licensing requirements, and bans, covering 838 HS6 products in 46 chapters of the tariff schedule, from chemicals, steel, and non-ferrous metals to fertilizers, coke, and agricultural goods. Thirty-five of these measures, covering roughly 600 products, were in force at the end of 2017, and the database lists the United States among the affected countries for nine in ten of the measures implemented over this period. The raw-material and rare-earth restrictions that the United States challenged at the WTO in 2009, 2012, and 2016 were the litigated part of this regime. The 2018 tariffs were not, however, a response to any of these measures. The Section 301 investigation that produced them concerned technology transfer and intellectual property, and the tariff lists covered mainly manufactured goods. What the United States had was a record: a decade of restrictions across hundreds of products, a subset of which it had litigated, and the 2010 embargo against Japan. From this record it drew the inference that the model attributes to Home: not that a restriction on any particular product was imminent, but that one could be imposed at a time of China's choosing. Contemporaneous policy documents state this inference explicitly. Executive Order 13817 of December 2017 stated that dependence on foreign sources of critical minerals "creates a strategic vulnerability for both its economy and military to adverse foreign government action" (Executive Office of the President, 2017). The interagency assessment of the defense industrial base of September 2018 warned that "when China needs to flex its soft power muscles by embargoing rare earths, it does not hesitate, as Japan learned in a 2010 maritime dispute," and that China's "willingness to use trade as a weapon" raised the risk of relying on it (Interagency Task Force in Fulfillment of Executive Order 13806, 2018). The Commerce Department's 2019 critical-minerals strategy posed the hypothetical directly: "if China or Russia were to stop exports to the United States and its allies for a prolonged period—similar to China's rare earths embargo in 2010" (U.S. Department of Commerce, 2019). Bown (2019) summarizes what the United States knew at the time: China "had already shown a willingness to exercise market power to its economic advantage," including through "WTO-illegal export restrictions on raw materials and rare earth metals." The restrictions that this policy anticipated arrived in 2023 and after. In the interim, the threat was made explicit: in May 2019, China's

official press asked whether rare earths would become "a counter weapon" against the United States and answered, "don't say we didn't warn you" (Reuters, 2019). Section 8.1 documents the same anticipation in the communications of the U.S. Trade Representative. A tariff set under this logic is, in the words of Bown (2025), an insurance policy whose "benefit pays off when a bad shock hits, but the costly premium has to be paid even in good states of the world." Such a tariff is set ex ante, against a distribution of outcomes, and it should be evaluated as such.

We take the theory to the data in three steps. First, we estimate the elasticity of the export-restriction hazard with respect to import dependence, ζ, using a newly assembled panel of export restrictions from the Global Trade Alert and OECD databases matched to bilateral trade flows over 2017–2023. The estimates lie between 0.1 and 0.3 across specifications: exporters that account for a larger share of an importer's purchases in a sector are more likely to restrict exports of that sector's products. Second, we calibrate the model to the United States and the European Union. The terms-of-trade benchmark yields an optimal tariff on the traded composite of 25.8 percent for the United States and 26.0 percent for the European Union; strategic risk raises these rates to 31.0 and 30.9 percent, so that the strategic-risk component accounts for roughly one sixth of the optimal tariff. Across sixteen sectors, the terms-of-trade component is uniform, whereas the strategic-risk component ranges from under two percentage points in services to more than eleven in the most import-dependent sectors. Third, we examine the 2018 U.S. tariffs on China from two angles: the rationale that the U.S. Trade Representative stated, and the structure of the tariffs that it imposed. In the universe of 2,983 USTR press releases issued between 2010 and 2025, the narratives that connect foreign exposure to the conduct of foreign governments and to strategic vulnerability roughly quadruple in the 2018 rollout year, from 6.0 to 26.6 percent and from 3.9 to 15.6 percent of releases, and the increase is concentrated between the March announcement and the July implementation of the first tariff list. Language about import dependence alone does not rise significantly, so the shift is specific to the risk of foreign-government action rather than to trade or China in general. Turning to the tariffs themselves, across HS10 tariff lines the growth of China's import share over 2013–2017 is a significant positive predictor of the 2018 tariff increase, with coefficients of 0.65 at the HS10 level and 0.92 at the HS4 level, whereas the export-supply and import-demand elasticities that govern terms-of-trade and political-economy tariffs have no predictive power. The same regression estimated on tariffs applied to non-China imports yields a coefficient of zero.

Related Literature

Optimal tariffs. The first component of our optimal tariff is the classical terms-of-trade tariff (Bickerdike, 1906; Johnson, 1953; Gros, 1987). In Ricardian and quantitative trade models, this component is uniform across sectors (Costinot et al., 2015), so that standard theory offers no rationale for targeting particular sectors. The second component provides such a rationale. Blanchard et al. (2026) show that discretionary tariffs decrease with the foreign value added embodied in domestic production and with the domestic value added embodied in foreign exports, so that global value chains lower optimal tariffs. Our model shares their input–output structure through roundabout production, but the strategic-risk component operates in the opposite direction and for a different reason: it depends not on the ownership of the value added in a trade flow but on how the foreign government's export policy responds to that flow.

Supply risk and resilience. A growing literature studies policy toward supply disruptions. Grossman et al. (2023) ask whether policy should promote international diversification or reshoring. In their model, firms choose between a home supplier, a foreign supplier, or both before disruptions are realized, and the disruption probabilities are exogenous, with the home source safer than the foreign source by assumption. The wedge between private and social sourcing decisions arises from the consumer-surplus and business-stealing externalities of monopolistic competition, which offset exactly under CES demand; in that case no intervention in supply chains is warranted. Grossman et al. (2024) (see also Grossman et al. 2023) study investments in resilience along a multitier supply chain in a closed economy, where bilateral bargaining leads firms to underinvest. Elliott et al. (2022) and Kopytov et al. (2024) study the formation of production networks under uncertainty and the substitution toward more reliable suppliers as uncertainty rises. Closest to our result in form is Adamopoulos and Leibovici (2025), whose optimal tariff is the sum of a terms-of-trade component and a risk premium. In their model, the probability of a trade disruption is fixed, and risk enters the optimal tariff by reweighting the terms-of-trade motive across states of the world. Leibovici and Santacreu (2020) study a dynamic model in which financial frictions generate underinvestment in essential goods, and the optimal policy combines an export tax, a production subsidy, and an import subsidy. In each of these papers the disruption is exogenous: its probability does not depend on the importer's reliance on the source, and the foreign government is not a decision maker. Our externality is therefore absent from all of them. In our model, the probability of a restriction responds to aggregate import dependence because the foreign government's gain from restricting increases with it. The externality operates through

the probability of a disruption rather than through the consequences of a given disruption, and it remains present under CES demand, where the demand-side externalities of Grossman et al. (2023) vanish.

Hold-up and sourcing externalities. The externality we identify has a clear antecedent in McLaren (1997). In that model, atomistic investors in a small country, anticipating trade liberalization with a large partner, sink resources into the export sector and thereby erode the country's bargaining position; the small country can be left worse off than under an anticipated trade war. Our strategic-dependence externality has the same structure: irreversible private decisions collectively determine a country's strategic position, and no individual agent internalizes the effect. The two models differ in the object that is degraded and in the policy response. In McLaren (1997), the degraded object is the threat point of a Nash bargain within an efficient agreement, so that dependence affects the division of surplus; in our model, it is the foreign government's incentive to restrict exports, so that dependence affects whether a disruption occurs. Moreover, in McLaren (1997) the government has no policy instrument prior to the investment decision, so the externality is identified but not corrected; in our model, the Home government commits to a durable tariff before firms choose their sourcing, so the externality can be corrected. The commitment role that Çelik et al. (2020) attribute to fast-track authority is played here by the import tariff.

Export restrictions and the trading system. A policy literature argues that the proliferation of export restrictions is itself driving the return of tariffs. Bown (2025) documents the rise of export restrictions on critical inputs and the weakness of multilateral disciplines on them, and argues that the fear of such restrictions "has driven . . . the increased use of tariffs" in importing countries. Ruhl (2026) makes a similar argument and emphasizes the difficulty of sustaining the required policy commitment. This paper provides a formal model of that argument and an empirical test of it. Our treatment of the cost of noncooperation builds on the theory of self-enforcing trade agreements (Bagwell and Staiger, 1999; Maggi and Staiger, 2011), in which a government's adherence to the rules is sustained by the cost of violating them; Bown (2004) shows empirically that this cost is determined largely by the complainant's capacity to retaliate. Bown and Hillman (2019) and Bown (2024) document China's accession commitments on export taxes and its compliance record, on which we draw above. On the empirical side, Bown and Crowley (2013) show that governments respond to import surges from a trading partner with new protection, and their specification anticipates our regression of the 2018 tariff changes on the growth of China's import share. The economic mechanism differs: in their model, foreign supply is passive and protection serves as a safety valve

that sustains a cooperative agreement, whereas in our model protection responds to the partner's incentive to restrict exports.

Geopolitical risk. Finally, a recent literature studies geopolitics directly, which is not the focus of this paper. Becko et al. (2025) derive the optimal tariff of a great power that values the political alignment of smaller states. In their model, each small country draws a non-pecuniary cost of aligning with the hegemon and aligns if the economic gain from preferential market access exceeds that cost; the hegemon's tariff raises this gain by lowering the price received by non-aligned exporters, and the optimal tariff exceeds the terms-of-trade level by the value of the alignment purchased at the margin. Our model shares this structure, with the roles of cooperation and defection reversed. They model the cost of cooperating and a tariff that increases the return to cooperation; we model the cost of defecting and a tariff that reduces the return to defection through the sourcing decisions of Home firms. In their model, there is no wedge between private and social incentives within the hegemon, and the tariff corrects no distortion; in ours, the tariff corrects a wedge among Home's own firms. Clayton et al. (2026a) and Clayton et al. (2024) study a hegemon that coordinates threats across economic relationships to extract concessions, and Clayton et al. (2026b) show that private firms are easier to coerce than their governments because they "internalize less than their own government the equilibrium consequences of acquiescing to the hegemon's coercion." The externality that they identify as a source of the hegemon's leverage is the externality that our tariff corrects. These papers solve the problem of the coercing country under complete information, with threats that are always effective and never executed; we solve the problem of the target country, with a probability of restriction that increases in the target's own import dependence. Becko and O'Connor (2025) is the most closely related paper to ours. In their model, capital producers fail to internalize how their investment affects the rival's incentives, as in ours. However, they show that state-contingent trade threats render a corrective instrument unnecessary, and they assign trade taxes to the rival's dependence rather than to Home's own. Each of these papers places a political objective, whether alignment, coercion, or the rival's geopolitical action, at the center of the analysis. In our model, Foreign restricts exports because doing so is profitable, and the only cost it weighs against that gain is the cost of violating the rules of the trading system.

2 Background: Shifting Trade Policy Landscape

We set the stage for our theory by highlighting two policy developments that are either overlooked in traditional optimal tariff theories or difficult to reconcile with them. The first of these is the remarkable rise in export restrictions over the past couple of decades. The second is the most recent wave of U.S. tariffs, the structure of which is difficult to explain with traditional theories.

Fact 1. *Systematic shift in global export policy: the frequency of export restrictions on critical inputs increased four-fold from 2009 to 2023.*

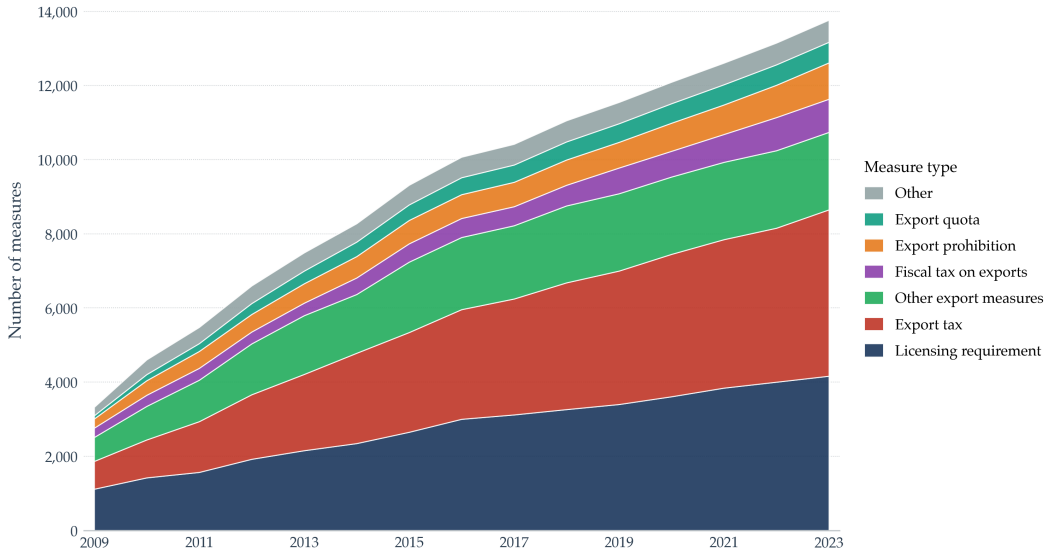
Concerns about export restrictions are often cited as justification for protectionist policies aimed at strengthening economic resilience and hedging strategic risk. The key question is whether these restrictions have become a greater cause for concern in recent years. A natural way to examine this is to look at the number of export restrictions on raw materials. Historically, these goods have been among the most susceptible to restrictions and have driven much of the anxiety around supply chain risk.

To this end, we examine the OECD Inventory of Export Restrictions on Industrial Raw Materials, which covers export taxes, licensing requirements, quotas, prohibitions, and fiscal levies. Figure 1 plots the number of active export restrictions by type from 2009 to 2023. The number of active restrictions quadrupled over this period, with export taxes representing the most prominent form of restriction. This pattern reaffirms the claim made at the outset of this paper that export restrictions have proliferated over the past decade, fueling a potential shift in trade policy sentiment. We next highlight a high-profile example of such a policy shift, noting that it is difficult to reconcile with traditional optimal tariff theories based on terms-of-trade manipulation or political economy pressures.

Fact 2. *Systematic shift in US import policy: Recent USTR tariffs reflect an aversion to excess dependence on risky foreign suppliers. The 2018 U.S. tariffs on China targeted precisely the products whose import dependence on China had been growing in 5-year leading to the policy, which is difficult to explain with traditional theories.*

The 2018 tariff war between the US and China marks another pivotal shift in the trade policy landscape. What's revealing about this shift is the type of goods that were targeted. The China-specific Section 301 tariff hikes targeted goods for which China's share of U.S. imports, a measure of China's leverage as a supplier, had risen, while exempting

Figure 1: Export restrictions on industrial raw materials, 2009–2023



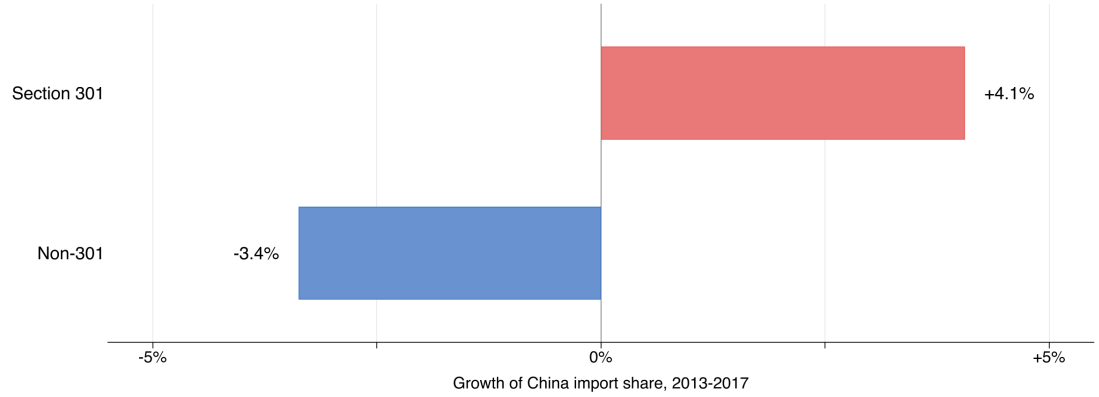
Note: The data is from the OECD Inventory of Export Restrictions on Industrial Raw Materials, which covers export taxes, licensing requirements, quotas, prohibitions, and fiscal levies. The figure displays the stock of active measures by year, where a measure is the unique combination of reporting country, six-digit HS product, measure type, and date of introduction that has been notified to the OECD and not yet eliminated. Between 2009 and 2023 the stock of export restrictions more than quadrupled, rising from 3,323 to 13,766 measures.

goods for which that leverage had fallen. Figure 2 displays this divergence: over 2013–2017, products on the Section 301 list saw China's import share grow by an average of 4.1 percent, while the remaining products saw it decline by 3.4 percent. Later, we examine this pattern more closely and find that Section 301 goods were not meaningfully different from non-targeted goods in the underlying export supply and import demand elasticities emphasized by traditional optimal tariff theory. Instead, the targeting echoes the political rhetoric surrounding these tariffs: the administration framed its objective in terms of strategic vulnerability and national security rather than import price manipulation.

This pattern of targeting is difficult to reconcile with classic optimal tariff theory. The textbook optimal tariff formula asserts that $t^* = 1/\varepsilon^S$, where ε^S is the elasticity of foreign export supply. This elasticity reflects the home country's monopsony power, which bears little relation to sector-level import dependence.[2] We later show that this elasticity

[2]In Ricardian models the terms of trade-optimal tariff is indeed uniform across sectors, making the model irreconcilable with targeting. The specific factors variant of the terms of trade theory in which each sector uses fixed industry-specific inputs generates optimal targeting based on the sector-level export supply elasticity. But this elasticity has no clear structural relationship to importer-side demand shares. Instead, it reflect buyer market power and is regulated by exporter-side supply shares.

Figure 2: The 2018 tariffs targeted goods with growing China import dependence



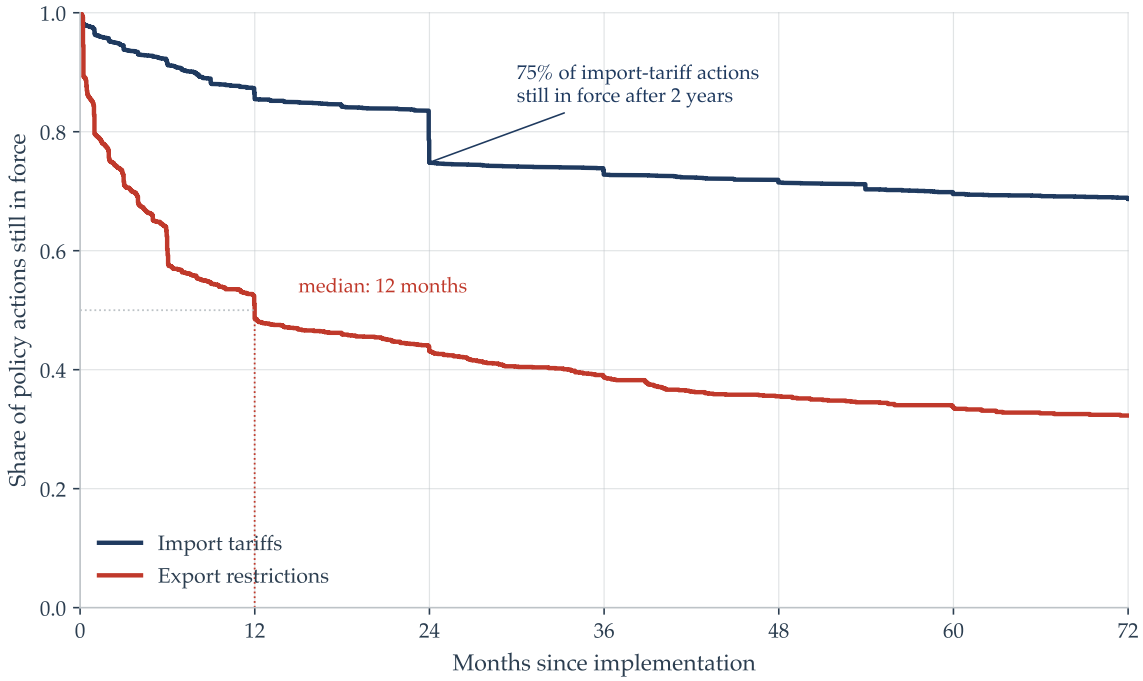
Notes: This paper reports the growth in China import dependence among goods targeted by Section 301 tariffs and the non-targeted goods. Each bar is the 2013–2017 growth of a group's aggregate China import share in percent. Corresponding average growth rates across individual lines are reported in Appendix J.2. The sample consists of 8,520 HS10 import lines (5,244 Section 301, 3,276 Non-301) with positive China imports in both 2013 and 2017 and an observed December 2017–December 2018 change in the statutory China tariff; the Non-301 group includes the tariff lines under the global Section 232 and 201 tariffs. Data on statutory tariffs and trade values is sourced from Fajgelbaum et al. (2020).

has no explanatory power for whether sectors were targeted by Section 301 tariffs. The political economy theory of optimal tariffs proposed by Grossman and Helpman (1994) cannot rationalize the targeting structure either. In the *protection for sale* framework, the political benefits of protection are proportional to the rents accruing to domestic producers, while the social cost rises with the volume of imports. The model thus predicts that protection tracks the import demand elasticity facing each good, and that controlling for that elasticity the rationale for protection would fade in sectors where import dependence is growing. The 2018 tariffs did precisely the opposite, targeting the goods whose dependence on China had been growing.

Fact 3. *Export restrictions are frequently raised and removed, whereas import tariffs are less frequently introduced and have high levels of persistence.*

An often overlooked regularity in the trade policy literature is that export and import restrictions operate on different time horizons. We show this asymmetry using data from the Global Trade Alert (GTA). Import tariffs in the GTA database tend to be far more persistent than export restrictions. As illustrated Figure (3), the median export restriction in the is removed within a year, suggesting that these restrictions are adopted in response to a temporary shock like the pandemic or war and then reversed shortly after. The median import tariff, by contrast, remains in force six years later. The asymmetry holds even when the same country applies both types of measures to similar goods (see

Figure 3



Appendix XXX).

Overall, these patterns suggest that import tariffs are durable commitments, whereas export restrictions are used as flexible instruments. This distinction is at the root of how set up our trade policy game in the following section: we assume that the Home country makes a durable import tariff commitment, anticipating that its trading partner will flexibly adjust its export taxes as circumstances change.

3 Model

The global economy consists of two countries: Home (H) and Foreign (F). Let $\mathcal{C} \equiv \{H, F\}$ denote the set of countries, with generic elements $i, n \in \mathcal{C}$. Each country is endowed with a fixed supply of labor, denoted L_i for country i, which is supplied inelastically. Let w_i denote the wage rate in country i.

3.1 Final Good Aggregator

Consider a representative household in destination country $i \in \mathcal{C}$. The household derives utility from consuming differentiated varieties produced by firms located across both

countries. Let Ω_n denote the set of firms located in origin country n, with $\bar{N}_n \equiv |\Omega_n|$ representing the measure of firms in country n. Each firm $\omega \in \Omega_n$ produces a unique differentiated variety.

Let $c_{ni}(\omega) \geq 0$ denote the quantity of the variety produced by firm $\omega \in \Omega_n$ that is consumed by the representative household in country i. The household aggregates varieties through a CES utility function

$$C_i = \left(\sum_{n \in \mathcal{C}} C_{ni}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}}, \qquad with \quad C_{ni} = \left[\int_{\Omega_n} c_{ni}(\omega)^{\frac{1}{\mu}} d\omega \right]^{\mu} \tag{2}$$

where $\sigma > 1$ is the Armington elasticity of substitution across national-level varieties and $\mu \geq 1$ is the degree of firm-level market power, determined by the sub-national elasticity of substitution between firm varieties:

$$\gamma \equiv \frac{\mu}{\mu - 1}$$

Consumer Demand. Let $\tilde{p}_{ni}(\omega) > 0$ denote the *consumer price* in destination country i for one unit of the variety produced by firm $\omega \in \Omega_n$. This price is inclusive of all trade costs and taxes, as specified below. Let Y_i denote the *net expendable income* of the representative household in country i, defined as wage income plus rebated trade-tax revenue. The household maximizes utility (2) subject to the budget constraint:

$$\sum_{n \in \mathcal{C}} \int_{\Omega_n} \tilde{p}_{ni}(\omega)\, c_{ni}(\omega)\, d\omega \leq Y_i.$$

The household's optimization yields the following CES demand function for each variety $\omega \in \Omega_n$:

$$c_{ni}(\omega) = \left(\frac{\tilde{p}_{ni}(\omega)}{\tilde{P}_{ni}} \right)^{-\gamma} \left(\frac{\tilde{P}_{ni}}{\tilde{P}_i} \right)^{-\sigma} \frac{Y_i}{\tilde{P}_i}, \tag{3}$$

where $\tilde{P}_{ni}$ and $\tilde{P}_i$ denote aggregate price indices in country i:

$$\tilde{P}_{ni} \equiv \left(\int_{\Omega_n} \tilde{p}_{ni}(\omega)^{1-\gamma} d\omega \right)^{\frac{1}{1-\gamma}} \qquad \tilde{P}_i = \left[\sum_{n \in \mathcal{C}} \tilde{P}_{ni}^{1-\sigma} \right]^{\frac{1}{1-\sigma}}.$$

The *indirect utility* of the representative household equals real income, and pins down social welfare, denote by $W_i = Y_i / \tilde{P}_i$.

3.2 Production Technology

Each country $n \in \mathcal{C}$ hosts a fixed set Ω_n of differentiated firms that compete under monopolistic competition. Firms produce using labor and a composite intermediate input, with a constant-returns-to-scale Cobb-Douglas technology. Let $\ell_n(\omega)$ denote the labor input and $M_n(\omega)$ denote the composite intermediate input used by firm $\omega \in \Omega_n$. Each firm is characterized by a productivity parameter $\varphi_n(\omega) > 0$, which is fixed and heterogeneous across firms. The production function is:

$$q_n(\omega) = A_n \left(\frac{\ell_n(\omega)}{\beta}\right)^{\beta} \left(\frac{M_n(\omega)}{1-\beta}\right)^{1-\beta}, \qquad \beta \in (0,1], \tag{4}$$

where $q_n(\omega)$ denotes the output of firm ω and β represents the *value-added share* in production, restricted to $\beta \geq 1/2$ in theory.

The composite intermediate input $M_n(\omega)$ aggregates over the same differentiated varieties that enter final consumption, with the identical CES aggregation structure:

$$M_n(\omega) = \left(\sum_{j \in \mathcal{C}} M_{jn}(\omega)^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}, \qquad \textit{with} \quad M_{jn}(\omega) = \left[\int_{\Omega_j} m_{jn}(\omega,\omega')^{\frac{1}{\mu}} d\omega'\right]^{\mu} \tag{5}$$

where $m_{jn}(\omega,\omega')$ denotes the quantity of variety $\omega' \in \Omega_j$ used as an intermediate input by firm $\omega \in \Omega_n$.

Cost Minimization and Optimal Pricing. Given the Cobb-Douglas production technology (4) and the CES intermediate input aggregator (5), cost minimization yields the unit input cost:

$$\mathscr{C}_n(w_n, \tilde{P}_n) = w_n^{\beta}\, \tilde{P}_n^{1-\beta}, \tag{6}$$

where $\tilde{P}_n$ represents the intermediate input price index, which is the same as the consumer price index, since the intermediate composite uses the same aggregation structure as final consumption.

Under monopolistic competition, each firm charges a factory-gate price equal to a constant markup equal to μ over marginal cost:

$$p_n(\omega) \;=\; \frac{\mu}{A_n}\mathscr{C}_n(w_n, \tilde{P}_n) \;=\; \frac{\mu}{A_n} w_n^{\beta}\, \tilde{P}_n^{1-\beta}. \tag{7}$$

3.3 Trade Costs and Policy Instruments

International trade is subject to both natural trade frictions and policy-induced barriers. Let $d_{ni} \geq 1$ denote the *iceberg trade cost* for shipping goods from origin n to destination i: delivering one unit in country i requires producing d_{ni} units in country n. We normalize domestic trade costs to unity: $d_{ii} = 1$ for all $i \in \mathcal{C}$.

Countries may impose ad valorem trade taxes. Let $t_{ni} \geq 0$ denote the *import tariff* levied by destination country i on imports from origin $n \neq i$, with $t_{ii} = 0$. Let $x_{ni} \geq 0$ denote the *export tax* levied by origin country n on exports to destination $i \neq n$, with $x_{ii} = 0$. Define the corresponding gross tax factors:

$$\tau_{ni}^{M} \equiv 1 + t_{ni}, \qquad \tau_{ni}^{X} \equiv 1 + x_{ni}. \tag{8}$$

3.4 Policy-Implied Buyer Prices

The consumer price in destination i for variety $\omega \in \Omega_n$ incorporates the factory-gate price, iceberg trade costs, export taxes, and import tariffs. We assume that the export tax applies to the after-iceberg price $d_{ni} p_n(\omega)$, and the import tariff applies on top of the export-tax-inclusive price. This yields:

$$\tilde{p}_{ni}(\omega) = \tau_{ni}^{M} \, \tau_{ni}^{X} \, d_{ni} \, p_n(\omega). \tag{9}$$

For domestic varieties, this simplifies to $\tilde{p}_{ii}(\omega) = p_i(\omega)$ since $\tau_{ii}^{X} = \tau_{ii}^{M} = d_{ii} = 1$.

3.5 Total Demand and Gross Expenditure

Cost minimization by firms implies that the demand for intermediate input varieties follows the same structure as final demand. Specifically, the demand for variety $\omega' \in \Omega_j$ as an intermediate input by firm $\omega \in \Omega_n$ is:

$$m_{jn}(\omega, \omega') = \left(\frac{\tilde{p}_{jn}(\omega')}{\tilde{P}_{jn}} \right)^{-\gamma} \left(\frac{\tilde{P}_{jn}}{\tilde{P}_n} \right)^{-\sigma} M_n(\omega), \tag{10}$$

where $M_n(\omega)$ is the firm's intermediate input aggregate, so that $\tilde{P}_n M_n(\omega)$ is its total intermediate input expenditure. Aggregating across all firms in country n, total intermediate input expenditure equals:

$$\tilde{P}_n M_n = \frac{1 - \beta}{\beta} w_n L_n, \tag{11}$$

Total demand facing firm $\omega \in \Omega_n$ in destination market i consists of *(i)* final consumption demand $c_{ni}(\omega)$ and *(ii)* intermediate input demand from all firms located in i. Define total demand as:

$$q_{ni}(\omega) \equiv c_{ni}(\omega) + \int_{\Omega_i} m_{ni}(\omega', \omega) d\omega'. \tag{12}$$

Given the common CES structure, total demand simplifies to:

$$q_{ni}(\omega) = \left(\frac{\tilde{p}_{ni}(\omega)}{\tilde{P}_{ni}}\right)^{-\gamma} \left(\frac{\tilde{P}_{ni}}{\tilde{P}_i}\right)^{-\sigma} \frac{E_i}{\tilde{P}_i}, \tag{13}$$

where E_i denotes *gross expenditure* in country i, defined as the sum of final consumption expenditure and intermediate input expenditure:

$$E_i \equiv \tilde{P}_i C_i + \tilde{P}_i M_i = Y_i + \frac{1-\beta}{\beta} w_i L_i. \tag{14}$$

The second lines imposes that final consumption expenditure equals net income Y_i which is the sum of wage income and profits rebates, which will be characterized shortly.

3.6 General Equilibrium

Aggregate Trade Shares. Define the *aggregate expenditure share* on goods sourced from origin n in destination i as $\lambda_{ni} \equiv \int_{\Omega_n} (\tilde{p}_{ni}(\omega) q_{ni}(\omega) d\omega) / E_i$. Substituting the demand function (13) and the pricing equation (7), and using the consumer price relationship (9), the aggregate trade share becomes:

$$\lambda_{ni} = \frac{\left(\frac{1}{\bar{\varphi}_n} \tau_{ni}^M \tau_{ni}^X d_{ni} w_n^\beta \tilde{P}_n^{1-\beta}\right)^{-\epsilon}}{\sum_{\ell \in \mathcal{C}} \left(\frac{1}{\bar{\varphi}_\ell} \tau_{\ell i}^M \tau_{\ell i}^X d_{\ell i} w_\ell^\beta \tilde{P}_\ell^{1-\beta}\right)^{-\epsilon}}, \tag{15}$$

where $\epsilon \equiv \sigma - 1$ is the trade elasticity and $\bar{\varphi}_n$ is a constant measure of aggregate productivity: $\bar{\varphi}_n \equiv \bar{N}_n^{\frac{1}{\epsilon}} A_n$. These shares satisfy the adding up constraint $\sum_{n \in \mathcal{C}} \lambda_{ni} = 1$ for all $i \in \mathcal{C}$. Henceforth, we impose the restrict ion $\frac{L_H}{L_F} \leq \left(\frac{\bar{\varphi}_F}{\bar{\varphi}_H}\right)^{1/\beta}$ to ensure that that market share satisfy a plausibility constraint, $\lambda_{HF} \leq \lambda_{FF}$, under Laissez-faire.

Aggregate Price Index. The consumer price index in country i can be expressed in terms of aggregate trade costs and unit costs:

$$\tilde{P}_i = \left[\sum_{n \in \mathcal{C}} \bar{\varphi}_n^{\epsilon} \left(\tau_{ni}^M \, \tau_{ni}^X \, d_{ni} \, w_n^{\beta} \, \tilde{P}_n^{1-\beta} \right)^{-\epsilon} \right]^{-\frac{1}{\epsilon}}. \tag{16}$$

This equation, together with its counterpart for the other country, forms a system that determines the equilibrium price indices as functions of wages and policy variables.

We also denote

$$p_n = \frac{1}{A_n} \, w_n^{\beta} \, \tilde{P}_n^{1-\beta} \qquad\qquad \tilde{p}_{ni} = \tau_{ni}^M \, \tau_{ni}^X \, d_{ni} \, p_n$$

Factor Market Clearing. Labor market clearing requires that total wage payments in country n equal the labor share of producer revenue across all destinations. The delivered producer-value revenue from origin n to destination i (net of both taxes) is $\lambda_{ni} E_i / \tau_{ni}^M \tau_{ni}^X$. Summing across destinations and using the labor share β yields:

$$w_n L_n = \frac{\beta}{\mu} \sum_{i \in \mathcal{C}} \frac{1}{\tau_{ni}^M \tau_{ni}^X} \, \lambda_{ni} E_i, \qquad \forall n \in \mathcal{C}. \tag{17}$$

Profits are given by $\Pi_n = (\mu - 1) \frac{1}{\beta} w_n L_n$. Net income is $Y_n = w_n L_n + \Pi_n$.

Tax Revenue and Net Income. Import tariff revenue collected by destination country i equals the tariff fraction of buyer spending on imports:

$$T_i^M = \sum_{n \neq i} \frac{\tau_{ni}^M - 1}{\tau_{ni}^M} \lambda_{ni} E_i. \tag{18}$$

Export tax revenue collected by origin country n equals the export-tax fraction of the net-of-import-tariff value of exports:

$$T_n^X = \sum_{i \neq n} \frac{\tau_{ni}^X - 1}{\tau_{ni}^X} \cdot \frac{1}{\tau_{ni}^M} \lambda_{ni} E_i. \tag{19}$$

Net income in country i is wage income plus profits plus rebated trade-tax revenue:

$$Y_i = (1 + \frac{\mu - 1}{\beta}) w_i L_i + T_i^M + T_i^X. \tag{20}$$

Combining (14) and (20) yields an alternative expression for gross expenditure:

$$E_i = \frac{\mu}{\beta} w_i L_i + T_i^M + T_i^X. \tag{21}$$

Definition 1. *[General Equilibrium]* For a given set of trade taxes $\{t_{ni}, x_{ni}\}$, a general equilibrium is a collection of prices $\{w_i, \tilde{P}_i\}_{i \in \mathcal{C}}$, trade shares $\{\lambda_{ni}\}_{n,i \in \mathcal{C}}$, tax revenues $\{T_i^M, T_i^X\}_{i \in \mathcal{C}}$, and expenditures $\{Y_i, E_i\}_{i \in \mathcal{C}}$ such that: (i) *trade shares satisfy* (15); (ii) *price indexes satisfy* (16); (iii) *labor markets clear according to* (17); (iv) *tax revenues satisfy* (18)–(19); (v) *income and expenditure satisfy* (20)–(21); *and* (vi) *trade is balanced according to* (22).

Balanced Trade. There is no international lending or borrowing. The national-level budget constraint, therefore, implies that trade is balanced. That is, the net-of-import-tariff value of Home's imports from Foreign equals the net-of-import-tariff value of Foreign's imports from Home. The balanced trade condition can be formally written as:

$$\frac{1}{\tau_{FH}^M} \lambda_{FH} E_H = \frac{1}{\tau_{HF}^M} \lambda_{HF} E_F. \tag{22}$$

4 Optimal Import Tariff if Foreign is Passive

This section characterizes Home's optimal unilateral trade policy under the assumption that Foreign remains passive. Since, our goal is to highlight the rationales for trade policy intervention, we purposely abstract away from domestic market failures that require domestic policy interventions. The following assumption ensures that.

Assumption 1. *Firms possess no market power, $\mu = 1$. Equivalently, any firm-level market power is fully offset by domestic policy, so that the allocation is efficient absent trade taxes.*

Policy environment and notation. Without loss of generality, we restrict Home's policy to import tariffs and Foreign's policy to export taxes. The reason this restriction does not lose generality is due to the Lerner symmetry, whereby Home's import tariff could replicate any mix of import and export taxes and so can Foreign's export tax. To streamline the notation we use t to denote Home's ad valorem import tariff rate and x to denote Foreign's ad valorem export tax rate:

$$t \equiv t_{FH} = \tau_{FH}^M - 1, \qquad x \equiv x_{FH} = \tau_{FH}^X - 1$$

Under this choice of notation, the buyer price of Foreign goods in Home is

$$\tilde{p}_{FH} = (1+t)(1+x)d_{FH}p_F,$$

where p_F is Foreign's producer price determined in equilibrium. In the benchmark case analyzed here, Foreign is passive, so the export tax is constrained to equal zero: $x = 0$. Home's optimal policy problem therefore maximizes Home's aggregate welfare subject to equilibrium constraints and a passive foreign policy. Formally,

$$\max_{t \geq 0} \; W_H(x,t) \qquad s.t. \quad \begin{cases} \text{equilibrium constraints} \\ x = 0 \end{cases}, \tag{23}$$

where $W_H(x,t) = Y_H(x,t)/\tilde{P}_H(x,t)$ is the aggregate Home welfare under policy (t,x). The constraint $x = 0$ is what distinguishes the above benchmark from the strategic export policy case analyzed next. The optimal tariff formula implied by the above problem is traditionally stated in terms of import demand and export supply elasticities. So, before presenting the formula, we derive the noted elasticities within our framework under CES aggregation and roundabout production.

Export Supply and Import Demand Elasticity. We set $p_H \equiv 1$ by choice of numeriare herefater. Leveraging the balanced trade condition, we can write the foreign export supply function in term of $p_F \equiv p_F/p_H$ as

$$Q_{FH} = \bar{\varphi}_F^{\frac{1}{\beta}} \lambda_{FF}(p_F)^{-\frac{1-\beta}{\beta\epsilon}} \lambda_{HF}(p_F) L_F \tag{24}$$

where the trade share functions are defined as $\lambda_{FF}(p_F) := p_F^{-\epsilon}/(p_F^{-\epsilon} + d_{HF}^{-\epsilon})$ and, analogously, $\lambda_{HF}(p_F) = 1 - \lambda_{FF}(p_F)$. Equation 24 implicitly solves p_F in terms of the export supply Q_{FH}. Applying the Implicit Function Theorem to the above equation obtains the export supply elasticity:

$$\varepsilon_{FH}^S \equiv \frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln p_F} = \epsilon\lambda_{FF} + \frac{1-\beta}{\beta}\lambda_{HF} \tag{25}$$

The term $\epsilon\lambda_{FF}$ accounts for the standard substitution effect: as the relative price of Foreign goods rises, Foreign buyers substitute away from Foreign goods toward Home goods, freeing Foreign output for export. The second term $\frac{1-\beta}{\beta}\lambda_{HF}$ reflects the change in per unit labor cost due to roundabout production. Since the labor requirement per unit

of gross output is $\beta p_F/w_F$, a rise in p_F reflects a reduction in p_F/w_F, and thus the unit labor cost, by exactly $(1-\beta)/\beta)\lambda_{HF}$. Foreign's fixed labor force can therefore support more gross output, and this raises export supply. To regulate the relative magnitude of these terms, we assume that the trade elasticity and value added share satisfy the following condition, putting them within the range of commonly-used estimates in the literature.

Assumption 2. The trade elasticity and value added share satisfy: $\epsilon > \frac{1-\beta}{\beta}$

The import demand elasticity could be characterized by noting that the import demand function for foreign varieties is $Q_{FH} = \lambda_{FH} E_H / \tilde{p}_{FH}$. Taking derivatives from this equation, yields the following expression for the import demand elasticity

$$\varepsilon^D_{FH} \equiv -\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln \tilde{p}_{FH}} = 1 + \frac{\epsilon\lambda_{HH}}{1-\frac{t}{1+t}\lambda_{FH}} + \frac{1-\beta}{\beta}\lambda_{FH} \tag{26}$$

The first two terms are standard under CES demand. The third term captures the intermediate input cost feedback arising from roundabout production. The intuition is as follows: holding $p_H \equiv 1$ fixed, a higher import price $\tilde{p}_{FH}$ passes through to the price of the foreign intermediate input bundle $\tilde{P}_F$ via exports of embedded intermediate inputs. This effect multiplies the effective import price increase, thereby magnifying the resulting contraction in import demand.

Optimal Tariff Formula. When the Foreign government is passive, Home's optimal tariff problem is essentially the textbook problem governed by a familiar trade-off. A higher tariff improves Home's terms of trade by lowering the pre-tariff price of imports p_F. However, the resulting contraction in trade volume lowers consumer surplus. The optimal tariff that maximizes the net welfare gains from Home, equals the inverse of the export supply elasticity defined above. The following proposition states this result.

Proposition 1. *If Foreign is passive, $x = 0$, Home's unilaterally-optimal tariff is*

$$t^* = \frac{1}{\varepsilon^S_{FH}} = \frac{1}{\epsilon\lambda_{FF} + \frac{1-\beta}{\beta}\lambda_{HF}}. \tag{27}$$

While Proposition 1 reinstates the textbook optimal tariff logic, it has a novel element worth highlighting. Here, the export supply elasticity is derived under roundabout production, distinguishing it from existing formulas in the literature. The extent of

roundabout production is regulated by β, creating two interesting limiting cases:

$$t^* = \begin{cases} \frac{1}{\epsilon \lambda_{FF}} & \beta = 1 \\ 0 & \beta \to 0 \end{cases}$$

In the special case with no traded intermediate inputs, $\beta = 1$, the formula reduces to the familiar formula in Gros (1987). However, when β assumes intermediate values, roundabout production effects raise Foreign's export supply elasticity, lowering Home's optimal tariff rate. Notably, in the limit as $\beta \to 0$, Foreign's export supply becomes infinitely elastic and the optimal tariff converges to zero. This result formalizes the conventional wisdom that under intermediate input trade, the optimal tariff must be lower as it penalized domestic producers.

5 Optimal Tariff if Foreign is Strategic

This section moves beyond the traditional setup to characterize Home's optimal tariff when Foreign may adopt a non-cooperative regime and set its export tax strategically. This introduces a new role for tariffs that is absent in the textbook case. Here tariffs can mitigate the risk of export restrictions due to excess sourcing from foreign suppliers. The idea is that firms' private sourcing decisions can collectively shift Foreign's export tax schedule. However, because individual firms are infinitesimal, their private decisions do not internalize these collective effects. This gives rise to a *strategic-dependence externality* whereby excessive concentration of imports from Foreign suppliers raises the vulnerability of other domestic firms to shifts in Foreign's export restrictions.

Sequential Trade Policy Game. Following the Lerner symmetry logic stated above, we suppress all policy instruments except Home's import tariff $t \equiv t_{FH}$ on Foreign goods and Foreign's export tax $x \equiv x_{FH}$ on shipments to Home. These policy instruments are selected sequentially in the strategic policy game outlined below.

Timing and Payoffs. The policy environment is a Bayesian Stackelberg game, expressed in Harsanyi form. Nature draws Foreign's type $\eta \sim G$, where η is the cost that Foreign incurs if it abandons the cooperative regime and restricts exports. We interpret η as the cost of violating the rules of the trading system, which comprises litigation and retaliation under the WTO and the reputational cost of using trade as an instrument of coercion. As documented in the Introduction, for China this cost is real but small. The

results that follow do not depend on this interpretation; η may represent any cost of noncooperation that does not vary with Home's import dependence. Home observes G, but not η, and commits to an import tariff $t \equiv t_{FH}$ as the first mover. After observing t and its type η, Foreign chooses a regime $r \in \{\text{cooperative}, \text{noncooperative}\}$ which determines its export tax $x \equiv x_{FH}$. If $r =$ cooperative, Foreign sets the cooperative export tax, $x = 0$. If $r =$ noncooperative, Foreign chooses a noncooperative export tax $x > 0$ that maximizes its own welfare.

We solve best policy responses in the above game via backward induction, starting with Foreign's optimal policy in the second period. To make the timing explicit, wherein t is selected first and x is chosen given t, country i's log welfare is denoted by

$$v_i(x;t) \equiv \ln W_i(x;t)$$

The function v_i represents government i's continuation payoff in terms of the chosen strategies. We continue to normalize $p_H \equiv 1$ by choice of numeraire.

5.1 Foreign's optimal strategy

Foreign choses the export tax rate $x \equiv x_{FH}$ given Home's import tariff $t \equiv t_{FH}$ and the cost of noncooperation η. Specifically, it decides whether to remain cooperative, $x = 0$, or to impose a noncooperative export-tax equivalent $x > 0$. Given the realization of the cost of noncooperation η, Foreign solves

$$\max_{x \geq 0} \left\{ v_F(x;t) - \eta \, \mathbf{1}_{\{x>0\}} \right\}.$$

The draw η affects whether Foreign finds it worthwhile to adopt the noncooperative regime. Conditional on adopting that regime, however, the export tax is chosen to maximize $v_F(x;t)$, taking Home's tariff t as given. So, given t, Foreign's best response is summarized by the pair (x^*, π) as

$$x = \begin{cases} x^* & \textit{probability } \pi \\ 0 & \textit{probability } 1 - \pi \end{cases}$$

where x^* is the unilaterally optimal export tax that maximizes $v_F(x;t)$ given t and π denotes the probability that the realized η is sufficiently low such that $v_F(x^*;t) - \eta > v_F(0;t)$. In what follows, we characterize x^* and π and show that both are increasing in the degree of aggregate import dependence.

Foreign's Export Tax in the Noncooperative State. Conditional on imposing an export tax, Foreign chooses the unilaterally optimal tax rate that maximizes it own welfare:

$$x^*(t) \in \arg\max_{x \geq 0} v_F(x; t).$$

The implied Home aggregate import dependence under the noncooperative regime is therefore:

$$\Lambda(t) \equiv \lambda_{FH}(x^*(t); t).$$

It turns out, as we show shortly, that Λ can be used as sufficient statistic that regulates Foreign's policy posture. In particular, Home's choice of t influences Foreign policy through $\Lambda(t)$. Moreover, if $t \mapsto \Lambda(t)$ is one-to-one in the range of tariffs considered here, the exposition could be improved by characterizing policy variables and payoffs explicitly in terms of the Λ coordinate. In the case of the export tax, we can invert this dependence function to obtain $t(\Lambda)$, the supporting Home tariff, and specify the export tax as $x^*(\Lambda) \equiv x^*(t(\Lambda))$. The next lemma characterizes $x^*(\Lambda)$ with the formal derivation presented in Appendix (B).

Lemma 1. *Foreign's optimal noncooperative export tax is*

$$x^*(\Lambda) = \frac{1}{\epsilon(1-\Lambda)\delta(\Lambda) + \frac{1-\beta}{\beta}\Lambda}$$

where $\delta(\Lambda)$ is one plus the share of tariff revenue in Home's aggregate income, with $d\delta/d\Lambda < 0$.

Differentiating the equation for $x^*(\Lambda)$ and invoking Assumption 2 ($\epsilon > \frac{1-\beta}{\beta}$), Appendix (B) shows that Foreign's optimal export tax x^* under the noncooperative regime increases with aggregate import dependence Λ. The following lemma summarizes this point.

Lemma 2. *Foreign's optimal export tax is increasing in aggregate import dependence:*

$$\frac{dx^*(\Lambda)}{d\Lambda} > 0$$

Intuitively, greater aggregate import dependence confers more export market power on Foreign, enabling it to extract a higher revenue-maximizing markup on its exports. Private sourcing decisions take Λ as given and do not internalize this dependence. However, Home's tariff choice t in the first period can contain the Foreign export tax rate by regulating $\Lambda(t)$. The other policy margin that Home's tariff can influence is the ex ante likelihood of Foreign actually adopting x^*, which we examine next.

The export restriction hazard $\pi(\Lambda)$. Next, we characterize the ex ante probability of foreign adopting the noncooperative regime. For a given t, Foreign compares welfare under the noncooperative regime involving the unilaterally-optimal export tax $x^*(\Lambda) > 0$ to the cooperative regime with no export taxation, $x = 0$. Let $t(\Lambda)$ denote the Home tariff that supports the noncooperative regime leading to aggregate import dependence Λ. The gain for Foreign from adopting the noncooperative regime given this tariff is

$$\Delta(\Lambda) \equiv v_F(x^*(\Lambda); t(\Lambda)) - v_F(0; t(\Lambda)). \tag{28}$$

Foreign adopts the noncooperative regime if and only if $\eta \leq \Delta(\Lambda)$. Hence the export-restriction hazard can be written as

$$\pi(\Lambda) = \Pr\{\eta \leq \Delta(\Lambda)\} = G(\Delta(\Lambda)). \tag{29}$$

Thus $\pi(\Lambda)$ provides a direct parameterization of the structural hazard in terms of aggregate import dependence, Λ. Our goal here is to show that the export restriction hazard $\pi(.)$ is also increasing in aggregate import dependence, Λ. To this end, we present an intermediate result proven in Appendix (C). To present this result compactly, we use a subscripted partial derivative operator. For example, $\partial_x v_H(x;t) \equiv \partial v_H/\partial x$ denotes the derivative of v_H with respect to x, and $\partial_{xt} v_H(x;t) = \partial^2 v_H/\partial x \partial t$ denotes the cross derivative.

Lemma 3. *Foreign's welfare function is log sub-modular for $t \leq \frac{\epsilon\beta}{1-\beta} - 1$ and $x \in [0, x^*(t)]$. Specifically,*

$$\partial_{xt} v_F(x;t) < 0, \qquad v_F(x;t) \equiv \ln W_F(x;t). \tag{30}$$

Equivalently, a higher Home import tariff lowers Foreign's marginal payoff from adopting a noncooperative export-tax equivalent.

Lemma 3 states that Home tariffs and the Foreign export taxes are strategic substitutes. A Home tariff depresses Foreign's aggregate market share, lowering its marginal gains from export taxation. This result, therefore, complements Lemma 2, giving a complete picture of the strategic interdependence between t and x: The previous lemma stated that a higher t would lower Λ and, thus, lower Foreign's noncooperative export tax rate x^*. The above lemma additionally states that the higher tariff would also lower the marginal gains from export taxation for Foreign, thereby lowering the likelihood of Foreign adopting the non-cooperative regime. This result is stated next.

Lemma 4. *The gains from adopting the noncooperative regime is increasing in import dependence: $\Delta'(\Lambda) > 0$. Thus, the probability of Foreign adopting the non-cooperative regime and imposing x^* is increasing in import dependence:*

$$\frac{d\pi(\Lambda)}{d\Lambda} = g(\Delta(\Lambda))\Delta'(\Lambda) \geq 0,$$

where g is the density of G with strict inequality whenever $g(\Delta(\Lambda)) > 0$.

The logic behind Lemma 4 can be unpacked as follows. Along the noncooperative equilibrium locus, a higher Λ is supported by a lower Home tariff. Differentiating (28) gives

$$\Delta'(\Lambda) = \underbrace{\partial_x v_F(x^*(\Lambda); t(\Lambda))}_{=0} x^{*\prime}(\Lambda) + [\partial_t v_F(x^*(\Lambda); t(\Lambda)) - \partial_t v_F(0; t(\Lambda))]\, t'(\Lambda).$$

The first term drops out by the envelope theorem, because $x^*(\Lambda)$ is Foreign's optimal export tax that maximizes v_F given t. The bracketed term is negative by Lemma 3. Since higher aggregate import dependence Λ corresponds to a lower Home tariff, $t'(\Lambda) < 0$, then $\Delta'(\Lambda) > 0$. Put differently, higher ex post aggregate import dependence, Λ, raises the payoff from export restriction and the probability that Foreign adopts the noncooperative regime.

Notice that the primitive results that $\pi(\Lambda)$ and $x^*(\Lambda)$ are increasing in Λ hold without any parametric restrictions on the distribution, $G(\eta)$. However, moving forward, we impose a parametric restriction to streamline the optimal policy analysis.[3]

Assumption 3. *The CDF of the cost of noncooperation satisfies $G(\eta) = [\Delta^{-1}(\eta)]^{\zeta}$.*

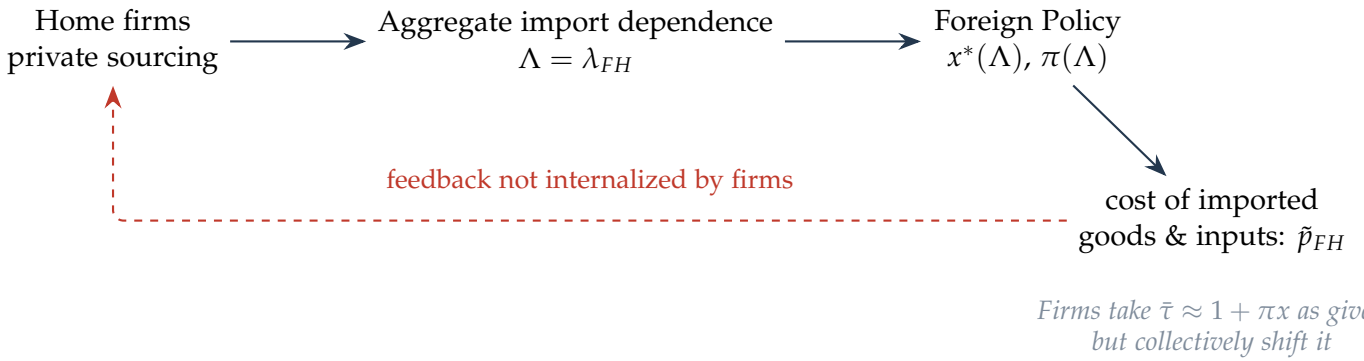
This assumption is purposely strong, as it implicitly ties the distribution G to the underlying tariff rate t. Meanwhile, it buys us a lot expositional simplicity going forward. First, it implies that the ex ante probability of noncooperation is $\pi(\Lambda) = \Lambda^{\zeta}$. Second, this assumption allows us to convert Home's ex ante Bayesian policy problem into a deterministic problem, when the risk elasticity ζ is sufficiently small.

5.2 Strategic Dependence Externality

We have established that the Foreign export policy schedule, (x^*, π), is determined by the economy's aggregate import dependence, $\Lambda = \lambda_{FH}(x^*; t)$. The difficulty is that no

[3]For completeness, we set $G(\eta) = 0$ for $\eta < \Delta(0)$ and $G(\eta) = 1$ for $\eta > \Delta(1)$.

Figure 4: Strategic-dependence externality leads to excessive import dependence



Notes: The diagram traces the causal chain linking individual sourcing decisions to the foreign export policy. Home firms make decentralized sourcing choices (left node), which aggregate into the economy-wide import share $\Lambda \equiv \lambda_{FH}$ (second node). Foreign observes Λ and sets $x^*(\Lambda)$ and $\pi(\Lambda)$ (third node), which determines the buyer prices$\tilde{p}_{FH}$ that Home firms face (right node). The dashed arrow shows the feedback loop: buyer prices shape sourcing decisions, which collectively determine the aggregate import dependence that drives Foreign policy. The externality arises because each firm is infinitesimal and takes $\bar{\tau} \approx 1 + \pi x$ as given when making sourcing decisions, ignoring its own marginal contribution to Λ and hence to Foreign's export incentives. Decentralized sourcing therefore generates excessive import concentration relative to the social optimum.

individual firm, in choosing its sourcing strategy, has reason to internalize the effect of that choice on Λ, and hence on the Foreign policy itself. To see this clearly, consider any firm $\omega \in \Omega_H$. The firm's demand for intermediate inputs and final goods is governed by the CES systems (3) and (10), taking buyer prices $\tilde{p}_{ni}(\cdot)$ as given. Most importantly, the firm treats the export policy functions $x^*(\Lambda)$ and $\pi(\Lambda)$ as fixed. Yet, Λ is an equilibrium share, determined in the aggregate by the collective decisions of all Home firms. This creates a strategic-dependence externality: As individual firms draw more heavily on Foreign suppliers, Λ rises. A higher Λ, in turn, increases the effective export wedge facing all other firms, by raising both the probability π that Foreign adopts a noncooperative export tax and the rate $x^*(\Lambda)$ of that tax, should it be imposed. So individual sourcing imposes a cost, in the form of an elevated risk of export restriction, that is not internalized by private sourcing decisions.

5.3 Home's optimal tariff

Having characterized Foreign's policy response (x^*, π), we next turn to characterizing Home's optimal tariff $t \equiv t_{FH}$ in the initial period of the sequential game. Home's ex ante payoff before Foreign's type is realized is

$$v_H^E(t) \equiv [1 - \pi(\Lambda(t))]v_H(0;t) + \pi(\Lambda(t))v_H(x^*(t);t). \tag{31}$$

The ex ante optimal policy problem on the tariff domain $\mathbb{T} = [0, \bar{t}]$ can be presented as

$$\max_{t \in \mathbb{T}} \; v_H^E(t) \qquad s.t. \quad \text{equilibrium constraints} \qquad [\text{P1}]$$

We tackle this problem in its Bayesian form and present the implied optimal tariff formula in Appendix F. However, we take a different route in this section to better highlight our strategic-risk mechanism. We show that P1 can alternatively be solved as a deterministic optimal tariff problem in which Foreign's deterministic export tax is increasing in Home's aggregate import share. To make this point, define the certainty-equivalent export tax as the deterministic tax rate $\bar{x}(t)$ that reproduces Home's expected payoff:

$$v_H(\bar{x}(t); t) = v_H^E(t). \tag{32}$$

Namely, $\bar{x}(t)$ is the export tax rate that leaves Home indifferent, at tariff t, to the expected welfare over aligned and noncooperative regimes. Since Home welfare is strictly decreasing in Foreign's export tax, (32) defines a unique $\bar{x}(t) \in [0, x^*(t)]$. The deterministic aggregate import dependence implied by this tax is denoted by

$$\bar{\Lambda}(t) \equiv \lambda_{FH}(\bar{x}(t); t).$$

The next lemma identifies the condition under which the Bayesian Stackelberg problem has a certainty-equivalent representation.[4]

Lemma 5. *If ζ is sufficiently small, then the certainty-equivalent export tax can be written as an increasing function of deterministic aggregate import dependence. There is a unique schedule $\bar{\tau}$ on $\bar{\Lambda}(\mathbb{T})$ such that*

$$\bar{\tau}(\bar{\Lambda}(t)) = 1 + \bar{x}(t), \qquad \bar{\tau}'(\bar{\Lambda}(t)) > 0. \tag{33}$$

Under this schedule, the deterministic economy with $1 + x = \bar{\tau}(\lambda_{FH})$ delivers the ex ante Home welfare $v_H^E(t)$ for each admissible t, thereby leaving Home's optimal tariff unchanged.

The proof is in Appendix D. The logic is straightforward. When Home raises its import tariff, CES demand lowers Home's aggregate import dependence. Lower aggregate import dependence reduces both Foreign's optimal export tax and the probability that Foreign finds noncooperation worthwhile. The certainty-equivalent export tax wedge therefore falls with Home's tariff. Since deterministic aggregate import dependence also

[4] As in Lemma 3, the tariff cap satisfies $(1 + \bar{t})\frac{1-\beta}{\beta} \leq \epsilon$. For $\beta = 1$, this condition is automatic and $\bar{t}$ may be any finite cap.

falls with Home's tariff, the deterministic export tax can be written as an increasing function of aggregate import dependence.

Lemma 5 is a payoff-equivalence result. It does not imply that the deterministic export tax reproduces every state-contingent allocation or Foreign's payoff. It implies only that the deterministic equilibrium gives Home the same payoff at every admissible tariff. Hence Home's ex ante problem can be written as the following deterministic optimal-policy problem:

$$\max_{t\in\mathbb{T}} \; v_H(x;t) \qquad s.t. \begin{cases} \text{equilibrium constraints} \\ x = \bar{\tau}(\lambda_{FH}) - 1 \end{cases} \qquad \text{[P2]}$$

Here, λ_{FH} is the deterministic equilibrium share obtained by solving the equilibrium system under x and t. The strategic-dependence externality described above takes an even simpler form here: Private Home firms take $\bar{\tau}(\lambda_{FH})$ as given when making sourcing decisions. But their decisions collectively determine λ_{FH} and shift the Foreign export policy rate. The Home government, by contrast, internalizes that increased λ_{FH} shifts Foreign's certainty-equivalent export-tax wedge, and adjusts the optimal tariff accordingly. The following proposition formalizes this point.

Proposition 2. *When Foreign is inclined to strategically use export taxes, Home's unilaterally optimal tariff becomes*

$$t^* \;=\; \frac{1}{\varepsilon^S_{FH}} \;+\; \underbrace{\frac{\epsilon\lambda_{HH}}{\varepsilon^D_{FH}} \frac{d\ln\bar{\tau}(\lambda_{FH})}{d\ln\lambda_{FH}}}_{\textit{strategic-risk correction}}.$$

where the export supply elasticity is $\varepsilon^S_{FH} = \frac{\epsilon\lambda_{FF}}{1-\frac{\bar{x}}{1+\bar{x}}\lambda_{HF}} + \frac{1-\beta}{\beta}\lambda_{HF}$ *and the import demand elasticity is given by Equation 26.*

The first term in Proposition 3 represents the standard terms-of-trade motive that also appeared under Proposition 1. The second term is new and addresses elevated strategic risk due to sourcing externalities. Atomistic firms make sourcing decisions *as if* $\mathrm{d}\ln\bar{\tau}/\mathrm{d}\ln\lambda_{FH} = 0$, thereby causing excess sourcing from Foreign and higher export taxes. A higher Home tariff corrects this by reducing λ_{FH}. And because $\bar{\tau}'(\lambda_{FH}) > 0$, this lowers the certainty-equivalent export tax facing Home firms.

Before concluding this section, we present a useful approximation that will be used later for quantification. For this we apply Taylor's theorem to equation (32), which traces the certainty equivalence schedule $\bar{\tau}$. Following the derivation presented in Appendix

D, we can write $\bar{\tau}$ as

$$\bar{\tau}(\lambda_{FH}) = 1 + \pi(\lambda_{FH})x^*(\lambda_{FH}) + O(x^{*2}), \tag{34}$$

where $x^* \equiv x^*(\lambda_{FH})$ is Foreign's optimal export tax implied by the deterministic share, λ_{FH}. This first-order formula is useful for simulating counterfactuals because it expresses the certainty-equivalent export tax factor directly in terms of the two primitive functions characterized earlier: the noncooperative export tax function $x^*(.)$ and the export restriction hazard $\pi(.)$.

5.4 Multiple Sectors

In this section we extend our optimal tariff framework to a multi-sector economy. Suppose firms can be partitioned into multiple sectors denoted by $k = 1, ..., K$. Suppose that the final good and intermediate input aggregator use a common Cobb-Douglas functional form: $\tilde{P}_i = \prod_k \tilde{P}_{i,k}^{\alpha_k}$ where α_k is the share of expenditure on sector k goods. The sector-level price index is

$$\tilde{P}_{i,k} = \left[\sum_{n \in \mathcal{C}} \frac{1}{\bar{\varphi}_{n,k}} \left(\tau_{ni,k}^M \, \tau_{ni,k}^X \, d_{ni,k} \, w_n^\beta \, \tilde{P}_n^{1-\beta} \right)^{-\epsilon_k} \right]^{-\frac{1}{\epsilon_k}},$$

where ϵ_k is the sector-level trade elasticity, with the corresponding trade share given by

$$\lambda_{ni,k} = \frac{\left(\frac{1}{\bar{\varphi}_{n,k}} \tau_{ni,k}^M \, \tau_{ni,k}^X \, d_{ni,k} \, w_n^\beta \, \tilde{P}_n^{1-\beta} \right)^{-\epsilon_k}}{\sum_{\ell \in \mathcal{C}} \left(\frac{1}{\bar{\varphi}_{\ell,k}} \tau_{\ell i,k}^M \, \tau_{\ell i,k}^X \, d_{\ell i,k} \, w_\ell^\beta \, \tilde{P}_\ell^{1-\beta} \right)^{-\epsilon_k}}.$$

We maintain an analog of Assumption 2 for the multi-sector framework; namely that $\epsilon_k > (1-\beta)/\beta$. The next lemma shows that under this assumption, Foreign's optimal export tax in each sector is increasing in the sector-level degree of import dependence. A formal proof for the lemma is provided in Appendix H.[5]

Lemma 6. *Given Home's sector-level import tariffs $\{t_k\}$, Foreign's optimal noncooperative ex-*

[5] Lemma 6 presents x_k^* for the case where Home's tariff revenue has no first-order effect on Foreign welfare at the optimum. Without this restriction, the formula becomes $x_k^* = \left(\epsilon_k (1 - \Lambda_k) \delta_k + \frac{1-\beta}{\beta} \Lambda_k \tilde{\delta}_k \right)^{-1}$, featuring an additional shifter δ_k, which is the sector-level analog of δ in Lemma 1. See Appendix H for details.

port tax on sector k is

$$x_k^*(\Lambda_k) = \frac{1}{\epsilon_k(1-\Lambda_k) + \frac{1-\beta}{\beta}\Lambda_k\,\tilde{\delta}_k}$$

where $\Lambda_k \equiv \lambda^*_{FH,k}$ *is Home's aggregate import dependence in sector k and* $\tilde{\delta}(\Lambda_k) \equiv \sum_g \frac{1+\epsilon_k(1-\Lambda_k)}{1+\epsilon_g(1-\Lambda_g)}\frac{\Lambda_g}{\Lambda_k}\alpha_g$.

The above lemma states that Foreign's optimal non-cooperative export tax in sector k increases with import dependence in that sector. If noncooperation is costless for Foreign, then Foreign will apply this tax with ex ante certainty. If noncooperation is costly, Appendix H shows that the ex ante probability of adopting x_k^* is increasing in Λ_k. Putting the pieces together, we model Foreign's certainty-equivalent export tax analogously to the single-sector benchmark. Specifically, from Home's perspective, the ex ante Foreign export tax response is modeled as

$$\bar{x}_k = \bar{\tau}_k(\lambda_{FH,k}) - 1,$$

where $\bar{\tau}_k(.)$ is an increasing function and $\{\lambda_{FH,k}\}$ is the sector-level import share schedule implied by $\{\bar{x}_k\}$. This representation tractably captures the core mechanism of the paper—that greater import dependence in sector k translates to a higher certainty-equivalent export tax in that sector. Moreover, because firms are infinitesimal, they do not internalize their collective influence on Foreign policy, $d\bar{\tau}_k/d\lambda_{FH,k}$, leading to excessive sourcing from Foreign within each sector. The optimal tariff seeks to correct the excessive exposure implied by private sourcing decisions.

Optimal sector-level tariff. Home's optimal tariff schedule is a vector $\{t_k\}$ that maximizes Home's welfare $W_H = Y_H/\tilde{P}_H$ subject to equilibrium constraints and the Foreign's certainty-equivalent export-tax wedge $\{\bar{x}_k\} = \{\bar{\tau}_k(\lambda_{FH,k}) - 1\}$. The next proposition characterizes the optimal tariff schedule, with a formal proof provided in Appendix G.

Proposition 3. *Suppose Foreign is inclined to strategically restrict exports on a sector-level basis. Then, Home's unilaterally import tariff on sector k is*

$$t_k^* = \frac{1}{\varepsilon^S_{FH}} + \frac{\epsilon_k(1-\lambda_{FH,k})}{\varepsilon^D_{FH,k}}\frac{d\ln\bar{\tau}_k(\lambda_{FH,k})}{d\ln\lambda_{FH,k}}. \tag{35}$$

where $\varepsilon^D_{FH,k}$ *is the sector-level import demand elasticity, and terms of trade (ToT) component of*

the tariff is uniform and determined by the aggregate export supply elasticity:

$$\varepsilon_{FH}^{S} = \sum_{k} \left[\frac{\lambda_{HF,k}\alpha_k}{\lambda_{HF}} \epsilon_k \lambda_{FF,k} \right] + \frac{1-\beta}{\beta} \lambda_{HF}.$$

This proposition provides a rationale for targeted tariffs that vary with sector-level import dependence. As noted earlier, recent US tariffs exhibit this form of targeting, which is difficult to reconcile with traditional optimal tariff theories. Case in point, the first term in Equation 35, which represents the standard terms of trade motive is non-targeted and uniform across sectors.[6] The second term, which corrects the strategic-dependence externality identified in this paper, is targeted and varies across sectors. The magnitude depends on the policy response, $d\bar{\tau}_k / d\lambda_{FH,k}$, not internalized by domestic firms' sourcing decisions. Intuitively, this term is larger in sectors with significant degrees of aggregate import dependence, and our quantitative exercise that follows confirms this.

6 Mapping Model to Data

This section maps our model to data. Our goal is to compute the unconstrained-optimal tariff schedule for the United States and the European Union as the two running examples for the home economy. We begin with describing the data sources and discuss how the export restriction hazard is empirically estimated. Then, we describe our calibration strategy and numerical algorithm for simulating the optimal tariff counterfactual.

6.1 Baseline Macroeconomic Data

We draw on the 2016–2022 Extended Edition of the OECD Inter-Country Input-Output (ICIO) database (OECD, 2021). The data cover 76 individual economies together with an aggregate of the Rest of World aggregate, and 50 ISIC Rev. 4 industries per country block. Each annual table is organized as a square flow matrix. The inter-industry block records gross intermediate purchases from every country-industry origin to every country-industry destination, yielding a transaction matrix Z of dimension $(85 \times 50) \times (85 \times 50)$. A final demand block records expenditure across six final good categories

[6] The exact uniformity of terms-of-trade optimal tariffs is a direct implication of the Ricardian structure whereby labor is perfectly mobile across sectors. But even with fixed sector-specific factors, it is often difficult to generate meaningful sectoral heterogeneity in optimal tariffs based on terms of trade motives alone. The variation would only be sizable if Home accounts for a significant share of global demand in a given sector, which in not true in practice.

for each country. We consolidate the industries into sixteen broad sectors: agriculture, mining, food processing, textiles, wood products, paper and printing, petroleum, chemicals, plastics and rubber, metals and fabricated metals, non-metallic minerals, electronics, machinery, transport equipment, other manufacturing, and services. These data points allow us to construct the baseline equilibrium values for 2018. For each choice of home economy $H \in \{\text{US}, \text{EU}\}$, we calibrate the import share based on the matrix of gross bilateral flows $X_{ni,k}$ in the non-service sectors. The sector-level and aggregate shares are more formally computed as

$$\lambda_{FH,k} = \frac{\sum_{n \neq H} X_{nH,k}}{\sum_n X_{nH,k}}, \qquad \lambda_{FH} = \frac{\sum_k \sum_{n \neq H} X_{nH,k}}{\sum_k \sum_n X_{nH,k}}$$

Table 1 reports the resulting baseline values for both home economies. Once intra-EU flows are excluded from the import measure, the two economies display comparable Foreign aggregate import shares: λ_{FH} equals 19.9 percent for the United States and 14.3 percent for the European Union. We calibrate the structural parameters β and α_k using the same ICIO tables. The value-added share β is calibrated separately for each home economy from the ICIO table as the ratio of aggregate value added to gross output, yielding $\beta = 0.55$ for the United States and $\beta = 0.48$ for the European Union. The Cobb-Douglas expenditure shares are calibrated to sector level expenditure divided by aggregate expenditure: $\alpha_k = \sum_n X_{ni,k}/E_i$. Lastly, the trade elasticity to $\epsilon = 4$, a value in the range commonly used in the quantitative trade literature (Simonovska and Waugh, 2014) For the macro-level analysis, we aggregate each economy into traded and non-traded sectors. The non-traded share of the economy is calibrated based on the share of expenditure on services.[7]

6.2 Estimating the Export Restriction Risk Elasticity

A key element of our model is the elasticity by which the risk of Foreign export restrictions increases with aggregate import dependence at the sector level. Specifically, the probability that sector k is subjected to an export restriction is parameterized as $\pi_k(\Lambda_k) = \Lambda_k^{\zeta}$, with ζ representing the elasticity of noncooperation risk with respect to import dependence. This normalization is the structural probability used in the counterfactual analysis; level differences across importer-sector cells are absorbed by fixed effects in estimation. A higher ζ corresponds to a steeper increase in disruption risk as

[7]The approach calibrates the weight of the non-traded sector to 79.71 percent for the United States and 73.19 percent for the European Union.

Table 1: Baseline Macroeconomic Data and Parameters (2018)

	United States	European Union	Source
Panel A. Macroeconomic Aggregates			
GDP (Trillions USD)	19.88	14.29	OECD ICIO
Imports / GDP (%)	14.32	19.57	
Services / GDP (%)	79.71	73.19	
Import Share $\lambda_{FH,T}$ (%)	35.71	37.42	
Panel B. Structural Parameters			
Trade Elasticity ϵ	4.00		Simonovska–Waugh
Value-added Share β	0.55	0.48	OECD ICIO
Certainty-equivalent fit $\hat{a}$	0.76	0.81	Model calibration

Notes: All macroeconomic aggregates are drawn from the 2018 OECD Inter-Country Input-Output (ICIO) tables. GDP is measured as aggregate value added; observed import-to-GDP ratios retain the actual Foreign-to-Home flow and treat intra-EU transactions as domestic for the European Union. The non-traded weights proxy latent expenditure shares with ISIC G–T service value-added shares; they do not classify measured service trade as non-traded. All observed cross-border deliveries enter the traded composite. Balanced aggregate and conditional traded shares are obtained from the zero-transfer exact-hat benchmark. The value-added shares β_H and β_F are economy-wide ratios of value added to gross output.

the home economy becomes more dependent on foreign supply.

Data Sources. We estimate ζ using a newly assembled data set on export restrictions matched to bilateral import flows from BACI for the 2017–2023 period.

Our main source on export restrictions is the Global Trade Alert (GTA) database (Evenett and Fritz, 2024), which provides a public record of state interventions affecting trade and investment. We use the subset pertaining to export restrictions and retain seven categories of trade restrictive measures: export taxes, export bans, export quotas, export licensing requirements, export price benchmarks, export tariff quotas, and local supply requirement for exports.[8] We further restrict the sample to interventions classified by GTA as Red or Amber, which indicate that the measure is discriminatory or likely discriminatory. The raw extract contains 4,157 export-related interventions across 146 economies. Restricting the sample to measures active during 2017–2023 yields 1,866 policy lines. We construct two measures for export restrictions: a flow measure that captures newly implemented restrictions and a stock measure that records measures in force in a given year. The resulting sample covers 126 exporters and 5,692 HS6 products in the flow panel, and 138 exporters and 5,709 HS6 products in the stock panel.

We supplement the GTA data with two OECD data sources: the Inventory of Export

[8]We exclude the residual GTA category "Export-related non-tariff measure, nes."

Restrictions on Industrial Raw Materials and the Database on Export Restrictions on Staple Crops. The industrial raw materials inventory tracks restrictions on 65 categories of industrial commodities and has been compiled since 2009. The staple-crops database, maintained by the OECD as a contribution to the Agricultural Market Information System, records six restrictive instruments applied to wheat, maize, rice, soybeans, and their combinations since 2007. In our harmonized 2017–2023 sample, the industrial inventory covers 39 exporters and 260 HS6 products in the flow panel, and 60 exporters and 485 HS6 products. The staple-crops database features 11 exporters and 34 HS6 products. Taken together, the two OECD sources contribute 1,979 flow observations and 22,059 stock observations.

Hazard and Import Dependence Measures. For each home economy i, exporter n, HS2 sector k, product $\kappa \in \mathcal{K}_k$, and year t, let

$$R_{n,\kappa,t} = \mathbf{1}\{\text{exporter } n \text{ restricts product } \kappa \text{ in year } t\}. \tag{36}$$

We calculate the empirical counterpart of the export restriction hazard as the weighted sum of applied restrictions

$$\pi_{ni,k,t} = \sum_{\kappa \in \mathcal{K}_k} \omega_{ni,\kappa,t} R_{n,\kappa,t}, \tag{37}$$

where the weight assigned to product κ within the importer-sector cell is $\omega_{ni,\kappa,t} = X_{ni,\kappa,t} / \sum_{\kappa' \in \mathcal{K}_k} X_{ni,\kappa',t}$, with X denoting trade values. This construction aggregates product-level policy events into sector-level bilateral restriction hazards. The logic is that under the law of large numbers the probability of imposing restrictions on a product line within sector k overlaps with the share of product lines that receive the restriction. Our measure of aggregate import dependence is the lagged market share of a given exporter within importer-sector cells. Namely, $\tilde{\lambda}_{ni,k,t-1} = \frac{X_{ni,k,t-1}}{\sum_{\ell} X_{\ell i,k,t-1}}$, where $X_{ni,k,t-1}$ denotes exporter n's exports to home i in sector k in year $t-1$. We use lagged trade values to circumvent the reverse causality between import values and export restrictions.

Estimating Equation. We estimate the export-restriction hazard elasticity ζ, under both the stock and flow hazard encodings. The stock hazard is our baseline because the model is disciplined by the probability that export restrictions are active when imports are realized, not merely by the arrival of new policy actions. The flow hazard provides

an event-style robustness check. For each encoding, we estimate the following equation:

$$\ln \pi_{ni,k,t} = \zeta \ln \tilde{\lambda}_{ni,k,t-1} + \Gamma_{ik} + \Gamma_t + u_{ni,k,t}, \tag{38}$$

where Γ_{ik} denotes home country-sector fixed effects and Γ_t denotes year fixed effects. We estimate equation (38) using OLS and Poisson pseudo-maximum likelihood (PPML) estimators. PPML can naturally handle the zero entries. When running the OLS specification, we assume a baseline risk set to the minimum inferred risk from existing measures in a given year.[9]

Results. Table 2 reports estimates under a range of specifications. Panel A presents the baseline stock-hazard specification. The pooled OLS estimate is $\hat{\zeta}$=0.265, whereas the pooled PPML estimate is 0.178. Standard errors are clustered at the sector level. The OLS estimates are generally larger than the PPML estimates, but the two estimators yield quantitatively similar results.

We also estimate the model separately for U.S. and E.U. samples. In each case, the region in question is treated as the home country, and the estimating equation identifies the elasticity of the hazard with respect to aggregate import dependence across trading partners. The resulting estimates are similar across specifications and are statistically significant across the board, with implied hazard elasticities in the 0.1 to 0.3 range.

Taken together, these findings support the view that greater aggregate import dependence increases the gains from export restrictions and thereby strengthens the incentive for foreign governments to impose them. Our estimates imply that exporters accounting for a larger share of the home country's imports are systematically more likely to adopt restrictive trade measures.

7 Computing the EU and US Optimal Tariffs

This section uses our calibrated model to compute the optimal tariff schedule for the US and the EU. We first describe the numerical algorithm for recovering optimal tariffs. We then compute the optimal macro tariff, which is applied to the traded composite while the non-traded wedge remains one. Lastly, we present the optimal sector-level tariffs and demonstrate that strategic de-risking motives generate significant variation in optimal tariffs across sectors based on import dependence.

[9]For the European Union, we aggregate imports across all 27 EU member states.

Table 2: Estimates of the Strategic-Risk Elasticity

	PPML			OLS		
	US (1)	EU (2)	Pooled (3)	US (4)	EU (5)	Pooled (6)
Panel A. Stock hazard						
ζ	0.146^{***}	0.218^{***}	0.177^{***}	0.287^{***}	0.230^{***}	0.267^{***}
	(0.010)	(0.014)	(0.010)	(0.014)	(0.014)	(0.013)
N	63,860	75,280	139,140	63,860	75,280	139,140
Panel B. Flow hazard						
ζ	0.157^{***}	0.292^{***}	0.204^{***}	0.174^{***}	0.121^{***}	0.151^{***}
	(0.012)	(0.013)	(0.010)	(0.011)	(0.006)	(0.008)
N	63,860	75,280	139,140	63,860	75,280	139,140

Notes: The dependent variable is the export-weighted restriction hazard at the HS2 sector level. The regressor is the log lagged exporter share in the importer's sector imports. Data span 2018 to 2023. Panel A reports the stock-based hazard encoding (baseline), and Panel B reports the flow-based hazard encoding (robustness). PPML estimates use Poisson pseudo-maximum likelihood; OLS estimates use log-transformed hazard with zeros winsorized to the year-specific minimum positive value. Fixed effects: HS2 sector and year (single-home), plus home (pooled). Standard errors clustered by HS2 sector in parentheses. The quantitative baseline uses the Panel A OLS estimates.

7.1 Quantitative Algorithm

his section outlines the quantitative method for recovering optimal tariffs. We focus the presentation on the one-sector model, since the same logic extends to the multi-sector case. We take a factual baseline equilibrium characterized by free trade ($\tau^M_{ni} = \tau^X_{ni} = 1$) and interpret Home as the United States or the European Union and Foreign as the rest of the world. We calibrate the baseline equilibrium variables to their observed counterparts in the data, as detailed earlier. We then perturb the tariff applied by Home t, update the export tax applied in Foreign $x(t)$ to reflect its best response, and leave all other taxes unchanged. Then, we me compute the counterfactual equilibrium in changes. For each variable z let z' denote the counterfactual value after the perturbation, and let $\hat{z} = z'/z$ denote the resulting change in the variable. The change in trade taxes is

$$\hat{\tau}^M_{ni}(t) = \begin{cases} 1+t & ni = FH \\ 1 & ni \neq FH \end{cases} \qquad \hat{\tau}^X_{ni}(t) = \begin{cases} 1+x(t) & ni = FH \\ 1 & ni \neq FH \end{cases}$$

For a given the Home tariff t, we set $x(t)$ to Foreign's certainty equivalence optimal export tax. As detailed earlier, the certainty equivalent optimal export tax is approximated

by $x(t) = a \cdot \pi\left(\lambda'_{FH}(t)\right) x^*\left(\lambda'_{FH}(t)\right)$, where $\lambda'_{FH}(t)$ denotes the counterfactual share under t, and $\pi(\Lambda) = \Lambda^{\zeta}$ and $x^*(\Lambda) = \left[\epsilon(1-\Lambda) + \frac{1-\beta}{\beta}\Lambda\right]^{-1}$, following Lemma 1. Thus, altogether, we can specify the Foreign's optimal export tax response as

$$x(t) = \underbrace{\frac{\hat{a}\left(\lambda'_{FH}(t)\right)^{\zeta}}{\epsilon(1-\lambda'_{FH}(t)) + \frac{1-\beta}{\beta}\lambda'_{FH}(t)}}_{\bar{\tau}(\lambda'_{FH}(t))-1}, \qquad \lambda'_{FH} = \lambda_{FH}\hat{\lambda}_{FH}(t) \tag{39}$$

We calibrate $\hat{a}$ to maximizes the goodness of the approximation. In particular, we construct of mesh with $\mathcal{J} = 121$ nodes, and calibrate $\hat{a}$ to minimize the leas squared error between the exact certainty equivalence export tax rate and the rate approximates by our closed-form formula:

$$\hat{a} = \arg\min_{a} \sum_{j=1}^{\mathcal{J}} \left[\bar{x}_j - a\,\pi(\bar{\lambda}_{FH,T,j}) x^*(\bar{\lambda}_{FH,T,j})\right]^2, \qquad 1 + t_j \in [1, 1.875].$$

The resulting coefficients are 0.76 for the United States and 0.81 for the European Union. The traditional terms of trade benchmark in which Foreign is passive, sets $x = 0$.

For given policy wedges, we solve the equilibrium in exact changes. Let $\alpha_{i,s}$ denote region i's expenditure weight and let P_i denote its aggregate input price index. The central price and share equations are

$$\widehat{p}_{i,s} = \widehat{w}_i^{\beta_i} \widehat{P}_i^{1-\beta_i}, \tag{40}$$

$$\widehat{P}_{i,s} = \left[\sum_n \lambda_{ni,s} \left(\widehat{\tau}^M_{ni,s} \widehat{\tau}^X_{ni,s} \widehat{p}_{n,s}\right)^{-\epsilon}\right]^{-1/\epsilon}, \tag{41}$$

$$\widehat{P}_i = \prod_{s \in \{N,T\}} \widehat{P}_{i,s}^{\alpha_{i,s}}, \tag{42}$$

$$\widehat{\lambda}_{ni,s} = \frac{\lambda_{ni,s}\left(\widehat{\tau}^M_{ni,s} \widehat{\tau}^X_{ni,s} \widehat{p}_{n,s}\right)^{-\epsilon}}{\sum_\ell \lambda_{\ell i,s}\left(\widehat{\tau}^M_{\ell i,s} \widehat{\tau}^X_{\ell i,s} \widehat{p}_{\ell,s}\right)^{-\epsilon}}. \tag{43}$$

Wage and expenditure market clearing incorporates intermediate demand and rebates tariff and export-tax revenue to the collecting region. The full equilibrium retains β_H and β_F in their respective production, income, and expenditure equations. This distinction affects general-equilibrium incidence even though the export-tax shortcut above depends on β_H.

We choose the Home tariff to maximize real income:

$$\max_{t} \quad \ln\left(\frac{\widehat{Y}_H(t)}{\widehat{P}_H(t)}\right) \qquad \text{subject to the equilibrium system,} \qquad 1 + x(t) = \bar{\tau}\left(\lambda'_{FH,T}(t)\right).$$

The sixteen-sector exercise applies the same logic to a vector of sectoral tariffs. Its sectors are observed ICIO aggregates rather than the macro model's latent composites. Consequently, measured service trade remains tariff-active in that exercise and should not be identified with the latent non-traded composite. We use each region's stock-OLS estimate of ζ and transport its aggregate $\hat{a}$ as a common approximation coefficient across sectors. Appendix I separately verifies the exact certainty-equivalent representation and measures the numerical error introduced by this fitted shortcut.

7.2 Quantitative Results

This section presents quantitative results from our optimal tariff simulation. We first characterize optimal macro tariffs, applied to the traded composite with a unit wedge on the non-traded composite, comparing the rates under the strategic sourcing externality with those in a traditional case where optimal tariffs are driven solely by terms-of-trade motives. We then examine optimal sector-level tariffs to highlight how the strategic sourcing externality creates a clear rationale for targeting. While terms-of-trade considerations call for a uniform tariff across sectors, the risk correction that addresses the strategic sourcing externality requires targeted tariffs, generating significant cross-sector heterogeneity in optimal rates similar to those observed in 2018.

Optimal Macro Tariffs. The optimal macro tariff is the rate applied to all imports represented by the traded composite. Table 3 reports the optimal macro tariffs for the US and EU. Column (1) reports the terms-of-trade optimal tariff when the Foreign government is passive ($x = 0$) and there is no scope for strategic-risk correction; this tariff corresponds to the one characterized in Proposition 1. The remaining columns present results when Home faces a strategic Foreign government, which creates a strategic sourcing externality that calls for correction. In this case, the optimal tariff corresponds to the formula in Proposition 2 and exhibits an additional strategic-risk correction term that addresses the strategic sourcing externality.

The terms-of-trade benchmark generates an optimal traded-composite tariff of 25.78 percent for the United States and 26.02 percent for the European Union. Their similarity reflects the common trade elasticity, while the small difference arises from the region-

Table 3: Optimal Macro Tariffs under Strategic Risk

	Benchmark: $\pi = 0$	Geopolitical risk: $\pi(\lambda_T)$		
	Tariff Rate (%)	Tariff Rate (%)	ToT Share	Risk Share
United States	25.78	31.01	83.14	16.86
European Union	26.02	30.93	84.12	15.88

The Benchmark column reports the optimal tariff in the absence of strategic risk ($\pi = 0$). The strategic-risk columns use $\pi(\lambda_T) = \lambda_T^{\zeta}$ and the fitted certainty-equivalent export wedge $\hat{a}\pi(\lambda_T)x^*(\lambda_T)$. The Decomposition columns report the share of the endogenous tariff attributable to terms-of-trade manipulation versus the strategic-risk component; these shares sum to 100 percent.

specific value-added shares and balanced trade structure.

Strategic risk raises the United States tariff to 31.01 percent and the European Union tariff to 30.93 percent. The corresponding corrections are 5.23 and 4.91 percentage points. Thus, the risk channel remains economically meaningful after using direction-specific bilateral flows, region-specific production shares, and a balanced traded-composite calibration.

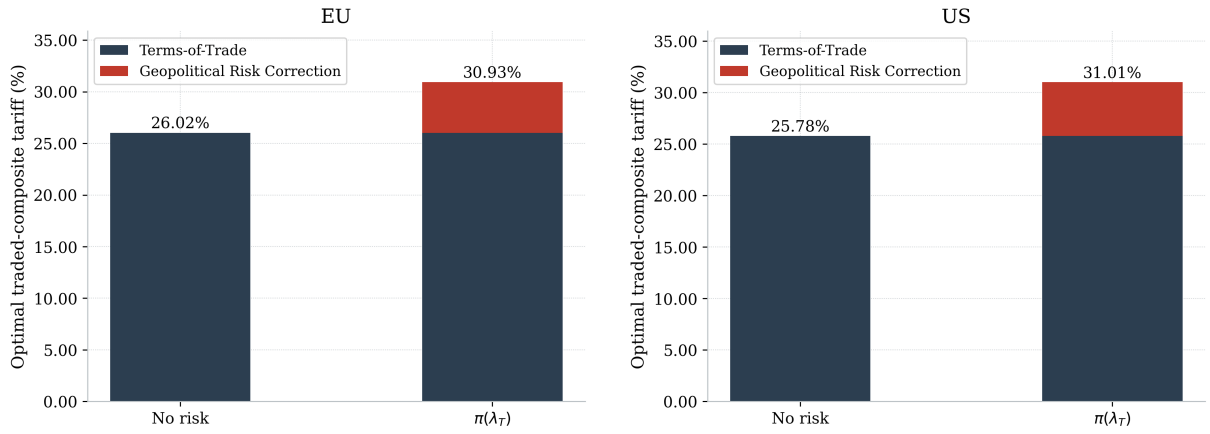
The model delivers a clean decomposition of the optimal tariff into the terms-of-trade component and the strategic-risk component (see Figure 5). For the United States, the terms-of-trade motive accounts for 83.14 percent of the optimal tariff, with the strategic-risk component contributing 16.86 percent. For the European Union, the corresponding shares are 84.12 and 15.88 percent. The two decompositions are similar because two offsetting forces are at work: the European Union exhibits greater aggregate import dependence, while the United States faces a higher estimated hazard elasticity.

These magnitudes are striking because they apply to a uniform macro tariff, not targeted sectoral tariffs. While skepticism often surrounds the political feasibility of high sectoral tariffs, our results suggest that even under the constraint of a uniform tariff policy, the structural incentives to hedge against strategic risk are powerful enough to justify roughly five percentage points of additional protection.

Optimal Sectoral Tariffs. We extend the analysis to sixteen observed ICIO sectors to study how strategic risk shapes the cross-sectoral distribution of optimal tariffs. This exercise differs conceptually from the macro aggregation. Its service sector records measured service production and trade; it is not the latent non-traded composite used to calibrate the macro expenditure share. All sixteen observed sectors, including services, are therefore tariff-active.

Sectoral import dependence varies widely in the 2018 data: Foreign's expenditure

Figure 5: Optimal Traded-Composite Tariff Decomposition



Notes: The figure decomposes the optimal traded-composite tariff into the terms-of-trade (ToT) component and the strategic-risk component. The endogenous-risk scenario uses $\pi(\lambda_T) = \lambda_T^{\zeta}$ and the fitted certainty-equivalent export wedge $\hat{a}\pi(\lambda_T)x^*(\lambda_T)$.

share ranges from 2.75 percent for United States services to 71.39 percent for United States textiles, and from 5.60 percent for European Union paper products to 80.19 percent for European Union mining. We maintain a common trade elasticity $\epsilon = 4.00$ and repeat each region's economy-wide value-added share across sectors. Noncooperation risk is $\pi_k(\lambda_{FH,k}) = \lambda_{FH,k}^{\zeta}$, using the regional stock-OLS estimates $\zeta_{US} = 0.29$ and $\zeta_{EU} = 0.23$. We transport the aggregate certainty-equivalent coefficients, $\hat{a}_{US} = 0.76$ and $\hat{a}_{EU} = 0.81$, uniformly across sectors. This is a tractability assumption rather than a sector-by-sector re-estimation of the approximation.

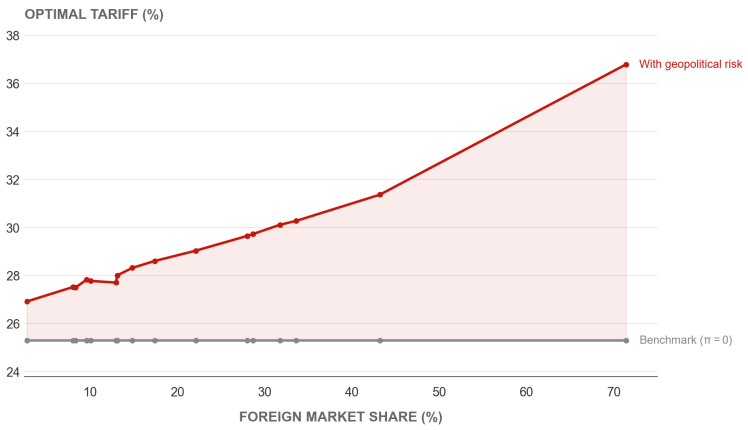
Absent strategic risk ($\pi = 0$), the common trade elasticity yields a uniform terms-of-trade tariff of 25.29 percent in every United States sector and 25.49 percent in every European Union sector. Under endogenous risk, the sectoral schedule becomes steeply increasing in import dependence. United States tariffs range from 26.92 percent in services to 36.79 percent in textiles. European Union tariffs range from 27.22 percent in wood products to 38.35 percent in mining; the tariff on services is 27.54 percent.

Figure 6 plots the two schedules. The no-risk schedule is flat, whereas the endogenous-risk schedule assigns the largest corrections to the most import-dependent sectors. The service observations remain near the lower end because measured service import shares are small, not because services are removed from the policy domain.

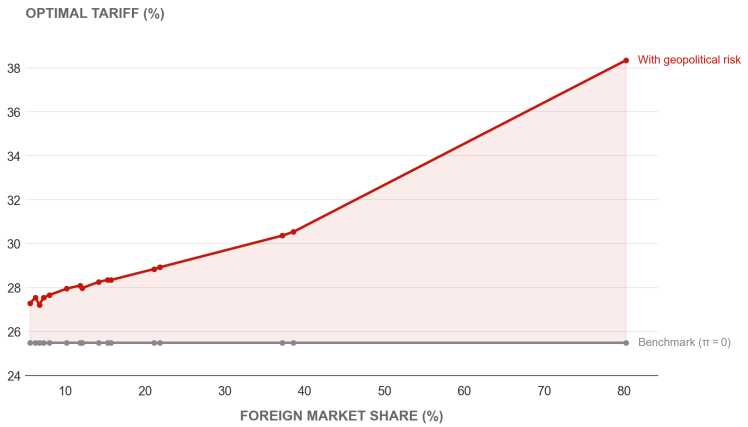
Strategic risk generates *targeted* protection. Unlike the terms-of-trade motive, which produces nearly uniform tariffs under a common trade elasticity, the strategic-risk component is largest precisely where aggregate import dependence is greatest. These are also the sectors where comparative advantage is most pronounced and the gains from

Figure 6: Optimal Sectoral Tariffs and Import Dependence

(a) United States



(b) European Union



Notes: Each point represents one of sixteen observed ICIO sectors, including measured services, ranked by Foreign's sectoral expenditure share. Solid line: endogenous strategic risk with $\pi_k(\lambda_{FH,k}) = \lambda_{FH,k}^{\zeta}$ and a common regional $\hat{a}$. Dashed line: terms-of-trade benchmark ($\pi = 0$).

trade are largest, so the efficiency cost of strategic de-risking falls disproportionately on the economy's most productive import relationships.

8 Empirical Evidence

In this section, we assess the model's empirical relevance by asking whether it can account for the 2018 shift in U.S. tariff policy. We do this in two steps. First, we examine USTR communications in the build up to the 2018 tariff war. The concerns expressed in these documents closely mimic the strategic-sourcing externality at the core of the model. Second, we study the pattern of the 2018 tariff hikes on Chinese imports. We argue that

these tariffs were not systematically related to traditional factors like the export-supply or import-demand elasticities, which underly the terms of trade and political economy theories. Instead, they were targeted towards products for which U.S. dependence on Chinese imports rose most sharply during the five years preceding the policy shift.

8.1 Textual Evidence from USTR Press Releases

We first present textual evidence that 2018 USTR tariffs were partly motivated by the de-risking mechanism this paper highlights. We analyze the text of USTR communications, the agency's complete public record of announcing and justifying its trade actions, and find that around the 2018 rollout, the agency's stated rationale for tariffs shifted sharply toward the strategic risk-mitigation logic that generates targeted protection in our model. The traditional terms-of-trade rationale offers no explanation for this shift, since it relies on foreign producers bearing part of the tariff incidence and is silent on the riskiness of foreign sourcing. Yet it is exactly the riskiness angle that policymakers began to invoke in the build up to the 2018 tariff war.

We measure official framing in the universe of 2,983 USTR press releases published between January 2010 and December 2025. Each release is segmented into short local passages, totaling 55,003, and a large language model classifies every passage into six narratives: (1) import dependence, (2) criticality, (3) vulnerability, (4) geopolitical risk, (5) mitigation or de-risking, and, (5) their conjunction into an explicit strategic vulnerability narrative. This approach follows the text-as-data tradition of Baker et al. 2016; Gentzkow et al. 2019; Hassan et al. 2019; Caldara and Iacoviello 2022; Ash and Hansen 2023. Passage classifications are aggregated to the release level along two margins: the extensive margin records whether a release contains at least one passage in a narrative, and mention intensity records the share of a release's passages assigned to that narrative. Appendix K reports the full ontology, the classification rules, worked examples, and the validation evidence. Appendix K.6 shows that every result below is reproduced in direction and timing by three alternative classifiers run on identical passages.

Two of the six narratives carry the mechanism. Dependence alone records that USTR discussed foreign sourcing, and criticality alone records that it discussed important sectors; neither is evidence of the strategic-risk channel. The narrative that corresponds most closely to the model is the strategic vulnerability narrative, which requires a passage to join foreign exposure explicitly to strategic importance and to a risky or coercible supply relationship, rather than asserting fragility alone as the vulnerability narrative does. The geopolitical risk narrative records the second element of the mechanism on its

own: language connecting trade exposure to state power, coercion, rivalry, or foreign-government conduct. USTR describes the risk in geopolitical terms, and geopolitical risk as USTR uses the term may encompass coercion, rivalry, and security concerns beyond the risk we model. What we observe, however, is not only the language but also the actions. The tariffs set under this rhetoric fall on the sectors in which the United States is most dependent on China (Section 8.2). The pattern of actions therefore identifies a particular component of the geopolitical risk that USTR invokes: the risk that a country on which the United States depends restricts its exports. That component need not be the only one, but the targeting shows that it is present. The classifier applies a strong-entailment standard: naming a strategic sector alongside a foreign country satisfies neither narrative. The June 15, 2018 release announcing the first tariff list describes "the unprecedented threat posed by China's theft of our intellectual property, the forced transfer of American technology, and its cyber attacks" and is coded as both geopolitical risk and the strategic vulnerability narrative; a release announcing that "U.S. Beef Gains New Market Access in Morocco" is coded as neither, as are bare tariff-list references and meeting announcements. Routine market-access disputes, WTO litigation, and anti-dumping language are hard negatives by construction, so ordinary trade-remedy discourse cannot mechanically generate the narratives of interest.

The shift in official language around the 2018 rollout is large, statistically decisive, and concentrated in precisely the narratives the mechanism singles out. Table 4 reports document-level extensive margins for all six narratives in the pre-2017 baseline (1,159 releases) and the 2018 rollout year (109 releases). The strategic vulnerability narrative rises from 3.9 to 15.6 percent of releases, and geopolitical-risk framing rises from 6.0 to 26.6 percent, an increase of 20.6 percentage points; both roughly quadruple. The gradient across narratives is as informative as the levels. Vulnerability and criticality, component narratives that the mechanism activates only in conjunction with foreign risk, rise by 8 to 10 percentage points. Bare import dependence rises by 4.4 percentage points and is not statistically distinguishable from baseline. Mitigation or de-risking language does not rise at all: it falls from 8.3 to 5.5 percent of releases. USTR did not simply talk more about trade, China, or import reliance in 2018, and it did not justify the tariffs in the vocabulary of de-risking. The language that surged is the language that joins exposure to the risk of foreign-government action. This is the pattern implied by a tariff set ex ante: the stated rationale is the diagnosis of a risk rather than the description of a remedy. The comparison with 2017 is particularly demanding because it holds the administration and its rhetorical register fixed: in 2017 the geopolitical-risk margin was 3.7 percent, below the long-run baseline, so the 2018 surge is not an artifact of a change in administration

voice.

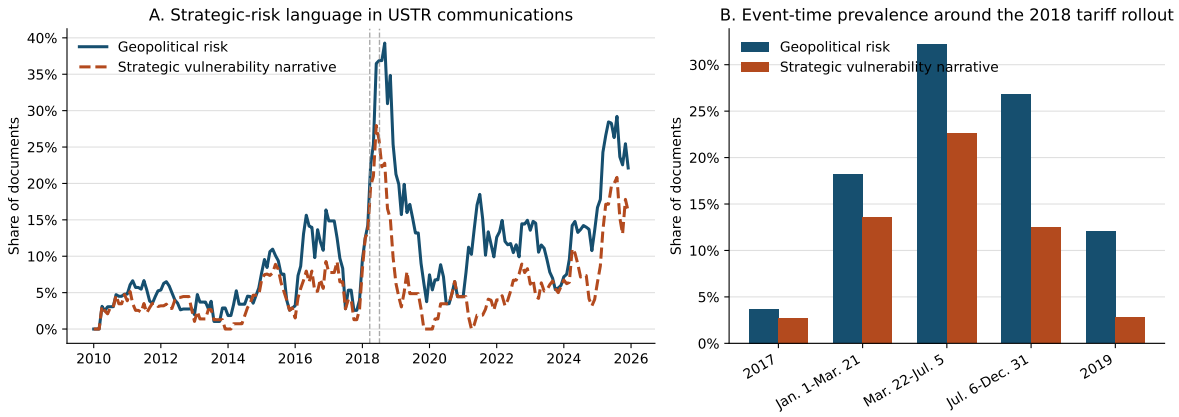
Table 4: Strategic-Risk Framing in USTR Communications: Pre-2017 versus 2018

Narrative	Pre-2017 baseline (percent)	2018 rollout (percent)	Change (pp)
Panel A. Mechanism-specific narratives			
Strategic vulnerability narrative	3.9	15.6	+11.7
	[2.9, 5.2]	[10.0, 23.6]	
Geopolitical risk	6.0	26.6	+20.6
	[4.8, 7.6]	[19.2, 35.6]	
Panel B. Component and response narratives			
Vulnerability	6.8	14.7	+7.9
	[5.5, 8.4]	[9.2, 22.5]	
Criticality	17.1	26.6	+9.5
	[15.0, 19.4]	[19.2, 35.6]	
Import dependence	6.6	11.0	+4.4
	[5.3, 8.2]	[6.4, 18.3]	
Mitigation or de-risking	8.3	5.5	−2.8
	[6.8, 10.0]	[2.5, 11.5]	

Notes: Entries are document-level extensive margins: the share of USTR releases containing at least one positive passage for the narrative. The pre-2017 baseline covers 1,159 releases (2010–2016) and the 2018 rollout covers 109 releases. Brackets report Wilson 95 percent confidence intervals. The final column reports two-sided Fisher exact tests of equal proportions across the two periods, treating releases as independent; Appendix K.4 discusses this assumption and reports the underlying counts.

The timing within 2018 ties the language to the policy event. Figure 7 plots the two mechanism-specific narratives over the full sample (Panel A) and in event time around the Section 301 action (Panel B). Framing begins to move in early 2018, when the steel and aluminum actions were announced and the Section 301 investigation, opened in August 2017, was nearing its findings. The concentrated jump occurs in the announcement-to-implementation window: among the 31 releases issued between the March 22 announcement of the China tariff action and the July 6 implementation of the first tariff list, 32.3 percent contain geopolitical-risk language and 22.6 percent contain the strategic vulnerability narrative, roughly nine and eight times the 2017 rates. The margins remain elevated through the second half of 2018 (26.8 and 12.5 percent), so the episode is not a single burst of announcement-day rhetoric. Because the announcement window and the post-implementation window are disjoint samples, each separately distinguishable from the 2017 baseline for geopolitical risk, the 2018 result does not hinge on any one cluster of related releases.

Figure 7: Textual evidence from USTR communications



Notes: Panel A reports six-month rolling document-level extensive margins for geopolitical risk and the strategic vulnerability narrative in USTR communications, 2010–2025. Panel B reports event-time extensive margins around the 2018 Section 301 tariff rollout. The vertical markers in Panel A denote the March 22, 2018 tariff announcement and the July 6, 2018 implementation of the first tariff list.

The path after 2018 is equally informative, and the two mechanism-specific narratives diverge in a way that fits their economic content. Background geopolitical-risk framing ratchets upward: it recedes only partially to 9.7 percent in 2019–2020, holds at 12.3 percent across 2021–2024, roughly double the pre-2017 baseline , and surges to 25.0 percent in 2025 as tariff escalation resumed.[10] The strategic vulnerability narrative instead behaves episodically: it returns to baseline in 2019–2020 (3.6 percent), edges up to 5.6 percent in 2021–2024, and spikes again in 2025. This is the pattern one would expect if the strategic vulnerability narrative is the vocabulary in which tariff action is publicly justified: it peaks in the two episodes of active escalation and subsides in between, while the background assessment of foreign-state risk, once raised, does not fully revert.

The same conclusions survive every reclassification and recomposition we examine, detailed in Appendix K.5. Measured on mention intensity rather than the extensive margin, geopolitical-risk language rises from 0.9 percent of passages per release before 2017 to 7.2 percent in 2018, so the result does not reflect isolated sentences in otherwise routine documents. Retaining only high-confidence positive classifications, a screen that can only remove positives, preserves the shift (4.5 to 18.3 percent for geopolitical risk). Excluding administrative scheduling releases leaves the 2018 margins unchanged, and within policy-action communications the 2018 margins reach 57.1 and 33.3 percent. Rerunning the entire classification with three alternative language models yields the same

[10]The 2025 releases lie outside the main estimation sample and function as an out-of-sample check: when tariff action returned, so did the framing, with the strategic vulnerability narrative reaching 18.3 percent of releases.

basic pattern. The models differ in how often they identify the two narratives tied most closely to the paper's strategic de-risking mechanism: geopolitical risk and strategic vulnerability. But all four find that both narratives rise sharply in 2018 and peak between the March tariff announcement and July implementation.[11]

8.2 Evidence from Applied Tariffs (Section 301)

We now turn to the applied Section 301 tariffs and the products they targeted. A natural starting point is to ask what motivated these tariffs. One hypothesis, grounded in the textbook terms-of-trade framework, is that by 2018 the U.S. government had concluded that the terms-of-trade gains from tariffs on China outweighed the costs that had deterred earlier administrations. This hypothesis carries a testable implication: tariff hikes should target goods with high export supply elasticities, where the incidence falls most heavily on foreign producers. An alternative hypothesis is that the tariffs reflected political pressure from lobbying sectors. If so, targeting should instead track import demand elasticities and we should see steeper hikes in sectors where China's import share was declining. Our analysis is mindful of both possibilities, yet the evidence supports neither. The significant predictor of the 2018 tariff hikes turns out to be whether a good had experienced a rapid rise in import dependence on China in the prior 5 years.

To provide some background, the 2018 tariff measures were rolled out in multiple phases. Two phases consisted of global safeguards that did not single out China: a Section 201 action in January 2018 on solar panels and washing machines, and a Section 232 action in March 2018 on steel and aluminum, applied to nearly every trading partner. The centerpiece of the policy shift was the China-specific Section 301 phase, implemented in three tranches over 2018: a first list covering roughly $34 billion of annual imports in July and a second covering about $16 billion in August, both subject to a 25 percent duty, followed by a much larger third list of almost $200 billion in September at an initial 10 percent rate, later raised to 25 percent in 2019.[12] We draw the monthly, product-level tariff and trade data from these phases from Fajgelbaum et al. (2020), with a detailed account of provided in Appendix J.

Our estimating equation regresses the statutory tariff change on a set of predictors

[11] Appendix K provides the ontology, examples, robustness checks, and event-time tables behind the classification.

[12] Because our tariff data end in early 2019, the changes we measure capture the three 2018 lists as implemented by year end, before the later escalation.

implied by theory:

$$\Delta \ln(1 + t_k) = \alpha + \beta \, \Delta \ln \lambda_k + \gamma \, \varepsilon_k^S + \delta \, \varepsilon_k^D + u_k. \tag{44}$$

Here, the outcome variable is the December 2017 to December 2018 change in statutory tariff on imports of product k from China. The main regressor of interest $\Delta \ln \lambda_k$ is the growth of China's import share in product k over the preceding five years, 2013 to 2017. This share is the empirical analogue of the good-level import dependence Λ_k in our model, and through the lens of our derisking theory, β must be positive. The controls ε_k^S and ε_k^D are the two elasticities emphasized by conventional theories of protection. The good-specific export supply elasticity ε_k^S is the central predictor under terms-of-trade theory, while the import demand elasticity ε_k^D is a key determinant of politically optimal tariffs (which are also negatively related to import shares). We source ε_k^S from Broda et al. (2008), who estimate foreign export-supply elasticities at the HS4 level, and ε^D from Broda and Weinstein (2006), who report import demand elasticities for each HS10 product line. We run this regression at two levels of aggregation, the HS10 tariff line and the HS4 sector heading, where the latter matches the level of aggregation for which export supply elasticities are available. The underlying assumption when running the HS10 level regression is that the export supply elasticity is common up to a random error term across HS10 lines within an HS4 heading.

As a falsification check, we re-estimate the regression on a placebo sample of tariffs. The outcome in regression 44 is replaced with the change in statutory tariffs applied to the same good k when imported from non-Chinese suppliers, holding the regressors unchanged. Because Section 301 duties are China-specific, growth in China dependence should not predict how tariffs on other partners were revised over this period. In other words, the placebo coefficients should be statistically indistinguishable from zero.

Table 5 reports the estimation results. The four columns pair the two aggregation levels, HS10 lines and HS4 headings, with the China and placebo outcomes, while the three panels vary the sample and specification. Across every column and panel, the two elasticity controls enter with small coefficients that are statistically insignificant in all but one cell, so the conventional determinants of the optimal tariff account for little of the cross-sectional variation in the 2018 tariff changes; the dependence-growth term carries the result. Our baseline, Panels A and B, excludes goods covered by the global Section 232 and 201 actions, which were not China-specific and are plausibly unrelated to the evolving U.S.-China trade relationship; Panel C restores them and shows that the China result is essentially unchanged. One caveat shapes what the table can say:

because a year-over-year tariff change is recorded only for goods that actually traded with the relevant origin in both Decembers, the China and placebo columns rest on different, unmatched samples. We therefore read each placebo column as a separate baseline rather than a matched counterfactual, and do not test cross-column coefficient differences.

Panel A presents the baseline contrast. The coefficient on prewar dependence growth is 0.649 at HS10 and 0.923 at HS4, both significant at the 1 percent level. A doubling of China's import share from 2013 to 2017 is associated with increases of about 0.45 and 0.64 points, respectively, in the resulting gross tariff hike in 2018. In the placebo non-China samples, the corresponding coefficients are statistically insignificant. Panel B interacts dependence growth with the sector's import-penetration ratio, measured from input-output accounts. The coefficient on the scaled regressor is larger in absolute magnitude and remains positive and statistically significant, while the non-China placebo coefficient remains indistinguishable from zero. Panel C restores goods covered by the global, all-origin Section 232 and 201 actions. The China coefficient remains positive and nearly unchanged at HS10; at HS4 it falls to 0.612 and is significant only at the 10 percent level. This attenuation is in fact reassuring: once China-specific and global tariff measures are pooled into the same sample, a China-specific determinant should lose explanatory power.

Taken together, Table 4 supports a narrow but informative conclusion: conditional on the two elasticity measures, tariffs on China rose more where reliance on China had been building before the trade war. Moreover, the placebo test confirms that the same pattern does not extend to tariff revisions on non-China partners. Read alongside the textual evidence of Section 8.1, which documents a parallel shift in USTR's stated rationale toward strategic derisking, the evidence accords well with our model's mechanism as an account of the shift in tariff policy that occurred circa 2018.

9 Conclusion

This paper develops a theory of optimal tariffs under strategic risk. We develop a model where a foreign supplier is more likely to impose export restrictions as aggregate import dependence grows. Under these circumstances, private sourcing decisions generate a negative externality: atomistic firms ignore how their collective import choices raise aggregate import dependence and exposure to foreign export restrictions, thereby amplifying aggregate risk. The optimal tariff, as a result, includes two components: a

Table 5: China Import Dependence and the 2018 Tariff Structure

	HS10 lines		HS4 headings	
	China $\Delta \ln \tau$ (1)	Placebo (2)	China $\Delta \ln \tau$ (3)	Placebo (4)
Panel A. Baseline				
Δ ln China import dependence	0.649***	0.021	0.923***	0.000
	(0.233)	(0.013)	(0.305)	(0.010)
Export supply elasticity ε^S	0.008	−0.001	0.007	−0.002
	(0.018)	(0.001)	(0.006)	(0.002)
Import demand elasticity ε^D	0.000	−0.000	0.045	0.001
	(0.003)	(0.000)	(0.028)	(0.001)
Observations	6,720	8,104	835	924
Panel B. Domestic absorption-scaled				
Δ ln China import dependence	1.921***	0.085*	5.797***	0.011
	(0.501)	(0.051)	(1.377)	(0.056)
Export supply elasticity ε^S	0.008	−0.001	0.007	−0.002
	(0.018)	(0.001)	(0.006)	(0.002)
Import demand elasticity ε^D	0.000	−0.000	0.041	0.001
	(0.003)	(0.000)	(0.027)	(0.001)
Observations	6,661	8,002	810	887
Panel C. All tariff lines including Section 232 & 201 goods				
Δ ln China import dependence	0.635***	−0.254	0.612*	−0.665**
	(0.220)	(0.291)	(0.329)	(0.259)
Export supply elasticity ε^S	0.010	0.026	0.012	0.012
	(0.018)	(0.033)	(0.008)	(0.011)
Import demand elasticity ε^D	0.000	0.000	0.052*	0.014
	(0.003)	(0.001)	(0.031)	(0.009)
Observations	6,886	8,469	859	954

Notes: Estimates of $\Delta \ln(1+\tau_k) \times 100 = \alpha + \beta \Delta \ln \lambda_k + \gamma \varepsilon_k^S + \delta \varepsilon_k^D + u_k$. The outcome is the December 2017 to December 2018 log change in the sector's value-weighted statutory tariff (m_stattariff1) on China imports (columns 1, 3) or non-China imports (placebo, columns 2, 4). $\Delta \ln \lambda_k$ is the 2013–2017 log growth of China's import share; ε^S is the Broda-Limão-Weinstein (2008) foreign export-supply elasticity (native HS4); ε^D is the Broda-Weinstein (2006) import-demand elasticity (native HS10); both enter in levels. Columns aggregate to the HS10 tariff line and the HS4 heading (the HS4 outcome, ε^D, and penetration are 2017-value-weighted means; the trajectory is the HS4 China-share log growth; ε^S is native). Panels A and B exclude Section 232 (steel/aluminum) and Section 201 (solar/washer) lines, both global (all-origin) tariffs; Panel C uses all goods. Panel A is the baseline; Panel B scales the trajectory by IO import penetration; Panel C is the baseline on all goods. No fixed effects. HS10 columns: standard errors clustered by HS2 chapter; HS4 columns: heteroskedasticity-robust (HC3).
*** Significant at the 1 percent level. ** 5 percent. * 10 percent.

standard terms-of-trade component and a strategic-risk component that internalizes the noted market failure. The strategic-risk component is absent when foreign is fully cooperative and rises with the elasticity of the export-restriction hazard with respect to aggregate import dependence.

The targeting implication is sharp. In the representative-sector model, the terms-of-trade motive generates nearly uniform tariffs. Sectoral heterogeneity in import concentration activates the sector-specific strategic-risk component and breaks this uniformity: the model predicts that optimal tariffs should be highest precisely in sectors where aggregate import dependence is greatest. These are also the sectors where comparative advantage is most pronounced. Thus, hedging against strategic risk calls for restricting trade in precisely the sectors that deliver the largest gains from trade.

We discipline the model using panel data on export restrictions and lagged exporter shares in sectoral imports for the US and the EU vis-a-vis trade partners. The stock-OLS estimates imply positive risk elasticities for both regions. In the calibrated macro model, strategic risk raises the optimal US tariff from 25.78 to 31.01 percent and the EU tariff from 26.02 to 30.93 percent. The risk correction accounts for 16.86 percent of the US tariff and 15.88 percent of the EU tariff. In the sixteen-sector model, endogenous-risk tariffs range from 26.92 to 36.79 percent in the United States and from 27.22 to 38.35 percent in the European Union, rising sharply with sectoral import concentration. This targeting pattern is consistent with the cross-sectoral incidence of the 2018 United States tariff increases on China.

References

Adamopoulos, T. and F. Leibovici (2025, November). Trade risk and food security. Working Paper 2024-004, Federal Reserve Bank of St. Louis. November 2025 version; the optimal-policy analysis appears as Section 6. 1

Ash, E. and S. Hansen (2023). Text algorithms in economics. *Annual Review of Economics 15*, 659–688. 8.1

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

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *Quarterly Journal of Economics 131*(4), 1593–1636. 8.1

Becko, J. S., G. M. Grossman, and E. Helpman (2025, August). Optimal tariffs with geopolitical alignment. Working Paper 34108, National Bureau of Economic Research. 1

Becko, J. S. and D. G. O'Connor (2025). Strategic (dis)integration. Working paper, January 2025 (first version March 2024). 1

Bickerdike, C. F. (1906). The theory of incipient taxes. *The Economic Journal 16*(64), 529–535. 1

Blanchard, E. J., C. P. Bown, and R. C. Johnson (2026). Global value chains and trade policy. *The Review of Economic Studies 93*(1), 181–214. Circulated as "Global Supply Chains and Trade Policy," NBER Working Paper 21883. 1

Bown, C. P. (2004). On the economic success of GATT/WTO dispute settlement. *Review of Economics and Statistics 86*(3), 811–823. 1, 1

Bown, C. P. (2019). The 2018 US–China trade conflict after forty years of special protection. *China Economic Journal 12*(2), 109–136. 1

Bown, C. P. (2023, June). The challenge of export controls. *Finance and Development 60*(2), 18–21. 1

Bown, C. P. (2024). Modern industrial policy and the World Trade Organization. *Annual Review of Economics 16*, 243–270. 1, 1

Bown, C. P. (2025). How export restrictions threaten economic security. *Journal of International Economic Law 28*(2), 278–292. PIIE Working Paper 25-11. 1, 1, 1

Bown, C. P. and M. A. Crowley (2013). Self-enforcing trade agreements: Evidence from time-varying trade policy. *American Economic Review 103*(2), 1071–1090. 1

Bown, C. P. and J. A. Hillman (2019). WTO'ing a resolution to the China subsidy problem. *Journal of International Economic Law 22*(4), 557–578. 1, 1

Bown, C. P. and S. Keynes (2020). Why Trump shot the sheriffs: The end of WTO dispute settlement 1.0. *Journal of Policy Modeling 42*(4), 799–819. 1

Broda, C., N. Limão, and D. E. Weinstein (2008). Optimal tariffs and market power: The evidence. *American Economic Review 98*(5), 2032–2065. 8.2

Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *Quarterly Journal of Economics 121*(2), 541–585. 8.2

Bronckers, M. and K. E. Maskus (2014). China–Raw Materials: A controversial step towards evenhanded exploitation of natural resources. *World Trade Review 13*(2), 393–408. 1

Caldara, D. and M. Iacoviello (2022). Measuring geopolitical risk. *American Economic Review 112*(4), 1194–1225. 8.1

Çelik, L., B. Karabay, and J. McLaren (2020). Fast-track authority: A hold-up interpretation. *Journal of International Economics 127*, 103392. NBER Working Paper 24427. 1

Clayton, C., M. Maggiori, and J. Schreger (2024, December). A theory of economic coercion and fragmentation. Working Paper 33309, National Bureau of Economic Research. Revised April 2025; also BIS Working Paper 1224. 1

Clayton, C., M. Maggiori, and J. Schreger (2026a). A framework for geoeconomics. *Econometrica 94*(1), 105–136. NBER Working Paper 31852. 1

Clayton, C., M. Maggiori, and J. Schreger (2026b, June). The targets of geoeconomic coercion. Working Paper 35343, National Bureau of Economic Research. Revised July 2026. 1

Costinot, A., D. Donaldson, J. Vogel, and I. Werning (2015). Comparative advantage and optimal trade policy. *The Quarterly Journal of Economics 130*(2), 659–702. 1

Eisenbarth, S. (2017). Is Chinese trade policy motivated by environmental concerns? *Journal of Environmental Economics and Management 82*, 74–103. 1

Elliott, M., B. Golub, and M. V. Leduc (2022). Supply network formation and fragility. *American Economic Review 112*(8), 2701–2747. 1

Evenett, S. J. and J. Fritz (2024). Global trade alert database. University of St. Gallen. Accessed: 2024. 6.2

Executive Office of the President (2017, December). Executive order 13817: A federal strategy to ensure secure and reliable supplies of critical minerals. Federal Register 82(246), 61033–61035. Signed December 20, 2017. 1

Fajgelbaum, P. D., P. K. Goldberg, P. J. Kennedy, and A. K. Khandelwal (2020). The return to protectionism. *Quarterly Journal of Economics 135*(1), 1–55. 2, 8.2, J.1, 10

Garred, J. (2018). The persistence of trade policy in China after WTO accession. *Journal of International Economics 114*, 130–142. 1

Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as data. *Journal of Economic Literature 57*(3), 535–574. 8.1

Gros, D. (1987). A note on the optimal tariff, retaliation and the welfare loss from tariff wars in a framework with intra-industry trade. *Journal of International Economics 23*(3–4), 357–367. 1, 4

Grossman, G. M. and E. Helpman (1994). Protection for sale. *American Economic Review 84*(4), 833–850. 2

Grossman, G. M., E. Helpman, and H. Lhuillier (2023). Supply chain resilience: Should policy promote international diversification or reshoring? *Journal of Political Economy 131*(12), 3462–3496. NBER Working Paper 29330. 1

Grossman, G. M., E. Helpman, and A. Sabal (2023, September). Resilience in vertical supply chains. Working Paper 31739, National Bureau of Economic Research. 1

Grossman, G. M., E. Helpman, and A. Sabal (2024). Optimal resilience in multitier supply chains. *The Quarterly Journal of Economics 139*(4), 2377–2425. NBER Working Paper 31739. 1

Handley, K. and N. Limão (2017). Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States. *American Economic Review 107*(9), 2731–2783. 1

Hassan, T. A., S. Hollander, L. van Lent, and A. Tahoun (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics 134*(4), 2135–2202. 8.1

Interagency Task Force in Fulfillment of Executive Order 13806 (2018, September). Assessing and strengthening the manufacturing and defense industrial base and supply chain resiliency of the United States. Report to president donald j. trump, U.S. Department of Defense. 1

Johnson, H. G. (1953). Optimum tariffs and retaliation. *The Review of Economic Studies 21*(2), 142–153. 1

Kopytov, A., B. Mishra, K. Nimark, and M. Taschereau-Dumouchel (2024). Endogenous production networks under supply chain uncertainty. *Econometrica 92*(5), 1621–1659. 1

Leibovici, F. and A. M. Santacreu (2020). Shortages of critical goods in a global economy: Optimal trade and industrial policy. Working Paper 2020-010, Federal Reserve Bank of St. Louis. Revised June 2025. 1

Maggi, G. and R. W. Staiger (2011). The role of dispute settlement procedures in international trade agreements. *Quarterly Journal of Economics 126*(1), 475–515. 1

McLaren, J. (1997). Size, sunk costs, and Judge Bowker's objection to free trade. *American Economic Review 87*(3), 400–420. 1

OECD (2021). Inter-country input-output tables. Paris: OECD Publishing. 6.1

Qin, J. Y. (2012). Reforming WTO discipline on export duties: Sovereignty over natural resources, economic development and environmental protection. *Journal of World Trade 46*(5), 1147–1190. 1

Reuters (2019). China ready to hit back at U.S. with rare earths: Ruling party newspaper. Reuters, May 29, 2019. Reporting the People's Daily commentary of May 29, 2019. 1

Ruhl, K. J. (2026, June). Rethinking free trade. *Finance and Development 63*(2), 10–12. 1

Ruta, M. and A. J. Venables (2012). International trade in natural resources: Practice and policy. *Annual Review of Resource Economics 4*, 331–352. 1

Simonovska, I. and M. E. Waugh (2014). The elasticity of trade: Estimates and evidence. *Journal of international Economics 92*(1), 34–50. 6.1

U.S. Department of Commerce (2019, June). A federal strategy to ensure secure and reliable supplies of critical minerals. Technical report, U.S. Department of Commerce. 1

Wu, M. (2017). China's export restrictions and the limits of WTO law. *World Trade Review 16*(4), 673–691. 1

Appendix

A Optimal Tariff with Passive Foreign

This appendix derives the optimal tariff formula presented under Proposition 1. Home chooses $t \equiv \tau_{HF}^M - 1 \geq 0$ with $\tau_{HF}^M = 1$ and $\tau_{ni}^X = 1$ for all (n, i). Define aggregate demand from Home for Foreign goods as $Q_{FH} \equiv \lambda_{FH} E_H / \tilde{p}_{FH}$, where $\tilde{p}_{FH} = (1+t) d_{FH} p_F$. Home income can be written as

$$Y_H = (1 + \frac{1}{\gamma\beta}) w_H L_H + (\tilde{p}_{FH} - d_{FH} p_F) Q_{FH}, \tag{45}$$

where the second term represents tariff revenue. Balanced trade implies $d_{FH} p_F Q_{FH} = \lambda_{HF} w_F L_F$, yielding:

$$w_F = \frac{d_{FH} p_F Q_{FH}}{\lambda_{HF} L_F}. \tag{46}$$

Foreign Prices and Trade Shares. The Foreign unit cost is $p_F = \frac{1}{\bar{\varphi}_F} w_F^{\beta} \tilde{P}_F^{1-\beta}$. With $d_{FF} = \tau_{FF}^M = \tau_{FF}^X = 1$ and $p_H \equiv 1$, the CES price index in Foreign is $\tilde{P}_F = p_F \lambda_{FF}^{1/\epsilon}$ where:

$$\lambda_{FF}(p_F) = \frac{p_F^{-\epsilon}}{p_F^{-\epsilon} + d_{HF}^{-\epsilon}}, \qquad \lambda_{HF}(p_F) = 1 - \lambda_{FF}(p_F).$$

Plugging $\tilde{P}_F = p_F \lambda_{FF}^{1/\epsilon}$ into the expression for p_F implies that $p_F = \bar{\varphi}_F^{-1/\beta} w_F \lambda_{FF}^{\frac{1-\beta}{\beta\epsilon}}$. Substituting the above expression back into (46) and rearranging yields:

$$Q_{FH} = \bar{\varphi}_F^{\frac{1}{\beta}} \lambda_{FF}(p_F)^{-\frac{1-\beta}{\beta\epsilon}} \lambda_{HF}(p_F) L_F. \tag{47}$$

This equation implicitly defines p_F as a function of Q_{FH}.

Inverse Export Supply Elasticity. Differentiating (47) logarithmically obtains,

$$d \ln Q_{FH} = -\frac{1-\beta}{\beta\epsilon} d \ln \lambda_{FF} + d \ln \lambda_{HF}, \tag{48}$$

where the CES demand system implies that $\frac{\mathrm{d} \ln \lambda_{FF}}{\mathrm{d} \ln p_F} = -\epsilon \lambda_{HF}$ and $\frac{\mathrm{d} \ln \lambda_{HF}}{\mathrm{d} \ln p_F} = \epsilon \lambda_{FF}$. Substituting $\mathrm{d} \ln \lambda_{FF} = -\epsilon \lambda_{HF} d \ln p_F$ and $\mathrm{d} \ln \lambda_{HF} = \epsilon \lambda_{FF} \mathrm{d} \ln p_F$ into 48 and rearranging

delivers:

$$\frac{1}{\varepsilon_{FH}^{S}} \equiv \frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}} = \frac{1}{-\frac{1-\beta}{\beta\epsilon}\cdot(-\epsilon\lambda_{HF})+\epsilon\lambda_{FF}} = \frac{1}{\epsilon\lambda_{FF}+\frac{1-\beta}{\beta}\lambda_{HF}}. \tag{49}$$

Income Derivative. Differentiating Home income (45) with respect to $\ln(1+t)$:

$$\begin{aligned}\frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln(1+t)} = \frac{1}{Y_H}\Bigg[&(1+\frac{1}{\gamma\beta})w_H L_H \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln(1+t)} + \tilde{p}_{FH}Q_{FH} \\ &+\left(\tilde{p}_{FH}-\left(1+\frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}}\right)d_{FH}p_F\right)Q_{FH}\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln(1+t)}\Bigg].\end{aligned} \tag{50}$$

Using $\tilde{p}_{FH} = (1+t)d_{FH}p_F$, this simplifies to:

$$\frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln(1+t)} = \left[\frac{w_H L_H}{E_H}\frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln(1+t)} + \lambda_{FH} + \left(1-\frac{1}{1+t}\left(1+\frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}}\right)\right)\lambda_{FH}\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln(1+t)}\right]\frac{E_H}{Y_H}. \tag{51}$$

Price Index Derivative. Taking the derivative of the CES price index in Home obtains: $\frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln(1+t)} = \lambda_{FH}$. The numeraire condition $p_H \equiv 1 = \frac{1}{\tilde{\varphi}_H} w_H^{\beta}\tilde{P}_H^{1-\beta}$ implies:

$$\beta\frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln(1+t)} + (1-\beta)\frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln(1+t)} = 0 \quad \Rightarrow \quad \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln(1+t)} = -\frac{1-\beta}{\beta}\lambda_{FH}. \tag{52}$$

First-Order Condition. Combining the expressions above yields the necessary first-order condition for optimality:

$$\begin{aligned}\frac{\mathrm{d}\ln W_H}{\mathrm{d}\ln(1+t)} =& \frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln(1+t)} - \frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln(1+t)} \\ =& \left[\frac{w_H L_H}{E_H}\frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln(1+t)} + \lambda_{FH} + \left(1-\frac{1}{1+t}\left(1+\frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}}\right)\right)\lambda_{FH}\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln(1+t)}\right]\frac{E_H}{Y_H} - \lambda_{FH} \\ =& \left[\frac{w_H L_H}{E_H}\left(\frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln(1+t)} + \frac{1-\beta}{\beta}\lambda_{FH}\right) + \left(1-\frac{1}{1+t}\left(1+\frac{1}{\varepsilon_{FH}^{S}}\right)\right)\lambda_{FH}\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln(1+t)}\right]\frac{E_H}{Y_H} = 0\end{aligned} \tag{53}$$

where the last lines uses $\lambda_{FH}\frac{E_H}{Y_H} - \lambda_{FH} = \frac{w_H L_H}{Y_H}\frac{1-\beta}{\beta}\lambda_{FH}$. Using (52), the first bracketed term vanishes. Thus, (53) reduces to $1-\frac{1}{1+t}(1+\frac{1}{\varepsilon_{FH}^{S}}) = 0$. Rearranging this equation delivers the optimal tariff formula: $t^* = 1/\varepsilon_{FH}^{S}$. Substituting(49) into this equation

produces the extended version of optimal tariff formula in Proposition 1:

$$t^* = \frac{1}{\varepsilon^S_{FH}} = \frac{1}{\epsilon\lambda_{FF} + \frac{1-\beta}{\beta}\lambda_{HF}}.$$

B Foreign's Optimal Export Tax

This appendix derives Foreign's unilaterally optimal export tax, proving Lemmas 1 and 2. As in the main text, Home's import tariff rate is t, Foreign's export tax rate is x, and the buyer price of Foreign goods in Home is $\tilde{p}_{FH} = (1+t)(1+x)d_{FH}p_F$. Throughout, we take Home labor as the numeraire, $w_H = 1$; because trade shares and welfare are homogeneous of degree zero in nominal prices, this choice is equivalent to the main text's normalization $p_H = 1$ and affects no statement below.

Tariffs range over $0 \leq t \leq \frac{\epsilon\beta}{1-\beta} - 1$, the domain on which Lemma 3 operates; by Assumption 2 this interval is nonempty, and it contains the policy domain $\mathbb{T}$ of the main text. We maintain the standing assumptions that the equilibrium exists, is unique, and is continuously differentiable in the policy pair, and that Foreign's noncooperative best response $x^*(t) \in \arg\max_{x\geq 0} v_F(x;t)$ is interior. The analysis below shows, as a byproduct, that this interior optimum is unique and globally optimal.

The argument has three parts. The first two subsections reduce the two-country equilibrium to a system of two equations in the import shares and record how those shares respond to the policy wedges. Lemma 1 then follows from differentiating Foreign's welfare and factoring the first-order condition; Lemma 2, from differentiating that condition along the noncooperative locus.

The equilibrium in import-share odds

Because domestic varieties face no policy wedges, the trade-share system (15) restricts relative, not absolute, shares. It is therefore convenient to work in the two import-share odds

$$\tilde{\lambda}_H \equiv \frac{\lambda_{FH}}{1-\lambda_{FH}} = \frac{\lambda_{FH}}{\lambda_{HH}}, \qquad \tilde{\lambda}_F \equiv \frac{\lambda_{HF}}{1-\lambda_{HF}} = \frac{\lambda_{HF}}{\lambda_{FF}}, \tag{54}$$

each measuring imports relative to domestic expenditure in the destination country. The odds separate the direct effects of the two policy wedges from the general-equilibrium adjustment of wages and price indices. In these coordinates the equilibrium reduces,

exactly, to two equations in the odds alone,

$$\tilde{\lambda}_H \tilde{\lambda}_F = \left[(1+t)(1+x)\, d_{FH} d_{HF}\right]^{-\epsilon}, \tag{55}$$

$$1 = \frac{L_H}{d_{HF} L_F} \left(\frac{\bar{\varphi}_H}{\bar{\varphi}_F}\right)^{1/\beta} \frac{\tilde{\lambda}_H \left(1+\tilde{\lambda}_H\right)^{\frac{1-\beta}{\beta\epsilon}}}{1+t+\tilde{\lambda}_H} \frac{1+\tilde{\lambda}_F/(1+x)}{\tilde{\lambda}_F^{1+1/\epsilon} \left(1+\tilde{\lambda}_F\right)^{\frac{1-\beta}{\beta\epsilon}}}. \tag{56}$$

The first equation is gravity: both policy wedges contract two-way trade symmetrically in the product of the odds, and because trade costs satisfy $d_{ni} \geq 1$ and the gross wedges are at least one, the product is at most one. The second equation is market clearing: it encodes the wage that balances trade and determines how the contraction splits between the two countries. Every remaining equilibrium variable is recovered algebraically from the pair $(\tilde{\lambda}_H, \tilde{\lambda}_F)$. The rest of this subsection derives the two equations and records the recovery.

Dividing the import share by the domestic share in each destination gives

$$\tilde{\lambda}_H = \left[(1+t)(1+x)\, d_{FH} \frac{\bar{\varphi}_H}{\bar{\varphi}_F} \frac{w_F^{\beta} \tilde{P}_F^{1-\beta}}{w_H^{\beta} \tilde{P}_H^{1-\beta}}\right]^{-\epsilon}, \qquad \tilde{\lambda}_F = \left[d_{HF} \frac{\bar{\varphi}_F}{\bar{\varphi}_H} \frac{w_H^{\beta} \tilde{P}_H^{1-\beta}}{w_F^{\beta} \tilde{P}_F^{1-\beta}}\right]^{-\epsilon}. \tag{57}$$

Multiplying the two ratios eliminates every price and delivers the gravity equation (55). The second equation comes from market clearing.

Evaluating the price index (16) against the domestic term of (15) ties each index to its own domestic share:

$$\tilde{P}_H = \bar{\varphi}_H^{-1/\beta}\, w_H\, \lambda_{HH}^{1/(\beta\epsilon)}, \qquad \tilde{P}_F = \bar{\varphi}_F^{-1/\beta}\, w_F\, \lambda_{FF}^{1/(\beta\epsilon)}. \tag{58}$$

Next, income accounting. Home levies only the import tariff and rebates it, so (14)–(20) give $E_H = \frac{1}{\beta} w_H L_H + \frac{t}{1+t} \lambda_{FH} E_H$, that is,

$$E_H = \frac{w_H L_H}{\beta\left[1 - \frac{t}{1+t}\lambda_{FH}\right]} = \frac{w_H L_H}{\beta}\, \delta, \qquad \delta \equiv \left[1 - \frac{t}{1+t}\lambda_{FH}\right]^{-1}. \tag{59}$$

Here δ is the revenue factor of Lemma 1—one plus the ratio of tariff revenue to Home's pre-rebate income—now evaluated at an arbitrary policy pair; along the noncooperative locus it coincides with $\delta(\Lambda)$. Foreign's accounts are the mirror image, with the export tax in place of the tariff. Its export-tax revenue (19) is $\frac{x}{1+x} \cdot \frac{1}{1+t} \lambda_{FH} E_H$, which equals

$\frac{x}{1+x}\lambda_{HF}E_F$ by balanced trade (22), so

$$E_F = \frac{w_F L_F}{\beta\left[1-\frac{x}{1+x}\lambda_{HF}\right]} = \frac{w_F L_F}{\beta}\,\delta_F, \qquad \delta_F \equiv \left[1-\frac{x}{1+x}\lambda_{HF}\right]^{-1}. \tag{60}$$

Net incomes follow from (14): $Y_H = \frac{w_H L_H}{\beta}\left[\delta-(1-\beta)\right]$ and $Y_F = \frac{w_F L_F}{\beta}\left[\delta_F-(1-\beta)\right]$. Dividing by the price indices (58) cancels the wage from each country's real income and gives, up to constants independent of policy,

$$v_H(x;t) = -\frac{1}{\beta\epsilon}\ln\lambda_{HH} + \ln\left[\delta-(1-\beta)\right], \tag{61}$$

$$v_F(x;t) = -\frac{1}{\beta\epsilon}\ln\lambda_{FF} + \ln\left[\delta_F-(1-\beta)\right]. \tag{62}$$

Each country's welfare is the sum of a gains-from-trade term, which rises as the country sources more from abroad, and a revenue term, which rises with the receipts its own wedge collects. Terms-of-trade gains are not lost in this representation: they accrue as revenue collected on imports bought at depressed producer prices, and hence live inside δ and δ_F.

Finally, the wage. Balanced trade (22) together with (59)–(60) gives $\lambda_{HF}\frac{w_F L_F}{\beta}\delta_F = \frac{1}{1+t}\lambda_{FH}\frac{w_H L_H}{\beta}\delta$, so

$$w_F = \frac{L_H}{L_F}\,\frac{\lambda_{FH}\,\delta}{(1+t)\,\lambda_{HF}\,\delta_F}\,w_H. \tag{63}$$

Substituting (58) and (63) into the second ratio of (57) yields

$$\tilde{\lambda}_F = \left[d_{HF}\left(\frac{\bar{\varphi}_F}{\bar{\varphi}_H}\right)^{1/\beta}\frac{L_F}{L_H}\,\frac{(1+t)\,\lambda_{HF}\,\delta_F}{\lambda_{FH}\,\delta}\left(\frac{\lambda_{HH}}{\lambda_{FF}}\right)^{\frac{1-\beta}{\beta\epsilon}}\right]^{-\epsilon}, \tag{64}$$

and expressing the shares and revenue factors in the odds—$\lambda_{FH}\delta = (1+t)\tilde{\lambda}_H/(1+t+\tilde{\lambda}_H)$, $\lambda_{HF}\delta_F = (1+x)\tilde{\lambda}_F/(1+x+\tilde{\lambda}_F)$, and $\lambda_{HH}/\lambda_{FF} = (1+\tilde{\lambda}_F)/(1+\tilde{\lambda}_H)$—turns (64) into the second equation (56). Conversely, from any positive pair solving (55)–(56) we recover the shares from (54), w_F from (63), the price indices from (58), and expenditures from (59)–(60); the recovered variables then satisfy the full equilibrium system of the main text. The reduction is therefore exact—not merely a necessary condition—and all derivatives below are computed from the two-equation system.

The response of import dependence to the two wedges

Totally differentiate (55) and (56) in logarithms, using the share form (64) of the second equation, so that the only inputs are the definitions of δ and δ_F and the identity $\mathrm{d}\lambda_{FH} = \lambda_{FH}\lambda_{HH}\,\mathrm{d}\ln\tilde{\lambda}_H$ (and its Foreign counterpart). Collecting terms,

$$\begin{aligned} \mathrm{d}\ln\tilde{\lambda}_H + \mathrm{d}\ln\tilde{\lambda}_F &= -\epsilon\left[\mathrm{d}\ln(1+t) + \mathrm{d}\ln(1+x)\right], \\ \varepsilon^S_{HF}\,\mathrm{d}\ln\tilde{\lambda}_H - \varepsilon^D_{HF}\,\mathrm{d}\ln\tilde{\lambda}_F &= \epsilon\lambda_{HH}\delta\,\mathrm{d}\ln(1+t) + \frac{\epsilon\lambda_{HF}\delta_F}{1+x}\,\mathrm{d}\ln(1+x), \end{aligned} \tag{65}$$

where

$$\varepsilon^S_{HF} = \epsilon\lambda_{HH}\delta + \frac{1-\beta}{\beta}\lambda_{FH} \tag{66}$$

is Home's export-supply elasticity—comparing with (26), it satisfies the identity $\varepsilon^S_{HF} = \varepsilon^D_{FH} - 1$—and

$$\varepsilon^D_{HF} \equiv 1 + \epsilon\lambda_{FF}\delta_F + \frac{1-\beta}{\beta}\lambda_{HF} \tag{67}$$

is Foreign's import demand elasticity, the exact counterpart of (26) with the export tax in the role of Foreign's Lerner-symmetric tariff. Solving the linear system (65) for the response of Home's import odds,

$$\mathrm{d}\ln\tilde{\lambda}_H = -\varepsilon^M\,\mathrm{d}\ln(1+t) - \varepsilon^X\,\mathrm{d}\ln(1+x), \qquad \varepsilon^M = \frac{\epsilon\left[\varepsilon^D_{HF} - \lambda_{HH}\delta\right]}{\varepsilon^S_{HF} + \varepsilon^D_{HF}}, \qquad \varepsilon^X = \frac{\epsilon\left[\varepsilon^D_{HF} - \frac{\lambda_{HF}\delta_F}{1+x}\right]}{\varepsilon^S_{HF} + \varepsilon^D_{HF}}, \tag{68}$$

where the superscripts follow the wedge notation of the main text: ε^M and ε^X are the equilibrium elasticities of Home's import odds $\tilde{\lambda}_H$ with respect to the gross tariff $\tau^M_{FH} = 1+t$ and the gross export tax $\tau^X_{FH} = 1+x$. Both are strictly positive: Foreign's $\varepsilon^D_{HF} > 1$ term by term, while

$$\lambda_{HH}\delta = \frac{(1+t)\lambda_{HH}}{1+t\lambda_{HH}} < 1, \qquad \frac{\lambda_{HF}\delta_F}{1+x} = \frac{\lambda_{HF}}{1+x\lambda_{FF}} < 1. \tag{69}$$

The first expression in (69) is Home's spending share on domestic goods measured gross of the rebated tariff; the second is its Foreign counterpart. A higher tariff and a higher export tax therefore each strictly lower Home's import dependence at every policy pair—the share signs invoked throughout the paper.

Both elasticities also lie strictly below the trade elasticity ϵ, since $\epsilon - \varepsilon^M = \epsilon\left[\varepsilon^S_{HF} + \lambda_{HH}\delta\right] / \left[\varepsilon^S_{HF} + \varepsilon\right.$ 0, and the same computation with $\lambda_{HF}\delta_F/(1+x)$ in place of $\lambda_{HH}\delta$ gives $\epsilon - \varepsilon^X > 0$: the general-equilibrium wage and price-index adjustment dampens but never overturns the

direct CES substitution response. The two gross-of-rebate shares are also ordered:

$$\lambda_{HH}\delta - \frac{\lambda_{HF}\delta_F}{1+x} = \frac{(1+t)(1+x)\lambda_{HH}\lambda_{FF} - \lambda_{FH}\lambda_{HF}}{(1+t\lambda_{HH})(1+x\lambda_{FF})} \geq 0, \tag{70}$$

because (55) bounds $\tilde{\lambda}_H\tilde{\lambda}_F$ by one while $(1+t)(1+x) \geq 1$; the inequality is strict whenever $t > 0$, $x > 0$, or trade costs exceed one. By (68), this ordering gives $\varepsilon^M \leq \varepsilon^X$, again strictly under the same conditions.

For later use, record the partial derivatives of the export-supply elasticity (66) with respect to its two arguments:

$$\partial_{\lambda_{FH}}\varepsilon^S_{HF} = \frac{1-\beta}{\beta} - \frac{\epsilon\delta^2}{1+t} < 0, \qquad \partial_{\ln(1+t)}\varepsilon^S_{HF}\Big|_{\lambda_{FH}} = \frac{\epsilon\lambda_{FH}\lambda_{HH}\delta^2}{1+t} > 0. \tag{71}$$

The first sign holds throughout the policy domain: at $t = 0$, $\delta = 1$ and the expression equals $\frac{1-\beta}{\beta} - \epsilon < 0$ by Assumption 2; for $t > 0$, $\delta > 1$ and $\epsilon\delta^2/(1+t) > \epsilon/(1+t) \geq \frac{1-\beta}{\beta}$ by the tariff cap. Economically, higher import dependence makes Home's residual demand for Foreign goods less elastic—the margin that gives dependence its strategic value—while, at a fixed share, a higher tariff makes it more elastic through the revenue channel.

Lemma 1: Foreign's optimal export tax given t

Fix t and differentiate (62) with respect to $\ln(1+x)$. Foreign's welfare responds to its export tax directly through the revenue factor δ_F and indirectly through the equilibrium contraction of its own import share. The partial derivatives of (62) are

$$\partial_{\ln(1+x)}v_F\Big|_{\lambda_{HF}} = \frac{\lambda_{HF}}{1+x}\frac{\delta_F^2}{\delta_F - (1-\beta)} > 0, \qquad \partial_{\lambda_{HF}}v_F = \frac{x}{1+x}\frac{\delta_F^2}{\delta_F - (1-\beta)} + \frac{1}{\beta\epsilon\lambda_{FF}} > 0, \tag{72}$$

and, from (55) and (68), the equilibrium response of Foreign's import share to its own tax at a fixed tariff is

$$\frac{\mathrm{d}\lambda_{HF}}{\mathrm{d}\ln(1+x)} = -\epsilon\lambda_{HF}\lambda_{FF}\frac{\varepsilon^S_{HF} + \frac{\lambda_{HF}\delta_F}{1+x}}{\varepsilon^S_{HF} + \varepsilon^D_{HF}} < 0. \tag{73}$$

Combining these three expressions and collecting terms over a common denominator, the total derivative factors exactly:

$$\frac{\mathrm{d}v_F}{\mathrm{d}\ln(1+x)} = \underbrace{\frac{\lambda_{HF}\left[(\epsilon+1)(1+x)\lambda_{FF} + \frac{1-\beta}{\beta}\lambda_{HF}\left(1+x\lambda_{FF}\right)\right]}{\beta\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]\left(1+x\lambda_{FF}\right)\left[\varepsilon^S_{HF}+\varepsilon^D_{HF}\right]}}_{>0} \times \left[1 - x\,\varepsilon^S_{HF}\right]. \tag{74}$$

The first factor collects market-size and elasticity terms and is strictly positive at every policy pair; the sign of the derivative is governed entirely by the bracket. The economics is the Lerner monopoly logic: by Lerner symmetry Foreign's export tax is its import tariff. The marginal revenue gain on inframarginal trade outweighs the marginal trade loss exactly while the tax remains below the inverse of the export-supply elasticity it faces. Foreign's optimal export tax therefore satisfies $x^*(t)\,\varepsilon^S_{HF}(x^*(t);t) = 1$, that is,

$$x^*(t) = \left[\frac{\epsilon\lambda_{HH}(x^*(t);t)}{1-\frac{t}{1+t}\lambda_{FH}(x^*(t);t)} + \frac{1-\beta}{\beta}\lambda_{FH}(x^*(t);t)\right]^{-1}. \tag{75}$$

Write $\Lambda(t) \equiv \lambda_{FH}(x^*(t);t)$ for aggregate import dependence under the noncooperative regime. Step 2 of the proof of Lemma 2 below establishes a bijection between t and Λ; it rests only on the share-form condition (75), so the reference is not circular. Using that bijection to define $\delta(\Lambda)$ along the locus, we obtain

$$x^*(\Lambda) = \left[\epsilon\left[1-\Lambda\right]\delta(\Lambda) + \frac{1-\beta}{\beta}\Lambda\right]^{-1}, \tag{76}$$

which is Lemma 1. Comparing (66) with Home's import demand elasticity (26) restates the optimality condition in Lerner form:

$$\varepsilon^D_{FH} = 1 + \epsilon\left[1-\Lambda(t)\right]\delta(\Lambda) + \frac{1-\beta}{\beta}\Lambda(t) = 1 + \frac{1}{x^*(t)}: \tag{77}$$

Foreign's unilaterally optimal export tax positions Home's residual import demand at the textbook monopoly point, where the demand elasticity equals one plus the inverse markup.

Two remarks complete the derivation. First, at $t = 0$ we have $\delta = 1$, and (76) reduces to $x^* = \left[\epsilon\lambda_{HH} + \frac{1-\beta}{\beta}\lambda_{FH}\right]^{-1}$, the Lerner-symmetric application of Proposition 1 stated in the main text. Second, the factorization (74) implies that the maintained interior optimum is unique and globally optimal: as x rises at a fixed tariff, λ_{FH} falls by (68), so

ε^S_{HF} rises by (71), and $x\,\varepsilon^S_{HF}$ increases strictly from zero through one; Foreign's welfare is single-peaked in its tax, rising before the crossing and falling after it.

Lemma 2: Monotonicity of $x^*(\Lambda)$

We now show that $x^*(\Lambda)$ is strictly increasing at every tariff in the policy domain. All equilibrium objects are evaluated along the noncooperative locus, on which Foreign plays its best response $x^*(t)$, and total derivatives account for the equilibrium adjustment of prices and of Foreign's tax. The argument proceeds in three steps. Step 1 establishes that Foreign's optimal tax falls with the tariff. Step 2 establishes that aggregate import dependence also falls, so that $t \mapsto \Lambda(t)$ is a strictly decreasing bijection onto its range. Step 3 combines the two monotonicity results to sign $\mathrm{d}x^*(\Lambda)/\mathrm{d}\Lambda$.

Step 1: Foreign's optimal tax falls with the tariff. Differentiate the optimality condition $x^*\varepsilon^S_{HF} = 1$ in logarithms along the locus. Because $1/x^* = \varepsilon^S_{HF}$, the Lerner form (77) converts the tax movement into gross-factor units via Home's demand elasticity, $\mathrm{d}\ln x^* = \frac{1+x^*}{x^*}\,\mathrm{d}\ln(1+x^*) = \varepsilon^D_{FH}\,\mathrm{d}\ln(1+x^*)$. The movement of the elasticity follows from its partials (71) together with the equilibrium share response (68), in which both wedges now move. Substituting and collecting terms,

$$\left[\varepsilon^S_{HF}\varepsilon^D_{FH} + \left(\frac{\epsilon\delta^2}{1+t} - \frac{1-\beta}{\beta}\right)\lambda_{FH}\lambda_{HH}\,\varepsilon^X\right]\frac{\mathrm{d}\ln(1+x^*)}{\mathrm{d}\ln(1+t)} = -\lambda_{FH}\lambda_{HH}\left[\left(\frac{\epsilon\delta^2}{1+t} - \frac{1-\beta}{\beta}\right)\varepsilon^M + \frac{\epsilon\delta^2}{1+t}\right]. \tag{78}$$

By (71) the round bracket is strictly positive on the policy domain, and therefore so are the left-hand bracket—the nonsingularity of the differentiated optimality condition—and the right-hand bracket. Hence

$$\frac{\mathrm{d}x^*(t)}{\mathrm{d}t} < 0 \qquad \text{for every } t \in \left[0, \tfrac{\epsilon\beta}{1-\beta} - 1\right]. \tag{79}$$

A higher tariff contracts import dependence, and lower dependence makes Home's residual import demand more elastic; the tariff also raises the elasticity directly through the revenue channel. Both forces compress Foreign's Lerner markup.

Step 2: Aggregate import dependence falls with the tariff. Insert (78) into (68), evaluated along the locus: $\mathrm{d}\ln[\Lambda/(1-\Lambda)] = -\varepsilon^M\,\mathrm{d}\ln(1+t) - \varepsilon^X\,\mathrm{d}\ln(1+x^*)$. The terms

involving $\frac{\epsilon\delta^2}{1+t} - \frac{1-\beta}{\beta}$ cancel exactly, leaving

$$\frac{\mathrm{d}}{\mathrm{d}\ln(1+t)} \ln \frac{\Lambda(t)}{1-\Lambda(t)} = -\frac{\varepsilon^M \varepsilon^S_{HF} \varepsilon^D_{FH} - \varepsilon^X \frac{\epsilon\lambda_{FH}\lambda_{HH}\delta^2}{1+t}}{\varepsilon^S_{HF}\varepsilon^D_{FH} + \left(\frac{\epsilon\delta^2}{1+t} - \frac{1-\beta}{\beta}\right)\lambda_{FH}\lambda_{HH}\,\varepsilon^X}. \tag{80}$$

The denominator coincides with the left-hand bracket of (78) and is therefore strictly positive by Step 1. For the numerator, note first the identity

$$\frac{\lambda_{FH}\lambda_{HH}\delta^2}{1+t} = \lambda_{HH}\delta\left[1-\lambda_{HH}\delta\right], \tag{81}$$

which follows from $1-\lambda_{HH}\delta = \lambda_{FH}\delta/(1+t)$. Next, by (68) and the ordering (70),

$$\frac{\varepsilon^M}{\varepsilon^X} = \frac{\varepsilon^D_{HF} - \lambda_{HH}\delta}{\varepsilon^D_{HF} - \frac{\lambda_{HF}\delta_F}{1+x}} \;\geq\; \frac{1-\lambda_{HH}\delta}{1-\frac{\lambda_{HF}\delta_F}{1+x}} \;>\; 1-\lambda_{HH}\delta, \tag{82}$$

where the first inequality follows by cross-multiplication—the difference of the two sides has the sign of $\left(\varepsilon^D_{HF}-1\right)\left[\lambda_{HH}\delta - \frac{\lambda_{HF}\delta_F}{1+x}\right] \geq 0$—and the second holds strictly because the denominator of its left-hand side lies in $(0,1)$ by (69). Finally, $\epsilon\lambda_{HH}\delta < \varepsilon^S_{HF}$ term by term and Home's $\varepsilon^D_{FH} > 1$. Chaining the three bounds,

$$\varepsilon^M\varepsilon^S_{HF}\varepsilon^D_{FH} \;>\; \varepsilon^X\left[1-\lambda_{HH}\delta\right]\varepsilon^S_{HF}\varepsilon^D_{FH} \;>\; \varepsilon^X\left[1-\lambda_{HH}\delta\right]\epsilon\lambda_{HH}\delta = \varepsilon^X\frac{\epsilon\lambda_{FH}\lambda_{HH}\delta^2}{1+t}.$$

The numerator of (80) is therefore strictly positive, and

$$\frac{\mathrm{d}\Lambda(t)}{\mathrm{d}t} < 0 \qquad \text{for every } t \in \left[0, \tfrac{\epsilon\beta}{1-\beta} - 1\right]. \tag{83}$$

In words, Foreign's equilibrium easing of its export tax dampens but never overturns the tariff's direct contraction of Home's import dependence. Three remarks are in order. First, setting the export tax to zero instead of $x^*(t)$ in (68) gives $\mathrm{d}\lambda_{FH}(0;t)/\mathrm{d}t = -\lambda_{FH}\lambda_{HH}\varepsilon^M/(1+t) < 0$: the cooperative locus slopes down as well. Second, $t \mapsto \Lambda(t)$ is a strictly decreasing, continuously differentiable bijection onto its range: the inverse $t(\Lambda)$ exists, $x^*(\Lambda) \equiv x^*(t(\Lambda))$ is well defined, and derivatives at $t=0$ and at the cap have their one-sided meaning. Because this step rests only on the share-form optimality condition (75) and the response system, and not on the Λ-indexed formula (76) that it underwrites, the forward reference in the proof of Lemma 1 is discharged without circularity. Third, the strictly positive left-hand bracket of (78) is precisely the nonsingularity

required by the implicit-function theorem along the locus.

Step 3: Monotonicity of x^*. Combining (79) and (83) through the bijection of Step 2,

$$\frac{\mathrm{d}x^*(\Lambda)}{\mathrm{d}\Lambda} = \frac{\mathrm{d}x^*(t)/\mathrm{d}t}{\mathrm{d}\Lambda(t)/\mathrm{d}t} > 0, \tag{84}$$

which is Lemma 2. Intuitively, moving down the locus to a lower tariff raises aggregate import dependence and, by (71), makes Home's residual import demand less elastic: the substitution component $\epsilon\lambda_{HH}\delta$ of (66) falls by more than the roundabout-production component $\frac{1-\beta}{\beta}\lambda_{FH}$ rises. Assumption 2, reinforced on the policy domain by the tariff cap, guarantees that inequality at every point. Foreign's market power grows, and its unilaterally optimal export tax rises with import dependence. Steps 1–3 establish Lemma 2; with the optimality condition and formula (76) of Lemma 1, the proofs of both lemmas are complete.

C Monotonicity of Strategic Risk π

This appendix proves Lemmas 3 and 4. We maintain Assumptions 1 and 2, the standing existence, interiority, and differentiability conditions, the policy domain $t, x \geq 0$, and the laissez-faire share restriction $\lambda_{HF}(0;0) \leq \lambda_{FF}(0;0)$. We allow the value-added share to take any value $\beta \in (0,1]$. The tariff ceiling in the lemma is most conveniently written without dividing by $1-\beta$:

$$(1+t)\frac{1-\beta}{\beta} \leq \epsilon. \tag{85}$$

For $\beta < 1$, this is equivalent to $t \leq \epsilon\beta/(1-\beta) - 1$; for $\beta = 1$, it imposes no additional upper restriction on t. We suppress the policy arguments $(x;t)$ on shares and elasticities whenever doing so causes no ambiguity.

Proof of Lemma 3

Maintain Assumptions 1 and 2, the trade policy domain $t, x \geq 0$, and the size restrictions that ensure $\lambda_{HF}(0;0) \leq \lambda_{FF}(0;0)$ under laissez-faire. We prove the stated result under the weaker tariff ceiling $(1+t)\frac{1-\beta}{\beta} \leq \epsilon + 1$. For $\beta < 1$, this is equivalent to $t \leq \frac{\beta(\epsilon+1)}{1-\beta} - 1$; for $\beta = 1$, it imposes no additional upper restriction on t. We suppress the policy arguments $(x;t)$ for shares and elasticity variables whenever doing so causes no ambiguity. As in the main text, ∂_t denotes an equilibrium derivative holding x fixed, and ∂_x denotes

an equilibrium derivative holding t fixed. With $p_H = 1$, the price paid by Home buyers for the Foreign good is

$$\widetilde{p}_{FH} = (1+t)(1+x)d_{FH}p_F.$$

The two-country adding-up identities are $\lambda_{HH} = 1 - \lambda_{FH}$ and $\lambda_{FF} = 1 - \lambda_{HF}$. The CES share system therefore gives

$$\frac{\partial \ln \lambda_{FH}}{\partial \ln \widetilde{p}_{FH}} = -\epsilon\lambda_{HH}, \qquad \frac{\partial \ln \lambda_{HF}}{\partial \ln p_F} = \epsilon\lambda_{FF}, \qquad \frac{\partial \ln \lambda_{FF}}{\partial \ln p_F} = -\epsilon\lambda_{HF}. \tag{86}$$

Marginal-cost pricing at the numeraire requires $\beta\, d\ln w_H + (1-\beta)\, d\ln \widetilde{P}_H = 0$. Because the CES price index satisfies $\partial \ln \widetilde{P}_H / \partial \ln \widetilde{p}_{FH} = \lambda_{FH}$,

$$\frac{\partial \ln w_H}{\partial \ln \widetilde{p}_{FH}} = -\frac{1-\beta}{\beta}\lambda_{FH}.$$

Home rebates tariff revenue and, under Assumption 1, gross expenditure is

$$E_H = \frac{w_H L_H}{\beta}\,\delta(x;t), \qquad \delta(x;t) = \left[1 - \frac{t}{1+t}\lambda_{FH}(x;t)\right]^{-1} \geq 1.$$

The inequality follows from $t \geq 0$ and the interiority of expenditure shares. Home import demand is $Q_{FH} = \lambda_{FH}E_H/\widetilde{p}_{FH}$. Differentiating this demand schedule with respect to the buyer price while holding t fixed gives

$$\begin{aligned}
\varepsilon^D_{FH} &\equiv -\left.\frac{\partial \ln Q_{FH}}{\partial \ln \widetilde{p}_{FH}}\right|_t \\
&= 1 + \epsilon\lambda_{HH} + \frac{1-\beta}{\beta}\lambda_{FH} + \frac{t}{1+t}\lambda_{FH}\delta\,\epsilon\lambda_{HH} \\
&= 1 + \epsilon\lambda_{HH}\delta + \frac{1-\beta}{\beta}\lambda_{FH} > 1.
\end{aligned} \tag{87}$$

The second equality uses $1 + [t/(1+t)]\lambda_{FH}\delta = \delta$. Thus the elasticity contains the mechanical unit response, the CES substitution response inclusive of the tariff-revenue feedback, and the roundabout-production response.

At a fixed buyer price, λ_{FH} and w_H do not change with the tariff. Gross expenditure then moves only because rebated tariff revenue changes. Since $d[t/(1+t)]/d\ln(1+t) = 1/(1+t)$, direct differentiation yields

$$\left.\frac{\partial \ln E_H}{\partial \ln(1+t)}\right|_{\widetilde{p}_{FH}} = \frac{\delta\lambda_{FH}}{1+t} = \frac{\lambda_{FH}}{1+t\lambda_{HH}} \in (0,1). \tag{88}$$

This is the outward shift in Home import demand generated by the tariff rebate.

We next derive Foreign's export-supply schedule using the same equilibrium accounting. Balanced trade implies

$$(1+x)d_{FH}p_F Q_{FH} = \lambda_{HF}E_F.$$

Foreign rebates export-tax revenue. Assumption 1 and the gross-expenditure identity therefore give $E_F = \frac{w_F L_F}{\beta} + \frac{x}{1+x}\lambda_{HF}E_F$. Solving this identity gives $E_F = (w_F L_F/\beta)(1+x)/(1+x\lambda_{FF})$. The Foreign wage relation implied by marginal-cost pricing is $\frac{w_F}{p_F} = \bar{\varphi}_F^{1/\beta}\lambda_{FF}^{-\frac{1-\beta}{\beta\epsilon}}$. Substituting these two relations into balanced trade gives the export-supply schedule $Q_{FH}(p_F, x) = \frac{\bar{\varphi}_F^{1/\beta}L_F}{\beta d_{FH}} \frac{\lambda_{HF}(p_F)\lambda_{FF}(p_F)^{-\frac{1-\beta}{\beta\epsilon}}}{1+x\lambda_{FF}(p_F)}$. Its relevant log derivatives are

$$\begin{aligned} \varepsilon_{FH}^S \equiv \left.\frac{\partial \ln Q_{FH}}{\partial \ln p_F}\right|_x &= \epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}} + \frac{1-\beta}{\beta}\lambda_{HF} > 0, \\ -\left.\frac{\partial \ln Q_{FH}}{\partial \ln(1+x)}\right|_{p_F} &= \frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}} \in (0,1). \end{aligned} \tag{89}$$

We now equate the changes in Home import demand and Foreign export supply. At a fixed export tax, the two schedules and the buyer-price identity imply

$$\begin{aligned} d\ln Q_{FH} &= -\varepsilon_{FH}^D\, d\ln \widetilde{p}_{FH} + \frac{\delta\lambda_{FH}}{1+t}\, d\ln(1+t) = \varepsilon_{FH}^S\, d\ln p_F, \\ d\ln \widetilde{p}_{FH} &= d\ln p_F + d\ln(1+t). \end{aligned}$$

Solving these two equations gives

$$\begin{aligned} \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+t)} &= \frac{\varepsilon_{FH}^S + \frac{\delta\lambda_{FH}}{1+t}}{\varepsilon_{FH}^S + \varepsilon_{FH}^D} \in (0,1), \\ \frac{\partial \ln p_F}{\partial \ln(1+t)} &= -\frac{\varepsilon_{FH}^D - \frac{\delta\lambda_{FH}}{1+t}}{\varepsilon_{FH}^S + \varepsilon_{FH}^D} < 0. \end{aligned} \tag{90}$$

The signs follow from $0 < \delta\lambda_{FH}/(1+t) < 1 < \varepsilon_{FH}^D$. Thus the tariff raises Home's buyer price but lowers Foreign's producer price.

At a fixed Home tariff, equating the demand and supply responses to the export tax gives

$$\frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} = \frac{\varepsilon_{FH}^S + \frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}}{\varepsilon_{FH}^S + \varepsilon_{FH}^D} \in (0,1). \tag{91}$$

Pass-through is positive because both terms in the numerator are positive. It is less than one because $(1+x)\lambda_{FF}/(1+x\lambda_{FF}) < 1 < \varepsilon^D_{FH}$. Consequently,

$$\frac{\partial \ln p_F}{\partial \ln(1+x)} = \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} - 1 < 0. \tag{92}$$

Equations (86), (90), (91), and (92) now sign all four share responses:

$$\partial_t \lambda_{FH} < 0, \qquad \partial_t \lambda_{HF} < 0, \qquad \partial_x \lambda_{FH} < 0, \qquad \partial_x \lambda_{HF} < 0. \tag{93}$$

For example, either policy raises Home's buyer price and therefore lowers λ_{FH}; either policy also lowers Foreign's producer price and therefore lowers λ_{HF}. The laissez-faire share restriction consequently extends to the entire policy region:

$$\lambda_{HF}(x;t) \le \lambda_{HF}(0;0) \le \lambda_{FF}(0;0) \le \lambda_{FF}(x;t), \qquad \lambda_{HF}(x;t) \le \lambda_{FF}(x;t). \tag{94}$$

The third inequality follows from $\lambda_{FF} = 1 - \lambda_{HF}$ and (93). We record the final ordering because it will sign the difficult multiplier inequality below.

We next derive Foreign's marginal payoff from the export tax, including the general-equilibrium response of p_F. Under Assumption 1, Foreign net income is labor income plus export-tax revenue. Combining this identity with the expenditure accounting above gives

$$Y_F = w_F L_F \frac{\beta + (1-\beta)\frac{x}{1+x}\lambda_{HF}}{\beta\left[1 - \frac{x}{1+x}\lambda_{HF}\right]}.$$

Marginal-cost pricing and the Foreign CES price index imply $w_F/\widetilde{P}_F \propto \lambda_{FF}^{-1/(\beta\epsilon)}$. Hence, up to a constant independent of (x,t),

$$v_F(x;t) = -\frac{1}{\beta\epsilon}\ln\lambda_{FF} + \ln\left[\beta + (1-\beta)\frac{x}{1+x}\lambda_{HF}\right] - \ln\left[1 - \frac{x}{1+x}\lambda_{HF}\right] \tag{95}$$

The export tax affects (95) directly through $x/(1+x)$ and indirectly through p_F and the two Foreign expenditure shares. The chain rule makes these two channels explicit:

$$\partial_x v_F = \left.\frac{\partial v_F}{\partial x}\right|_{p_F} + \left.\frac{\partial v_F}{\partial \ln p_F}\right|_{x} \frac{1}{1+x}\left[\frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} - 1\right]. \tag{96}$$

Direct differentiation, using $d[x/(1+x)]/dx = (1+x)^{-2}$ and (86), gives

$$\left.\frac{\partial v_F}{\partial x}\right|_{p_F} = \frac{\lambda_{HF}}{(1+x\lambda_{FF})\{\beta[1+x-x\lambda_{HF}]+x\lambda_{HF}\}}, \tag{97}$$

$$\left.\frac{\partial v_F}{\partial \ln p_F}\right|_{x} = \frac{\lambda_{HF}}{\beta} + \frac{\epsilon x\lambda_{FF}\lambda_{HF}(1+x)}{(1+x\lambda_{FF})\{\beta[1+x-x\lambda_{HF}]+x\lambda_{HF}\}}. \tag{98}$$

Substitute (97) and (98) into (96). The term $\frac{(1+x)\{\beta[1+x-x\lambda_{HF}]+x\lambda_{HF}\}}{\lambda_{HF}}$ is strictly positive by $\beta > 0$, $x \geq 0$, and interiority. Multiplying $\partial_x v_F$ by this quantity makes the collection of terms transparent. The coefficient on the equilibrium producer-price response becomes

$$\frac{\beta[1+x-x\lambda_{HF}]+x\lambda_{HF}}{\beta} + \frac{\epsilon x\lambda_{FF}(1+x)}{1+x\lambda_{FF}} = 1+x+x\varepsilon^S_{FH}, \tag{99}$$

where the equality uses (89). Equation (91) also implies

$$\frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} - 1 = -\frac{\varepsilon^D_{FH} - \frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}}{\varepsilon^S_{FH}+\varepsilon^D_{FH}}. \tag{100}$$

Putting the resulting terms over the common denominator $\varepsilon^S_{FH}+\varepsilon^D_{FH}$ yields the formal identity

$$\begin{aligned}&\frac{(1+x)\{\beta[1+x-x\lambda_{HF}]+x\lambda_{HF}\}}{\lambda_{HF}}\,\partial_x v_F(x;t)\\ &= \frac{1+x}{1+x\lambda_{FF}} - \left(1+x+x\varepsilon^S_{FH}\right)\frac{\varepsilon^D_{FH} - \frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}}{\varepsilon^S_{FH}+\varepsilon^D_{FH}} = \frac{\varepsilon^S_{FH} + \frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}}{\varepsilon^S_{FH}+\varepsilon^D_{FH}}\left[1 - x\left(\varepsilon^D_{FH}-1\right)\right].\end{aligned} \tag{101}$$

The fraction multiplying the markup gap in the final equality is precisely the export-tax pass-through in (91). Dividing both sides of (101) by the positive multiplier and noting that $\beta[1+x-x\lambda_{HF}]+x\lambda_{HF} = \beta\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]$ gives the Foreign's marginal payoff from export tax imposition

$$\partial_x v_F(x;t) = \underbrace{\frac{\lambda_{HF}}{\beta(1+x)\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]}\frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)}}_{\text{export tax base + income feedback + pass-through >0}} \times \underbrace{\left[1 - x\left(\varepsilon^D_{FH}-1\right)\right]}_{\text{optimal markup gap}} \tag{102}$$

The first term is the balanced trade representation of Foreign's export tax base after the income feedback from rebated tax revenue, multiplied by the fraction of the tax passed through to Home buyers. The second term is positive exactly when the current export

tax lies below Foreign's inverse-elasticity benchmark, and it vanishes at the interior optimum. Because the first term in (102) is strictly positive, an interior Foreign optimum satisfies

$$x^*(t)\left[\varepsilon^D_{FH}(x^*(t);t)-1\right]=1, \tag{103}$$

which is analogous to the optimal rate presented under Lemma 1. We now show that the optimal markup gap is nonnegative on $x\in[0,x^*(t)]$ and nonincreasing in t.

Regard (87) as a function of λ_{FH} and t. Differentiating first with respect to the share at a fixed tariff, and then with respect to the tariff at a fixed share, gives

$$\frac{\partial\varepsilon^D_{FH}}{\partial\lambda_{FH}}=\frac{1-\beta}{\beta}-\frac{\epsilon\delta^2}{1+t'},\qquad \left.\frac{\partial\varepsilon^D_{FH}}{\partial t}\right|_{\lambda_{FH}}=\frac{\epsilon\lambda_{HH}\lambda_{FH}\delta^2}{(1+t)^2}>0. \tag{104}$$

Under the stated tariff ceiling, the first partial derivative in (104) need not be nonpositive. What matters for the proof are two equilibrium properties: the response of ε^D_{FH} to the tariff and the monotonicity of $x(\varepsilon^D_{FH}-1)$ with respect to the export tax. We establish both directly.

Equation (86) and the tariff pass-through in (90) imply

$$\partial_t\lambda_{FH}=-\frac{\epsilon\lambda_{HH}\lambda_{FH}}{1+t}\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+t)}<0.$$

Substituting this share response and (104) into the chain rule gives the exact total equilibrium response

$$\begin{aligned}
\partial_t\varepsilon^D_{FH}&=\left(\frac{1-\beta}{\beta}-\frac{\epsilon\delta^2}{1+t}\right)\left[-\frac{\epsilon\lambda_{HH}\lambda_{FH}}{1+t}\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+t)}\right]+\frac{\epsilon\lambda_{HH}\lambda_{FH}\delta^2}{(1+t)^2}\\
&=\frac{\epsilon\lambda_{HH}\lambda_{FH}}{(1+t)^2}\left\{\delta^2\left[1+\epsilon\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+t)}\right]-(1+t)\frac{1-\beta}{\beta}\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+t)}\right\}\\
&\geq\frac{\epsilon\lambda_{HH}\lambda_{FH}}{(1+t)^2}\left\{1-\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+t)}\left[(1+t)\frac{1-\beta}{\beta}-\epsilon\right]\right\}\\
&\geq\frac{\epsilon\lambda_{HH}\lambda_{FH}}{(1+t)^2}\left[1-\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+t)}\right]>0.
\end{aligned} \tag{105}$$

The first inequality uses $\delta^2\geq 1$ and the positivity of tariff pass-through. The second uses $(1+t)(1-\beta)/\beta-\epsilon\leq 1$, which is the tariff ceiling, together with the positivity of tariff pass-through. The final inequality is strict because tariff pass-through lies strictly between zero and one by (90), and the remaining terms are positive. Thus, at a fixed

x, a higher Home tariff makes the residual import demand facing Foreign strictly more elastic. The direct increase in elasticity generated by the tariff-revenue feedback more than offsets any opposing effect operating through the equilibrium decline in λ_{FH}. This is an equilibrium derivative, not a derivative along Foreign's best response.

We next establish the monotonicity needed to sign the optimal markup gap. At a fixed tariff, equations (86) and (91) give

$$\partial_x \lambda_{FH} = -\frac{\epsilon \lambda_{HH} \lambda_{FH}}{1+x} \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} < 0.$$

Because ε^D_{FH} depends on x through λ_{FH} when t is held fixed, (104) therefore implies

$$\partial_x \left(\varepsilon^D_{FH} - 1\right) = -\frac{\epsilon \lambda_{HH} \lambda_{FH}}{1+x} \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} \left[\frac{1-\beta}{\beta} - \frac{\epsilon \delta^2}{1+t}\right]. \tag{106}$$

Applying the product rule and substituting (106) yields

$$\frac{\partial}{\partial x}\left\{x\left(\varepsilon^D_{FH} - 1\right)\right\} = \varepsilon^D_{FH} - 1 - \frac{x}{1+x} \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} \epsilon \lambda_{HH} \lambda_{FH} \left[\frac{1-\beta}{\beta} - \frac{\epsilon \delta^2}{1+t}\right]. \tag{107}$$

There are two cases. If $\frac{1-\beta}{\beta} - \frac{\epsilon\delta^2}{1+t} \leq 0$ then the second term on the right-hand side of (107) is nonnegative after accounting for its leading minus sign. The derivative is therefore strictly positive because $\varepsilon^D_{FH} - 1 > 0$. Suppose instead that $\frac{1-\beta}{\beta} - \frac{\epsilon\delta^2}{1+t} > 0$. Using $\delta \geq 1$, the tariff ceiling, and $t \geq 0$ gives

$$0 < \frac{1-\beta}{\beta} - \frac{\epsilon\delta^2}{1+t} \leq \frac{1-\beta}{\beta} - \frac{\epsilon}{1+t} = \frac{(1+t)\frac{1-\beta}{\beta} - \epsilon}{1+t} \leq \frac{1}{1+t} \leq 1.$$

Moreover, $x/(1+x) < 1$, and export-tax pass-through is strictly less than one by (91). Substituting these bounds into (107) gives

$$\begin{aligned}\frac{\partial}{\partial x}\left\{x\left(\varepsilon^D_{FH} - 1\right)\right\} &> \varepsilon^D_{FH} - 1 - \epsilon \lambda_{HH} \lambda_{FH} \\ &= \epsilon \lambda_{HH}\left(\delta - \lambda_{FH}\right) + \frac{1-\beta}{\beta} \lambda_{FH} > 0.\end{aligned}$$

The equality uses (87). The final inequality follows because Assumption 2 gives $\epsilon > 0$, interiority gives $\lambda_{HH} > 0$ and $\lambda_{FH} < 1$, $\delta \geq 1$, and $\beta \in (0,1]$ makes $(1-\beta)/\beta$ nonnegative. When $\beta = 1$, the expression in square brackets in (107) is strictly negative, so only the first case can arise.

Thus, in either case,

$$\frac{\partial}{\partial x}\left\{x\left(\varepsilon_{FH}^{D}-1\right)\right\}>0. \tag{108}$$

At $x=0$, the derivative in (107) reduces to $\varepsilon_{FH}^{D}-1>0$; thus the boundary is included. Thus $x(\varepsilon_{FH}^{D}-1)$ is strictly increasing even though $\partial\varepsilon_{FH}^{D}/\partial\lambda_{FH}$ need not be nonpositive. The expression starts at zero and, by (103), equals one at $x^{*}(t)$. Therefore,

$$1-x\left(\varepsilon_{FH}^{D}-1\right)\geq 0 \qquad \text{for every } x\in[0,x^{*}(t)]. \tag{109}$$

For $x>0$ in this interval, the same inequality gives $\varepsilon_{FH}^{D}\leq 1+\frac{1}{x}$. Finally, (105) implies

$$\partial_t\left[1-x\left(\varepsilon_{FH}^{D}-1\right)\right]=-x\,\partial_t\varepsilon_{FH}^{D}\leq 0, \tag{110}$$

with strict inequality whenever $x>0$.

It remains to show that the positive multiplier in (102) also falls with the Home tariff. Substituting (89) into the pass-through formula (91) formally identifies that multiplier as

$$\begin{aligned}&\frac{\lambda_{HF}}{\beta(1+x)\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]}\frac{\partial\ln\widetilde{p}_{FH}}{\partial\ln(1+x)}\\&\quad=\frac{\lambda_{HF}}{\beta(1+x)\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]}\times\frac{(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}}{\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon_{FH}^{D}}\end{aligned} \tag{111}$$

The left-hand side is exactly the first term in (102); the right-hand side replaces pass-through by its elasticity representation. Thus the first fraction on the right is the adjusted export-tax base, while the second is pass-through.

Holding x and λ_{HF} fixed, the multiplier in (111) is strictly decreasing in ε_{FH}^{D} because this elasticity enters only through the strictly positive denominator of the second fraction. We next show that, on $x\in[0,x^{*}(t)]$, holding x and ε_{FH}^{D} fixed, the same multiplier is strictly increasing in λ_{HF}. Using $\lambda_{FF}=1-\lambda_{HF}$, logarithmic differentiation of the

complete right-hand side of (111) gives

$$\frac{\partial}{\partial\lambda_{HF}}\ln\left\{\frac{\lambda_{HF}}{\beta(1+x)\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]}\frac{(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}}{\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon^D_{FH}}\right\}\Bigg|_{x,\varepsilon^D_{FH}}$$

$$=\frac{1+x}{\lambda_{HF}\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]}+\frac{\frac{1-\beta}{\beta}-(\epsilon+1)\frac{1+x}{(1+x\lambda_{FF})^2}}{(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}}-\frac{\frac{1-\beta}{\beta}-\epsilon\frac{1+x}{(1+x\lambda_{FF})^2}}{\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon^D_{FH}}. \tag{112}$$

The right-hand side of (112) can be rearranged as

$$\frac{1+x}{\lambda_{HF}\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]}+\frac{\frac{1-\beta}{\beta}\left(\varepsilon^D_{FH}-\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}\right)}{\left[(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}\right]\left[\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon^D_{FH}\right]}$$
$$-\frac{1+x}{(1+x\lambda_{FF})^2}\frac{(\epsilon+1)\varepsilon^D_{FH}+\frac{1-\beta}{\beta}\lambda_{HF}}{\left[(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}\right]\left[\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon^D_{FH}\right]}. \tag{113}$$

The middle term is nonnegative because $\varepsilon^D_{FH}>1>\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}$. Because the middle term is nonnegative, the expression (112) is bounded below by the sum of the remaining two terms, which is positive if

$$(1+x\lambda_{FF})^2\left[(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}\right]\times\left[\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon^D_{FH}\right]$$
$$>\lambda_{HF}\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]\left[(\epsilon+1)\varepsilon^D_{FH}+\frac{1-\beta}{\beta}\lambda_{HF}\right]. \tag{114}$$

The left-hand side of (114) satisfies

$$(1+x\lambda_{FF})^2\left[(\epsilon+1)\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}\right]\times\left[\epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}}+\frac{1-\beta}{\beta}\lambda_{HF}+\varepsilon^D_{FH}\right]$$
$$>\lambda_{FF}(\epsilon+1)(1+x)(1+x\lambda_{FF})\left[\varepsilon^D_{FH}+\frac{1-\beta}{\beta}\lambda_{HF}\right]$$
$$\geq\lambda_{HF}(\epsilon+1)(1+x)(1+x\lambda_{FF})\left[\varepsilon^D_{FH}+\frac{1-\beta}{\beta}\lambda_{HF}\right]. \tag{115}$$

The first inequality retains $(\epsilon+1)(1+x)\lambda_{FF}/(1+x\lambda_{FF})$ from the first bracket and

$\varepsilon^D_{FH} + [(1-\beta)/\beta]\lambda_{HF}$ from the second. It is strict because the omitted CES term $\epsilon(1+x)\lambda_{FF}/(1+x\lambda_{FF})$ is positive. The second inequality uses (94). Comparing (114) and (115), it is sufficient to prove that

$$(\epsilon+1)(1+x)(1+x\lambda_{FF})\left[\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right] \geq \left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]\left[(\epsilon+1)\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right]. \tag{116}$$

We prove (116) by exhausting the possible cases. If $x = 0$, the inequality holds trivially as $\epsilon\frac{1-\beta}{\beta}\lambda_{HF} \geq 0$. Let $x > 0$. First suppose $\frac{1-\beta}{\beta}\lambda_{HF} \leq (1+x)\lambda_{FF}$.Then

$$(1+x)(1+x\lambda_{FF}) \geq 1+x+\frac{1-\beta}{\beta}x\lambda_{HF}, \tag{117}$$

because the left-hand side minus the right-hand side is $x\left[(1+x)\lambda_{FF} - \frac{1-\beta}{\beta}\lambda_{HF}\right] \geq 0$. Likewise,

$$(\epsilon+1)\left[\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right] \geq (\epsilon+1)\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}, \tag{118}$$

because the difference is $\epsilon[(1-\beta)/\beta]\lambda_{HF} \geq 0$. All four terms in (117) and (118) are positive, so multiplying the two inequalities preserves their direction and proves (116) in this case.

It remains to consider the complementary case $\frac{1-\beta}{\beta}\lambda_{HF} > (1+x)\lambda_{FF}$. This case cannot arise when $\beta = 1$. The cross-term comparison used above no longer applies, so we evaluate the left-minus-right difference in (116) at the least favorable values allowed by Assumption 2 and the nonnegative optimal markup gap. This is solely an algebraic bounding device: hold x, β, and λ_{HF} fixed and regard the difference as a function of ϵ and ε^D_{FH}. Its derivative with respect to ϵ is

$$\begin{aligned}\frac{\partial}{\partial\epsilon}\Bigg\{&(\epsilon+1)(1+x)(1+x\lambda_{FF})\left[\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right] \\ &- \left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]\left[(\epsilon+1)\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right]\Bigg\} \\ &= \varepsilon^D_{FH}x(1+x)\lambda_{FF} + \frac{1-\beta}{\beta}\lambda_{HF}\left[(1+x)(1+x\lambda_{FF}) - x\varepsilon^D_{FH}\right] > 0.\end{aligned} \tag{119}$$

The bracket multiplying $[(1-\beta)/\beta]\lambda_{HF}$ in (119) is nonnegative because

$$(1+x)(1+x\lambda_{FF}) - x\varepsilon^D_{FH} \geq 1 + x - x\varepsilon^D_{FH} = 1 - x(\varepsilon^D_{FH} - 1) \geq 0, \tag{120}$$

where the final inequality is (109). The term $\varepsilon^D_{FH} x(1+x)\lambda_{FF}$ in (119) is strictly positive because $x > 0$, $\lambda_{FF} > 0$ by maintained interiority, and $\varepsilon^D_{FH} > 1$. The difference is therefore strictly increasing in ϵ. Assumption 2, $\epsilon > (1-\beta)/\beta$, implies that its maintained value is strictly larger than its value at the boundary $\epsilon = (1-\beta)/\beta$. At that boundary, $\epsilon = \frac{1-\beta}{\beta}$, the derivative of the same left-minus-right difference with respect to ε^D_{FH} is

$$\begin{aligned}
&\frac{\partial}{\partial \varepsilon^D_{FH}} \Bigg\{ (\epsilon+1)(1+x)(1+x\lambda_{FF}) \left[\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right] \\
&\qquad - \left[1 + x + \frac{1-\beta}{\beta} x\lambda_{HF}\right] \left[(\epsilon+1)\varepsilon^D_{FH} + \frac{1-\beta}{\beta}\lambda_{HF}\right] \Bigg\} \Bigg|_{\epsilon = \frac{1-\beta}{\beta}} \\
&\quad = \qquad \left(1 + \frac{1-\beta}{\beta}\right) x \left[(1+x)\lambda_{FF} - \frac{1-\beta}{\beta}\lambda_{HF}\right] < 0,
\end{aligned} \tag{121}$$

where the inequality follows from the case under consideration satisfying $\frac{1-\beta}{\beta}\lambda_{HF} > (1+x)\lambda_{FF}$. Because the difference is decreasing in ε^D_{FH} and $\varepsilon^D_{FH} \leq 1 + 1/x$, its value is bounded below by the value at $\varepsilon^D_{FH} = 1 + 1/x$. Substituting both boundary values gives

$$\begin{aligned}
&\left(1 + \frac{1-\beta}{\beta}\right)(1+x)(1+x\lambda_{FF}) \left[1 + \frac{1}{x} + \frac{1-\beta}{\beta}\lambda_{HF}\right] \\
&\quad - \left[1 + x + \frac{1-\beta}{\beta} x\lambda_{HF}\right] \left[\left(1 + \frac{1-\beta}{\beta}\right)\left(1 + \frac{1}{x}\right) + \frac{1-\beta}{\beta}\lambda_{HF}\right] \\
&\quad = \left[1 + x + \frac{1-\beta}{\beta} x\lambda_{HF}\right] \left\{\lambda_{FF}\left[1 + x + \frac{1-\beta}{\beta} x\right] + \frac{1-\beta}{\beta}(\lambda_{FF} - \lambda_{HF})\right\} > 0.
\end{aligned}$$

The term preceding the braces in the above equation is positive. Inside the braces, $\lambda_{FF} > 0$ by the maintained interiority condition and $\lambda_{FF} - \lambda_{HF} \geq 0$ by (94). Thus (116) holds in the complementary case and hence in every case.

Combining (116), (115), and (113) shows that the logarithmic derivative in (112) is strictly positive. Because the multiplier itself is positive, its level derivative with respect to λ_{HF} is also strictly positive. Thus, under the maintained conditions and on $x \in [0, x^*(t)]$, the positive multiplier in (111) is strictly increasing in λ_{HF} and strictly decreasing in ε^D_{FH}. A higher Home tariff lowers λ_{HF} by (93) and raises ε^D_{FH} by (105). At a fixed x, the right-hand side of (111) depends on t only through these two objects

because $\lambda_{FF} = 1 - \lambda_{HF}$. The chain rule therefore implies that the full positive multiplier is strictly decreasing in t. Finally, differentiate (102) with respect to t, holding x fixed. The product rule gives

$$\partial_{xt} v_F(x;t) = \partial_t \left\{ \frac{\lambda_{HF}}{\beta(1+x)\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]} \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} \right\} \left[1 - x(\varepsilon^D_{FH} - 1)\right] - \frac{x\lambda_{HF}}{\beta(1+x)\left[1+x+\frac{1-\beta}{\beta}x\lambda_{HF}\right]} \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} \partial_t \varepsilon^D_{FH} < 0. \tag{122}$$

The derivative in braces in the first term is strictly negative, while the optimal markup gap is nonnegative by (109). The coefficient multiplying $-x\partial_t \varepsilon^D_{FH}$ is positive, and $\partial_t \varepsilon^D_{FH} > 0$ by (105). At $x = 0$, the optimal markup gap equals one, so the first term is strictly negative. For $0 < x < x^*(t)$, the second term is strictly negative. At $x = x^*(t)$, the optimal markup gap is zero, but the second term remains strictly negative. Hence,

$$\partial_{xt} v_F(x;t) < 0 \qquad \text{for every } x \in [0, x^*(t)].$$

□

Proof of Lemma 4

Along the non-cooperative equilibrium branch,

$$\Lambda = \lambda_{FH}\big(x^*(\Lambda); t(\Lambda)\big), \tag{123}$$

and Foreign's gain from entering the noncooperative regime is

$$\Delta(\Lambda) \equiv v_F\big(x^*(\Lambda); t(\Lambda)\big) - v_F\big(0; t(\Lambda)\big). \tag{124}$$

Our goal is yo show that $\Delta'(\Lambda) > 0$. Consequently, if $\pi(\Lambda) = G(\Delta(\Lambda))$, then, wherever G has density f,

$$\pi'(\Lambda) = f(\Delta(\Lambda))\Delta'(\Lambda) \geq 0, \tag{125}$$

with strict inequality whenever $f(\Delta(\Lambda)) > 0$. To clarify, the equation defining the non-cooperative branch, (123), is not an additional equilibrium condition. It is the non-cooperative equilibrium written in Λ coordinates. The cooperative-regime counterfactual in (124) is evaluated at the same Home tariff $t(\Lambda)$; its realized import share need not

equal Λ.

First we show that higher noncooperative Λ is supported by a lower Home tariff. Following Lemma 2, holding the other policy wedge fixed, both Home's import tariff and Foreign's export tax raise the buyer price $\tilde{p}_{FH}$:

$$\frac{\partial \ln \tilde{p}_{FH}}{\partial \ln(1+t)} > 0, \qquad \frac{\partial \ln \tilde{p}_{FH}}{\partial \ln(1+x)} > 0. \tag{126}$$

The CES demand specification gives $\frac{\partial \ln \lambda_{FH}}{\partial \ln \tilde{p}_{FH}} = -\epsilon \lambda_{HH} < 0$. Combining this with (126), yields

$$\frac{\partial \lambda_{FH}(x;t)}{\partial t} < 0, \qquad \frac{\partial \lambda_{FH}(x;t)}{\partial x} < 0. \tag{127}$$

Next differentiate (123). With all partial derivatives evaluated at $(x^*(\Lambda), t(\Lambda))$, we obtain

$$1 = \lambda_x\big(x^*(\Lambda); t(\Lambda)\big) x^{*\prime}(\Lambda) + \lambda_t\big(x^*(\Lambda); t(\Lambda)\big) t'(\Lambda). \tag{128}$$

Solving for $t'(\Lambda)$ gives

$$t'(\Lambda) = \frac{1 - \lambda_x\big(x^*(\Lambda); t(\Lambda)\big)\, x^{*\prime}(\Lambda)}{\lambda_t\big(x^*(\Lambda); t(\Lambda)\big)} < 0. \tag{129}$$

The denominator is negative by (127). The numerator is positive because $\lambda_x < 0$ and $x^{*\prime}(.) > 0$ as shown earlier. Thus higher noncooperative-regime aggregate import dependence corresponds to a lower supporting Home tariff.

Now differentiate (124):

$$\begin{aligned} \Delta'(\Lambda) = {} & v_{F,x}\big(x^*(\Lambda); t(\Lambda)\big)\, x^{*\prime}(\Lambda) \\ & + \big[v_{F,t}\big(x^*(\Lambda); t(\Lambda)\big) - v_{F,t}\big(0; t(\Lambda)\big)\big]\, t'(\Lambda). \end{aligned} \tag{130}$$

Because Foreign's non-cooperative choice is interior, the envelope result, $v_{F,x}\big(x^*(\Lambda); t(\Lambda)\big) = 0$, could be invoked to eliminate the first term in (130). The bracketed term is negative by Lemma 3. Formally, applying the sub-modularity ($v_{F,xt}(x;t) < 0$) established by Lemma 3 at the tariff $t(\Lambda)$,

$$\begin{aligned} & v_{F,t}\big(x^*(\Lambda); t(\Lambda)\big) - v_{F,t}\big(0; t(\Lambda)\big) \\ & \qquad = \int_0^{x^*(\Lambda)} v_{F,xt}\big(x; t(\Lambda)\big)\, dx < 0. \end{aligned} \tag{131}$$

Combining (129), (130), and (131) yields $\Delta'(\Lambda) > 0$. Finally, $\pi(\Lambda) = G(\Delta(\Lambda))$. Since G

is weakly increasing and Δ is strictly increasing, π is weakly increasing in Λ. Wherever G has density g, the chain rule gives

$$\pi'(\Lambda) = f(\Delta(\Lambda))\Delta'(\Lambda) \geq 0,$$

with strict inequality whenever $f(\Delta(\Lambda)) > 0$.

D Deterministic Optimal Tariff Problem

This appendix establishes the certainty-equivalent representation used in the main text. The argument separates three results. First, a unique certainty-equivalent export tax exists for every positive value of the risk elasticity. Second, when the risk elasticity is below one-half, the import share induced by this tax falls with Home's tariff. Third, for all sufficiently small values of the risk elasticity, the certainty-equivalent export tax also falls with Home's tariff. The last two signs imply that the export tax can be written as an increasing function of the deterministic import share. We maintain Assumptions 2, whereby $\epsilon > (1-\beta)/\beta$. We also maintain the existence and uniqueness conditions used throughout the paper: for every (t, x), the equilibrium exists, is unique and interior; and Foreign's noncooperative best response is unique, interior, and continuously differentiable; and the laissez-faire share restriction holds. Home's tariff belongs to $\mathbb{T} = [0, \bar{t}]$, where

$$(1+\bar{t})\frac{1-\beta}{\beta} \leq \epsilon.$$

For $\beta = 1$, this restriction is automatic. We also maintain the assumption that $0 < \zeta < 1/2$. Existence holds for every $\zeta > 0$. The one-half bound is used only to ensure that tariffs order the certainty-equivalent import share. The stronger sufficiently small ζ condition then makes the export-tax schedule increasing and delivers a unique deterministic equilibrium.

D.1. Setup and existence of the certainty equivalent

As in the draft, $t \equiv t_{FH}$ denotes Home's import tariff, $x \equiv x_{FH}$ denotes Foreign's export tax, and $v_H(x;t) \equiv \ln W_H(x;t)$ denotes Home's log welfare. We retain the draft's notation

$$\lambda(x;t) \equiv \lambda_{FH}(x;t), \qquad \mathscr{T} \equiv \frac{t}{1+t}.$$

Foreign's noncooperative export tax is $x^*(t)$, and the corresponding Home import share is

$$\Lambda(t) \equiv \lambda(x^*(t); t).$$

The other trade objects retain their draft meanings: λ_{HF} and λ_{FF} are Foreign's import and domestic expenditure shares, and $\widetilde{p}_{FH}$ is Home's tariff-inclusive price of Foreign goods. The preceding lemmas establish

$$\lambda_x(x;t) < 0, \qquad \lambda_t(x;t) < 0, \qquad \frac{d\Lambda(t)}{dt} < 0, \qquad \frac{dx^*(t)}{dt} < 0,$$

as well as the global share ordering $\lambda_{HF}(x;t) \leq \lambda_{FF}(x;t)$. They also give

$$x^*(t) = \left[\frac{\epsilon[1-\Lambda(t)]}{1-\mathscr{T}\Lambda(t)} + \frac{1-\beta}{\beta}\Lambda(t)\right]^{-1}. \tag{132}$$

Unless a variable is explicitly held fixed, derivatives with respect to a policy rate are equilibrium derivatives. Derivatives with respect to $\ln(1+t)$ have the same sign as derivatives with respect to t.

Home welfare admits the reduced form already derived in the draft:

$$W_H(x;t) = \bar{W}\,\mathcal{W}\big(\lambda(x;t), \mathscr{T}\big), \qquad \mathcal{W}(\lambda, \mathscr{T}) = (1-\lambda)^{-1/(\beta\epsilon)}\frac{\beta + (1-\beta)\mathscr{T}\lambda}{1-\mathscr{T}\lambda}, \tag{133}$$

where $\bar{W} > 0$ is independent of policy. Holding the tariff-revenue wedge fixed, differentiation with respect to the import share gives

$$\frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda} = \frac{1}{\beta\epsilon(1-\lambda)} + \frac{(1-\beta)\mathscr{T}}{\beta + (1-\beta)\mathscr{T}\lambda} + \frac{\mathscr{T}}{1-\mathscr{T}\lambda} > 0. \tag{134}$$

Every denominator is positive because $\lambda \in (0,1)$, $\mathscr{T} \in [0,1)$, and $\beta, \epsilon > 0$. Combining (134) with $\lambda_x < 0$ yields

$$\partial_x v_H(x;t) < 0.$$

Assumption 3 implies that the probability of the noncooperative regime is $\pi(t;\zeta) = \Lambda(t)^\zeta$. Home's ex ante log welfare is therefore

$$v_H^E(t;\zeta) = \big[1-\Lambda(t)^\zeta\big]v_H(0;t) + \Lambda(t)^\zeta v_H(x^*(t);t). \tag{135}$$

For every $\zeta > 0$, interiority of $\Lambda(t)$ implies $0 < \Lambda(t)^\zeta < 1$. Equation (135) is thus a strict

convex combination, and

$$v_H(x^*(t);t) < v_H^E(t;\zeta) < v_H(0;t).$$

Continuity and strict monotonicity of $v_H(\cdot;t)$ now imply that there is exactly one $\bar{x}(t;\zeta) \in (0, x^*(t))$ satisfying

$$v_H(\bar{x}(t;\zeta);t) = v_H^E(t;\zeta). \tag{136}$$

The resulting deterministic import share is

$$\bar{\Lambda}(t;\zeta) \equiv \lambda(\bar{x}(t;\zeta);t).$$

Since $0 < \bar{x}(t;\zeta) < x^*(t)$ and $\lambda_x < 0$, $\Lambda(t) < \bar{\Lambda}(t;\zeta) < \lambda(0;t)$. The derivative of the left-hand side of (136) with respect to $\bar{x}$ is $\partial_x v_H < 0$. The implicit function theorem and the maintained unique, continuously differentiable equilibrium therefore imply that $\bar{x}$ and $\bar{\Lambda}$ are continuously differentiable on $\mathbb{T} \times [0, 1/2]$, with one-sided derivatives at the boundary. At $\zeta = 0$, the continuous extension of (135) places probability one on the noncooperative regime. Hence

$$\bar{x}(t;0) = x^*(t), \qquad \bar{\Lambda}(t;0) = \Lambda(t). \tag{137}$$

D.2. Comparing tariff derivatives

The proof of share monotonicity requires a comparison between two tariff derivatives of Home welfare. We first compute the direct effect of a change in $\ln(1+t)$ while holding the import share fixed. From (133),

$$\left.\frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \ln(1+t)}\right|_{\lambda} = \frac{\lambda(1-\mathscr{T})}{(1-\mathscr{T}\lambda)[\beta + (1-\beta)\mathscr{T}\lambda]}. \tag{138}$$

This derivative is positive. It is also strictly increasing in the share, since

$$\frac{\partial}{\partial \lambda}\left\{\left.\frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \ln(1+t)}\right|_{\lambda}\right\} = \frac{(1-\mathscr{T})[\beta + (1-\beta)\mathscr{T}^2\lambda^2]}{(1-\mathscr{T}\lambda)^2[\beta + (1-\beta)\mathscr{T}\lambda]^2} > 0.$$

We next compute the total tariff derivative when Foreign's export tax is held fixed. The import-demand and export-supply elasticities established in the preceding appendix

are

$$\varepsilon_{FH}^{D} = 1 + \frac{\epsilon(1-\lambda)}{1-\mathscr{T}\lambda} + \frac{1-\beta}{\beta}\lambda,$$
$$\varepsilon_{FH}^{S} = \epsilon\frac{(1+x)\lambda_{FF}}{1+x\lambda_{FF}} + \frac{1-\beta}{\beta}\lambda_{HF}. \tag{139}$$

At a fixed x, tariff incidence satisfies

$$\frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+t)} = \frac{\varepsilon_{FH}^{S} + \dfrac{\lambda(1-\mathscr{T})}{1-\mathscr{T}\lambda}}{\varepsilon_{FH}^{S} + \varepsilon_{FH}^{D}}. \tag{140}$$

Combining (140) with the CES share equation and (134) gives

$$\left.\frac{\partial v_H(x;t)}{\partial \ln(1+t)}\right|_x = \frac{\lambda\left[(\epsilon+1)\frac{1-\lambda}{1-\mathscr{T}\lambda} + \frac{1-\beta}{\beta}\lambda\right]\left\{1-\mathscr{T}[1+\varepsilon_{FH}^{S}]\right\}}{[\beta+(1-\beta)\mathscr{T}\lambda][\varepsilon_{FH}^{D}+\varepsilon_{FH}^{S}]}. \tag{141}$$

All factors other than the last term in braces are positive. Thus, the sign of the fixed-export-tax derivative depends on whether the direct fiscal gain from a higher tariff exceeds the induced deterioration in Foreign's producer price.

The next auxiliary result compares the tariff derivative in the aligned regime with the direct derivative evaluated at the certainty-equivalent share. This is the main technical step in the proof.

Lemma 1. *For every $t \in \mathbb{T}$ and every $\zeta > 0$,*

$$\left[1-\Lambda(t)^{\zeta}\right]\frac{dv_H(0;t)}{d\ln(1+t)} + \Lambda(t)^{\zeta}\left.\frac{\partial \ln \mathcal{W}(\lambda,\mathscr{T})}{\partial \ln(1+t)}\right|_{\lambda=\Lambda(t)} < \left.\frac{\partial \ln \mathcal{W}(\lambda,\mathscr{T})}{\partial \ln(1+t)}\right|_{\lambda=\bar{\Lambda}(t;\zeta)}. \tag{142}$$

Proof. The argument proceeds in three steps. We first bound the aligned-regime tariff derivative. We then establish the curvature of this bound. Finally, we combine these properties with the certainty-equivalence condition.

Step 1: bound the aligned-regime tariff derivative. At $x = 0$, equation (139) becomes

$$\varepsilon_{FH}^{S}(0;t) = \epsilon\lambda_{FF}(0;t) + \frac{1-\beta}{\beta}\lambda_{HF}(0;t).$$

The ordering $\lambda_{HF}(0;0) \leq \lambda_{FF}(0;0)$ and adding up constraint imply $\lambda_{HF}(0;t) \leq 1/2$.

Therefore,

$$1+\varepsilon_{FH}^{S}(0;t)\geq\frac{1}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)>1. \tag{143}$$

Indeed, the difference between the two sides of the weak inequality is

$$\left(\epsilon-\frac{1-\beta}{\beta}\right)\left[\frac{1}{2}-\lambda_{HF}(0;t)\right]\geq 0.$$

For fixed (λ,t), the right-hand side of (141) is strictly decreasing in $1+\varepsilon_{FH}^{S}$. Differentiating with respect to this argument gives

$$-\frac{\lambda\left[(\epsilon+1)\frac{1-\lambda}{1-\mathscr{T}\lambda}+\frac{1-\beta}{\beta}\lambda\right]\left[1+\mathscr{T}\left\{\epsilon\frac{1-\lambda}{1-\mathscr{T}\lambda}+\frac{1-\beta}{\beta}\lambda\right\}\right]}{[\beta+(1-\beta)\mathscr{T}\lambda]\left[\epsilon\frac{1-\lambda}{1-\mathscr{T}\lambda}+\frac{1-\beta}{\beta}\lambda+1+\varepsilon_{FH}^{S}\right]^{2}}<0. \tag{144}$$

Equations (143) and (144) imply the upper bound

$$\frac{dv_H(0;t)}{d\ln(1+t)}\leq\left.\frac{\lambda\left[(\epsilon+1)\frac{1-\lambda}{1-\mathscr{T}\lambda}+\frac{1-\beta}{\beta}\lambda\right]\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]}{[\beta+(1-\beta)\mathscr{T}\lambda]\left[\epsilon\frac{1-\lambda}{1-\mathscr{T}\lambda}+\frac{1-\beta}{\beta}\lambda+\frac{1}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]}\right|_{\lambda=\lambda(0;t)}. \tag{145}$$

We distinguish two cases. Suppose first that

$$1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\leq 0. \tag{146}$$

The upper bound in (145) is then nonpositive. Since the direct derivative in (138) is positive and strictly increasing in λ, while $\Lambda(t)<\bar{\Lambda}(t;\zeta)$, we obtain

$$\begin{aligned}&[1-\Lambda(t)^{\zeta}]\frac{dv_H(0;t)}{d\ln(1+t)}+\Lambda(t)^{\zeta}\left.\frac{\partial\ln\mathcal{W}}{\partial\ln(1+t)}\right|_{\lambda=\Lambda(t)}\\&\quad\leq\Lambda(t)^{\zeta}\left.\frac{\partial\ln\mathcal{W}}{\partial\ln(1+t)}\right|_{\lambda=\Lambda(t)}<\left.\frac{\partial\ln\mathcal{W}}{\partial\ln(1+t)}\right|_{\lambda=\bar{\Lambda}(t;\zeta)}.\end{aligned}$$

This proves the result in the first case.

Step 2: establish the curvature of the upper bound. Suppose next that the inequality in (146) is reversed. We record the curvature properties of the upper bound in (145). Before

evaluation at $\lambda(0;t)$, its right-hand side can be factored as

$$\frac{\lambda}{\beta+(1-\beta)\mathscr{T}\lambda} \times \frac{2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]\left[(\epsilon+1)(1-\lambda)+\frac{1-\beta}{\beta}\lambda(1-\mathscr{T}\lambda)\right]}{2\left[(\epsilon+1)(1-\lambda)+\frac{1-\beta}{\beta}\lambda(1-\mathscr{T}\lambda)\right]+\epsilon+\frac{1-\beta}{\beta}+2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]\lambda}. \quad (147)$$

The first factor is positive and increasing, with derivatives

$$\begin{aligned}\frac{d}{d\lambda}\frac{\lambda}{\beta+(1-\beta)\mathscr{T}\lambda} &= \frac{\beta}{[\beta+(1-\beta)\mathscr{T}\lambda]^2} > 0,\\ \frac{d^2}{d\lambda^2}\frac{\lambda}{\beta+(1-\beta)\mathscr{T}\lambda} &= -\frac{2\beta(1-\beta)\mathscr{T}}{[\beta+(1-\beta)\mathscr{T}\lambda]^3} \leq 0.\end{aligned} \quad (148)$$

The second factor in (147) is positive, decreasing, and strictly concave. To verify these signs, first observe that the bracket

$$(\epsilon+1)(1-\lambda)+\frac{1-\beta}{\beta}\lambda(1-\mathscr{T}\lambda) \quad (149)$$

is positive. Its first derivative is $-[\epsilon-(1-\beta)/\beta+1]-2[(1-\beta)/\beta]\mathscr{T}\lambda < 0$, and its second derivative is $-2[(1-\beta)/\beta]\mathscr{T} \leq 0$. Direct differentiation of the second factor gives

$$\begin{aligned}&\frac{2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]}{\left\{2\left[(\epsilon+1)(1-\lambda)+\frac{1-\beta}{\beta}\lambda(1-\mathscr{T}\lambda)\right]+\epsilon+\frac{1-\beta}{\beta}+2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]\lambda\right\}^2}\\ &\times\left\{\left[\epsilon+\frac{1-\beta}{\beta}+2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]\lambda\right]\left[-\epsilon+\frac{1-\beta}{\beta}-1-2\frac{1-\beta}{\beta}\mathscr{T}\lambda\right]\right.\\ &\left.-2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]\left[(\epsilon+1)(1-\lambda)+\frac{1-\beta}{\beta}\lambda(1-\mathscr{T}\lambda)\right]\right\} < 0.\end{aligned} \quad (150)$$

The expression in braces is strictly negative. Its derivative is

$$-2\frac{1-\beta}{\beta}\mathscr{T}\left[\epsilon+\frac{1-\beta}{\beta}+2\left[1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]\lambda\right] \leq 0. \quad (151)$$

The derivative of the denominator of the second factor in (147)—the quantity squared in

the first line of (150)—is

$$-2\left[\epsilon-\frac{1-\beta}{\beta}+2\frac{1-\beta}{\beta}\mathscr{T}\lambda+\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\right]<0. \tag{152}$$

One more application of the quotient rule therefore gives a strictly negative second derivative: its numerator is the derivative in (151) times a positive denominator, minus twice the negative expression in braces times the negative derivative in (152).

The denominator of the second factor in (147) exceeds twice the positive bracket in (149). Hence this factor is strictly smaller than

$$1-\frac{\mathscr{T}}{2}\left(\epsilon+\frac{1-\beta}{\beta}+2\right)\leq 1-\mathscr{T}\leq\frac{1-\mathscr{T}}{1-\mathscr{T}\lambda}, \tag{153}$$

where the first weak inequality uses $[\epsilon+(1-\beta)/\beta+2]/2>1$. The final term in (153) is the second factor in the direct derivative (138); it is increasing in λ. Equations (148)–(153) imply three properties. The upper bound in (145), viewed as a function of λ, is strictly concave; it lies below the direct derivative in (138); and its first derivative is strictly smaller than the derivative in (138). These conclusions follow from the product rule: the first factor is positive, increasing, and weakly concave, while the second factor of the upper bound is positive, decreasing, and strictly concave.

We also need convexity of Home log welfare as a function of the import share in this case. Twice differentiating (133) gives

$$\frac{\partial^2\ln\mathcal{W}(\lambda,\mathscr{T})}{\partial\lambda^2}=-\frac{(1-\beta)^2\mathscr{T}^2}{[\beta+(1-\beta)\mathscr{T}\lambda]^2}+\frac{\mathscr{T}^2}{(1-\mathscr{T}\lambda)^2}+\frac{1}{\beta\epsilon(1-\lambda)^2}. \tag{154}$$

The maintained case implies

$$\mathscr{T}<\frac{2}{\epsilon+(1-\beta)/\beta+2}. \tag{155}$$

Moreover,

$$\left(\epsilon+\frac{1-\beta}{\beta}+2\right)^2\geq 4\epsilon\left(\frac{1-\beta}{\beta}+2\right)>\frac{4\epsilon[(1-\beta)/\beta]^2}{1+(1-\beta)/\beta}. \tag{156}$$

The first inequality is equivalent to $\{\epsilon-[(1-\beta)/\beta+2]\}^2\geq 0$; the second follows because $[(1-\beta)/\beta+2]>[(1-\beta)/\beta]^2/[1+(1-\beta)/\beta]$. Equations (155) and (156) imply $\left(\frac{1-\beta}{\beta}\right)^2\mathscr{T}^2<\frac{1}{\beta\epsilon}$. The negative term in (154) is therefore smaller in absolute value than

$1/(\beta\epsilon)$, whereas the final positive term is strictly larger than this number. It follows that

$$\frac{\partial^2 \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda^2} > 0. \tag{157}$$

Fix any $\lambda \in (\Lambda(t), \lambda(0;t))$. Apply the tangent inequality to the concave upper bound in (145), and then use the level and slope comparisons established above. Strict convexity in (157) gives the final inequality below:

$$\begin{aligned}\frac{dv_H(0;t)}{d\ln(1+t)} < & \left.\frac{\partial \ln \mathcal{W}}{\partial \ln(1+t)}\right|_{\lambda} \\ & + \frac{\dfrac{\partial}{\partial \lambda}\left\{\left.\dfrac{\partial \ln \mathcal{W}}{\partial \ln(1+t)}\right|_{\lambda}\right\}}{\dfrac{\partial \ln \mathcal{W}}{\partial \lambda}} \left[\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\lambda, \mathscr{T})\right].\end{aligned} \tag{158}$$

Strict convexity enters through the tangent inequality

$$\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\lambda, \mathscr{T}) > \frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda}[\lambda(0;t) - \lambda].$$

To complete the comparison, divide the gap in tariff slopes by the associated welfare gap. Direct differentiation gives

$$\begin{aligned}&\frac{d}{d\lambda}\left\{\frac{\dfrac{dv_H(0;t)}{d\ln(1+t)} - \left.\dfrac{\partial \ln \mathcal{W}}{\partial \ln(1+t)}\right|_{\lambda}}{\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\lambda, \mathscr{T})}\right\} \\ &= \frac{\begin{array}{l}\dfrac{\partial \ln \mathcal{W}}{\partial \lambda}\left[\dfrac{dv_H(0;t)}{d\ln(1+t)} - \left.\dfrac{\partial \ln \mathcal{W}}{\partial \ln(1+t)}\right|_{\lambda}\right] \\ \quad - \dfrac{\partial}{\partial \lambda}\left\{\left.\dfrac{\partial \ln \mathcal{W}}{\partial \ln(1+t)}\right|_{\lambda}\right\}\left[\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\lambda, \mathscr{T})\right]\end{array}}{\left[\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\lambda, \mathscr{T})\right]^2} < 0.\end{aligned} \tag{159}$$

The denominator is positive, and the numerator is negative by (158).

Step 3: impose certainty equivalence. The certainty-equivalence condition implies

$$\begin{aligned}&\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\bar{\Lambda}(t;\zeta), \mathscr{T}) \\ &\quad = \Lambda(t)^{\zeta}\left[\ln \mathcal{W}(\lambda(0;t), \mathscr{T}) - \ln \mathcal{W}(\Lambda(t), \mathscr{T})\right].\end{aligned} \tag{160}$$

Because $\bar{\Lambda}(t;\zeta) > \Lambda(t)$, the expression inside the derivative in (159) is smaller at $\bar{\Lambda}(t;\zeta)$

than at $\Lambda(t)$. Substituting (160) and rearranging yields (142). □

D.3. Bounding the welfare loss between regimes

The remaining term in the derivative of expected welfare comes from the response of the probability $\Lambda(t)^{\zeta}$. The next auxiliary result bounds the welfare loss from moving from the aligned regime to the noncooperative regime.

Lemma 2. *For every $t \in \mathbb{T}$,*

$$v_H(0;t) - v_H(x^*(t);t) < 2\Lambda(t) \left. \frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda} \right|_{\lambda=\Lambda(t)}. \tag{161}$$

Proof. The proof proceeds in three steps. We first establish the required concavity of Home welfare. We then bound the distance between the aligned and noncooperative shares. We finally combine these two results.

Step 1: establish concavity in the domestic-share coordinate. At a fixed tariff, regard $\ln \mathcal{W}$ as a function of $-\ln(1-\lambda)$. Applying the chain rule twice gives

$$\begin{aligned} &\frac{\partial^2 \ln \mathcal{W}}{\partial [-\ln(1-\lambda)]^2} \\ &= -(1-\lambda)\left[\frac{(1-\beta)\mathscr{T}[\beta + (1-\beta)\mathscr{T}]}{[\beta + (1-\beta)\mathscr{T}\lambda]^2} + \frac{\mathscr{T}(1-\mathscr{T})}{(1-\mathscr{T}\lambda)^2}\right] \leq 0. \end{aligned} \tag{162}$$

Thus, Home log welfare is concave in $-\ln(1-\lambda)$.

Step 2: bound the distance between the endpoint shares. We next bound the distance between the aligned and noncooperative import shares. The CES share equation and the export-tax pass-through result from the preceding appendix imply

$$-\frac{d}{d\ln(1+x)} \ln \frac{\lambda(x;t)}{1-\lambda(x;t)} = \epsilon \frac{\partial \ln \widetilde{p}_{FH}}{\partial \ln(1+x)} < \epsilon. \tag{163}$$

The last inequality uses export-tax pass-through strictly below one. Integrating (163) from $x = 0$ to $x = x^*(t)$ and then using (132) gives

$$\begin{aligned} \ln\left[\frac{\lambda(0;t)}{1-\lambda(0;t)} \frac{1-\Lambda(t)}{\Lambda(t)}\right] &< \epsilon \ln[1 + x^*(t)] < \epsilon x^*(t) \\ &= \left[\frac{1-\Lambda(t)}{1-\mathscr{T}\Lambda(t)} + \frac{1-\beta}{\beta\epsilon}\Lambda(t)\right]^{-1} \leq \frac{1}{1-\Lambda(t)}. \end{aligned} \tag{164}$$

The second inequality uses $\ln(1+x^*) < x^*$ for $x^* > 0$. The last inequality follows because the expression inside square brackets is at least $1 - \Lambda(t)$.

We use the scalar inequality

$$\ln \frac{1 + z\exp(1+z)}{1+z} < 2z \qquad \text{for every } z > 0. \tag{165}$$

To verify it, exponentiate and rearrange. Equation (165) is equivalent to

$$(1+z)e^{z-1} - z - e^{-1-z} > 0.$$

The left-hand side equals zero at $z = 0$. Its derivative at zero is $3/e - 1 > 0$, and its second derivative is

$$e^{z-1}(z+3) - e^{-1-z} > 0 \qquad \text{for } z > 0.$$

The first derivative is therefore positive for every $z \geq 0$, which proves (165).

Equation (164) implies that the aligned-regime import odds are strictly smaller than

$$\frac{\Lambda(t)}{1-\Lambda(t)} \exp\left(1 + \frac{\Lambda(t)}{1-\Lambda(t)}\right).$$

Evaluating (165) at $z = \Lambda(t)/[1-\Lambda(t)]$ therefore yields

$$\ln \frac{1-\Lambda(t)}{1-\lambda(0;t)} < 2\frac{\Lambda(t)}{1-\Lambda(t)}.$$

Step 3: translate the share bound into a welfare bound. Finally, concavity in (162) gives

$$\begin{aligned} v_H(0;t) - v_H(x^*(t);t) &\leq [1-\Lambda(t)] \left.\frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda}\right|_{\lambda=\Lambda(t)} \ln \frac{1-\Lambda(t)}{1-\lambda(0;t)} \\ &< 2\Lambda(t) \left.\frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda}\right|_{\lambda=\Lambda(t)}. \end{aligned}$$

This is (161). □

D.4. Monotonicity of the certainty-equivalent import share

We now combine the two auxiliary results. Differentiate (135) with respect to $\ln(1+t)$. The product and chain rules give

$$\begin{aligned}\frac{dv_H^E(t;\zeta)}{d\ln(1+t)} &= [1-\Lambda(t)^\zeta]\frac{dv_H(0;t)}{d\ln(1+t)} \\ &\quad + \Lambda(t)^\zeta \left.\frac{\partial \ln \mathcal{W}(\lambda,\mathscr{T})}{\partial \ln(1+t)}\right|_{\lambda=\Lambda(t)} \\ &\quad + \frac{\Lambda(t)^\zeta}{\Lambda(t)[1-\Lambda(t)]}\frac{d\Lambda(t)}{d\ln(1+t)}\Bigg\{\Lambda(t)[1-\Lambda(t)]\left.\frac{\partial \ln \mathcal{W}}{\partial \lambda}\right|_{\lambda=\Lambda(t)} \\ &\qquad\qquad - \zeta[1-\Lambda(t)][v_H(0;t)-v_H(x^*(t);t)]\Bigg\}. \qquad (166)\end{aligned}$$

The last line combines the change in the regime probability,

$$\frac{d\Lambda(t)^\zeta}{d\ln(1+t)} = \zeta\Lambda(t)^{\zeta-1}\frac{d\Lambda(t)}{d\ln(1+t)},$$

and movement along the noncooperative branch,

$$\left.\frac{\partial \ln \mathcal{W}}{\partial \lambda}\right|_{\lambda=\Lambda(t)}\frac{d\Lambda(t)}{d\ln(1+t)}.$$

Suppose $0<\zeta<1/2$. The endpoint welfare-loss bound implies that the expression in braces in (166) is strictly positive:

$$\begin{aligned}&\Lambda(1-\Lambda)\left.\frac{\partial \ln \mathcal{W}}{\partial \lambda}\right|_{\lambda=\Lambda} - \zeta(1-\Lambda)[v_H(0;t)-v_H(x^*;t)] \\ &\quad > (1-\Lambda)(1-2\zeta)\Lambda\left.\frac{\partial \ln \mathcal{W}}{\partial \lambda}\right|_{\lambda=\Lambda} > 0.\end{aligned}$$

Here and only here, Λ and x^* are written without their argument t to keep the display readable. Since $d\Lambda(t)/d\ln(1+t)<0$, the final term in (166) is strictly negative. The tariff-slope comparison therefore gives

$$\frac{dv_H^E(t;\zeta)}{d\ln(1+t)} < \left.\frac{\partial \ln \mathcal{W}(\lambda,\mathscr{T})}{\partial \ln(1+t)}\right|_{\lambda=\bar{\Lambda}(t;\zeta)}. \qquad (167)$$

The certainty-equivalence condition can also be written as

$$\ln \bar{W} + \ln \mathcal{W}(\bar{\Lambda}(t;\zeta), \mathscr{T}) = v_H^E(t;\zeta). \tag{168}$$

Differentiating (168) with respect to $\ln(1+t)$ gives

$$\frac{dv_H^E(t;\zeta)}{d\ln(1+t)} = \frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \ln(1+t)}\bigg| + \frac{\partial \ln \mathcal{W}(\lambda, \mathscr{T})}{\partial \lambda}\bigg|_{\lambda=\bar{\Lambda}(t;\zeta)} \frac{d\bar{\Lambda}(t;\zeta)}{d\ln(1+t)}. \tag{169}$$

The share derivative in the second line is strictly positive by (134). Comparing (167) and (169) proves

$$\frac{d\bar{\Lambda}(t;\zeta)}{dt} < 0 \qquad \text{for every } t \in \mathbb{T} \text{ and every } 0 < \zeta < \frac{1}{2}. \tag{170}$$

At $t = 0$ and $t = \bar{t}$, the derivatives are one-sided.

D.5. The certainty-equivalent schedule and uniqueness

Fix $\zeta \in (0, 1/2)$ and suppress its effect on the schedule's notation. Equation (170) implies that $t \mapsto \bar{\Lambda}(t;\zeta)$ is a continuously differentiable, strictly decreasing bijection from $\mathbb{T}$ onto

$$\bar{\Lambda}(\mathbb{T};\zeta) = [\bar{\Lambda}(\bar{t};\zeta), \bar{\Lambda}(0;\zeta)]. \tag{171}$$

It therefore defines a unique schedule $\bar{\tau}$ on this range through

$$\bar{\tau}(\bar{\Lambda}(t;\zeta)) = 1 + \bar{x}(t;\zeta). \tag{172}$$

The inverse function theorem implies that this schedule is continuously differentiable on the interior of the range, with one-sided derivatives at its endpoints. Equation (172) also gives

$$\bar{\Lambda}(t;\zeta) = \lambda(\bar{\tau}(\bar{\Lambda}(t;\zeta)) - 1; t). \tag{173}$$

Thus, for every tariff, the deterministic economy admits a certainty-equivalent branch with export tax $\bar{x}(t;\zeta)$ and import share $\bar{\Lambda}(t;\zeta)$. On this branch, Home's log welfare is $v_H^E(t;\zeta)$ by (136). For a general $\zeta \in (0, 1/2)$, however, neither monotonicity of $\bar{\tau}$ nor uniqueness of the deterministic fixed point has yet been established.

It remains to prove the stronger conclusions for sufficiently small ζ. Differentiating

(136) with respect to t and solving for the total derivative of $\bar{x}$ gives

$$\frac{d\bar{x}(t;\zeta)}{dt} = \frac{\dfrac{dv_H^E(t;\zeta)}{dt} - \dfrac{\partial v_H(x;t)}{\partial t}\Big|_{x=\bar{x}(t;\zeta)}}{\dfrac{\partial v_H(\bar{x}(t;\zeta);t)}{\partial x}}. \tag{174}$$

The partial derivative in the numerator holds the export tax fixed. The right-hand side is continuous on the compact rectangle $\mathbb{T} \times [0, 1/2]$. Indeed, the noncooperative share remains in a compact subset of $(0, 1)$; the map $(t, \zeta) \mapsto \Lambda(t)^\zeta$ is continuously differentiable; and the denominator is continuous, strictly negative, and therefore bounded away from zero on the rectangle.

At $\zeta = 0$, equation (137) implies

$$\frac{d\bar{x}(t;0)}{dt} = \frac{dx^*(t)}{dt} < 0.$$

The derivative dx^*/dt is continuous and strictly negative on the compact tariff domain, so its maximum is strictly negative. Uniform continuity of (174) therefore implies that there is a $\bar{\zeta} \in (0, 1/2)$ such that

$$\frac{d\bar{x}(t;\zeta)}{dt} < 0 \qquad \text{for every } (t, \zeta) \in \mathbb{T} \times (0, \bar{\zeta}]. \tag{175}$$

Combining (175) with (170) and differentiating (172) yields

$$\bar{\tau}'(\bar{\Lambda}(t;\zeta)) = \frac{d\bar{x}(t;\zeta)/dt}{d\bar{\Lambda}(t;\zeta)/dt} > 0 \qquad \text{for } 0 < \zeta \leq \bar{\zeta}.$$

This derivative is taken on the range in (171); it is one-sided at the two endpoints.

For these sufficiently small values of ζ, extend $\bar{\tau}$ from its original range to $[0, 1]$. Below $\bar{\Lambda}(\bar{t};\zeta)$, set the schedule equal to $\bar{\tau}(\bar{\Lambda}(\bar{t};\zeta))$; above $\bar{\Lambda}(0;\zeta)$, set it equal to $\bar{\tau}(\bar{\Lambda}(0;\zeta))$; and leave it unchanged on its original range. The extension is continuous and weakly increasing on $[0, 1]$, and it is strictly increasing on the original range. In general it has a kink at each join, so no differentiability claim is made for the extended schedule there.

Given $t \in \mathbb{T}$, a deterministic equilibrium import share solves

$$\lambda = \lambda(\bar{\tau}(\lambda) - 1; t). \tag{176}$$

The right-hand side is continuous and non-increasing in λ because $\bar{\tau}$ is nondecreasing

and $\lambda_x < 0$. By interiority of CES shares,

$$0 - \lambda(\bar{\tau}(0) - 1; t) < 0, \qquad 1 - \lambda(\bar{\tau}(1) - 1; t) > 0.$$

The intermediate value theorem gives existence. For uniqueness, if $\lambda_2 > \lambda_1$, then

$$\lambda(\bar{\tau}(\lambda_2) - 1; t) \leq \lambda(\bar{\tau}(\lambda_1) - 1; t),$$

whereas the left-hand side of (176) is strictly larger at λ_2. Hence (176) has at most one solution. Equation (173) shows that this solution is $\bar{\Lambda}(t; \zeta)$. Together with the maintained uniqueness of the equilibrium under each exogenous policy pair, uniqueness of the share gives uniqueness of the deterministic general equilibrium.

At this equilibrium,

$$v_H(\bar{\tau}(\bar{\Lambda}(t; \zeta)) - 1; t) = v_H(\bar{x}(t; \zeta); t) = v_H^E(t; \zeta).$$

Thus, the deterministic and Bayesian policy problems have the same objective at every tariff in $\mathbb{T}$. Since both problems maximize over the same compact tariff domain, their sets of optimal tariffs coincide. This proves all parts of the revised lemma.

The economic logic is as follows. A higher Home tariff lowers the import share in the noncooperative regime and therefore lowers the probability of that regime. The decline in this probability reduces the certainty-equivalent export tax, which partially offsets the direct decline in import dependence. The endpoint welfare-loss bound shows that this offset cannot reverse the share response when $\zeta < 1/2$. For sufficiently small ζ, the certainty-equivalent locus also remains close enough to the noncooperative locus to inherit its negative export-tax slope. The export tax and the import share then move in the same direction along the locus, making $\bar{\tau}$ increasing in import dependence.

D.6. First-order formula used in the quantitative analysis

We retain the approximation used in the quantitative analysis. This approximation requires one additional regularity condition: Home welfare and the equilibrium share must be twice continuously differentiable on the relevant compact policy set. This extra smoothness is not needed for the revised lemma. Fix $\zeta \in (0, 1/2)$ and suppress only its argument. All $O(\cdot)$ terms below are uniform in $t \in \mathbb{T}$: their absolute values are bounded by a common finite constant times the displayed power of $x^*(t)$. The first-order interpretation is relevant when the noncooperative export tax is small.

Taylor's theorem applied to (136) around $x = 0$ gives

$$v_H(0;t) + \partial_x v_H(0;t)\bar{x}(t) + O(x^*(t)^2) = v_H(0;t) + \Lambda(t)^{\zeta}\left[\partial_x v_H(0;t)x^*(t) + O(x^*(t)^2)\right]. \quad (177)$$

The remainders are of order $x^*(t)^2$ because $0 < \bar{x}(t) < x^*(t)$. The derivative $\partial_x v_H(0;t)$ is continuous and strictly negative on the compact tariff domain, so it is bounded away from zero. Rearranging (177) yields

$$\bar{x}(t) = \Lambda(t)^{\zeta} x^*(t) + O(x^*(t)^2). \quad (178)$$

We next replace the noncooperative share $\Lambda(t)$ by the deterministic share $\bar{\Lambda}(t)$. The mean value theorem, boundedness of λ_x on the compact region between the two loci, and $0 < \bar{x}(t) < x^*(t)$ imply

$$|\bar{\Lambda}(t) - \Lambda(t)| = O(x^*(t)). \quad (179)$$

Both shares remain in a compact subset of $(0,1)$. Hence the derivative of $\lambda \mapsto \lambda^{\zeta}$ is bounded on the relevant range, and

$$\Lambda(t)^{\zeta} = \bar{\Lambda}(t)^{\zeta} + O(x^*(t)). \quad (180)$$

It remains to change the argument of the noncooperative export-tax schedule. We use the draft's existing schedule

$$x^*(\lambda) = \left[\epsilon(1-\lambda)\delta(\lambda) + \frac{1-\beta}{\beta}\lambda\right]^{-1}, \qquad \delta(\lambda) = \left[1 - \frac{t(\lambda)}{1+t(\lambda)}\lambda\right]^{-1}, \quad (181)$$

where $t(\lambda)$ is the inverse of the noncooperative locus. Since $\bar{\Lambda}(t)$ may exceed $\Lambda(0)$, extend the schedule above $\Lambda(0)$ by setting $t(\lambda) = 0$ and $\delta(\lambda) = 1$. This is a continuous extension because $t(\Lambda(0)) = 0$.

Where the extended schedule is differentiable, direct differentiation of (181) gives

$$x^{*\prime}(\lambda) = \left[\epsilon\delta(\lambda) - \epsilon(1-\lambda)\delta'(\lambda) - \frac{1-\beta}{\beta}\right] x^*(\lambda)^2.$$

The bracket is bounded on each side of $\Lambda(0)$. Equivalently,

$$\frac{d}{d\lambda}\frac{1}{x^*(\lambda)} = -\epsilon\delta(\lambda) + \epsilon(1-\lambda)\delta'(\lambda) + \frac{1-\beta}{\beta}.$$

The reciprocal schedule is continuous at the join and has bounded one-sided derivatives there. It is therefore Lipschitz on the relevant range. Applying the mean value theorem separately on the two sides of the join and using (179) gives

$$\frac{1}{x^*(\bar{\Lambda}(t))} = \frac{1}{x^*(\Lambda(t))} + O(x^*(t)).$$

Since $x^*(\Lambda(t)) = x^*(t)$, inverting the last display gives

$$x^*(\Lambda(t)) = x^*(\bar{\Lambda}(t)) + O(x^*(t)^3). \tag{182}$$

Substituting (180) and (182) into (178) yields

$$\bar{\tau}(\bar{\Lambda}(t)) = 1 + \bar{\Lambda}(t)^{\zeta} x^*(\bar{\Lambda}(t)) + O(x^*(t)^2).$$

This is the first-order formula used in the main text. It evaluates both the hazard and the noncooperative export tax at the deterministic equilibrium share; the coordinate changes affect only the remainder.

E Optimal Tariff with Strategic Foreign government

The only difference in this case with the baseline derivation is that the foreign price includes an additional wedge $p_{FH} \equiv \bar{\tau}(\lambda_{FH}) d_{FH} p_F$. Thus, the expression for net income at home must be updated as

$$Y_H = w_H L_H + (\tilde{p}_{FH} - \bar{\tau}(\lambda_{FH}) d_{FH} p_F) Q_{FH}$$

To better handle this additional policy wedge, we formulate the optimal tariff problem as one where the government directly selects price $\tilde{p}_{FH}$ rather that τ^M_{FH}. Extrapolating

from our baseline derivation:

$$\frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH}} = \frac{1}{Y_H}\left[w_H L_H \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH}} + \tilde{p}_{FH}Q_{FH} + \left(\tilde{p}_{FH} - \left(1 + \frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}}\right)p_{FH}\right)Q_{FH}\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln \tilde{p}_{FH}} - p_{FH}Q_{FH}\frac{\partial \ln \bar{\tau}(\lambda_{FH})}{\partial \ln \lambda_{FH}}\frac{\mathrm{d}\ln \lambda_{FH}}{\mathrm{d}\ln \tilde{p}_{FH}}\right].$$

Henceforth we apply our compact notation for the export supply and import demand elasticities:

$$\varepsilon^D_{FH} \equiv -\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln \tilde{p}_{FH}} \qquad \frac{1}{\varepsilon^S_{FH}} \equiv \frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}}$$

Following the same logic we can show that

$$\frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln \tilde{p}_{FH}} = \lambda_{FH}, \qquad \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH}} = -\frac{1-\beta}{\beta}\lambda_{FH}$$

Noting that $1 + t = \tilde{p}_{FH}/p_{FH}$ and $\lambda_{FH} = \tilde{p}_{FH}Q_{FH}/E_H$, we can write the first-order condition by extrapolating from the baseline derivation:

$$\begin{aligned}\frac{\mathrm{d}\ln W_H}{\mathrm{d}\ln \tilde{p}_{FH}} &= \frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH}} - \frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln \tilde{p}_{FH}} \\ &= -\left[\left(1 - \frac{1}{\tau^M_{FH}}\left(1 + \frac{1}{\varepsilon^S_{FH}} - \frac{\frac{\mathrm{d}\ln \lambda_{FH}}{\mathrm{d}\ln \tilde{p}_{FH}}}{\varepsilon^D_{FH}}\frac{\partial \ln \bar{\tau}}{\partial \ln \lambda_{FH}}\right)\right)\lambda_{FH}\varepsilon^D_{FH}\right]\frac{E_H}{Y_H} = 0,\end{aligned}$$

The trade share derivative can be obtained by noting that $\lambda_{FH} = \frac{\tilde{p}^{-\epsilon}_{FH}}{\tilde{p}^{-\epsilon}_{FH} + P^{-\epsilon}_H}$, where $p_H = 1$ by choice of numeriare. More specifically,

$$\frac{d\ln \lambda_{FH}}{d\ln \tilde{p}_{FH}} = -\epsilon(1 - \lambda_{FH}) = -\epsilon\lambda_{HH}$$

The import demand elasticity is

$$\varepsilon^D_{FH} \equiv -\frac{d\ln Q_{FH}}{d\ln \tilde{p}_{FH}} = -\frac{\partial \ln(\lambda_{FH}E_H/\tilde{p}_{FH})}{\partial \ln \tilde{p}_{FH}} = -\frac{d\ln \lambda_{FH}}{d\ln \tilde{p}_{FH}} - \frac{d\ln E_H}{d\ln \tilde{p}_{FH}} + 1 \approx 1 + \epsilon\lambda_{HH} + \frac{1-\beta}{\beta}\lambda_{FH}$$

where the last line follows from the simplifying assumption, $T^M_H/E_H \approx 0$, whereby

$$\frac{d\ln E_H}{d\ln \tilde{p}_{FH}} \approx \frac{d\ln w_H}{d\ln \tilde{p}_{FH}} = -\frac{1-\beta}{\beta}\lambda_{FH}.$$

Below, we relax this assumption and derive ε^D_{FH} without the noted restriction. Plugging the expression $\frac{d\ln\lambda_{FH}}{d\ln\tilde{p}_{FH}}$ into the first-order condition, the optimal tariff solves

$$1-\frac{1}{\tau^M_{FH}}\left(1+\frac{1}{\varepsilon^S_{FH}}+\frac{\epsilon\lambda_{HH}}{\varepsilon^D_{FH}}\frac{\partial\ln\bar{\tau}}{\partial\ln\lambda_{FH}}\right)=0$$

which delivers:

$$t^*=\frac{1}{\varepsilon^S_{FH}}+\frac{\epsilon\lambda_{HH}}{\varepsilon^D_{FH}}\frac{\partial\ln\bar{\tau}}{\partial\ln\lambda_{FH}}$$

Following the assumption that $T_i/E_i\approx 0$,the export supply and import demand elasticities are

$$\varepsilon^S_{FH}=\epsilon\lambda_{FF}+\frac{1-\beta}{\beta}\lambda_{HF},\qquad\qquad \varepsilon^D_{FH}=1+\epsilon\lambda_{HH}+\frac{1-\beta}{\beta}\lambda_{FH}$$

Import Demand Elasticity with Large Trade Tax Revenues. Previously, we assumed that $T_H/E_H\approx 0$, which implied that $d\ln E_H=d\ln w_H$. Now we relax this assumption, without which gross expenditure is precisely given by

$$E_H=\frac{w_HL_H}{\beta}+\frac{t}{1+t}\lambda_{FH}E_H=\frac{w_HL_H/\beta}{1-\frac{t}{1+t}\lambda_{FH}}$$

Taking log derivatives from the above equation, yields

$$\frac{d\ln E_H}{d\ln\tilde{p}_{FH}}=\frac{d\ln w_H}{d\ln\tilde{p}_{FH}}-\frac{d\ln(1-\frac{t}{1+t}\lambda_{FH})}{d\ln\tilde{p}_{FH}}=\frac{d\ln w_H}{d\ln\tilde{p}_{FH}}+\frac{\frac{t}{1+t}\lambda_{FH}}{1-\frac{t}{1+t}\lambda_{FH}}d\ln\lambda_{FH}$$

Plugging the last line back into the previously-derived expression

$$\varepsilon^D_{FH}=1-\frac{d\ln\lambda_{FH}}{d\ln\tilde{p}_{FH}}-\frac{d\ln E_H}{d\ln\tilde{p}_{FH}}$$

and noting that wage derivative is $\frac{d\ln w_H}{d\ln\tilde{p}_{FH}}=-\frac{1-\beta}{\beta}\lambda_{FH}$, delivers the precise expression for the import demand elasticity:

$$\varepsilon^D_{FH}=1+\frac{\epsilon\lambda_{HH}}{1-\frac{t}{1+t}\lambda_{FH}}+\frac{1-\beta}{\beta}\lambda_{FH}$$

The previous formulation results from the approximation, $\frac{t}{1+t}\lambda_{FH}\approx 0$.

F Optimal Tariff: Bayesian Representation

Maintain the paper's notation and assumptions. In particular, Home sets $\tau_{FH}^M = 1 + t$ before Foreign's type is realized, Foreign sets $x = 0$ with probability $1 - \pi(t)$ and $x = x^*(t)$ with probability $\pi(t)$. Under the non-cooperative regime the import share and export tax are denoted by

$$\Lambda \equiv \lambda_{FH}(x^*(t); t), \qquad x^*(t) = \mathrm{x}(\Lambda).$$

Home's ex ante problem is therefore

$$\max_t \; v_H^E(t) \equiv [1 - \pi(t)] v_H(0; t) + \pi(t) v_H(x^*(t); t).$$

The point of the derivation below is to solve the Bayesian optimal policy problem *directly*, rather than converting the problem into an equivalent deterministic problem. Specifically, we show that the optimal tariff, t_{FH}^*, satisfies

$$\begin{aligned} t^* = & \frac{(1-\pi)\Omega_0}{(1-\pi)\Omega_0 + \pi\Omega_1} \cdot \frac{1}{\varepsilon_{FH}^S(0;t)} \\ & + \frac{\pi\Omega_1}{(1-\pi)\Omega_0 + \pi\Omega_1} \left[\frac{1}{\varepsilon_{FH}^S(x^*(t);t)} + \frac{\epsilon(1-\Lambda)}{\varepsilon_{FH}^D(x^*(t);t)} \frac{d\ln(1+\mathrm{x}(\Lambda))}{d\ln\Lambda} \right] \\ & - \frac{\left[v_H(0;t) - v_H(x^*(t);t) \right]}{(1-\pi)\Omega_0 + \pi\Omega_1} \frac{d\pi(t)}{d\ln(1+t)} \quad , \end{aligned} \tag{183}$$

where every object is evaluated at $t = t^*$ and Ω_0 and Ω_1 are defined as

$$\Omega_0 \equiv -\frac{\lambda_{FH}(0;t) E_H(0;t)}{(1+t) Y_H(0;t)} \frac{d\ln Q_{FH}(0;t)}{d\ln(1+t)}, \qquad \Omega_1 \equiv -\frac{\Lambda E_H(x^*(t);t)}{(1+t) Y_H(x^*(t);t)} \frac{d\ln Q_{FH}(x^*(t);t)}{d\ln(1+t)}, \tag{184}$$

both of which are positive because a higher Home tariff lowers imports independent of the chosen regime.

To derive this formula, we obtain the first-order condition by differentiating Home's objective, $v_H^E(t)$, with respect to $\ln(1+t)$. At an interior optimum,

$$[1 - \pi(t)] \frac{dv_H(0;t)}{d\ln(1+t)} + \pi(t) \frac{dv_H(x^*(t);t)}{d\ln(1+t)} + \frac{d\pi(t)}{d\ln(1+t)} \left[v_H(x^*(t);t) - v_H(0;t) \right] = 0. \tag{185}$$

Next, we compute the derivative of log Home welfare v_H in each realized state. Fix any

differentiable state-contingent Foreign export-tax schedule $1 + x(\Lambda)$. Let

$$p_{FH} \equiv (1+x)d_{FH}p_F, \qquad x = 1 + \mathbb{1}_r x(\lambda_{FH}), \qquad \tilde{p}_{FH} = (1+t)p_{FH}.$$

where $\mathbb{1}_r$ is equals to 1 under the non-cooperative regime and zero otherwise. Home income is

$$Y_H = w_H L_H + (\tilde{p}_{FH} - p_{FH})Q_{FH}.$$

Differentiating with respect to $\ln \tilde{p}_{FH}$ gives

$$\frac{\partial \ln Y_H}{\partial \ln \tilde{p}_{FH}} = \frac{1}{Y_H}\left[w_H L_H \frac{\partial \ln w_H}{\partial \ln \tilde{p}_{FH}} + \tilde{p}_{FH}Q_{FH} + \left(\tilde{p}_{FH} - p_{FH}\left[1 + \tfrac{d \ln p_{FH}}{d \ln Q_{FH}}\right]\right)Q_{FH}\frac{\partial \ln Q_{FH}}{\partial \ln \tilde{p}_{FH}}\right]. \tag{186}$$

Because $\tilde{p}_{FH}q_{FH} = \lambda_{FH}E_H$ and $p_{FH}q_{FH} = \lambda_{FH}E_H/(1+t)$, (186) becomes

$$\frac{\partial \ln Y_H}{\partial \ln \tilde{p}_{FH}} = \left[\frac{w_H L_H}{E_H}\frac{\partial \ln w_H}{\partial \ln \tilde{p}_{FH}} + \lambda_{FH} + \left(1 - \frac{1}{1+t}\left[1 + \frac{d \ln p_{FH}}{d \ln Q_{FH}}\right]\right)\lambda_{FH}\frac{\partial \ln Q_{FH}}{\partial \ln \tilde{p}_{FH}}\right]\frac{E_H}{Y_H}. \tag{187}$$

Next, $\partial \ln \tilde{P}_H / \partial \ln \tilde{p}_{FH} = \lambda_{FH}$, and the numeraire condition $p_H = \bar{\varphi}_H^{-1} w_H^{\beta} \tilde{P}_H^{1-\beta} \equiv 1$ implies $\frac{\partial \ln w_H}{\partial \ln \tilde{p}_{FH}} = -\frac{1-\beta}{\beta}\lambda_{FH}$. Using $E_H - Y_H = \frac{1-\beta}{\beta} w_H L_H$, the wage term cancels against the part of $\lambda_{FH}E_H/Y_H$ that exceeds λ_{FH}. Hence

$$\frac{\partial \ln W_H}{\partial \ln \tilde{p}_{FH}} = -\left(1 - \frac{1}{1+t}\left[1 + \frac{d \ln p_{FH}}{d \ln Q_{FH}}\right]\right)\lambda_{FH}\varepsilon^D_{FH}\frac{E_H}{Y_H}, \tag{188}$$

where $\varepsilon^D_{FH} \equiv -\partial \ln Q_{FH} / \partial \ln \tilde{p}_{FH}$. Now decompose the elasticity of the pre-tariff import price. Since $p_{FH} = (1+x)d_{FH}p_F$,

$$\frac{d \ln p_{FH}}{d \ln Q_{FH}} = \frac{d \ln p_F}{d \ln Q_{FH}} + \frac{d \ln(1+x)}{d \ln \lambda_{FH}}\frac{d \ln \lambda_{FH}}{d \ln Q_{FH}}. \tag{189}$$

By definition, $d \ln p_F / d \ln Q_{FH} = 1/\varepsilon^S_{FH}$. Moreover,

$$\frac{d \ln \lambda_{FH}}{d \ln q_{FH}} = \frac{d \ln \lambda_{FH} / d \ln \tilde{p}_{FH}}{d \ln Q_{FH} / d \ln \tilde{p}_{FH}} = \frac{\epsilon \lambda_{HH}}{\varepsilon^D_{FH}},$$

so (189) becomes

$$\frac{d \ln p_{FH}}{d \ln Q_{FH}} = \frac{1}{\varepsilon^S_{FH}} + \frac{\epsilon \lambda_{HH}}{\varepsilon^D_{FH}}\frac{d \ln(1+x)}{d \ln \lambda_{FH}}. \tag{190}$$

Finally, multiply (188) by $\frac{d\ln\tilde{p}_{FH}}{d\ln(1+t)}$ and use $\frac{d\ln Q_{FH}}{d\ln(1+t)} = -\varepsilon^D_{FH}\frac{d\ln\tilde{p}_{FH}}{d\ln(1+t)}$. This gives the generic state derivative

$$\frac{d\ln W_H}{d\ln(1+t)} = \left[1 - \frac{1}{1+t}\left(1 + \frac{1}{\varepsilon^S_{FH}} + \frac{\epsilon\lambda_{HH}}{\varepsilon^D_{FH}}\frac{d\ln(1+x)}{d\ln\lambda_{FH}}\right)\right]\lambda_{FH}\frac{d\ln Q_{FH}}{d\ln(1+t)}\frac{E_H}{Y_H}. \tag{191}$$

Apply (191) to the aligned state first. There $x = 0$, so the corrective term is absent, and

$$\frac{dv_H(0;t)}{d\ln\tau^M_{FH}} = -\Omega_0\left(t - \frac{1}{\varepsilon^S_{FH}(0;t)}\right). \tag{192}$$

Next apply (191) to the non-aligned state. There $x = x^*(t) = x^*(\Lambda)$ and $\lambda_{FH} = \Lambda$, so

$$\frac{dv_H(x^*(t);t)}{d\ln(1+t)} = -\Omega_1\left(t - \frac{1}{\varepsilon^S_{FH}(x^*(t);t)} - \frac{\epsilon(1-\Lambda)}{\varepsilon^D_{FH}(x^*(t);t)}\frac{d\ln(1+x^*(\Lambda))}{d\ln\Lambda}\right). \tag{193}$$

Substituting (192) and (193) into the Bayesian first-order condition (185) and solving for t yields (183). The economic initiation underlying the Bayesian formula is as follows. The first line of (183) represents the terms-of-trade optimal tariffs under the cooperative regime. The second line is the optimal tariff under the non-cooperative regime. This line includes the standard terms of trade plus the corrective term that appears because Foreign's non-cooperative export tax rises with aggregate import dependence. The third line accounts for foreign export policy adjustment along the extensive margin. This term is strictly positive: by construction $v_H(0;t) > v_H(x^*(t);t)$, Lemma 4 in the paper implies $d\pi/d\Lambda \geq 0$, so considering that a higher Home tariff lowers Λ, then it follows that $d\pi/d\ln(1+t) < 0$.

G Optimal Tariffs with Multiple Sectors

(i) Multi-Sector Export Supply Elasticity

In the multi-sector framework, the sector-level price is

$$p_{ni,k} = \tilde{d}_{ni,k}p_n, \qquad p_n = \frac{1}{A_n}w_n^{\beta}\tilde{P}_n^{1-\beta}$$

where $\tilde{d}$ collects iceberg costs and sector-specific productivity shifter, and $\tilde{P}_n = \prod_k \tilde{P}_{n,k}^{\alpha_k}$, with α_k representing the constant sector-level expenditure share. With multiple sectors, we can use the fact that the sector-level CES demand system implies, $\tilde{P}_{F,k} = \tilde{\psi}_{FF,k}p_F\lambda_{FF,k}^{1/\epsilon_k}$,

yielding

$$\tilde{P}_F = \prod_k \tilde{P}_{F,k}^{\alpha_k} = \psi_F p_F \prod_k \lambda_{FF,k}^{\frac{\alpha_k}{\epsilon_k}}$$

where $\psi_F \equiv \sum_k \tilde{d}_{FF,k}$, The balance trade condition is given by

$$\tilde{d}_{FH,k} p_F Q_{FH} = \lambda_{HF}(1+\delta_F) w_F L_F$$

where $Q_{ni} \equiv \sum_k \tilde{d}_{ni,k} q_{ni,k}$ is aggregate trade quantity and $\delta_i \equiv T_i / w_i L_i$. Noting that $\delta_F \approx 0$, we cab plug $p_F = \frac{1}{A_F} w_F^{\beta} \tilde{P}_F^{1-\beta}$ and $\tilde{P}_{F,k} = \psi_F p_F \prod_k \lambda_{FF,k}^{\frac{\alpha_k}{\epsilon_k}}$ into the balanced trade condition to obtain

$$Q_{FH} = \bar{\varphi}_F^{\frac{1}{\beta}} \prod_k \left(\lambda_{FF,k}(p_F)^{\frac{\alpha_k}{\epsilon_k}} \right)^{-\frac{1-\beta}{\beta}} \lambda_{HF}(p_F)\, L_F \tag{194}$$

where $\varphi_F \equiv A_F / \psi_F^{1-\beta}$ and the aggregate shares are now a weighted sum of sector-level shares:

$$\lambda_{FF}(p_F) = \sum_k \lambda_{FF,k}(p_F)\alpha_k, \qquad \lambda_{HF}(p_F) = 1 - \lambda_{FF}(p_F).$$

Taking log derivatives from Equation 194, derives the following expression for the inverse export supply elasticity:

$$\frac{1}{\varepsilon_{FH}^S} \equiv \frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}} = \frac{1}{-\frac{1-\beta}{\beta}\sum_k \frac{\alpha_k}{\epsilon_k}\frac{d\ln\lambda_{FF,k}}{d\ln p_F} + \frac{d\ln\lambda_{HF}}{d\ln p_F}}. \tag{195}$$

The first term in the denominator can be simplified as

$$\sum_k \frac{\alpha_k}{\epsilon_k}\frac{d\ln\lambda_{FF,k}}{d\ln p_F} = \sum_k \frac{\lambda_{FF,k}\alpha_k}{\lambda_{FF}}\frac{d\ln\lambda_{FF,k}}{d\ln p_F} = -\sum_k \frac{\alpha_k}{\epsilon_k}\epsilon_k \lambda_{HF,k} = \lambda_{HF}$$

The second term is obtained by taking the derivatives of the above-listed equation for the aggregate shares:

$$\frac{d\ln\lambda_{HF}}{d\ln p_F} = \sum_k \frac{\lambda_{HF,k}\alpha_k}{\lambda_{HF}}\frac{d\ln\lambda_{HF,k}}{d\ln p_F} = \sum_k \frac{\lambda_{HF,k}\alpha_k}{\lambda_{HF}}\epsilon_k \lambda_{FF,k}.$$

Plugging the above two expressions back into Equation 195, yields

$$\frac{1}{\varepsilon_{FH}^S} \equiv \frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}} = \frac{1}{\sum_k \left[\frac{\lambda_{HF,k}\alpha_k}{\lambda_{HF}}\epsilon_k \lambda_{FF,k}\right] + \frac{1-\beta}{\beta}\lambda_{HF}}.$$

(i) First-Order Condition

The government chooses the vector $\{\tilde{p}_{FH,k}\}_k$ of sector-level buyer prices for imported varieties. The first-order condition w.r.t. $\tilde{p}_{FH,k}$ is

$$\frac{\mathrm{d}\ln W_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = \frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} - \frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = 0$$

Note that

$$Y_H = w_H L_H + \sum_k (\tilde{p}_{FH} - \bar{\tau}_k(\lambda_{FH,k})\tilde{d}_{FH,k} p_F) Q_{FH,k}$$

Following the logic of our baseline analysis, the income derivative becomes

$$\frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = \frac{1}{Y_H}\left[w_H L_H \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} + \tilde{p}_{FH,k} Q_{FH,k} - p_{FH,k} Q_{FH,k} \frac{\partial \ln \bar{\tau}_k(\lambda_{FH,k})}{\partial \ln \lambda_{FH,k}} \frac{\mathrm{d}\ln \lambda_{FH,k}}{\mathrm{d}\ln \tilde{p}_{FH,k}} + \sum_g \left[\left(\tilde{p}_{FH,g} - \left(1 + \frac{\mathrm{d}\ln p_F}{\mathrm{d}\ln Q_{FH}} \frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln Q_{FH,k}}\right) \sum_{g'} p_{FH,g'} Q_{FH,g'} \right) \frac{\mathrm{d}\ln Q_{FH,g}}{\mathrm{d}\ln \tilde{p}_{FH,k}} \right] \right].$$

Which considering that $\frac{\mathrm{d}\ln Q_{FH}}{\mathrm{d}\ln Q_{FH,k}} = \frac{p_{FH,g} Q_{FH,g}}{\sum_{g'} p_{FH,g'} Q_{FH,g'}}$ yields

$$\frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = \left[\frac{w_H L_H}{E_H} \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} + \lambda_{FH,k} + \frac{\lambda_{FH,k}}{\tau^M_{FH,k}} \frac{\partial \ln \bar{\tau}_k}{\partial \ln \lambda_{FH,k}} \epsilon_k \lambda_{HH,k} + \sum_g \left[\left(1 - \frac{1}{1+t_g}(1 + \frac{1}{\varepsilon^S_{FH}}) \right) \lambda_{FH,g} \varepsilon^D_{FH,g} \right] \right] \frac{E_H}{Y_H}.$$

With $\delta_H \approx 0$ and Cobb-Douglas aggregator across sectors, the cross-effects reduce to zero and the above equation simplifies to

$$\frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = \left[\frac{w_H L_H}{E_H} \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} + \lambda_{FH,k} + \left(1 - \frac{1}{\tau^M_{FH,k}} [1 + \frac{1}{\varepsilon^S_{FH}} + \frac{\epsilon_k \lambda_{HH,k}}{\varepsilon^D_{FH,k}} \frac{\partial \ln \bar{\tau}_k}{\partial \ln \lambda_{FH,k}}] \right) \lambda_{FH,k} \varepsilon^D_{FH,k} \right] \frac{E_H}{Y_H}.$$

The price index and wage derivatives can be obtained following the same steps as in the baseline derivation. In particular, noting that $p_H = 1$, we get

$$\frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = \lambda_{FH,k}, \qquad \frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} = -\frac{1-\beta}{\beta} \lambda_{FH,k}$$

The above two equations ensure that $\frac{w_H L_H}{E_H}\frac{\mathrm{d}\ln w_H}{\mathrm{d}\ln \tilde{p}_{FH,k}}\frac{E_H}{Y_H} - (\frac{E_H}{Y_H} - 1)\lambda_{FH,k} = 0$, which simplifies the first-order condition a

$$\begin{aligned}\frac{\mathrm{d}\ln W_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} &= \frac{\mathrm{d}\ln Y_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} - \frac{\mathrm{d}\ln \tilde{P}_H}{\mathrm{d}\ln \tilde{p}_{FH,k}} \\ &= \left[\left(1 - \frac{1}{1+t_k}\left(1 + \frac{1}{\varepsilon^S_{FH}} + \frac{\epsilon_k \lambda_{HH,k}}{\varepsilon^D_{FH,k}}\frac{\partial \ln \bar{\tau}_k}{\partial \ln \lambda_{FH,k}}\right)\right)\lambda_{FH,k}\frac{\mathrm{d}\ln Q_{FH,k}}{\mathrm{d}\ln \tilde{p}_{FH,k}}\right]\frac{E_H}{Y_H} = 0.\end{aligned}$$

The solution to this system yields the optimal tariff formula:

$$t^*_k = \frac{1}{\varepsilon^S_{FH}} + \frac{\epsilon_k \lambda_{HH,k}}{\varepsilon^D_{FH,k}}\frac{\partial \ln \bar{\tau}_k(\lambda_{FH,k})}{\partial \ln \lambda_{FH,k}}$$

H Foreign's Optimal *Sector-Level* Export Policy

This appendix derives the optimal export tax schedule for Foreign under the non-cooperative regime in the multi-sector framework. Building on this result, we then derive an approximate univariate expression for Foreign's sector-level export restriction hazard π_k as a function of sector-level import dependence, Λ_k. To set the stage, we begin with a formal statement of Foreign's optimal export tax problem: The optimal export tax vector $\mathbf{x}^* = \{x^*_k\}$ given Home tariffs $\mathbf{t} = \{t_k\}$ is

$$\mathbf{x}^* \equiv \arg\max v_F(\mathbf{x};\mathbf{t}) = \ln Y_F(\mathbf{x};\mathbf{t}) - \ln \tilde{P}_F(\mathbf{x};\mathbf{t})$$

with foreign income given by $Y_F = w_F L_F + T^X_F$ and $\tilde{P}_F$ only depending on $(\mathbf{x};\mathbf{t})$ through p_F since $p_H \equiv 1$ by choice of numeraire. The first-order condition that solves for $\mathbf{x}^*$ is

$$Y_F\frac{\partial v_F}{\partial \ln(1+x_k)} = w_F L_F \frac{\partial \ln w_F}{\partial \ln(1+x_k)} + \frac{\partial T^X_F}{\partial \ln(1+x_k)} - Y_F\frac{\partial \ln \widetilde{P}_F}{\partial \ln(1+x_k)} = 0 \tag{196}$$

In what follows, we derive the three terms in (196). First, we characterize $\partial T^X_F/\partial \ln(1+x_k)$. The Foreign export-tax revenue is $T^X_F = E_H \sum_g \frac{x_g \alpha_g \lambda_{FH,g}}{(1+t_g)(1+x_g)}$. To derive, the derivative of Home's CES import share $\lambda_{FH,g}$ w.r.t. export taxes, we note that the Home buyer price of Foreign goods in sector g is $\widetilde{p}_{FH,g} = (1+t_g)(1+x_g)d_{FH,g}p_F$. Thus, using the CES import demand system be obtain

$$\frac{\partial \ln \lambda_{FH,g}}{\partial \ln(1+x_k)} = -\epsilon_g \lambda_{HH,g}\left[\mathbf{1}\{g=k\} + \frac{\partial \ln p_F}{\partial \ln(1+x_k)}\right], \tag{197}$$

Differentiating $T_F^X = E_H \sum_g \frac{x_g \alpha_g \lambda_{FH,g}}{(1+t_g)(1+x_g)}$ and using (197) gives

$$\begin{aligned}\frac{1}{E_H}\frac{\partial T_F^X}{\partial \ln(1+x_k)} =& \frac{\alpha_k \lambda_{FH,k}}{(1+t_k)(1+x_k)} - \frac{x_k}{(1+t_k)(1+x_k)}\alpha_k \lambda_{FH,k}\epsilon_k \lambda_{HH,k} \\ &- \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g \frac{x_g}{(1+t_g)(1+x_g)}\alpha_g \lambda_{FH,g}\epsilon_g \lambda_{HH,g} + \frac{T_F^X}{E_H}\frac{\partial \ln E_H}{\partial \ln(1+x_k)}.\end{aligned} \tag{198}$$

To complete the derivation we must substitute for the Home expenditure derivative, $\partial \ln E_H / \partial \ln(1+x_k)$. For this, note that Home's total expenditure is $E_H = \delta_H w_H L_H / \beta$, with $\delta_H \equiv \left[1 - \sum_g \frac{t_g}{1+t_g}\alpha_g \lambda_{FH,g}\right]^{-1}$. Thus, $\frac{\partial \ln E_H}{\partial \ln(1+x_k)} = \frac{\partial \ln w_H}{\partial \ln(1+x_k)} + \frac{\partial \ln \delta_H}{\partial \ln(1+x_k)}$. First, we calculate $\frac{\partial \ln w_H}{\partial \ln(1+x_k)}$ by appealing to the fact that Home's producer price is normalized to one. Differentiating the Home pricing equation $p_H = w_H^\beta \tilde{P}_H^{1-\beta}$ implies

$$\frac{\partial \ln w_H}{\partial \ln(1+x_k)} = -\frac{1-\beta}{\beta}\frac{\partial \ln \widetilde{P}_H}{\partial \ln(1+x_k)} = \alpha_k \lambda_{FH,k} + \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g \alpha_g \lambda_{FH,g}.$$

where the last line substitutes for $\frac{\partial \ln \widetilde{P}_H}{\partial \ln(1+x_k)}$ by noting that the Home price index is $\ln \tilde{P}_H = -\sum_g \frac{\alpha_g}{\epsilon_g} \ln\left(1 + [(1+x_g)(1+t_g)d_{FH,g}p_F]^{-\epsilon_g}\right)$. The derivative $\frac{\partial \ln \delta_H}{\partial \ln(1+x_k)}$ is obtained by invoking and invoking $\frac{\partial \ln \delta_H}{\partial \ln(1+x_k)}$. Combining the these derivatives, delivers

$$\begin{aligned}\frac{\partial \ln E_H}{\partial \ln(1+x_k)} =& \frac{\partial \ln w_H}{\partial \ln(1+x_k)} + \frac{\partial \ln \delta_H}{\partial \ln(1+x_k)} = -\frac{1-\beta}{\beta}\left[\alpha_k \lambda_{FH,k} + \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g \alpha_g \lambda_{FH,g}\right] \\ &- \delta_H \sum_g \left(\frac{t_g}{1+t_g}\alpha_g \lambda_{FH,g}\epsilon_g \lambda_{HH,g}\left[\mathbf{1}\{g=k\} + \frac{\partial \ln p_F}{\partial \ln(1+x_k)}\right]\right)\end{aligned}$$

Substituting the above equation into (198) separates the direct sector k terms from the general equilibrium terms multiplied by $\partial \ln p_F / \partial \ln(1+x_k)$:

$$\begin{aligned}\frac{1}{E_H}\frac{\partial T_F^X}{\partial \ln(1+x_k)} =& \left[\frac{1}{(1+t_k)(1+x_k)} - \frac{x_k \epsilon_k \lambda_{HH,k}}{(1+t_k)(1+x_k)} - \frac{1-\beta}{\beta}\frac{T_F^X}{E_H} - \frac{T_F^X \delta_H}{E_H}\frac{t_k \epsilon_k \lambda_{HH,k}}{1+t_k}\right]\alpha_k \lambda_{FH,k} \\ &- \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g \left[\frac{x_g \epsilon_g \lambda_{HH,g}}{(1+t_g)(1+x_g)} + \frac{1-\beta}{\beta}\frac{T_F^X}{E_H} + \frac{T_F^X \delta_H}{E_H}\frac{t_g \epsilon_g \lambda_{HH,g}}{1+t_g}\right]\alpha_g \lambda_{FH,g}.\end{aligned} \tag{199}$$

Next, we characterize the The remaining terms in the first-order condition, (196), corresponding to Foreign wage and price index derivatives. Foreign's sectoral price indexes move with x_k only through the common producer price term p_F, since $p_H \equiv 1$ by choice of numeraire. Specifically,

$$\frac{\partial \ln \widetilde{P}_F}{\partial \ln(1+x_k)} = \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g \alpha_g \lambda_{FF,g}. \tag{200}$$

The Foreign pricing equation $p_F = w_F^{\beta} \tilde{P}_F^{1-\beta}$ gives $\frac{\partial \ln w_F}{\partial \ln(1+x_k)} = \frac{1}{\beta} \frac{\partial \ln p_F}{\partial \ln(1+x_k)} - \frac{1-\beta}{\beta} \frac{\partial \ln \widetilde{P}_F}{\partial \ln(1+x_k)}$. Using this expression, (200), and $E_F = Y_F + \frac{1-\beta}{\beta} w_F L_F$, and the Foreign labor market clearing, $\frac{w_F L_F}{\beta} = \sum_g \left[\alpha_g \lambda_{FF,g}\right] E_F + \sum_g \left[\frac{\alpha_g \lambda_{FH,g}}{(1+t_g)(1+x_g)}\right] E_H$, one obtains

$$w_F L_F \frac{\partial \ln w_F}{\partial \ln(1+x_k)} - Y_F \frac{\partial \ln \widetilde{P}_F}{\partial \ln(1+x_k)} = E_H \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g \frac{\alpha_g \lambda_{FH,g}}{(1+t_g)(1+x_g)}. \tag{201}$$

Substituting (199) and (201) into (196) yields simplified first-order condition in simplified form

$$\frac{\partial v_F}{\partial \ln(1+x_k)} = \frac{E_H}{Y_F} \left[M_k + \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g M_g \right] = 0 \tag{202}$$

where M_g is defined as follows for each sector g:

$$M_g \equiv \frac{\alpha_g \lambda_{FH,g} \left(1 - x_g \epsilon_g \lambda_{HH,g}\right)}{(1+t_g)(1+x_g)} - \frac{T_F^X}{E_H} \alpha_g \lambda_{FH,g} \left[\frac{1-\beta}{\beta} + \delta_H \frac{t_g}{1+t_g} \epsilon_g \lambda_{HH,g} \right]. \tag{203}$$

Optimal export tax formula. At an interior optimum, the first-order conditions are $M_k + \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g M_g = 0$ for all k. If $1 + \sum_k \frac{\partial \ln p_F}{\partial \ln(1+x_k)} \neq 0$, summing these equations gives $\sum_g M_g = 0$, and hence $M_k = 0$ in each sector. Thus the sector-level optimality condition becomes $\frac{1 - x_k \epsilon_k \lambda_{HH,k}}{(1+t_k)(1+x_k)} = \frac{T_F^X}{E_H} \left[\frac{1-\beta}{\beta} + \frac{\delta_H t_k \epsilon_k \lambda_{HH,k}}{1+t_k} \right]$, rearranging which gives

$$x_k^* = \frac{1 - \frac{T_F^X}{E_H} \left[\frac{1-\beta}{\beta} (1+t_k) + \delta_H t_k \epsilon_k \lambda_{HH,k} \right]}{\epsilon_k \lambda_{HH,k} + \frac{T_F^X}{E_H} \left[\frac{1-\beta}{\beta} (1+t_k) + \delta_H t_k \epsilon_k \lambda_{HH,k} \right]}. \tag{204}$$

Equivalently, we can express the optimal export tax more compactly as

$$x_k^* = \left[\epsilon_k \lambda_{HH,k} \delta_k + \frac{1-\beta}{\beta} \lambda_{FH,k} \widetilde{\delta}_k \right]^{-1},$$

with the sector-level shifters $\widetilde{\delta}_k$ and $\tilde{\delta}_k$ defined as

$$\widetilde{\delta}_k \equiv \frac{(1+\epsilon_k\lambda_{HH,k})(1+t_k)\frac{T_F^X}{E_H}}{1-\left[\frac{1-\beta}{\beta}(1+t_k)+\delta_H t_k\epsilon_k\lambda_{HH,k}\right]}, \qquad \delta_k \equiv 1+\delta_H\frac{t_k}{1+t_k}\lambda_{FH,k}\delta_H\widetilde{\delta}_k$$

A useful special case is $\mathbf{t}=\mathbf{0}$ in which case the export tax formula reduces to:

$$x_k^* = \left[\epsilon_k\lambda_{HH,k}+\frac{1-\beta}{\beta}\lambda_{FH,k}\sum_g\frac{1+\epsilon_k\lambda_{HH,k}}{1+\epsilon_g\lambda_{HH,g}}\frac{\lambda_{FH,g}}{\lambda_{FH,k}}\alpha_g\right]^{-1}$$

To see this, note that the primitive representation of Foreign export tax revenues gives $\frac{T_F^X}{E_H}=\sum_g\frac{x_g^*}{1+x_g^*}\frac{\lambda_{FH,g}}{1+t_g}$. Substituting the optimal export tax rate x^* implied by (204) into this expression and rearranging gives

$$\frac{T_F^X}{E_H} = \sum_g\frac{\lambda_{FH,g}}{(1+t_g)(1+\epsilon_g\lambda_{HH,g})}\Bigg/\left(1+\sum_g\frac{\left[\frac{1-\beta}{\beta}(1+t_g)+\delta_H t_g\epsilon_g\lambda_{HH,g}\right]\lambda_{FH,g}}{(1+t_g)(1+\epsilon_g\lambda_{HH,g})}\right)$$

With zero Home tariffs ($t_k=0$ for all k), we get $\delta_H=1$ and $\delta_k=1$, and the above equation implies $\frac{T_F^X/E_H}{1-\frac{1-\beta}{\beta}T_F^X/E_H}=\sum_g\frac{\lambda_{FH,g}}{1+\epsilon_g\lambda_{HH,g}}$, so that $\tilde{\delta}_k=\sum_g\frac{1+\epsilon_k\lambda_{HH,k}}{1+\epsilon_g\lambda_{HH,g}}\frac{\lambda_{FH,g}}{\lambda_{FH,k}}\alpha_g$.

Deriving π_k as a function Λ_k

Now we show that the sector-level Foreign export restriction hazard is a function of the sector-level import dependence in the non-cooperative state: $\Lambda_k \equiv \lambda_{FH,k}(\mathbf{x}^*;\mathbf{t})$. To demonstrate this note that Foreign's gains from adopting the non-cooperative export tax in sector k relative to Laissez-faire is $\Delta_k \equiv v_F(x_k^*,\mathbf{0}_{-k};\mathbf{t})-v_F(0,\mathbf{0}_{-k};\mathbf{t})$, which following Equation (202) can be calculated as

$$\Delta_k = \int_0^{x_k^*}\frac{E_H}{Y_F}\left[M_k+\frac{\partial\ln p_F}{\partial\ln(1+x_k)}\sum_g M_g\right]d\ln(1+x_k), \tag{205}$$

with every object inside the integral evaluated at $(x_k,\mathbf{0}_{-k};\mathbf{t})$. The direct term takes a simple from around $\mathbf{x}=\mathbf{0}$ since $T_F^X=0$. Specifically, equation (203) gives $M_k|_{\mathbf{x}=\mathbf{0}}=\frac{\alpha_k\lambda_{FH,k}(\mathbf{0};\mathbf{t})}{1+t_k}$. It remains to control the common producer-price term in (205). The first-order conditions imply that $\frac{\partial\ln p_F}{\partial\ln(1+x_k)}\sum_g M_g=0$ at $\mathbf{x}^*$. Thus applying the the mean-value theorem between the point $(x_k,\mathbf{0}_{-k};\mathbf{t})$ on the integration path and the non-cooperative

point $(\mathbf{x}^*; \mathbf{t})$, obtains

$$\frac{\partial \ln p_F}{\partial \ln(1+x_k)} \sum_g M_g \Big|_{(x_k, \mathbf{0}_{-k}; \mathbf{t})} = O\left(\max_g x_g^*\right)$$

Combining these points gives and noting that $\ln(1+x^*) \approx x^*$, yields $\Delta_k = \frac{E_H}{Y_F} \frac{\alpha_k \lambda_{FH,k}}{1+t_k} x_k^* + O(x_k^* \max_g x_g^*)$ with all the variables evaluated at $(\mathbf{x} = \mathbf{0}; \mathbf{t})$. Since the CES share system is continuously differentiable in export taxes, $\lambda_{FH,k}(\mathbf{0}; \mathbf{t}) = \Lambda_k + O(\max_g x_g^*)$, which delivers

$$\Delta_k = \frac{E_H}{Y_F}\Big|_{\mathbf{x}=\mathbf{0}} \frac{\alpha_k \Lambda_k x_k^*(\Lambda_k)}{(1+t_k)(1+x_k^*(\Lambda_k))} + O\left(x_k^* \max_g x_g^*\right). \tag{206}$$

where we invoke the previously derived optimal export tax formula, whereby $x_k^* = x_k^*(\Lambda_k)$. Foreign applies the noncooperative export tax in sector k whenever the cost of noncooperation satisfies $\eta_k \leq \Delta_k$. Formally,

$$\pi_k(\Lambda_k) = G_k(\Delta_k(\Lambda_k)) \approx G_k\left(\frac{\alpha_k \Lambda_k x_k^*(\Lambda_k)}{(1+t_k)(1+x_k^*(\Lambda_k))}\right).$$

Under the restriction $\epsilon_k > (1-\beta)/\beta$, the optimal export tax $x_k^*(\Lambda_k)$ is increasing in Λ_k. Thus, since $x_k^*/(1+x_k^*)$ is increasing in x_k^*, the leading term in (206) is increasing in Λ_k given t. At points where G_k has density g_k, we can state this as

$$\frac{\partial \pi_k(\Lambda_k)}{\partial \Lambda_k} = g_k(\Delta_k(\Lambda_k)) \frac{\partial \Delta_k(\Lambda_k)}{\partial \Lambda_k} \geq 0,$$

with strict inequality whenever $g_k(\Delta_k(\Lambda_k)) > 0$. Altogether, these equations establish that to a reasonable approximation, π_k can be expressed as an increasing univariate function of Λ_k, mirroring the aggregate relationship in the baseline model.

I Validating the Certainty Equivalence Representation

This appendix evaluates the certainty-equivalent representation using the same calibrated environment as the macro counterfactual. The economy contains latent non-traded and traded composites. The ISIC G–T service value-added share proxies each region's non-traded expenditure weight, but all observed cross-border deliveries, including measured service trade, enter the traded composite. We retain the two observed bilateral directions and then construct the zero-transfer exact-hat benchmark. Home and Foreign have separate economy-wide value-added shares, and both policy wedges are

applied only to the traded composite. The probability of noncooperation is $\pi(\lambda_T) = \lambda_T^{\zeta}$, using the stock-OLS estimates $\zeta_{US} = 0.29$ and $\zeta_{EU} = 0.23$.

We distinguish the exact representation from the fitted shortcut used in the counterfactual algorithm. For any Home traded-composite tariff t, define the exact certainty-equivalent Foreign export tax $\bar{x}(t)$ by

$$v_H(\bar{x}(t); t) = [1 - \pi(t)]v_H(0; t) + \pi(t)v_H(x^*(t); t), \qquad \pi(t) = \lambda_{FH,T}(x^*(t); t)^{\zeta}.$$

This equation is an exact welfare identity. If the induced schedule $\bar{x}(\bar{\lambda}_{FH,T})$ is single-valued and increasing, Lemma 5 implies that the Bayesian problem and the deterministic policy problem using that exact schedule have the same Home objective. We evaluate this condition at 121 gross tariff rates, $1 + t \in [1, 1.875]$. For both regions, the exact schedule is strictly increasing throughout the grid.

The numerical counterfactual replaces the exact schedule with

$$\widehat{x}(t) = \hat{a}\, \pi(\bar{\lambda}_{FH,T}(t))x^*(\bar{\lambda}_{FH,T}(t)).$$

We estimate the coefficient separately for each region by

$$\hat{a} = \arg\min_{a} \sum_{j=1}^{121} \left[\bar{x}_j - a\, \pi(\bar{\lambda}_{FH,T,j})x^*(\bar{\lambda}_{FH,T,j})\right]^2.$$

The fitted values are $\hat{a} = 0.76$ for the United States and 0.81 for the European Union. Figure 8 compares the exact and fitted schedules. The maximum absolute fitted-tax errors are 0.24 percentage points for the United States and 0.29 percentage points for the European Union. These errors quantify the fitted shortcut; they do not qualify the exact certainty-equivalence result.

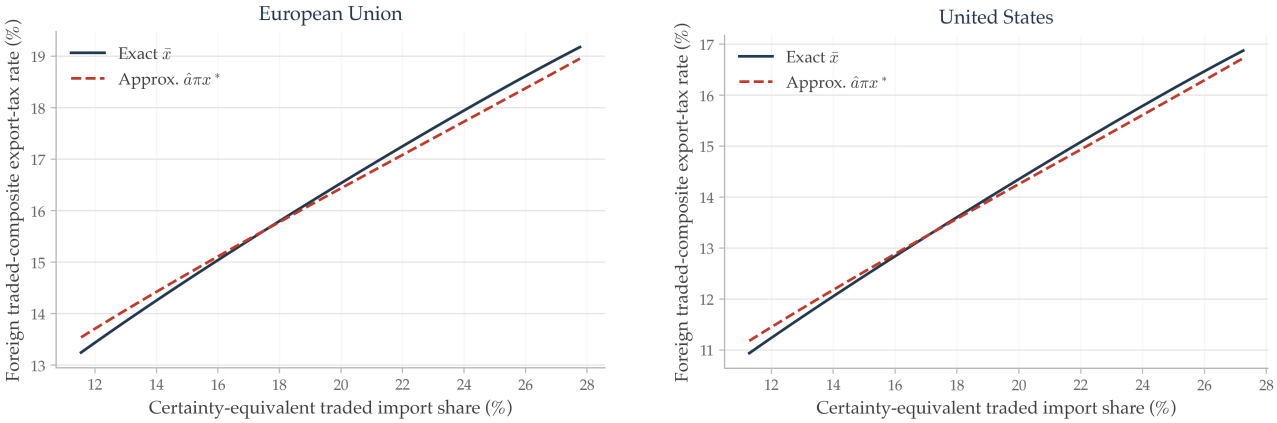
We next compare the Bayesian objective with the deterministic objective that uses the fitted export tax. Specifically,

$$W^B(t) = [1 - \pi(t)]v_H(0; t) + \pi(t)v_H(x^*(t); t), \qquad W^D(t) = v_H(\widehat{x}(t); t).$$

The exact certainty-equivalent schedule characterized in Figure 8 would make these objectives coincide point by point. Figure 9 instead measures the additional discrepancy created by substituting the fitted shortcut.

For the United States, the fitted deterministic optimum is 31.01 percent and the Bayesian optimum is 31.50 percent, an absolute gap of 0.50 percentage points. Rela-

Figure 8: Exact and Fitted Certainty-Equivalent Traded Export Taxes

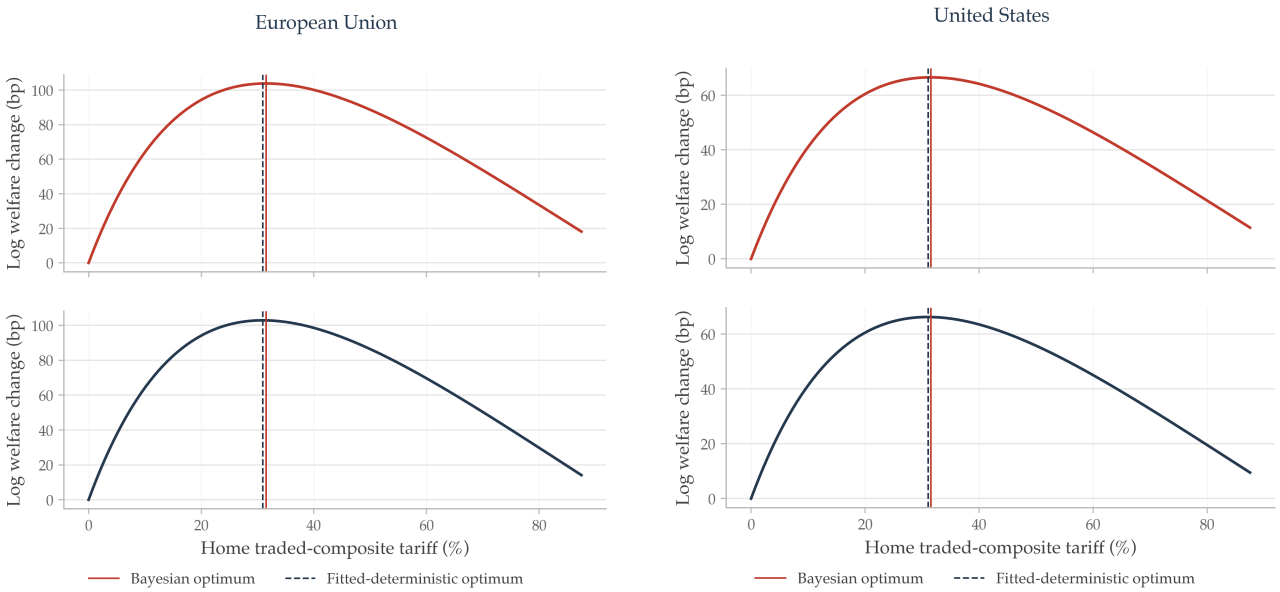


Notes: The figure plots the exact certainty-equivalent Foreign traded-composite export tax and the fitted shortcut $\hat{a}\pi(\bar{\lambda}_T)x^*(\bar{\lambda}_T)$ against the certainty-equivalent traded import share for the European Union (left) and the United States (right). Solid lines show the exact schedules; dashed lines show the fitted schedules. Both exact schedules are strictly increasing on the 121-point gross-tariff grid $1 + t \in [1, 1.875]$. The largest absolute fitted-tax error is 0.29 percentage points for the European Union and 0.24 percentage points for the United States.

tive to the zero-tariff benchmark, the Bayesian-optimal tariff raises welfare by 0.67 percent. Evaluated under exact Bayesian welfare, using the fitted-deterministic tariff instead produces approximation-induced policy regret of9.44×10^{-5} percent. For the European Union, the corresponding tariff rates are 30.93 and 31.51 percent, an absolute gap of 0.59 percentage points, and the Bayesian-optimal tariff raises welfare by 1.04 percent. The approximation-induced policy regret is2.06×10^{-4} percent. Thus, the approximation is not an exact statement about tariff rates; its practical accuracy comes from the flatness of welfare near the optimum.

The exercises therefore support two distinct conclusions. First, the exact certainty-equivalent schedules are strictly increasing at every evaluated tariff node, providing numerical support for the monotonicity condition required by Lemma 5. Second, the fitted formula used in the higher-dimensional counterfactual introduces small but nonzero export-tax and optimal-tariff errors. Those errors have negligible welfare consequences at the calibrated optima, which supports use of the fitted representation for computation while preserving the conceptual distinction between exact equivalence and numerical approximation.

Figure 9: Bayesian and Fitted-Deterministic Optimal Traded Tariffs



Notes: The figure plots log welfare relative to free trade, in basis points, against the Home traded-composite tariff for the European Union (left) and the United States (right). The upper panels show exact Bayesian expected log welfare; the lower panels show deterministic welfare under the fitted certainty-equivalent export tax. Vertical lines mark the Bayesian and fitted-deterministic optima. The fitted-deterministic and Bayesian tariffs are 30.93 and 31.51 percent for the European Union, and 31.01 and 31.50 percent for the United States.

J The 2018 Tariff Episode: Data and Robustness

This appendix provides the institutional and data background for the 2018 tariff episode studied in the main text. The first part describes the underlying trade and tariff data and the three 2018 tariff programs, organized around Table 6. The second part probes, line by line, the robustness of the prewar dependence-growth pattern in Figure 2.

J.1 Background: Section 301 Tariffs

The data are the Fajgelbaum et al. (2020) replication flow: monthly U.S. imports by partner country and ten-digit HS product line, from January 2013 through April 2019. Each record carries the import value, the statutory tariff rate applied that month, and flags marking whether the product falls under each 2018 tariff program. China is one partner among many. The flow records a statutory rate only in months when the line actually traded with the partner, a fact that drives the measurability limits documented below.

From this flow, we keep the 17,891 HS10 lines with positive total U.S. imports in 2017.

Percentages in Table 6 use this universe, or a stated subset, as the denominator; trade values enter only as shares or weights, and dollar levels are never reported.

Three programs raised the 2018 tariffs, and they differ sharply in size. Section 301 is the China-specific program and by far the largest: 10,703 lines, or 59.8 percent of the universe. Section 232 (steel and aluminum) and Section 201 (solar panels and washing machines) are global—they tax the product from every origin—and small: 808 lines together, or 4.5 percent. The remaining 6,380 lines (35.7 percent) faced no 2018 program. The three programs do not overlap; each line belongs to at most one. Table 6 summarizes the programs in two blocks: lines and measurability, then realized statutory changes.

Table 6: The 2018 Tariff Programs: Lines Covered, Measurability, and Realized Statutory Changes

	Section 301	Sections 232/201	No program	All lines
Tariff lines and measurability				
Tariff lines importing in 2017	10,703	808	6,380	17,891
as % of universe	59.8%	4.5%	35.7%	100.0%
China trade in 2013 and 2017 ($\Delta \ln \lambda$ computable)	7,276	474	4,000	11,750
China trade in both Decembers (Δt computable)	5,868	222	3,601	9,691
Realized statutory change, Dec. 2017 to Dec. 2018				
Statutory China rate increased	5,675	222	12	5,909
Rate unchanged	193	0	3,586	3,779
Rate decreased	0	0	3	3
Mean statutory tariff change $100 \times \Delta \ln(1+t)$	11.92%	19.93%	0.04%	7.69%

Notes: The universe is 17,891 HS10 lines with positive 2017 import value in the Fajgelbaum–Goldberg–Kennedy–Khandelwal (2020) flow. Program flags are time-invariant product attributes: m_china_hit identifies Section 301, steel/aluminum flags identify Section 232, and solar/washer flags identify Section 201. The three groups are mutually exclusive in these data, as asserted in construction. $\Delta \ln \lambda$ is computable when China imports are positive in both 2013 and 2017. Δt is the December-to-December change in the line's import-value-weighted mean statutory China rate t and requires China-origin December trade in both years. The realized-change block is restricted to those lines. Unchanged Section 301 lines are targeted lines whose statutory rate never moved through April 2019, including proposed-list removals and Chapter 99 carve-outs. The reported mean is $100 \times \Delta \ln(1+t)$, in percent.

The Section 301 flag marks targeting, not implementation. A flagged line was swept into the Section 301 process, but the statutory rate actually moves only for lines that survived to a final list and received Chapter 99 duty language (heading 9903.88.xx) covering the ten-digit line. The gap between the two is the public-comment bargaining: USTR removed hundreds of proposed lines before implementation, and the Chapter 99 text maps imperfectly onto HS10 children in some product areas. In the data, 5,868 flagged lines have a measurable statutory change; of these, 5,675 rose, none fell, and 193 stayed flat through the end of the data. The 193 flat lines are targets that were bargained out, not measurement noise: they cluster in the chapters the list process cut—headgear,

hand tools, and a bloc of food products. Their December rows do carry statutory rates, and their China dependence had been growing even faster than that of the tariffed lines. We therefore keep them in the treated group: membership on a Section 301 list is the treatment, and dropping or reclassifying lines by whether the tariff ultimately landed would condition on the outcome. Dropping them anyway leaves the results essentially unchanged.

Two requirements thin the universe, and the first block of Table 6 shows both. To measure the five-year growth in a line's China share, the line must have positive China imports in both 2013 and 2017—true for 11,750 lines (7,276 of them Section 301). To measure the December 2017 to December 2018 statutory tariff change, the line must have traded with China in both Decembers—true for 9,691 lines (5,868 Section 301). The figures need both at once, which holds for 8,520 lines: 5,244 Section 301, 3,068 with no program, and 208 under Sections 232 and 201. This is the sample of Figure 2.

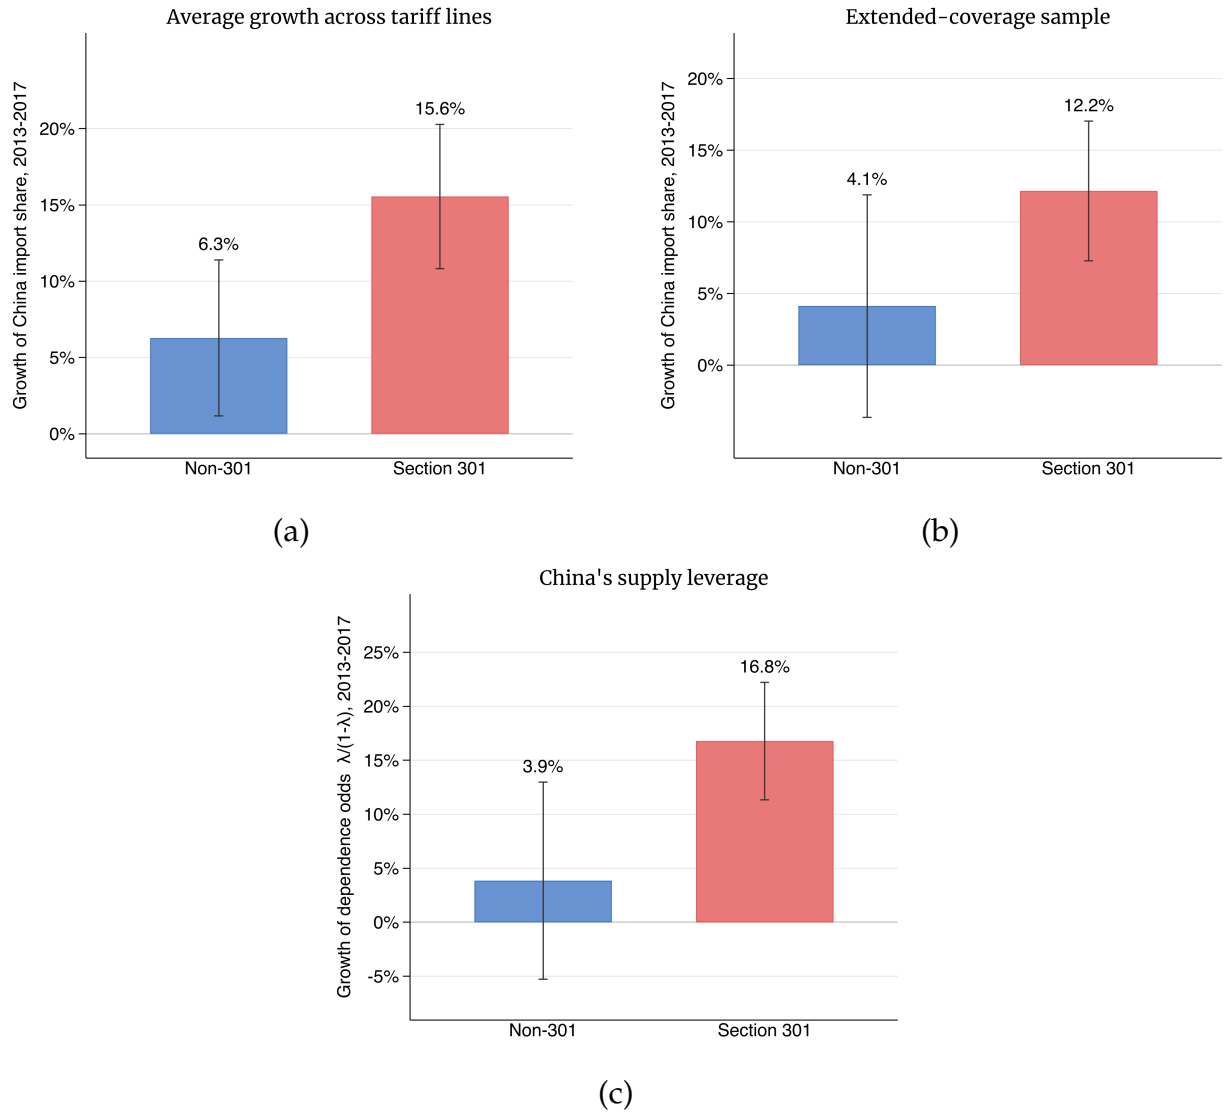
In every exhibit, the comparison group is the set of lines outside the Section 301 lists, labeled Non-301. It contains the Section 232 and 201 lines, which are themselves tariffed—their statutory China rate rose about 20 log points on average—so the group cannot honestly be called untariffed. Placing tariffed goods in the comparison group works against finding a gap, which makes the convention conservative.

The second block of Table 6 shows that, among lines with a measurable change, the contrast across programs is sharp. Under Section 301, 5,675 lines rose, 193 stayed flat, and none fell; the mean statutory change—measured as $100 \times \Delta \ln(1+t)$—is 11.92 percent. Under Sections 232 and 201, all 222 measurable lines rose, by 19.93 percent on average, reflecting the 25-percent-plus steel and solar rates; the global programs had no comment-period escape hatch. Among no-program lines, 3,586 stayed flat, 12 rose, and 3 fell, for a mean of 0.04 percent—essentially unchanged, as they should be.

J.2 Robustness of the Motivating Evidence

Figure 2 presents the prewar dependence-growth pattern as an aggregate contrast between Section 301 goods and the rest (Fact 2 of Section 2). The three panels of Figure 10 ask the same question—did Section 301 goods gain China dependence faster than the rest?—but measure it line by line rather than in the aggregate, so each estimate carries a confidence interval. In each panel a bar is the unweighted mean, across a group's HS10 lines, of the 2013–2017 growth in China dependence ($100 \times \Delta \ln$, in percent); the whiskers are 95 percent confidence intervals from standard errors clustered by HS2 chapter. The groups are Section 301 lines and all others ("Non-301").

Figure 10: Line-level growth of China import dependence, 2013–2017



Notes: Each bar is the unweighted mean, across a group's HS10 lines, of the 2013–2017 growth in China dependence ($100 \times \Delta \ln$, in percent); whiskers are 95 percent confidence intervals from standard errors clustered by HS2 chapter. The "Non-301" group includes the few lines under the global Section 232 and 201 tariffs. Panel (a): dependence is λ, China's share of a line's imports, on the sample of Figure 2 (8,520 lines: 5,244 Section 301, 3,276 Non-301). Panel (b): the tariff change is measured over schedule-stable windows—the 2017 average against the December 2018–April 2019 average—widening the sample to 10,727 lines (6,622 and 4,105). Panel (c): dependence is the odds $\lambda/(1-\lambda)$; the 68 lines sourced entirely from China are dropped, leaving 8,452 lines (5,229 and 3,223). Source: Fajgelbaum et al. (2020).

Panel (a) is the baseline. Dependence is λ, China's share of a line's imports. On the 8,520 lines with China trade in both years and an observed statutory tariff change—the sample of Figure 2—Section 301 lines grew by 15.6 percent against 6.3 percent for the rest, a difference of 9.3 percentage points ($p = 0.005$). The main-text contrast is

therefore not an artifact of aggregation: the average targeted line was deepening its China dependence significantly faster than the average non-targeted line.

Panel (b) applies the same measure to a wider sample. The statutory tariff change can be read only for lines that traded with China in both Decembers, which drops lines that shipped in other months. To recover them, panel (b) measures the tariff change over schedule-stable windows—the 2017 average against the December 2018–April 2019 average, across which the statutory rate does not move—widening the sample to 10,727 lines. The gap holds at 8.0 percentage points (12.2 versus 4.1 percent, $p = 0.068$). The recovered lines trade with China intermittently, so their share growth is measured from thinner flows; the gap attenuates slightly and loses precision, as expected.

Panel (c) recasts dependence as leverage. It replaces the share λ with the dependence odds $\lambda/(1-\lambda)$, the quantity that governs an exporter's market power in the model: as the domestic alternative $1-\lambda$ shrinks, the odds rise steeply. The 68 lines sourced entirely from China ($\lambda = 1$, undefined odds) are dropped. Among the remaining 8,452 lines, Section 301 lines grew by 16.8 percent against 3.9 percent for the rest, a gap of 12.9 percentage points ($p = 0.012$)—the targeted goods were climbing the steep, high-dependence part of the curve, where the model says leverage grows fastest.

Across the cuts examined—December or extended-coverage sample, share or odds—the Section 301 gap in dependence growth ranges from 8 to 13 percentage points and never changes sign. The December, line-level, share-based cut in panel (a) is the cleanest benchmark: its sample comprises the goods that trade with China continuously, which is also where a China-dependence share is an economically meaningful object.

K Textual Classification and Validation Evidence

This appendix documents the textual classification used in Section 8.1. The object of the exercise is official public framing: whether USTR's communications around the 2018 tariff rollout invoked the strategic-risk mechanism, at what dates, and with what intensity. The classifications are not a reconstruction of private deliberations and are not used as structural inputs to the tariff model. The appendix proceeds from corpus construction (Appendix K.1) through classification design (Appendix K.2), aggregation (Appendix K.3), sampling precision (Appendix K.4), robustness (Appendix K.5), replication across classifiers and sensitivity to GPT-OSS decoding temperature (Appendix K.6), validation status (Appendix K.7), and the limitations that bound interpretation (Appendix K.8).

K.1 Corpus and Unit of Analysis

The corpus is the universe of 2,983 press releases published by the Office of the United States Trade Representative between January 2010 and December 2025. The sample is partitioned into six periods: a pre-2017 baseline (1,159 releases), the 2017 transition (109), the 2018 tariff rollout (109), the 2019–2020 follow-through (196), the 2021–2024 period (1,306), and a 2025 post-sample window (104). The unit submitted to the classifier is a local passage, an overlapping two-sentence window within a release, so that claims split across a sentence boundary are seen intact; the corpus yields 55,003 such passages, an average of about 18 per release. Classification occurs strictly at the passage level: the classifier sees the passage, not the document's date, title, or period, so no information about the 2018 episode enters the labeling of any individual passage beyond the passage's own text.

K.2 Classification Design

Table 7 reports the six-narrative ontology. The first three narratives record exposure, importance, and fragility separately: import dependence, criticality, and vulnerability. A passage can therefore mention foreign sourcing without being coded as strategic risk, and can mention an important sector without linking it to foreign exposure. The two mechanism-specific narratives are geopolitical risk, covering references to state power, coercion, rivalry, security risk, and foreign-government conduct, and the strategic vulnerability narrative, covering the conjunction of foreign exposure with strategic importance or vulnerability. The distinction is central for the paper: the model is not about dependence per se but about the policy value of reducing exposure when foreign sourcing creates strategic vulnerability. The sixth narrative, mitigation or de-risking, records proposed remedies and serves as a comparison narrative.

Three design choices prevent the classifier from conflating ordinary trade-policy discourse with strategic-vulnerability discourse. First, labels satisfy a strong-entailment standard: a narrative is positive only when an ordinary reader would treat the claim as part of the passage's message without importing adversarial intent, strategic importance, or supply logic from background knowledge. Sector names, country mentions, and generic tariff references do not by themselves satisfy the standard. Second, three hard carve-outs bound the negative space. Pure U.S. outbound export controls carry a hard zero for import dependence, vulnerability, and the strategic vulnerability narrative (the geopolitical-risk narrative may still apply when the control is framed as a response to rival-state conduct). Foreign barriers to U.S. exports do not by themselves establish

Table 7: Textual Ontology for USTR Communications

Narrative	Classification target
Import dependence	A passage links a U.S. sector, input, market, or capability to reliance on foreign supply, foreign production, or foreign firms.
Criticality	A passage identifies a good, sector, technology, input, or capability as important for national welfare, industrial capacity, infrastructure, innovation, health, or security.
Vulnerability	A passage describes exposure to disruption, coercion, unfair practices, shortages, supply interruption, or other fragility in the relevant economic relationship.
Geopolitical risk	A passage connects trade or investment exposure to state power, strategic rivalry, coercive policy, security risk, or foreign-government behavior.
Mitigation or de-risking response	A passage describes policy actions intended to reduce exposure, diversify supply, rebuild capacity, or otherwise manage a strategic trade vulnerability.
Strategic vulnerability narrative	A passage combines foreign exposure with strategic importance or vulnerability, so that reliance on the foreign relationship is framed as a policy-relevant risk rather than only a commercial fact.

Notes: The mechanis
tive. Import depend
The strategic-risk ch
state behavior, rivalr

any of the key narratives. Routine market-access disputes, WTO cases, and standard anti-dumping or countervailing-duty language are hard negatives unless the same passage also describes dominance, restriction, lock-in, or a strategic supply-risk mechanism. Third, the strategic vulnerability narrative is the most demanding: it requires explicit textual linkage between strategic importance and a risky or coercible supply relationship, and multiple positive component labels within a passage do not imply the strategic vulnerability narrative unless the linkage is itself explicit.

Table 8 illustrates both boundaries with classified passages. The positive examples show that the mechanism-specific narratives require language tying exposure to coercive foreign policy, state-backed competition, or vulnerability in strategically important sectors. The hard negatives show that country names, meeting announcements, tariff-list references, and ordinary market-access language do not identify the mechanism.

K.3 Aggregation

Let d index releases, w index passages within a release, and ℓ index narratives. For each passage-narrative pair, $a_{\ell w} = 1$ if passage w is classified positive for narrative ℓ and zero otherwise. With N_d passages in release d , the two document-level outcomes are

$$E_{\ell d} = \mathbf{1}\Big\{ \sum_{w \in d} a_{\ell w} > 0 \Big\}, \qquad I_{\ell d} = \frac{1}{N_d} \sum_{w \in d} a_{\ell w}.$$

Table 8: Selected Positive Examples and Hard Negatives

Passage type	Passage excerpt	Classification logic
Positive: geopolitical risk and strategic vulnerability narrative	"unprecedented threat posed by China's theft of our intellectual property, the forced transfer of American technology, and its cyber attacks"	The passage ties foreign economic exposure to coercive state behavior and to U.S. leadership in technology and innovation.
Positive: geopolitical risk and strategic vulnerability narrative	"policies that coerce American companies into transferring their technology and intellectual property to domestic Chinese enterprises"	The passage treats technology transfer as a foreign-government policy risk, not simply as ordinary import dependence.
Positive: geopolitical risk and strategic vulnerability narrative	"seizing economic dominance of certain advanced technology sectors"	The passage frames advanced technology sectors as strategically important and links exposure to foreign industrial policy.
Hard negative	"Beijing, China Closed Press"	The passage names a country and an event location but contains no substantive claim about exposure, vulnerability, or strategic risk.
Hard negative	"Click here to view the final tariff list."	The passage is an administrative reference to a policy document rather than a statement about the rationale for protection.
Hard negative	"U.S. Beef Gains New Market Access in Morocco"	The passage concerns market access and export promotion, not strategically consequential exposure to foreign supply or coercion.

Notes: Positive e
within the annou
tariff references,
sification.

The extensive margin is the period average of $E_{\ell d}$ and mention intensity is the period average of $I_{\ell d}$. The first measure asks whether a release contains a narrative at all; the second asks how central the narrative is within the release. Because both measures are defined at the release level, the overlap between adjacent passages affects neither the sign nor the interpretation of the period contrasts.

A conservative administrative gate routes 101 of the 2,983 releases to all-zero labels before classification. The gate fires only when a release simultaneously exhibits at least three administrative-signal patterns, has at least 80 percent purely administrative passages, and contains no instrument, sector, risk, or mitigation hints. Gated releases remain in the universe and enter every denominator, so the gate can only attenuate measured framing. No release was gated in 2017 or 2018; the 101 gated releases fall in the pre-2017 baseline (62) and the 2021–2024 period (39), where administrative scheduling releases are most prevalent.

K.4 Sampling Precision of the Key Contrasts

The analysis reports descriptive shares. Because the underlying counts are small integers over known numbers of releases, exact finite-sample statements are available. Table 9 reports them for the contrasts used in the main text: the count and share of positive releases in each period, the difference in percentage points, and a two-sided Fisher exact test of equal proportions. Three features stand out. First, the central contrasts are decisive at any conventional level: the 2018 geopolitical-risk margin differs from the pre-2017 baseline at $p = 2.5 \times 10^{-10}$, and the Wilson intervals of the two periods, [19.2, 35.6] against [4.8, 7.6], are separated by nearly twelve percentage points. Second, the event-window contrasts are decisive despite resting on only 31 releases, because the baseline rates are so low. Third, the placebo-like contrasts match the composition result in Table 4: bare import dependence and mitigation do not move significantly in 2018, so the shift is specific to the mechanism narratives rather than a general inflation of trade language.[13]

Table 10 reports the complete period-level evidence for all six narratives and both margins, and Table 11 reports the event-time evidence in finer bins. The March 22 to July

[13]The tests treat releases as independent draws within period. Releases are not literally independent, since one policy episode generates several related communications. Three considerations bound the concern. The contrasts of interest exceed twenty percentage points on baseline shares below seven percent, so they survive substantial variance inflation. The announcement-to-implementation and post-implementation windows are disjoint samples, and each is separately distinguishable from the 2017 baseline for geopolitical risk at $p < 10^{-4}$. And the composition of the shift across narratives, which no clustering structure can manufacture, matches the mechanism's prediction.

Table 9: Counts, Confidence Intervals, and Exact Tests for the Key Contrasts

Narrative	Contrast	Counts	Shares (%)	Change (pp)	Fisher p
Panel A. The 2018 rollout					
Geopolitical risk	2018 vs. pre-2017	29/109 vs. 70/1159	26.6 vs. 6.0	+20.6	2.5×10^{-10}
Strategic vulnerability narrative	2018 vs. pre-2017	17/109 vs. 45/1159	15.6 vs. 3.9	+11.7	6.4×10^{-6}
Geopolitical risk	2018 vs. 2017	29/109 vs. 4/109	26.6 vs. 3.7	+22.9	2.1×10^{-6}
Strategic vulnerability narrative	2018 vs. 2017	17/109 vs. 3/109	15.6 vs. 2.8	+12.8	0.002
Panel B. Event windows within 2018 (vs. 2017 baseline)					
Geopolitical risk	Mar. 22–Jul. 5	10/31 vs. 4/109	32.3 vs. 3.7	+28.6	4.1×10^{-5}
Strategic vulnerability narrative	Mar. 22–Jul. 5	7/31 vs. 3/109	22.6 vs. 2.8	+19.8	0.001
Geopolitical risk	Jul. 6–Dec. 31	15/56 vs. 4/109	26.8 vs. 3.7	+23.1	2.6×10^{-5}
Strategic vulnerability narrative	Jul. 6–Dec. 31	7/56 vs. 3/109	12.5 vs. 2.8	+9.7	0.032
Panel C. Persistence and out-of-sample check (vs. pre-2017)					
Geopolitical risk	2021–2024	160/1306 vs. 70/1159	12.3 vs. 6.0	+6.2	1.1×10^{-7}
Strategic vulnerability narrative	2021–2024	73/1306 vs. 45/1159	5.6 vs. 3.9	+1.7	0.058
Geopolitical risk	2025	26/104 vs. 70/1159	25.0 vs. 6.0	+19.0	6.2×10^{-9}
Strategic vulnerability narrative	2025	19/104 vs. 45/1159	18.3 vs. 3.9	+14.4	2.0×10^{-7}
Panel D. Placebo-like contrasts, 2018 vs. pre-2017					
Import dependence	2018 vs. pre-2017	12/109 vs. 77/1159	11.0 vs. 6.6	+4.4	0.113
Mitigation or de-risking	2018 vs. pre-2017	6/109 vs. 96/1159	5.5 vs. 8.3	−2.8	0.362

Notes: Counts report positive releases over total releases in each period. Fisher p -values are two-sided exact tests of equal proportions treating releases as independent draws; the footnote in Appendix K.4 discusses this assumption. All counts come from the canonical GPT-5.4-mini run.

5 window is the period in which the tariff action had been announced but the first list had not yet taken effect; it is the natural window in which an administration explains a policy it has committed to but not yet executed, and it is where both mechanism-specific margins peak.

Figure 11 plots the full 2010–2025 monthly extensive margin for the four main narratives with a six-month rolling average and reference lines at the March 22 and July 6, 2018 dates. The eight-year pre-2017 history shows that strategic-risk language was not a recurring feature of USTR communications that happened to peak in 2018: the 2018 spike has no precedent in the baseline period. Figure 12 shows the same pattern on the intensity margin, which rises only when mechanism-specific language becomes central within releases rather than incidental to them.

K.5 Robustness

Table 12 reports two families of checks on the mechanism-specific narratives. Panel A retains only positive passage classifications with classifier confidence of at least 0.90. This is a one-sided screen: it can remove positives but never reinterprets uncertain negatives, so it should be read as a conservative positive-evidence bound rather than a calibrated

Table 10: Full-Period Textual Evidence Across All Frames

Period	Narrative	Documents	Windows	Avg. windows/doc.	Extensive margin	Mention intensity
Pre-2017 baseline	Import dependence	1159	24716	21.3	6.6%	0.5%
Pre-2017 baseline	Criticality	1159	24716	21.3	17.1%	3.1%
Pre-2017 baseline	Vulnerability	1159	24716	21.3	6.8%	1.0%
Pre-2017 baseline	Geopolitical risk	1159	24716	21.3	6.0%	0.9%
Pre-2017 baseline	Mitigation / de-risking response	1159	24716	21.3	8.3%	0.9%
Pre-2017 baseline	Strategic vulnerability narrative	1159	24716	21.3	3.9%	0.4%
2017 transition	Import dependence	109	1931	17.7	7.3%	0.6%
2017 transition	Criticality	109	1931	17.7	11.9%	2.6%
2017 transition	Vulnerability	109	1931	17.7	3.7%	0.7%
2017 transition	Geopolitical risk	109	1931	17.7	3.7%	0.8%
2017 transition	Mitigation / de-risking response	109	1931	17.7	5.5%	0.7%
2017 transition	Strategic vulnerability narrative	109	1931	17.7	2.8%	0.5%
2018 tariff rollout	Import dependence	109	2268	20.8	11.0%	1.6%
2018 tariff rollout	Criticality	109	2268	20.8	26.6%	5.7%
2018 tariff rollout	Vulnerability	109	2268	20.8	14.7%	2.4%
2018 tariff rollout	Geopolitical risk	109	2268	20.8	26.6%	7.2%
2018 tariff rollout	Mitigation / de-risking response	109	2268	20.8	5.5%	0.4%
2018 tariff rollout	Strategic vulnerability narrative	109	2268	20.8	15.6%	2.4%
2019-2020 follow-through	Import dependence	196	3357	17.1	6.6%	0.8%
2019-2020 follow-through	Criticality	196	3357	17.1	17.9%	4.3%
2019-2020 follow-through	Vulnerability	196	3357	17.1	7.7%	1.2%
2019-2020 follow-through	Geopolitical risk	196	3357	17.1	9.7%	2.4%
2019-2020 follow-through	Mitigation / de-risking response	196	3357	17.1	7.1%	1.3%
2019-2020 follow-through	Strategic vulnerability narrative	196	3357	17.1	3.6%	0.9%
2021-2024 resilience period	Import dependence	1306	20738	15.9	2.7%	0.3%
2021-2024 resilience period	Criticality	1306	20738	15.9	19.7%	5.0%
2021-2024 resilience period	Vulnerability	1306	20738	15.9	11.9%	2.3%
2021-2024 resilience period	Geopolitical risk	1306	20738	15.9	12.3%	3.4%
2021-2024 resilience period	Mitigation / de-risking response	1306	20738	15.9	14.4%	3.2%
2021-2024 resilience period	Strategic vulnerability narrative	1306	20738	15.9	5.6%	0.9%
2025 post-sample	Import dependence	104	1993	19.2	4.8%	0.1%
2025 post-sample	Criticality	104	1993	19.2	38.5%	10.8%
2025 post-sample	Vulnerability	104	1993	19.2	19.2%	4.8%
2025 post-sample	Geopolitical risk	104	1993	19.2	25.0%	6.5%
2025 post-sample	Mitigation / de-risking response	104	1993	19.2	21.2%	4.4%
2025 post-sample	Strategic vulnerability narrative	104	1993	19.2	18.3%	4.2%

Notes: The table reports release counts, passage counts, average passages per release, document-level extensive margins, and mention intensities for all six narratives and all period bins, from the canonical GPT-5.4-mini run.

Table 11: Event-Time Textual Evidence Around the 2018 Section 301 Tariffs

	Extensive margin				Mention intensity			
Period	Import dependence	Vulnerability	Geopolitical risk	Strategic vulnerability narrative	Import dependence	Vulnerability	Geopolitical risk	Strategic vulnerability narrative
2017 baseline	7.3	3.7	3.7	2.8	0.6	0.7	0.8	0.5
Jan. 1–Mar. 21, 2018	9.1	9.1	18.2	13.6	1.3	0.7	1.2	0.9
Mar. 22–Jul. 5, 2018	19.4	22.6	32.3	22.6	3.1	5.3	11.8	5.3
Jul. 6–Dec. 31, 2018	7.1	12.5	26.8	12.5	1.0	1.4	7.0	1.4
2019 follow-through	6.5	6.5	12.0	2.8	1.0	1.4	3.2	0.8

Notes: Entries are percentages. The extensive margin is the share of USTR communications with at least one positive local passage for a given narrative. Mention intensity is the within-document share of local passages classified as positive. The event bins contain 109, 22, 31, 56, and 108 releases, respectively.

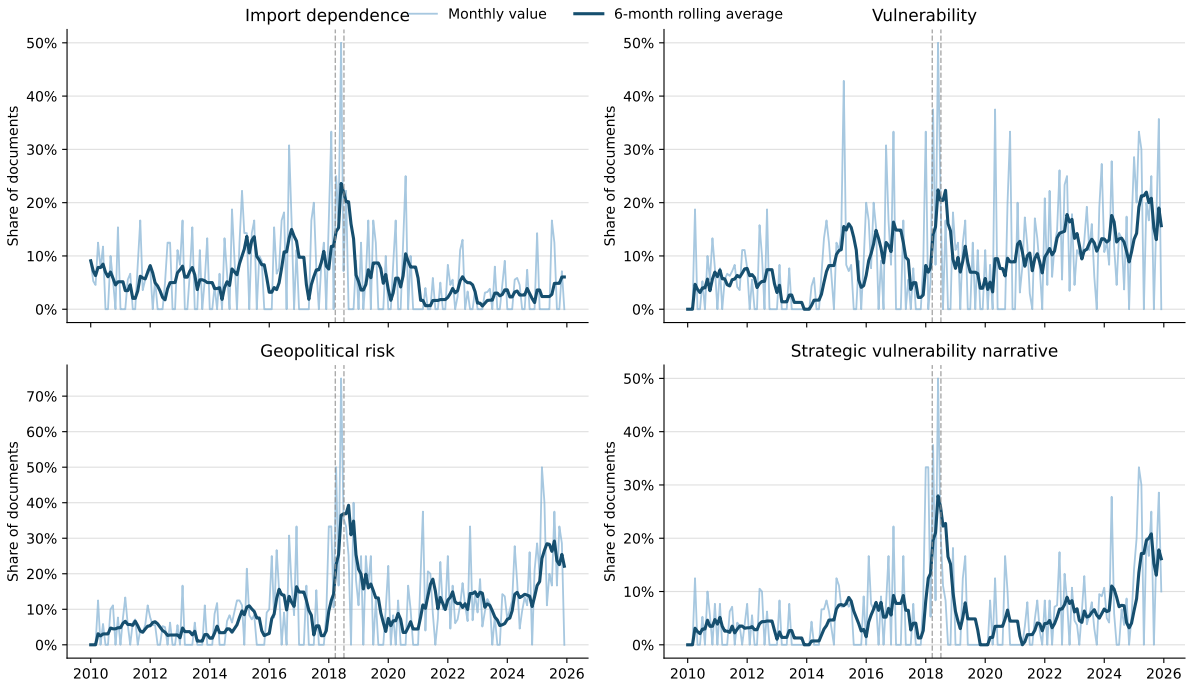


Figure 11: Full-Period Extensive Margin of Strategic-Risk Language in USTR Communications

Notes: The figure reports the monthly share of USTR documents containing at least one positive local passage in each of the four main narratives (import dependence, vulnerability, geopolitical risk, and the strategic vulnerability narrative), together with a six-month rolling average. Vertical reference lines mark the March 22, 2018 tariff announcement and the July 6, 2018 initial implementation date.

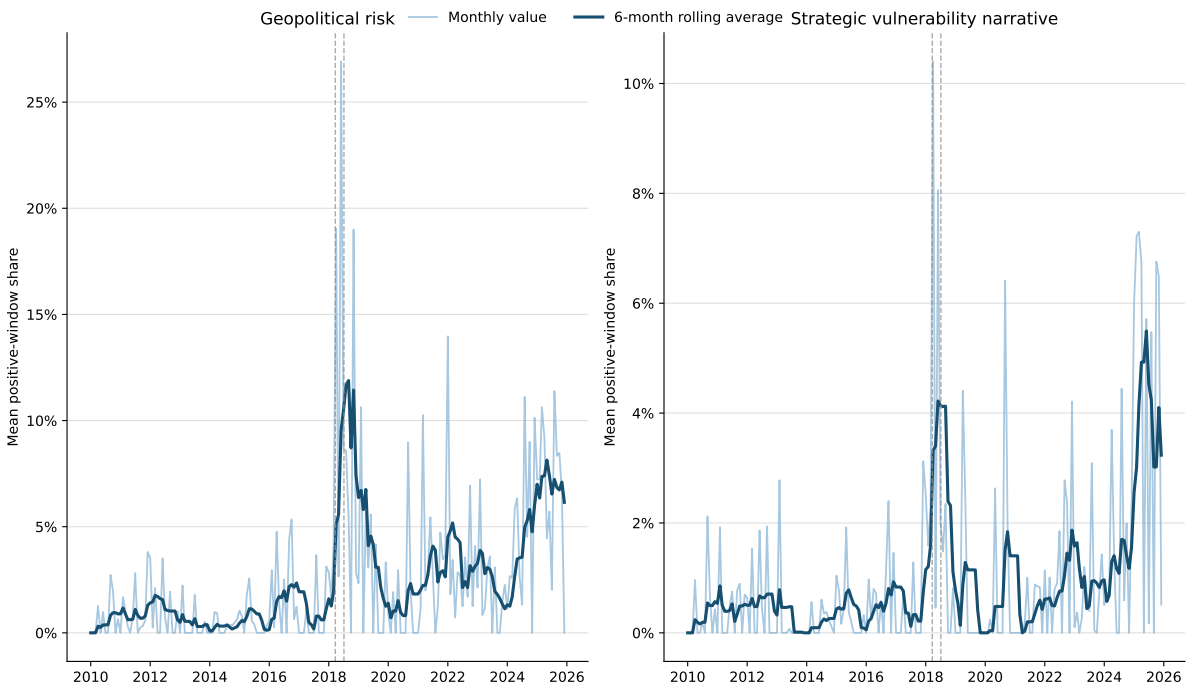


Figure 12: Full-Period Mention Intensity of Mechanism-Specific Strategic-Risk Language

Notes: The figure reports the monthly document-weighted mean of the within-document positive-passage share for geopolitical risk and the strategic vulnerability narrative, together with a six-month rolling average. Vertical reference lines mark the March 22, 2018 tariff announcement and the July 6, 2018 initial implementation date.

estimate of narrative prevalence. The 2018 shift survives it intact, with geopolitical-risk releases rising from 4.5 percent before 2017 to 18.3 percent in 2018 and releases invoking the strategic vulnerability narrative from 2.9 to 10.1 percent.

Panel B confronts the alternative that the 2018 increase reflects a change in the composition of USTR output rather than a change in framing, for instance fewer administrative scheduling notices in a busy year. Excluding administrative schedules leaves the 2018 margins unchanged at 26.6 and 15.6 percent. Restricting to policy-action communications, the release type in which a tariff rationale would have to appear, raises the 2018 margins to 57.1 and 33.3 percent: a majority of 2018 policy-action releases invoked geopolitical risk. Excluding administrative, meeting, and dialogue releases together yields 35.4 and 21.5 percent. Under every composition, the 2018 margins are several times their pre-2017 counterparts within the same restricted sample, so composition shifts do not generate the result.

Table 12: Robustness of Mechanism-Specific Textual Evidence

Sample or classification rule	Narrative	Pre-2017	2018
Panel A. High-confidence positive classifications			
Probability threshold of 0.90	Geopolitical risk	4.5	18.3
Probability threshold of 0.90	Strategic vulnerability narrative	2.9	10.1
Panel B. Release-composition restrictions			
All USTR communications	Geopolitical risk	6.0	26.6
All USTR communications	Strategic vulnerability narrative	3.9	15.6
Excluding administrative schedules	Geopolitical risk	7.2	26.6
Excluding administrative schedules	Strategic vulnerability narrative	5.2	15.6
Policy-action communications only	Geopolitical risk	17.6	57.1
Policy-action communications only	Strategic vulnerability narrative	13.3	33.3
Excluding administrative, meeting, and dialogue releases	Geopolitical risk	10.0	35.4
Excluding administrative, meeting, and dialogue releases	Strategic vulnerability narrative	7.0	21.5

Notes: Entries report document-level extensive margins in percent. The high-confidence screen retains positive passage classifications only when classifier confidence is at least 0.90 and does not reinterpret uncertain negatives, so Panel A is a conservative positive-evidence bound. Panel B shows that the increase is not driven by short administrative schedules or meeting announcements.

K.6 Stability Across Classifiers

The canonical results use GPT-5.4-mini. One concern is that the magnitude or timing of the 2018 contrast is specific to this baseline classifier. Figure 13 addresses this concern by repeating the analysis with GPT-5.4-mini, DeepSeek-v4-pro, GPT-OSS-120b, and GLM-5.2.

Absolute prevalence differs considerably across classifiers. From the pre-2017 period to 2018, the geopolitical-risk extensive margin rises by 20.3 to 23.3 percentage points across the four classifiers, while the corresponding increase for the strategic vulnerability narrative ranges from 4.5 to 11.7 percentage points. Each classifier also places its highest 2018 prevalence for both narratives between the March announcement and July implementation. GLM-5.2 lies toward the conservative end of the level estimates, yet records increases of 20.3 and 7.3 percentage points, respectively. Classifier choice therefore changes estimated levels and the magnitude of the 2018 shift, but not its direction or timing.

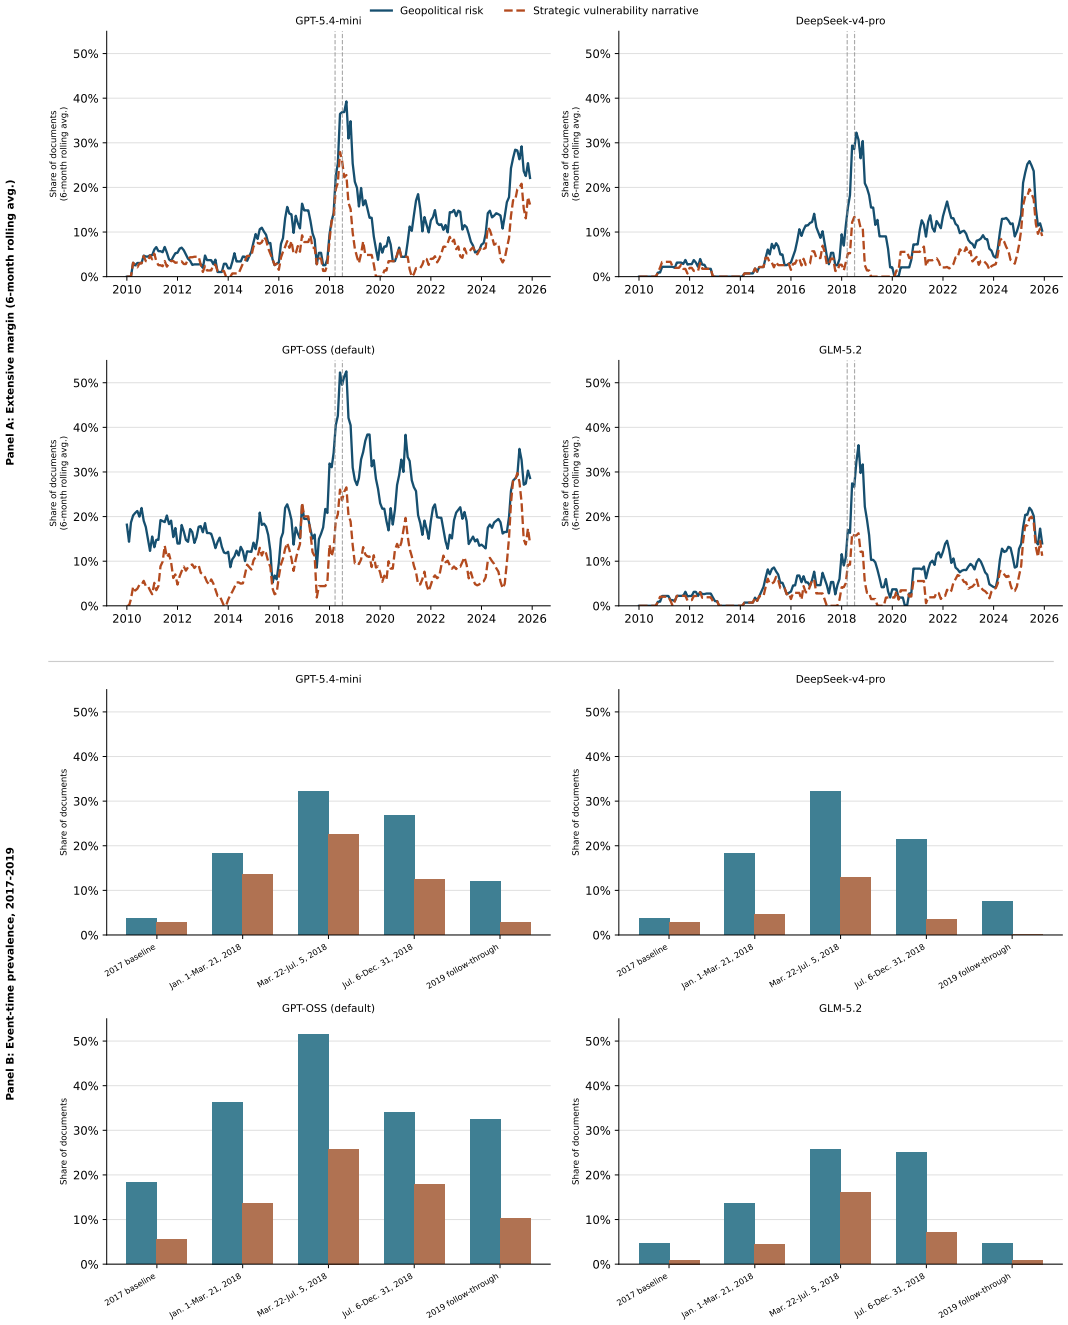


Figure 13: Stability of the Mechanism-Specific Narratives Across Classifiers

Notes: Panel A reports six-month rolling averages of the monthly share of USTR releases containing at least one positive passage for geopolitical risk or the strategic vulnerability narrative. Vertical lines mark the March 2018 tariff announcement and the July 2018 initial implementation date. Panel B reports the corresponding document-level extensive margins for the 2017 baseline, three subperiods within 2018, and the 2019 follow-through period. GPT-5.4-mini, DeepSeek-v4-pro, GPT-OSS-120b, and GLM-5.2 apply the same six-narrative ontology to the same 55,003 passages from 2,983 releases. All outcomes are binary at the document level, and confidence scores are neither pooled nor compared across classifiers. Axis scales are held fixed across classifiers within each panel.

Temperature sensitivity within GPT-OSS. Model choice and decoding temperature are distinct sources of sensitivity. Figure 14 isolates the latter by holding the GPT-OSS-120b model, prompt, input passages, and document-level aggregation fixed while comparing the server-default setting with an explicit temperature of 0.2. In 2018, the temperature-0.2 run yields lower annual prevalence: the geopolitical-risk extensive margin falls from 39.4 to 33.9 percent, and the strategic vulnerability narrative falls from 19.3 to 11.9 percent. The timing, however, remains stable. Both settings place the 2018 peak between the March announcement and July implementation; in this interval, geopolitical-risk language appears in 51.6 and 48.4 percent of documents under the default and 0.2 settings, respectively, while the strategic vulnerability narrative appears in 25.8 percent under both settings. The temperature specification thus affects measured prevalence, especially for the strategic vulnerability narrative, but not the concentration of the shift during the tariff rollout.

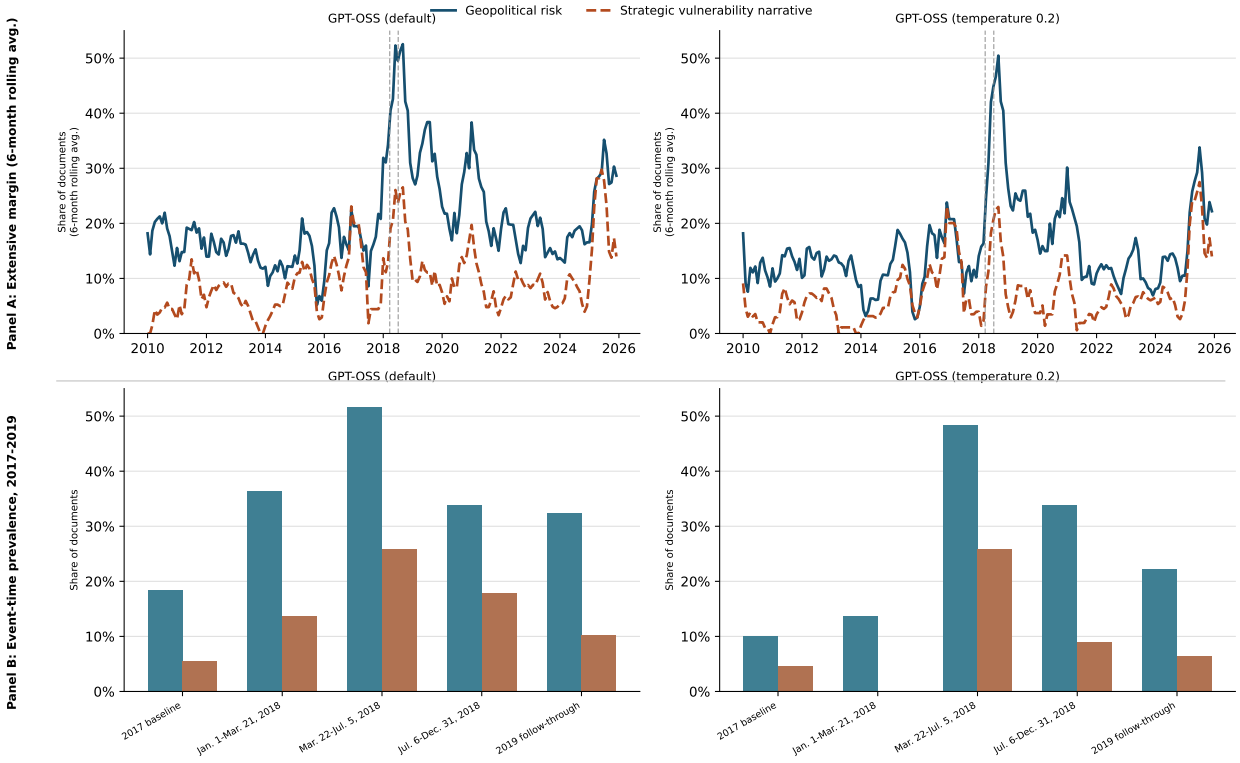


Figure 14: Sensitivity of GPT-OSS Classifications to Decoding Temperature

Notes: The figure repeats the binary document-level extensive-margin analysis in Figure 13 using GPT-OSS-120b only. The left column uses the server-default temperature, while the right column fixes temperature at 0.2. The model, high-reasoning setting, prompt, six-narrative ontology, 55,003 input passages, and document-level aggregation are otherwise held fixed. Axis scales are held fixed across temperature settings within each panel.

K.7 Validation Status and Audit Protocol

The validation evidence currently in place consists of four layers: hand inspection of selected positives and hard negatives against the ontology (Table 8), the high-confidence positive screen, the release-composition restrictions (Table 12), and replication across classifiers and, within GPT-OSS, across decoding-temperature specifications (Appendix K.6). These checks support the interpretation that the classifier captures the intended public framing: positives combine foreign exposure with state-directed risk, and negatives exclude country mentions, administrative tariff references, and ordinary market-access language that do not articulate the mechanism.

One planned validation layer is outstanding, and we flag it rather than bury it: a blinded human-coded audit with quantitative agreement statistics has been designed but not yet completed. The design stratifies the audit sample by period, release type, classifier-positive and classifier-negative status, and narrative, with deliberate oversampling of geopolitical risk, the strategic vulnerability narrative, and the 2018 rollout window; independent coders classify passages under the ontology of Table 7 without access to model labels; and the audit report will present pre-adjudication inter-coder agreement, adjudicated labels, and narrative-specific precision and recall for the model classifications. Until those statistics are in hand, the textual results should be read as a description of official public framing, measured by the conservative baseline classifier and cross-checked against alternative classifiers and the within-GPT-OSS temperature comparison, rather than as a human-adjudicated gold standard.

K.8 Limitations and Interpretation

Four limitations bound what the textual evidence can establish. First, the corpus is one agency's public record. USTR is the agency that announces and administers the tariff instruments, which makes its communications the natural place to look for the policy's stated rationale, but statements by other parts of the government are outside the sample. Second, the narrative assignments are produced by language models, and estimated levels vary across classifiers. The analysis therefore rests on within-classifier contrasts across periods. Four distinct models reproduce the direction and timing of those contrasts, and the within-GPT-OSS comparison shows that the pattern persists under both temperature specifications. Third, the period contrasts are descriptive. The exact tests in Appendix K.4 quantify sampling uncertainty under an independence assumption, and the disjoint-window replication limits the scope for episode-level clustering to overturn them, but no causal claim about what produced the language is available or

needed. Fourth, public framing is a necessary rather than sufficient implication of the mechanism: rhetoric chosen for political packaging could resemble it. That alternative, however, does not predict the cross-sector structure of the tariffs themselves, while the model's strategic-risk component predicts the targeting gradient and the framing shift jointly, with the observed timing.

Within these boundaries, the textual record supports a specific and falsifiable reading of the 2018 episode. Official communications shifted toward describing foreign economic exposure as a strategic liability. This shift was concentrated in the two narratives that express the model's mechanism rather than in generic trade language, and it peaked between announcement and implementation. The pattern survives the high-confidence positive screen, the release-composition restrictions, replication across four classifiers, and the within-GPT-OSS temperature comparison; it also re-emerged out of sample when tariff escalation resumed in 2025. Had USTR's language not moved, or moved only in generic dependence and criticality narratives, the strategic-risk interpretation of the tariff regressions would have been strained. Taken together, these patterns provide corroboration rather than identification.