Tax Incidence with Endogenous Exit

Jackson Mejia*

February 2026

Measured tax incidence is biased when agents differentially select out of the observed tax base. When low-passthrough agents are more likely to exit, measured incidence overstates the true domestic burden. I show that correcting this bias requires only the exit share. Applying this framework to tariffs, where rerouting generates detectable exit, I find that a ten percentage point tariff rate increase raises exit probability by 10-30 percentage points. Correcting for selection implies that the incidence of tariffs on the U.S. during the 2018-19 trade war was 85% as large as commonly supposed, and 75% as large for intermediate business inputs. Existing tax provisions can discourage exit—corporate tax deductibility reduces the return to avoidance—but a major reform in 2017 weakened this channel, reducing tariff revenue by $2-9 billion.

JEL-Classification: H22, H26, F14

Keywords: tax incidence, sufficient statistics, tax avoidance, corporate taxation, tariffs

*Massachusetts Institute of Technology, jpmejia@mit.edu. First Draft: September 2025. I am especially grateful to Kim Ruhl for extensive discussions and comments on the development of these ideas, as well as to Maxx Cook, Andrew Johnston, Pedro Martinez-Bruera, Steve Miran, and Jim Poterba for useful conversations and comments. Refine.ink was used to proofread this paper for consistency and clarity. The views and conclusions contained in this document are mine and should not be interpreted as representing the official policies, either expressed or implied, of any employer or affiliated organization.

1 Introduction

Chetty (2009) establishes that selection into tax avoidance biases the elasticity of taxable income, the key sufficient statistic for measuring deadweight loss. This paper derives the analogous result for tax incidence: when agents can exit the observed tax base, measured incidence overstates the burden on those who remain. The logic is simple. Agents who would bear the highest incidence—those who cannot pass the tax forward—have the strongest incentive to pay the cost of exit. When they leave, only high-passthrough "survivors" remain in the data. Standard estimates therefore average over a selected sample, mechanically overstating incidence.

Correcting this bias requires only measuring how much of the tax base exits, a sufficient statistic I denote θ. Under monotone selection, where low-passthrough agents are more likely to exit, θ bounds true incidence from below without requiring identification of which agents exit or why. The intuition is simple: if the agents who exit are those who would have absorbed the most tax, then averaging over survivors overstates incidence, and the share that exits tells us by how much. The correction does not depend on the structural parameters generating selection: not the demand system that determines passthrough heterogeneity, not the fixed costs that govern exit decisions, not the distribution of firm productivity. It requires only that low-passthrough agents are more likely to leave, and a measure of how much of the tax base they take with them. The same correction applies wherever differential exit is possible, and that portability is this paper's core theoretical contribution.

The central empirical contribution is an application to tariffs, where exit takes the form of rerouting goods through third countries to avoid duties. Tariffs offer an ideal setting for three reasons. First, rerouting leaves a detectable footprint in bilateral trade data. Second, the 2018 U.S. tariffs on China provide sharp policy variation, with rates rising from near zero to around 20 percent on thousands of products. Third, existing estimates present a puzzle that selection can resolve. A large body of work using customs data finds near-complete passthrough, implying U.S. importers bore the full incidence of the 2018 tariffs—a burden which makes its welfare cost comparable to top income taxes (Finkelstein and Hendren 2020).[1] This is striking: for an economy as large as the United States, we would expect foreign exporters to absorb some of the burden. But the customs data typically used to study passthrough conditions on firms and products that continue shipping directly; exporters who reroute exit the sample. When goods are tracked from true origin to final destination regardless of routing, the picture changes. Evidence from washing machines suggests that eliminating low-cost avoidance options can substantially raise measured passthrough (Flaaen, Hortaçsu, and Tintelnot 2020). My framework reconciles these findings: customs-based estimates correctly measure passthrough among survivors, but

1. See, for example, Amiti, Redding, and Weinstein (2019, 2020), Cavallo et al. (2021), and Fajgelbaum et al. (2020).

survivors are selected. Given that markups are heterogeneous in practice (Gopinath and Itskhoki 2010; De Loecker, Eeckhout, and Unger 2020; Edmond, Midrigan, and Xu 2023), the monotone selection condition is plausible, and measuring exit provides the correction.

I measure exit by extending the trade footprint methodology of Freund (2024), which screens bilateral trade data for patterns consistent with rerouting. The approach yields lower and upper bounds on the share of trade that exits the direct channel. By 2023, 30–40 percent of tariffed products exhibit evidence of rerouting, with aggregate exit reaching 15 percent of baseline China-U.S. trade. These estimates of θ imply true passthrough falls from near 100 percent to approximately 85 percent in aggregate. Since exit is concentrated in business inputs, the correction is largest there—approximately 80 percent for intermediate and capital goods, near 100 percent for consumer goods. I extend the methodology to estimate the elasticity of exit with respect to tariff incentives, the key input for welfare analysis. The behavioral response is large: a 10 percentage point increase in the tariff raises the probability of exit by 10–30 percentage points on the extensive margin and increases the exiting share by 2–11 percentage points on the intensive margin. Identification comes from within-product variation in tariff exposure; event studies show flat pre-trends, and placebo tests using European and Canadian destinations show no response to the U.S. tariff wedge.

These estimates allow me to re-evaluate the welfare cost of tariffs using the marginal cost of public funds (MCPF), which measures the welfare cost to society of raising one dollar of government revenue. The MCPF combines domestic incidence in the numerator (who bears the burden) with the fiscal externality in the denominator (how much the tax base shrinks when rates rise). For a lump-sum tax that induces no behavioral distortions, the MCPF equals one. Existing estimates place the MCPF of the 2018 tariffs between 1.2 and 1.6 (Finkelstein and Hendren 2020; Jaccard 2021), comparable to top income taxes. My correction affects both components: the numerator falls because true incidence is 85 percent, not 100 percent; the denominator rises because exit erodes the tax base. These forces partially offset, with sharp heterogeneity by end-use. For consumer goods, where exit is minimal, the MCPF remains near one throughout the post-tariff period. For intermediate and capital goods, the high exit elasticity generates a substantial fiscal externality, pushing the upper bound to around 1.5 by 2023, though the lower bound remains near one.

The welfare cost depends on how much exit occurs, which in turn depends on policy. Alessandria et al. (2025) show that one lever is revenue use: deploying tariff revenue to offset distortionary taxes lowers the net welfare cost. I identify a different lever: the domestic tax code affects the incentive to exit in the first place. When importers can deduct tariff-inclusive costs from taxable income, their effective burden falls. This compresses the wedge between compliant and non-compliant channels, discouraging exit. Higher corporate tax rates therefore serve as implicit enforcement, reducing avoidance even without explicit penalties. The 2017 Tax Cuts and Jobs Act cut the corporate rate from 35 to 21 percent, widening the wedge and increasing the incentive to reroute. I estimate this reform cost the U.S. between 2 and 9 billion dollars in foregone tariff revenue.

The sufficient statistics approach is deliberately narrow and partial equilibrium in scope. Transshipment is one channel of tariff avoidance among many; firms also misreport value to fall under de minimis thresholds (Fajgelbaum and Khandelwal 2024), break up shipments into smaller parcels, or misclassify goods (Fisman and Wei 2004). My estimates are therefore conservative lower bounds on total exit, and the correction applies regardless of whether exit reflects physical rerouting, supply chain reorganization, or shifts in sourcing. The welfare analysis is partial equilibrium, abstracting from general equilibrium effects on wages, prices, and terms of trade. The 85 percent passthrough reflects import prices paid by U.S. firms, but ultimate consumer incidence depends on whether these costs are absorbed by domestic firms or passed through to retailers (Amiti, Itskhoki, and Konings 2019; Flaaen et al. 2025). These limitations suggest the true correction could be larger, and the welfare implications more nuanced, than my baseline estimates indicate.

Related Literature. This paper relates to three strands of literature.

First, this paper contributes to the public finance literature on how tax avoidance biases measured behavioral parameters. Chetty (2009) established that selection into avoidance biases the elasticity of taxable income, requiring a correction for welfare analysis. I derive the analogous result for tax incidence: when agents can exit the observed tax base, measured passthrough overstates domestic burden, and the exit share provides a sufficient statistic for correction. The closest empirical predecessor is Kopczuk et al. (2016), who show that moving diesel tax collection upstream raises both passthrough and revenue because evasion capacity differs across agents. Their finding demonstrates that statutory incidence can matter when avoidance opportunities are heterogeneous, but they do not derive a general correction formula. Benzarti and Carloni (2019) document that VAT incidence varies substantially across firms, with owners capturing most of the benefit from rate cuts—further evidence that passthrough heterogeneity is empirically important. My sufficient statistic approach provides the correction these papers lack. The deductibility channel I identify extends the enforcement literature (Allingham and Sandmo 1972; Slemrod and Yitzhaki 2002; Slemrod 2019) by showing that the tax code contains implicit enforcement margins operating automatically through deductibility, without explicit audits or penalties.

Second, this paper builds on a methodology for inferring illicit trade from official statistics. Fisman and Wei (2004) introduced the trade discrepancy approach, comparing exporter and importer reports to quantify evasion; Javorcik and Narciso (2008) extended it to show evasion is concentrated in differentiated products where customs verification is difficult. Fisman, Moustakerski, and Wei (2008) provided early intuition that indirect trade via Hong Kong rises with Chinese tariff rates due to costly evasion; I generalize their framework to allow heterogeneous markups and selection bias in measured passthrough. My empirical approach extends the screening methodology of Freund (2024), who flag products showing trade patterns consistent with rerouting. Recent work confirms that rerouting is empirically prevalent: Do et al. (2025) use bill of lading data to show that 70 percent of U.S. containerized imports are transshipped through hubs; Deng et al. (2025) estimate that $5.5 billion of the 2018 surge in U.S. imports from Mexico was circumvention; Iyoha et al. (2025) find that rerouting accounts for nearly 9 percent of Vietnam's export growth to the United States. Alfaro and Chor (2025) distinguish transshipment from genuine supply chain reorganization, noting that China's value added in third-country exports remains substantial. My sufficient statistics correction applies regardless of mechanism—whether exit reflects physical rerouting, supply chain integration, or sourcing shifts, the bias in measured incidence is the same.[2]

Third, this paper reconciles findings from the tariff incidence literature. Studies using bilateral customs data classified by country of origin find near-complete passthrough of the 2018 tariffs (Amiti, Redding, and Weinstein 2019, 2020; Fajgelbaum et al. 2020). These estimates correctly measure passthrough among firms that continue shipping directly, but when firms reroute, they exit the sample. Cavallo et al. (2021) use BLS microdata that tracks identical goods over time and also find near-complete passthrough at the border; their sample is weighted toward consumer goods, where I find essentially no transshipment, so their estimates are consistent with my framework. Similarly, Flaaen, Hortaçsu, and Tintelnot (2020) study washing machines—another consumer good—and find that country-specific tariffs allowed firms to relocate, muting price effects, but a global tariff that closed this option caused passthrough to exceed 100 percent. My contribution is to formalize the selection mechanism, show that the bias is concentrated in business inputs where transshipment is prevalent, and derive the sufficient statistic correction for welfare analysis.

Roadmap. Section 2 develops a partial equilibrium model of avoidance and passthrough. Section 3 describes the data and methodology used to measure transshipment, which Section 4 uses to present the main empirical results, including the magnitude of avoidance and the elasticity estimates. Section 5 uses these estimates to re-evaluate the tariff's welfare cost and conduct the TCJA counterfactual. Section 6 concludes.

2. Mishra, Subramanian, and Topalova (2008) study tariff evasion in India and find a substantially smaller elasticity than my estimates, likely reflecting differences in enforcement capacity and the type of evasion.

2 A Model of Endogenous Exit

This section presents the economic logic behind the paper's main results: (i) when exit from the observed tax base is costly, low-passthrough agents are more likely to exit, (ii) this selection biases measured incidence upward, and (iii) the exit share is a sufficient statistic for correcting the bias. Section 2.1 sets up the environment and Section 2.2 characterizes pricing and passthrough heterogeneity. Section 2.3 introduces the exit margin and derives the selection result. Section 2.4 shows how selection biases measured incidence and derives the sufficient statistic correction. Section 2.5 maps the framework to welfare analysis. Finally, Section 2.6 shows how other tax instruments affect the return to exit.

2.1 Environment

Consider a static partial equilibrium setting with monopolistic competition over differentiated varieties indexed by $i \in \mathcal{I}$. Each seller sets a price p_i and faces an ad valorem tax $\tau^d \geq 0$. The gross-of-tax price is $\tilde{p}_i = (1 + \tau^d)p_i$.

Buyers. Buyers submit demand for each variety based on its gross price $\tilde{p}_i$ and a price aggregator $\mathbf{P}$. Formally, the demand for variety i takes the form

$$q_i = D_i(\tilde{p}_i \mid \mathbf{P}), \quad \mathbf{P} \equiv \mathbf{P}(\mathbf{P}_{-i}; \Theta).$$

$\mathbf{P}$ is exogenous at the firm level, corresponding to a standard monopolistic competition setup in which each firm is small relative to the market and takes the price distribution of rivals as given. Denoting the elasticity of demand as $\varepsilon_i > 1$, I assume that demand is weakly Kimball (1995) with superelasticity

$$\kappa_i(\tilde{p}_i \mid \mathbf{P}) \equiv -\frac{\partial \ln \varepsilon_i(\tilde{p}_i \mid \mathbf{P})}{\partial \ln \tilde{p}_i} \leq 0.$$

When $\kappa < 0$, elasticity rises with price. Under CES ($\kappa = 0$), elasticity is constant. Additional regularity conditions are imposed in Appendix A.1.

Sellers. Each firm $i \in \mathcal{I}$ has constant marginal cost m_i, drawn independently from a continuous distribution G with support $[\underline{m}, \bar{m}] \subset (0, \infty)$ and density $g(m) > 0$. The distribution is independent of policy and of the aggregator $\mathbf{P}$. Each firm solves

$$\pi_i(m_i) = \max_{p_i \geq m_i} (p_i - m_i) D_i((1 + \tau^d)p_i \mid \mathbf{P}).$$

The framework applies to any ad valorem tax where agents can exit the observed base. I develop the model in general terms but use tariffs as the running example throughout: sellers are foreign exporters, buyers are domestic importers, and exit takes the form of rerouting through third countries.

2.2 Pricing and Passthrough

Before analyzing behavioral responses to taxation, I first characterize the firm's pricing decision and how tax passthrough varies with marginal cost. These results establish the heterogeneity that drives selection when exit becomes possible.

The firm's first-order condition yields a standard Lerner condition: the optimal markup is inversely proportional to the demand elasticity. Appendix A.2 verifies that the profit function is strictly concave and establishes the following comparative statics: firms with higher marginal costs charge higher prices, sell lower quantities, earn lower profits, and have lower markups.

The key object for incidence is passthrough: how much of a tax increase is reflected in the gross price.

Proposition 1 (Passthrough heterogeneity)**.** *Holding* **P** *fixed, the passthrough of the tax onto the gross price of variety* i *is*

$$\beta_i \equiv \frac{d \ln \tilde{p}_i}{d\tau^d} = \frac{1}{1+\tau^d} \cdot \frac{\varepsilon(\tilde{p}_i) - 1}{\varepsilon(\tilde{p}_i) - 1 - \kappa(\tilde{p}_i)}, \tag{1}$$

where $\tilde{p}_i = (1+\tau^d)p_i^*(m_i)$*. When* $\kappa = 0$ *(CES), passthrough is homogeneous across firms. When* $\kappa < 0$ *(Kimball), passthrough is monotonically increasing in marginal cost:* $\partial\beta_i/\partial m_i > 0$*.*

Proof. See Appendix A.4. □

Equation (1) decomposes passthrough into a mechanical factor $1/(1+\tau^d)$ and a curvature factor $(\varepsilon - 1)/(\varepsilon - 1 - \kappa)$.[3] The intuition for heterogeneous passthrough under Kimball demand is as follows. Low-cost firms charge high prices and earn high markups. At these high prices, they face elastic demand (since $\kappa < 0$ implies elasticity rises with price). When taxes increase, maintaining their price would push them further up the demand curve where elasticity is even higher. To avoid losing market share, they must absorb part of the tax by compressing their markup, resulting in low passthrough. In the tariff context, these are the most productive foreign exporters—they absorb the largest share of a tariff increase. High-cost firms, by contrast, charge

3. Proposition 1 characterizes firm-level passthrough taking the price index **P** as given, standard in monopolistic competition where individual firms are atomistic. Industry-level comparative statics, where aggregate tariff changes affect **P**, involve additional equilibrium effects: higher tariffs raise the price index, which shifts demand toward all varieties and puts upward pressure on markups.

lower prices where demand is less elastic. They can pass through most of the tax without losing as much volume. These are the marginal exporters, who pass most of the tariff on to domestic buyers.

Under CES demand ($\kappa = 0$), elasticity is constant regardless of price, so all firms have identical passthrough $\beta = 1/(1 + \tau^d)$ and no heterogeneity arises. Variable markups and heterogeneous passthrough arise naturally under non-CES demand, as in Gopinath and Itskhoki (2010) and Melitz and Ottaviano (2008). The Kimball specification provides a tractable parameterization that nests CES as a special case.

To summarize: low-cost firms have high markups, high profits, and low passthrough; high-cost firms have low markups, low profits, and high passthrough. This heterogeneity is central to what follows. In the next section, I introduce the possibility that firms can exit the observed tax base—and show that it is precisely the low-passthrough firms who choose to do so.

2.3 Selection into Exit

I now introduce the possibility that firms can exit the observed tax base. Exiting requires paying a fixed cost $F > 0$ but allows the firm to face a lower effective tax rate $\tau^a \in [0, \tau^d)$. The fixed cost captures setup expenses for exit. With tariffs, that may include establishing transshipment routes, contracting with intermediaries, or relabeling goods. While F likely varies across products (a possibility explored in Appendix A.5), it is best interpreted as a sunk entry cost. This interpretation implies hysteresis: once the cost is paid and the network is built, firms may not revert to direct shipping even if tariffs are later removed. Furthermore, F may decline endogenously over time as firms learn optimal routing strategies and intermediaries specialize in avoidance services, suggesting that long-run exit elasticities could exceed short-run responses.

The parameter τ^a reflects the expected cost of operating outside the observed base, whether through enforcement risk, explicit taxes in the exit channel, or other frictions. The key object is the wedge $\tau^d - \tau^a$: a larger wedge makes exit more attractive.

Given the results of Section 2.2, we can characterize selection by comparing optimized profits across channels. For any marginal cost m, define the optimized value in a channel with gross-price multiplier $1 + \tau^r$:

$$V(m, 1 + \tau^r) \equiv \max_{p \geq m} (p - m) D((1 + \tau^r)p), \quad \pi^d(m) = V(m, 1 + \tau^d), \quad \pi^a(m) = V(m, 1 + \tau^a) - F,$$

and the exit-vs-stay difference

$$H(m; \tau^d, \tau^a) \equiv \pi^a(m) - \pi^d(m) = V(m, 1 + \tau^a) - V(m, 1 + \tau^d) - F.$$

Proposition 2 (Selection into exit). *Under D1–D4 (see Appendix A.1), there exists a unique cutoff* $\widehat{m}(\tau^d, \tau^a) \in [\underline{m}, \bar{m}]$ *such that*

$$m < \widehat{m}(\tau^d, \tau^a) \Rightarrow \textit{Exit}, \quad m \geq \widehat{m}(\tau^d, \tau^a) \Rightarrow \textit{Stay}.$$

If the cutoff is interior, then

$$\frac{\partial \widehat{m}}{\partial \tau^d} > 0 \quad \textit{and} \quad \frac{\partial \widehat{m}}{\partial \tau^a} < 0.$$

Proof. See Appendix A.3. □

The intuition is straightforward. From Section 2.2, low-cost firms have high markups, high profits, and low passthrough. When the tax increases, these firms must absorb most of the increase themselves, cutting directly into their profit margins. They therefore have the strongest incentive to pay F and exit. High-cost firms, by contrast, pass through most of the tax to buyers. The tax hits their profits less severely, so the gain from exit does not justify the fixed cost.

Let $\theta(\tau^d, \tau^a) \equiv \Pr(m < \widehat{m}(\tau^d, \tau^a))$ denote the share of firms that exit; the remaining $1 - \theta$ are *survivors*. By Proposition 2, $\partial\theta/\partial\tau^d > 0$ and $\partial\theta/\partial\tau^a < 0$: widening the wedge increases exit. The magnitude of this response is captured by the exit elasticity:

$$\psi \equiv \frac{\partial \ln \theta}{\partial \Delta} > 0, \quad \Delta \equiv \ln \frac{1 + \tau^d}{1 + \tau^a}. \tag{2}$$

The exit elasticity ψ depends on demand curvature. When κ is more negative, passthrough heterogeneity is greater, so the profit differential between low-cost and high-cost firms under taxation is larger. This amplifies the response to wedge changes: more curvature means more selection.

Figure 1 illustrates the mechanism. The distribution $g(m)$ shows marginal costs; the upward-sloping line shows passthrough $\beta(m)$ under Kimball demand. The cutoff $\widehat{m}$ partitions the distribution: low-cost firms (shaded) exit, while high-cost firms stay. A tax increase shifts the cutoff rightward, causing additional low-passthrough firms to exit.

In the tariff context, exit takes the form of rerouting through third countries. F captures setup costs such as establishing transshipment routes and contracting with intermediaries. τ^a reflects expected enforcement costs or tariffs in the hub country. If rerouting also involves higher per-unit costs—longer shipping routes, intermediary fees—this reinforces rather than undermines the selection story, since rerouted varieties exit the observed sample entirely. For income taxes, exit might mean offshore sheltering; for sales taxes, cross-border shopping. The sufficient statistic θ is common across settings.

Figure 1: Selection into Exit under Heterogeneous Passthrough

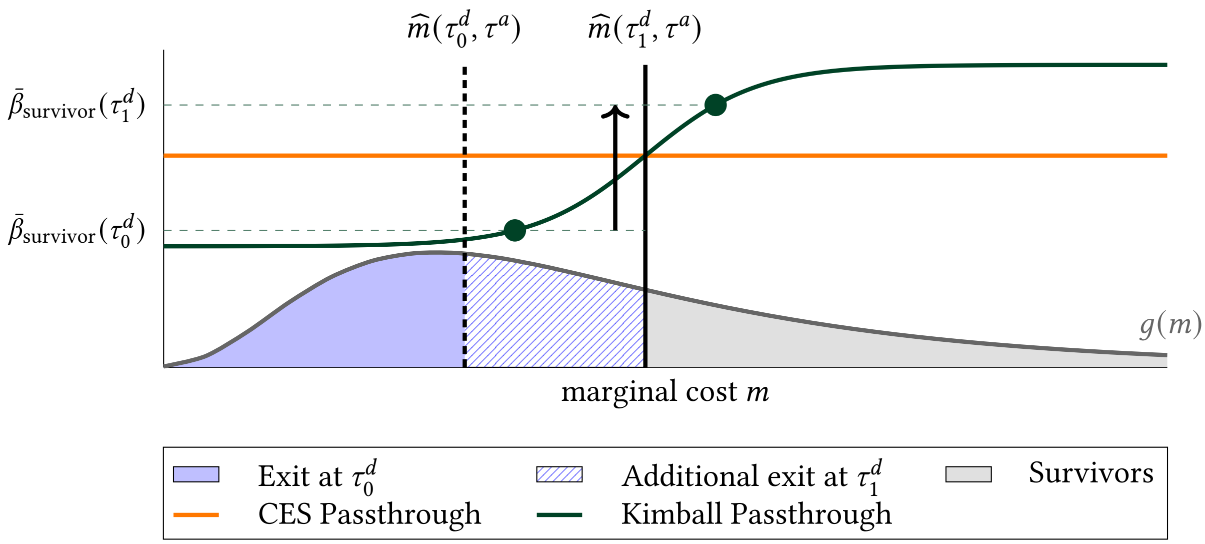

Note: The distribution $g(m)$ shows the density of marginal costs. Under CES (flat line), passthrough is constant. Under Kimball (upward-sloping line), passthrough increases with cost. The exit cutoff $\widehat{m}(\tau^d, \tau^a)$ partitions the distribution: low-cost firms (shaded area) exit, while high-cost firms stay in the observed base. A tax increase shifts the cutoff rightward (dashed to solid line), causing additional low-passthrough firms to exit.

2.4 Selection Bias in Measured Incidence

Propositions 1 and 2 establish that (i) low-cost firms have low passthrough and (ii) low-cost firms are more likely to exit. Combined, these results imply that incidence estimates based on observed transactions are biased upward. This section formalizes the bias and derives a sufficient statistic correction.

The central issue is that standard data sources capture only agents who remain in the observed tax base after the policy change. Agents who exit are unobserved and mechanically drop out of the estimating sample. This creates a gap between what the regression identifies and the cohort-level parameter of interest.

To formalize the bias, let $S_i \in \{0, 1\}$ indicate whether agent i remains in the observed base ($S_i = 1$) or exits ($S_i = 0$). Agents with $S_i = 1$ are *survivors*; those with $S_i = 0$ are *exiters*. Consider a regression of price changes on tax changes, run on observed transactions:

$$\Delta \ln \tilde{p}_i = \alpha + \beta \, \Delta\tau_i^d + \varepsilon_i. \tag{3}$$

Since $\Delta \ln \tilde{p}_i = \beta_i \Delta\tau^d$ for an agent who stays, the OLS coefficient identifies the survivor average:

$$\widehat{\beta}_{\text{survivor}} = \mathbb{E}[\beta_i \mid S_i = 1].$$

The object of interest is $\beta_{\text{cohort}} \equiv \mathbb{E}[\beta_i]$, the average passthrough for the pre-policy cohort. The following identity makes the selection bias transparent.

Proposition 3 (Survivor vs. cohort bias)**.** *Let* $\Pr(S_i = 1)$ *be the survivor share. Then*

$$\beta_{survivor} \equiv \mathbb{E}[\beta_i \mid S_i = 1] = \underbrace{\mathbb{E}[\beta_i]}_{\beta_{cohort}} + \frac{\text{Cov}(\beta_i, S_i)}{\Pr(S_i = 1)}. \tag{4}$$

Proof. $\mathbb{E}[\beta_i \mid S_i = 1] = \mathbb{E}[\beta_i S_i]/\mathbb{E}[S_i] = \big(\mathbb{E}[\beta_i]\mathbb{E}[S_i] + \text{Cov}(\beta_i, S_i)\big)/\Pr(S_i = 1)$. □

The selection term $\text{Cov}(\beta_i, S_i)/\Pr(S_i = 1)$ is positive whenever low-passthrough agents are more likely to exit. Two benchmarks clarify when bias arises. Under CES ($\kappa = 0$), all firms have identical passthrough $\beta_i = 1/(1 + \tau^d)$, so $\text{Cov}(\beta_i, S_i) = 0$ and $\beta_{\text{survivor}} = \beta_{\text{cohort}}$. Under Kimball demand ($\kappa < 0$), passthrough increases with cost, low-cost firms exit, and $\text{Cov}(\beta_i, S_i) > 0$ generates upward bias. Given heterogeneous markups in practice (De Loecker, Eeckhout, and Unger 2020), survivor-based estimates likely overstate the burden on domestic buyers.

A sufficient statistic correction. The magnitude of the bias can be bounded directly. By the law of total expectation,

$$\beta_{\text{cohort}} = (1 - \theta)\mathbb{E}[\beta_i \mid S_i = 1] + \theta\mathbb{E}[\beta_i \mid S_i = 0].$$

Using $\beta_{\text{survivor}} = \mathbb{E}[\beta_i \mid S_i = 1]$ and non-negative passthrough,

$$\beta_{\text{cohort}} \in \big[(1 - \theta)\beta_{\text{survivor}},\ \beta_{\text{survivor}}\big].$$

This bounds the bias using only the exit share θ.[4] Correcting measured incidence requires quantifying how much exit occurs, not the structural parameters generating selection. Importantly, the correction requires only θ, not identification of which agents exit or why. If low-passthrough agents differentially exit for any reason, measured incidence is biased upward, and θ corrects it.

This bound has a key methodological property: it requires only θ, not identification of the selection mechanism. The correction does not require observing which agents exit, estimating structural parameters such as the fixed cost F or demand curvature κ, or testing whether low-passthrough agents are more likely to exit. If monotone selection holds for any reason—whether through the mechanism formalized above or some alternative—the bound applies. This portability

4. This sufficient statistic parallels the share function in Fisman, Moustakerski, and Wei (2008), where the aggregate evasion share enters welfare calculations. Here, θ plays an analogous role but also determines the sample-selection bias in measured passthrough, linking evasion to incidence through the same parameter.

is a feature of the sufficient statistics approach: the same object corrects incidence bias in any setting where agents can differentially exit the observed tax base.

The theoretical θ is the mass of firms selecting into exit. Whether this is the relevant empirical object depends on how passthrough is estimated. If the regression weights observations equally, θ is the share of firms that exit. If the regression is value-weighted—as is standard in the tariff literature—then θ is the share of trade value that exits. The sufficient statistic correction applies in either case; what changes is the appropriate measure. In the empirical analysis, I measure θ as the share of trade value rerouted, matching the value-weighted regressions in the literature.

Application to tariffs. In the tariff literature, survivor-based estimates come from customs unit values or import price indices constructed from direct shipments. Studies using these data estimate β_{survivor} by construction, since firms that reroute through third countries exit the sample when tariffs induce origin switching. This describes the estimand in Amiti, Redding, and Weinstein (2019), Fajgelbaum et al. (2020), and the unit-value components of Cavallo et al. (2021).

When is there no selection bias? When firm-product transactions are tracked over time with the cohort fixed at the micro level, regardless of routing. This requires observing the same exporter-importer pair before and after the tariff whether it ships direct or reroutes. Parts of Cavallo et al. (2021) and Flaaen, Hortaçsu, and Tintelnot (2020) implement this approach. The findings from Flaaen, Hortaçsu, and Tintelnot (2020) illustrate the broader point that exit options shape measured passthrough. For earlier country-specific tariffs on washing machines, they document extensive country-hopping through relocation, resulting in minimal or even negative passthrough. For the 2018 global tariff, which eliminated the benefit of relocation, passthrough exceeded 100 percent. When low-cost exit is available, firms use it; when it is not, they absorb or pass through the tariff. The contrast is consistent with exit shaping measured incidence, though the specific mechanism in their setting—relocation under oligopolistic competition—differs from the transshipment margin I study here.

2.5 Sufficient Statistics for Welfare

The theoretical results clarify how selection into exit affects the welfare cost of taxation. The marginal cost of public funds (MCPF) combines domestic incidence in the numerator and the fiscal externality from behavioral responses in the denominator. I now show how the model's parameters map into this formula.

Numerator: Domestic incidence. Hendren (2016) shows that the welfare effect of a tax is the willingness to pay to avoid it. Applying the envelope theorem, Finkelstein and Hendren (2020) show this is simply domestic incidence, or passthrough. Proposition 3 shows that true cohort passthrough is bounded below by $(1-\theta)\beta_{\text{survivor}}$ and above by β_{survivor}. This provides a portable correction for the numerator of the MCPF.

Denominator: Fiscal externality. The total fiscal externality, η_{total}, is the elasticity of the tax base with respect to the tax rate. It has two components: the externality from standard behavioral responses such as substitution, η_{other}, and the externality from exit, η_{exit}. The total is the sum: $\eta_{\text{total}} = \eta_{\text{other}} + \eta_{\text{exit}}$.

The exit externality depends on how rapidly the tax base (the survivor share $1-\theta$) shrinks as the effective wedge $\Delta = \ln \frac{1+\tau^d}{1+\tau^a}$ increases. This is a function of the exit elasticity ψ.[5]

Combining these components:

$$\text{MCPF} \in \left[\frac{\beta_{\text{survivor}}(1-\theta)}{1+\eta_{\text{total}}}, \frac{\beta_{\text{survivor}}}{1+\eta_{\text{total}}} \right] = \left[\frac{\beta_{\text{survivor}}(1-\theta)}{1+\eta_{\text{other}} - \frac{\theta}{1-\theta}\psi}, \frac{\beta_{\text{survivor}}}{1+\eta_{\text{other}} - \frac{\theta}{1-\theta}\psi} \right], \tag{5}$$

where η_{other} captures all other fiscal externalities, taken as given from other studies. A larger ψ increases the MCPF through the denominator. When agents are highly responsive to the effective wedge, the base erodes rapidly, so raising rates generates less revenue per unit of welfare cost. Appendix A.6 provides a self-contained derivation of this formula from the social planner's problem, showing how domestic incidence β and the fiscal externality η emerge from the first-order conditions of welfare maximization subject to a revenue constraint.

This expression clarifies what must be measured. To correct the numerator, we need the exit share θ. To correct the denominator, we need the exit elasticity ψ. Section 5 combines these estimates to evaluate the MCPF for tariffs.

2.6 Deductibility as Implicit Enforcement

The return to exit depends not only on the focal tax but also on the broader tax environment. When payments under one tax are deductible against another, that deductibility compresses the effective wedge between staying and exiting. This creates policy spillovers: reforms to one instrument affect compliance with another.

I develop this point in the context of tariffs and the corporate income tax. When U.S. firms import intermediate or capital goods, they deduct the tariff-inclusive cost from taxable income.

5. The fiscal externality from exit is the elasticity of the tax base with respect to the effective wedge:

$$\eta_{\text{exit}} \equiv \frac{\partial(1-\theta)}{\partial\Delta}\frac{1}{1-\theta} = -\frac{\partial\theta}{\partial\Delta}\frac{1}{1-\theta} = -\frac{\theta}{1-\theta}\psi.$$

This deductibility compresses the effective wedge between direct shipments and rerouting, making exit less attractive. The corporate tax code thus acts as implicit enforcement, with higher corporate tax rates τ^c or greater deductibility z reducing the incentive to reroute.

Tax Deductibility and the Exit Wedge. Suppose the importer faces corporate tax rate $\tau^c \in [0, 1)$ with partial deductibility $z \in [0, 1]$ of tariff-inclusive input costs, where z is the net present value of a dollar of deductions. For intermediate inputs expensed immediately, $z = 1$; for capital expensed over time, $z < 1$. After-tax unit costs in each channel are

$$C_d = (1 - \tau^c z)(1 + \tau^d)p, \qquad C_a = (1 - \tau^c z)p,$$

yielding a cost gap

$$\Delta C_{\text{ded}} = (1 - \tau^c z)\tau^d p.$$

In the general two-wedge model, the gap is $\Delta C_{\text{wedge}} = (\tau^d - \tau^a)p$. Equating these yields a simple isomorphism.

Corollary 1 (Deductibility-Wedge Isomorphism)**.** *When imported inputs are deductible, the exit problem is behaviorally equivalent to a two-wedge problem with effective exit wedge*

$$\tau^a = z\tau^c\tau^d.$$

With no deductibility, $\tau^a = 0$.

The isomorphism implies that corporate tax policy and trade enforcement are linked. A higher corporate tax rate τ^c or greater expensing z raises the effective exit wedge τ^a, discouraging rerouting. Conversely, cutting corporate taxes lowers implicit enforcement, increasing exit. Section 5 quantifies this effect, showing that the 2017 TCJA rate cut from 35% to 21% reduced tariff revenue by \$2–9 billion through increased rerouting.

Combining Corollary 1 with Proposition 2 yields comparative statics for domestic tax policy:

$$\frac{d\widehat{m}}{d\tau^c} < 0 \quad \text{and} \quad \frac{d\widehat{m}}{dz} < 0.$$

Higher corporate tax rates or greater deductibility reduce exit. The cutoff $\widehat{m}$ shifts left, the survivor set expands, and the marginal firm is less inclined to exit because deductibility loads a share $z\tau^c$ of the direct tariff into the exit channel.

3 Measuring Tariff-Induced Rerouting

The theoretical framework in Section 2 established that selection out of the observed tax base biases measured incidence upward, with the exit share θ and exit elasticity ψ as sufficient statistics for correction. The framework is general; this section and the next specialize to tariffs, where exit takes the form of rerouting through third countries. Rerouting leaves a detectable footprint in bilateral trade data, unlike other exit margins such as misclassification or value misreporting, making it possible to construct θ from public sources. It is also quantitatively significant—Do et al. (2025) document that the majority of U.S. containerized imports pass through transshipment hubs—and has attracted substantial policy attention, with U.S. Customs and Border Protection launching enforcement actions against hub countries.

Two empirical tasks follow from the theory. First, measuring the level of rerouting θ to correct the numerator of the MCPF. Second, estimating the rerouting elasticity ψ to quantify the fiscal externality in the denominator. This section develops the methodology; Section 4 presents the empirical estimates.

While the model characterizes firm-level routing decisions, the empirical analysis operates at the HS6 product level, where trade flows aggregate across many exporters. Appendix B formalizes the link: under the model, the product-level rerouting share θ_k reflects the mass of firms below the marginal cost cutoff, weighted by their trade volumes. The mapping is not exact because it depends on the distribution of firm sizes within each product, but the qualitative predictions carry through: θ rises with the wedge, and the composition of survivors shifts toward higher-cost, higher-passthrough firms. This justifies using HS6-level variation in the empirical analysis.

I exploit the 2018 U.S. tariffs on China as a natural experiment, extending the methodology of Freund (2024) to construct bounded measures of the share of goods originating from China but rerouted through third countries to the United States over 2012–2023.

3.1 Data

I use annual data on bilateral trade flows from CEPII's BACI HS-12 vintage from 2012–2023 at the 6-digit product (HS6) level.[6] For each year t, country pair (i, j), and HS6 code k, I observe the customs value of shipments $v_{i\to j,k,t}$. The goal is to identify when goods originate in China, flow to a hub country on the first leg, and then arrive in the United States on the second leg. I restrict potential hubs to 36 countries selected for geographic proximity to China, port infrastructure, and pre-tariff trade relationships with both the U.S. and China (Table C.1). Leave-one-out analysis

6. BACI provides consistent bilateral flows at annual frequency with comprehensive country coverage. Monthly UN Comtrade data have gaps, reporting lags, and reconciliation issues across partner countries that would complicate the screening methodology.

confirms no single hub drives the results. I deflate all values to 2017 dollars using the BLS end-use import price index.

The tariff wedge Δ operationalizes the theory's routing incentive, which is the gap between the effective cost of direct shipping versus rerouting:

$$\Delta_{k,t} \equiv \log \frac{1+\tau^d_{k,t}}{1+\tau^a_{k,t}} = \log \frac{1+\tau^d_{k,t}}{1+z_{k,t}\tau^c_t\tau^d_{k,t}}.$$

Constructing this wedge requires data on two components: the direct tariff τ^d and the effective cost in the rerouting channel $\tau^a = z\tau^c\tau^d$, which Section 2.6 shows depends on domestic tax policy.

For the direct tariff, I use Amiti, Redding, and Weinstein (2020) for monthly HS10 China-specific rates through 2017, aggregated to annual HS6 using 2017 import shares as weights.[7] From 2018–2023, I use the Global Tariff Database from Rodríguez-Clare, Ulate, and Vasquez (2025), which provides bilateral HS6 tariffs based on the methodology in Teti (2025).

For the rerouting wedge τ^a, I use the UN's Broad Economic Categories (BEC) to classify each product as consumption, intermediate, or capital. For consumption and intermediate goods, which are fully expensed ($z_k = 1$), Corollary 1 implies $\tau^a = \tau^c\tau^d$.[8] For capital goods, $z_{k,t}$ is the net present value of depreciation deductions under MACRS.[9] I map each capital good to its IRS tax life and compute $z_{k,t}$ accordingly. Figure D.1 plots the distribution of the resulting wedge $\Delta_{k,t}$ by end-use category over time.

Products that span multiple BEC categories are included once per category, with regression weights equal to the product's BEC category share multiplied by its share of 2017 China-U.S. imports. This ensures each product contributes to estimates for all relevant end-use categories in proportion to its actual composition.

3.2 Constructing the Exit Share

Measuring rerouting is inherently difficult: unlike domestic tax evasion, there is no audit trail or third-party reporting. I construct bounds on θ by screening bilateral trade data for patterns consistent with Chinese goods flowing through third countries en route to the United States. The approach extends Freund (2024), who develops conditions to identify transshipment in bilateral

7. Trade-weighted aggregation reflects the tariff rate on goods actually traded. Results are robust to alternative aggregation schemes including time-weighting and simple averages.

8. The effective tax rate may differ from statutory due to firm-level tax planning, and not all importers face the corporate rate (e.g., pass-through entities). This introduces measurement error in Δ. Under classical measurement error, this attenuates elasticity estimates toward zero, making the estimates conservative.

9. For an asset with tax life T years and depreciation schedule $\{d_s\}_{s=0}^T$, the NPV of deductions at discount rate r is $z = \sum_{s=0}^T (1+r)^{-s} d_s$. With bonus depreciation share b, this becomes $z_{\text{bonus}} = b + (1-b)z$. Under 100% bonus depreciation (post-TCJA), $z = 1$.

trade data; I tighten these conditions and calibrate thresholds to control the false positive rate. The methodology complements firm-level evidence from Vietnamese customs records (Iyoha et al. 2025) and bill-of-lading data (Do et al. 2025).

The methodology imposes five conditions that a product-hub-year triplet must satisfy simultaneously. First, the product must have faced U.S. tariff increases on China, which restricts attention to goods with a policy-induced incentive to reroute. Second, China's share of U.S. imports must have fallen below its pre-trend while a hub country's share rose above its pre-trend, which is a reallocation pattern consistent with rerouting. Third, China's exports to the rest of the world must have remained stable relative to pre-trends, ruling out the alternative that China lost global competitiveness in the product. Fourth, the hub must face a lower U.S. tariff than China on that specific product, ensuring an economic incentive to reroute through that hub. Fifth, China-to-hub trade flows must be large enough to plausibly supply the hub-to-U.S. surge. If Vietnamese firms were expanding domestic production rather than re-exporting Chinese goods, we would not observe a corresponding increase in Chinese shipments to Vietnam for the same products in the same years.

The screens are stringent by design. Each condition is evaluated against product-specific pre-trends estimated from 2012–2017, so a triplet is flagged only when trade patterns deviate sharply from historical norms. I construct two versions of the measure: a conservative version that flags only above-trend growth on both legs of the route, and a liberal version that also captures scaling of pre-existing routes. Both versions are calibrated against a placebo sample of products that never faced tariff increases. I select thresholds such that the implied exit share among these never-treated products does not exceed 1% in any pre-tariff year. This calibration deliberately prioritizes specificity over sensitivity: the methodology is designed to minimize false positives at the cost of missing some genuine rerouting.

For the sufficient statistics correction developed in Section 2, the precise mechanism of exit does not matter. Selection bias in survivor-based incidence estimates arises whenever firms leave the direct channel, regardless of whether exit reflects physical rerouting, supply chain reorganization, or sourcing shifts. What matters is the share of trade that leaves the observed sample, not how it leaves. The measure θ should therefore be interpreted as tariff-induced exit from the direct China-U.S. channel. It is conservative: it omits other avoidance margins such as value misreporting and product misclassification, and the stringent screening criteria yield a lower bound on total exit. Implementation details are in Appendix C.

4 Measured Transshipment

The theoretical framework established the two key parameters needed to evaluate the welfare cost of tariffs: the level of avoidance θ and the avoidance elasticity. This section provides empirical estimates for both. While the model in Section 2 is characterized at the firm level, its predictions about selection and passthrough aggregate to the observable product-level trade flows used in the analysis that follows (see Appendix B for a detailed discussion).

4.1 The Magnitude of Transshipment

Figure 2 plots transshipment on the extensive and intensive margins. Prior to the 2018 tariffs, both margins are low: around 10 percent of HS6 codes show any evidence of rerouting, and transshipment volume is near zero as a share of baseline trade. Following the tariffs, both margins rise sharply. By 2023, nearly 40 percent of HS6 codes exhibit positive transshipment under the liberal screen, and around 35 percent under the conservative screen. On the intensive margin, the liberal screen identifies transshipment equal to roughly 15 percent of 2017 baseline trade, while the conservative screen identifies around 5 percent.

Figure 2: Aggregate Transshipment on the Intensive and Extensive Margins

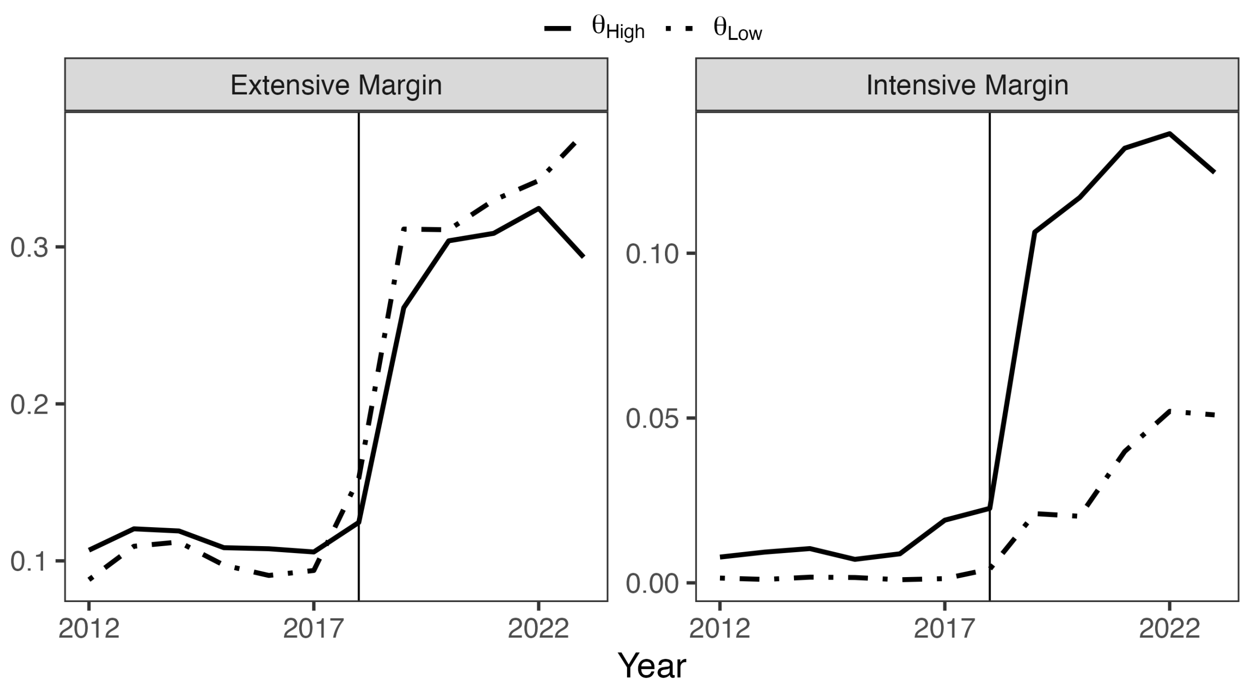

Notes: The left panel plots the share of HS6 codes which exhibit positive transshipment (extensive margin). The right panel plots the transshipment share of 2017 baseline trade (intensive margin). Conservative (low) and liberal (high) bounds correspond to the growth-based and levels-based screens, respectively.

These magnitudes are economically large and consistent with independent estimates. Iyoha et al. (2025) document similar patterns in Vietnamese customs records using firm-level data, and Freund (2024) finds comparable product-level trends. Since transshipment is only one avoidance channel, these estimates provide a lower bound on total exit. The aggregate patterns mask heterogeneity by end-use. Appendix Figure D.2 shows that transshipment is concentrated in capital and intermediate goods, with far lower magnitudes for consumption goods. This pattern is consistent with the theoretical prediction that exit responds to the net-of-deduction tariff wedge, which is larger for business inputs.

4.2 Identification

Figure 2 suggests that the tariffs caused transshipment to increase. But aggregate time series cannot establish causality: transshipment might have been rising for unrelated reasons, or the screens might mechanically flag any trade reallocation. This section addresses both concerns.

Event Study Design

I estimate an event study that tests whether products with larger tariff exposure experienced more transshipment, and whether this relationship emerged only after the tariffs took effect. For each HS6 code k, I define exposure as the product's tariff wedge in 2019, $\Delta_{k,2019}$—the first year of full implementation.[10] The specification is:

$$f(\theta_{k,t}) = \alpha_k + \delta_{t\times e(k)} + \sum_{r\neq -1} \beta_r \,\mathbf{1}\{\text{rel}_t = r\} \times \Delta_{k,2019} + \varepsilon_{k,t}, \tag{6}$$

where rel_t denotes years relative to 2018, α_k are HS6 fixed effects, and $\delta_{t\times e(k)}$ are year×end-use fixed effects. The outcome is either the inverse hyperbolic sine of transshipment or an indicator for any transshipment.[11] Observations are weighted by 2017 import shares with standard errors are clustered by HS6. Each coefficient β_r answers: in year 2018+r, how much more transshipment occurred in high-exposure products than in low-exposure products, relative to the baseline year 2017?

10. I anchor exposure at 2019 so that coefficients measure the response per unit of the long-run wedge rather than the time-varying contemporaneous exposure. Results are robust to alternative anchoring years (Figure S1).

11. The inverse hyperbolic sine, $\sinh^{-1}(\theta) = \ln(\theta + \sqrt{\theta^2 + 1})$, approximates the log for large values while accommodating zeros.

Results

Figure 3 plots the coefficients. Two patterns stand out. First, pre-trends are flat. From 2013 through 2017, the coefficients are small and statistically insignificant for all four outcomes. High-exposure and low-exposure products were not diverging before the tariffs. Block-permutation tests (Table E.2) confirm this: for three of the four measures, pre-period transshipment does not differ significantly across exposure levels (Holm-adjusted p-values > 0.2). The conservative intensive margin shows a small positive gap ($p = 0.005$), but as Figure 3 shows, this reflects a level difference, not a trend. Second, the post-2018 response is sharp and persistent. Both margins jump discretely when the tariffs take effect and remain elevated through 2023. The intensive margin rises immediately; the extensive margin increases in 2018–2019, stabilizes, and rises again in later years. This persistence indicates durable supply chain restructuring rather than temporary adjustment.

Figure 3: Event Study: Transshipment Response to Tariff Exposure

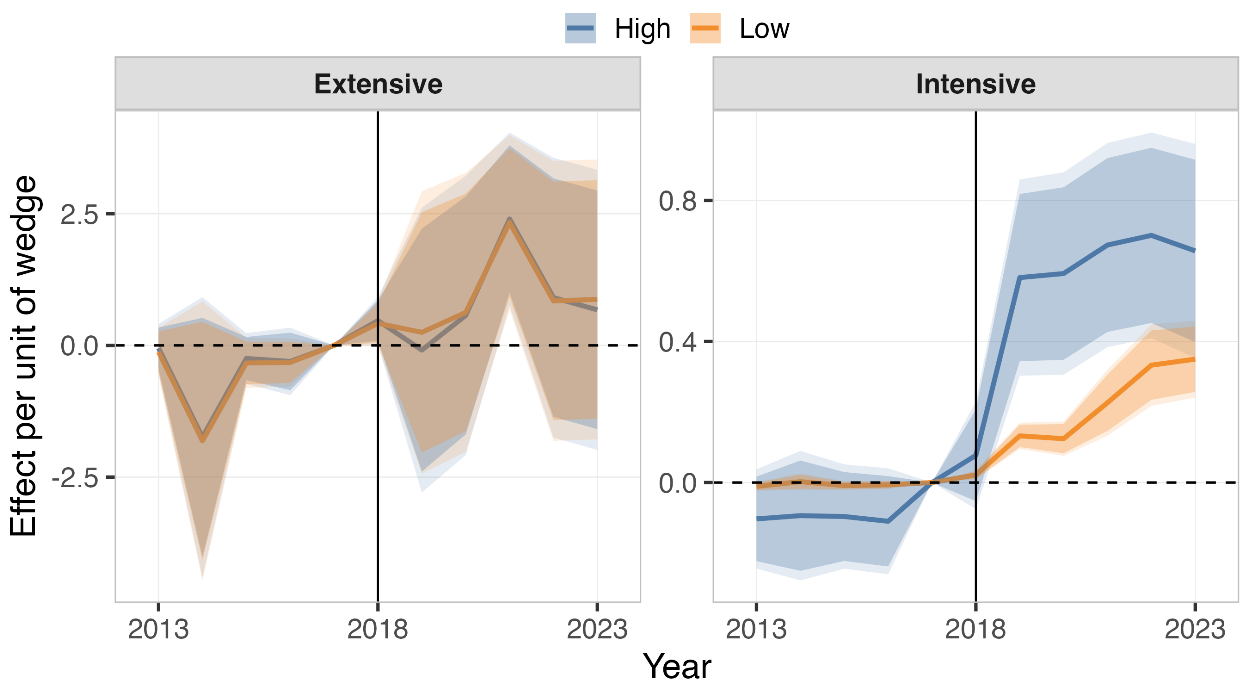

Note: Coefficients β_r from equation (6). Extensive margin (left) and intensive margin (right), with 90% and 95% confidence intervals. Conservative bounds in orange; liberal bounds in blue.

A natural concern is that this design is circular. The screens detect trade patterns consistent with rerouting, so any post-2018 increase might seem mechanical: tariffs shift trade, the screens flag it, and the event study recovers this by construction. This concern confuses detection with identification. The screens determine *whether* a product-hub-year triplet gets flagged. The event study asks a different question: *among flagged products*, did those with larger tariff wedges experience more transshipment? The coefficient β_r measures this differential—high-exposure versus low-exposure products in each year—not the aggregate level shift. Flat pre-trends rule out confounding on this differential. If high-exposure products were trending toward more transshipment before 2018 for reasons unrelated to tariffs, we would see $\beta_r > 0$ in the pre-period. We do not. The combination of flat pre-trends and a sharp post-2018 break establishes that tariff exposure causally determines the magnitude of exit.

Selection in the Trade Data

The model predicts that transshipment should reshape observable trade flows through two channels: low-passthrough firms exit the direct-shipping sample (extensive margin), and the composition of surviving exporters shifts toward higher-cost, higher-passthrough firms (intensive margin). Table 1 tests both predictions by regressing deviations from pre-2018 trends in direct China→US trade on measured exit intensity. Columns 1–2 show quantity effects. Products with higher θ exhibit significantly larger declines in direct import volumes relative to their pre-2018 trends. The coefficients are close to negative one, indicating that exit substitutes nearly one-for-one for direct trade: what disappears from the direct channel reappears in the rerouting channel. Columns 3–4 test whether exit is selective. Proposition 2 predicts that low-cost, high-markup firms are most likely to exit—these firms have the highest passthrough and thus the strongest incentive to avoid the tariff. If such firms also charge higher prices (due to quality premia or market power), their exit should lower average prices among surviving direct shippers. This is what we find: higher θ is associated with lower Chinese export prices within HS6-year cells.

Three interpretations are consistent with the price pattern. First, low-cost, high-markup firms—which Proposition 2 predicts are most likely to exit—may also charge high prices due to quality premia, so their exit mechanically lowers average unit values. Second, surviving firms may downgrade quality to preserve market share under tariff pressure. Third, firms may preferentially reroute higher-value varieties to maximize tariff savings. Distinguishing these mechanisms requires firm-product panel data tracking the same varieties across channels. For this paper's purposes, all three support the central claim: selection into exit biases measured passthrough upward, and θ bounds the magnitude of this bias.[12] Together, these patterns confirm that the firms exiting differ systematically from those that remain, exactly as the model predicts.

12. If quality downgrading is significant, the welfare cost includes not just reduced trade volumes but also quality losses borne by consumers—a margin not captured in the sufficient statistics framework.

Table 1: Exit Intensity and Direct China–US Trade

	$\Delta^{\text{trend}}(\text{CN}\rightarrow\text{US})_{k,t}$			
	Quantities		Prices	
	(1)	(2)	(3)	(4)
$\sinh^{-1}(\theta_{\text{low}})$	-0.952***		-0.945**	
	(0.306)		(0.386)	
$\sinh^{-1}(\theta_{\text{high}})$		-0.732***		-0.448***
		(0.142)		(0.139)
R^2	0.46	0.47	0.34	0.34
Observations	53,485	53,485	53,134	53,134
HS6 & Year×Use FE	✓	✓	✓	✓
Weighted	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Regressions of $\Delta^{\text{trend}}_{\text{CN}\rightarrow\text{US}}$—the deviation of HS6-level direct China→US imports from their 2012–2017 linear trend—on exit intensity. Columns 1–2: log import values. Columns 3–4: China FOB unit values. All specifications include HS6 and year×end-use fixed effects; standard errors clustered by HS6.

4.3 From θ to Passthrough: A Portable Correction

The transshipment shares from Section 4.1 provide a direct correction for the bias in survivor-based passthrough estimates. Proposition 3 showed that when low-passthrough firms select into exit, survivor-based estimates overstate the true cohort passthrough:

$$\beta_{\text{cohort}} \in \left[(1-\theta_{\text{high}})\beta_{\text{survivor}},\ \beta_{\text{survivor}}\right].$$

This bound requires only the exit share θ—a sufficient statistic under monotone selection—and does not require estimating a structural selection model.

Figure 4 applies this correction to existing survivor-based estimates of tariff passthrough. Prior work finds that U.S. importers bore the full incidence of the 2018 China tariffs, with passthrough coefficients indistinguishable from one (Amiti, Redding, and Weinstein 2019, 2020; Cavallo et al. 2021; Fajgelbaum et al. 2020). Taking $\beta_{\text{survivor}} = 1$ as the empirical benchmark and applying the exit shares from Figure 2, the corrected passthrough is approximately 0.85 in the post-2019 period.

The correction varies substantially by end-use. For consumption goods, where exit is minimal, corrected passthrough remains near unity—consistent with the cohort-design findings of Flaaen, Hortaçsu, and Tintelnot (2020) and Cavallo et al. (2021). For capital and intermediate goods, where exit is concentrated, corrected passthrough falls to approximately 0.80. The shaded regions in Figure 4 show the feasible range given the conservative (θ_{low}) and liberal (θ_{high}) exit bounds, with the hatched band marking the conservative floor.

Figure 4: Corrected Passthrough Coefficient by End Use

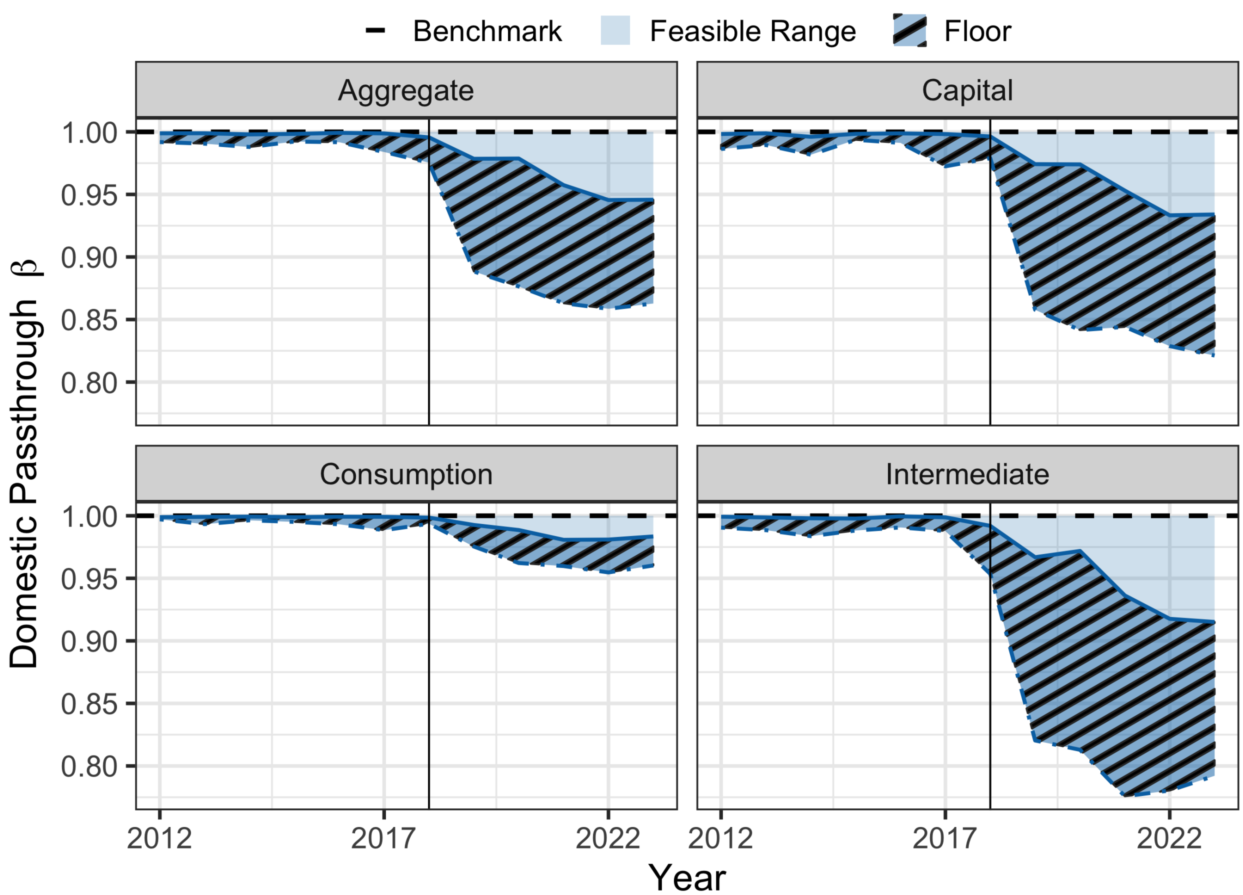

Note: Shaded regions show corrected domestic passthrough relative to the survivor-based benchmark. The light blue area indicates the full feasible range implied by the model, $\beta_{\text{cohort}} \in [(1-\theta_{\text{high}})\beta_{\text{survivor}}, \beta_{\text{survivor}}]$. The hatched band marks the conservative floor using the two screens, $[(1-\theta_{\text{high}})\beta_{\text{survivor}}, (1-\theta_{\text{low}})\beta_{\text{survivor}}]$. The dashed line shows the survivor-based benchmark $\beta_{\text{survivor}} = 1$.

This convergence between the corrected survivor-based estimates and cohort-design findings validates both approaches. Cohort designs correctly measure passthrough for direct shippers by following the same firms before and after the tariff. Survivor-based designs capture the selected sample of firms that remain in the direct channel. The correction shows that once we account for who exits, the two approaches yield consistent estimates of true incidence.

The remaining wedge between 0.80 and 1.0 for business inputs suggests that even after correcting for exit, some burden falls on foreign exporters. Two explanations are possible. First, exporters may absorb part of the tariff directly. Second, transshipment may understate total exit because it omits other avoidance margins such as value misreporting and product misclassification. Under the second interpretation, 0.80 is an upper bound on true corrected passthrough for these goods—the actual incidence on U.S. importers may be lower still.

4.4 The Tariff Elasticity of Exit

Correcting measured incidence requires only the level of exit θ. But evaluating the welfare cost of tariffs through the marginal cost of public funds also requires characterizing the fiscal externality: how much the tariff base erodes as rates rise. This section estimates the exit elasticity with respect to the effective wedge using within-HS6 variation in tariff exposure.

The baseline specification is

$$f(\theta_{k,t}) = \alpha_k + \delta_{t\times e(k)} + \psi \cdot \Delta_{k,t} + \varepsilon_{k,t}, \tag{7}$$

where $\Delta_{k,t} \equiv \log[(1+\tau^d_{k,t})/(1+\tau^a_{k,t})]$ is the effective tariff wedge, α_k are HS6 fixed effects, $\delta_{t\times e(k)}$ are year×end-use fixed effects, and standard errors are clustered by HS6. The outcome $f(\theta_{k,t})$ is either $\sinh^{-1}(\theta_{k,t})$ (intensive margin) or $\mathbf{1}(\theta_{k,t} > 0)$ (extensive margin). Observations are weighted by 2017 China→US import shares.

The year×end-use fixed effects are motivated by the heterogeneous fixed-cost extension in Appendix A.5, which shows that the exit elasticity ψ varies with product characteristics that affect the cost of rerouting. By allowing consumption, capital, and intermediate goods to follow different aggregate time paths, these fixed effects absorb common shocks within each category while preserving the identifying variation from differential exposure to the wedge within end-use groups. Table S12 in the appendix confirms this intuition, showing that products with pre-existing hub networks, lower bulkiness, or electronics content exhibit larger wedge elasticities along the intensive margin, although that is not true along the extensive margin (Table S13).

Identification comes from within-product changes in the effective wedge induced by the 2018–2019 tariff increases. The product fixed effect α_k absorbs time-invariant product characteristics, while the year×end-use fixed effect absorbs common shocks within categories. The coefficient ψ identifies the causal elasticity under a parallel-trends assumption: conditional on these fixed effects, high-wedge and low-wedge products would have followed the same exit trajectory absent the tariff. The event study in Section 4.2 provides direct evidence for this assumption—pre-2018 coefficients are flat and jointly insignificant, indicating that high- and low-exposure products were not differentially trending before the tariffs took effect.

Because the conservative and liberal exit measures provide bounds rather than point estimates, I report joint Imbens and Manski (2004) confidence sets that are guaranteed to contain the true effect under either specification.[13] Figure 5 shows these bounds scaled to represent the impact of a 10 percentage point increase in the tariff wedge.

A 10 percentage point increase in the effective tariff wedge raises the probability that a product exhibits any exit by approximately 10–30 percentage points (extensive margin). On the intensive margin, it increases the transformed exit share by 0.03 to 0.11 in $\sinh^{-1}$ units—corresponding roughly to a 3 to 12 percentage point increase in the exit share for products near the sample mean.[14] These effects are largest for capital and intermediate goods and smallest for consumption goods (Appendix Tables S12–S13), consistent with the heterogeneity in exit levels documented in Section 4.1.

Figure 5: The Tariff Elasticity of Exit

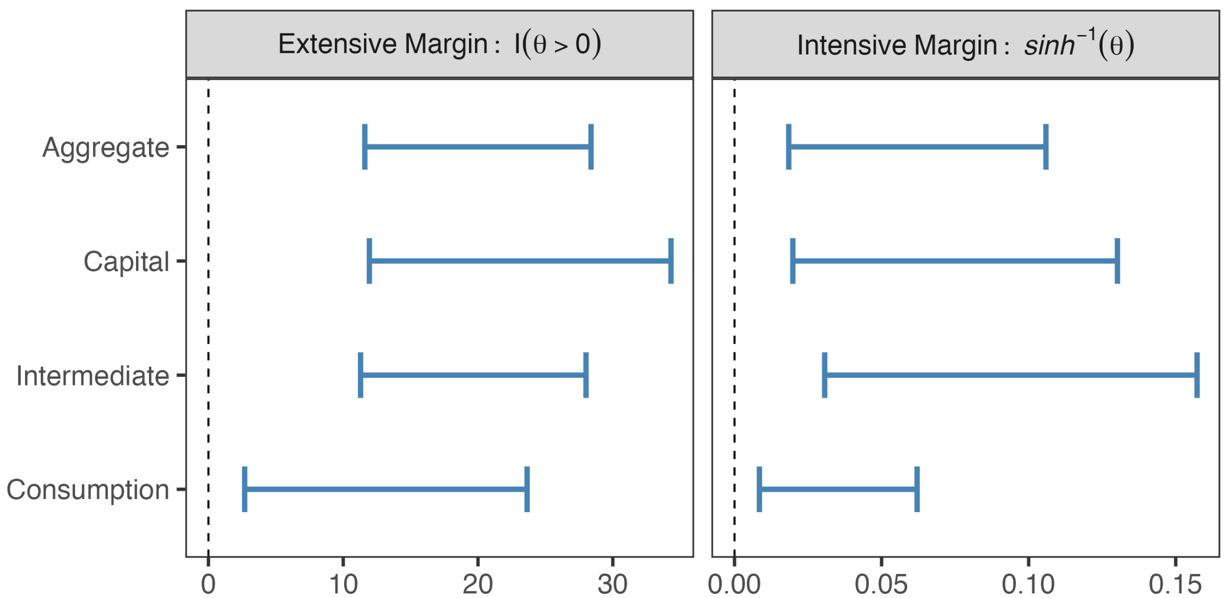

Note: Estimated elasticity of exit with respect to the effective tariff wedge, scaled to show the impact of a 10 percentage point increase. Left panel: effect on the probability of any exit (extensive margin, in percentage points). Right panel: effect on exit share (intensive margin, in $\sinh^{-1}(\theta)$ units). Horizontal bars are 95% joint confidence sets from equation (7), constructed using Imbens and Manski (2004) to account for upper and lower bound estimates.

13. The bounds take the form $\left[\widehat{\beta}_{\text{Low}} - c^\star \, \text{se}(\widehat{\beta}_{\text{Low}}),\ \widehat{\beta}_{\text{High}} + c^\star \, \text{se}(\widehat{\beta}_{\text{High}})\right]$, where $c^\star$ is an adjusted critical value that depends on the estimated correlation $\rho \in (-1, 1)$ between the two slopes. The constant $c^\star$ is obtained by integrating the bivariate normal distribution to achieve joint coverage of $1 - \alpha$.

14. For small θ, $\sinh^{-1}(\theta) \approx \theta$, so the coefficient approximates the level effect. For products with higher baseline exit, the mapping is nonlinear.

The large elasticities indicate that the tariff base erodes substantially as rates rise. This has direct implications for the fiscal externality term in the marginal cost of public funds: when exit is highly elastic, raising tariff rates generates less revenue per unit of welfare cost, increasing the MCPF. Section 5 combines these elasticity estimates with the level corrections from the previous subsection to evaluate the overall welfare cost of the 2018 tariffs.

4.5 Robustness and Falsification

The estimated elasticities pass two sharp falsification tests that directly address potential confounds.

First, a natural concern is that these elasticities reflect global trade dynamics rather than U.S.-specific tariff-induced exit. If the observed relationship between the U.S. effective wedge $\Delta_{k,t}$ and measured exit were driven by broader shifts in China's export patterns—such as capacity constraints, changing comparative advantage, or global supply chain restructuring—we should observe similar effects when applying the exit screens to other destinations. Table E.1 tests this alternative. For both the EU-27+UK and Canada, the coefficients on the U.S. tariff wedge are small and generally statistically indistinguishable from zero across all four outcomes. A few coefficients are marginally significant, but they show no consistent pattern and are substantially smaller than the main estimates for the U.S. This confirms that the wedge-exit relationship is specific to the U.S. market where the tariffs were imposed, ruling out the concern that the results reflect coincidental timing between U.S. tariff implementation and unrelated changes in China's global export networks.

Second, I directly test whether the network structure underlying exit matters. For each year, I randomly reassign hub→US shipments across HS6 codes within HS4 product groups, breaking the directional China→hub→US link while preserving each product's total China exports and each hub's total US-bound shipments. The permutation also preserves the exposure wedge and all fixed effects. This "broken network" placebo retains any mechanical correlation between trade volumes and the wedge but eliminates the exit channel. Figure E.1 shows that across 5,000 such permutations, the resulting coefficients center far below the true estimates, with the observed elasticities falling in the extreme right tail of the null distribution. This confirms that only the intact two-leg network structure generates the large responses observed in the data.

Additional robustness checks in the Supplemental Appendix show that the results are insensitive to pre-trend adjustment (Table S7), long-difference specifications (Table S8), HS4-specific trends (Table S9), entropy balancing (Table S11), and alternative wedge definitions excluding the deductibility channel (Table S4). Leave-one-out analysis excluding individual hubs one at a time (Table S14) confirms that no single country drives the results, though plausible transshipment hubs like Canada, Japan, and Mexico are the most influential.

5 Re-evaluating the Welfare and Revenue Effects of Tariffs

Given the parameters estimated in the previous section, I now re-evaluate the marginal cost of public funds for tariffs. The MCPF provides a welfare-relevant metric for comparing tax instruments by measuring the welfare cost to society of raising one additional dollar of government revenue. A natural benchmark is a hypothetical lump-sum tax, which does not distort economic behavior. For such a tax, the welfare cost of raising one dollar is exactly one dollar, yielding an MCPF of one. Real-world taxes like tariffs, by contrast, distort behavior and generate deadweight losses. The MCPF captures these costs through two channels, represented in the numerator and denominator of the following expression:[15]

$$\text{MCPF} = \frac{\beta}{1+\eta}. \tag{8}$$

The numerator, β, represents the *domestic incidence* of the tariff—the share of the tariff burden borne by domestic consumers and firms rather than foreign exporters. When $\beta = 1$, domestic agents bear the full incidence; when $\beta = 0$, the tariff is effectively a transfer from foreign producers. The denominator, $1 + \eta$, captures the *fiscal externality* from behavioral responses. The elasticity $\eta \leq 0$ reflects how the tax base shrinks as economic agents change their behavior to reduce their tax burden. A larger behavioral response (more negative η) means the government collects less revenue per unit of welfare cost, raising the MCPF.

Existing estimates of the tariff MCPF following the 2018 U.S.-China trade war place it between 1.2 and 1.6 (Finkelstein and Hendren 2020; Jaccard 2021), suggesting tariffs are more costly than lump-sum taxation but comparable to top income taxes. However, these estimates do not fully account for selection into exit. In the numerator, survivor-based estimates overstate domestic incidence because low-passthrough firms exit the sample. In the denominator, the fiscal externality η in prior studies aggregates across multiple behavioral margins—substitution to domestic goods, quality downgrading, and exit—but does not explicitly decompose the exit component or account for how selection biases its measurement.

The net effect on the MCPF is *a priori* ambiguous: lower true incidence reduces the numerator, but a larger exit elasticity (once properly measured) may increase or decrease the denominator depending on how it compares to the implicit exit response embedded in prior estimates. Which

15. As derived in Appendix A.6. Hendren (2016) and Finkelstein and Hendren (2020) provide a thorough overview of the marginal *value* of public funds. Because I focus on taxes rather than transfers, I refer to it as the marginal *cost* for clarity.

effect dominates is an empirical question that this section addresses.

Three caveats qualify the forthcoming welfare analysis. First, the framework is partial equilibrium, abstracting from general equilibrium effects on wages, the terms of trade, and sectoral reallocation.[16] Second, the resource costs of exit—additional shipping, relabeling, intermediary fees—represent genuine deadweight losses not captured in the sufficient statistics approach. Third, because rerouting is only one exit channel, the 85% passthrough correction likely overstates the true burden, making the welfare cost estimates conservative upper bounds.

5.1 Correcting Measured Domestic Incidence

The empirical findings allow a direct correction to existing estimates of the tariff's welfare cost. As shown in Section 2.4, the true domestic incidence can be approximated by scaling the biased survivor-based estimate by $(1 - \theta)$. The literature finds $\beta_{\text{survivor}} \approx 1.0$ for the 2018 tariffs. Applying the aggregate estimate for the level of exit ($\theta_{\text{high}} \approx 0.15$), the corrected scaling factor is approximately 0.85. This implies a new, lower range for the MCPF of roughly 1.0–1.6, depending on the magnitude of the fiscal externality in the denominator. This is a notable result: after correcting for selection bias in the numerator alone, we cannot reject that the welfare cost of the 2018 tariffs approaches that of a distortion-free lump-sum tax.

However, the MCPF also depends on the fiscal externality in the denominator: how rapidly the tax base erodes when rates rise. In reality, the total fiscal externality aggregates across multiple margins: substitution to domestic goods ($\eta_{\text{substitution}}$), quality adjustments ($\eta_{\text{quality}}$), and exit ($\eta_{\text{exit}}$), such that $\eta_{\text{total}} = \eta_{\text{substitution}} + \eta_{\text{quality}} + \eta_{\text{exit}}$. Prior estimates do not decompose these components or explicitly model the exit elasticity. Section 4.4 implies that the elasticity of exit is high. These two forces—lower incidence (pushes MCPF down) versus high exit elasticity (pushes MCPF up)—work in opposite directions. To isolate the welfare cost attributable specifically to exit, I construct an "exit-only" MCPF in the next subsection.

5.2 The Exit-Based MCPF

To isolate the welfare cost of exit, I construct an MCPF using only the exit component of the fiscal externality. This exercise asks: holding fixed the other behavioral responses (substitution, quality adjustment), what is the incremental contribution of exit to the tariff's welfare cost? I use the exit level θ to correct the incidence term and the exit elasticity ψ to construct the exit-specific

16. In particular, I do not model how foreign exporters adjust supply in response to U.S. tariffs, nor do I account for potential terms-of-trade improvements if the U.S. has market power.

fiscal externality, $\eta_{\text{exit}} \approx -\frac{\theta}{1-\theta}\psi$. The exit-based MCPF is then

$$\text{MCPF}_{\text{exit}} \in \left[\frac{(1-\theta)\beta_{\text{survivor}}}{1-\frac{\theta}{1-\theta}\psi}, \frac{\beta_{\text{survivor}}}{1-\frac{\theta}{1-\theta}\psi} \right]. \tag{9}$$

Using the Imbens–Manski bounds on ψ from Section 4.4, I compute a 95% confidence set for this exit-based MCPF, propagating the uncertainty from the elasticity estimates. Figure 6 plots the resulting time series. Before the 2018 trade war, with little exit, the MCPF for all goods is approximately one. After 2018, the story diverges sharply by end-use. For consumption goods, where exit is minimal (Section 4.1), the MCPF remains near one throughout the post-tariff period. The corrected incidence is close to unity (Figure 4), and the fiscal externality from exit is negligible, leaving the welfare cost indistinguishable from a lump-sum tax.

Figure 6: The Marginal Cost of Public Funds by End Use

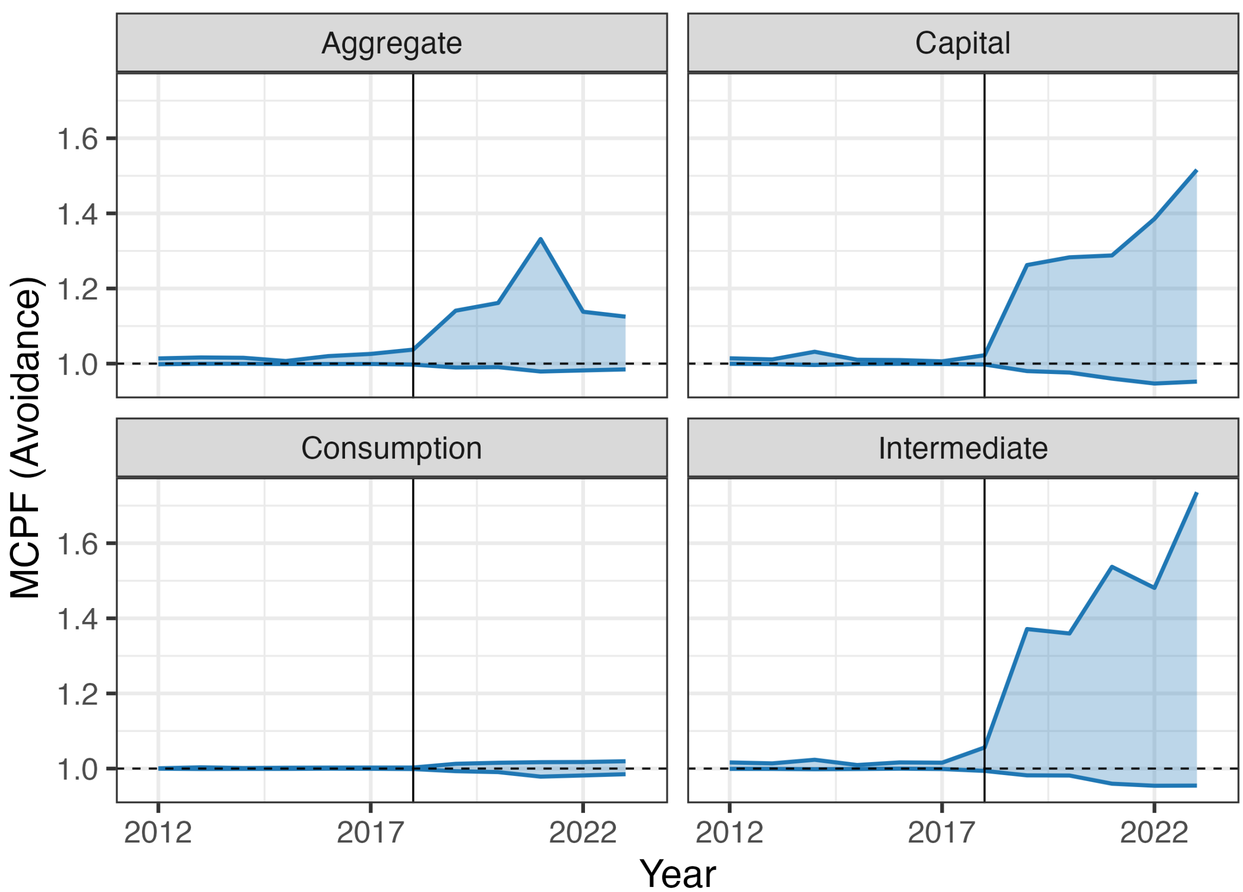

Note: Shaded regions show the exit-based marginal cost of public funds evaluated in each year given the corresponding regressions. Uncertainty is propagated via the Imbens–Manski bands. The dashed horizontal line marks MCPF = 1, the benchmark for a lump-sum tax.

For capital and intermediate goods, however, the picture is starkly different. The high exit elasticity documented in Section 4.4 generates a substantial fiscal externality. This pushes the upper bound of the exit-based MCPF to around 1.5 by 2023. While the lower bound of the estimates remains near one, the wide confidence interval indicates that we cannot rule out a substantial welfare cost for these goods, driven entirely by the exit margin. The mechanism is straightforward: when the tariff base erodes rapidly in response to rate increases (large negative η_{exit}), each dollar of revenue requires greater welfare sacrifice. Combined with the lower corrected incidence for business inputs ($\beta \approx 0.80$ from Figure 4), the net effect leaves the MCPF for capital and intermediate goods meaningfully above one.

This concentration of the welfare cost in production-related goods is consistent with the findings in Alessandria et al. (2025). In their dynamic general equilibrium model, tariffs are especially distortionary for investment because capital goods have a high import share and tariffs distort the intertemporal margin. My paper provides a complementary mechanism: these are precisely the goods where the fiscal externality from exit is largest. Both papers thus underscore the critical interaction between trade and fiscal policy.

Alessandria et al. (2025) focus on how tariff revenue is spent—using it to offset distortionary taxes lowers the net welfare cost. My mechanism highlights how the domestic tax system itself alters the incentive to *pay* the tariff in the first place. Through deductibility, the corporate tax code acts as an implicit enforcement penalty: higher corporate tax rates compress the effective wedge between direct and rerouted shipments, discouraging exit. This interaction was fundamentally altered by the 2017 Tax Cuts and Jobs Act (TCJA), which cut the corporate tax rate from 35% to 21% just before the 2018 tariffs were imposed. The next subsection quantifies the revenue consequences of this policy change.

5.3 The Tax Cuts and Jobs Act's Impact on Tariff Revenue

Recall that the tariff wedge on good k is

$$\Delta_k = \log(1 + \tau_k^d) - \log\left(1 + \tau_k^d \times \underbrace{\tau^c \times z_k}_{\text{Domestic Taxes}}\right),$$

where τ_k^d is the tariff rate on good k, τ^c is the business income tax rate, and z_k is the present value of depreciation deductions. The theoretical monotonicity results and the empirical work together show that when either z_k or τ^c rises, exit declines and tariff revenue rises. Just prior to the 2018 trade war, the 2017 Tax Cuts and Jobs Act increased z_k for all imported capital goods through full expensing (100% bonus depreciation), but the corporate tax rate τ^c fell from 35% to 21%. In principle, both should affect exit and revenue. I construct two counterfactuals to isolate these effects:

- In the first counterfactual, I consider what would have happened to exit and revenue had TCJA not been passed. Under this counterfactual, τ^c remains at 35% and bonus depreciation follows its statutory path from 40% in 2018 to 30% in 2019, and 0% from 2020 onward.

- Under a second counterfactual, I examine the effect of full bonus depreciation for equipment, but maintaining $\tau^c = 0.35$.

In each scenario, I compute the change in estimated exit and the resulting increase in tariff revenue given the larger tax base, comparing to the observed exit and revenue given TCJA's passage. Figure 7 plots the predicted change in the aggregate exit share (top panel) and cumulative tariff revenue (bottom panel) over 2018–2023 under the two counterfactual tax regimes, with Imbens–Manski bounds reflecting uncertainty in the intensive margin elasticity.

Figure 7: Counterfactual Impact of the TCJA on Exit and Tariff Revenue

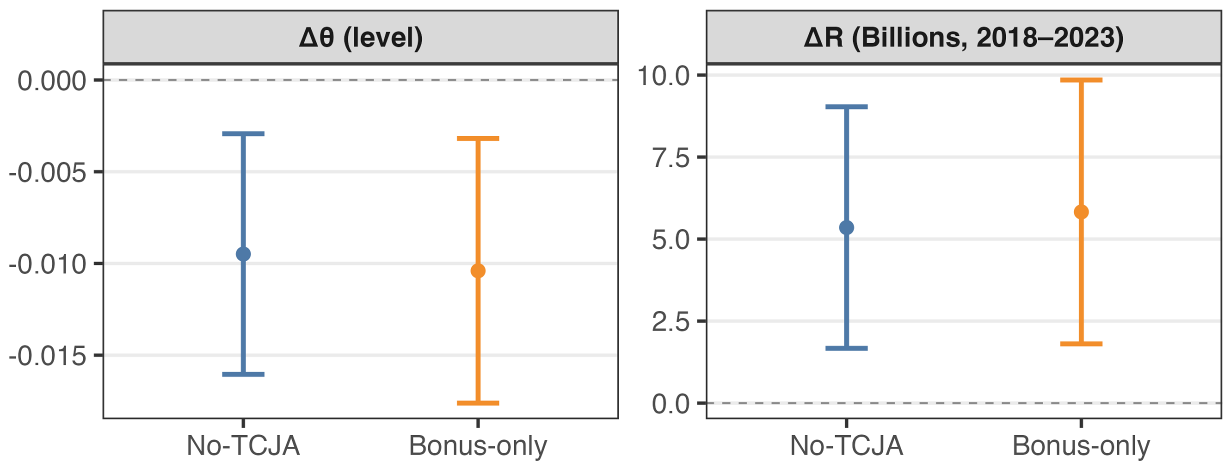

Note: This figure plots the estimated change in the aggregate exit share ($\Delta\theta$) and cumulative tariff revenue (ΔR) over 2018–2023 under two counterfactual scenarios relative to the observed outcome under the TCJA. "No-TCJA" assumes the corporate tax rate remained at 35% and bonus depreciation followed its pre-TCJA phase-out schedule. "Bonus-only" assumes a 35% corporate tax rate but with 100% bonus depreciation. The points represent the central estimates, and the vertical bars show 95% joint confidence sets constructed using the method of Imbens and Manski (2004).

Without the reform ("No-TCJA"), the aggregate exit share would have been lower, implying that firms would have shifted 0.3–1.6 percentage points less of imported value through rerouting channels. Cumulatively, tariff revenue would have been \$2–\$9 billion higher, reflecting both a higher effective tariff base and reduced exit. The second scenario layers in full expensing ($z_k = 1$)

while holding τ^c fixed at 35%. This marginally raises revenue and lowers exit compared to the first scenario. The effect is small because the reform only affects imported equipment, which had a high expensing rate relative to the statutory baseline.

This demonstrates a meaningful and quantitatively significant interaction between the corporate tax base and tariff revenue. Domestic tax policy serves as a powerful, if unintentional, trade enforcement tool. The magnitude—$2 to $9 billion in foregone revenue—is economically substantial and suggests that policymakers evaluating trade enforcement strategies should account for spillovers from the domestic tax system. While the public finance literature emphasizes explicit enforcement tools like audits and penalties, the deductibility of imported inputs provides an implicit enforcement margin that operates automatically through the tax code. Spillovers from domestic tax policy may lessen the need for explicit penalties on rerouting and other forms of exit.

6 Conclusion

This paper shows that tariff avoidance generates economically significant selection bias in measured incidence. When avoidance is costly, low-passthrough firms exit the direct-shipping channel, leaving only high-passthrough survivors in customs data. Measuring this exit in bilateral trade flows, I find 30–40 percent of tariffed products exhibit evidence of rerouting by 2023, implying true aggregate passthrough of approximately 85 percent. The bias is concentrated in intermediate and capital goods where avoidance is highest; consumer goods, where transshipment is minimal, show near-complete passthrough consistent with prior estimates.

The corporate tax code is an implicit enforcement mechanism. Because firms deduct tariff-inclusive import costs, higher corporate tax rates compress the wedge between compliant and non-compliant channels, discouraging avoidance. The 2017 tax cut from 35 to 21 percent widened this wedge, increasing transshipment and costing the U.S. an estimated $2–9 billion in tariff revenue. This spillover operates automatically, without audits or penalties, and has gone unrecognized in both the tax and trade literatures.

Transshipment is one margin among many, and my estimates are conservative lower bounds. The welfare analysis abstracts from general equilibrium effects, and to the extent enforcement on one margin displaces avoidance to others, single-channel policies will prove ineffective. But the selection framework generalizes: wherever agents can exit observed transactions to avoid taxes, measured behavioral parameters will be biased. The sufficient statistic approach developed here provides a methodology for bounding that bias without requiring identification of who avoids or why.

References

Alessandria, George, Jiaxiaomei Ding, Shafaat Khan, and Carter Mix. 2025. *The Tariff Tax Cut: Tariffs as Revenue.* Technical report. Cambridge, MA: National Bureau of Economic Research, May. https://doi.org/10.3386/w33784.

Alfaro, Laura, and Davin Chor. 2025. *An Anatomy of the Great Reallocation in US Supply Chain Trade.* Technical report. Cambridge, MA: National Bureau of Economic Research, November. https://doi.org/10.3386/w34490.

Allingham, Michael G., and Agnar Sandmo. 1972. "Income tax evasion: a theoretical analysis." *Journal of Public Economics* 1, nos. 3-4 (November): 323–338. ISSN: 00472727. https://doi.org/10.1016/0047-2727(72)90010-2.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2019. "International Shocks, Variable Markups, and Domestic Prices." *The Review of Economic Studies* 86, no. 6 (November): 2356–2402. ISSN: 0034-6527. https://doi.org/10.1093/restud/rdz005.

Amiti, Mary, Stephen J. Redding, and David E. Weinstein. 2019. "The Impact of the 2018 Tariffs on Prices and Welfare." *Journal of Economic Perspectives* 33, no. 4 (November): 187–210. ISSN: 0895-3309. https://doi.org/10.1257/jep.33.4.187.

———. 2020. "Who's Paying for the US Tariffs? A Longer-Term Perspective." *AEA Papers and Proceedings* 110. ISSN: 2574-0768. https://doi.org/10.1257/pandp.20201018.

Benzarti, Youssef, and Dorian Carloni. 2019. "Who Really Benefits from Consumption Tax Cuts? Evidence from a Large VAT Reform in France." *American Economic Journal: Economic Policy* 11, no. 1 (February): 38–63. ISSN: 1945-7731. https://doi.org/10.1257/pol.20170504.

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang. 2021. "Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy." *American Economic Review: Insights* 3, no. 1 (March): 19–34. ISSN: 2640-205X. https://doi.org/10.1257/aeri.20190536.

Chetty, Raj. 2009. "Is the Taxable Income Elasticity Sufficient to Calculate Deadweight Loss? The Implications of Evasion and Avoidance." *American Economic Journal: Economic Policy* 1, no. 2 (July): 31–52. ISSN: 1945-7731. https://doi.org/10.1257/pol.1.2.31.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications*." *The Quarterly Journal of Economics* 135, no. 2 (May): 561–644. ISSN: 0033-5533. https://doi.org/10.1093/qje/qjz041.

Deng, Jianpeng, Jialin Li, Joseph Mai, Yanmin Shi, and Linke Zhu. 2025. "Trade circumvention in free trade areas." *Journal of International Money and Finance* 150 (February): 103232. ISSN: 02615606. https://doi.org/10.1016/j.jimonfin.2024.103232.

Do, Anh D., Sharat Ganapati, Woan Foong Wong, and Oren Ziv. 2025. "Transshipment Hubs, Trade, and Supply Chains."

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2023. "How Costly Are Markups?" *Journal of Political Economy* 131, no. 7 (July): 1619–1675. ISSN: 0022-3808. https://doi.org/10.1086/722986.

Fajgelbaum, Pablo, and Amit Khandelwal. 2024. *The Value of De Minimis Imports.* Technical report. Cambridge, MA: National Bureau of Economic Research, June. https://doi.org/10.3386/w32607.

Fajgelbaum, Pablo D, Pinelopi K Goldberg, Patrick J Kennedy, and Amit K Khandelwal. 2020. "The Return to Protectionism*." *The Quarterly Journal of Economics* 135, no. 1 (February): 1–55. ISSN: 0033-5533. https://doi.org/10.1093/qje/qjz036.

Finkelstein, Amy, and Nathaniel Hendren. 2020. "Welfare Analysis Meets Causal Inference." *Journal of Economic Perspectives* 34, no. 4 (November): 146–167. ISSN: 0895-3309. https://doi.org/10.1257/jep.34.4.146.

Fisman, Raymond, Peter Moustakerski, and Shang-Jin Wei. 2008. "Outsourcing Tariff Evasion: A New Explanation for Entrepôt Trade." *Review of Economics and Statistics* 90, no. 3 (August): 587–592. ISSN: 0034-6535. https://doi.org/10.1162/rest.90.3.587.

Fisman, Raymond, and Shang-Jin Wei. 2004. "Tax Rates and Tax Evasion: Evidence from "Missing Imports" in China." *Journal of Political Economy* 112, no. 2 (April): 471–496. ISSN: 0022-3808. https://doi.org/10.1086/381476.

Flaaen, Aaron, Ali Hortacsu, Felix Tintelnot, Nicolas Urdaneta, and Daniel Yi Xu. 2025. "Who Pays for Tariffs Along the Supply Chain? Evidence from European Wine Tariffs."

Flaaen, Aaron, Ali Hortaçsu, and Felix Tintelnot. 2020. "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines." *American Economic Review* 110, no. 7 (July): 2103–2127. ISSN: 0002-8282. https://doi.org/10.1257/aer.20190611.

Freund, Caroline. 2024. *The China Wash: Tracking Products To Identify Tariff Evasion Through Transshipment.* Technical report.

Gopinath, Gita, and Oleg Itskhoki. 2010. "Frequency of Price Adjustment and Pass-Through." *Quarterly Journal of Economics* 125, no. 2 (May): 675–727. ISSN: 0033-5533. https://doi.org/10.1162/qjec.2010.125.2.675.

Hendren, Nathaniel. 2016. "The Policy Elasticity." *Tax Policy and the Economy* 30, no. 1 (January): 51–89. ISSN: 0892-8649. https://doi.org/10.1086/685593.

Imbens, Guido W., and Charles F. Manski. 2004. "Confidence Intervals for Partially Identified Parameters." *Econometrica* 72, no. 6 (November): 1845–1857. ISSN: 0012-9682. https://doi.org/10.1111/j.1468-0262.2004.00555.x.

Iyoha, Ebehi, Edmund Malesky, Jaya Wen, and Sung-Ju Wu. 2025. *Exports in Disguise? Trade Rerouting during the US-China Trade War.* Technical report.

Jaccard, Torsten. 2021. "Who Pays for Protectionism? The Welfare and Substitution Effects of Tariffs." *SSRN Electronic Journal,* ISSN: 1556-5068. https://doi.org/10.2139/ssrn.3967700.

Javorcik, Beata S., and Gaia Narciso. 2008. "Differentiated products and evasion of import tariffs." *Journal of International Economics* 76, no. 2 (December): 208–222. ISSN: 00221996. https://doi.org/10.1016/j.jinteco.2008.07.002.

Kimball, Miles S. 1995. "The Quantitative Analytics of the Basic Neomonetarist Model." *Journal of Money, Credit and Banking* 27, no. 4 (November): 1241. ISSN: 00222879. https://doi.org/10.2307/2078048.

Kopczuk, Wojciech, Justin Marion, Erich Muehlegger, and Joel Slemrod. 2016. "Does Tax-Collection Invariance Hold? Evasion and the Pass-Through of State Diesel Taxes." *American Economic Journal: Economic Policy* 8, no. 2 (May): 251–286. ISSN: 1945-7731. https://doi.org/10.1257/pol.20140271.

Melitz, Marc J, and Gianmarco I. P. Ottaviano. 2008. "Market Size, Trade, and Productivity." *Review of Economic Studies* 75, no. 1 (January): 295–316. ISSN: 0034-6527. https://doi.org/10.1111/j.1467-937X.2007.00463.x.

Mishra, Prachi, Arvind Subramanian, and Petia Topalova. 2008. "Tariffs, enforcement, and customs evasion: Evidence from India." *Journal of Public Economics* 92, nos. 10-11 (October): 1907–1925. ISSN: 00472727. https://doi.org/10.1016/j.jpubeco.2008.04.017.

Rodríguez-Clare, Andrés, Mauricio Ulate, and Jose Vasquez. 2025. *The 2025 Trade War: Dynamic Impacts Across U.S. States and the Global Economy.* Technical report. Cambridge, MA: National Bureau of Economic Research, May. https://doi.org/10.3386/w33792.

Slemrod, Joel. 2019. "Tax Compliance and Enforcement." *Journal of Economic Literature* 57, no. 4 (December): 904–954. ISSN: 0022-0515. https://doi.org/10.1257/jel.20181437.

Slemrod, Joel, and Shlomo Yitzhaki. 2002. "Tax Avoidance, Evasion, and Administration," 1423–1470. https://doi.org/10.1016/S1573-4420(02)80026-X.

Teti, Feodora. 2025. "Missing Tariffs." https://doi.org/10.2139/ssrn.5097020.

A Theoretical Appendix

A.1 Assumptions about Demand

I impose several regularity conditions on demand:

Assumption 1 (Demand)**.** *The demand for each* variety i *depends on its own delivered price* $\tilde{p}_i$ *and on a price aggregator* $\mathbf{P}$ *summarizing rivals. Fixing* $\mathbf{P}$ *as parametric at the variety level, the primitives satisfy:*

D1*.** ***Downward sloping own demand and outward shifts in rivals: $\dfrac{\partial D_i}{\partial \tilde{p}_i} < 0$ *and* $\dfrac{\partial D_i}{\partial \mathbf{P}} > 0$.

D2*.** ***Smoothness: $D_i(\cdot \mid \mathbf{P}) \in C^2$ *in own price.*

D3*.** ***Elastic demand: $\varepsilon_i(\tilde{p}_i \mid \mathbf{P}) \equiv -\dfrac{\tilde{p}_i}{D_i}\dfrac{\partial D_i}{\partial \tilde{p}_i} > 1$.

D4*.** ***Profit concavity / local stability: *For either channel* $c \in \{D, T\}$ *with delivered price* $\tilde{p}_i^c$,

$$\varepsilon(\tilde{p}_i^c \mid \mathbf{P}) - 1 - \kappa(\tilde{p}_i^c \mid \mathbf{P}) > 0.$$

D5*.** ***Kimball curvature: $\kappa_i(\tilde{p}_i \mid \mathbf{P}) \equiv -\dfrac{\partial \ln \varepsilon_i(\tilde{p}_i \mid \mathbf{P})}{\partial \ln \tilde{p}_i} \le 0$ *(equals* 0 *under CES). We also impose two extra conditions:*

(a) $\dfrac{\partial \kappa(\tilde{p}_i \mid \mathbf{P})}{\partial \ln \tilde{p}_i} \le 0$.

(b) $\kappa'(\tilde{p}) \equiv \partial \kappa(\tilde{p})/\partial \ln \tilde{p} \ge -\frac{\varepsilon(\tilde{p})}{\varepsilon(\tilde{p})-1}\kappa(\tilde{p})^2$.

Conditions **D1**–**D4** are standard and ensure well-behaved monopolistic competition with elastic demand and unique profit-maximizing prices. Condition **D5** introduces Kimball-type demand, where $\kappa \le 0$ means the demand elasticity ε is weakly increasing in price. Intuitively, consumers become more price-sensitive as prices rise, making high-priced (high-markup) goods face flatter residual demand curves. This is the source of heterogeneous passthrough: when a tax raises delivered prices, high-markup firms face larger elasticity increases and therefore compress their markups more, leading to lower passthrough. The technical conditions **D5(a)–(b)** ensure that this passthrough is monotone in marginal cost, which is critical for the sorting result in Proposition 2. Without these curvature restrictions, passthrough could be non-monotone, and the selection mechanism would be more complex. Empirically, **D5** is a weak condition: the literature consistently finds that markups vary across firms (Gopinath and Itskhoki 2010; Amiti, Itskhoki, and Konings 2019; De Loecker, Eeckhout, and Unger 2020; Edmond, Midrigan, and Xu 2023), and Kimball demand is a tractable way to microfound this heterogeneity.

A.2 Lerner Condition

Lemma 1 (Lerner condition and monotone comparative statics)**.** *For either channel $r \in \{d, a\}$ and corresponding wedges τ^r, under **D1**–**D4**, any $p_i^*(m_i; r)$ is interior and satisfies*

$$\frac{p_i^* - m_i}{p_i^*} = \frac{1}{\varepsilon_i(\tilde{p}_i^r \mid \mathbf{P})}, \qquad \tilde{p}_i^r = (1+\tau^r)p_i^*.$$

Moreover: (i) p_i^* strictly increases in m_i; (ii) p_i^* is weakly decreasing in the wedge τ^r (constant under CES); and (iii) q_i^* strictly decreases in m_i and τ^r.

Proof. Interior and FOC. By **D2**–**D4**, $\Phi_i(\cdot)$ is strictly concave in p_i; the unique maximizer is interior and characterized by the FOC. Let the delivered price be

$$\tilde{p}_i \equiv \tilde{p}_i(p_i, \tau^r) = \begin{cases} (1+\tau^a)p_i & \text{(exit channel)} \\ (1+\tau^d)\,p_i & \text{(direct channel).} \end{cases}$$

Then $\partial\tilde{p}_i/\partial p_i \in \{1+\tau^a,\ 1+\tau^d\}$, and the FOC is

$$0 = \frac{\partial \Phi_i}{\partial p_i} = D_i(\tilde{p}_i) + (p_i - m_i)\, D_i'(\tilde{p}_i)\, \frac{\partial \tilde{p}_i}{\partial p_i}.$$

With $\varepsilon_i(\tilde{p}_i) \equiv -\tilde{p}_i D_i'(\tilde{p}_i)/D_i(\tilde{p}_i)$,

$$\frac{p_i^* - m_i}{p_i^*} = \frac{\tilde{p}_i}{(\partial\tilde{p}_i/\partial p_i)\, p_i^*} \cdot \frac{-D_i(\tilde{p}_i)}{\tilde{p}_i D_i'(\tilde{p}_i)} = \frac{1}{\varepsilon_i(\tilde{p}_i)}.$$

This is the Lerner condition.

Monotonicity in m_i. Define

$$g(p, m; \tau^r) \equiv D(\tilde{p}) + (p - m)\, D'(\tilde{p})\, \frac{\partial \tilde{p}}{\partial p}, \qquad \tilde{p} = \tilde{p}(p, \tau^r).$$

At the optimum, $g(p^*, m; \tau^r) = 0$. Partials:

$$\frac{\partial g}{\partial m} = -\,D'(\tilde{p})\,\frac{\partial \tilde{p}}{\partial p} > 0 \quad (\mathbf{D1}), \qquad \frac{\partial g}{\partial p} = 2\,D'(\tilde{p})\,\frac{\partial \tilde{p}}{\partial p} + (p-m)\,D''(\tilde{p})\left(\frac{\partial \tilde{p}}{\partial p}\right)^2.$$

Strict concavity of profits (**D4**) implies $\partial g/\partial p < 0$ at p^*. By the implicit function theorem, $\partial p^*/\partial m = -(\partial g/\partial m)/(\partial g/\partial p) > 0$.

Monotonicity in τ^r. For the direct channel, $\tilde{p} = (1+\tau^r)p$. Write

$$G(p, \tau^r) \;\equiv\; \frac{p-m}{p} - \frac{1}{\varepsilon((1+\tau^r)p)} \;=\; 0.$$

Then
$$\frac{\partial G}{\partial \tau^r} = \frac{\varepsilon'((1+\tau^r)p)}{\varepsilon((1+\tau^r)p)^2}\, p \;\geq 0, \qquad \frac{\partial G}{\partial p} = \frac{m}{p^2} + \frac{\varepsilon'((1+\tau^r)p)}{\varepsilon((1+\tau^r)p)^2}\,(1+\tau^r).$$
Under Kimball demand, $\varepsilon'(\cdot) \geq 0$, so $\partial G/\partial \tau^r \geq 0$. Profit concavity (D4, equivalently $\varepsilon - 1 - \kappa > 0$ at the optimum) implies $\partial G/\partial p > 0$. Hence, by the IFT,

$$\frac{\partial p^*}{\partial \tau^r} = -\frac{\partial G/\partial \tau^r}{\partial G/\partial p} \;\leq\; 0,$$

with equality under CES ($\varepsilon' \equiv 0$).

Quantities. In either channel $\tilde{p} = (1+\tau^r)p^*$ and

$$\frac{d\big((1+\tau^r)p^*\big)}{d\tau^r} = p^* + (1+\tau^r)\frac{dp^*}{d\tau^r} = p^*\left(1 + \frac{\kappa(\tilde{p})}{\varepsilon(\tilde{p}) - 1 - \kappa(\tilde{p})}\right) = p^*\frac{\varepsilon(\tilde{p}) - 1}{\varepsilon(\tilde{p}) - 1 - \kappa(\tilde{p})} \;>\; 0,$$

since $\varepsilon(\tilde{p}) > 1$ and the denominator is positive by D4. With $D'(\cdot) < 0$, $q_i^* = D(\tilde{p})$ strictly falls in τ^r; similarly, q_i^* strictly falls in m because p^* rises in m and demand slopes down. This proves (i)–(iii). □

Lemma 1 establishes two key results. First, it confirms the standard Lerner condition: firms set their markup inversely to the elasticity of demand they face. This makes all comparative statics run through how a given wedge, τ^r, shifts the delivered price and how elasticity varies with that price. Second, the lemma establishes the monotone comparative statics that are essential for the paper's sorting mechanism.

Part (i) shows that a higher marginal cost pushes a firm's optimal pre-wedge price upward, providing the clean ordering in m necessary to derive a unique sorting cutoff. Part (ii) highlights that any wedge acts as a demand shifter. Finally, Part (iii) records the quantity implications: a higher marginal cost or a higher wedge τ^r reduces the quantity sold. This decline in profitability within a channel is what ultimately drives selection when wedges differ, a mechanism we formalize next.

A.3 Proof of Proposition 2

Proof. By **D2**–**D4**, the maximizer in $V(m, 1+\tau^r)$ is interior and unique; envelope arguments apply. The primitive $(p-m)D((1+\tau^r)p)$ has increasing differences in $(m, 1+\tau^r)$ because $\partial^2[(p-m)D((1+$

$\tau^r)p)]/\partial m\,\partial(1+\tau^r) = -pD'((1+\tau^r)p) > 0$ by **D1**. Hence V inherits increasing differences in $(m, 1+\tau^r)$ by Topkis. Fix (τ^d, τ^a). Then

$$\frac{\partial H}{\partial m} = V_m(m, 1+\tau^a) - V_m(m, 1+\tau^d) < 0$$

since $1+\tau^a < 1+\tau^d$ and V has increasing differences; continuity of V yields a unique threshold $\hat{m}$ solving $H(\hat{m}; \tau^d, \tau^a) = 0$. For the comparative statics, differentiate the indifference condition:

$$\frac{d\hat{m}}{d\tau^d} = -\frac{\partial H/\partial \tau^d}{\partial H/\partial m}, \qquad \frac{d\hat{m}}{d\tau^a} = -\frac{\partial H/\partial \tau^a}{\partial H/\partial m}.$$

By the envelope theorem at the direct and exit optima p_d^* and p_a^*,

$$\frac{\partial V(m, 1+\tau^r)}{\partial(1+\tau^r)} = (p^* - m)\, p^*\, D'\big((1+\tau^r)p^*\big) \le 0,$$

so $\partial V(m, 1+\tau^d)/\partial \tau^d \le 0$ and $\partial V(m, 1+\tau^a)/\partial \tau^a \le 0$. Hence

$$\frac{\partial H}{\partial \tau^d} = -\frac{\partial V(m, 1+\tau^d)}{\partial \tau^d} \;\ge\; 0, \qquad \frac{\partial H}{\partial \tau^a} = \frac{\partial V(m, 1+\tau^a)}{\partial \tau^a} \;\le\; 0.$$

Since $\partial H/\partial m < 0$ at an interior cutoff, the stated signs follow. □

Figure A.1: An increase in taxes raises the cutoff

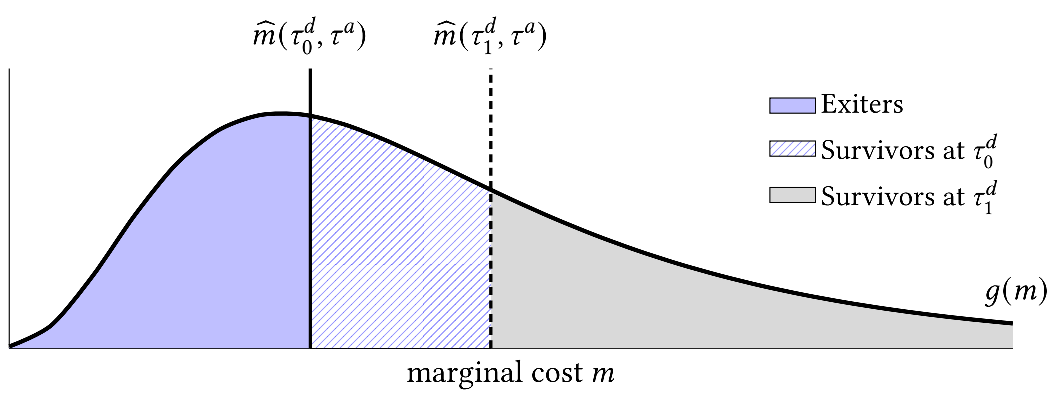

Note: Distribution of marginal costs and wedge-induced selection. The cutoff $\widehat{m}(\tau^d, \tau^a)$ partitions the cohort: firms with $m < \widehat{m}$ exit, while $m \ge \widehat{m}$ stay. A wider wedge gap (e.g., a higher τ^d or a lower τ^a) shifts the cutoff right, increasing the exit share θ and shrinking the survivor set.

A.4 Proof of Proposition 1

Proof. Write the first-order condition for a firm in the direct channel as $G(p, \tau^d) \equiv (p - m)/p - 1/\varepsilon((1 + \tau^d)p) = 0$. Totally differentiate and use

$$\frac{\partial G}{\partial \tau^d} = \frac{\varepsilon'((1 + \tau^d)p)}{\varepsilon((1 + \tau^d)p)^2}\, p, \qquad \frac{\partial G}{\partial p} = \frac{m}{p^2} + \frac{\varepsilon'((1 + \tau^d)p)}{\varepsilon((1 + \tau^d)p)^2}\,(1 + \tau^d).$$

Using $\kappa(u) = -u\,\varepsilon'(u)/\varepsilon(u)$ and the Lerner condition $m = p^*(\varepsilon - 1)/\varepsilon$,

$$\frac{dp^*}{d\tau^d} = -\frac{\partial G/\partial \tau^d}{\partial G/\partial p} = \frac{\kappa(\tilde{p}^d)}{1 + \tau^d} \cdot \frac{p^*}{\varepsilon(\tilde{p}^d) - 1 - \kappa(\tilde{p}^d)}.$$

Finally, the passthrough elasticity is

$$\beta_i = \frac{d \ln\big((1 + \tau^d)p^*\big)}{d\tau^d} = \frac{1}{1 + \tau^d} + \frac{1}{p^*}\frac{dp^*}{d\tau^d} = \frac{1}{1 + \tau^d} \cdot \frac{\varepsilon(\tilde{p}^d) - 1}{\varepsilon(\tilde{p}^d) - 1 - \kappa(\tilde{p}^d)}.$$

Under Assumption **D4**, strict profit concavity at the optimum is equivalent to $\varepsilon - 1 - \kappa > 0$, which ensures the denominator is positive. Moreover, under Assumption **D5(b)**, $\partial\beta/\partial \ln \tilde{p} \geq 0$. Since Lemma 1 implies $\tilde{p}$ increases in m, it follows that $\partial\beta/\partial m \geq 0$. □

A.5 Heterogeneous Fixed Costs of Exit

This appendix extends the baseline model by allowing the fixed cost of exit to vary across firms or products. We show that the selection logic and the survivor–cohort bias results of Propositions 1–3 remain valid under heterogeneity, and we clarify how cross-category differences in fixed costs rationalize the patterns in Section 4.

Corollary 2 (Heterogeneous fixed costs of exit)**.** *Let Assumptions **D1–D5** hold. For each unit i, let the fixed cost of exit F_i be drawn from a continuous distribution G with support $[F_{\min}, F_{\max}]$ and density $g > 0$, independent of the unit's marginal cost m_i. For $r \in \{d, a\}$, write*

$$V(m, F; 1 + \tau^r) \equiv \max_{p \geq m}\, (p - m)\, D\big((1 + \tau^r)p \,\big|\, P\big) \; - \; 1\{r = a\}\, F,$$

and define the exit advantage

$$H(m, F; \tau^d, \tau^a) \equiv V(m, F; 1 + \tau^a) - V(m, F; 1 + \tau^d).$$

Then:

(i) ***Cutoff and comparative statics.*** *For every $F \in [F_{\min}, F_{\max}]$ there exists a unique cutoff $m_b(F;\tau^d,\tau^a)$ solving $H(m_b(F;\tau^d,\tau^a), F;\tau^d,\tau^a) = 0$ such that*

$$m < m_b(F;\tau^d,\tau^a) \Rightarrow \textit{Exit}, \qquad m \geq m_b(F;\tau^d,\tau^a) \Rightarrow \textit{Stay}.$$

Moreover,

$$\frac{\partial m_b}{\partial \tau^d} > 0, \qquad \frac{\partial m_b}{\partial \tau^a} < 0, \qquad \frac{\partial m_b}{\partial F} < 0.$$

(ii) ***Monotone selection and survivor bias.*** *Let $S_i = 1\{m_i \geq m_b(F_i;\tau^d,\tau^a)\}$ denote survival in the direct (survivor) sample. Under Kimball curvature (**D5**), direct-channel passthrough is increasing in marginal cost: $\beta_i = \beta(m_i)$ with $\beta'(m) \geq 0$. Because S_i is weakly increasing in m_i for any realization of F_i, it follows that*

$$\mathrm{Cov}(\beta_i, S_i) \;\geq\; 0,$$

so that the survivor estimate exceeds the cohort average:

$$\beta_{survivor} \;=\; \mathbb{E}\big[\beta_i \mid S_i = 1\big] \;\geq\; \mathbb{E}[\beta_i] \;=\; \beta_{cohort}.$$

(iii) ***Incidence bound unchanged.*** *With $\theta \equiv \Pr\big(m < m_b(F;\tau^d,\tau^a)\big)$ denoting the exit share, the cohort passthrough remains bounded by*

$$\beta_{cohort} \;\in\; \big[(1-\theta)\,\beta_{survivor},\; \beta_{survivor}\big].$$

(iv) ***Cross-category heterogeneity.*** *If F is deterministic by end-use category k, all conclusions hold* within *k:*

$$\beta_{survivor,k} \geq \beta_{cohort,k}, \qquad \beta_{cohort,k} \in \big[(1-\theta_k)\,\beta_{survivor,k},\; \beta_{survivor,k}\big],$$

where $\theta_k \equiv \Pr\big(m < m_b(F_k;\tau^d,\tau^a)\big)$. Lower F_k raises m_b and thus increases θ_k, providing a structural rationale for larger exit among capital/intermediate goods even when their statutory wedge Δ_k is smaller (e.g., because of deductibility).

Proof. Proposition 2 implies that $H(m, F;\tau^d,\tau^a)$ is continuous and strictly decreasing in m, with $\partial H/\partial \tau^d > 0$ and $\partial H/\partial \tau^a < 0$. For any F, define $m_b(F;\tau^d,\tau^a)$ by the indifference condition $H(m_b, F;\tau^d,\tau^a) = 0$. Uniqueness and the signs of the partial derivatives follow from the Implicit Function Theorem:

$$\frac{\partial m_b}{\partial \tau^d} = -\frac{\partial H/\partial \tau^d}{\partial H/\partial m} > 0, \qquad \frac{\partial m_b}{\partial \tau^a} = -\frac{\partial H/\partial \tau^a}{\partial H/\partial m} < 0,$$

$$\frac{\partial m_b}{\partial F} = \frac{1}{\partial H/\partial m} < 0, \quad \text{since } \frac{\partial H}{\partial F} = -1 \text{ and } \frac{\partial H}{\partial m} < 0.$$

Monotonicity of $S_i = 1\{m_i \geq m_b(F_i; \cdot)\}$ in m_i and of $\beta(m)$ in m (Proposition 1) under Kimball demand implies $\mathrm{Cov}(\beta_i, S_i) \geq 0$ by Chebyshev's association inequality; the survivor–cohort ordering and identity of Proposition 3 then deliver $\beta_{\text{survivor}} \geq \beta_{\text{cohort}}$. The bound $\beta_{\text{cohort}} \in [(1-\theta)\beta_{\text{survivor}}, \beta_{\text{survivor}}]$ follows from the law of total expectation and does not depend on whether F is degenerate. Category-specific statements are immediate by conditioning on k. □

Remarks. Independence of F and m is sufficient, not necessary. The results continue to hold whenever the survival probability is weakly increasing in m after integrating over F, e.g., if $F \mid m$ shifts in first-order stochastic dominance toward larger F as m increases (a positive association that makes survival more likely at higher m since $\partial m_b/\partial F < 0$). A sufficiently strong negative association could, in principle, overturn monotone selection.

The heterogeneity in fixed costs can be interpreted as *endogenous* in reduced form. For example, let $F_i = F_0(m_i) - \phi\,\theta$ with $\phi \geq 0$, where θ is the aggregate exit share (network effects or thicker intermediation reduce setup costs). The cutoff solves

$$H(m_b, F; \tau_d, \tau_a) = 0 \quad \text{where } F = F_0(m_b) - \phi\,\theta,$$

and the aggregate share is the fixed point

$$\theta \;=\; T(\theta) \;\equiv\; \Pr\big(m < m_b(F_0(m) - \phi\theta; \tau^d, \tau^a)\big).$$

Since $\partial m_b/\partial F < 0$, we have $\partial m_b/\partial\theta = -\phi\,\partial m_b/\partial F > 0$, so T is monotone; a fixed point exists by Tarski's theorem. Uniqueness holds under a mild slope/contraction condition (e.g., $\sup_\theta |T'(\theta)| < 1$, which obtains for sufficiently small ϕ or under standard log-concavity/MLR shape restrictions). Under these conditions, the comparative statics $\partial m_b/\partial\tau^d > 0$ and $\partial m_b/\partial\tau^a < 0$ carry through, and the survivor–cohort ordering $\beta_{\text{survivor}} \geq \beta_{\text{cohort}}$ is unchanged. Moreover, $\phi > 0$ steepens T and makes $\theta(\tau^d)$ more concave (stronger diffusion/plateau), reinforcing the dynamics in Section 4.

A.6 MCPF Derivation

Suppose a planner has access to a vector of linear tax instruments $\boldsymbol{\tau}$ (with individual elements τ_i). Each tax instrument has a corresponding base $B_i(\cdot)$. For simplicity, assume no spillovers across instruments, so B_i depends only on its own instrument. The government chooses instruments to maximize a well-behaved welfare function $W(\boldsymbol{\tau})$ subject to a revenue constraint:

$$\max_{\boldsymbol{\tau}} W(\boldsymbol{\tau}) \quad \text{subject to} \quad G = R(\boldsymbol{\tau}) \equiv \sum_{i=1}^{N} \tau_i \, B_i(\tau_i). \tag{SPP}$$

Denote λ as the shadow value of public funds:

$$\lambda \equiv -\frac{\partial W/\partial \tau_i}{\partial R/\partial \tau_i} \quad \iff \quad \frac{\partial W/\partial \tau_i}{\partial R/\partial \tau_i} = \frac{\partial W/\partial \tau_j}{\partial R/\partial \tau_j} \;\; \forall\, i, j.$$

With no spillovers,

$$\frac{\partial R}{\partial \tau_i} = B_i(\tau_i) + \tau_i \frac{\partial B_i}{\partial \tau_i} = B_i(\tau_i) \underbrace{\left(1 + \tau_i \frac{(\partial B_i/\partial \tau_i)}{B_i(\tau_i)}\right)}_{1+\eta_i}, \qquad \eta_i \equiv \frac{\partial \ln B_i}{\partial \ln \tau_i}.$$

Under quasi-linear preferences and in partial equilibrium, the marginal welfare cost of increasing each instrument equals the domestic willingness to pay to avoid it, i.e.

$$-\frac{\partial W}{\partial \tau_i} = \beta_i \, B_i(\tau_i),$$

where β_i is the (cohort) domestic incidence per unit increase in instrument i. Hence

$$\text{MCPF}_{\tau_i} = \frac{-(\partial W/\partial \tau_i)}{\partial R/\partial \tau_i} = \frac{\beta_i \, B_i(\tau_i)}{B_i(\tau_i)\,(1+\eta_i)} = \frac{\beta_i}{1+\eta_i}.$$

Specialization to the wedge. In our application the policy instrument is the effective wedge $\Delta \equiv \ln\frac{1+\tau_d}{1+\tau_a}$. Identifying $i = \Delta$ and using $\beta_\Delta = \beta_{\text{cohort}} = (1-\theta)\beta_{\text{border}}$ and $\eta_\Delta = \eta_{\text{other}} - \frac{\theta}{1-\theta}\psi$, the expression above delivers the MCPF used in the main text.

B From Firm-Level Selection to Product-Level Observables

The model developed in Section 2 is stated in general terms, but the empirical application focuses on tariffs. This appendix bridges the firm-level theory and product-level trade data by showing how individual selection decisions manifest in observable HS6-level statistics. Throughout, I specialize to the tariff context: exit takes the form of rerouting through third countries, the direct channel corresponds to recorded China-to-US shipments, and the exit channel corresponds to shipments routed through transshipment hubs.

Consider an HS6 product code k as defining a market populated by a measure-one continuum of monopolistically competitive foreign exporters, each drawing marginal cost m_i from the distribution $F_m(\cdot)$ with density $f_m(\cdot)$. Before the tariff increase, all firms ship directly. After the tariff increase, the routing cutoff $\widehat{m}(\tau^d, \tau^a)$ from Proposition 2 partitions this continuum: firms with $m_i < \widehat{m}$ reroute through hubs, while firms with $m_i \geq \widehat{m}$ continue direct shipping. The product-level rerouting share θ_k is simply the mass of firms below the cutoff:

$$\theta_k = \Pr(m_i < \widehat{m}) = F_m(\widehat{m}(\tau_k^d, \tau_k^a)). \tag{A.1}$$

This aggregation has three immediate implications for product-level observables. First, direct trade volumes fall not only because of standard substitution effects but also because the extensive margin of firms exits the direct channel. For product k, observed direct imports are

$$\text{Imports}_k^{\text{direct}} = \int_{\widehat{m}}^{\overline{m}} q_i^*(m_i; \tau_k^d)\, dF_m(m_i), \tag{A.2}$$

which mechanically declines as $\widehat{m}$ rises. This is the quantity effect documented in Table 1, Columns (1)–(2): products with higher θ_k exhibit larger declines in direct trade flows relative to pre-tariff trends.

Second, and more subtly, the composition of surviving direct exporters shifts toward higher-cost, higher-passthrough firms. Since price is increasing in marginal cost (Lemma 1), the average unit value of direct shipments is

$$\overline{p}_k^{\text{direct}} = \frac{\int_{\widehat{m}}^{\overline{m}} p_i^*(m_i; \tau_k^d) \cdot q_i^*(m_i; \tau_k^d)\, dF_m(m_i)}{\int_{\widehat{m}}^{\overline{m}} q_i^*(m_i; \tau_k^d)\, dF_m(m_i)}. \tag{A.3}$$

As the cutoff $\widehat{m}$ rises with tariff exposure, low-cost firms with low prices reroute, mechanically raising the weighted-average unit value. However, under Kimball demand (**D5**), *passthrough* is also increasing in marginal cost (Proposition 1), so the rerouters are precisely the low-passthrough, high-markup firms. This creates a composition effect that works in the *opposite* direction: the remaining direct shippers are higher-cost but also better able to pass through cost increases.

The net effect on observed unit values depends on which force dominates. Proposition 2 implies that rerouters are low-m_i firms, which under standard pricing (p_i increasing in m_i) have *lower* prices. Their exit should *raise* average unit values. However, Table 1, Columns (3)–(4) documents the opposite: products with higher θ_k exhibit *lower* observed unit values for direct shipments. This negative price-rerouting correlation is consistent with the theoretical prediction that compositional shifts toward higher-cost survivors are overwhelmed by other adjustments. One interpretation is that the rerouters were high-*quality* varieties (which may have low marginal cost but high prices due to quality premia), so their departure lowers average unit values even as average marginal costs rise. Alternatively, remaining firms may engage in quality downgrading to preserve market share, a margin not modeled here. Distinguishing between these mechanisms is an important avenue for future work, but the key point for this paper is that both the quantity decline and the price pattern are consistent with selection driving compositional change in the direct-shipping sample.[17]

Third, the product-level passthrough regression in Equation 3 averages over the selected survivor distribution. Discretizing for exposition, suppose product k consists of N_k firms in the pre-tariff period. Post-tariff, $(1-\theta_k)N_k$ firms ship directly, each with passthrough $\beta_i(m_i)$. The survivor-based regression estimates

$$\widehat{\beta}_{k,\text{survivor}} = \frac{1}{(1-\theta_k)N_k} \sum_{i \in \text{Survivors}} \beta_i(m_i) = \mathbb{E}[\beta_i \mid m_i \geq \widehat{m}], \tag{A.4}$$

which is the survivor-weighted average passthrough. Under Kimball demand, β_i is increasing in m_i (Proposition 1), so $\mathbb{E}[\beta_i \mid m_i \geq \widehat{m}] > \mathbb{E}[\beta_i]$, confirming that the survivor estimate exceeds the true cohort passthrough. The product-level correction in Proposition 3 therefore applies: the bias equals $\text{Cov}(\beta_i, S_i)/\Pr(S_i = 1)$, where now $S_i = \mathbf{1}\{m_i \geq \widehat{m}\}$ is the firm-level survival indicator.

In summary, the product-level rerouting share θ_k aggregates firm-level routing decisions, and observable trade statistics—quantities, unit values, and passthrough estimates—reflect the composition of firms that select into each channel. The empirical strategy in Section 3 exploits this link by measuring θ_k at the HS6 level and using within-product variation in tariff exposure to identify the rerouting elasticity ψ.

C Screening Implementation Details

This appendix provides complete documentation of the methodology used to construct the exit share θ. The goal is to identify tariff-induced rerouting of Chinese goods through third countries to the United States.

C.1 Data and Notation

I use annual bilateral trade flows from CEPII's BACI HS-12 vintage (2012–2023) at the HS6 product level. For each year t, country pair (i, j), and HS6 code k, I observe the customs value of shipments $v_{i \to j,k,t}$.

For each HS6 product k and hub i in year t, define the two potential legs of rerouting:

17. An additional consideration is that rerouters may re-label higher-value goods to evade detection, which would also contribute to the negative correlation between θ_k and observed unit values in the direct channel.

- First leg (China $\longrightarrow$ hub): $\ell_{k,i,t} \equiv v_{\mathrm{CHN}\to i,k,t}$
- Second leg (hub $\longrightarrow$ United States): $x_{k,i,t} \equiv v_{i\to\mathrm{USA},k,t}$

Define market shares as follows. Let $s^{\mathrm{US}}_{k,i,t}$ denote country i's share of U.S. imports in product k:

$$s^{\mathrm{US}}_{k,i,t} = \frac{v_{i\to\mathrm{USA},k,t}}{\sum_j v_{j\to\mathrm{USA},k,t}}.$$

Let $s^{\mathrm{ROW}}_{k,i,t}$ denote country i's share of exports to the rest of the world (excluding the U.S.):

$$s^{\mathrm{ROW}}_{k,i,t} = \frac{\sum_{j\neq\mathrm{USA}} v_{i\to j,k,t}}{\sum_{j\neq\mathrm{USA}}\sum_m v_{m\to j,k,t}}.$$

C.2 Hub Countries

I restrict potential hubs to 36 countries with geographic proximity to China, port infrastructure, and established trade relationships with both China and the U.S. (Table C.1). This excludes implausible routes and reduces false positives from noise in global trade data.

Table C.1: Potential Transshipment Hubs (ISO-3 codes)

ARE	BHR	CAN	CHL	DEU	DOM
EGY	HKG	IDN	IND	ISR	JOR
JPN	KAZ	KHM	LAO	LKA	MAC
MAR	MEX	MYS	NLD	NPL	OMN
PAK	PAN	PHL	POL	RUS	SAU
SGP	THA	TUR	VNM	ZAF	TWN

Note: In BACI, Taiwan is recorded as "Other Asia" with country code S19.

C.3 Trend Estimation

I estimate product- and partner-specific trends in the pre-tariff period (2012–2017) with time centered at 2014.5.

Share trends (logit). For China and hub shares in U.S. imports and in ROW exports:

$$\mathrm{logit}(s_{k,i,t}) = \alpha_{k,i} + \beta_{k,i}(t-2014.5) + \varepsilon_{k,i,t}, \quad \widehat{s}_{k,i,t} = \mathrm{invlogit}(\widehat{\alpha}_{k,i} + \widehat{\beta}_{k,i}(t-2014.5)).$$

If estimation is unreliable (fewer than 3 non-zero observations), I use the pre-period mean $\bar{s}_{k,i,\mathrm{pre}}$.

Level trends (log-linear). For trade levels $v \in \{\ell_{k,i,t}, x_{k,i,t}\}$:

$$\log(\max(v,1)) = a_{k,i} + g_{k,i}(t - 2014.5) + \nu_{k,i,t}, \quad \widehat{v}_{k,i,t} = \exp(\widehat{a}_{k,i} + \widehat{g}_{k,i}(t - 2014.5)).$$

If estimation is unreliable, I use the pre-period mean $\bar{v}_{k,i,\text{pre}}$.

Surprise growth. Define above-trend growth, floored at zero:

$$\tilde{\ell}_{k,i,t} = \max\{\ell_{k,i,t} - \widehat{\ell}_{k,i,t}, 0\} \quad \text{and} \quad \tilde{x}_{k,i,t} = \max\{x_{k,i,t} - \widehat{x}_{k,i,t}, 0\}.$$

C.4 Screening Criteria

To be flagged as rerouting, a product-hub-year triplet (k, i, t) must simultaneously satisfy five conditions.

Screen 1: Tariff Exposure. Product k must have been exposed to U.S. tariff increases on China:

$$\tau^d_{k,2019} > \tau^d_{k,2017},$$

where $\tau^d_{k,t}$ is the U.S. ad valorem tariff on Chinese product k in year t.

Screen 2: U.S. Reallocation. China's share of U.S. imports must fall below its pre-trend while the hub's share rises above its pre-trend:

$$s^{\text{US}}_{k,\text{CHN},t} < \widehat{s}^{\text{US}}_{k,\text{CHN},t} \quad \text{and} \quad s^{\text{US}}_{k,i,t} > \widehat{s}^{\text{US}}_{k,i,t}.$$

Screen 3: Rest-of-World Specificity. China's ROW export share must remain stable or rise relative to the hub's. Specifically, China's above-trend gain must weakly exceed the hub's above-trend gain:

$$\left(s^{\text{ROW}}_{k,\text{CHN},t} - \widehat{s}^{\text{ROW}}_{k,\text{CHN},t}\right) \geq \gamma \left(s^{\text{ROW}}_{k,i,t} - \widehat{s}^{\text{ROW}}_{k,i,t}\right),$$

where $\gamma = 1$. This rules out cases where China is losing global competitiveness in product k.

Screen 4: Tariff Feasibility. The hub country's tariff on product k must be lower than the U.S. tariff on China:

$$\tau^{\text{hub}}_{k,i,t} < \tau^d_{k,t}.$$

Hub tariffs are MFN rates from Teti (2025). This ensures there is an economic incentive to reroute through hub i.

Screen 5: Quantitative Consistency. The China-to-hub flow must be large enough to plausibly supply the hub-to-U.S. flow. I impose two versions:

(a) **Conservative (growth-based):** Uses above-trend growth on both legs. The condition is:

$$\tilde{\ell}_{k,i,t} \geq \rho^* \max(\tilde{x}_{k,i,t}, x_{\text{hinge},\Delta}).$$

The flagged quantity is the bottleneck: $\min(\tilde{\ell}_{k,i,t}, \tilde{x}_{k,i,t})$.

(b) **Liberal (levels-based):** Uses trade levels directly. The conditions are:

$$\min(\ell_{k,i,t}, x_{k,i,t}) \geq x_{\text{hinge,levels}} \quad \text{and} \quad \ell_{k,i,t} \geq \phi^* x_{k,i,t}.$$

The flagged quantity is the bottleneck: $\min(\ell_{k,i,t}, x_{k,i,t})$.

C.5 Parameter Calibration

Table C.2 lists fixed parameters and calibrated thresholds.

Table C.2: Screening Parameters

Parameter	Value	Description
γ	1.0	ROW specificity multiplier (Screen 3).
$x_{\text{hinge},\Delta}$	0	Hinge for $\tilde{x}$ in growth rule (thousands of current USD).
$x_{\text{hinge,levels}}$	2000	Minimum size for $\min(\ell, x)$ in levels rule (thousands of current USD).
ρ^*	0	Growth-rule proportionality threshold (Screen 5a).
ϕ^*	0.325	Levels-rule proportionality threshold (Screen 5b).

The calibrated values ρ^* and ϕ^* are selected via placebo discipline. I search over a grid and choose the smallest values such that the implied exit share for never-treated products (those with $\tau^d_{k,2019} = \tau^d_{k,2017}$) does not exceed 1% in any pre-tariff year (2012–2017). This yields $\rho^* = 0$ (any positive above-trend growth on both legs qualifies) and $\phi^* = 0.325$ (the China-to-hub leg must be at least 32.5% of the hub-to-U.S. leg).

C.6 Aggregation and Capping

For each product k and year t, I aggregate flagged values across all hubs passing Screens 1–5:

$$\theta_{k,t}^{\text{raw}} = \sum_{i \in \mathcal{H}} \mathbf{1}[\text{Screens 1–5 pass for } (k, i, t)] \times q_{k,i,t},$$

where $q_{k,i,t}$ is the flagged quantity from Screen 5 (either $\min(\tilde{\ell}_{k,i,t}, \tilde{x}_{k,i,t})$ for the conservative version or $\min(\ell_{k,i,t}, x_{k,i,t})$ for the liberal version).

I then deflate to 2017 dollars and cap at the 2017 baseline:

$$\theta_{k,t} = \min\left(\frac{\theta_{k,t}^{\text{raw}} / P_t}{v_{\text{CHN}\to\text{USA},k,2017}}, 1\right),$$

where P_t is the BLS end-use import price index for all commodities, normalized so $P_{2017} = 1$.

C.7 Persistence Filter

To reduce noise from temporary fluctuations, I zero out single-year spikes. Flagged exit must persist for at least two consecutive years:

$$\theta_{k,t} \leftarrow \theta_{k,t} \times \mathbf{1}\left[\theta_{k,t-1}^{\text{raw}} > 0 \text{ or } \theta_{k,t+1}^{\text{raw}} > 0\right].$$

C.8 Tariff Data Sources

U.S. tariffs on China (τ^d) are from Amiti, Redding, and Weinstein (2020) for monthly HS10 rates through 2017, aggregated to annual HS6 using 2017 import shares as weights. From 2018–2023, I use the Global Tariff Database from Teti (2025) and Rodríguez-Clare, Ulate, and Vasquez (2025), which provides bilateral HS6 tariffs. Hub country tariffs (τ^{hub}) are MFN rates from the same source.

D Empirical Results

D.1 Results

Table D.1: Summary Statistics

	Pre		Post			N		
Variable	Mean	SD	Mean	SD	Post−Pre	$HS6_{\text{Pre}}$	$HS6_{\text{Post}}$	$HS6_{\cap}$
$\Delta_{k,t}$	0.025	0.066	0.121	0.062	0.096	4246	4246	4246
θ_{low}	0.014	0.092	0.045	0.148	0.031	4246	4246	4246
θ_{high}	0.037	0.160	0.079	0.212	0.042	4246	4246	4246
Share($\theta_{\text{low}} > 0$)	0.097	0.296	0.295	0.456	0.197	4246	4246	4246
Share($\theta_{\text{high}} > 0$)	0.110	0.313	0.263	0.440	0.153	4246	4246	4246
Consumption share (2017)	0.203	0.384	0.203	0.385	0.000	4246	4246	4246
Capital share (2017)	0.151	0.346	0.149	0.348	-0.002	4246	4246	4246
Intermediate share (2017)	0.646	0.462	0.648	0.463	0.003	4246	4246	4246

Notes: Variable construction is described in the main text.

Figure D.1: Interquartile Range of the Tariff Wedge by End Use

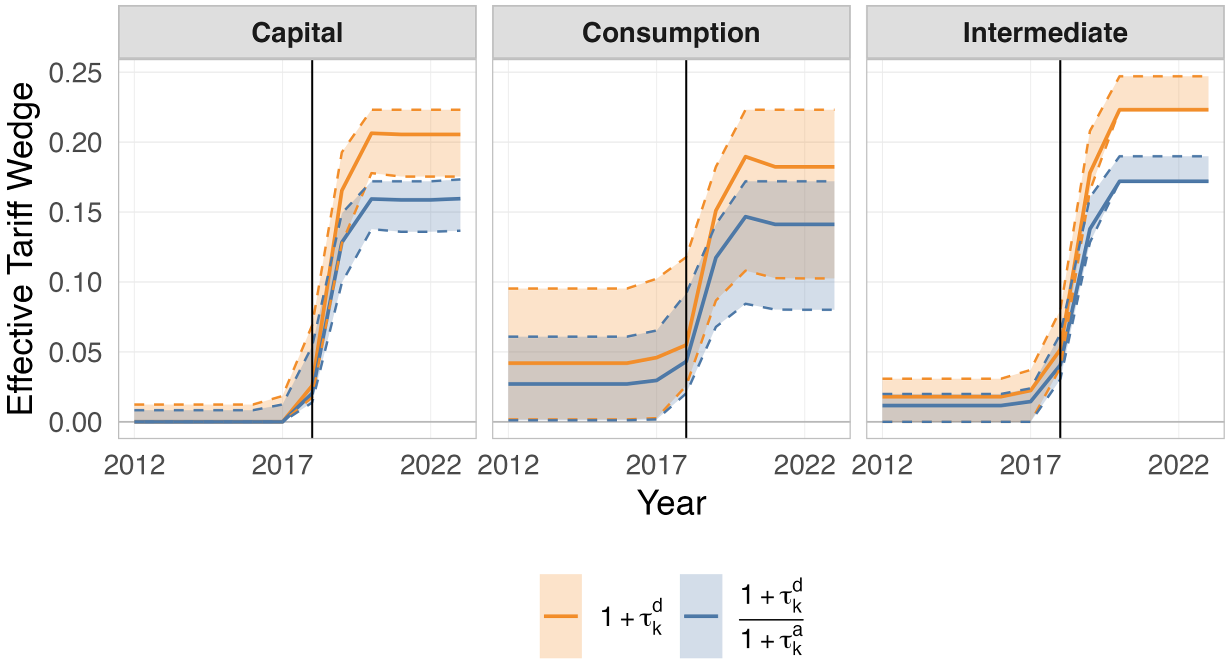

Note: Each panel displays the interquartile range of effective tariff rates for imported goods categorized by end use. The orange line plots the uncorrected effective tariff rate $1 + \tau^d$, while the blue line plots the tariff rate corrected for import deductibility: $(1 + \tau^d)/(1 + \tau^a)$. Since intermediates are fully tax-deductible, $\tau^a_{k,t} = \tau^c_t \times \tau^d_{k,t}$. Prior to the 2017 Tax Cuts and Jobs Act, $\tau^c_t = 0.35$; after TCJA $\tau^c_t = 0.21$. Capital goods are expensed over time. Denoting the present-value of such deductions as z_t, the effective tariff wedge for capital goods is $\tau^a_{k,t} = z_{k,t}\tau^c_{k,t} \times \tau^d_{k,t}$. After TCJA, $z = 1$. Prior to TCJA, z was slightly less than one for most imported capital goods. I obtain z by mapping the imported capital good HS6 codes into the corresponding IRS tax lives and calculating the present value of deductions with a discount rate of 0.06.

Figure D.2: Transshipment by End Use Category

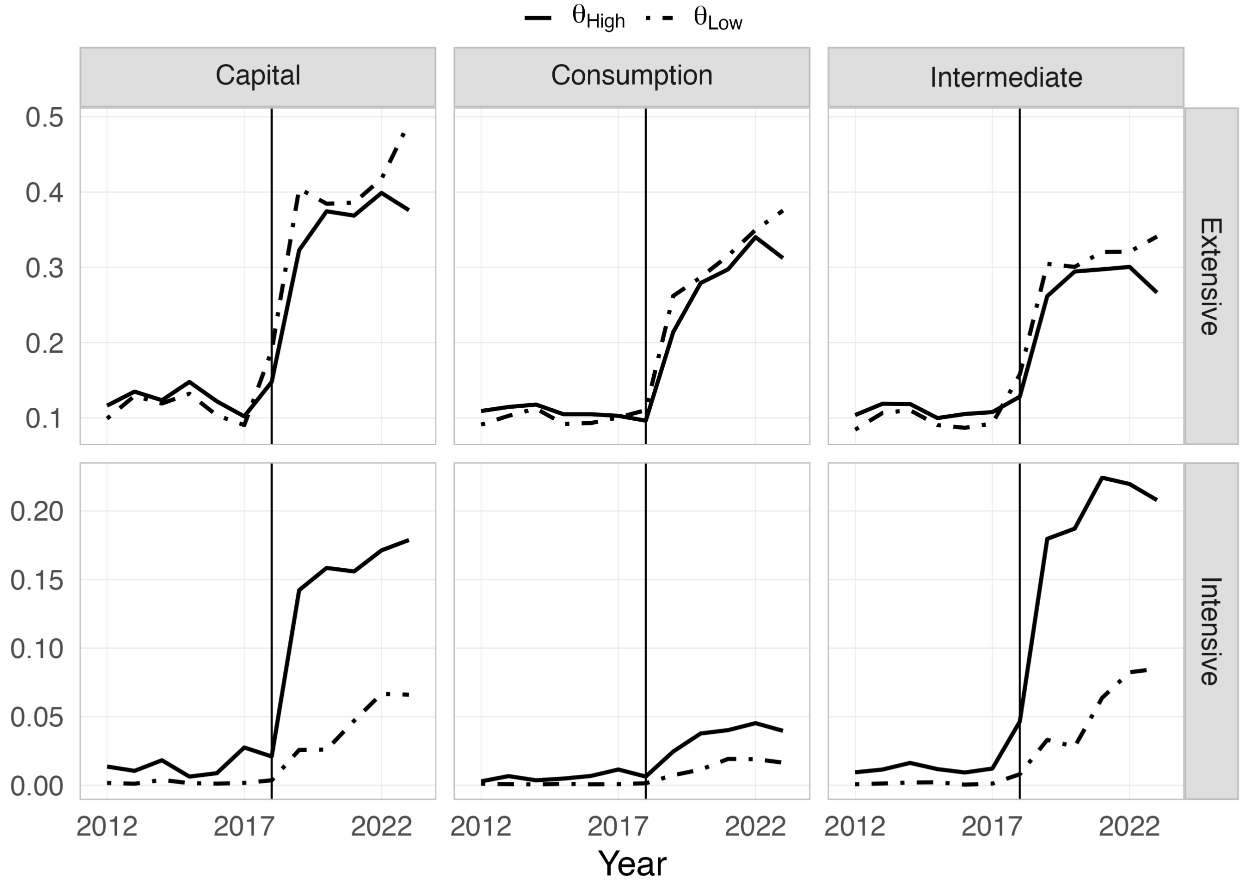

Note: The top panels plot the share of HS6 codes within each end use category which have detected transshipment. The bottom panels display a weighted average transshipment share by end use. Each HS6 code is weighted by its share of 2017 imports from China.

E Robustness

Table E.1: EU-27+UK and Canada Destination Placebo Results

EU-27+UK	$\sinh^{-1}(\theta_{\text{low}})$ (1)	(2)	$\sinh^{-1}(\theta_{\text{high}})$ (3)	(4)	$\mathbf{1}(\theta_{\text{low}} > 0)$ (5)	(6)	$\mathbf{1}(\theta_{\text{high}} > 0)$ (7)	(8)
$\Delta_{k,t}$	0.018*	0.021	0.030	0.034	0.110	0.134	0.047	0.106
	(0.010)	(0.014)	(0.020)	(0.026)	(0.158)	(0.121)	(0.158)	(0.104)
$\Delta_{k,t} \times \mathbf{1}(t \geq 2018)$		-0.003		-0.004		-0.027		-0.067
		(0.006)		(0.011)		(0.131)		(0.141)
Observations	54,564	54,564	54,564	54,564	54,564	54,564	54,564	54,564
R^2	0.29	0.29	0.32	0.32	0.33	0.33	0.37	0.37
Canada	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta_{k,t}$	0.059***	0.053*	0.058**	0.012	0.242	0.095	0.162	0.057
	(0.018)	(0.027)	(0.024)	(0.040)	(0.186)	(0.215)	(0.187)	(0.202)
$\Delta_{k,t} \times \mathbf{1}(t \geq 2018)$		0.006		0.053		0.170		0.121
		(0.017)		(0.035)		(0.323)		(0.318)
Observations	52,716	52,716	52,716	52,716	52,716	52,716	52,716	52,716
R^2	0.19	0.19	0.23	0.23	0.29	0.29	0.31	0.31
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Entries report two–way fixed–effects regressions of placebo diversion measures on the U.S. tariff wedge, where the wedge is the same as in the main specification and is constructed from statutory U.S. HS6 tariffs and the tax–shield (see text). Outcomes are diversion measures computed with (i) the EU–27+UK treated as a single destination (left block) and (ii) Canada as the destination (right block). For each destination and year, diversion is built from HS6 flows using the same screening steps as in the main analysis but with the destination replaced by EU–27+UK or Canada. For comparability with the U.S. baseline, the exposure gate uses the U.S. HS6 tariff panel and diversion is capped within HS6 by the 2017 China→U.S. base so that $\sum_i \theta_{k,t} \leq 1$. Columns (1)–(2) and (5)–(6) use inverse–hyperbolic–sine outcomes $\sinh^{-1}(\theta_{\text{low}})$ and $\sinh^{-1}(\theta_{\text{high}})$; columns (3)–(4) and (7)–(8) use indicators $\mathbf{1}\{\theta_{\text{low}} > 0\}$ and $\mathbf{1}\{\theta_{\text{high}} > 0\}$. All specifications include HS6 and year×use fixed effects with standard errors clustered by HS6. The bottom panel repeats the exercise for Canada.

Figure E.1: Permutation Null Distributions of β under a Broken China–Hub–US Network Mapping

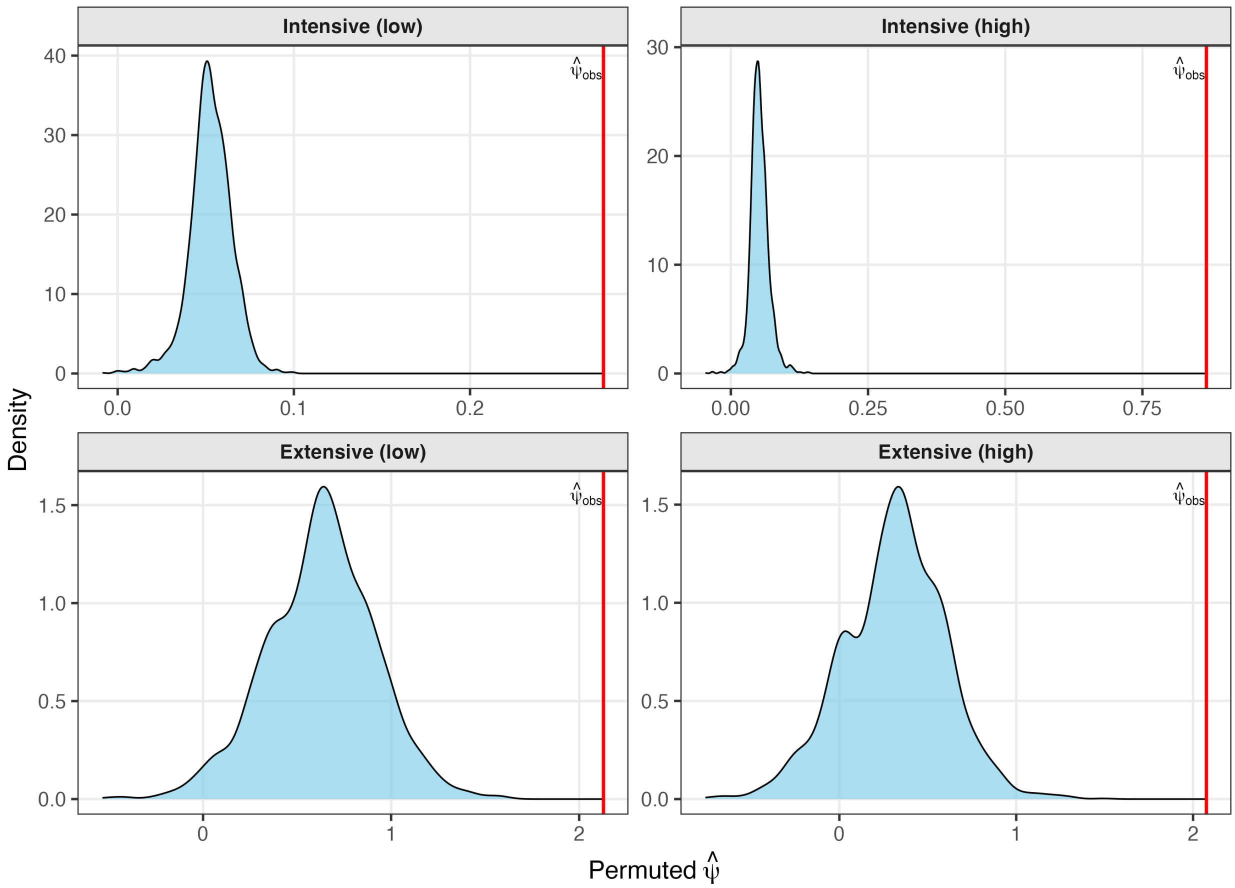

Note: This figure displays the permutation-based placebo distributions for the estimated elasticity of diversion with respect to the tax wedge, β, under a broken network mapping. For each replication, the pairing between China→hub and hub→US shipment legs is randomly deranged within HS4–year product groups while keeping the marginal leg volumes, pre-trend screens, and denominator structure fixed. Each permutation thus preserves the observed scale and composition of trade at the HS4 level but eliminates the structural correspondence between the two legs of the network. The distributions shown in blue correspond to the resulting estimates of β from 5,000 such random permutations, re-estimated separately for intensive and extensive outcomes. The vertical red lines indicate the coefficients obtained under the intact network mapping. The top panels use the inverse-hyperbolic-sine transformations of the diversion measures $\sinh^{-1}(\theta_{\text{low}})$ and $\sinh^{-1}(\theta_{\text{high}})$, while the bottom panels use binary indicators for positive diversion $1\{\theta_{\text{low}} > 0\}$ and $1\{\theta_{\text{high}} > 0\}$. All specifications include HS6 and year fixed effects and are weighted by 2017 China–US import shares.

Table E.2: Permutation Tests for Pre-Period Balance

Outcome	Observed Gap	P (left)	P (right)	P (two-sided)	P (Holm)
$\sinh^{-1}(\theta_{\text{high}})$	0.006	0.960	0.040	0.079	0.238
$\sinh^{-1}(\theta_{\text{low}})$	0.001	0.999	0.001	0.001	0.005
$\mathbf{1}\{\theta_{\text{high}} > 0\}$	-0.007	0.096	0.904	0.193	0.386
$\mathbf{1}\{\theta_{\text{low}} > 0\}$	0.001	0.132	0.868	0.263	0.386

Notes: Each row reports the observed weighted mean difference (treated minus placebo) in the specified outcome during the pre-period (2012–2017), along with permutation-based p-values. For each outcome, 5,000 block permutations were performed by reassigning treated HS6 codes within HS4 product groups while preserving the overall number of treated observations. For each permutation draw, the treated–placebo gap in the weighted mean outcome (weighted by 2017 China–US import shares) was recalculated. The empirical p-values are computed as the fraction of permuted gaps that are as or more extreme than the observed gap, with Holm–Bonferroni adjustments applied for multiple outcomes. All outcomes are defined at the HS6–year level and correspond to the conservative and liberal diversion measures in the main text, expressed as either inverse hyperbolic sine or binary indicators.

Supplemental Appendix

Table S1: Transshipment Response to Tariffs

	Intensive Margin				Extensive Margin			
	$\sinh^{-1}(\theta_{\text{low}})$		$\sinh^{-1}(\theta_{\text{high}})$		$\mathbf{1}(\theta_{\text{low}} > 0)$		$\mathbf{1}(\theta_{\text{high}} > 0)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta_{k,t}$	0.276***		0.867***		2.13***		2.08***	
	(0.045)		(0.135)		(0.533)		(0.528)	
$\Delta_{k,t} \times$ Cap		0.264***		0.897***		2.05***		1.99***
		(0.064)		(0.211)		(0.671)		(0.632)
$\Delta_{k,t} \times$ Con		0.158***		0.389***		1.36*		1.23
		(0.045)		(0.132)		(0.815)		(0.833)
$\Delta_{k,t} \times$ Int		0.502***		1.60***		3.60***		3.71***
		(0.085)		(0.182)		(0.405)		(0.474)
R^2	0.40	0.41	0.49	0.50	0.52	0.52	0.52	0.52
Observations	54,912	54,912	54,912	54,912	54,912	54,912	54,912	54,912
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the coefficients for regressions of (7). Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share multiplied by the HS6 code's share of imports from China to the United States, and standard errors are clustered by HS6.

Table S2: Transshipment Response to Tariffs (Clustered by HS4 × year)

	Intensive Margin				Extensive Margin			
	$\sinh^{-1}(\theta_{\text{low}})$		$\sinh^{-1}(\theta_{\text{high}})$		$\mathbf{1}(\theta_{\text{low}} > 0)$		$\mathbf{1}(\theta_{\text{high}} > 0)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta_{k,t}$	0.276***		0.867***		2.13***		2.08***	
	(0.065)		(0.144)		(0.576)		(0.589)	
$\Delta_{k,t}$ × Cap		0.264***		0.897***		2.05**		1.99**
		(0.077)		(0.223)		(0.702)		(0.676)
$\Delta_{k,t}$ × Con		0.158***		0.389**		1.36		1.23
		(0.048)		(0.135)		(0.766)		(0.815)
$\Delta_{k,t}$ × Int		0.502***		1.60***		3.60***		3.71***
		(0.125)		(0.183)		(0.502)		(0.532)
R^2	0.40	0.41	0.49	0.50	0.52	0.52	0.52	0.52
Observations	54,912	54,912	54,912	54,912	54,912	54,912	54,912	54,912
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the coefficients for regressions of (7). Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share multiplied by the HS6 code's share of imports from China to the United States, and standard errors are clustered by HS4 and year.

Table S3: Transshipment Response to Tariffs (unweighted)

	Intensive Margin				Extensive Margin			
	$\sinh^{-1}(\theta_{\text{low}})$		$\sinh^{-1}(\theta_{\text{high}})$		$\mathbf{1}(\theta_{\text{low}} > 0)$		$\mathbf{1}(\theta_{\text{high}} > 0)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta_{k,t}$	0.131***		0.202***		0.612***		0.529***	
	(0.019)		(0.030)		(0.071)		(0.067)	
$\Delta_{k,t} \times$ Cap		0.125***		0.145*		0.404***		0.295**
		(0.038)		(0.078)		(0.126)		(0.120)
$\Delta_{k,t} \times$ Con		0.044*		0.092***		0.408***		0.347***
		(0.025)		(0.028)		(0.097)		(0.092)
$\Delta_{k,t} \times$ Int		0.186***		0.295***		0.828***		0.742***
		(0.031)		(0.048)		(0.119)		(0.112)
R^2	0.23	0.23	0.27	0.27	0.31	0.31	0.32	0.32
Observations	54,912	54,912	54,912	54,912	54,912	54,912	54,912	54,912
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the coefficients for regressions of (7). Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share (since we are stacking by end-use, this is the same as unweighted), and standard errors are clustered by HS6.

Table S4: Transshipment Response to Tariffs (No Deductibility)

	$\sinh^{-1}(\theta_{\text{low}})$		$\sinh^{-1}(\theta_{\text{high}})$	
	(1)	(2)	(3)	(4)
$1+\tau^d_{k,t}$	0.128^{***}		0.430^{***}	
	(0.039)		(0.121)	
$1+\tau^d_{k,t} \times$ Cap		0.113^{**}		0.394^{**}
		(0.054)		(0.175)
$1+\tau^d_{k,t} \times$ Con		0.094^{***}		0.226^{***}
		(0.029)		(0.081)
$1+\tau^d_{k,t} \times$ Int		0.257^{***}		0.981^{***}
		(0.048)		(0.137)
R^2	0.39	0.40	0.48	0.49
Observations	54,912	54,912	54,912	54,912
HS6 & Year×Use FE	✓	✓	✓	✓
Weighted	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table repeats the intensive margin exercise from Table S1, but uses the tariff wedge gross of the deductibility penalty as a regressor. Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share multiplied by the HS6 code's share of imports from China to the United States, and standard errors are clustered by HS6.

Table S5: Transshipment Response to Tariffs (with HS4×Year clustering and HS6, year, and HS2×Year Fixed Effects)

	Intensive Margin				Extensive Margin			
	$\sinh^{-1}(\theta_{\text{low}})$		$\sinh^{-1}(\theta_{\text{high}})$		$\mathbf{1}(\theta_{\text{low}} > 0)$		$\mathbf{1}(\theta_{\text{high}} > 0)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta_{k,t}$	0.383***		1.26***		2.48***		2.42***	
	(0.087)		(0.166)		(0.673)		(0.689)	
$\Delta_{k,t}$ × Cap		0.327***		1.09***		2.57**		2.52**
		(0.079)		(0.196)		(0.890)		(0.913)
$\Delta_{k,t}$ × Con		0.266***		0.742***		1.80**		1.78**
		(0.069)		(0.177)		(0.631)		(0.663)
$\Delta_{k,t}$ × Int		0.458***		1.53***		2.56***		2.46***
		(0.107)		(0.199)		(0.577)		(0.596)
R^2	0.42	0.43	0.52	0.53	0.58	0.58	0.58	0.58
Observations	54,912	54,912	54,912	54,912	54,912	54,912	54,912	54,912
HS6 & Year & Year×HS2 FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the coefficients for regressions of (7). Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6, year, and year × HS2 fixed effects; weights are pre-period end-use share (since we are stacking by end-use, this is the same as unweighted), and standard errors are clustered by HS4 and year.

Table S6: Pre-Period Test

	Intensive Margin		Extensive Margin	
	$\sinh^{-1}(\theta_{\text{low}})$	$\sinh^{-1}(\theta_{\text{high}})$	$\mathbf{1}(\theta_{\text{low}} > 0)$	$\mathbf{1}(\theta_{\text{high}} > 0)$
	(1)	(2)	(3)	(4)
$\Delta_{k,t}$	0.281^{***}	0.882^{***}	2.07^{***}	1.98^{***}
	(0.040)	(0.124)	(0.576)	(0.580)
$\Delta_{k,t} \times \mathbf{1}(t < 2018)$	-0.052	-0.142	0.556	0.864
	(0.070)	(0.205)	(0.957)	(0.943)
R^2	0.40	0.49	0.52	0.52
Observations	54,912	54,912	54,912	54,912
HS6 & Year×Use FE	✓	✓	✓	✓
Weighted	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column repeats the main regression specification for both intensive and extensive margins, but controls for the pre-period with an indicator variable for pre-tariff years. Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share multiplied by the HS6 code's share of imports from China to the United States, and standard errors are clustered by HS6.

Table S7: Slope-adjusted difference-in-differences

	$\sinh^{-1}(\tilde{\theta}_{\text{low}})$	$\sinh^{-1}(\tilde{\theta}_{\text{high}})$	$\mathbf{1}(\tilde{\theta}_{\text{low}} > 0)$	$\mathbf{1}(\tilde{\theta}_{\text{high}} > 0)$
	(1)	(2)	(3)	(4)
$\Delta_{k,t}$	0.264^{***}	0.766^{***}	2.08^{***}	2.09^{***}
	(0.047)	(0.177)	(0.692)	(0.703)
R^2	0.40	0.47	0.52	0.51
Observations	54,912	54,912	54,912	54,912
Weighted	✓	✓	✓	✓
HS6 & Year×Use FE	✓	✓	✓	✓

* p < 0.1, ** p < 0.05, *** p < 0.01

Notes: Each outcome is first residualized by its HS6-specific linear trend estimated over 2012–2017, removing product-level pre-slopes before estimation. We then re-estimate the baseline specification with HS6 and year×end-use fixed effects, weighted by 2017 China→US import shares multiplied by pre-period end-use share and clustered by HS6. Coefficients therefore capture the effect of the tariff wedge on deviations from each product's pre-trend.

Table S8: Long-differenced Specification

	$\sinh^{-1}(\theta_{\text{Low}})$	$\sinh^{-1}(\theta_{\text{High}})$	$\mathbf{1}(\theta_{\text{low}} > 0)$	$\mathbf{1}(\theta_{\text{high}} > 0)$	$\mathbf{1}(\theta^{\text{Act}}_{\text{low}} > 0)$	$\mathbf{1}(\theta^{\text{Act}}_{\text{high}} > 0)$
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta_{k,\text{Post}} - \Delta_{k,\text{Pre}}$	0.224^{***}	0.684^{***}	1.48^{***}	1.50^{***}	1.53^{**}	1.48^{**}
	(0.050)	(0.135)	(0.474)	(0.474)	(0.607)	(0.591)
Observations	4,246	4,246	4,246	4,246	2,640	2,513
R^2	0.18	0.24	0.12	0.11	0.16	0.15
Weighted	✓	✓	✓	✓	✓	✓
Use FE	✓	✓	✓	✓	✓	✓

* p < 0.1, ** p < 0.05, *** p < 0.01

Notes: For each HS6 we compute pre (2012–2017) and post (2018–2023) means of the outcome and the effective wedge, form long differences, and regress the long-differenced outcomes on the long-differenced wedge with end-use fixed effects. All regressions are weighted by 2017 import shares and standard errors are clustered by HS6. Columns (5)-(6) estimate an "entry" of the extensive margin in which we restrict the sample to HS6 codes with zero evidence of pre-period transshipment.

Table S9: Baseline Specification with HS4 Trends

	$\sinh^{-1}(\theta_{\text{low}})$	$\sinh^{-1}(\theta_{\text{high}})$	$\mathbf{1}(\theta_{\text{low}} > 0)$	$\mathbf{1}(\theta_{\text{high}} > 0)$
	(1)	(2)	(3)	(4)
$\Delta_{k,t}$	0.183***	0.813***	2.85***	2.76***
	(0.030)	(0.132)	(0.819)	(0.862)
R^2	0.54	0.61	0.59	0.59
Observations	54,912	54,912	54,912	54,912
HS6 & Year×Use FE	✓	✓	✓	✓
Weighted	✓	✓	✓	✓
HS4 Trend	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table repeats the baseline specification but adds linear trends for HS4 codes. Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share multiplied by the HS6 code's share of imports from China to the United States, and standard errors are clustered by HS6.

Figure S1: Robustness to Anchoring Differently

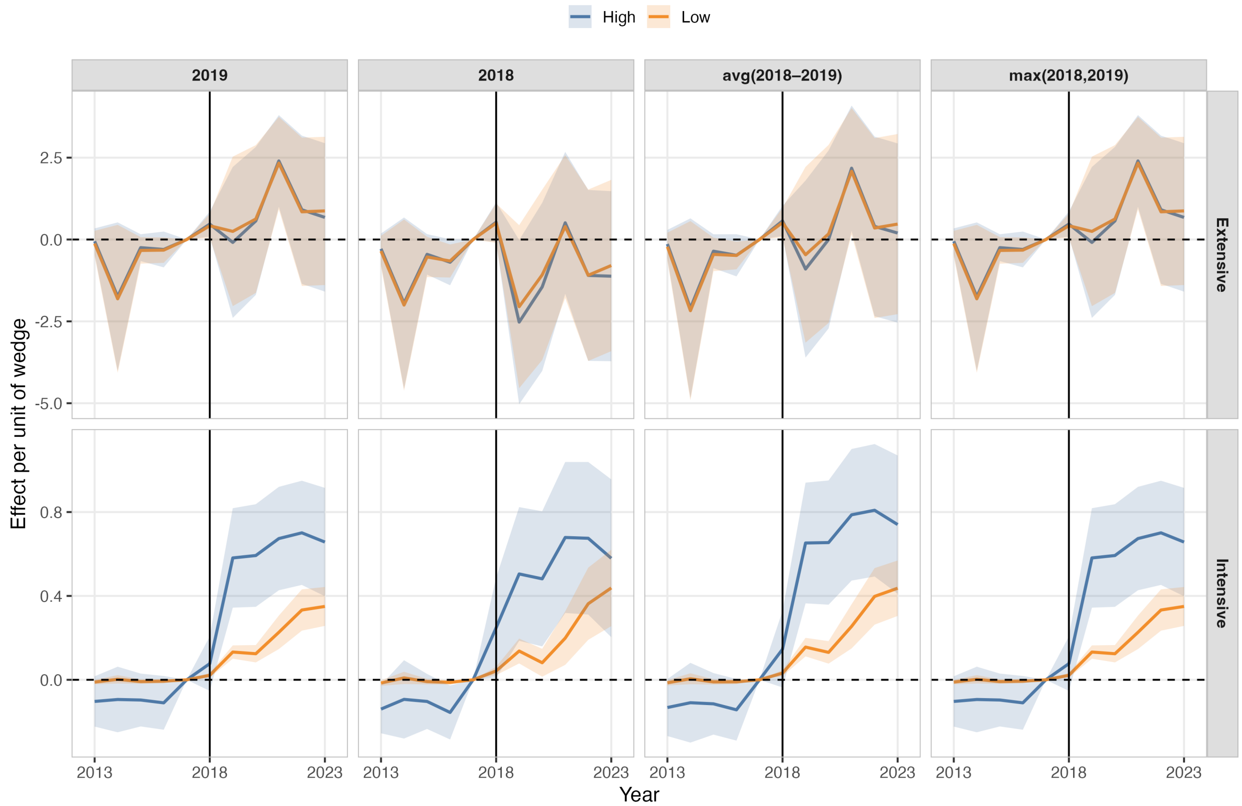

Note: This figure estimates the event–study under four alternative definitions of the exposure anchor used to scale coefficients: (i) the 2019 effective wedge (baseline), (ii) the 2018 wedge, (iii) the average of the 2018–2019 wedges, and (iv) the maximum of the two. The dynamic profiles are nearly identical across anchors, with flat pre–trends through 2017, a sharp rise in 2018–2019, and persistent positive coefficients thereafter. The only exception is the *extensive* margin when exposure is anchored to 2018, where the coefficients appear attenuated or slightly negative immediately after 2018.

This discrepancy is mechanical. The 2018 wedge captures only the partial–year effect of tariff exposure, since most Section 301 tariffs were announced in the spring and implemented between July and September 2018. The exposure variable therefore understates the full shock in that year, while the dependent variable—the share of HS6 products exhibiting positive transshipment—responds to the cumulative policy change by the end of the year. In other words, the denominator (the wedge) is too small relative to the behavioral response, biasing the per–unit effect downward. When exposure is defined using the 2019 wedge, or alternatively using the average or maximum of the 2018–2019 wedges to account for the phase–in of tariffs, the extensive–margin coefficients align closely with the baseline pattern. The intensive–margin results are virtually unaffected across all anchors. Overall, the event–study evidence is robust to how the exposure anchor is defined, and the muted 2018–anchor response is fully consistent with the mid–year timing of tariff implementation and the partial measurement of exposure in that year.

Table S10: Transshipment Response to Tariffs (2016-2019)

	Intensive Margin				Extensive Margin			
	$\sinh^{-1}(\theta_{\text{low}})$		$\sinh^{-1}(\theta_{\text{high}})$		$\mathbf{1}(\theta_{\text{low}} > 0)$		$\mathbf{1}(\theta_{\text{high}} > 0)$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\Delta_{k,t}$	0.168^{***}		0.854^{***}		1.63		1.40	
	(0.028)		(0.165)		(1.08)		(1.10)	
$\Delta_{k,t} \times$ Cap		0.157^{***}		0.822^{***}		1.32		1.15
		(0.043)		(0.260)		(1.31)		(1.32)
$\Delta_{k,t} \times$ Con		0.094^{***}		0.264^{**}		0.669		0.302
		(0.024)		(0.125)		(1.41)		(1.45)
$\Delta_{k,t} \times$ Int		0.311^{***}		1.85^{***}		3.87^{***}		3.68^{***}
		(0.046)		(0.245)		(0.480)		(0.490)
R^2	0.35	0.36	0.44	0.45	0.49	0.49	0.47	0.47
Observations	22,880	22,880	22,880	22,880	22,880	22,880	22,880	22,880
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓	✓	✓
Weighted	✓	✓	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: This table reports the coefficients for regressions of (7). Each HS6–year observation is fractionally allocated across end-uses according to its UN-BEC shares. We estimate on the stacked HS6–year–end-use panel with HS6 and year × end-use fixed effects; weights are pre-period end-use share multiplied by the HS6 code's share of imports from China to the United States, and standard errors are clustered by HS6.

Figure S2: Aggregate Transshipment Comparing Treated to Untreated Codes

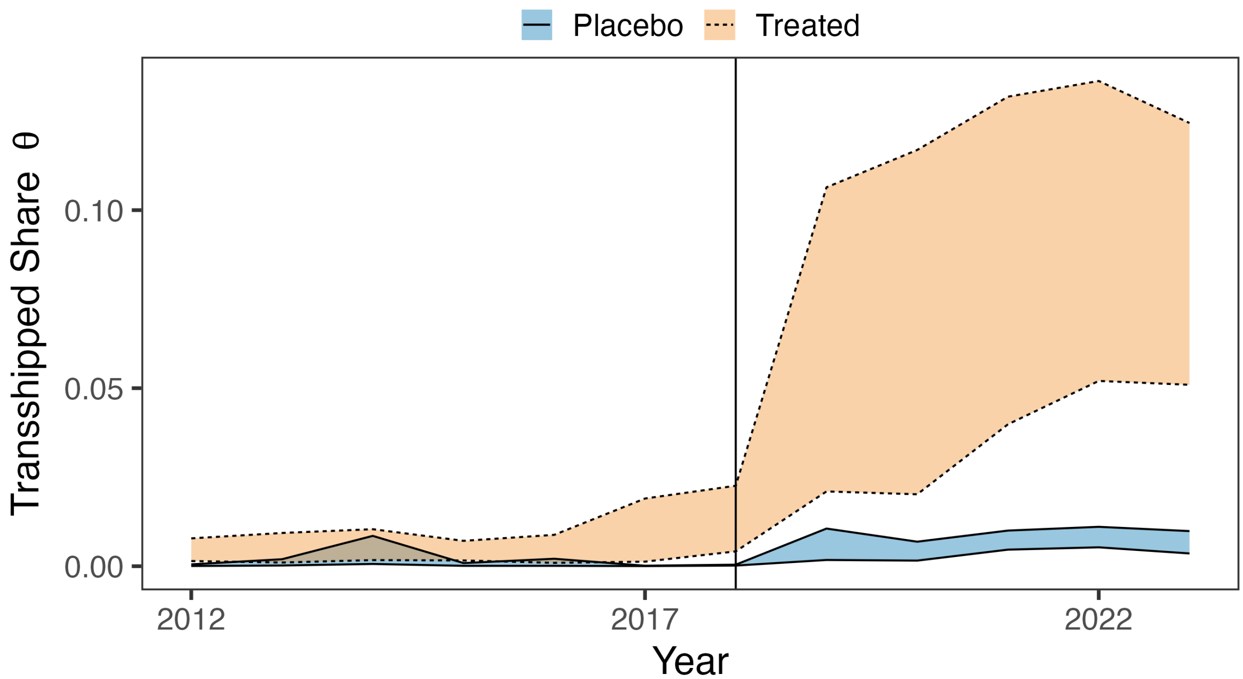

Notes: This shows the intensive margin of transshipment for both treated and untreated HS6 codes following the construction as described in the text.

Table S11: The Entropy-Weighted Effect of Tariffs on Transshipment

	Intensive Margin		Extensive Margin	
	$\sinh^{-1}(\theta_{\text{low}})$	$\sinh^{-1}(\theta_{\text{high}})$	$\mathbf{1}(\theta_{\text{low}} > 0)$	$\mathbf{1}(\theta_{\text{high}} > 0)$
	(1)	(2)	(3)	(4)
$\Delta_{k,t}$	0.275***	0.862***	2.12***	2.06***
	(0.045)	(0.135)	(0.532)	(0.527)
Observations	54,912	54,912	54,912	54,912
R^2	0.40	0.49	0.52	0.52
HS6 & Year×Use FE	✓	✓	✓	✓
Weighted	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: We re-weight HS6 units via entropy balancing so that treated (top quartile of pre-period wedge) and control have identical pre-period moments of diversion (winsorized and standardized). The panel regressions are re-estimated with weights = 2017 exposure × EB weights × pre-period end-use share, HS6 and year×end-use fixed effects, and HS4-clustered SEs. Coefficients are per unit of the wedge. Estimates are effectively unchanged, indicating the results are not driven by pre-period imbalances.

Table S12: Heterogeneous Fixed Costs Along the Intensive Margin

	$\sinh^{-1}(\theta_{\text{low}})$			$\sinh^{-1}(\theta_{\text{high}})$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta_{k,t}$	0.266***	0.266***	0.208***	0.803***	0.802***	0.614***
	(0.044)	(0.043)	(0.035)	(0.134)	(0.132)	(0.101)
$\Delta_{k,t}$ × Pre-2017 Hub Intensity	0.059			0.374**		
	(0.052)			(0.150)		
$\Delta_{k,t}$ × Bulkiness		-0.045			-0.304**	
		(0.030)			(0.151)	
$\Delta_{k,t}$ × Electronics			0.089***			0.331***
			(0.028)			(0.083)
R^2	0.40	0.40	0.41	0.50	0.49	0.51
Observations	54,912	54,912	54,912	54,912	54,912	54,912
Weighted	✓	✓	✓	✓	✓	✓
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: Each column reports regressions of $\sinh^{-1}(\theta_{k,t})$ on the effective wedge and its interaction with a pre-2017 product characteristic. Hub intensity is the standardized share of China's pre-2017 exports of product k routed through third-country hubs. "Bulkiness" is kilograms per U.S. dollar of import value. "Electronics" is an indicator for HS2 codes 84, 85, or 90. HS6 and year×use fixed effects are included; observations are weighted by 2017 China→US exposure multiplied by pre-period end-use share; standard errors are clustered by HS6. Coefficients are per unit of the wedge (log gross factor). The interaction terms measure how the wedge effect varies across products with different pre-existing hub exposure, bulkiness, or electronics content.

Table S13: Heterogeneous Fixed Costs Along the Extensive Margin

	$\mathbf{1}(\theta_{\text{low}} > 0)$			$\mathbf{1}(\theta_{\text{high}} > 0)$		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta_{k,t}$	2.07***	2.14***	1.92***	1.99***	2.04***	1.83***
	(0.533)	(0.522)	(0.541)	(0.530)	(0.518)	(0.539)
$\Delta_{k,t}$ × Pre-2017 Hub Intensity	0.345			0.493		
	(0.376)			(0.405)		
$\Delta_{k,t}$ × Bulkiness		0.043			-0.188	
		(0.260)			(0.296)	
$\Delta_{k,t}$ × Electronics			0.280			0.317
			(0.219)			(0.238)
R^2	0.52	0.52	0.52	0.52	0.52	0.52
Observations	54,912	54,912	54,912	54,912	54,912	54,912
Weighted	✓	✓	✓	✓	✓	✓
HS6 & Year×Use FE	✓	✓	✓	✓	✓	✓

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: The procedure is the same as in Table S12, except we now have the extensive margin as the dependent variable.

Table S14: Percent difference between baseline and leave-one-out: $100 \times \frac{\beta_{LOO}-\beta}{\beta}$

Country Code	Extensive (High)	Extensive (Low)	Intensive (High)	Intensive (Low)
CAN	-17.741	-17.074	17.887	19.343
JPN	-3.000	-4.686	30.124	29.223
MEX	4.916	4.173	31.988	25.131
DEU	-8.019	-7.489	23.321	13.826
SGP	-17.408	-16.486	2.176	3.188
TWN	-10.480	-11.643	3.447	5.798
MYS	-1.852	-2.032	3.420	9.883
VNM	-45.038	-44.523	0.136	4.061
HKG	-13.705	-12.496	0.476	1.888
IDN	-1.840	-1.933	1.845	5.406
IND	-0.433	0.745	7.272	8.685
THA	-1.719	-0.346	3.393	7.414
KHM	-1.019	1.706	1.225	6.550
POL	-1.388	-0.611	2.198	2.402
TUR	1.561	1.202	1.371	1.862

Notes: For each regression specification, I leave out one hub and recompute the coefficient. After that, I compute the percent difference between the leave-one-out estimate and the baseline regression. For example, the leave-one-out coefficient when excluding Canada from the sample is 17 percent smaller than the baseline intensive margin coefficient from Column 1 of Table S1. The table shows the most influential fifteen hubs.

Figure S3: Local Projections of Transshipment from Tariffs

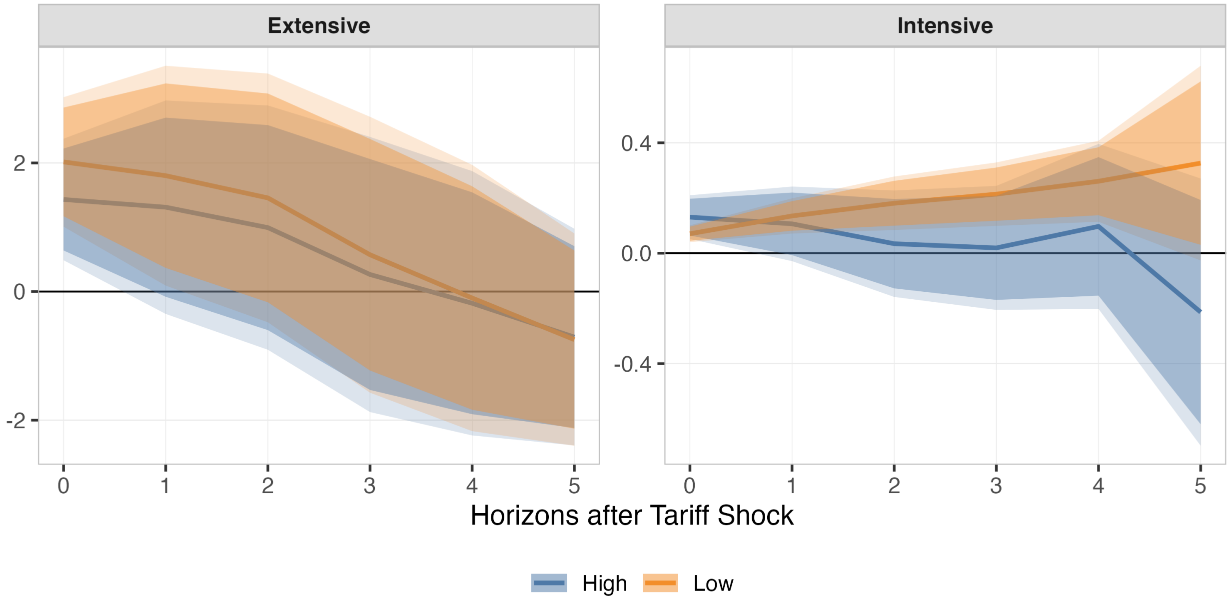

Note: This figure estimates a local projection of the form

$$f(\theta_{k,t+h}) - f(\theta_{k,t-1}) = \alpha_i + \delta_{t\times e(k)} + \psi_{t+h}\Delta_{k,t} + \varepsilon_{k,t+h}.$$

I do this for both the extensive margin, using the linear probability model from the main text, as well as an intensive margin with the inverse hyperbolic sine transformation. Standard errors are clustered by HS6 code.