The Contribution of Imports to Domestic Prices*

Omar Barbiero — FRB Boston
Alvaro Silva — FRB Boston
Hillary Stein — FRB Boston

June 30, 2026

Preliminary Draft – Comments Welcome

Abstract

We develop a multisector open-economy framework for measuring how border price shocks pass through to US producer and consumer prices. The model nests domestic input-output propagation, pass-through in levels, oligopolistic competition with foreign suppliers, and within-industry heterogeneity in exposure and pass-through. We show how to map the model to public US data using BEA input-output accounts and Census tariff data. We apply the framework to the 2018 and 2025 tariff episodes, constructing model-implied tariff exposures for PPI industries and PCE expenditure categories and comparing them with observed price changes using local-projection difference-in-differences regressions. The baseline production-network model explains substantial cross-sectional variation in downstream price responses, but its performance differs across episodes and price indices. We discuss how different wedges can explain the different effects of tariff episodes.

Keywords: Tariffs, inflation, import prices, indirect imports
JEL Codes: F40, E65, E31

*We thank Philippe Andrade, Anthoine Berthou (discussant), Şebnem Kalemli-Özcan, Robert Minton (discussant), Giovanni Olivei, Mariano Somale (discussant), Jenny Tang, Egon Zakrajsek, ASSA 2026, BIS, BEAR Conference 2026, and SCIEA 2026 seminar participants for comments and suggestions. Caleb Fitzpatrick provided outstanding research assistance. This paper was previously circulated under the title "The Contribution of Imports to PCE Inflation." The views expressed here are solely those of the authors and do not reflect the opinions of the Federal Reserve Bank of Boston, or the Federal Reserve System.
Omar.Barbiero@bos.frb.org, asilvub@gmail.com, Hillary.Stein@bos.frb.org

1 Introduction

How do border price shocks ripple through the US economy? Suppose a tariff is imposed on a product from a country at the border; its downstream effect depends on how imported goods enter domestic production, how cost shocks move through input-output linkages, how retailers will react to higher finished goods prices, and how firms set prices when they compete with foreign suppliers. Because of the interaction between these forces, the economic literature has struggled to reach a consensus on the extent of the effects of foreign price pass-through.

We build a multisector open-economy model with oligopolistically competitive retailers and intermediate-goods producers. The pricing block is based on a task-assignment problem with discrete firm choice, which delivers pass-through in levels and allows us to compare it with full pass-through in logs. Firms and retailers choose prices under Bertrand competition, thus allowing foreign prices to affect domestic companies both via marginal costs and via their reference price. Firms and retailers are also heterogeneous in terms of marginal cost exposure to foreign price shocks, thus creating incomplete pass-through according to the interaction of a firm's exposure to the shock and its pass-through. We derive a first-order mapping from foreign price shocks to domestic producer prices that depends on three objects: a markup/pass-through term, an oligopolistic competition term, and a covariance term capturing within-industry heterogeneity in exposure and pass-through.

We show how to take this model to official, publicly available US data. The framework maps product-country tariff shocks into BEA commodities, BEA commodities into domestic producer industries, and domestic producer prices into PCE expenditure categories. The central empirical objects are sensitivity matrices that measure how much each producer or consumer price category should respond to a change in the border price of each imported commodity. These sensitivities separate direct exposure, through imported goods purchased by consumers, from indirect exposure, through imported inputs used by domestic producers and propagated through the input-output network. We also provide a mapping on how oligopolistic forces or firm exposure heterogeneity should plausibly affect such mapping. This gives a transparent accounting of how a dollar of foreign price pressure can become domestic inflation.

We apply the framework to the 2018 and 2025 tariff episodes and compare model-implied exposure with observed PPI and PCE price changes using local-projection difference-in-differences regressions. We consider this exercise both a a pass-through estimate and a diagnostic tool. By starting from the simplest observable version of the model and then adding richer channels, we can ask which assumptions are needed to match the data, which ones are weakly identified, and which ones behave differently across episodes.

This approach helps organize a literature that has produced relatively consistent evidence on border pass-through but more dispersed estimates of pass-through into domestic producer and consumer prices. Micro-based studies of the 2018 tariffs generally find that US importers faced higher border prices but that this increase did not result in meaningful consumer price inflation. Other studies exploiting industry data on PPI or PCE find instead very large pass-through estimates (Amiti, Redding and Weinstein, 2019; Flaaen and Pierce, 2024; Minton and Somale, 2025). Some of the disagreement across studies can be reconciled by different assumptions and by conditioning on different channels. This is our attempt to spell out some canonical channels that can generate pass-through heterogeneity in the context of industry-level pass-through estimation of foreign price shocks.

The first channel in our model concerns the markup assumption. Under pass-through in levels, a cost increase raises prices dollar-for-dollar, so the percentage price response is attenuated by the share of costs in revenues. Under full pass-through in logs, firms maintain percentage markups. This makes the implied price response larger for two reasons. First, when price includes a positive markup over cost, maintaining a percentage markup mechanically scales up a given cost increase. Second, in a production network, this scaling can compound as upstream cost increases are passed downstream through successive stages of production. The same tariff can therefore imply a larger consumer-price effect under full pass-through in logs than under pass-through in levels, especially in high-margin sectors and in categories with large distribution margins.

The second channel is oligopolistic competition with foreign suppliers. If domestic firms compete with imported varieties, a tariff on foreign products can raise domestic prices even when domestic marginal costs do not change. In the model, this force is governed by an industry-level pass-through parameter. When oligopolistic forces are week, domestic prices move only with marginal costs. When an industry has strong oligopolistic forces, domestic prices also respond to the industry reference price, which includes foreign competitors. This channel captures the protectionist effect emphasized in some of the tariff literature: tariffs may raise domestic prices not only because domestic firms pay more for imported inputs, but also because foreign competitors become more expensive.

The third channel is within-industry exposure heterogeneity. Input-output tables measure average industry exposure, but the firms exposed to a tariff need not be representative of the firms setting the relevant domestic prices. If the firms most exposed to a tariff also have high pass-through, the industry price response will be larger than one would predict from average exposure alone. If they have low pass-through, the response will be smaller. We summarize this force with a covariance term, which captures the relationship between firm-level exposure and firm-level pass-through within an industry. This channel may be especially relevant when tariff shocks are concentrated by source country, as in the 2018-19 tariffs, because firms in the same industry may differ substantially in where they source their inputs. Viewed through this lens, the paper brings the firm-level markup-adjustment logic of Amiti, Itskhoki and Konings (2019) and generalizes it in the case of a tariff-shock with full border pass-through and input-output network. In this case foreign price pressure can reach domestic prices through imported-input costs, through competition with foreign suppliers, and through the within-industry allocation of those exposures across firms with different pass-through elasticities.

We bring the framework to the data in two steps. First, we construct model-implied tariff exposures under the observable part of the model: direct and indirect exposure under pass-through in levels and full pass-through in logs, with no oligopolistic or covariance wedge. We use effective tariff rates, measured from duties paid relative to import values, and aggregate them from product-country cells to BEA commodities using fixed import shares. This gives a predicted tariff shock for each PCE category and each PPI industry. Second, we compare these predicted shocks with observed price changes using local-projection difference-in-differences regressions, separately for the 2018 and 2025 tariff episodes.

The baseline results show that the model captures an important part of the cross-sectional incidence of tariffs, but that its fit differs across episodes and price systems. For PCE prices, the estimated response to the 2018 tariffs rises toward full pass-through and exceeds one at longer horizons, suggesting that the baseline model somewhat understates the realized consumer-price response. For the 2025 tariffs, the same type of exposure measure predicts the direction of price changes but overstates their magnitude, with longer-horizon coefficients below one. The data do not sharply distinguish pass-through in levels from full pass-through in logs: the log specification fits somewhat better in 2018, while the level specification fits somewhat better in 2025, but the difference is not statistically significant.

The PCE estimates are driven primarily by the direct component of exposure, meaning the expenditure component driven by finished goods imports that are directly consumed. This does not mean that indirect import exposure via intermediate inputs is economically unimportant. Rather, the indirect component is propagated through the production network and then aggregated into broad NIPA expenditure categories, which compresses the cross-sectional variation available for identification. Producer-price data provide a cleaner setting for evaluating the indirect channel because PPI industries are closer to the production network and are measured at a less aggregated level. In the PPI regressions, the model-implied marginal-cost exposure performs well, especially in 2025, where longer-horizon pass-through estimates are close to one. The 2018 PPI response is larger than the baseline model predicts, leaving a residual wedge concentrated in that episode.

We then ask which model extension can account for this residual wedge. A revealed-oligopoly exercise can fit the 2018 PPI response by increasing oligopoly power, thereby increasing the role of foreign competitor prices. But the same channel has little role in 2025, where the baseline production-network model already fits producer prices relatively well. This cross-episode instability is informative. If foreign competition is interpreted as a stable feature of market structure, it should not be important in one tariff episode and nearly irrelevant in another. We believe the covariance channel offers a more natural interpretation since it is a more flexible channel that hinges more on firm exposure heterogeneity rather than market structure per se. For instance, the 2018 tariffs were concentrated on one source country, so within-industry differences in sourcing could generate a large covariance between exposure and pass-through. The 2025 tariffs were broader across source countries, so firms that differ mainly in sourcing geography may have been affected more similarly. In that case, the same industry can display a large covariance wedge in 2018 and a smaller one in 2025 without requiring a change in its underlying competitive structure.

2 Literature

Our paper contributes to three strands of research.

The first is the empirical literature on tariff pass-through. A large body of work studies the 2018–19 U.S. tariff episode. Amiti, Redding and Weinstein (2019), Fajgelbaum et al. (2020), and Cavallo et al. (2021) find that the tariffs were largely passed through to U.S. import prices at the border, with limited absorption by foreign exporters. A related set of papers studies how these border-price increases affected downstream domestic prices. Cavallo et al. (2021) trace tariff effects into retail prices for selected product categories, while **?** use barcode-level scanner data matched to country-of-origin and tariff information to study retail pass-through, product replacement, and shrinkflation during the 2018–19 tariff episode. Flaaen and Pierce (2024) estimate pass-through to producer prices in manufacturing. More recent work studies the 2025 tariff episode. Minton and Somale (2025) and Minton, Ray and Somale (2026) develop real-time methods for detecting tariff effects in consumer prices using public data. Cavallo, Llamas and Vazquez (2025) use high-frequency retail microdata matched to product-level tariff rates and countries of origin to measure the short-run retail-price effects of the 2025 tariffs. Hacıoğlu-Hoke, Malladi and Feler (2026) document the gradual retail-price response using item-level spending data and information on countries of production. Amiti et al. (2026) study the incidence of the 2025 tariffs at the border, and Fajgelbaum and Khandelwal (2026) and Gopinath and Neiman (2026) analyze the broader short-run effects and incidence of the recent tariff increases.

Relative to this literature, our contribution is to provide a framework that clarifies which downstream price response is being estimated. A tariff can affect consumer prices through imported final goods, through imported intermediate inputs, through domestic input-output propagation, and through the pricing response of domestic competitors. Empirical designs that include different subsets of these channels need not estimate the same object. Our framework maps tariff shocks from product-country cells to producer industries and PCE expenditure categories, separates direct from indirect exposure, and then asks which additional pricing channels are needed to reconcile model-implied exposure with observed PPI and PCE price changes.

The second strand is the literature on production networks, inflation, and shock propagation. Baqaee and Farhi (2019) and Baqaee and Rubbo (2023) develop general-equilibrium frameworks for studying how shocks propagate through input-output linkages and affect aggregate outcomes. La'O and Tahbaz-Salehi (2022) study optimal monetary policy in production networks. Related work emphasizes that supply-chain shocks can affect inflation with delays and through heterogeneous downstream exposure; for example, Minton and Wheaton (2023) document delayed inflation in supply chains. Our paper builds on the first-order network approximation logic in this literature, but focuses on a partial-equilibrium object that is directly useful for tariff and import-price scenarios: the mapping from a vector of border price shocks to predicted PPI and PCE price changes, accounting for oligopolistic competition and heterogeneous firm exposure. We also show how to construct this mapping with public U.S. national-accounts data, including BEA input-output tables, import-use matrices, the Make table, and the PCE bridge.

The third strand is the literature on markup adjustment, imperfect competition, and pass-through. Classic contributions such as Feenstra (1989) and Goldberg and Knetter (1997) show that pass-through is generally incomplete under imperfect competition. Atkeson and Burstein (2008) emphasize strategic complementarities in oligopolistic markets, where firms' price responses depend on the prices of their competitors. More recently, Sangani (2026) argues that much of the evidence that appears to be incomplete pass-through in percentages is consistent with complete pass-through in levels. This distinction is central for our paper. Under pass-through in levels, a cost increase raises prices dollar-for-dollar, so percentage pass-through is attenuated by the cost share of price. We find that accounting for pass-through in level is not always enough to achieve full pass-through in this context.

We extend this pricing literature in two ways. First, we embed pass-through in levels and full pass-through in logs in a common multisector open-economy environment with input-output linkages and retailers. Second, we add two channels that are difficult to study in a purely aggregate network model: oligopolistic competition with foreign suppliers and within-industry covariance between firm-level exposure and firm-level pass-through. The former captures the protectionist effect of tariffs on domestic competitors' prices; the latter captures the possibility that industry-average input-output exposure is not representative of the firms that actually drive the price index. In this sense, the paper connects the theoretical literature on pass-through and markup dynamics to an empirical measurement exercise at the PPI and PCE levels.

The closest paper in this strand is Amiti, Itskhoki and Konings (2019), who show that aggregate pass-through depends not only on average exposure to international shocks, but also on the covariance between firm-level exposure and markup adjustment. Our framework builds on this insight, but generalizes the framework and changes the object of measurement. Rather than estimating exchange-rate pass-through into a sectoral price index, we condition on post-border foreign price changes and trace their effect through domestic producer and consumer prices. This distinction separates two forces that are combined in a sectoral exchange-rate pass-through exercise. Foreign price shocks can affect domestic firms through imported-input, and they can also affect domestic firms by raising the prices of foreign competitors. The latter is the protection channel. We then embed both forces in an input-output system and allow the domestic price response to depend on the covariance between domestic firms' exposure to the relevant shock and their pass-through.

Finally, our paper is closely related to Barbiero and Stein (2025), which provides a policy-oriented methodology for translating import-price shocks into PCE inflation using public data. The present paper substantially extends that work. We develop a formal model that nests pass-through in levels and logs, domestic input-output propagation, oligopolistic competition with foreign suppliers, and within-industry exposure heterogeneity. We apply the framework to both the 2018 and 2025 tariff episodes, compare PPI and PCE evidence, and use the model as a diagnostic tool to evaluate which channels can account for the residual gap between predicted tariff exposure and observed price responses.

3 Model

Notation. There is a set $\mathcal{I}$ of industries. There is also a set $\mathcal{E}$ of expenditure categories.

We use i and j, and occasionally k, to index industries. We use e to index expenditure categories.

Each industry i contains a discrete number of firms $\mathcal{N}_i$. The number of firms comprises a set of domestic firms $\mathcal{D}_i$ and a single foreign firm F_i, so that $|\mathcal{N}_i| = |\mathcal{D}_i| + 1$. We use f, g, h to index firms.

Each expenditure category e also contains a discrete number of firms $\mathcal{N}_e$. In contrast to industries, all firms within expenditure categories are domestic. As in the industry case, we use f, g, h to index firms within expenditure categories as well.

3.1 Intermediate Input Producers

3.1.1 Procurement

To generate pass-through in levels in intermediate-input markets, we model a procurement assignment problem for an industry entity that requires a continuum of tasks to be completed by a discrete number of firms to produce an aggregate good. This aggregated good serves as a reference for the oligopolistic competition setup that follows. Our setup thus belongs to the class of *shift-invariance* models analyzed in Sangani (2026). Here, we provide a microfoundation for the possibility of incomplete pass-through in levels within intermediate-input markets.

The problem of the industry entity is to choose what tasks to assign to firms. Each task is assigned to a single firm, but firms may complete more than one task. Thus, for each task, this is a discrete choice problem of choosing the firm that produces it.

A well-known result in the discrete choice literature is that this problem can be written either as a standard discrete choice maximization problem or as an expected minimization problem with additional entropy-regularized cost. To ease exposition, we proceed with the second route. For additional details on the exact mathematical foundation of this duality, see Section 1.3 in Galichon (2026).

The industry entity i chooses task shares to allocate to each firm f (π_{fi}) that minimize an expected cost problem with an additional entropy-regularized procurement cost as follows:

$$p_i^{eff} = \min_{\pi_{fi} \in \Delta_i} \left\{ b_i + \sum_f \pi_{fi}(p_{fi} - z_{fi}) + \sigma_i \sum_f \pi_{fi} \log \pi_{fi} \right\}, \tag{1}$$

where $\Delta_i = \{\pi_{fi} : \sum_f \pi_{fi} = 1\}$ is the simplex set.

p_i^{eff} represents the *effective price index* of industry i. b_i captures a baseline cost of satisfying an industry-i task. An example is a common physical or administrative cost of completing an industry i requirement: handling, distribution, assembly, etc. It is common across all alternatives. Its sole role is to pin down the level of the effective price index, as defined below, and ensure that it is positive. In contrast, z_{fi} is a firm-specific *deterministic* advantage in serving industry-i tasks, measured in money units. We can think of it as quality, reliability, and delivery speed, among others. σ_i represents the elasticity of substitution across tasks/firms within an industry.

The solution to the entropy-regularized problem delivers *task-shares* π_{fi}

$$\pi_{fi} = \frac{\exp\left(\frac{z_{fi}-p_{fi}}{\sigma_i}\right)}{\sum\limits_{g\in i}\exp\left(\frac{z_{gi}-p_{gi}}{\sigma_i}\right)} \tag{2}$$

Importantly, π_{fi} is not a quantity; it is an assignment share. It tells us the share of industry i tasks, orders, or procurements assigned to firm f.

Replacing equation (2) into the effective price problem delivers the *effective cost index of industry* i as

$$p_i^{eff} = b_i - \sigma_i \log\left[\sum_{f\in i}\exp\left(\frac{z_{fi}-p_{fi}}{\sigma_i}\right)\right]. \tag{3}$$

The effective cost index is not just a function of p_{fi}, but it is also a function of z_{fi} and b_i. As a result, it is more akin to a *quality-adjusted price* than to a standard price index. Using this definition, we can rewrite the assignment share as

$$\pi_{fi} = \exp\left(\frac{z_{fi} - b_i + p_i^{eff} - p_{fi}}{\sigma_i}\right) \tag{4}$$

The quantity produced by firm f in industry i is:

$$q_{fi} = \pi_{fi} q_i, \tag{5}$$

where q_i is the real composite output of industry i from the supply-side.

Sales/Revenues of firm f in industry i are then

$$R_{fi} = p_{fi} q_{fi} = p_{fi} \pi_{fi} q_i. \tag{6}$$

Remark 1. *Under the logit structure, we have*

$$p_i^{eff} q_i \neq \sum_{f\in i} p_{fi} q_{fi}. \tag{7}$$

Remark 1 highlights that we cannot separate a quantity and a price index as we can under constant elasticity of substitution (CES) aggregation.

Using the definition of sales, we can construct sales shares for each firm in the industry

(s_{fi}) as :

$$s_{fi} = \frac{R_{fi}}{\sum\limits_{g \in i} R_{gi}} = \frac{p_{fi}\pi_{fi}}{\sum\limits_{g \in i} p_{gi}\pi_{gi}} = \frac{p_{fi}\exp\left(\frac{z_{fi}-p_{fi}}{\sigma_i}\right)}{\sum\limits_{g \in i} p_{gi}\exp\left(\frac{z_{gi}-p_{gi}}{\sigma_i}\right)}. \tag{8}$$

Remark 2. *Under the logit structure, sales shares are no longer homogeneous of degree 0 in prices.*

Remark 2 highlights a key property of shift-invariant models: sales shares are no longer homogeneous of degree zero in firms' prices within the industry.

For future reference, it will prove useful to define the sales share of the foreign firm (s_{Fi}) as a separate entity:

$$\sum_{f \in \mathcal{D}_i} s_{fi} + s_{Fi} = 1. \tag{9}$$

3.1.2 Within Industry Competition

Each domestic firm $f \in \mathcal{D}_i$ solves the following program:

$$\Pi_{fi} = \max_{p_{fi}} \quad (p_{fi} - mc_{fi})\pi_{fi}q_i \quad \text{s.t} \quad (2). \tag{10}$$

Note that we assume Bertrand competition in what follows, since firms choose prices directly.

At this point, we need to make the following assumption:

Assumption 1. *Firms consider their effect on the task share π_{fi}, but not the impact of their pricing decisions on the inputs they source.*

Assumption 1 rules out input market power and allows us to focus on the role of competition across firms. This is a strong assumption that we are forced to make given our data availability. The implication of this assumption is that we can take firms' marginal costs as given when making the pricing decision.

The solution to this problem delivers the pricing condition:

$$p_{fi} = mc_{fi} + \mu_{fi}, \tag{11}$$

where $\mu_{fi} = \frac{\sigma_i}{1-\pi_{fi}}$ is firm f's markup. Importantly, the price level p_{fi} is linear in the markup and the marginal cost (μ_{fi}, mc_{fi}). Markups can react in *levels* in our setup, so incomplete pass-through in levels is a possibility.

The production function of firms is constant returns to scale and takes the following form

$$q_{fi} = q_{fi}(\ell_{fi}, x_{fi}, x^*_{fi}; a_{fi}), \tag{12}$$

where ℓ_{fi} is labor demand, x_{fi} is a domestic intermediate input bundle, x^*_{fi} is an imported intermediate input and a_{fi} is a firm-specific productivity.

The domestic intermediate input bundle is a composite of industry-specific bundles $x_{fi,j}$

$$x_{fi} = x_{fi}(\{x_{fi,j}\}_{j\in\mathcal{I}}). \tag{13}$$

The industry-specific bundles are themselves an aggregation across *domestic* firms within each industry

$$x_{fi,j} = x_{fi,j}(\{x_{fi,gj}\}_{g\in\mathcal{D}_j}). \tag{14}$$

The imported intermediate input bundle is a composite of industry-specific import bundles $x^*_{fi,j}$:

$$x^*_{fi} = x^*_{fi}(\{x^*_{fi,j}\}_{j\in\mathcal{I}}). \tag{15}$$

The price of each industry-specific import bundles is p^*_j and is taken as given.

Given our setup, marginal costs satisfy

$$mc_{fi}(w, p^x_{fi}, p^{x^*}_{fi}; a_{fi}), \tag{16}$$

where w is the wage, p^x_{fi} is the price of the domestic intermediate input bundle, $p^{x^*}_{fi}$ is the price of the imported intermediate input bundle and a_{fi} is firm's productivity. Note that for now we leave both prices of intermediate input bundles as firm-industry specific.

3.2 Final Demand

3.2.1 Assignment Problem

We model final consumers in a manner similar to that of input markets. We consider a discrete set of expenditure categories $e \in \mathcal{E}$. Within each expenditure category e there is a continuum of varieties $\omega \in [0, c_e]$. A discrete number of firms $\mathcal{N}_e$ produce these varieties. The household's problem is to decide the share of varieties to buy from each firm in each category. For each category e, the household problem is then:

$$p_e^{eff} = \min_{\pi_{fe} \in \Delta_e} \left\{ b_e + \sum_f \pi_{fe}(p_{fe} - z_{fe}) + \sigma_e \sum_f \pi_{fe} \log \pi_{fe} \right\}, \tag{17}$$

where $\Delta_e = \{\pi_{fe} : \sum_f \pi_{fe} = 1\}$ is the simplex set.

As before, the solution to this problem delivers:

$$\pi_{fe} = \frac{\exp\left(\frac{z_{fe} - p_{fe}}{\sigma_e}\right)}{\sum_{g \in e} \exp\left(\frac{z_{ge} - p_{ge}}{\sigma_e}\right)}, \tag{18}$$

which is the share of varieties that households buy from firm f in category e. The demand for a firm's output is thus

$$c_{fe} = \pi_{fe} c_e. \tag{19}$$

3.2.2 Retailers

We model the problem of firms within each category as a Bertrand competition. The problem of each retailer is thus

$$\Pi_{fe} = \max_{p_{fe}} \, (p_{fe} - m_{fe}) \pi_{fe} c_e \quad \text{s.t.} \quad (18) \tag{20}$$

This delivers the pricing condition:

$$p_{fe} = mc_{fe} + \frac{\sigma_e}{1 - \pi_{fe}}. \tag{21}$$

The retailers production function is constant returns to scale

$$c_{fe} = c_{fe}(\ell_{fe}, \{x_{fe,gj}\}_{g \in \mathcal{D}_j, j \in \mathcal{I}}, \{x^*_{fe,j}\}_{gj \in \mathcal{I}}; a_{fe}), \tag{22}$$

where ℓ_{fe} is labor usage, $x_{fe,gj}$ is domestic intermediate input demand, $x^*_{fe,j}$ is imported intermediate input demand and a_{fe} is a productivity.

The production function delivers a marginal cost function for the retailers

$$mc_{fe} = mc_{fe}(w, \{p_{gj}\}_{g \in \mathcal{D}_j, j \in \mathcal{I}}, \{p^*_j\}_{j \in \mathcal{I}}; a_{fe}) \tag{23}$$

3.3 Market Clearing

Total sales of firm f in industry i must satisfies the following market clearing condition

$$p_{fi}q_{fi} = \sum_{j\in\mathcal{I}}\sum_{g\in\mathcal{D}_j} p_{fi}x_{gj,fi} + \sum_{e\in\mathcal{E}}\sum_{h\in e} p_{fi}x_{he,fi}, \tag{24}$$

where the right hand side is demand by intermediate producing firms plus demand by retailers.

Total labor payments of the household satisfies:

$$w\bar{L} = \sum_{j\in\mathcal{I}}\sum_{g\in\mathcal{D}_j} w\ell_{gj} + \sum_{e\in\mathcal{E}}\sum_{h\in e} w\ell_{he}. \tag{25}$$

3.4 Price Indices Definitions

3.4.1 Industry level price indices

We define the sales weighted average price index change of industry i as

$$\mathrm{d}\log p_i = \sum_{f\in\mathcal{D}_i} s_{fi}\mathrm{d}\log p_{fi} + s_{Fi}\mathrm{d}\log p_{Fi}. \tag{26}$$

This index is observed in the data but it does not correspond to any market price in the usual sense.

The *observed* price index change of industry i, $\mathrm{d}\log p_i$, is linked to the effective price change as

$$\mathrm{d}p_i^{eff} = \frac{R_i}{q_i}\mathrm{d}\log p_i \tag{27}$$

Define domestic sales as $R_i^d = \sum_{f\in\mathcal{D}_i} R_{fi}$ and domestic sales shares as

$$s_{fi}^d = \frac{R_{fi}}{R_i^d}. \tag{28}$$

Let the industry-level variable costs be

$$VC_i^d = \sum_{g\in\mathcal{D}_i} mc_{gi}q_{gi}; \qquad VC_i = \sum_{g\in i} mc_{gi}q_{gi}, \tag{29}$$

where VC_i^d is variable costs across domestic firms and VC_i is the industry's total variable costs, inclusive of those of the foreign firm.

We also define the domestic price index of industry i as a domestic-sales weighted average across domestic firms $\mathcal{D}_i$

$$\mathrm{d}\log p_i^d = \sum_{f \in \mathcal{D}_i} s_{fi}^d \mathrm{d}\log p_{fi} \tag{30}$$

Equation (30) defines the producer price index (PPI) of industry i as in the data.

3.4.2 Consumer's price indices

We define the price index change in a category e as

$$\mathrm{d}\log p_e = \sum_{f \in e} s_{fe} \mathrm{d}\log p_{fe}, \tag{31}$$

where $s_{fe} = \frac{p_{fe} c_{fe}}{\sum_{g \in e} p_{ge} c_{ge}}$ is the expenditure share on firm f within category e.

Using this definition, we further define the consumer price index change as in the data:

$$\mathrm{d}\log p^C = \sum_{e \in \mathcal{E}} s_e \mathrm{d}\log p_e, \tag{32}$$

where $s_e = \frac{\sum_{g \in e} p_{ge} c_{ge}}{\sum_{e \in \mathcal{E}} \sum_{g \in e} p_{ge} c_{ge}}$ is the expenditure share in category e.

3.5 The Transmission of Import Prices

In this section, we derive the key equations that we bring to the data. To do so, we consider a first-order approximation around an initial equilibrium.

3.5.1 Pass-through

The change in the price of a firm f-industry i can be written as a function of both marginal cost changes and changes in competitors' prices. The following proposition characterizes this result in levels:

Proposition 1 (Pass-through in levels). *Up to a first-order approximation, the level change in the price of firm f in industry i can be written as*

$$\mathrm{d}p_{fi} = (1 - \pi_{fi})\mathrm{d}mc_{fi} + \pi_{fi}\mathrm{d}p_{-f,i}, \tag{33}$$

where

$$\mathrm{d}p_{-f,i} = \sum_{h\in i\backslash f} \tilde{\pi}_{hi}\mathrm{d}p_{hi} \qquad \textit{and} \qquad \tilde{\pi}_{hi} = \frac{\pi_{hi}}{\sum\limits_{g\in i\backslash f}\pi_{gi}}, \tag{34}$$

is the competitor's price index for firm f and $\tilde{\pi}_{hi}$ is the probability that firm h in industry i performs task within domestic firms.

Proof. See Appendix B.1. □

Proposition 1 shows that changes in the firm's price can be decomposed between an own marginal cost change and a competitors' price index. The complete pass-through in level is achieved whenever $\pi_{fi} = 0$ and corresponds to the case analyzed in Sangani (2026). In this case, the fraction of tasks a particular firm performs is negligible relative to the aggregate, an intuition similar to that in monopolistic competition models. By varying π_{fi}, we can nest different cases from the literature, but under pass-through in levels.

The next proposition rewrites the firm-level pass-through in percentage changes[1] :

Proposition 2 (Pass-through in percentage changes)**.** *Up to a first-order approximation, the percentage change in the price of firm f in industry i can be written as*

$$\mathrm{d}\log p_{fi} = (1-\alpha_{fi})(1-\pi_{fi})\mathrm{d}\log mc_{fi} + \pi_{fi}\frac{p_{-f,i}}{p_{fi}}\mathrm{d}\log p_{-f,i}, \tag{35}$$

where

$$1-\alpha_{fi} = \frac{mc_{fi}}{p_{fi}} \in [0,1] \tag{36}$$

is the share of price due to marginal costs.

Proof. See Appendix B.2. □

Proposition 2 shows that when there is complete pass-through in levels, $\pi_{fi} = 0$, there is still incomplete pass-through in percentage changes since $1-\alpha_{fi} \leq 1$. More generally, $\pi_{fi} \neq 1$. Thus, the presence of oligopolistic competition when firms face a logit-demand structure adds an additional mechanism through which changes in marginal costs can affect firm prices captured by the competitor's price channel $\pi_{fi}\frac{p_{-f,i}}{p_{fi}}\mathrm{d}\log p_{-f,i}$.

We are now ready to examine how changes in import prices affect PPI. In our model, changes in import prices enter into two places: (i) through marginal cost, as an input in

[1] Note that to write $\mathrm{d}\log mc_{fi} = \mathrm{d}\log mc_i$ we require all firms to have the same initial marginal cost. This can still lead to different market and task shares provided that there is heterogeneity in z_{fi}.

the production function, and (ii) as changes in the price of competitors in an industry (as it contains $\mathrm{d}\log p_{Fi}$). The following proposition introduces a perturbation $(\mathrm{d}\log mc_{fi}, \mathrm{d}\log p_{Fi})$ and provides an expression for industry level PPI changes:

Proposition 3. *Changes in the producer price index after a perturbation* $(\mathrm{d}\log mc_{fi}, \mathrm{d}\log p_{Fi})$ *satisfy*

$$\mathrm{d}\log p_i^d = \frac{1}{\left(1-(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)\right)}\left(\upsilon_i^d\left(\sum_{f\in\mathcal{D}_i}\tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi}\right)\right), \tag{37}$$
$$+\frac{1}{\left(1-(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)\right)}\left(\frac{(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)}{(1-s_{Fi})}s_{Fi}\mathrm{d}\log p_{Fi}\right)$$

where

$$\tilde{\delta}_{fi}=\frac{mc_{fi}q_{fi}}{VC_i^d};\qquad \upsilon_i^d=\frac{VC_i^d}{R_i^d};\qquad \rho_{fi}=\frac{(1-\pi_{fi})^2}{1-\pi_{fi}+\pi_{fi}^2}\in[0,1]$$
$$\tilde{\pi}_{fi}=\frac{\pi_{fi}}{\sum_{h\in\mathcal{D}_i}\pi_{hi}};\qquad 1-\pi_{Fi}=\sum_{f\in\mathcal{D}_i}\pi_{fi}$$

Proof. See Appendix B.3. □

There are some parameters that deserve discussion from the proposition. The parameter ν_i^d represents the share of domestic variable costs in domestic revenues; this is the industry level counterpart of $1-\alpha_{fi}$. The parameter ρ_{fi} captures a "firm cost pass-through in levels" as it is an increasing function of $(1-\pi_{fi})$, the level marginal cost firm's pass-through we found in Proposition 1.

The term $\sum_{f\in\mathcal{D}_i}\tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi}$ is an industry-level shifter in marginal cost changes of all firms appropriately weighted. This industry level component have two sources of firm-level heterogeneity (beyond the idiosyncratic cost shifters): (i) initial cost shares of each firm in total industry costs $(\tilde{\delta}_{fi})$ and (ii) the cost pass-through in levels parameters (ρ_{fi}). The industry-level marginal cost shifter thus places more weight on larger firms (in terms of costs) and firms with higher cost pass-through in levels. This will play an important role in our empirical application. The direct cost pass-through in percentage changes at the industry level is thus the numerator $\nu_i^d\sum_{f\in\mathcal{D}_i}\tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi}$.

The second row shows the effect of an increase in foreign competition, captured by a change in the foreign firm's price, $\mathrm{d}\log p_{Fi}$. This effect can be understood by looking at two elasticities. The first elasticity is how domestic prices react to a change in the industry-level price; this effect is captured by $\frac{(1-\pi_{Fi})}{1-s_{Fi}} \sum\limits_{f\in\mathcal{D}_i} \tilde{\pi}_{fi}(1-\rho_{fi})$ that depends on task shares π_{fi}, foreign sales shares s_{Fi}, and cost pass-through ρ_{fi}. Of course, the pass-through of the industry price index on the producer price index will be larger whenever the competitors' effect is stronger (higher $1-\rho_{fi}$), (ii) sales to foreigners are larger (higher s_{Fi}) and (iii) tasks assigned to domestic firms are an important fraction $(1-\pi_{Fi})$. The second elasticity is the elasticity of the industry price index with respect to the foreign firm's price. This elasticity is simply $s_{Fi}\mathrm{d}\log p_{Fi}$, the first-order effect of the foreign firm price on the aggregate industry price index.

The total effect of both perturbations also needs to account for the fact that changes in the producer price index affect the industry-level price index, which in turn affects the producer price index, and so on. The first term multiplying both perturbations can be written as

$$\frac{1}{1-(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)} = \sum_{t=0}^{\infty}\left((1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)\right)^t$$

and thus captures all these higher-order rounds.

We conclude this subsection by noting that Proposition 3 also delivers a corollary for the change in prices of expenditure categories. The sole difference between the models is that $\pi_{Fe}=0$ for all e, as there is no foreign competition in retailers. The following corollary captures this:

Corollary 1. *Changes in the price index of category e, $\mathrm{d}\log p_e$, after a change in firm level marginal costs $(\mathrm{d}\log mc_{fe})$ satisfy*

$$\mathrm{d}\log p_e = \frac{1}{\left(1-\sum\limits_{f\in e}\pi_{fe}(1-\rho_{fe})\right)}\upsilon_e\left(\sum_{f\in e}\delta_{fe}\rho_{fe}\mathrm{d}\log mc_{fe}\right) \tag{38}$$

where

$$\delta_{fe}=\frac{mc_{fe}c_{fe}}{VC_e}; \qquad \upsilon_e=\frac{VC_e}{R_e}; \qquad \rho_{fe}=\frac{(1-\pi_{fe})^2}{1-\pi_{fe}+\pi_{fe}^2}\in[0,1]$$

3.5.2 The Role of Industry Level Input-Output Networks

The preceding section provided a perturbation on the space $(\mathrm{d}\log mc_{fi}, \mathrm{d}\log p_{Fi})$. In our model, firms' marginal costs depends of firms' prices in other industries as well. We now provide an extension where we perturb intermediate import prices $\mathrm{d}\log p_i^*$ and the price of foreign firms $\mathrm{d}\log p_{Fi}$ and study how this propagates through the producer price index.

To characterize the producer price index response we make some simplifying assumptions. The role of these simplifying assumptions is to allow us to map our formulas to industry-level data. Appendix A.1 provides a formal discussion on the implication of each assumption. We make the following assumptions:

Assumption 2 (Expenditure shares across firms within industries)**.** *Expenditure shares across varieties of a given industry are the same regardless of the buyer origin* $\Omega_{fi,gj} = \Omega_{\cdot,gj}$ *for all* f, i *pairs.*

Assumption 3 (Expenditure shares across industries)**.** *There is no heterogeneity in expenditure shares across firms within an industry:* $\Omega_{fi,j} = \Omega_{i,j}$ *for all* $f \in i$*.*

Assumption 4. *For each firm-industry pair,* $\Omega_{\cdot,fi} = s^d_{fi}$*.*

Assumption 2 ensures that all firms pay the same price when sourcing varieties from another industry because and so it allows us to use a single intermediate-input industry price index change that is common across buyers.

Assumption 3 implies we only need information at the industry level to construct the industry-level domestic intermediate input price index changes. This also implies that marginal cost changes across firms within industries will be common $\mathrm{d}\log mc_{fi} = \mathrm{d}\log mc_i$.

Assumption 4 ensures that the intermediate input price index change paid by industries coincides with the producer price index change.

Proposition 4 (Producer price indices and input-output networks)**.** *Consider a perturbation of import prices* $(\mathrm{d}\log \boldsymbol{p}^*, \mathrm{d}\log \boldsymbol{p}_F)$ *keeping constant wages* $\mathrm{d}\log w = 0$*. The response of producer price indices is such that*

$$\mathrm{d}\log \boldsymbol{p}^d = \underbrace{\boldsymbol{\Phi}^* \mathrm{d}\log \boldsymbol{p}^*}_{\textit{Intermediate Import Price Channel}} + \underbrace{\boldsymbol{\Phi}^F \mathrm{d}\log \boldsymbol{p}_F}_{\textit{Foreign Competition Channel}}, \tag{39}$$

where

$$\boldsymbol{\Phi}^* = \boldsymbol{G}\boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^* \tag{40}$$

$$\boldsymbol{\Phi}^F = \boldsymbol{G}(\boldsymbol{I} - \boldsymbol{S}_F)^{-1}(\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d)\boldsymbol{S}_F \tag{41}$$

$$\boldsymbol{G} = (\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^d - (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d))^{-1} \tag{42}$$

and

$$\boldsymbol{\Upsilon}^d = diag(\boldsymbol{\upsilon}^d); \quad \boldsymbol{\rho}^d = diag(\{\rho_i^d\}_{i\in\mathcal{I}}); \quad \boldsymbol{S}_F = diag(\boldsymbol{s}_F); \quad \boldsymbol{\Pi}_F = diag(\boldsymbol{\pi}_F),$$

are diagonal matrices and

$$\rho_i^d = \sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi},$$

is a cost-weighted average of the cost pass-through parameters.

Proof. See Appendix B.4. □

Proposition 4 provides an industry level counterpart to Proposition 3. While Proposition 3 required detailed firm-level information, Proposition 4 only requires industry level data. Moreover, Proposition 4 decomposes changes in industry level producer price indices into an intermediate import price channel and a foreign competition channel, while the earlier proposition was agnostic about the particular sources of variation on marginal costs.

Although it is heavy on notation, Proposition 4 provides a natural extension to our earlier results. There are two key differences. First, note that the higher order terms now take a matrix form as now producer price indices are linked across sectors via marginal costs. This production network structure is captured in $\boldsymbol{\Upsilon}^d\boldsymbol{\rho}^d\boldsymbol{\Omega}^d$ in the propagation matrix $\boldsymbol{G}$, where $\boldsymbol{\Omega}^d$ captures the exposure of industries to producer price indices in all other industries. The remaining objects in $\boldsymbol{G}$ are just the natural extension to matrices of the higher order rounds in Proposition 3. Apart from this modification, the direct effect of foreign competition on producer price indices is identical to the one Proposition 4.

Second, Proposition 4 directly links the effect of imported intermediate input prices on producer price indices, which is captured by the matrix $\boldsymbol{\Phi}^*$. This matrix has a direct component and a propagation component ($\boldsymbol{G}$). As we already explained the propagation component, we focus on the direct effect. This direct effect combines two components: (i) the direct marginal cost exposure to imported intermediate inputs $\boldsymbol{\Omega}^*$, (ii) the cost pass-through in percentage changes matrix $\boldsymbol{\Upsilon}^d\boldsymbol{\rho}^d$. Intuitively, a rise in imported intermediate input prices raises industry's marginal cost by Ω_i^* (the ith row of the $\mathcal{I}\times\mathcal{I}$ matrix $\boldsymbol{\Omega}^*$). This marginal cost change then translates into a change in the producer price index of the industry of $\upsilon_i^d\rho_i^d$. This change in the producer price index of industry i then affect both producer price indices in all the other industries as well as the industry price level (inclusive of foreign firms), which in turn affect industries' marginal costs and so on. The matrix $\boldsymbol{\Phi}^*$ is capturing essentially these transmission mechanisms.

3.6 Preview of Empirical Strategy

Equation (37) shows how domestic price indices react to a change in the foreign price $\mathrm{d}\log p_{Fi}$ and firms' marginal costs $\mathrm{d}\log mc_{fi}$. We can rewrite equation (37) as[2]:

$$d\log p_i^d = \frac{\upsilon_i^d\left(\mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[\boldsymbol{\rho}_i]\,\mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[\mathrm{d}\log \boldsymbol{mc}_i] + \mathbb{C}\mathrm{ov}_{\tilde{\boldsymbol{\delta}}_i}\big(\boldsymbol{\rho}_i, d\log \mathbf{mc}_i\big)\right) + (1-\mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[\boldsymbol{\rho}_i])s_{Fi}d\log p_{Fi}}{\mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[\boldsymbol{\rho}_i] + (1-\mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[\boldsymbol{\rho}_i])s_{Fi}}, \quad (43)$$

where we define the expectation operator $\mathbb{E}_\omega[x] = \sum_i \omega_i x_i$ and $\mathbb{C}\mathrm{ov}_\omega[x,y] = E_\omega[xy] - \mathbb{E}_\omega[x]\,\mathbb{E}_\omega[y]$.

To streamline notation, we re-express the equation above in an equivalent but more compact form:

$$d\log p_i^d = \frac{\upsilon_i^d\left(\rho_i^d\, d\log mc_i + \mathbb{C}\mathrm{ov}_{\tilde{\boldsymbol{\delta}}_i}\big(\boldsymbol{\rho}_i, d\log \mathbf{mc}_i\big)\right) + (1-\rho_i^d)s_{Fi}d\log p_{Fi}}{\rho_i^d + (1-\rho_i^d)s_{Fi}}, \quad (44)$$

There are three channels that can generate a wedge between industry-level price changes $d\log p_i^d$ and marginal cost changes $d\log \boldsymbol{mc}_i$:

1. **Markup assumption (levels vs logs).** Captured by ν_i^d, which attenuates the pass-through of marginal cost shocks.

2. **Competition channel.** When $0 < \rho_i^d < 1$, firms respond to movements in the reference price, which itself depends on foreign competitors through s_{Fi}.

3. **Within-industry exposure heterogeneity.** Captured by the covariance term $\mathbb{C}\mathrm{ov}_{\tilde{\boldsymbol{\delta}}_i}\big(\boldsymbol{\rho}_i, d\log \mathbf{mc}_i\big)$, which reflects the interaction between firm-level pass-through and firm-level exposure to shocks.

The covariance term is the analogue of the aggregate markup-adjustment term in Amiti, Itskhoki and Konings (2019). In their exchange-rate setting, firms differ both in their exposure to the exchange-rate shock and in the extent to which they absorb cost shocks through markups. In our setting, we condition directly on border price changes and decompose firms' exposure into imported-input marginal-cost exposure and foreign-competitor price exposure. The covariance term captures whether the domestic firms most exposed to the foreign cost shock are also the firms with high pass-through.[3]

[2] We approximate this equation under the assumption that $\pi_{Fi} \approx s_{Fi}$ and $\tilde{\pi}_{fi} \approx \tilde{\delta}_{fi}$.

[3] Amiti, Itskhoki and Konings (2019) write the covariance in terms of markup adjustment, whereas we write it in terms of own-cost pass-through. Since these objects are complements, the sign convention differs mechanically.

3.6.1 Markup Assumption

We approximate firm-level marginal cost of firms within an industry as

$$d\log \text{mc}_{fi} = \sum_j \Omega^d_{fi,j} d\log p^d_j + \sum_j \Omega^*_{fi,j} d\log p^*_j, \tag{45}$$

Aggregating across firms using domestic cost shares $\tilde{\delta}_{fi}$ yields

$$d\log \text{mc}^d_i = \sum_f \tilde{\delta}_{fi} d\log \text{mc}_{fi} \tag{46}$$

$$= \sum_j \Omega^d_{ij} d\log p^d_j + \sum_j \Omega^*_{ij} d\log p^*_j, \tag{47}$$

where

$$\Omega^d_{ij} = \sum_f \tilde{\delta}_{fi} \Omega^d_{fi,j}, \qquad \Omega^*_{ij} = \sum_f \tilde{\delta}_{fi} \Omega^*_{fi,j}$$

are the observed industry-level cost shares.

We use the same approximation for the firm-level marginal cost of firms within expenditure categories:

$$d\log \text{mc}_{fe} = \sum_j \theta^d_{fe,j} d\log p^d_j + \sum_j \theta^*_{fe,j} d\log p^*_j, \tag{48}$$

Aggregating across firms using cost-shares δ_{fe} yields

$$d\log \text{mc}_e = \sum_f \delta_{fe} d\log \text{mc}_{fe} \tag{49}$$

$$= \sum_j \theta^d_{ej} d\log p^d_j + \sum_j \theta^*_{ej} d\log p^*_j, \tag{50}$$

where

$$\theta^d_{ej} = \sum_f \delta_{fe} \theta^d_{fe,j}, \qquad \theta^*_{ej} = \sum_f \delta_{fe} \theta^*_{fe,j}$$

are the observed expenditure-level cost shares.

Under the benchmark with full own-cost pass-through and no covariance wedge, the domestic producer-price system can be written as

$$\mathrm{d}\log \boldsymbol{p}^d = \boldsymbol{\Upsilon}^d \left(\boldsymbol{\Omega}^d \mathrm{d}\log \boldsymbol{p}^d + \boldsymbol{\Omega}^* \mathrm{d}\log \boldsymbol{p}^* \right). \tag{51}$$

Solving for domestic producer prices gives

$$\mathrm{d}\log \boldsymbol{p}^d = \left(\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^d\right)^{-1} \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^* \mathrm{d}\log \boldsymbol{p}^*. \tag{52}$$

The matrix $\boldsymbol{\Omega}^*$ maps border-price changes into the marginal costs of domestic producers that purchase imported intermediate inputs. The inverse matrix $\left(\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^d\right)^{-1}$ then captures the propagation of these cost changes through the domestic production network. The diagonal matrix $\boldsymbol{\Upsilon}^d$ implements the selected markup convention. Its empirical construction under pass-through in levels and pass-through in logs is described in Section 4.

We next map domestic producer prices into expenditure-category prices. Let

$$\boldsymbol{\Theta}^d = \left[\theta^d_{ej}\right]_{e\in\mathcal{E},\, j\in\mathcal{I}}, \qquad \boldsymbol{\Theta}^* = \left[\theta^*_{ej}\right]_{e\in\mathcal{E},\, j\in\mathcal{I}}$$

denote the expenditure-category sourcing matrices defined above, and let $\mathrm{d}\log \mathbf{mc}^r$ denote the vector of retailer marginal-cost changes. Aggregating equation (48) across retailers within each expenditure category yields

$$\mathrm{d}\log \mathbf{mc}^r = \boldsymbol{\Theta}^d \mathrm{d}\log \boldsymbol{p}^d + \boldsymbol{\Theta}^* \mathrm{d}\log \boldsymbol{p}^*. \tag{53}$$

The two matrices capture distinct sources of retailer costs. The matrix $\boldsymbol{\Theta}^d$ measures retailers' exposure to domestically produced commodities, whose prices already incorporate the propagation of imported-input shocks through the domestic production network. The matrix $\boldsymbol{\Theta}^*$ measures retailers' direct exposure to imported finished commodities.

This final-demand layer differs from the domestic production block in one important respect. All firms within an expenditure category are domestic retailers. Imported finished commodities have already crossed the border when they enter this layer: they are purchased as inputs by domestic retailers rather than supplied by foreign retailers competing with domestic retailers. There is therefore no foreign-competitor-price term in equation (53).

Under the benchmark of Corollary 1, the expenditure-category price vector satisfies

$$\mathrm{d}\log \boldsymbol{p} = \boldsymbol{\Upsilon}^r \mathrm{d}\log \mathbf{mc}^r, \tag{54}$$

where $\boldsymbol{\Upsilon}^r$ is the diagonal markup-adjustment matrix for the distribution layer. Combining equations (53) and (54) gives

$$\mathrm{d}\log \boldsymbol{p} = \boldsymbol{\Upsilon}^r \left(\boldsymbol{\Theta}^d \mathrm{d}\log \boldsymbol{p}^d + \boldsymbol{\Theta}^* \mathrm{d}\log \boldsymbol{p}^*\right). \tag{55}$$

Finally, substituting equation (52) into equation (55) yields

$$\mathrm{d}\log \boldsymbol{p} = \boldsymbol{\Upsilon}^r \left[\boldsymbol{\Theta}^d \left(\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^d\right)^{-1} \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^* + \boldsymbol{\Theta}^*\right] \mathrm{d}\log \boldsymbol{p}^*. \tag{56}$$

Equation (56) decomposes the response of expenditure-category prices into two channels. The indirect channel maps border-price changes into domestic producer prices through imported intermediate inputs and the domestic input-output network, and then maps domestic producer prices into retailer costs. The direct channel maps border-price changes in imported finished commodities directly into retailer costs. All the objects in equation (56) are observable, so the equation provides the baseline sensitivity measure used in the descriptive analysis.

3.6.2 Oligopolistic Competition Channel

To study this channel we will focus on domestic producer prices since retailers do not directly compete with foreign companies. We will study this channel on top of the markup assumption but shutting down the covariance channel. This is because we want to gradually enrich the model and find out which channels are necessary, while ideally maintaining a parsimonious specification.

In this case (44) becomes a weighted average between the marginal cost impact and the impact of the tariffs via foreign competitors:

$$d\log p_i^d = \upsilon_i^d \tilde{\rho}_i^d \, \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[\mathrm{d}\log \boldsymbol{mc}_i] + (1-\tilde{\rho}_i^d) d\log p_{Fi} \,, \tag{57}$$

where $\tilde{\rho}_i^d = \rho_i^d/(\rho_i^d + (1-\rho_i^d)s_{Fi})$ Now domestic prices move more or less than marginal cost according to the sales share of foreign competitors within i and according to the tariff hike in i.

The object ρ_i^d is not observable since it is a function of the industry structure and the elasticity of substitution between products. However, given the amount of shocks and the amount of observables in the model it can be estimated non-parametrically. With this 'revealed oligopoly' estimation, we can evaluate its performance and plausibility in explaining observed inflation.

3.6.3 Covariance channel

Using (45), the covariance term can be written as

$$\mathbb{C}\text{ov}_{\tilde{\boldsymbol{\delta}}_i}\big(\boldsymbol{\rho}_i, d\log \mathbf{mc}_i\big) = \sum_j \mathbb{C}\text{ov}_{\tilde{\boldsymbol{\delta}}_i}\big(\boldsymbol{\rho}_i, \boldsymbol{\Omega}^d_{i,j}\big) d\log p^d_j + \sum_j \mathbb{C}\text{ov}_{\tilde{\boldsymbol{\delta}}_i}\big(\boldsymbol{\rho}_i, \boldsymbol{\Omega}^*_{i,j}\big) d\log p^*_j \tag{58}$$

$$= \sum_j \mathbb{C}^d_{ij} d\log p^d_j + \sum_j \mathbb{C}^*_{ij} d\log p^*_j, \tag{59}$$

where

$$\mathbb{C}^*_{ij} = \sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\left(\rho_{fi} - \rho^d_i\right)\Omega^*_{fi,j}. \tag{60}$$

Define the exposure-weighted pass-through among firms exposed to input j:

$$\rho^*_{ij} \equiv \frac{\sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi}\Omega^*_{fi,j}}{\sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\Omega^*_{fi,j}} = \frac{\sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi}\Omega^*_{fi,j}}{\Omega^*_{ij}}. \tag{61}$$

Then (60) can be rewritten as

$$\mathbb{C}^*_{ij} = \Omega^*_{ij}\left(\rho^*_{ij} - \rho^d_i\right). \tag{62}$$

Restrictions. We now impose the following identification assumptions:

1. *Within industry i, the difference between exposure-weighted pass-through and average pass-through is constant across input sectors:*

$$\rho^*_{ij} - \rho^d_i = \gamma_i \qquad \forall j. \tag{63}$$

2. $\mathbb{C}^d = \boldsymbol{I}$

Under the first assumption, (62) simplifies to

$$\mathbb{C}^*_{ij} = \gamma_i\Omega^*_{ij}. \tag{64}$$

The first restriction implies that within industry i, firms that are more exposed to imported inputs systematically have pass-through that differs from the industry average by a constant amount γ_i, independently of which input sector generates the exposure.

When $\gamma_i > 0$, firms that are more exposed to import costs also tend to have higher pass-through, amplifying the effect of foreign shocks. When $\gamma_i < 0$, the opposite holds. This is simply a plausible assumption to allow for limited degrees of freedom in the nonparametric estimation. [4] The second restriction is also necessary to reduce degrees of freedom. Since $\mathbb{C}^d$ affects prices indirectly, this assumption should not have a large impact.

Under such restrictions equation (44), in vector form, becomes

$$d\log \mathbf{p}^d = \Big[I - (I-\boldsymbol{\rho}^d)(I-\boldsymbol{S}_F) - \boldsymbol{\rho}^d\boldsymbol{\Upsilon}^d\boldsymbol{\Omega}^d\Big]^{-1}\Big[\left(\boldsymbol{\rho}^d+\boldsymbol{\gamma}\right)\boldsymbol{\Upsilon}^d\boldsymbol{\Omega}^* + \left(I-\boldsymbol{\rho}^d\right)\boldsymbol{S}_F\Big] d\log \mathbf{p}^* \quad (65)$$

Equation (65) implies that a covariance term would possibly amplify or attenuate the effects of marginal costs increases, in accordance to an estimated diagonal matrix $\boldsymbol{\gamma}$ that summarizes how the distribution of markups and affected firms within the industries are affected by the tariffs. Effectively, under our restriction assumption, the matrix $\boldsymbol{\gamma}$ can further disentangle the link between predicted and observed marginal costs.

The issue with this equation is that both the matrix $\boldsymbol{\rho}$ and $\boldsymbol{\gamma}$ are unobserved. Therefore this is a system with N equations and $2N$ parameters. In practice, we aim at parameterizing such values with industry observable characteristics such as HHI index, markup size, foreign shares, and others. Of course equation (65) also implies that any model estimating ρ without accounting for the potential impact of γ would be biased.

Note that this channel cannot be active in the intermediate sector unless $\boldsymbol{\rho}^d \neq I$. Under the normalization $\boldsymbol{\rho}^d = \boldsymbol{I}$, the covariance wedge $\mathbb{C}^*$ collapses to zero. This is because if the average ρ_i^d in an industry is 1, then all $\rho_{fi} = 1$ for all firms f in i. But this implies that the covariance term between ρ_{fi} and the marginal cost changes is zero.

Retail-side covariance wedge. A related covariance wedge may arise in the final-demand layer. Unlike domestic producers, retailers do not compete with foreign retailers. Imported finished goods have already crossed the border when they enter this layer and are purchased as inputs by domestic retailers. There is therefore no retail analogue of the foreign-competition channel governed by $\boldsymbol{\rho}^d$. Retailers may nevertheless differ both in their own-cost pass-through and in their exposure to domestic and imported commodities.

To make this distinction explicit, recall from equation (48) that retailer-level marginal-cost changes satisfy

$$\mathrm{d}\log \mathrm{mc}_{fe} = \sum_j \theta^d_{fe,j}\,\mathrm{d}\log p^d_j + \sum_j \theta^*_{fe,j}\,\mathrm{d}\log p^*_j. \quad (66)$$

Under the approximation $\pi_{fe} \approx \delta_{fe}$, define the cost-weighted average own-cost pass-through within expenditure category e as

$$\rho^r_e \equiv \sum_{f\in e} \delta_{fe}\rho_{fe}. \quad (67)$$

[4] We have also allowed for $\mathbb{C}$ to be a free diagonal matrix and our non-parametric estimation has shown that the cost share is strongly correlated to the diagonal elements.

Corollary 1 can then be rewritten as

$$\mathrm{d}\log p_e = \upsilon_e \left[\mathrm{d}\log \mathrm{mc}_e + \frac{\mathbb{C}\mathrm{ov}_{\boldsymbol{\delta}_e}\left(\boldsymbol{\rho}_e, \mathrm{d}\log \mathbf{mc}_e\right)}{\rho^r_e} \right]. \tag{68}$$

Using equation (66), define the two retail-side covariance matrices by

$$\mathbb{C}^{r,d}_{ej} \equiv \frac{\mathbb{C}\mathrm{ov}_{\boldsymbol{\delta}_e}\left(\boldsymbol{\rho}_e, \boldsymbol{\theta}^d_{e,j}\right)}{\rho^r_e}, \tag{69}$$

$$\mathbb{C}^{r,*}_{ej} \equiv \frac{\mathbb{C}\mathrm{ov}_{\boldsymbol{\delta}_e}\left(\boldsymbol{\rho}_e, \boldsymbol{\theta}^*_{e,j}\right)}{\rho^r_e}. \tag{70}$$

The general retail-price equation is therefore

$$\mathrm{d}\log \boldsymbol{p} = \boldsymbol{\Upsilon}^r \left[\left(\boldsymbol{\Theta}^d + \mathbb{C}^{r,d}\right) \mathrm{d}\log \boldsymbol{p}^d + \left(\boldsymbol{\Theta}^* + \mathbb{C}^{r,*}\right) \mathrm{d}\log \boldsymbol{p}^* \right]. \tag{71}$$

Let $\boldsymbol{\Phi}^{d,*}$ denote the mapping from imported prices into domestic producer prices under the chosen producer-side specification:

$$\mathrm{d}\log \boldsymbol{p}^d = \boldsymbol{\Phi}^{d,*} \mathrm{d}\log \boldsymbol{p}^*. \tag{72}$$

Under the baseline specification in equation (52),

$$\boldsymbol{\Phi}^{d,*} = \left(\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^d\right)^{-1} \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^*. \tag{73}$$

Substituting equation (72) into equation (71) yields

$$\mathrm{d}\log \boldsymbol{p} = \boldsymbol{\Upsilon}^r \left[\left(\boldsymbol{\Theta}^d + \mathbb{C}^{r,d}\right) \boldsymbol{\Phi}^{d,*} + \boldsymbol{\Theta}^* + \mathbb{C}^{r,*} \right] \mathrm{d}\log \boldsymbol{p}^*. \tag{74}$$

Equation (74) makes clear that retailer heterogeneity can affect both the indirect channel, through $\mathbb{C}^{r,d}$, and the direct channel, through $\mathbb{C}^{r,*}$.

Direct-import restriction. The two retail-side covariance matrices cannot be estimated unrestrictedly from expenditure-category price variation alone. As a baseline empirical specification, we therefore impose

$$\mathbb{C}^{r,d} = \mathbf{0}, \qquad \mathbb{C}^{r,*} = \boldsymbol{\Gamma}^r \boldsymbol{\Theta}^*, \qquad \boldsymbol{\Gamma}^r = \mathrm{diag}\left(\{\gamma^r_e\}_{e\in\mathcal{E}}\right). \tag{75}$$

Under this restriction, equation (74) becomes

$$\mathrm{d}\log \boldsymbol{p} = \boldsymbol{\Upsilon}^r \left[\boldsymbol{\Theta}^d \boldsymbol{\Phi}^{d,*} + \left(\boldsymbol{I} + \boldsymbol{\Gamma}^r\right) \boldsymbol{\Theta}^*\right] \mathrm{d}\log \boldsymbol{p}^*. \tag{76}$$

This restriction is not a theoretical implication of the model. It is a parsimonious normalization that isolates the covariance between retailer-level own-cost pass-through and direct imported-final-good exposure, which is the component that generates the sharpest cross-sectional variation across expenditure categories.

Alternative sourcing-substitution specification. As an alternative, suppose that retailers within an expenditure category differ in whether they source a given commodity domestically or internationally, but not in their total exposure to that commodity:

$$\theta^d_{fe,j} + \theta^*_{fe,j} = \theta_{e,j} \qquad \text{for all } f \in e. \tag{77}$$

Under this assumption,

$$\mathbb{C}^{r,d} = -\mathbb{C}^{r,*}. \tag{78}$$

Combining equation (78) with the proportionality restriction

$$\mathbb{C}^{r,*} = \boldsymbol{\Gamma}^r \boldsymbol{\Theta}^* \tag{79}$$

gives

$$\begin{aligned} \mathrm{d}\log \boldsymbol{p} &= \boldsymbol{\Upsilon}^r \left[\left(\boldsymbol{\Theta}^d - \boldsymbol{\Gamma}^r \boldsymbol{\Theta}^*\right) \boldsymbol{\Phi}^{d,*} + \left(\boldsymbol{I} + \boldsymbol{\Gamma}^r\right) \boldsymbol{\Theta}^*\right] \mathrm{d}\log \boldsymbol{p}^* \\ &= \boldsymbol{\Upsilon}^r \left[\boldsymbol{\Theta}^d \boldsymbol{\Phi}^{d,*} + \boldsymbol{\Theta}^* + \boldsymbol{\Gamma}^r \boldsymbol{\Theta}^* \left(\boldsymbol{I} - \boldsymbol{\Phi}^{d,*}\right)\right] \mathrm{d}\log \boldsymbol{p}^*. \end{aligned} \tag{80}$$

This alternative specification gives $\boldsymbol{\Gamma}^r$ a different interpretation. Retailers with higher own-cost pass-through are disproportionately exposed to imported rather than domestically sourced versions of the same commodities. The same wedge therefore raises the direct-import contribution and correspondingly lowers the contribution of domestic sourcing.

4 Mapping the Model to the Data

As a contribution in its own right, this study carefully maps each model object to official US public data. This section provides intuition about the available data and the economic concepts behind them.

4.1 Direct and Indirect Import Sensitivities

We map the model to data using the Bureau of Economic Analysis (BEA) Benchmark Input-Output accounts, which provide a comprehensive, commodity-by-industry picture of intermediate input flows, commodity output shares, import penetration, and final demand in the US economy. We use the Benchmark accounts for 2007, 2012, and 2017 and interpolate them to annual frequency using the BEA Annual Industry Accounts, yielding a time-varying panel of input-output matrices disaggregated to 402 BEA commodity codes, corresponding to 212 NIPA expenditure categories. The Use table records intermediate input purchases by commodity and industry; the Make table records commodity output shares by industry; the Import matrix identifies the import share of each intermediate input; and the PCE bridge maps commodity-level production to the 360 NIPA expenditure categories at purchaser prices, decomposed into producer, retail, wholesale, and transportation margin components.[5]

Equation (56) in our model defines the import price sensitivity of expenditure category j as the sum of a direct and an indirect sensitivity matrix. Specifically, the direct sensitivity is:

$$\Delta^{\text{Direct}} = \boldsymbol{\Upsilon}^r \boldsymbol{\Theta}^*, \tag{81}$$

and the indirect sensitivity is:

$$\Delta^{\text{Indirect}} = \boldsymbol{\Upsilon}^r \boldsymbol{\Theta}^d (\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^d)^{-1} \boldsymbol{\Upsilon}^d \boldsymbol{\Omega}^*. \tag{82}$$

Here, Δ^{Direct} is a $J \times C$ matrix mapping import commodity groups c to NIPA expenditure categories j; each entry $\Delta^{\text{Direct}}_{j,c}$ is the percentage increase in the consumer price of category j per 1 percent increase in the border price of commodity c, operating through the direct import content alone. The indirect matrix $\Delta^{\text{Indirect}}_{j,c}$ captures the additional pass-through through the input-output network.

Constructing equations (81) and (82) from data requires mapping each of the six model objects, $\boldsymbol{\Theta}^d$, $\boldsymbol{\Theta}^*$, $\boldsymbol{\Omega}^d$, $\boldsymbol{\Omega}^*$, $\boldsymbol{\Upsilon}^r$, and Υ, to an observable counterpart in the BEA accounts.

The direct import share matrix, $\boldsymbol{\Theta}^*$ The entry Θ^*_{en} is the share of NIPA expenditure category e's domestic supply sourced directly from imported BEA commodity n.[6] We construct it in two steps. First, for each BEA commodity, we compute the direct import penetration ratio as the ratio of import value, taken from the Import matrix, to total commodity supply, taken from the commodity output totals in the BEA Input-Output accounts. Second, we

[5] In all cases, we use the "after redefinition" versions of the tables.

[6] In this paper, we use the terms "commodity" and "final good" interchangeably—both are the direct inputs into the retail sectors, but "commodity" is the term used in the BEA data.

aggregate across the commodities contributing to each NIPA category using the PCE bridge weights, which record the commodity composition of each expenditure category at basic prices. The resulting $E \times N$ matrix relates each NIPA category to each import commodity group, capturing the foreign sourcing intensity of final consumer expenditure before distribution margins are applied.

The domestic input share matrix, $\mathbf{\Theta}^d$ The entry Θ^d_{en} records the share of the marginal cost of NIPA expenditure category e attributable to the domestically produced output of BEA commodity n, prior to the indirect supply-chain amplification. We construct it from the domestic portion of the BEA Use table, specifically total intermediate purchases net of the Import use matrix, normalized at the industry level by gross output. We convert from industry to commodity basis using the BEA Make matrix, which records each industry's production shares across commodities, and then map to the E NIPA categories through the PCE bridge.

The input coefficient matrices, $\mathbf{\Omega}^d$ and $\mathbf{\Omega}^*$ Both matrices are constructed from the BEA Use table, normalized by gross output, and converted from industry to commodity basis via the Make matrix.[7] Under both markup assumptions they are always gross-output-normalized, with the level/log distinction entering through $\mathbf{\Upsilon}^d$ rather than through the $\mathbf{\Omega}^d$ and $\mathbf{\Omega}^*$ matrices themselves. $\mathbf{\Omega}^d$ is the $N \times N$ domestic input coefficient matrix: entry Ω^d_{kl} records the value of domestic commodity l used as an intermediate per dollar of gross output of producing commodity k, obtained by subtracting the BEA Import use matrix from the total Use table to isolate domestic intermediate flows. $\mathbf{\Omega}^*$ is the $N \times N$ import intermediate use matrix: entry Ω^*_{kl} records the value of imported commodity l used as an intermediate per dollar of gross output of producing commodity k, obtained from the import rows of the Use table. Together, they govern equation (82): $\mathbf{\Omega}^d$ drives the Leontief inverse that propagates cost shocks round-by-round through the domestic production network, while $\mathbf{\Omega}^*$ identifies the exposure of each domestic production stage to foreign cost shocks, generating the indirect pass-through channel even for expenditure categories that import little directly.

The retail margin matrix, $\mathbf{\Upsilon}^r$ The diagonal entry Υ^r_{ee} converts the producer-price sensitivity of expenditure category e into a purchaser-price sensitivity. Under level pass-through, the retail margin matrix $\mathbf{\Upsilon}^r$ is the identity matrix. Under log pass-through, it is

[7] Note that all theoretical matrices follow the general convention in the literature to have outputs in the row and inputs in the columns of the factor share matrix $\mathbf{\Omega}^d$. The BEA instead specifies inputs in the rows and outputs in the columns of the use tables. In addition, the BEA distinguishes between commodity and industry in its input–output table.

built from the PCE bridge by netting employee compensation from the value added of each distribution sector (since distribution margins scale with variable cost under log pricing) and then expressing the remaining producer value as a share of total purchaser expenditure. Categories with large distribution margins, such as food at home and motor vehicles, carry smaller $\boldsymbol{\Upsilon}^r$ entries under log pass-through, so that a given producer-price movement generates a proportionally smaller consumer price change. The two assumptions diverge most sharply in high-margin categories.

The markup adjustment matrix, $\boldsymbol{\Upsilon}^d$ The matrix $\boldsymbol{\Upsilon}^d$ is diagonal with entry Υ_{nn} for each BEA producing commodity n. Under level pass-through, $\boldsymbol{\Upsilon}^d$ is the identity matrix, and equations (81) and (82) reduce to the standard Leontief pass-through formulas with input coefficients normalized by gross output. Under log pass-through, Υ^d_{nn} is the ratio of gross output to total variable cost for commodity n,[8] where both are obtained from the BEA Use Table. This markup factor converts gross-output-normalized coefficients into variable-cost-normalized equivalents, amplifying the sensitivity of high-markup sectors.

We compute sensitivities under both markup assumptions. Later, we also estimate commodity-level scaling of $\boldsymbol{\Upsilon}^d$ as a logistic-linear function of industry concentration and supply-chain characteristics, with coefficients recovered by minimum distance, and we evaluate sensitivities over a continuous grid of Υ values spanning the level and log extremes to characterize how aggregate pass-through and its channel decomposition respond to the markup assumption.

The total sensitivity matrix $\Delta = \Delta^{\text{Direct}} + \Delta^{\text{Indirect}}$ is computed for each year for which we have input-output data (2017 and 2024), separately under the log and level pass-through assumptions. Later, we additionally vary the competition-channel parameter ρ continuously over $[0, 1]$ and characterize how the predicted tariff impact on each industry responds to ρ under both markup assumptions. Additionally, we estimate commodity-level ρ values.

Our import sensitivities under the assumption of pass-through in levels can be reinterpreted as the share of US PCE that is imported.[9] Under level pass-through and 2017 IO weights, the PCE-weighted aggregate import sensitivity is approximately 11 percentage points, of which roughly 6 percentage points reflect direct import content and 5 percentage points reflect indirect supply-chain exposure.[10]

[8]Total variable costs are the sum of salaries and cost of materials.

[9]Under the pass-through in levels assumption, markups can be thought of as a domestic factor of production.

[10]This is consistent with the estimates of Hale et al. (2019).

4.2 Additional Data

Tariff Data We calculate effective monthly tariff rates by dividing the total duties paid in month m on HS-10 product-level code p coming from country of origin c by the total imports of that product in that month from that country:

$$\tau_{cp,m} = \frac{\text{Duties Paid}_{cp,m}}{\text{Imports}_{cp,m}} \tag{83}$$

The data behind this calculation comes from publicly-available US Census data. Calculating this effective rate accounts for exclusions and exemptions.

We aggregate tariff shocks from the HS-10 level to BEA commodity categories using country-level import value shares from the Census Bureau's Foreign Trade Statistics (FTS), measured in 2017 for the 2018–2019 episode and in 2024 for the 2025 episode. The FTS reports US merchandise imports by HS-10 product and country of origin at monthly frequency; we collapse these to annual import-share weights at the BEA commodity-country level. We hold these shares fixed over the episode in order to isolate price effects from shifts in import composition. This results in a monthly panel of import-weighted tariff shocks at the BEA commodity level, separately by source country, that serves as the primary input to the marginal cost calculations in Section 5.

Concentration Data We measure industry concentration using the Herfindahl-Hirschman Index (HHI) at the BEA commodity level, drawing on two complementary sources. The first is the Census Bureau's Economic Census, which reports establishment-level concentration statistics for manufacturing, retail trade, and selected service industries, disaggregated by NAICS industry code, for benchmark years including 2012 and 2017. The second is the National Establishment Time-Series (NETS) database (henceforth NETS), which provides annual establishment-level revenue and employment data derived from Dun & Bradstreet records, enabling us to construct sales-weighted firm-level HHI at annual frequency for a broader set of industries and years. We map both sources to BEA commodity codes via a three-layer concordance connecting NAICS 2022 to NAICS 2017 to BEA codes.[11] We compute HHI using 2017 data for the 2018 episode and the most recent available year for the 2025 episode.

[11]We resolve many-to-many links by assigning each industry code to its connected component in the bipartite classification graph. In commodity categories where Census concentration data are unavailable, we impute HHI from NETS firm-level shares.

PCE and PPI Price Data For our pass-through estimations, we combine two monthly price series. For consumer prices, we use the Bureau of Labor Statistics (BLS) Personal Consumption Expenditures price indices at the NIPA expenditure-category level (henceforth PCE prices). For producer prices, we use the BLS Producer Price Index (PPI), which measures average changes in prices received by domestic producers at the first stage of commercial transaction. We match PPI series to BEA commodity codes via a BLS-BEA concordance that links NAICS-based PPI industry codes to the BEA 402-commodity classification, using connected components to resolve many-to-many NAICS-BEA overlaps.[12]

5 Impact of Tariffs on Marginal Costs

In this section, we use our methodology to estimate the inflationary effect of specific tariff episodes—the 2018–19 tariffs and the 2025 tariffs (through December 2025). In doing so, we also are able to decompose the cumulated headline PCE impact along two dimensions: the NIPA expenditure category of final consumption and the source country of the imported good.

In order to build our estimates, we first build a concordance matrix between product codes available from border declarations and BEA commodity codes.[13] We define F as a matrix containing the share of country-product specific value of import for each BEA commodity. Each column represents a country-HTSUS combination, while each row represents one of the 402 BEA commodities. We then use a country-product-specific vector of tariffs τ_t at any given time to predict the aggregate tariff effect as

$$d \log \mathbf{p} = \mathbf{1}^\top \boldsymbol{W}^c (\Delta^{\text{Direct}} + \Delta^{\text{Indirect}}) \boldsymbol{F} \boldsymbol{\tau}, \tag{84}$$

where $\boldsymbol{W}^c$ is a diagonal matrix with diagonal values defined as the share of each PCE category in total PCE. Aggregating the importance-weighted sensitivity by imported commodity (that is, pre-multiplying by $\mathbf{1}^\top$) gives the aggregate sensitivity of personal consumption expenditures to each imported commodity. The foreign trade concordance matrix $\boldsymbol{F}$ offers a convenient bridge to simulate the effects of tariffs or border shocks τ starting at the most disaggregated level possible.[14] Such level of disaggregation in the tariff scenario τ allows us to account for exemptions, which are usually set at the 10-digit HTSUS-country level by

[12] Where PPI series contain gaps of two months or fewer, for example during the 2018–2019 partial government shutdown, we interpolate using adjacent observations; longer gaps are left as missing.

[13] We use the US Census Bureau Foreign Trade import data, from which we calculate the total dollar value imported into the United States of each HTSUS commodity code from each foreign country in each year.

[14] As in the previous sections, all these matrices are year-specific, but we omit the time subscript for clarity.

law. Our methodology allows for an exempted category to have both a direct (null) effect on consumption purchases via Δ^{Direct} and an indirect (null) effect via the domestic supply chain.[15]

Our estimates assume that the full burden of the tariffs is paid by US importers, consistent with what was observed with the tariff increases imposed in 2018 (Fajgelbaum et al., 2020; Amiti, Redding and Weinstein, 2019; Cavallo et al., 2021). Throughout this section, we use the 2017 IO structure for the 2018 episode and the 2024 IO structure for the 2025 episode. We set $\rho = 1$ as the baseline; we later examine sensitivity to ρ.

We estimate, using pass-through in levels, that headline PCE rose 0.20 percentage points in the 2018 episode and 0.81 percentage points in the 2025 episode. Using pass-through in logs increases the estimates to 0.36 percentage points in the 2018 episode and 1.42 percentage points in the 2025 episode.

Figure 1 plots expected increases in headline PCE due to each month's tariff increase.[16] We can see that the 2025 tariff increases had a substantially higher predicted impact, due to their larger size. The figure helps us understand where the variation in our shock variables used in the next section come from.

Figure 2 decomposes the overall PCE impact of each episode into source country. We can see that in the 2018–2019 tariff episode, almost the entire increase in headline PCE came from tariffs on China. The 2025 tariff episode, however, was much broader in that it targeted a large variety of countries, although China received the highest tariffs. In line with this the increase in tariffs on headline PCE was more distributed across a variety of countries. China remained the largest contributor, but India and Vietnam, as well as a host of other countries, were also important contributors.

Figures 3 and 4 decompose the overall PCE impact of each episode into NIPA categories. The size of the bars in these figures reflect both the importance of a country to US PCE and the effect of tariffs on that category. In line with the fact that China is a particularly large exporter of finished apparel to the US, the garments category stands out as the largest category in both episodes, and it is dominated by the direct import channel. However, there is considerable distribution across categories in both episodes. Additionally, the figures highlight how affects are larger under pass-through in logs, primarily due to an increased indirect affect. The intuition is that an upstream marginal cost increase due to a tariff is amplified along the supply chain, as each producer charges a constant-percentage markup on

[15] For example, say that mechanical components are exempted from a certain tariff but steel is not exempted. This methodology would account for the final price increase of mechanical components due to the increased price of steel, to the extent that steel is used to produce mechanical components.

[16] Note, this is a different exercise than the expected inflation impact in that month, since we do not model the fact that inflation may appear with a lag.

Figure 1: Monthly Tariff Shock: 2018–19 vs 2025

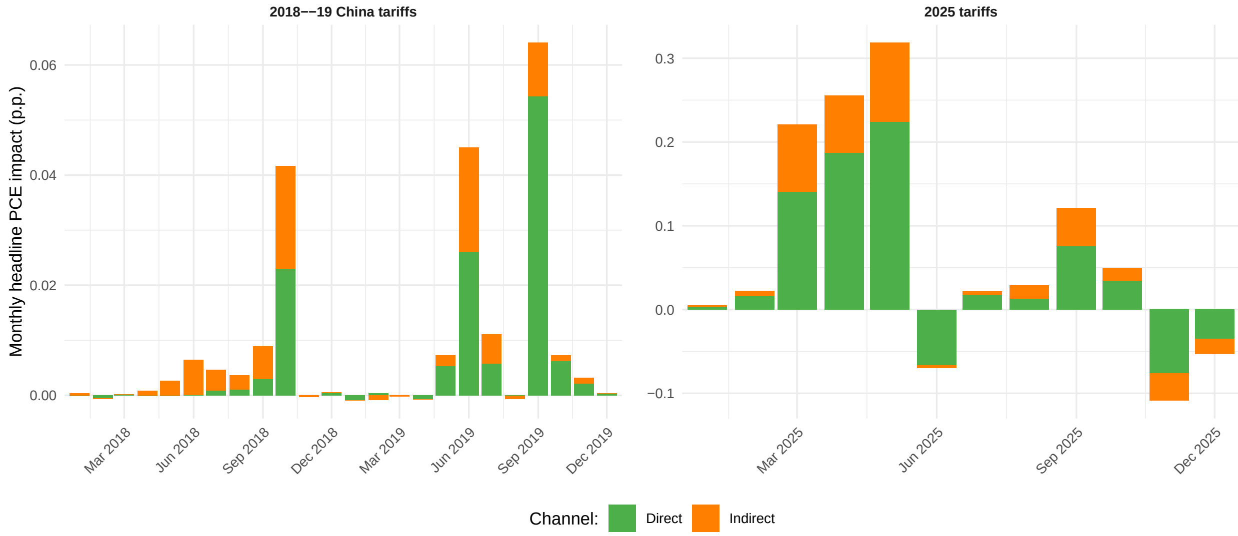

Notes: This figure shows the predicted headline-PCE impact by month of tariff change under pass-through in levels at $\rho = 1$. Each bar is stacked into a direct and an indirect channel. The left panel shows the 2018–19 tariff episode. The right panel shows the 2025 tariff episode (through December 2025). Y-axes are independent. *Sources*: BEA, BLS, authors' calculations.

top of each respective input purchase.

Figures 3 and 4 show the expected price increase of coarse expenditure categories due to tariff increases in each episode. The figures show coarse expenditure categories to elucidate patterns in the data; however, our baseline empirical analysis considers the expected effect of narrow PCE categories, split by tariff changes each month. This is the unit of analysis we examine in the next section, when we turn to understanding how the *expected* price increases at the expenditure category-month level, based solely off our model-implied IO structure and calculated tariff changes, fed through to *actual* price changes, allowing us to test the validity of our model.

Figure 2: Source-Country Decomposition of Headline PCE Impact

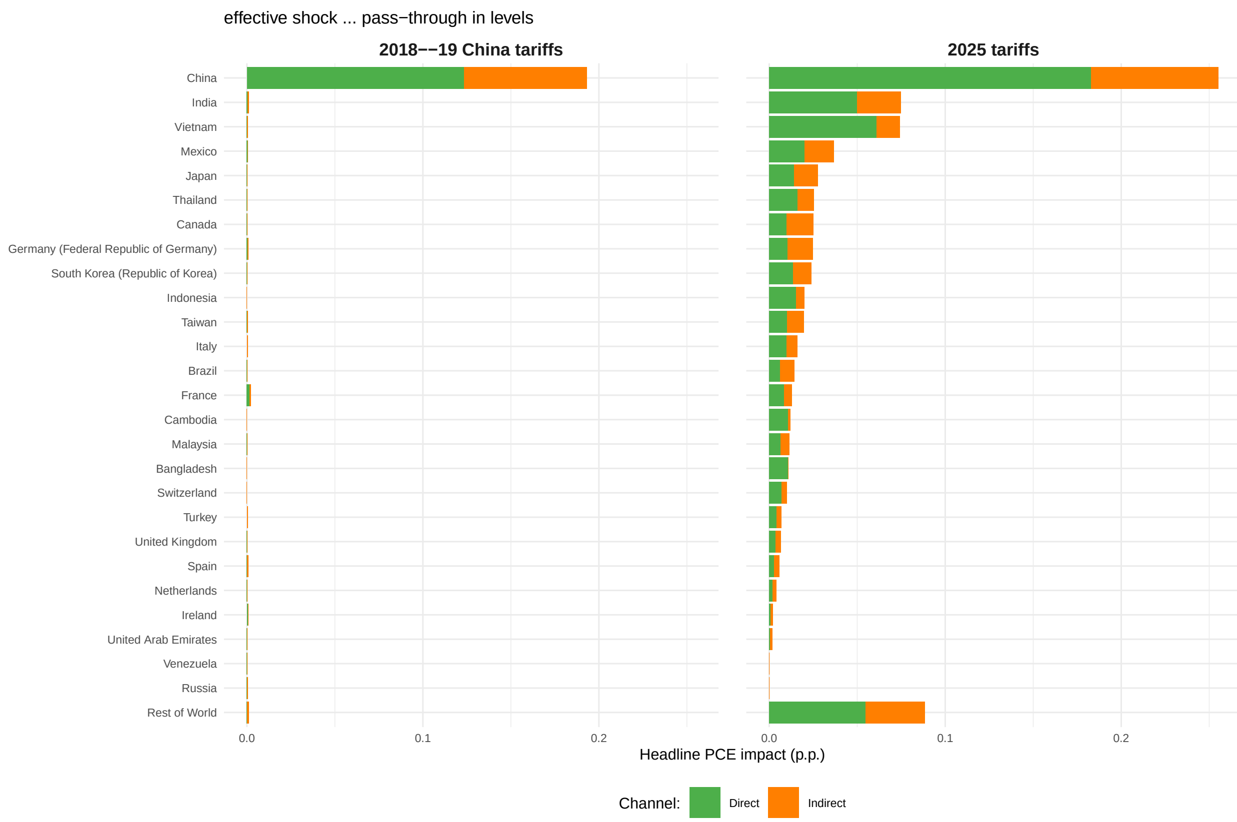

Notes: This figure decomposes the overall headline PCE impact of each tariff episode by source countries. The left panel shows the 2018–19 tariff episode; the right panel shows the 2025 episode. The countries shown are the union of the top 20 source countries of each episode, plus a residual "Rest of World." Countries share the same vertical axis and are ordered by the right panel (2025). Bars are stacked into a direct and an indirect channel. The x-axis is the cumulated contribution to headline PCE inflation, in percentage points.

Figure 3: NIPA Decomposition of Headline PCE Impact—2018 Episode

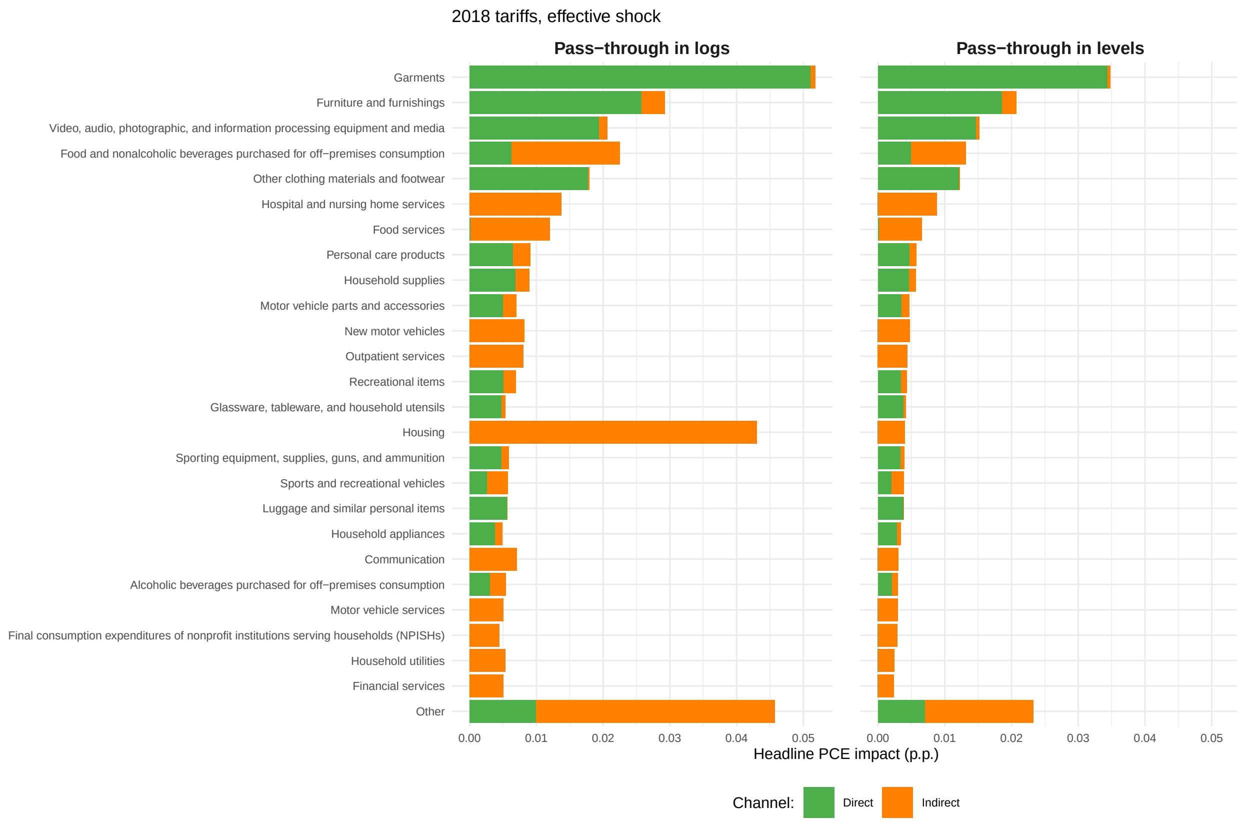

Notes: This figure decomposes the overall headline PCE impact of the 2018–19 tariffs. Each row is one coarse NIPA expenditure category (level D, ≈ 95 groups) with very small categories collapsed into "Other." The left panel shows results under the assumption of pass-through in logs; the right panel shows results under the assumption of pass-through in levels. Categories share the same vertical axis and are ordered by the right panel (levels). Bars are stacked into a direct and an indirect channel. The x-axis is the cumulated contribution to headline PCE inflation, in percentage points.

Figure 4: NIPA Decomposition of Headline PCE Impact—2025 Episode

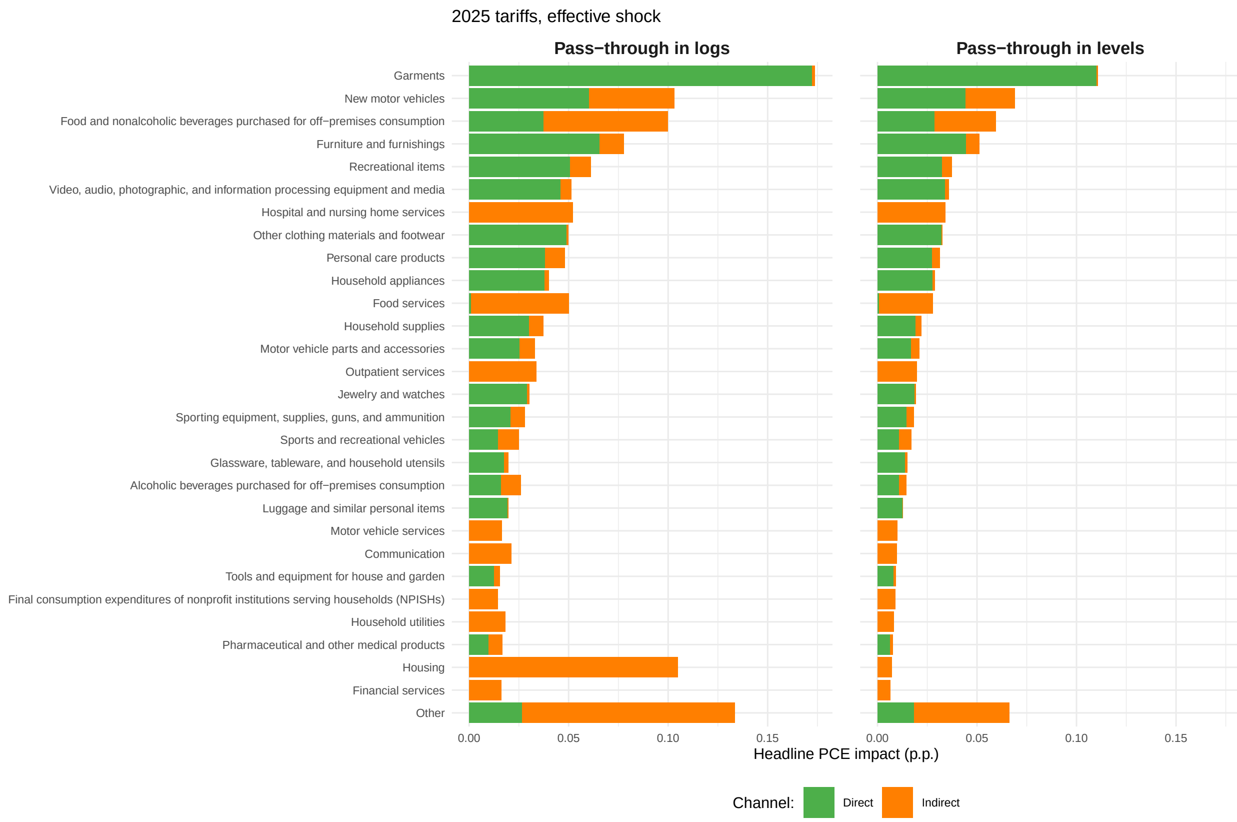

Notes: This figure decomposes the overall headline PCE impact of the 2025 tariffs. Each row is one coarse NIPA expenditure category (level D, ≈ 95 groups) with very small categories collapsed into "Other." The left panel shows results under the assumption of pass-through in logs; the right panel shows results under the assumption of pass-through in levels. Categories share the same vertical axis and are ordered by the right panel (levels). Bars are stacked into a direct and an indirect channel. The x-axis is the cumulated contribution to headline PCE inflation, in percentage points.

6 Results

6.1 Pass-through of marginal costs

Our model is static in nature and therefore it predicts the long term effect of a tariff shock at any given time. We follow Dube et al. (2025) and run a local projection diff-in-diff estimation to study the dynamic convergence toward the predicted effects of our model under different assumptions. For each horizon $h \in \{1, \ldots, 7\}$ we estimate, separately by tariff episode, the regression

$$\Delta_h \log p_{i,t+h} \;=\; \alpha_i + \delta_t + \beta_h \, s_{i,t} + \sum_{k=-6, k\neq 0}^{6} \gamma_{h,k} \, s_{i,t-k} + \lambda_h \, \Delta \log w_{i,t-1} + \varepsilon_{i,t+h}, \qquad (85)$$

where $\Delta_h \log p_{i,t+h} \equiv \log p_{i,t+h} - \log p_{i,t-1}$ is the cumulated log-price change from $t-1$ to $t+h$, $s_{i,t}$ is the predicted tariff shock for item (or industry) i, α_i and δ_t are item and date fixed effects, and $\Delta \log w_{i,t-1}$ is one lag of wage growth. The sample for the 2018-19 tariff episode includes all months between 2016 and end of 2019. The sample for 2025 includes all months between 2024 and February 2026. We study both PPI and PCE prices.

The figures in this section show the estimates β_h across episodes and under full-pass through in levels and logs. Standard errors are clustered by item; 95% confidence intervals are shaded. Each figure places the 2018 and 2025 episodes side-by-side; the two panels are estimated from disjoint samples. We report passthrough in *level* (constant-dollar markup) and in *logs* (constant-percentage markup) on the same panels.

The PCE estimates in Figure 5 show that the model captures an important part of the cross-sectional incidence of tariff shocks, but also that the fit differs across episodes. For the 2018 tariffs, the dynamic response converges toward full pass-through within roughly five months and then exceeds one at longer horizons. The natural interpretation is that the model-implied shock is somewhat too small relative to the realized PCE response. For the 2025 tariffs, the sign and ranking of the predicted exposure remain informative, but the long-run coefficients are substantially below one, implying that the model-implied shock is too large relative to observed inflation.

Neither markup assumption uniformly dominates. The constant-percent-markup specification is closer to full pass-through in 2018, while the constant-dollar-markup specification is closer in 2025. Yet these differences are small relative to the sampling uncertainty. Taken together, the estimates provide little basis for sharply distinguishing the two markup assumptions, but they suggest a consistent cross-episode pattern: the baseline model tends to understate the 2018 response and overstate the 2025 response.

Figure 5: PCE Total Tariff Pass-through — Level vs Logs

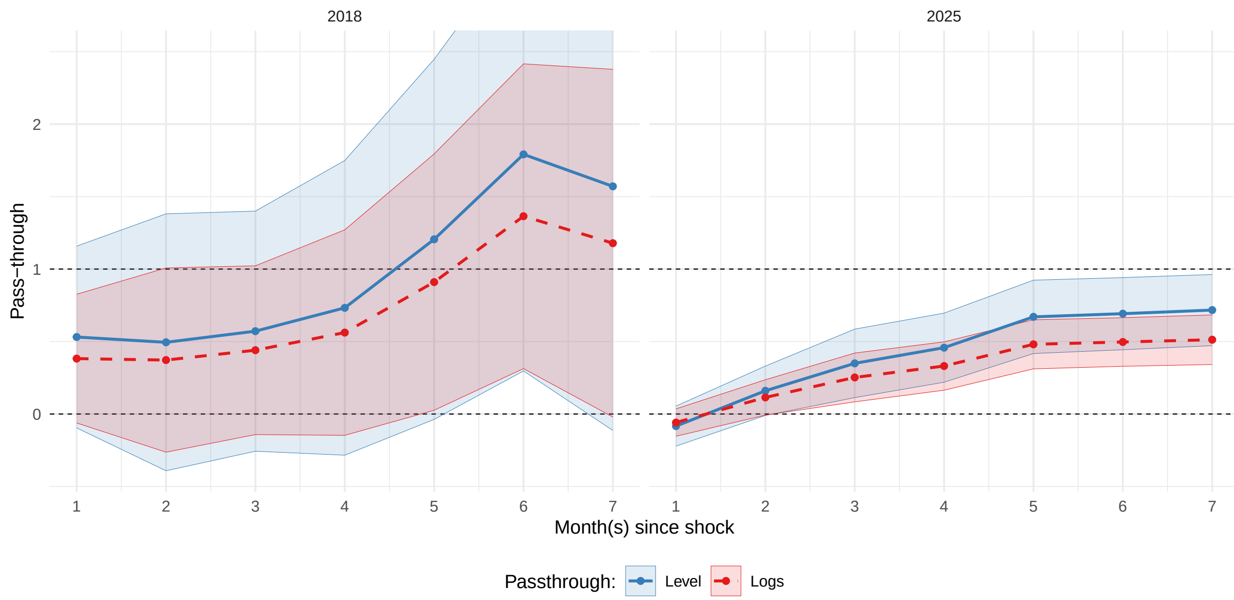

Notes: Pass-through of the predicted total (direct + indirect) tariff shock to PCE log-price growth, under the level (constant-dollar markup) and logs (constant-percentage markup) assumptions. The 2018 and 2025 episodes share a common y-axis capped at 2.5; the black dashed line marks full pass-through (pass-through = 1). Specification: equation (85). *Sources*: BEA, BLS, authors' calculations.

The comparison in Figure 6 shows that the PCE pass-through estimates are driven almost entirely by the direct component of tariff exposure. Replacing total exposure with direct exposure produces very similar dynamic coefficients in both episodes and under both markup assumptions. The result should not be read as evidence that supply-chain exposure is small. It instead indicates that the PCE specification has limited independent variation with which to separately identify the indirect channel.

There is a simple reason for this. The direct component preserves the sharp cross-sectional differences in final import content across expenditure categories. The indirect component, by contrast, is filtered through the domestic input-output network and then mapped into aggregate NIPA categories. This procedure is economically important for measuring total exposure, but it also compresses the relevant cross-sectional variation. For this reason, producer-price data, which are closer to the domestic production network and available at a less aggregated industry level, provide a more suitable setting for assessing the indirect channel.

The PPI estimates in Figure 7 are useful because they evaluate the model at the stage where tariff-induced input-cost changes first affect domestic production. The same asymmetry across episodes is still visible: relative to the data, the model fits the 2025 episode better

Figure 6: PCE Pass-through — Total and Direct Shocks, Level vs Logs

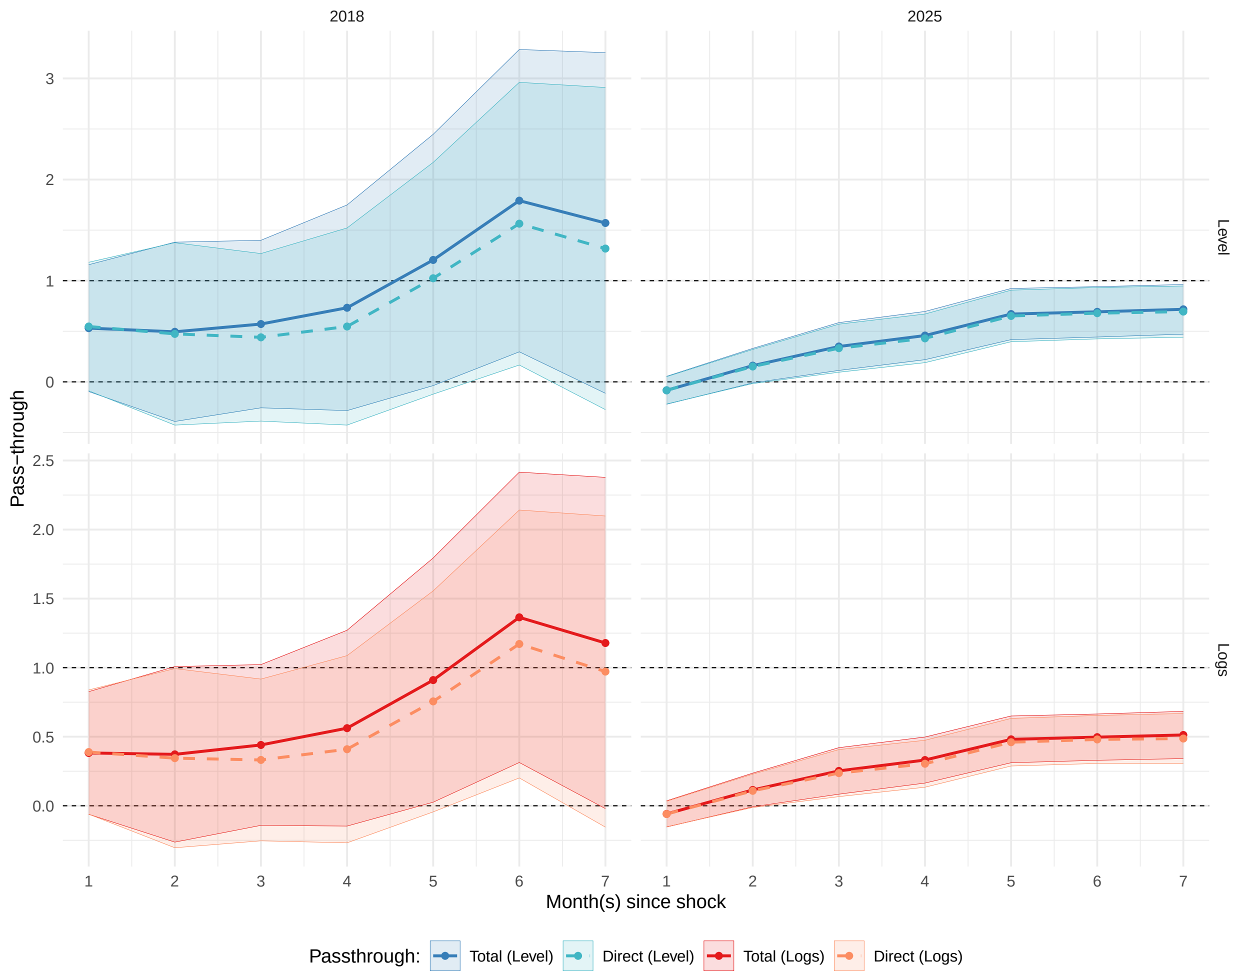

Notes: Pass-through of the predicted *total* and *direct* tariff shocks to PCE log-price growth, estimated from separate regressions and overlaid. Rows: level (constant-dollar markup, cool hues) and logs (constant-percentage markup, warm hues). Columns: 2018 and 2025 episodes. Within each panel, the Total shock is drawn as a solid line and the Direct shock as a dashed line in a companion hue. Specification: equation (85). *Sources*: BEA, BLS, authors' calculations.

than the 2018 episode. The performance of the PPI model for 2025 at longer-horizon are close to one for 2025, suggesting that the model's predicted marginal-cost exposure is a good approximation to the producer-price response in this case.

The level and log specifications are also closer to each other in the PPI regressions. This contrasts with the PCE results, where retail margins amplify the empirical relevance of the markup assumption. For producer prices, the choice between constant-dollar and constant-percent markups changes the exposure measure less, and the data do not provide a sharp distinction between the two.

Appendix E examines the sensitivity of the local-projection results to alternative samples,

Figure 7: PPI Tariff Pass-through — Level vs Logs

Notes: Pass-through of the predicted tariff shock to PPI log-price growth, under the level (constant-dollar markup) and logs (constant-percentage markup) assumptions. The 2018 and 2025 episodes share a common y-axis capped at 2.5; the black dashed line marks full pass-through (pass-through = 1). Specification: equation (85), with item i denoting a BEA industry. *Sources*: BEA, BLS, authors' calculations.

weights, and shock construction. We first restrict the PCE sample to goods categories, where tariff exposure is more direct. We then compare the unweighted estimates with specifications weighted by inverse price volatility and by economic size. Finally, we construct a detail-level by-region shock that allows the source-country composition of imports to differ across commodities, rather than imposing a homogeneous import-share profile. The estimates are not identical across specifications, but they support the same qualitative interpretation as the baseline results. In particular, the consumer-price evidence remains largely tied to direct exposure, while the producer-price evidence is more informative about input-output propagation and continues to display the same cross-episode asymmetry.

6.2 Sensitivity to the conditional pass-through ρ

Several studies find that 2018 tariffs had a protectionist effect on US domestic prices (Amiti, Redding and Weinstein, 2019; Flaaen and Pierce, 2024). We modeled this channel by allowing for $0 < \rho < 1$ and we evaluate whether such channel can explain the heterogeneity in pass-through estimates observed between 2018 and 2025.

In principle, different tariff episodes can produce higher or lower than marginal costs effects on an industry. Recall equation (57) can be rewritten as

$$d \log p_i^d = \nu_i^d \; \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[d \log \mathbf{mc}_i] + (1 - \tilde{\rho}_i)\big[d \log p_{Fi} - \nu_i^d \; \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[d \log \mathbf{mc}_i]\big] \;, \tag{86}$$

Allowing for partial pass-through of marginal costs via ρ can in principle lower or increase the impact of tariffs on an industry, according to the sign of $d \log p_{Fi} - \nu_i^d \; \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[d \log \mathbf{mc}_i]$. And the deviation between the own industry tariff price increase and the tariff-implied marginal cost increase depends both on the distribution of the tariff shock across industries and the share of imports of each sector from all other sectors.

Figure 8: ρ-sensitivity of PPI impact — Selected industries

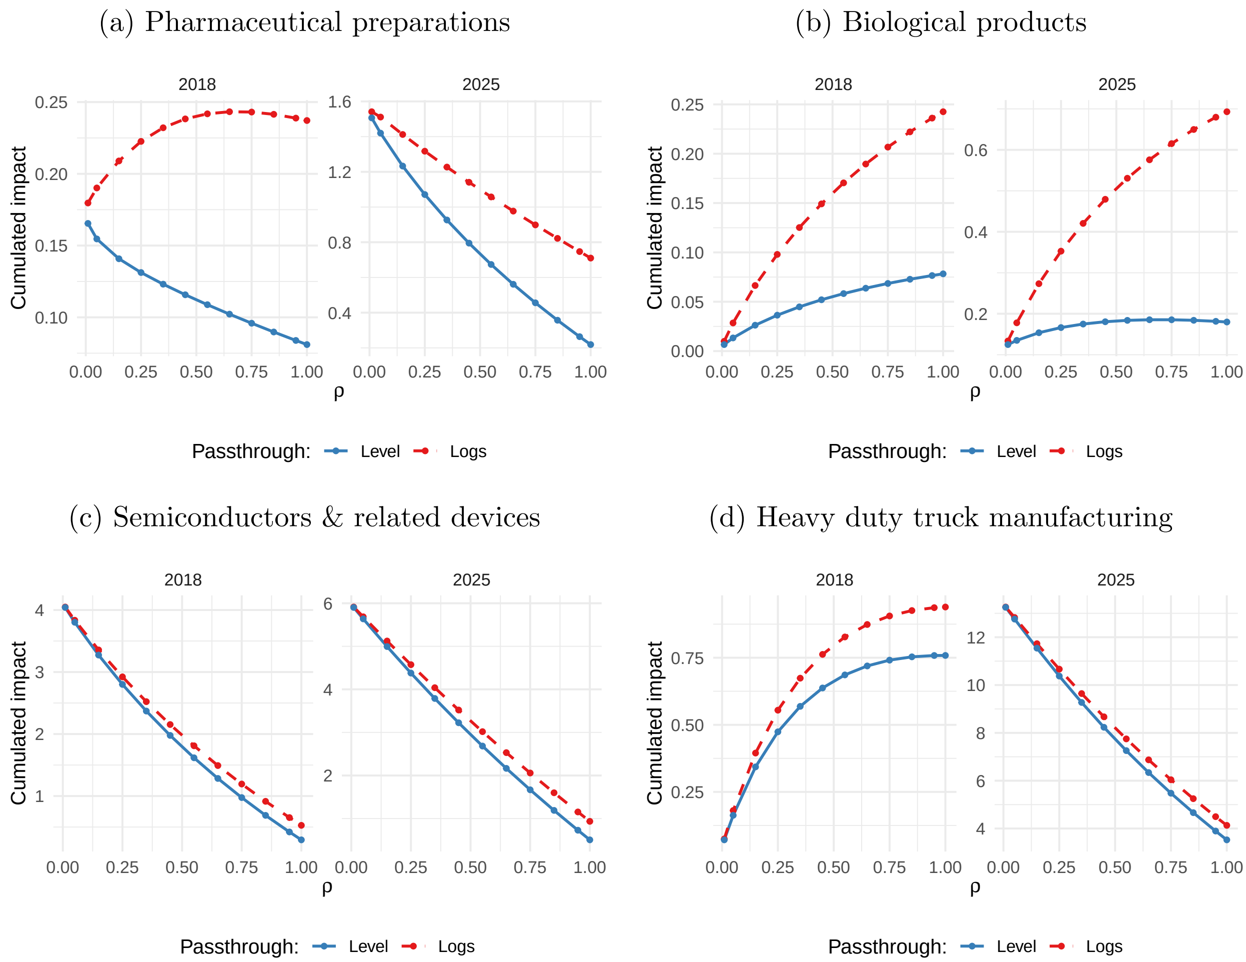

Notes: Cumulated PPI impact as a function of ρ, for four manufacturing commodities of interest in the 2018 and 2025 episodes. Each panel has an independent y-axis and overlays the level (solid) and logs (dashed) assumptions. Homogeneous (non-byregion) shock. *Sources*: BEA, BLS, authors' calculations.

Figure 8 illustrates this point. Reading the four panels through equation (57), the slope of the cumulated impact in ρ is controlled by the balance between the own-tariff term

$(1-\tilde{\rho}_i)\, d\log p_{Fi}$ and the cost-channel term $\tilde{\rho}_i\, \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}[d\log \mathbf{mc}_i]$, and by how much s_{Fi} lets $\tilde{\rho}_i$ vary with ρ. Table D.1 in Appendix D reports these four ingredients for each commodity and episode. For biological products the foreign competitors share s_{Fi} is large (≈ 0.50) but the own tariff is small in both episodes, so $s_{Fi}\, d\log p_{Fi}$ never matters and the cost channel — which grows from 8 to 18 pp in levels — carries the curve upward identically in shape across 2018 and 2025. For semiconductors $s_{Fi} \approx 0.16$ combines with a very large own tariff (3.6 pp in 2018, 5.1 pp in 2025) and a comparably large cost shock. On net the direct-foreign contribution $s_{Fi}\, d\log p_{Fi}$ dominates the marginal cost shock. Therefore, as oligopolistic competitions matters more (ρ decreases) the impact increases in both episodes. For heavy-duty trucks s_{Fi} is small, but the own tariff $d\log p_{Fi}$ moves from 0.05 pp in 2018 to 15.4 pp in 2025; the 2018 curve is cost-channel-dominated and rises in ρ, while the 2025 curve is dominated by the very large $d\log p_{Fi}$ (together with a cost shock that itself grows roughly fivefold) and falls sharply in ρ. For pharmaceuticals the foreign competitors share s_{Fi} is essentially zero, which silences the competitor channel entirely; what differs across episodes is the size of the upstream cost shock — small enough in 2018 that the logs curve humps through the pharma markup multiplier in (52), and large enough in 2025 that both levels and logs decline monotonically.

But what kind of relationship between ρ and the tariff impact is more representative of the US industries under the two tariff scenarios? In practice, most industries look more like semiconductors in Figure 8. Figure 9 pools all manufacturing commodities, normalizing each path by its value at $\rho = 1$; the median is overlaid with a gradient of quantile ribbons (inner bands span 45–55%, 35–65%, 25–75%, 15–85%, and 5–95% of the cross-commodity distribution). Clearly, regardless of the tariff episode, a stronger oligopoly channel leads to *higher* tariff impact. The reason why the curve is steeper in 2025 is that tariffs were more spread out across manufacturing industries, making the oligopolistic channel stronger.

6.3 Revealed pass-through ρ under the oligopolistic channel

In this subsection we use the local-projection evidence to recover the conditional pass-through ρ_j that governs the oligopolistic competition channel. Recall from Section 3.6.2 that ρ_j is not directly observable: it is a reduced-form object that depends on industry structure and elasticities of substitution. Conditional on $\mathbb{C}^* = 0$, however, the domestic-price response to tariff shocks only depends on ρ and on observable input-output shares, so the pass-through parameter can be estimated as a "revealed" parameter from the LP at horizon $h = 7$. We estimate ρ_j across industries in two ways: under parametrization, that is using industry characteristics that can best predict the ρ_j that fits the inflation responses, and a

Figure 9: Cross-commodity ρ-sensitivity of PPI impact — Level vs Logs

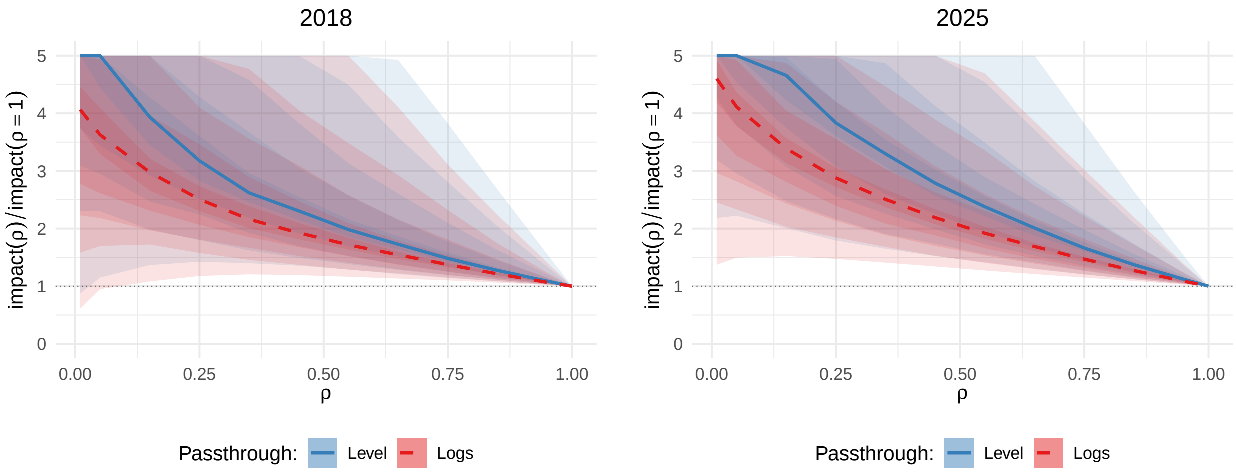

(a) 2018 episode (b) 2025 episode

Notes: Cumulated PPI impact at each ρ, divided by the impact at $\rho = 1$, across all 232 BEA manufacturing commodities. Ratios are clipped to $[0, 5]$. Lines are cross-commodity medians (Level solid, Logs dashed); shaded bands are stacked inter-quantile ribbons at 45–55%, 35–65%, 25–75%, 15–85%, and 5–95%, fading away from the median. Homogeneous (non-byregion) shock. *Sources*: BEA, BLS, authors' calculations.

non-parametric estimation.

We parameterize ρ_j as a logistic-linear function of three observable commodity-level characteristics — the Herfindahl index HHI_j, producer markup μ_j = gross output/variable cost, and imported-input share s_j^*:

$$\rho_j(\alpha) = \Lambda\big(\alpha_{HHI}\, HHI_j^c + \alpha_\mu\, \mu_j^c + \alpha_{s^*}\, s_j^{*c}\big), \qquad \Lambda(x) = \frac{1}{1+e^{-x}}, \tag{87}$$

where the superscript c denotes de-meaning over active manufacturing commodities. This reduces 222 per-commodity parameters (all manufacturing commodities present in the offical IO table) to 3 coefficients and guarantees $\rho_j \in (0, 1)$. Inactive commodities (non-manufacturing or non-importing) are pinned at $\rho_j = 1$, consistent with the rest of the paper.

The complementary nonparametric benchmarkestimates a step-function specification with one $\rho_g \in (0, 1)$ per active NAICS-3 manufacturing subsector ($G \leq 21$). Both parameterizations use the same estimator; the nonparametric version serves as a test of whether the three observables are sufficient to summarize the between-commodity heterogeneity.

Restricted profile least squares. For a candidate α (or per-group vector $\{\rho_g\}$), we construct the model-implied per-period exposure $m_{it}(\alpha)$ and impose the structural normalization that its contemporaneous LP coefficient equals one. Substituting m_{it} into the left-hand side of the LP at $h = 7$ and concentrating out the nuisance coefficients on leads, lags, fixed effects, and wages gives the objective

$$\widehat{\alpha} = \arg\min_{\alpha} \sum_{i,t} \widehat{\varepsilon}_{it}(\alpha)^2, \tag{88}$$

where $\widehat{\varepsilon}_{it}(\alpha)$ are the residuals from

$$y_{i,t+7} - y_{i,t-1} - m_{it}(\alpha) = \sum_{\ell \in \mathcal{L},\, \ell \neq 0} \psi_\ell\, m_{i,t+\ell}(\alpha) + \alpha_i + \delta_t + \eta \log w_{it-1} + \varepsilon_{it}(\alpha),$$

with $\mathcal{L} = \{-6, \ldots, -1, +1, \ldots, +6\}$, item and date fixed effects, and lagged log-wage as the only linear control. The LP sample retains industries with direct or input-output exposure to at least one active manufacturing commodity. As a diagnostic, after convergence we relax the $\beta_0 = 1$ restriction and rerun the LP with a free contemporaneous coefficient $\widehat{\lambda}$; deviations from $\widehat{\lambda} = 1$ indicate that the restriction is binding against the data.

Figure 10 reports the commodity-level distribution of the implied $\widehat{\rho}_j$ from the parametric specification, separately by (episode, assumption). The distributions make the instability across episodes visible: in 2018, $\widehat{\rho}_j$ is broadly dispersed over (0.2, 0.9) with substantial mass below 0.5, while in 2025 the distribution concentrates tightly near 1. A similar pattern is evident from the nonparametric estimation in Figure F.1 in appendix.

The skewness of the 2018 estimates toward zero is intuitive. In 2018, the baseline full pass-through model ($\rho = \mathbf{I}$, $\mathbb{C}^* = 0$) *undershoots* the observed PPI response to tariffs — so the estimator lowers ρ in order to raise $(1 - \rho)$, shifting weight onto direct foreign-price pass-through and boosting the model-implied impact to match the data. In 2025, the baseline is already close to the observed response and the estimator has little work to do, pulling $\widehat{\rho}_j$ toward 1 for most commodities. Table F.1 in appendix also shows that the parametric estimation does not show full consistency over time e.g. by having the sign on the HHI coefficient flip sign between 2018 and 2025.

Taken together, the evidence suggests that ρ is not behaving as a stable structural primitive in these data. It absorbs a large part of the 2018 impact heterogeneity — where the baseline is furthest from the data — but it does not improve the fit nor the interpretation of the 2025 episode. This lack of cross-episode invariance is a caveat for this channel: if ρ were a deep parameter of market structure, it should be comparably identified in both shock episodes. Instead, the estimator behaves as a residual-absorbing device, picking up the larger signal in 2018 and almost nothing in 2025.

Figure 10: Implied $\widehat{\rho}_j$ across the 222 active manufacturing commodities from the parametric RPLS estimator.

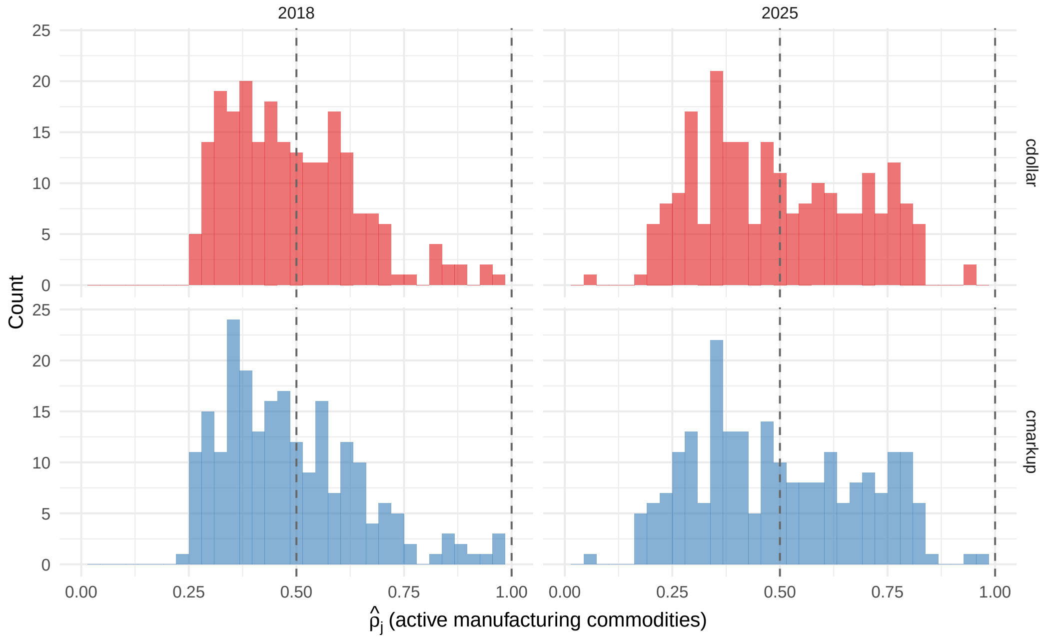

Notes: Rows are markup assumptions; columns are episodes. Dashed lines at $\widehat{\rho} = 0.5$ and $\widehat{\rho} = 1$. The 2018 distribution is spread toward lower values of ρ; the 2025 distribution piles close to the upper bound $\widehat{\rho} = 1$.

6.4 Pooled estimation of ρ and the network wedge γ [TO BE COMPLETED]

This subsection is work in progress and should be read as outlining the estimation strategy we plan to implement next, rather than as a final set of results. The evidence so far points to a residual wedge, but not necessarily to the same wedge in every part of the price system. For consumer prices, the main discrepancy is that PCE prices appear to have responded, on average, more than the baseline model predicts in 2018 and less than it predicts in 2025. For producer prices, the relevant discrepancy arises inside the domestic production block and is concentrated primarily in the 2018 episode, where prices respond more than the model-implied marginal-cost exposure would predict. We therefore plan to separate the estimation problem into a PCE exercise, focused on a revealed retail-side covariance wedge, and a PPI exercise, where both the oligopolistic channel ρ and the network covariance wedge γ may matter.

This separation also clarifies the identification problem. In the PCE exercise, there is no direct analogue of the industry-level ρ parameter. The estimation is therefore not trying to distinguish between two sectoral objects that may move together. In the PPI exercise, by contrast, both ρ and γ can alter the same model-implied exposure. A lower ρ increases the weight on the foreign competitor price, while a nonzero γ changes the effective imported-input exposure of the firms that matter most for domestic prices. Since both objects may be related to concentration, markups, and import intensity, the PPI estimates must be interpreted with greater caution.

The empirical patterns above so far favor the covariance channel as an interpretation. The revealed-ρ estimates are useful for fitting the 2018 PPI response, where the baseline model undershoots the data. In 2025, however, the same channel has little role to play, because the baseline production-network model already fits producer prices reasonably well. If ρ is interpreted as a stable feature of market structure, this cross-episode instability is difficult to reconcile. The covariance wedge, by contrast, is inherently about the interaction between a particular shock and the firms exposed to it. When tariffs are concentrated on one source country, as in 2018, firms within an industry may differ substantially in exposure depending on where they source their inputs. When tariffs are broad-based across source countries, as in 2025, those within-industry differences are likely to be smaller. The same industry can therefore display a large covariance wedge in one episode and a small wedge in another without requiring a change in its underlying competitive structure.

This interpretation remains provisional. Deviations from the model may also reflect the limits of the first-order approximation, differences in timing, or unobserved shocks that are correlated with tariff exposure in one episode but not the other. The value of the covariance specification is that it gives these deviations an economically interpretable structure: the model need not fail because average import exposure is mismeasured; it may fail because the firms carrying the exposure are not representative of the firms setting the relevant prices.

[TBC]

7 Conclusion

This paper develops a framework for tracing how foreign price shocks pass through to domestic producer and consumer prices. The model nests several channels that are often studied separately: direct import exposure, input-output propagation, pass-through in levels versus logs, oligopolistic competition with foreign suppliers, and heterogeneity in the firms exposed to the shock. Mapping these objects to public US data allows us to construct model-implied tariff exposures for PPI industries and PCE expenditure categories and to compare those exposures with observed price dynamics.

The empirical results suggest that production-network exposure explains an important part of downstream tariff incidence, but also that no single pass-through coefficient summarizes the evidence. PCE responses are identified primarily from direct import exposure, while PPI data provide a cleaner test of the indirect marginal-cost channel. The baseline model fits the 2025 PPI response relatively well but undershoots the 2018 response, leaving a residual wedge. A revealed oligopolistic channel can rationalize part of this wedge, but its lack of stability across episodes points instead toward the importance of within-industry heterogeneity in exposure and pass-through.

References

Amiti, Mary, Chris Flanagan, Sebastian Heise, and David E. Weinstein. 2026. "Who Is Paying for the 2025 U.S. Tariffs?" *Liberty Street Economics*. Federal Reserve Bank of New York Liberty Street Economics Blog.

Amiti, Mary, Oleg Itskhoki, and Jozef Konings. 2019. "International Shocks, Variable Markups, and Domestic Prices." *The Review of Economic Studies*, 86(6): 2356–2402.

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(4): 187–210.

Atkeson, Andrew, and Ariel Burstein. 2008. "Pricing-to-Market, Trade Costs, and International Relative Prices." *American Economic Review*, 98(5): 1998–2031.

Baqaee, David, and Elisa Rubbo. 2023. "Micro Propagation and Macro Aggregation." *Annual Review of Economics*, 15(1): 91–123.

Baqaee, David Rezza, and Emmanuel Farhi. 2019. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica*, 87(4): 1155–1203.

Barbiero, Omar, and Hillary Stein. 2025. "The Impact of Tariffs on Inflation." *Federal Reserve Bank of Boston Current Policy Perspectives 25-2.*

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(1): 19–34.

Cavallo, Alberto, Paola Llamas, and Franco M. Vazquez. 2025. "Tracking the Short-Run Price Impact of U.S. Tariffs." National Bureau of Economic Research Working Paper 34496.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor. 2025. "A Local Projections Approach to Difference-in-Differences." *Journal of Applied Econometrics*, 40(7): 741–758.

Fajgelbaum, Pablo D., and Amit Khandelwal. 2026. "Tariffs in 2025: Short-Run Impacts on the U.S. Economy." National Bureau of Economic Research Working Paper 35064.

Fajgelbaum, Pablo D., Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. 2020. “The Return to Protectionism.” *The Quarterly Journal of Economics*, 135(1): 1–55.

Feenstra, Robert C. 1989. “Symmetric Pass-Through of Tariffs and Exchange Rates under Imperfect Competition: An Empirical Test.” *Journal of International Economics*, 27(1–2): 25–45.

Flaaen, Aaron, and Justin Pierce. 2024. “Disentangling the Effects of the 2018–2019 Tariffs on a Globally Connected U.S. Manufacturing Sector.” *The Review of Economics and Statistics*, 1–45. Advance publication.

Galichon, Alfred. 2026. *Discrete Choice Models: Mathematical Methods, Econometrics, and Data Science.* Princeton University Press.

Goldberg, Pinelopi Koujianou, and Michael M. Knetter. 1997. “Goods Prices and Exchange Rates: What Have We Learned?” *Journal of Economic Literature*, 35(3): 1243–1272.

Gopinath, Gita, and Brent Neiman. 2026. “The Incidence of Tariffs: Rates and Reality.” National Bureau of Economic Research Working Paper 34620.

Hacıoğlu-Hoke, Sinem, Sarojini Malladi, and Leo Feler. 2026. “The Slow Climb: How Tariffs Gradually Raised Retail Prices in 2025.” *FEDS Notes.*

Hale, Galina, Bart Hobijn, Fernanda Nechio, and Doris Wilson. 2019. “How Much Do We Spend on Imports?” *FRBSF Economic Letter*, 2019-10.

La’O, Jennifer, and Alireza Tahbaz-Salehi. 2022. “Optimal Monetary Policy in Production Networks.” *Econometrica*, 90(3): 1295–1336.

Minton, Robbie, and Mariano Somale. 2025. “Detecting Tariff Effects on Consumer Prices in Real Time.” *FEDS Notes.*

Minton, Robert, and Brian Wheaton. 2023. “Delayed Inflation in Supply Chains: Theory and Evidence.” *Available at SSRN.*

Minton, Robert, Madeleine Ray, and Mariano Somale. 2026. “Detecting Tariff Effects on Consumer Prices in Real Time – Part II.” *FEDS Notes.*

Sangani, Kunal. 2026. “Complete Pass-Through in Levels.” *The Quarterly Journal of Economics*, 141(2): 1077–1135.

A Model Details

A.1 Price indices

A.1.1 Intermediate input price indices

$$p^x_{fi}(\{p^x_{fi,j}\}_{j\in\mathcal{I}}) = \min_{\{x_{fi,j}\}_{j\in\mathcal{I}}} \left\{ \sum_{j\in\mathcal{I}} p^x_{fi,j} x_{fi,j} \quad \text{s.t.} \quad x_{fi}(\{x_{fi,j}\}_{j\in\mathcal{I}}) \geq 1 \right\}, \tag{A.1}$$

$$p^x_{fi,j}(\{p_{gj}\}_{g\in\mathcal{D}_j}) = \min_{\{x_{fi,gj}\}_{g\in\mathcal{D}_j}} \left\{ \sum_{g\in\mathcal{D}_j} p_{gj} x_{fi,gj} \quad \text{s.t.} \quad x_{fi,j}(\{x_{fi,gj}\}_{g\in\mathcal{D}_j}) \geq 1 \right\}, \tag{A.2}$$

Both aggregators define price indices as a function of firms' prices:

$$p^x_{fi,j} = p^x_{fi,j}(\{p_{gj}\}_{g\in\mathcal{D}_j}), \tag{A.3}$$

which is the price paid by firm f in industry i for inputs from j and

$$p^x_{fi} = p^x_{fi}(\{p^x_{fi,j}\}_{j\in\mathcal{I}}), \tag{A.4}$$

is the *domestic* intermediate input bundle paid by firm f in industry i.

Up to a first-order approximation both price indices [equations A.3 and A.4] satisfies

$$\mathrm{d}\log p^x_{fi,j} = \sum_{g\in\mathcal{D}_j} \frac{p_{gj} x_{fi,gj}}{p^x_{fi,j} x_{fi,j}} \mathrm{d}\log p_{gj} = \sum_{g\in\mathcal{D}_j} \Omega_{fi,gj} \mathrm{d}\log p_{gj} \tag{A.5}$$

$$\mathrm{d}\log p^x_{fi} = \sum_{j\in\mathcal{I}} \frac{p^x_{fi,j} x_{fi,j}}{p^x_{fi} x_{fi}} \mathrm{d}\log p^x_{fi,j} = \sum_{j\in\mathcal{I}} \Omega_{fi,j} \mathrm{d}\log p^x_{fi,j} \tag{A.6}$$

In the main text, we listed the assumptions we needed to empirically implement our model with the available data. We repeat the assumptions here so that this appendix is self-contained.

Assumption A.1 (Expenditure shares across varieties)**.** *Expenditure shares across varieties of a given industry are the same regardless of the buyer origin* $\Omega_{fi,gj} = \Omega_{\cdot,gj}$ *for all* f, i *pairs.*

Assumption A.1 ensures that all firms pay the same price when sourcing varieties from another industry because

$$\mathrm{d}\log p^x_{fi,j} = \sum_{g\in j} \Omega_{fi,gj} \mathrm{d}\log p_{gj} = \sum_{g\in j} \Omega_{\cdot,gj} \mathrm{d}\log p_{gj} = \mathrm{d}\log p^x_j. \tag{A.7}$$

Hence, Assumption A.1 allows us to use a single intermediate-input industry price index change that is common across buyers.

Assumption A.2 (Expenditure shares across industries)**.** *There is no heterogeneity in expenditure shares across firms within an industry:* $\Omega_{fi,j} = \Omega_{i,j}$ *for all* $f \in i$.

Assumption A.2 implies we only need information at the industry level to construct the industry-level domestic intermediate input price index changes. Formally,

$$\mathrm{d}\log p^x_{fi} = \sum_{j\in\mathcal{I}} \Omega_{fi,j} \mathrm{d}\log p^x_{fi,j} = \sum_{j\in\mathcal{I}} \Omega_{i,j} \mathrm{d}\log p^x_j, \tag{A.8}$$

where we also use the result implied by Assumption A.1.

The final result we need is that $\mathrm{d}\log p^x_j = \mathrm{d}\log p^d_j$ for all $j \in \mathcal{I}$. That is, the intermediate input price index paid by each industry precisely coincides with the producer price index. The next assumption makes this holds within our model:

Assumption A.3. *For each firm-industry pair,* $\Omega_{.,fi} = s^d_{fi}$.

The producer price index change $\mathrm{d}\log p^d_i$ coincides with that of the intermediate input price index paid by industries $\mathrm{d}\log p^x_i$ provided that Assumption A.3.

We highlight that Assumptions A.1 to A.3 are not innocuous: any degree of heterogeneity in these shares would imply modifications to our equations and affect measurement. Unfortunately, eliminating these assumptions requires information on bilateral expenditure shares at the firm level.

B Proofs

B.1 Proof of Proposition 1

Start from the pricing condition

$$p_{fi} = mc_{fi} + \frac{\sigma_i}{1-\pi_{fi}}. \tag{B.1}$$

Differentiating this expression

$$\mathrm{d}p_{fi} = \mathrm{d}mc_{fi} + \frac{\mu_{fi}}{1-\pi_{fi}} \mathrm{d}\pi_{fi}.$$

Differentiating the share, we get

$$\mathrm{d}\pi_{fi} = -\frac{\pi_{fi}}{\sigma_i}\left(\mathrm{d}p_{fi} - \sum_{g\in i}\pi_{gi}\mathrm{d}p_{gi}\right)$$
$$\mathrm{d}\pi_{fi} = -\frac{\pi_{fi}}{\sigma_i}(1-\pi_{fi})\left(\mathrm{d}p_{fi} - \sum_{g\in i\setminus f}\tilde{\pi}_{gi}\mathrm{d}p_{gi}\right)$$
$$\mathrm{d}\pi_{fi} = -\frac{\pi_{fi}}{\sigma_i}(1-\pi_{fi})\left(\mathrm{d}p_{fi} - \mathrm{d}p_{-f,i}\right),$$

where we defined the competitors' price index of firm f as

$$\mathrm{d}p_{-f,i} = \sum_{g\in i\setminus f}\tilde{\pi}_{gi}\mathrm{d}p_{gi}, \tag{B.2}$$

with $\tilde{\pi}_{gi} = \frac{\pi_{gi}}{\sum\limits_{h\in i\setminus f}\pi_{hi}}$.

Replacing into the differentiated pricing condition

$$\mathrm{d}p_{fi} = \mathrm{d}mc_{fi} - \frac{\mu_{fi}}{1-\pi_{fi}}\frac{\pi_{fi}}{\sigma_i}(1-\pi_{fi})\left(\mathrm{d}p_{fi} - \mathrm{d}p_{-f,i}\right)$$
$$\mathrm{d}p_{fi} = \mathrm{d}mc_{fi} - \frac{\pi_{fi}}{1-\pi_{fi}}\left(\mathrm{d}p_{fi} - \mathrm{d}p_{-f,i}\right)$$

Rearranging, we get

$$\mathrm{d}p_{fi} = (1-\pi_{fi})\mathrm{d}mc_{fi} + \pi_{fi}\mathrm{d}p_{-f,i}, \tag{B.3}$$

which is the expression in the main text.

B.2 Proof of Proposition 2

Divide the final expression of Proposition 1 by p_{fi} and rearrange to get

$$\mathrm{d}\log p_{fi} = (1-\pi_{fi})\frac{mc_{fi}}{p_{fi}}\mathrm{d}\log mc_{fi} + \pi_{fi}\frac{p_{-f,i}}{p_{fi}}\mathrm{d}\log p_{-f,i}. \tag{B.4}$$

B.3 Proof of Proposition 3

Start from the definition of the producer price index

$$
\begin{aligned}
\mathrm{d}\log p_i^d &= \sum_{f\in\mathcal{D}_i} s_{fi}^d \mathrm{d}\log p_{fi} \\
\mathrm{d}\log p_i^d &= \sum_{f\in\mathcal{D}_i} s_{fi}^d (1-\alpha_{fi}) \left[\rho_{fi}\mathrm{d}\log mc_{fi} + (1-\rho_{fi})\frac{R_i}{mc_{fi}q_i}\mathrm{d}\log p_i\right] \\
\mathrm{d}\log p_i^d &= \left(\sum_{g\in\mathcal{D}_i} s_{gi}^d (1-\alpha_{gi})\right)\left(\sum_{g\in\mathcal{D}_i}\frac{s_{gi}^d(1-\alpha_{gi})}{\left(\sum_{g\in\mathcal{D}_i} s_{gi}^d(1-\alpha_{gi})\right)}\rho_{gi}\right)\sum_{f\in\mathcal{D}_i}\frac{s_{fi}^d(1-\alpha_{fi})\rho_{fi}}{\sum_{g\in\mathcal{D}_i} s_{gi}^d(1-\alpha_{gi})\rho_{gi}}\mathrm{d}\log mc_{fi} \\
&\quad + \left[\sum_{f\in\mathcal{D}_i}\frac{s_{fi}^d}{mc_{fi}}(1-\alpha_{fi})(1-\rho_{fi})\right]\frac{R_i}{q_i}\mathrm{d}\log p_i \\
\mathrm{d}\log p_i^d &= \frac{VC_i^d}{R_i^d}\rho_i^d \sum_{f\in\mathcal{D}_i}\gamma_{fi}\mathrm{d}\log mc_{fi} + \left[\sum_{f\in\mathcal{D}_i}\frac{s_{fi}^d}{mc_{fi}}(1-\alpha_{fi})(1-\rho_{fi})\right]\frac{R_i}{q_i}\mathrm{d}\log p_i \\
\mathrm{d}\log p_i^d &= \frac{VC_i^d}{R_i^d}\rho_i^d \sum_{f\in\mathcal{D}_i}\gamma_{fi}\mathrm{d}\log mc_{fi} + \frac{R_i}{R_i^d}\left(\sum_{f\in\mathcal{D}_i}\pi_{fi}(1-\rho_{fi})\right)\mathrm{d}\log p_i \\
\mathrm{d}\log p_i^d &= \frac{VC_i^d}{R_i^d}\rho_i^d \sum_{f\in\mathcal{D}_i}\gamma_{fi}\mathrm{d}\log mc_{fi} + \frac{R_i}{R_i^d}(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)\mathrm{d}\log p_i
\end{aligned}
$$

We first solve for $\mathrm{d}\log p_{fi}$ as a function of marginal costs and the industry price index. To do so, recall that the level change in the markup is

$$
\mathrm{d}\mu_{fi} = -\mu_{fi}\frac{\pi_{fi}}{\sigma_i(1-\pi_{fi})}\left(\mathrm{d}p_{fi} - \sum_{g\in i}\pi_{gi}\mathrm{d}p_{gi}\right) = -\frac{\pi_{fi}}{(1-\pi_{fi})^2}\left(\mathrm{d}p_{fi} - \mathrm{d}p_i^{eff}\right),
$$

where we use the definition of the change in the effective price index

$$
\mathrm{d}p_i^{eff} = \sum_{g\in i}\pi_{gi}\mathrm{d}p_{gi}. \tag{B.5}
$$

Replacing the change in the markup into the pricing condition

$$\mathrm{d}p_{fi} = \mathrm{d}mc_{fi} + \mathrm{d}\mu_{fi}$$
$$\mathrm{d}p_{fi} = \mathrm{d}mc_{fi} - \frac{\pi_{fi}}{(1-\pi_{fi})^2}\left(\mathrm{d}p_{fi} - \mathrm{d}p_i^{eff}\right)$$
$$\mathrm{d}p_{fi} = \frac{1}{1+\frac{\pi_{fi}}{(1-\pi_{fi})^2}}\mathrm{d}mc_{fi} + \frac{\frac{\pi_{fi}}{(1-\pi_{fi})^2}}{1+\frac{\pi_{fi}}{(1-\pi_{fi})^2}}\mathrm{d}p_i^{eff}.$$

This implies

$$\mathrm{d}p_{fi} = \rho_{fi}\mathrm{d}mc_{fi} + (1-\rho_{fi})\mathrm{d}p_i^{eff}, \tag{B.6}$$

where

$$\rho_{fi} = \frac{1}{1+\frac{\pi_{fi}}{(1-\pi_{fi})^2}} = \frac{(1-\pi_{fi})^2}{(1-\pi_{fi})^2+\pi_{fi}} = \frac{(1-\pi_{fi})^2}{1-\pi_{fi}+\pi_{fi}^2}$$

We can link the effective price index to the industry price index change $\mathrm{d}\log p_i$:

$$\mathrm{d}p_i^{eff} = \sum_{f\in i} p_{fi}\pi_{fi}\mathrm{d}\log p_{fi},$$
$$= \sum_{f\in i}\frac{p_{fi}q_{fi}}{q_i}\mathrm{d}\log p_{fi},$$
$$= \frac{R_i}{q_i}\sum_{f\in i}\frac{p_{fi}q_{fi}}{R_i}\mathrm{d}\log p_{fi},$$
$$\mathrm{d}p_i^{eff} = \frac{R_i}{q_i}\mathrm{d}\log p_i.$$

Using this equality and expressing the pricing condition in percentage terms

$$\mathrm{d}p_{fi} = \rho_{fi}\mathrm{d}mc_{fi} + (1-\rho_{fi})\frac{R_i}{q_i}\mathrm{d}\log p_i,$$
$$\mathrm{d}\log p_{fi} = \rho_{fi}(1-\alpha_{fi})\mathrm{d}\log mc_{fi} + (1-\rho_{fi})\frac{R_i}{p_{fi}q_i}\mathrm{d}\log p_i.$$

Multiplying by s^d_{fi} and adding across domestic firms we get

$$\mathrm{d}\log p^d_i = \sum_{f\in\mathcal{D}_i} s^d_{fi}\left(\rho_{fi}(1-\alpha_{fi})\mathrm{d}\log mc_{fi} + (1-\rho_{fi})\frac{R_i}{p_{fi}q_i}\mathrm{d}\log p_i\right)$$
$$\mathrm{d}\log p^d_i = \sum_{f\in\mathcal{D}_i} s^d_{fi}\rho_{fi}(1-\alpha_{fi})\mathrm{d}\log mc_{fi} + \sum_{f\in\mathcal{D}_i} s^d_{fi}(1-\rho_{fi})\frac{R_i}{p_{fi}q_i}\mathrm{d}\log p_i$$
$$\mathrm{d}\log p^d_i = \sum_{f\in\mathcal{D}_i} \frac{p_{fi}q_{fi}}{R^d_i}\rho_{fi}(1-\alpha_{fi})\mathrm{d}\log mc_{fi} + \sum_{f\in\mathcal{D}_i} \frac{p_{fi}q_{fi}}{R^d_i}(1-\rho_{fi})\frac{R_i}{p_{fi}q_i}\mathrm{d}\log p_i$$
$$\mathrm{d}\log p^d_i = \sum_{f\in\mathcal{D}_i} \frac{mc_{fi}q_{fi}}{R^d_i}\rho_{fi}\mathrm{d}\log mc_{fi} + \sum_{f\in\mathcal{D}_i} \pi_{fi}(1-\rho_{fi})\frac{R_i}{R^d_i}\mathrm{d}\log p_i$$
$$\mathrm{d}\log p^d_i = \frac{VC^d_i}{R^d_i}\sum_{f\in\mathcal{D}_i} \frac{mc_{fi}q_{fi}}{VC^d_i}\rho_{fi}\mathrm{d}\log mc_{fi} + \sum_{f\in\mathcal{D}_i} \pi_{fi}(1-\rho_{fi})\frac{R_i}{R^d_i}\mathrm{d}\log p_i$$
$$\mathrm{d}\log p^d_i = \upsilon^d_i\sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi} + (1-\pi_{Fi})\sum_{f\in\mathcal{D}_i} \tilde{\pi}_{fi}(1-\rho_{fi})\frac{1}{s^d_i}\mathrm{d}\log p_i$$

Now, we use the definition of $\mathrm{d}\log p_i = s^d_i\mathrm{d}\log p^d_i + s_{Fi}\mathrm{d}\log p_{Fi}$, to get

$$\mathrm{d}\log p^d_i = \upsilon^d_i\sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi} + (1-\pi_{Fi})\sum_{f\in\mathcal{D}_i} \tilde{\pi}_{fi}(1-\rho_{fi})\frac{1}{s^d_i}\left(s^d_i\mathrm{d}\log p^d_i + s_{Fi}\mathrm{d}\log p_{Fi}\right).$$

Rearranging, we get

$$\mathrm{d}\log p^d_i = \frac{1}{\left(1-(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)\right)}\left(\upsilon^d_i\left(\sum_{f\in\mathcal{D}_i}\tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi}\right)\right), \tag{B.7}$$
$$+\frac{1}{\left(1-(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)\right)}\left(\frac{(1-\pi_{Fi})\left(\sum_{f\in\mathcal{D}_i}\tilde{\pi}_{fi}(1-\rho_{fi})\right)}{(1-s_{Fi})}s_{Fi}\mathrm{d}\log p_{Fi}\right),$$

which is the expression in the main text.

B.4 Proof of Proposition 4

From the earlier section, we have an expression for the PPI change of an industry i:

$$\mathrm{d}\log p^d_i = \upsilon^d_i\sum_{f\in\mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi}\mathrm{d}\log mc_{fi} + (1-\pi_{Fi})\sum_{f\in\mathcal{D}_i} \tilde{\pi}_{fi}(1-\rho_{fi})\frac{1}{s^d_i}\left(s^d_i\mathrm{d}\log p^d_i + s_{Fi}\mathrm{d}\log p_{Fi}\right) \tag{B.8}$$

Given our assumptions, we have $\mathrm{d}\log mc_{fi} = \mathrm{d}\log mc_i$ for all $f \in \mathcal{D}_i$. Moreover, we defined:

$$\rho_i^d = \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}(\boldsymbol{\rho}_i) = \sum_{f \in \mathcal{D}_i} \tilde{\delta}_{fi}\rho_{fi},$$

as the industry level pass-through of marginal cost to industry prices in levels, weighted according to domestic cost-shares $\tilde{\boldsymbol{\delta}}_i = \{\delta_{fi}\}_{f \in \mathcal{D}_i}$.

Using these results, we can write equation (B.8) as

$$\mathrm{d}\log p_i^d = \upsilon_i^d \rho_i^d \mathrm{d}\log mc_i + (1-\pi_{Fi})(1-\mathbb{E}_{\tilde{\boldsymbol{\pi}}_i}(\boldsymbol{\rho}_i))(\mathrm{d}\log p_i^d + (1-s_{Fi})^{-1}s_{Fi}\mathrm{d}\log p_{Fi}), \quad \text{(B.9)}$$

which follows since $\sum_{f \in \mathcal{D}_i} \tilde{\pi}_{fi} = 1$ and $\mathbb{E}_{\tilde{\boldsymbol{\pi}}_i}(\boldsymbol{\rho}_i) = \sum_{f \in \mathcal{D}_i} \tilde{\pi}_{fi}\rho_{fi}$ denotes the expected marginal cost pass-through in levels when weighted according to domestic task shares $\tilde{\boldsymbol{\pi}}_i = \{\tilde{\pi}_{fi}\}_{f \in \mathcal{D}_i}$.

To arrive at the proposition in the main text, we assume that $\mathbb{E}_{\tilde{\boldsymbol{\pi}}_i}(\boldsymbol{\rho}_i) \approx \mathbb{E}_{\tilde{\boldsymbol{\delta}}_i}(\boldsymbol{\rho}_i) = \rho_i^d$. Hence, industry level marginal cost pass-through in levels when weighted according to either domestic task shares or domestic cost shares are not too dissimilar. This assumption allows us to write:

$$\mathrm{d}\log p_i^d = \upsilon_i^d \rho_i^d \mathrm{d}\log mc_i + (1-\pi_{Fi})(1-\rho_i^d)(\mathrm{d}\log p_i^d + (1-s_{Fi})^{-1}s_{Fi}\mathrm{d}\log p_{Fi}). \quad \text{(B.10)}$$

Equation (B.10) in vector form is

$$\mathrm{d}\log \boldsymbol{p}^d = \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \mathrm{d}\log \boldsymbol{mc} + (\boldsymbol{I}-\boldsymbol{\Pi}_F)(\boldsymbol{I}-\boldsymbol{\rho}^d)(\mathrm{d}\log \boldsymbol{p}^d + (\boldsymbol{I}-\boldsymbol{S}_F)^{-1}\boldsymbol{S}_F\mathrm{d}\log \boldsymbol{p}_F), \quad \text{(B.11)}$$

where

$$\boldsymbol{\Upsilon}^d = diag(\boldsymbol{\upsilon}^d); \quad \boldsymbol{\rho}^d = diag(\{\rho_i^d\}_{i \in \mathcal{I}}); \quad \boldsymbol{S}_F = diag(\boldsymbol{s}_F); \quad \boldsymbol{\Pi}_F = diag(\boldsymbol{\pi}_F).$$

Up to a first-order approximation, marginal cost changes are only a function of input price changes:

$$\mathrm{d}\log mc_{fi} = \sum_j \Omega_{fi,j}\mathrm{d}\log p_j^d + \sum_j \Omega_{fi,j}^*\mathrm{d}\log p_j^*. \quad \text{(B.12)}$$

Under Assumptions A.1 and A.2, this can be written omitting the firm subscript

$$\mathrm{d}\log mc_i = \sum_j \Omega_{i,j}\mathrm{d}\log p_j^d + \sum_j \Omega_{i,j}^*\mathrm{d}\log p_j^*, \quad \text{(B.13)}$$

or in vector form

$$\mathrm{d}\log \boldsymbol{mc} = \boldsymbol{\Omega}^d \mathrm{d}\log \boldsymbol{p}^d + \boldsymbol{\Omega}^* \mathrm{d}\log \boldsymbol{p}^*, \tag{B.14}$$

with $\boldsymbol{\Omega}^d = \{\Omega_{i,j}\}_{i,j\in\mathcal{I}}$ and $\boldsymbol{\Omega}^* = \{\Omega^*_{i,j}\}_{i,j\in\mathcal{I}}$.

Replacing (B.14) into (B.11), we get

$$\mathrm{d}\log \boldsymbol{p}^d = \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \left(\boldsymbol{\Omega}^d \mathrm{d}\log \boldsymbol{p}^d + \boldsymbol{\Omega}^* \mathrm{d}\log \boldsymbol{p}^*\right) + (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d)(\mathrm{d}\log \boldsymbol{p}^d + (\boldsymbol{I} - \boldsymbol{S}_F)^{-1} \boldsymbol{S}_F \mathrm{d}\log \boldsymbol{p}_F).$$

Solving for $\mathrm{d}\log \boldsymbol{p}^d$ we get

$$\begin{aligned}\mathrm{d}\log \boldsymbol{p}^d = {} & (\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^d - (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d))^{-1} \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^* \mathrm{d}\log \boldsymbol{p}^*, \\ & + (\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^d - (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d))^{-1} (\boldsymbol{I} - \boldsymbol{S}_F)^{-1} (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d) \boldsymbol{S}_F \mathrm{d}\log \boldsymbol{p}_F.\end{aligned}$$

Letting

$$\boldsymbol{\Phi}^* = \boldsymbol{G} \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^*, \tag{B.15}$$

$$\boldsymbol{\Phi}^F = \boldsymbol{G}(\boldsymbol{I} - \boldsymbol{S}_F)^{-1} (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d) \boldsymbol{S}_F, \tag{B.16}$$

$$\boldsymbol{G} = (\boldsymbol{I} - \boldsymbol{\Upsilon}^d \boldsymbol{\rho}^d \boldsymbol{\Omega}^d - (\boldsymbol{I} - \boldsymbol{\Pi}_F)(\boldsymbol{I} - \boldsymbol{\rho}^d))^{-1}, \tag{B.17}$$

we arrive at

$$\mathrm{d}\log \boldsymbol{p}^d = \boldsymbol{\Phi}^* \mathrm{d}\log \boldsymbol{p}^* + \boldsymbol{\Phi}^F \mathrm{d}\log \boldsymbol{p}_F, \tag{B.18}$$

which is the expression in the main text.

C Additional Figures

C.1 Tariff Impact under the Statutory-Tariff Shock

The main-text tariff-impact figures use the effective tariff change, collected duties over value). Here we reproduce the same decompositions under the statutory tariff change (duties over dutiable value).

Figure C.1: NIPA Decomposition of Headline PCE Impact—2018 Episode (2018–19 tariffs, statutory shock.)

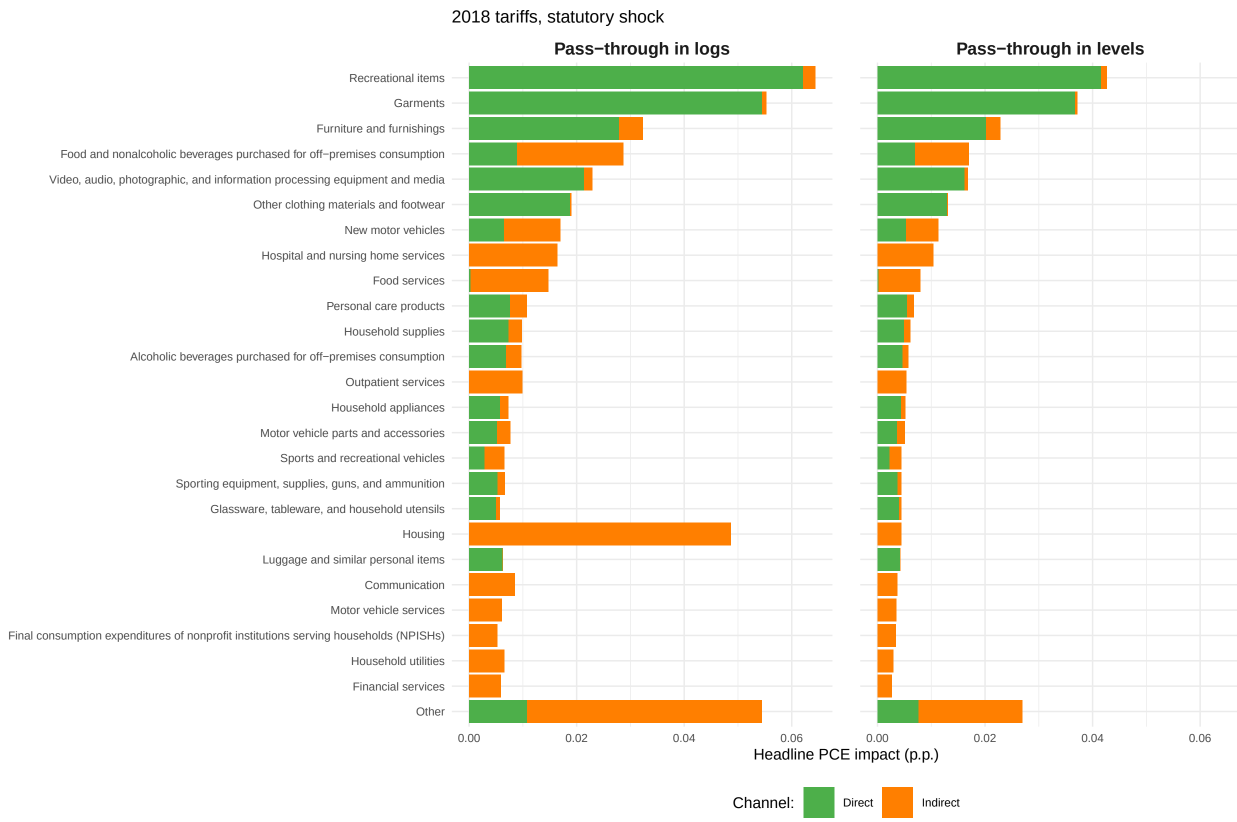

Notes: NIPA decomposition of the headline PCE impact for the 2018 tariff episode, 2018–19 tariffs, statutory shock. Bars are stacked into direct and indirect channels; left panel uses pass-through in logs, right panel uses pass-through in levels.

Figure C.2: NIPA Decomposition of Headline PCE Impact—2025 Episode (2025 tariffs, statutory shock.)

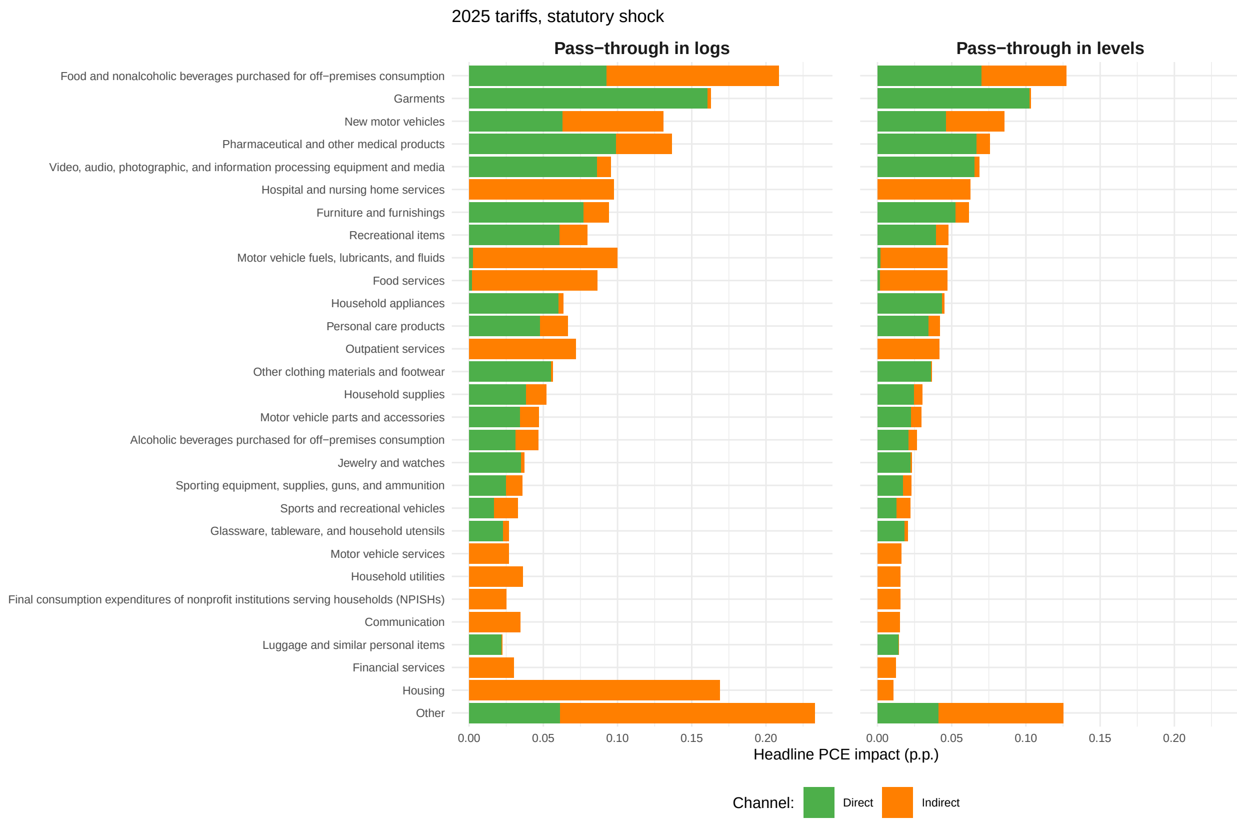

Notes: NIPA decomposition of the headline PCE impact for the 2025 tariff episode, 2025 tariffs, statutory shock. Bars are stacked into direct and indirect channels; left panel uses pass-through in logs, right panel uses pass-through in levels.

Figure C.3: Source-Country Decomposition of Headline PCE Impact (statutory, homog)

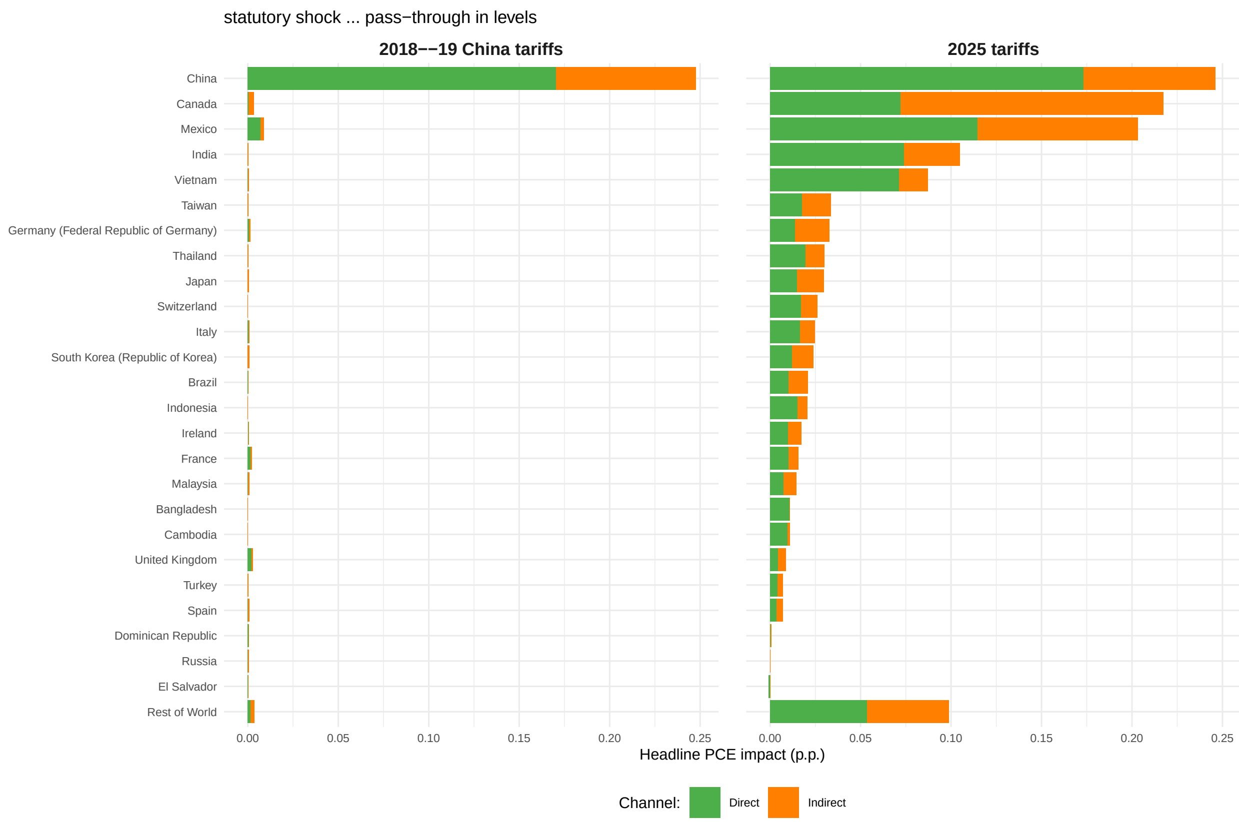

Notes: Source-country decomposition of the headline PCE impact. Statutory-tariff shock. Bars are stacked into direct and indirect channels.

C.2 Tariff Impact under the By-region Estimation

The main-text and preceding appendix figures use the homogeneous estimation, in which ρ is common to all commodities and import shares are pooled across source regions. Here we reproduce the same decompositions under the *byregion* estimation, which splits imports into seven BEA source regions and computes the network propagation separately for each commodity-region cell. Byregion outputs are at NIPA level G (headline, ≈ 76 categories).

Figure C.4: NIPA Decomposition of Headline PCE Impact—2018 Episode (2018–19 tariffs, effective shock, byregion.)

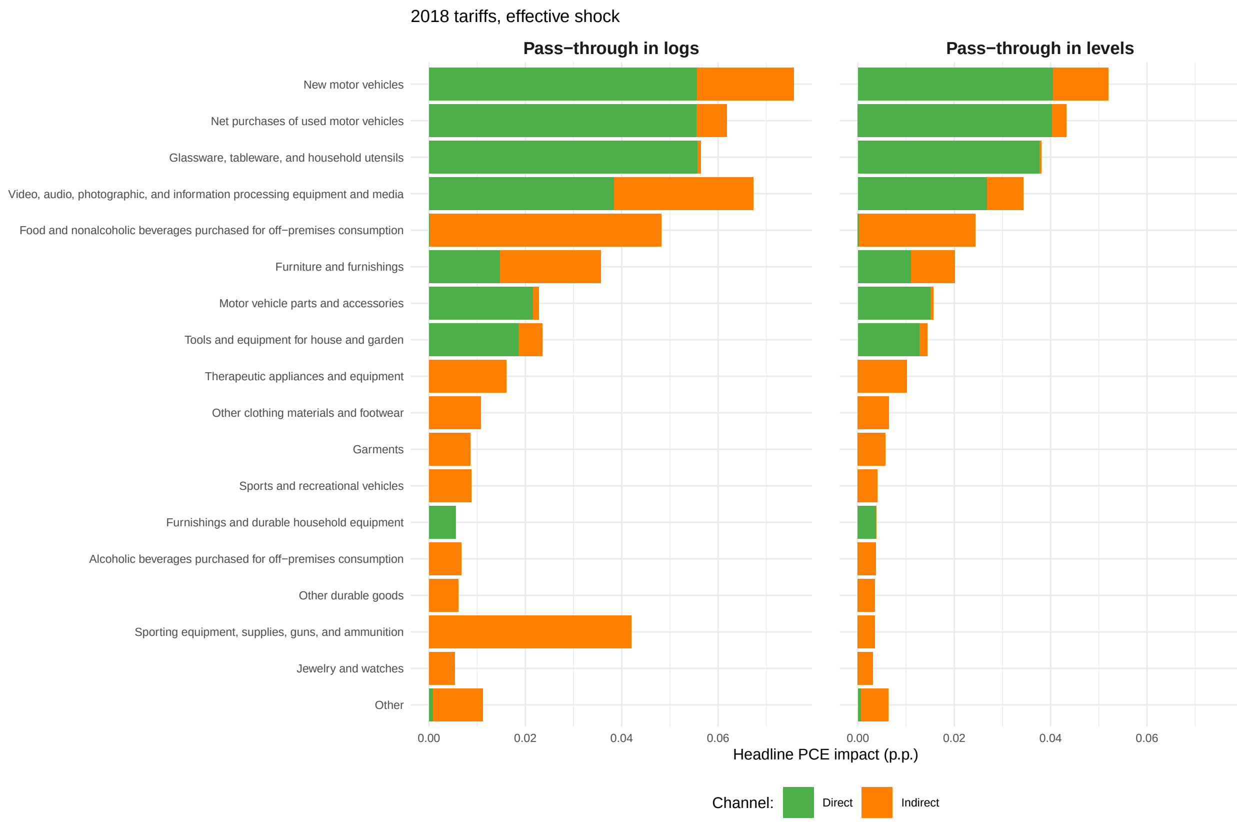

Notes: NIPA decomposition of the headline PCE impact for the 2018 tariff episode, 2018–19 tariffs, effective shock, byregion. Bars are stacked into direct and indirect channels; left panel uses pass-through in logs, right panel uses pass-through in levels.

Figure C.5: NIPA Decomposition of Headline PCE Impact—2025 Episode (2025 tariffs, effective shock, byregion.)

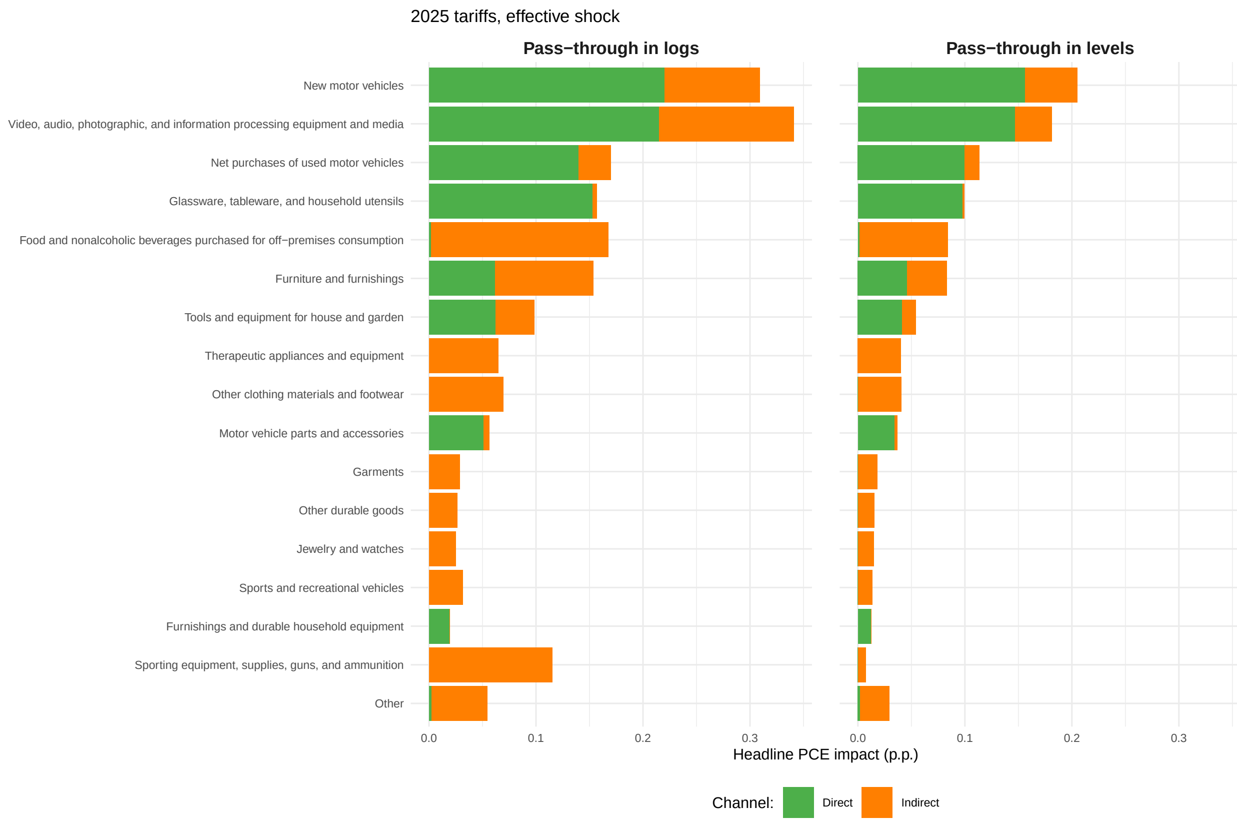

Notes: NIPA decomposition of the headline PCE impact for the 2025 tariff episode, 2025 tariffs, effective shock, byregion. Bars are stacked into direct and indirect channels; left panel uses pass-through in logs, right panel uses pass-through in levels.

Figure C.6: Source-Country Decomposition of Headline PCE Impact (effective, byregion)

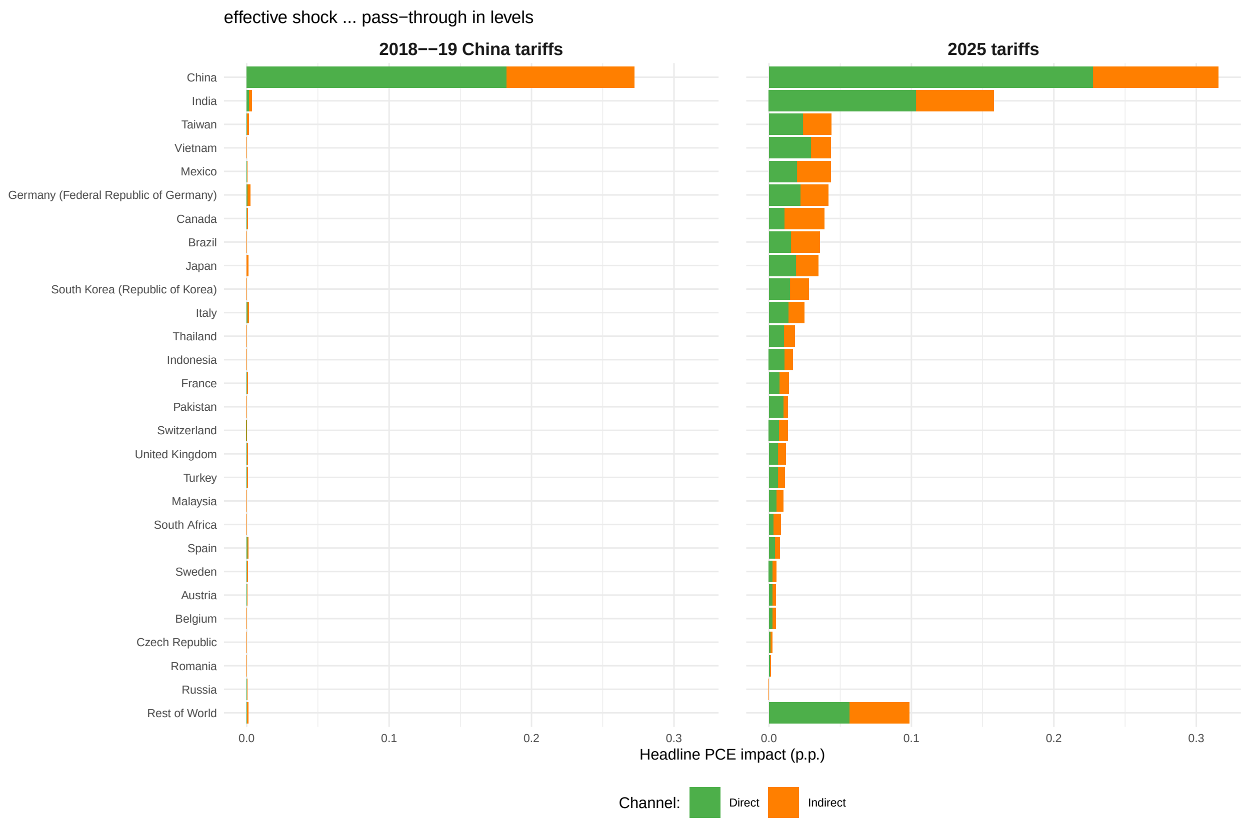

Notes: Source-country decomposition of the headline PCE impact. Effective-tariff shock, byregion. Bars are stacked into direct and indirect channels.

Figure C.7: NIPA Decomposition of Headline PCE Impact—2018 Episode (2018–19 tariffs, statutory shock, byregion.)

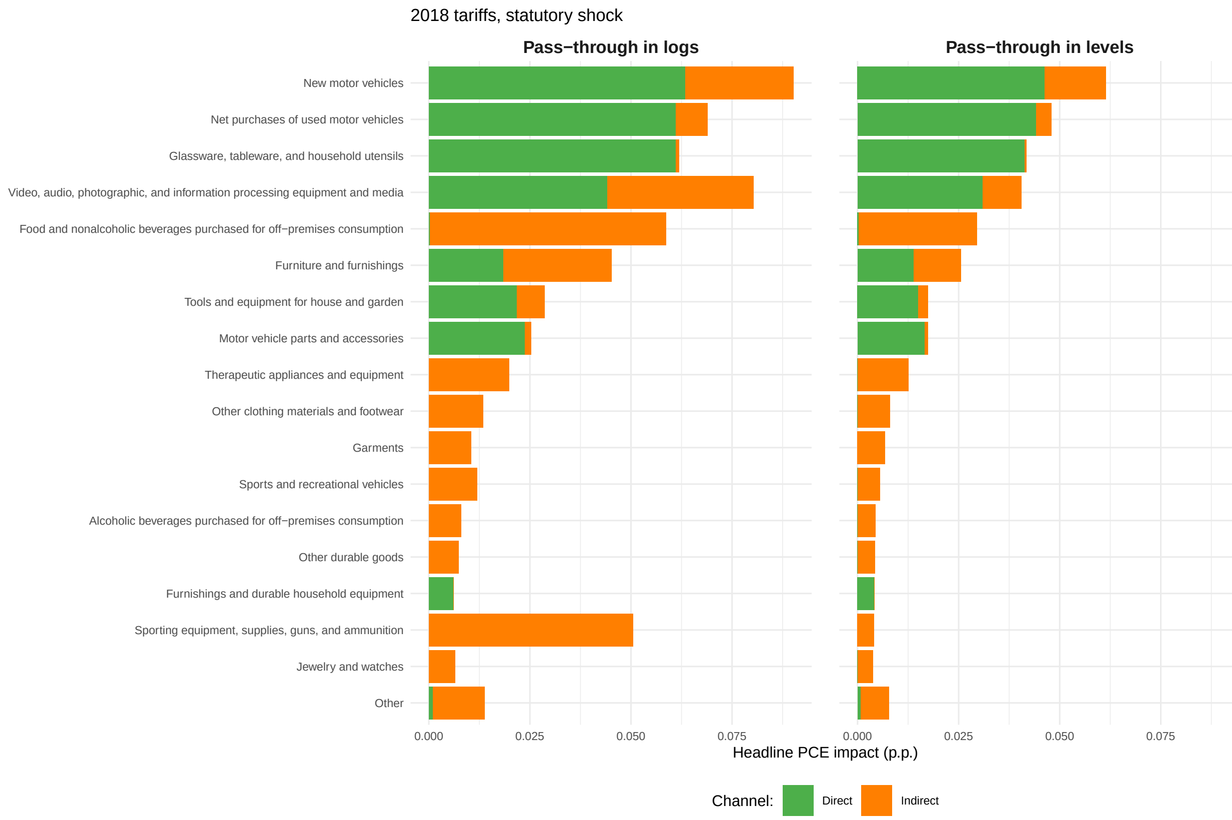

Notes: NIPA decomposition of the headline PCE impact for the 2018 tariff episode, 2018–19 tariffs, statutory shock, byregion. Bars are stacked into direct and indirect channels; left panel uses pass-through in logs, right panel uses pass-through in levels.

Figure C.8: NIPA Decomposition of Headline PCE Impact—2025 Episode (2025 tariffs, statutory shock, byregion.)

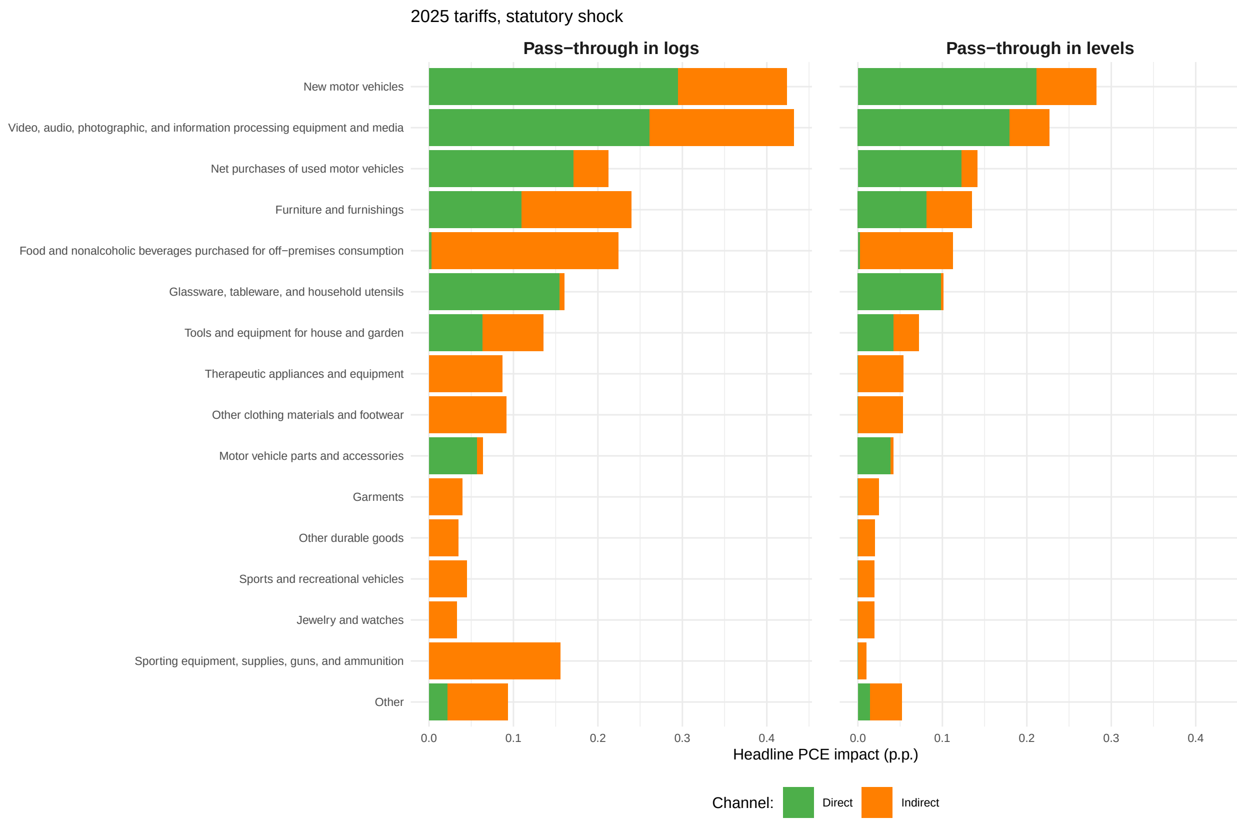

Notes: NIPA decomposition of the headline PCE impact for the 2025 tariff episode, 2025 tariffs, statutory shock, byregion. Bars are stacked into direct and indirect channels; left panel uses pass-through in logs, right panel uses pass-through in levels.

Figure C.9: Source-Country Decomposition of Headline PCE Impact (statutory, byregion)

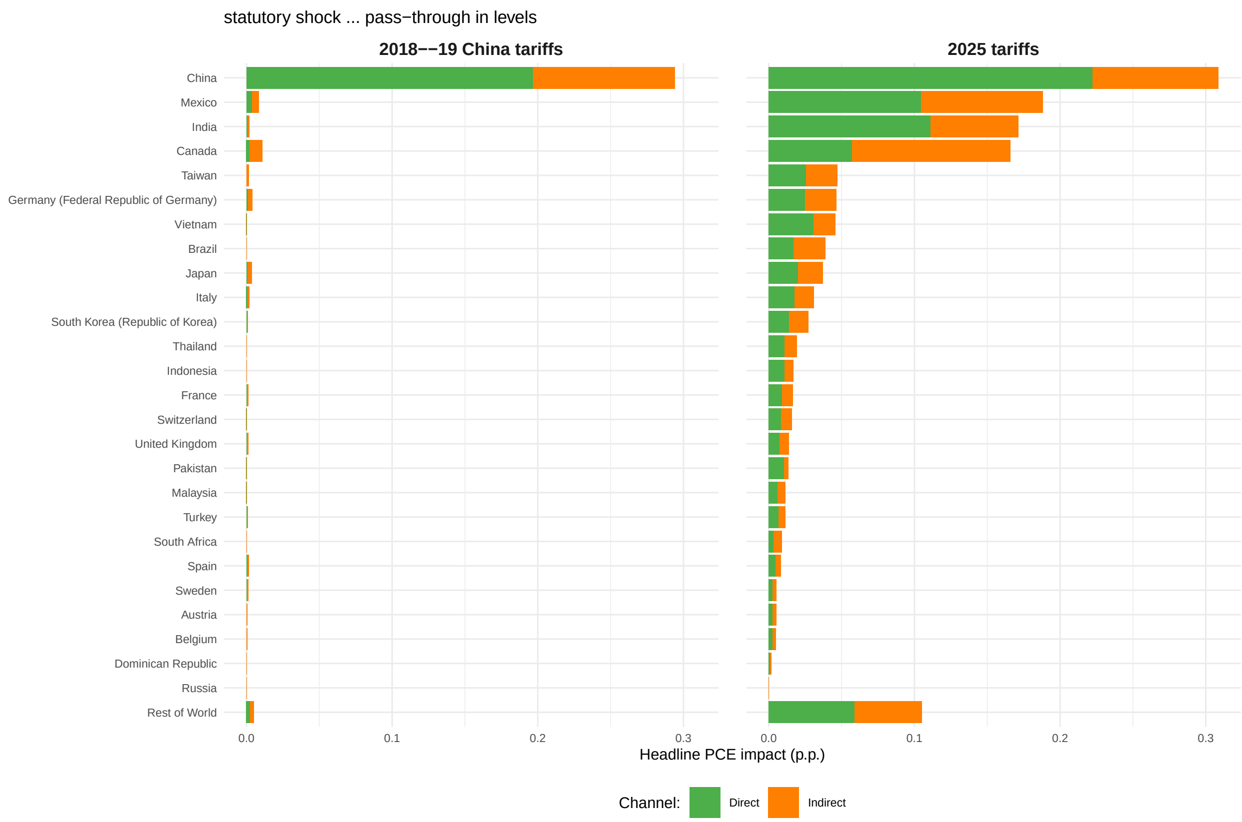

Notes: Source-country decomposition of the headline PCE impact. Statutory-tariff shock, byregion. Bars are stacked into direct and indirect channels.

D Decomposition of the ρ-Sensitivity Commodities

Table D.1 reports the four ingredients of equation (57) for each of the four commodities in Figure 8, for each episode. $d\log p_i^*$ is the cumulated own effective tariff change, built by share-weighting the monthly HS-country tariff changes through the Census HS×country composition matrix into the BEA commodity. s_i^* is the direct import share from the IO matrix (`imports_shares_fromdomproduction_disagg.csv`). The product $s_i^*\, d\log p_i^*$ is the contribution of the direct-foreign channel to the commodity's own marginal cost. $d\log \bar{\text{mc}}_i^d$ is the cumulated upstream marginal-cost increase obtained by evaluating the full network in equation (52) at $\rho = 1$; we report it under both the levels (constant-dollar markup) and the logs (constant-percentage markup) assumptions.

Table D.1: ρ-sensitivity: decomposition of the four highlighted commodities by episode.

Commodity	2018 episode				2025 episode			
	$d\log p_{Fi}$	s_{Fi}	$s_{Fi} d\log p_{Fi}$	$d\log \text{mc}_i^d$ (lvl / log)	$d\log p_{Fi}$	s_{Fi}	$s_{Fi} d\log p_{Fi}$	$d\log \text{mc}_i^d$ (lvl / log)
Pharmaceutical preparations (325412)	0.07	0.03	0.00	8.1 / 23.7	0.77	0.03	0.02	21.9 / 71.0
Biological products (325414)	0.004	49.57	0.002	7.8 / 24.3	0.12	49.57	0.06	18.0 / 69.3
Semiconductors & related (334413)	3.56	16.06	0.57	29.3 / 52.7	5.09	16.06	0.82	50.5 / 93.5
Heavy duty truck mfg. (336120)	0.05	0.86	0.00	75.9 / 93.9	15.40	0.86	0.13	352.7 / 413.2

Notes: All values are in percent and cumulated across the months of each episode (2018–2019 and 2025 through December). $d\log p_{Fi}$ is the industry's own effective tariff change (`diff_eff`), share-weighted through the Census HS×country import composition to the BEA commodity level. s_{Fi} is the direct import share of commodity i (diagonal of the domestic-imports matrix at 2017 for the 2018 episode and 2023 for the 2025 episode; values are essentially identical across the two vintages and we report the 2017 figure). $d\log \bar{\text{mc}}_i^d$ reports the cumulated upstream marginal-cost increase from evaluating the network solution in (52) at $\rho = 1$, under the levels and logs pass-through assumptions. *Sources*: BEA, Census, tariff database, authors' calculations.

E Robustness of the Pass-through Estimates

This appendix reports three robustness exercises for the local-projection pass-through estimates in the Pass-through estimation subsection of the Results. Throughout, the specification is that of equation (85) (baseline `fwdtreat6`): horizons $h \in \{1, \ldots, 7\}$, six lags and six leads of the tariff shock, one lag of wage growth as control, and item plus date fixed effects with clustered standard errors.

E.1 Goods-only PCE

The main-text PCE estimates pool all NIPA consumption categories. Here we restrict attention to goods (NIPA items 3–149), dropping services. Figure E.1 shows the resulting pass-through.

Figure E.1: PCE Total Tariff Pass-through — Goods only

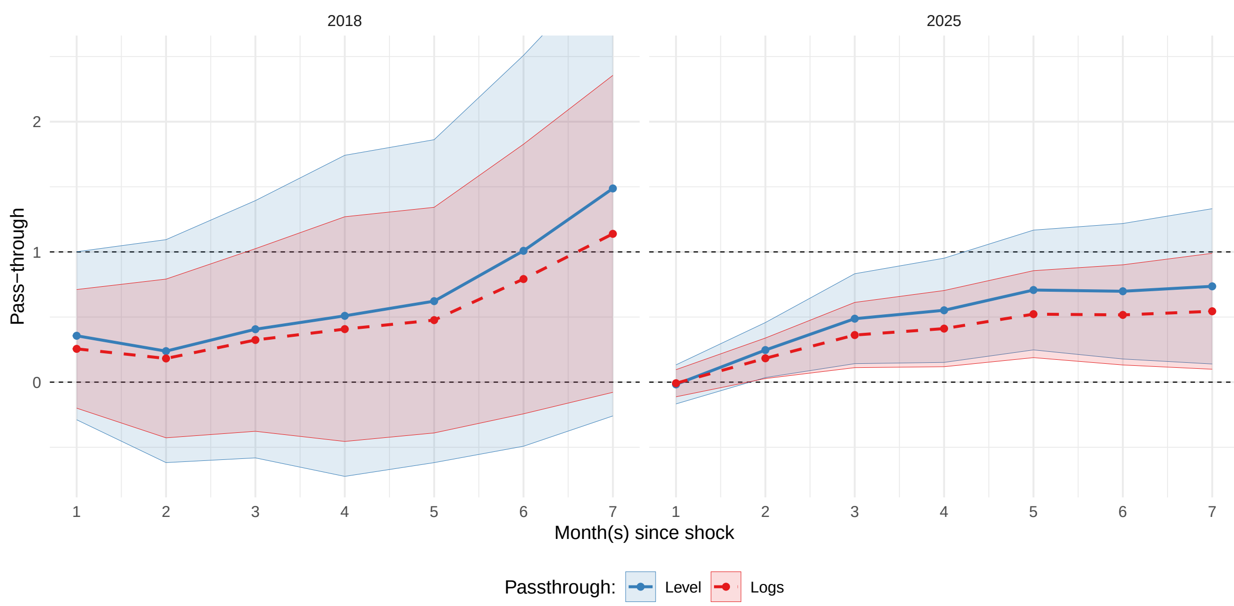

Notes: Same specification and layout as Figure 5, restricted to the goods subset of the NIPA panel (items 3–149). *Sources*: BEA, BLS, authors' calculations.

E.2 Weight robustness

The main-text estimates are unweighted. Figures E.2 and E.3 show the pass-through under three weighting schemes: unweighted (baseline), inverse of the item's monthly inflation volatility, and relative final-demand size (PCE value for PCE, gross output for PPI).

Figure E.2: PCE Total Tariff Pass-through — Weight robustness

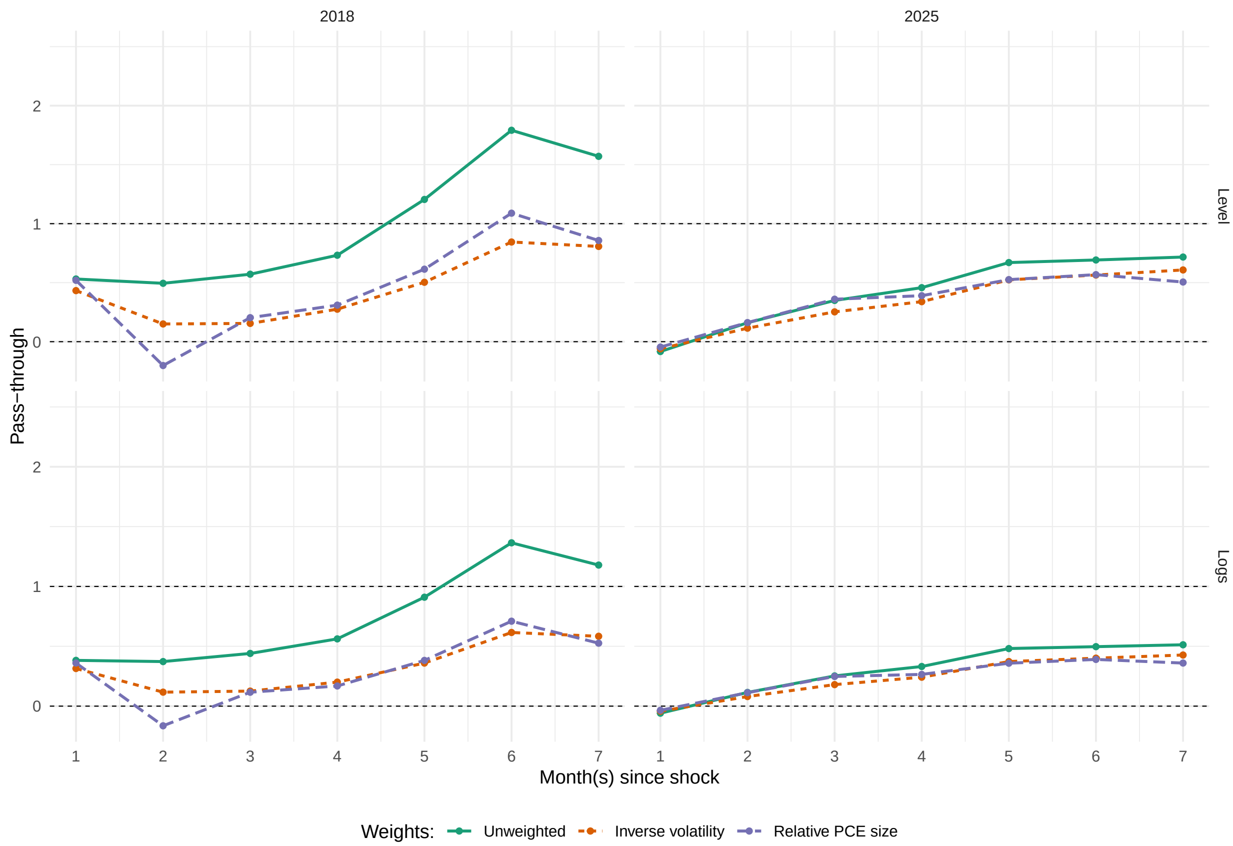

Notes: Pass-through of the predicted total tariff shock to PCE log-price growth under three weighting schemes. Rows: level (top) and logs (bottom) pass-through assumptions. Columns: 2018 and 2025 episodes. Rows share a y-axis capped at 2.5; the black dashed line marks full pass-through. *Sources*: BEA, BLS, authors' calculations.

E.3 By-region shock (Detail-level)

The main-text estimates use the homogeneous shock, which imposes a single import-share profile on every commodity. The by-region pipeline instead allows the import-country mix to vary by commodity and is run at the BEA Detail level (402 commodities). Figures E.4 and E.5 report the PCE and PPI pass-through using the Detail by-region shock.

Figure E.3: PPI Tariff Pass-through — Weight robustness

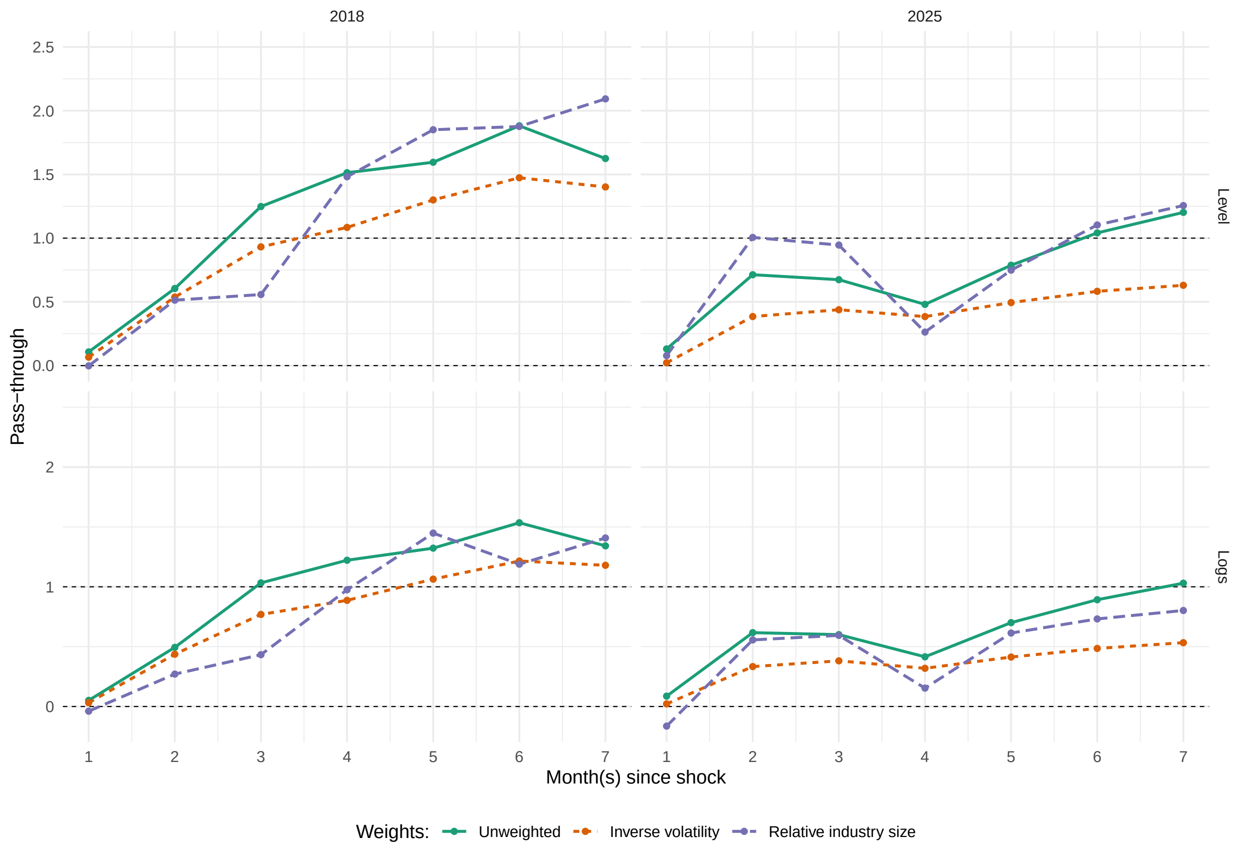

Notes: Pass-through of the predicted tariff shock to PPI log-price growth under three weighting schemes. Layout matches Figure E.2. *Sources*: BEA, BLS, authors' calculations.

Figure E.4: PCE Total Tariff Pass-through — By-region (Detail)

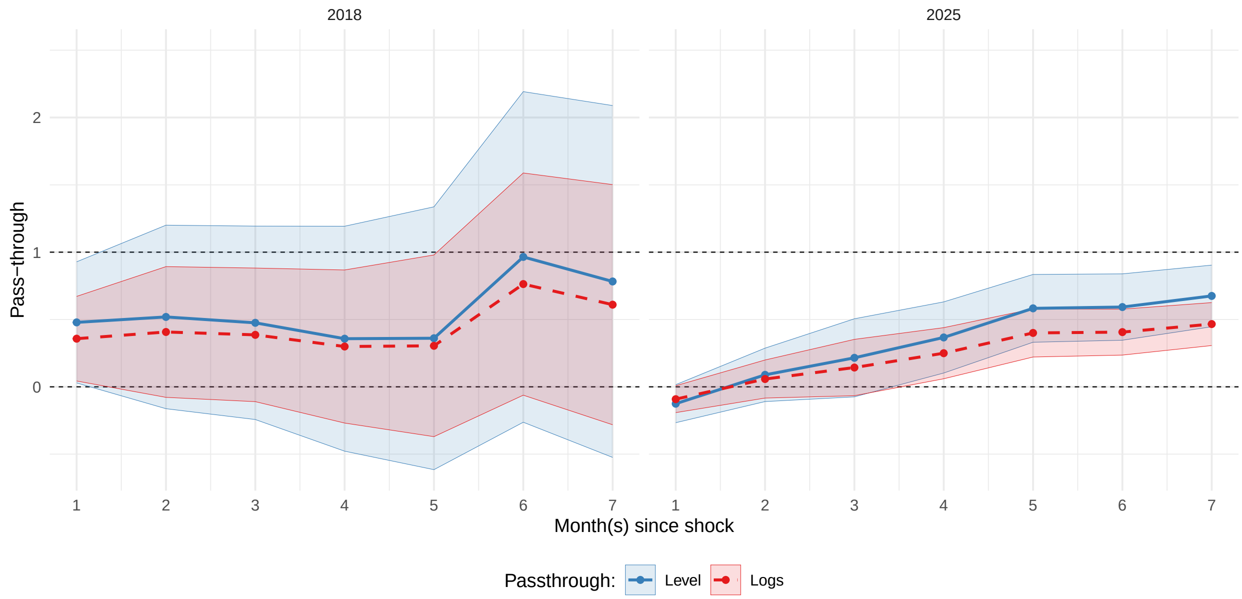

Notes: Same specification and layout as Figure 5, using the BEA Detail-level by-region tariff shock in place of the homogeneous shock. *Sources*: BEA, BLS, authors' calculations.

Figure E.5: PPI Tariff Pass-through — By-region (Detail)

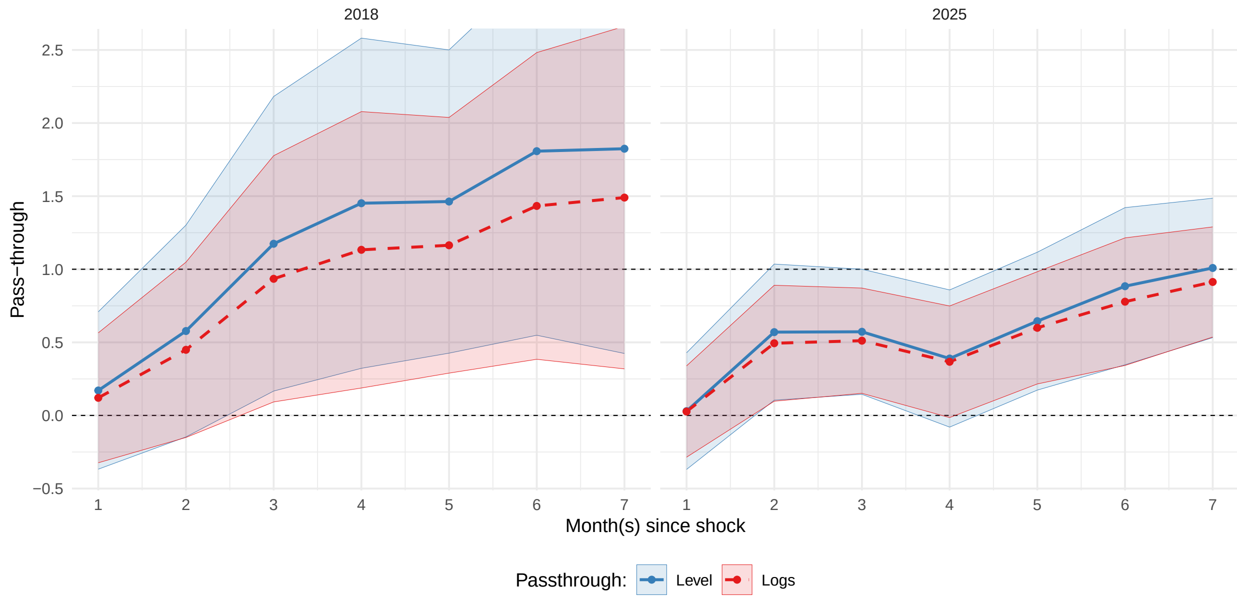

Notes: Same specification and layout as Figure 7, using the BEA Detail-level by-region tariff shock in place of the homogeneous shock. *Sources*: BEA, BLS, authors' calculations.

F Oligopolistic channel Estimation

As a nonparametric check on the parametric RPLS pass-through estimates in Section 6.3, we rerun the restricted-profile estimator replacing the logistic-linear form (87) with a step function on the $G \leq 21$ NAICS-3 manufacturing subsectors present in the active sample. Figure F.1 visualizes the estimated $\widehat{\rho}_g$ across (episode, markup assumption) cells; the diagnostic panel of Table F.2 reports restricted SSE and $\widehat{\lambda}$ for the same specification alongside the parametric reference.

Table F.1 reports the three coefficients of the parametric specification (87) across the two episodes and the two markup assumptions (`cd` = constant-dollar, `cm` = constant-markup). The signs are stable across markup assumptions within each episode but *differ across episodes*, especially on HHI: in 2018 the HHI loading is small and positive (slightly higher pass-through in concentrated industries), while in 2025 it is strongly negative (large attenuation in concentrated industries). The markup loading is positive in both episodes and the import-share loading is uniformly positive.

Table F.1: Revealed pass-through coefficients from the restricted profile-least-squares estimator with $\rho_j = \Lambda(\alpha_{HHI}\, HHI_j^c + \alpha_\mu\, \mu_j^c + \alpha_{s^*}\, s_j^{*c})$. Columns index (episode)/(markup assumption); positive coefficients push ρ_j above the logistic midpoint. Coefficients are in logit-link units and can be converted to marginal effects on ρ_j at the population mean by dividing by four.

	$\hat{\alpha}$ on ρ-logit index			
Observable	2018/cd	2018/cm	2025/cd	2025/cm
HHI	0.406	0.503	-2.850	-3.194
s^*	1.308	1.193	1.999	2.064
μ	2.858	3.453	1.205	1.294

Figure F.1: Nonparametric $\widehat{\rho}_g$ by NAICS-3 subsector.

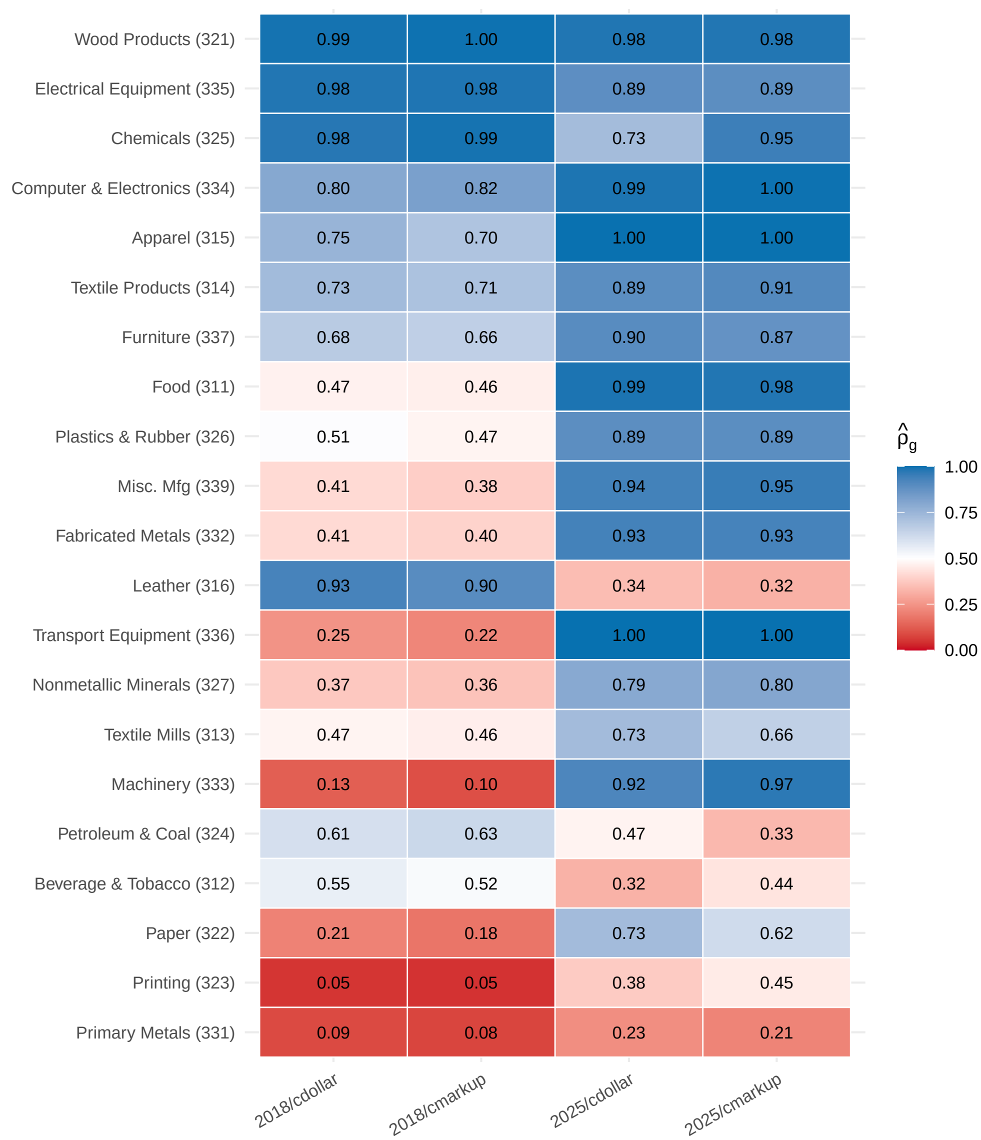

Notes: Step-function RPLS estimate with one $\widehat{\rho}_g \in (0, 1)$ per active NAICS-3 manufacturing subsector (G≤21), estimated by the same restricted profile-least-squares objective as the parametric headline in Section 6.3. Rows are subsectors sorted by their mean across all (episode, assumption) combos; columns are (episode, markup assumption). Diverging red-white-blue color scale centered at 0.5 (the logistic midpoint). The 2018 columns show a much wider spread of $\widehat{\rho}_g$ across subsectors than the 2025 columns, consistent with the main-text interpretation that ρ is absorbing episode-specific impact heterogeneity rather than a stable structural primitive. *Sources*: BEA, BLS, authors' calculations.

Table F.2: Restricted-estimator SSE and diagnostic $\widehat{\lambda}$ (unrestricted contemporaneous LP coefficient at $\widehat{\alpha}$) from the parametric ($G = 3$) and nonparametric NAICS-3 ($G = 21$) specifications. A smaller SSE at $G = 21$ is mechanical — a richer parameterization must fit weakly better. The informative contrast is $\widehat{\lambda}$: both specifications face the same $\beta_0 = 1$ restriction, and their divergence from unity is a pure measure of how binding that restriction is against the data.

		Parametric ($G = 3$)			Nonparametric ($G = 21$)		
Episode	Assm	SSE^R	$\hat{\lambda}$	M	SSE^R	$\hat{\lambda}$	G
2018	cdollar	12.145	0.981	9739	12.040	0.777	21
2018	cmarkup	12.144	0.909	9739	12.044	0.729	21
2025	cdollar	4.268	0.442	4059	4.139	0.537	21
2025	cmarkup	4.274	0.418	4059	4.154	0.549	21