The Effect of Tariffs on U.S. Producer Prices*

Simon Gilchrist[†] Bin Wei[‡] Vivian Z. Yue[§] Egon Zakrajšek[¶]

September 9, 2026

Abstract

This paper estimates the pass-through of import tariffs into U.S. producer prices, using the Census Bureau's Longitudinal Firm Trade Transactions Database linked with Bureau of Labor Statistics producer price microdata. For the 2018–19 U.S.–China trade war we build firm-level effective input tariffs from each firm's own product–country import portfolio, and instrument them with Bartik measures that hold the portfolio at its pre-trade-war composition and move statutory rates. A one percentage point increase in a firm's effective input tariff raises its output price within a year by 0.5–0.6 percentage points per unit of imported-input intensity. Effective tariffs rise by only about two-thirds of the statutory shock, because firms reallocate away from tariffed origins. A model of endogenous supplier choice with oligopolistic pricing and input–output propagation shows that a within-industry regression identifies relative-cost pass-through while a sector-level regression identifies the common component plus propagation. Counting propagation through domestic suppliers, the China tariffs raised manufacturing producer prices by 0.44 percentage points at complete pass-through; a universal 15 percent tariff on imported inputs would raise them by 3.0.

JEL Codes: F13, F14, E31, L11
Keywords: tariffs, pass-through, producer prices, trade war, intermediate inputs, LFTTD, Bartik instrument

*The views expressed in this paper are solely those of the authors and should not be interpreted as reflecting the views of the Federal Reserve Banks of Atlanta and Boston, the Federal Reserve System, the U.S. Census Bureau, or the Bureau of Labor Statistics (BLS). The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was conducted with restricted access to BLS data and was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2878 (CBDRB-FY25-0470).

[†]Department of Economics, New York University, and NBER. Email: sg40@nyu.edu.
[‡]Research Department, Federal Reserve Bank of Atlanta. Email: bin.wei@atl.frb.org.
[§]Department of Economics, Emory University, NBER, and CEPR. Email: vyue@emory.edu.
[¶]Federal Reserve Bank of Boston and CEPR. Email: egon.zakrajsek@bos.frb.org.

1 Introduction

A central question in international trade policy is how far import tariffs travel into domestic prices. The 2018–2019 trade war between the United States and China produced some of the largest tariff increases in modern U.S. history and a natural experiment for answering it. A rapidly growing literature has established that the tariffs were borne almost entirely by U.S. importers: pass-through to import prices was near-complete (Amiti et al., 2019, 2020; Fajgelbaum et al., 2020; Cavallo et al., 2021). Import prices, however, are only the first link in the domestic price chain. Tariffs on intermediate inputs raise the costs of the domestic firms that use them, and how those costs propagate into the prices that domestic producers charge for their own output has received far less attention.

This paper provides firm-level evidence on that link. We combine two confidential microdata sets: the Census Bureau's Longitudinal Firm Trade Transactions Database (LFTTD), which records the universe of U.S. import and export transactions at the firm–HTS10–country–month level together with the duties actually paid, and the producer price microdata of the Bureau of Labor Statistics (BLS), which track the prices that individual manufacturing firms charge for specific products over time. The linkage lets us observe, for each firm in each month, both its complete portfolio of import transactions—products, origins, values, and tariff rates—and the prices of its domestic output. To our knowledge, this is the first paper to estimate the pass-through from input tariffs to output prices at the firm level using linked trade and price microdata.

Our empirical approach extends the industry-level design of Amiti, Redding and Weinstein (2019) in two ways. First, whereas Amiti, Redding and Weinstein measure tariff exposure at the industry level from public trade data and input–output tables, we exploit the granularity of the LFTTD to construct *firm-level* effective input tariffs: for each firm in each month, the import-weighted average of product–country tariff rates, using the firm's own rolling twelve-month import values as weights. Two firms in the same six-digit industry can import very different bundles from very different origins and therefore face very different tariffs; that variation is invisible in industry-level measures, and it is what identifies our estimates. Second, we construct the tariff rates themselves from the customs records in two variants, the de facto rate actually paid and the de jure statutory rate, and combine them with industry-level measures of imported-input intensity built from the Census of Manufactures, the Longitudinal Business Database, and the LFTTD.

An effective tariff built from contemporaneous import weights is endogenous. A firm facing higher tariffs on Chinese inputs can source the same products elsewhere, replace imports with domestic substitutes, or shift its import mix toward products that face lower rates, and each response alters the weights from which its effective tariff is computed. We therefore construct Bartik-style instruments that hold each firm's import composition fixed at its 2017, pre-trade-war level and let statutory tariff rates vary: the share component is the firm's 2017 portfolio, determined before the Section 301 lists existed, and the shift component is the change in de jure rates at the product–country level.

Our empirical findings are four. First, the pass-through of input tariffs into producer prices is positive, statistically significant, and economically meaningful in every specification. Per unit of imported-input intensity—that is, for a firm whose entire material cost is imported—a one percentage point increase in the effective input tariff raises the output price by 0.5 to 0.6 percentage points within twelve months; for the average importing manufacturer, whose imported inputs are about a fifth of variable cost, the response is 0.12 to 0.20 percentage points. Second, the instrumented estimates exceed their least-squares counterparts by half or more—0.58 against 0.35 in the exposure-interaction specification—and the first-stage coefficient is well below one, between 0.64 and 0.74. Both patterns are what endogenous sourcing implies: an effective tariff built from current import shares is a Paasche bound on the true cost increase and understates it by exactly the substitution the tariff induces, so least squares attenuates and the pre-determined instrument moves the effective tariff less than one for one. Third, the estimates are essentially unchanged by industry-by-month fixed effects, so they reflect differences in exposure between firms in the same industry rather than industry-wide shocks; they are also robust to the choice of de jure or de facto rates, to weighting by BLS importance weights, and to extending the sample through 2022. Fourth, a twenty-four-month window yields estimates of similar magnitude, so most of the pass-through is complete within a year.

To read these estimates, we develop a model in which the firm of Halpern et al. (2015)—producing with domestic and imported materials, and choosing the set of origins it imports from subject to a fixed cost per origin—is embedded in an economy with roundabout production, oligopolistic pricing, and an input–output network. The model delivers four results in closed form. The empirical regressor, exposure times the import-weighted change in the tariff factor, is exactly the first-order change in marginal cost at fixed prices and a fixed sourcing set, so a regression of price changes on it has population coefficient one under constant markups; every shortfall from one is a channel to be identified rather than a magnitude to be interpreted. The duty–cost wedge is bounded by a Laspeyres and a Paasche index, and a firm that abandons importing altogether forfeits the entire cost advantage of imported inputs, a jump that is a large multiple of its direct exposure. A firm that drops one taxed origin while retaining others forfeits only the part of that advantage the origin supplied, and stays in the sample with a lower effective tariff. Under Cournot competition among a sector's firms, in the manner of Atkeson and Burstein (2008), a cost shock common to the sector passes into the sector's price index one for one whatever the markups, while a firm-specific shock passes into the firm's relative price at a rate below one. And the change in a firm's domestic share of materials is a sufficient statistic for the effect of the entire import environment on its cost, in the sense of Blaum et al. (2018).

These results say what a firm-level and a sector-level regression identify, and the answers differ. A within-industry comparison identifies from *relative* shocks—one firm's imported inputs become dearer than its rivals'—and a relative cost shock is passed through at the rate at which an oligopolist protects its market share, discounted further because rivals' costs also rise through the price of domestic inputs. Combining the two, the model's benchmark for the within-industry coefficient is 0.68 to 0.82, against a reduced form of 0.375; the residual is what the model does not produce, and whether it reflects nominal rigidity or measurement is a question for the price data. A sector-level regression identifies instead from the component of the shock that is *common* to a sector's firms, which passes through in full and carries propagation through the network on top. The model predicts a sector-level coefficient equal to one plus the projection of propagated price changes on direct exposure across sectors, a number that can be computed from the input–output table without any price and that equals 1.33 on the 2007 table. Held out of everything else in the paper, this prediction is set against the 1.87 of Amiti, Redding and Weinstein (2019): the network accounts for the 0.33, and the remainder is what a sector-level regressor loads on that the cost channel excludes, principally the protection of a sector's own suppliers, which the firm-level design removes by construction because two firms in the same industry face the same upstream protection and different import bills.

We then use the model to quantify aggregate producer-price responses, which requires the propagation that a firm-level regression cannot see. At complete pass-through, the 2018–2019 tariffs on China raised manufacturing producer prices by 0.44 percentage points inclusive of propagation through domestic input prices, and by 0.20 to 0.24 if the residual reflects nominal rigidity within the year. A universal tariff of fifteen percent on imported inputs—the counterfactual motivated by the shift in U.S. trade policy during 2025–2026—raises them by 3.0 percentage points, or 1.4 to 1.7 with the same damping: roughly seven times the China tariffs, for two reasons the model makes explicit. Weighted by imported-input exposure, China accounted for only 16 percent of manufacturers' imports, because the most import-intensive industries source relatively little from China, whereas a universal tariff reaches all of them; and a universal tariff raises every domestic supplier's price at once, so the network multiplier is larger. Allowing sourcing sets to respond would raise these figures, because a firm that leaves an origin forfeits the cost advantage that origin supplied. The sample sees part of this and not the rest: a firm that drops one origin and keeps others remains in it with a lower effective tariff, while a firm that ceases importing altogether leaves it, and it is the second group whose cost jump the data cannot observe. The model bounds the size of that margin by the density of firms at the sourcing cutoff, an object the import data measure.

This paper contributes to several literatures. We contribute most directly to the literature on the effects of the 2018–2019 trade war (Amiti et al., 2019, 2020; Fajgelbaum et al., 2020; Cavallo et al., 2021; Flaaen et al., 2020) by tracing the propagation of tariff costs one step further through the production chain—from import prices to domestic output prices—and by saying what industry-level and firm-level designs each identify. We also contribute to the cost pass-through literature (Goldberg and Knetter, 1997; Gopinath et al., 2010; Amiti et al., 2014): a tariff shifts only the cost side, so it identifies cost pass-through free of the revenue-side effects that complicate exchange-rate pass-through. Finally, our paper draws on the literature that uses U.S. firm-level trade data to study trade dynamics (Bernard et al., 2009) and trade adjustment margins (Monarch, 2021). We add to this body of work by linking LFTTD import transactions directly to BLS producer price microdata, a combination that—to the best of our knowledge—has not been exploited before.

Our model extends Halpern et al. (2015), in which a firm combines domestic and imported inputs and pays a fixed cost for each origin it sources from, so that its set of origins is a choice; related structures with an extensive margin of sourcing appear in Gopinath and Neiman (2014), Antràs et al. (2017), Blaum et al. (2018) and Lu et al. (2024). Where that literature measures the gains from input trade, we use the structure to say what a regression of output prices on a tariff measure recovers when sourcing responds. On the aggregate side we draw on the quantitative trade literature that evaluates changes in trade costs and foreign shocks (Eaton and Kortum, 2002; Arkolakis et al., 2012; Autor et al., 2013; Caliendo and Parro, 2015), most directly on Caliendo and Parro (2015), whose multi-sector model with input–output linkages and tariffs is the antecedent of our propagation block. Our exercise is narrower: it computes producer-price responses through the input–output structure rather than welfare.

The remainder of the paper is organized as follows. Section 2 describes the data and the construction of the tariff and exposure measures. Section 3 presents the empirical specifications and the estimates, with robustness checks. Section 4 develops the model. Section 5 states what the two regression designs identify, reads the estimates through the model, tests the sector-level prediction against Amiti, Redding and Weinstein (2019), and reports the aggregate effects of the China tariffs and of a universal tariff. Section 6 concludes.

2 Data and Measurement

Our analysis requires linking firms' import transactions to their output prices, which we accomplish by combining four confidential datasets (Section 2.1). A central contribution of this paper is the construction of tariff measures at multiple levels of aggregation: we build product-country tariff rates in both de facto and de jure variants from the LFTTD microdata, aggregate them to industry-level input and output tariff measures following Amiti, Redding and Weinstein, and construct firm-level effective tariff measures that exploit each firm's own import portfolio (Section 2.3). We then describe the baseline and extended samples used in estimation (Section 2.2).

2.1 Data Sources

We combine four datasets. The *Longitudinal Firm Trade Transactions Database* (LFTTD), maintained by the U.S. Census Bureau, contains transaction-level records from U.S. customs forms, covering the universe of U.S. merchandise imports since 1992. Each import transaction records the customs value and quantity of the shipment, a 10-digit Harmonized Tariff Schedule (HTS10) product code (comprising approximately 17,000 product categories), the country of origin, and identifiers for both the U.S. importing firm and the foreign exporter. Critically for our purposes, the LFTTD also records the duties collected on each transaction, enabling the construction of transaction-level measures of the effective tariff rate as the ratio of duties paid to customs value.

The *Producer Price Index* (PPI) microdata, published by the Bureau of Labor Statistics (BLS), provide transaction-based prices at the individual good level that are representative of the entire U.S. production sector (see, e.g., Nakamura and Steinsson, 2008; Gilchrist et al., 2017). The BLS samples goods produced by each firm and uniquely identifies them according to their price-determining characteristics, including the type of buyer, the type of market transaction, the method of shipment, the size and units of shipment, the freight type, and the day of the month of the transaction. This granularity ensures that the PPI captures pure price changes—changes attributable solely to market factors—rather than variation arising from shifts in product characteristics or transaction terms. Following Gilchrist et al. (2017), we construct the dependent variable as the h-month log (quality-adjusted) price change, $\Delta_h \log(p_{j,k,i,t}) = \log(p_{j,k,i,t}) - \log(p_{j,k,i,t-h})$.[1] In the paper, we consider both $h = 12$ and $h = 24$. The 12-month change is our baseline specification, and we refer to it simply as $\Delta \log(p_{j,k,i,t})$.

The *Census of Manufactures* (CMF), conducted every five years as part of the U.S. Economic Census, covers the universe of manufacturing establishments with one or more paid employees and provides detailed data on the total value of shipments, annual payroll, production worker wages and hours, cost of materials consumed (by material code), capital expenditures, and energy costs. We merge the CMF with the *Longitudinal Business Database* (LBD) by establishment identifier to aggregate

[1]A distinctive feature of the PPI micro data is the BLS's quality-adjustment procedures. To ensure that the index measures only pure price changes, the BLS excludes any change in the recorded price that is attributable to a change in the "base price," such as physical modifications to the sampled item or changes in the way it is sold. The BLS employs several quality-adjustment methods depending on the type of change and the information available from the respondent (US Bureau of Labor Statistics, 2014). Letting $p^*_{j,k,i,t}$ denote the recorded transaction price of good k produced by firm i and $p^b_{j,k,i,t}$ the corresponding base price, the quality-adjusted good-level price is defined as $p^*_{j,k,i,t} \equiv p^*_{j,k,i,t}/p^b_{j,k,i,t}$ (Gilchrist et al., 2017).

establishment-level records to the firm level and construct measures of firm-level sales, wage bills, and material costs. The 2017 CMF—the most recent vintage available prior to the onset of the 2018–2019 U.S.–China trade war—provides a snapshot of firms' cost structures immediately before the tariff shock.

2.2 Sample

Our baseline sample covers January 2017 to December 2019, the period of active tariff escalation and implementation. The extended sample runs from January 2017 to December 2022, incorporating the partial tariff rollbacks and the COVID-19 pandemic period.

We merge the PPI and LFTTD databases using a multi-step name-matching algorithm designed to maximize match rates while minimizing false positives. Because the two databases do not share a common firm identifier, we rely on business names as the primary matching variable (see Appendix A for more details).

The resulting sample is a panel of approximately 290,000–310,000 firm-product-month observations from roughly 2,900–3,200 unique firms in the baseline sample, and 540,000–580,000 observations from 3,400–3,600 firms in the extended sample. These firms span a broad range of manufacturing industries, including food products, chemicals, metals, machinery, electrical equipment, and transportation equipment.

2.3 Tariff Measures

We construct tariff measures at three levels of aggregation. First, we build product-country-month tariff rates in both de facto and de jure variants directly from the LFTTD microdata. These product-country rates are the building blocks for all subsequent measures. Second, we aggregate them to industry-level input and output tariff changes following the methodology of Amiti, Redding and Weinstein, using BEA input-output tables to trace the propagation of tariffs through the production network. Third, and most novel, we construct firm-level effective tariff measures that exploit each firm's own product-country import portfolio observed in the LFTTD, providing a substantially richer source of cross-sectional variation than the industry-level measures.

Product-Country Tariff Rates

A key innovation of this paper is the construction of firm-level effective and de jure tariff measures. Both rates are constructed from the LFTTD microdata rather than from external tariff schedule databases.

De facto (effective) tariff rates. The LFTTD records, for each import transaction by firm j of HTS10 product h from country c in month t, both the import value $v^{\text{import}}_{j,h,c,t}$ and the duty payment $v^{\text{import}}_{j,h,c,t} \cdot \tau_{j,h,c,t}$, where $\tau_{j,h,c,t}$ is the ad valorem tariff rate effectively applied to that shipment. Because different firms importing the same product-country pair may face different effective rates—due to tariff exemptions, duty drawback programs, product exclusions, or differences in rate provision classifications—we aggregate across firms to obtain the product-country-month effective tariff rate:

$$\tau^{\text{eff}}_{h,c,t} = \frac{\sum_j v^{\text{import}}_{j,h,c,t}\, \tau_{j,h,c,t}}{\sum_j v^{\text{import}}_{j,h,c,t}}. \tag{1}$$

This import-value-weighted average measures the tariff rate that firms *actually pay* on a given product from a given origin in a given month. It reflects all exemptions, exclusions, and administrative adjustments as they are realized in practice.

De jure (statutory) tariff rates. The de jure tariff rate aims to capture the officially announced statutory rate, abstracting from firm-specific exemptions that reduce the tariff burden below the announced level. To construct this measure, we exploit the *rate provision code* recorded in the LFTTD for each transaction. The rate provision code classifies each import entry according to the specific tariff provision under which it enters the United States—including standard duty rates, various preferential programs, temporary duty suspensions, and tariff exemptions (such as foreign trade zone entries and duty-free provisions).[2] We exclude transactions that enter under exemption provisions and focus on the maximum tariff rate observed across the remaining non-exempt transactions within each (h, c, t) cell. This maximum rate is taken as the de jure tariff rate $\tau^{\text{dejure}}_{h,c,t}$.

The distinction between de facto and de jure rates is economically important. During the 2018–2019 trade war, the U.S. government announced large tariff increases on Chinese imports, but simultaneously implemented a product exclusion process that

[2] See https://www.census.gov/foreign-trade/reference/codes/rp.html.

exempted certain transactions from the new duties. As a result, the effective tariff rate that firms actually paid could be substantially lower than the announced statutory rate. By constructing both measures from the LFTTD, we can assess whether our pass-through estimates are sensitive to this distinction. As we show below, the de jure and de facto Bartik instruments yield very similar second-stage estimates, indicating that our findings are not driven by the specific treatment of tariff exemptions.

Lastly, we backfill and forward fill missing effective tariff rates to create a balanced panel and compute 12-month changes:

$$\Delta \log \left(1 + \tau_{h,c,t}^{eff}\right) = \log \left(1 + \tau_{h,c,t}^{eff}\right) - \log \left(1 + \tau_{h,c,t-12}^{eff}\right). \tag{2}$$

Similarly, we compute 12-month changes in de jure tariff rates and refer to them as $\Delta \log \left(1 + \tau_{h,c,t}^{dejure}\right)$.

Output Tariffs

Before we measure input tariffs, we construct an industry-level *output* tariff measure that captures the *competition channel*: the degree to which import tariffs shield domestic producers from foreign competitors selling similar products. When tariffs are imposed on finished goods that compete directly with a domestic industry's output, they raise the effective price of foreign substitutes, reducing competitive pressure and allowing domestic firms to raise their own prices—even absent any change in input costs. To measure this channel, we follow Amiti, Redding and Weinstein and compute, for each 6-digit NAICS industry i, the import-weighted average change in tariff rates on products classified within that industry:

$$\Delta \log \left(1 + \tau_{i,t}^{\text{output, eff}}\right) \equiv \sum_{(h,c),\, h \in i} \frac{v_{h,c,0}^{\text{import}}}{\sum_{(h',c'),\, h' \in i} v_{h',c',0}^{\text{import}}} \Delta \log \left(1 + \tau_{h,c,t}^{eff}\right), \tag{3}$$

where $v_{h,c,0}^{\text{import}} \equiv \sum_j v_{j,h,c,0}^{\text{import}}$ denotes total 2017 U.S. imports of HS product h from country c, and the summation in the denominator aggregates over all product–country pairs (h', c') with h' classified within industry i. The weights are held fixed at their 2017 (pre-trade-war) values to avoid endogeneity from contemporaneous import adjustments. We map HS codes to 6-digit NAICS industries using the Census Bureau's official concordance files.[3]

The output tariff $\Delta \log(1 + \tau_{i,t}^{\text{output, eff}})$ varies only at the industry-by-month level, in contrast to the firm-level input tariff. An increase in this measure indicates that industry i's domestic products face less foreign competition because tariffs have raised the landed cost of competing imports. In our regression framework, a positive coefficient on the output tariff captures this competition effect: domestic producers exploit the tariff-induced reduction in import competition by raising markups and, consequently, output prices.

Input Tariffs

Import tariffs on intermediate inputs—raw materials, components, and semi-finished goods that domestic manufacturers purchase from abroad—raise the cost of production for firms that rely on imported intermediates. Unlike output tariffs, which affect competitive pressure from foreign rivals, input tariffs operate through the *cost channel*: they increase firms' marginal costs and, to the extent that these cost increases are passed through to customers, raise domestic producer prices. Measuring the input tariff burden that each firm faces is therefore central to estimating the inflationary consequences of trade policy.

A key innovation of this paper is the construction of *firm-level* effective input tariff measures that exploit the granularity of the LFTTD. Because the LFTTD records the universe of U.S. import transactions at the firm-product-country-month level, we can observe each firm's complete import portfolio and compute a tariff exposure measure that reflects the specific products each firm imports, the countries from which it sources, and the relative importance of each product–country pair in the firm's total import bill.

Using the LFTTD, we first compute the effective (de facto) tariff rate at the product-country-month level as the import-value-weighted average across firms, defined in equation (1). For firm j in month t, we measure the rolling 12-month import value $v_{(h,c)\to j,t}^{(12m)}$ of good h from country c between months $t-11$ and t, and compute

[3] Available at https://www.census.gov/foreign-trade/reference/codes/index.html; we use the "HTS_Import_Export_Concordance" file.

the effective input tariff level as:

$$\sum_{(h,c)} \frac{v^{(12m)}_{(h,c)\to j,t}}{\sum_{(h,c)} v^{(12m)}_{(h,c)\to j,t}} \log\left(1+\tau^{eff}_{h,c,t}\right).$$

Constructing the analogous quantity at $t-12$ using import values between $t-23$ and $t-12$, the 12-month change in the firm's effective input tariff is:

$$\begin{aligned}\Delta\log\left(1+\tau^{\text{input,eff}}_{j,t}\right) = &\sum_{(h,c)} \frac{v^{(12m)}_{(h,c)\to j,t}}{\sum_{(h,c)} v^{(12m)}_{(h,c)\to j,t}} \log\left(1+\tau^{eff}_{h,c,t}\right) \\ &- \sum_{(h,c)} \frac{v^{(12m)}_{(h,c)\to j,t-12}}{\sum_{(h,c)} v^{(12m)}_{(h,c)\to j,t-12}} \log\left(1+\tau^{eff}_{h,c,t-12}\right).\end{aligned} \tag{4}$$

The rolling 12-month window smooths out seasonal fluctuations in trade patterns while remaining responsive to medium-term shifts in sourcing. Two firms classified in the same NAICS industry may import very different bundles of inputs from different countries and thus face very different effective tariff rates—variation that is invisible in industry-level measures.

It is instructive to compare our firm-level input tariff measures with the industry-level approach pioneered by Amiti et al. (2019). Because Amiti, Redding and Weinstein work with publicly available data, they cannot observe individual firms' import portfolios. Instead, they rely on the BEA's 2007 input-output (I-O) table to approximate each industry's exposure to tariffs on imported intermediates. Let $v^{\text{input}}_{i\to i^o}$ denote the value of intermediate inputs that downstream industry i^o sources from upstream industry i. The ARW input tariff rate change for industry i^o is:

$$\Delta\log\left(1+\tau^{\text{input, ARW}}_{i^o,t}\right) \equiv \sum_i \frac{v^{\text{input}}_{i\to i^o}}{v^{\text{input}}_{i^o}} \Delta\log\left(1+\tau^{\text{output}}_{i,t}\right), \tag{5}$$

where $v^{\text{input}}_{i^o} \equiv \sum_i v^{\text{input}}_{i\to i^o}$ is total intermediate inputs sourced by industry i^o from all upstream industries, and $\Delta\log(1+\tau^{\text{output}}_{i,t})$ is the change in the output tariff for upstream industry i. Intuitively, this measure assigns each downstream industry a weighted average of its upstream industries' output tariffs, with weights proportional to the I-O table's expenditure shares.

Our firm-level measures offer three advantages over this industry-level approach.

First, the I-O table captures *average* inter-industry linkages but cannot reflect the substantial within-industry heterogeneity in sourcing patterns. Two firms in the same 6-digit NAICS industry may source intermediates from entirely different countries—one primarily from China, the other from Germany—and therefore face very different tariff exposures during the trade war. Our firm-level measures capture this variation directly. Second, the I-O table does not distinguish between domestically produced and imported intermediates within each upstream industry, nor does it identify the country of origin of imports. Our measures, built from actual firm-level import transactions, incorporate both the product composition and the geographic origin of each firm's imports. Third, the BEA I-O table is updated infrequently (the version used by Amiti, Redding and Weinstein dates to 2007), so the industry-level weights may not accurately reflect the production network at the time of the trade war. Our firm-level weights are constructed from contemporaneous or near-contemporaneous LFTTD transactions, ensuring that the measured tariff exposure reflects actual sourcing patterns during the relevant period.

Figure 1 plots the value-weighted average ad valorem duty rate on U.S. merchandise imports from 2001 to 2022, broken out by major source countries. From 2001 through early 2018, tariff rates were low and remarkably stable across all source countries, generally hovering below 5 percent. As successive rounds of Section 301 tariffs took effect, tariff rates on imports from China surged from roughly 3 percent to approximately 15–20 percent by late 2019, although no other major trading partner experienced a comparable increase.

Tariff Exposure

To measure the degree to which each industry relies on imported intermediates, we construct two exposure variables: the import share (IS_i) and the imported input intensity (III_i). Both are computed at the industry level using 2017 data from the LFTTD and CMF.

The import share captures the fraction of domestic absorption accounted for by imports:

$$IS_i = \frac{\text{mvalue}_i}{\text{mvalue}_i + \text{sales}_i - \text{xvalue}_i}, \tag{6}$$

where mvalue_i and xvalue_i denote total import and export values for industry i in 2017 from the LFTTD, and sales_i denotes total shipments from the CMF. The denominator

Figure 1: Ad Valorem Duty Rates by Source Country

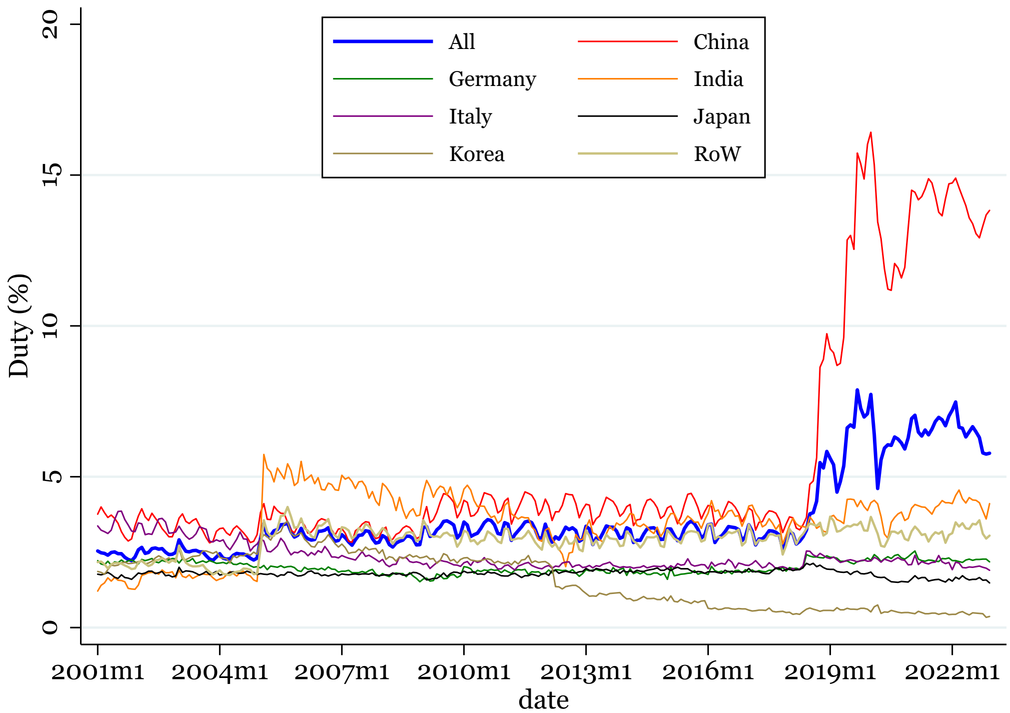

NOTE: This figure plots the value-weighted average ad valorem import duty rate (duties collected as a percentage of customs value) for U.S. merchandise imports from 2001 through 2022, disaggregated by major source country: China, Germany, India, Italy, Japan, Korea, and the rest of the world (RoW). The "All" line shows the simple average across all source countries. The sharp increase in the China series beginning in mid-2018 reflects the successive rounds of Section 301 tariffs imposed during the U.S.–China trade war.
SOURCE: Authors' calculations using the U.S. Census Bureau's Longitudinal Foreign Trade Transactions Database.

is industry-level absorption: domestic shipments net of exports plus imports.

The imported input intensity scales the import share by the share of total variable cost attributable to materials:

$$III_i = \frac{IS_i \times \text{materialcost}_i}{\text{materialcost}_i + \text{wagebill}_i}, \tag{7}$$

where materialcost_i and wagebill_i are total material costs and total wages for industry i from the CMF. The denominator approximates total variable cost (materials plus labor), so III_i measures the fraction of variable cost that comes from imported intermediates.

Summary Statistics. Table 1 reports summary statistics for key variables used in our empirical analysis. From the table, 12-month changes in both input and output tariffs are shown to be right-skewed: the median firm experiences zero or near-zero tariff change, while the 75th percentile shows meaningful positive changes of 0.45 and 0.53 percentage points, respectively, reflecting the concentrated nature of the trade war shock. The table documents substantial variation in import exposure across industries. The mean industry-level imported input intensity is 20.6 percent, indicating that roughly one-fifth of the average manufacturing industry's material costs are sourced from abroad. The interquartile ranges are wide, confirming that import exposure varies substantially across industries. For the interaction terms in the last two rows, their medians are at or near zero, consistent with the fact that many observations experience negligible tariff-exposure products, while the right tail captures firms in import-intensive industries hit by the tariff shock.

Table 1: Selected Summary Statistics

	Mean	StdDev	P25	P50	P75
Δ_{12} Input tariffs	0.46	3.29	−0.15	0.0	0.45
Δ_{12} Output tariffs	0.54	1.79	−0.02	0.01	0.53
Δ_{12} log PPI (industry)	3.66	11.20	0.0	0.0	5.11
Import share, ind. (IS_i)	30.0	25.2	8.4	23.6	45.0
Imp. input intensity, ind. (III_i)	20.6	16.9	6.6	17.3	31.5
$III_i \times \Delta_{12}$ input tariff	0.11	0.85	−0.02	0.0	0.07
$IS_i \times \Delta_{12}$ output tariff	0.17	0.64	0.00	0.0	0.09

NOTE: This table reports summary statistics for selected variables in the merged LFTTD–PPI panel dataset during 2017–2022. It summarize the regression variables used in Sections 3.1 and 3.2; percentile columns report P25, P50, and P75. Tariff changes are 12-month log changes, expressed in percentage points. Industry-level measures of import share (IS_i) and imported input intensity (III_i), expressed in percent, are constructed from the 2017 CMF, LBD, and LFTTD.

3 Empirical Methodology and Results

This section presents our empirical framework and main results. We organize the analysis in four parts. Section 3.1 introduces our baseline specifications—which estimate the unconditional average pass-through from input tariffs to output prices—and reports OLS and IV results. Section 3.2 extends the analysis to exposure-interaction specifications that allow the tariff effect to scale with each industry's reliance on imported intermediates. Section 3.3 discusses the distinct economic interpretations of the two sets of results. Section 3.4 presents robustness checks.

3.1 Baseline Pass-Through Estimates

Our baseline regressions relate 12-month changes in firm-level output prices to changes in effective input tariffs. We consider two specifications that vary in whether the output tariff competition channel is included or absorbed by fixed effects.

Our first specification estimates the pass-through of both input and output tariffs to producer prices:

$$(I): \quad \Delta \log\left(p_{j,k,i,t}\right) = \beta\, \Delta \log\left(1 + \tau_{j,t}^{\text{input,eff}}\right) + \gamma\, \Delta \log\left(1 + \tau_{i,t}^{\text{output,eff}}\right) + \eta_j + \delta_t + \epsilon_{j,k,i,t},$$

where $\Delta \log(p_{j,k,i,t})$ is the 12-month log price change for product k of firm j in industry i in month t; $\Delta \log(1 + \tau_{j,t}^{\text{input,eff}})$ is the 12-month change in the firm's effective input tariff; $\Delta \log(1 + \tau_{i,t}^{\text{output,eff}})$ is the change in the industry-level output tariff; η_j denotes firm fixed effects; and δ_t denotes time fixed effects. The dependent and independent variables are winsorized at the 0.1% and 99.9% levels. The coefficient β captures the input cost channel (how much of a tariff increase on imported intermediates is passed through to the firm's output price), while γ captures the competition channel (how much the firm raises its output price when tariffs shield it from foreign competition).

The firm fixed effects η_j absorb time-invariant firm characteristics that influence price-setting behavior—including the firm's average inflation rate, its product mix, its market position, and other persistent features of its cost and demand environment. Since many firms in our sample produce multiple products, the firm fixed effect also controls for the fact that a given firm's prices for different products k may share a common level or trend. Identification of β therefore comes from *within-firm* variation in tariff exposure over time: how the same firm's output prices change as the tariffs on its imported inputs change. Our second specification replaces the time fixed effects and the output tariff variable with industry-by-month fixed effects:

$$(II): \quad \Delta \log\left(p_{j,k,i,t}\right) = \beta\, \Delta \log\left(1 + \tau_{j,t}^{\text{input,eff}}\right) + \eta_j + \delta_{i,t} + \epsilon_{j,k,i,t},$$

where $\delta_{i,t}$ denotes industry-by-month fixed effects. This is a substantially more demanding specification. The industry-by-month fixed effects absorb *all* time-varying industry-level shocks—including demand shifts, technology changes, sectoral price trends, and, importantly, the output tariff itself (which varies only at the industry-by-month level). Identification of β in this specification comes entirely from *within-industry* variation in firm-level tariff exposure: how firms within the same industry that face different input tariffs (because they import different products from different countries) adjust their output prices differently. A positive and significant β in Specification (II) therefore rules out the possibility that our results are driven by industry-wide confounders correlated with tariff changes.

For each specification, we estimate both unweighted and weighted regressions. The weighted regressions use the BLS importance weight assigned to each price quote in the Producer Price Index program. This weight reflects the item's share in the total value of shipments within its product category, so that price quotes from economically larger product lines receive proportionally more influence in the estimation. The unweighted regressions treat each firm-product-month observation equally, giving equal voice to small and large product lines, and therefore estimate the pass-through for the *typical* firm-product observation in the sample. The weighted regressions, by contrast, estimate the pass-through for the *economically representative* firm-product observation—the one that matters most for the aggregate PPI. If pass-through differs systematically by firm size or product importance (for example, because larger firms have more diversified sourcing or greater market power to absorb cost shocks into margins), the weighted and unweighted estimates will diverge. A finding that the weighted estimates are modestly smaller than the unweighted estimates would suggest that economically larger product lines exhibit somewhat lower pass-through, consistent with the prediction that firms with larger market shares charge variable markups and absorb a greater fraction of cost shocks (Atkeson and Burstein, 2008; Gopinath et al., 2010). Standard errors are clustered at the firm level throughout to account for the correlation of pricing decisions across products within the same firm and over time.

OLS Results. In specifications (I)–(II), the coefficient β on $\Delta \log(1 + \tau_{j,t}^{\text{input,eff}})$ measures the *unconditional average* pass-through: by how many percentage points the output price rises when a firm's effective input tariff increases by 1 percentage point, averaging across all importing firms in the sample regardless of their degree of reliance on imported inputs.

Table 2: OLS Regression Results

	Specification (I)		Specification (II)	
	Unwt.	Wt.	Unwt.	Wt.
$\Delta \log(1 + d^{\text{input,eff}})$	0.135	0.131	0.100	0.090
	(0.026)	(0.025)	(0.023)	(0.024)
$\Delta \log(1 + d^{\text{output,eff}})$	0.264	0.287		
	(0.068)	(0.073)		
R^2	0.152	0.177	0.211	0.260
Observations	299k	299k	319k	319k
Firms	3,100	3,100	3,300	3,300
Firm FE	Yes	Yes	Yes	Yes
Industry×Month FE	No	No	Yes	Yes

Notes: Standard errors clustered at the firm level in parentheses. Observation counts are truncated for disclosure avoidance. The dependent variable is the 12-month log change in producer price.

Table 2 presents OLS estimates for the 12-month horizon using the 2017–2019 sample. The coefficient $\hat{\beta}$ on the input tariff variable is positive and statistically significant across all specifications. In specification (I) with firm and time fixed effects, the point estimate is 0.135 (s.e. 0.026) unweighted and 0.131 (0.025) weighted. These estimates imply that, on average, a 10 percentage point increase in a firm's effective input tariff raises its output price by roughly 1.3 percent over the following 12 months. Moving to specification (II), which replaces time fixed effects with industry-by-month fixed effects $\delta_{i\times t}$, the estimates decline modestly to 0.100 (0.023) and 0.090 (0.024). The attenuation is expected: industry-by-month fixed effects absorb all time-varying industry-level shocks, including the output tariff channel, so that identification comes entirely from within-industry, across-firm variation in tariff exposure. The fact that $\hat{\beta}$ remains positive and highly significant in this demanding specification indicates that the result is not driven by industry-level confounders such as sector-wide demand shifts or common cost trends.

The output tariff coefficient $\hat{\gamma}$ in specification (I) is 0.264 (0.068) unweighted and 0.287 (0.073) weighted, providing evidence for a competition channel: tariffs on goods that directly compete with a firm's products also raise domestic producer prices, as import protection reduces competitive pressure from foreign suppliers. This coefficient is not identified in specification (II), where it is absorbed by the industry-by-month fixed effects.

Endogeneity of Import Weights and IV Strategy

While the OLS estimates provide a useful benchmark, there are reasons to believe they understate the true causal pass-through. A key identification challenge is that the effective tariff rate $\Delta \log(1 + \tau_{j,t}^{\text{input,eff}})$ is constructed from contemporaneous import weights, which reflect the firm's endogenous sourcing decisions. As we formalize in Section 4, firms facing higher tariffs on imports from a given country may substitute toward alternative suppliers, shift to domestic inputs, or adjust the product composition of their imports—all of which alter the measured effective tariff rate. Because these sourcing adjustments are themselves responses to tariff changes, OLS estimation of Specifications (I)–(II) produces downward-biased estimates of the causal pass-through: firms with the largest tariff increases are precisely those that substitute most aggressively, compressing the measured tariff change and attenuating $\hat{\beta}$ toward zero.

To address this endogeneity, we construct Bartik-style instrumental variables that hold each firm's import composition fixed at its pre-trade-war (2017) level while allowing product-country tariff rates to vary over time. The "share" component is firm j's 2017 import portfolio—determined before the trade war was announced—and the "shift" component is the subsequent change in product-country tariff rates.

Our instrument uses de jure tariff rates as the shift component:

$$\Delta \log\left(1 + \tau_{j,t}^{\text{IV,dejure}}\right) \equiv \sum_{(h,c)} \frac{v_{(h,c)\to j}^{(2017)}}{\sum_{(h,c)} v_{(h,c)\to j}^{(2017)}} \Delta \log\left(1 + \tau_{h,c,t}^{\text{dejure}}\right). \tag{8}$$

This instrument has two sources of variation: cross-sectional variation arising from differences in firms' pre-trade-war import portfolios (which products they import, and from which countries), and time-series variation arising from the staggered implementation of tariff increases across product categories and trade-war lists. We favor the de jure instrument because the statutory tariff rates are directly announced by trade policy authorities and are therefore free of the firm-specific exemptions, exclusions, and administrative adjustments that affect de facto rates. The de jure rates thus provide a cleaner measure of the exogenous policy shock.

In the first stage, we regress the endogenous regressor on the corresponding IV, controlling for the same fixed effects as in the second stage. For the exposure-interaction specifications, the first stage regresses $III_i \times \Delta \log(1+\tau_{j,t}^{\text{input,eff}})$ on $III_i \times \Delta \log(1+\tau_{j,t}^{\text{IV,dejure}})$, along with firm fixed effects and either time or industry-by-time fixed effects.

Identification assumptions. The validity of the Bartik instrument requires two conditions: relevance and exogeneity. Regarding the first condition of relevance, the pre-determined tariff exposure must be a strong predictor of the contemporaneous effective tariff. This is confirmed by the first-stage F-statistics, which uniformly exceed 200 for the de jure IV and 400 for the de facto IV (see Tables 3 and B-1), far above the Stock and Yogo (2005) critical values for weak instruments.

Regarding the second condition of exogeneity, the 2017 import shares must be uncorrelated with subsequent shocks to producer prices, conditional on firm and time fixed effects. This condition is plausible for two reasons. On the "share" side, the import composition was determined before the trade war tariffs were announced (the first Section 301 tariff list was not published until June 2018), so the 2017 import portfolios could not have anticipated the tariff shock. On the "shift" side, the tariff changes are set at the broad product-country level by trade policy authorities in response to geopolitical and macroeconomic considerations, not in response to individual firm pricing decisions. This identification argument follows the framework of Borusyak, Hull and Jaravel (2022), who show that Bartik instruments are valid when the shifts (here, tariff changes) are as-good-as-randomly assigned, even if the shares (here, firm-level import portfolios) are endogenous in levels.

IV Results. Table 3 reports 2SLS estimates using the de jure Bartik IV. The IV estimates are systematically larger than OLS, as the duty–cost wedge of Proposition 1(c) implies: an effective tariff built from contemporaneous import shares is a Paasche bound on the cost increase the firm actually experiences, understating it by exactly the substitution that the tariff induces. In specification (I), the IV estimate rises to 0.203 (0.035) unweighted and 0.197 (0.036) weighted—roughly 50 percent larger than the corresponding OLS estimates. In specification (II), the IV estimates

Table 3: IV Regression Results, De Jure Tariffs

Panel A: Second-stage results				
	Specification (I)		Specification (II)	
	Unwt.	Wt.	Unwt.	Wt.
$\Delta \log(1 + d^{\text{input,eff}})$	0.203	0.197	0.144	0.121
	(0.035)	(0.036)	(0.033)	(0.036)
$\Delta \log(1 + d^{\text{output,eff}})$	0.244	0.268		
	(0.069)	(0.074)		
Observations	296k	296k	316k	316k
Firms	3,000	3,000	3,200	3,200
Panel B: First-stage results				
$\Delta \log(1 + d^{\text{IV,dejure}})$	0.738	0.740	0.725	0.728
	(0.028)	(0.027)	(0.026)	(0.025)
R^2	0.661	0.659	0.717	0.718
F-statistic	694.1	755.7	806.5	874.0
Firm FE	Yes	Yes	Yes	Yes
Industry×Month FE	No	No	Yes	Yes

Notes: Panel A reports 2SLS estimates using the de jure Bartik IV from equation (8). Panel B reports first-stage results. Standard errors clustered at the firm level. F-statistic is the Kleibergen-Paap weak-instrument test.

are 0.144 (0.033) and 0.121 (0.036). The IV-to-OLS ratio is approximately 1.5 across specifications, consistent with the prediction that endogenous import substitution attenuates OLS toward zero.

The first-stage results (Panel B) confirm a powerful relationship between the Bartik instrument and the endogenous tariff measure. The first-stage coefficient on $\Delta \log(1 + d^{\text{IV,dejure}})$ is approximately 0.73–0.74, with R^2 values of 0.66–0.72 and F-statistics of 694–874, far exceeding conventional thresholds for instrument relevance. That the first-stage coefficient is below unity is itself informative: it implies that the contemporaneous effective tariff moves less than one-for-one with the pre-determined Bartik instrument, exactly as predicted by endogenous import substitution.

3.2 Pass-Through with Import Exposure

We further refine the analysis by interacting the tariff variables with measures of import exposure. Following Amiti and Heise (2025), we construct industry-level exposure measures from the CMF, LBD, and LFTTD. We merge the CMF and LBD by establishment identifier and retain observations with positive sales in 2017. From the CMF we obtain a firm's wage bills as the sum of salaries and wages (sw) and labor compensation (lc), as well as total material cost (cm). We convert them from thousands of dollars to dollars and refer to them as wagebill_j and materialcost_j, respectively, for firm j. Negative values are replaced by zero.

The exposure interactions allow the tariff effect to scale with each industry's reliance on imported intermediates, separating the *exposure margin* (how much of the firm's cost base is affected) from the *pass-through margin* (how much of the cost increase is reflected in output prices). The exposure-interaction counterparts of Specifications (I) and (II) are:

$$
\begin{aligned}
(III): \quad \Delta \log (p_{j,k,i,t}) = {} & \beta \cdot III_i \times \Delta \log \left(1 + \tau_{j,t}^{\text{input,eff}}\right) \\
& + \gamma \cdot IS_i \times \Delta \log \left(1 + \tau_{i,t}^{\text{output,eff}}\right) + \eta_j + \delta_t + \epsilon_{j,k,i,t},
\end{aligned}
$$

and

$$
(IV): \quad \Delta \log (p_{j,k,i,t}) = \beta \cdot III_i \times \Delta \log \left(1 + \tau_{j,t}^{\text{input,eff}}\right) + \eta_j + \delta_{i,t} + \epsilon_{j,k,i,t}.
$$

In Specification (III), the coefficient β on the interaction $III_i \times \Delta \log(1 + \tau_{j,t}^{\text{input,eff}})$ answers the question: for a firm whose entire material cost base consists of imported intermediates ($III = 1$), what is the output price response to a 1 percentage point tariff increase? The second regressor, $IS_i \times \Delta \log(1 + \tau_{i,t}^{\text{output,eff}})$, captures the competition channel scaled by the industry's import share. Specification (IV) replaces time fixed effects with industry-by-month fixed effects $\delta_{i,t}$, which absorb all time-varying industry-level shocks and therefore also absorb the output tariff interaction. As before, all regressions are estimated both unweighted and weighted by the BLS importance weight, and standard errors are clustered at the firm level throughout.

OLS Results. Table 4 presents OLS estimates with the exposure interaction. In specification (III), $\hat{\beta}$ is 0.349 (0.106) unweighted and 0.313 (0.099) weighted. In speci-

Table 4: OLS with Exposure Interaction

	Specification (III)		Specification (IV)	
	Unwt.	Wt.	Unwt.	Wt.
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.349	0.313	0.315	0.252
	(0.106)	(0.099)	(0.102)	(0.102)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.691	0.767		
	(0.181)	(0.183)		
R^2	0.151	0.175	0.219	0.262
Obs.	299k	299k	298k	298k
Firms	3,100	3,100	3,100	3,100
Firm FE	Yes	Yes	Yes	Yes
Ind.×Month FE	No	No	Yes	Yes

Notes: Standard errors clustered at the firm level in parentheses. Observation counts truncated for disclosure avoidance. In columns (III)–(IV) the exposure measures III and IS are at the industry level. The dependent variable is the 12-month log change in producer price.

fication (IV) with industry-by-month fixed effects, the estimates are 0.315 (0.102) and 0.252 (0.102). The output tariff interaction $IS_i \times \Delta \log(1 + \tau_{i6,t}^{\text{output,eff}})$ is positive and significant in specification (III), with coefficients of 0.691 (0.181) and 0.767 (0.183), indicating that the competition channel also operates more strongly in industries with higher import penetration.

IV Results. Table 5 reports 2SLS estimates using the de jure Bartik IV. The IV estimates are substantially and systematically larger than OLS. In specification (III), $\hat{\beta}$ is 0.577 (0.143) unweighted and 0.531 (0.141) weighted—65 percent larger than the corresponding OLS estimates. With industry-by-month fixed effects in specification (IV), the IV estimates are 0.579 (0.149) and 0.492 (0.160). The stability of the input tariff coefficient across specifications (III) and (IV) is notable: adding the demanding industry-by-month fixed effects barely changes the point estimate (0.577 versus 0.579), providing strong evidence that the estimated pass-through reflects firm-level variation in tariff exposure rather than industry-level confounders.

The first-stage results (Panel B) show a coefficient of 0.669 on the Bartik IV interaction with F-statistics of 229–278 for industry-level exposure. These large F-statistics rule out weak-instrument concerns.

Table 5: IV with Exposure Interaction, De Jure Tariffs

Panel A: Second-stage results	Specification (III)		Specification (IV)	
	Unwt.	Wt.	Unwt.	Wt.
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.577	0.531	0.579	0.492
	(0.143)	(0.141)	(0.149)	(0.160)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.612	0.694		
	(0.187)	(0.199)		
Obs.	296k	296k	296k	296k
Firms	3,000	3,000	3,000	3,000
Panel B: First-stage results				
$III \times \Delta \log(1 + d^{\text{IV,dejure}})$	0.669	0.669	0.647	0.644
	(0.040)	(0.041)	(0.042)	(0.043)
R^2	0.671	0.667	0.719	0.718
F-statistic	278.2	264.3	239.9	228.7
Firm FE	Yes	Yes	Yes	Yes
Ind.×Month FE	No	No	Yes	Yes

Notes: Panel A reports 2SLS estimates using the de jure Bartik IV from equation (8). Panel B reports first-stage results. Standard errors clustered at the firm level. F-statistic is the Kleibergen-Paap weak-instrument test.

3.3 Interpreting the Two Sets of Results

The two sets of results—without and with exposure interactions—answer related but distinct economic questions.

The non-interacted specifications (I)–(II) estimate the *unconditional average* pass-through across the population of importing firms. This coefficient blends two margins: the *exposure margin*—how much of the firm's cost base is affected by the tariff—and the *pass-through margin*—how much of the cost increase is reflected in output prices. The IV estimates from Table 3 are therefore directly informative about the aggregate inflationary impact of a tariff shock: they tell us how much the typical importing manufacturer raises its output price in response to a 1 percentage point increase in its effective input tariff, inclusive of the fact that the typical firm in the sample sources only a fraction of its inputs from abroad.

The interacted specifications (III)–(IV) instead isolate the *structural* pass-through rate conditional on the degree of import exposure, separating the exposure and pass-through margins. The coefficient β on $III \times \Delta \log(1 + d^{\text{input,eff}})$ answers: for a firm whose entire material cost base is imported ($III = 1$), what is the output price response to a 1 percentage point tariff increase? The IV estimates from Table 5 are roughly three times larger than the non-interacted estimates, consistent with the fact that the average III_i in our sample is well below unity. Once the coefficient is expressed per unit of import exposure rather than as an unconditional average, it mechanically scales up.

The difference in magnitudes between the two sets of specifications reflects the average exposure at which the interacted coefficient is evaluated, and the relevant average is not the unweighted one. If prices respond to III_i times the tariff change, the non-interacted coefficient equals $\hat{\beta}^{\text{interaction}}$ times a weighted mean of III_i in which each observation enters in proportion to the square of its tariff change, so that industries contribute in proportion to the identifying variation they supply. For the de jure IV the ratio $0.203/0.577 \approx 0.35$ recovers that weighted mean. It exceeds the unweighted industry mean of 0.21 in Table 1 because the tariff increases of 2018–2019 fell disproportionately on import-intensive industries; the two are averages of the same object under different weights and are not in conflict.

This decomposition has a practical policy implication. For assessing the aggregate inflationary impact of a tariff shock across the manufacturing sector, the non-interacted estimates (0.12–0.21 per percentage point of tariff increase, across the de jure and de facto instruments) are the directly relevant numbers. For understanding the cost pass-through mechanism and predicting the inflationary impact on *specific* firms or industries with known import intensities, the interacted estimates (0.49–0.60 per unit of III) are more informative, as they can be scaled by the firm's or industry's actual III to obtain a tailored pass-through prediction.

3.4 Robustness

We subject our main findings to a battery of robustness checks along five dimensions: the price-change horizon, the sample period, and the choice of tariff instrument. Throughout, we focus on the exposure-interaction specifications (III)–(IV) with industry-level III and IS. The key finding—positive, significant, and economically meaningful pass-through from input tariffs to output prices, with IV estimates systematically exceeding OLS—is robust across all variations.

Table 6: Regression Results with Exposure Interaction (24-month)

	Specification (III)		Specification (IV)	
	Unwt.	Wt.	Unwt.	Wt.
Panel A: OLS				
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.482	0.464	0.313	0.270
	(0.107)	(0.112)	(0.103)	(0.109)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.555	0.565		
	(0.168)	(0.178)		
R^2	0.223	0.249	0.285	0.328
Obs.	246k	246k	245k	245k
Firms	2,600	2,600	2,600	2,600
Panel B: IV (de jure)				
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.624	0.619	0.389	0.326
	(0.152)	(0.172)	(0.163)	(0.182)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.507	0.513		
	(0.167)	(0.177)		
Obs.	244k	244k	243k	243k
Firms	2,600	2,600	2,600	2,600
Panel C: IV (de facto)				
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.688	0.670	0.424	0.356
	(0.147)	(0.169)	(0.160)	(0.179)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.484	0.496		
	(0.167)	(0.177)		
Obs.	244k	244k	243k	243k
Firms	2,600	2,600	2,600	2,600
Firm FE	Yes	Yes	Yes	Yes
Ind.×Month FE	No	No	Yes	Yes

Notes: See notes to Table 4. Sample period: 2017–2019.

Twenty-four-month price-change horizon. Our baseline measures the output price response at a 12-month horizon. Table 6 extends the analysis to a 24-month horizon, in which the dependent variable is the price change between month $t-11$ and month $t+12$, while the right-hand-side tariff variable remains the 12-month change $\Delta_{12} \log(1 + \tau_{j,t}^{\text{input,eff}})$.

The 24-month IV estimates are remarkably similar to their 12-month counterparts.

In the de jure IV, specification (III), the 24-month estimate is 0.624 (0.152) compared to 0.577 (0.143) at 12 months; in the de facto IV, it is 0.688 (0.147) compared to 0.600 (0.145). These estimates are statistically indistinguishable from the 12-month results. The 24-month results confirm that the vast majority of tariff pass-through is completed within the first year. The modest increase in point estimates at 24 months is consistent with a small residual adjustment beyond 12 months, but the difference is well within sampling uncertainty.

Two additional patterns deserve note. First, the IV-to-OLS gap persists at the longer horizon: in specification (III), the de jure IV estimate of 0.624 exceeds the OLS estimate of 0.482 by approximately 30%, confirming that endogenous import substitution continues to attenuate OLS at longer horizons. Second, the output tariff coefficient $\hat{\gamma}$ in specification (III) is 0.507 (0.167) under the de jure IV, comparable to the 12-month estimate of 0.612 (0.187), indicating that the competition channel also operates at similar magnitudes over the longer horizon.

Extended sample: 2017–2022. Table 7 extends the sample from 2017–2019 to 2017–2022, incorporating the partial tariff rollbacks under the Phase One trade agreement and the COVID-19 pandemic. This extension serves two purposes: it tests whether the pass-through relationship is stable beyond the initial tariff escalation period, and it nearly doubles the sample size (from approximately 296k to 548k observations).

All qualitative patterns are preserved. The IV estimates remain positive and statistically significant, and they continue to exceed OLS systematically. In the de jure IV, specification (III), the estimate is 0.556 (0.154) for 2017–2022 versus 0.577 (0.143) for 2017–2019—a modest attenuation of about 4% that is well within the confidence intervals of both estimates.

The attenuation is expected for two reasons. First, the post-2019 period saw partial tariff rollbacks and the granting of additional product exclusions, which reduced the cross-sectional variation in tariff exposure that identifies the pass-through coefficient. Second, the COVID-19 pandemic introduced large supply and demand shocks—container shipping disruptions, lockdown-induced demand shifts, commodity price swings—that added substantial noise to the pricing relationship and may have temporarily altered firms' pricing behavior. That the estimates remain significant and of similar magnitude despite this dilution speaks to the robustness of the underlying

Table 7: Regression Results with Exposure Interaction (2017–2022)

	Specification (III)		Specification (IV)	
	Unwt.	Wt.	Unwt.	Wt.
Panel A: OLS				
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.213	0.220	0.167	0.142
	(0.096)	(0.090)	(0.081)	(0.083)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.359	0.382		
	(0.153)	(0.162)		
R^2	0.209	0.217	0.322	0.357
Obs.	565k	565k	565k	565k
Firms	3,600	3,600	3,600	3,600
Panel B: IV (de jure)				
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.556	0.510	0.500	0.384
	(0.154)	(0.150)	(0.156)	(0.161)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.255	0.300		
	(0.162)	(0.171)		
Obs.	548k	548k	548k	548k
Firms	3,400	3,400	3,400	3,400
Panel C: IV (de facto)				
$III \times \Delta \log(1 + d^{\text{input,eff}})$	0.489	0.445	0.446	0.352
	(0.142)	(0.140)	(0.149)	(0.154)
$IS \times \Delta \log(1 + d^{\text{output,eff}})$	0.279	0.322		
	(0.161)	(0.170)		
Obs.	548k	548k	548k	548k
Firms	3,400	3,400	3,400	3,400
Firm FE	Yes	Yes	Yes	Yes
Ind.×Month FE	No	No	Yes	Yes

Notes: See notes to Table 4. Sample period: 2017–2022.

pass-through mechanism.

Alternative IV: de facto Tariffs. As a robustness check, we also construct an instrument using de facto (effectively collected) tariff rates, which uses the same 2017 import shares but replaces the statutory tariff rates with the import-value-weighted average rates actually collected at the border. Because the de facto rates reflect product exclusions and other administrative adjustments that were partly determined after the trade war began, this instrument is arguably less clean than the de jure version. We include it primarily to verify that our results are not sensitive to the distinction between announced and collected tariff rates—a concern given the extensive exclusion process that accompanied the Section 301 tariffs. The near-identical second-stage estimates across all specifications, reported in Tables B-1 and B-2 in Appendix B, provide reassuring evidence that our findings are not sensitive to the specific treatment of tariff exceptions.

4 A Model of Endogenous Supplier Country Choice

In this section, we develop a model with endogenous supplier country choice. We build on Halpern et al. (2015) to model the cost structure of the firm and its choice to import intermediate inputs.

4.1 Environment

Sectors, firms and technology. The home country has S sectors indexed by s, and sector s contains a finite set $\mathcal{F}_s$ of firms, $n_s \equiv |\mathcal{F}_s| \geq 2$; sums over $\mathcal{F}_s$ are read as integrals where the distinction matters. It trades with origins $c \in \mathcal{C} = \{1, \ldots, N\}$. Firm $f \in \mathcal{F}_s$ produces a differentiated good with labor L_f and a composite material input X_f according to

$$Y_f = A_f \, X_f^{\phi_s} L_f^{1-\phi_s}, \tag{9}$$

where A_f is productivity and $\phi_s \in (0, 1)$ the intermediate input share. Materials combine a domestic composite Z_f and an import composite M_f with elasticity $1+\zeta$, and the import composite combines the origins in the firm's sourcing set J_f with elasticity $1 + \eta$:

$$X_f = \left(a_D^{\frac{1}{1+\zeta}} Z_f^{\frac{\zeta}{1+\zeta}} + a_M^{\frac{1}{1+\zeta}} M_f^{\frac{\zeta}{1+\zeta}} \right)^{\frac{1+\zeta}{\zeta}}, \qquad M_f = \left(\sum_{c \in J_f} \xi_c^{\frac{1}{1+\eta}} M_{f,c}^{\frac{\eta}{1+\eta}} \right)^{\frac{1+\eta}{\eta}}. \tag{10}$$

Here ξ_c denotes the attractiveness of origin c, and $\zeta, \eta > 0$ are elasticities of substitution less one.

Imports from c cost $p_c^M = p_c^*(1 + \tau_c)$, where p_c^* is the f.o.b. price, and $\tau_c \geq 0$ is the ad valorem tariff rate. Throughout, $\widehat{\tau}_c \equiv d \ln(1 + \tau_c)$ denotes a change in the log tariff factor of origin c.

The domestic composite. The domestic composite is assembled from the goods of other domestic firms in two layers whose parameters are common to the buying sector:

$$Z_f = \prod_{s'} Z_{f,s'}^{\omega_{s,s'}}, \qquad Z_{f,s'} = \left(\sum_{f' \in \mathcal{F}_{s'}} q_{f,f'}^{\frac{\sigma_{s'}-1}{\sigma_{s'}}} \right)^{\frac{\sigma_{s'}}{\sigma_{s'}-1}}, \qquad \sum_{s'} \omega_{s,s'} = 1, \tag{11}$$

where $q_{f,f'}$ is the quantity of good produced by firm f' in sector s' that firm f (in sector s) uses as an input — the same good the household buys — and $\omega_{s,s'}$ is the share of supplying sector s' in the domestic materials of buying sector s, that is, the domestic part of the input–output use table. Firms are atomistic as buyers. Sectors are therefore linked in a roundabout structure, and $\Omega \equiv [\omega_{s,s'}]$ is row-stochastic.

Preferences, demand and market structure. Preferences are Cobb–Douglas across sectors and symmetric CES within them, defined over domestic firm goods only:

$$W = \prod_s C_s^{\beta_s}, \qquad C_s = \left[\sum_{f \in \mathcal{F}_s} \left(q_f^H \right)^{\frac{\sigma_s - 1}{\sigma_s}} \right]^{\frac{\sigma_s}{\sigma_s - 1}}, \tag{12}$$

where $\sum_s \beta_s = 1$ and $\sigma_s > 1$ is the elasticity of substitution. No imported good enters (12): imports reach the economy only as firm inputs, which is the channel the empirical analysis measures. The household supplies $\bar{L}$ units of labor inelastically at the wage w, owns the firms, and receives tariff revenue T and an exogenous transfer D; its expenditure is E. Export demand is an exogenous value $\mathcal{X}_s \geq 0$ per sector, spent through the same nest. Firms within a sector compete in quantities, taking rivals' quantities and total spending on the sector's goods as given; monopolistic competition obtains in the limit in which every revenue share is zero.

The cost of sourcing. Importing from an origin costs $\varepsilon_f C_F$ per period and per origin, where $C_F = w\bar{f}$ is paid in labor and ε_f is a draw from a distribution F with unit mean. The draw belongs to the firm and is common to its origins; it is what leads two firms with the same technology to choose different sourcing sets. Define

$$\tilde{\xi}_c \equiv \xi_c \left(p_c^M \right)^{-\eta} > 0, \qquad \Xi_f \equiv \sum_{c \in J_f} \tilde{\xi}_c, \tag{13}$$

the contribution of origin c and the firm's *sourcing capability*. We order origins by attractiveness, $\tilde{\xi}_1 > \tilde{\xi}_2 > \cdots > \tilde{\xi}_N$, and write Ξ_k for the capability of the top-k set.

Total variable cost and marginal cost. Standard calculations imply that there is an import price index given by

$$P_f^M(J_f) = \Xi_f^{-1/\eta}, \tag{14}$$

where $\Xi_f = \sum_{c \in J_f} \xi_c \left(p_c^M\right)^{-\eta}$ denotes the firm's sourcing capability. Conditional on J_f, this price index is exogenous from the firm's point of view. Next, given the CES production structure between domestic and foreign inputs, the price index for intermediate inputs is given by

$$P_f^X(J_f) = \left[a_D p_{D,s}^{-\zeta} + a_M P_f^M(J_f)^{-\zeta}\right]^{-1/\zeta}, \tag{15}$$

where $p_{D,s}$ denotes the price of domestic inputs.

Proposition 1 collects the cost structure implied by (9)–(13).

Proposition 1 (Cost). *Fix the sourcing set J_f, the domestic input price $p_{D,s}$, the wage w, and the duty-inclusive origin prices p_c^M. Then:*

(a) *The import price index, the materials price index and the total variable cost function are*

$$TVC_f(Y_f \mid J_f) = \frac{C_s^*}{B_f^{\phi_s}} \cdot \frac{Y_f}{A_f}, \tag{16}$$

where $C_s^* \equiv \left(a_D^{-1/\zeta} p_{D,s}/\phi_s\right)^{\phi_s} (w/(1-\phi_s))^{1-\phi_s}$ is the unit cost of a firm with unit productivity that sources no imports, and

$$B_f = \left(1 + \frac{a_M}{a_D} p_{D,s}^{\zeta}\, \Xi_f^{\zeta/\eta}\right)^{1/\zeta} \equiv B(\Xi_f) = (1-\alpha_f)^{-1/\zeta} \tag{17}$$

is the cost-reduction factor from importing, with B strictly increasing, $B(0) = 1$, and $B_f > 1$ if and only if $J_f \neq \varnothing$. The sourcing set enters cost only through the scalar Ξ_f.

(b) *The import share of materials and the within-import origin expenditure shares* *are*

$$\alpha_f = \frac{(a_M/a_D)\left(P_f^M/p_{D,s}\right)^{-\zeta}}{1+(a_M/a_D)\left(P_f^M/p_{D,s}\right)^{-\zeta}}, \qquad s_{f,c}^M = \frac{\tilde{\xi}_c}{\Xi_f}, \tag{18}$$

with α_f decreasing in the relative import price $P_f^M/p_{D,s}$, and the firm's imported share of variable cost, which we term its exposure, *is*

$$III_f \equiv \frac{\sum_{c\in J_f} p_c^M M_{f,c}}{TVC_f} = \phi_s\,\alpha_f. \tag{19}$$

(c) *Holding the sourcing set fixed, marginal cost responds to tariffs, the domestic input price and the wage according to*

$$d\ln MC_f = III_f \sum\nolimits_{c\in J_f} s_{f,c}^M\,\widehat{\tau}_c + (\phi_s - III_f)\,d\ln p_{D,s} + (1-\phi_s)\,d\ln w. \tag{20}$$

For a finite change with the set fixed, the exact change in the import price index lies between the two share-weighted averages,

$$\sum\nolimits_c s_{f,c}^{M,\mathrm{pre}}\widehat{\tau}_c \;\geq\; \Delta\ln P_f^M = -\frac{1}{\eta}\ln\sum\nolimits_c s_{f,c}^{M,\mathrm{pre}} e^{-\eta\widehat{\tau}_c} \;\geq\; \sum\nolimits_c s_{f,c}^{M,\mathrm{post}}\widehat{\tau}_c. \tag{21}$$

(d) *Between any two equilibria — whatever the change in tariffs, in foreign prices, or in the firm's sourcing set — with $s_f^D \equiv 1-\alpha_f$ the domestic share of materials,*

$$\Delta\ln MC_f = \phi_s\,\Delta\ln p_{D,s} + (1-\phi_s)\,\Delta\ln w + \frac{\phi_s}{\zeta}\,\Delta\ln s_f^D. \tag{22}$$

Proposition 1 has three important implications. First, part (b) identifies the regressor of the empirical specification — exposure times the import-share-weighted change in the tariff factor — with the cost shock itself. By part (c) it is not an approximation to that shock but exactly its first-order value at fixed prices and a fixed sourcing set. It is against this benchmark that the estimates are read.

Second, part (c) contains the duty–cost wedge in its simplest form. The left-hand bound in (21) is a Laspeyres index constructed on pre-shock shares and the right-hand bound a Paasche index constructed on post-shock shares, with the firm's cost lying between them. A duty index built from *current* shares is the Paasche bound and therefore understates the cost increase; the gap between them is genuine substitution.

Third, part (d) is the sufficiency result of Blaum et al. (2018): conditional on the prices of domestic factors, the entire effect of the import environment on a firm's cost — the tariff, the substitution it induces, and any change in the set of origins — is summarized by the change in a single observable share.

4.2 The firm at fixed prices

In this subsection we hold the domestic input price, the wage and the firm's output fixed, and ask what a tariff does to the cost of a firm that can change where it buys.

The sourcing ladder. The firm chooses J_f to minimize $TVC(J) + |J|\,\varepsilon_f C_F$. By Proposition 1(a), variable cost depends on J only through $\Xi(J)$ and is strictly decreasing in it, while the fixed-cost bill depends on J only through its cardinality. Among sets of a given size the top-k set is therefore optimal.

Order the origins by attractiveness, $\tilde{\xi}_1 > \tilde{\xi}_2 > \cdots > \tilde{\xi}_N$, write $\Xi_k \equiv \sum_{j \le k} \tilde{\xi}_j$ for the capability of the top-k set and $TVC_f^{(k)}$ for variable cost on it. To reduce variable costs, firm f has an incentive to import from as many supplier countries as possible. However, it also faces a fixed cost $\varepsilon_f\, C_F$ per period for each country from which it imports. The firm therefore chooses to import from exactly k suppliers if and only if the realized fixed cost draw satisfies

$$TVC_f^{(k)} - TVC_f^{(k+1)} < \varepsilon_f\, C_F < TVC_f^{(k-1)} - TVC_f^{(k)}.$$

The first inequality states that the additional cost reduction from importing from country $k+1$ is insufficient to cover the fixed cost, while the second ensures that the marginal benefit from importing from country k is large enough to justify it. For $k \in \{1, \cdots, N\}$, define the cutoffs

$$\varepsilon_f^{(k)} \equiv \frac{TVC_f^{(k-1)} - TVC_f^{(k)}}{C_F}, \tag{23}$$

which is the saving from the k-th origin in units of the fixed cost. We denote $\varepsilon^{(0)} \equiv +\infty$. Write $f(\Xi) \equiv Y_f C_s^* A_f^{-1} B(\Xi)^{-\phi_s}$ for variable cost as a function of capability, so that $\varepsilon_f^{(k)} = [f(\Xi_{k-1}) - f(\Xi_k)]/C_F$.

We impose throughout the parametric restriction

$$\eta \geq \zeta, \tag{24}$$

that is, imports across origins are at least as substitutable as home and foreign inputs. It is straightforward to show that under this condition the function $f(\Xi)$ is strictly convex at every $\Xi > 0$.

Proposition 2 (The sourcing ladder). *Given its draw ε_f, it is optimal for the firm to source from a top-k set for some k. Under the condition (24), the cutoffs are strictly decreasing, $\varepsilon_f^{(1)} > \cdots > \varepsilon_f^{(N)}$, and the optimal count is the largest k satisfying $\varepsilon_f \leq \varepsilon_f^{(k)}$:*

$$k(\varepsilon_f) = \max\{k \geq 0 : \varepsilon_f \leq \varepsilon^{(k)}\}, \tag{25}$$

Sourcing sets are then nested top-k lists ordered by ε_f.

Expected cost and the extensive margin. We now construct the expected total cost function, taking expectations over ε_f. By Proposition 2, a firm drawing ε sources from the top $k(\varepsilon)$ origins, so the probability of employing exactly k suppliers is $F\left(\varepsilon_f^{(k)}\right) - F\left(\varepsilon_f^{(k+1)}\right)$, and we can write the expected variable cost and the expected fixed cost as follows

$$\begin{aligned}
\mathbb{E}[TVC_f] &= \left(1 - F\left(\varepsilon_f^{(1)}\right)\right) TVC_f^0 + \sum\nolimits_{k=1}^{N} \left(F\left(\varepsilon_f^{(k)}\right) - F\left(\varepsilon_f^{(k+1)}\right)\right) TVC_f^{(k)}, \\
\mathbb{E}[TFC_f] &= \sum\nolimits_{k=1}^{N} \left(F\left(\varepsilon_f^{(k)}\right) - F\left(\varepsilon_f^{(k+1)}\right)\right) \mathbb{E}\left(\varepsilon_f \middle| \varepsilon_f^{(k+1)} \leq \varepsilon_f < \varepsilon_f^{(k)}\right) (kC_F).
\end{aligned}$$

We can further show that the expected total and marginal cost functions can be expressed as a combination of TVC_f^0, the total variable cost of a non-importing firm, minus the cost reduction from importing, which can be expressed as a function of the extensive-margin cutoffs $\varepsilon_f^{(k)}$. This is stated in the following proposition:

Proposition 3 (Expected Marginal Cost). *Let F be continuous on $[0, \infty)$ with $F(0) = 0$, and let $G(x) \equiv \int_0^x (x - \varepsilon_f)\, dF(\varepsilon_f) = \int_0^x F(\varepsilon_f)\, d\varepsilon_f$, so that $G' = F$. Write $MC_f^{(k)} \equiv TVC_f^{(k)}/Y_f$ for marginal cost on rung k, so that $MC_f^{(0)} = c_f^0$, and $c_f^F \equiv C_F/Y_f$ for* *the fixed cost per unit of output; adopt the conventions* $F\left(\varepsilon_f^{(0)}\right) \equiv 1$ *and* $\varepsilon_f^{(N+1)} \equiv 0$*. Then, at fixed output* Y_f*,*

$$\mathbb{E}[TC_f] = TVC_f^0 - C_F \sum\nolimits_{k=1}^{N} G\left(\varepsilon_f^{(k)}\right), \tag{26}$$

and expected marginal cost is (both routes below give the same object)

$$\mathbb{E}[MC_f] = \frac{\partial\, \mathbb{E}[TC_f]}{\partial Y_f} = c_f^0 - c_f^F \sum_{k=1}^{N} \varepsilon_f^{(k)} F\left(\varepsilon_f^{(k)}\right) = \sum_{k=0}^{N} MC_f^{(k)} \left[F\left(\varepsilon_f^{(k)}\right) - F\left(\varepsilon_f^{(k+1)}\right)\right]. \tag{27}$$

A tariff moves the cutoffs. Consider a small tariff increase on the origin c, $\widehat{\tau}_c > 0$, and assume that the increase is small enough that there is no re-ranking. Since $\tilde{\xi}_c = \xi_c (p_c^M)^{-\eta}$ and $p_c^M = p_c^*(1+\tau_c)$, we have $d\tilde{\xi}_c = -\eta\, \tilde{\xi}_c\, \widehat{\tau}_c < 0$, and hence $d\Xi_k = d\tilde{\xi}_c$ for $k \geq c$ while $d\Xi_k = 0$ for $k < c$. Then under (24):

$$d\varepsilon_f^{(k)} \begin{cases} = 0, & k < c, \\ = -\dfrac{f'(\Xi_c)}{C_F}\, d\tilde{\xi}_c < 0, & k = c, \\ = \dfrac{f'(\Xi_{k-1}) - f'(\Xi_k)}{C_F}\, d\tilde{\xi}_c > 0, & k > c, \end{cases} \tag{28}$$

the tariffed origin's own cutoff falls (marginal importers of c drop it) while every *lower-ranked* cutoff rises — tariffing China pushes marginal firms to add Vietnam.[4]

Proposition 3 and the cutoff responses (28) combine into the elasticity of expected marginal cost with respect to the tariff. We state it for a general fixed-cost distribution first, because the generality costs nothing and isolates exactly which distributional property the sign of the response depends on.

For a tariff vector $\{\widehat{\tau}_c\}$ define the direct effect on rung k and its level-weighted counterpart,

$$\Lambda_k \equiv \phi_s\, \alpha(\Xi_k) \sum\nolimits_{c \leq k} \frac{\tilde{\xi}_c}{\Xi_k} \widehat{\tau}_c, \qquad \Psi_k \equiv MC_f^{(k)} \Lambda_k, \qquad \Psi_0 \equiv 0, \tag{29}$$

[4] For $k < c$ neither argument moves. For $k = c$ only Ξ_c moves: $d\varepsilon_f^{(c)} = -f'(\Xi_c)\, d\Xi_c / C_F$, and $-f' > 0$, $d\Xi_c < 0$ give $d\varepsilon_f^{(c)} < 0$. For $k > c$ both arguments move by the same $d\tilde{\xi}_c < 0$: $d\varepsilon_f^{(k)} = [f'(\Xi_{k-1}) - f'(\Xi_k)]\, d\tilde{\xi}_c / C_F$; convexity makes $f'(\Xi_{k-1}) < f'(\Xi_k)$ (both negative, f' increasing), so the bracket is negative and the product with $d\tilde{\xi}_c < 0$ is positive.

so that Λ_k is the fixed-set shock of Proposition 1(c) evaluated at the top-k set. Write $F_k \equiv F(\varepsilon_f^{(k)})$ and $v_k \equiv \varepsilon_f^{(k)} F'(\varepsilon_f^{(k)})$, with $F_{N+1} \equiv v_{N+1} \equiv 0$.

Proposition 4 (The extensive margin)**.** *Assume the cutoffs are strictly decreasing and interior, and hold w, $p_{D,s}$, Y_f and the ranking of origins fixed. Then, for any tariff vector:*

(a) (Decomposition.)

$$d\,\mathbb{E}[MC_f] = \underbrace{\sum\nolimits_{k=1}^{N} \Psi_k\,(F_k - F_{k+1})}_{intensive} + \underbrace{\sum\nolimits_{k=1}^{N} \Psi_k\,(v_k - v_{k+1})}_{extensive}, \tag{30}$$

and the intensive term equals $Y_f^{-1} d\,\mathbb{E}[TC_f]$ exactly: the extensive-margin terms vanish from total *cost by the envelope theorem, because switching firms are indifferent at the margin.*

(b) (Uniform fixed costs.) *If F is uniform on $(0, \bar{\varepsilon})$ with every cutoff interior, then $v_k = F_k$ for every k, the extensive margin equals the intensive margin exactly, and*

$$d\,\mathbb{E}[MC_f] = 2\sum\nolimits_{k=1}^{N} \Psi_k\,(F_k - F_{k+1}) = \frac{2}{Y_f}\,d\,\mathbb{E}[TC_f], \tag{31}$$

whatever the pattern of tariffs across origins.

(c) (Universal tariff.) *If $\widehat{\tau}_c = \widehat{\tau} > 0$ for all c and $\alpha(\Xi_N) < \zeta/(\zeta + \phi_s)$, then every cutoff falls: marginal firms retreat toward domestic inputs rather than diversify. The extensive term is then non-negative for every distribution F, and strictly positive whenever F has positive density at some cutoff.*

4.3 Pricing

Propositions 1–4 are statements about cost. To convert a cost change into a price change the model requires a pricing rule, which this subsection supplies. The analysis remains partial equilibrium and concerns a single sector: total spending on the sector's goods is taken as given, and the wage and the domestic input price are held fixed as in subsection 4.2. The sector price index below is the CES aggregate over the sector's own firms, so nothing here uses the multi-sector structure of (11).

Demand and the constant-markup benchmark. Buyers allocate R_s across the sector's goods through the nest of (12), so that

$$q_f = \left(\frac{p_f}{P_s}\right)^{-\sigma_s} \frac{R_s}{P_s}, \qquad R_f \equiv p_f q_f = \varsigma_f R_s, \qquad \varsigma_f \equiv \frac{p_f^{1-\sigma_s}}{\sum_{g \in \mathcal{F}_s} p_g^{1-\sigma_s}}, \tag{32}$$

where $P_s \equiv [\sum_g p_g^{1-\sigma_s}]^{1/(1-\sigma_s)}$ and $\sum_f \varsigma_f = 1$. Holding R_s fixed is the statement that demand for the sector's composite has unit elasticity, which is what the Cobb–Douglas preferences of (12) and a fixed value of export demand deliver in subsection 4.5.

Under monopolistic competition each firm treats P_s as given and prices at the constant markup $\mu_s \equiv \sigma_s/(\sigma_s - 1)$ over marginal cost, so that $\Delta \ln p_f = \Delta \ln MC_f$ exactly. Combining this with Proposition 1(c) at fixed $p_{D,s}$ and w yields the benchmark against which the empirical estimates are read.

Proposition 5 (The fixed-sourcing benchmark)**.** *With sourcing sets fixed and constant markups,*

$$d \ln p_f = d \ln MC_f = III_f \sum\nolimits_{c \in J_f} s^M_{f,c} \, \widehat{\tau}_c, \tag{33}$$

so that a regression of price changes on exposure times the import-share-weighted change in the tariff factor has population coefficient one.

An estimate below one is therefore not a magnitude to be interpreted directly, but a shortfall to be accounted for. The model identifies the channels through which such a shortfall can arise, and this subsection supplies the first of them.

Variable markups. With finitely many firms the constant markup is an approximation. The alternative we adopt, following Atkeson and Burstein (2008), is Cournot competition: each firm chooses its output taking rivals' outputs and total spending R_s as given. A firm that expands raises the sector's quantity index and lowers its price index, and internalizes both effects, so that the inverse demand elasticity it perceives is

$$\frac{1}{\epsilon_f} \equiv -\frac{d \ln p_f}{d \ln q_f} = \varsigma_f + \frac{1 - \varsigma_f}{\sigma_s}, \tag{34}$$

the share-weighted average of unity, the inverse elasticity of demand for the composite, and $1/\sigma_s$, the inverse elasticity of substitution among the sector's goods. A firm with a negligible share faces the elasticity σ_s, and a firm constituting the entire sector faces the unit-elastic demand for the composite. We write

$$\gamma_f \equiv \frac{(\sigma_s - 1)\,\varsigma_f}{1 - \varsigma_f}, \qquad \psi_f \equiv \frac{1}{1 + \gamma_f} = \frac{1 - \varsigma_f}{1 + (\sigma_s - 2)\,\varsigma_f} \in (0, 1], \qquad \tilde{\varsigma}_g \equiv \frac{\varsigma_g \psi_g}{\sum_{h \in \mathcal{F}_s} \varsigma_h \psi_h}, \tag{35}$$

for, respectively, the elasticity of the firm's markup with respect to its relative price, the pass-through of a firm-specific cost change into that relative price, and a second set of weights over the sector's firms that tilts toward those with low shares.

Proposition 6 (Cournot pricing)**.** *Let the $n_s \geq 2$ firms of the sector compete in quantities in the nest (32), each taking rivals' outputs and R_s as given. Then:*

(a) (Markups.) *In equilibrium $p_f = \mu_f MC_f$ with*

$$\mu_f = \frac{\epsilon_f}{\epsilon_f - 1} = \frac{\sigma_s}{(\sigma_s - 1)\,(1 - \varsigma_f)}, \tag{36}$$

increasing in the firm's revenue share; the constant markup $\sigma_s/(\sigma_s - 1)$ is the limit as $\varsigma_f \to 0$.

(b) (Responses.) *For any vector of marginal-cost changes $\{d \ln MC_g\}_{g \in \mathcal{F}_s}$, the equilibrium price changes satisfy*

$$d \ln p_f - d \ln P_s = \psi_f \,(d \ln MC_f - d \ln P_s)\,, \qquad d \ln P_s = \sum\nolimits_{g \in \mathcal{F}_s} \tilde{\varsigma}_g \, d \ln MC_g. \tag{37}$$

(c) (Aggregation.) *Writing $\mathbb{E}_\varsigma$ for the revenue-weighted mean,*

$$d \ln P_s = \mathbb{E}_\varsigma \left[d \ln MC_f\right] + \frac{\mathrm{Cov}_\varsigma\,(\psi_f, d \ln MC_f)}{\mathbb{E}_\varsigma\,[\psi_f]}, \tag{38}$$

so that the sector index departs from the revenue-weighted mean cost change only through the covariance of pass-through with exposure.

Two features of (37) are central to what follows. The first is a neutrality property. The sector's price index is a weighted mean of its firms' cost changes, so that the common component of any shock — a change in $p_{D,s}$ or in the wage, or a tariff that raises every firm's import costs in the same proportion — passes through completely whatever the markups and whatever the shares. Variable markups therefore leave the response of the sector as a whole where constant markups place it, exactly so under equal shares and otherwise up to the covariance term in (38). This property is what permits the multi-sector layer of subsection 4.4 to be constructed at a constant markup and the present subsection to be added to it afterwards. It also implies that a micro pass-through estimate below unity does not license a proportional reduction of an aggregate response.

The second concerns where variable markups bind. They bind on the firm-specific component of the cost shock, which is precisely the variation that a comparison of firms within an industry isolates. Under Cournot pricing the within-sector deviation of a firm's price is ψ_f times the deviation of its cost, so that the benchmark of Proposition 5 becomes ψ_f rather than unity, and the cost jumps of firms that drop or add an origin in Proposition 4 enter firm-level comparisons scaled by ψ_f as well, whereas the sector index absorbs them in full. Under equal shares $\psi_s = (n_s - 1)/(n_s + \sigma_s - 2)$, which is 0.75 at $\sigma_s = 4$ with ten firms and 0.905 at $\sigma_s = 3$ with twenty.

4.4 Propagation

We now allow the price of domestic inputs to respond. We take the wage as numeraire, so that every price change below is relative to the wage.

Standard arguments show The minimal cost of one unit of Z_f is common to every firm in sector s,

$$p_{D,s} = \prod\nolimits_{s'} \left(\frac{P_{s'}}{\omega_{s,s'}} \right)^{\omega_{s,s'}}, \tag{39}$$

with conditional demands $Z_{f,s'} = \omega_{s,s'}(p_{D,s}/P_{s'})Z_f$ and $q_{f,f'} = (p_{f'}/P_{s'})^{-\sigma_{s'}} Z_{f,s'}$. In particular, between any two equilibria, $\Delta \ln p_{D,s} = \sum_{s'} \omega_{s,s'} \Delta \ln P_{s'}$ exactly.

As a result, the domestic input price now enters the cost-reduction factor (17): a rise in $p_{D,s}$ makes imports relatively cheaper, raises B_f, and moves every cutoff.

The sector price index.

Proposition 7 (Sector prices)**.** *Let the set $\mathcal{F}_s$ be fixed and markups constant. Between two equilibria with domestic shares $\{s_f^D\}$ and $\{s_f^{D\prime}\}$,*

$$\Delta \ln P_s = \phi_s \, \Delta \ln p_{D,s} + (1 - \phi_s) \, \Delta \ln w + \Lambda_s, \tag{40}$$

where, exactly,

$$\Lambda_s \equiv \frac{1}{1-\sigma_s} \ln\left[\sum\nolimits_{f\in\mathcal{F}_s} \varsigma_f \left(\frac{s_f^{D\prime}}{s_f^D}\right)^{-\chi_s}\right], \qquad \chi_s \equiv \frac{\phi_s\left(\sigma_s - 1\right)}{\zeta}, \tag{41}$$

with ς_f the pre-change revenue shares. To first order in the share changes, $\Lambda_s = (\phi_s/\zeta)\sum_f \varsigma_f \Delta \ln s_f^D$; and to first order in a tariff change, holding $p_{D,s}$, w and the sourcing sets fixed,

$$\Lambda_s = \Lambda_s^\tau \equiv \sum\nolimits_{f\in\mathcal{F}_s} \varsigma_f\, III_f \sum\nolimits_{c\in J_f} s_{f,c}^M\, \widehat{\tau}_c, \tag{42}$$

the revenue-weighted average over the sector of the firm-level regressor of Proposition 5.

The inequality reflects demand substituting toward the firms whose costs rose least, so that the index rises by less than the mean cost. A calculation that assigns every firm the sector's average share — the aggregate-data approach — is biased in a direction determined by χ_s: it understates the effect of input trade when $\chi_s < 1$ and overstates it when $\chi_s > 1$ (Blaum et al., 2018).

The input–output fixed point.

Proposition 8 (Propagation)**.** *Write $\Gamma \equiv \operatorname{diag}(\phi_s)$ and with the wage as numeraire:*

(a) (Measurement mode.) *Given the realised change in every firm's domestic share, the vector of sector price changes solves, exactly,*

$$\Delta \ln P_s = \Lambda_s + \phi_s \sum\nolimits_{s'} \omega_{s,s'} \Delta \ln P_{s'}, \qquad \text{that is,} \qquad \Delta \ln \mathbf{P} = \left(I - \Gamma\Omega\right)^{-1} \Lambda. \tag{43}$$

With a single sector, $\Delta \ln P = \Lambda/(1-\phi)$.

(b) (Counterfactual mode.) *To first order in the tariff vector, with sourcing sets fixed,*

$$\Delta \ln \mathbf{P} = \mathcal{M}\Lambda^\tau, \qquad \mathcal{M} \equiv \left[I - \Gamma \operatorname{diag}(1-\bar{\alpha}_s)\,\Omega\right]^{-1}, \tag{44}$$

where $\bar{\alpha}_s \equiv \sum_{f\in\mathcal{F}_s} \varsigma_f \alpha_f$ and $\rho \le \max_s \phi_s(1-\bar{\alpha}_s)$.

(c) (Aggregate producer prices.) *With revenue weights ϑ_s,*

$$\Delta \ln P^{PPI} = \vartheta' \mathcal{M} \Lambda^\tau = \sum\nolimits_s \vartheta_s \, m_s \, \Lambda_s^\tau, \qquad m_s \equiv \frac{[\mathcal{M}\Lambda^\tau]_s}{\Lambda_s^\tau} \geq 1, \tag{45}$$

the revenue-weighted regressor times a sector-specific propagation multiplier, with $m_s = 1$ when $\Omega = 0$ and $m_s = 1/(1-\phi(1-\bar{\alpha}))$ with a single sector.

Parts (a) and (b) are the same system written at two information sets, and combining them double-counts. In measurement mode the realised share $s_f^{D\prime}$ already embeds the firm's substitution toward imports when $p_{D,s}$ rises, so that the feedback matrix carries the full ϕ_s; in counterfactual mode that substitution is computed by the model, which moves the factor $1 - \bar{\alpha}_s$ out of Λ and into the matrix. Two limitations should be recorded. Prices are relative to the wage, and with sector-specific labor the wage would rise where demand rises and the price response would be larger. And the Neumann series counts rounds of input use rather than calendar time, so it carries no implication for the speed of adjustment.

4.5 Closure

The preceding subsections take total spending on a sector's goods as given. We now determine it and state the equilibrium. The closure is used in the counterfactual analysis; the identification results above do not depend on it beyond the unit elasticity of sector demand.

Given expenditure E, the household spends $\beta_s E$ on sector s, with demands and price indices

$$q_f^H = \left(\frac{p_f}{P_s}\right)^{-\sigma_s} \frac{\beta_s E}{P_s}, \qquad P_s = \left[\sum\nolimits_{f \in \mathcal{F}_s} p_f^{1-\sigma_s}\right]^{\frac{1}{1-\sigma_s}}; \tag{46}$$

the ideal consumer price index is $P = \prod_s (P_s/\beta_s)^{\beta_s}$ and indirect utility is $W = E/P$. Every buyer — the household, every intermediate buyer, and the foreign buyer — demands firm f's good through the same nest, so that total revenue is $R_f = \varsigma_f R_s$ with

$$R_s = \beta_s E + \sum\nolimits_{s'} \omega_{s',s} \sum\nolimits_{f' \in \mathcal{F}_{s'}} p_{D,s'} Z_{f'} + \mathcal{X}_s. \tag{47}$$

Proposition 9 (Revenue and expenditure)**.** *Let markups be constant and let*

$$\nu_s \equiv \frac{\sigma_s - 1}{\sigma_s}\,\phi_s\,(1 - \bar{\alpha}_s) \tag{48}$$

denote the share of the sector's revenue that returns to domestic firms as input purchases. Then:

(i) *Sector revenues solve* $R = (I - K)^{-1}(\beta E + \mathcal{X})$ *with* $K_{s,s'} \equiv \omega_{s',s}\nu_{s'}$ *and* $\rho(K) \leq \max_s \nu_s < 1$*, so that the inverse exists and its Neumann series converges.*

(ii) *With* $\varpi_s \equiv 1/\sigma_s$ *the profit share of revenue,* $\Phi = C_F \sum_f |J_f|\,\varepsilon_f$ *the fixed-cost bill and* T *tariff revenue, household expenditure is*

$$E = \frac{w\bar{L} - \Phi + T + D + \varpi'(I - K)^{-1}\mathcal{X}}{1 - \varpi'(I - K)^{-1}\beta}, \tag{49}$$

and the denominator lies strictly in $(0, 1)$*.*

Equilibrium. Given tariffs, foreign prices, export values and the draws $\{\varepsilon_f\}$, an equilibrium is a collection of sourcing sets $\{J_f\}$, prices $\{p_f\}$, indices $\{P_s, p_{D,s}\}$, a wage w, expenditure E and a transfer D such that firms minimize cost and choose J_f as in Propositions 1–2 given $(p_{D,s}, w)$; prices satisfy Proposition 6; demands, indices and revenues satisfy (39) and (46)-(47); E satisfies (49); the transfer equals the trade deficit; and the labor market clears. The wage is the numeraire. Summing the household budget, the firms' revenue identities and goods-market clearing shows that with D the deficit the labor-market condition is implied by the others; were D held fixed instead, that identity would determine the wage.

5 Model Implications and Quantification

This section takes the model of Section 4 to the data in five steps. Subsection 5.1 states what a firm-level and a sector-level regression identify in the economy of Section 4. Subsection 5.2 lists the inputs, all of which are measured on the input–output table or taken from independent estimates. Subsection 5.3 reads the paper's own firm-level estimates through the model: those coefficients are inputs to the exercise, and nothing in this section is estimated. Subsection 5.4 turns to the sector-level evidence of Amiti et al. (2019), which enters as a *prediction*: no sector-level coefficient is used in subsections 5.2–5.3, so the comparison is held out. Subsection 5.5 reports aggregates and counterfactuals. The order is a dependency rather than a preference: the sector-level prediction uses the propagation objects that subsection 5.2 measures, and the aggregates use the damping factor that subsection 5.3 establishes.

5.1 What the regressions identify

Proposition 5 fixed the benchmark for a firm-level regression at unity: with sourcing sets fixed and constant markups the regressor is the cost shock itself. The proposition below says how a firm-level and a sector-level design differ in what they identify once the price of domestic inputs responds. It is a statement about estimators rather than about the economy, which is why it is placed here rather than with the model.

Firm-level and sector-level regressions. Input-tariff pass-through into producer prices has been estimated at two levels of aggregation: at the level of the six-digit industry by Amiti et al. (2019), and at the level of the firm in the empirical analysis of this paper. The model determines what each regression identifies. Write $\widehat{\tau}_f \equiv \sum_{c\in J_f} s_{f,c}^{M,\mathrm{pre}}\widehat{\tau}_c$ for the import-weighted tariff change on firm f's imports and $z_f \equiv III_f\widehat{\tau}_f$ for the firm-level regressor, the shares being those of the initial equilibrium because the first-order response of a cost function to a price is the initial cost share. For sector s let $III_s \equiv \sum_f \varsigma_f III_f$ and let $\bar{\tau}_s$ be the import-weighted average of $\widehat{\tau}_f$ over the sector's firms. Let $\tilde{\omega}_{s,s'}$ be the share of supplying sector s' in sector s's total material inputs, domestic and imported, and let $\bar{\tau}_{s'}^{out}$ be the import-weighted tariff change on all imports of s'-goods; the sector-level regressor is then

$$X_s \equiv III_s \sum_{s'} \tilde{\omega}_{s,s'}\bar{\tau}_{s'}^{out}. \tag{50}$$

We say that *proportionality* holds if each sector's imports are spread across supplying sectors as its total inputs are and face the economy-wide tariff change on those goods, so that $\sum_{s'} \tilde{\omega}_{s,s'}\bar{\tau}_{s'}^{out} = \bar{\tau}_s$. This is the assumption under which the use table's import matrix is constructed.

Proposition 10 (Firm-level and sector-level regressions)**.** *Under the assumptions of Proposition 8(b):*

(a) (Firm prices.)

$$\Delta \ln p_f = III_f \left(\widehat{\tau}_f - \Delta \ln p_{D,s} \right) + \phi_s \Delta \ln p_{D,s}, \tag{51}$$

exposure times the change in the price of imported relative *to domestic inputs, plus a sectoral term.*

(b) (Firm-level regression.) *The pooled within-sector regression of $\Delta \ln p_f$ on z_f, which is the population counterpart of a specification with industry-by-month fixed effects, has coefficient*

$$\beta^{\text{firm}} = 1 - \frac{\sum_s \Delta \ln p_{D,s} \, C_s}{\sum_s V_s}, \quad C_s \equiv \sum\nolimits_f \left(z_f - \bar{z}_s \right) \left(III_f - \overline{III}_s \right), \quad V_s \equiv \sum\nolimits_f \left(z_f - \bar{z}_s \right)^2 . \tag{52}$$

If the tariff on imports is common to the firms of each sector, $\widehat{\tau}_f = \widehat{\tau}_s$, this becomes

$$\beta^{\text{firm}} = 1 - \sum\nolimits_s w_s \frac{\Delta \ln p_{D,s}}{\widehat{\tau}_s}, \qquad w_s \equiv \frac{\widehat{\tau}_s^2 D_s}{\sum_{s'} \widehat{\tau}_{s'}^2 D_{s'}}, \qquad D_s \equiv \sum\nolimits_f \left(III_f - \overline{III}_s \right)^2 , \tag{53}$$

which is strictly below one whenever some sector with dispersed exposure buys from a sector with positive direct exposure. Under Cournot pricing with equal within-sector shares each sector's term is multiplied by ψ_s, so that $\beta^{\text{firm}} = \sum_s w_s \, \psi_s \left(1 - \Delta \ln p_{D,s}/\widehat{\tau}_s\right)$.

(c) (Sector-level regression.) *We have $\Lambda_s^\tau = III_s \bar{\tau}_s$ exactly, and under proportionality $X_s = \Lambda_s^\tau$. The cross-sectional regression of $\Delta \ln P_s$ on X_s with an intercept then has coefficient*

$$\beta^{\text{sec}} = 1 + \frac{\text{Cov}\left(\left[(\mathcal{M} - I) \Lambda^\tau \right]_s , \Lambda_s^\tau \right)}{\text{Var}\left(\Lambda_s^\tau \right)}, \tag{54}$$

the direct effect plus the projection of the propagated price change on direct exposure.

(d) (The gap.) *If propagation is non-negatively correlated with exposure across sectors, then $\beta^{\text{firm}} < 1 \leq \beta^{\text{sec}}$; and with a common multiplier m and a common discount $\rho \equiv \Delta \ln p_{D,s}/\widehat{\tau}_s$, $\beta^{\text{sec}}/\beta^{\text{firm}} = m/(1-\rho)$, or $m/[\psi(1-\rho)]$ under Cournot pricing.*

Three implications follow. First, the two regressions estimate different objects, and outside one knife-edge case neither delivers the multiplier that Proposition 8(c) requires. Industry-by-month fixed effects absorb the sectoral term in (51) but not the relative-price correction $-III_f \Delta \ln p_{D,s}$, which is collinear with the regressor, so that the firm-level coefficient is direct pass-through discounted by the propagated rise in domestic input prices. The sector-level coefficient is subtler, because a cross-sectional regression carries an intercept and that intercept absorbs the average propagated price change. Rearranging (54),

$$\beta^{\text{sec}} - 1 = \sum\nolimits_s w_s^{\text{sec}} \left(m_s - 1\right), \qquad w_s^{\text{sec}} \equiv \frac{\Lambda_s^\tau \left(\Lambda_s^\tau - \bar{\Lambda}^\tau\right)}{\sum_{s'} \left(\Lambda_{s'}^\tau - \bar{\Lambda}^\tau\right)^2},$$

weights that sum to one but are negative for sectors below mean exposure, whereas (45) requires the revenue-weighted mean, whose weights are non-negative. The coefficient β^{sec} is therefore not an average of the m_s and can fall below one.

Second, β^{sec} is a diagnostic for the network rather than a measurement of propagation. By (54), $\beta^{\text{sec}} - 1$ is the slope of the propagated price change on direct exposure across sectors, and is thus a statement about the input–output network's assortativity in exposure — whether exposed sectors buy from exposed sectors — and about nothing else in the closure. Above one, exposed sectors source from exposed suppliers; equal to one, propagation is unrelated to own exposure and the intercept absorbs all of it; below one, exposed sectors source from sheltered ones. Because the slope is a property of the network and of the pattern of exposure across sectors, it can be computed from the use table without reference to any price, and set against a regression estimate that computation constitutes a test of the closure.

Third, part (d) orders the two designs, and the ordering has an economic reading. A within-industry comparison identifies from shocks that are *relative*: one firm's imported inputs become dearer than its rivals'. A relative cost shock is passed through at the rate ψ_f, because a firm that raises its price relative to its rivals loses market share, and it is discounted by $1 - \rho_s$, because the rivals' costs also rise through the price of domestic inputs. A sector-level regression identifies from the component of the shock that is *common* to a sector's firms, which by Proposition 6(b) passes into the sector index one for one whatever the markups, and which carries propagation on top. Hence $\beta^{\text{firm}} < 1 \leq \beta^{\text{sec}}$, and the two differ by the factor $m/[\psi(1-\rho)]$ at common values. Neither coefficient mismeasures the other. The firm-level design isolates the cost channel because two firms in the same industry face the same sectoral conditions and different import bills; the sector-level design measures the total response of an industry's prices to tariffs on its suppliers' goods, through whichever channel operates.

5.2 Inputs

The quantification uses the 2007 and 2017 vintages of the BEA use table as processed by Amiti et al. (2019), with wage bills from the Economic Census, and the schedule of tariffs on China in place in December 2018. The finer 2007 table, at 231 combined industries, serves the sector-level comparison; the 2017 table, at nineteen summary manufacturing sectors, serves the multiplier. China's share of 2017 imports is 0.24 overall and 0.25 in manufacturing; weighted by imported-input exposure it is 0.16, because the most import-intensive sectors — vehicles, aircraft, pharmaceuticals, refining — source relatively little from China.

Three objects are computed on the table under the model and enter the comparisons below; none uses a price or a regression. The discount $\rho_s \equiv \Delta \ln p_{D,s}/\widehat{\tau}_s$, the propagated rise in the price of a sector's domestic inputs per unit of the tariff on its imports, is computed on the 2007 network with the 2017 China shares; across manufacturing industries its median is 0.095, with an interquartile range of 0.06 to 0.15. The multipliers m_s of Proposition 8(c) depend on how much of the economy is allowed to respond: 1.41 when only the nineteen summary manufacturing sectors of the 2017 table are endogenous, 1.6 to 1.7 for the manufacturing block of the finer 2007 table, and 1.6 to 1.9 when the prices of mining, agriculture and services move as well — 1.64 for the China tariffs, 1.81 for a universal tariff, and 1.88 for the schedule in place in December 2018. Across the 190 manufacturing industries with imported-input intensity above five percent the industry multipliers have a median of 1.75 and an interquartile range of 1.57 to 1.97. The assortativity slope of (54), the projection of the propagated price change on direct exposure across sectors, is 0.329 on the 2007 table with the December-2018 schedule and $\phi_s(1-\bar{\alpha}_s)$ measured industry by industry.

The markup damping ψ_s is the one input not read off the table. We use the equal-shares result $\psi_s = (n_s - 1)/(n_s + \sigma_s - 2)$ of Proposition 6, so that it is determined by the number of competitors and the within-sector elasticity of substitution. With σ_s between three and four and revenue shares between five and ten percent — the range the Census of Manufactures gives for a typical six-digit industry — ψ_s lies between 0.75 and 0.905. The producer-price sample tilts toward large firms, whose shares are higher and whose ψ_f is accordingly lower, so the upper end of that range is the less likely one. Table 8 collects these values together with the within-industry estimates that subsection 5.3 reads through the model.

Table 8: Inputs to the quantification

Input	Value	Basis
Use table	2007 (231 industries); 2017 (19 manufacturing sectors)	BEA, as processed by Amiti et al. (2019); wage bills from the Economic Census
Tariff schedules	China lists in place in December 2018; universal 15% on all imported inputs	Statutory rates
China share of imports	0.24 all; 0.25 manufacturing; 0.16 weighted by imported-input exposure	2017 imports
Discount ρ_s	median 0.095; interquartile range 0.06–0.15	2007 network, 2017 China shares
Markup damping ψ_s	0.75–0.905	Proposition 6; $\sigma_s \in [3, 4]$, revenue shares 5–10%
Multiplier m	1.41; 1.6–1.7; 1.64–1.88	2017 manufacturing; 2007 manufacturing block; all sectors endogenous
Assortativity slope	0.329	2007 table, December-2018 schedule
Within-industry estimates	instrumented 0.579; first stage 0.647; reduced form 0.375	Empirical analysis, industry-by-month fixed effects, industry intensity

5.3 The within-industry estimates

The empirical analysis regresses twelve-month log changes in firm producer prices on the effective input tariff interacted with the industry's imported-input intensity III_i, with firm and industry-by-month fixed effects, instrumenting the effective tariff with the statutory Bartik measure built on pre-shock import shares. The design is firm-level in its unit of observation and in its fixed effects; the exposure weight is the industry's. In the specification with industry-by-month effects the instrumented coefficient is 0.579, the first stage is 0.647, and the reduced form — the coefficient of the price change on the statutory Bartik regressor, which in a just-identified design is the product of the two — is 0.375.

Two results of Section 4 fix what such a regression should deliver, and we use only those two. Proposition 5 sets the starting point at unity: with sourcing sets fixed and constant markups the Bartik regressor is the cost shock itself. Proposition 10(b) discounts it by $1 - \rho_s$, because industry-by-month effects remove the sectoral term in (51) but not the relative-price correction that is collinear with the regressor; at the median discount of 0.095 the factor is 0.905. Proposition 6(b) multiplies it by

ψ_s, because the fixed effects leave only the firm-specific part of a cost shock and an oligopolist absorbs part of that in its markup. The benchmark for the within-industry coefficient is therefore $\psi_s(1-\rho_s)$: 0.68 at $\psi_s = 0.75$ and 0.82 at $\psi_s = 0.905$, the lower value being the more likely for a sample that tilts toward large firms.

The object the benchmark applies to is the reduced form. The coefficient of Proposition 10(b) is the population slope on z_f, a regressor built from pre-shock shares and the change in the tariff factor; that is the Bartik instrument of the empirical analysis, not the effective tariff, whose current shares embed the substitution of Proposition 1(c). Against a reduced form of 0.375 the benchmark leaves a residual

$$\kappa^{\mathrm{res}} \equiv \frac{\widehat{\beta}^{RF}}{\psi_s\,(1-\rho_s)}, \tag{55}$$

which is 0.55 at $\psi_s = 0.75$ and 0.46 at $\psi_s = 0.905$: the part of the shortfall from complete pass-through that the model does not itself produce. Two qualifications bound its interpretation. First, the exposure measure in the regressor is the industry's imported-input intensity rather than the firm's, whereas the regressor of Proposition 10 carries the firm's own. The substitution does not thereby attenuate the coefficient: were a firm's intensity its industry's plus an independent firm-specific deviation, that deviation would be of the Berkson kind and would leave the population slope unchanged. What it changes is the object the coefficient is measured per unit of, the industry's imported-input intensity being a larger and differently constructed quantity than the firm's own import bill. Whatever part of κ^{res} this accounts for is a property of the regressor rather than of pass-through, and does not transfer to the aggregate. Second, the benchmark holds sourcing sets fixed. Proposition 4(b) says that in the population of importers, with uniform fixed costs, the extensive margin doubles the response. The sample contains no firm that ceased importing, since the effective tariff is constructed from current imports, so the largest cost jumps of Proposition 2(d) are absent from it. It does, however, contain firms that dropped a taxed origin while retaining others, and their smaller jumps are in the price data. The fixed-set benchmark therefore understates what applies to this sample and the doubled benchmark overstates it; the two bracket the residual, which is 0.46 to 0.55 against the fixed-set benchmark and 0.23 to 0.28 against the doubled one.

5.4 The sector-level evidence

Proposition 10(c) says what a sector-level regression identifies, and Amiti et al. (2019) have run one. In their Table 4, twelve-month changes in six-digit producer price indices are regressed on the industry's imported-input intensity times an input tariff built from the BEA use table — the input-value-weighted average, over supplying industries, of the import-weighted tariff change on each supplying industry's goods — with industry and time fixed effects, and the coefficient is 1.87. Nothing in subsections 5.2–5.3 uses that coefficient, so the comparison is held out.

The starting point is again unity: a regression of sector price changes on direct exposure would have coefficient one if there were no network. With one, a sector's price also moves with its suppliers' prices, and because that propagated movement is correlated with the sector's own exposure it loads on the coefficient. By Proposition 10(c) the loading is exactly the assortativity slope,

$$\beta^{\text{sec}} = 1 + A, \qquad A \equiv \frac{\operatorname{Cov}\left([(\mathcal{M} - I)\Lambda^{\tau}]_s, \Lambda^{\tau}_s\right)}{\operatorname{Var}\left(\Lambda^{\tau}_s\right)},$$

which subsection 5.2 puts at 0.329. The prediction is therefore 1.33, and Cournot pricing does not move it: by Proposition 6(b) markups damp only within-sector deviations, while both sides of a sector-level regression are sector-level objects. Nor is it the multiplier: propagation that were unrelated to own exposure would leave the coefficient at one however large the multiplier, and the value 0.329 reflects that on the actual table exposed sectors buy from exposed sectors.

Set against 1.87, the network accounts for 0.33 of the 0.87 by which the estimate exceeds one. The remainder measures what the sector-level regressor loads on that the closure sets to zero. The leading candidate is the competition channel one step upstream: a tariff on s'-goods lets domestic producers of s'-goods raise their prices — Amiti et al. (2019) estimate 0.40 per unit of import share for a sector's own output — and the resulting rise in the cost of domestic inputs from s' reaches sector s through the same use-table weights that build (50). The two are the same weighted sum of the same tariff changes, scaled by import intensity in one case and by import share in the other, and cannot be told apart at the sector level. In the closure imported goods do not compete with domestic goods, so the channel is absent by construction, and the firm-level design removes it on purpose: two firms in the same industry face the same upstream protection and different import bills. The comparison is like for like, since the ARW regressor holds import weights at their pre-shock values and moves statutory rates and is therefore a reduced form; a sector's imports, unlike a marginal firm's, do not fall to zero when the tariff arrives, so the sector-level first stage is close to one and the instrumented and reduced-form objects nearly coincide there. Differences in window — the ARW sample ends in December 2018, five months into the China lists, and covers all origins — and in sample composition are second to the upstream channel.

Proposition 10(d) then predicts a ratio of the sector-level to the firm-level coefficient of $1.33/[\psi_s(1-\rho_s)]$, between 1.6 and 2.0; the data give 3.2 against the instrumented estimate and 5.0 against the reduced form. Propagation and markups account for a factor of about two out of three to five, and the upstream competition channel, with the differences in window and sample, accounts for the rest. The sector-level coefficient is accordingly a total response through whichever channel operates, and not an estimate of imported-input cost pass-through.

5.5 Aggregates and counterfactuals

Proposition 8(c) turns the direct exposure of each industry into an aggregate producer-price response through the network: $\Delta \ln P^{PPI} = \vartheta' \mathcal{M} \Lambda^{\tau}$, the revenue-weighted direct exposure times the sector multipliers. The inputs are the direct exposures Λ_s^{τ} built from the tariff schedule and the China shares, the matrix $\mathcal{M}$ built from the use table with $\phi_s(1-\bar{\alpha}_s)$ measured industry by industry, and revenue weights from the Economic Census. This subsection reports the response for the 2018–2019 tariffs on China and for a universal tariff at complete pass-through, then applies the damping factor of subsection 5.3, and finally records what the extensive margin adds.

At complete pass-through the tariffs on China raised manufacturing producer prices by 0.44 percentage points inclusive of propagation, and a universal fifteen-percent tariff on imported inputs would raise them by 3.0. The multiplier behind these figures lies between 1.4 and 1.9 depending on how much of the economy is allowed to respond, as subsection 5.2 records, and it differs between the two shocks, 1.64 for the China tariffs and 1.81 for the universal tariff. This is where the amplification of a universal tariff beyond its wider coverage lives: not in the sourcing margin but in the network, because a universal tariff raises every supplier's price at once.

These figures hold at complete pass-through of every realized cost change. The firm-level estimate does not enter them, and it should not enter as $\widehat{\beta}$: the discount and the markup damping that lower the within-industry coefficient are not aggregate attenuators, because the common component of a sector's cost shock passes into its index one for one by Proposition 6(b), and propagation is already carried by $\mathcal{M}$. Scaling the aggregate by $\widehat{\beta}$ would count the discount and the markups twice. What can enter is the residual κ^{res} of (55), and whether it does depends on the force behind it. If the residual is measurement in the regressor, it attenuates the slope and leaves the aggregate untouched. If it is nominal rigidity within the twelve-month window, or adjustment that continues past it, then a fraction κ of every realized cost change reaches prices within the window, costs received through domestic suppliers included; the fixed point becomes $\kappa[I - \kappa\,\Gamma\,\text{diag}(1 - \bar{\alpha}_s)\,\Omega]^{-1}\Lambda^{\tau}$, every round of propagation is damped along with the direct term, and the one-sector multiplier falls from $1/(1-a)$ to $1/(1-\kappa a)$ with $a \equiv \phi(1-\bar{\alpha})$. Two numbers should therefore be reported. The model-consistent response, 0.44 and 3.0 percentage points, is the response if the residual is measurement. That response times κ^{res}, between 0.20 and 0.24 percentage points for the China tariffs and between 1.4 and 1.7 for the universal tariff, is an upper bound on the within-window response if the residual is nominal rigidity, since it damps the direct term but not the rounds of propagation. Which force is behind the residual is a question for the price microdata, which can separate them, and not for the model, which cannot.

Finally, the figures hold sourcing sets fixed. Proposition 8(b) replaces Λ_s^{τ} by $\kappa_s^{\text{ext}}\Lambda_s^{\tau}$ when sets respond, and under uniform fixed costs $\kappa^{\text{ext}} = 2$ whatever the pattern of tariffs across origins, by Proposition 4(b): in the population of importers, including the firms that leave the taxed origin and forfeit the premium of Proposition 2(d), the responses would be twice those reported, 0.9 and 6.0 percentage points at complete pass-through. The firm-level sample disciplines this factor only in part: it contains firms that dropped a taxed origin, whose cost jumps are in the price data, but not firms that ceased importing, whose jumps are the largest. The doubling is in any case exact only for the uniform distribution. Proposition 4(d) gives the condition on F under which the response is positive, and the size of the extensive term is set by the mass of firms at the taxed origin's cutoff net of the mass at the lower-ranked cutoffs they move to. Both masses are observable in the import data as exit and entry rates by exposure, which is where the factor is to be measured.

Read together, the section establishes three numbers and one qualification. The within-industry coefficient the model predicts with sourcing sets fixed is $\psi_s(1-\rho_s)$, between 0.68 and 0.82, against a reduced form of 0.375; the residual of 0.46 to 0.55 is what the model does not produce, and its interpretation is for the price data to settle. The sector-level coefficient the model predicts is 1.33, against 1.87 in Amiti et al. (2019); the network accounts for the 0.33, and the remainder is the upstream competition channel that the firm-level design removes by construction. The aggregate response of manufacturing producer prices is 0.44 percentage points for the tariffs on China and 3.0 for a universal fifteen-percent tariff at complete pass-through; the residual, if it is nominal rigidity, cuts these by roughly half within the window, and the extensive margin, if fixed costs are uniform, doubles them in the population of importers. The qualification is that the two factors the firm-level sample cannot discipline — the force behind the residual and the size of the extensive margin — are the two that the price and import microdata can.

6 Conclusion

This paper has traced the tariffs of the 2018–2019 trade war one link further along the domestic price chain than the existing literature, from the prices of imported inputs to the prices that U.S. manufacturers charge for their output. Linking the universe of customs transactions in the LFTTD to BLS producer price microdata, we construct firm-level effective input tariffs from each firm's own import portfolio, instrument them with Bartik measures that hold the portfolio at its pre-trade-war composition, and estimate that a one percentage point increase in a firm's effective input tariff raises its output price within a year by 0.5 to 0.6 percentage points per unit of imported-input intensity, or by 0.12 to 0.20 percentage points for the average importing manufacturer.

Four findings deserve emphasis. First, the first stage is well below one: a firm's effective input tariff rises by about two-thirds of the statutory shock on its pre-shock portfolio. That is a measured property of the index rather than an inference, and it says firms reallocate away from taxed origins. The instrumented estimates also exceed least squares by half or more. This is consistent with the same reallocation but does not follow from it — understatement that is uniform across firms raises both coefficients without opening a gap between them, and a gap requires understatement that varies across firms in a way correlated with the shock — so we read the gap as suggestive rather than as identifying the mechanism. Section 4 adds that a firm which stops importing altogether forfeits the entire cost advantage of imported inputs, a jump the sample cannot observe because such a firm has no effective tariff, while a firm that merely drops a taxed origin forfeits less and remains in the sample.

Second, a within-industry coefficient below one is not by itself evidence that firms absorb tariffs in their margins. Industry-by-month fixed effects isolate the firm-specific component of the cost shock, and the model predicts that this component passes through at less than one for one for two reasons that have nothing to do with rigidity: an oligopolist that raises its price relative to its rivals loses market share, and its rivals' costs also rise through the price of domestic inputs. The model's benchmark for the within-industry coefficient is 0.68 to 0.82, against a reduced form of 0.375. The residual is what the model does not produce, and its source—nominal rigidity within the year, or measurement in an exposure variable that is observed at the industry level—is a question the price microdata can settle and the model cannot.

Third, a firm-level and a sector-level regression identify different objects, and neither mismeasures the other. The firm-level design identifies from relative shocks and isolates the cost channel, because two firms in the same industry face the same upstream protection and different import bills. The sector-level design of Amiti, Redding and Weinstein (2019) identifies from the component common to a sector's firms, which passes through in full and carries propagation on top; the model predicts a coefficient of 1.33 from the input–output table alone, against their 1.87, and attributes the difference to the protection of a sector's suppliers, which a sector-level regressor cannot separate from input costs.

Fourth, at complete pass-through the China tariffs raised manufacturing producer prices by 0.44 percentage points inclusive of propagation through domestic input prices, or by 0.20 to 0.24 if the residual reflects nominal rigidity within the year; a universal fifteen-percent tariff on imported inputs would raise them by 3.0 percentage points, or by 1.4 to 1.7 with the same damping. The universal tariff is roughly seven times as inflationary because it reaches all imported inputs rather than the 16 percent of manufacturers' imports, weighted by exposure, that came from China, and because it raises every domestic supplier's price at once, so that propagation through the network is stronger. These figures exclude the competition channel, the wage response that a balanced-trade closure would add, and the further step from producer to consumer prices.

The policy content is twofold. Tariffs on intermediate inputs do not stop at the border: a large part of their cost reaches the prices of domestic producers within a year, and the input–output network multiplies the direct effect by a factor between 1.4 and 1.9. And the inflationary cost of protection rises with its breadth by more than the coverage it adds, because a tariff on every origin leaves firms no untaxed origin to switch to and raises every supplier's price at once; the cost per dollar of imported inputs covered is higher for a universal tariff than for a targeted one.

Two objects the firm-level sample cannot discipline set the agenda for future work, and both are measurable in the microdata we use. The force behind the residual can be separated in the price data, through the frequency and size of the price changes that follow a tariff shock. The size of the extensive margin can be measured in the import data, through the rates at which firms drop a taxed origin and add untaxed ones as a function of their exposure. Beyond these, extending the closure to a wage response under balanced trade and to the competition channel, and tracing the chain one step further from producer to consumer prices, would complete the accounting of the aggregate price effects of tariffs.

References

Amiti, Mary and Sebastian Heise, "U.S. Market Concentration and Import Competition," *Review of Economic Studies*, 2025, *92* (2), 737–771.

_ **, Oleg Itskhoki, and Jozef Konings**, "Importers, Exporters, and Exchange Rate Disconnect," *American Economic Review*, 2014, *104* (7), 1942–1978.

_ **, Stephen J. Redding, and David E. Weinstein**, "The Impact of the 2018 Tariffs on Prices and Welfare," *Journal of Economic Perspectives*, 2019, *33* (4), 187–210.

_ **,** _ **, and** _ , "Who's Paying for the US Tariffs? A Longer-Term Perspective," *AEA Papers and Proceedings*, 2020, *110*, 541–546.

Antràs, Pol, Teresa C. Fort, and Felix Tintelnot, "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 2017, *107* (9), 2514–2564.

Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

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

Autor, David H., David Dorn, and Gordon H. Hanson, "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 2013, *103* (6), 2121–2168.

Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott, "Importers, Exporters, and Multinationals: A Portrait of Firms in the U.S. That Trade Goods," in "Producer Dynamics: New Evidence from Micro Data," University of Chicago Press, 2009.

Blaum, Joaquin, Claire Lelarge, and Michael Peters, "The Gains from Input Trade with Heterogeneous Importers," *American Economic Journal: Macroeconomics*, 2018, *10* (4), 77–127.

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

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

Cavallo, Alberto, Gita Gopinath, Brent Neiman, and Jenny Tang, "Tariff Pass-Through at the Border and at the Store: Evidence from US Trade Policy," *American Economic Review: Insights*, 2021, *3* (1), 19–34.

Eaton, Jonathan and Samuel Kortum, "Technology, Geography, and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

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

Flaaen, Aaron, Ali Hortacsu, and Felix Tintelnot, "The Production Relocation and Price Effects of US Trade Policy: The Case of Washing Machines," *American Economic Review*, 2020, *110* (7), 2103–2127.

Gilchrist, Simon, Raphael Schoenle, Jae Sim, and Egon Zakrajšek, "Inflation Dynamics during the Financial Crisis," *American Economic Review*, 2017, *107* (3), 785–823.

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

Gopinath, Gita and Brent Neiman, "Trade Adjustment and Productivity in Large Crises," *American Economic Review*, 2014, *104* (3), 793–831.

_ , Oleg Itskhoki, and Roberto Rigobon, "Currency Choice and Exchange Rate Pass-Through," *American Economic Review*, 2010, *100* (1), 304–336.

Halpern, László, Miklós Koren, and Adam Szeidl, "Imported Inputs and Productivity," *American Economic Review*, 2015, *105* (12), 3660–3703.

Lu, Dan, Asier Mariscal, and Luis-Fernando Mejía, "How firms accumulate inputs: Evidence from import switching," *Journal of International Economics*, 2024, *148*, 103847.

Monarch, Ryan, ""It's Not You, It's Me": Prices, Quality, and Switching in U.S.-China Trade Relationships," *Review of Economics and Statistics*, 2021, *104* (5), 909–928.

Nakamura, Emi and Jón Steinsson, "Five Facts about Prices: A Reevaluation of Menu Cost Models," *Quarterly Journal of Economics*, 2008, *123* (4), 1415–1464.

Stock, James H. and Motohiro Yogo, "Testing for Weak Instruments in Linear IV Regression," in "Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg," Cambridge University Press, 2005, pp. 80–108.

US Bureau of Labor Statistics, "Quality Adjustment in the Producer Price Index," Technical Report, US Bureau of Labor Statistics, Washington, DC 2014.

Supplementary Material

— For Online Publications Only —

A Merging PPI and LFTTD Databases

We merge the PPI and LFTTD databases using a multi-step name-matching algorithm designed to maximize match rates while minimizing false positives. Because the two databases do not share a common firm identifier, we rely on business names as the primary matching variable.[5]

The matching procedure proceeds in several stages. First, we conduct an iterative word match routine that is repeated on two versions of the business name: (i) the raw original text from the PPI records and (ii) a cleaned version of the name (converted to uppercase, with punctuation such as "&" and "-" removed, and common business suffixes such as "INC," "CO," "LLC," "CORP," and "LTD" standardized and dropped).

The word match routine is carried out in three steps. In the first step, we directly match the full cleaned name string against the LFTTD database. This step captures exact or near-exact name matches. In the second step, the name string is tokenized into its constituent words, and we attempt to match the first five words of the name against the LFTTD. This step captures cases where business names differ only in trailing words (e.g., division names or location identifiers). The third step is identical to the second step except that we use a maximum of four words only, capturing additional matches where names share a common root but differ in specificity.

After exhausting the deterministic matching steps, we apply a fuzzy name matching procedure to the remaining unmatched records. Specifically, we use the SAS DQMATCH function with a sensitivity parameter of 95 to generate a standardized version of the cleaned business name. We retain matches to a unique firm identifier. For cases in which a PPI business name matches to multiple LFTTD firm identifiers, we resolve the ambiguity by matching the name strings word by word and retaining only those matches for which at least two constituent words agree, starting from the maximum possible number of matched words.

[5] In constructing our sample from the LFTTD, we require: (i) positive and nonmissing import or export value; (ii) nonmissing shipping costs and duties; (iii) nonmissing firm identifier; and (iv) we drop warehouse entries.

To further validate the matches, we impose several consistency checks. We require that matched firms belong to the same broad industry group (at the 2-digit NAICS level) and, where geographic information is available, that they are located in the same state. These filters help eliminate spurious matches arising from common business names (e.g., "American Manufacturing" or "National Products").

The resulting matched sample contains approximately 6,000 unique manufacturing firms that appear in both the PPI and LFTTD databases during the 2005–2022 period. These firms account for a substantial share of total U.S. manufacturing imports, suggesting that our sample is representative of the population of importing manufacturers.

B Empirical Results using Alternative IV

As a robustness check, we also construct an instrument using de facto (effectively collected) tariff rates:

$$\Delta \log \left(1 + d_{j,t}^{\mathrm{IV,eff}}\right) \equiv \sum_{(h,c)} \frac{v_{(h,c)\to j}^{(2017)}}{\sum_{(h,c)} v_{(h,c)\to j}^{(2017)}} \Delta \log \left(1 + d_{h,c,t}^{eff}\right). \tag{B-1}$$

Table B-1 reports 2SLS estimates using the de facto Bartik IV based on Specifications (I) and (II). The results closely parallel the de jure IV: 0.211 (0.035) versus 0.203 (0.035) in specification (I) unweighted, and 0.152 (0.034) versus 0.144 (0.033) in specification (II). The first-stage is even stronger, with coefficients of 0.82 and F-statistics exceeding 830. The near-identical second-stage estimates across the two IV constructions—despite the de jure and de facto tariff rates differing by the extent of product exclusions and administrative exemptions—provide reassuring evidence that our findings are not sensitive to the specific treatment of tariff exceptions.

Table B-2 confirms these patterns using the de facto Bartik IV. In specification (III), the IV estimate is 0.600 (0.145) unweighted and 0.544 (0.143) weighted—very close to the de jure results of 0.577 and 0.531. In specification (IV), the estimates are 0.589 (0.151) and 0.497 (0.160). The first-stage is stronger than for the de jure instrument, with coefficients of 0.76 and F-statistics exceeding 450. The robust agreement between the de jure and de facto IV results across all specifications reinforces the conclusion that our pass-through estimates are not artifacts of how tariff exemptions are measured.

Table B-1: IV Regression Results, De Facto Tariffs

	Specification (I)		Specification (II)	
Panel A: Second-stage results				
	Unwt.	Wt.	Unwt.	Wt.
$\Delta \log(1 + d^{\text{input,eff}})$	0.211	0.201	0.152	0.127
	(0.035)	(0.036)	(0.034)	(0.037)
$\Delta \log(1 + d^{\text{output,eff}})$	0.242	0.267		
	(0.070)	(0.074)		
Observations	296k	296k	316k	316k
Firms	3,000	3,000	3,200	3,200
Panel B: First-stage results				
$\Delta \log(1 + d^{\text{IV,eff}})$	0.819	0.819	0.815	0.816
	(0.028)	(0.027)	(0.025)	(0.024)
R^2	0.700	0.695	0.748	0.748
F-statistic	830.6	931.6	1039.0	1160.0
Firm FE	Yes	Yes	Yes	Yes
Industry×Month FE	No	No	Yes	Yes

Notes: Panel A reports 2SLS estimates using the de facto Bartik IV from equation (B-1). Panel B reports first-stage results. Standard errors clustered at the firm level. F-statistic is the Kleibergen-Paap weak-instrument test.

The first-stage regressions in Tables 5 and B-2 display two systematic patterns. First, the first-stage coefficients are uniformly below one (0.53–0.76), confirming the structural prediction that endogenous import substitution compresses the contemporaneous effective tariff relative to the pre-determined Bartik instrument. Second, the de facto IV yields larger first-stage coefficients (0.75–0.76) than the de jure IV (0.65–0.67), consistent with the de facto measure absorbing less measurement noise because it reflects actually collected duties rather than announced statutory rates.

C Proofs

Proof of Proposition 1. (a) Step 1 (origins). Minimize $\sum_{c \in J_f} p_c^M M_{f,c}$ subject to $M_f = 1$ in (10). Writing $S \equiv \sum_c \xi_c^{1/(1+\eta)} M_{f,c}^{\eta/(1+\eta)}$, so $M_f = S^{(1+\eta)/\eta}$, the first-order condition $p_c^M = \lambda\, \partial M_f / \partial M_{f,c} = \lambda \xi_c^{1/(1+\eta)} (M_f / M_{f,c})^{1/(1+\eta)}$ gives $M_{f,c} = \xi_c M_f (\lambda / p_c^M)^{1+\eta}$. Substituting into the constraint, each summand of S becomes $\lambda^\eta \xi_c (p_c^M)^{-\eta}$, so $1 =$

Table B-2: IV with Exposure Interaction, De Facto Tariffs

Panel A: Second-stage results	Specification (III)		Specification (IV)	
	Unwt.	Wt.	Unwt.	Wt.
$III \times \Delta\log(1+d^{\text{input,eff}})$	0.600	0.544	0.589	0.497
	(0.145)	(0.143)	(0.151)	(0.160)
$IS \times \Delta\log(1+d^{\text{output,eff}})$	0.604	0.690		
	(0.187)	(0.199)		
Obs.	296k	296k	296k	296k
Firms	3,000	3,000	3,000	3,000
Panel B: First-stage results				
$III \times \Delta\log(1+d^{\text{IV,eff}})$	0.761	0.763	0.748	0.749
	(0.035)	(0.034)	(0.035)	(0.032)
R^2	0.712	0.708	0.753	0.753
F-statistic	461.6	507.1	458.1	532.2
Firm FE	Yes	Yes	Yes	Yes
Ind.×Month FE	No	No	Yes	Yes

Notes: Panel A reports 2SLS estimates using the de facto Bartik IV from equation (B-1). Panel B reports first-stage results. Standard errors clustered at the firm level. F-statistic is the Kleibergen-Paap weak-instrument test.

$(\lambda^\eta \Xi_f)^{(1+\eta)/\eta}$ and $\lambda = \Xi_f^{-1/\eta}$. By Euler's theorem the minimized cost is $\lambda M_f = \lambda$, giving $P_f^M = \Xi_f^{-1/\eta}$; dividing the demand by $P_f^M M_f$ gives $s_{f,c}^M = \tilde{\xi}_c/\Xi_f$. For $J_f = \varnothing$ set $\Xi_f = 0$, $P_f^M = +\infty$, $M_f = 0$.

Step 2 (materials). Minimize $p_{D,s}Z_f + P_f^M M_f$ subject to $X_f = 1$. This is the same CES problem with two goods at prices $(p_{D,s}, P_f^M)$, weights (a_D, a_M) and exponent ζ, so $P_f^X = [a_D p_{D,s}^{-\zeta} + a_M (P_f^M)^{-\zeta}]^{-1/\zeta}$, the conditional demands are $Z_f = a_D X_f (P_f^X/p_{D,s})^{1+\zeta}$ and $M_f = a_M X_f (P_f^X/P_f^M)^{1+\zeta}$, and the import share of materials is $\alpha_f = P_f^M M_f/(P_f^X X_f) = a_M (P_f^X/P_f^M)^\zeta$, which rearranges to (18) and is decreasing in $P_f^M/p_{D,s}$ because $\zeta > 0$.

Step 3 (labor and materials). Minimize $wL_f + P_f^X X_f$ subject to (9). With multiplier λ the first-order conditions give $X_f = \lambda\phi_s Y_f/P_f^X$ and $L_f = \lambda(1-\phi_s)Y_f/w$; substituting into the technology and cancelling Y_f,

$$\lambda = \frac{1}{A_f}\Big(\frac{P_f^X}{\phi_s}\Big)^{\phi_s}\Big(\frac{w}{1-\phi_s}\Big)^{1-\phi_s}, \qquad \text{(C-1)}$$

and $wL_f + P_f^X X_f = \lambda Y_f$, so λ is marginal cost, $P_f^X X_f = \phi_s TVC_f$ and $wL_f = (1-\phi_s)TVC_f$.

Step 4. From Step 2, $1-\alpha_f = a_D p_{D,s}^{-\zeta}(P_f^X)^{\zeta}$, so $P_f^X = a_D^{-1/\zeta} p_{D,s}(1-\alpha_f)^{1/\zeta} = a_D^{-1/\zeta} p_{D,s}/B_f$ with $B_f \equiv (1-\alpha_f)^{-1/\zeta}$. Substituting into (C-1) and multiplying by Y_f gives (16).

Step 5. From Steps 1–2, $\alpha_f/(1-\alpha_f) = (a_M/a_D)p_{D,s}^{\zeta}\Xi_f^{\zeta/\eta}$, so $1/(1-\alpha_f) = 1 + (a_M/a_D)p_{D,s}^{\zeta}\Xi_f^{\zeta/\eta}$ and $B_f = B(\Xi_f)$ as in (17), strictly increasing in Ξ with $B(0) = 1$ and $B_f > 1$ iff $\Xi_f > 0$. In particular the set enters only through Ξ_f.

(b) By Euler's theorem in Step 1, $\sum_{c\in J_f} p_c^M M_{f,c} = P_f^M M_f$; by Step 2, $P_f^M M_f = \alpha_f P_f^X X_f$; by Step 3, $P_f^X X_f = \phi_s TVC_f$. Hence imports over variable cost equal $\phi_s\alpha_f$.

(c) By (a), $\ln MC_f = \ln C_s^* - \ln A_f - \phi_s \ln B(\Xi_f)$. Write $u \equiv (a_M/a_D)p_{D,s}^{\zeta}\Xi_f^{\zeta/\eta}$, so that $B = (1+u)^{1/\zeta}$ and $\alpha_f = u/(1+u)$ by (18). Then $d\ln B = \zeta^{-1}u\, d\ln u/(1+u) = (\alpha_f/\zeta)\, d\ln u$ and $d\ln u = \zeta\, d\ln p_{D,s} + (\zeta/\eta)\, d\ln \Xi_f$, so

$$d\ln B = \alpha_f\, d\ln p_{D,s} + \frac{\alpha_f}{\eta}\, d\ln \Xi_f.$$

With the set fixed, $\ln\tilde{\xi}_c = \ln\xi_c - \eta\ln p_c^* - \eta\ln(1+\tau_c)$ gives $d\ln\Xi_f = \sum_c s_{f,c}^M d\ln\tilde{\xi}_c = -\eta\sum_c s_{f,c}^M\widehat{\tau}_c$. Since $\ln C_s^*$ has elasticities ϕ_s and $1-\phi_s$ with respect to $p_{D,s}$ and w,

$$\begin{aligned} d\ln MC_f &= \phi_s d\ln p_{D,s} + (1-\phi_s)d\ln w - \phi_s\Big[\alpha_f d\ln p_{D,s} + \frac{\alpha_f}{\eta}d\ln\Xi_f\Big] \\ &= III_f \sum\nolimits_c s_{f,c}^M\widehat{\tau}_c + (\phi_s - III_f)\, d\ln p_{D,s} + (1-\phi_s)\, d\ln w, \end{aligned}$$

using $III_f = \phi_s\alpha_f$. For the finite change, with the set fixed,

$$\Xi_f' = \sum\nolimits_c \tilde{\xi}_c e^{-\eta\widehat{\tau}_c} = \Xi_f \sum\nolimits_c s_{f,c}^{M,\text{pre}} e^{-\eta\widehat{\tau}_c}, \qquad \text{so} \qquad \Delta\ln P_f^M = -\frac{1}{\eta}\ln\sum\nolimits_c s_{f,c}^{M,\text{pre}} e^{-\eta\widehat{\tau}_c}.$$

Since $x \mapsto e^{-\eta x}$ is convex, Jensen's inequality gives $\sum_c s_{f,c}^{M,\text{pre}} e^{-\eta\widehat{\tau}_c} \geq e^{-\eta\sum_c s_{f,c}^{M,\text{pre}}\widehat{\tau}_c}$, which is the Laspeyres bound. The post-shock shares are

$$s_{f,c}^{M,\text{post}} = \frac{s_{f,c}^{M,\text{pre}} e^{-\eta\widehat{\tau}_c}}{\sum_{c'} s_{f,c'}^{M,\text{pre}} e^{-\eta\widehat{\tau}_{c'}}}, \qquad \text{so} \qquad \sum\nolimits_c s_{f,c}^{M,\text{post}} e^{\eta\widehat{\tau}_c} = \Big(\sum\nolimits_c s_{f,c}^{M,\text{pre}} e^{-\eta\widehat{\tau}_c}\Big)^{-1},$$

whence $\Delta\ln P_f^M = \eta^{-1}\ln\sum_c s_{f,c}^{M,\text{post}} e^{\eta\widehat{\tau}_c} \geq \sum_c s_{f,c}^{M,\text{post}}\widehat{\tau}_c$ by Jensen again, the Paasche bound. These weighted-*log* bounds use the CES form with $\eta \geq 0$. For a general homothetic aggregator only the ratio form

$$\ln \sum\nolimits_c s_{f,c}^{M,\mathrm{pre}} e^{\widehat{\tau}_c} \geq \Delta \ln P_f^M \geq -\ln \sum\nolimits_c s_{f,c}^{M,\mathrm{post}} e^{-\widehat{\tau}_c}$$

survives, because the old bundle at the new prices is feasible and symmetrically.

(d) By Step 4, $B_f = (1-\alpha_f)^{-1/\zeta}$, so $\ln MC_f = \ln C_s^* - \ln A_f + (\phi_s/\zeta)\ln(1-\alpha_f)$, and $\Delta \ln C_s^* = \phi_s \Delta \ln p_{D,s} + (1-\phi_s)\Delta \ln w$ while A_f, a_D and ϕ_s do not change. Nothing in Steps 1–4 restricts J_f, the p_c^M or the τ_c between the two equilibria, so (22) holds for any change, including a change of set. □

Proof of Proposition 2. First, by Proposition 1(a), $TVC(J) = f(\Xi(J))$ with f strictly decreasing and $\Xi(J) = \sum_{c\in J} \tilde{\xi}_c$ additive, while the fixed-cost bill $|J|\varepsilon_f C_F$ depends on J only through $|J|$. Among sets of size k, replacing an included origin by an excluded one with higher $\tilde{\xi}$ raises Ξ and lowers cost; repeating the exchange yields the top-k set. Hence the optimum lies among the $N+1$ prefixes and comparing $TVC_f^{(k)} + k\varepsilon_f C_F$ over $k = 0, \ldots, N$ selects it exactly. No property of f beyond monotonicity is used.

Second, with f convex on $[0, \Xi_N]$ and $\tilde{\xi}_{k+1} < \tilde{\xi}_k$,

$$C_F \varepsilon_f^{(k+1)} = f(\Xi_k) - f(\Xi_k + \tilde{\xi}_{k+1}) < f(\Xi_k) - f(\Xi_k + \tilde{\xi}_k) < f(\Xi_{k-1}) - f(\Xi_{k-1} + \tilde{\xi}_k) = C_F \varepsilon_f^{(k)},$$

the first inequality because f is decreasing and $\tilde{\xi}_{k+1} < \tilde{\xi}_k$, the second because a decreasing convex function has decrements that shrink as the base rises and $\Xi_{k-1} < \Xi_k$. Both comparisons are over the intervals $[\Xi_{k-1}, \Xi_{k+1}]$, which is why convexity is needed there and not only at the rungs. Total cost on the top-k set satisfies $TC(k) - TC(k-1) = C_F(\varepsilon_f - \varepsilon_f^{(k)})$, negative below the cutoff and positive above it; with the cutoffs strictly decreasing the increments change sign once, TC is quasi-convex in k, and the minimum is at the largest k with $\varepsilon_f \leq \varepsilon_f^{(k)}$. □

Proof of Proposition 3. By Proposition 2(b) the firm holds exactly k origins when $\varepsilon_f \in (\varepsilon_f^{(k+1)}, \varepsilon_f^{(k)}]$, an event of probability $F_k - F_{k+1}$. Telescoping the optimised total cost,

$$TC(\varepsilon) = TVC_f^0 - C_F \sum\nolimits_{k=1}^{N} \left(\varepsilon_f^{(k)} - \varepsilon\right) \mathbf{1}\{\varepsilon \leq \varepsilon_f^{(k)}\},$$

since each origin k that is taken contributes a saving $C_F \varepsilon_f^{(k)}$ and costs $C_F \varepsilon$. Taking expectations and using $\mathbb{E}[(x-\varepsilon)\mathbf{1}\{\varepsilon \leq x\}] = \int_0^x (x-\varepsilon)dF = \int_0^x F = G(x)$ by parts gives (26). For (27), constant returns make $TVC_f^{(k)} = MC_f^{(k)} Y_f$ with $MC_f^{(k)}$ independent of Y_f, so every cutoff is proportional to output, $\partial \varepsilon_f^{(k)} / \partial Y_f = \varepsilon_f^{(k)} / Y_f$; differentiating (26) with $G' = F$ gives the middle expression. The right-hand expression is the direct probability-weighted average, and the two agree because $c_f^0 - MC_f^{(k)} = c_f^F \sum_{j \leq k} \varepsilon_f^{(j)}$ and Abel summation. □

Proof of Proposition 4. (a) Differentiate the right-hand expression in (27). Two kinds of term arise: the change in $MC_f^{(k)}$ holding the masses fixed, and the movement of the boundaries. For the first, by Proposition 1(c) applied at the top-k set, $d \ln MC_f^{(k)} = \Lambda_k$ at fixed $p_{D,s}, w$, so $dMC_f^{(k)} = MC_f^{(k)} \Lambda_k = \Psi_k$; weighting by the masses gives the intensive term. For the second, differentiating $\sum_k MC_f^{(k)}[F_k - F_{k+1}]$ through the F_k and using $d\varepsilon_f^{(k)} = (\Psi_{k-1} - \Psi_k)/c_f^F$ — which follows from (23) after dividing by Y_f — and then summing by parts gives $\sum_k \Psi_k (v_k - v_{k+1})$ with $v_k = \varepsilon_f^{(k)} F'(\varepsilon_f^{(k)})$. That the intensive term equals $Y_f^{-1} d\mathbb{E}[TC_f]$ follows by differentiating (26) and using $G' = F$: the boundary terms cancel because a firm at a cutoff is indifferent, which is the envelope theorem.

(b) For F uniform on $(0, \bar{\varepsilon})$ with interior cutoffs, $F_k = \varepsilon_f^{(k)} / \bar{\varepsilon}$ and $F' = 1/\bar{\varepsilon}$, so $v_k = \varepsilon_f^{(k)} / \bar{\varepsilon} = F_k$ for every k. The two sums in (30) are then identical, giving (31). Nothing in the argument restricts the tariff vector.

(c) With $\widehat{\tau}_c = \widehat{\tau}$ for all c, $d \ln \Xi_k = -\eta \widehat{\tau}$ for every k, so $d\varepsilon_f^{(k)} \propto f'(\Xi_{k-1})\Xi_{k-1} - f'(\Xi_k)\Xi_k$. Note that $-f'(\Xi)\Xi = \phi_s \alpha(\Xi) f(\Xi)/\eta$, whose logarithmic derivative in Ξ is $[\zeta(1-\alpha) - \phi_s \alpha]/\eta$, positive exactly when $\alpha < \zeta/(\zeta + \phi_s)$. Under that condition $-f'(\Xi)\Xi$ is increasing, so the bracket is negative and every cutoff falls. Then $v_k - v_{k+1} \geq 0$ requires only $F' \geq 0$ together with the ordering of the cutoffs, and the extensive term is a non-negative combination of the $\Psi_k > 0$, strictly positive if $F' > 0$ at some cutoff. □

Proof of Proposition 5. Immediate from Proposition 1(c) at $d \ln p_{D,s} = d \ln w = 0$ together with $p_f = \mu_s MC_f$ and μ_s constant, so that $d \ln p_f = d \ln MC_f$. The population regression of $d \ln p_f$ on $z_f = III_f \sum_c s_{f,c}^M \widehat{\tau}_c$ is then a regression of z_f on itself. The second claim is Proposition 4(b). □

Proof of Proposition 6. (a) Let $S \equiv \sum_g q_g^r$ with $r \equiv (\sigma_s - 1)/\sigma_s$. Inverting (32), $p_f = R_s q_f^{-1/\sigma_s}/S$ and $\varsigma_f = q_f^r/S$. Holding rivals' quantities and R_s fixed, $d\ln S/d\ln q_f = r\varsigma_f$, so

$$-\frac{d\ln p_f}{d\ln q_f} = \frac{1}{\sigma_s} + r\varsigma_f = \varsigma_f + \frac{1-\varsigma_f}{\sigma_s} = \frac{1}{\epsilon_f},$$

which is (34). Marginal revenue is $p_f(1 - 1/\epsilon_f)$, so the first-order condition $MC_f = p_f(1 - 1/\epsilon_f)$ gives (36) after substituting $1/\epsilon_f$ and simplifying. Marginal revenue is strictly decreasing in q_f and runs from $+\infty$ to 0, so the solution is unique and is the global maximum.

(b) Differentiating (36) at fixed σ_s, $d\ln\mu_f = \frac{\varsigma_f}{1-\varsigma_f} d\ln\varsigma_f$; and from (32), $d\ln\varsigma_f = (1-\sigma_s)(d\ln p_f - d\ln P_s)$. Hence $d\ln\mu_f = -\gamma_f(d\ln p_f - d\ln P_s)$. Differentiating $\ln p_f = \ln\mu_f + \ln MC_f$ and collecting,

$$(1+\gamma_f)\, d\ln p_f = d\ln MC_f + \gamma_f d\ln P_s,$$

which is the first equality in (37) after dividing by $1+\gamma_f$ and subtracting $d\ln P_s$. Multiplying by ς_f, summing, and using $\sum_f \varsigma_f d\ln p_f = d\ln P_s$ gives $(\sum_f \varsigma_f\psi_f) d\ln P_s = \sum_f \varsigma_f\psi_f d\ln MC_f$, which is the second equality. A common cost change $d\ln MC_g = \delta$ gives $d\ln P_s = \delta$ since the $\tilde{\varsigma}$ sum to one, and then $d\ln p_f = \delta$ for every f with markups unchanged. The three alternative derivatives follow by substituting, respectively, $d\ln P_s = \tilde{\varsigma}_f d\ln MC_f$; $d\ln P_s = \varsigma_f d\ln p_f$; and, along the best response, $d\ln p_f = -\vartheta_f d\ln q_f$ with $\vartheta_f = 1/\epsilon_f$ together with $d\ln\varsigma_f = r(1-\varsigma_f)d\ln q_f$.

(c) Write $\bar{\psi}_s \equiv \mathbb{E}_\varsigma[\psi_f]$. By the definition of $\tilde{\varsigma}$,

$$\sum_g \tilde{\varsigma}_g\, d\ln MC_g = \frac{\mathbb{E}_\varsigma\left[\psi_f\, d\ln MC_f\right]}{\bar{\psi}_s}, \qquad \mathbb{E}_\varsigma\left[\psi_f\, d\ln MC_f\right] = \bar{\psi}_s \mathbb{E}_\varsigma\left[d\ln MC_f\right] + \mathrm{Cov}_\varsigma\left(\psi_f, d\ln MC_f\right),$$

which together give (38). □

Proof of Proposition 7. With constant markups $p_f = \mu_s MC_f$, so by Proposition 1(d),

$$\Delta\ln p_f = \phi_s\Delta\ln p_{D,s} + (1-\phi_s)\Delta\ln w + \frac{\phi_s}{\zeta}\Delta\ln s_f^D.$$

Substituting into $P_s = [\sum_f p_f^{1-\sigma_s}]^{1/(1-\sigma_s)}$ and factoring out the terms common to the sector,

$$\Delta \ln P_s = \phi_s \Delta \ln p_{D,s} + (1-\phi_s)\Delta \ln w + \frac{1}{1-\sigma_s} \ln \sum\nolimits_f \varsigma_f \exp\Big[(1-\sigma_s)\frac{\phi_s}{\zeta}\Delta \ln s_f^D\Big],$$

where ς_f are the pre-change revenue shares; the last term is Λ_s of (41) with $\chi_s = \phi_s(\sigma_s - 1)/\zeta$. A first-order expansion of the log-sum in the share changes gives $\Lambda_s = (\phi_s/\zeta)\sum_f \varsigma_f \Delta \ln s_f^D$. For (42), hold $p_{D,s}$, w and the sets fixed and use Proposition 1(c), which gives $\Delta \ln MC_f = III_f \sum_c s_{f,c}^M \widehat{\tau}_c$ to first order; averaging with weights ς_f gives Λ_s^τ. The inequality is Jensen applied to the convex function $x \mapsto e^{-\chi_s x}$ inside the log-sum, with equality iff $\Delta \ln s_f^D$ is constant across f. □

Proof of Proposition 8. (a) Substitute (39) into (40) with $\Delta \ln w = 0$: $\Delta \ln P_s = \Lambda_s + \phi_s \sum_{s'} \omega_{s,s'} \Delta \ln P_{s'}$, which in matrix form is (43). Row s of $\Gamma\Omega$ has non-negative entries summing to ϕ_s, so $\|\Gamma\Omega\|_\infty = \max_s \phi_s < 1$ and $\rho(\Gamma\Omega) \leq \|\Gamma\Omega\|_\infty$; the Neumann series therefore converges geometrically at that rate and the norm bound follows from $\|(I-A)^{-1}\| \leq 1/(1-\|A\|)$.

(b) To first order in tariffs with sets fixed, Proposition 1(c) gives $d \ln MC_f = \Lambda_f^\tau + (\phi_s - III_f) d \ln p_{D,s}$, so the revenue-weighted average over the sector is $\Lambda_s^\tau + \phi_s(1-\bar{\alpha}_s) d \ln p_{D,s}$, using $\sum_f \varsigma_f III_f = \phi_s \bar{\alpha}_s$. Substituting (39) gives (44), and the row sums of $\Gamma \mathrm{diag}(1-\bar{\alpha}_s)\Omega$ are $\phi_s(1-\bar{\alpha}_s)$. The statement about κ_s^{ext} is Proposition 4 applied firm by firm and averaged.

(c) $\Delta \ln P^{PPI} = \sum_s \vartheta_s \Delta \ln P_s = \vartheta' \mathcal{M} \Lambda^\tau$ by (b); writing $[\mathcal{M}\Lambda^\tau]_s = m_s \Lambda_s^\tau$ defines m_s. Since $\mathcal{M} = \sum_{n \geq 0} (\Gamma \mathrm{diag}(1-\bar{\alpha})\Omega)^n$ is a non-negative matrix with unit diagonal at $n = 0$ and $\Lambda^\tau \geq 0$, $[\mathcal{M}\Lambda^\tau]_s \geq \Lambda_s^\tau$, so $m_s \geq 1$ wherever $\Lambda_s^\tau > 0$. The two special cases are immediate. □

Proof of Proposition 9. The household's Cobb–Douglas problem gives $P_s C_s = \beta_s E$, and the inner CES problem gives (46); the ideal index and indirect utility follow by substitution. Intermediate buyers use the same nest by (11) and foreign buyers by assumption, so all demands for firm f's good are proportional to $(p_f/P_s)^{-\sigma_s}$ and aggregate to $R_f = \varsigma_f R_s$ with R_s as in (47).

For Proposition 9(i), by Proposition 1 and constant markups a firm spends $\phi_s(1-\alpha_f)$ of variable cost on domestic materials and variable cost is R_f/μ_s, so sector s's domestic-material purchases are $\nu_s R_s$ with ν_s as in (48); by (39) a share $\omega_{s,s'}$ of that goes to sector s'. Substituting into (47) gives $R = \beta E + KR + \mathcal{X}$ with the stated K, and $\rho(K) \leq \|K\|_1 = \max_s \nu_s < 1$ since the columns of K sum to $\nu_{s'}$. Part (ii) is $\mathbf{1}'(I-K) = (\mathbf{1}-\nu)'$ transposed, and the decomposition of $1-\nu_s$ is the firm's budget: profits $1/\sigma_s$, labor $\frac{\sigma_s-1}{\sigma_s}(1-\phi_s)$ and duty-inclusive imports $\frac{\sigma_s-1}{\sigma_s}\phi_s\bar{\alpha}_s$. For (iii), household income is $w\bar{L} - \Phi + T + D$ plus profits $\varpi' R$; substituting R from (i) and solving for E gives (49), and the denominator is in $(0,1)$ because $\varpi'(I-K)^{-1}\beta$ is a share of final spending accruing as profit. $\square$

Proof of Proposition 10. *(a)* By Proposition 1(c) with $d\ln w = 0$ and constant markups, $\Delta \ln p_f = III_f \widehat{\tau}_f + (\phi_s - III_f)\Delta \ln p_{D,s}$, which is (51) after grouping.

(b) Within sector s, demeaning removes the term $\phi_s \Delta \ln p_{D,s}$, which is constant across the sector's firms, leaving

$$\Delta \ln p_f - \overline{\Delta \ln p}_s = (z_f - \bar{z}_s) - \Delta \ln p_{D,s}\left(III_f - \overline{III}_s\right).$$

The pooled within regression coefficient is $\sum_s \sum_f (z_f - \bar{z}_s)(\Delta \ln p_f - \overline{\Delta \ln p}_s)/\sum_s V_s$, which on substituting gives (52). If $\widehat{\tau}_f = \widehat{\tau}_s$ then $z_f = III_f \widehat{\tau}_s$, so $C_s = \widehat{\tau}_s D_s$ and $V_s = \widehat{\tau}_s^2 D_s$ with $D_s \equiv \sum_f (III_f - \overline{III}_s)^2$, and the stated weighted form follows. It is below one whenever $\Delta \ln p_{D,s} > 0$ for some sector with $D_s > 0$, which holds when a sector with dispersed exposure buys from a sector with positive direct exposure.

(c) That $\Lambda_s^\tau = III_s \bar{\tau}_s$ is the definition of $\bar{\tau}_s$ as the import-bill-weighted mean of $\widehat{\tau}_f$: $\sum_f \varsigma_f III_f \widehat{\tau}_f = III_s \sum_f (\varsigma_f III_f / III_s)\widehat{\tau}_f$, and the weights $\varsigma_f III_f/III_s$ are proportional to import bills. Under proportionality $X_s = III_s \bar{\tau}_s = \Lambda_s^\tau$. Regressing $\Delta \ln P_s = [\mathcal{M}\Lambda^\tau]_s = \Lambda_s^\tau + [(\mathcal{M}-I)\Lambda^\tau]_s$ on Λ_s^τ with an intercept gives a slope of 1 from the first term plus $\mathrm{Cov}([(\mathcal{M}-I)\Lambda^\tau]_s, \Lambda_s^\tau)/\mathrm{Var}(\Lambda_s^\tau)$ from the second, which is (54). If $m_s = m$ for all s then $[(\mathcal{M}-I)\Lambda^\tau]_s = (m-1)\Lambda_s^\tau$ and the covariance ratio is $m-1$; if the propagated term is constant its covariance with anything is zero.

(d) $\beta^{\text{firm}} < 1$ by (b) and $\beta^{\text{sec}} \geq 1$ by (c) when the covariance is non-negative. With a common multiplier and a common discount, (b) gives $\beta^{\text{firm}} = 1 - \rho$ and (c) gives $\beta^{\text{sec}} = m$, so the ratio is $m/(1-\rho)$; under Cournot pricing the denominator carries the extra factor ψ. $\square$