

# Higher Tariffs, Lower Productivity: Evidence from U.S. Manufacturing

Luca Guerrieri\*

Preliminary — September 20, 2026

## Abstract

Between 2018 and 2019 the United States raised tariffs on roughly \$350 billion of imports. I ask how this protection affected the productivity of the targeted manufacturing industries, and find that it lowered it. Because tariffs are not assigned at random, I exploit the trade war’s design: U.S. authorities front-loaded intermediate and capital goods and delayed consumer goods, making an industry’s consumer-good share a strong instrument for its tariff exposure. The quartiles of the 2017–19 increase in the statutory tariff are 6.8 percentage points apart, and across that gap an industry at the top quartile ends 2019 with total factor productivity about 3.6 percent below one at the bottom quartile. This is a relative comparison across industries, not an aggregate. A second design, which requires no instrument, puts the same contrast at 2.4 percent. The effect operates through the own-product tariff rather than input costs.

---

\*The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

# 1 Introduction

Between 2018 and 2019 the United States raised tariffs on roughly \$350 billion of imports, the largest increase in U.S. protection since the 1930s, with the stated aim of reviving domestic manufacturing. This paper asks what the tariffs did to the *productivity* of the protected industries, and finds that they lowered it. Industries more exposed to the 2018–19 tariff increases experienced relative declines in total factor productivity (TFP) that grow over the following years.

A simple example crystallizes the effects of higher tariffs according to the standard trade model (Krugman et al., 2012, ch. 6). Take an economy that produces two goods, and a tariff that raises the home price of the second. Production shifts toward the protected good along the economy’s production possibility frontier, as in Figure 1(a). The new equilibrium is inefficient, but technology is untouched and the economy stays on its production frontier (Kehoe and Ruhl, 2008). Extensions of the standard trade model with heterogeneous firms do let trade costs move industry productivity (Melitz, 2003; Melitz and Trefler, 2012), but the evidence behind them comes mostly from trade liberalizations. Following the standard trade model, which holds total factor productivity fixed, empirical work on the 2018–19 tariffs has measured their costs in prices, trade flows, employment and welfare (Amiti et al., 2019; Fajgelbaum et al., 2020; Flaaen and Pierce, 2024). What I find is the case the example omits. As illustrated in Figure 1(b), production falls inside the production frontier, so the same inputs yield less output; in other words, higher tariffs compress TFP.


Protection is not handed out at random. Tariffs go to industries for reasons that are themselves correlated with productivity, among them import pressure and decline, so comparing more- and less-protected industries confounds the effect of the tariff with whatever earned the industry its protection. The endogeneity is visible in my data: the leads of the tariff change predict current productivity, so protected industries were already on distinct paths.

My identification exploits a feature of *how* the 2018–19 tariffs were assigned. The U.S. Trade Representative deliberately front-loaded intermediate and capital goods (Section 301 Lists 1–3, imposed in 2018) and delayed consumer goods (List 4, 2019) to limit the pass-through to consumer prices (Fajgelbaum et al., 2020; Cavallo et al., 2021; Bown, 2021).<sup>1</sup> For reasons of consumer-price exposure rather than of the industry’s own productivity, an industry’s tariff exposure therefore depended heavily on its *consumer-good share*: the fraction

---

<sup>1</sup>The sequencing was explicit. When the consumer-goods wave (List 4) was finalized in August 2019, the administration carved out cellphones, laptops, video-game consoles, and much apparel and footwear and delayed their tariffs to December 15, 2019, with the President stating the delay was “for the Christmas season” so as not to raise prices on holiday shoppers (Bown, 2021).

Figure 1: A tariff in the standard trade model, and the margin this paper adds.



*Notes:* A schematic, not an estimate. The curve is a production possibility frontier over two goods. In panel (a), a tariff raises the home price of good 2; the price line pivots and production moves from A to B along the production frontier, with technology unchanged (Kehoe and Ruhl, 2008). The new equilibrium is inefficient for the world as a whole, even if one country can gain, and it changes welfare, real income and labor supply in ways the figure does not show. Panel (b) adds the margin this paper finds: production also moves from B to C, *inside* the production frontier, so a given bundle of inputs yields less output.

of output sold to final consumers rather than used as an input. I use the consumer-good share, measured from the Bureau of Economic Analysis (BEA) input-output accounts, as an instrument for tariff exposure.

In the clean pre-pandemic window, a one-percentage-point increase in the statutory output tariff lowers an industry's TFP by about 0.5 percent. The instrument is strong, with a first-stage  $F \approx 28$ , and the estimate is sharp. Since the quartiles of the 2017–19 tariff increase are 12.7 and 19.4 points, that amounts to an industry at the top quartile of exposure ending 3.6 percent below one at the bottom. The comparison is deliberately relative: a cross-sectional instrument identifies the *difference* between more- and less-exposed industries, while anything common to all of manufacturing sits in the constant. The estimate grows over windows ending in later years, but those windows overlap the COVID-19 pandemic, so I rest the magnitude on the pre-pandemic figure. Two checks support the design: the consumer-good share does *not* predict industries' *pre-war* productivity changes (the exclusion restriction passes its natural placebo), and the estimate is unchanged when I control for foreign retaliation, to which the instrument is uncorrelated. A second design, following Flaaen and Pierce (2024), identifies the same effect without an instrument at all: exposed and less-exposed industries move together through 2018 and separate sharply from 2019, and the two designs bracket the magnitude.

Even within narrowly defined industries, plants differ strikingly in productivity: in a typical four-digit U.S. manufacturing industry the establishment at the 90th percentile produces about twice as much output from the same inputs as the one at the 10th (Syverson, 2011). Where in this distribution does the tariff-induced decline occur? Separately, and identified off the full panel rather than the trade war, I find that the within-industry *dispersion* of productivity compresses: the 90–10 and interquartile spreads narrow, while neither tail’s own spread responds. Because the underlying statistics are industry-year percentiles rather than a panel in which individual plants can be followed, they identify *where* in the distribution the narrowing occurs, not which establishments moved. Every statistic is a

emphgap between two percentiles, so a spread that does not change is consistent with both of its ends shifting together. The level decline and the distributional compression are distinct findings, but together they are consistent with a downward leveling of productivity concentrated among the better performers.

What is the mechanism? The evidence points most clearly to competition. It is the *output* (own-product) tariff, which shields an industry’s home market from foreign rivals, that drives the TFP decline, while the tariff an industry pays on its inputs does not. This is the signature of a competition channel. Theory predicts that exposure to foreign rivals raises industry productivity by shifting output toward more productive firms (Melitz, 2003; Melitz and Trefler, 2012). The evidence on total factor productivity points the same way: after trade liberalization, Chilean plants facing import competition raised their own productivity, while reallocation toward more productive plants accounted for most of the aggregate gain (Pavcnik, 2002), and Chinese import competition raised patenting, IT intensity and TFP within European firms (Bloom et al., 2016). Protection that dulls this discipline can therefore let productivity slip. Two further questions bear on that reading, and the next paragraphs take them up: what increases in competition from other episodes did to productivity, and whether an input-cost channel could account for the decline instead.

If the tariffs act through competition, evidence on changes in competition from other episodes bears directly on them, and two strong designs supply it. The first is the China shock of Autor et al. (2013): an industry’s import penetration from China, instrumented by China’s exports to other high-income countries to strip out U.S.-specific demand. The second is the tariff-uncertainty shock of Pierce and Schott (2016a): when the United States granted China Permanent Normal Trade Relations (PNTR) in 2000, it made permanent the low “normal” tariff rates that Congress had until then renewed one year at a time, and industries with the widest gap between those rates and the punitive ones that would otherwise have applied lost the most manufacturing jobs. Both designs have been used chiefly to study employment; I extend them to productivity. On my data both reproduce

the employment losses, but neither detects a significant response of productivity or of its dispersion (Sections 8.1 and 8.2). Two readings are open. The increases in competition in those episodes may not have been sharp enough, relative to the noise in industry productivity, to register a response with statistical significance; the China-shock productivity estimates are in fact positive, the direction the competition channel predicts, and in other settings rising import competition has raised measured TFP (Pavcnik, 2002; Bloom et al., 2016). Or the effects of competition may be asymmetric: an increase may be absorbed on the labor margin, while a decrease, such as the protection of 2018–19, reaches the productivity frontier.

An alternative to a competition channel is an input-cost channel: tariffs could lower productivity by raising the cost and narrowing the variety of imported inputs, the channel Amiti and Konings (2007) and Topalova and Khandelwal (2011) find dominant when tariffs *fall*. Were it operative here, the *input* tariff (the input-cost-weighted average of the tariffs an industry’s suppliers face) should carry the effect. It does not: the input tariff is insignificant, and remains so when it is weighted by actual imported-input use from the BEA import matrix rather than by total input use (Section 8.3). The action is on the own-product tariff, not on the tariffs an industry pays for its inputs. The result is therefore the movement inside the production frontier of Figure 1(b) rather than the movement along it that the standard trade model allows.

The paper speaks to several literatures, laid out in Section 2. It adds to the evidence on trade policy and firm productivity, which comes mostly from trade liberalizations, often in developing economies. Amiti and Konings (2007) and Topalova and Khandelwal (2011) find that the *input*-tariff channel raises productivity when tariffs *fall*; I study a tariff *increase* in an advanced economy and find the action is on the *output* tariff, with the input channel absent. It adds a productivity margin to the growing literature on the 2018–19 trade war, which has focused on prices, trade flows, and welfare. It extends two influential designs for increases in competition, the China shock and the PNTR tariff-uncertainty shock, from employment to productivity. And it connects the level and dispersion of productivity to trade policy, complementing the misallocation literature.

The rest of the paper proceeds as follows. Section 3 builds the tariff and productivity measures; Section 4 sets out two channels through which a tariff can operate and the three designs used to identify them. Section 5 is the central result: it explains why the uninstrumented association is not reported as a finding (its leads predict the outcome) and then identifies the effect from the trade war’s own design, instrumenting the 2017–19 tariff increase with the industry’s consumer-good share. Section 6 identifies the same effect a second time without an instrument, borrowing the differential-exposure design of Flaaen and Pierce (2024): exposed and less-exposed industries move together through 2018 and separate from

2019, and the two designs bracket the magnitude. Section 7 asks where in the within-industry distribution the decline occurs, shows that the compression survives dropping the war years entirely, and subjects it to a timing placebo that the level specification fails and this one passes. Section 8 turns to mechanism. It extends the Autor et al. (2013) China-shock and Pierce and Schott (2016a) PNTR designs, both increases in competition, to productivity: both move employment, neither detects a significant productivity response. It shows the input tariff null under a variety of alternative plausible constructions. Section 9 weighs two readings of that missing response, namely changes in competition not sharp enough to register or an asymmetry between rising and falling competitive pressure, and offers an effort-ceiling conjecture for the second. Further robustness is in Appendix B.

## 2 Related literature

This paper connects five strands of work on trade and productivity.

**Output versus input tariffs and productivity.** The decomposition at the heart of my analysis, which separates the tariff an industry receives on its own output from the tariff it pays on its inputs, originates in the firm-level productivity literature. Amiti and Konings (2007) for Indonesia and Topalova and Khandelwal (2011) for India find that cutting *input* tariffs raises productivity by more than cutting output tariffs, a pattern they attribute to access to cheaper, higher-quality, and more varied imported inputs (Goldberg et al., 2010; Halpern et al., 2015). De Loecker et al. (2016) caution that such gains need not surface in measured productivity: in the same Indian reform, falling input tariffs lowered marginal costs, but firms raised markups and passed little of the decline into prices. I borrow the input-tariff construction of Amiti and Konings, an input-cost-weighted average of supplier output tariffs, but apply it to tariff *increases* in a developed economy and to industry total factor productivity, and find the opposite: the action is on the *output* tariff, while the input channel is absent.

**Trade costs, the input mix, and the productivity frontier.** A second, more foundational strand concerns *how* a tariff is presumed to cost the economy. In the standard trade model a tariff is a price phenomenon (Krugman et al., 2012, chs. 6 and 9). By raising the home price of the taxed good, it leads households and firms to substitute away from it, whether they use it as an input to consumption or as an input to production, and leads domestic producers to shift output toward it. The new equilibrium is inefficient for the world as a whole, and it changes welfare and real income. Crucially, this happens with technology

held fixed: total factor productivity is exogenous, and a tariff, like a terms-of-trade shock, has no first-order effect on it (Kehoe and Ruhl, 2008). In this view protection moves the economy *along* a fixed production frontier and distorts the mix of inputs, to consumption and to production alike; it does not move the production frontier itself.<sup>2</sup> My evidence both confirms and departs from this benchmark. The input mix does respond as the standard trade model predicts; the Pierce and Schott (2016a) quasi-experiment shows a trade shock reallocating inputs across factors (Section 8.2). However, I also find what the benchmark rules out by construction: as tariffs rise, TFP itself contracts. This is a movement *inside* the production possibility frontier, in which the same inputs yield less output. That consequence of tariffs lies outside the standard trade model, and it comes on top of the inefficiency the model does describe, a distorted mix of inputs at unchanged technology. Extensions with heterogeneous firms do contain such a channel, through reallocation across firms and innovation (Melitz and Trefler, 2012), but the evidence for it comes from liberalizations rather than from increases in protection.

**The 2018–19 trade war.** My setting is the recent U.S. tariff episode. The closest paper is the one by Flaaen and Pierce (2024), who decompose four-digit NAICS tariff exposure into import protection, input costs, and retaliation and show that the input-cost and retaliation channels overwhelm a small positive protection effect, lowering manufacturing employment and raising producer prices. I study the same tariffs and the same decomposition but for the level and dispersion of productivity rather than employment. The price pass-through that underlies the input-cost channel is documented by Amiti et al. (2019) and Fajgelbaum et al. (2020), the latter of which also supplies the Section 301 product lists I use.

**Protection, competition, and productivity.** Whether shielding producers from competition lowers productivity is an old and unsettled question. In the selection model of Melitz (2003), surveyed with its extensions by Melitz and Trefler (2012), lower trade costs raise industry productivity by shifting output toward more productive firms. Trade liberalization raised plant TFP among Chilean import-competing plants and reallocated output toward more productive plants (Pavcnik, 2002); in Canada, industry labor productivity rose, at least half of it through the exit and contraction of low-productivity plants (Trefler, 2004). Bloom et al. (2016) find that Chinese import competition raised patenting, IT intensity and TFP within European firms and shifted employment toward more technologically advanced firms.

---

<sup>2</sup>I take from Kehoe and Ruhl (2008) only this economic point: standard models treat a trade-cost shock as a price phenomenon leaving technology unchanged. Their paper’s main concern, the national-accounts measurement question of whether a terms-of-trade shock shows up in real-GDP-based TFP under chain weighting, is separate from my industry-level analysis and I do not engage it.

Yet [Autor et al. \(2020\)](#) show that the same competition *reduced* patenting among exposed U.S. firms, and [Pierce and Schott \(2016a\)](#) trace large U.S. manufacturing employment losses to reduced trade-policy uncertainty. My finding that the output tariff is associated with lower TFP, and with a compression of within-industry dispersion, speaks to this contested margin.

**Misallocation and productivity dispersion.** Finally, I read within-industry productivity dispersion through the misallocation lens of [Hsieh and Klenow \(2009\)](#) and [Syverson \(2011\)](#). [Kilumelume et al. \(2025\)](#) link import tariffs to capital misallocation and markup dispersion among South African firms. [Berthou et al. \(2021\)](#) give the link its theoretical form: in an open-economy model with heterogeneous firms and firm-specific distortions, trade shocks move measured aggregate productivity by reallocating activity across firms, with technology at each firm unchanged, and the direction of the move depends on the distortions already in place. Their European evidence is the mirror of mine, since import competition raises measured aggregate productivity there, and it also widens the dispersion of productivity across firms, while [Ding et al. \(2016\)](#) find compression in China. Relative to this work, my contribution is to connect trade policy to *both* the level of productivity and its within-industry dispersion, a link that, to my knowledge, has not previously been drawn for the U.S. trade war. The output tariff lowers the level and compresses the spread, while the input channel is absent.

### 3 Data

The analysis is built on a yearly panel indexed by four-digit NAICS manufacturing industry. The cross-section is the 86 manufacturing industries (NAICS 31–33) for which the establishment-dispersion statistics described below are published; the time dimension runs from 1989 to 2025, though individual series cover shorter windows, summarized in [Table 1](#). Two distinct objects are measured for each industry-year: the tariff protection the industry receives, and the productivity of the establishments within it. All series are placed on a single, longitudinally consistent 2012 NAICS basis, so that codes are comparable across the five-yearly NAICS revisions and across datasets. [Appendix A.2](#) sets out that harmonization, and [Appendix A](#) the construction and validation of every series.

Table 1: Data sources and coverage

| Series                                  | Concept                            | Source                                                                                                              | Coverage     |
|-----------------------------------------|------------------------------------|---------------------------------------------------------------------------------------------------------------------|--------------|
| Output & input tariff, effective        | Duties over customs value          | <a href="#">Schott (2025)</a>                                                                                       | 1989–2024    |
| Output & input tariff, statutory        | Schedule rate (MFN + Sec. 232/301) | <a href="#">Feenstra et al. (2002)</a> ; USITC DataWeb; <a href="#">Fajgelbaum et al. (2020)</a>                    | 1989–2023    |
| Input–output weights                    | 2012 benchmark detail use table    | BEA ( <a href="#">U.S. Bureau of Economic Analysis, n.d.</a> )                                                      | 2012 (fixed) |
| Total factor productivity               | Detailed-industry TFP index        | BLS ( <a href="#">U.S. Bureau of Labor Statistics, 2026</a> )                                                       | 1987–2023    |
| Labor productivity                      | Real sectoral output per hour      | BLS                                                                                                                 | 1987–2025    |
| Total factor productivity (cross-check) | Five-factor TFP index              | NBER–CES ( <a href="#">Becker et al., 2021</a> )                                                                    | 1958–2016    |
| Productivity dispersion                 | 90–10 establishment spread         | BLS–Census DiSP ( <a href="#">Cunningham et al., 2022</a> ; <a href="#">U.S. Bureau of Labor Statistics, 2024</a> ) | 1989–2021    |

*Notes:* All series are at the four-digit NAICS level for the 86 manufacturing industries (NAICS 31–33) and are placed on a longitudinally consistent 2012 NAICS basis. The effective and statutory tariffs are each built in an output (own-industry) and an input (cost-weighted) version; the input version uses the fixed BEA input–output weights shown. The statutory rate is the most-favored-nation rate plus the Section 232 and 301 add-ons; the Column 2 (non-NTR) rate is carried only for the NTR gap of Section 8.2. Construction and validation are detailed in Appendix A.7.

### 3.1 Tariff measures

I measure an industry’s *output tariff*, the tariff on the goods it sells, in two complementary ways: a **statutory** rate based on the tariff schedule, and the **realized** rate actually collected at the border. Every regression reported in the paper uses the statutory rate, so I describe its construction first and then define the realized rate, which the paper uses descriptively and in robustness checks. A third subsection (Appendix A.5) recasts them as the *input tariff*, the tariff an industry pays on the goods it buys. Section 6 draws on a different measure of exposure to the same tariffs, the share of an industry’s market covered by the 2018 actions, built by Flaaen and Pierce (2024).

**The statutory (schedule) rate.** Tariffs are specified by product lines, whereas productivity data is available not by product but by industry. To bridge this difference, I need to compute statutory tariff rates at the industry level.

I start from the legal ad valorem schedule rate: the most-favored-nation rate, listed in Column 1 of the U.S. Harmonized Tariff Schedule and applied to imports from countries with normal trade relations, plus the 2018–19 Section 232 and Section 301 add-ons, averaged across the eight-digit tariff lines in each industry. Two problems stand in the way. The first is the remapping just described, from product lines onto the industry classification. The second is that the 2018–19 duties were imposed by the executive under special authorities and fall on goods of a particular origin, entering the schedule through its Chapter 99 rather than through the rate on the line itself. I take the two in that order.

For industry  $j$  in year  $t$ ,

$$\tau_{jt}^O = \frac{\sum_{\ell \in L_j} \kappa_{\ell j} (m_{\ell t} + a_{\ell t}^{232} + a_{\ell t}^{301})}{\sum_{\ell \in L_j} \kappa_{\ell j}}, \quad (1)$$

where  $\tau_{jt}^O$  is the statutory output tariff of industry  $j$  in year  $t$ , in percentage points;  $\ell$  indexes tariff lines and  $L_j$  is the set of eight-digit lines mapped to industry  $j$ ;  $m_{\ell t}$  is the Column 1 rate carried by line  $\ell$  in year  $t$ ;  $a_{\ell t}^{232}$  and  $a_{\ell t}^{301}$  are the Section 232 and Section 301 add-ons on that line in force that year, zero before 2018 and weighted within the year an action takes effect by the days it applies; and  $\kappa_{\ell j}$  is the share of line  $\ell$  assigned to industry  $j$ , which equals one when the line maps to a single industry. The numerator therefore sums the total rate a line carries, weighted by the share of the line that belongs to the industry, and the denominator sums those shares, so  $\tau_{jt}^O$  is a weighted mean of line rates.

Tariffs are legislated at the tariff line, the eight-digit heading of the Harmonized Tariff Schedule at which a rate is written; the ten-digit codes beneath it carry statistical detail

but no separate rate. The Section 232 and 301 actions are specified the same way, as lists of lines, so the add-ons in (1) are line-level too. Productivity, by contrast, is reported for four-digit NAICS industries, so the two can be merged only after each line is assigned to an industry, and the set  $L_j$  is the result of that assignment. The industries the paper works with are therefore four-digit NAICS industries, and how many lines each contains varies widely: in 2019 the median industry draws on 82 tariff lines, with a range from 5 for forging and stamping to 1,314 for basic chemicals. Most lines belong to one industry alone, but a few straddle two or more, so the assignment has to say where a line goes and how it is divided.

The assignment itself exists at ten digits. The HS–NAICS concordance of [Pierce and Schott \(2012\)](#) was built for trade data, which are reported at ten digits, so it maps ten-digit import codes to NAICS industries, whereas the rate to be averaged belongs to the eight-digit line above them. I therefore carry the labels up to the line: an eight-digit line inherits the industries of its ten-digit children, and  $\kappa_{\ell j}$  records how those children divide across them, in proportion to their number. Averaging at ten digits instead would weight a line by how finely the statistical codes beneath it happen to be split, which is a reporting choice rather than a policy one. The weight matters only for the few lines that straddle industries. In 2019 only about one eight-digit line in forty straddles more than one four-digit industry, so for the rest  $\kappa_{\ell j} = 1$  and (1) is a plain mean over the industry’s lines. The line HS 8522.90.80, parts for sound- and video-recording apparatus, is one of the exceptions: its two ten-digit children fall one into audio and video equipment (NAICS 3343) and one into electronic components (NAICS 3344), so it enters each industry’s average with  $\kappa_{\ell j} = 1/2$ .

Every line enters with equal weight, so the measure averages the legal rates on an industry’s lines rather than the duties its imports attract, and an add-on that covers part of an industry raises  $\tau^O$  by the covered share times the add-on. Equation (1) is the construction of [Pierce and Schott \(2016a\)](#), who define an industry’s rate as the average across the eight-digit lines belonging to it, and it keeps the weights exogenous to trade. The same line average underlies the industry tariffs of the input-tariff literature: [Amiti and Konings \(2007\)](#) take “a simple average of the HS 9-digit codes within each 5-digit industry code”, and [Topalova and Khandelwal \(2011\)](#) average HS6 line rates within Indian industries. Because the average runs over the lines in a given year’s schedule, revisions to the schedule move it, a point Appendix A.3 documents along with the split weights  $\kappa_{\ell j}$ .

The weighting determines which lines enter the average, and the origin of the goods determines which rate each line carries. A tariff line is an origin-neutral product code, but the rate applied to it is not: the same line carries the Column 1 rate for partners with normal trade relations and the Column 2 rate for the few countries outside them. The 2018–19 duties are attached to it through separate Chapter 99 headings, so a listed line draws

the Section 301 add-on only when the goods are of Chinese origin, while the Section 232 add-on applies to covered lines from all origins except those granted country exemptions, which this build does not model. Each term in (1) is therefore a rate conditional on origin:  $m_{\ell t}$  is what any normal-trade-relations partner pays,  $a_{\ell t}^{232}$  what a covered origin pays, and  $a_{\ell t}^{301}$  what China pays. China has paid most-favored-nation rates throughout the sample, so from 2018 on the sum is the schedule rate facing Chinese-origin goods. The measure applies no origin weighting, so two industries whose lines are equally covered receive the same  $\tau^O$  however much of the product each buys from China.

Defining the rate as the Column 1 rate plus the two add-ons leaves every other trade barrier outside it. The measure covers all four Section 301 waves, Lists 1 through 3 and List 4A, and the Section 232 duties on steel and aluminum. It omits the Section 201 safeguards, the derivative steel and aluminum products added to the Section 232 action in February 2020, and antidumping and countervailing duties, and it does not net out exclusions granted to individual importers. What remains outside is either late or outside the schedule by construction. The derivative products arrive after the 2017–19 changes the design uses, so they move the level of the series and not the variation I exploit. Exclusions are granted to individual importers, whereas  $\tau^O$  is the rate a covered line carries; the realized rate does net them out, and it gives the same sign as the headline estimate (Section 5). The Column 2 rate is no part of  $\tau^O$ , and matters only for the NTR gap of Section 8.2. Schedule rates come from Feenstra et al. (2002) for 1989–2001 and the USITC DataWeb (U.S. International Trade Commission, 2025) for 2002–2023, with the product lists for Section 301 Lists 1 through 3 drawn from the replication data of Fajgelbaum et al. (2019) and the List 4A list from Annex A of Office of the United States Trade Representative (2019).

An origin-weighted alternative exists. Flaaen and Pierce (2024) measure exposure to the same actions as the share of an industry’s market covered by the newly tariffed products, valued at pre-war imports and crossed with the countries to which each action applied. That measure is weighted by trade and specific to origin where (1) is neither, and Section 6 puts it to work in an alternative regression design.

**The realized (effective) rate.** The realized rate brings origin in the other way, ex post, through the duties actually paid. It is collected duty as a share of the value imported,

$$\tau_{jt}^R = \frac{\sum_{i \in j} D_{it}}{\sum_{i \in j} V_{it}}, \quad (2)$$

where  $D_{it}$  and  $V_{it}$  are the calculated duties and customs value of the HS import lines  $i$  assigned to industry  $j$ , from the U.S. import database of Schott (2025) (1989–2024). It

averages over imports from every origin, so the Section 301 duties in it are diluted by imports from elsewhere, and it nets out duty-free entry under free-trade agreements and preference programs. On both counts it sits well below the schedule rate. Appendix A.4 gives the construction.

The two answer different questions: duties paid versus the policy on the books. Therefore, every regression reported in the paper uses the statutory rate, for the reason given in Section 4; the realized rate enters only descriptively, in Table 2, Figure 2(a), and the comparison of realized and statutory input tariffs in Section 8.3.

**The input tariff.** A tariff protects the industry that makes the taxed good but raises costs for the industries that buy it. I therefore pair each industry’s output tariff  $\tau_{jt}^O$  with an *input tariff* that summarizes the tariffs on the goods it purchases,

$$\tau_{jt}^I = \sum_c \omega_{cj} \tau_{ct}^O, \quad (3)$$

the input-cost-weighted average of supplying industries’ output tariffs, following [Amiti and Konings \(2007\)](#) and [Topalova and Khandelwal \(2011\)](#).

The weights  $\omega_{cj}$  are the share of commodity  $c$  in industry  $j$ ’s intermediate-input purchases, taken from the 2012 benchmark detail input–output table of the BEA ([U.S. Bureau of Economic Analysis, n.d.](#)) and *held fixed* so that all variation in  $\tau^I$  comes from tariffs rather than from a shifting input mix. The output tariff captures the *competition* channel (shelter from import rivalry in an industry’s own market), while the input tariff captures the *cost* channel (the tariffs an industry pays on its materials). Separating the two is the central measurement task of the paper, and both the statutory and realized measures are constructed in output and input versions.

## 3.2 Productivity measures

My primary productivity measure is the BLS detailed-industry total factor productivity (TFP) index ([U.S. Bureau of Labor Statistics, 2026](#)): sectoral output divided by a Törnqvist index of combined capital, labor, energy, and intermediate inputs, normalized to 2017=100 and taken in logs. It is published for exactly the same 86 four-digit NAICS manufacturing industries as the dispersion data, so the two merge with no concordance, and it spans 1987–2023. The same source supplies the intermediate-input, capital, and labor cost shares used in the mechanism test of Section 5. Like any productivity measure built from revenues deflated by an industry price index, it moves with markups as well as with technology, so a change in it need not be a change in physical productivity alone ([De Loecker, 2011](#); [De Loecker and](#)

Warzynski, 2012). As an alternative outcome with a longer recent reach I also use BLS labor productivity (real sectoral output per hour), available through 2025. Appendix A.6 lists the three productivity series and their coverage.

I validate the BLS series against the NBER–CES Manufacturing Industry Database (Becker et al., 2021), an independently constructed five-factor TFP index. It reaches back to 1958, far earlier than any other series used here, though its TFP index is populated only through 2016; on the overlapping industry-years the two series correlate at 0.77 in within-industry variation, and I report results on both.

The dispersion outcome is the within-industry spread of establishment productivity from the BLS–Census Dispersion Statistics on Productivity (DiSP) of Cunningham et al. (2022); U.S. Bureau of Labor Statistics (2024). For each industry-year I use the 90–10 log total-factor-productivity gap,  $d9010$ , available for the 86 industries over 1989–2021. It measures how much more productive a 90th-percentile establishment is than a 10th-percentile one in the same industry. DiSP also publishes the analogous labor-productivity dispersion, which I use as a contrast in Appendix B.2. Where the TFP indexes measure the *level* of average productivity, DiSP measures its *dispersion* across establishments; the paper asks how tariffs move both.


### 3.3 Descriptive statistics

Table 2 provides summary statistics for the panel used in my analysis. The statutory output tariff averages 5.9 percentage points and the input tariff 5.3; both are right-skewed, and the effective (realized) rates are roughly half as large. The 90–10 productivity spread averages just over one log point: a 90th-percentile establishment is about twice as productive as a 10th-percentile one within the same industry.

Figure 2(a) plots the cross-industry mean of each tariff series. Statutory tariffs drift gently downward over the liberalization era and then jump sharply in 2018–19: the Section 301 and 232 actions raise the mean statutory output tariff from about 3.6 percent in 2017 to 20 percent in 2020, a rate facing Chinese goods, since the Section 301 add-ons apply to China alone. The realized rates, which average over imports from all origins, rise far less, from about 2.0 to 4.4 percent. After 2021 both level off, high but no longer rising.

However, the flat average hides important underlying variation. Figure 2(b) shows the cross-industry distribution of the year-over-year *change* in the statutory output tariff. Even in the liberalization era the rates churn: broad MFN reductions move most industries’ rates in the mid-1990s and again around 2016, interspersed with quiet years, before the 2018–19 actions raise them across the board. It is this within- and across-industry, and over-time

Figure 2: Output and input tariffs, 1989–2024.



*Notes:* Panel (a): cross-industry mean of each tariff series. Statutory rates (solid, through 2023) are schedule rates: the most-favored-nation rate plus the Section 232 and 301 add-ons, which from 2018 makes them the rates facing Chinese goods. Realized rates (dashed, through 2024) are duties collected over customs value on imports from all origins. The statutory series dips after 2021 because the 2022 Harmonized System revision adds tariff lines to the average while the Section 301 lists are held at their 2018–19 codes, not because policy changed (Appendix A.3). Panel (b): the median and the 10th–90th-percentile range, across industries, of the year-over-year change in the statutory output tariff, on a symmetric-log scale. Panel (c): statutory output-tariff paths for four example industries (3311 iron and steel mills, 3361 motor vehicles, 3162 footwear, and 3254 pharmaceuticals); the dotted vertical lines mark the Uruguay Round phase-ins, the 2002 Section (Sec.) 201 steel safeguard (which the statutory rate does not include), the 2016 Information Technology Agreement (ITA) expansion, and the 2018 Section 232 and 301 actions. Tariff rates are in percentage points (pp); the four-digit codes are North American Industry Classification System (NAICS) industries.

Table 2: Summary statistics, NAICS-4 manufacturing, 1989–2021

|                                             | N     | Mean   | SD     | p10    | Median | p90     |
|---------------------------------------------|-------|--------|--------|--------|--------|---------|
| <i>Levels</i>                               |       |        |        |        |        |         |
| Output tariff $\tau^O$ (statutory, pp)      | 2,805 | 5.999  | 7.781  | 0.718  | 3.644  | 15.445  |
| Input tariff $\tau^I$ (statutory, pp)       | 2,805 | 5.391  | 5.650  | 1.982  | 3.661  | 10.170  |
| Output tariff $\tau^O$ (effective, pp)      | 2,805 | 2.679  | 3.142  | 0.087  | 1.622  | 6.379   |
| Input tariff $\tau^I$ (effective, pp)       | 2,805 | 2.194  | 1.499  | 0.917  | 1.655  | 4.262   |
| 90–10 productivity dispersion (log pts)     | 2,805 | 1.060  | 0.401  | 0.690  | 0.976  | 1.512   |
| Total factor productivity (index, 2017=100) | 2,805 | 96.002 | 14.128 | 81.475 | 97.646 | 109.706 |
| Labor productivity (index, 2017=100)        | 2,805 | 91.811 | 24.109 | 62.350 | 94.359 | 114.073 |
| <i>First differences (year over year)</i>   |       |        |        |        |        |         |
| $\Delta$ output tariff (pp)                 | 2,720 | 0.516  | 2.683  | -0.374 | 0.000  | 0.781   |
| $\Delta$ input tariff (pp)                  | 2,720 | 0.574  | 2.444  | -0.307 | -0.001 | 0.670   |
| $\Delta$ 90–10 dispersion (log pts)         | 2,720 | 0.008  | 0.248  | -0.229 | 0.008  | 0.257   |
| TFP growth ( $\Delta \ln$ )                 | 2,720 | 0.007  | 0.051  | -0.047 | 0.006  | 0.061   |
| Labor-productivity growth ( $\Delta \ln$ )  | 2,720 | 0.015  | 0.077  | -0.062 | 0.015  | 0.092   |

*Notes:* Pooled over 86 four-digit NAICS manufacturing industries and 1989–2021 (industry-years with a defined tariff). Tariffs in percentage points. Statutory rates are the most-favored-nation rate plus the Section 232 and 301 add-ons (from 2018, the rates facing Chinese goods); effective rates are duties collected over customs value on imports from all origins; dispersion is the 90th-minus-10th-percentile spread of within-industry log TFP; productivity indexes are 2017=100.


variation that provides information for my empirical analysis.

Figure 2(c) makes the variation concrete by focusing on four industries: steel’s rate falls through the liberalization era and then leaps with the Section 232 action, motor vehicles jump with Section 301 and partly reverse by 2023, footwear stays persistently high yet untouched by the trade war, and pharmaceuticals sit near zero throughout.

Figure 3 shows the productivity side: On average, total factor and labor productivity both rise and then plateau after about 2010 (labor productivity faster, reflecting capital deepening), while the within-industry dispersion of productivity widens secularly, from about 0.95 to 1.2 log points.


Finally, Figure 4 plots the output against the input tariff across industries in a representative pre-war year and a trade-war year. The trade war shifts the entire input-tariff distribution upward, placing most industries above the 45-degree line; within a given year the two tariffs are only weakly correlated across industries, which helps separate the two channels in the analysis that follows. The seams between the datasets, and the checks I run on the merged panel, are in Appendix A.7.

Figure 3: Productivity and its dispersion, 1989–2025.



*Notes:* Cross-industry mean of the total factor productivity index (through 2023) and the labor productivity index (through 2025), 2017=100 (left), and the mean 90–10 within-industry dispersion of log TFP (through 2021, right).

Figure 4: Output versus input tariff across industries, 2010 and 2019.



*Notes:* Each point is a four-digit NAICS manufacturing industry: the statutory output tariff against the statutory input tariff, in a pre-war year (2010) and a trade-war year (2019). The dotted line is the 45-degree line.

## 4 Two channels, and how to identify them

A tariff can lower an industry’s productivity through two distinct channels, and the empirical work that follows is organized around separating them. This section defines the two, explains why a straightforward regression cannot identify either, and lays out the three designs the paper uses instead. The first two designs identify the effect on the *level* of productivity from the 2018–19 trade war, by an instrument and then without one; the third identifies the effect on the within-industry *distribution* from the full panel.

**The two channels.** The *output tariff*  $\tau_{jt}^O$  is the tariff levied on industry  $j$ ’s own product; under the competition hypothesis (H1) it shields domestic producers, dulling the pressure to improve. The *input tariff*  $\tau_{jt}^I = \sum_c \omega_{cj} \tau_{ct}^O$  is the input-cost-weighted output tariff of  $j$ ’s supplying industries, with weights  $\omega_{cj}$  fixed at their 2012 input–output shares so that all of its time variation comes from tariffs rather than from the mix shift H2 predicts; under the input-cost hypothesis (H2) it raises the price of  $j$ ’s inputs and distorts its production mix. In every regression the paper reports, both are built from statutory (schedule) rates, and  $\tau^O$  and  $\tau^I$  denote those statutory measures hereafter. The schedule rate is set by legislation and trade actions rather than by trade outcomes, whereas the realized rate (duties collected over customs value) has an import-composition denominator that responds to the tariff itself.

**Why an association is not enough.** Protection is not assigned at random. It is granted for reasons correlated with an industry’s prospects, and the data show it: replacing the causal lag of the tariff change with its *leads* (Appendix Table A.9), the leads predict the current change in TFP. A first-difference regression of productivity on the lagged tariff change therefore fails its own timing placebo: future tariffs “predict” today’s productivity, the signature of protected industries already on distinct paths. For that reason the paper does not report the uninstrumented level association as a result. It needs variation in tariff exposure that is orthogonal to those paths.

**Three designs.** The paper supplies that variation three ways. Section 5 exploits the assignment rule of the 2018–19 trade war itself: U.S. authorities delayed tariffs on consumer goods, so an industry’s consumer-good share predicts its exposure for reasons unrelated to its productivity, and serves as an instrument in a cross-section of long differences. Section 6 drops the instrument and borrows the differential-exposure design of Flaaen and Pierce (2024), which identifies from exposure differences together with fixed effects, exposure-by-year controls, and an explicit pre-trend correction. Section 7 turns to the within-industry distribution, where the object of interest is a percentile spread rather than a level; there

the first-difference panel regression *is* read as a result, because, unlike its level counterpart, it passes the timing placebo. The table notes report the estimation years of every panel regression, meaning the years  $t$  of the change in the outcome; because the regressor is the tariff change lagged one year and the tariff series start in 1989, the lagged specifications are estimated from 1991.

## 5 Tariffs and the level of productivity

Protection lowered the productivity of the industries it shielded, and the way the 2018–19 tariffs were assigned identifies by how much. This section delivers that estimate, the effect of the output tariff on the level of industry productivity, identified from the assignment rather than from who ended up protected. The instrument is the industry’s consumer-good share; the finding is that across the 6.8-percentage-point gap between the quartiles of the tariff increase, an industry at the top quartile ends 2019 with TFP about 3.6 percent below one at the bottom. The rest of the section states the instrument, tests its exclusion restriction, and checks that foreign retaliation does not confound it.

**A within-war instrument: the consumer-good share.** The 2018–19 trade war supplies it, through how the tariffs were assigned. To limit the pass-through to consumer prices, the U.S. Trade Representative front-loaded intermediate and capital goods (Section 301 Lists 1–3, imposed in 2018) and delayed consumer goods to List 4 in 2019 (Fajgelbaum et al., 2020; Cavallo et al., 2021; Bown, 2021). An industry’s *consumer-good share*, the fraction of output sold to final consumers (personal consumption expenditures over total commodity output, from the BEA 2012 detail Use table), therefore predicts its 301 exposure for reasons tied to consumer prices rather than to productivity. The design follows Pierce and Schott (2016a) in using a *pre-determined* industry characteristic to predict policy exposure: there the non-NTR rates fixed by the Smoot–Hawley Act of 1930, here the consumer-good share. My exogeneity claim, however, rests on the stated motive for the sequencing, the containment of consumer-price increases, rather than on seventy years’ distance. The prediction is strong: industries in the bottom half of the consumer-good-share distribution saw their statutory output tariff rise by about 18 points between 2017 and 2019, against about 13 points for the top half (first-stage  $F \approx 28$ ). Figure 5 shows the first stage, over 2017–19, and the reduced form, over 2017–2021.


I instrument the 2017–19 change in the statutory output tariff (the most-favored-nation rate plus the Section 301 and 232 add-ons) with the consumer-good share.<sup>3</sup> Table 3 reports

---

<sup>3</sup>The realized rate cannot stand in for the statutory rate in this design. Instrumented by the consumer-

Figure 5: The consumer-good-share instrument for the 2018–19 tariffs.

The consumer-good-share instrument for the 2018–19 tariffs



*Notes:* Across the 79 four-digit NAICS manufacturing industries with a defined consumer-good share. Left, first stage: the change in the statutory output tariff over 2017–2019 (the Section 301/232 increase) against the consumer-good share. Right, reduced form: the change in log BLS TFP over 2017–2021 against the consumer-good share. Red lines are OLS fits. The consumer-good share is personal consumption expenditures over total commodity output, from the BEA 2012 detail Use table.

the two-stage least-squares estimates over outcome windows ending in 2019 through 2022. I anchor on the clean *pre-pandemic* window, which measures the change in log TFP from 2017 to 2019. There the instrumented effect is  $-0.0053$  ( $p = 0.008$ ): a one-point increase in the output tariff lowers TFP by about 0.53 percent, with a 95 percent confidence interval from 0.14 to 0.91 percent. The interval is wide, but even its lower end is economically large. The OLS estimate, 0.16 percent per point, sits at the bottom edge of that interval. The instrumented estimate is a local average treatment effect for the consumer-good-share compliers. Intuitively, it compares industries whose exposure was settled by the sequencing of the lists rather than by their own prospects, so it speaks for the industries that the

good share, the 2017–19 change in the realized output tariff has a weak first stage ( $F = 2.9$ , against 27.8 for the statutory change). Section 301 duties apply only to goods from China, so the realized change depends on how much an industry imports from China: across the 79 industries it correlates at 0.61 with the statutory change weighted by the industry’s 2017 share of imports from China, computed from the [Schott \(2025\)](#) import data, against 0.37 with the statutory change alone. Weighting the statutory change by that import share does not rescue the design. The consumer-good share does not predict the weighted change ( $F = 0.01$ ), nor the version that weights only the Section 301 component, Section 232 being owed on goods of every origin ( $F = 0.87$ ). Intuitively, the industries whose products appeared on the later lists are also the ones that buy more from China, and the two effects offset in the product. What the weighting does not change is the association itself: ordinary least squares gives a fall of 1.05 percent of TFP at an interquartile move of the unweighted change, against 0.96 percent for the weighted one. The Anderson–Rubin 95 percent confidence set for the realized-rate coefficient excludes zero, and every value between  $-0.011$  and  $0.198$ , but is unbounded; ordinary least squares on the realized change gives the same sign as the headline estimate ( $-0.0042$  in the 2019 window). Origin is therefore a question for a different design, and Section 6 takes one up whose Protection measure counts Section 301 flows from China alone.

sequencing actually moved.

Although the regression is in changes, I can recast it to deliver a comparison of levels. The BLS index is normalized to 100 in 2017, so the change in log TFP from 2017 to 2019 is the log level of TFP in 2019, measured against each industry’s own 2017 base. Because the Section 301 and 232 add-ons were zero in 2017, the 2017–19 change in the statutory tariff is, up to small movements in the most-favored-nation rate, the level of protection that the war added. The coefficient therefore maps a gap in tariff levels into a gap in TFP levels. To size that gap, I compare two industries at the quartiles of the tariff increase, whose interquartile range is 6.8 points (12.7 at the 25th percentile, 19.4 at the 75th). By 2019 the industry at the top quartile has TFP about 3.6 percent below the one at the bottom, relative to where each stood in 2017, with a 95 percent confidence interval from 0.9 to 6.2 percent. Figure 6 shows the estimate, its interval, and the OLS benchmark on this scale.

This is a relative comparison between two industries inside the sample, not an aggregate effect. The constant absorbs whatever hit all of manufacturing, and the estimate carries no general-equilibrium correction. Nor should it be rescaled to the full 13–18 point exposure, which would extrapolate a local average treatment effect beyond its support. Section 6 offers a comparison from a design that needs no instrument. Applied to the same 6.8-point gap, that design places the top-quartile industry 2.4 percent below the bottom one by 2019, with a 95 percent confidence interval from 0.1 to 4.6 percent, inside the interval of the instrumented estimate.

Extending the outcome window to 2020, 2021, and 2022, the estimate grows steadily more negative, to  $-0.0091$  by 2021. I read that growth cautiously: those later windows overlap the COVID-19 pandemic, which struck consumer- and intermediate-good industries very differently, so the building profile blends any dynamic effect of the tariff with differential pandemic exposure. The shape is suggestive of a genuine dynamic, but I do not rest on it.

**Exclusion.** The estimate’s credibility rests on the consumer-good share affecting productivity only through the tariff. The concern is that consumer-good and intermediate-good industries may differ in ways that move productivity regardless. The natural test asks whether the consumer-good share predicts *pre-war* productivity changes. It does not: the 2013–2017 change in TFP is unrelated to the instrument ( $t = 0.1$ ), so the design does not pick up a consumer-versus-intermediate trend. I cannot, however, corroborate the estimate with the NBER–CES TFP series, which ends in 2016 and so does not reach the war.


**Retaliation.** The tariff war entailed a second shock alongside higher U.S. tariffs. Foreign governments retaliated against U.S. exports, and that retaliation depressed activity on its

Table 3: The 2018–19 tariffs and TFP: the consumer-good-share instrument

|                                      | Change in log BLS TFP, 2017 to: |                       |                       |                       |
|--------------------------------------|---------------------------------|-----------------------|-----------------------|-----------------------|
|                                      | (1)<br>2019                     | (2)<br>2020           | (3)<br>2021           | (4)<br>2022           |
| $\Delta$ statutory output tariff, IV | -0.0053***<br>(0.0020)          | -0.0064**<br>(0.0030) | -0.0091**<br>(0.0037) | -0.0091**<br>(0.0042) |
| same, OLS                            | -0.0016                         | -0.0023               | -0.0031               | -0.0034               |
| First-stage $F$                      | 27.8                            | 27.8                  | 27.8                  | 27.8                  |
| Observations                         | 79                              | 79                    | 79                    | 79                    |

*Notes:* Cross-sectional two-stage least-squares estimates across the 79 four-digit NAICS manufacturing industries with a defined consumer-good share. The endogenous regressor is the change in the statutory output tariff over 2017–2019 (the Section 301/232 increase); the instrument is the industry’s consumer-good share (personal consumption expenditures over total commodity output, BEA 2012 detail Use table). Each column is the change in log BLS TFP from 2017 to the indicated year; the 2019 column is the clean pre-pandemic window on which I anchor the magnitude, while the 2020–2022 windows overlap the COVID-19 pandemic and are read cautiously. “Same, OLS” is the corresponding ordinary-least-squares coefficient. The consumer-good share does not predict the 2013–2017 change in TFP ( $t = 0.13$ ), and the implied effect on the 90–10 dispersion spread is insignificant ( $t = 1.3$ ). The 2019 estimate is unchanged ( $-0.0056$ ) when retaliatory-tariff exposure (Fajgelbaum et al., 2020) is added as a control, and the instrument is uncorrelated with that exposure. Robust standard errors. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Figure 6: The magnitude: the TFP gap between the top and bottom quartile of exposure, 2019.



*Notes:* The 2019-window estimates of Table 3, each multiplied by the interquartile range of the 2017–19 increase in the statutory output tariff (6.8 points, across the 79 industries of the instrumented sample). The axis reads as the percent by which TFP in 2019, measured against each industry’s 2017 level, is lower in an industry at the 75th percentile of the increase than in one at the 25th. Filled marker and bar: the instrumented estimate and its 95 percent confidence interval, from robust standard errors. Hollow marker: the corresponding OLS estimate.

own (Flaaen and Pierce, 2024). My exclusion restriction would fail if the industries the consumer-good share sorts were also the industries retaliation hit. Retaliation does not run through my instrument. An industry’s retaliatory-tariff exposure is the trade-weighted retaliatory tariff that foreign governments placed on its exports, from the Fajgelbaum et al. (2020) data. It is essentially uncorrelated with the consumer-good share (correlation 0.09) and with the 2017–19 change in its statutory output tariff (correlation 0.05): retaliation fell on agricultural and export-oriented sectors, the 301 tariffs on import-competing manufacturers. Furthermore, adding retaliatory exposure as a control leaves the instrumented output-tariff effect unchanged ( $-0.0056$  in the 2019 window), while retaliation carries its own, separable, and also negative effect on TFP ( $-0.0046$ , significant); it is a distinct channel rather than a confound.

The level decline is thus identified, sharply and with its magnitude stated as a contrast between industries rather than as an aggregate.<sup>4</sup> It rests on one exclusion restriction: that the consumer-good share moves TFP only through tariff exposure. The next section removes it, identifying the same effect from a design that needs no instrument at all.

## 6 A second design: identification without an instrument

The estimate of Section 5 rests on one exclusion restriction. That restriction is testable and it passes its natural test, but it cannot be proved. This section identifies the same effect a second time, using a design that requires no instrument at all: the differential-exposure framework of Flaaen and Pierce (2024), who study the same tariffs at the same level of aggregation.

Two features make their design worth borrowing rather than merely citing. First, it identifies from exposure differences across industries together with fixed effects and pre-trend removal, so it fails in ways unrelated to the failure modes of an instrument. Second, it measures exposure differently (as the *share* of an industry’s market or costs covered by the new tariffs, not as an *ad valorem* rate), so agreement between the two designs cannot be an artifact of a shared measure. The share also respects the origin of imports: it counts Section 301 flows from China alone, whereas the statutory rate applies a listed product’s add-on whatever the country it comes from, and so may overstate the protection of industries that buy little from China.

---

<sup>4</sup>Appendix B.1 reports the full timing placebo and its robustness, the value-added-weighted estimates, and the same specification run on the independently constructed NBER–CES series and on labor productivity as an alternative outcome.

Flaen and Pierce (2024) do not study productivity. Their outcomes are employment, industrial production and the producer price index; no productivity measure appears in their paper. I borrow their design and put it to a question their paper does not address.

paragraphExposure as a share of the market. Let  $\Omega$  denote the set of product–country pairs newly covered by a 2018 U.S. tariff action and  $\Omega^*$  the pairs hit by foreign retaliation. Writing  $M$ ,  $X$  and  $Q$  for an industry’s imports, exports and own shipments, all fixed at their 2017 (pre-war) values,

$$\begin{aligned}\text{Protection}_j &= \frac{\sum_{pc \in \Omega} M_{j,pc}}{Q_j + M_j - X_j}, \\ \text{Retaliation}_j &= \frac{\sum_{pc \in \Omega^*} X_{j,pc}}{Q_j}, \\ \text{InputCost}_j &= \sum_c \omega_{cj} \frac{\sum_{pc \in \Omega} M_{c,pc}}{Q_c + M_c},\end{aligned}\tag{4}$$

where  $\omega_{cj}$  is commodity  $c$ ’s share of industry  $j$ ’s costs from the BEA 2012 detail Use table. The denominator of Protection,  $Q_j + M_j - X_j$ , is domestic absorption, the size of the home market. Exports are inside  $Q$ , so subtracting them removes the part of domestic output that left the country rather than a duplicate of imports.

The three measures use different denominators, each matched to the channel it captures. Protection scales newly tariffed imports by absorption, the size of the home market in which the industry competes. Retaliation scales newly tariffed exports by output, since an export channel is measured against what an industry makes rather than against what its home market absorbs. The input term scales newly tariffed imports of each commodity by its domestic supply,  $Q_c + M_c$ , so that it measures the share of the commodity’s total supply that is imported and newly tariffed. Appendix A.8 documents the rebuild of these measures and validates them against the published aggregates.

**Specification.** The annual counterpart of Flaen and Pierce’s equation (7) regresses the outcome on each exposure measure interacted with a full set of year dummies. Because every exposure measure is time-invariant, a year-subscripted coefficient on a fixed regressor *is* that full set of interactions, so for industry  $j$  in year  $t$ ,

$$\begin{aligned}\tilde{y}_{jt} &= \gamma_t \text{Protection}_j + \theta_t \text{InputCost}_j + \lambda_t \text{Retaliation}_j \\ &\quad + \omega_t \text{ImpShare}_j + \varphi_t \text{ExpShare}_j + \delta_j + \delta_t + \varepsilon_{jt},\end{aligned}\tag{5}$$

estimated on 85 industries over 2013–2022 with standard errors clustered by industry. Each coefficient is therefore a full path of year-specific slopes, normalized to zero in the 2017 base

year. The outcome  $\tilde{y}_{jt}$  is log TFP with each industry’s own 2013–17 linear trend removed; this detrending is [Flaen and Pierce](#)’s first pre-trend correction.

The term that this paper’s own specification lacks is the second line.  $\text{ImpShare}_j$  and  $\text{ExpShare}_j$  are the industry’s pre-war import share of absorption and export share of shipments, each interacted with *every* year dummy. A time fixed effect absorbs an aggregate shock only if that shock hits every industry equally; these terms let trade-exposed industries follow their own arbitrary path, so a dollar movement or a foreign demand slowdown that loads on exactly the industries that were protected is absorbed non-parametrically rather than attributed to the tariff.

This is a design without an instrument, and it is being introduced in the same section that declined to report an uninstrumented association. The two are not the same object. The association of Section 5 relates a one-year change to a one-year change with no control for differential trends, and its leads predict the outcome. Equation (5) carries industry and year fixed effects, the exposure-by-year controls, and an explicit pre-trend removal. The flat pre-war coefficients reported below are what that removal imposes rather than a test the design passes, since the trend is fitted on the very years those coefficients cover. The identification therefore rests on [Flaen and Pierce](#)’s equation (8), which measures the post-war separation against the pre-war drift. Appendix B.1.4 reports the same specification with no trend removed, so the pre-war paths can be read directly.

**Result.** Figure 7 plots the year-by-year coefficient on Protection, scaled so the vertical axis reads as the TFP gap between industries at the top and bottom quartile of exposure. The left panel uses [Flaen and Pierce](#)’s coverage-share measure of Protection; the right substitutes the 2017–19 change in the statutory output tariff, the regressor instrumented in Section 5, into the same specification. In both panels InputCost and Retaliation remain coverage shares. Both are flat through 2018 (no coefficient before the tariffs reaches  $|t| = 1.1$ ), and both turn down from 2019.

Table 4 reports the path. On the coverage-share measure the 2019 coefficient is  $-0.192$  ( $t = -2.53$ ), growing to  $-0.335$  by 2021 and  $-0.440$  by 2022; netting out the pre-war path in the manner of [Flaen and Pierce](#)’s equation (8) gives a cumulative  $-0.252$  ( $t = -2.82$ ) at the 2019 endpoint. Intuitively, that statistic asks how far the more and less exposed industries have separated since 2017 once the drift they were already on is subtracted. The statutory-rate column tracks it closely.


**The two designs bracket the magnitude.** The natural comparison is in a common unit. The interquartile gap in the 2017–19 increase in the statutory output tariff is 6.8 percentage

Table 4: The Flaaen–Pierce design: coefficient path on the Protection channel

|              | Detrended log TFP      |                      |
|--------------|------------------------|----------------------|
|              | (1)                    | (2)                  |
|              | Share of market        | Statutory rate       |
| 2013         | 0.0183<br>(0.0327)     | -0.0228<br>(0.0385)  |
| 2014         | -0.0460<br>(0.0553)    | -0.0747<br>(0.0972)  |
| 2015         | -0.0599<br>(0.0561)    | -0.1334<br>(0.1031)  |
| 2016         | -0.0095<br>(0.0548)    | -0.1203<br>(0.0892)  |
| 2018         | -0.0445<br>(0.0455)    | -0.0482<br>(0.0772)  |
| 2019         | -0.1923**<br>(0.0760)  | -0.2164*<br>(0.1231) |
| 2020         | -0.1938*<br>(0.1133)   | -0.1978<br>(0.1760)  |
| 2021         | -0.3353**<br>(0.1293)  | -0.3817<br>(0.2390)  |
| 2022         | -0.4396***<br>(0.1602) | -0.4938*<br>(0.2827) |
| Observations | 850                    | 850                  |

*Notes:* Coefficients  $\gamma_t$  from equation (5), 2017 omitted as the base year. Column 1 measures Protection as the share of the home market newly covered by a 2018 tariff action; column 2 substitutes the change in the statutory output tariff over 2017–19 into the same specification. Outcome: log BLS TFP detrended on each industry’s 2013–17 path; Appendix Table A.14 reports the same specification with no trend removed, where the pre-war paths can be read directly. The equation (8) statistic is symmetric about the base year and is therefore identical in the two. InputCost and Retaliation, measured as coverage shares in both columns, and the two baseline exposure controls are included and interacted with every year dummy. Industry and year fixed effects; SEs clustered by industry. Sample: 85 industries, 2013–2022. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Figure 7: The Flaaen–Pierce design: TFP by exposure, relative to 2017.



*Notes:* Year-by-year coefficients on the Protection channel from equation (5), with InputCost and Retaliation held fixed; 90 percent bands. The outcome is log BLS TFP detrended on each industry’s own 2013–17 path, normalized to zero in 2017. Coefficients are scaled by each measure’s own interquartile move, so the axis reads as the gap between an industry at the 75th and one at the 25th percentile of exposure: 8.4 points of home market in the left panel, 6.8 points of statutory output tariff in the right. The dotted line marks the first 2018 tariffs. The 2020–22 window overlaps the COVID-19 pandemic. Sample: 85 industries, 2013–2022.

points: the 25th percentile industry saw its tariff rise 12.7 points, the 75th 19.4. Applying each design’s coefficient to that same gap puts all three on one scale, as Table 5 does.

Table 5: The same object from three designs

|                        | $\Delta \log \text{TFP per pp}$ | $t$   | % of TFP, quartile gap | 95% CI     |
|------------------------|---------------------------------|-------|------------------------|------------|
| Differenced panel, OLS | -0.0012                         | -3.63 | 0.8                    | [0.4, 1.3] |
| Flaaen–Pierce design   | -0.0035                         | -2.09 | 2.4                    | [0.1, 4.6] |
| Consumer-share IV      | -0.0053                         | -2.66 | 3.6                    | [0.9, 6.2] |

*Notes:* Effect of the statutory output tariff on log BLS TFP, per percentage point and scaled to the interquartile gap of 6.8 points in its 2017–19 increase. Row 1 regresses the change in log TFP on the once-lagged change in the statutory output tariff in the annual panel, with industry and year fixed effects; row 2 is equation (5) with the 2017–19 change in the statutory output tariff as the Protection measure, cumulated in the manner of Flaaen and Pierce’s equation (8); row 3 instruments the 2017–19 change in the statutory output tariff with the consumer-good share (Table 3, 2019 window). The last column is the 95 percent confidence interval for the quartile gap, from the normal distribution for row 3 (robust standard errors) and from the  $t$  distribution for rows 1 and 2 (standard errors clustered by industry). Samples: row 1, the annual panel over 1991–2023; row 2, 85 industries over 2013–2022; row 3, the 2017–19 changes across 79 industries.

The instrumented estimate is the largest, the uninstrumented panel association the smallest, and the Flaaen–Pierce design falls between them and closer to the instrument. That ordering is what one expects if protection is directed at industries already deteriorating: the association understates the damage, and correcting for the direction of assignment (whether by an instrument or by a design that nets out differential trends and exposure) moves the

estimate away from zero. Both designs deliver the same kind of statement: a gap in 2019 TFP levels, measured against 2017, between industries at the top and bottom quartile of exposure. Two designs with different exposure measures, different samples and different identifying assumptions place that gap between 2 and 4 percent of TFP: 3.6 percent with the instrument (95 percent confidence interval 0.9 to 6.2) and 2.4 percent without one (0.1 to 4.6). The two intervals overlap, and neither includes zero. What separates them is the design: substituting the statutory output tariff for the coverage share in equation (5) changes the answer by about three tenths of a percentage point, while changing the design moves it by more than a percentage point.

Two cautions remain. The post-2019 portion of the path overlaps the pandemic, for the same reason the instrumented estimate is anchored on 2019. And the InputCost channel in equation (5) is poorly behaved before the war: its pre-war coefficients are significant even after detrending. Its null here should therefore not be read as independent evidence on the input-cost hypothesis; Section 8.3 tests that directly.

The level decline is now established twice over, by an instrument and without one, with its magnitude bracketed between 2 and 4 percent of TFP for an interquartile move in exposure. The question turns from how much to where: inside a protected industry, which plants bear the decline?

## 7 The decline across the productivity distribution

Sections 5 and 6 establish that the output tariff lowers the average level of industry productivity. This section asks where inside the industry that decline occurs, using a different object and a different design: the outcome is a percentile spread of plant-level productivity from DiSP, and the design is a first-difference regression on the full 1989–2021 panel rather than the trade war alone. The finding is that the within-industry spread *compresses*, that the compression sits between the quartiles rather than in either tail, that it survives dropping the war years entirely, and that this specification, unlike the level association, passes its timing placebo. What the finding does not license is a claim about which plants moved, and the section is explicit about that limit.

### 7.1 The compression

Protection compresses the spread of productivity within an industry, and the first-difference panel regression foreshadowed in Section 4 is what establishes it. Differencing removes each industry’s level and leaves an industry-specific drift (the industry fixed effect in the

differenced equation), so each industry is detrended by its own trend; year fixed effects  $\delta_t$  absorb common shocks. For a percentile spread  $Y_{jt}$  of within-industry productivity,

$$\Delta Y_{jt} = \alpha_j + \delta_t + \beta^O \Delta \tau_{j,t-1}^O + \beta^I \Delta \tau_{j,t-1}^I + \varepsilon_{jt}, \quad (6)$$

where  $\tau^O$  and  $\tau^I$  are the statutory output and input tariffs. I use the once-lagged change because contemporaneous changes are uninformative (tariffs act with a delay). Standard errors are clustered by industry. Intuitively, the specification asks whether a change in protection is followed by a change in the spread inside the same industry, not whether more protected industries carry wider spreads.

Table 6 relates the change in the 90–10 spread of within-industry log TFP to the lagged change in the statutory output and input tariffs. The output tariff *compresses* the spread (−0.0058, significant); the input tariff has no detectable effect. The coefficient is essentially unchanged when I weight industries by value added, and holds when the statutory output tariff is instrumented with its Section 301/232 add-on alone (Appendix Table A.15).

Table 6: Statutory output and input tariffs and the change in 90–10 dispersion

|                                         | (1)<br>Output         | (2)<br>Input        | (3)<br>Both           |
|-----------------------------------------|-----------------------|---------------------|-----------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | −0.0058**<br>(0.0025) |                     | −0.0059**<br>(0.0024) |
| $\Delta$ input tariff $\tau^I$ , lag 1  |                       | −0.0009<br>(0.0099) | 0.0013<br>(0.0095)    |
| Observations                            | 2,635                 | 2,635               | 2,635                 |
| Within $R^2$                            | 0.038                 | 0.037               | 0.038                 |
| Industry FE                             | Yes                   | Yes                 | Yes                   |
| Year FE                                 | Yes                   | Yes                 | Yes                   |

*Notes:* First-differences regressions of the change in the 90–10 dispersion spread on the once-lagged change in the indicated tariff, with industry and year fixed effects; standard errors clustered by industry. The input tariff  $\tau^I$  is the 2012-I-O input-cost-weighted statutory rate of supplying industries; the output tariff is the own-industry statutory rate (both in percentage points). Estimation years: 1991–2021.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Because the level results are identified from the war years, it is worth showing directly that the compression is not. Re-estimated on the years through 2017, which removes all Section 301 and 232 variation, the specification gives the output tariff a coefficient of −0.0047 with both tariffs entered (Table 7), against −0.0059 in the full sample (Table 6). The compression is a feature of the whole sample period rather than an artifact of the recent episode.

Table 7: Output and input tariffs and 90–10 dispersion, pre-trade-war subsample (1991–2017)

|                                         | (1)<br>Output         | (2)<br>Input       | (3)<br>Both           |
|-----------------------------------------|-----------------------|--------------------|-----------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0043**<br>(0.0018) |                    | -0.0047**<br>(0.0018) |
| $\Delta$ input tariff $\tau^I$ , lag 1  |                       | 0.0140<br>(0.0211) | 0.0194<br>(0.0211)    |
| Observations                            | 2,295                 | 2,295              | 2,295                 |
| Within $R^2$                            | 0.039                 | 0.038              | 0.039                 |
| Industry FE                             | Yes                   | Yes                | Yes                   |
| Year FE                                 | Yes                   | Yes                | Yes                   |

*Notes:* Table 6 re-estimated on years up to 2017, before any Section 301 or 232 tariff enters the lagged change. Dependent variable: change in the 90–10 dispersion spread; output and input statutory tariffs, lagged one year; industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2017. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 7.2 Where in the distribution

Table 8 traces the effect across the distribution. The output tariff narrows the overall 90–10 spread and the 75–25 interquartile range (both significant at five percent), while the very-upper (99–90) tail is imprecise ( $t = -1.1$ ) and the lower (10–1) tail is flat ( $t = 0.7$ ): neither tail’s own spread responds. Read literally, the 90th percentile falls relative to the 10th and the narrowing is concentrated between the quartiles. It cannot be placed more finely than that: there is no median among these statistics, so the 75–25 compression cannot be split into its halves.

**What this does and does not show.** Set beside the fall in the average level, the pattern is consistent with a downward leveling concentrated among the better performers, the distributional counterpart of the decline in the mean. But that is an interpretation, and three features of the data stop it from being a measurement. First, every column of Table 8 is a *gap* between two percentiles, and DISP publishes industry-year percentiles rather than a panel in which plants can be followed; a flat coefficient never establishes that either end of a gap held still, and no coefficient here says which establishments moved. Second, the compression is in the *unweighted* (establishment-count) dispersion: in the activity-weighted spread the output tariff gives a precise zero (+0.0005, s.e. 0.0015; Appendix Table A.18, column 2), so it sits among plants that carry little of the sector’s output, whereas the level result of Section 5 is an output-weighted aggregate. Third, the two results are identified off different variation

(the level from the trade war, the compression from the full panel), and the instrument does not separately identify a dispersion effect, which is insignificant in the war cross-section. The tempting syllogism, that a falling level and a narrowing spread together imply the top came down, therefore does not go through on these data: the plants doing the compressing are not the plants moving the measured level. The downward-leveling reading stays a reading. What the sign of the compression can be set against is the cross-country evidence on trade exposure and dispersion, which points both ways: [Berthou et al. \(2021\)](#) find that import competition widens dispersion in Europe, which under protection would compress it as it does here, while [Ding et al. \(2016\)](#) find that import competition compresses dispersion in China.

Table 8: Output and input tariffs across the productivity distribution (first differences)

|                                         | (1)<br>90–10          | (2)<br>99–90        | (3)<br>10–1         | (4)<br>75–25           |
|-----------------------------------------|-----------------------|---------------------|---------------------|------------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0059**<br>(0.0024) | -0.0025<br>(0.0029) | 0.0015<br>(0.0020)  | -0.0060***<br>(0.0020) |
| $\Delta$ input tariff $\tau^I$ , lag 1  | 0.0013<br>(0.0095)    | 0.0149<br>(0.0113)  | -0.0011<br>(0.0047) | 0.0058<br>(0.0038)     |
| Observations                            | 2,635                 | 2,606               | 2,607               | 2,635                  |
| Within $R^2$                            | 0.038                 | 0.027               | 0.039               | 0.037                  |
| Industry FE                             | Yes                   | Yes                 | Yes                 | Yes                    |
| Year FE                                 | Yes                   | Yes                 | Yes                 | Yes                    |

*Notes:* Each column is a first-differences regression of the change in the indicated percentile spread on the lagged change in the statutory output and input tariffs, with industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2021. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

### 7.3 The placebo on tariff timing

If the lagged change in the statutory output tariff captures a delayed *effect* of protection, then *future* tariff changes should not predict today’s change in dispersion. Table 9 replaces the causal lag with leads of the statutory change. The placebo requires care, because the statutory change is itself serially correlated: within industry its first-order autocorrelation is about 0.4, while its autocorrelation at two or more years is negligible. The one-year lead is therefore mechanically linked to the contemporaneous change and is *not* a clean placebo, whereas the two- and three-year leads are. Indeed, at those clean horizons the placebo passes: the lead is economically small and statistically indistinguishable from zero (columns 3–4), so the lagged association for dispersion is not an artifact of a slow-moving trend. This is in

contrast to the level regression, whose leads remain significant. The one-year lead is positive and significant and survives controlling for the contemporaneous change (column 5); I read this narrow, single-year pattern as a limited reverse-causality or anticipation window rather than a pervasive pre-trend, since it disappears at longer horizons.

The output tariff therefore lowers the level of productivity and compresses its spread, and the compression is not a trend in disguise. What neither result yet says is why. The next section asks through which channel protection operates, by testing the two leading candidates against evidence from outside the trade war.

Table 9: Placebo test: lead progression of the statutory output-tariff change (change in 90–10 dispersion)

|                                      | (1)<br>Lag (base)     | (2)<br>Lead 1         | (3)<br>Lead 2      | (4)<br>Lead 3       | (5)<br>Lead 1 +ctmp   |
|--------------------------------------|-----------------------|-----------------------|--------------------|---------------------|-----------------------|
| $\Delta$ statutory, lag 1 (causal)   | -0.0058**<br>(0.0025) |                       |                    |                     |                       |
| $\Delta$ statutory, contemporaneous  |                       |                       |                    |                     | -0.0050**<br>(0.0019) |
| $\Delta$ statutory, lead 1 (placebo) |                       | 0.0065***<br>(0.0020) |                    |                     | 0.0064***<br>(0.0021) |
| $\Delta$ statutory, lead 2 (placebo) |                       |                       | 0.0017<br>(0.0047) |                     |                       |
| $\Delta$ statutory, lead 3 (placebo) |                       |                       |                    | -0.0029<br>(0.0020) |                       |
| Observations                         | 2,635                 | 2,635                 | 2,550              | 2,465               | 2,635                 |
| Within $R^2$                         | 0.038                 | 0.039                 | 0.037              | 0.039               | 0.040                 |
| Industry FE                          | Yes                   | Yes                   | Yes                | Yes                 | Yes                   |
| Year FE                              | Yes                   | Yes                   | Yes                | Yes                 | Yes                   |

*Notes:* First-differences regressions of the change in the 90–10 dispersion spread on the indicated changes in the statutory output tariff, with industry and year fixed effects; standard errors clustered by industry in parentheses. Column (1) is the maintained (causal) lagged change. Columns (2)–(4) replace it with the change one, two, and three years in the future: a clean placebo is indistinguishable from zero. Because  $\Delta\tau$  has first-order autocorrelation  $\approx 0.4$  within industry but negligible autocorrelation beyond, only the one-year lead is mechanically linked to the contemporaneous change; column (5) adds that contemporaneous change to column (2). Estimation years: 1991–2021 in column 1, 1990–2020 in columns 2 and 5, 1990–2019 in column 3, and 1990–2018 in column 4. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

## 8 Mechanisms

Sections 5 and 7 establish that the output tariff lowers the level of industry TFP and compresses its within-industry dispersion. This section takes up the channel through which it

does so. As argued in the introduction, the evidence is most consistent with a weakening of competitive discipline in the protected market, where the own-product tariff carries the effect and the tariff an industry pays on its inputs does not. I proceed in two steps. First, if the tariffs act through competition, evidence from episodes in which competition rose bears directly on them: Sections 8.1 and 8.2 extend two such designs, the Autor et al. (2013) China shock and the Pierce and Schott (2016a) PNTR shock, from employment to productivity. Both move employment; neither detects a significant productivity response, and Section 9 weighs what that does and does not imply. Second, Section 8.3 tests the alternative of an input-cost channel and finds the input tariff null under a variety of alternative plausible constructions.

## 8.1 The China shock: an increase in competition

China’s share of U.S. manufacturing imports rose from 9 percent in 2000 to 20 percent in 2007 and 23 percent by 2014, computed from the Schott (2025) import data. Autor et al. (2013) attribute that rise to China’s own expansion of supply rather than to U.S. demand, and they trace its consequences in local labor markets, where the most exposed places lost manufacturing employment. The episode has become the standard natural experiment on import competition in the United States, and it is the mirror of the trade war: the shift in competitive pressure is large, it falls on the same four-digit industries, and its sign is the opposite of a tariff increase. If protection lowers productivity by dulling competition, this episode shows what competition moves when it rises. Autor et al. and most of the work that follows study employment; I extend the design to productivity.

**Design.** I instrument the change in US import penetration from China with the change in China’s penetration of *eight other high-income countries* (Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland). The logic is that China’s surging exports reflect a supply-side push (falling costs and rising productivity in China) that is common across destination markets and orthogonal to demand or productivity shocks specific to a US industry. China’s exports to those countries by HS6 product come from the CEPII BACI database (Gaulier and Zignago, 2010; CEPII, 2026) and are mapped to NAICS-4 with the same HS–NAICS concordance used to build the US import data. Import penetration is China’s value divided by base-period US absorption (shipments plus imports); I estimate the regression in stacked seven-year differences (2000–2007 and 2007–2014, the China-shock era) with period fixed effects and standard errors clustered by industry. No tariff measure enters this regression; the treatment is import penetration.

**Result.** Table 10 reports the estimates. The instrument is strong: the first-stage  $F$  is about 114, and China’s penetration of the other markets tracks US China penetration with a correlation of 0.88. The shock bites, and it bites on *employment*: a one-percentage-point rise in Chinese import penetration lowers industry employment by about 1.9 percent ( $-0.0186$ ,  $t = -3.0$ ), the familiar Autor et al. (2013) result recovered at the industry level. On productivity, by contrast, the second stage is null. The dispersion estimate is *positive* ( $+0.0027$ ), the direction the competition story predicts: more competition should *widen* the spread, the mirror image of protection compressing it. It is, however, statistically indistinguishable from zero ( $t = 0.5$ ). The TFP-level estimates are also positive (the pro-competitive direction) but not significant (BLS,  $t = 1.5$ ). The employment response shows that the design can detect a change of that size when one is present. A productivity estimate indistinguishable from zero is a statement about what these data resolve, not a finding that competition leaves productivity unchanged, and resolving power is specific to the outcome: industry productivity is measured with more noise than employment, so the same design pins down the one and not the other. For BLS TFP the 95 percent interval runs from  $-0.0010$  to  $0.0080$ , which is consistent with a productivity gain of nearly one percent for every percentage point of Chinese penetration. That is not a precise zero.

Table 10: The China shock, employment and productivity: a test of the competition channel

|                                   | (1)<br>Employment           | (2)<br>Dispersion 90–10 | (3)<br>TFP (BLS)     | (4)<br>TFP (NBER–CES) |
|-----------------------------------|-----------------------------|-------------------------|----------------------|-----------------------|
| $\Delta$ China import penetration | $-0.0186^{***}$<br>(0.0062) | $0.0027$<br>(0.0058)    | $0.0035$<br>(0.0023) | $0.0017$<br>(0.0027)  |
| same, OLS                         | $-0.0154$                   | $-0.0019$               | $0.0026$             | $0.0010$              |
| First-stage $F$                   | 113.8                       | 113.8                   | 113.8                | 113.8                 |
| Observations                      | 170                         | 170                     | 170                  | 170                   |
| Period FE                         | Yes                         | Yes                     | Yes                  | Yes                   |

*Notes:* Two-stage least-squares estimates of the change in the outcome (log NBER–CES employment, the 90–10 TFP dispersion, or the BLS or NBER–CES log TFP level) on the change in US import penetration from China, instrumented by the change in China’s penetration of eight other high-income countries (Autor et al., 2013); China’s exports to those countries are from CEPII BACI (Gaulier and Zignago, 2010). Stacked seven-year differences (2000–2007, 2007–2014); penetration in percentage points of base-period US absorption; period fixed effects; SEs clustered by industry. “Same, OLS” is the corresponding ordinary-least-squares coefficient. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Accordingly, the China shock moves employment clearly, but any reaction of total factor productivity is not detectable. This is the same pattern the Pierce–Schott design in Section 8.2 shows from a different episode. It does not follow that competition leaves productivity untouched. The change in competition may not have been sharp enough to register a productivity response, or the effects of rising and falling competition may differ; Section 9

returns to both readings.

## 8.2 A Pierce–Schott quasi-experiment: which margin absorbs a competition shock?

A second episode raises competition and moves employment without moving productivity: the quasi-experiment of [Pierce and Schott \(2016a\)](#). It shifts an industry along the *same* axis of competitive pressure as the 2018–19 tariffs, only in the opposite direction: toward more competition rather than less. [Pierce and Schott \(2016a\)](#) use it to study employment; I extend it to productivity and to the other inputs, to learn *which margin* of the firm a change in competition moves.

**Identification.** When the United States granted China Permanent Normal Trade Relations (PNTR) in 2000, it removed a standing threat: that U.S. tariffs on Chinese goods would revert from the low NTR (Column 1) rates actually applied to the high non-NTR (Column 2) rates set under the Smoot–Hawley Act of 1930. Industries differ in their *NTR gap*, the distance between the two rates, which is large where the threatened increase was large. Because the non-NTR rates were fixed seventy years earlier, the gap is plausibly exogenous to an industry’s later trajectory, and the threat’s removal at a known date (2001) makes high-gap industries the treated group in a generalized difference-in-differences,

$$Y_{it} = \theta (\text{NTR gap}_i \times \text{post-PNTR}_t) + \text{post}_t \times \mathbf{X}'_i \gamma + \mathbf{X}'_{it} \lambda + \delta_i + \delta_t + \varepsilon_{it}, \quad (7)$$

with industry ( $\delta_i$ ) and year ( $\delta_t$ ) fixed effects, the PS control battery, 1990-employment weights, and clustering by industry, over 1990–2007. The treatment is an *increase* in Chinese import competition; [Pierce and Schott \(2016a\)](#) find  $\theta < 0$  for employment: high-gap industries shed disproportionately more jobs after PNTR. I hold their specification fixed and swap the outcome  $Y_{it}$ , first across my productivity measures and then across the real input quantities, validating throughout by reproducing the employment result.

**Differences from Pierce–Schott.** My replication uses only public data, which forces three departures: I use the public NBER–CES employment that PS report reproduces their confidential-LBD result; I aggregate the NTR gap and controls of their published replication data ([Pierce and Schott, 2016b](#)) from their six-digit “family” industries to my 86 four-digit NAICS industries; and I omit two restricted controls (Chinese production subsidies and export-licensing requirements), retaining the rest. The NTR gap is built from the same Feenstra–Romalis–Schott schedule rates and its 1999 manufacturing mean of 31.5 percentage

points matches theirs. The treatment is therefore a gap between two statutory schedules; the paper’s own output and input tariffs do not enter equation (7).

**No significant productivity response.** Table 11 swaps the outcome across employment and three productivity measures. Column 1 recovers Pierce–Schott: the NTR-gap- $\times$ -post coefficient on employment is negative and significant ( $-0.0081$ ; an interquartile move in the gap implies about a 0.14 log-point relative employment decline, against their 0.08). By contrast, the three productivity columns show no significant response: the coefficients are close to zero for both TFP series and for the 90–10 dispersion, although their confidence intervals do not exclude modest effects. The sizable cross-sectional association (the “no fixed effects” row) *vanishes* once the fixed effects that identify the employment result are included. The same design that delivers a strong employment effect thus finds no significant productivity effect; that it recovers the employment result also indicates, in passing, that the public-data build is sound.

Table 11: Pierce–Schott NTR-gap diff-in-differences: employment versus productivity

|                            | (1)<br>Employment      | (2)<br>TFP (NBER–CES) | (3)<br>TFP (BLS)   | (4)<br>Dispersion 90–10 |
|----------------------------|------------------------|-----------------------|--------------------|-------------------------|
| NTR gap $\times$ post-PNTR | -0.0081***<br>(0.0028) | -0.0002<br>(0.0010)   | 0.0008<br>(0.0020) | -0.0027<br>(0.0047)     |
| same, no fixed effects     | -0.0067                | 0.0003                | 0.0029             | 0.0044                  |
| Observations               | 1,491                  | 1,491                 | 1,491              | 1,491                   |
| Industry FE                | Yes                    | Yes                   | Yes                | Yes                     |
| Year FE                    | Yes                    | Yes                   | Yes                | Yes                     |
| PS controls                | Yes                    | Yes                   | Yes                | Yes                     |

*Notes:* Each column is the [Pierce and Schott \(2016a\)](#) generalized difference-in-differences (7) with the indicated outcome: the post-PNTR (2001) change in that outcome for industries with larger NTR gaps. NTR gap in percentage points (1999); industry and year fixed effects; controls are capital- and skill-intensity (1990) interacted with post-PNTR, plus Nunn contract intensity, Chinese import tariffs, MFA exposure, the NTR rate, and union membership; weighted by 1990 employment; SEs clustered by industry; 1990–2007. “Same, no fixed effects” is the NTR-gap- $\times$ -post coefficient from the corresponding regression without industry and year fixed effects. The Chinese subsidy and export-licensing controls are omitted (restricted); employment is from NBER–CES. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

**Only labor adjusts.** Where, then, does the competition shock register? Table 12 swaps the outcome across the four real input quantities. *Employment* (column 1) and production-worker *hours* (column 4) both fall sharply ( $-0.0081$  and  $-0.0085$ , each significant at one percent): the labor adjustment runs on both the extensive (headcount) and the intensive (hours) margins. *Capital*, the real perpetual-inventory capital stock (column 2), is essentially unchanged ( $-0.0016$ , insignificant), as one expects of a sunk, slow-moving input over a hori-

zon this short. *Intermediate inputs* (column 3) do not move either ( $-0.0004$ , insignificant). I measure their real volume as each industry’s nominal materials cost deflated by the NBER–CES materials price index, summed from the six-digit children to the four-digit industry, so the coefficient is a quantity rather than an expenditure response: high-gap industries do not cut their material throughput. The PNTR competition shock is therefore absorbed almost entirely on the *labor* margin, while capital, materials, and measured productivity are left where they were. The labor adjustment is consistent with the offshoring and labor-shedding that [Pierce and Schott \(2016a\)](#) document.

Table 12: The PNTR shock and the input mix: which inputs adjust?

|                            | (1)<br>Employment      | (2)<br>Capital      | (3)<br>Materials    | (4)<br>Hours           |
|----------------------------|------------------------|---------------------|---------------------|------------------------|
| NTR gap $\times$ post-PNTR | -0.0081***<br>(0.0028) | -0.0016<br>(0.0015) | -0.0004<br>(0.0031) | -0.0085***<br>(0.0030) |
| same, no fixed effects     | -0.0067                | 0.0004              | -0.0019             | -0.0081                |
| Observations               | 1,491                  | 1,491               | 1,491               | 1,491                  |
| Industry FE                | Yes                    | Yes                 | Yes                 | Yes                    |
| Year FE                    | Yes                    | Yes                 | Yes                 | Yes                    |
| PS controls                | Yes                    | Yes                 | Yes                 | Yes                    |

*Notes:* The same [Pierce and Schott \(2016a\)](#) difference-in-differences (7) as Table 11, with the outcome swapped across four real input quantities from the NBER–CES Manufacturing Industry Database, 1990–2007: log employment, log real capital stock, log real intermediate inputs, and log production-worker hours. Real intermediate inputs are nominal materials cost deflated by the NBER–CES materials price index, summed from the six-digit children to the four-digit NAICS industry. NTR gap in percentage points (1999); the same PS control battery, industry and year fixed effects, 1990-employment weights, and clustering by industry as in Table 11. “Same, no fixed effects” is the NTR-gap- $\times$ -post coefficient without the fixed effects. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

The Pierce–Schott shock and the China shock therefore agree, from two unrelated episodes: a rise in competitive pressure shows up clearly on the labor margin and not detectably in measured productivity. What that implies for the headline, a *fall* in pressure that does reach the production frontier, is taken up in Section 9, after the input-cost channel has been tested.

### 8.3 The input-cost channel: a contrast with Amiti–Konings

The output/input tariff decomposition originates with [Amiti and Konings \(2007\)](#), who find that for Indonesian plants the *input* channel dominates: cheaper and more varied imported intermediates raise productivity by roughly twice as much as the competition effect of output-tariff cuts. [Topalova and Khandelwal \(2011\)](#) report the same for India, though [De Loecker et al. \(2016\)](#) show for the same Indian reform that firms kept much of the input-cost saving

as markups rather than passing it into prices, which mutes what a productivity measure built on revenues can detect. By contrast, in my setting the input channel is *absent*: across specifications the input tariff is statistically insignificant, while the output tariff carries the result. This section documents that the absence is robust to how the input tariff is constructed, and discusses why the channel is muted here.

**Two refinements of the input tariff.** The input tariff  $\tau_j^I = \sum_c \omega_{cj} \tau_c^O$ , an average of suppliers’ statutory output tariffs, already excludes the own industry ( $c = j$ ; Appendix A.5). I subject it to the two refinements most likely to revive the channel. First, and most important, tariffs apply only to *imported* inputs, so the cost-channel logic calls for weighting each supplier by its share of  $j$ ’s *imported* intermediate use rather than its share of total use. I therefore replace the fixed-2012 cost shares with import-use shares  $\omega_{cj}^M$  from the *BEA 2012 benchmark detail Import Matrix* (before redefinitions), which records the value of each imported commodity used by each industry. Second, to rule out that the input tariff merely inherits the output tariff’s variation through vertically-linked, similarly-protected neighbors, I build a “leave-out” version that drops every supplier in  $j$ ’s own three-digit NAICS sector and renormalizes (dropping the own sector removes little: the median industry retains 0.96 of its input weight). Both refinements change only the weights; the supplier tariffs remain statutory. The import-use weights differ materially from total-use weights (their correlation across supplier–buyer cells is 0.61), and outside the trade war the leave-out-three-digit measure’s change correlates only 0.06 with the output-tariff change (against 0.17 for the baseline). Within the war all of these measures co-move with the output tariff, because the 2018–19 increases raised tariffs across nearly the whole sector at once.

**Result.** Table 13 reports first-difference regressions on the lagged changes in the statutory output tariff and the import-weighted input tariff, with the total-use and leave-out-three-digit coefficients shown as memos. The input-cost channel is absent under a variety of alternative plausible constructions: the input-tariff coefficient is insignificant (or, for BLS TFP, marginal at best) for all three outcomes, whether weighted by total use, by actual imported use, or restricted to far-upstream suppliers. Moreover, the output-tariff coefficient is unchanged throughout, at about  $-0.006$  for dispersion and  $-0.001$  for each TFP level, so the headline result does not depend on how the input tariff is measured. The absence of an Amiti–Konings input channel is thus not an artifact of the weighting scheme: it survives the very import-weighting that channel’s logic demands. The years before the war add little on this channel: in the pre-war subsample the input tariff’s coefficient on dispersion is insignificant and imprecise, with a standard error nearly three times its full-sample value

(Table 7).

**Why the channel is muted here.** Several features distinguish my setting from [Amiti and Konings \(2007\)](#). First, the variation I exploit is a tariff *increase* (the 2018–19 war), not a liberalization; the input-variety and quality gains that drive their result run in reverse and may be asymmetric or slow. Second, realized input tariffs rose far less than the statutory ones that enter the regressions. Firms shifted sourcing and obtained exclusions, so the effective cost shock was small (the effective input tariff rose about 2.6 points over 2017–2019 against 16.5 for the statutory measure). Measuring the input tariff at those realized rates, instrumented by the statutory rate, leaves the dispersion result null (Appendix Table A.17). Third, my productivity data are at the four-digit NAICS industry level, which cannot see the plant-level imported-input-variety margin that operates in their micro data. The firm-level TFP evidence for the channel also comes from developing and emerging economies: Indonesia ([Amiti and Konings, 2007](#)), India ([Topalova and Khandelwal, 2011](#)) and Hungary ([Halpern et al., 2015](#)). Finally, the window is short. The reading is not that input tariffs are irrelevant in general, but that the margin Amiti–Konings isolate is not active in a brief, partly-avoided tariff *increase* observed at the industry level.

Table 13: The input-cost channel under alternative input-tariff constructions

|                                                      | (1)<br>Dispersion 90–10 | (2)<br>TFP (BLS)       | (3)<br>TFP (NBER–CES)  |
|------------------------------------------------------|-------------------------|------------------------|------------------------|
| $\Delta$ output tariff                               | -0.0057**<br>(0.0024)   | -0.0013***<br>(0.0004) | -0.0012***<br>(0.0003) |
| $\Delta$ input tariff (import-weighted)              | -0.0023<br>(0.0093)     | -0.0021*<br>(0.0012)   | -0.0021<br>(0.0041)    |
| $\Delta$ input tariff (total-use, own 4-digit excl.) | 0.0013                  | -0.0020                | -0.0037                |
| $\Delta$ input tariff (leave-out own 3-digit)        | -0.0014                 | -0.0009                | 0.0029                 |
| Observations                                         | 2,635                   | 2,720                  | 2,210                  |
| Industry FE                                          | Yes                     | Yes                    | Yes                    |
| Year FE                                              | Yes                     | Yes                    | Yes                    |

*Notes:* First-differences regressions of the change in the outcome (the 90–10 TFP dispersion, or the BLS or NBER–CES log TFP level) on the once-lagged changes in the statutory output tariff and the *import-weighted* input tariff. Every input-tariff construction averages suppliers’ statutory output tariffs. The import-weighted input tariff uses fixed-2012 shares of imported intermediate use from the BEA benchmark detail Import Matrix (before redefinitions), with the own industry excluded. Memo rows report the input-tariff coefficient under the baseline total-use weighting (own four-digit excluded) and under a leave-out-own-three-digit-sector construction. Industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2021 in column 1, 1991–2022 in column 2, and 1991–2016 in column 3. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

The input channel, then, is absent under a variety of alternative plausible constructions of the input tariff. With both alternatives tested, Section 9 draws the three subsections

together.

## 9 Conclusion

Between 2018 and 2019 the United States raised tariffs on roughly \$350 billion of imports, and this paper has asked what that protection did to the productivity of the industries it shielded. It lowered it. Because tariffs are not handed out at random, I identify the effect from the design of the trade war itself: the U.S. Trade Representative front-loaded intermediate and capital goods and delayed consumer goods, so an industry’s consumer-good share is a strong instrument for its tariff exposure that is unrelated to its productivity path. The quartiles of the 2017–19 increase in the statutory tariff are 6.8 percentage points apart, and across that gap an industry at the top quartile ends 2019 with total factor productivity about 3.6 percent below one at the bottom. This estimate is robust to controlling for retaliation, and it is reproduced at 2.4 percent by a second design that uses no instrument.

The decline is not spread evenly across the productivity distribution: the within-industry dispersion of productivity compresses, with the 90th percentile falling relative to the 10th and the narrowing concentrated between the quartiles rather than in either tail. These are industry-year percentiles rather than a panel of plants, so they locate the narrowing without identifying which establishments moved; the pattern is consistent with a downward leveling concentrated among the better performers. The two results rest on different variation, and deliberately so: the level decline is a trade-war finding, identified twice over, while the compression is a full-sample association that survives dropping the war years entirely.

The *output* tariff, the rate that shields an industry’s own market, produces both the decline in the level and the compression of the spread, while the tariff an industry pays on its inputs produces neither. Extending the China-shock and Pierce–Schott designs, two increases in competition from other episodes, to productivity, I find that both move employment but neither detects a significant productivity response. The reason is either that those changes in competition were not sharp enough to register, or that rising and falling competition act asymmetrically. Evidence from other settings fits the first reading: import competition raised plant TFP in Chile after its liberalization ([Pavcnik, 2002](#)) and within European firms exposed to Chinese imports ([Bloom et al., 2016](#)). The Pierce–Schott design also shows a competition shock absorbed almost entirely on the labor margin, with capital and materials unchanged. The input-cost channel emphasized by Amiti–Konings, meanwhile, is absent under a variety of alternative plausible constructions of the input tariff.

One mechanism would produce such an asymmetry, an *effort ceiling*. Competitive discipline already holds firms close to the best they can do with what they have, so a further

rise in pressure finds little slack to squeeze and moves scale instead, which is what the labor response in both episodes shows. A fall in pressure is not the mirror image, because nothing bounds slack from above: effort, monitoring intensity and the tolerance of waste can all drift back, and none of them is locked in the way innovation is by its up-front investment (Melitz and Trefler, 2012). Firms fall inside the production frontier faster than they climb back to it. Further tests of this conjecture are beyond the scope of this paper.

This places my result outside the standard trade model (Krugman et al., 2012, chs. 6 and 9), in which protection distorts the input mix but leaves technology untouched: a movement *along* the production frontier (Kehoe and Ruhl, 2008). I see that input-mix reallocation, in the Pierce and Schott (2016a) quasi-experiment of Section 8.2, but I also find that tariffs lower TFP itself: a movement *inside* the production possibility frontier, in which the same inputs yield less output. The factor reallocation the standard trade model isolates is present, and the larger margin is the contraction of the production frontier itself, which lies outside the model and comes on top of the inefficiency it does describe.

## References

- Mary Amiti and Jozef Konings. Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review*, 97(5):1611–1638, 2007.
- Mary Amiti, Stephen J. Redding, and David E. Weinstein. The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives*, 33(4):187–210, 2019.
- David Autor, David Dorn, Gordon Hanson, Gary Pisano, and Pian Shu. Foreign competition and domestic innovation: Evidence from U.S. patents. *American Economic Review: Insights*, 2(3):357–374, 2020.
- David H. Autor, David Dorn, and Gordon H. Hanson. The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168, 2013.
- Randy A. Becker, Wayne B. Gray, and Jordan Marvakov. NBER-CES manufacturing industry database. *National Bureau of Economic Research*, 2021. 1958–2018, version 2021a, NAICS 2012 version (nberces5818v1\_n2012), <https://www.nber.org/research/data/nber-ces-manufacturing-industry-database>.
- Antoine Berthou, John Jong-Hyun Chung, Kalina Manova, and Charlotte Sandoz. Trade, productivity and (mis)allocation. Working paper, Banque de France, Auburn University, University College London and International Monetary Fund, 2021. Earlier version issued as IMF Working Paper 20/163, August 2020.
- Nicholas Bloom, Mirko Draca, and John Van Reenen. Trade induced technical change? the impact of Chinese imports on innovation, IT and productivity. *Review of Economic Studies*, 83(1):87–117, 2016.
- Chad P. Bown. The US–China trade war and Phase One agreement. *Journal of Policy Modeling*, 43(4):805–843, 2021.
- Alberto Cavallo, Gita Gopinath, Brent Neiman, and Jenny Tang. Tariff passthrough at the border and at the store: Evidence from US trade policy. *American Economic Review: Insights*, 3(1):19–34, 2021.
- CEPII. BACI: International trade database at the product-level. Version 202601, HS92 nomenclature, 2026. Accessed January 2026. [http://www.cepii.fr/CEPII/en/bdd\\_modele/bdd\\_modele\\_item.asp?id=37](http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=37).

- Cindy Cunningham, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. Dispersion in dispersion: Measuring establishment-level differences in productivity. BLS Working Paper 530, U.S. Bureau of Labor Statistics, 2022.
- Jan De Loecker. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451, 2011.
- Jan De Loecker and Frederic Warzynski. Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471, 2012.
- Jan De Loecker, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. Prices, markups, and trade reform. *Econometrica*, 84(2):445–510, 2016.
- Sai Ding, Wei Jiang, and Puyang Sun. Import competition, dynamic resource allocation and productivity dispersion: Micro-level evidence from China. *Oxford Economic Papers*, 68(4):994–1015, 2016.
- Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. Replication data for: The return to protectionism. Harvard Dataverse, V1, 2019. Accessed June 2026. <https://doi.org/10.7910/DVN/KSOVSE>.
- Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. The return to protectionism. *Quarterly Journal of Economics*, 135(1):1–55, 2020.
- Robert C. Feenstra, John Romalis, and Peter K. Schott. U.S. imports, exports, and tariff data, 1989–2001. NBER Working Paper 9387, National Bureau of Economic Research, 2002. Ad valorem equivalent NTR (MFN) and non-NTR tariff rates at the 8-digit HS level; distributed as `tar_val.dta` on Schott’s data page.
- Aaron Flaaen and Justin R. Pierce. Disentangling the effects of the 2018–2019 tariffs on a globally connected U.S. manufacturing sector. *Review of Economics and Statistics*, 2024. [https://doi.org/10.1162/rest\\_a\\_01498](https://doi.org/10.1162/rest_a_01498).
- Guillaume Gaulier and Soledad Zignago. BACI: International trade database at the product-level. Working Paper 2010-23, CEPII, 2010.
- Pinelopi K. Goldberg, Amit K. Khandelwal, Nina Pavcnik, and Petia Topalova. Imported intermediate inputs and domestic product growth: Evidence from India. *Quarterly Journal of Economics*, 125(4):1727–1767, 2010.
- László Halpern, Miklós Koren, and Adam Szeidl. Imported inputs and productivity. *American Economic Review*, 105(12):3660–3703, 2015.

- Chang-Tai Hsieh and Peter J. Klenow. Misallocation and manufacturing TFP in China and India. *Quarterly Journal of Economics*, 124(4):1403–1448, 2009.
- Timothy J. Kehoe and Kim J. Ruhl. Are shocks to the terms of trade shocks to productivity? *Review of Economic Dynamics*, 11(4):804–819, 2008.
- Michael Kilumelume, Bruno Morando, Carol Newman, and John Rand. Tariffs, productivity and resource misallocation. *World Bank Economic Review*, 2025. <https://doi.org/10.1093/wber/lhaf010>.
- Paul R. Krugman, Maurice Obstfeld, and Marc J. Melitz. *International Economics: Theory and Policy*. Pearson Addison-Wesley, Boston, 9th edition, 2012.
- Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.
- Marc J. Melitz and Daniel Trefler. Gains from trade when firms matter. *Journal of Economic Perspectives*, 26(2):91–118, 2012.
- Office of the United States Trade Representative. Notice of modification of section 301 action: China’s acts, policies, and practices related to technology transfer, intellectual property, and innovation. 84 Fed. Reg. 43304 (August 20, 2019), Annex A, 2019. Product list for Section 301 List 4A. <https://www.govinfo.gov/content/pkg/FR-2019-08-20/pdf/2019-17865.pdf>.
- Nina Pavcnik. Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *Review of Economic Studies*, 69(1):245–276, 2002.
- Justin R. Pierce and Peter K. Schott. A concordance between ten-digit U.S. Harmonized System codes and SIC/NAICS product classes and industries. *Journal of Economic and Social Measurement*, 37(1–2):61–96, 2012.
- Justin R. Pierce and Peter K. Schott. The surprisingly swift decline of US manufacturing employment. *American Economic Review*, 106(7):1632–1662, 2016a.
- Justin R. Pierce and Peter K. Schott. Replication data for: The surprisingly swift decline of US manufacturing employment. American Economic Association, openICPSR, 2016b. Accessed June 2026. <https://doi.org/10.3886/E112965V1>.
- Peter K. Schott. U.S. HS-level imports and exports, 1989–2024. Data, version v2025.4.21, [https://sompks4.github.io/sub\\_data.html](https://sompks4.github.io/sub_data.html), 2025. Annual import-detail

- files `imp_det1_<YYYY>_12n.dta`, converted from U.S. Census Bureau records. Accessed June 2026.
- Chad Syverson. What determines productivity? *Journal of Economic Literature*, 49(2): 326–365, 2011.
- Petia Topalova and Amit Khandelwal. Trade liberalization and firm productivity: The case of india. *Review of Economics and Statistics*, 93(3):995–1009, 2011.
- Daniel Trefler. The long and short of the Canada–U.S. free trade agreement. *American Economic Review*, 94(4):870–895, 2004.
- U.S. Bureau of Economic Analysis. Input-output accounts: Use table, before redefinitions, producers’ prices, detail, and import matrices, before redefinitions, detail. Industry Economic Accounts, n.d. Benchmark years 2007, 2012 and 2017; the 2012 tables are used. Accessed June 2026. <https://apps.bea.gov/industry/Release/XLSX/>.
- U.S. Bureau of Labor Statistics. Dispersion statistics on productivity (DiSP). Experimental data product, U.S. Bureau of Labor Statistics and U.S. Census Bureau, 2024. Updated September 12, 2024. Covers 86 four-digit NAICS manufacturing industries, 1987–2021, on a longitudinally consistent 2017 NAICS basis. <https://www.bls.gov/productivity/articles-and-research/dispersion-statistics-on-productivity/>.
- U.S. Bureau of Labor Statistics. Total factor productivity for manufacturing and transportation detailed industries. Office of Productivity and Technology, detailed-industry tables, 2026. Data released August 26, 2026. <https://www.bls.gov/productivity/tables/>.
- U.S. International Trade Commission. Annual tariff data. USITC DataWeb, 2025. Annual tariff database files, 2002–2023. Accessed March 2025. <https://dataweb.usitc.gov/tariff/annual>.

# Appendix for Online Distribution

to

“Higher Tariffs, Lower Productivity:  
Evidence from U.S. Manufacturing”

Luca Guerrieri

Draft — September 20, 2026

## Abstract

This appendix has two parts. The first documents the construction of an industry-year panel that pairs U.S. import tariff rates with the Dispersion Statistics on Productivity (DiSP) of [Cunningham et al. \(2022\)](#); [U.S. Bureau of Labor Statistics \(2024\)](#). We build *two* conceptually distinct tariff measures at the four-digit NAICS level for 1989–2021: a **realized** (effective) rate—duties actually collected as a share of customs value—and a **statutory** (schedule) rate—the legal ad valorem rate, including the Section 232 and Section 301 actions of 2018–19. Section [A.2](#) describes the common industry-classification and harmonization framework; Section [A.3](#) builds the statutory rates and Section [A.4](#) the realized rates; Section [A.5](#) the input tariff and Section [A.6](#) the productivity series; Section [A.7](#) merges them with DiSP and reports the validation checks; and Section [A.8](#) documents the [Flaaen and Pierce \(2024\)](#) exposure measures. The second part, Section [B](#), reports additional empirical results: robustness for the productivity level, the value-added-weighted regressions and the placebo-timing tests, the [Flaaen and Pierce](#) design estimated without removing the pre-war trend, and robustness for the dispersion result.

## A Data and measurement

This appendix documents how the analysis panel is built and validated, and is organized to follow the build itself. After an overview of the panel and the two tariff concepts, I describe the shared industry classification and harmonization; the statutory and then the realized tariff measures; the input tariff; the productivity measures; the merge with DiSP and its validation; and the [Flaaen and Pierce \(2024\)](#) exposure measures. Reconstruction steps, the log of downloaded inputs, and known limitations are recorded in `README.txt` at the root of the repository. The unifying choice throughout is that every series is placed on a single, longitudinally consistent 2012 NAICS basis, so the tariff, productivity, and dispersion data are directly comparable over time and with one another.

### A.1 Overview

The target dataset is a panel indexed by four-digit NAICS manufacturing industry  $j$  and year  $t$ , covering the 86 NAICS manufacturing industries for which DiSP reports productivity dispersion. For each cell I observe DiSP’s dispersion moments alongside two tariff concepts:

- the **realized (effective) rate**  $\tau_{jt}^R$  = duties/customs value, an ex-post, import-value-weighted average that falls below schedule rates because of duty-free entry under preference programs (Section [A.4](#));
- the **statutory (schedule) rate**  $\tau_{jt}^S$ , the legal ad valorem rate: the most-favored-nation (NTR) rate plus the 2018–19 Section 232 and 301 add-ons, which from 2018 makes it the rate facing Chinese goods (Section [A.3](#)). The non-NTR (Column 2) rate is carried separately and enters only the NTR gap.

Both are placed on a single, longitudinally consistent **2012 NAICS** basis so they are comparable over time and with DiSP. The two measures answer different questions and are designed to sit side by side: the realized rate reflects duties actually paid (net of exclusions and origin), while the statutory rate reflects the policy on the books. Every regression reported in the paper uses the statutory rate; the realized rate is documented here for comparison and enters the paper only descriptively.

### A.2 Industry classification and harmonization

Both tariff measures share one classification framework, described here once.

### A.2.1 Industries covered

The panel covers the 86 four-digit NAICS manufacturing industries (sectors 31–33) for which DiSP reports productivity dispersion. With their 2012 NAICS titles, they are split across Tables A.2 and A.3 by leading two-digit sector. NAICS 3328 (coating, engraving, and heat treating, a contract-service industry with essentially no traded product) is retained for completeness but carries a missing tariff in every year by construction.

### A.2.2 The NAICS vintage problem

NAICS is revised every five years (1997, 2002, 2007, 2012, 2017). Trade and tariff data carry the vintage contemporaneous with each year, listed in Table A.1, so codes are not directly comparable across revision boundaries (the largest manufacturing reclassifications occur at 2002 → 2007 and 2007 → 2012, e.g. in computers and electronics, NAICS 334). DiSP, by contrast, is already longitudinally consistent on a single 2017 NAICS basis, using the establishment-level recoding of Fort (2013), published as Fort and Klimek (2018).

### A.2.3 The EFSY time-consistent concordance

I harmonize the trade data with the public counterpart of that same method: the time-consistent NAICS concordance of Eckert et al. (2020) (the “EFSY” concordance, by Eckert, Fort, Schott, and Yang). Using it applies *the same classification logic, by overlapping authors*, that underlies DiSP, rather than an unrelated crosswalk. The concordance supplies pairwise links `naics97` → 02 → 07 → 12 (and SIC links used for the early years); each is a complete six-digit crosswalk whose weights sum to one per input code. The target basis is 2012; at the four-digit manufacturing level the 2012 and 2017 vintages coincide, so the choice of common target is immaterial to the merge.

### A.2.4 One aggregation method for levels and rates

I aggregate by carrying two quantities through the concordance and the six- to four-digit rollup,

$$R_i = \sum_{\ell \in i} r_\ell w_\ell, \quad M_i = \sum_{\ell \in i} w_\ell, \quad (8)$$

and forming the cell value  $R/M$ . For the *realized* rate,  $r_\ell$  and  $w_\ell$  are the calculated duties and customs value of the underlying records, so  $R/M$  is the customs-value-weighted average tariff and totals are conserved. For a *statutory* rate,  $r_\ell$  is a tariff-line rate and  $w_\ell$  a line weight (Section A.3.2). In both cases  $(R, M)$  are aggregated like levels (through the same NAICS vintage chain and the same rollup), so a single piece of code serves both measures.

Table A.1: NAICS vintage carried by the trade/tariff data, by year

| Years     | NAICS vintage        |
|-----------|----------------------|
| 1989–1996 | 1997 (back-assigned) |
| 1997–2001 | 1997                 |
| 2002–2006 | 2002                 |
| 2007–2011 | 2007                 |
| 2012–2016 | 2012                 |
| 2017–2021 | 2017                 |

### A.3 Statutory (schedule) tariff rates

The statutory rate is the legal schedule rate. Unlike the realized rate, it has no collected duty to sum and so must be averaged across tariff lines. After defining it and the within-industry aggregation choice, I describe the two source eras ([Feenstra et al. \(2002\)](#) for 1989–2001 and the USITC DataWeb thereafter) and the two 2018–19 add-ons, Section 232 and Section 301.

#### A.3.1 Definition

The statutory output tariff is the Column 1 / NTR (most-favored-nation) ad valorem rate of the Harmonized Tariff Schedule plus the 2018–19 Section 232 and Section 301 additional duties. China has paid Column 1 rates since 1980, and the Section 301 add-ons apply to Chinese goods alone, so from 2018 the statutory rate is the schedule rate facing Chinese-origin goods. The build also carries the Column 2 / non-NTR rate, but that column enters no regression other than the NTR gap of Section 8.2. Statutory rates lie *above* realized rates because the latter are pulled down by duty-free entry and, from 2018, by imports from origins other than China, on which no Section 301 duty is due; the two are complementary. The statutory rate omits Section 301 List 4A, the 2002 and 2018 Section 201 safeguards, the February 2020 Section 232 derivative products, and antidumping and countervailing duties, and it does not net out exclusions granted to individual importers.

#### A.3.2 Within-industry aggregation

Unlike the realized rate, a statutory rate has no collected duty to sum; the line rates must be averaged, and the weight is a modeling choice. My default is the *simple (unweighted) mean* of line rates, following [Pierce and Schott \(2016\)](#), who define an industry’s rate as “the average ... across the eight-digit HS tariff lines belonging to that industry”; this definition keeps the weights exogenous to trade. An *import-weighted* variant (comparable to  $\tau^R$  but biased down, since high tariffs depress their own weights) is also supported. Aggregation

reuses (8) with  $r_\ell$  the line rate and  $w_\ell \equiv 1$  (simple) or the line’s import value (weighted). The concordance is defined at the HS10 level, so each HS8 line is assigned to the NAICS codes of its HS10 children and split across them in proportion to the number of children in each, with shares summing to one per line and year. Lines that map only to non-industry codes are dropped and X-filled codes are kept at their four-digit prefix; the build records the weight that fails to match.

### A.3.3 1989–2001: Feenstra–Romalis–Schott

NTR and non-NTR ad valorem equivalent rates at the eight-digit HS level come from [Feenstra et al. \(2002\)](#) (the `tar_val` file). Each tariff line is assigned to NAICS as described in Section A.3.2, using the HS10→NAICS concordance of [Pierce and Schott \(2012\)](#); the line file already carries import value for the weighted variant. Specific and compound duties enter as ad valorem equivalents.

### A.3.4 2002–2023: USITC DataWeb

Schedule rates for 2002–2023 come from the USITC DataWeb ([U.S. International Trade Commission, 2025](#)) annual tariff files (`mfn_ad_val_rate` for Column 1, `col2_ad_val_rate` for Column 2), stitched to [Feenstra et al. \(2002\)](#) at 2001/2002. The files vary in format (pipe-delimited through 2016, comma-delimited thereafter) and use 9999.99 as a placeholder for non-ad-valorem (specific-duty) lines, which I set to missing. Purely specific-duty lines are therefore currently excluded pending ad valorem equivalent conversion. The exclusion is small: of the 11,111 ordinary lines in the 2021 schedule (that is, outside the Chapter 98–99 special provisions, which the build drops in any case), 153, or 1.4 percent, carry no ad valorem rate, and they sit mainly in watches and clocks (HS 91), sugar (HS 17) and apparel (HS 61–62) rather than across consumer goods generally. Because the industry rate is a simple mean over the lines in that year’s schedule, a revision to the schedule changes the denominator: the 2022 Harmonized System revision raised the count of ordinary lines from 11,111 to 11,871 while the set of lines matched to the Section 301 action was unchanged at 9,277, so measured coverage and the mean add-on fall in 2022–23 with no change in policy.

### A.3.5 Section 232 (steel and aluminum)

The Section 232 add-on is built directly from [The White House \(2018\)](#): “steel articles” at the HS six-digit ranges 7206.10–7216.50, 7216.99–7301.10, 7302.10, 7302.40–7302.90, 7304.10–7306.90 (25%) and “aluminum articles” in HS 7601 and 7604–7609 plus castings/forgings 7616.99.51.60/.70 (10%), both effective 2018-03-23, expanded to the eight-digit universe.

The overlay is country-agnostic (it records the rate applicable to a covered origin; the proclamations’ Canada/Mexico and later quota-country carve-outs are left to the user’s origin weighting) and omits the February 2020 “derivative” products. As for Section 301, the build reports a companion coverage share, `share_232`, which is the (extensive-margin) fraction of an industry’s tariff lines subject to the steel/aluminum action and is distinct from the add-on *rate* `statutory_232`; the two are linked by  $\text{statutory\_232} = \text{share\_232} \times$  (average add-on on covered lines), exactly as for Section 301 (Table A.5).

### A.3.6 Section 301 (China)

Section 301 applies only to China-origin goods, so it is kept in a separate column rather than folded into a country-agnostic schedule rate. Product membership for Lists 1–3 comes from the replication data of Fajgelbaum et al. (2019), which identify each list by its statutory effective date for China. Because their panel ends 2019-04, I attach the documented rate *schedule* rather than rely on post-2018 values: List 1 (25%, 2018-07-06) and List 2 (25%, 2018-08-23) held open, and List 3 (10% from 2018-09-24, raised to 25% on 2019-05-10), carried forward to the end of the statutory panel. List 4A lies outside their panel, so its membership comes from Annex A of the USTR notice that established it (Office of the United States Trade Representative, 2019): 3,233 eight-digit subheadings, dutiable at 15% from 2019-09-01 and at 7.5% from 2020-02-14. Lines already carrying a List 1–3 add-on and lines subject to Section 232 are excluded from it, leaving 2,894 additions, so no line is counted twice. List 4B, the December 2019 tranche, was suspended before it took effect and is not included. Each covered line’s add-on is date-annualized (a rate in force for part of a year contributes in proportion to the days it applies) and aggregated by (8) into `statutory_301`, with a companion coverage share. The carry-forward is verified against DataWeb’s Chapter 99 provisions (Section A.7.5), which implement the action as separate Chapter 99 headings that add a stated percentage to the duty of the underlying subheading and apply only to articles the product of China. List membership is held at the 2018–19 HS vintage and is not remapped to later revisions, so from 2023 some renumbered lines no longer match: 9,277 lines match the action in 2021 and 2022 but 8,892 in 2023. Coverage in 2023 is therefore understated, which matters for the level of the series rather than for the 2017–19 changes the paper uses.

## A.4 Realized (effective) tariff rates

The realized rate is the duty actually collected as a share of import value. I set out its source and definition, the collapse-and-aggregate construction, and the harmonization to the 2012

NAICS basis.

#### A.4.1 Source and definition

The realized rate uses the annual import-detail files `imp_det1_<YYYY>_12n.dta` of [Schott \(2025\)](#), in which each record is an HS commodity  $\times$  country  $\times$  district observation carrying calculated duties (`cal_dut_yr`), the customs value of imports for consumption (`con_val_yr`), and a pre-assigned six-digit NAICS code ([Pierce and Schott, 2012](#); the data citation is [Schott, 2008](#)). The realized rate is duties divided by customs value. The import-detail files run from 1989 through 2024, one year beyond the statutory series, and the merge with DiSP in Section [A.7](#) trims the panel to 2021.

#### A.4.2 Construction: collapse to NAICS and aggregate

For year  $t$  and six-digit code  $i$  I sum duties  $D_{it}$  and customs value  $V_{it}$  over all records, and for four-digit industry  $j$  form

$$\tau_{jt}^R = \frac{\sum_{i \in j} D_{it}}{\sum_{i \in j} V_{it}}, \quad (9)$$

the customs-value-weighted average ad valorem rate. I aggregate the levels and divide once; six-digit ratios are never averaged. Special codes (980000/990000, NN0000) and missing codes are excluded; “X”-filled partial codes are retained in the four-digit rollup via their numeric prefix (omitting them would, e.g., understate NAICS 3363 by roughly 13% in 1989).

#### A.4.3 Harmonization to 2012 NAICS-4

The native-vintage six-digit levels are mapped to the 2012 basis through the EFSY chain (Section [A.2](#)) and rolled to four digits via (8)–(9).

### A.5 The input tariff

To separate the input-cost channel from the competition channel I build an *input tariff* (the tariff an industry faces on the goods it buys) following the effective-rate-of-protection logic of [Amiti and Konings \(2007\)](#) and [Topalova and Khandelwal \(2011\)](#):

$$\tau_{j,t}^I = \sum_c \omega_{cj} \tau_{c,t}^O, \quad (10)$$

where  $\tau_{c,t}^O$  is the statutory output tariff of supplying industry  $c$  (Section [A.3](#)) and  $\omega_{cj}$  is the share of commodity  $c$  in industry  $j$ ’s intermediate-input purchases.

**Weights.** The  $\omega_{ej}$  come from the 2012 benchmark *detail* use table of the BEA ([U.S. Bureau of Economic Analysis, n.d.](#)) (“Use of Commodities by Industries, Before Redefinitions, Producers’ Value,” ~402 industries). I read the commodity  $\times$  industry intermediate-use block and map both BEA detail codes to NAICS-4 with the workbook’s own “NAICS Codes” concordance tab. Weights are held *fixed at 2012* so that all time variation in  $\tau^I$  comes from tariffs, not from the production-mix shift that hypothesis H2 predicts (which I instead test directly). An industry’s own output is excluded from its input bundle, so  $\tau^I$  reflects purchased inputs only.

**Coverage and conventions.** Inputs with no output tariff in my manufacturing panel (services, and the small agricultural and mining inputs) carry no rate; I renormalize  $\tau^I$  over the priced inputs, so it is the average tariff on the *tariffable* inputs an industry buys, and the choice of weight denominator cancels. BEA detail represents all apparel (NAICS 315) and all leather (NAICS 316) as single three-digit industries; their six four-digit children inherit the parent’s input mix, so  $\tau^I$  is identical within those two groups. BEA publishes no imported-use (import) matrix at 2012 detail, so the baseline weights total (not imports-only) input use; an imported-input-weighted variant can be approximated from the Schott import penetration and is left as a refinement. Resulting coverage is all 86 manufacturing industries, 1989–2023 for the statutory version and 1989–2024 for the realized one, and the  $\tau^I$  arithmetic is unit-tested. As an economic check, the highest 2018 input tariffs fall on the metal-fabrication industries most exposed to the Section 232 steel and aluminum tariffs (forging and stamping, fabricated metal products), and  $\tau^I$  rises even more than  $\tau^O$  during the 2018–19 trade war; the two are correlated at 0.58, leaving ample independent variation to separate the channels.

## A.6 Productivity measures

The productivity-*level* analysis uses two independently constructed NAICS-4 panels of total factor productivity (TFP).

**Primary: BLS detailed-industry TFP.** The Bureau of Labor Statistics publishes TFP indexes (sectoral output over a Tornqvist index of combined capital, labor, and intermediate inputs) for detailed manufacturing industries ([U.S. Bureau of Labor Statistics, 2026](#)). The `MachineReadable` tab is filtered to four-digit, all-workers rows; the TFP index (base 2017=100) is taken in logs. These are exactly the same 86 four-digit NAICS manufacturing industries as `DISP`, so the TFP panel merges onto the tariff panel with no concordance. The same source supplies the intermediate-input, capital, and labor cost shares (reported as fractions) used in the mechanism test of Section [A.5](#). Coverage is 1987–2023.

**Cross-check: NBER–CES.** The NBER–CES Manufacturing Industry Database ([Becker](#)

et al., 2021), NAICS-2012 vintage (364 six-digit industries; dollar magnitudes 1958–2018, the `tfp5` index populated through 2016), gives an independent TFP index (`tfp5`, five-factor) and the dollar magnitudes (shipments, materials, energy, payroll, value added) from which I form four-digit aggregates: dollar magnitudes are summed to NAICS-4 and the four-digit TFP index is the shipments-weighted mean of the six-digit `tfp5`; materials, energy, and labor shares are formed as ratios to shipments and capital as the residual. Because both TFP series are within-industry indexes with different base years, I compare them on the *within-industry* variation (removing each industry’s mean): the BLS and NBER–CES log-TFP series correlate at 0.77 over the 2,580 overlapping industry-years, validating the primary measure.

## A.7 The merged panel and its validation

The harmonized tariff measures are joined to DiSP on (`naics4, year`), keeping the DiSP panel complete over the covered years (an industry with no imports retains a missing tariff rather than being dropped). The resulting file, `tariff_disp_full.csv`, has **2,838** rows: **1989–2021**  $\times$  the **86** manufacturing industries, carrying  $\tau^R$ , the statutory NTR/non-NTR/301/232 columns and their coverage shares, and all DiSP dispersion moments. NAICS 3328 (a contract-service industry with essentially no traded products) carries a missing tariff in every year by construction.

Because the panel splices independent sources, I inspect each seam and check the result against external aggregates rather than smoothing the joins: the cross-dataset seams, the agreement with known economic magnitudes, the two measures for example industries, the 2018–19 trade-war movements, and a reconciliation of the carried-forward Section 301 schedule against the published Fajgelbaum et al. (2020) data.

### A.7.1 Cross-dataset seams

The panel splices independent sources, so I inspect each seam rather than smooth it. *Statutory, 2001/2002*: across the FRS/DataWeb seam the realized rate is continuous (manufacturing mean 2.07% in both 2001 and 2002), whereas statutory NTR steps from 4.62% to 3.97%. This gap is expected: FRS carries ad valorem equivalents for specific-duty lines while DataWeb currently drops them, and those lines are disproportionately high-tariff agriculture/textile/footwear products. Closing it is a planned refinement.

### A.7.2 Economic checks

Summing calculated duties across all codes reproduces aggregate U.S. customs revenue: for 1989 I obtain \$16.1B in duties and a 3.44% aggregate rate, matching published 1989 collec-

tions. The harmonization *conserves* total customs value exactly (concordance weights sum to one), and in every year zero percent of genuine-industry customs value is unmatched to NAICS. The realized manufacturing tariff declines from about 4.9% (1989) to 2.0% (mid-2000s), tracking the Uruguay Round and subsequent liberalization, before the 2018–19 increase.

### A.7.3 Example industries

Table [A.4](#) contrasts the two measures for four illustrative industries. The realized and statutory rates diverge in economically sensible ways: apparel is highly protected under both; motor vehicles and semiconductors carry low realized rates (free-trade agreements; the Information Technology Agreement) below their statutory NAICS rates; and the 2018–19 actions load onto different statutes by industry (steel via Section 232 and electronics/vehicles via Section 301).

Table A.2: Manufacturing industries in the panel: NAICS 31–32 (food, textiles, apparel, paper, chemicals, and related)

| NAICS | Industry                                              | NAICS | Industry                                                               |
|-------|-------------------------------------------------------|-------|------------------------------------------------------------------------|
| 3111  | Animal food manufacturing                             | 3211  | Sawmills and wood preservation                                         |
| 3112  | Grain and oilseed milling                             | 3212  | Veneer, plywood, and engineered wood product manufacturing             |
| 3113  | Sugar and confectionery products                      | 3219  | Other wood products                                                    |
| 3114  | Fruit and vegetable preserving and specialty food     | 3221  | Pulp, paper, and paperboard mills                                      |
| 3115  | Dairy products                                        | 3222  | Converted paper products                                               |
| 3116  | Animal slaughtering and processing                    | 3231  | Printing and related support activities                                |
| 3117  | Seafood product preparation and packaging             | 3241  | Petroleum and coal products                                            |
| 3118  | Bakeries and tortilla manufacturing                   | 3251  | Basic chemicals                                                        |
| 3119  | Other food products                                   | 3252  | Resin, synthetic rubber, and artificial synthetic fibers and filaments |
| 3121  | Beverages                                             | 3253  | Pesticides, fertilizers, and other agricultural chemicals              |
| 3122  | Tobacco manufacturing                                 | 3254  | Pharmaceuticals and medicine                                           |
| 3131  | Fiber, yarn, and thread mills                         | 3255  | Paint, coatings, and adhesives                                         |
| 3132  | Fabric mills                                          | 3256  | Soaps, cleaning compounds, and toilet preparations                     |
| 3133  | Textile and fabric finishing and fabric coating mills | 3259  | Other chemical products and preparations                               |
| 3141  | Textile furnishings mills                             | 3261  | Plastics products                                                      |
| 3149  | Other textile product mills                           | 3262  | Rubber products                                                        |
| 3151  | Apparel knitting mills                                | 3271  | Clay products and refractories                                         |
| 3152  | Cut and sew apparel                                   | 3272  | Glass and glass products                                               |
| 3159  | Apparel accessories and other apparel manufacturing   | 3273  | Cement and concrete products                                           |
| 3161  | Leather and hide tanning and finishing                | 3274  | Lime and gypsum products                                               |
| 3162  | Footwear manufacturing                                | 3279  | Other nonmetallic mineral products                                     |
| 3169  | Other leather and allied product manufacturing        |       |                                                                        |

Table A.3: Manufacturing industries in the panel: NAICS 33 (metals, machinery, electronics, transportation equipment, and related)

| NAICS | Industry                                                     | NAICS | Industry                                                   |
|-------|--------------------------------------------------------------|-------|------------------------------------------------------------|
| 3311  | Iron and steel mills and ferroalloy production               | 3342  | Communications equipment                                   |
| 3312  | Steel products from purchased steel                          | 3343  | Audio and video equipment manufacturing                    |
| 3313  | Alumina and aluminum production and processing               | 3344  | Semiconductors and other electronic components             |
| 3314  | Nonferrous metal (except aluminum) production and processing | 3345  | Electronic instruments                                     |
| 3315  | Foundries                                                    | 3346  | Manufacturing and reproducing magnetic and optical media   |
| 3321  | Forging and stamping                                         | 3351  | Electric lighting equipment                                |
| 3322  | Cutlery and handtools                                        | 3352  | Household appliances                                       |
| 3323  | Architectural and structural metals                          | 3353  | Electrical equipment                                       |
| 3324  | Boilers, tanks, and shipping containers                      | 3359  | Other electrical equipment and components                  |
| 3325  | Hardware                                                     | 3361  | Motor vehicles                                             |
| 3326  | Spring and wire products                                     | 3362  | Motor vehicle bodies and trailers                          |
| 3327  | Machine shops; turned products; and screws, nuts, and bolts  | 3363  | Motor vehicle parts                                        |
| 3328  | Coating, engraving, heat treating, and allied activities     | 3364  | Aerospace products and parts                               |
| 3329  | Other fabricated metal products                              | 3365  | Railroad rolling stock manufacturing                       |
| 3331  | Agriculture, construction, and mining machinery              | 3366  | Ship and boat building                                     |
| 3332  | Industrial machinery                                         | 3369  | Other transportation equipment                             |
| 3333  | Commercial and service industry machinery                    | 3371  | Household and institutional furniture and kitchen cabinets |
| 3334  | HVAC and commercial refrigeration equipment                  | 3372  | Office furniture (including fixtures)                      |
| 3335  | Metalworking machinery                                       | 3379  | Other furniture related products                           |
| 3336  | Engine, turbine, and power transmission equipment            | 3391  | Medical equipment and supplies                             |
| 3339  | Other general purpose machinery                              | 3399  | Other miscellaneous manufacturing                          |
| 3341  | Computer and peripheral equipment                            |       |                                                            |

Table A.4: Realized and statutory tariff rates for example industries (percent)

| Industry                            | Year | Realized | Stat. NTR | Stat. 301 | Stat. 232 |
|-------------------------------------|------|----------|-----------|-----------|-----------|
| Apparel (3152)                      | 1991 | 17.9     | 13.7      | 0.0       | 0.0       |
|                                     | 2017 | 14.0     | 10.9      | 0.0       | 0.0       |
|                                     | 2021 | 14.7     | 11.3      | 7.1       | 0.0       |
| Motor vehicles (3361)               | 1991 | 1.6      | 4.2       | 0.0       | 0.0       |
|                                     | 2017 | 1.1      | 5.4       | 0.0       | 0.0       |
|                                     | 2021 | 1.3      | 5.6       | 18.7      | 0.0       |
| Semiconductors & electronics (3344) | 1991 | 1.8      | 4.4       | 0.0       | 0.0       |
|                                     | 2017 | 0.1      | 1.5       | 0.0       | 0.0       |
|                                     | 2021 | 2.3      | 1.4       | 24.1      | 0.0       |
| Iron & steel mills (3311)           | 1991 | 4.4      | 5.2       | 0.0       | 0.0       |
|                                     | 2017 | 0.1      | 0.2       | 0.0       | 0.0       |
|                                     | 2021 | 5.0      | 0.3       | 2.6       | 21.6      |

#### A.7.4 The 2018–19 China tariff war

Table A.5 reports manufacturing means around the trade war. Both measures rise, but they capture different objects: the realized rate reflects duties actually collected on imports from all origins (so its Section 301 component is net of exclusions and diluted by non-Chinese imports), while `statutory_301` is the schedule add-on on covered lines, the rate a Chinese exporter of those products faces. Section 301 covers about 86% of manufacturing tariff lines once List 4A is included; its manufacturing-mean add-on jumps from zero (2017) to 14.4% (2019, after List 3’s escalation to 25% and the arrival of List 4A in September) and 17.5% (2021). The realized rate rises in parallel (2.0%  $\rightarrow$  4.6%), confirming the schedule changes show up in collected duties.

#### A.7.5 Reconstruction following the QJE data

Differencing China’s statutory rate in Fajgelbaum et al. (2020) against its 2017 baseline recovers the policy timeline exactly: add-ons of +25% at 2018-07-06 (List 1), +25% at 2018-08-23 (List 2), and +10% at 2018-09-24 (List 3), alongside the +25% steel action at 2018-03-23. Each matches the known effective dates and rates. I further verify the carried-forward Section 301 schedule against the USITC DataWeb Chapter 99 provisions, whose `mfn_text_rate` confirms “the applicable subheading + 25%” for the List 1–3 headings (9903.88.01/03/04) in 2021 and “+ 7.5%” for List 4A (9903.88.15). These checks validate both the membership extracted from the replication data and the schedule applied to it.

Table A.5: Manufacturing-mean tariffs around the 2018–19 trade war (percent)

| Year | Realized | Stat. NTR | Stat. 301 | Stat. 232 | 301 coverage |
|------|----------|-----------|-----------|-----------|--------------|
| 2016 | 2.00     | 3.58      | 0.00      | 0.00      | 0            |
| 2017 | 1.97     | 3.56      | 0.00      | 0.00      | 0            |
| 2018 | 2.72     | 3.56      | 2.92      | 0.47      | 63           |
| 2019 | 4.15     | 3.55      | 14.35     | 0.61      | 86           |
| 2020 | 4.40     | 3.55      | 17.75     | 0.61      | 86           |
| 2021 | 4.63     | 3.74      | 17.54     | 0.61      | 86           |

*Note:* The first four columns are ad valorem tariff *rates* (percent), averaged across the manufacturing (NAICS 31–33) industries. The final column, set off by the vertical rule, is a different object: **301 coverage** is the *extensive* margin, that is, the share of an industry’s eight-digit tariff lines that are subject to *any* Section 301 tariff (here counted by line and averaged across industries), as opposed to the size of the duty. The two are linked by **statutory\_301** = **301 coverage**  $\times$  (average add-on on the covered lines); for example, in 2021 manufacturing,  $17.54\% = 86\% \times 20.4\%$ , the average add-on across covered lines being a mix of the 25% on Lists 1–3 and the 7.5% on List 4A. A low **statutory\_301** with high coverage thus means broad exposure to a modest add-on, while the reverse means a few heavily taxed lines.

## A.8 Flaaen–Pierce exposure measures

The three exposure measures of equation (4) are rebuilt from the HS-10  $\times$  country trade flows of Fajgelbaum et al. (2020), the BEA 2012 detail Use table and NBER–CES shipments, all at 2017 values.

One construction detail matters. The tariff flags in the source files are defined at the *product* level: a Section 301 flag is set for every origin of a listed product, not only for China. The product–country set  $\Omega$  therefore cannot be read off them directly. It is rebuilt by crossing each flag with the countries the action actually applied to: Section 301 against China alone, Sections 232 (steel and aluminum) and 201 (solar panels and washing machines) against all origins. Foreign retaliation is crossed the same way, partner by partner, for China, the European Union, Mexico, Canada, Turkey and Russia.

Two validations. In 2017 the manufacturing sector imported \$2,065 billion, of which the rebuild assigns \$300 billion to the 2018 tariff actions, against a published figure of roughly \$283 billion; the gap is the Section 232 country exemptions, which are not modelled here and which therefore overstate protection slightly for steel- and aluminum-using industries. Manufacturing exports were \$1,199 billion, of which \$97 billion were newly retaliated against, against a published \$121 billion that includes agriculture, which lies outside this panel. Separately, eight of Flaaen and Pierce’s ten most-protected industries appear among the ten most protected here, a check on the product–country set industry by industry rather than only in aggregate.

Five industries report exports exceeding shipments, because the Census export figures include re-exports that were never part of domestic output. None produces a non-positive absorption, so Protection is defined for every industry; the export-share control is winsorized at its 95th percentile before entering equation (5).

## B Additional empirical results

This appendix collects the results that support but do not belong in the main text, organized by the object they concern: the *level* of productivity and its within-industry *distribution*. The same pattern recurs across all three: the own-product (output) tariff lowers productivity and compresses its dispersion, while the input tariff does not. This pattern is shown to survive alternative productivity series, value-added weighting, instruments, placebo timing, and trend controls.

### B.1 Results for the productivity level (Sections 5 and 6)

The main text identifies the level effect with the consumer-good-share instrument. Here I gather the corroborating level evidence in three groups: further results that vary the productivity series, the outcome, and the identifying variation; value-added-weighted estimates; and a battery of timing-placebo tests. Every variant preserves the negative output-tariff effect.

#### B.1.1 Measurement cross-checks for the TFP level

The main text identifies the level decline from the trade war, twice over, and does *not* report an uninstrumented full-sample association as a result. This section collects the checks that bear on *measurement* rather than on identification (whether the association appears in an independently constructed series and in an alternative productivity concept), together with the input-mix mechanism and the timing placebo whose failure is the reason the uninstrumented specification is not read as causal. The regressions here are first-difference associations and should be read as such.

Table A.6 repeats the exercise with the independently constructed NBER–CES five-factor TFP index. The output-tariff coefficient is essentially unchanged ( $-0.0012$  to  $-0.0013$ , significant) and the input-tariff coefficient is again negative and larger ( $-0.0037$  to  $-0.0050$ ), though the shorter NBER–CES sample, whose productivity index ends in 2016, renders it imprecise. The sign and ordering of the two channels do not depend on which productivity series I use.

Total factor productivity nets out all measured inputs; a related alternative outcome is labor productivity (output per hour), which extends two years further than the TFP index. Table A.7 re-runs the statutory specification with the change in log BLS labor productivity, which is available through 2025; the statutory tariff ends in 2023, so the estimation sample runs through 2024, one year beyond the TFP sample and five years past the trade war. The output (competition) tariff enters more strongly than for TFP ( $-0.0015$ , significant at 5 percent when it enters alone) and survives these post-war years. The input (cost) tariff, by contrast, is essentially zero ( $+0.0004$  alone and  $+0.0008$  beside the output tariff, both insignificant), unlike its negative effect on TFP. This is what the input-cost channel implies: by raising the price of intermediates and shifting the input mix, the input tariff lowers the output an industry obtains from all its inputs taken together, which is what total factor productivity measures, without necessarily lowering output per *hour*, because labor productivity absorbs the offsetting change in input intensity. The input channel is thus specific to total factor productivity, while the competition channel operates on both.

Table A.8 probes the H2 mechanism directly: if firms shift away from tariffed inputs, the materials share of output should respond to the input tariff. In first differences it does not: the coefficient is small and insignificant, so support for H2 rests on the reduced-form input-tariff effect rather than on a demonstrated shift in the input mix.

Table A.9 applies a timing placebo: if the lagged tariff change drives TFP, future tariff changes should not. As in the dispersion analysis, the one-year lead is contaminated by the autocorrelation of tariff changes and is uninformative; the uncontaminated test is the two- and three-year leads. Here, unlike for dispersion, those leads do *not* cleanly vanish; several remain significant, and for the input tariff they are same-signed as the lag. I read this as a symptom of the persistence of the 2018–19 trade war: because tariffs rose over several consecutive years, changes at neighboring horizons co-move with the same multi-year movement in TFP, which a lead-based placebo cannot separate from a causal lag. The level results are therefore more exposed to this common-episode concern than the dispersion results, and I read their causal interpretation with corresponding caution.

### B.1.2 Value-added-weighted regressions

The headline TFP regressions in the main text weight every industry-year equally. Here I re-estimate the two main TFP specifications weighting each industry by its value added averaged over all sample years, which gives larger industries more influence. The weight is fixed and time-invariant, so it does not move with the outcome. Table A.10 is the value-added-weighted statutory specification. The qualitative pattern is unchanged: both tariffs lower TFP. Weighting sharpens the input tariff, which reaches 5 percent when it enters alone

Table A.6: Robustness: statutory tariffs and NBER–CES TFP (first differences)

|                                         | (1)<br>Output          | (2)<br>Input        | (3)<br>Both            |
|-----------------------------------------|------------------------|---------------------|------------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0013***<br>(0.0003) |                     | -0.0012***<br>(0.0003) |
| $\Delta$ input tariff $\tau^I$ , lag 1  |                        | -0.0050<br>(0.0045) | -0.0037<br>(0.0045)    |
| Observations                            | 2,210                  | 2,236               | 2,210                  |
| Within $R^2$                            | 0.196                  | 0.200               | 0.196                  |
| Industry FE                             | Yes                    | Yes                 | Yes                    |
| Year FE                                 | Yes                    | Yes                 | Yes                    |

*Notes:* First-differences regressions of the change in the log NBER–CES five-factor TFP index (Becker et al., 2021; ends 2016) on the once-lagged change in the statutory output and input tariffs, with industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2016. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

(−0.0029) and 10 percent beside the output tariff (−0.0024), while the output tariff stays close to its unweighted value (−0.0017 alone, −0.0016 with the input tariff). These are first-difference associations of the kind Section B.1 opens by flagging, and the causal reading of the input channel rests on Section 8.3 rather than here.

### B.1.3 Placebo timing tests: robustness

The timing placebo of Table A.9 replaces the once-lagged tariff change with its one-, two-, and three-year *leads*: if the lagged change reflects a causal effect rather than a pre-existing trend, future tariff changes should not predict current TFP. I report two robustness variants of that test. Table A.11 drops the 2018–19 China trade war: the panel is truncated at 2017, so the leads cannot reach the war years. Table A.12 weights industries by their sample-average value added. In both, the maintained once-lagged effect (column 1) is the object the placebo is meant to discipline.

The placebo leads at horizons two and three are individually significant. If this were mechanical, with the lead merely standing in for the autocorrelated causal change that acts around  $t$ , then entering the lead jointly with the maintained once-lagged change should drive the lead toward zero. Table A.13 shows it does not: in columns 3 and 5 the lead retains essentially its full magnitude and significance alongside the lag, for both channels. The leads therefore carry variation that is not explained by the lagged causal change, consistent with anticipation or with tariffs responding to the industry’s productivity path rather than with simple serial-correlation contamination.

Table A.7: Statutory tariffs and industry labor productivity (first differences)

|                                         | (1)<br>Output         | (2)<br>Input       | (3)<br>Both          |
|-----------------------------------------|-----------------------|--------------------|----------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0015**<br>(0.0007) |                    | -0.0015*<br>(0.0008) |
| $\Delta$ input tariff $\tau^I$ , lag 1  |                       | 0.0004<br>(0.0018) | 0.0008<br>(0.0019)   |
| Observations                            | 2,890                 | 2,924              | 2,890                |
| Within $R^2$                            | 0.120                 | 0.121              | 0.120                |
| Industry FE                             | Yes                   | Yes                | Yes                  |
| Year FE                                 | Yes                   | Yes                | Yes                  |

*Notes:* Dependent variable: change in log BLS labor productivity (sectoral output per hour), 86 four-digit NAICS manufacturing industries, as an alternative outcome to TFP. Output and input statutory tariffs, lagged one year; industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2024. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

#### B.1.4 The Flaaen–Pierce design without detrending

The outcome in equation (5) is log TFP with each industry’s own 2013–17 line removed, so the flat pre-war coefficients of Table 4 are in part what that correction imposes. Table A.14 therefore reports the same specification on the raw series, as Flaaen and Pierce (2024) do in their own appendix.

The pre-war path is sloped. Protection carries coefficients of  $-0.268$  in 2013 and  $-0.203$  in 2015, both significant, so the more exposed industries were gaining on the less exposed ones through 2017. After 2017 the gain stops, and no post-war coefficient in the Protection column is distinguishable from zero. The detrended path in Table 4 describes the same history: a series that flattens while its fitted trend keeps rising reads as a widening decline.

Two considerations recommend reading the design through the detrended path rather than through the raw one. First, the statistic that carries the inference does not depend on the correction. Flaaen and Pierce’s equation (8) places the pre-period and the post-period symmetrically about the base year, so an industry-specific linear trend differences out of it: the Protection entry in the last row of Table A.14,  $-0.2523$  with a standard error of  $0.0893$ , is the number the detrended specification gives, to four decimals. The year-by-year detrended coefficients carry a caveat that this statistic escapes, since the fitted trend rests on five annual points per industry and is extrapolated as much as five years ahead, and the clustered standard errors do not carry its sampling error. Flaaen and Pierce fit theirs on twelve monthly observations inside 2017 and extrapolate at most twenty months, so their first step is the better identified one. Neither first step enters equation (8).

Table A.8: H2 mechanism: change in the input-cost share and the input tariff

|                                         | (1)<br>Interm. share | (2)<br>Materials sh.  |
|-----------------------------------------|----------------------|-----------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | 0.0002<br>(0.0002)   | 0.0004***<br>(0.0001) |
| $\Delta$ input tariff $\tau^I$ , lag 1  | -0.0005<br>(0.0008)  | 0.0007<br>(0.0018)    |
| Observations                            | 2,805                | 2,380                 |
| Within $R^2$                            | 0.104                | 0.138                 |
| Industry FE                             | Yes                  | Yes                   |
| Year FE                                 | Yes                  | Yes                   |

*Notes:* Dependent variable: change in the intermediate-input cost share (BLS, col. 1) or the materials share (NBER-CES, col. 2), on the once-lagged change in the statutory output and input tariffs. A response to  $\tau^I$  is direct evidence that the production mix shifts with input tariffs (H2). Industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2023 in column 1 and 1991–2018 in column 2. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Second, [Flaen and Pierce](#) read pre-war paths of this kind as short-run shocks rather than as a durable industry characteristic that also determines exposure, and the pattern here supports that reading. The three channels approach 2017 from different directions: Protection and InputCost from below, Retaliation from above. A characteristic that both shaped the tariff lists and moved productivity would move the three paths together.

## B.2 Robustness of the dispersion result (Section 7)

The paper’s anchor result is that the output tariff compresses within-industry productivity dispersion. Table [A.15](#) reproduces it under four cuts: the baseline, dropping the 2018–19 China trade war, weighting industries by their sample-average value added, and identifying off the Section 301/232 war add-on alone (the statutory output tariff instrumented by its Section 301/232 add-on, with the statutory input tariff kept as a control). The output-tariff coefficient is stable between  $-0.005$  and  $-0.006$  and stays significant in the first three; the war-IV column is the same sign and magnitude but imprecise, because the dispersion data end in 2021 and so cover only the first years of the war (and, with 85 of 86 industries eventually exposed, there is no never-treated control group). The input tariff is insignificant throughout.

Table [A.16](#) reports the value-added-weighted dispersion regressions in full (the output tariff alone, the input tariff alone, and both together). It is the weighted counterpart of the headline dispersion table. The output tariff continues to compress the spread (about

$-0.0048$ , significant) whether entered alone or with the input tariff, and the input tariff remains insignificant throughout, so the channel split is unchanged by weighting.

The input-tariff null does not depend on measuring the input tariff at statutory rates. Firms paid far less than the schedule implies (Section 8.3), and the input-cost logic concerns the costs they actually bore, so Table A.17 replaces the statutory input tariff with its realized counterpart: the same fixed-2012 cost-weighted average, taken over suppliers' realized rates. Because the realized rate's import-composition denominator responds to the tariff, I instrument it with the statutory input tariff; the first stage is strong ( $F \approx 143$ ). The realized input tariff remains insignificant ( $-0.0025$ , s.e.  $0.0267$ ). The coefficient is per point of the realized rate, which moves far less than the statutory rate, so its standard error is correspondingly larger.

The dispersion outcome throughout is *total factor* productivity dispersion. DiSP also publishes the dispersion of *labor* productivity (output per hour), and each productivity concept in an unweighted and an activity-weighted version. Table A.18 shows that the output-tariff compression is specific to unweighted TFP dispersion: it is absent in labor-productivity dispersion, and absent in the activity-weighted (large-establishment) version of both. The compression appears in unweighted TFP dispersion and not in the labor-productivity measure, and its concentration in the unweighted version locates it among the mass of smaller establishments rather than the large plants that dominate activity. I do not read the first of those contrasts as settling whether the compression is technological. Neither productivity concept is adjusted for capacity utilization, and neither uses establishment-specific prices, so utilization, prices and markups enter these measures as they enter any revenue-based index (Basu et al., 2006; De Loecker and Warzynski, 2012). A fall in utilization concentrated among smaller establishments would compress the unweighted TFP spread with no change in technology, and I cannot separate the two with these data.

Finally, the output-tariff effect is not an artifact of smooth pre-existing industry trends. The headline first-difference specification already absorbs industry fixed effects, i.e. industry-specific linear drift in the level; Table A.19 adds industry-specific linear trends to the differenced equation (industry-specific curvature in the level) for all three outcomes. The output-tariff coefficient is essentially unchanged ( $-0.006$  for dispersion and about  $-0.001$  for each TFP level, all significant, against memo baselines of  $-0.006$  and  $-0.001$ ). The compression of dispersion and the decline in the TFP levels therefore survive absorbing each industry's own trajectory. The input tariff stays insignificant for all three outcomes. I stress, however, that industry trends do *not* dispose of the timing concern: the leads of the output-tariff change remain jointly significant predictors of the current outcome even with industry-specific trends included (for all three outcomes,  $p < 0.01$ ). The pre-trend therefore reflects the endogenous

*timing* of protection rather than a smooth industry trajectory, which is why identification rests on the instrumented and Flaaen–Pierce designs rather than on trend controls.

## References

- Mary Amiti and Jozef Konings. Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. *American Economic Review*, 97(5):1611–1638, 2007.
- Susanto Basu, John G. Fernald, and Miles S. Kimball. Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448, 2006.
- Randy A. Becker, Wayne B. Gray, and Jordan Marvakov. NBER-CES manufacturing industry database. *National Bureau of Economic Research*, 2021. 1958–2018, version 2021a, NAICS 2012 version (nberces5818v1\_n2012), <https://www.nber.org/research/data/nber-ces-manufacturing-industry-database>.
- Cindy Cunningham, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. Dispersion in dispersion: Measuring establishment-level differences in productivity. BLS Working Paper 530, U.S. Bureau of Labor Statistics, 2022.
- Jan De Loecker and Frederic Warzynski. Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471, 2012.
- Fabian Eckert, Teresa C. Fort, Peter K. Schott, and Natalie J. Yang. Imputing missing values in the US Census Bureau’s county business patterns. NBER Working Paper 26632, National Bureau of Economic Research, 2020. Distributes the public time-consistent NAICS concordance (“EFSY”); panel placed on a 2012 NAICS basis.
- Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. Replication data for: The return to protectionism. Harvard Dataverse, V1, 2019. Accessed June 2026. <https://doi.org/10.7910/DVN/KSOVSE>.
- Pablo D. Fajgelbaum, Pinelopi K. Goldberg, Patrick J. Kennedy, and Amit K. Khandelwal. The return to protectionism. *Quarterly Journal of Economics*, 135(1):1–55, 2020.
- Robert C. Feenstra, John Romalis, and Peter K. Schott. U.S. imports, exports, and tariff data, 1989–2001. NBER Working Paper 9387, National Bureau of Economic Research, 2002. Ad valorem equivalent NTR (MFN) and non-NTR tariff rates at the 8-digit HS level; distributed as `tar_val.dta` on Schott’s data page.

- Aaron Flaaen and Justin R. Pierce. Disentangling the effects of the 2018–2019 tariffs on a globally connected U.S. manufacturing sector. *Review of Economics and Statistics*, 2024. [https://doi.org/10.1162/rest\\_a\\_01498](https://doi.org/10.1162/rest_a_01498).
- Teresa C. Fort. NAICS 2002 code assignments to the longitudinal business database. Unpublished. Method underlying the longitudinally consistent NAICS codes in the DiSP data; see also Fort and Klimek (2018)., 2013.
- Teresa C. Fort and Shawn D. Klimek. The effects of industry classification changes on US employment composition. CES Working Paper CES-WP-18-28, U.S. Census Bureau, Center for Economic Studies, 2018.
- Office of the United States Trade Representative. Notice of modification of section 301 action: China’s acts, policies, and practices related to technology transfer, intellectual property, and innovation. 84 Fed. Reg. 43304 (August 20, 2019), Annex A, 2019. Product list for Section 301 List 4A. <https://www.govinfo.gov/content/pkg/FR-2019-08-20/pdf/2019-17865.pdf>.
- Justin R. Pierce and Peter K. Schott. A concordance between ten-digit U.S. Harmonized System codes and SIC/NAICS product classes and industries. *Journal of Economic and Social Measurement*, 37(1–2):61–96, 2012.
- Justin R. Pierce and Peter K. Schott. The surprisingly swift decline of US manufacturing employment. *American Economic Review*, 106(7):1632–1662, 2016.
- Peter K. Schott. The relative sophistication of Chinese exports. *Economic Policy*, 23(53): 5–49, 2008. Reference cited by Schott for the U.S. import/export data used here.
- Peter K. Schott. U.S. HS-level imports and exports, 1989–2024. Data, version v2025.4.21, [https://sompks4.github.io/sub\\_data.html](https://sompks4.github.io/sub_data.html), 2025. Annual import-detail files `imp_det1_<YYYY>_12n.dta`, converted from U.S. Census Bureau records. Accessed June 2026.
- The White House. Presidential proclamations 9705 (steel) and 9704 (aluminum): Adjusting imports under section 232. Federal Register 83 FR 11625 and 83 FR 11619; effective 2018-03-23 (steel 25%, aluminum 10%), 2018.
- Petia Topalova and Amit Khandelwal. Trade liberalization and firm productivity: The case of india. *Review of Economics and Statistics*, 93(3):995–1009, 2011.

U.S. Bureau of Economic Analysis. Input-output accounts: Use table, before redefinitions, producers' prices, detail, and import matrices, before redefinitions, detail. Industry Economic Accounts, n.d. Benchmark years 2007, 2012 and 2017; the 2012 tables are used. Accessed June 2026. <https://apps.bea.gov/industry/Release/XLSX/>.

U.S. Bureau of Labor Statistics. Dispersion statistics on productivity (DiSP). Experimental data product, U.S. Bureau of Labor Statistics and U.S. Census Bureau, 2024. Updated September 12, 2024. Covers 86 four-digit NAICS manufacturing industries, 1987–2021, on a longitudinally consistent 2017 NAICS basis. <https://www.bls.gov/productivity/articles-and-research/dispersion-statistics-on-productivity/>.

U.S. Bureau of Labor Statistics. Total factor productivity for manufacturing and transportation detailed industries. Office of Productivity and Technology, detailed-industry tables, 2026. Data released August 26, 2026. <https://www.bls.gov/productivity/tables/>.

U.S. International Trade Commission. Annual tariff data. USITC DataWeb, 2025. Annual tariff database files, 2002–2023. Accessed March 2025. <https://dataweb.usitc.gov/tariff/annual>.

Table A.9: Timing placebo: leads of the tariff change and the change in TFP  
*Panel A. Output tariff*

|                                        | (1)<br>Lag (base)      | (2)<br>Lead 1         | (3)<br>Lead 2         | (4)<br>Lead 3          |
|----------------------------------------|------------------------|-----------------------|-----------------------|------------------------|
| $\Delta$ output tariff, lag 1 (causal) | -0.0013***<br>(0.0004) |                       |                       |                        |
| $\Delta$ output tariff, lead 1         |                        | 0.0015***<br>(0.0004) |                       |                        |
| $\Delta$ output tariff, lead 2         |                        |                       | 0.0020***<br>(0.0006) |                        |
| $\Delta$ output tariff, lead 3         |                        |                       |                       | -0.0019***<br>(0.0005) |
| Observations                           | 2,805                  | 2,890                 | 2,890                 | 2,805                  |
| Within $R^2$                           | 0.186                  | 0.168                 | 0.163                 | 0.138                  |
| Industry FE                            | Yes                    | Yes                   | Yes                   | Yes                    |
| Year FE                                | Yes                    | Yes                   | Yes                   | Yes                    |

*Panel B. Input tariff*

|                                       | (1)<br>Lag (base)     | (2)<br>Lead 1        | (3)<br>Lead 2          | (4)<br>Lead 3         |
|---------------------------------------|-----------------------|----------------------|------------------------|-----------------------|
| $\Delta$ input tariff, lag 1 (causal) | -0.0024**<br>(0.0012) |                      |                        |                       |
| $\Delta$ input tariff, lead 1         |                       | 0.0026**<br>(0.0012) |                        |                       |
| $\Delta$ input tariff, lead 2         |                       |                      | -0.0035***<br>(0.0012) |                       |
| $\Delta$ input tariff, lead 3         |                       |                      |                        | -0.0023**<br>(0.0009) |
| Observations                          | 2,838                 | 2,924                | 2,924                  | 2,838                 |
| Within $R^2$                          | 0.190                 | 0.172                | 0.165                  | 0.139                 |
| Industry FE                           | Yes                   | Yes                  | Yes                    | Yes                   |
| Year FE                               | Yes                   | Yes                  | Yes                    | Yes                   |

*Notes:* First-differences regressions of the change in log BLS TFP on the lagged change (column 1, the maintained effect) or on the change one, two, and three years in the future (columns 2–4) of the indicated statutory tariff, output or input, with industry and year fixed effects; SEs clustered by industry. A clean placebo lead is indistinguishable from zero. Because  $\Delta\tau$  is autocorrelated at one year, the two- and three-year leads are the uncontaminated horizons. Estimation years, in both panels: 1991–2023 in column 1, 1989–2022 in column 2, 1988–2021 in column 3, and 1988–2020 in column 4.  
 \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.10: Statutory tariffs and industry TFP, value-added-weighted (first differences)

|                                         | (1)<br>Output          | (2)<br>Input          | (3)<br>Both            |
|-----------------------------------------|------------------------|-----------------------|------------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0017***<br>(0.0006) |                       | -0.0016***<br>(0.0006) |
| $\Delta$ input tariff $\tau^I$ , lag 1  |                        | -0.0029**<br>(0.0012) | -0.0024*<br>(0.0013)   |
| Observations                            | 2,805                  | 2,838                 | 2,805                  |
| Within $R^2$                            | 0.189                  | 0.189                 | 0.190                  |
| Industry FE                             | Yes                    | Yes                   | Yes                    |
| Year FE                                 | Yes                    | Yes                   | Yes                    |

*Notes:* The headline statutory specification (change in log BLS TFP on the once-lagged change in the output and input statutory tariffs, with industry and year fixed effects and SEs clustered by industry), weighting each industry by its sample-average value added. Estimation years: 1991–2023. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.11: Timing placebo excluding the 2018–19 China trade war (change in TFP)

| <i>Panel A. Output tariff</i>          |                        |                        |                       |                        |
|----------------------------------------|------------------------|------------------------|-----------------------|------------------------|
|                                        | (1)<br>Lag (base)      | (2)<br>Lead 1          | (3)<br>Lead 2         | (4)<br>Lead 3          |
| $\Delta$ output tariff, lag 1 (causal) | -0.0011***<br>(0.0004) |                        |                       |                        |
| $\Delta$ output tariff, lead 1         |                        | 0.0013**<br>(0.0006)   |                       |                        |
| $\Delta$ output tariff, lead 2         |                        |                        | 0.0022***<br>(0.0007) |                        |
| $\Delta$ output tariff, lead 3         |                        |                        |                       | -0.0022***<br>(0.0004) |
| Observations                           | 2,295                  | 2,380                  | 2,380                 | 2,295                  |
| Within $R^2$                           | 0.131                  | 0.137                  | 0.140                 | 0.138                  |
| Industry FE                            | Yes                    | Yes                    | Yes                   | Yes                    |
| Year FE                                | Yes                    | Yes                    | Yes                   | Yes                    |
| <i>Panel B. Input tariff</i>           |                        |                        |                       |                        |
|                                        | (1)<br>Lag (base)      | (2)<br>Lead 1          | (3)<br>Lead 2         | (4)<br>Lead 3          |
| $\Delta$ input tariff, lag 1 (causal)  | -0.0023<br>(0.0043)    |                        |                       |                        |
| $\Delta$ input tariff, lead 1          |                        | -0.0145***<br>(0.0043) |                       |                        |
| $\Delta$ input tariff, lead 2          |                        |                        | -0.0007<br>(0.0034)   |                        |
| $\Delta$ input tariff, lead 3          |                        |                        |                       | -0.0019<br>(0.0034)    |
| Observations                           | 2,322                  | 2,408                  | 2,408                 | 2,322                  |
| Within $R^2$                           | 0.135                  | 0.143                  | 0.140                 | 0.138                  |
| Industry FE                            | Yes                    | Yes                    | Yes                   | Yes                    |
| Year FE                                | Yes                    | Yes                    | Yes                   | Yes                    |

*Notes:* The timing placebo estimated on the pre-trade-war panel (years through 2017, truncated so the tariff-change leads do not reach the 2018–19 actions). Panel A uses the statutory output tariff, Panel B the statutory input tariff; column 1 is the maintained once-lagged change and columns 2–4 the one-, two-, and three-year leads. Dependent variable: change in log BLS TFP; industry and year fixed effects; SEs clustered by industry. Estimation years, in both panels: 1991–2017 in column 1, 1989–2016 in column 2, 1988–2015 in column 3, and 1988–2014 in column 4. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.12: Timing placebo, value-added-weighted (change in TFP)

*Panel A. Output tariff*

|                                        | (1)<br>Lag (base)      | (2)<br>Lead 1         | (3)<br>Lead 2         | (4)<br>Lead 3          |
|----------------------------------------|------------------------|-----------------------|-----------------------|------------------------|
| $\Delta$ output tariff, lag 1 (causal) | -0.0017***<br>(0.0006) |                       |                       |                        |
| $\Delta$ output tariff, lead 1         |                        | 0.0013***<br>(0.0004) |                       |                        |
| $\Delta$ output tariff, lead 2         |                        |                       | 0.0025***<br>(0.0005) |                        |
| $\Delta$ output tariff, lead 3         |                        |                       |                       | -0.0018***<br>(0.0006) |
| Observations                           | 2,805                  | 2,890                 | 2,890                 | 2,805                  |
| Within $R^2$                           | 0.189                  | 0.174                 | 0.176                 | 0.151                  |
| Industry FE                            | Yes                    | Yes                   | Yes                   | Yes                    |
| Year FE                                | Yes                    | Yes                   | Yes                   | Yes                    |

*Panel B. Input tariff*

|                                       | (1)<br>Lag (base)     | (2)<br>Lead 1         | (3)<br>Lead 2         | (4)<br>Lead 3        |
|---------------------------------------|-----------------------|-----------------------|-----------------------|----------------------|
| $\Delta$ input tariff, lag 1 (causal) | -0.0029**<br>(0.0012) |                       |                       |                      |
| $\Delta$ input tariff, lead 1         |                       | 0.0049***<br>(0.0012) |                       |                      |
| $\Delta$ input tariff, lead 2         |                       |                       | -0.0034**<br>(0.0015) |                      |
| $\Delta$ input tariff, lead 3         |                       |                       |                       | -0.0026*<br>(0.0014) |
| Observations                          | 2,838                 | 2,924                 | 2,924                 | 2,838                |
| Within $R^2$                          | 0.189                 | 0.178                 | 0.172                 | 0.150                |
| Industry FE                           | Yes                   | Yes                   | Yes                   | Yes                  |
| Year FE                               | Yes                   | Yes                   | Yes                   | Yes                  |

*Notes:* The timing placebo weighting each industry by its sample-average value added. Panel A uses the statutory output tariff, Panel B the statutory input tariff; column 1 is the maintained once-lagged change and columns 2–4 the one-, two-, and three-year leads. Dependent variable: change in log BLS TFP; industry and year fixed effects; SEs clustered by industry. Estimation years, in both panels: 1991–2023 in column 1, 1989–2022 in column 2, 1988–2021 in column 3, and 1988–2020 in column 4.

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.13: Conditional placebo: tariff-change leads with and without the maintained causal lag (change in TFP)

| <i>Panel A. Output tariff</i>          |                        |                        |                        |                        |                        |
|----------------------------------------|------------------------|------------------------|------------------------|------------------------|------------------------|
|                                        | (1)<br>Lag             | (2)<br>Lead 2          | (3)<br>Lead 2 +lag     | (4)<br>Lead 3          | (5)<br>Lead 3 +lag     |
| $\Delta$ output tariff, lag 1 (causal) | -0.0013***<br>(0.0004) |                        | -0.0011*<br>(0.0006)   |                        | -0.0016***<br>(0.0005) |
| $\Delta$ output tariff, lead 2         |                        | 0.0020***<br>(0.0006)  | 0.0019***<br>(0.0007)  |                        |                        |
| $\Delta$ output tariff, lead 3         |                        |                        |                        | -0.0019***<br>(0.0005) | -0.0021***<br>(0.0005) |
| Observations                           | 2,805                  | 2,890                  | 2,635                  | 2,805                  | 2,550                  |
| Within $R^2$                           | 0.186                  | 0.163                  | 0.166                  | 0.138                  | 0.143                  |
| Industry FE                            | Yes                    | Yes                    | Yes                    | Yes                    | Yes                    |
| Year FE                                | Yes                    | Yes                    | Yes                    | Yes                    | Yes                    |
| <i>Panel B. Input tariff</i>           |                        |                        |                        |                        |                        |
|                                        | (1)<br>Lag             | (2)<br>Lead 2          | (3)<br>Lead 2 +lag     | (4)<br>Lead 3          | (5)<br>Lead 3 +lag     |
| $\Delta$ input tariff, lag 1 (causal)  | -0.0024**<br>(0.0012)  |                        | -0.0029**<br>(0.0014)  |                        | -0.0024*<br>(0.0014)   |
| $\Delta$ input tariff, lead 2          |                        | -0.0035***<br>(0.0012) | -0.0035***<br>(0.0012) |                        |                        |
| $\Delta$ input tariff, lead 3          |                        |                        |                        | -0.0023**<br>(0.0009)  | -0.0021**<br>(0.0009)  |
| Observations                           | 2,838                  | 2,924                  | 2,666                  | 2,838                  | 2,580                  |
| Within $R^2$                           | 0.190                  | 0.165                  | 0.168                  | 0.139                  | 0.141                  |
| Industry FE                            | Yes                    | Yes                    | Yes                    | Yes                    | Yes                    |
| Year FE                                | Yes                    | Yes                    | Yes                    | Yes                    | Yes                    |

*Notes:* Each lead is entered alone and then jointly with the maintained once-lagged change. Panel A uses the statutory output tariff, Panel B the statutory input tariff. Columns: (1) causal lag alone; (2) lead 2 alone; (3) lead 2 with the lag; (4) lead 3 alone; (5) lead 3 with the lag. A lead that were merely proxying for the autocorrelated causal change would collapse toward zero in columns 3 and 5. Dependent variable: change in log BLS TFP; industry and year fixed effects; SEs clustered by industry. Estimation years, in both panels: 1991–2023 in column 1, 1988–2021 in column 2, 1991–2021 in column 3, 1988–2020 in column 4, and 1991–2020 in column 5. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.14: The Flaaen–Pierce design estimated without removing the pre-war trend

|                    | Log TFP, not detrended |                        |                     |
|--------------------|------------------------|------------------------|---------------------|
|                    | Protection             | InputCost              | Retaliation         |
| 2013               | -0.2676***<br>(0.0856) | -0.1015<br>(0.1916)    | -0.2148<br>(0.5516) |
| 2014               | -0.2603**<br>(0.1032)  | 0.0470<br>(0.2223)     | 0.1470<br>(0.6515)  |
| 2015               | -0.2028***<br>(0.0757) | -0.3479**<br>(0.1521)  | 0.2080<br>(0.5044)  |
| 2016               | -0.0809<br>(0.0610)    | -0.6352***<br>(0.1225) | 0.4309<br>(0.3051)  |
| 2017               |                        | 0 (base year)          |                     |
| 2018               | 0.0269<br>(0.0427)     | 0.1634<br>(0.1191)     | 0.0468<br>(0.2636)  |
| 2019               | -0.0494<br>(0.0806)    | -0.1030<br>(0.1756)    | -0.0241<br>(0.6347) |
| 2020               | 0.0206<br>(0.1259)     | -0.4256**<br>(0.1791)  | 0.6694<br>(0.8245)  |
| 2021               | -0.0495<br>(0.1200)    | -0.1457<br>(0.2387)    | -0.3637<br>(0.7447) |
| 2022               | -0.0824<br>(0.1544)    | -0.3929<br>(0.2615)    | -0.4053<br>(1.0345) |
| Equation (8), 2019 | -0.2523***<br>(0.0893) | -0.4509<br>(0.2880)    | 0.1839<br>(0.4085)  |
| Observations       |                        | 850                    |                     |

*Notes:* Year-by-year coefficients from equation (5), estimated on log BLS TFP with no trend removed, 2017 omitted as the base year. All three channels are measured as coverage shares, and the specification carries the baseline import- and export-share controls interacted with every year dummy, industry and year fixed effects, and standard errors clustered by industry. The last row is [Flaaen and Pierce](#)’s equation (8) at the 2019 endpoint,  $\gamma_{2019} + \gamma_{2015}$ , which nets the post-war move against the pre-war one; because the two years sit symmetrically about the base year, it is invariant to the detrending and reproduces the detrended estimate of Table 4 exactly. Estimation years: 2013–2022, 85 industries. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.15: Robustness of the output-tariff effect on 90–10 productivity dispersion

|                                         | (1)<br>Baseline       | (2)<br>Excl. war      | (3)<br>VA-weighted    | (4)<br>War-IV         |
|-----------------------------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0059**<br>(0.0024) | -0.0047**<br>(0.0018) | -0.0050**<br>(0.0020) | -0.0098**<br>(0.0050) |
| $\Delta$ input tariff $\tau^I$ , lag 1  | 0.0013<br>(0.0095)    | 0.0194<br>(0.0211)    | 0.0051<br>(0.0092)    | 0.0027<br>(0.0095)    |
| Estimator                               | OLS                   | OLS                   | OLS                   | IV                    |
| First-stage $F$                         |                       |                       |                       | 12966.10              |
| Observations                            | 2,635                 | 2,295                 | 2,635                 | 2,635                 |
| Industry FE                             | Yes                   | Yes                   | Yes                   | Yes                   |
| Year FE                                 | Yes                   | Yes                   | Yes                   | Yes                   |

*Notes:* First-differences regressions of the change in the 90–10 dispersion of within-industry log TFP on the once-lagged change in the statutory output and input tariffs, with industry and year fixed effects and SEs clustered by industry. Column 1 is the baseline; column 2 drops years after 2017 (the trade war); column 3 weights by sample-average value added; column 4 instruments the statutory output tariff with its Section 301/232 add-on (the statutory input tariff is treated as exogenous). Estimation years: 1991–2021, except 1991–2017 in column 2. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.16: Output and input tariffs and 90–10 productivity dispersion, value-added-weighted (first differences)

|                                         | (1)<br>Output         | (2)<br>Input       | (3)<br>Both           |
|-----------------------------------------|-----------------------|--------------------|-----------------------|
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0048**<br>(0.0023) |                    | -0.0050**<br>(0.0020) |
| $\Delta$ input tariff $\tau^I$ , lag 1  |                       | 0.0029<br>(0.0097) | 0.0051<br>(0.0092)    |
| Observations                            | 2,635                 | 2,635              | 2,635                 |
| Within $R^2$                            | 0.059                 | 0.057              | 0.059                 |
| Industry FE                             | Yes                   | Yes                | Yes                   |
| Year FE                                 | Yes                   | Yes                | Yes                   |

*Notes:* The dispersion specification weighting each industry by its sample-average value added. Dependent variable: change in the 90–10 dispersion of within-industry log TFP, on the once-lagged change in the statutory output and input tariffs; industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2021. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.17: The input tariff at realized rates and 90–10 productivity dispersion (first differences)

|                                        | (1)<br>Statutory    | (2)<br>Realized     |
|----------------------------------------|---------------------|---------------------|
| $\Delta$ statutory input tariff, lag 1 | -0.0009<br>(0.0099) |                     |
| $\Delta$ realized input tariff, lag 1  |                     | -0.0025<br>(0.0267) |
| Estimator                              | OLS                 | IV                  |
| First-stage $F$                        |                     | 143.19              |
| Observations                           | 2,635               | 2,635               |
| Industry FE                            | Yes                 | Yes                 |
| Year FE                                | Yes                 | Yes                 |

*Notes:* Dependent variable: change in the 90–10 dispersion of within-industry log TFP. Column 1 regresses it on the once-lagged change in the statutory input tariff. Column 2 replaces that with the once-lagged change in the realized input tariff (the same fixed-2012 input-cost-weighted average, taken over suppliers’ realized rates), instrumented by the once-lagged change in the statutory input tariff. Industry and year fixed effects; SEs clustered by industry. Estimation years: 1991–2021. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.18: Output tariff and 90–10 productivity dispersion: TFP vs. labor productivity, unweighted vs. activity-weighted

|                                         | TFP dispersion        |                     | Labor-prod. dispersion |                     |
|-----------------------------------------|-----------------------|---------------------|------------------------|---------------------|
|                                         | (1)<br>Unweighted     | (2)<br>Activity-wtd | (3)<br>Unweighted      | (4)<br>Activity-wtd |
| $\Delta$ output tariff $\tau^O$ , lag 1 | -0.0059**<br>(0.0024) | 0.0005<br>(0.0015)  | 0.0035<br>(0.0028)     | 0.0039<br>(0.0028)  |
| $\Delta$ input tariff $\tau^I$ , lag 1  | 0.0013<br>(0.0095)    | -0.0026<br>(0.0032) | 0.0137<br>(0.0091)     | 0.0016<br>(0.0053)  |
| Observations                            | 2,635                 | 2,635               | 2,635                  | 2,635               |
| Within $R^2$                            | 0.038                 | 0.033               | 0.071                  | 0.033               |
| Industry FE                             | Yes                   | Yes                 | Yes                    | Yes                 |
| Year FE                                 | Yes                   | Yes                 | Yes                    | Yes                 |

*Notes:* First-differences regressions of the change in the 90–10 dispersion of log TFP (columns 1–2) or of log labor productivity (columns 3–4) on the once-lagged change in the statutory output and input tariffs; industry and year fixed effects; SEs clustered by industry. “Unweighted” counts each establishment equally; “activity-weighted” weights by the productivity denominator (the input index for TFP, hours for labor productivity). Estimation years: 1991–2021. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .

Table A.19: Robustness to industry-specific trends

|                                | (1)<br>Dispersion 90–10 | (2)<br>TFP (BLS)       | (3)<br>TFP (NBER–CES)  |
|--------------------------------|-------------------------|------------------------|------------------------|
| $\Delta$ output tariff         | -0.0063**<br>(0.0025)   | -0.0010***<br>(0.0004) | -0.0012***<br>(0.0003) |
| $\Delta$ input tariff          | 0.0008<br>(0.0097)      | -0.0019<br>(0.0012)    | -0.0031<br>(0.0049)    |
| output-tariff coef., no trends | -0.0059                 | -0.0012                | -0.0012                |
| Observations                   | 2,635                   | 2,805                  | 2,210                  |
| Industry FE                    | Yes                     | Yes                    | Yes                    |
| Year FE                        | Yes                     | Yes                    | Yes                    |
| Industry-specific trends       | Yes                     | Yes                    | Yes                    |

*Notes:* First-differences regressions of the change in the outcome (the 90–10 TFP dispersion, or the BLS or NBER–CES log TFP level) on the once-lagged change in the statutory output and input tariffs, with industry and year fixed effects, industry-specific linear trends, and SEs clustered by industry. The memo row reports the output-tariff coefficient from the corresponding specification *without* industry-specific trends. Estimation years: 1991–2021 in column 1, 1991–2023 in column 2, and 1991–2016 in column 3. \*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$ .