

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

Luca Guerrieri

Preliminary — September 14, 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 7.5 percentage points apart, and across that gap an industry at the top quartile ends 2019 with total factor productivity about 3.3 percent below one at the bottom. This is a relative comparison across industries, not an aggregate. A second design, which requires no instrument, puts the same contrast at 2.1 percent. The effect operates through the own-product tariff rather than input costs.

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 standard trade theory. 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). Inputs are misallocated and the bundle is worth less, but technology is untouched and the economy stays on its frontier (Kehoe and Ruhl, 2008). With total factor productivity held fixed in this way, 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 leaves out. As illustrated in Figure 1(b), production falls inside the frontier, so the same inputs yield less output; in other words, higher tariffs compress TFP.



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 frontier. The loss is one of allocative efficiency: technology is unchanged, and this is the whole cost of protection when total factor productivity is exogenous (Kehoe and Ruhl, 2008). Panel (b) adds the margin this paper finds: production also moves from B to C, *inside* the frontier, so a given bundle of inputs yields less output.

Protection is not handed out at random. Tariffs go to industries for reasons (import

pressure, political organization, decline) that are themselves correlated with productivity, 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 and the attendant political backlash (Fajgelbaum et al., 2020; Cavallo et al., 2021; Bown, 2021).¹ For reasons of political optics rather than of the industry’s own productivity, an industry’s tariff exposure therefore depended heavily on its *consumer-good share*: the fraction 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.

The instrument is strong (first-stage $F \approx 29$) and the estimate is sharp: in the clean pre-pandemic window a one-percentage-point increase in the statutory output tariff lowers an industry’s TFP by about 0.4 percent. Since the quartiles of the 2017–19 tariff increase are 11.6 and 19.1 points, that amounts to an industry at the top quartile of exposure ending 3.3 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 land? Separately, and

¹The 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).

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 *gap* between two percentiles, so a spread that does not change is consistent with both of its ends shifting together; what the estimates locate is the narrowing, not the position of any percentile. 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. A large literature finds that exposure to foreign rivals disciplines domestic producers, reallocating output toward more productive firms and pressing the survivors to cut slack (Melitz, 2003; Pavcnik, 2002; Trefler, 2004), and raising innovation and technology adoption (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 (2016): 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. 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 frontier of Figure 1(b) rather than the movement along it that the standard model allows.

The paper speaks to several literatures, laid out in Section 2. It studies trade liberalization and firm productivity. [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* 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 the two channels a tariff can operate through 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 [Flaaten 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 lands, 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 \(2016\)](#) 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. It also weighs two readings of the productivity nulls, 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 the appendix.

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](#)). I borrow their input-tariff construction, 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 workhorse trade model a tariff is a price phenomenon: by raising the relative price of imported inputs it leads firms to substitute away from them and onto a more costly input mix than they would otherwise choose, so production costs rise and real income, demand, and output fall. 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 the cost of protection is a movement *along* a fixed production frontier to a distorted input mix, not a movement *of* the frontier.² My evidence both confirms and departs from this benchmark. The input mix does respond as the standard model predicts; the [Pierce and Schott \(2016\)](#) quasi-experiment shows a trade shock reallocating inputs across factors (Section 8.2). But 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 a given bundle of inputs yields less output. Protection is costly not only through the misallocation the standard model describes (the wrong bundle of inputs, at unchanged

²I 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.

technology) but through a second channel it does not contain: less output from the bundle actually used.

The 2018–19 trade war. My setting is the recent U.S. tariff episode. The closest paper is [Flaen 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. Trade liberalization raised plant productivity and reallocated output toward more productive producers in [Pavcnik \(2002\)](#), [Trefler \(2004\)](#), and the selection model of [Melitz \(2003\)](#); [Bloom et al. \(2016\)](#) find that Chinese import competition raised European firms’ productivity, patenting, and IT adoption. Yet [Autor et al. \(2020\)](#) show that the same competition *reduced* patenting among exposed U.S. firms, and [Pierce and Schott \(2016\)](#) 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. 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 2024, 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; the construction, harmonization, and validation are detailed in the data appendix.

Table 1: Data sources and coverage

Series	Concept	Source	Coverage
Output & input tariff, effective	Duties over customs value	Schott (2025)	1989–2024
Output & input tariff, statutory	Schedule rate (MFN, Col. 2, Sec. 232/301)	Feenstra et al. (2002) ; USITC DataWeb; Fajgelbaum et al. (2020)	1989–2023
Input–output weights	2012 benchmark detail use table	BEA (U.S. Bureau of Economic Analysis, 2018)	2012 (fixed)
Total factor productivity	Detailed-industry TFP index	BLS (U.S. Bureau of Labor Statistics, 2025)	1987–2022
Labor productivity	Real sectoral output per hour	BLS	1987–2024
Total factor productivity (cross-check)	Five-factor TFP index	NBER–CES (Becker et al., 2021)	1958–2016
Productivity dispersion	90–10 establishment spread	BLS–Census DiSP (Cunningham et al., 2022 ; U.S. Bureau of Labor Statistics, 2024)	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. Construction and validation are detailed in the data appendix.

3.1 Tariff measures

I measure an industry’s *output tariff*, the tariff on the goods it sells, in two complementary ways. The **effective** (or realized) rate is the duty actually collected as a share of the value imported,

$$\tau_{jt}^R = \frac{\sum_{i \in j} D_{it}}{\sum_{i \in j} V_{it}}, \quad (1)$$

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). Because it covers imports from all origins and nets out duty-free entry under free-trade agreements and preference programs, the effective rate sits well below the schedule rate. The **statutory** rate is the legal ad valorem schedule rate: the most-favored-nation (Column 1) rate plus the 2018–19 Section 232 and Section 301 add-ons, as a simple average across the eight-digit tariff lines in each industry. China has paid most-favored-nation rates throughout the sample, and the Section 301 add-ons apply only to Chinese goods, so from 2018 on the statutory rate is the schedule rate facing Chinese-origin goods. The Section 232 add-on is applied to all origins, ignoring country exemptions. The statutory rate omits Section 301 List 4A, the Section 201 safeguards, and antidumping and countervailing duties, and it does not net out exclusions granted to individual importers. The Column 2 (non-NTR) rate does not enter it; it 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 for 2002–2023, with the Section 301 product lists drawn from [Fajgelbaum et al. \(2020\)](#). The two answer different questions: duties paid versus the policy on the books. Every regression reported in the paper uses the statutory rate, for the reason given in Section 4; the effective rate enters only descriptively, in Table 2, Figure 2(a), and the comparison of realized and statutory input tariffs in Section 8.3.

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 τ_{jt}^O with an *input tariff* that summarizes the tariffs on the goods it purchases,

$$\tau_{jt}^I = \sum_c \omega_{cj} \tau_{ct}^O, \quad (2)$$

the input-cost-weighted average of supplying industries’ output tariffs, following [Amiti and Konings \(2007\)](#) and [Topalova and Khandelwal \(2011\)](#). The weights ω_{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, 2018](#)) and *held fixed* so that all variation in τ^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 effective and statutory 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, 2025](#)): 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 1987–2022. The same source supplies the intermediate-input, capital, and labor cost shares used in the mechanism test of Section 5. As an alternative outcome with a longer recent reach I also use BLS labor productivity (real sectoral output per hour), available through 2024.

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

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

	N	Mean	SD	p10	Median	p90
<i>Levels</i>						
Output tariff τ^O (statutory, pp)	2,805	5.856	7.607	0.718	3.626	13.653
Input tariff τ^I (statutory, pp)	2,805	5.305	5.444	1.982	3.661	10.124
Output tariff τ^O (effective, pp)	2,805	2.679	3.142	0.087	1.622	6.379
Input tariff τ^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>						
Δ output tariff (pp)	2,720	0.463	2.571	-0.364	0.000	0.615
Δ input tariff (pp)	2,720	0.542	2.340	-0.303	-0.000	0.670
Δ 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.

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 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

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. 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.

Figure 3: Productivity and its dispersion, 1989–2024.



Notes: Cross-industry mean of the total factor productivity index (through 2022) and the labor productivity index (through 2024), 2017=100 (left), and the mean 90–10 within-industry dispersion of log TFP (through 2021, right).


channels in the analysis that follows.

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* τ_{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 ω_{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 τ^O and τ^I denote those statutory measures from here on. 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

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.

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.10), 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.

5 Tariffs and the level of productivity

This section delivers the paper’s central estimate: the effect of the output tariff on the level of industry productivity, identified from the way the 2018–19 tariffs were assigned rather than from who ended up protected. The instrument is the industry’s consumer-good share; the finding is that across the 7.5-percentage-point gap between the quartiles of the tariff increase, an industry at the top quartile ends 2019 with TFP about 3.3 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 and the political backlash, 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 of political optics rather than productivity. The design follows Pierce and Schott (2016) 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 political motive for the sequencing 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 17 points between 2017 and 2019, against about 12 points for the top half (first-stage $F \approx 29$). Figure 5 shows the first stage and the reduced form.

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.³ Table 3 reports

³The realized rate cannot stand in for the statutory rate in this design. Instrumented by the consumer-good share, the 2017–19 change in the realized output tariff has a weak first stage ($F = 2.9$, against 29.3 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 0.64 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, and the consumer-good share predicts neither that import share nor the weighted change. 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). Section 6 returns to the origin of imports.

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 -0.0044 ($p = 0.008$; 95% confidence interval $[-0.0076, -0.0011]$), two to three times the corresponding OLS estimate. A one-point output-tariff increase thus lowers TFP by about 0.4 percent, a local average treatment effect for the consumer-good-share compliers. Carrying the actual tariff change: the interquartile range of the 2017–19 increase is 7.5 points (from 11.6 at the 25th percentile to 19.1 at the 75th), so the industry at the top quartile of exposure ends 3.3 percent below the one at the bottom, with the confidence interval running from 0.9 to 5.7 percent. This is a contrast between two industries inside the sample and should not be rescaled to the full 12–17 point exposure: the estimate is relative rather than aggregate, carries no general-equilibrium correction, and is a local average treatment effect extrapolated beyond its support if pushed that far. Extending the outcome window to 2020, 2021, and 2022, the estimate grows steadily more negative, to -0.0075 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

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

	Change in log BLS TFP, 2017 to:			
	(1) 2019	(2) 2020	(3) 2021	(4) 2022
Δ statutory output tariff, IV	-0.0044*** (0.0017)	-0.0053** (0.0025)	-0.0075** (0.0031)	-0.0075** (0.0034)
same, OLS	-0.0012	-0.0017	-0.0024	-0.0029
First-stage F	29.3	29.3	29.3	29.3
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 (-0.0046) when retaliatory-tariff exposure (Fajgelbaum et al., 2020) is added as a control, and the instrument is uncorrelated with that exposure. Robust standard errors. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 is not merely capturing 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. A war-specific threat to the exclusion restriction is that industries hit by U.S. protection were simultaneously exposed to foreign *retaliation*, which depressed activity in its own right (Flaen and Pierce, 2024). 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 0.03): retaliation fell on agricultural and export-oriented sectors, the 301 tariffs on import-competing manufacturers. Adding retaliatory exposure as a control leaves the instrumented output-tariff effect unchanged (-0.0046 in the 2019 window), while retaliation carries its own, separable, and also negative effect on TFP (-0.005 , significant); it is a distinct channel rather than a confound.

Appendix B.1.1 collects the full timing placebo and two measurement cross-checks: the same specification on the independently constructed NBER–CES series and on labor productivity as an alternative outcome.

The level decline is thus identified, sharply and with its magnitude stated as a contrast between industries rather than as an aggregate. It rests on one exclusion restriction: that the consumer-good share moves TFP only through tariff exposure. That restriction passes the placebo available to it and survives the retaliation check, but it cannot be proved. 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.

It is worth noting what Flaaen and Pierce (2024) do not study. Their outcomes are employment, industrial production and the producer price index; no productivity measure appears in their paper. The design is theirs; the question it is put to here is not.

Exposure as a share, not as a rate. Let Ω denote the set of product–country pairs newly covered by a 2018 U.S. tariff action and Ω^* 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,

$$\begin{aligned}
\text{Protection}_j &= \frac{\sum_{pc \in \Omega} M_{j,pc}}{Q_j + M_j - X_j}, \\
\text{Retaliation}_j &= \frac{\sum_{pc \in \Omega^*} X_{j,pc}}{Q_j}, \\
\text{InputCost}_j &= \sum_c \omega_{cj} \frac{\sum_{pc \in \Omega} M_{c,pc}}{Q_c + M_c},
\end{aligned} \tag{3}$$

where ω_{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.7.3 documents the rebuild of these measures and validates them against the published aggregates.

Specification. The annual counterpart of Flaaen 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 ,

$$\begin{aligned}
\tilde{y}_{jt} &= \gamma_t \text{Protection}_j + \theta_t \text{InputCost}_j + \lambda_t \text{Retaliation}_j \\
&\quad + \omega_t \text{ImpShare}_j + \varphi_t \text{ExpShare}_j + \delta_j + \delta_t + \varepsilon_{jt},
\end{aligned} \tag{4}$$

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 Flaaen and Pierce's first pre-trend correction.


The term that this paper's own specification lacks is the second line. ImpShare_j and 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 (4) carries industry and year fixed effects, the exposure-by-year controls, and an explicit pre-trend removal. As the next paragraph shows, it displays no pre-trend to correct.

Result. Figure 6 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 Flaaen 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.

Figure 6: The Flaaen–Pierce design: TFP by exposure, relative to 2017.



Notes: Year-by-year coefficients on the Protection channel from equation (4), 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, 7.5 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.

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 Flaaen and Pierce’s equation (8) gives a cumulative -0.252 ($t = -2.82$) at the 2019 endpoint. The statutory-rate column tracks it closely.

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 (0.0327)	-0.0115 (0.0304)
2014	-0.0460 (0.0553)	-0.0571 (0.0776)
2015	-0.0599 (0.0561)	-0.0980 (0.0845)
2016	-0.0095 (0.0548)	-0.0801 (0.0714)
2018	-0.0445 (0.0455)	-0.0415 (0.0622)
2019	-0.1923** (0.0760)	-0.1766* (0.0974)
2020	-0.1938* (0.1133)	-0.1509 (0.1411)
2021	-0.3353** (0.1293)	-0.3124 (0.1941)
2022	-0.4396*** (0.1602)	-0.4291* (0.2246)
Observations	850	850

Notes: Coefficients γ_t from equation (4), 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. 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 7.5 percentage points: the 25th percentile industry saw its tariff rise 11.6 points, the 75th 19.1. 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
Differenced panel, OLS	-0.0012	-3.78	0.9
Flaaen–Pierce design	-0.0027	-2.08	2.1
Consumer-share IV	-0.0044	-2.65	3.3

Notes: Effect of the statutory output tariff on log BLS TFP, per percentage point and scaled to the interquartile gap of 7.5 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 (4) 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 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. Two designs with different exposure measures, different samples and different identifying assumptions place the effect between 2 and 3 percent of TFP for an interquartile move in exposure. It is the design, not the measure, that separates them: substituting the statutory output tariff for the coverage share in equation (4) changes the answer by 0.2 percentage points, while changing the design moves it by more than a percentage point.

Two cautions carry over. 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 (4) 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 3 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 Where the decline lands: 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 lands, 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

The specification is the first-difference panel regression foreshadowed in Section 4. 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 δ_t absorb common shocks. For a percentile spread Y_{jt} of within-industry productivity,

$$\Delta Y_{jt} = \alpha_j + \delta_t + \beta^O \Delta \tau_{j,t-1}^O + \beta^I \Delta \tau_{j,t-1}^I + \varepsilon_{jt}, \quad (5)$$

where τ^O and τ^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.

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 (−0.0050, 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).

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, against −0.0052 in the full sample (Table 7). The compression is a feature of the whole sample period rather than an artifact of the recent episode.

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

	(1) Output	(2) Input	(3) Both
Δ output tariff τ^O , lag 1	-0.0050** (0.0021)		-0.0052** (0.0020)
Δ input tariff τ^I , lag 1		0.0016 (0.0082)	0.0035 (0.0079)
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 τ^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). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Output and input tariffs and 90–10 dispersion, pre-trade-war subsample (1989–2017)

	(1) Output	(2) Input	(3) Both
Δ output tariff τ^O , lag 1	-0.0043** (0.0018)		-0.0047** (0.0018)
Δ input tariff τ^I , lag 1		0.0140 (0.0211)	0.0194 (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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 things 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 (+0.0009, s.e. 0.0013; 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.

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. At those clean horizons the placebo passes: the lead is economically small and statistically indistinguishable from zero (columns 3–4), so

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

	(1) 90–10	(2) 99–90	(3) 10–1	(4) 75–25
Δ output tariff τ^O , lag 1	-0.0052** (0.0020)	-0.0031 (0.0028)	0.0014 (0.0019)	-0.0054** (0.0021)
Δ input tariff τ^I , lag 1	0.0035 (0.0079)	0.0105 (0.0105)	-0.0008 (0.0040)	0.0047 (0.0032)
Observations	2,635	2,606	2,607	2,635
Within R^2	0.038	0.026	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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 two years out.

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.

8 Mechanisms

Sections 5 and 7 establish that the output tariff lowers the level of industry TFP and compresses its within-industry dispersion. Through what channel? As argued in the introduction, the evidence is most consistent with a weakening of competitive discipline in the protected market: it is the own-product tariff, not the tariff an industry pays on its inputs, that carries the effect. This section does two things. 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 (2016) PNTR shock, from employment to productivity. Both move employment; neither detects a significant productivity response, and Section 8.4 weighs what that does and does not imply. Second, Section 8.3 tests the alternative of an input-cost channel and finds

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

	(1) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3	(5) Lead 1 +ctmp
Δ statutory, lag 1 (causal)	-0.0050** (0.0021)				
Δ statutory, contemporaneous					-0.0050** (0.0019)
Δ statutory, lead 1 (placebo)		0.0052*** (0.0018)			0.0054*** (0.0019)
Δ statutory, lead 2 (placebo)			0.0019 (0.0043)		
Δ statutory, lead 3 (placebo)				-0.0018 (0.0018)	
Observations	2,635	2,635	2,550	2,465	2,635
Within R^2	0.038	0.038	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 ≈ 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). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the input tariff null under a variety of alternative plausible constructions.

8.1 The China shock: an increase in competition

If protection lowers productivity by dulling competition, a large, well-identified *increase* in competition from another episode is directly relevant: it shows what competition moves when it rises. The [Autor et al. \(2013\)](#) China shock supplies one, from the supply side.

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](#)) 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 instrument has bite, but it does not establish that productivity is unaffected. Statistical power is specific to the outcome, industry productivity is noisier than employment, and the intervals are wide: for BLS TFP the 95 percent interval runs from -0.0010 to 0.0080 , which leaves room for a pro-competitive effect of meaningful size.

The China shock thus moves employment clearly and productivity not detectably. This is the same pattern the Pierce–Schott design in Section 8.2 shows from a different episode. It

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

	(1) Employment	(2) Dispersion 90–10	(3) TFP (BLS)	(4) TFP (NBER–CES)
Δ China import penetration	-0.0186*** (0.0062)	0.0027 (0.0058)	0.0035 (0.0023)	0.0017 (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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 8.4 returns to both readings.

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

A second design for an increase in competition, from a different episode, is the quasi-experiment of Pierce and Schott (2016). 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 (2016) 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,

$$Y_{it} = \theta (\text{NTR gap}_i \times \text{post-PNTR}_t) + \text{post}_t \times \mathbf{X}'_i \gamma + \mathbf{X}'_{it} \lambda + \delta_i + \delta_t + \varepsilon_{it}, \quad (6)$$

with industry (δ_i) and year (δ_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 \(2016\)](#) 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 their NTR gap and controls 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 (6).

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). 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 sizeable 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.

Only labor adjusts. Where, then, does the competition shock land? 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

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

	(1) Employment	(2) TFP (NBER–CES)	(3) TFP (BLS)	(4) Dispersion 90–10
NTR gap $\times$ post-PNTR	-0.0081*** (0.0028)	-0.0002 (0.0010)	0.0008 (0.0020)	-0.0027 (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 \(2016\)](#) generalized difference-in-differences (6) 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

unchanged (-0.0016 , insignificant), as one expects of a sunk, slow-moving input over a horizon 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 \(2016\)](#) document.

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 frontier) is taken up in Section 8.4, 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. 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

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

	(1) Employment	(2) Capital	(3) Materials	(4) Hours
NTR gap $\times$ post-PNTR	-0.0081*** (0.0028)	-0.0016 (0.0015)	-0.0004 (0.0031)	-0.0085*** (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 \(2016\)](#) difference-in-differences (6) 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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 ($c = j$; Section A.7.2). 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 ω_{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. The output-tariff coefficient is unchanged throughout, at -0.005 for dispersion and about -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. 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.

The input channel, then, is absent under a variety of alternative plausible constructions of the input tariff. With both alternatives tested, what remains is to say what the three subsections add up to.

8.4 What the mechanism evidence adds up to

Three findings share a pattern. The China shock and the PNTR uncertainty shock, two increases in competitive pressure from other episodes, move employment and show no significant productivity response. The one decrease in pressure the paper studies, the 2018–19 output tariff, lowers measured productivity and leaves the input mix largely alone. And the input tariff, which would carry the effect if the mechanism were cost rather than competition,

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

	(1) Dispersion 90–10	(2) TFP (BLS)	(3) TFP (NBER–CES)
Δ output tariff	-0.0049** (0.0020)	-0.0011*** (0.0004)	-0.0012*** (0.0003)
Δ input tariff (import-weighted)	-0.0005 (0.0084)	-0.0021* (0.0012)	-0.0021 (0.0041)
Δ input tariff (total-use, own 4-digit excl.)	0.0035	-0.0021	-0.0037
Δ input tariff (leave-out own 3-digit)	0.0007	-0.0011	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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

is null under a variety of alternative plausible constructions. The evidence points to competition. What it does not settle is why rising competition did not register in productivity; this subsection weighs two readings, offers a conjecture for the second, and places the result against the standard model.

Two readings of the productivity nulls. The designs sit on the same axis of competitive pressure, and the natural worry is that they contradict the headline. They need not. The China-shock productivity estimates are positive, the direction the headline implies, and no estimate in either episode points significantly the other way.

The first reading is statistical. The increases in competition in these episodes may not have been sharp enough, relative to the noise in industry productivity, to register a response: the intervals leave room for pro-competitive effects of meaningful size, and a strong employment response is no guarantee of power for an outcome as noisy as TFP. On this reading the two episodes are simply uninformative about productivity.

The second reading is economic: the effects of competition may be asymmetric in the direction of the change. Two shocks that raise competition move labor and not detectably the frontier; one that lowers it moves the frontier. These data cannot separate the two readings, but the second calls for a mechanism.

An effort ceiling. My conjecture is an *effort ceiling*. Competitive discipline already holds firms close to the best they can do with what they have. A further *rise* in pressure therefore finds little slack left to squeeze: what remains adjustable is scale, and labor is the margin that actually moves, consistent with what PNTR and the China shock show. 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 they show up as a decline in measured total factor productivity. The frontier can be lost more easily than it can be gained.

This is a conjecture offered to reconcile evidence from separate episodes, not something these data test. Two things recommend it over the alternative I considered, namely that protection forecloses the escape valves of exit and offshoring, so adjustment is forced inward. First, that account sits badly with Section 8.3, which finds that firms facing higher input tariffs re-sourced freely; sourcing flexibility and a closed offshoring valve are hard to hold at once. Second, the effort-ceiling account does not require the two episodes to be symmetric in anything except the direction of pressure, whereas the escape-valve account has to explain why the valves were open in one episode and shut in the other. What both accounts share, and what the evidence does support, is the observation that the margin of adjustment differs: under PNTR and the China shock labor moves and the frontier does not; under the output tariff the frontier moves and the input mix does not.

Inside the frontier. The mechanism evidence maps onto the distinction drawn in Section 2. The input mix responds much as the standard trade model predicts: the [Pierce and Schott \(2016\)](#) quasi-experiment shows a trade shock reallocating inputs across factors (Section 8.2). In the standard model that would be the end of it: technology unchanged, the cost of protection a movement *along* the production frontier ([Kehoe and Ruhl, 2008](#)). But it is not the end of it: the central result is that protection also lowers TFP (Section 5), so the economy produces less from a given bundle of inputs. That is a movement *inside* the frontier rather than along it. The factor reallocation the standard model isolates is present; the novel and larger margin is the contraction of the frontier itself.

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 7.5 percentage points apart, and across that gap an industry at the top quartile ends 2019 with total factor productivity about 3.3 percent below one at the bottom. This estimate is robust to controlling for retaliation, and it is reproduced at 2.1 percent by a second design that uses no instrument and shows no pre-trend. 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 effect is carried by the *output* tariff, the rate that shields an industry’s own market, not by the tariff an industry pays on its inputs. 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. 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. This places my result outside the standard model of tariff costs, in which protection distorts the input mix but leaves technology untouched: a movement *along* the production frontier. I see that input-mix reallocation, but I also find that protection lowers TFP itself: a movement *inside* the frontier, in which a given bundle of inputs yields less output. Protection is costly not only through the misallocation the standard model describes, but through a second channel it does not contain: less output from the bundle of inputs actually used.

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. Version 2018a (1958–2018), <https://www.nber.org/research/data/nber-ces-manufacturing-industry-database>.
- 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.
- 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.
- 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.
- 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>.
- Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.
- 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. The surprisingly swift decline of US manufacturing employment. *American Economic Review*, 106(7):1632–1662, 2016.
- Peter K. Schott. U.s. HS-level imports and exports, 1989–2023. Data, https://sompks4.github.io/sub_data.html, 2025. Annual import-detail files `imp_det1_<YYYY>_12n.dta`, converted from U.S. Census Bureau records.
- 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 Treffer. 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: 2012 benchmark detail use table, before redefinitions, producers’ prices. Industry Economic Accounts, 2018. <https://www.bea.gov/industry/input-output-accounts-data>.
- 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, 2025. <https://www.bls.gov/productivity/tables/>.

Data Appendix

to

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

Luca Guerrieri

Draft — September 14, 2026

Abstract

This appendix 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 realized rates and Section [A.4](#) the statutory rates; Section [A.5](#) merges them with DISP; Section [A.6](#) reports validation checks; and Section [A.8.1](#) lists known limitations and planned refinements.

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 realized and the statutory tariff measures; the merge with DiSP and its validation; the productivity and input-tariff measures; and the limitations and reproduction materials. 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** τ_{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.3);
- the **statutory (schedule) rate** τ_{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.4). 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, 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,

$$R_i = \sum_{\ell \in i} r_\ell w_\ell, \quad M_i = \sum_{\ell \in i} w_\ell, \quad (7)$$

and forming the cell value R/M . For the *realized* rate, r_ℓ and w_ℓ 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_ℓ is a tariff-line rate and w_ℓ a line weight (Section A.4.2). 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.

A.2.5 Code types

Each six-digit code is tagged: **ok** (complete six-digit industry); **partial** (an “X”-filled code resolvable only to four digits, e.g. 33631X, kept in the four-digit rollup via its numeric prefix); **special** (Chapters 98–99 provisions and sector placeholders NN0000 such as 910000/980000, excluded); and **missing** (excluded). These categories coincide with the non-industry “base-roots” flagged by [Pierce and Schott \(2012\)](#).

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 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, the harmonization to the 2012 NAICS basis, and the extension to the earliest years.

A.3.1 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.

A.3.2 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

$$\tau_{jt}^R = \frac{\sum_{i \in j} D_{it}}{\sum_{i \in j} V_{it}}, \quad (8)$$

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 (Section A.2.5); “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.3.3 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 (7)–(8).

A.3.4 Early years, 1987–1988

DiSP begins in 1987 but Schott (2025) begins in 1989. The 1987–1988 extension uses the commodity-level U.S. import files of Feenstra (1996) (the duty-bearing `IMPYR_*/IMP89_*` files, not the SIC-aggregated `impsic` files, which lack duties), aggregated to the 1987 SIC code they carry and concorded `sic87→naics97` with the EFSY weights, then run through the chain. This extension is held in reserve in the current panel: it requires the commodity files, and the 1987 source was unavailable at the time of writing, so the realized panel currently spans 1989–2021.

A.4 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.4.1 Definition

Statutory rates are the legal ad valorem rates from the Harmonized Tariff Schedule: the Column 1 / NTR (most-favored-nation) rate, the Column 2 / non-NTR rate, and, separately, the 2018–19 Section 232 and Section 301 additional duties. 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 output tariff used in the analysis is the Column 1 rate plus the two add-ons. 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 Column 2 rate enters only the NTR gap of Section 8.2. The statutory rate omits Section 301 List 4A, the 2002 and 2018 Section 201 safeguards, and antidumping and countervailing duties.

A.4.2 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 τ^R but biased down, since high tariffs depress their own weights) is also supported. Aggregation reuses (7) with r_ℓ the line rate and $w_\ell \equiv 1$ (simple) or the line’s import value (weighted).

A.4.3 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 inherits the NAICS code of its HS line via the complete 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.4.4 2002–2021: USITC DataWeb

Schedule rates for 2002–2021 come from the USITC DataWeb 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 (Section [A.8.1](#)).

A.4.5 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 `statutory_232 = share_232 ×` (average add-on on covered lines), exactly as for Section 301 (Table A.5).

A.4.6 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 Fajgelbaum et al. (2020) replication data, 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 through 2021. 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 (7) into `statutory_301`, with a companion coverage share. The carry-forward is verified against DataWeb’s Chapter 99 provisions (Section A.6.5). List 4A is not yet included (Section A.8.1).

A.5 Merge with DiSP and the resulting panel

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** × the **86** manufacturing industries, carrying τ^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.

A.6 Validation checks

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.6.1 Cross-dataset seams

The panel splices independent sources, so I inspect each seam rather than smooth it. (i) *Re-estimated, 1988/1989*: Feenstra (1996) overlaps Schott (2025) in 1989–1994, allowing a direct

four-digit cross-check of the SIC-based and NAICS-based series (the early-years extension also changes the HS→NAICS route). (ii) *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 (Section [A.8.1](#)).

A.6.2 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 collections. 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.6.3 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	0.4	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	24.0	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.6.4 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 63% of manufacturing tariff lines; its manufacturing-mean add-on jumps from zero (2017) to 13.2% (2019, after List 3’s escalation to 25%) and 15.9% (2021). The realized rate rises in parallel (2.0% $\rightarrow$ 4.6%), confirming the schedule changes show up in collected duties.

A.6.5 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	13.22	0.61	63
2020	4.40	3.55	15.85	0.61	63
2021	4.63	3.74	15.85	0.61	63

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** × (average add-on on the covered lines); for example, in 2021 manufacturing, 15.85% = 63% × 25%. 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.7 Productivity and input-tariff measures

Two further measures are built alongside the tariffs and used as the outcome and the key control in the regressions: the productivity series and the input tariff that separates the cost channel from competition. I describe the two independently constructed total factor productivity series first, then the input-tariff construction from the BEA input–output accounts.

A.7.1 Total factor productivity data

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, 2025](#)). 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.7.2](#). Coverage is 1987–2022; `build_tfp.py`.

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.7.2 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 \(2011\)](#):

$$\tau_{j,t}^I = \sum_c \omega_{cj} \tau_{c,t}^O, \quad (9)$$

where $\tau_{c,t}^O$ is the statutory output tariff of supplying industry c (Section [A.4](#)) and ω_{cj} is the share of commodity c in industry j ’s intermediate-input purchases.

Weights. The ω_{cj} come from the 2012 benchmark *detail* use table of the BEA ([U.S. Bureau of Economic Analysis, 2018](#)) (“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 τ^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 τ^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 τ^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 τ^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, 1989–2021; `build_input_tariff.py`, whose τ^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 τ^I rises even more than τ^O during the 2018–19 trade war; the two are correlated at 0.58, leaving ample independent variation to separate the channels.

A.7.3 Flaaen–Pierce exposure measures

The three exposure measures of equation (3) 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 Ω 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 (4).

A.8 Limitations, reproduction, and provenance

Finally, I record what the current build does not yet capture and how to rebuild it: the known limitations and planned refinements, a step-by-step reconstruction checklist, and a log of every downloaded input and its source.

A.8.1 Limitations and planned refinements

1. **Section 301 List 4A.** Its $\sim 3,200$ -line product list (effective 2019-09-01, 15% then 7.5% from 2020-02-14) lies beyond the [Fajgelbaum et al. \(2020\)](#) window and is not yet in `statutory_301`; its rate is confirmed in DataWeb (heading 9903.88.15) but its line membership requires HTS U.S. Note 20. `statutory_301` is therefore a slight *under*-count for 2020–2021.
2. **Specific-duty AVEs (`statutory`, 2002–2021).** DataWeb’s non-ad-valorem lines are currently dropped; computing ad valorem equivalents from `imp_det1` duty/value would include them and tighten the 2001/2002 seam.
3. **Import-weighted statutory rates, 2002–2021.** DataWeb tariff files carry no import value, so the weighted variant for these years needs an HS-8 value join from `imp_det1`.
4. **Realized early years (1987–1988).** Pending the duty-bearing Feenstra commodity files; 1987 was unavailable at the time of writing.
5. **Establishment vs. aggregate recoding.** DiSP recodes individual establishments ([Fort, 2013](#)); the trade data can only be recoded at the level of NAICS aggregates ([Eckert et al., 2020](#)). The two share a target and a methodology but differ in granularity; this residual gap cannot be closed with public data.

A.8.2 Reconstruction checklist

1. *Realized:* collapse `imp_det1` to NAICS-6 by summing `cal_dut_yr/con_val_yr`; drop special/missing codes; apply the EFSY chain to 2012 NAICS; roll to NAICS-4 via (8).
2. *Statutory:* assign HS-8 lines to NAICS ([Pierce and Schott, 2012](#)); aggregate line rates by (7) (default simple mean); overlay Section 232 (proclamations) and Section 301 (FGKK membership + schedule); harmonize and roll up.
3. *Merge:* join both tariff measures to DiSP on (`naics4, year`) over 1989–2021.

A.8.3 Data access log

For reproducibility I record the provenance of each downloaded input.

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, its within-industry *distribution*, and the *dynamics* of the response. The same pattern recurs across all three: the own-product

Table A.6: Downloaded inputs

Input	Source	Local path
Import-detail files	Schott (2025) , <code>imp_det1.<YYYY>_12n.dta</code> , 1989–2021	<code>Trade_data/Schott_data/</code>
EFSY NAICS/SIC concordance	Eckert et al. (2020) , GitHub release <code>cbp-data</code>	<code>concordance/</code>
HS10→NAICS concordance	Pierce and Schott (2012) , <code>hs_sic_naics_imports.*</code>	<code>concordance/</code>
FRS statutory rates	Feenstra et al. (2002) , <code>tar_val.dta</code> , 1989–2001	<code>Trade_data/Schott_data/</code>
DataWeb annual tariffs	USITC DataWeb, 2002–2021	<code>DataWeb_tariff_data/</code>
FGKK replication	Fajgelbaum et al. (2020) , Harvard Dataverse KSOVSE	<code>QJE_replication/</code>
DiSP panel	U.S. Bureau of Labor Statistics (2024)	<code>TFP_data.xls</code>

(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

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.7](#) 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 three years further than the TFP index. Table A.8 re-runs the statutory specification with the change in log BLS labor productivity, which is available through 2024 and so extends the panel three years past the trade war. The output (competition) tariff enters more strongly than for TFP (-0.0016 , significant) and survives these post-war years. The input (cost) tariff, by contrast, is essentially zero ($+0.0007$, 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 efficiency with which *all* inputs are converted into output (TFP) 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.9 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.10 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.11 is the value-added-weighted statutory specification. The qualitative pattern is unchanged: both tariffs lower TFP, with the output (competition) channel the more robust of the two.

Table A.7: Robustness: statutory tariffs and NBER–CES TFP (first differences)

	(1) Output	(2) Input	(3) Both
Δ output tariff τ^O , lag 1	-0.0013*** (0.0003)		-0.0012*** (0.0003)
Δ input tariff τ^I , lag 1		-0.0050 (0.0045)	-0.0037 (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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B.1.3 Placebo timing tests: robustness

The timing placebo of Table A.10 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.12 drops the 2018–19 China trade war: the panel is truncated at 2017, so the leads cannot reach the war years. Table A.13 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.14 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.

B.2 Robustness of the dispersion result

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

Table A.8: Statutory tariffs and industry labor productivity (first differences)

	(1) Output	(2) Input	(3) Both
Δ output tariff τ^O , lag 1	-0.0016** (0.0007)		-0.0016** (0.0007)
Δ input tariff τ^I , lag 1		0.0003 (0.0015)	0.0007 (0.0016)
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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 at about -0.005 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 84 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 -0.0045 , 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 ($F \approx 70$). The realized input tariff remains insignificant (0.0049, s.e. 0.0261). 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

Table A.9: H2 mechanism: change in the input-cost share and the input tariff

	(1) Interm. share	(2) Materials sh.
Δ output tariff τ^O , lag 1	0.0001 (0.0002)	0.0004*** (0.0001)
Δ input tariff τ^I , lag 1	-0.0002 (0.0008)	0.0007 (0.0018)
Observations	2,720	2,380
Within R^2	0.099	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 τ^I is direct evidence that the production mix shifts with input tariffs (H2). Industry and year fixed effects; SEs clustered by industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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 *real* TFP dispersion and not in the labor-productivity measure that would move with capacity utilization, prices, or markups. This pattern points to a real-efficiency channel rather than a measurement artifact; its concentration in the unweighted measure locates it among the mass of smaller establishments rather than the large plants that dominate activity.

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 -0.005 and -0.001). The compression of dispersion and the decline in the TFP levels therefore survive absorbing each industry's own trajectory, and the input tariff remains insignificant throughout. 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 local-projection 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.
- Randy A. Becker, Wayne B. Gray, and Jordan Marvakov. NBER-CES manufacturing industry database. *National Bureau of Economic Research*, 2021. Version 2018a (1958–2018), <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.
- 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. The return to protectionism. *Quarterly Journal of Economics*, 135(1):1–55, 2020.
- Robert C. Feenstra. U.S. imports, 1972–1994: Data and concordances. NBER Working Paper 5515, National Bureau of Economic Research, 1996. Data distributed by the UC Davis Center for International Data, <https://cid.ucdavis.edu/usixd>; import files IMPSIC<YY>.ASC by 4-digit SIC.
- 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.

- 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.
- 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, 1989–2023. Data, https://sompks4.github.io/sub_data.html, 2025. Annual import-detail files `imp_det1_<YYYY>_12n.dta`, converted from U.S. Census Bureau records.
- 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: 2012 benchmark detail use table, before redefinitions, producers’ prices. Industry Economic Accounts, 2018. <https://www.bea.gov/industry/input-output-accounts-data>.
- 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, 2025. <https://www.bls.gov/productivity/tables/>.

Table A.10: Timing placebo: leads of the tariff change and the change in TFP
Panel A. Output tariff

	(1) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3
Δ output tariff, lag 1 (causal)	-0.0012*** (0.0004)			
Δ output tariff, lead 1		0.0013*** (0.0004)		
Δ output tariff, lead 2			0.0019*** (0.0006)	
Δ output tariff, lead 3				-0.0016*** (0.0005)
Observations	2,720	2,890	2,890	2,805
Within R^2	0.168	0.168	0.163	0.137
Industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Panel B. Input tariff

	(1) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3
Δ input tariff, lag 1 (causal)	-0.0025** (0.0012)			
Δ input tariff, lead 1		0.0019 (0.0012)		
Δ input tariff, lead 2			-0.0026** (0.0011)	
Δ input tariff, lead 3				-0.0018** (0.0009)
Observations	2,752	2,924	2,924	2,838
Within R^2	0.172	0.171	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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.11: Statutory tariffs and industry TFP, value-added-weighted (first differences)

	(1) Output	(2) Input	(3) Both
Δ output tariff τ^O , lag 1	-0.0015*** (0.0005)		-0.0014*** (0.0005)
Δ input tariff τ^I , lag 1		-0.0027* (0.0015)	-0.0021 (0.0015)
Observations	2,720	2,752	2,720
Within R^2	0.174	0.175	0.175
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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.12: Timing placebo excluding the 2018–19 China trade war (change in TFP)

<i>Panel A. Output tariff</i>				
	(1) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3
Δ output tariff, lag 1 (causal)	-0.0011*** (0.0004)			
Δ output tariff, lead 1		0.0013** (0.0006)		
Δ output tariff, lead 2			0.0022*** (0.0007)	
Δ output tariff, lead 3				-0.0022*** (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) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3
Δ input tariff, lag 1 (causal)	-0.0023 (0.0043)			
Δ input tariff, lead 1		-0.0145*** (0.0043)		
Δ input tariff, lead 2			-0.0007 (0.0034)	
Δ input tariff, lead 3				-0.0019 (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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.13: Timing placebo, value-added-weighted (change in TFP)

Panel A. Output tariff

	(1) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3
Δ output tariff, lag 1 (causal)	-0.0015*** (0.0005)			
Δ output tariff, lead 1		0.0011** (0.0005)		
Δ output tariff, lead 2			0.0024*** (0.0005)	
Δ output tariff, lead 3				-0.0017*** (0.0006)
Observations	2,720	2,890	2,890	2,805
Within R^2	0.174	0.174	0.176	0.151
Industry FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes

Panel B. Input tariff

	(1) Lag (base)	(2) Lead 1	(3) Lead 2	(4) Lead 3
Δ input tariff, lag 1 (causal)	-0.0027* (0.0015)			
Δ input tariff, lead 1		0.0039*** (0.0009)		
Δ input tariff, lead 2			-0.0021 (0.0013)	
Δ input tariff, lead 3				-0.0017 (0.0013)
Observations	2,752	2,924	2,924	2,838
Within R^2	0.175	0.177	0.171	0.149
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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.14: Conditional placebo: tariff-change leads with and without the maintained causal lag (change in TFP)

<i>Panel A. Output tariff</i>					
	(1) Lag	(2) Lead 2	(3) Lead 2 +lag	(4) Lead 3	(5) Lead 3 +lag
Δ output tariff, lag 1 (causal)	-0.0012*** (0.0004)		-0.0010* (0.0005)		-0.0015*** (0.0004)
Δ output tariff, lead 2		0.0019*** (0.0006)	0.0018*** (0.0007)		
Δ output tariff, lead 3				-0.0016*** (0.0005)	-0.0019*** (0.0006)
Observations	2,720	2,890	2,635	2,805	2,550
Within R^2	0.168	0.163	0.166	0.137	0.142
Industry FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
<i>Panel B. Input tariff</i>					
	(1) Lag	(2) Lead 2	(3) Lead 2 +lag	(4) Lead 3	(5) Lead 3 +lag
Δ input tariff, lag 1 (causal)	-0.0025** (0.0012)		-0.0029** (0.0012)		-0.0023* (0.0014)
Δ input tariff, lead 2		-0.0026** (0.0011)	-0.0027** (0.0011)		
Δ input tariff, lead 3				-0.0018** (0.0009)	-0.0017** (0.0008)
Observations	2,752	2,924	2,666	2,838	2,580
Within R^2	0.172	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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

	(1) Baseline	(2) Excl. war	(3) VA-weighted	(4) War-IV
Δ output tariff τ^O , lag 1	-0.0052** (0.0020)	-0.0047** (0.0018)	-0.0047** (0.0018)	-0.0069 (0.0042)
Δ input tariff τ^I , lag 1	0.0035 (0.0079)	0.0194 (0.0211)	0.0076 (0.0079)	0.0042 (0.0080)
Estimator	OLS	OLS	OLS	IV
First-stage F				19829.19
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). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

	(1) Output	(2) Input	(3) Both
Δ output tariff τ^O , lag 1	-0.0043** (0.0021)		-0.0047** (0.0018)
Δ input tariff τ^I , lag 1		0.0051 (0.0084)	0.0076 (0.0079)
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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

	(1) Statutory	(2) Realized
Δ statutory input tariff, lag 1	0.0016 (0.0082)	
Δ realized input tariff, lag 1		0.0049 (0.0261)
Estimator	OLS	IV
First-stage F		69.58
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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

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

	TFP dispersion		Labor-prod. dispersion	
	(1) Unweighted	(2) Activity-wtd	(3) Unweighted	(4) Activity-wtd
Δ output tariff τ^O , lag 1	-0.0052** (0.0020)	0.0009 (0.0013)	0.0032 (0.0026)	0.0037 (0.0027)
Δ input tariff τ^I , lag 1	0.0035 (0.0079)	-0.0028 (0.0034)	0.0140* (0.0081)	-0.0008 (0.0055)
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). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.19: Robustness to industry-specific trends

	(1) Dispersion 90–10	(2) TFP (BLS)	(3) TFP (NBER–CES)
Δ output tariff	-0.0058** (0.0022)	-0.0010*** (0.0003)	-0.0012*** (0.0003)
Δ input tariff	0.0032 (0.0083)	-0.0016 (0.0012)	-0.0031 (0.0049)
output-tariff coef., no trends	-0.0052	-0.0011	-0.0012
Observations	2,635	2,720	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. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.