Tariffs, Productivity, and the Welfare Effects of Protectionism

Luca Guerrieri*

October 2026

Abstract

In 2018–19 the United States raised tariffs on more than $350 billion of imports. I ask how these tariffs affected the productivity of the targeted manufacturing industries and what that implies for welfare. The added tariff was nearly uniform, 25 percent on most covered goods, but the share of an industry's home market supplied by newly tariffed imports differed widely across industries. Using a difference-in-differences design, I find that the more exposed industries broke from their pre-war productivity trend relative to the less exposed ones, with total factor productivity about 2 percent lower by 2019 for an interquartile difference in exposure. To draw the welfare implications, I extend the workhorse macro-trade model of Ghironi and Melitz (2005) to sectors with different exposure to the tariffs. The tariffs shift production toward less productive firms, but the resulting gap in measured productivity falls far short of the evidence. Matching it requires an exogenous technology shock that can reverse the welfare gain from the terms-of-trade improvement generated by the tariffs.

JEL codes: F13, F14, D24, L60. *Keywords*: tariffs, protectionism, total factor productivity, trade war.

*Federal Reserve Board, 20th and C Streets NW, Washington, DC 20051, United States. Corresponding author. Email: luca.guerrieri@frb.gov. The views expressed in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

In open-economy macroeconomic models with heterogeneous firms, a unilateral tariff lowers physical productivity, yet home welfare rises. These models build in a well-documented fact, that exporters are more productive than firms that sell only at home (Bernard et al., 2007). A unilateral import tariff appreciates the real exchange rate, curbs the exports of the most productive firms and shifts production toward less productive ones. In the workhorse model of Ghironi and Melitz (2005), with a tariff added, physical productivity, the average log productivity of the firms weighted by the labor each employs in production, therefore falls after a home tariff, although no firm's own productivity changes. The terms-of-trade effects of the tariff are nonetheless strong enough that home welfare rises, as long as other countries do not retaliate. The purpose of this paper is to stress test that result with an empirically grounded exploration.

The welfare case for a unilateral tariff rests on market power. A country large enough to influence world prices can, by taxing its imports, lower the price it pays for them. The tariff thereby improves its terms of trade and transfers income from foreign exporters to domestic households. Johnson (1951) derived the optimal tariff, which equals the reciprocal of the foreign elasticity of supply of imports. The argument holds only under strict conditions: the country must be large enough to move world prices, and the terms-of-trade gain must outweigh the distortion of domestic consumption and production. Retaliation can erase the gain, although Johnson (1953) shows that a country can still gain when its partners retaliate, for instance when its own demand for imports is much more elastic than theirs. Countries do appear to exploit such market power, setting higher tariffs on goods whose foreign supply is inelastic (Broda et al., 2008), and quantitative trade models place the terms-of-trade motive at the center of their welfare accounting (Costinot and Rodríguez-Clare, 2014). Itskhoki and Mukhin (2026) study what shapes the optimal unilateral tariff in a calibrated model of the U.S. economy, and find that the valuation losses an appreciation imposes on the country's large external dollar liabilities cut the optimal tariff from 35 to 9 percent.

The U.S. tariffs of 2018–19 provide a natural setting to test whether a productivity loss can overturn the welfare gain from a unilateral tariff. In 2018 the United States raised tariffs on $303 billion of imports, 12.7 percent of the annual total, a return to protectionism unprecedented in the postwar era (Fajgelbaum et al., 2020); by the end of 2019 the tariffs on Chinese goods alone covered roughly $350 billion of imports (Fajgelbaum and Khandelwal, 2022). The empirical assessments of the episode find that the cost fell on U.S. buyers (Amiti et al., 2019; Fajgelbaum et al., 2020) and that the aggregate loss was small (Fajgelbaum et al., 2020; Fajgelbaum and Khandelwal, 2022). The assessments are short-run and static

by design, by their authors' own account. They hold technology and the set of producers fixed, so they cannot register a cost that builds up through the productivity of the protected industries.[1] I ask whether the productivity of the protected industries did change, and what the answer implies for the welfare accounting of the tariffs.

The paper makes three contributions. The first is evidence. I measure how the total factor productivity (TFP) of the protected industries moved relative to its pre-war trend after a large tariff increase in an advanced economy, and find a break by 2019. The second is a way to read such cross-sectional evidence through a macroeconomic model of trade. A multisector version of the model of Ghironi and Melitz (2005), with sector TFP measured as the statistical agency measures it, delivers the same contrast between more and less exposed industries as the empirical specification, so the model and the evidence are compared like for like. The third concerns the models and the welfare accounting. The gap in measured TFP that the tariffs open in the model falls far short of the estimate. A decline of the estimated size, if it persists, turns the welfare gain from the tariffs into a loss. If tariffs lower productivity in the industries they reach, a tariff on all imports lowers it in every sector, and then even a small fraction of the estimated decline brings the optimal tariff to zero. The first contribution requires measuring the break in productivity, and the other two require a model that, like the evidence, distinguishes the industries that the tariffs hit to different degrees.

I borrow the difference-in-differences design of Flaaen and Pierce (2026), who study the same tariffs at the same level of aggregation, with employment, industrial production and producer prices as outcomes. The design compares industries by how much the tariffs shielded them from foreign competition. The measure of exposure, import protection, is the added duty that the 2018 actions placed on an industry's imports, as a share of the value of the industry's home market. Because almost every covered import faced the same additional 25 percent tariff, the measure is in effect roughly a quarter of the share of the home market supplied by newly tariffed imports. It is large where imports supplied much of the market and the tariffs covered many of those imports, and small where either was small. I rebuild the measure from public sources and apply the design to the BLS measure of TFP, the output an industry obtains from all its inputs taken together.

The more exposed industries broke from their relative productivity trend in 2019, the

[1] Fajgelbaum et al. (2020) analyze "the short-run impact" of the tariffs (p. 1) and describe their aggregate result as "a ballpark estimate" (p. 4); Amiti et al. (2019, p. 200) note that their approach "makes a number of simplifying assumptions, including partial equilibrium and perfect competition"; and Fajgelbaum and Khandelwal (2022, p. 222) list among the possible explanations of the gap between stock-market and static estimates that "static models do not incorporate potentially important mechanisms, such as dynamic losses or growing uncertainty."

first full year in which all the tariffs applied. An industry at the 75th percentile of import protection, where the added duty was 1.94 percent of the value of its home market, against 0.12 percent at the 25th, is associated with TFP about 2.3 percent lower in 2019, relative to its pre-war path ($t = -7.63$). Substituting the change in the statutory tariff for import protection gives a similar contrast. Before the war the more exposed industries were gaining on the others; after 2017 the gain stopped. A tariff that slowed productivity in the protected industries would produce that pattern, and so would mean reversion after a pre-war run-up, which I cannot exclude with this design. By Flaaen and Pierce's own test for pre-trends the pre-war gain is significant, which by their logic is what calls for removing it. The break is very unlikely to reflect a chance assignment of exposure, correlation among industries in the same sector or the business cycle, and it is not accompanied by a fall in utilization. Per point of tariff, the decline lies between the estimates of the two other studies of trade policy whose productivity measures are comparable with mine. It is smaller than the fall in physical productivity that Pierce (2011) finds under antidumping duties and larger than the gain in industry TFP that Bernard et al. (2006) find after trade costs fell, both measured at longer horizons.

The empirical evidence challenges the estimates of the tariffs' static cost but cannot by itself overturn them. By comparing more and less exposed industries I identify only a relative effect: nothing about manufacturing as a whole, the general-equilibrium adjustment of wages and prices, or the terms-of-trade gain that is the classic case for a unilateral tariff. Nor can I separate the sources of a change in measured TFP: a change in technology, a shift of production toward less productive firms, and the way output is valued in the published series all move the measure together. In a model I can separate the last two, because I can compute measured TFP there as the statistical agency does.

The workhorse model of Ghironi and Melitz (2005), once a tariff is added, suits the question. The tariff improves the terms of trade, which is the source of the welfare gain from a unilateral tariff. It also lowers physical productivity, by shifting production across heterogeneous firms and by changing which of them export. Countries also borrow and lend, so the welfare effect of an announced tariff path can be computed in full.

Under the CES demand of the published model, however, measured TFP cannot register the fall in physical productivity, because BLS values each firm's output at its own price. Under CES every firm prices at the same markup over its marginal cost, so a more productive firm charges a proportionally lower price. What it gains in physical output per worker it gives up in price, and every firm earns the same revenue per production worker. Moving workers from one firm to another therefore lowers physical output per worker but leaves output valued at producer prices unchanged. The only movement in the measure comes from the

labor that exporters spend on the fixed cost of exporting. The tariff reduces that labor, so measured TFP edges up even as physical productivity falls. My baseline therefore replaces CES with translog demand, which Bilbiie et al. (2012) use in a closed economy with identical firms and which Rodríguez-López (2011) brings to heterogeneous firms in two countries to study exchange-rate pass-through, so that the endogenous response of productivity reaches measured TFP as well. Under translog, larger and more productive firms charge higher markups, so a shift of production toward less productive firms moves output toward lower-markup firms, which pulls measured TFP down and offsets part of the rise from lower export overhead.

In a three-sector extension of the model, a decline of the estimated size reverses the welfare gain from the tariffs. The three sectors stand for the most exposed quarter of industries, the middle half and the least exposed quarter, and I calibrate each to its markups, trade shares and tariff path, so that the model reproduces both the aggregate tariff increase and the gap in tariffs between the most and least exposed industries. Because I measure sector TFP in the model as BLS measures it, the comparison with the estimate is like for like, and with the model I translate the measured decline into the physical decline on which welfare depends. With the tariffs alone, home welfare rises by 0.30 percent of consumption through the familiar terms-of-trade effect, as the tariffs raise the relative price of home goods and transfer income from foreign producers to home households. The tariffs alone, however, open a gap in measured TFP of only 0.01 percent between the most and least exposed sectors in 2019, against an estimated decline of 5.05 percent in the most exposed quarter of industries relative to the least exposed. To match the estimated gap, I add a contractionary technology shock to the most exposed sector. With the estimated decline, home welfare changes by -0.59 percent, and a decline of a little over a third of the estimate is enough to erase the gain.

The reversal rests on two conditions that I cannot settle with the evidence: how long the decline persists, and how much of the estimated gap is a decline. The terms-of-trade effects in the model are strong. When the decline fades at 5 percent a quarter, re-sized to reproduce the same measured gap, home welfare with the decline is 0.11 percent of consumption, about two-fifths of the gain from the tariffs alone, instead of -0.59 percent. For the 2018–19 tariffs, which persist, the decline turns the gain into a loss only if its half-life exceeds about 27 quarters (Section 5.1.4). With the evidence through 2019, the last year before the pandemic, I cannot tell whether the protected industries regain the lost ground. On the second condition, I identify the gap between more and less exposed industries, not its level. The loss deepens if the decline also reached the less exposed industries, as the input-cost and retaliation channels of Flaaen and Pierce (2026) would suggest. Because the level is not identified, I cannot rule

out that part of the gap is instead a rise in productivity in the industries the tariffs did not reach, although no obvious mechanism would produce such a gain. The loss disappears once such a rise accounts for more than about 8 percent of the gap (Section 5.2). The model also omits foreign retaliation and the cost of tariffs on intermediate inputs, the channels that Fajgelbaum et al. (2020) find decisive for the aggregate loss (retaliation) and that Flaaen and Pierce (2026) find dominant for manufacturing (retaliation and input costs), so its gain from the tariffs alone is, if anything, generous (Section 5.1).

Persistence is precisely what a model of endogenous productivity would determine. In the models of medium-term fluctuations of Comin and Gertler (2006), research and development and the adoption of new technologies both respond to economic conditions, so a shock that depresses activity, even one that is not technological, leaves productivity below its previous path for years. Such a model could make the decline an outcome of the tariffs and discipline its persistence. Bringing that class of models to the evidence, however, would require extending it to many countries and many sectors, which is beyond the scope of this paper. Here the decline comes from outside the model, and I find no sign that it is sharper where a competition channel would predict. I size its consequences without explaining it.

Firm heterogeneity does not reverse the optimal tariff of Itskhoki and Mukhin (2026), but a productivity response to the tariff can bring it to zero. They find that the optimal U.S. tariff is 35 percent when cross-border positions are left out. In the one-sector model with heterogeneous firms and home a quarter of the world, the optimal tariff is 36 percent. Their finding, like this one, holds technology fixed. That assumption carries much weight. A decline of the estimated size in the most exposed sector alone turns the gain from the 2018–19 tariffs into a loss (Section 5.1.2). The optimal tariff, however, is a tariff on all imports, not a tariff on particular goods. The 2018–19 tariffs were of the second kind, so the evidence measures how much technology fell, per point of tariff, in the industries they reached. If a tariff on all imports lowered technology by as much per point in every sector, a decline far smaller than that would bring the optimal tariff to zero, as long as it persisted (Section 5.1.4).

The rest of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 builds the exposure and productivity measures, sets out the design and measures the break in productivity relative to its pre-war trend; Appendix A documents the data and Appendix B reports further empirical results. Section 4 sets out the model, the mechanics of a tariff and the optimal unilateral tariff, and Section 5 extends the model to three sectors, differentially exposed to the tariffs, interprets the empirical evidence through the lens of that extension and assesses the welfare consequences; Appendix C states the one-sector model in full and the changes the three-sector extension makes to it. Section 6 concludes.

2 Related literature

I build on three strands of work: studies of the 2018–19 tariffs, models of tariffs with heterogeneous firms, and evidence on trade policy and productivity.

Studies of the 2018–19 tariffs find a small aggregate cost while holding the productivity of the protected industries fixed. Estimates of the static cost find that the tariffs passed through almost fully to U.S. import prices and that the loss of real income was small. Fajgelbaum et al. (2020) put it at about 0.04 percent of GDP for the 2018 tariffs, Fajgelbaum and Khandelwal (2022) at 0.1 percent with those of 2019, and Amiti et al. (2019) measure deadweight losses of about $1.4 billion a month by the end of 2018. In these estimates the sign of the welfare effect turns on the terms of trade. The closest paper to mine is Flaaen and Pierce (2026), who decompose tariff exposure into import protection, input costs and retaliation, and find that the input-cost and retaliation channels overwhelm a small positive effect of import protection, lowering manufacturing employment. I borrow their design and decomposition with productivity as the outcome, then carry the estimate into a model in which the tariff also improves the terms of trade.

Models with heterogeneous firms let a tariff change industry productivity, but only by shifting output across firms whose technology is given. Exporters are larger and more productive than firms that sell only at home (Bernard et al., 2007), so trade costs change industry productivity by shifting output between more and less productive firms (Melitz, 2003). Demidova and Rodríguez-Clare (2009) show that an import tariff then raises welfare even though productivity falls, in a small economy that takes the prices of its imports as given; my two-country model lets the tariff lower them, which is the source of the terms-of-trade gain that is central to the quantitative trade models used for welfare assessment (Costinot and Rodríguez-Clare, 2014).[2] In dynamic models with endogenous entry in the tradition of Ghironi and Melitz (2005), Barattieri et al. (2021) find tariffs recessionary, in part because tariffs "reallocate domestic market share toward less efficient domestic producers, lowering aggregate productivity" (p. 2). I show how to bring cross-industry evidence to a model of this class, and find that the reallocation it generates falls far short of the measured decline.

Most of the evidence on whether shielding producers from competition lowers their productivity measures revenue productivity and comes from liberalizations in developing economies. Holmes and Schmitz (2010) review the evidence from industry studies and trade liberalizations and find that in most of them increased competition raised industry productiv-

[2]Terms-of-trade effects and productivity effects are usually studied apart. In the words of Costinot and Rodríguez-Clare (2014, p. 251), terms-of-trade effects "are central" to quantitative trade models, while they "play no direct role in productivity studies, whose main goal is to identify the consequences of trade liberalization on productivity at the firm- and industry-level."

ity; Goldberg and Pavcnik (2016) survey the evidence on trade policy. Revenue productivity moves with a producer's price and markup as well as with its physical productivity (Foster et al., 2008). Evidence on physical or industry productivity is scarce. Bernard et al. (2006) link falling U.S. trade costs to faster growth of industry TFP, and Pierce (2011) finds that antidumping duties lowered the physical productivity of the U.S. plants it shielded while their revenue productivity barely moved. Revenue productivity rose under antidumping measures at the average protected European firm (Konings and Vandenbussche, 2008) and after liberalization in Chilean import-competing plants (Pavcnik, 2002). Chinese import competition raised TFP within European firms, estimated with a correction for imperfect competition (Bloom et al., 2016), although it reduced patenting among exposed U.S. firms (Autor et al., 2020). By contrast, my evidence comes from a tariff increase in an advanced economy, and I measure industry TFP built from product-level deflators, which divide out the industry's own output price; per point of tariff, it lies between the estimates of Pierce (2011) and Bernard et al. (2006) (Section 3.5).

3 Tariffs and the productivity of the protected industries

The 2018 tariffs are associated with relative declines in productivity. The industries most exposed to them broke from their relative productivity trend in 2019. I first describe how exposure and productivity are measured, then set out the design with which I identify the break, report it as a contrast between more and less exposed industries, and examine how far it can be trusted.

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) covered by the BLS detailed-industry productivity accounts described below; the time dimension runs from 1987 to 2025, with statutory tariffs from 2002, though individual series cover shorter windows, summarized in Appendix Table A.1. Two objects are measured for each industry: its exposure to the 2018 tariffs, and its productivity. 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.[3]

I work with industries rather than establishments because the industry series are public, so that every number in the paper can be rebuilt from the tariff schedule, the input–output

[3] I set out the harmonization of the NAICS codes in Appendix A.2 and the construction and validation of every series in Appendix A, where I compare the BLS TFP series with the NBER–CES index (Appendix A.4) and the rebuilt protection measure with Flaaen and Pierce's (Appendix A.6).

tables and the BLS releases, without restricted-access data. The cost is that I observe industry aggregates. The results therefore concern industry TFP, not the establishments within each industry.

3.1 Exposure to the 2018 tariffs

Following Flaaen and Pierce (2026), I measure an industry's exposure to the 2018 tariffs by the added duty on its markets and costs. Let Ω denote the set of product–country pairs newly covered by a 2018 U.S. tariff action, a_{pc} the rate the action added on pair pc, and Ω^* and a^*_{pc} the pairs hit by foreign retaliation and the rates it added.[4] The protection measure, Protection in Flaaen and Pierce's notation, is the added duty on an industry's imports as a share of those imports, measured on 2017 trade, times the share of imports in the industry's home market, measured as Flaaen and Pierce (2026) do on 2016 trade and shipments. Writing M, X and Q for an industry's general imports, domestic exports and own shipments, with superscripts for the year,

$$\begin{aligned}
\text{Protection}_j &= \frac{\sum_{pc\in\Omega} a_{pc} M^{2017}_{j,pc}}{M^{2017}_j} \cdot \frac{M^{2016}_j}{Q^{2016}_j + M^{2016}_j - X^{2016}_j}, \\
\text{Retaliation}_j &= \frac{\sum_{pc\in\Omega^*} a^*_{pc} X^{2017}_{j,pc}}{X^{2017}_j} \cdot \frac{X^{2016}_j}{Q^{2016}_j}, \\
\text{InputCost}_j &= \sum_c s_{cj}\, m_c \, \frac{\sum_{pc\in\Omega} a_{pc} M^{2017}_{c,pc}}{M^{2017}_c},
\end{aligned} \tag{1}$$

where s_{cj} is commodity c's share of industry j's costs, intermediate inputs and labor compensation, in the BEA 2012 detail Use table (U.S. Bureau of Economic Analysis, n.d.), and m_c is the share of imports in the domestic absorption of commodity c in 2016. The denominator of the second factor of the protection measure, $Q_j + M_j - X_j$, 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. Because the two years of imports differ by only one, the protection measure is, to a close approximation, the added duty on an industry's imports as a share of the value of its home market.

The three measures use different denominators, each matched to the channel it captures.

[4] Not every 2018 tariff action was aimed at China. The Section 232 tariffs on steel and aluminum (The White House, 2018b) and the Section 201 safeguard tariffs on solar cells and large residential washers (The White House, 2018a) covered particular products and applied to imports from most countries of origin, while the Section 301 tariffs of Lists 1 to 3 (Office of the United States Trade Representative, 2018a,b,c) covered a broad range of products from China alone. Because I sum the added duty over product–country pairs, the protection measure counts the duties of all these actions on the imports they covered, whatever their origin. An industry's exposure to the tariffs therefore differs from its exposure to imports from China.

The protection measure scales the added duty on imports by absorption, the size of the home market in which the industry competes. Retaliation scales the added duty on 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 weights the added duty on each imported commodity by the commodity's import share and by its share of the industry's costs, so that it measures the rise in the industry's costs if the duties passed through in full. I rebuild all three measures from public sources and the code of Flaaen and Pierce (2024), at the level of the 86 industries (Appendix A.6).

Exposure is concentrated. The added duty came to 4.0 percent of the value of the home market, on average, in the most exposed quarter of industries and 0.05 percent in the least exposed quarter.

Because nearly every covered import faced the same additional 25 percent tariff, the protection measure varies mostly with how much of an industry's home market the tariffs reached. In the data of Flaaen and Pierce (2024), the average added tariff on an industry's covered imports is 25 percent at the 10th percentile of industries and 25 percent at the 90th, so the measure ranks industries almost exactly as the share of the home market newly covered does (rank correlation 1.00). I report the estimate with that share as the measure of exposure in Section 3.6.

3.2 Productivity measures

My primary productivity measure is the BLS detailed-industry TFP index (U.S. Bureau of Labor Statistics, 2026c), 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 86 four-digit NAICS manufacturing industries, which define the cross-section of the panel, and it spans 1987–2023. It is a quantity-based measure. BLS distributes nominal sectoral output across detailed product categories and deflates each by that product's own producer price index before aggregating, so an industry's own output price is divided out by construction (Cunningham et al., 2023). This distinguishes it from revenue-based productivity measures, which deflate revenue by a common industry index and so retain establishment-specific prices; the distinction between physical and revenue productivity, and the bias a common deflator imparts, are those of Foster et al. (2008). Three caveats nonetheless apply to the level measure. The combined-input index weights by nominal cost shares in which the cost of capital is a residual, so economic profit is loaded onto the measured capital share and a change in markups moves the weights (Hall, 1988; U.S. Bureau of Labor Statistics, 2026a). Markups therefore enter the level measure through the input weights rather than

through revenue deflation, a narrower channel than the one that afflicts firm-level revenue productivity (De Loecker, 2011). The deflators inherit the coverage and quality-adjustment limitations of the underlying price indexes, and BLS notes that existing techniques "may not fully take into account changes in the quality of goods and services produced, which could result in reduced measures of output" (U.S. Bureau of Labor Statistics, 2026a). And deflating at the product-class level accounts for the changing mix of *existing* product classes but not for goods that appear or disappear, so a policy that changes the set of available varieties moves measured real output for reasons unrelated to technology (Feenstra, 1994). That last point bears on the tariff setting directly, and in Section 4 I take it up with a model in the tradition of Ghironi and Melitz (2005) in which the set of traded varieties is endogenous and measured TFP is built as BLS builds it. As an alternative outcome with a longer recent reach I also use BLS labor productivity (real sectoral output per hour), available through 2025 (U.S. Bureau of Labor Statistics, 2026c).[5]

3.3 The design

With a regression of productivity on an industry's tariff alone I cannot identify the effect of the tariffs, because tariffs are not assigned at random. Adão et al. (2026) estimate that redistributive motives account for about a third of the variation across U.S. tariff lines in 2017, with the highest welfare weights on apparel, motor vehicles and metals. Tariffs are granted for reasons correlated with an industry's prospects, so the protected industries may already be on distinct paths when the tariffs arrive. The industries most exposed in 2018 were, since before the war they were gaining on the others (Appendix Table B.1). I therefore need variation in tariff exposure that is orthogonal to those paths. This is the logic Adão et al. (2025) formalize for testing causal predictions in general equilibrium. An observed change is the causal effect of the tariff plus the effect of every other shock, and only variation in the tariff that is independent of those other shocks separates the two.

The difference-in-differences design of Flaaen and Pierce (2026) supplies that variation. It draws on differences in exposure across industries together with fixed effects, exposure-by-year controls and an explicit pre-trend correction, which removes those distinct paths. Flaaen and Pierce apply it to employment, industrial production and producer prices; no productivity measure appears in their paper, and I apply the design to productivity.

[5] I document the two TFP series in Appendix A.4; Appendix Table A.1 gives the coverage of all three productivity series, and Appendix Figure A.1 plots the average paths of the two BLS series. I validate the BLS TFP series against the NBER–CES Manufacturing Industry Database (Becker et al., 2021), an independently constructed five-factor TFP index. The NBER–CES index 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.

In the annual counterpart of Flaaen and Pierce's event-study regression, I regress 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,

$$\begin{aligned}\tilde{y}_{jt} \;&= b_t^P \,\text{Protection}_j \;+\; b_t^I \,\text{InputCost}_j \;+\; b_t^R \,\text{Retaliation}_j \\ &\quad+\; b_t^M \,\text{ImpShare}_j \;+\; b_t^X \,\text{ExpShare}_j \;+\; b_t^C \,\text{ICShare}_j \;+\; b_t^K \ln(K/L)_j \;+\; d_j \;+\; d_t \;+\; \varepsilon_{jt},\end{aligned} \tag{2}$$

estimated on 86 industries over 2013–2022[6] with standard errors clustered by three-digit sector, as in Flaaen and Pierce (2026).[7] 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 follows the alternative pre-trend correction of Flaaen and Pierce (2026), who remove an industry-specific linear trend fitted over the last year before the tariffs. On annual data I fit it over 2013–17.

The second line holds the four controls of Flaaen and Pierce (2026), each interacted with *every* year dummy. ImpShare_j and ExpShare_j are the industry's pre-war import share of absorption and export share of shipments, measured on 2016 trade and shipments. $\text{ICShare}_j = \sum_c s_{cj} m_c$ is the share of the industry's costs spent on imported inputs, with the cost and import shares of equation (1), and $\ln(K/L)_j$ is log capital per worker in 2016 (Becker et al., 2021). 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.

Because the outcome is net of each industry's own 2013–17 trend, the flat pre-war coefficients reported below are what that removal imposes rather than a test the design passes,

[6] Coating, engraving and heat treating (NAICS 3328) imports nothing and has no tariff lines, so its protection measure, Retaliation and trade shares are zero, as in the data of Flaaen and Pierce (2024). I set its 2017–19 change in the statutory output tariff to zero, which matches their rate-weighted protection measure of zero for it.

[7] I follow the construction of Flaaen and Pierce (2026) for every series of the design: their formulas, data sources and years, their rate-weighted measures of exposure, their four baseline controls and their clustering by three-digit sector. I depart from it in four respects. First, I keep apart the 86 four-digit industries of the BLS productivity accounts. Flaaen and Pierce combine 17 of them into seven groups because their employment data do not separate them, and I rebuild their measures at four digits from the same sources (Appendix A.6). Second, I take the products covered by each 2018 action from the Federal Register notices and proclamations, which correct omissions in their lists (Appendix Table A.3). Third, because productivity is annual, I measure the break with the annual counterpart of their pre-trend correction, equation (3), and with 21 sectors I base inference on the wild cluster bootstrap rather than on the usual approximation. Fourth, I weight every industry equally, whereas Flaaen and Pierce weight industries by value added. I examine robustness to these differences in sensitivity analysis.

since the trend is fitted on the very years those coefficients cover. The identification therefore rests on the pre-trend-adjusted statistic of Flaaen and Pierce (2026), which follows Finkelstein (2007) in measuring the post-war separation against the pre-war drift. They compute it on monthly data, netting the change from just before the tariffs to the end of 2019 against the change over the preceding ten months, with a factor that adjusts for the different lengths of the two windows (their equation 7). On annual data I measure the drift over 2014–17 and scale it to the two-year horizon of the post-war change, so that, with the coefficients normalized to zero in 2017,

$$\Delta_{2019} = (b^P_{2019} - b^P_{2017}) - \tfrac{2}{3}(b^P_{2017} - b^P_{2014}) = b^P_{2019} + \tfrac{2}{3} b^P_{2014}. \tag{3}$$

The factor $\frac{2}{3}$ scales the three-year drift to the two-year horizon of the post-war change, so the statistic removes any industry-specific linear trend exactly and takes the same value whether or not the outcome is detrended first.[8]

3.4 The 2019 break

The estimates of equation (2) trace, year by year, how the TFP of the more exposed industries moved relative to that of the less exposed, and they show a break in 2019. Table 1 reports the coefficients on the protection measure. After detrending, they are flat through 2018, with no pre-tariff coefficient significant at conventional levels, and the gap opens from 2019. The 2019 coefficient on the protection measure is -1.146 ($t = -4.47$), growing to -1.599 by 2021 and -2.154 by 2022; netting out the 2014–17 pre-war drift as in equation (3) gives the pre-trend-adjusted statistic $\Delta_{2019} = -1.292$ ($t = -7.63$; the "Equation (3)" row of Table 1).[9] Intuitively, Δ_{2019} measures how far the more and less exposed industries have separated since 2017 once the drift they were already on is subtracted. The second column replaces the protection measure with the 2017–19 change in the statutory output tariff, the legal rate on an industry's products (the most-favored-nation rate plus the Section 232 and 301 add-ons, averaged over the industry's tariff lines). Because it is built from the tariff schedule rather than from trade flows, it is a check on the protection measure, and it tracks the first column closely.[10]

I state the break on two scales, one for comparison with Flaaen and Pierce (2026) and one

[8] I report equation (2) with no trend removed in Appendix B.1, where the pre-war paths can be read directly.

[9] In Table 1 I cluster the standard errors by three-digit sector, as Flaaen and Pierce (2026) do. With only 21 sectors, I base inference on wild cluster bootstrap p-values by sector (Cameron et al., 2008), which the table reports.

[10] I build the statutory output tariff from the tariff schedule for 2002–2023 and document its construction and the assignment of tariff lines to industries in Appendix A.3.

Table 1: The Flaaen–Pierce design: coefficient path on the measure of exposure

	Detrended log TFP	
	(1) Protection measure	(2) Statutory rate
2013	0.1272	-0.0225
	(0.1115)	(0.0404)
2014	-0.2179	-0.0711
	(0.3417)	(0.1013)
2015	-0.1019	-0.0869
	(0.3179)	(0.0853)
2016	0.0364	-0.1161
	(0.2654)	(0.0856)
2018	-0.3265	-0.1158*
	(0.1934)	(0.0642)
2019	-1.1463***	-0.3390***
	(0.2565)	(0.0931)
2020	-1.0522***	-0.3591**
	(0.3323)	(0.1265)
2021	-1.5994***	-0.5054**
	(0.5258)	(0.1956)
2022	-2.1535***	-0.5869**
	(0.7303)	(0.2525)
Equation (3), 2014–17 pre-trend, 2019	-1.2916***	-0.3864***
	(0.1692)	(0.1165)
Wild bootstrap p, 2019	0.015	0.001
Wild bootstrap p, equation (3)	0.000	0.005
Observations	860	860

Notes: Coefficients b_t^P from equation (2), 2017 omitted as the base year. In column 1 the protection measure is the added duty of the 2018 tariff actions as a share of the value of the home market; column 2 substitutes the change in the statutory output tariff over 2017–19 into the same specification, measured as a fraction (0.01 is one percentage point), so its coefficient is per 100 points. Outcome: log BLS TFP detrended on each industry's 2013–17 path; Appendix Table B.1 reports the same specification with no trend removed, where the pre-war paths can be read directly. The statistic Δ_{2019} removes any industry-specific linear trend exactly and is therefore the same whether or not the pre-war trend is removed (Appendix Table B.1). InputCost, Retaliation and the four controls of Flaaen and Pierce (2026), rebuilt at four digits, are included in both columns and interacted with every year dummy. Industry and year fixed effects; SEs clustered by three-digit sector. The "Equation (3)" row is $b^P_{2019} + \frac{2}{3} b^P_{2014}$, the 2017–19 separation net of the 2014–17 drift scaled to two years; the wild bootstrap rows are p-values of a wild cluster bootstrap by three-digit sector (null imposed, 999 replications) for the 2019 coefficient and for that contrast. Sample: 86 industries, 2013–2022. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

for the model exercise of Section 5. In the left panels of Figure 1 I follow their approach, multiplying each year's coefficient on the protection measure, b_t^P in Table 1, by the interquartile range of the measure. The vertical axis then reads as the TFP gap between industries at the 75th and 25th percentiles of exposure.[11] In the right panels I compare instead the average industries of the top and bottom quartiles of exposure, as I do in Section 5. The top right panel shows the gap in the import-weighted tariff change between the two quartiles. The bottom right panel shows the gap in TFP, measured at each horizon as in equation (3).[12]

The timing follows the tariffs. In the top left panel I plot the gap in the statutory output tariff at the interquartile range, estimated with the same specification, and in the top right panel the gap in the import-weighted tariff change between the quartiles. Less than half of the 2019 tariff gap was in place in 2018 (0.5 of 1.7 points at the interquartile range, 2.5 of 6.6 points between the quartiles), and the TFP gap barely moved that year; it opened in 2019, when the tariff gap was complete. After 2019 the tariff gap moves little, while the TFP gap narrows slightly in 2020 and widens in 2021–22. Those years overlap the pandemic. With controls for each industry's sensitivity to the cycle and its reliance on semiconductors, the path is slightly steeper (Appendix B.5), although other features of the pandemic remain uncontrolled.

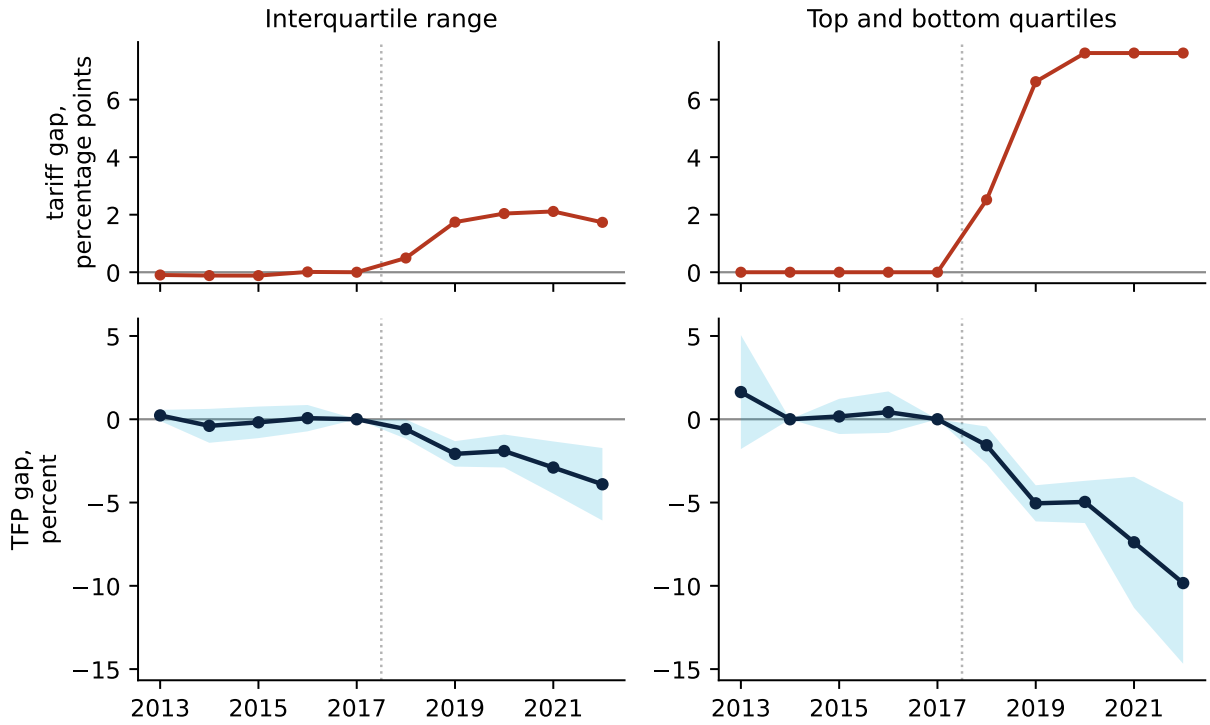
The size of Δ_{2019} from equation (3) comes mostly from subtracting the pre-war drift. In the undetrended series of Appendix Table B.1, the more exposed industries were gaining on the others before the war (a coefficient on the protection measure of -0.819 in 2015 relative to 2017), and after 2017 the gain stopped. No post-war coefficient differs from zero, and the plain 2017–19 change is -0.4291 with a standard error of 0.3332 (Appendix Table B.2). Most of Δ_{2019} is therefore the pre-war gain that did not continue. What I can identify is a break in the relative trend of the more exposed industries at the time of the tariffs, not a fall in their level. Mean reversion after a pre-war run-up would produce the same break. The tests speak against a general tendency of these industries to give back their gains. A break as large appears in 0 of 22 earlier two-year windows (Appendix Table B.4). Nor is the break significantly larger in the industries whose TFP rose fastest before the war (Appendix Table B.7). I therefore cannot rule out a reversion specific to this episode, but I find no sign of a general one.

I apply Flaaen and Pierce's direct test of whether such a drift should be removed, which

[11] The industries at the 25th and 75th percentiles of exposure are forging and stamping (NAICS 3321) and other furniture related products (NAICS 3379), where the added duties amounted to 0.12 and 1.94 percent of the value of the home market.

[12] The gap in TFP at h years after 2017 is the change in log TFP from 2017 net of the 2014–17 drift scaled to h years, regressed across industries on the regressors of equation (2), times the difference between the average protection of the two quartiles, 3.9 points. It is zero by construction in 2014 and 2017.

Figure 1: The Flaaen–Pierce design: tariffs and TFP by exposure, relative to 2017



Notes: Left panels: year-by-year coefficients on the protection measure, the share of the home market newly tariffed weighted by the added tariff rate, from equation (2), with InputCost, Retaliation and the four controls held fixed, scaled by the interquartile range of the protection measure, 1.8 points, between the industries at its 25th and 75th percentiles. Bottom left: the outcome is log BLS TFP detrended on each industry's own 2013–17 path, normalized to zero in 2017, with a 90 percent confidence band from standard errors clustered by three-digit sector. Top left: the same specification with the statutory output tariff, not detrended, as the outcome. Right panels: the top quartile of exposure against the bottom quartile, the comparison of Section 5. Top right: the import-weighted tariff change of the top quartile less that of the bottom quartile, annual averages of the quarterly paths of Appendix Table C.1, held at the 2019Q4 level after 2019. Bottom right: for each year, the contrast of equation (3) at that horizon, the change in log TFP from 2017 net of the 2014–17 drift scaled to the same span, times the quartile-average gap of the protection measure, 3.9 points, with a 90 percent confidence band from standard errors clustered by three-digit sector. It is zero by construction in 2014 and 2017. Its 2018 and 2019 points are the targets of Section 5. Appendix Figure B.1 repeats the left panels with the statutory tariff rate as the measure of exposure. Each row has one vertical scale. The dotted line marks the first 2018 tariffs; the 2020–22 window overlaps the COVID-19 pandemic. Sample: 86 industries, 2013–2022.

estimates the subtracted pre-war term on its own (their online appendix D.6). Regressed on the regressors of equation (2), the pre-war term $\frac{2}{3}(y_{2017} - y_{2014})$ is positive and significant for the protection measure, at 0.8626 ($t = 3.02$), with a randomization p-value of 0.004 and a wild cluster bootstrap p-value by sector of 0.016 (Appendix Table B.3). It accounts for 67 percent of Δ_{2019}, and the post-war change on its own is not significant ($t = -1.29$). For the statutory rate the pre-war term is a smaller part of the statistic, 46 percent, and only borderline ($p = 0.074$ by randomization), while the post-war change comes close to significance on its own ($p = 0.065$ by randomization). By Flaaen and Pierce's logic, a significant pre-trend is what calls for the correction, so the test supports removing the drift. The test also makes plain that the estimate rests on the assumption that without the tariffs the protected industries would have continued their pre-war relative rise. In Flaaen and Pierce's employment and output, by contrast, the more protected industries were on no distinct path before the war, which is why the pre-trend question weighs more for productivity.

3.5 The size of the break and the role of chance

An interquartile shift in exposure is associated with a relative decline in TFP of about 2.3 percent by 2019, measured against the pre-war trend. The decline is sizable by two benchmarks that need no model. The first is these industries' own productivity growth. Over the decade to 2007 the average industry's TFP rose by about 11.8 percent, or 1.2 percent a year, while over the decade to 2017 it changed by −1.5 percent in total. A relative gap of about 2 percent that opens within two years is therefore close to two years of the earlier growth, and larger than the change of the whole decade before the war. The second is the evidence from the two other studies of trade policy whose productivity measures are comparable with mine. Pierce (2011) finds that U.S. antidumping duties lowered the physical productivity of the plants they protected, by more per point of duty than the decline I estimate per point of tariff, while their revenue productivity barely moved. Bernard et al. (2006) find that falling trade costs raised the TFP of U.S. manufacturing industries, measured much as I measure it, by less per point than the decline I estimate. The per-point figure for my estimate comes from the statutory-rate row of Table 2, which scales the same contrast by the gap in statutory tariffs, and I use it only to put the estimate in the unit of these studies. Neither horizon matches mine. Pierce's estimate averages over the first one to nine years after the duty, and that of Bernard et al. (2006) cumulates over the five years after a five-year decline in trade costs, while mine is measured within two years (Appendix B.6). Because the horizons differ, I read the two studies as placing my estimate within the range

of earlier evidence, not as confirming its size.[13]

The 2019 break is not an artifact of chance in any of the senses I can test. These tests go beyond those of Flaaen and Pierce (2026), who cluster standard errors by three-digit sector and run one placebo test before the tariffs, comparing 2015–16 with 2014 (their online appendix D.7.2). Productivity is a noisier outcome than their employment and output, and its pre-war path is not flat. The exposure measures are built from a handful of list-level actions, so identification rests on the exogeneity of the industries' exposure to them (Goldsmith-Pinkham et al., 2020), and the natural test re-assigns the measures at random. A break as large appears in about one draw in a thousand for the protection measure ($p = 0.001$). The wild cluster bootstrap by three-digit sector, which allows for correlation among industries within sectors, rejects a zero effect ($p = 0.001$). Controls for each industry's sensitivity to the cycle leave the estimate within a hundredth of its baseline value, so the break is not a cyclical response to the 2019 slowdown in manufacturing. Where Flaaen and Pierce test one earlier period, I recompute the statistic of equation (3) at every earlier base year the data allow, from 1990 to 2011. I consider the two-year change after each base year, net of the drift over the three years prior. For the protection measure a break as large appears in 0 of 22 windows, and for the statutory rate in none. With at most four non-overlapping windows to a set, I cannot reach a conventional significance level from the ranking alone, but 2019 shows the largest break in every set, for both measures.[14]

3.6 Sensitivity of the estimate

The 2019 break does not depend on how I measure exposure to the 2018 tariffs (Table 2). Each row of Panel A changes one aspect of the measure and keeps the rest of the design fixed. The first row is the baseline, the added duty on an industry's imports as a share of the value of its home market, as Flaaen and Pierce (2026) define it. The second row drops the tariff rate and counts only the share of the home market newly covered. Because almost every covered import faced the same additional 25 percent tariff, the two measures rank industries almost identically, and the break is -1.93 percent against -2.34 in the baseline. The third row uses the measure Flaaen and Pierce (2026) publish, built from their product lists and their groups of industries. It tests whether my reconstruction matters, and it gives -2.20 percent. The fourth row divides the baseline by the industry's import share, so that exposure becomes the average tariff increase on its imports, whatever their weight in its home market.

[13] Firm-level studies of trade policy measure revenue productivity, which moves with prices and markups as well as with physical productivity, so I compare them with the model rather than with my estimate. In Appendix C.4 I set each beside the model's revenue productivity at the study's own window.

[14] I set out each test against chance in Appendix B.2 and collect them in Appendix Table B.4.

This is the tariff that drives the model of Section 5, and it gives −4.00 percent. The fifth row replaces the trade-based measures with the change in the statutory output tariff from 2017 to 2019, an average of schedule rates that uses no trade data. It gives −2.66 percent.

In Panel B I change three choices of the design, one at a time: I weight industries by value added, as Flaaen and Pierce (2026) do; I keep only the two trade-share controls; and I cap the trade shares at their 95th percentiles. The break is −3.34 percent with the weights, −2.23 percent with the two controls and −2.37 percent with the caps.

Table 2: The relative decline in TFP by 2019 is significant under every measure of exposure to the 2018 tariffs

	Coefficient	t, by sector	Wild bootstrap p	Interquartile effect, percent
A. Measure of exposure				
Protection measure, rate-weighted (baseline)	-1.2916	-7.64	0.001	-2.34
Protection measure, coverage share	-0.2675	-4.42	0.007	-1.93
Flaaen and Pierce's published measure	-1.3078	-4.48	<0.001	-2.20
Import-weighted tariff change	-0.6767	-3.71	0.010	-4.00
Change in the statutory output tariff	-0.3864	-3.32	0.007	-2.66
B. Baseline measure, one choice of the design changed				
Industries weighted by value added (2017)	-1.8440	-4.73	0.002	-3.34
Trade shares as the only controls	-1.2327	-6.22	0.002	-2.23
Trade shares capped at the 95th percentile	-1.3098	-6.33	0.002	-2.37

Notes: Equation (3) at the 2019 endpoint. Each row reports the coefficient on the measure of exposure, its t-statistic with standard errors clustered by three-digit sector, the wild cluster bootstrap p-value by sector, and the effect of an interquartile move in percent of TFP. Panel A, the measure of exposure: the baseline, the added duty on an industry's imports as a share of the value of its home market, rebuilt from the Federal Register lists; the share of the home market newly covered, without the rate; the rate-weighted measure of Flaaen and Pierce (2024), with each member of one of their industry groups assigned the group's value; the import-weighted tariff change, the baseline over the industry's import share; and the 2017–19 change in the statutory output tariff. Panel B, the baseline with one choice changed. Every regression includes Retaliation, InputCost, the import share of absorption, the export share of shipments, the share of input costs coming from imported goods and capital per worker, and gives every industry the same weight, except where Panel B says otherwise. Sample: 86 industries.

The break does not reflect idle capacity rather than productivity. BLS computes TFP as output over measured inputs, with capital entering as a stock, so an industry that runs its plants below capacity records lower TFP with no change in its technology (Basu et al., 2006). Following Basu et al. (2006), I proxy unobserved utilization by hours per worker, which firms adjust before they adjust headcount, and run the same contrast on it (Panel D of Appendix Table B.4). Hours per worker did not fall in the more exposed industries, so falling utilization cannot account for the decline in measured TFP. On the rebuilt protection measure the point estimate is a rise of about 1.8 percent for an interquartile move ($t = 2.01$), but the rise does not survive a change of measure. On two other versions of the measure, built with different product lists and with 17 industries pooled into seven groups, it is 0.4 percent ($t = 0.43$) and 0.5 percent ($t = 0.45$), and on the 69 industries outside those groups

it is essentially zero. On the statutory-rate measure it is imprecise. I therefore read the evidence as ruling out a fall in utilization, not as establishing a rise.

The break also appears in output per hour, the BLS measure of labor productivity, when I run the same design on it.[15] Flaaen and Pierce estimate the design separately for industrial production and employment, but output per hour cannot be read off their two estimates, because a pre-trend correction is specific to each outcome and their estimates differ in sample and weights. An interquartile move in the protection measure is associated with output per hour about 3.1 percent lower in 2019, relative to its pre-war path (a coefficient of -1.6922, standard error 0.5933), and random assignment of exposure rarely produces a break as large ($p = 0.011$; Appendix Table B.6). The measure is noisier than TFP against the industries' own history, however, since 5 of the 22 pre-war placebo windows produce a break at least as large (Appendix Table B.6) and the pre-war term carries 69 percent of the statistic (Appendix Table B.3). I therefore read output per hour as confirming the direction of the TFP result rather than as separate evidence on its size.[16]

Whatever the measure, the post-2019 portion of the path overlaps the pandemic, which is why I rest the magnitude on 2019. I ask in Appendix B.3 whether the decline was larger where the tariffs shielded industries from more competition, and in Section 5 what it implies for welfare.

4 A tariff in an open-economy model with heterogeneous firms

The case for a unilateral tariff weighs a gain against a cost. A large country that taxes its imports lowers the world price of what it buys, and that gain in the terms of trade can outweigh the distortion the tax creates. This optimal-tariff argument rests on terms-of-trade effects, which the quantitative trade literature puts at the center of its welfare accounting (Costinot and Rodríguez-Clare, 2014). A relative decline in productivity of the size I estimate adds a cost that this accounting omits. The estimates of the static cost of the 2018–19 tariffs hold technology fixed and find a small aggregate loss (Fajgelbaum et al., 2020; Fajgelbaum and Khandelwal, 2022), driven largely by foreign retaliation in the 2018 estimates of Fajgelbaum et al. (2020). To ask whether a decline in productivity of the estimated size overturns the terms-of-trade gain, I need a model in which a tariff both

[15] Output per hour is a finer measure of labor productivity than output per worker, because firms adjust hours as well as employment, so output per worker misstates productivity whenever hours per worker move.

[16] I also report the same contrast for a revenue-based productivity measure, beside the quantity measures, in Appendix Table B.2.

improves the terms of trade and moves productivity. In this section I describe the model, trace the mechanics of a tariff in its one-sector version, and locate the optimal unilateral tariff against that of Itskhoki and Mukhin (2026); Section 5 extends the model to three sectors, differentially exposed to the tariffs, and interprets the empirical evidence through the lens of that extension.

4.1 The Ghironi–Melitz model

Ghironi and Melitz (2005) build a dynamic general equilibrium model of two countries, home and foreign, that share preferences, technology and parameters but need not be in the same state at any date. Each country has an infinitely lived representative household and a continuum of monopolistically competitive firms, each producing one differentiated variety with labor as the only factor. Firms differ in productivity, drawn once on entry. What distinguishes the model from a standard international business cycle model is that the *set* of goods produced and traded is endogenous and moves with shocks.

Every good is differentiated and made by a single firm, so all trade is intra-industry trade in varieties. Shipping abroad is costly in two distinct ways: an iceberg cost melts part of the cargo, and a fixed cost must be paid every period for selling in the foreign market at all. The fixed cost is the model's central feature. Because it does not depend on how much a firm sells, only firms productive enough to clear it find exporting worthwhile, so at any date only a subset of each country's varieties is traded and the rest are, in effect, nontraded. The size of that traded set is not a parameter. It is determined in equilibrium by a productivity cutoff that moves with the real exchange rate, with foreign demand and with trade costs, and the exporters are the most productive firms, as in the data (Bernard et al., 2007).

Entry is free but sunk. A prospective firm pays an entry cost, then draws its productivity from a Pareto distribution, then produces from the next period until an exogenous death shock removes it. Free entry drives the expected value of entering down to its cost, which pins the mass of firms. The stock of producing firms is the model's main endogenous state besides net foreign assets, accumulated by entry and depreciating with exit, so it plays the role a capital stock plays in other models. Because the number of producers and the number of exporters both move, the extensive margins of trade and of product variety respond to shocks alongside the usual adjustment of prices and quantities.

Households consume a basket of whatever varieties happen to be available, supply labor inelastically, and save in two kinds of assets. They own a mutual fund of their own country's firms, so the entry margin is financed domestically and equity is not traded across borders. Internationally they trade exactly two risk-free bonds, one paying in units of each country's

consumption basket, which are the only channel for international risk sharing and the only reason a current account exists. Because these bonds cannot span the shocks, net foreign assets would follow a unit root; a small quadratic fee on bond holdings removes it and makes the economy stationary. I add one instrument the published model does not have, an ad valorem import tariff in each country, whose revenue is rebated to households as a lump sum.

In the model technology responds to no shock, since each firm keeps the productivity it draws at entry. A tariff therefore changes physical productivity only by reallocating production across firms. Productivity could move along three margins. The first is the intensive margin, the reallocation of production among firms whose productivity is given. The exporters, the most productive firms, sell less abroad, and every firm, including the least productive, sells more at home, so the average productivity of production falls although no firm's own productivity changes. The second is the extensive margin of exporting. When exporting becomes less profitable, the least productive exporters stop exporting, which raises the average productivity of the firms that still export and offsets part of the first margin. Entry moves the number of firms as well, but entrants draw their productivity from the same distribution as incumbents and, in the calibration, all of them sell at home, so entry changes the number of varieties rather than average productivity. The third margin is technology itself, aggregate technology Z_t and the distribution from which firms draw their productivity, and in the model it does not respond to a tariff. I call physical productivity the outcome of all three. It is the productivity the firms experience, the physical output they obtain from their inputs, measured in logs as log technology plus the average log productivity of the firms, weighted by the labor they use in production. It is the productivity concept on which welfare depends. It differs from TFP as statistical agencies measure it, which values each firm's output at its price. Where prices move inversely with productivity, as they do when every firm charges the same markup over marginal cost, a shift of production between firms changes physical output per worker without changing the value of that output, and the measured statistic misses the change in physical productivity.

The published model is unsuited to the question of this paper because of its CES demand. Under CES every firm charges the same markup over marginal cost, so a more productive firm charges a proportionally lower price, and its revenue per unit of production labor equals that of any other firm. Valued at producer prices, the concept behind the BLS measure, output per unit of labor is then blind to how production is allocated across firms. Accordingly, a tariff that shifts production from the most productive firms, the exporters, toward less productive firms that sell only at home leaves it unchanged, even though physical productivity falls. What moves the measure is the labor that exporters spend on the fixed cost of exporting.

A tariff makes exporting less profitable, fewer firms pay that cost, and the labor released from export overhead goes into production, so under CES measured TFP at producer prices rises after a tariff (Section 4.4), the opposite of the sign the evidence of Section 3 points to. My baseline therefore replaces CES with translog demand, as in Bilbiie et al. (2012), whose closed-economy model has identical firms, and as in Rodríguez-López (2011), who brings it to a two-country model with heterogeneous firms in the tradition of Ghironi and Melitz (2005) to study exchange-rate pass-through; I use that combination to study tariffs. Under translog a firm's demand becomes less elastic as its market share grows, so larger, more productive firms charge higher markups, and demand has a choke price above which a firm sells nothing. A shift of production toward less productive firms then moves labor toward lower-markup firms and pulls measured TFP down, so the reallocation the tariff causes registers in the measure the evidence is about. Throughout I compare the baseline with the published CES model.[17]

4.2 The households' and firms' problems

The equations follow the description above, first under the CES demand of the published model and then with the change that translog demand makes.[18]

Each country has a representative household and a continuum of monopolistically competitive firms, each producing one variety with labor alone. A home firm with productivity draw z produces $Z_t z$ units per unit of labor, where Z_t is aggregate technology. Entry requires a sunk cost of f_E effective labor units; after paying it the entrant draws z once and for all from a Pareto distribution $G(z) = 1 - (z_{\min}/z)^k$ on $[z_{\min}, \infty)$, and it starts producing the next period. Firms die at the exogenous rate δ, independently of z. There is no fixed cost of production, so every firm that has entered produces, and G is also the distribution of producers. Exporting requires an iceberg cost $\tau \geq 1$ per unit delivered and a fixed cost of f_X effective labor units per period, so only the more productive firms export. The household supplies L units of labor inelastically, owns the home firms through a mutual fund, and trades a home and a foreign bond, paying a small quadratic fee on bond holdings that makes net foreign assets stationary. The foreign country mirrors the home country. Each government levies an ad valorem tariff on its imports and rebates the revenue lump sum: t_t at home, t_t^* abroad.

[17] I state both models in full in Appendix C.

[18] I complete the statement of the model, with the equilibrium conditions, the steady state and the measurement of productivity and welfare, in Appendix C.1 under CES and in Appendix C.2 under translog.

4.2.1 Households

The household's problem is the standard one, with the bond fee as its only departure. The home household maximizes $E_t \sum_{s\geq t} \beta^{s-t} C_s^{1-\gamma}/(1-\gamma)$, where consumption is a CES aggregate over the set Ω_t of varieties available at home,

$$C_t = \left[\int_{\Omega_t} c_t(\omega)^{\frac{\theta-1}{\theta}} d\omega\right]^{\frac{\theta}{\theta-1}}, \qquad P_t = \left[\int_{\Omega_t} p_t(\omega)^{1-\theta} d\omega\right]^{\frac{1}{1-\theta}}, \tag{4}$$

with $\theta > 1$ the elasticity of substitution and P_t the welfare-based price index, so that demand for a variety priced at p is $c = (p/P)^{-\theta} C$. Writing B_{t+1} and $B_{*,t+1}$ for holdings of home and foreign bonds, x_t for shares in the mutual fund of the $N_{H,t} = N_{D,t} + N_{E,t}$ home firms in existence, $\tilde{v}_t$ for the price of a claim to a firm and $\tilde{d}_t$ for average profit per producer, the budget constraint in units of home consumption is

$$\begin{aligned} &B_{t+1} + Q_t B_{*,t+1} + \tfrac{\eta}{2} B_{t+1}^2 + \tfrac{\eta}{2} Q_t B_{*,t+1}^2 + \tilde{v}_t N_{H,t} x_{t+1} + C_t \\ &\quad = (1 + r_t) B_t + Q_t (1 + r_t^*) B_{*,t} + (\tilde{d}_t + \tilde{v}_t) N_{D,t} x_t + w_t L + T_t^f + T_t, \end{aligned} \tag{5}$$

where Q_t is the real exchange rate in units of home consumption per unit of foreign consumption, w_t the real wage, r_t and r_t^* the real returns on the two bonds, T_t^f the rebated bond fee and T_t the rebated tariff revenue. With $\lambda_t = C_t^{-\gamma}$, the first-order conditions for shares and the two bonds are

$$\tilde{v}_t = \beta(1-\delta) E_t \left[\left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \left(\tilde{v}_{t+1} + \tilde{d}_{t+1} \right) \right], \tag{6}$$

$$C_t^{-\gamma}(1 + \eta B_{t+1}) = \beta(1 + r_{t+1}) E_t C_{t+1}^{-\gamma}, \qquad C_t^{-\gamma}(1 + \eta B_{*,t+1}) = \beta(1 + r_{t+1}^*) E_t \left[\frac{Q_{t+1}}{Q_t} C_{t+1}^{-\gamma} \right]. \tag{7}$$

A claim to a firm pays only if the firm survives, which is why $1-\delta$ discounts the share Euler equation. The fee is a wedge in the bond Euler equations and, being rebated, not a resource cost.

4.2.2 Firms

The firms' problem determines prices, profits and entry. A firm takes the aggregates as given and sets a price in each market it serves. Its real marginal cost of serving the home market is $w_t/(Z_t z)$ and, per unit delivered abroad, $\tau w_t/(Z_t z)$. Under CES the demand elasticity is θ in every market and for every firm, so the markup over marginal cost is the constant

$A \equiv \theta/(\theta-1)$ and, with prices relative to the destination's price index,

$$\rho_{D,t}(z) = A\frac{w_t}{Z_t z}, \qquad \rho_{X,t}(z) = \frac{\tau}{Q_t}\,\rho_{D,t}(z). \tag{8}$$

A foreign tariff multiplies the demand for a home export by $(1+t_t^*)^{-\theta}$ and leaves its elasticity unchanged, so it does not move the producer price. Profit is revenue over θ in each market, net of the fixed export cost,

$$d_{D,t}(z) = \frac{1}{\theta}\rho_{D,t}(z)^{1-\theta}C_t, \qquad d_{X,t}(z) = \frac{Q_t}{\theta}(1+t_t^*)^{-\theta}\rho_{X,t}(z)^{1-\theta}C_t^* - \frac{w_t f_X}{Z_t}, \tag{9}$$

and a firm exports if and only if $d_{X,t}(z) \geq 0$. Gross export profit rises with $z^{\theta-1}$ while the overhead does not depend on z, so exporters are the firms above a cutoff $z_{X,t}$. Free entry equates the value of a firm with the entry cost, and the mass of producers evolves with entry and exit,

$$\tilde{v}_t = \frac{w_t f_E}{Z_t}, \qquad N_{D,t} = (1-\delta)\big(N_{D,t-1} + N_{E,t-1}\big). \tag{10}$$

Translog demand changes the demand each firm faces. Under translog the expenditure share of a variety with consumer price p, in a market where a measure N_t of varieties is sold, is

$$s = \frac{1}{N_t} + \sigma\big(\bar{L}_t - \ln p\big), \qquad \bar{L}_t = \frac{1}{N_t}\int_{\text{sold}} \ln p_j \, dj, \tag{11}$$

with $\sigma > 0$ a preference parameter. The share reaches zero at a finite choke price, $\ln p^{\max} = \bar{L}_t + 1/(\sigma N_t)$, and a firm taking N_t and $\bar{L}_t$ as given faces a perceived demand elasticity $1 + \sigma/s$, so its markup is

$$\mu = 1 + \frac{s}{\sigma}. \tag{12}$$

Markups rise with market share, because a larger, more productive firm faces less elastic demand and keeps part of its cost advantage as margin instead of passing it all into price. Substituting (12) into (11), a firm's share solves $s + \sigma\ln(1+s/\sigma) = 1/N_t + \sigma(\bar{L}_t - c)$, where c is its log marginal cost in the destination, $\ln[w_t/(Z_t z)]$ at home and $\ln[(1+t_t^*)\tau w_t/(Q_t Z_t z)]$ abroad. A firm's behavior therefore depends on its own cost and on two aggregates of the market it sells into, N_t and $\bar{L}_t$. Its profit in a market with real expenditure C is $d(z) = \frac{C}{1+t}\,s(z)^2/(\sigma + s(z))$. The choke price creates a domestic cutoff, $\ln w_t - \ln(Z_t z_{D,t}) = \bar{L}_t + 1/(\sigma N_t)$, which CES lacks; at the calibration below it lies below $z_{\min}$, so every entrant sells at home. The export cutoff is the familiar one, gross export profit equal to $w_t f_X/Z_t$.

4.2.3 Aggregation and resource constraints

Under Pareto draws and constant markups every aggregate is a mass of firms times the firm-level function evaluated at an average productivity: $\tilde{z}_D = \nu z_{\min}$ for all producers and $\tilde{z}_{X,t} = \nu z_{X,t}$ for exporters, with $\nu = [k/(k-\theta+1)]^{1/(\theta-1)}$, which requires $k > \theta - 1$. A tilde marks a function evaluated at the relevant average. The share of producers that export and the zero-profit condition at the cutoff take the forms

$$\frac{N_{X,t}}{N_{D,t}} = \left(\frac{z_{\min}}{z_{X,t}}\right)^k, \qquad \tilde{d}_{X,t} = \frac{w_t f_X}{Z_t}\,\frac{\theta-1}{k-(\theta-1)}. \tag{13}$$

Four conditions close the model. The home price index, normalized to one, is built from the home producers and the foreign exporters, whose consumer price carries the home tariff; labor is used in production, in export overhead and in entry; each bond is in zero net supply; and combining the two budget constraints gives the law of motion of home net foreign assets:

$$N_{D,t}\tilde{\rho}_{D,t}^{1-\theta} + N_{X,t}^{*}\big((1+t_t)\tilde{\rho}_{X,t}^{*}\big)^{1-\theta} = 1, \tag{14}$$

$$L = \frac{\theta-1}{w_t}\big(N_{D,t}\tilde{d}_{D,t} + N_{X,t}\tilde{d}_{X,t}\big) + \frac{1}{Z_t}\big(\theta N_{X,t} f_X + N_{E,t} f_E\big), \tag{15}$$

$$B_{t+1} + B_{t+1}^{*} = 0, \qquad B_{*,t+1} + B_{*,t+1}^{*} = 0, \tag{16}$$

$$\begin{aligned} B_{t+1} + Q_t B_{*,t+1} = {} & (1+r_t)B_t + Q_t(1+r_t^{*})B_{*,t} + \tfrac{1}{2}\big(w_t L - Q_t w_t^{*} L^{*}\big) + \tfrac{1}{2}\big(N_{D,t}\tilde{d}_t - Q_t N_{D,t}^{*}\tilde{d}_t^{*}\big) \\ & - \tfrac{1}{2}\big(N_{E,t}\tilde{v}_t - Q_t N_{E,t}^{*}\tilde{v}_t^{*}\big) + \tfrac{1}{2}\big(T_t - Q_t T_t^{*}\big) - \tfrac{1}{2}\big(C_t - Q_t C_t^{*}\big), \end{aligned} \tag{17}$$

with $\tilde{d}_t = \tilde{d}_{D,t} + (N_{X,t}/N_{D,t})\tilde{d}_{X,t}$ and tariff revenue $T_t = t_t(1+t_t)^{-\theta} N_{X,t}^{*}(\tilde{\rho}_{X,t}^{*})^{1-\theta} C_t$. The overhead enters the labor market once directly and $\theta - 1$ times through the production labor implied by $\tilde{d}_{X,t}$, which is net of it, so f_X carries the factor θ.

Under translog, the aggregation changes because markups differ across firms. Under CES every firm's price is the same markup over its marginal cost, so prices, sales and labor are powers of productivity, and with Pareto draws each aggregate reduces to the mass of firms times a function of one average productivity. Under translog a firm's markup depends on its market share, and its market share depends on its price, so the share of each firm has to be solved for, firm by firm, and no single average summarizes the distribution. Aggregates are instead sums over the productivity distribution, which I compute at a fixed set of productivity nodes. Each market needs three moments of its log prices: the measure of varieties, the mean and the second moment, N_t, $\bar{L}_t N_t = \int \ln p$ and $M_{2,t} N_t = \int (\ln p)^2$, the integrals running over the home producers and the foreign exporters at consumer prices. The translog price index,

normalized to one, replaces (14) and pins the real wage,

$$\bar{L}_t + \frac{1}{2\sigma N_t} - \frac{\sigma N_t}{2}\left(M_{2,t} - \bar{L}_t^2\right) = 0, \tag{18}$$

where I take the limit in which the measure of potential varieties is unbounded. The labor market replaces (15) with revenue over markup in each market,

$$L = N_{D,t} E\left[\frac{s_D C_t}{\mu_D w_t}\right] + N_{X,t} E\left[\frac{Q_t s_X C_t^*}{(1+t_t^*)\mu_X w_t}\right] + \frac{N_{X,t} f_X}{Z_t} + \frac{N_{E,t} f_E}{Z_t}, \tag{19}$$

the expectations taken over producers and over exporters. Every expectation over Pareto draws above a cutoff z_c is evaluated by Gauss–Laguerre quadrature. The substitution $y = k\ln(z/z_c)$ turns dG into $\exp(-y)\,dy$, so $\int_{z_c}^{\infty} f\,dG = (z_{\min}/z_c)^k \sum_i \omega_i f\left(z_c \exp(t_i/k)\right)$ with fixed nodes and weights, and the nodes move with the cutoff, which keeps the extensive margin differentiable. Each node's share is an endogenous variable whose defining equation is part of the model.

4.3 Solution and calibration

I solve every model exactly, by perfect foresight. Every experiment shocks the economy in the first period only, with announced paths entering through news shocks dated then, so each is a deterministic path back to the steady state. I stack the equilibrium conditions over 6905 quarters, long enough for a shock with persistence 0.999 to decay to a thousandth of its size, and solve them by Newton's method with the perfect-foresight solver of Dynare (Adjemian et al., 2024). Newton's method needs a starting guess for the path of every variable over the whole horizon. I start from the path implied by the linear decision rules rather than from the steady state, and I treat a solution that does not converge in a few Newton steps as a failure rather than something to rescue. Welfare is conditional welfare in the first period, from that path, as a consumption equivalent.

I keep the calibration of Ghironi and Melitz (2005): $\beta = 0.99$, $\gamma = 2$, $\delta = 0.025$, $\theta = 3.8$, $k = 3.4$, $\tau = 1.3$, $f_E = z_{\min} = L = Z = 1$, f_X set so that 21 percent of firms export, and a bond fee $\eta = 0.0025$. The tariff follows $t_t = 0.999\, t_{t-1} + \varepsilon_t^t$, nearly permanent, as the tariff of Itskhoki and Mukhin (2026) is permanent. Under translog, all parameters shared with the CES model keep their values, and two are recalibrated: σ, so that the expenditure-weighted average markup equals the CES markup $\theta/(\theta-1) = 1.357$, which gives $\sigma = 0.565$; and f_X, so that 21 percent of firms export. For comparability with the published model, I give the two economies the same average markup and the same share of exporters. Bilbiie et al. (2012) make a similar choice when they compare translog with CES. With identical firms,

they set the translog parameter so that the steady-state markup and the number of firms both equal their CES values, because the data offer only one value of each. The switch to translog changes how markups are distributed across firms. Under CES every firm charges $\theta/(\theta-1)$, while under translog larger and more productive firms charge higher markups.

For the comparison with the three-sector model of Section 5, I also calibrate the one-sector model with home a quarter of the world (Table 3), setting the foreign labor force to 3 times home's.[19] Three parameters are reset so that the economy keeps the moments of the model with countries of equal size. The translog parameter σ keeps the home expenditure-weighted markup at 1.357, τ keeps the import share of absorption at 0.246, and f_X keeps the share of exporting firms at 21 percent. A fourth parameter is new. The foreign market has its own translog parameter σ^*, set so that its expenditure-weighted markup equals home's. Modeled as one integrated market, a foreign economy three times larger would offer each buyer about three times as many varieties, and under translog demand its markups would fall well below home's. Intuitively, with σ^* the rest of the world behaves like many markets about as competitive as home's, while its total spending remains three times larger. With countries of equal size $\sigma^* = \sigma$, so the symmetric model is the published one.

4.4 A tariff in the one-sector model

A persistent home tariff has three effects in the model (Figure 2): it improves the home terms of trade, it shifts production toward less productive firms, and it raises measured TFP at producer prices even as physical productivity falls. The first is the terms of trade. The tariff appreciates the home real exchange rate, which moves by -0.506 percent per point on impact under translog, and makes home labor dearer relative to foreign labor; home consumption rises and foreign consumption falls. The second is the allocation of production across firms. The appreciation makes exporting less profitable, so the most productive firms, which are the exporters, produce less for foreign markets, while home demand shifts toward domestic sales, which every firm makes, including the least productive. Production shifts toward less productive firms, and physical productivity changes on impact by -0.138 percent per point under translog and by -0.0814 percent under CES, although no firm's own productivity changes. Under translog the intensive margin accounts for more than the whole fall, -0.159 percent, and the extensive margin of exporting offsets a small part of it, 0.0209 percent, as the least productive exporters stop exporting (Appendix C.2 defines the split). The third is measurement, where the two demand systems differ in magnitude.

Whether measured TFP rises or falls after the tariff depends on how output is valued.

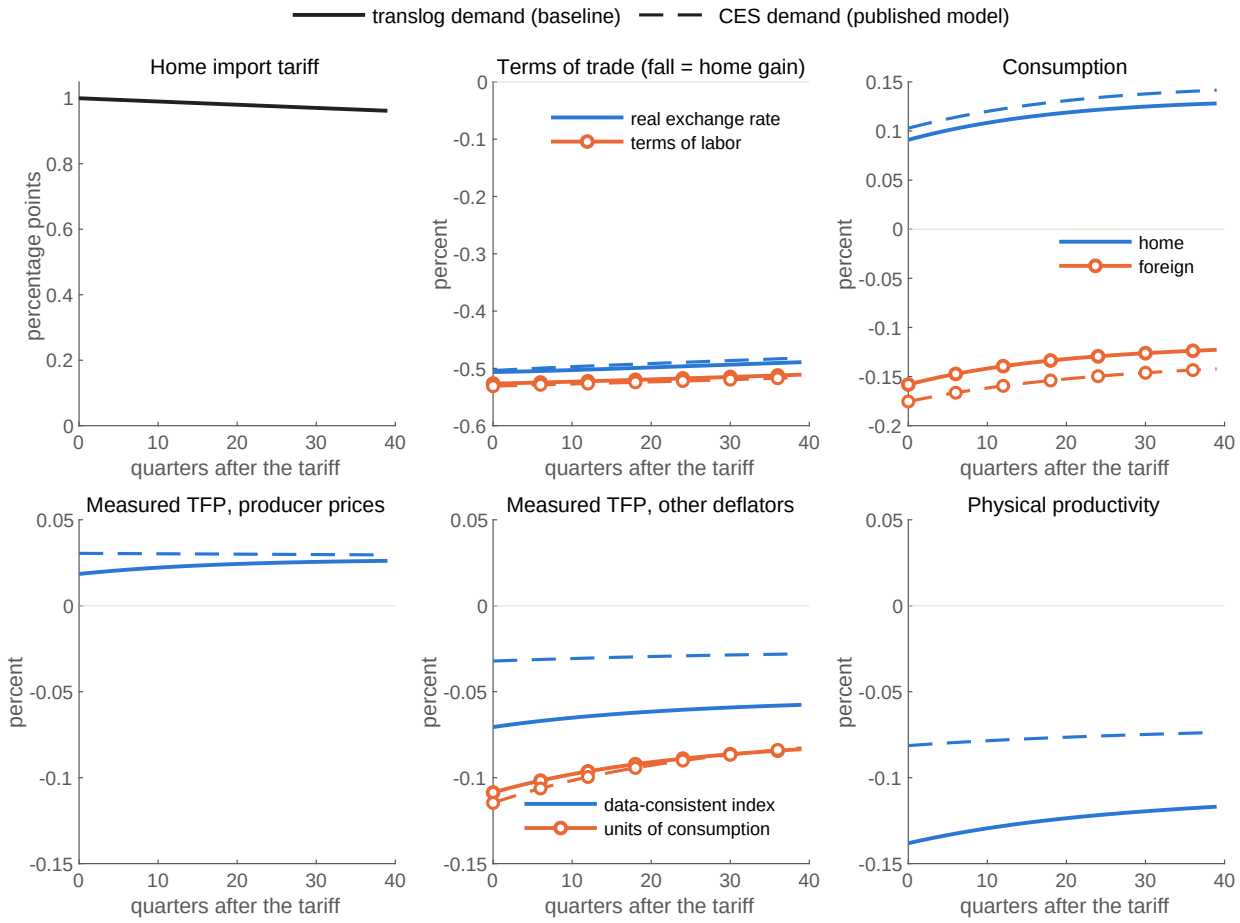
[19] The United States accounted for 23.9 percent of world GDP in 2017 (World Bank, 2026).

Table 3: Calibration of the one-sector translog model

Parameter	Controls	Target	Value	
			Equal size	Home a quarter
β, γ, δ	discounting, risk aversion, exit	published values	0.99, 2, 0.025	0.99, 2, 0.025
$f_E, z_{\min}, L$	normalizations	published values	1, 1, 1	1, 1, 1
L^*	size of the foreign economy	U.S. share of world GDP, about a quarter	1	3
σ	markup	expenditure-weighted markup 1.357, the CES markup $\theta/(\theta-1)$, $\theta = 3.8$	0.5654	0.5252
σ^*	competition abroad	foreign markup equal to home's	0.5654	0.1896
k	dispersion of firm productivity	the published value	3.4	3.4
$\tau = \tau^*$	trade cost	import share of absorption, 0.246 (the published $\tau = 1.3$ at equal size)	1.3	1.654
f_X	fixed export cost	share of firms that export, 0.21	0.01035	0.006533

Notes: Equal size is the calibration of Ghironi and Melitz (2005), with σ set, following Bilbiie et al. (2012), so that the expenditure-weighted markup equals the CES markup. With home a quarter of the world, the foreign labor force is 3 times home's, and σ, σ^*, τ and f_X are set so that the home and foreign markups, the import share of absorption and the share of exporting firms equal their values at equal size. No domestic or export cutoff binds in either calibration.

Figure 2: A nearly permanent one percent home tariff in the one-sector model



Notes: Solid lines: translog demand, the baseline; dashed lines: the published CES model. Within a panel, hollow circles mark the second series of the legend. The tariff is the same in both and follows an AR(1) with coefficient 0.999; quarter 0 is the quarter it is imposed. Terms of trade: the real exchange rate and the terms of labor, the cost of foreign relative to home effective labor; a fall is a home gain. Measured TFP: revenue over the labor of the home producers, valued at each firm's steady-state price (producer prices, the concept BLS measures), deflated by the data-consistent price index, or in units of consumption. Physical productivity: the production-labor-weighted mean of log productivity. Paths from the exact deterministic solution (perfect foresight), percent deviations from the steady state.

At producer prices, as BLS values it, measured TFP rises under both demand systems, while in units of consumption it falls with the real wage (Table 4 reports the impact response term by term). When each firm's output is valued at its own steady-state price, as a product-by-product deflator does, the measure moves with the share of labor absorbed by export overhead and, under translog, with the allocation of labor across firms with different markups. The shift toward less productive, lower-markup firms contributes −0.0080 percent per point, a term that is exactly zero under CES, where a more productive firm charges a proportionally lower price. Under CES, measured TFP at producer prices therefore rises, by 0.0305 percent per point, entirely because export overhead shrinks, while physical productivity changes by −0.0814 percent. Under translog the markup term offsets about a third of the overhead effect, and the measure rises by only 0.0185 percent. Valuing output in units of consumption adds the fall in the real wage, the main reason measured TFP in units of consumption falls after a tariff. Deflating by the data-consistent price index, which averages the prices of the varieties available instead of crediting their number, offsets part of that fall.

Table 4: Measured TFP in the quarter a nearly permanent one percent tariff is imposed, percent deviation from the steady state

	Translog	CES
Technology of the firms	0.000	0.000
Labor share outside export overhead, $1 - s_f$	0.027	0.030
Labor weighted by steady-state markups	−0.008	0.000
= TFP at producer prices	0.018	0.030
+ real wage, units of consumption	−0.135	−0.145
+ change in markups	0.008	0.000
= TFP in units of consumption	−0.109	−0.115
+ variety term	0.038	0.083
= TFP deflated by the data-consistent index	−0.071	−0.032
Labor-weighted mean of $\ln z$, all labor	−0.133	−0.040
Labor-weighted mean of $\ln z$, production labor	−0.138	−0.081

Notes: s_f is the share of the producers' labor absorbed by export overhead. The last row is physical productivity (Section 4.1). Producer prices value each firm's output at its steady-state price; the data-consistent index is the average price of the varieties available at home, which does not credit changes in their number (Ghironi and Melitz, 2005). Formulas in Appendices C.1.3 and C.2.2.

The gap between what statistical agencies can measure and physical productivity in the model is not a feature of the model alone. Ghironi and Melitz (2005) note that the average prices that correspond to empirical price indexes omit the variety effects that welfare-based indexes contain, and that "statistical agencies typically do not adjust price indexes to reflect differences in the price levels of new goods." The BLS documents the same limits for the

producer prices it uses to deflate industry output: "it is not always possible to obtain a value for a quality adjustment" (U.S. Bureau of Labor Statistics, 2026b), and its productivity measures may "not fully take into account changes in the quality of goods and services produced" (U.S. Bureau of Labor Statistics, 2026a). The empirical estimate is of measured TFP, with these imperfections. I therefore match the model to the evidence on measured TFP, built as BLS builds it, and with the model translate the measured change into the physical change on which welfare depends.

Table 5: Welfare effect of a nearly permanent one percent tariff, consumption equivalent, percent

	Translog		CES	
	Home	Foreign	Home	Foreign
Home tariff only	+0.1209	−0.1191	+0.1341	−0.1359
Both countries levy the tariff	+0.0002	+0.0002	−0.0037	−0.0037
Offsetting technology decline, percent per point	0.116		0.123	

Notes: Conditional welfare when the tariff arrives, reported as the permanent change in consumption with the same value, $\lambda = (\mathcal{W}/\mathcal{W}^{ss})^{1/(1-\gamma)} - 1$ (equation (C.7)), from the exact deterministic solution. Offsetting technology decline: the persistent decline in home technology, at the tariff's persistence, whose welfare cost equals the home gain from the tariff.

Welfare, by contrast, follows the terms of trade (Table 5). A unilateral tariff raises home welfare by 0.121 percent of consumption per point[20] under translog and 0.134 percent under CES, and the gain comes mostly at the foreign country's expense, whose welfare changes by −0.119 and −0.136 percent. When the foreign country levies the same tariff, the transfers cancel and the home gain disappears. The last row restates the gain in productivity units, as the persistent decline in home technology of 0.116 percent per point of tariff under translog, and 0.123 percent under CES, that would cost as much as the tariff gains. These welfare figures already include the reallocation the tariff causes, so a productivity loss can reverse the welfare effect of a tariff only if it comes on top of that reallocation, from a margin the model holds fixed, and persists.

[20]Welfare is the home household's conditional lifetime utility in the quarter the tariffs are announced, $\mathcal{W} = E_t \sum_{s\geq 0} \beta^s C_{t+s}^{1-\gamma}/(1-\gamma)$, which the model carries as the recursion $\mathcal{W}_t = C_t^{1-\gamma}/(1-\gamma) + \beta E_t \mathcal{W}_{t+1}$ and which I compute from the exact deterministic solution of the model, by perfect foresight. As in the long-run model of Itskhoki and Mukhin (2026), the setup is deterministic, so welfare leaves aside how a tariff changes the welfare cost of fluctuations. The consumption equivalent is the permanent proportional change λ in steady-state consumption $\bar{C}$ that would give the same welfare, $\sum_{s\geq 0} \beta^s [(1+\lambda)\bar{C}]^{1-\gamma}/(1-\gamma) = \mathcal{W}$, so that $\lambda = (\mathcal{W}/\mathcal{W}^{ss})^{1/(1-\gamma)} - 1$ with $\mathcal{W}^{ss} = \bar{C}^{1-\gamma}/[(1-\gamma)(1-\beta)]$; with $\gamma = 2$, $\lambda = \mathcal{W}^{ss}/\mathcal{W} - 1$. Tariff revenue is rebated lump sum and so enters through consumption. Appendix C.1.3 gives the details.

4.5 The optimal unilateral tariff

Firm heterogeneity does not reverse the optimal tariff that Itskhoki and Mukhin (2026) find. I compute the optimal tariff for the experiment they consider, a tariff imposed by surprise and kept in place, which I model as nearly permanent, with a persistence of 0.999 a quarter. For such a tariff home welfare first rises with the tariff, because the terms of trade improve, and then falls, because the efficiency loss grows faster than the gain. Two objects summarize the curve: the gain per point at a zero tariff, a, which governs the first tariffs and the threshold below, and the optimal tariff t^{opt}, the tariff at which welfare peaks. I compute welfare exactly, from the deterministic solution of the model by perfect foresight, and locate t^{opt} by search (Appendix C.3.5). The exact solution matters for the optimum. A second-order approximation, accurate for small tariffs, overstates how fast welfare falls at large ones and places the optimum in the one-sector translog model at 20 percent, about half its exact value of 39. Itskhoki and Mukhin compute it in a two-country Armington model, in which each country produces a single national good. Firms there do not differ in productivity, so a tariff cannot shift production toward less productive firms. I first ask how that channel changes the optimal tariff in the one-sector model; with the three-sector model below I bring the evidence to bear on it.

With heterogeneous firms the optimal tariff stays large (Table 6). It is 38 percent in the published CES model and 39 percent under translog. Itskhoki and Mukhin (2026) put the United States at a fifth of world output. In the one-sector model with home a quarter of the world (Section 4.3), the counterpart of their economy without cross-border positions, the exact optimal tariff is 36 percent, close to their 35 percent. To gauge the role of heterogeneity, I solve versions of the one-sector model in which every firm has the same productivity and every firm exports, so that a tariff cannot reallocate production across firms. Each identical-firm model is calibrated to the moments its heterogeneous-firm counterpart matches: the import share of spending and, under translog, the expenditure-weighted markup (Appendix C.5). From these comparisons I cannot sign the effect of heterogeneity on the optimum. Under CES, heterogeneity lowers the optimal tariff from 47 to 38 percent. Under translog, the identical-firm optimum is 45 percent with σ set to match the markup and 26 percent with σ held at the heterogeneous-firm value, on either side of 39. The benchmarks differ in the elasticity of substitution between domestic and foreign goods as well as in heterogeneity, so these comparisons do not isolate the reallocation. The external position, by contrast, cuts the optimum by more than half. I compute it in the model with countries of equal size, where the gross position of the end of 2017 lowers the optimum from 39 to 18 percent, against their cut to a 9 percent tariff.

The comparison with Itskhoki and Mukhin (2026) is made at an elasticity close to theirs.

In their analysis the optimal tariff is governed by the elasticity of substitution between domestic and foreign goods, which they set to 4. In the Ghironi–Melitz model the same parameter also sets the substitution across varieties and the markup, and the baselines imply an elasticity between domestic and foreign goods of 3.80 under CES and 4.59 under translog (Appendix C.5), on either side of their value. Under translog the elasticity also varies across firms. The demand elasticity a home exporter faces abroad ranges from 8.16 at the export cutoff to 2.90 at the 90th percentile of exporters, with an export-sales-weighted mean of 4.27. Because the translog calibration raises the elasticity between domestic and foreign goods relative to CES, Table 6 also reports the translog model recalibrated so that this elasticity equals the CES value, with the same import share and share of exporting firms. The gain per point rises from 0.124 to 0.129, closing about a third of the gap with CES, and the optimal tariff barely moves, from 39.2 to 38.6 percent. Much of the difference between the two demand systems therefore reflects how markups are distributed across firms rather than the elasticity, although the recalibration also raises the average markup (Appendix C.5). A nested demand system, in which an outer elasticity between a domestic and a foreign bundle differs from the inner elasticity across varieties, would let the two move apart. Feenstra et al. (2018) find such a gap for between a quarter and a third of the goods they study, where the micro elasticity is significantly higher than the macro elasticity.

5 The empirical evidence through the lens of the model

To interpret the estimate with a model, I need one that distinguishes more from less exposed industries. In a one-sector model the relative decline would have to stand for a decline in manufacturing as a whole, an assumption the evidence cannot support, because the empirical specification compares industries and differences out everything common to them. I therefore extend the model to three sectors and compare home welfare under the tariffs alone and under the tariffs with the estimated decline (Section 5.1). I then ask how that welfare comparison depends on how much of the estimated gap is a decline in the most exposed industries rather than a rise in the others (Section 5.2).

5.1 The three-sector model

In the three-sector model the 2018 tariffs fall unevenly across sectors, as they fell across the industries in the data. The estimate fixes the gap in measured TFP between the most and the least exposed sectors, not how that gap divides between them. In the baseline I attribute the whole gap to a decline in the technology of the most exposed sector, an assumption I relax

Table 6: The optimal unilateral tariff in the one-sector models

	Gain per point, a	Optimal tariff, percent	Welfare at the optimum	Home–foreign elasticity
A. Demand system, heterogeneous firms				
CES (published model)	0.1373	38.1	2.026	3.80
Translog, markup of CES (baseline)	0.1237	39.2	1.860	4.59
Translog, home–foreign elasticity of CES	0.1288	38.6	1.929	3.80
B. Identical firms				
CES	0.1325	47.3	2.325	3.80
Translog	0.1759	44.8	3.414	5.75
C. External positions, translog with heterogeneous firms				
Gross positions of 2017	0.0766	17.7	0.602	
Net position only, in home currency	0.1130	31.9	1.454	
D. Country size, translog with heterogeneous firms				
Home a quarter of the world	0.1148	35.9	1.586	

Notes: One-sector models, labor supply inelastic. An unanticipated home tariff with persistence 0.999; welfare as the exact consumption equivalent from the exact deterministic solution (perfect foresight). The gain per point is the slope of welfare at a zero tariff, from tariffs of 1 and 2 percent; the optimal tariff maximizes welfare, and the third column is welfare there, percent of consumption. Identical firms: every firm has the lowest productivity and the fixed export cost is zero, so every firm exports; the iceberg cost matches the import share of the heterogeneous-firm model, 0.267 under CES and 0.246 under translog, and under translog σ matches its expenditure-weighted markup. Home–foreign elasticity: the elasticity of substitution between domestic and foreign goods, equation (C.12), comparable with the value of 4 in Itskhoki and Mukhin (2026). Translog, home–foreign elasticity of CES: σ, f_X and τ are set so that the elasticity of substitution between domestic and foreign goods equals the CES value, 3.80, with the import share and the share of exporting firms of the translog baseline (Appendix C.5). External positions: the translog model with three identical sectors, which reproduces the one-sector model exactly; the gross and net positions of the end of 2017 (Section 5.1.3). Home a quarter of the world: the same model with a foreign labor force 3 times home's, calibrated as in Table 3.

in Section 5.2. I first describe how the model departs from the one-sector model and how I calibrate its sectors to the industries of Figure 1 (Section 5.1.1). I then compare the tariffs alone with the tariffs and the estimated decline (Section 5.1.2), check how that comparison depends on the model's assumptions and on the external position (Section 5.1.3), and ask how a productivity response changes the optimal tariff (Section 5.1.4).

5.1.1 Environment and calibration

Each country now has three manufacturing sectors instead of one, standing for the quarter of industries most exposed to the 2018 tariffs, the middle half and the least exposed quarter, ranked by the protection measure as in the right panels of Figure 1 (Appendix Table C.4 lists the industries of each sector). Households spend a fixed share of their budget on each sector, and within a sector they buy varieties according to a translog expenditure function with its own parameter, so markups differ across firms and sectors. Sectors also differ in the dispersion of productivity draws and in iceberg costs that differ by direction of trade. The labor force of each sector is fixed. Only the home country levies tariffs ($t^*_{j,t} = 0$ in every experiment), so the model has no retaliation, and production uses labor alone, so it has no input-cost channel.

Relative to the one-sector model of Section 4.2, the household's problem changes in two places and the firms' problem in its sector-specific parameters.[21] First, consumption is a Cobb–Douglas aggregate of three sector baskets, so the household spends a fixed share α_j of its budget on each sector, and the aggregate price index is the numeraire:

$$C_t = \prod_{j=1}^{3} \left(\frac{C_{j,t}}{\alpha_j} \right)^{\alpha_j}, \qquad P_{j,t} C_{j,t} = \alpha_j C_t, \qquad \sum_j \alpha_j \ln P_{j,t} = 0. \tag{20}$$

Within sector j the basket is a translog aggregate with its own parameter σ_j, so the share and markup of a variety, (11) and (12), hold sector by sector, with the measure and mean log price of the sector's varieties in place of the market's and spending $\alpha_j C_t$ in place of C_t:

$$s = \frac{1}{N_{j,t}} + \sigma_j \big(\bar{L}_{j,t} - \ln p \big), \qquad \mu = 1 + \frac{s}{\sigma_j}. \tag{21}$$

Second, the labor force of each sector is fixed, so the sector's wage $w_{j,t}$ adjusts to clear its labor market.[22] The bond Euler equations (7) are unchanged, and the share Euler equation

[21] Appendix C.3 lists every change relative to the one-sector model, and the model notes in the replication package state the three-sector model in full, with every first-order condition, the aggregation and adding-up conditions and the definition of equilibrium.

[22] In the computations, labor can move across sectors at a quadratic cost set so high that no more than

(6) and free entry (10) hold sector by sector. In the firms' problem, a firm of sector j draws its productivity from a Pareto distribution with its own shape k_j, produces $Z_t Z_{j,t} z$ units per unit of labor, where $Z_{j,t}$ is sector technology, and pays iceberg costs specific to the sector and to the direction of trade, τ_j on home goods shipped abroad and τ_j^* on foreign goods shipped home, which lets a sector's import share differ from its export share. Its log marginal costs at home and abroad become

$$c_{j,t}^D(z) = \ln \frac{w_{j,t}}{Z_t Z_{j,t} z}, \qquad c_{j,t}^X(z) = \ln \frac{(1 + t_{j,t}^*)\, \tau_j\, w_{j,t}}{Q_t Z_t Z_{j,t} z}, \tag{22}$$

and its pricing, export and entry decisions are those of Section 4.2, sector by sector, with a fixed export cost $f_{X,j}$. As in the one-sector model with home a quarter of the world, the foreign labor force is 3 times home's. The foreign market of each sector has its own translog parameter σ_j^*, which replaces σ_j in the share and markup of every variety sold there.

The three-sector model carries the one-sector economy with home a quarter of the world over to the industries of the empirical exercise. I calibrate it only in that asymmetric version (Table 7). The parameters the sectors share keep their values in the one-sector model of that size, apart from the bond fee, which I rescale to the size of the economy. I adapt the sectoral parameters so that each sector reproduces, on average, a feature of its industries in 2016, the year on which Flaaen and Pierce (2026) measure trade shares.

I build every sector target from industry data, weighting each industry by the quantity of which the statistic is a share. The sector's expenditure share α_j is its share of the absorption of the 86 industries. The import cost τ_j^* sets its import share of absorption, the absorption-weighted average of its industries' import shares. The export cost τ_j sets its export share of shipments, weighted by shipments; in the middle sector, balanced trade implies the export share. The fixed export cost $f_{X,j}$ sets the share of firms that export, weighted by the number of employer firms, in the middle and least exposed sectors. The most exposed sector keeps the fixed export cost of the one-sector model.[23] Because the model imposes balanced trade while U.S. manufacturing imports exceed exports, I scale every import share down by a common factor and every ratio of exports to domestic sales up by the same factor, set at 1.28 so that exports equal imports.

I set markups by the one-sector rule and the dispersion of firm sales by the calibration of Ghironi and Melitz (2005), sector by sector. The translog parameter σ_j sets the sector's

0.02 percent of the labor force changes sector in any experiment I report (Appendix C.3.2).

[23] In the most exposed sector, 24 percent of firms export, yet exports are 0.364 of shipments and imports 0.400 of absorption after the scaling to balanced trade. At the dispersion of sales of the published calibration, so few exporters can carry that much trade only if a trade cost falls to one. The share of exporting firms in the most exposed sector is therefore an outcome of the model, 33 percent.

average markup to the CES markup that would reproduce its labor share of value added relative to the middle sector, which keeps the published $\theta = 3.8$ there. This adapts the rule of Bilbiie et al. (2012) to each sector. The Pareto shape k_j sets the standard deviation of log domestic sales to the value implied by the calibration of Ghironi and Melitz (2005). The foreign parameter σ_j^* sets the average markup of the foreign market equal to home's in each sector, which gives σ_j^*/σ_j between 0.34 and 0.38.

Each sector's tariff is the import-weighted tariff change of its industries, the added duty on their 2017 imports over those imports, at the dates of the 2018 actions. I build it from the same lists and rates as the protection measure (Section 3.1). In a model with two countries, every import of a sector comes from abroad and pays the sector's tariff, so the sector's protection measure is its import share times this tariff. The model therefore reproduces both the manufacturing-wide increase of 3.13 points and the 7.62-point gap between the most and least exposed sectors.

I measure sector TFP as BLS measures it, applying (C.8) to each sector, and size the technology shock below so that this measured gap, the concept the data measure, matches the estimates; I report the gap in physical productivity beside it. The model's statistic is the gap between the most and least exposed sectors, in annual averages. Its empirical counterparts are the 2018 and 2019 points in the bottom right panel of Figure 1, the contrast of equation (3) at each horizon times the gap between the average exposures of the top and bottom quartiles. The target is -1.56 percent in 2018 and -5.05 percent in 2019. The 2019 target exceeds the 2.3 percent of the interquartile comparison because the model's outer sectors stand for every industry in the outer quartiles, whose average exposures lie further apart than those of the industries at the quartile edges.

5.1.2 The tariffs and the estimated decline

I run two experiments, each announced in the first quarter of 2018 and known thereafter. The first introduces the tariffs alone. The second adds a decline in the technology of the most exposed sector, persistent and phased in over 2018 and 2019, announced in 2018Q1 as news of the declines of the two years and sized so that the model's gap in measured TFP matches the estimates in both years. The technology shock is not a model outcome. It stands for whatever lowered the productivity of the protected industries, and by comparing the two experiments I size what a decline of the estimated size does to the outcomes of the tariffs. Welfare is the home household's conditional welfare from 2018 on, expressed as the permanent change in consumption that would yield the same welfare.

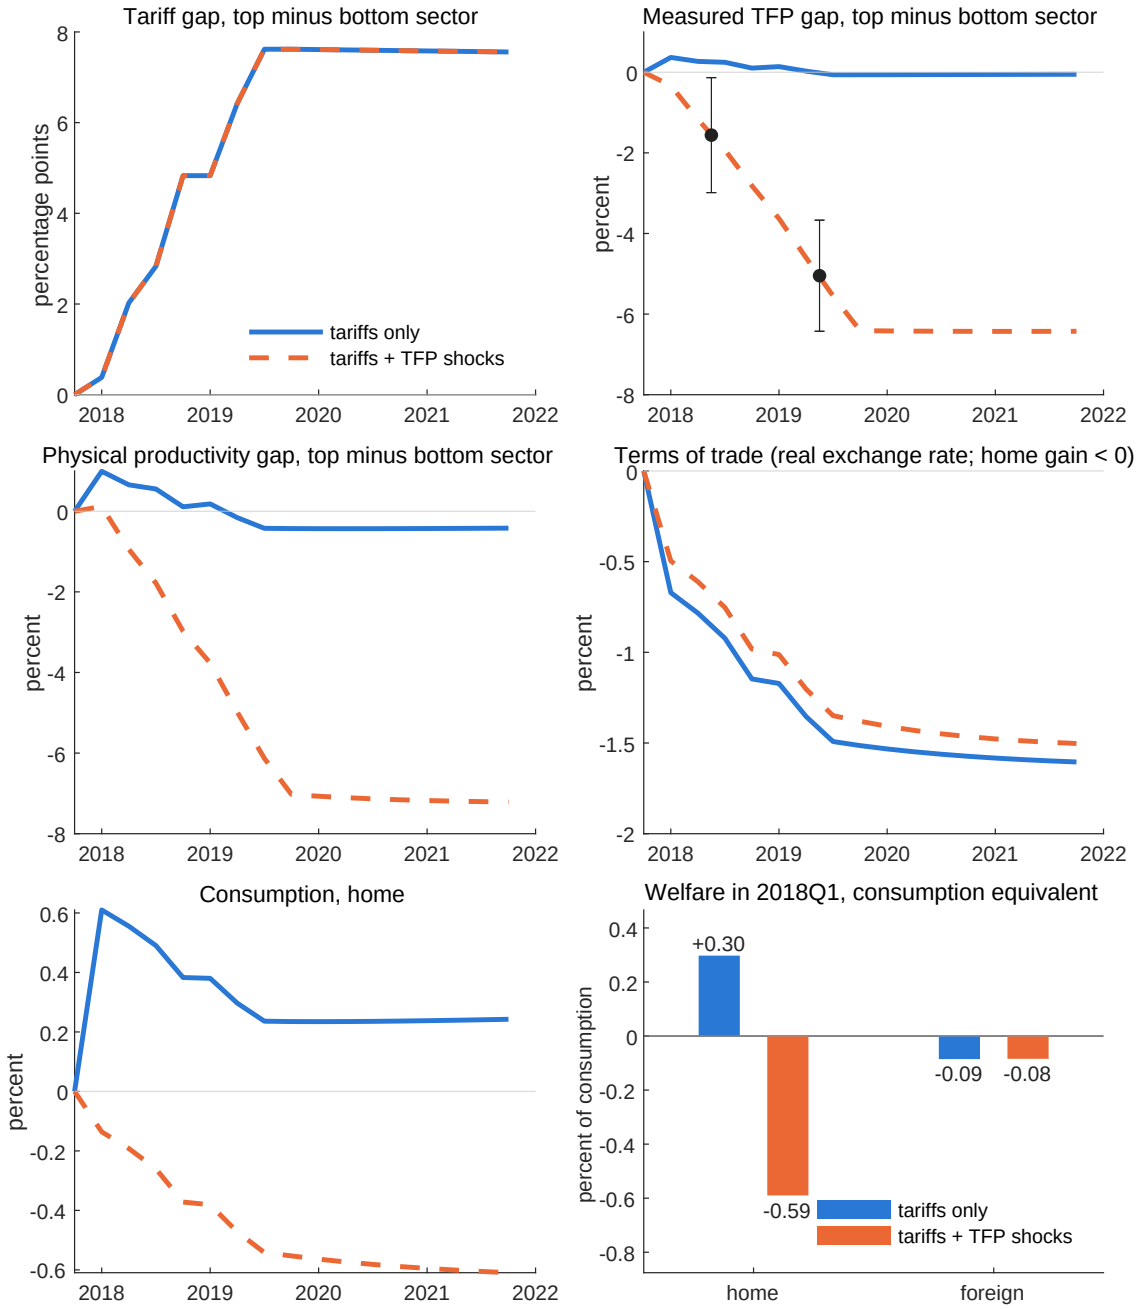
The tariffs alone leave the gap in measured TFP far short of the estimate (Figure 3). With the tariff gap matched, the model's measured gap between the most and least exposed

Table 7: Calibration of the three-sector model

Parameter	Controls	Target	Data and weighting within the sector	Value by sector
$\beta, \gamma, \delta, f_E, z_{\min}, L, L^*, \eta$	as in the one-sector model	none	the one-sector values	$L^* = 3$
α_j	size of the sector	share of absorption	sum of the industries' absorption over all 86	0.149 / 0.467 / 0.385
τ_j^*	cost of home imports	import share of absorption, scaled to balanced trade	industry import shares weighted by absorption	1.198 / 1.671 / 8.949
τ_j	cost of home exports	export share of shipments, scaled; middle sector by balanced trade	industry export shares weighted by shipments	1.351 / 1.671 / 5.551
$f_{X,j}$	fixed export cost	share of firms that export (middle and least exposed); one-sector value (most exposed)	shares weighted by employer firms	0.0065 / 0.0041 / 0.0007
σ_j	markup	CES-equivalent markup from the labor share of value added, $\theta_2 = 3.8$	labor shares weighted by shipments	6.346 / 1.165 / 0.253
σ_j^*	competition abroad	foreign markup equal to home's	none	2.188 / 0.431 / 0.095
k_j	dispersion of firm productivity	standard deviation of log domestic sales, 0.82	the value implied by the published CES calibration, $(\theta - 1)/k$	4.768 / 3.297 / 1.171

Notes: The sectors are the quartiles of the protection measure over the 86 industries of the empirical sample, as in the right panels of Figure 1. Values are listed most exposed / middle / least exposed. I solve for the sector parameters jointly with the steady state. Appendix Table C.2 reports the targets beside the data before the scaling to balanced trade.

Figure 3: Tariffs alone and tariffs with the estimated TFP decline, three-sector model



Notes: Solid lines: tariffs only; dashed lines: tariffs plus the technology shock to the most exposed sector. Top row: the tariff gap between the most and least exposed sectors, in percentage points, and the gap in measured TFP, in percent, with the estimates from the design of Section 3 for 2018 and 2019 and their 95 percent intervals, which the second experiment matches. Middle row: the gap in physical productivity (Section 4.1), on the same scale as the measured gap, and the terms of trade, which here move with the real exchange rate (a fall is a home gain), in percent deviation from the steady state. Bottom row: home consumption, and conditional welfare in 2018Q1, home and foreign, as a percent of steady-state consumption. Translog demand; quarterly paths from the exact deterministic solution (perfect foresight), labor fixed in each sector.

sectors is 0.25 percent in 2018 and 0.01 percent in 2019, against −1.56 and −5.05 percent in the data. The reallocation within the protected sector is real, but it is small, and the overhead its exporters release works the other way. Matching the estimate takes technology news to the most exposed sector of −0.75 percent a quarter in 2018 and −0.86 percent in 2019, which changes its technology by −6.41 percent by the end of 2019. With the measured gap matched, the physical gap is −1.40 percent in 2018 and −5.48 percent in 2019. By 2019 the measured statistic thus understates the physical decline, because a product-by-product deflator does not register all of the shift of labor toward low-markup firms.

The estimated decline also reverses the welfare gain from the tariffs. On their own the tariffs are a terms-of-trade gain for the home country. The real exchange rate appreciates, moving by −1.38 percent by 2019, the terms of labor, the cost of foreign relative to home effective labor, improve, moving by −1.31 percent, home consumption rises by 0.51 percent in 2018, and home welfare rises by 0.30 percent of consumption. The technology decline leaves the real exchange rate where the tariffs put it, but it cuts the improvement in the terms of labor to −0.36 percent and turns the consumption gain into a loss, −0.49 percent by 2019, so that home welfare changes by −0.59 percent instead of rising. Intuitively, the tariffs redirect home spending toward the very sector whose productivity falls, so the household pays more for the goods it has been induced to buy at home. Foreign welfare changes by −0.09 and −0.08 percent in the two experiments.

With labor fixed, the protected sector's relative wage absorbs the tariff, and with the technology decline its price index rises relative to the least exposed sector and its relative wage falls (Appendix Figure C.1). The anticipated tariffs also draw a rush of entrants into the protected sector, and with the sector's labor fixed the labor they use comes out of production, which is why relative real output changes by −4.52 percent in 2018 with the technology decline, against −2.54 percent in the data.

The reversal does not require the full estimate. Scaling the technology decline down, the home welfare gain from the tariffs vanishes once the decline reproduces 35 percent of the estimated gap in measured TFP. The end of the 95 percent interval for the 2019 gap closest to zero is about 73 percent of the point estimate, above that threshold, so even a decline at the edge of what the data allow leaves welfare close to zero.

5.1.3 Robustness and the external position

The reversal depends on the persistence of the decline (Table 8). If the technology shock fades at 5 percent a quarter, home welfare with the decline is 0.11 percent, about two-fifths of the gain from the tariffs alone. I cannot tell the two cases apart with the design of Section 3, since the estimates that are not confounded with the pandemic end in 2019; the

Table 8: The welfare comparison under alternative assumptions, three-sector model

	Home welfare		TFP gap 2019		Real output gap	
	tariffs	plus TFP	measured	physical	2018	2019
Baseline: persistent shock	0.30	-0.59	0.01	-5.48	-4.52	-2.43
Transitory technology shock ($\rho_S = 0.95$)	0.30	0.11	0.01	-5.38	-4.33	-2.16
Gross external position, end of 2017	0.22	-0.71	0.00	-5.38	-4.29	-2.48
Net external position only, in home currency	0.28	-0.66	0.00	-5.42	-4.26	-2.47
Data			-5.05		-2.54	-6.07

Notes: Home welfare is conditional welfare in 2018Q1, percent of steady-state consumption, with the tariffs alone and with the technology shock that matches the estimated gaps in measured TFP. TFP gaps are most minus least exposed sector, percent: measured, with the tariffs alone, and physical, with the technology shock; real output gaps are with the technology shock. Each row changes one assumption from the baseline and recalibrates the technology shock. The transitory shock has an autoregressive coefficient of 0.95 a quarter instead of 0.999. The external positions are the U.S. positions at the end of 2017 as shares of annual GDP (U.S. Bureau of Economic Analysis, 2026a,b): gross liabilities of 181 percent, in home currency, and gross assets of 141 percent, in foreign currency; the last row holds only the net position, in home currency. Data: the 2018 and 2019 points of the right panels of Figure 1.

welfare consequence of the decline therefore turns on whether the productivity the protected industries lost is lost for good.

The rows with the external position add the force Itskhoki and Mukhin (2026) emphasize. With gross dollar liabilities and foreign-currency assets of the size the United States held at the end of 2017, the tariffs alone raise home welfare by 0.22 percent instead of 0.30, and with the estimated decline welfare changes by −0.71 percent. Two forces account for the smaller gain. A net debtor, which must run a trade surplus to service its debt, gains less from the terms of trade, as the row with the net position alone shows. And the appreciation the tariffs cause transfers wealth to the rest of the world, the valuation channel of Itskhoki and Mukhin (2026). They describe it as a rise in the value of the country's foreign liabilities; measured in home currency, as here, it is a fall in the value of the foreign-currency assets. Both make the gain the technology decline has to overturn smaller, and both reinforce the reversal.

The external position matters much less for the tariffs actually imposed than for a permanent tariff imposed by surprise, the experiment of Section 4.5 (Table 9 reports the path of the 2018–19 tariffs beside a permanent tariff imposed by surprise). The 2018–19 tariffs were announced and phased in, so the real exchange rate moves little when they are announced, −0.72 percent, and appreciates gradually afterwards. Because bonds are short-term and returns adjust to the expected path of the exchange rate, only the surprise at announcement revalues the external position, and the valuation effect is −0.059 percent of consumption, close to the impact revaluation of the foreign-currency assets, −0.052 percent. A permanent tariff imposed by surprise moves the exchange rate by −1.73 percent at once, the valuation

effect rises to −0.143 percent, and the gross position cuts the gain from the tariff from 0.333 to 0.153 percent. The difference from the magnitudes Itskhoki and Mukhin (2026) report therefore lies in the design of the tariff, not in the model. Both tariffs are nearly permanent, and the valuation channel bites when a tariff comes as a surprise, whereas the 2018–19 tariffs were announced and phased in.

Table 9: The valuation channel: the 2018–19 tariffs and a permanent unanticipated tariff

	2018–19 tariffs, announced, phased in	Uniform permanent tariff, unanticipated
Real exchange rate when the tariff arrives, no position, percent	-0.67	-1.85
Real exchange rate when the tariff arrives, gross positions, percent	-0.72	-1.73
Real exchange rate after two years, gross positions, percent	-1.51	-1.71
Home welfare, percent of consumption		
No external position	0.298	0.333
Net position only, in home currency	0.276	0.296
Gross positions of the end of 2017	0.217	0.153
Valuation effect, percent of consumption		
In the model: gross minus net	-0.059	-0.143
Impact revaluation of foreign-currency assets	-0.052	-0.123

Notes: Three-sector translog model, tariffs alone, labor fixed. The first column is the announced, phased-in path of the 2018–19 tariffs; the second an unanticipated, nearly permanent (persistence 0.999) tariff of 3.13 percent on every sector, the manufacturing-wide increase of the data. Gross positions: liabilities of 181 percent of annual GDP in home currency and assets of 141 percent in foreign currency (U.S. Bureau of Economic Analysis, 2026a,b); the net position alone is held in home currency. The impact revaluation is $(1-\beta)\bar{B}_*(Q_1-\bar{Q})/\bar{C}$, the consumption equivalent of the change in the home-currency value of the foreign-currency assets when the tariff arrives.

The terms-of-trade gain in the model is, if anything, generous by the standard of the evidence, so the reversal is not an artifact of a small gain. Fajgelbaum et al. (2020) offer what they call a ballpark estimate of the short-run aggregate loss in U.S. real income from the 2018 tariffs, about 0.04 percent of GDP, with a confidence interval that includes zero, and essentially no change without retaliation; Fajgelbaum and Khandelwal (2022) put the losses implied by static trade models, which by construction omit dynamic losses, at 0.01 to 0.1 percent of GDP; and Adão et al. (2025) cannot reject the welfare predictions of the model of Fajgelbaum et al. (2020). Import prices rose nearly one for one with the tariffs (Amiti et al., 2019; Fajgelbaum et al., 2020), which leaves little room for a decline in foreign export prices product by product. As Fajgelbaum et al. (2020, p. 3) caution, however,

their estimation controls for country-time and product-time effects, so it cannot capture terms-of-trade effects that operate through relative wages across countries or sectors. The model's gain runs through that channel. The tariff improves the terms of labor, the cost of foreign relative to home effective labor (Appendix Figure C.1), and the model therefore complements the product-level estimates rather than contradicting them. Even so, its gain of 0.30 percent of consumption, without retaliation and without an input-cost channel, sits above the estimates, and a model with those channels would set the technology decline against a smaller gain. With the estimated decline, the model's welfare change of -0.59 percent is a loss several times the size of those that static trade models attribute to the tariffs, and the difference is the technology margin that those models hold fixed.

5.1.4 The optimal unilateral tariff

The optimal unilateral tariff of Section 4.5 carries over to the three-sector model, for a uniform tariff on every sector imposed by surprise and nearly permanent (Table 10). Without a productivity response it is 44 percent, and with the external position of the end of 2017 it falls to about 15 percent. The external position cuts the optimal tariff by more than half in the three-sector model, as it does in the one-sector model of Section 4.5, where the optimum falls from 39 to 18 percent.

The optimum is nonetheless highly sensitive to a productivity response, because the efficiency cost of a tariff rises slowly. The gain from one more point of tariff therefore declines slowly as the tariff rises, and the optimum lies where that gain is finally exhausted. If the tariff lowers technology by a given percentage per point, the marginal tariff stops paying once that loss, valued at the welfare cost of a nearly permanent technology decline, exceeds the gain a from the first point. Beyond that threshold the optimal unilateral tariff is zero. Short of it the optimum falls quickly. With a decline of half the threshold per point, it is 14 percent rather than 44. Intuitively, a loss that is the same for every point of tariff lowers the marginal gain at every tariff by the same amount, and because the marginal gain declines slowly, even a small downward shift moves a long way the tariff at which it reaches zero.

The threshold is small because the gain at the margin is small relative to the cost of lower technology. Starting from free trade, the marginal gain from a point of tariff is 0.114 percent of consumption, while a nearly permanent decline of 1 percent in every sector's technology lowers welfare by 1.16 percent (Table 10). The marginal tariff therefore stops paying at a decline of 0.10 percent per point, about 12 percent of the response that reproduces the estimate (0.84 percent per point), and 6 percent with the external position. Read as a decline in every sector, a persistent response of the size estimated is far beyond the threshold, and

the optimal unilateral tariff is not positive. Read this way, however, the relative estimate becomes a decline in the level, which I cannot identify from differences across industries (Section 5.2). A decline in the level of about an eighth of the estimate would be enough.

Table 10: The optimal unilateral tariff and the productivity threshold

	No external position	Gross positions, 2017
Welfare gain per point of tariff, a	0.1143	0.0564
Optimal tariff without a productivity response, percent	44.2	14.7
Welfare at the optimal tariff, percent of consumption	1.848	0.375
Welfare cost of a 1 percent technology decline, percent of consumption	1.157	1.168
Threshold technology decline per point of tariff, percent	0.099	0.048
as a share of the estimated response	0.12	0.06

Notes: Three-sector translog model, labor fixed. A uniform tariff on every sector, unanticipated and nearly permanent (persistence 0.999); welfare as the exact consumption equivalent from the exact deterministic solution (perfect foresight). The gain per point a is the slope of welfare at a zero tariff, from tariffs of 1 and 2 percent, and the optimal tariff maximizes welfare. The welfare cost of a 1 percent technology decline applies it to every sector with the same persistence. The threshold is a divided by that cost; the estimated response is the decline in the most exposed sector's technology per point of tariff gap that reproduces the measured TFP gap in Section 5.1. Gross positions as in Table 8.

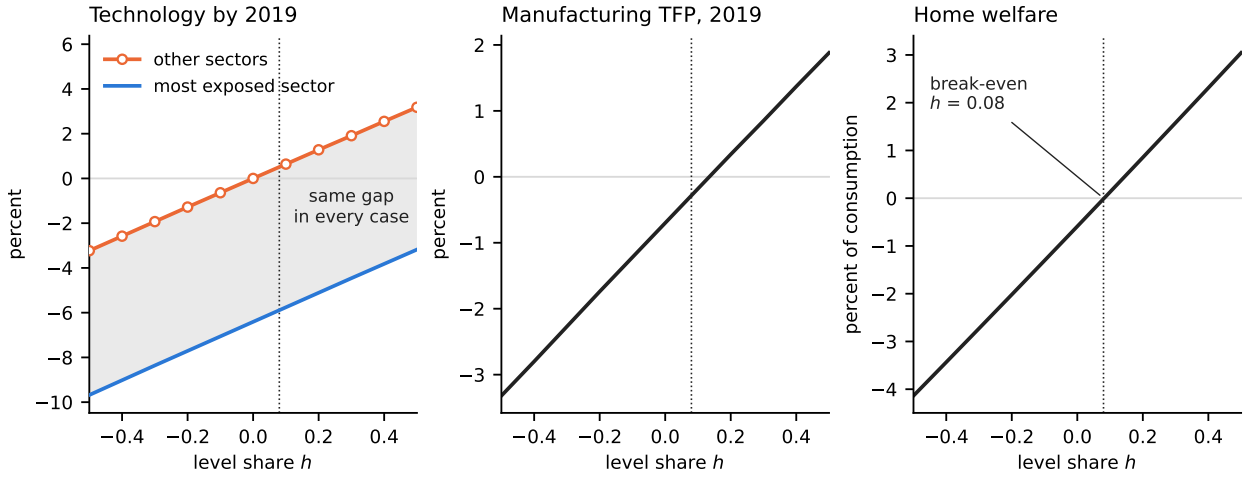
The decline overturns the case for a tariff only if it lasts. A less persistent decline must start larger to reproduce the measured gap of 2018–19, and it fades faster. When I re-size the decline in this way for each persistence, the optimal unilateral tariff stays at zero as long as the decline persists at more than 0.918 per quarter, a half-life of about 8 quarters. For the 2018–19 tariffs themselves the bar is higher. Their welfare gain becomes a loss only once the persistence exceeds 0.974, a half-life of about 27 quarters. How long the decline lasts I cannot settle with evidence through 2019.

The conclusion that the optimal tariff is not positive has limits. It fails if the decline is transitory, if the decline in the level is less than the threshold share of the estimate, if productivity responds only beyond some tariff size, or if tariffs may differ by sector; and it assumes that a uniform tariff lowers every sector's technology in proportion, whereas a response proportional to each sector's import share would lower the aggregate loss and raise the threshold.

5.2 Sensitivity to how the productivity decline is distributed across sectors

The three-sector results rest on an assumption about the evidence that I cannot test with the difference-in-differences design of Section 3. I identify the gap in productivity between more and less exposed industries, not its level. The same gap could reflect a decline in the most exposed industries, a rise in the others, or any mix of the two. The baseline of Section 5.1

Figure 4: The welfare effect of the 2018–19 tariffs depends on how much of the estimated gap is a decline



Notes: Three-sector translog model, labor nearly fixed, the 2018–19 tariff paths. The level share h moves part of the technology gap between the most exposed sector and the other two: the most exposed sector's technology changes by $1-h$ times the news of 2018 and 2019 and the other sectors' by $-h$ times it, with the news re-sized for every h to reproduce the estimated measured TFP gap of 2018 and 2019; $h=0$ is the baseline of Section 5.1. Left: technology by 2019, with the gap between the two groups of sectors shaded. Middle: measured TFP for manufacturing as a whole, sectors weighted by their expenditure shares. Right: home welfare, percent of consumption, from the exact deterministic solution. The dotted line marks the level share at which home welfare changes sign.

attributes the whole gap to a decline in the most exposed sector and leaves the technology of the other two sectors unchanged. In this subsection I relax that assumption and ask how the welfare result depends on it.

I keep the gap and vary its level. A level share h moves part of the gap between the most exposed sector and the other two. The technology of the most exposed sector changes by $1-h$ times the technology news of 2018 and 2019, and the technology of the other two sectors by $-h$ times it. At $h=0$ the model is the baseline. A positive h attributes part of the gap to a rise in the less exposed sectors; a negative h lets the decline reach them as well, with the most exposed sector falling the most. For every h the news is re-sized to reproduce the estimated gap in measured TFP in 2018 and 2019 exactly, the tariffs are those of 2018–19, and welfare comes from the exact deterministic solution. Figure 4 plots, for h between $-\frac{1}{2}$ and $\frac{1}{2}$, technology by 2019 in the two groups of sectors, the implied change in measured TFP for manufacturing as a whole, with sectors weighted by their expenditure shares, and home welfare.

The sign of the welfare result can change in one direction only. A decline that also reaches the less exposed sectors deepens the loss, while a modest rise in those sectors restores the gain. At $h=-\frac{1}{2}$ home welfare changes by -4.14 percent of consumption, against -0.59

percent in the baseline. In the other direction the gain from the tariffs returns quickly, and welfare turns positive once about 8 percent of the gap is a rise in the less exposed sectors. A small share suffices because the most exposed sector accounts for only a small part of manufacturing spending, an expenditure share of 0.149 (Table 7). A modest rise spread across the rest of manufacturing therefore weighs as much in the aggregate as a large fall concentrated in the most exposed sector, and the terms-of-trade gain from the tariffs covers what remains. The level share matters because it moves productivity in manufacturing as a whole, against which the terms-of-trade gain is set. The reversal holds as long as the change in measured TFP for manufacturing as a whole in 2019 is below its break-even value, -0.27 percent; in the baseline it is -0.70 percent. A positive level share amounts to a rise in the technology of the sectors that the tariffs protected least, which together account for most of manufacturing, and the welfare gain at positive h depends on that rise, which the tariffs do not cause. Such a rise needs a mechanism that raises productivity where the tariffs did not reach, and the tariffs themselves supply none. A negative level share, by contrast, follows naturally from the input-cost and retaliation channels of Flaaen and Pierce (2026), which reach industries across manufacturing whatever their import protection, although that common component is equally beyond the reach of a comparison across industries.

Because the empirical approach rests on differences across industries, the level of the decline is not identified. Everything common to manufacturing drops out of those differences, and recovering the level would require an assumption about the common counterfactual that lies outside this paper. The welfare verdict therefore rests on two quantities that remain open, the level of the decline and its persistence. In this subsection and in Section 5.1 I size how much of each the reversal needs.

6 Conclusion

The U.S. tariffs of 2018–19, on more than $350 billion of imports, coincided with a break in the relative productivity of the industries they shielded. Because tariffs are not handed out at random, I estimate the association with the difference-in-differences design of Flaaen and Pierce (2026), which nets out the pre-war paths of more and less exposed industries and lets trade-exposed industries follow their own course. The more exposed industries break from their relative productivity trend in 2019. Before the war they were gaining on the others, and the gain stopped. An industry at the 75th percentile of import protection is associated with TFP about 2.3 percent lower by 2019 than one at the 25th, relative to the pre-war path (Tables 1 and 2).

A decline of this size bears directly on the gains from protectionism, but I cannot assess

those gains from the estimate alone, because it is a relative effect and leaves out the terms-of-trade gain that is the classic case for a unilateral tariff. I therefore carry the estimate into the model of Ghironi and Melitz (2005), extended with a tariff, which improves the terms of trade. I measure productivity in the model as statistical agencies measure it in the data, and with the model I translate the measured decline into the physical decline on which welfare depends. In a three-sector extension, calibrated to the more and less exposed industries, the tariffs raise home welfare through the terms of trade, and the additional technology decline needed to reproduce the estimated gap turns that gain into a loss; a decline a little over a third as large suffices. The gross external position of the United States reinforces the reversal, because the appreciation the tariffs bring lowers the home-currency value of its foreign-currency assets, which shrinks the gain from the tariffs alone.

The terms-of-trade effects are strong enough, however, that the reversal requires a persistent decline. If the decline fades with a half-life shorter than about 27 quarters, the welfare loss disappears. Similarly, only a lasting decline brings the optimal unilateral tariff to zero. Firm heterogeneity alone does not reverse the optimal tariff of Itskhoki and Mukhin (2026). With heterogeneous firms and home a quarter of the world, it is 36 percent, close to their 35 percent. Because the efficiency cost of a tariff rises slowly, the optimum is highly sensitive to a productivity response. If the tariff lowered technology in every sector, the technology shock needed to bring the optimal unilateral tariff to zero would be a fraction of the shock in the baseline case. Both conclusions rest as well on how much of the estimated gap is a decline rather than a rise in the less exposed industries, which the empirical approach cannot tell apart. The case for the tariffs on terms-of-trade grounds therefore turns on a margin that the standard welfare assessments hold fixed, and on how long it stays depressed.

I stress test a result of open-economy models with heterogeneous firms and, in doing so, assess the model itself. Once a tariff is added to the model of Ghironi and Melitz (2005), it lowers physical productivity along the intensive margin of Section 4.1. Production is reallocated among firms whose productivity is given, away from the foreign sales of exporters, toward firms that sell only at home and, under translog demand, toward firms with lower markups. Set against the evidence, that margin is too small. With the tariffs alone, the three-sector model opens a gap in measured productivity between the most and least exposed sectors of 0.01 percent by 2019, against -5.05 percent in the data, and matching the estimate requires a technology decline from outside the model. The reason lies in what the model holds fixed. Every firm draws its physical productivity once, from a distribution that does not respond to shocks. A tariff therefore moves productivity only by reallocating production among firms, by changing which firms export, and, through the fixed cost of exporting, by changing how much labor goes into export overhead, which moves measured but not physical

productivity; it does not change the technology with which any firm produces.

A model in which the tariff also changes the technology available to firms could produce the rest of the decline. Import competition raised revenue productivity in Chilean plants after liberalization (Pavcnik, 2002) and TFP within European firms exposed to Chinese imports (Bloom et al., 2016). Tariffs could work the other way, by dulling the discipline of foreign rivals and with it the incentive to improve technology. Two classes of extensions would let the tariff change the technology available to firms. In the first, the distribution of productivity across firms changes in response to shocks, through entry, exit and the diffusion of technology among firms (Perla et al., 2021). In the second, firms invest to raise their own productivity, through innovation or the adoption of new technologies (Atkeson and Burstein, 2010; Akcigit et al., 2026), or, as in Comin and Gertler (2006), through the creation and adoption of new technologies across the economy. Models of either kind could also determine the two conditions on which the welfare reversal rests and which I cannot settle with the evidence: how long a decline persists, and whether it is a fall in the protected industries or a rise in the others. Such models are worth building because the reallocation the tariffs cause in the calibrated workhorse model falls short of the evidence, and the margin they add is the one needed to close the gap.

References

Rodrigo Adão, Arnaud Costinot, and Dave Donaldson. Putting quantitative models to the test: An application to the US–China trade war. *Quarterly Journal of Economics*, 140(2): 1471–1524, 2025.

Rodrigo Adão, John Sturm Becko, Arnaud Costinot, and Dave Donaldson. Why is trade not free? A revealed preference approach. *Econometrica*, 2026. Forthcoming.

Stéphane Adjemian, Michel Juillard, Frédéric Karamé, Willi Mutschler, Johannes Pfeifer, Marco Ratto, Normann Rion, and Sébastien Villemot. Dynare: Reference manual, version 6. Dynare Working Papers 80, CEPREMAP, 2024.

Ufuk Akcigit, Sina T. Ates, and Giammario Impullitti. Innovation and trade policy in a globalized world. *Journal of Political Economy Macroeconomics*, 4(2):211–255, 2026.

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.

Andrew Atkeson and Ariel T. Burstein. Innovation, firm dynamics, and international trade. *Journal of Political Economy*, 118(3):433–484, 2010.

David H. Autor, David Dorn, Gordon H. 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.

Alessandro Barattieri, Matteo Cacciatore, and Fabio Ghironi. Protectionism and the business cycle. *Journal of International Economics*, 129:103417, 2021.

Susanto Basu, John G. Fernald, and Miles S. Kimball. Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448, 2006.

Randy A. Becker, Wayne B. Gray, and Jordan Marvakov. [dataset] NBER-CES manufacturing industry database. National Bureau of Economic Research, 2021. 1958–2018, version 2021a, NAICS 2012 version (`nberces5818v1_n2012`), https://www.nber.org/research/data/nber-ces-manufacturing-industry-database.

Andrew B. Bernard, J. Bradford Jensen, and Peter K. Schott. Trade costs, firms and productivity. *Journal of Monetary Economics*, 53(5):917–937, 2006.

Andrew B. Bernard, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. Firms in international trade. *Journal of Economic Perspectives*, 21(3):105–130, 2007.

Florin O. Bilbiie, Fabio Ghironi, and Marc J. Melitz. Endogenous entry, product variety, and business cycles. *Journal of Political Economy*, 120(2):304–345, 2012.

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.

Christian Broda, Nuno Limão, and David E. Weinstein. Optimal tariffs and market power: The evidence. *American Economic Review*, 98(5):2032–2065, 2008.

A. Colin Cameron, Jonah B. Gelbach, and Douglas L. Miller. Bootstrap-based improvements for inference with clustered errors. *Review of Economics and Statistics*, 90(3):414–427, 2008.

Diego Comin and Mark Gertler. Medium-term business cycles. *American Economic Review*, 96(3):523–551, 2006.

Arnaud Costinot and Andrés Rodríguez-Clare. Trade theory with numbers: Quantifying the consequences of globalization. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, editors, *Handbook of International Economics*, volume 4, pages 197–261. Elsevier, 2014.

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. *Review of Income and Wealth*, 69(4):999–1032, 2023. Earlier version: BLS Working Paper 530 (2022).

Jan De Loecker. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451, 2011.

Svetlana Demidova and Andrés Rodríguez-Clare. Trade policy under firm-level heterogeneity in a small economy. *Journal of International Economics*, 78(1):100–112, 2009.

Pablo D. Fajgelbaum and Amit K. Khandelwal. The economic impacts of the US–China trade war. *Annual Review of Economics*, 14:205–228, 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. New product varieties and the measurement of international prices. *American Economic Review*, 84(1):157–177, 1994.

Robert C. Feenstra, Philip Luck, Maurice Obstfeld, and Katheryn N. Russ. In search of the Armington elasticity. *Review of Economics and Statistics*, 100(1):135–150, 2018.

Amy Finkelstein. The aggregate effects of health insurance: Evidence from the introduction of Medicare. *Quarterly Journal of Economics*, 122(1):1–37, 2007.

Aaron Flaaen and Justin R. Pierce. [dataset] Replication data for: Disentangling the effects of the 2018–2019 tariffs on a globally connected U.S. manufacturing sector. Harvard Dataverse, https://doi.org/10.7910/DVN/NUDHEA, 2024. Accessed October 2026.

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*, 108 (5):1289–1303, 2026.

Lucia Foster, John Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1): 394–425, 2008.

Fabio Ghironi and Marc J. Melitz. International trade and macroeconomic dynamics with heterogeneous firms. *Quarterly Journal of Economics*, 120(3):865–915, 2005.

Pinelopi K. Goldberg and Nina Pavcnik. The effects of trade policy. In Kyle Bagwell and Robert W. Staiger, editors, *Handbook of Commercial Policy*, volume 1A, chapter 3, pages 161–206. Elsevier, 2016.

Paul Goldsmith-Pinkham, Isaac Sorkin, and Henry Swift. Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8):2586–2624, 2020.

Robert E. Hall. The relation between price and marginal cost in U.S. industry. *Journal of Political Economy*, 96(5):921–947, 1988.

Thomas J. Holmes and James A. Schmitz, Jr. Competition and productivity: A review of evidence. *Annual Review of Economics*, 2:619–642, 2010.

Oleg Itskhoki and Dmitry Mukhin. The optimal macro tariff. Working paper, Harvard University and London School of Economics, March 2026.

Harry G. Johnson. Optimum welfare and maximum revenue tariffs. *Review of Economic Studies*, 19(1):28–35, 1951.

Harry G. Johnson. Optimum tariffs and retaliation. *Review of Economic Studies*, 21(2): 142–153, 1953.

Jozef Konings and Hylke Vandenbussche. Heterogeneous responses of firms to trade protection. *Journal of International Economics*, 76(2):371–383, 2008.

Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.

Office of the United States Trade Representative. Notice of action and request for public comment concerning proposed determination of action pursuant to section 301: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 83 Fed. Reg. 28710 (June 20, 2018), Annexes A and B, 2018a. Product list for Section 301 List 1, 818 eight-digit subheadings. https://www.govinfo.gov/content/pkg/FR-2018-06-20/pdf/2018-13248.pdf.

Office of the United States Trade Representative. Notice of action pursuant to section 301: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 83 Fed. Reg. 40823 (August 16, 2018), Annex A, 2018b. Product list for Section 301 List 2, 279 eight-digit subheadings. https://www.govinfo.gov/content/pkg/FR-2018-08-16/pdf/2018-17709.pdf.

Office of the United States Trade Representative. Notice of modification of section 301 action: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 83 Fed. Reg. 47974 (September 21, 2018), Annex A, notes 20(f) and 20(g), 2018c. Product list for Section 301 List 3, 5,745 full and partial tariff lines. https://www.govinfo.gov/content/pkg/FR-2018-09-21/pdf/2018-20610.pdf.

Nina Pavcnik. Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *Review of Economic Studies*, 69(1):245–276, 2002.

Jesse Perla, Christopher Tonetti, and Michael E. Waugh. Equilibrium technology diffusion, trade, and growth. *American Economic Review*, 111(1):73–128, 2021.

Justin R. Pierce. Plant-level responses to antidumping duties: Evidence from U.S. manufacturers. *Journal of International Economics*, 85(2):222–233, 2011.

José Antonio Rodríguez-López. Prices and exchange rates: A theory of disconnect. *Review of Economic Studies*, 78(3):1135–1177, 2011.

The White House. Presidential proclamations 9693 (solar cells) and 9694 (large residential washers): To facilitate positive adjustment to competition from imports. Federal Register 83 FR 3541 and 83 FR 3553 (January 25, 2018), 2018a. https://www.govinfo.gov/content/pkg/FR-2018-01-25/pdf/2018-01592.pdf, https://www.govinfo.gov/content/pkg/FR-2018-01-25/pdf/2018-01604.pdf.

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%), 2018b.

U.S. Bureau of Economic Analysis. [dataset] International investment position, table 1.1. U.S. net international investment position at the end of the period. Interactive Data Application, release of September 24, 2026, 2026a. https://apps.bea.gov/iTable/, accessed October 25, 2026.

U.S. Bureau of Economic Analysis. [dataset] National income and product accounts, table 1.1.5. gross domestic product. Interactive Data Application, last revised September 30, 2026, 2026b. https://apps.bea.gov/iTable/, accessed October 25, 2026.

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

U.S. Bureau of Labor Statistics. Industry productivity measures. Handbook of Methods, 2026a. https://www.bls.gov/opub/hom/inp/calculation.htm.

U.S. Bureau of Labor Statistics. Producer price indexes: Calculation. Handbook of Methods, 2026b. https://www.bls.gov/opub/hom/ppi/calculation.htm.

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

World Bank. [dataset] GDP (current US$), indicator NY.GDP.MKTP.CD, United States and World. World Development Indicators, 2026. https://api.worldbank.org/v2/country/USA;WLD/indicator/NY.GDP.MKTP.CD, accessed October 27, 2026.

Appendix for Online Distribution

to

"Tariffs, Productivity, and the Welfare Effects of Protectionism"

Luca Guerrieri

October 2026

Abstract

This appendix has three parts. Appendix A documents the construction of the industry–year panel: the statutory (schedule) import tariff rate for the 86 four-digit NAICS manufacturing industries of the BLS productivity accounts, on a single 2012 NAICS basis, including the Section 232 and Section 301 actions of 2018–19 and running through 2023; the productivity series; and the Flaaen and Pierce (2026) exposure measures. Appendix B reports additional empirical results: the Flaaen and Pierce design without the pre-war detrending, with alternative outcomes and the pre-trend test, the tests of the 2019 break against chance, the break by pre-war industry characteristics, the statutory rate as the measure of exposure, with the effect in a common unit, and controls for the pandemic years, and compares the estimate with other evidence on tariffs and productivity. Appendix C states the models of the paper: the one-sector model of Ghironi and Melitz (2005) under CES and translog demand, its three-sector extension, and the versions with identical firms.

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 tariff measure, I describe the industry classification and harmonization; the statutory tariff measure; the productivity measures; the merged panel; and the Flaaen and Pierce (2026) exposure measures. Reconstruction steps, the log of downloaded inputs, and known limitations are recorded in the replication package. The unifying choice throughout is that every series is placed on a single, longitudinally consistent 2012 NAICS basis, so the tariff and productivity data are directly comparable over time and with one another.

A.1 Overview

The analysis dataset is a panel indexed by four-digit NAICS manufacturing industry j and year t, covering the 86 NAICS manufacturing industries of the BLS detailed-industry productivity accounts. For each cell I observe productivity alongside the **statutory (schedule) rate** τ_{jt}^{O}, the Column 1 (most-favored-nation, NTR) rate in ad valorem terms plus the 2018–19 Section 232 and 301 add-ons, which from 2018 makes it the rate facing Chinese goods (Section A.3). The rate is placed on a single, longitudinally consistent **2012 NAICS** basis so that it is comparable over time and with the productivity data (Table A.2 lists the NAICS vintage the raw data carry in each year).

Table A.1 lists the sources and coverage of every series. Figure A.1 shows the productivity side of the panel. On average, total factor productivity rises until the mid-2000s and then moves sideways, while labor productivity rises faster, reflecting capital deepening, peaks around 2013 and drifts down after that.

A.2 Industry classification and harmonization

I place the tariff and productivity data on one industry classification, the 2012 vintage of four-digit NAICS, as the subsections below describe.

A.2.1 Industries covered

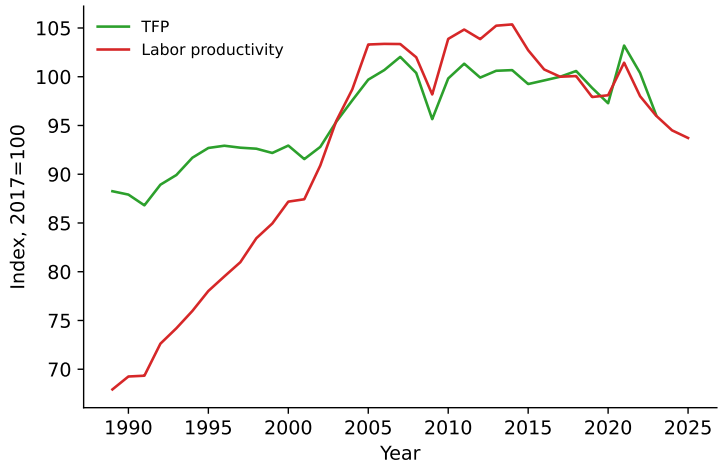
The panel covers the 86 four-digit NAICS manufacturing industries (sectors 31–33) of the BLS detailed-industry productivity accounts. NAICS 3328 (coating, engraving, and heat treating, a contract-service industry with essentially no traded product) is retained for completeness but carries a missing tariff in every year by construction; the design of Section 3, which needs a 2017–19 change for every industry, sets its change to zero (Appendix A.6).

Table A.1: Data sources and coverage

Series	Concept	Source	Coverage
Output tariff, statutory	Schedule rate (MFN + Sec. 232/301)	USITC DataWeb (U.S. International Trade Commission, 2025); Fajgelbaum et al. (2019); Office of the United States Trade Representative (2019a); The White House (2018b)	2002–2023
Input–output weights	2012 benchmark detail use table	BEA (U.S. Bureau of Economic Analysis, n.d.)	2012 (fixed)
Total factor productivity	Detailed-industry TFP index	BLS (U.S. Bureau of Labor Statistics, 2026c)	1987–2023
Labor productivity, hours, employment and nominal output	Real sectoral output per hour; hours worked; employment; current-dollar sectoral output	BLS (U.S. Bureau of Labor Statistics, 2026c)	1987–2025
Total factor productivity (cross-check)	Five-factor TFP index	NBER–CES (Becker et al., 2021)	1958–2016
Imports by product and origin	HS10 import values by country, for Protection and the import share	Schott (2025)	2016–2017
Exports by product	HS10 export values, for the export share	Schott (2018)	2016
Shipments	Value of shipments, for the import and export shares	NBER–CES (Becker et al., 2021)	2016
Retaliation, InputCost, imported share of input costs, capital per worker	Flaaen and Pierce's measures, rebuilt at four digits	Flaaen and Pierce (2024); BEA (U.S. Bureau of Economic Analysis, n.d.); NBER–CES (Becker et al., 2021)	2012, 2016–2017

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 fixed BEA input–output weights enter the semiconductor-share control of Appendix B.5. The statutory rate is the most-favored-nation rate plus the Section 232 and 301 add-ons. Construction is detailed in Appendices A.3–A.6 and validation in Appendices A.4 and A.6.

Figure A.1: Productivity indexes, 1989–2025



Notes: Cross-industry mean of the total factor productivity index (through 2023) and the labor productivity index (through 2025), 2017=100.

A.2.2 The NAICS vintage problem

NAICS is revised every five years (1997, 2002, 2007, 2012, 2017 and 2022); I carry the 2017 vintage for every year from 2017 on. Trade and tariff data carry the vintage contemporaneous with each year, listed in Table A.2, so codes are not directly comparable across revision boundaries (the largest manufacturing reclassifications occur at 2002 → 2007 and 2007 → 2012, for instance in computers and electronics, NAICS 334). The BLS productivity accounts, by contrast, are published on a single NAICS basis, that of 2012, for the whole period, and the Census establishment microdata behind U.S. productivity statistics have been recoded to a consistent basis at the establishment level (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 that establishment-level recoding, 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,

$$R_j = \sum_{\ell \in j} r_\ell \, w_\ell, \qquad M_j = \sum_{\ell \in j} w_\ell, \tag{A.1}$$

and forming the cell value R/M. For the statutory rate, r_ℓ is a tariff-line rate and w_ℓ a line weight (Section A.3.2). The pair (R, M) is aggregated like a level, through the NAICS vintage chain and the six- to four-digit rollup; the BLS productivity series is published on the 2012 basis and needs no concordance (Appendix A.4).

Table A.2: NAICS vintage of the statutory tariff panel, by year

Years	NAICS vintage
2002–2006	2002
2007–2011	2007
2012–2016	2012
2017–2023	2017

A.3 Statutory (schedule) tariff rates

The statutory rate is the legal schedule rate, set line by line, so it must be averaged across the tariff lines of each industry. After defining it and the within-industry aggregation choice, I describe the source, the USITC DataWeb schedules of 2002–2023, and the two 2018–19 add-ons, Section 232 and Section 301.

A.3.1 Definition

The statutory output tariff is the Column 1 (NTR, most-favored-nation) rate of the Harmonized Tariff Schedule plus the 2018–19 Section 232 and Section 301 additional duties. The Column 1 rate is ad valorem or, for a specific, compound or formula duty, its ad valorem equivalent. 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.

For industry j in year t,

$$\tau^{O}_{jt} \;=\; \frac{\sum_{\ell \in L_{jt}} \kappa_{\ell j} \left(m_{\ell t} + a^{232}_{\ell t} + a^{301}_{\ell t} \right)}{\sum_{\ell \in L_{jt}} \kappa_{\ell j}}, \tag{A.2}$$

where τ^O_{jt} is the statutory output tariff of industry j in year t, in percent; ℓ indexes tariff lines and L_{jt} is the set of eight-digit lines of year t's schedule that map to industry j and carry a Column 1 rate that year, one set for all three components; $m_{\ell t}$ is the Column 1 rate of line ℓ in year t, in ad valorem terms; $a^{232}_{\ell t}$ and $a^{301}_{\ell t}$ are the Section 232 and Section 301 add-ons on that line in force that year, zero before 2018 and weighted within the year an action takes effect by the days it applies; and $\kappa_{\ell j}$ is the share of line ℓ assigned to industry j, which equals one when the line maps to a single industry. The numerator therefore sums the total rate a line carries, weighted by the share of the line that belongs to the industry, and the denominator sums those shares, so τ^O_{jt} is a weighted mean of line rates.

Defining the rate as the Column 1 rate plus the two add-ons leaves every other trade barrier outside it. The measure omits the 2002 and 2018 Section 201 safeguards, the derivative steel and aluminum products added to the Section 232 action in February 2020, and antidumping and countervailing duties, and it does not net out exclusions granted to individual importers. The derivative products arrive after the 2017–19 changes the paper's designs use, so they move the level of the series and not the variation exploited; exclusions are granted to individual importers, whereas τ^O is the rate a covered line carries. The Column 2 rate is no part of τ^O. An origin-weighted alternative exists in the coverage shares of Flaaen and Pierce (2026), which Section 3 of the main text puts to work.

A.3.2 Within-industry aggregation

To average the line rates within an industry I must choose a weight. My default is the *simple (unweighted) mean* of line rates, the construction of Pierce and Schott (2016, p. 1638). They compute NTR gaps from the ad valorem equivalent rates of Feenstra et al. (2002) at the eight-digit line, and "the gap for industry i is the average NTR gap across the eight-digit HS tariff lines belonging to that industry." I apply the same construction to the Column 1 rate and add the Section 232 and 301 duties to each line's rate before averaging, a step that their period does not require. Unlike a mean weighted by imports, the simple mean gives a line no less weight because its tariff discourages imports of it. Trade still enters the measure in two places: the equivalents of specific, compound and formula duties are computed from unit import values, and a line without a rate in a year leaves that year's average. Aggregation reuses (A.1) with r_ℓ the line rate and $w_\ell = \kappa_{\ell j}$, the share of the line assigned to the industry, which is one for a line that maps to a single industry. The concordance is defined at the HS10 level, so each HS8 line is assigned to the NAICS codes of its HS10 children and split across them in proportion to the number of children in each, with shares summing to one per line and year. Lines that map only to non-industry codes are dropped and X-filled codes are kept at their four-digit prefix; I record the weight that fails to match.

How many lines each industry contains varies widely. In 2019 the median industry draws on 92 tariff lines, with a range from 5 for forging and stamping to 1,313 for basic chemicals. About 2.5 percent of eight-digit lines straddle more than one four-digit industry. The line HS 8522.90.80, parts for sound- and video-recording apparatus, is one of the exceptions. Its two ten-digit children fall one into audio and video equipment (NAICS 3343) and one into electronic components (NAICS 3344), so it enters each industry's average with a split weight of one half. Averaging at ten digits instead would weight a line by how finely the statistical codes beneath it happen to be split, which is a reporting choice rather than a policy one.

A.3.3 The USITC DataWeb schedules, 2002–2023

Schedule rates for 2002–2023 come from the USITC DataWeb (U.S. International Trade Commission, 2025) annual tariff files. For 2019 I use DataWeb's 2019 database, since the text file in the 2019 folder is a snapshot of November 2018. Each file states a line's Column 1 duty in words (`mfn_text_rate`), with a code for how the duty is computed and the rate in numeric fields, and I read the rate from the words. The numeric fields are truncated to two decimals in several years (in the files of 2016–2020 and 2022–2023 a rate of 7.3 percent is stored as 0.07), and they hold a placeholder, 9999.99, where a duty is not a fixed rate. The text of two lines is in error in some years; I correct it from the same line's text in neighboring years and from the numeric fields, and the README lists both corrections.

A specific or compound duty enters at its ad valorem equivalent, computed as Feenstra et al. (2002) compute theirs, "using observed unit import values." The equivalent is the duty on the line's imports for consumption over their customs value, $a + (sQ_1 + oQ_2)/V$, where a is the ad valorem part of the duty, s and o its specific rates per unit, Q_1 and Q_2 the quantities on which they are levied and V the customs value, each summed over the line's ten-digit commodities and every origin in the import files of Schott (2025). I depart from Feenstra et al. (2002) in one respect. They use each year's unit values, so the equivalent of an unchanged specific duty falls as prices rise. I hold the unit values at 2017, the last year before the tariffs, so that a change in τ^O reflects a change in the schedule. A line with no imports in 2017 takes the nearest year with imports (134 lines in the 2021 schedule), and a line with no imports in any year has no equivalent and leaves the average. With each year's unit values, the method reproduces the equivalents DataWeb publishes for 2002 to a median difference of 0.05 percentage points.

Some lines state their duty only as a formula in words: a rate on the lead content of an ore, a sugar duty that varies with the degree of polarization, or the rate applicable to each garment of a set. Most of them are watch and clock parts (HS 91). For these lines I use the duties collected on the line's imports at the general rate (Census rate provision 61) over their

dutiable value, from origins other than China, whose duties include the Section 301 add-ons from 2018, at the 2017 values as for the specific duties. On lines with an ad valorem rate this ratio equals the schedule rate, and for 2002, the one year in which DataWeb publishes equivalents for these lines, it matches DataWeb's.

Of the 11,111 ordinary lines in the 2021 schedule (that is, outside the Chapter 98–99 special provisions, which I drop in any case), 743 carry a specific duty, 316 a compound duty and 157 a formula, and 1,197 of these 1,216 lines have an equivalent. Because the industry rate is a simple mean over the lines in that year's schedule, a revision to the schedule changes the denominator. The 2022 Harmonized System revision raised the count of ordinary lines from 11,111 to 11,871 while the set of lines matched to the Section 301 action was unchanged at 9,276, so measured coverage and the mean add-on fall in 2022–23 with no change in policy.

A.3.4 Section 232 (steel and aluminum)

The Section 232 add-on is built directly from The White House (2018b): "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 percent) and "aluminum articles" in HS 7601 and 7604–7609 plus castings/forgings 7616.99.51.60/.70 (10 percent), both effective March 23, 2018, 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 not modeled) and omits the February 2020 "derivative" products. I also compute a coverage share, the fraction of an industry's tariff lines subject to the steel and aluminum action, which is distinct from the add-on *rate*. As for Section 301, the rate equals the coverage share times the average add-on on the covered lines.

A.3.5 Section 301 (China)

Section 301 applies only to China-origin goods, so it is kept in a separate column rather than folded into a country-agnostic schedule rate. Product membership for Lists 1–3 comes from the replication data of Fajgelbaum et al. (2019), which identify each list by its statutory effective date for China. Because their panel ends in April 2019, I attach the documented rate *schedule* rather than rely on their later values. Lists 1 and 2 (25 percent, from July 6 and August 23, 2018) are held open, and List 3 (10 percent from September 24, 2018, raised to 25 percent on May 10, 2019) is carried forward to the end of the statutory panel. List 4A, the first of the two tranches into which the U.S. Trade Representative split the fourth round of Section 301 tariffs, on about $300 billion of imports from China, took effect on September 1, 2019. It lies outside their panel, so its membership comes from Annex A of the

USTR notice that established it (Office of the United States Trade Representative, 2019a), 3,233 eight-digit subheadings.[24] Lines already carrying a List 1–3 add-on and lines subject to Section 232 are excluded from it, leaving 2,894 additions, so no line is counted twice. List 4B, the December 2019 tranche, was suspended before it took effect and is not included. Each covered line's add-on is date-annualized (a rate in force for part of a year contributes in proportion to the days it applies) and aggregated by (A.1) into `statutory_301`, with a companion coverage share. The carry-forward is verified against DataWeb's Chapter 99 provisions, which implement the action as separate Chapter 99 headings that add a stated percentage to the duty of the underlying subheading and apply only to articles the product of China. List membership is held at the 2018–19 HS vintage and is not remapped to later revisions, so from 2023 some renumbered lines no longer match. Of the 9,291 eight-digit lines the lists name (794, 280 and 5,323 for Lists 1–3 in Fajgelbaum et al. (2019), against the 818, 279 and 5,745 the notices name and the protection measure uses, Section A.6, plus the 2,894 List 4A additions), 9,278 appear in the 2019 schedule, 9,277 in 2021 and 2022, and 8,892 in 2023. The overlay matches one line fewer in every year, because the castings line 7616.99.51, which List 3 names, falls under Section 232. Coverage in 2023 is therefore understated, which matters for the level of the series rather than for the 2017–19 changes the paper uses.

A.4 Productivity measures

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

Primary: BLS detailed-industry TFP. The Bureau of Labor Statistics publishes TFP indexes (sectoral output over a Törnqvist index of combined capital, labor, energy and intermediate inputs) for detailed manufacturing industries (U.S. Bureau of Labor Statistics, 2026c). The `MachineReadable` tab is filtered to four-digit, all-workers rows; the TFP index (base 2017=100) is taken in logs. These 86 four-digit NAICS manufacturing industries define the cross-section of the panel, so the TFP panel merges onto the tariff panel with no concordance. Coverage is 1987–2023.

Cross-check: NBER–CES. The NBER–CES Manufacturing Industry Database (Becker et al., 2021), NAICS-2012 vintage (364 six-digit industries; dollar magnitudes 1958–2018, the `tfp5` index populated through 2016), gives an independent TFP index (`tfp5`, five-factor) and the dollar magnitudes (shipments, materials, energy, payroll, value added) from which

[24]The USTR notice that established List 4A set an additional duty of 10 percent from September 1, 2019. A later notice raised the rate to 15 percent before it took effect (Office of the United States Trade Representative, 2019b), and a third cut it to 7.5 percent from February 14, 2020 (Office of the United States Trade Representative, 2020). The statutory tariff applies 15 percent to List 4A lines from September 1, 2019 and 7.5 percent from February 14, 2020.

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), and the BLS and NBER–CES log-TFP series correlate at 0.77 over the 2,580 overlapping industry-years, validating the primary measure.

A.5 The merged panel

The harmonized tariff rates are attached to the 86 industries of the BLS productivity accounts on (`naics4`, `year`), keeping the industry panel complete over the covered years (an industry with no imports retains a missing tariff rather than being dropped). The panel covers 2002–2023 for the **86** manufacturing industries, carrying the statutory Column 1, Section 301 and Section 232 columns and their coverage shares. NAICS 3328 (a contract-service industry with essentially no traded products) carries a missing tariff in every year by construction; the design of Section 3 sets its 2017–19 change to zero.

A.6 Flaaen–Pierce exposure measures

I build the protection measure of equation (1) from public sources and turn to the replication package of Flaaen and Pierce (2024) only afterwards, to check the result; Retaliation and InputCost are taken from that package, as described at the end of this section. The protection measure rests on the Census import and export records distributed by Schott (2025) and Schott (2018), on NBER–CES shipments (Becker et al., 2021), and on the Federal Register texts that define each 2018 action: the three Section 301 notices (Office of the United States Trade Representative, 2018a,b,c), the Section 232 proclamations (The White House, 2018b) and the Section 201 safeguards (The White House, 2018a). Where the sources leave a construction choice open and Flaaen and Pierce's choice is reasonable, I follow it, so that the comparison isolates the differences that the legal texts decide.

The product lists come from the notices themselves. Parsing their annexes yields 818, 279 and 5,745 lines for Lists 1, 2 and 3, the counts each notice states, and Section 301 applies to imports from China alone. Sections 232 and 201 apply to all origins; like the code of Flaaen and Pierce (2024), I do not model the country exemptions, quotas and product exclusions granted after the proclamations, so the protection measure is mildly overstated for the steel- and aluminum-producing industries. Two cases need a decision. List 3 covers eleven eight-digit lines only in part, excepting named statistical lines, and one of the exceptions,

8517.62.0090, was created in the 2018 schedule out of the 2017 line 8517.62.0050, which carried $22.8 billion of imports from China in 2017. I treat as covered the share that the line's covered successor, 8517.62.0020, took of China's imports of the two successors in the fourth quarter of 2018, when both first appear, 26 percent. The solar safeguard covers cells and modules, which have statistical lines of their own, and cells attached to generators and batteries, which the trade data cannot tell apart from the ordinary generators and batteries of the same lines; I count the first and not the second.

The protection measure factors exactly into the covered share of an industry's imports and the import share of its home market. I measure the first on 2017 general imports at customs value and the second on 2016 general imports, domestic exports and shipments, as Flaaen and Pierce do. A dollar of imports covered by two actions counts once. NAICS codes come from the Census records, and coating, engraving and heat treating (NAICS 3328), which records no imports and has no tariff lines, enters with a zero protection measure and zero Retaliation, as in the data of Flaaen and Pierce (2024); I set its 2017–19 change in the statutory tariff to zero, which matches their rate-weighted protection measure of zero for it.

The comparison with Flaaen and Pierce's own measure, which I compute by running their code through its exposure step, is close. On the 83 industries the two constructions share at the stage where the code of Flaaen and Pierce (2024) builds the exposure measures, at which it pools only the three leather industries, and with NAICS 3328, which imports nothing, left out, the import share of absorption agrees in every industry, and the protection measure correlates at 0.997. The rebuild assigns $268.5 billion of the $2,022 billion of 2017 manufacturing imports to the 2018 actions, against $262.5 billion in theirs. I attribute the whole difference in covered imports, cell by cell, to five decisions, with nothing left unexplained (Table A.3).

Each of the five differences has a reason in the legal texts. The first, and the largest, is the partly covered line 8517.62.0050. Flaaen and Pierce subtract it in full as a List 3 exception, which removes the covered part along with the excepted part, so the rebuild assigns $5.9 billion more to communications equipment (NAICS 3342), whose protection measure rises from 0.030 to 0.070. The second is that their List 3 file omits 339 of the 5,734 full lines of note 20(f), mostly food, minerals and chemicals ($1.3 billion). The third is that their steel list omits 24 of the six-digit subheadings that note 16(b) of Proclamation 9705 covers, mainly flat-rolled products ($1.4 billion, almost all in iron and steel mills, NAICS 3311). The fourth runs the other way, because they match the solar safeguard on four eight-digit lines in full, which brings in $1.4 billion of generators, batteries and other photosensitive devices. Most of these are not solar products, and the proclamation does not tax them. The fifth also runs the other way, because they add up the coverage of overlapping actions, so a Chinese

product on a Section 301 list and under a safeguard counts twice ($1.0 billion). In short, the first three differences follow the legal texts where a list file departs from them, and the last two avoid counting imports the actions did not cover, or counting the same imports twice.

With the same texts I also resolve the cases in which the product flags of Fajgelbaum et al. (2019), the natural alternative source, differ from Flaaen and Pierce's lists, and in each of them the texts side with Flaaen and Pierce. Five lines flagged on the List 2 date appear on no final list (three were on the proposed list of June 2018 and were dropped), seven seat lines that note 20(g) of List 3 excepts are flagged as covered, the aluminum castings and forgings of 7616.99.51 and the rail products of 7302.90 are not flagged, and the washer safeguard reaches only four statistical lines of finished machines.

I rebuild every measure of Flaaen and Pierce (2026) at the four-digit level by running the steps of their code without the step that combines industries into groups. The measures are their retaliation and input-cost channels, the share of input costs coming from imported goods, and capital per worker. Run with their groups, my code reproduces the values their code saves, with a largest difference of 5×10^{-7}. Two limitations remain, both in the source data. The detailed input-output table records apparel (NAICS 315) and leather (NAICS 316) only at three digits, so their four-digit industries share one input-cost share. The export file of Flaaen and Pierce (2024) records the three leather industries together, so they share one measure of retaliation. The import data of Schott (2025) do separate the leather industries, and the protection measure differs sharply among them. Other leather goods (NAICS 3169) had 48 percent of their home market newly covered in 2018, and footwear (NAICS 3162) none.

I depart from their code in three respects. For InputCost I take the products covered by each action from the Federal Register lists, as for the protection measure. I count the Section 201 safeguards on solar cells and washers among the duties on input commodities, as their protection measure does, although their code leaves them out of the input side. I keep each input commodity's import share of absorption between zero and one; in their data it falls outside that range for several commodities, chief among them scrap, which makes the imported share of input costs negative for the primary-metal industries. For Retaliation I keep their lists of the products that other countries targeted, which I have not checked against those countries' notices. One industry, audio and video equipment (NAICS 3343), reports domestic exports above its 2016 shipments, so its export share exceeds one; absorption is positive for every industry, so the protection measure is defined throughout.

Table A.3: The rebuilt protection measure against the Flaaen–Pierce original

	Decision	$ billion	Industries most affected
(a)	8517.62.0050 from China apportioned List 3, note 20(g): 8517.62 covered except 8517.62.0090, a 2018 line	+5.86	3342
(b)	List 3 lines missing from Flaaen and Pierce's list file List 3, note 20(f): 5,734 full lines	+1.29	3251, 3114, 3113, 3117
(c)	Steel lines missing from their list Proclamation 9705, note 16(b)	+1.36	3311
(d)	Solar limited to cells and modules Proclamation 9693: solar cells and solar panels, codes 8541.40.6020 and .6030. The proclamation also taxes solar cells sold attached to a battery or a generator, but the trade data record these together with ordinary batteries and generators, so they are left out	-1.41	3359, 3353, 3344
(e)	A product on two lists counted once A dollar of imports is covered once	-1.02	3344, 3359, 3315, 3353
	Total difference in covered imports	+6.08	
	Unexplained	0.00	

Notes: Difference in 2017 manufacturing imports covered by the 2018 tariff actions, rebuild minus Flaaen and Pierce (2024), in billions of dollars of general imports at customs value, attributed to the decision behind it by HS-10 × origin cell; positive entries are imports the rebuild counts as covered and theirs does not. Industries are the four-digit NAICS industries with the largest differences in each row. The legal source is given under each decision: Office of the United States Trade Representative (2018c) for rows (a) and (b), Proclamation 9705 (The White House, 2018b) for row (c), Proclamation 9693 (The White House, 2018a) for row (d).

B Additional empirical results

This appendix collects results that support the main text, in the order in which the paper draws on them. The first five sections concern the Flaaen and Pierce design of Section 3: the design without the pre-war detrending, with a revenue-based outcome, output per hour and the pre-trend test (Section B.1); the tests of the 2019 break against chance (Section B.2); the break by pre-war industry characteristics (Section B.3); the statutory rate as the measure of exposure, with the effect in a common unit (Section B.4); and controls for the pandemic years (Section B.5). The last section compares the estimate with the two studies of trade policy whose productivity measures are comparable with mine (Section B.6).

B.1 The Flaaen–Pierce design without detrending, alternative outcomes and the pre-trend test (Section 3)

The outcome in equation (2) is log TFP with each industry's own 2013–17 line removed, so the flat pre-war coefficients of Table 1 are in part what that correction imposes. Table B.1 therefore reports the same specification on the raw series, as Flaaen and Pierce (2026) do in their own appendix.

The pre-war path is sloped. The protection measure carries coefficients of −1.307 in 2013 and −0.819 in 2015, both significant, so the more exposed industries were gaining on the less exposed ones through 2017. After 2017 the gain stops, and no post-war coefficient in the column for the protection measure is distinguishable from zero. The detrended path in Table 1 is the same history, seen against the fitted trend. A series that flattens while its fitted trend keeps rising reads as a widening decline.

The choice between the detrended and the raw path matters little for inference, for two reasons. First, the statistic on which I base inference does not depend on the correction.

Equation (3) weights the post-war and pre-war changes so that an industry-specific linear trend differences out of it exactly, with a two-year pre-war window and with the three-year window of the main text. The entry for the protection measure in the last row of Table B.1, −1.2916 with a standard error of 0.1692, is the number the detrended specification gives, to four decimals. The year-by-year detrended coefficients carry a caveat that this statistic escapes, since the fitted trend rests on five annual points per industry and is extrapolated as much as five years ahead, and the clustered standard errors do not carry its sampling error. Flaaen and Pierce fit theirs on thirteen monthly observations, from January 2017 to January 2018, and extrapolate at most twenty months, so their first step is the better identified one. Neither first step enters equation (3).

Second, the pattern here is consistent with the argument of Flaaen and Pierce (2026) that the tariff lists were chosen on strategic grounds rather than on industries' short-run trends. The three channels approach 2017 along different paths: the protection measure from below, Retaliation from above, and InputCost from above in 2013–14 and from below in 2015–16. A characteristic that both shaped the tariff lists and moved productivity would move the three paths together.

That the BLS series measures quantities matters when the result is set beside tests of models in which technology is fixed. Adão et al. (2026) test a general-equilibrium model in which technology is fixed by checking that it predicts the 2017–19 changes in value added per worker across U.S. region-sector pairs (their Table 1). The closest public counterpart at the four-digit industry level is nominal sectoral output per hour, which moves with prices and composition as well as with physical productivity. Table B.2 reports the contrast of equation (3) for it, for real output per hour, and for TFP. The coefficient on the protection measure for nominal output per hour is −1.138 ($t = -1.51$), on real output per hour −1.692 ($t = -2.85$), and on TFP −1.292 ($t = -7.64$). The point estimates are of one size; what differs is precision, and the revenue measure is the least significant of the three because it adds price movements to a quantity response. The contrast is therefore not sharp, and I read it as consistent with the paper's distinction rather than as independent evidence for it. A model with fixed technology can match an earnings measure that prices move, while the quantity measure the paper uses responds in the same direction with less noise.

Flaaen and Pierce (2026) test for pre-trends by estimating the subtracted pre-period term on its own (their online appendix D.6, Table D12). Table B.3 reports the same test for equation (3). I split the statistic into the post-war change, $y_{2019} - y_{2017}$, and the pre-war term it subtracts, $\frac{2}{3}(y_{2017} - y_{2014})$, regress each on the regressors of equation (3), and add a two-year pre-war window. For TFP the pre-war term is significant for the protection measure and carries about two-thirds of the statistic; for the statutory rate it is smaller and

borderline. Output per hour shows the same pattern for the protection measure, with less precision.

B.2 Could the 2019 break arise by chance? (Section 3)

In this section I set out the tests of the 2019 break against chance that I summarize in Section 3 (Table B.4). Three possibilities must be excluded. The first is that the exposure measures are unrelated to the break, with the estimate a draw from the noise in industry productivity. The second is that industries within a three-digit sector are not independent draws, since the tariff lists were written by product group. The third is that the most exposed industries are ones whose productivity moves sharply in any year, so that a break of this size is ordinary for them.

In the first test I re-assign the three exposure measures together across industries at random, 9,999 times, and recompute equation (3) at the 2019 endpoint. A coefficient as large as the actual one appears in about one draw in a thousand for the protection measure ($p = 0.001$), and in about five in a thousand for the statutory rate ($p = 0.005$). Re-assigning the measures only within three-digit sectors, which holds each sector's mix of exposures fixed and so tests only the variation inside sectors, gives a similar answer ($p < 0.001$ and $p = 0.001$).

In the second test I cluster the standard errors by three-digit sector. With 21 clusters the asymptotic approximation is unreliable, so I read the wild cluster bootstrap by sector, with the null imposed, which rejects a zero effect for the protection measure ($p = 0.001$) and for the statutory rate ($p = 0.007$).

In the third test I hold each industry's 2017 exposure fixed and compute the same contrast over earlier windows, none of which involves a tariff war. Each window spans six years, from three years before its base year to two after, and windows that share years are not independent draws, so I use only windows that overlap neither each other nor the actual 2014–19 window. From 1987, the first year of the BLS series, that gives 6 sets of three or four windows. Across all 6 sets, no window matches the break for the protection measure; for the statutory rate it is the largest in 6, and it is matched or exceeded only in none. With four windows to a set, the smallest attainable p-value is 0.2, so with this comparison I can place 2019 at the top of the industries' own history but cannot by itself reject chance at conventional levels. Table B.5 reports every window.

In Table B.5, the industries with the largest statutory increases also took large relative swings in the windows with base years none, around the Great Recession. One explanation for the 2019 break would then be cyclical, since manufacturing output fell in 2019, and industries

that respond strongly to the cycle would have fallen with it. I test that explanation with two controls fixed before estimation. The first is each industry's sensitivity to the cycle, the slope of its TFP growth on the growth of manufacturing output outside the industry, estimated on the years before the war and, for each placebo date, without the years of the window being tested. The second is the change in each industry's TFP over 2007–09. Neither control moves the 2019 estimate, which stays within a hundredth of its baseline value (Panel C of Table B.4). Cyclical sensitivity is nearly unrelated to tariff exposure across industries, correlating at 0.15 with the protection measure, which is why it cannot account for the break.

The first control also sharpens the comparison for the statutory rate. Once each industry's cyclical sensitivity is held fixed, the 2019 break is the largest in 6 of the 6 sets of windows, and only the window based in none matches it.

Table B.5 lists, window by window, the placebo comparisons behind the third test of Table B.4. Each is the equation (3) contrast, with a three-year pre-window, around each pre-war base year from 1990 to 2011 that enters the non-overlapping placebo sets, on the same 2017 exposure measures; the second column gives the set. Each row is a separate cross-sectional regression on the specification of that table; the last row is the actual 2019 estimate.

Table B.6 reports the tests of Table B.4 run on output per hour, with the cyclical controls built from output per hour in place of TFP.

B.3 The break by pre-war industry characteristics

The exposure I measure is the tariffs' coverage of an industry's home market, so a competition channel would predict a sharper break where foreign rivals mattered more before the war: in industries whose imports came more from China, where imports took a larger share of the home market, where margins were thin, or where Chinese competition had been rising. I test that prediction by interacting the protection measure in the equation (3) regression with one characteristic at a time, each measured before the war, either as an indicator for an above-median value or standardized, and by re-estimating the regression on each half of the sample, with every estimate scaled to the interquartile move in the protection measure (Table B.7).

None of the 5 characteristics, with TFP growth measured at two horizons, sharpens the break. No interaction is significant, the smallest randomization p-value is 0.059, and the three measures of import competition carry the sign opposite to the one a competition channel predicts. The break is, if anything, weaker in the industries where imports, and Chinese imports in particular, mattered more. The price–cost margin has the predicted sign, but its

interaction is a small fraction of the baseline effect. The one characteristic along which the break vanishes in one half of the sample is the industry's own TFP growth over 2012–17. The break is concentrated in the industries that grew faster before the war (-1.95 percent, $t = -3.86$, against -0.80 percent, $t = -0.52$). That characteristic overlaps the pre-war drift that equation (3) subtracts, so the split describes the statistic rather than testing a channel, but it fits the reading of Section 3, that the break is a pre-war gain that stopped. I therefore cannot attribute the decline to weaker competition with these data, and I leave the channel open.

B.4 The statutory rate as the measure of exposure, and the estimate in a common unit

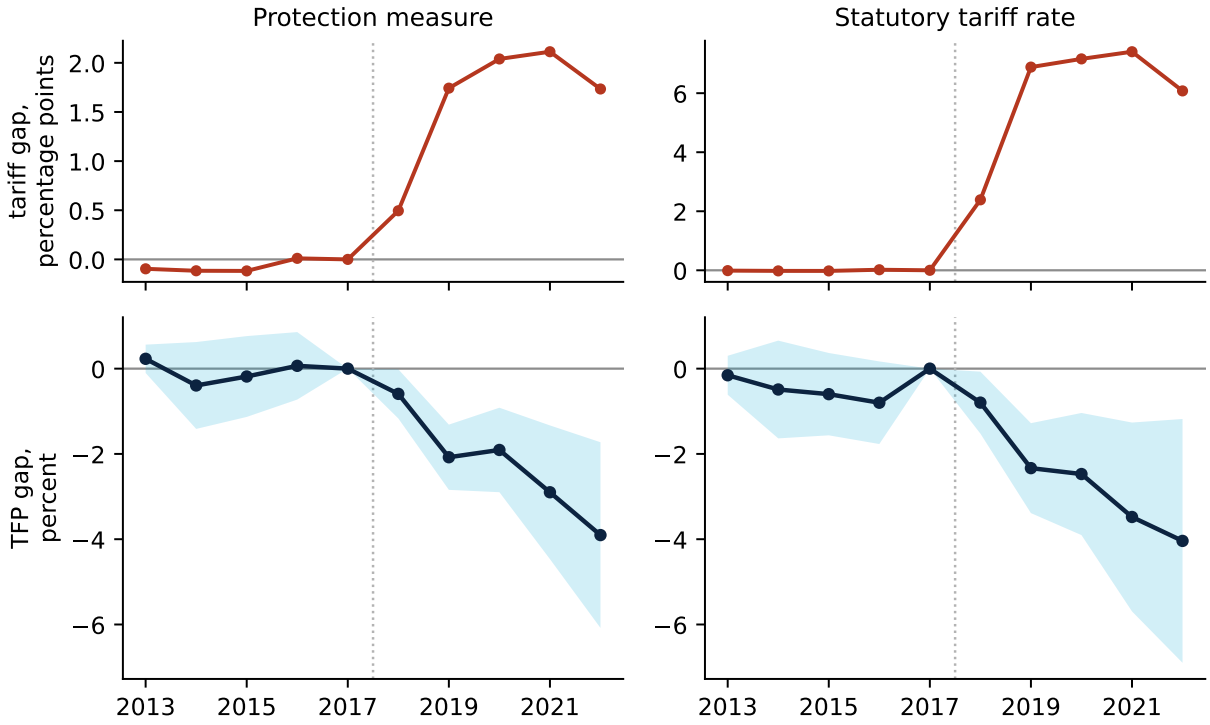
Figure B.1 plots the year-by-year path for both measures of exposure. I put the design's estimate with the statutory output tariff on one scale. The interquartile gap in the 2017–19 increase in the statutory output tariff is 6.9 percentage points; the 25th-percentile industry saw its tariff rise 11.9 points, the 75th 18.8. Applied to that gap, the coefficient gives the percent by which TFP in 2019 is lower at the 75th percentile of the statutory increase than at the 25th (Table B.8).

The design places that gap at 2.7 percent of TFP, with a 95 percent confidence interval from 1.0 to 4.3 percent. On the coverage-share scale of Figure 1, an interquartile move in the protection measure (1.8 points of the home market) gives 2.3 percent, with a 95 percent confidence interval from 1.7 to 3.0, so the choice of exposure measure changes the answer little.

B.5 Controls for the pandemic years (Section 3.4)

With controls for two features of the pandemic years, the path keeps its shape and grows slightly in size (Table B.9). The first control is each industry's sensitivity to the manufacturing cycle, the slope of its annual TFP growth on the growth of manufacturing real output outside the industry, estimated over 1988–2012 so that no year of the event-study panel enters it. The second is the share of semiconductors and other electronic components (NAICS 3344) in the industry's intermediate purchases in the 2012 input–output tables, aimed at the chip shortage of 2021–22. This direct share misses chips embedded in purchased parts.

Figure B.1: The Flaaen–Pierce design with two measures of exposure, relative to 2017



Notes: Bottom row: year-by-year coefficients on the exposure measure of each column from equation (2), with InputCost and Retaliation held fixed. The outcome is log BLS TFP detrended on each industry's own 2013–17 path, normalized to zero in 2017; shaded areas are 90 percent confidence bands, from standard errors clustered by three-digit sector. These are the year-by-year coefficients, not the equation (3) statistic of Table 1, which nets out the 2014–17 drift instead. Top row: the same specification with the statutory output tariff, not detrended, as the outcome, giving the average tariff gap, in percentage points, between industries at the 75th and 25th percentiles of exposure; it is drawn without a band because the tariffs are observed for every industry. In both rows coefficients are scaled by each measure's own interquartile move, so the axes read as gaps between an industry at the 75th and one at the 25th percentile of exposure: 1.8 points of home market in the left panel, 6.9 points of statutory output tariff in the right. The dotted line marks the first 2018 tariffs. The 2020–22 window overlaps the COVID-19 pandemic. Sample: 86 industries, 2013–2022.

Both controls are fixed before the pandemic. With them, equation (2) becomes

$$\tilde{y}_{jt} = b_t^P \,\text{Protection}_j + b_t^I \,\text{InputCost}_j + b_t^R \,\text{Retaliation}_j + b_t^M \,\text{ImpShare}_j + b_t^X \,\text{ExpShare}_j + b_t^C \,\text{Cycle}_j + b_t^S \,\text{Chips}_j + d_j + d_t + \varepsilon_{jt}, \tag{B.1}$$

where Cycle_j is the industry's cyclical sensitivity and Chips_j its semiconductor share. The other terms are those of equation (2). Like the other coefficients, b_t^C and b_t^S are year-specific slopes normalized to zero in 2017, so a control correlated with the protection measure before 2020 would show in the pre-war coefficients on the protection measure. None of them becomes significant. With the controls, Δ_{2019} is -1.417 ($t = -7.34$), against -1.292 without them,

and the 2022 coefficient is −2.576, against −2.153. The widening after 2019 is therefore not the cyclical swing of 2020 or the chip shortage, as far as these two controls measure them. Other features of the pandemic remain uncontrolled, which is why I rest the magnitude on 2019.

B.6 The estimate against other evidence on tariffs and productivity (Section 3.5)

Two studies of trade policy measure productivity in a form I can compare with my estimate, and I compare it with these two alone. The first measures physical productivity. Pierce (2011) sets U.S. plants in products granted antidumping duties against plants in products whose petitions failed, in the Census of Manufactures for 1987, 1992 and 1997. For the plants that report output in physical units, he measures physical TFP as in Foster et al. (2008). In that sample, the coefficient on the interaction of the post-duty period with the effective antidumping duty implies a change in physical TFP of −3.12 percent per percentage point of duty (his Table 7, column 3). My estimate is −0.39 percent per point of the statutory tariff (Table 2). In the same plants, his revenue productivity, which deflates revenue by industry price indexes, barely moves. The coefficient implies −0.05 percent per point and is not significant (his Table 6, column 7). Pierce reads the gap between the two measures as the work of prices, concluding that "antidumping protection is more likely to lead to higher prices than higher productivity" (p. 232).

The two estimates refer to different horizons. Mine is the 2017–19 change, net of the pre-war drift, and so measures the gap within two years of the tariffs imposed in 2018. Pierce's post-duty indicator equals one in every census year after the final determination of the investigation, and the determinations fall between 1988 and 1996. His estimate therefore averages over the first one to nine years after the duty, and he reports no estimate by years since the duty was imposed. Moreover, his duties fall on narrowly defined products and his outcome is the productivity of the plants that make them, while my exposure is industry-wide and my outcome is industry TFP built from BLS deflators. These differences in design and horizon aside, physical productivity has fallen under antidumping duties before, in Pierce's plants, by more per point than I estimate.

The second study measures industry TFP much as I do. Bernard et al. (2006) regress the average annual growth of five-factor industry TFP from the NBER–CES database over the five years after a census year on the annualized change in trade costs, the ad valorem duty plus freight and insurance, over the five years before it, for U.S. manufacturing industries. Their Table 2 is titled 1982–1997, while its notes say the regressions cover 1972–1996. Their

coefficients, −0.152 without and −0.190 with two-digit industry fixed effects, are significant at the 10 percent level. They state both variables in percentage points per year: the change in trade costs as the change in the ad valorem rate over five years, divided by five (p. 922), and TFP growth as an average annual percent change (p. 924). A cut of one point in trade costs spread over five years therefore raises industry TFP by 0.152 to 0.190 percent over the following five years. Their horizon reaches up to ten years from the start of the cut, their policy is a liberalization rather than an increase in tariffs, and their exposure includes freight. Per point, my estimate is more than twice as large, and it opens within two years.

Because both studies measure at longer horizons, I read them as placing my estimate per point within the range of earlier evidence, not as confirming its size. Other studies of trade policy and productivity measure the revenue productivity of plants, which moves with prices and markups as well as with physical productivity, so I do not set their magnitudes beside mine; Appendix C.4 compares them with the model instead.

Table B.1: The Flaaen–Pierce design estimated without removing the pre-war trend

	Log TFP, not detrended		
	Protection measure	InputCost	Retaliation
2013	-1.3074***	0.9694	-0.3276
	(0.2017)	(1.4906)	(3.3013)
2014	-1.2938***	0.1086	1.6729
	(0.4292)	(1.4853)	(4.6805)
2015	-0.8192**	-2.6467**	1.1875
	(0.3499)	(0.9590)	(3.7923)
2016	-0.3222	-3.6603***	3.7180
	(0.2717)	(0.8074)	(2.3996)
2017		0 (base year)	
2018	0.0321	1.2845*	-1.6605
	(0.1924)	(0.7341)	(1.2170)
2019	-0.4291	0.7061	-2.5597
	(0.3335)	(1.2894)	(3.3450)
2020	0.0237	-0.6747	2.4738
	(0.4131)	(1.2144)	(3.8869)
2021	-0.1649	-0.2320	-5.1007
	(0.5638)	(2.0072)	(5.1773)
2022	-0.3603	-1.6270	-5.6184
	(0.7943)	(1.9518)	(6.8173)
Equation (3), 2014–17 pre-trend, 2019	-1.2916***	0.7785	-1.4445
	(0.1692)	(1.6174)	(2.5661)
Observations		860	

Notes: Year-by-year coefficients from equation (2), estimated on log BLS TFP with no trend removed, 2017 omitted as the base year. All three channels are rate-weighted, and the specification carries the four controls of Flaaen and Pierce (2026) interacted with every year dummy, industry and year fixed effects, and standard errors clustered by three-digit sector. InputCost, one of the three channels of Flaaen and Pierce (2026), has significant pre-war coefficients even after detrending, so I do not interpret its path. The last row is equation (3) at the 2019 endpoint with the 2014–17 pre-war window, $b^P_{2019} + \frac{2}{3}b^P_{2014}$, which nets the post-war move against the pre-war drift; because it removes any linear trend exactly, it is invariant to the detrending and reproduces the detrended estimate of Table 1 exactly. The cross-sectional regressions of Tables B.2 to B.7 report the same point estimate. Estimation years: 2013–2022, 86 industries. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$.

Table B.2: Revenue and quantity productivity measures in the Flaaen–Pierce cross-section

	(1) Nominal output per hour	(2) Real output per hour	(3) TFP
Equation (3), $(y_{2019} - y_{2017}) - \frac{2}{3}(y_{2017} - y_{2014})$			
Protection measure	-1.1383	-1.6922***	-1.2916***
	(0.7523)	(0.5933)	(0.1690)
InputCost	3.3367	1.9881	0.7785
	(2.2776)	(2.3371)	(1.6162)
2017–19 change, $y_{2019} - y_{2017}$			
Protection measure	-0.0356	-0.5214	-0.4291
	(0.5335)	(0.4733)	(0.3332)
Industries	86	86	86

Notes: Cross-sectional regressions across 86 industries of the indicated contrast in the log outcome on the protection measure, InputCost, Retaliation and the four controls of Flaaen and Pierce (2026), with standard errors clustered by three-digit sector in parentheses. Nominal output per hour is BLS current-dollar sectoral output divided by hours, formed as ln(nominal output) − ln(real output) + ln(output per hour); real output per hour is the BLS labor-productivity index; TFP is the BLS index. The upper panel is equation (3) at the 2019 endpoint with the 2014–17 pre-war window, as in Table B.1; the lower panel is the 2017–19 change. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B.3: The pre-trend test of the Flaaen–Pierce design

	TFP		Output per hour	
	Protection measure	Statutory rate	Protection measure	Statutory rate
Pre-war term, $\frac{2}{3}(y_{2017} - y_{2014})$				
Coefficient	0.8626	0.0018	1.1708	0.0034
	(0.2859)	(0.0014)	(0.5617)	(0.0020)
Wild cluster bootstrap by sector, p	0.016	0.281	0.158	0.140
Randomization across industries, p	0.004	0.074	0.022	0.051
Post-war change, $y_{2019} - y_{2017}$				
Coefficient	-0.4291	-0.0021	-0.5214	-0.0021
	(0.3332)	(0.0010)	(0.4733)	(0.0012)
Wild cluster bootstrap by sector, p	0.260	0.050	0.245	0.124
Randomization across industries, p	0.161	0.065	0.242	0.197
Equation (3), post minus pre				
Coefficient	-1.2916	-0.0039	-1.6922	-0.0055
	(0.1690)	(0.0012)	(0.5933)	(0.0020)
Wild cluster bootstrap by sector, p	0.001	0.007	0.064	0.040
Randomization across industries, p	0.001	0.005	0.011	0.016
Two-year pre-war change, $y_{2017} - y_{2015}$				
Coefficient	0.8192	0.0022	1.2637	0.0043
	(0.3496)	(0.0014)	(0.8175)	(0.0025)
Wild cluster bootstrap by sector, p	0.031	0.121	0.229	0.136
Randomization across industries, p	0.012	0.053	0.045	0.058
Pre-war term as a share of equation (3)	0.67	0.46	0.69	0.62
Industries / three-digit sectors	86 / 21			

Notes: Cross-sectional regressions across 86 industries of the indicated term, in the undetrended log outcome, on the exposure measure of the column (the protection measure, or the 2017–19 change in the statutory output tariff, per percentage point, where Tables 1 and 2 report it per 100 points), InputCost, Retaliation and the four controls of Flaaen and Pierce (2026), all fixed at 2017; standard errors clustered by three-digit sector in parentheses. The post-war change minus the pre-war term equals the equation (3) statistic exactly. The wild cluster bootstrap is by three-digit sector, with the null imposed; randomization re-assigns the three exposure measures together across industries; 9,999 draws each, seed 2026. Output per hour is BLS real sectoral output per hour worked.

Table B.4: Could the 2019 break arise by chance?

	(1) Protection measure	(2) Statutory rate
A. Equation (3), 2014–17 pre-trend, 2019		
Coefficient	-1.2916	-0.0039
	(0.1690)	(0.0012)
B. Tests against chance		
Randomization across industries, p	0.001	0.005
Randomization within three-digit sectors, p	<0.001	0.001
Clustered by three-digit sector, t	-7.64	-3.32
Wild cluster bootstrap by sector, p	0.001	0.007
Placebo windows 2011, 2005, 1999, 1993: at least as large	0 of 4	0 of 4
Sets of non-overlapping placebo windows in which 2019 is the largest	6 of 6	6 of 6
Placebo windows with a break at least as large	0 of 22	0 of 22
C. Controlling for cyclical sensitivity		
With the cyclical beta	-1.2817	-0.0038
	(0.1861)	(0.0012)
sets in which 2019 is the largest	6 of 6	6 of 6
With the 2007–09 change in TFP	-1.2892	-0.0039
	(0.1773)	(0.0012)
D. Utilization: log hours per worker, equation (3)		
Coefficient	1.0182	0.0013
	(0.5054)	(0.0015)
Randomization across industries, p	0.007	0.335
Sets of non-overlapping placebo windows in which 2019 is the largest	6 of 6	0 of 6
Interquartile effect, percent (t), by version of the protection measure:		
rebuilt, as above	1.85 (2.01)	
Flaaen and Pierce's coverage share	0.43 (0.43)	
Flaaen and Pierce's rate-weighted measure	0.54 (0.45)	
rebuilt, 69 industries they do not combine	-0.13 (-0.11)	
Industries / three-digit sectors	86 / 21	86 / 21

Notes: The equation (3) contrast at the 2019 endpoint, $(y_{2019} - y_{2017}) - \frac{2}{3}(y_{2017} - y_{2014})$ with y log BLS TFP, regressed across industries on the exposure measure of the column (the protection measure, or the 2017–19 change in the statutory output tariff, per percentage point, where Tables 1 and 2 report it per 100 points), InputCost, Retaliation and the four controls of Flaaen and Pierce (2026), all fixed at 2017; standard errors clustered by three-digit sector in parentheses. Panel B: randomization re-assigns the three exposure measures together across industries, or only within three-digit NAICS sectors, 9,999 times, and reports the share of draws with a coefficient at least as large in absolute value as the actual one; the wild cluster bootstrap is by sector, with the null imposed, 9,999 draws; placebo windows use the same contrast around earlier base years on the same 2017 exposures, restricted to windows that overlap neither each other nor the actual one: base years from 1990 to 2011, spaced six years apart, in six sets of three or four (Table B.5); the first set is 2011, 2005, 1999 and 1993. Panel C adds a control: the cyclical beta is the slope of the industry's TFP growth on the growth of manufacturing real output outside the industry (weights: lagged nominal sectoral output), estimated over 1988–2016 without the changes inside the window tested; the 2007–09 change in TFP is the change in log BLS TFP between those years. Panel D runs the same contrast on log hours per worker (BLS hours worked over employment), the proxy for unobserved utilization of Basu et al. (2006), with the randomization and placebo tests of Panel B, and reports its interquartile effect on three further versions of the protection measure: the coverage share built from the final estimation sample of Flaaen and Pierce (2024) and the rate-weighted measure they publish, which pools 17 industries into seven groups, with each group's value assigned to its members, and the rebuilt measure on the 69 industries they do not combine. Seed 2026.

Table B.5: The Flaaen–Pierce contrast around placebo base years

		Protection measure		Statutory rate	
Base year	Set	Coefficient	t	Coefficient	t
1990	4	-0.2009	-0.42	-0.0001	-0.07
1991	3	-0.1881	-0.32	-0.0002	-0.12
1992	2	-0.1741	-0.85	0.0003	0.25
1993	1	0.4909	1.11	0.0017	0.82
1994	6	0.6413	1.00	0.0017	0.87
1995	5	0.5248	0.87	0.0025	1.52
1996	4	0.1042	0.17	0.0022	1.58
1997	3	-0.4617	-0.69	-0.0011	-0.65
1998	2	-0.4768	-0.67	-0.0027	-1.42
1999	1	0.1888	0.36	-0.0029	-1.40
2000	6	-0.7126	-0.73	-0.0018	-0.59
2001	5	-0.7092	-1.02	-0.0008	-0.40
2002	4	0.6900	1.75	0.0001	0.06
2003	3	1.0346	3.46	0.0008	0.39
2004	2	-0.3150	-0.98	-0.0012	-0.92
2005	1	0.1544	0.42	0.0026	1.63
2006	6	0.0882	0.13	-0.0005	-0.25
2007	5	-0.4348	-0.64	-0.0034	-2.01
2008	4	-0.6393	-0.92	0.0000	0.01
2009	3	-0.1645	-0.34	0.0032	2.24
2010	2	-0.3977	-0.36	-0.0006	-0.20
2011	1	-0.2336	-0.20	-0.0025	-0.90
2017 (actual)		-1.2916	-7.64	-0.0039	-3.32

Notes: For base year b, the cross-sectional regression across 86 industries of $(y_{b+2} - y_b) - \frac{2}{3}(y_b - y_{b-3})$, with y log BLS TFP, on the exposure measure of the column, InputCost, Retaliation and the four controls of Flaaen and Pierce (2026), all fixed at 2017; t-statistics from standard errors clustered by three-digit sector. The statutory-rate coefficient is per percentage point of the 2017–19 change (Tables 1 and 2 report it per 100 points). Every year used by a placebo row is before the 2018 tariff actions.

Table B.6: Could the 2019 break in output per hour arise by chance?

	(1) Protection measure	(2) Statutory rate
A. Equation (3), 2014–17 pre-trend, 2019		
Coefficient	-1.6922	-0.0055
	(0.5933)	(0.0020)
B. Tests against chance		
Randomization across industries, p	0.011	0.016
Randomization within three-digit sectors, p	<0.001	0.004
Clustered by three-digit sector, t	-2.85	-2.72
Wild cluster bootstrap by sector, p	0.064	0.040
Placebo windows 2011, 2005, 1999, 1993: at least as large	1 of 4	0 of 4
Sets of non-overlapping placebo windows in which 2019 is the largest	1 of 6	5 of 6
Placebo windows with a break at least as large	5 of 22	1 of 22
C. Controlling for cyclical sensitivity		
With the cyclical beta	-1.6920	-0.0059
	(0.5937)	(0.0021)
sets in which 2019 is the largest	2 of 6	3 of 6
With the 2007–09 change in output per hour	-1.6969	-0.0054
	(0.6210)	(0.0021)
Industries / three-digit sectors	86 / 21	86 / 21

Notes: As Table B.4, with log BLS output per hour worked in place of log TFP in the equation (3) contrast, in the placebo windows and in both cyclical controls (the cyclical beta is the slope of the industry's growth in output per hour on the growth of manufacturing real output outside the industry). The utilization panel of Table B.4 does not depend on the outcome and is not repeated.

Table B.7: The 2019 break by pre-war industry characteristics

	Split at the median		Above-median indicator		Standardized	
Characteristic	Below	Above	Interaction	RI p	Interaction	RI p
No interaction	-2.34 (-7.64)					
China's share of imports, 2016–17	-4.57	-2.24	0.71	0.619	0.27	0.729
	(-2.52)	(-3.77)	(0.55)		(0.53)	
Import share of absorption, 2016	-2.67	-2.45	1.61	0.250	0.03	0.955
	(-1.44)	(-5.42)	(0.59)		(0.06)	
Price–cost margin, 2012–16	-2.81	-1.74	0.85	0.553	1.54	0.059
	(-5.76)	(-1.68)	(0.93)		(1.92)	
Change in China's import share, 2007–16	-2.51	-1.15	1.74	0.187	0.21	0.758
	(-4.94)	(-1.33)	(1.76)		(0.64)	
TFP growth, 2012–14	-3.60	-2.04	0.56	0.702	0.02	0.980
	(-2.30)	(-3.92)	(0.51)		(0.04)	
TFP growth, 2012–17†	-0.80	-1.95	-1.21	0.384	0.42	0.577
	(-0.52)	(-3.86)	(-1.09)		(1.16)	

Notes: The equation (3) contrast at the 2019 endpoint on the protection measure, with InputCost, Retaliation and the four controls of Flaaen and Pierce (2026), as in Table B.4. Coefficients are scaled by the interquartile range of the protection measure and multiplied by 100, so they read as the percent TFP gap between industries at the 75th and 25th percentile of exposure; t-statistics from standard errors clustered by three-digit sector in parentheses. Split at the median: the regression re-estimated on the industries below and above the median of the characteristic. Interaction: the coefficient on the protection measure times an indicator for an above-median value, or times the standardized characteristic, with the characteristic itself as a control (the import share of absorption is already a control). RI p: one plus the number of the 2,999 random permutations of the characteristic across industries with an interaction at least as large in absolute value, divided by 3,000. China's import shares are from Schott (2025); the price–cost margin is shipments less payroll and materials over shipments, from Becker et al. (2021). †Overlaps the 2014–17 pre-war drift of equation (3). Seed 2026.

Table B.8: The estimate in a common unit

	Δlog TFP per pp	t	% of TFP, interquartile gap	95% CI
Flaaen–Pierce design	-0.0039	-3.32	2.7	[1.0, 4.3]

Notes: Association of the statutory output tariff with log BLS TFP, per percentage point and scaled to the interquartile gap of 6.9 points in its 2017–19 increase across the 86 industries. The estimate is equation (2) with the 2017–19 change in the statutory output tariff as the protection measure, net of the 2014–17 pre-war drift as in equation (3). The last column is the 95 percent confidence interval for the interquartile gap, from the t distribution with standard errors clustered by three-digit sector. Sample: 86 industries over 2013–2022.

Table B.9: The Flaaen–Pierce design with controls for the pandemic years

	Protection measure		Statutory rate	
	(1) Baseline	(2) Controls	(3) Baseline	(4) Controls
2013	0.1272	0.1763	-0.0225	-0.0197
	(0.1115)	(0.1218)	(0.0404)	(0.0416)
2014	-0.2179	-0.2112	-0.0711	-0.0753
	(0.3417)	(0.3644)	(0.1013)	(0.0938)
2015	-0.1019	-0.0340	-0.0869	-0.0932
	(0.3179)	(0.3423)	(0.0853)	(0.1083)
2016	0.0364	0.1414	-0.1161	-0.1147
	(0.2654)	(0.2748)	(0.0856)	(0.0979)
2018	-0.3265	-0.4294*	-0.1158*	-0.1517**
	(0.1934)	(0.2135)	(0.0642)	(0.0713)
2019	-1.1463***	-1.2762***	-0.3390***	-0.3448***
	(0.2565)	(0.2073)	(0.0931)	(0.0992)
2020	-1.0522***	-1.1653***	-0.3591**	-0.3239**
	(0.3323)	(0.3020)	(0.1265)	(0.1293)
2021	-1.5994***	-1.9385***	-0.5054**	-0.5596**
	(0.5258)	(0.4441)	(0.1956)	(0.2243)
2022	-2.1535***	-2.5760***	-0.5869**	-0.6231**
	(0.7303)	(0.6136)	(0.2525)	(0.2963)
Equation (3), 2014–17 pre-trend, 2019	-1.2916***	-1.4171***	-0.3864***	-0.3951***
	(0.1692)	(0.1932)	(0.1165)	(0.1196)
Wild bootstrap p, 2019	0.015	0.002	0.001	0.001
Wild bootstrap p, equation (3)	0.000	0.000	0.005	0.006
Observations	860	860	860	860

Notes: Coefficients on the measure of exposure from equation (2), 2017 omitted, as in Table 1; columns 1 and 3 repeat its two columns. Columns 2 and 4 estimate equation (B.1), which adds two controls, each interacted with every year dummy: the industry's cyclical sensitivity (the slope of its annual TFP growth on the growth of manufacturing real output outside the industry, lagged nominal-output weights, 1988–2012) and the share of NAICS 3344 in its intermediate purchases (2012 input–output tables). Outcome: log BLS TFP detrended on each industry's 2013–17 path. InputCost, Retaliation and the four controls of Flaaen and Pierce (2026) are included and interacted with every year dummy. Industry and year fixed effects; SEs clustered by three-digit sector. Sample: 86 industries, 2013–2022. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

C The model

I state the one-sector model of Ghironi and Melitz (2005) with CES demand in full, and for the translog baseline and the three-sector extension of Sections 4 and 5 I list only what changes. A further subsection defines the averages of firms' revenue productivity that I set beside the firm-level evidence, and another states the versions of the one-sector model with identical firms that Section 4.5 uses to size the role of firm heterogeneity in the optimal tariff. Full derivations and the checks of the implementation are in the model notes in the replication package, together with the code.

C.1 The one-sector model with CES demand

The model is the two-country model of Ghironi and Melitz (2005) with international trade in two risk-free bonds, to which I add an ad valorem import tariff in each country. Section 4.2 states the environment, the households' and firms' problems and the aggregation and resource constraints, and Section 4.3 the calibration; here I complete the statement with the equilibrium, the steady state and the measurement of productivity and welfare, because the translog and multisector versions that follow change only parts of it, and the changes are easiest to read against a complete statement.

C.1.1 Equilibrium

Given initial masses of producers and bond holdings and processes for technology and the two tariffs, an equilibrium is a process for $\{C_t, w_t, \tilde{d}_t, \tilde{v}_t, N_{D,t}, N_{E,t}, N_{X,t}, z_{X,t}, r_t, B_{t+1}, B_{*,t+1}\}$, their foreign counterparts and Q_t such that households solve their problems, delivering (6)–(7); firms price by (8) and export when (9) is non-negative, which gives (13); entry satisfies (10); the price index (14) holds in each country; labor and bond markets clear, (15)–(16); net foreign assets evolve by (17); and the transversality conditions hold.

C.1.2 Steady state

In the symmetric steady state with zero tariffs and zero bond holdings, $Q = 1$, $1 + r = 1/\beta$, and the share Euler equation and free entry fix average profit at $\tilde{d} = \frac{r+\delta}{1-\delta} w f_E$. Substituting the cutoff condition and the export share into the definition of $\tilde{d}$ eliminates the real wage and leaves one equation in the export cutoff,

$$f_X \left[\left(\frac{\nu \tau z_{\min}}{z_X} \right)^{\theta-1} + \left(\frac{z_{\min}}{z_X} \right)^{k} \frac{\theta - 1}{k - (\theta - 1)} \right] = \frac{r + \delta}{1 - \delta} f_E, \tag{C.1}$$

whose left side falls monotonically in z_X, so the root is unique. Everything else follows in closed form; the calibration is in Section 4.3.

C.1.3 Measured productivity and welfare

Measured TFP in the model is what a statistical agency would compute for the home producers, real output over the labor they employ, which is production labor L_t^p plus export overhead L_t^f, with the labor used to create new firms excluded because it belongs to no producer. From (9),

$$L_t^p = \frac{\theta - 1}{w_t}\big(N_{D,t}\tilde{d}_{D,t} + N_{X,t}\tilde{d}_{X,t}\big) + (\theta - 1)\frac{N_{X,t}f_X}{Z_t}, \qquad L_t^f = \frac{N_{X,t}f_X}{Z_t}, \tag{C.2}$$

and I write $s_{f,t} = L_t^f/(L_t^p + L_t^f)$ for the overhead share. Every firm's revenue is A times its production labor costs, so revenue per unit of production labor is Aw_t for every firm whatever its draw, because a more productive firm charges a proportionally lower price. Three measures then differ only in the deflator. Valuing each firm's output at its own steady-state price, as a product-by-product deflator does, gives the measure that corresponds to the BLS series, which I write as TFP without a superscript,

$$\mathrm{TFP}_t \propto Z_t(1 - s_{f,t}), \tag{C.3}$$

which moves only with technology and with the overhead share. Valuing output in units of consumption gives $\mathrm{TFP}_t^C = Aw_t(1 - s_{f,t})$, which also moves with the real wage. And deflating revenue by the data-consistent price index $\tilde{P}_t$, the average price of the varieties available at home, which under CES satisfies $P_t = N_t^{1/(1-\theta)}\tilde{P}_t$ with $N_t = N_{D,t} + N_{X,t}^*$, gives $\mathrm{TFP}_t^A = \mathrm{TFP}_t^C/N_t^{1/(\theta-1)}$. Technology, $Z_t\nu z_{\min}$, does not move after a tariff.

The shift of production toward less productive firms shows in physical terms. Production labor in the home market is distributed across firms in proportion to $z^{\theta-1}$, so under Pareto draws its distribution of $\ln z$ has mean $\ln z_{\min} + 1/(k - \theta + 1)$, and likewise above the export cutoff; overhead labor is spread evenly across exporters, whose mean $\ln z$ is $\ln z_{X,t} + 1/k$. Writing L_t^D and L_t^X for production labor serving the home and foreign markets, physical productivity, log technology plus the mean of $\ln z$ over production labor, is

$$\ln Z_t + \frac{L_t^D\big(\ln z_{\min} + \frac{1}{k-\theta+1}\big) + L_t^X\big(\ln z_{X,t} + \frac{1}{k-\theta+1}\big)}{L_t^p}, \tag{C.4}$$

the last row of Table 4. The mean over all labor, which adds the export overhead, is

$$\overline{\ln z_t} = \frac{L_t^D\left(\ln z_{\min} + \frac{1}{k-\theta+1}\right) + L_t^X\left(\ln z_{X,t} + \frac{1}{k-\theta+1}\right) + L_t^f\left(\ln z_{X,t} + \frac{1}{k}\right)}{L_t^p + L_t^f}. \tag{C.5}$$

The labor-weighted mean of z itself does not exist at this calibration, because $k < \theta$.

Revenue productivity is the measure that firm-level studies of trade policy observe when they deflate a firm's revenue by an industry price index, the measure Foster et al. (2008) call traditional TFP. Unlike the measures above it is computed firm by firm. A firm's revenue productivity is its revenue over the labor it employs, production labor and, for an exporter, the export overhead, because a statistical agency observes total employment and not its split. In units of consumption, a firm's revenue is $R_t(z) = R_{D,t}(z) + R_{X,t}(z)$, with $R_{D,t}(z) = \theta d_{D,t}(z)$ and, for an exporter, $R_{X,t}(z) = \theta\big[d_{X,t}(z) + w_t f_X/Z_t\big]$ from (9), and its production labor is $R_t(z)/(Aw_t)$. Therefore

$$\mathrm{TFPR}_t(z) = \frac{R_t(z)}{R_t(z)/(Aw_t) + \mathbf{1}[z \geq z_{X,t}]\, f_X/Z_t} = \frac{Aw_t}{1 + \mathbf{1}[z \geq z_{X,t}]\, Aw_t f_X/\big(Z_t R_t(z)\big)}. \tag{C.6}$$

Every firm that does not export records Aw_t whatever its draw, because its price undoes its physical productivity exactly. Accordingly, $\mathrm{TFPR}_t(z) = \rho_{D,t}(z)\, Z_t z$, the identity of Foster et al. (2008) in which revenue productivity is the firm's price times its physical productivity; with a single sector price index as the deflator, their two revenue measures differ only by a constant. Exporters record less, by an amount that shrinks as their revenue grows and spreads the overhead. Under CES, revenue productivity therefore carries no information about a firm's draw beyond whether it exports, and a tariff moves it only through the real wage and through the overhead of exporters. Section C.4 states the deflator and the averages across firms.

Welfare is conditional welfare when the tariff arrives, carried in the model as $\mathcal{W}_t = C_t^{1-\gamma}/(1-\gamma) + \beta E_t \mathcal{W}_{t+1}$, and reported as the permanent change in consumption with the same value, the λ that solves $\sum_t \beta^t[(1+\lambda)\bar{C}]^{1-\gamma}/(1-\gamma) = \mathcal{W}$, that is

$$\lambda = \left(\frac{\mathcal{W}}{\mathcal{W}^{ss}}\right)^{\frac{1}{1-\gamma}} - 1, \qquad \mathcal{W}^{ss} = \frac{\bar{C}^{1-\gamma}}{(1-\gamma)(1-\beta)}, \tag{C.7}$$

with $\mathcal{W}$ read in the period the tariff is announced, computed from the exact deterministic solution by perfect foresight. The technology equivalent of the welfare gain is the persistent decline in Z_t, at the tariff's persistence, that would cost the home household as much as the tariff gains it. Section 4 reports all these objects.

C.2 Translog demand

The translog version replaces the CES aggregator with the symmetric translog expenditure function of Bergin and Feenstra (2000), in the continuum-of-goods form that Feenstra (2003) generalizes to a varying number of goods. Bilbiie et al. (2012) use it in a closed-economy model with identical firms, and Rodríguez-López (2011) in a two-country model with heterogeneous firms; here it is carried over to the model of Section C.1. Technology, entry and exit, trade costs, the asset structure and the household's intertemporal problem are unchanged, and so are (5), (6), (7), (10), (16) and the form of (17). Only the demand system changes, and with it everything that depends on how firms price. The translog demand, markups, price index and labor market are stated in Section 4.2.

C.2.1 Steady state and calibration

The symmetric steady state reduces to three unknowns, the domestic cutoff, the export cutoff and the mass of producers, pinned by consistency of the mean log price, the zero-profit condition of the marginal exporter, and free entry; labor market clearing is then implied and serves as a check. The calibration is in Section 4.3.

C.2.2 Measured productivity under translog

With firm-specific markups, revenue per unit of production labor is $\mu(z)w_t$ and no longer the same across firms, so the shift of production toward less productive firms reaches measured TFP. Writing $l(z)$ for a firm's production labor in a market and $\bar{\mu}(z)$ for the markup it would charge at steady-state aggregates, the product-by-product measure becomes

$$\text{TFP}_t \propto \frac{Z_t}{\bar{Z}} \frac{N_{D,t}E[\bar{\mu}_D l_D] + N_{X,t}E[\bar{\mu}_X l_X]}{L_t^p + L_t^f}, \tag{C.8}$$

which falls when labor moves toward low-markup, less productive firms, and reduces to (C.3) when markups are common. Measured TFP in units of consumption is revenue over $L_t^p + L_t^f$. The variety elasticity of the translog price index, under the normalization (18), is $\varepsilon_V = \bar{L} + 1/(\sigma N)$, half its CES value $1/(\theta - 1)$ when σ matches the CES average markup, so the data-consistent measure is $\text{TFP}_t^A = \text{TFP}_t^C (N_t/\bar{N})^{-\varepsilon_V}$. Physical productivity is log technology plus the mean of $\ln z$ over production labor, as under CES; under translog the mean has no closed form, and I compute it using a quadrature method. The mean of z itself exists under translog, because markups bound the size of the largest firms, but not under CES with $k < \theta$, so I use the log measure to compare the two models. Physical productivity falls after a tariff although no firm's productivity changes. I split its impact

response into two margins. The intensive margin evaluates (C.4) at the impact values of every variable except the export cutoff, which stays at its steady-state value, so it captures the reallocation of production labor among the firms that sell in each market; the extensive margin of exporting is the remainder, the contribution of the move in the export cutoff.

Revenue productivity is defined as under CES, as revenue over total labor, but markups now differ across firms. A firm's revenue in each market is its expenditure share times expenditure there, in home consumption units and, for exports, net of the foreign tariff, and its production labor in that market is revenue over $\mu(z)w_t$, so

$$\mathrm{TFPR}_t(z) = \frac{R_{D,t}(z) + R_{X,t}(z)}{\dfrac{R_{D,t}(z)}{\mu_{D,t}(z)w_t} + \dfrac{R_{X,t}(z)}{\mu_{X,t}(z)w_t} + \mathbf{1}[z \geq z_{X,t}]\dfrac{f_X}{Z_t}}, \tag{C.9}$$

which for a firm that does not export is $\mu_{D,t}(z)w_t$. Because the markup rises with market share, a more productive firm now records higher revenue productivity, though by less than its physical advantage. Intuitively, the part of its cost advantage that the firm keeps as margin reaches revenue productivity, and the part it passes into its price does not. A tariff moves revenue productivity through the real wage, through each firm's markup and through the overhead of exporters.

C.3 The multisector extension

The multisector model extends the translog model of Section C.2 in one direction, giving each country three manufacturing sectors instead of one, standing for the quarter of industries most exposed to the 2018 tariffs, the middle half and the least exposed quarter. The full statement, with every problem, first-order condition and equilibrium condition, is in the model notes in the replication package; here I list what changes relative to the one-sector translog model.

C.3.1 What changes

The household spends a fixed share α_j of its budget on each sector, through a Cobb–Douglas aggregate of three sector baskets, $C_t = \prod_j (C_{j,t}/\alpha_j)^{\alpha_j}$, with the aggregate price index the numeraire, $\sum_j \alpha_j \ln P_{j,t} = 0$. Within each sector the basket is a translog aggregate with its own parameter σ_j, so the share equation (11), the markup (12) and the sector price index (18) hold sector by sector, with the moments of log prices computed over the varieties of the sector sold in each country and expenditure $\alpha_j C_t$ in place of C_t. Firms of sector j draw productivity from a Pareto distribution with shape k_j and produce $Z_t Z_{j,t} z$ units per unit

of labor, where $Z_{j,t}$ is sector technology. Iceberg costs are specific to the sector and to the direction of trade, τ_j for home goods shipped abroad and τ_j^* for foreign goods shipped home, which lets a sector's import share differ from its export share. The fixed export cost is $f_{X,j}$. The foreign labor force is $L^* = 3$ times home's. The foreign market of each sector has its own translog parameter σ_j^*, which replaces σ_j in (11), (12) and (18) for every variety sold there, home exports included. Each sector has its own labor market, with wage $w_{j,t}$ and a fixed labor force.[25] The mutual fund holds the incumbents of each sector, the share Euler equation and free entry hold sector by sector, and the net foreign asset condition sums labor income, profits and investment over sectors. Each sector has its own domestic cutoff from its choke price; at the calibration below every one of them is slack, so every entrant sells at home and the domestic quadrature nodes are fixed, as in the one-sector model.

Tariffs and the technology of the most exposed sector are announced in 2018Q1. A single news shock drives the sector tariffs, with loadings equal to the quarterly innovations of each sector's phased-in tariff path, and two news shocks drive $\ln Z_{1,t}$, one for the innovations of 2018 and one for those of 2019, with persistence $\rho_S = 0.999$. Sector TFP is measured as BLS measures it, by (C.8) applied to each sector, and the technology shocks are sized so that the model's gap in measured TFP between the most and least exposed sectors matches the estimates. Welfare is conditional welfare in 2018Q1.

C.3.2 Calibration

The calibration starts from the model with home a quarter of the world and proceeds in four steps. The first keeps the parameters of that model that are not specific to a sector: $\beta = 0.99$, $\gamma = 2$, $\delta = 0.025$, $f_E = z_{\min} = L = 1$, a foreign labor force $L^* = 3$, the bond fee, rescaled by the ratio of steady-state consumption so that it bears the same relation to the economy, and the fixed export cost of that model, 0.0065, in the most exposed sector. In the second step I read the sector parameters backward from the steady state, so that the model reproduces, for each group of industries, the targets of Table C.2. Each target pins one parameter:

1. the expenditure share α_j is the group's share of absorption;

[25] I hold each sector's labor force fixed by letting labor move across sectors only at the cost $\Gamma_{j,t} = \frac{\varphi}{2}\omega_t L_{j,t-1}(L_{j,t}/L_{j,t-1} - 1)^2$ units of consumption, where ω_t is the shadow wage on $\sum_j L_{j,t} = L$. The household's allocation of labor then satisfies

$$\omega_t = w_{j,t} - \varphi\,\omega_t(g_{j,t} - 1) + \beta E_t\left[\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\frac{\varphi}{2}\omega_{t+1}(g_{j,t+1}^2 - 1)\right], \qquad g_{j,t} = \frac{L_{j,t}}{L_{j,t-1}}.$$

I set $\varphi = 10^5$, so that no more than 0.02 percent of the labor force changes sector in any experiment I report.

2. σ_j sets the expenditure-weighted markup of the sector equal to the markup a CES sector would need to reproduce the group's labor share of value added relative to the middle group, $\theta_j/(\theta_j-1)$, with $\theta_2 = 3.8$ as in the one-sector model; this is the rule that calibrates σ in the one-sector translog model, applied sector by sector;

3. k_j sets the standard deviation of log domestic sales across producers at 0.82, the value the published calibration implies, $(\theta-1)/k = 2.8/3.4$;

4. τ_j^* sets the import share of absorption, and τ_j the export share of shipments in the two outer sectors, with $\tau_2 = \tau_2^*$ in the middle sector;

5. $f_{X,j}$ sets the share of firms that export in the middle and least exposed sectors;

6. σ_j^*, the translog parameter of the foreign market, sets its expenditure-weighted markup equal to home's.

I impose balanced trade in the model, while in the data manufacturing imports exceed exports. I therefore split the imbalance equally, in logs, between the two sides of trade. Every export-to-domestic-sales ratio is scaled up, and every import share scaled down, by the common factor 1.28 that makes exports equal imports. The scaling keeps the ranking of the sectors and their relative trade intensities, and the middle sector's export share is then implied by balanced trade. In the most exposed sector the share of firms that export is an outcome, which Table C.3 reports beside the data.

In the third step I set the tariff paths (Table C.1). Each sector's tariff follows the group's import-weighted tariff change at the dates of the 2018 actions, so the model reproduces the manufacturing-wide increase of 3.13 points and the 7.62-point gap between the outer groups. In the fourth step I size the technology shocks to the estimated gap in measured TFP, −1.56 percent in 2018 and −5.05 percent in 2019.

Table C.1: Quarterly tariff paths, percentage points

	2018Q1	Q2	Q3	Q4	2019Q1	Q2	Q3	Q4
Most exposed	0.38	2.02	2.91	4.96	4.96	6.59	7.81	7.81
Middle	0.00	0.03	0.58	1.32	1.32	1.92	2.37	2.37
Least exposed	0.00	0.00	0.08	0.13	0.13	0.17	0.20	0.20

Notes: The solar and washer safeguards from February 7, 2018; steel and aluminum from March 23; List 1 from July 6; List 2 from August 23; List 3 from September 24 at 10 percent and from May 10, 2019, at 25 percent.

The calibration is found by continuation from an exact three-sector copy of the one-sector model, in which $\sigma_j = \sigma/\alpha_j$ and the mass of incumbents of sector j is α_j times the one-sector

Table C.2: Calibration targets of the three-sector model

	Most exposed	Middle	Least exposed	Pins
Share of absorption	0.149	0.467	0.385	α_j
Labor share of value added (data)	0.473	0.435	0.276	
CES-equivalent markup, $\theta_j/(\theta_j-1)$	1.249	1.357	2.142	σ_j
Standard deviation of log domestic sales	0.824	0.824	0.824	k_j
Share of firms that export	0.237	0.223	0.160	$f_{X,j}$ (middle, least exposed)
Import share of absorption (data)	0.512	0.319	0.214	
scaled to balanced trade	0.400	0.250	0.168	τ_j^*
Export share of shipments (data)	0.309	0.219	0.135	
scaled to balanced trade	0.364	(0.264)	0.166	τ_j
Import-weighted tariff change, points	7.81	2.37	0.20	tariff paths

Notes: The 86 industries of the empirical sample by quartile of the protection measure (Appendix Table C.4). Absorption, import and export shares from the trade and shipments data behind the protection measure, aggregated with absorption and shipment weights; tariff changes from the coverage and rates in the replication package of Flaaen and Pierce (2024), applied to each industry's 2017 imports; labor shares are BLS labor cost shares of gross output divided by the NBER–CES ratio of value added to shipments, 2016; the CES-equivalent markup is $\theta_j/(\theta_j-1)$, with θ_j set so that $(\theta_j-1)/\theta_j$ stands to the middle group's 2.8/3.8 as the group's labor share stands to the middle group's. The dispersion of log domestic sales is the value implied by the published calibration, $(\theta-1)/k = 2.8/3.4$. Trade shares are scaled to balanced trade by the common factor 1.28 (export-to-domestic-sales ratios up, import shares down); the middle group's export share, in parentheses, is implied by balanced trade. Shares of exporting firms are those of Appendix Table C.4, weighted by the number of employer firms; the most exposed sector's is not targeted.

Table C.3: Calibrated sector parameters and implied steady state, translog model

	Most exposed	Middle	Least exposed
Translog parameter σ_j	6.346	1.165	0.253
Translog parameter abroad σ_j^*	2.188	0.431	0.095
Pareto shape k_j	4.768	3.297	1.171
Iceberg cost, home exports τ_j	1.351	1.671	5.551
Iceberg cost, home imports τ_j^*	1.198	1.671	8.949
Fixed export cost $f_{X,j}$	0.0065	0.0041	0.0007
Expenditure-weighted markup	1.249	1.357	2.142
Expenditure-weighted markup abroad	1.249	1.357	2.142
Import share of absorption	0.400	0.250	0.168
Export share of shipments	0.364	0.264	0.166
Share of home firms exporting	0.330	0.223	0.160
Standard deviation of log domestic sales	0.824	0.824	0.824
Domestic cutoff, $z_{D,j}/z_{\min}$	0.967	0.950	0.839
Export cutoff, $z_{X,j}/z_{\min}$	1.262	1.576	4.777
Varieties sold at home	1.05	4.22	6.55

Notes: Parameters solved jointly with the steady state so that the model reproduces the targets of Table C.2. The fixed export cost of the most exposed sector is set at the one-sector value, and that sector's share of exporting firms is an outcome. In the data the three shares are 23.7, 22.3 and 16.0 percent. Every domestic cutoff lies below $z_{\min}$, so every entrant sells at home.

mass, so that every markup, cutoff and firm revenue is unchanged. Each firm's share of the spending on its sector is $1/\alpha_j$ times its share of spending in the one-sector model. Two features of the translog calibration differ from what a CES sector would deliver. The most exposed sector needs a large σ_1 to reach its low markup, because it has few varieties and translog markups fall with the number of competitors. When I first built the calibration, with two countries of equal size, the sectors could not trade as intensively as the data require at the data's own trade shares, because markups compress the distribution of sales and leave exporters a smaller share of revenue than under CES. For that reason I split the trade imbalance across both sides instead of loading it on exports alone, and I keep the split with home a quarter of the world. The least exposed sector has the lowest labor share of value added, so the markup rule gives it the highest markup, 2.142. The sector reaches that markup with a small translog parameter, a Pareto shape close to one and high iceberg costs (Table C.3).

C.3.3 The industries by quartile of the protection measure

Table C.4 lists the 86 industries of the empirical sample by quartile of the protection measure, the grouping of the right panels of Figure 1. The panels hold the 22 most exposed industries, the 42 of the middle half and the 22 least exposed, and each closes with the values of the group as a whole. The most exposed quarter accounts for 14.9 percent of the industries' absorption and draws 51.2 percent of it from imports, against 38.5 and 21.4 percent for the least exposed quarter. Across all 86 industries, 21.2 percent of employer firms export, close to the 21 percent of the one-sector calibration, and the share falls from 23.7 percent in the most exposed quarter to 16.0 percent in the least exposed.

Table C.4: The industries of the empirical sample by quartile of the protection measure

NAICS	Industry	Protection (% of home market)	Absorption share (%)	Import share (%)	Exporting firms (%)
Most exposed quarter of the protection measure (22 industries)					
3169	Other leather and allied product manufacturing	12.05	0.20	98.2	21.3
3351	Electric lighting equipment	7.42	0.41	57.4	42.5
3371	Household and institutional furniture and kitchen cabinets	6.01	1.04	43.6	7.8

Table C.4, continued

NAICS	Industry	Protection (% of home market)	Absorption share (%)	Import share (%)	Exporting firms (%)
3159	Apparel accessories and other apparel manufacturing	5.84	0.09	87.7	14.1
3311	Iron and steel mills and ferroalloy production	5.46	1.49	26.1	36.9
3344	Semiconductors and other electronic components	5.34	2.03	58.9	51.3
3352	Household appliances	4.33	0.67	64.2	42.5
3359	Other electrical equipment and components	3.71	1.03	49.2	42.5
3332	Industrial machinery	3.62	0.57	63.8	39.4
3322	Cutlery and handtools	3.59	0.25	53.2	17.4
3313	Alumina and aluminum production and processing	3.01	0.68	31.3	36.9
3325	Hardware	2.98	0.21	55.2	17.4
3132	Fabric mills	2.85	0.23	43.8	31.3
3372	Office furniture (including fixtures)	2.80	0.51	25.2	7.8
3353	Electrical equipment	2.74	0.88	57.4	42.5
3151	Apparel knitting mills	2.59	0.06	79.1	14.1
3329	Other fabricated metal products	2.39	1.33	38.4	17.4
3341	Computer and peripheral equipment	2.29	1.48	94.0	51.3
3369	Other transportation equipment	2.03	0.32	31.3	35.2
3336	Engine, turbine, and power transmission equipment	2.02	0.72	49.4	39.4
3272	Glass and glass products	1.98	0.45	24.9	12.0
3379	Other furniture related products	1.94	0.19	18.2	7.8
Sector		3.92	14.9	51.2	23.7
Middle half of the protection measure (42 industries)					
3326	Spring and wire products	1.85	0.17	28.0	17.4
3334	HVAC and commercial refrigeration equipment	1.82	0.80	27.2	39.4
3333	Commercial and service industry machinery	1.78	0.47	35.0	39.4

Table C.4, continued

NAICS	Industry	Protection (% of home market)	Absorption share (%)	Import share (%)	Exporting firms (%)
3342	Communications equipment	1.75	2.16	86.4	51.3
3271	Clay products and refractories	1.68	0.20	51.7	12.0
3312	Steel products from purchased steel	1.64	0.32	10.3	36.9
3149	Other textile product mills	1.60	0.26	41.6	11.0
3339	Other general purpose machinery	1.51	1.84	38.7	39.4
3262	Rubber products	1.42	0.86	36.4	38.4
3212	Veneer, plywood, and engineered wood product manufacturing	1.33	0.44	23.1	9.8
3279	Other nonmetallic mineral products	1.23	0.45	23.0	12.0
3343	Audio and video equipment manufacturing	1.22	0.54	101.9	51.3
3335	Metalworking machinery	1.13	0.60	34.0	39.4
3133	Textile and fabric finishing and fabric coating mills	1.10	0.13	17.0	31.3
3363	Motor vehicle parts	0.99	5.29	34.9	35.2
3345	Electronic instruments	0.91	2.51	34.0	51.3
3331	Agriculture, construction, and mining machinery	0.82	1.30	32.6	39.4
3219	Other wood products	0.79	0.86	10.4	9.8
3117	Seafood product preparation and packaging	0.78	0.25	16.3	10.1
3261	Plastics products	0.66	3.27	14.9	38.4
3161	Leather and hide tanning and finishing	0.63	0.02	60.9	21.3
3324	Boilers, tanks, and shipping containers	0.61	0.51	15.0	17.4
3222	Converted paper products	0.59	1.68	8.8	34.9
3251	Basic chemicals	0.58	3.28	23.8	40.9
3141	Textile furnishings mills	0.54	0.43	52.4	11.0
3315	Foundries	0.50	0.49	4.6	36.9
3131	Fiber, yarn, and thread mills	0.46	0.10	11.3	31.3
3362	Motor vehicle bodies and trailers	0.46	0.55	9.7	35.2

Table C.4, continued

NAICS	Industry	Protection (% of home market)	Absorption share (%)	Import share (%)	Exporting firms (%)
3256	Soaps, cleaning compounds, and toilet preparations	0.46	1.29	14.6	40.9
3252	Resin, synthetic rubber, and artificial synthetic fibers and filaments	0.45	1.25	22.3	40.9
3259	Other chemical products and preparations	0.45	0.79	28.0	40.9
3114	Fruit and vegetable preserving and specialty food	0.44	1.22	13.0	10.1
3323	Architectural and structural metals	0.39	1.42	6.4	17.4
3327	Machine shops; turned products; and screws, nuts, and bolts	0.35	1.03	7.6	17.4
3253	Pesticides, fertilizers, and other agricultural chemicals	0.33	0.55	23.5	40.9
3365	Railroad rolling stock manufacturing	0.33	0.23	7.2	35.2
3273	Cement and concrete products	0.25	0.94	3.1	12.0
3399	Other miscellaneous manufacturing	0.19	2.07	64.8	19.6
3314	Nonferrous metal (except aluminum) production and processing	0.18	0.89	70.5	36.9
3231	Printing and related support activities	0.15	1.34	6.6	6.7
3152	Cut and sew apparel	0.15	1.34	92.6	14.1
3364	Aerospace products and parts	0.14	2.57	31.3	35.2
Sector		0.74	46.7	31.9	22.3
Least exposed quarter of the protection measure (22 industries)					
3321	Forging and stamping	0.12	0.50	2.4	17.4
3255	Paint, coatings, and adhesives	0.10	0.61	4.7	40.9
3346	Manufacturing and reproducing magnetic and optical media	0.10	0.14	97.0	51.3
3113	Sugar and confectionery products	0.09	0.66	19.1	10.1
3111	Animal food manufacturing	0.09	0.88	2.4	10.1

Table C.4, continued

NAICS	Industry	Protection (% of home market)	Absorption share (%)	Import share (%)	Exporting firms (%)
3361	Motor vehicles	0.08	7.75	41.1	35.2
3221	Pulp, paper, and paperboard mills	0.06	1.20	14.0	34.9
3211	Sawmills and wood preservation	0.06	0.54	20.7	9.8
3391	Medical equipment and supplies	0.05	1.62	34.6	19.6
3366	Ship and boat building	0.05	0.53	5.8	35.2
3112	Grain and oilseed milling	0.04	1.38	12.1	10.1
3119	Other food products	0.04	1.75	7.8	10.1
3118	Bakeries and tortilla manufacturing	0.03	1.19	6.9	10.1
3241	Petroleum and coal products	0.03	6.54	10.4	23.2
3116	Animal slaughtering and processing	0.02	3.18	4.9	10.1
3122	Tobacco manufacturing	0.02	0.70	3.2	10.9
3274	Lime and gypsum products	0.01	0.12	2.3	12.0
3115	Dairy products	0.01	1.76	2.7	10.1
3121	Beverages	0.01	2.04	17.1	10.9
3254	Pharmaceuticals and medicine	0.00	4.52	40.3	40.9
3162	Footwear manufacturing	0.00	0.43	95.1	21.3
3328	Coating, engraving, heat treating, and allied activities	0.00	0.41	0.0	17.4
Sector		0.04	38.5	21.4	16.0

Notes: The 86 four-digit NAICS (2012) industries of the empirical sample, with BLS industry titles, ranked by the protection measure, the added duty of the 2018 actions on an industry's 2017 imports as a share of the value of its 2016 home market (Section 3.1); quartiles over the 86 industries, with NAICS 3321 and 3379 at the 25th and 75th percentiles. Absorption is 2016 shipments (Becker et al., 2021) plus imports less domestic exports (Schott, 2025, 2018); the absorption share is a percent of the 86 industries' total, and the import share is imports over absorption, the second factor of the protection measure. It exceeds 100 percent for audio and video equipment (NAICS 3343), whose recorded domestic exports exceed its shipments. The share of exporting firms takes identified exporters in 2016 by three-digit company NAICS (U.S. Census Bureau, 2019, Table 7b), less the exporters without known employment at the rate for all manufacturers in its Table 1b (76 percent of manufacturing exporters have known employment), and divides them by employer firms by three-digit NAICS in the Statistics of U.S. Businesses (U.S. Census Bureau, 2018). No public source reports the share at four digits, so each industry carries the value of its three-digit subsector. Group rows: the protection measure and the import share weighted by absorption, absorption shares summed, and the share of exporting firms weighted by the number of employer firms in each four-digit industry. Three checks are not shown in the table. The BEA 2017 detail Use table (U.S. Bureau of Economic Analysis, n.d.) gives import shares of 48.1, 29.7 and 23.7 percent for the three groups; the same exporting measure for 2017 (U.S. Census Bureau, 2020, 2021) gives 24.4, 22.5 and 15.8 percent; and Bernard et al. (2007, Table 2), for 2002, gives 19.6, 16.5 and 14.2 percent.

C.3.4 Sector and national detail of the two experiments

Figure C.1 reports the sector and national paths behind the two experiments of Section 5.1.2, the tariffs alone and the tariffs with the estimated decline.

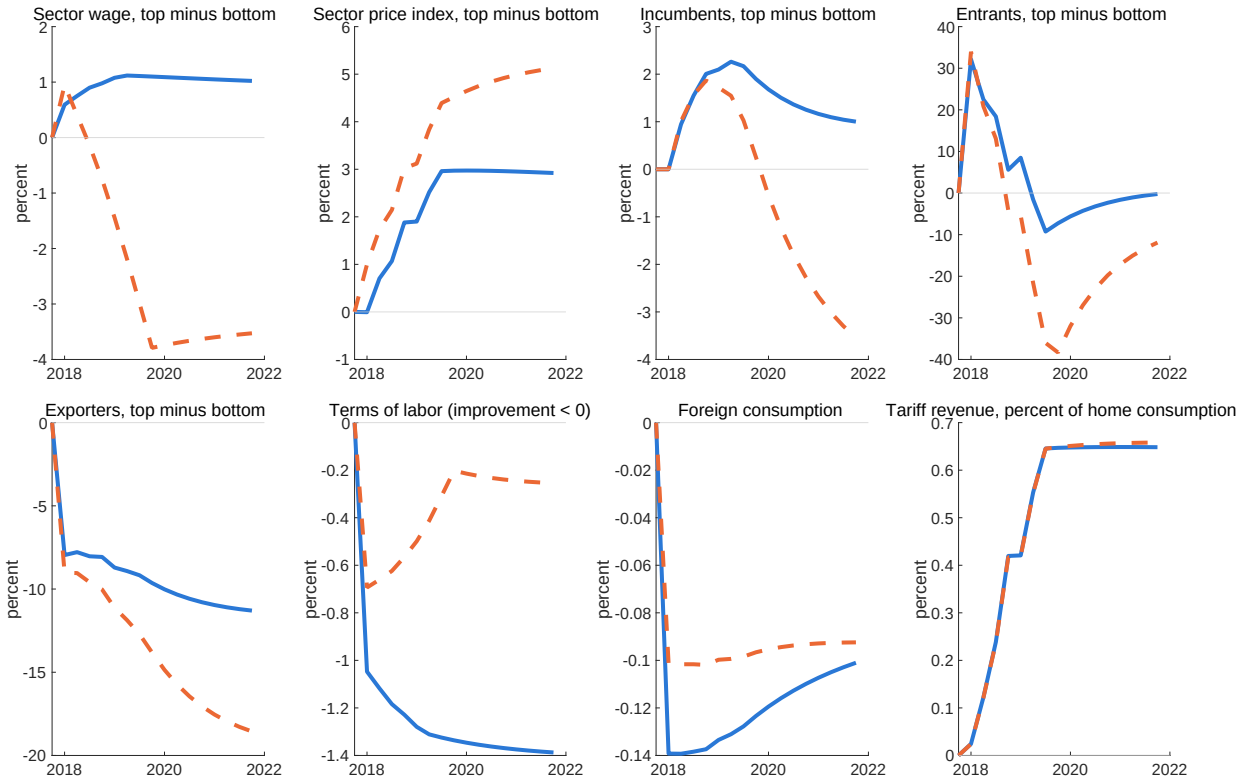
C.3.5 The optimal unilateral tariff

For a uniform tariff t on every sector, imposed by surprise with persistence 0.999, home welfare $\lambda(t)$, as a consumption equivalent, comes from the exact deterministic solution of the model, computed by perfect foresight with the first guess set to the path implied by the linear decision rules. The gain per point at a zero tariff is $a = \lambda'(0)$, which I compute from $t = 1$ and $t = 2$ percent, $a = 2\lambda(1) - \lambda(2)/2$. The optimal tariff t^{opt} maximizes $\lambda(t)$; I find it by evaluating λ at steps of 10 points until it falls and refining by golden-section search with parabolic interpolation. If the tariff also lowers every sector's technology by c percent per point, with the same persistence, the slope at a zero tariff becomes $a - \kappa c$, where κ is the welfare cost of a 1 percent technology decline; because λ is concave in the tariff, the optimal tariff is zero whenever $c > a/\kappa$. Section 5.1 reports a/κ against the estimated response.

C.3.6 External positions

The baseline has zero bond holdings in the steady state, because the quadratic fee on bond holdings is centered on zero. For the experiments with external positions I center the fee on target positions, so that the steady state holds gross positions, home-currency liabilities

Figure C.1: Sector and national detail of the two experiments, three-sector model



Notes: Solid lines: tariffs only; dashed lines: tariffs plus the technology shock. Sector panels: wages, sector price indexes, incumbents, entrants and exporters, most minus least exposed sector, percent. National panels: the terms of labor (a fall is a home gain), foreign consumption and tariff revenue as a percent of steady-state home consumption. Translog demand; quarterly paths from the exact deterministic solution (perfect foresight), labor nearly fixed across sectors.

$\bar{B} < 0$ and foreign-currency assets $\bar{B}_* > 0$. The home budget constraint replaces the fee terms with $\frac{\eta}{2}(B_{t+1} - \bar{B})^2 + \frac{\eta}{2} Q_t (B_{*,t+1} - \bar{B}_*)^2$, the foreign household's fees are centered on $-\bar{B}$ and $-\bar{B}_*$, its holdings when bond markets clear, and the bond Euler equations become

$$\begin{aligned} C_t^{-\gamma}\big(1 + \eta(B_{t+1} - \bar{B})\big) &= \beta(1 + r_{t+1}) E_t C_{t+1}^{-\gamma}, \\ C_t^{-\gamma}\big(1 + \eta(B_{*,t+1} - \bar{B}_*)\big) &= \beta(1 + r^*_{t+1}) E_t \left[\frac{Q_{t+1}}{Q_t} C_{t+1}^{-\gamma}\right]. \end{aligned} \tag{C.10}$$

In the steady state the fees are zero at the targets, so $1 + r = 1 + r^* = 1/\beta$, and the home country earns net interest $r(\bar{B} + Q\bar{B}_*)$ each period. Its steady-state budget constraint is $C = \sum_j w_j L_j + \sum_j (N_j^{\text{in}} \tilde{d}_j - N_{E,j} \tilde{v}_j) + r(\bar{B} + Q\bar{B}_*)$, with N_j^{in} the mass of incumbent firms in sector j, so a country with a negative net position runs a trade surplus equal to its interest payments. I set the positions to the U.S. international investment position at the end of 2017, as shares of annual GDP (U.S. Bureau of Economic Analysis, 2026a,b): liabilities of 181 percent and assets of 141 percent. As in Itskhoki and Mukhin (2026), the liabilities are in home currency and the assets in foreign currency. The net-position variant holds the same net position entirely in home currency. Because returns on assets and liabilities are equal in the model, a negative net position implies a steady-state trade surplus, whereas the United States runs a deficit; Itskhoki and Mukhin reconcile the two through a return differential, which this model omits. The one-sector experiments with positions use the three-sector model with three identical sectors, which reproduces the one-sector model exactly.

C.4 Revenue productivity in the model

To set the model beside the firm-level evidence on revenue productivity, I compute the averages that those studies estimate, with the deflator they use, from the firm-level measures (C.6) and (C.9). Three choices define them: the price index that deflates revenue, the weights that average across firms and, in the three-sector model, the comparison across sectors.

The first choice is the deflator. Firm-level studies deflate revenue by an industry producer price index, and I build that index in the model as the BLS builds the PPI, from matched items: the prices of the domestic sales of all home producers and the export prices of the firms that export both in the steady state and in the period measured, each item weighted by its steady-state revenue share, with log changes averaged geometrically. Prices are producer prices in home consumption units, so an export price is net of the foreign tariff and converted at the real exchange rate. Writing ϖ_i for the steady-state revenue weight of matched item

i, $p_{i,t}$ for its price and p_i for its steady-state price, the deflated measure is

$$\ln \mathrm{TFPR}_t^P(z) = \ln \mathrm{TFPR}_t(z) - \sum_i \varpi_i \big(\ln p_{i,t} - \ln p_i \big). \tag{C.11}$$

Under CES every producer price is A times marginal cost, so the index moves with w_t/Z_t and the deflated revenue productivity of a firm that does not export moves only with technology. Under translog the index also moves with the revenue-weighted average change in markups, so a firm's deflated revenue productivity moves with its own markup relative to that average. I also compute the measure in units of consumption, undeflated, which the replication package reports.

The second choice is the average across firms. Most firm-level regressions in this literature weight every firm equally, so the baseline is the unweighted mean of log revenue productivity over the home firms that produce, taken over the Pareto distribution of draws. I also compute the means weighted by each firm's revenue and by its total labor. Following Olley and Pakes (1996), I split the revenue-weighted mean into the simple mean and a term that is larger when the more productive firms have larger shares of revenue. They use productivity in levels and shares of output; I use the log of revenue productivity and shares of revenue. That term measures how much of the weighted mean comes from larger firms recording higher revenue productivity. Under CES it is negative, because the larger firms are exporters and record less; under translog it is positive, because larger firms charge higher markups. In the one-sector steady state it is -0.3 log points under CES and 11.6 under translog. The ratio of total revenue to total labor across producers is TFP_t^C of Section C.1.3, which checks the firm-level computation against the aggregate one.

The third choice applies to the three-sector model. The evidence compares more and less exposed producers, so for each measure I report the change in the most exposed sector minus the change in the least exposed, with each sector's revenue deflated by its own sector index built as in (C.11), in 2018 and 2019 relative to the steady state, with the tariffs alone. With the technology decline that matches the estimated gap, the revenue-weighted gap in deflated revenue productivity, -5.03 percent in 2019, is close to the gap in measured TFP. For my own estimate I use the 2019 value, two years after the announcement, which matches its horizon; for the other studies the rules below read the model at each study's own window.

In Table C.5, I set the firm-level evidence on revenue productivity beside the model. Three rules make the comparison like for like. First, I read the model at each study's own window and in the form the study estimates: the effect a given number of years after the tariff, a difference over a given number of years, and the mean over the first nine years after the duty for Pierce (2011). Second, the common unit is percent per percentage point of tariff,

which for the model is the gap between the most and least exposed sectors over the gap in their tariffs; where a study's exposure has no size in points, an antidumping dummy or a trade orientation, I compare only the sign and the time profile. Third, for the liberalization studies I flip the sign of the model's response to a tariff increase, which treats the model as linear in a change of this size.

In the model, the tariffs alone raise the revenue productivity of the protected producers by about a tenth of a percent per point in the unweighted mean, an effect that barely changes with the horizon; in the revenue-weighted mean the effect is about as large in the first year and close to zero after it. The sign matches the antidumping evidence of Konings and Vandenbussche (2008). By contrast, the liberalization studies find that revenue productivity rises after tariffs fall. In Amiti and Konings (2007) the rise is about the same over two and five years and much larger over nine years, while the model's unweighted mean predicts a small decline. The model's response is also small next to my own estimate. Per point of tariff, the measured TFP of Section 3 falls by more than fifty times as much as measured TFP rises in the model. On revenue productivity, as on measured TFP, the reallocation the model generates on its own falls far short of the evidence.

C.5 One-sector models with identical firms

Section 4.5 compares the optimal tariff of the one-sector models with that of versions in which firms do not differ in productivity, the setting of the Armington model of Itskhoki and Mukhin (2026). Two changes make firms identical, and everything else is kept. First, every firm draws the same productivity $z_{\min}$, the limit of the Pareto distribution as $k \to \infty$. Second, the fixed export cost is zero. With identical firms and a positive fixed cost, either every firm exports or none does, and the case in which every firm exports is the one closest to an Armington economy, whose whole output is traded. Households, their assets and the external position, entry and exit, the iceberg cost and the tariff are those of the one-sector models. Physical productivity is $Z_t z_{\min}$, so a tariff cannot move it. One difference from Itskhoki and Mukhin remains, because the number of varieties is endogenous through entry.

Under CES, every incumbent produces and exports, so the masses of producers and exporters both equal the mass of incumbents N_t^{in}, and the cutoff condition and the share of exporters drop out. The home price index becomes $N_t^{\text{in}} \rho_{D,t}^{1-\theta} + N_t^{\text{in}*} \big((1+t_t)\rho_{X,t}^*\big)^{1-\theta} = 1$, with $\rho_{D,t} = \frac{\theta}{\theta-1}\frac{w_t}{Z_t z_{\min}}$ and $\rho_{X,t}^*$ the foreign firms' price at home before the tariff, and profit per firm is the sum of profit from the two markets, each a share $1/\theta$ of revenue. The steady state has a closed form. The import share of spending is $\tau^{1-\theta}/(1+\tau^{1-\theta})$, so the iceberg cost that matches an import share Λ^M is $\tau = \big(\Lambda^M/(1-\Lambda^M)\big)^{1/(1-\theta)}$.

Table C.5: Revenue productivity: the evidence and the model at each study's window

Study	Sample	Exposure	Window	Estimate	Model, same window	
					unweighted	revenue-weighted
Pierce (2011)	plants, quantity sample	antidumping duty, per point	years 1–9 after the duty, pooled	−0.05	+0.088	+0.008
Konings and Vandenbussche (2008)	firms	antidumping (dummy)†	1 year after the duty	+4.2***	+0.112	+0.094
			2 years after the duty	+4.9***	+0.082	+0.003
			4 years after the duty	+3.9***	+0.084	−0.004
Amiti and Konings (2007)	plants	output-tariff cut, per point	2-year difference	+0.01 to +0.14***	−0.082	−0.003
			5-year difference	+0.09* to +0.20***	−0.085	+0.004
			9-year difference	+0.30*** to +0.39***	−0.086	+0.002
Pavcnik (2002)	plants	import-competing after liberalization†	1982 (3 years after)	+3.0*	−0.082	+0.005
			1985 (6 years after)	+10.4**	−0.086	+0.003
This paper	industries (BLS TFP)	statutory tariff, per point	2017–19	−0.39***	+0.002 (measured TFP)	

Notes: Estimates are the studies' own, in percent, restated per percentage point of tariff where the exposure is measured in points (the record and each transformation are in the replication package). Pierce: revenue TFP in the plants that report quantities, per point of effective antidumping duty, not significant. Amiti and Konings: Table 9, the range from the output tariff alone to the output tariff with input tariffs, lower value first; an end without stars is not significant. †Exposure without a size in points (an antidumping dummy; import-competing relative to nontraded plants): sign and time profile only. Pavcnik: Chile lowered its tariffs to a uniform 10 percent by 1979 and raised them temporarily after the 1982–83 recession, to 35 percent in 1984 and 20 percent from mid-1985 (her pp. 246–247), so her 1985 estimate follows a partial reversal of the liberalization; she judges that reversal very small next to the reform of the late 1970s. Model: the three-sector translog model with the tariffs alone, revenue productivity deflated by the sector's matched-item producer price index (Section C.4), the gap between the most and least exposed sectors over their tariff gap at the same window; for the liberalization studies the sign is flipped to give the effect of a cut. For this paper the model column is measured TFP as BLS measures it, and the stars are from the t-statistic clustered by three-digit sector of Table 2. ***, **, *: significant at 1, 5 and 10 percent in the study.

Under translog, each market has two prices, the home firms' and the foreign exporters', and the integrals over the productivity distribution collapse to sums over these two. The share of each firm solves equations (11) and (12) together, with its consumer price equal to its markup times its marginal cost, $w_t/(Z_t z_{\min})$ for a home firm at home and $(1 + t_t)\tau^* Q_t w_t^*/(Z_t^* z_{\min})$ for a foreign firm, the varieties sold at home are the home and foreign incumbents, $N_t^{\text{in}} + N_t^{\text{in}*}$, and the moments of log prices and the price index follow from the two prices. A tariff still lowers the foreign exporters' share at home and with it their markup, so its pass-through to home prices remains incomplete; what the identical-firm model removes is the reallocation of production across firms of different productivity.

Each identical-firm model is calibrated to the moments its heterogeneous-firm counterpart matches. Under CES, θ is kept and the iceberg cost, $\tau = 1.434$, matches the import share of the published model, 0.267. Under translog, the iceberg cost and σ jointly match the import share of the heterogeneous-firm translog model, 0.246, and its expenditure-weighted markup, A, which gives $\tau = 1.672$ and $\sigma = 0.229$. The value of σ is lower than in the baseline because with identical firms markups differ only between home and foreign firms, so the large, high-markup firms that raise the average markup in the baseline are absent. Holding σ at its heterogeneous-firm value instead, and matching only the import share, raises the trade elasticity of the identical-firm model, the response of the ratio of imported to domestic spending to the iceberg cost, to 5.06, against 3.25 with heterogeneous firms, and its elasticity of substitution between domestic and foreign goods to 8.81, and lowers its optimal tariff to 26 percent, below the heterogeneous-firm 39. The comparison of optima across firm structures therefore depends on how the identical-firm benchmark is calibrated, because the benchmark also changes how fast welfare falls at large tariffs; the comparison of the gain per point at a zero tariff does not. For the external-position rows of Table 6, I use the translog model with three identical sectors, which reproduces the one-sector model exactly. Without positions, its optimal tariff equals the one-sector one.

The elasticity of substitution between domestic and foreign goods, the parameter that governs the optimal tariff in Itskhoki and Mukhin (2026), is a property of preferences, and I compute it at the steady state of each model as the response of relative spending on foreign and domestic goods to their relative consumer price, holding the set of firms and every other price fixed. Under CES all varieties enter one aggregator, so it is $\theta = 3.8$ with heterogeneous and with identical firms. Under translog, raising the log consumer price of every foreign variety sold at home by x raises $\bar{L}$ by $(N_F/N)x$, so each home variety's share rises by $\sigma(N_F/N)x$ and each foreign variety's falls by $\sigma(N_H/N)x$, where N_H and N_F are the home and foreign varieties sold at home and $N = N_H + N_F$. With S_H and $S_F = 1 - S_H$ the

spending shares of the two groups,

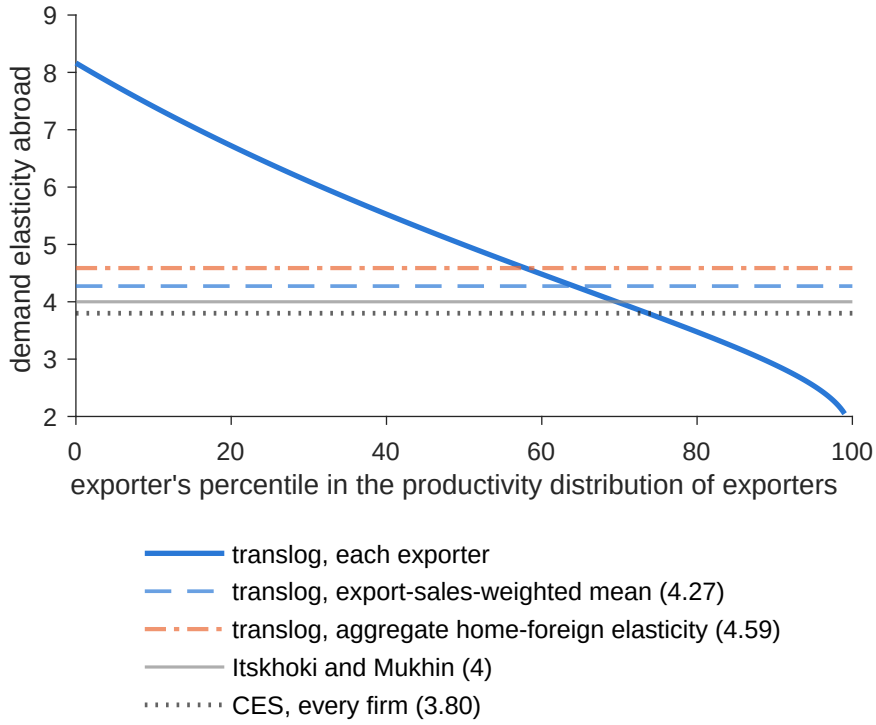
$$\varepsilon_{HF} = 1 - \frac{d\ln(S_F/S_H)}{d\ln(p_F/p_H)} = 1 + \frac{\sigma N_H N_F}{N\, S_H S_F}, \tag{C.12}$$

which in the symmetric case, $N_H = N_F$ and $S_H = S_F$, is the $1 + \sigma N$ of Bilbiie et al. (2012). In the heterogeneous-firm translog model it is 4.59, against 4 in Itskhoki and Mukhin (2026).

The aggregate elasticity hides a wide dispersion across firms, which Figure C.2 shows. With a continuum of varieties, the elasticity of substitution between one exporter's variety and the goods it competes with in the foreign market is the demand elasticity the exporter faces, $1+\sigma/s_X(z)$, the firm-level counterpart of the single elasticity between home and foreign goods in Itskhoki and Mukhin (2026). Under CES it is θ for every firm. Under translog it falls with market share. In the steady state it is 8.16 for the exporter at the export cutoff, 4.99 for the median exporter and 2.90 at the 90th percentile of exporters. Weighted by export sales, which put most of the weight on the large exporters, the mean is 4.27, close to the value of 4 in Itskhoki and Mukhin (2026). The translog calibration therefore matches their elasticity on average while building in the dispersion across firms that makes markups, and with them measured TFP, respond to the reallocation of production.

The difference between CES and translog in Table 6 could reflect the higher elasticity of substitution between domestic and foreign goods that the translog calibration implies, rather than the distribution of markups across firms. To separate the two, I recalibrate the translog model so that ε_{HF} equals the CES value of 3.80, keeping the import share and the share of exporting firms at their baseline values. This requires $\sigma = 0.421$, $\tau = 1.237$ and $f_X = 0.0181$, and it raises the expenditure-weighted markup to 1.395. The welfare results move little. The gain per point at a zero tariff rises from 0.124 to 0.129, which closes about a third of the gap with CES (Table 6); the optimal tariff is 38.6 percent, against 39.2 in the baseline; and a nearly permanent one percent tariff raises home welfare by 0.126 percent of consumption, about as much as in the baseline (Table 5). The productivity mechanics change more. On impact, physical productivity changes by -0.124 percent per point, against -0.138 in the baseline, with the intensive margin at -0.139 and the extensive margin of exporting at 0.015, and TFP at producer prices rises by 0.036 percent. The fall in physical productivity stays larger than under CES, -0.0814 percent, so the conclusion that translog demand lets reallocation lower productivity more does not rest on the higher elasticity.

Figure C.2: The demand elasticity faced by home exporters in the foreign market, one-sector translog model, steady state



Notes: The solid line is $1 + \sigma/s_X(z)$, the demand elasticity faced by a home exporter of productivity z in the foreign market, against the exporter's percentile in the productivity distribution of exporters (Pareto above the export cutoff). Reference lines: its export-sales-weighted mean; the aggregate elasticity of substitution between home and foreign goods, equation (C.12); the CES value, $\theta = 3.8$ for every firm; and the value of 4 in Itskhoki and Mukhin (2026). Baseline calibration of Section 4.3.

References

Rodrigo Adão, John Sturm Becko, Arnaud Costinot, and Dave Donaldson. Why is trade not free? A revealed preference approach. *Econometrica*, 2026. Forthcoming.

Mary Amiti and Jozef Konings. Trade liberalization, intermediate inputs, and productivity: Evidence from Indonesia. *American Economic Review*, 97(5):1611–1638, 2007.

Susanto Basu, John G. Fernald, and Miles S. Kimball. Are technology improvements contractionary? *American Economic Review*, 96(5):1418–1448, 2006.

Randy A. Becker, Wayne B. Gray, and Jordan Marvakov. [dataset] NBER-CES manufacturing industry database. National Bureau of Economic Research, 2021. 1958–2018, version 2021a, NAICS 2012 version (`nberces5818v1_n2012`), https://www.nber.org/research/data/nber-ces-manufacturing-industry-database.

Paul R. Bergin and Robert C. Feenstra. Staggered price setting, translog preferences, and endogenous persistence. *Journal of Monetary Economics*, 45(3):657–680, 2000.

Andrew B. Bernard, J. Bradford Jensen, and Peter K. Schott. Trade costs, firms and productivity. *Journal of Monetary Economics*, 53(5):917–937, 2006.

Andrew B. Bernard, J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott. Firms in international trade. *Journal of Economic Perspectives*, 21(3):105–130, 2007.

Florin O. Bilbiie, Fabio Ghironi, and Marc J. Melitz. Endogenous entry, product variety, and business cycles. *Journal of Political Economy*, 120(2):304–345, 2012.

Fabian Eckert, Teresa C. Fort, Peter K. Schott, and Natalie J. Yang. Imputing missing values in the US Census Bureau's county business patterns. 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. [dataset] Replication data for: The return to protectionism. Harvard Dataverse, V1, 2019. Accessed June 2026. https://doi.org/10.7910/DVN/KSOVSE.

Robert C. Feenstra. A homothetic utility function for monopolistic competition models, without constant price elasticity. *Economics Letters*, 78(1):79–86, 2003.

Robert C. Feenstra, John Romalis, and Peter K. Schott. U.S. imports, exports, and tariff data, 1989–2001. 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. [dataset] Replication data for: Disentangling the effects of the 2018–2019 tariffs on a globally connected U.S. manufacturing sector. Harvard Dataverse, https://doi.org/10.7910/DVN/NUDHEA, 2024. Accessed October 2026.

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*, 108 (5):1289–1303, 2026.

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.

Lucia Foster, John Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1): 394–425, 2008.

Fabio Ghironi and Marc J. Melitz. International trade and macroeconomic dynamics with heterogeneous firms. *Quarterly Journal of Economics*, 120(3):865–915, 2005.

Oleg Itskhoki and Dmitry Mukhin. The optimal macro tariff. Working paper, Harvard University and London School of Economics, March 2026.

Jozef Konings and Hylke Vandenbussche. Heterogeneous responses of firms to trade protection. *Journal of International Economics*, 76(2):371–383, 2008.

Office of the United States Trade Representative. Notice of action and request for public comment concerning proposed determination of action pursuant to section 301: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 83 Fed. Reg. 28710 (June 20, 2018), Annexes A and B, 2018a. Product list for Section 301 List 1, 818 eight-digit subheadings. https://www.govinfo.gov/content/pkg/FR-2018-06-20/pdf/2018-13248.pdf.

Office of the United States Trade Representative. Notice of action pursuant to section 301: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 83 Fed. Reg. 40823 (August 16, 2018), Annex A, 2018b. Product list for Section 301 List 2, 279 eight-digit subheadings. https://www.govinfo.gov/content/pkg/FR-2018-08-16/pdf/2018-17709.pdf.

Office of the United States Trade Representative. Notice of modification of section 301 action: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 83 Fed. Reg. 47974 (September 21, 2018), Annex A, notes 20(f) and 20(g), 2018c. Product list for Section 301 List 3, 5,745 full and partial tariff lines. https://www.govinfo.gov/content/pkg/FR-2018-09-21/pdf/2018-20610.pdf.

Office of the United States Trade Representative. Notice of modification of section 301 action: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 84 Fed. Reg. 43304 (August 20, 2019), Annex A, 2019a. Product list for Section 301 List 4A. https://www.govinfo.gov/content/pkg/FR-2019-08-20/pdf/2019-17865.pdf.

Office of the United States Trade Representative. Notice of modification of section 301 action: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 84 Fed. Reg. 45821 (August 30, 2019), 2019b. Raises the additional duty on List 4A from 10 to 15 percent from September 1, 2019.

Office of the United States Trade Representative. Notice of modification of section 301 action: China's acts, policies, and practices related to technology transfer, intellectual property, and innovation. 85 Fed. Reg. 3741 (January 22, 2020), 2020. Reduces the additional duty on List 4A to 7.5 percent from February 14, 2020. https://www.govinfo.gov/content/pkg/FR-2020-01-22/pdf/2020-00904.pdf.

G. Steven Olley and Ariel Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297, 1996.

Nina Pavcnik. Trade liberalization, exit, and productivity improvements: Evidence from Chilean plants. *Review of Economic Studies*, 69(1):245–276, 2002.

Justin R. Pierce. Plant-level responses to antidumping duties: Evidence from U.S. manufacturers. *Journal of International Economics*, 85(2):222–233, 2011.

Justin R. Pierce and Peter K. Schott. The surprisingly swift decline of US manufacturing employment. *American Economic Review*, 106(7):1632–1662, 2016.

José Antonio Rodríguez-López. Prices and exchange rates: A theory of disconnect. *Review of Economic Studies*, 78(3):1135–1177, 2011.

Peter K. Schott. [dataset] U.S. HS-level exports, annual detail files. Data, version v2018.10.15, https://faculty.som.yale.edu/peterschott/international-trade-data/, 2018. File `exp_detl_yearly_116n.dta`, domestic exports identified by the domestic/foreign code. Accessed October 2026.

Peter K. Schott. [dataset] U.S. HS-level imports and exports, 1989–2024. Data, version v2025.4.21, https://sompks4.github.io/sub_data.html, 2025. Annual import-detail files `imp_detl_<YYYY>_12n.dta`, converted from U.S. Census Bureau records. Accessed June 2026.

The White House. Presidential proclamations 9693 (solar cells) and 9694 (large residential washers): To facilitate positive adjustment to competition from imports. Federal Register 83 FR 3541 and 83 FR 3553 (January 25, 2018), 2018a. https://www.govinfo.gov/content/pkg/FR-2018-01-25/pdf/2018-01592.pdf, https://www.govinfo.gov/content/pkg/FR-2018-01-25/pdf/2018-01604.pdf.

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%), 2018b.

U.S. Bureau of Economic Analysis. [dataset] International investment position, table 1.1. U.S. net international investment position at the end of the period. Interactive Data Application, release of September 24, 2026, 2026a. https://apps.bea.gov/iTable/, accessed October 25, 2026.

U.S. Bureau of Economic Analysis. [dataset] National income and product accounts, table 1.1.5. gross domestic product. Interactive Data Application, last revised September 30, 2026, 2026b. https://apps.bea.gov/iTable/, accessed October 25, 2026.

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

U.S. Bureau of Labor Statistics. Industry productivity measures. Handbook of Methods, 2026a. https://www.bls.gov/opub/hom/inp/calculation.htm.

U.S. Bureau of Labor Statistics. Producer price indexes: Calculation. Handbook of Methods, 2026b. https://www.bls.gov/opub/hom/ppi/calculation.htm.

U.S. Bureau of Labor Statistics. [dataset] Total factor productivity for manufacturing and transportation detailed industries. Office of Productivity and Technology,

detailed-industry tables, 2026c. Data released August 26, 2026. https://www.bls.gov/productivity/tables/.

U.S. Census Bureau. [dataset] statistics of U.S. businesses, 2016: U.S., 6-digit NAICS. Annual data tables, released December 18, 2018, 2018. https://www2.census.gov/programs-surveys/susb/tables/2016/us_6digitnaics_2016.xlsx, accessed October 27, 2026.

U.S. Census Bureau. [dataset] a profile of U.S. importing and exporting companies, 2016–2017. Release CB19-143, September 26, 2019, 2019. https://www.census.gov/foreign-trade/Press-Release/edb/edbrel2017.pdf, accessed October 27, 2026.

U.S. Census Bureau. [dataset] a profile of U.S. importing and exporting companies, 2017–2018. Release CB20-55, April 2, 2020, 2020. https://www.census.gov/foreign-trade/Press-Release/edb/edbrel2018.pdf, accessed October 27, 2026.

U.S. Census Bureau. [dataset] statistics of U.S. businesses, 2017: U.S. and states, 6-digit NAICS. Annual data tables, disclosure approval CBDRB-FY21-188, 2021. https://www2.census.gov/programs-surveys/susb/tables/2017/us_state_6digitnaics_2017.xlsx, accessed October 27, 2026.

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