

Where Do Tariff Shocks Propagate? Evidence on Production, Prices, and Shipments from Section 301 China Tariffs*

Alfredo Sosa

2026-07-15

1 Abstract

Section 301 China tariffs matter because they represent a major recent intervention in U.S. manufacturing trade policy, with potential implications for domestic production, prices, and shipments. According to textbook trade theory, tariffs are intended to support the domestic economy by raising the relative cost of foreign goods and services, thereby protecting import-competing domestic producers from foreign competition. This paper uses the 2018-2019 Section 301 Tariffs on Chinese imports as a source of exogenous variation to estimate the effects of these tariffs on production, prices, and shipments in the US economy. The paper uses a triple-difference design at the NAICS-4¹ level, with world import penetration as the control dimension, and analyzes three outcome panels: BLS-PPI, FRED-IP, and Census-M3.²

*Independent Economist. Correspondence: alfredosos@gmail.com. The author declares no competing financial interests related to this research. All errors are the author's own.

¹NAICS-4 is the four-digit North American Industry Classification System industry-group level, which groups establishments into detailed industries based on similar production processes.

²BLS-PPI is the producer price index obtained through the [BLS API](#); FRED-IP is the industrial production index obtained through the [FRED API](#); and Census-M3 is nominal manufacturers' shipments obtained from the Census Bureau's [Manufacturers' Shipments, Inventories, and Orders program](#).

The estimation using the FRED-IP panel indicates that tariff exposure reduced manufacturing activity: $\beta = -0.783$ and $\text{boot-p} = 0.005$, with an attenuation ratio of 1.016. Practically, at the mean exposure level among treated industries, the estimated effect corresponds to an approximate 17.7 percent reduction in industrial production after the tariff shock (see Appendix A7).

Methodologically, I use a three-difference-in-differences approach to compare high-exposure vs low-exposure industries before and after the tariff shocks; the third difference corresponds to the level of penetration of non-Chinese goods to account for pre-trends across a family of commodities. When the effect of tariffs on prices was examined, the results indicated a positive but imprecise effect, which may reflect insufficient variation at the industry level; nevertheless, the clustered standard errors detected a statistically significant relationship, though this relationship became imprecise under bootstrap inference. In the primary specification, Census-M3 shows a negative effect on nominal shipments (attenuation ratio of -0.961), although this estimate is imprecisely estimated.

The paper contributes evidence on how tariff shocks propagate across production, prices, and shipments within manufacturing industries. In doing so, it complements research on the price and welfare incidence of the 2018–2019 tariffs by adding production-side evidence at the NAICS-4 level, contributes to the broader literature on how tariff and import-competition shocks propagate through domestic manufacturing, and offers a triple-difference design with a continuous world-import-penetration control as an alternative to more conventional binary treatment definitions.

2 Introduction

This period of comparative policy stability sits within a broader wave of U.S. tariff actions that began in 2025 and whose legal footing has shifted substantially since. In February 2026, the U.S. Supreme Court held, by a 6-3 vote, in *Learning Resources, Inc. v. Trump* (consolidated with *Trump v. V.O.S. Selections, Inc.*), that the International Emergency Economic Powers Act (IEEPA) does not authorize the President to impose tariffs, invalidating the broad country-level tariffs previously imposed under that authority. The administration responded the same day by invoking Section 122 of the Trade Act of 1974 to impose a temporary global import surcharge, initially 10 percent and subsequently raised to 15 percent. That authority carries a hard 150-day statutory limit; the Section 122 surcharge expired by operation of law on July 24, 2026, and was immediately followed by new Section 301 duties targeting a number of trading partners under a separate forced-labor enforcement investigation.³ This continued volatility in the legal basis for tariff policy underscores the value of studying the 2018 Section 301 episode, whose statutory authority was never successfully challenged and whose timing is unambiguous.

The 2018 Section 301 tariffs on China, by contrast, offer a comparatively stable and well-documented policy shock: a discrete, product-level tariff schedule with clear implementation dates and a tranche structure that has remained largely in place for several years, making it better suited to the identification strategy developed in this paper. They covered approximately \$350 billion⁴ in goods across multiple tranches and arrived after a long period in

³*Learning Resources, Inc. v. Trump*, 607 U.S. ____ (2026) (consolidated with *Trump v. V.O.S. Selections, Inc.*, No. 24-1287), decided February 20, 2026. Proclamation invoking Section 122 of the Trade Act of 1974 issued the same day; Section 122 surcharge in effect February 24–July 24, 2026 under its 150-day statutory ceiling.

⁴Tariff coverage and tranche information are sourced from the U.S. International Trade Commission’s

which U.S. manufacturing firms had become deeply integrated with imported intermediate inputs and global supply chains. Because the tariff schedules were defined at the product level, U.S. manufacturing industries were differentially exposed depending on their product mix and their reliance on tariffed inputs. That feature makes the policy useful for studying how a common national trade shock translated into uneven industry-level pressure.

The central policy question is whether tariff-induced import costs resulted in significant changes in domestic industrial production, producer prices, and nominal shipments. That question matters not only for the incidence of trade policy, but also for manufacturing employment, industrial output, and price stability. If tariffs mainly raise costs without stimulating real domestic activity, their effects can appear as a production-side drag rather than as a straightforward boost to protected industries. Existing studies show that the trade war affected prices, welfare, and globally connected manufacturing sectors, but the industry-level production consequences remain an important empirical margin (Amiti et al. 2019; Fajgelbaum et al. 2020; Flaaen and Pierce 2019).

Estimating those consequences is difficult because tariff exposure was not randomly assigned across industries. A naive difference-in-differences comparison can confound tariff-driven supply shocks with secular trends in import-competing industries, especially if more exposed industries were already following different paths before the policy shock. Industries facing Chinese import competition may also differ in technology, demand conditions, and supply-chain organization, so a simple exposed-versus-less-exposed comparison can attribute too much to tariffs. This paper addresses that identification challenge with a triple-difference design

[USITC DataWeb](#), the project source used for Section 301 tariff and import data.

at the NAICS-4 level. The design uses world, non-China import penetration as the control dimension, separating tariff exposure from broader import-competition trends that may have affected manufacturing outcomes independently of the Section 301 treatment. This distinction is important because the exposure measure is negatively correlated with world import penetration: $\text{corr}(E_i, \text{world_pen}) = -0.315$ for FRED-IP, $\text{corr}(E_i, \text{world_pen}) = -0.420$ for BLS-PPI, and $\text{corr}(E_i, \text{world_pen}) = -0.388$ for Census-M3. These correlations indicate that the most tariff-exposed industries were not simply the industries with the highest general import penetration, which supports the use of a continuous DDD control dimension. The identifying comparison therefore asks whether industries with higher tariff exposure changed differentially after the shock, conditional on the import-penetration path that might otherwise proxy for longer-run competitive pressure. As an additional robustness exercise, the paper also uses a Bartik or shift-share instrument, following the logic that such designs require transparent interpretation of exposure shares and shocks (Goldsmith-Pinkham et al. 2020). The Bartik specification has a first-stage $F = 1,981.6$, so the large first-stage F-statistic reduces concerns about weak first-stage relevance in this specification. Such a large first-stage statistic is expected because construction of the Bartik instrument overlaps substantially with the definition of the exposure variable itself.

The analysis combines three monthly NAICS-4 outcome panels over approximately 2015–2023. The first is BLS-PPI, which measures the log producer price index and includes 77 clusters. The second is FRED-IP, which measures the log industrial production index and includes 31 clusters. FRED-IP is the primary anchor outcome because it is the most direct measure of real manufacturing activity and has the strongest signal-to-noise for the production-side

question. The third is Census-M3, which measures log nominal shipments and includes 85 clusters. Studying these panels together makes it possible to distinguish movements in real activity from movements in producer prices and nominal shipments, which need not respond in the same direction after a tariff shock. That distinction is central because a tariff can raise input costs, change output quantities, and affect nominal revenues at the same time.

The results suggest that tariff exposure reduced real manufacturing activity in the FRED-IP panel. In the primary FRED-IP D1b specification, the estimated coefficient is -0.783 , with $\text{boot-p} = 0.005$ and a boot 95% CI of $[-1.387, -0.176]$. The FRED-IP attenuation ratio is 1.016 , indicating that the DDD control does not absorb the estimated treatment effect. This estimate implies a substantial production response: greater tariff exposure after the policy change is associated with lower real manufacturing activity even after accounting for world import penetration. The BLS-PPI evidence is suggestive but inconclusive: its attenuation ratio is 0.694 and its boot-p is 0.301 . Suppliers appear to have increased prices, but the industry panel may lack the granularity needed to measure that response with sufficient precision. Clustered standard errors indicate a statistically significant relationship between exposure and pricing, but the relationship becomes imprecise under bootstrap inference. In the primary specification, Census-M3 shows a negative effect on nominal shipments, with a coefficient of -0.101 and an attenuation ratio of -0.961 , although this estimate is imprecisely estimated.

The Bartik IV estimate is similar in sign and magnitude, with a coefficient estimate of -0.8685 and first-stage $F = 1,981.6$. Its agreement with the primary estimate strengthens the paper by applying the transparent exposure-share interpretation emphasized in the Bartik

and shift-share literature (Goldsmith-Pinkham et al. 2020). However, the high alignment between the DDD and the IV estimates stems partly from the fact that both measures overlap in their definition. The preferred exposure specification, `exp_dtariff_tv`, is a continuous, time-varying measure built directly from realized NAICS-4 tariff tranches and import shares. The Bartik instrument is instead constructed as a shift-share: a fixed, pre-period industry exposure share interacted with the aggregate national timing of tariff implementation, so that only the timing component is treated as exogenous variation (Goldsmith-Pinkham et al. 2020). Because both measures draw on the same underlying tariff-tranche and import-share data, their close numerical agreement partly reflects shared construction rather than fully independent identification, which is why the Bartik estimate is presented as a corroborating check (Appendix A6) rather than as the primary identification strategy.

Back-of-the-envelope calculations imply an annual loss of $-\$51.9$ billion, a cumulative loss of $-\$311.4$ billion, and implied jobs of -168 thousand. These values are illustrative, not structural, and should be read as scale calculations rather than a complete welfare accounting. Appendix A7 provides the underlying scale assumptions and derivation of these figures.

The paper contributes to three related literatures. First, it adds causal evidence on the production-side effects of tariff exposure at NAICS-4 granularity, complementing work that focuses on price incidence and welfare. Amiti, Redding, and Weinstein (2019) and Fajgelbaum, Goldberg, Kennedy, and Khandelwal (2020) provide influential estimates of how the 2018 tariffs passed through to import prices and consumer welfare, drawing primarily on trade-flow and price data. This paper complements that work by studying industry-level production, producer-price, and shipments responses directly, rather than inferring production-side

consequences from trade flows and prices alone.

Second, it contributes to research on trade and domestic production by showing how tariff shocks propagate across real output, producer prices, and nominal shipments within manufacturing industries. This joint view helps explain why evidence can look stronger in real activity than in nominal shipments. It also keeps the policy interpretation tied to observable industry outcomes. Flaaen and Pierce (2019) find that U.S. manufacturing industries more exposed to the 2018–2019 tariffs experienced relative declines in employment, as protection gains were offset by higher input costs and retaliation, while Flaaen, Hortaçsu, and Tintelnot (2020) document similar cost pass-through and production-relocation effects in the washing machine industry specifically. This paper differs by estimating a triple-difference effect on real industrial production, producer prices, and nominal shipments jointly, using a continuous world-import-penetration control to separate tariff-driven variation from the slower-moving import-competition trends studied in the longer-run China-shock literature (e.g., Autor et al. 2016).

Third, it contributes methodologically by using a continuous triple-difference control dimension as an alternative to more binary instrument approaches. The design is not meant to eliminate every concern about tariff assignment, but it provides a disciplined comparison that separates targeted tariff exposure from broader import-penetration dynamics and clarifies where the strongest evidence lies. The resulting interpretation is deliberately cautious: the strongest signal is in the primary real-activity panel, while the price and shipment panels bound how broadly that signal should be generalized.

3 US-China Trade

China is one of the largest single-country trading partners of the United States, and the bilateral relationship is a central reference point for any study of U.S. trade policy. U.S. goods-only trade with China totaled an estimated \$414.7 billion in 2025, consisting of \$106.3 billion in U.S. exports to China and \$308.4 billion in U.S. imports from China, and produced a goods trade deficit of \$202.1 billion (Office of the United States Trade Representative 2026). On a broader goods-and-services basis, the U.S. Bureau of Economic Analysis ranks China among the United States' largest trading partners. The scale of this relationship means that changes in the tariff treatment of Chinese imports have the potential to affect a wide range of U.S. manufacturing industries that rely on Chinese inputs, compete with Chinese products, or sell into the Chinese market.

China's importance to the U.S. economy is not limited to the size of the trade flows themselves. Chinese-made components and materials are embedded in many U.S. manufacturing supply chains, so tariffs on Chinese imports can raise input costs for domestic producers even in industries that do not compete directly with Chinese output. At the same time, China has historically been an important export destination for some U.S. industries, so retaliatory tariffs imposed by China in response to Section 301 actions can reduce demand for U.S. goods abroad. This combination of import-cost exposure and export-retaliation exposure is part of the motivation for measuring tariff exposure broadly, at the industry level, rather than focusing only on industries whose final output appeared on a tariff list, and it is a central reason why the production-side effects of the 2018–2019 tariffs are an important empirical question independent of their effect on prices paid by consumers.

Identifying precisely which industries were most affected by the 2018–2019 tariffs would require reporting NAICS-4-level point estimates rather than the pooled DDD coefficients that are the focus of this paper, and the current estimation output does not support that breakdown without additional analysis. Descriptively, the institutional record summarized in the Institutional Background section indicates that the four Section 301 tariff tranches targeted machinery, electrical equipment and components, and, in later tranches, a broader set of consumer and industrial goods; automobile manufacturers were less directly named in the original tranches than machinery and electronics producers, though vehicle assembly relies heavily on globally sourced parts and so remains exposed through input-cost pass-through. A full industry-by-industry decomposition of the estimated production, price, and shipments effects would be a natural extension for a future appendix, but it is not part of the current estimation output and should not be asserted without re-running the specification at the individual NAICS-4 level.

4 Literature Review

The first strand is the recent literature on the Section 301 tariffs, which emphasizes that the 2018 trade war was not a simple transfer from foreign producers to domestic firms. Studies of the tariff episode show substantial pass-through into U.S. prices, changes in sourcing and trade flows, and welfare costs borne by U.S. consumers and firms (Amiti et al. 2019; Fajgelbaum et al. 2020). Work on globally connected manufacturing further stresses that tariffs can affect domestic producers through imported intermediate inputs, retaliation, and disrupted supply chains, rather than only through protection from final-good import competition (Handley

et al. 2020; Flaaen and Pierce 2019). This paper builds on that literature by shifting the outcome focus toward real industrial production at the NAICS-4 level. The question is not only whether tariffs raised prices or changed trade flows, but whether tariff exposure translated into measurable changes in domestic manufacturing activity, producer prices, and nominal shipments.

The closest tariff-specific evidence is especially useful for framing the empirical problem. Flaaen and Pierce study the 2018-2019 tariffs through an input-output exposure logic, emphasizing how globally connected industries can face higher costs even when policy is intended to protect domestic production (Flaaen and Pierce 2019). This paper is complementary but distinct. Rather than organizing the design around input-output exposure alone, it uses a triple-difference framework organized around tariff exposure, production outcomes, and time-varying post-2018 shocks. That design helps separate the tariff treatment from broader import-competition trends and asks whether more exposed industries changed differentially after the policy shock. The approach therefore connects the tariff-incidence literature to a production-side question that is closer to industrial adjustment than to price pass-through alone.

A second strand of work studies how trade exposure affects domestic manufacturing more broadly. The China shock literature shows that import competition can have persistent effects on labor markets and manufacturing activity, and related work links changes in trade policy regimes to sharp employment and production adjustments (Autor et al. 2013; J. Pierce and Schott 2016). More recent evidence on supply-chain shocks shows that input disruptions can propagate through domestic production networks, affecting firms and industries that are not

directly targeted by the original shock (Boehm et al. 2019). These studies motivate treating manufacturing exposure as multidimensional: industries differ not only in final demand conditions, but also in input use, import dependence, and their position in production networks.

This paper is complementary to that broader trade literature because it studies tariff-driven variation rather than long-run import penetration levels. Long-run import exposure can be difficult to separate from technological change, sectoral demand shifts, and preexisting differences across industries. A tariff shock creates a more discrete policy event, but exposure to that event is still correlated with industry characteristics that may influence outcomes independently. The DDD design addresses this problem by comparing tariff-exposed industries to observably less-exposed sectors while preserving the post-2018 timing variation. The design does not make exposure random, and the interpretation remains cautious, but it reduces a key endogeneity concern: that industries with high tariff exposure may have been on systematically different trajectories even before the tariffs took effect.

A third strand is the methodological literature on difference-in-differences, which also motivates the paper’s empirical choices. Recent work has shown that conventional two-way fixed effects estimators can be difficult to interpret when treatment effects are heterogeneous or when treatment timing varies across units (Sun and Abraham 2021). Related DID estimators clarify how to form comparisons that are more robust to staggered adoption and dynamic treatment effects (Callaway and Sant’Anna 2021). Although the Section 301 setting is not a textbook staggered-adoption design, the same concerns are relevant because treatment intensity varies continuously across industries and post-period effects can evolve over time.

The paper’s stacked DID and DDD specifications are therefore best understood as cautious attempts to discipline the comparison set, not as a claim that fixed effects alone solve all identification concerns.

A fourth strand concerns the continuous DDD control dimension, which is central to that strategy. By using world, non-China import penetration as a control dimension, the design compares tariff-exposed industries with observably less-exposed industries while accounting for broader import-penetration dynamics. This helps distinguish a tariff shock from the longer-run exposure patterns emphasized in the import-competition literature. The paper also uses a Bartik or shift-share robustness exercise, which is interpreted in light of recent guidance on how such designs identify effects through exposure shares and shocks (Goldsmith-Pinkham et al. 2020). Following Callaway and Sant’Anna (2021), the robustness checks provide an additional check on the DID interpretation without changing the paper’s cautious causal framing.

A fifth strand comes from quantitative trade models that embed tariff shocks in general-equilibrium frameworks with input-output linkages. This literature emphasizes that a tariff on a given product can affect downstream industries that use it as an intermediate input, so the incidence of a tariff shock need not be confined to the directly targeted sector. That mechanism is consistent with the paper’s decision to measure exposure at the industry level using import shares rather than restricting attention to a narrow set of directly tariffed product categories, and it motivates reading the NAICS-4 exposure measure as capturing input-cost pressure transmitted through supply chains rather than only final-good protection. This paper does not estimate a full general-equilibrium model, but the reduced-form design

is informed by the same logic: exposure is defined broadly enough to capture indirect cost channels, not only industries whose own output was named on a tariff list. The paper operationalizes this by defining tariff exposure, `exp_dtariff_tv`, from NAICS-4 import shares subject to the Section 301 tariff schedules, rather than restricting the treatment variable to industries whose own final output appeared on a tariff list. Combined with the continuous world-import-penetration control, this allows industries that rely heavily on tariffed intermediate inputs, but were not themselves named in a tariff tranche, to register as exposed through their import-share weighting rather than being classified as untreated.

A sixth strand places the 2018 Section 301 episode in the context of earlier U.S. tariff actions, including the 2002 steel safeguard tariffs and antidumping and countervailing duty actions studied in the trade-policy literature. Work on these earlier episodes generally finds that protective tariffs raise costs for downstream users of the protected input, and that any output gains for directly protected producers can be partially or fully offset by losses among industries that consume the tariffed good. That body of evidence is one reason this paper does not take for granted that tariff protection raises domestic industrial production even in the directly protected industries, and it foreshadows the paper’s finding that tariff-exposed manufacturing industries experienced lower, not higher, real production after the 2018 shock. Situating Section 301 alongside these earlier episodes also underscores why a well-identified, product-level exposure measure is useful: aggregate before-and-after comparisons of past tariff episodes have often struggled to separate the tariff shock from concurrent macroeconomic and sector-specific shocks, which is the same identification problem this paper’s triple-difference design is built to address. For example, Pierce (2011) finds that U.S. plants receiving

antidumping protection do not show genuine productivity gains once price and quality effects are accounted for, and that protection can slow the reallocation of resources away from less productive uses, which complicates the presumption that protection straightforwardly benefits protected industries. Downstream-cost studies of the 2002 steel safeguard tariffs, such as Francois and Baughman (2003), estimate that job losses among steel-consuming manufacturers exceeded the employment gains realized in protected steel production. Bown (2021) further documents that the 2018–2019 Section 301 episode itself unfolded through repeated rounds of tariffs, exclusions, and retaliation, underscoring why a well-identified, tranche-level exposure measure is preferable to a simple before-and-after comparison.

Taken together, these strands define the paper’s contribution. The tariff literature establishes that Section 301 affected prices, sourcing, and welfare (Amiti et al. 2019; Fajgelbaum et al. 2020; Flaaen and Pierce 2019), while the trade-and-manufacturing literature shows that import exposure and input shocks can have lasting domestic consequences (J. Pierce and Schott 2016; Acemoglu et al. 2016; Autor et al. 2016). The DID literature clarifies why the comparison design matters when exposure is heterogeneous and industries differ systematically. This paper links those insights by studying whether tariff exposure affected real manufacturing production, producer prices, and nominal shipments in a NAICS-4 monthly panel. Its contribution is not to claim a first result on tariffs or manufacturing, but to provide complementary evidence on how a recent tariff shock propagated across production-side outcomes using a DDD design tailored to the structure of the policy. In sum, the gap this paper aims to fill is a joint, production-side account of a single, well-identified tariff shock: existing work has established that the 2018–2019 tariffs raised prices and reduced trade flows,

and separately that longer-run import-competition shocks depress domestic manufacturing employment and output, but comparatively little evidence directly connects a discrete Section 301 tariff shock to production, prices, and shipments together within the same identification framework.

5 Institutional Background and Policy Timeline

Section 301 of the Trade Act of 1974 grants the United States Trade Representative (USTR) authority to investigate and respond to foreign trade practices deemed unfair or discriminatory. In August 2017, the USTR initiated a formal Section 301 investigation into China’s policies related to intellectual property protection, forced technology transfer, and state-directed industrial practices. The investigation concluded in March 2018, finding that China’s actions imposed unreasonable burdens on U.S. commerce. These findings provided the statutory basis for the imposition of retaliatory tariffs (Office of the United States Trade Representative 2018; Bown 2021).

Between July 2018 and September 2019, the United States implemented four major waves of Section 301 tariffs targeting imports from China. The first wave, List 1, took effect in July 2018 and covered a relatively narrow set of intermediate goods. Subsequent waves, Lists 2, 3, and 4A, expanded tariff coverage substantially to include a broader range of intermediate inputs, capital goods, and selected final goods. Tariff rates also escalated over time, with ad valorem rates rising from 10 percent in some tranches to as high as 25 percent in others. Although some tariff increases were delayed or partially rolled back following the Phase One trade agreement in early 2020, the majority of the tariff structure remained in place

throughout the estimation window.

A key feature of the policy rollout is that tariff coverage was implemented according to predetermined lists of HS-coded goods, which this paper maps to NAICS-4 manufacturing industries. Because individual products entered tariff coverage at different points in time, industries varied substantially in exposure depending on their import composition and reliance on tariffed intermediate inputs. This product-to-industry mapping is the basis for the paper’s exposure measure, `exp_dtariff_tv`, which combines USITC tranche dates with industry import shares to construct a continuous, time-varying measure of tariff exposure at the NAICS-4 level.

This staggered implementation of tariff coverage provides useful variation for empirical analysis. Because different products became subject to tariffs at different points between July 2018 and September 2019, industries linked to more heavily affected inputs experienced larger and earlier changes in exposure than industries linked to inputs that entered coverage later or not at all. This variation is what allows the paper’s triple-difference design to compare more- and less-exposed NAICS-4 industries within the same calendar months.

These institutional features motivate the empirical strategy used in the paper. By combining product-level tariff exposure with monthly NAICS-4 outcome data on production, prices, and shipments, the analysis traces how industry-level outcomes move with changes in tariff exposure over time. The staggered timing of tariff implementation, together with the continuous exposure measure, is what supports both the DDD specification and the event-study evidence reported in the Results section.

6 Data

The empirical analysis combines six data sources into monthly NAICS-4 industry panels. Producer prices come from the BLS API and are measured as `log_ppi`, covering NAICS-4 manufacturing industries at monthly frequency. Real production comes from the FRED API and is measured as `log_ip`, also at the NAICS-4 monthly level. Nominal shipments come from Census M3 and are measured as `log_shipments_nom`, covering NAICS-4 manufacturing industries at monthly frequency through the Census API. These three outcome panels allow the analysis to compare price, real activity, and nominal revenue responses to the tariff shock. They also discipline interpretation across outcomes: BLS-PPI captures producer-price adjustment, FRED-IP captures real production, and Census-M3 captures nominal shipments that can combine price and quantity movements.

The treatment and control variables are assembled from trade and industry data. Section 301 tariff rates and tranche dates come from USITC DataWeb, with coverage mapped to NAICS-4 industries using a manual pull. The resulting treatment variable is `exp_dtariff_tv`, a continuous, time-varying tariff exposure measure at NAICS-4. It is constructed from USITC tranche dates and industry import shares, so it varies across industries according to product exposure and over time according to the implementation of tariff tranches. The main control dimension is `world_pen`, defined as world, non-China import penetration. It is measured as the ratio of non-China imports to domestic shipments at NAICS-4. The import component comes from WTO Tariff Data through the portal, and the industry shipments denominators come from the BEA API. The alternative variable `china_pen` measures China import penetration at NAICS-4. Existing panel construction defines it as scaled China

imports divided by absorption, where absorption is 2017 shipments plus scaled world imports. It is used only in the D3 stress-test⁵ specification as an alternative control dimension. Using non-China import penetration separates general import exposure from exposure to the China-specific tariff shock, which is central to the DDD design. Full mathematical definitions of the exposure and penetration variables are provided in the Variable Construction Appendix (Appendix A8).

The unit of observation is a NAICS-4 industry by calendar month. The sample period is approximately 2015 through 2023, which provides a pre-policy window, the tariff implementation period, and the later pandemic period. `Post = 1` from July 2018, aligning the treatment period with the first Section 301 tranche. `CovidPost = 1` from March 2020, allowing specifications to distinguish the initial tariff period from the later pandemic period. All specifications include NAICS-4 industry fixed effects and calendar-month fixed effects throughout, so identification comes from differential exposure across industries after absorbing time-invariant industry differences and common monthly shocks. The resulting panel is therefore designed to compare industries facing different tariff exposure within the same calendar months, rather than comparing aggregate manufacturing before and after the policy. Panel coverage differs across outcomes because the source series are available for different sets of NAICS-4 industries. The BLS-PPI panel contains 77 clusters, 8,316 full-sample observations, and 4,774 truncated observations through February 2020. The FRED-IP panel contains 31 clusters, 3,348 full-sample observations, and 1,922 truncated observations through February 2020. The Census-M3 panel contains 85 clusters, 9,180 full-sample observations, and

⁵D3 is the DDD stress-test that uses `china_pen` as the alternative control dimension in place of `world_pen`.

5,270 truncated observations through February 2020. These differences matter for inference because the FRED-IP panel is the smallest by cluster count, but it is also the most direct measure of real manufacturing activity. For that reason, FRED-IP serves as the primary anchor outcome, while BLS-PPI and Census-M3 provide complementary evidence on producer prices and nominal shipments. The truncated samples are used for falsification checks that end before the pandemic period, while the full samples preserve the main post-tariff window. Appendix Tables A1 and A2 provide supplementary descriptive statistics and correlation evidence for the NAICS-4 panels.

7 Empirical Identification

The baseline specification is a continuous-exposure difference-in-differences model at the NAICS-4 industry-month level. Let y_{it} denote an outcome for industry i in calendar month t , E_i denote tariff exposure, and Post_t indicate months beginning in July 2018. Cell C1 estimates (the naming convention for C1 and the other specification codes used throughout the tables is defined immediately below):⁶

$$y_{it} = \alpha_i + \gamma_t + \beta \cdot E_i \cdot \text{Post}_t + \varepsilon_{it} \quad (1)$$

This specification includes NAICS-4 industry fixed effects and calendar-month fixed effects, so the coefficient compares changes after the tariff shock across industries with different

⁶D1b is the parsimonious triple-difference specification (31 clusters) that serves as the primary specification for the FRED-IP panel; D1 is the corresponding primary specification for the BLS-PPI and Census-M3 panels. The full naming convention for specifications (C1, C4, D1, D1b, D3, and related labels) is described later in this section and is preserved in Tables 1A–1C and the appendix tables.

exposure levels. The baseline DID is useful as a starting point because it provides the most transparent exposure-by-post comparison. Its limitation is that it can still confound tariff effects with broader trends in import-competing industries if tariff exposure is correlated with unobserved industry trajectories. In that sense, C1 is interpreted as a benchmark rather than the strongest causal design.

Throughout the tables and figures, specifications are labeled with a short code rather than a descriptive name. The letter indicates the broad family of specification: “C” specifications refer to the baseline continuous-exposure model introduced above, and “D” specifications refer to DDD (triple-difference) variants that add a control dimension, such as world import penetration, to that baseline. The number distinguishes among variants within a family (for example, D1 is the baseline DDD specification, D1b is the parsimonious 31-cluster variant used as the primary specification for the FRED-IP panel, and D3 substitutes `china_pen` for `world_pen` as an alternative control dimension). These codes are retained here, rather than replaced with descriptive labels, because they are the same codes used to identify columns and rows in Tables 1A–1C and in the appendix tables (including Appendix A6); readers who want to trace a specific estimate back to its table row can do so using this code.

The specification map is as follows. C1 is the baseline DID in the full sample, while C4 is the corresponding baseline DID in the sample truncated at February 2020. D1b and D1 are the two DDD control-block variants: D1b uses the parsimonious `world_pen × {Post, CovidPost}` block and is primary for FRED-IP, whereas D1 uses the full `world_pen × month` block and is primary for BLS-PPI and Census-M3; the other variant serves as the secondary row for each outcome. D2 is the truncated-sample DDD falsification, and D3 is the stress-test

that substitutes `china_pen` for `world_pen`. Tables 1A-1C report all six cells for each outcome. Table 2 compares C1 with each outcome’s primary DDD specification, and Table 3 reports D2 and D3 as columns in its one-row-per-outcome causal-readiness summary.

The primary identification strategy therefore uses a triple-difference design that adds world, non-China import penetration as a control dimension. Let W_i denote `world_pen`. The primary DDD specification is:

$$y_{it} = \alpha_i + \gamma_t + \beta \cdot E_i \cdot \text{Post}_t + \delta \cdot W_i \cdot \text{Post}_t + \phi \cdot E_i \cdot W_i \cdot \text{Post}_t + \varepsilon_{it} \quad (2)$$

For FRED-IP, Table 1B, row 3 reports D1b as the primary specification, using the parsimonious `world_pen × {Post, CovidPost}` control block; row 4 reports D1 with the full `world_pen × month` control block as the secondary specification. D1b is primary for FRED-IP because it avoids excessive saturation in the smallest outcome panel, which contains 31 clusters. For BLS-PPI in Table 1A and Census-M3 in Table 1C, row 3 instead reports D1 as the primary specification and row 4 reports D1b as the secondary specification, reversing the FRED-IP primary/secondary assignment. Across all three outcomes, row 5 reports the D2 truncated-sample falsification and row 6 reports the D3 `china_pen` stress-test. The coefficient of interest remains the tariff-exposure-by-post effect after accounting for the applicable import-penetration path. Table 2 provides the cross-outcome comparison of C1 with each outcome’s primary DDD specification, while Table 3 summarizes causal readiness with D2 and D3 shown as columns.

Inference is clustered at the NAICS-4 level throughout. The paper reports cluster-robust p-

values alongside wild cluster bootstrap p-values and bootstrap 95% confidence intervals. The bootstrap uses 999 replications, Rademacher weights, and the seed rule `301 + boot_order`. This is especially important for FRED-IP, which has 31 clusters, so tight degrees of freedom motivate bootstrap inference rather than relying only on asymptotic p-values. The bootstrap evidence is therefore treated as a central part of inference, not as a cosmetic add-on to the fixed-effects estimates.

The exclusion logic is that world import penetration should not be correlated with unobserved industry trends driving the DID confound after fixed effects and common time shocks are absorbed. The empirical correlations support this interpretation: $\text{corr}(E_i, \text{world_pen}) = -0.315$ for FRED-IP and $\text{corr}(E_i, \text{world_pen}) = -0.420$ for BLS-PPI. These negative correlations indicate that high tariff exposure is not simply a proxy for high general import penetration. The D2 falsification design further uses the truncated sample ending February 2020; D2 coefficients are near zero across all panels, with exact values reported in the Results section. This falsification check is useful because a large pre-pandemic effect in the truncated window would weaken the interpretation that the main post-period estimates reflect the tariff shock. The paper also uses a shift-share or Bartik instrument as a robustness check, interpreted in light of guidance on exposure-share designs (Goldsmith-Pinkham et al. 2020). The first-stage $F = 1,981.6$, reducing concerns about weak first-stage relevance in that robustness exercise. Formal pre-trend tests have limited power given small cluster counts, so they are interpreted cautiously.

The stacked DID structure handles staggered tariff tranche timing while preserving the post-2018 policy variation central to the design. Following Callaway and Sant’Anna (2021),

the robustness checks provide an additional check on the DID interpretation without changing the paper’s cautious causal framing. Overall, the empirical strategy is intended to triangulate evidence across the baseline DID, the primary DDD, falsification tests, and the Bartik robustness design rather than to rely on a single specification.

8 Results

The baseline DID estimates provide the starting comparison for the Results section, but they are best read as context for the triple-difference design rather than as the paper’s strongest causal evidence. C1 measures the full-sample exposure-by-post DID without the `world_pen` control dimension. In Table 1A, row 1, the BLS-PPI C1 estimate is positive: $\beta = 0.332$, cluster- $p = 0.022$, boot- $p = 0.031$, and the 95% CI is $[0.028, 0.650]$. In Table 1B, row 1, FRED-IP moves in the opposite direction, with $\beta = -0.771$, cluster- $p = 0.012$, boot- $p = 0.011$, and a 95% CI of $[-1.375, -0.133]$. In Table 1C, row 1, Census-M3 is much closer to zero, with $\beta = 0.105$, cluster- $p = 0.475$, boot- $p = 0.470$, and a 95% CI of $[-0.182, 0.387]$. This mixed sign pattern is substantively useful because the three panels measure different margins: producer prices, real industrial production, and nominal shipments. It also motivates the DDD control dimension. A simple exposure-by-post comparison can blend the tariff shock with broader import-penetration dynamics, so the baseline DID evidence sets the empirical terrain rather than settling the interpretation.

The primary evidence comes from FRED-IP, where the DDD estimates indicate a negative production response among more tariff-exposed manufacturing industries. In Table 1B, row 3, the D1b primary specification gives $\beta = -0.783$, $SE = 0.264$, cluster- $p = 0.003$, boot- p

= 0.005, and a 95% CI of [-1.387, -0.176]. Because FRED-IP is the most direct measure of real manufacturing activity, this is the headline estimate for the paper’s production-side interpretation. The magnitude is close to the baseline DID estimate, but the DDD specification is more informative because it conditions on the world import-penetration control dimension. Interpreted at the mean exposure level among treated industries (0.2487), the estimate corresponds to an approximate 17.7 percent decrease in industrial production after the tariff shock, holding the `world_pen` control path fixed (Appendix A7).

The robustness specifications support the same reading. Table 1B, row 4 reports D1 with $\beta = -0.809$, $SE = 0.275$, $\text{cluster-p} = 0.003$, $\text{boot-p} = 0.041$, and a 95% CI of [-1.594, -0.029]. Table 1B, row 6 reports D3 with $\beta = -0.874$, $SE = 0.257$, $\text{cluster-p} = 0.001$, $\text{boot-p} = 0.003$, and a 95% CI of [-1.495, -0.285]; D3 is also reported in Table 3’s FRED-IP columns. These estimates are directionally and statistically consistent with D1b. By contrast, Table 1B, row 5 reports a D2 falsification estimate near zero, with $\beta = -0.060$, $SE = 0.190$, $\text{cluster-p} = 0.752$, $\text{boot-p} = 0.756$, and a 95% CI of [-0.414, 0.287]; D2 is also reported in Table 3’s FRED-IP columns. The persistence of the negative estimate under the full control block and the alternative `china_pen` dimension indicates that the production effect is not tied to a single control construction. Its disappearance in the pre-pandemic falsification window is the expected policy pattern and weighs against a preexisting differential decline. The attenuation ratio is 1.016, which implies that adding the control dimension does not absorb the treatment effect. Taken together, the FRED-IP results support a cautious causal interpretation: higher Section 301 tariff exposure appears associated with lower real industrial production after the policy shock.

The BLS-PPI DDD evidence points in a positive direction but remains inconclusive. In Tables 1A-1C, the D1 coefficient is $\beta = 0.231$, while $\text{boot-p} = 0.301$. The attenuation ratio is 0.694, indicating that the DDD control dimension reduces the baseline DID estimate but does not eliminate its positive sign. Substantively, that pattern is consistent with some producer-price pass-through or cost pressure in more exposed industries. Statistically, however, the evidence is not strong enough to support a claim of a reliable price effect in this specification. For that reason, BLS-PPI is best interpreted as complementary context for the production result: it is suggestive of upward price pressure, but the estimate is too imprecise to carry the paper’s causal interpretation.

The Census-M3 DDD result is different again. In Tables 1A-1C, the D1 estimate is $\beta = -0.101$ with $\text{boot-p} = 0.730$, and the attenuation ratio is -0.961. The results show a negative effect on nominal shipments in the primary specification, although this estimate is imprecisely estimated.

The FRED-IP event-study overlay (Figure 3) provides dynamic evidence consistent with the main DDD estimates. The displayed window is truncated to $k = -18$ to $+19$, while the full ± 24 -month results are retained in the underlying table. The plotted normalization also requires care: $k = -1$ is not zero by construction because of the renormalization convention used to compare the C1 baseline and D1b DDD paths. The bootstrap confidence intervals are pointwise, so they should be read as coefficient-by-coefficient uncertainty rather than as a simultaneous band over the whole event-study path. With those qualifications, the figure reinforces the main result. The post-treatment coefficients are consistently negative and significant, and the DDD path remains below zero after the tariff shock. That pattern is

important because it shows that the FRED-IP result is not driven only by a single static coefficient. The timing evidence is instead broadly aligned with a production-side decline after the policy change, while the cautions about windowing, normalization, and pointwise inference keep the event-study interpretation appropriately bounded.

The attenuation-ratio comparison in Figure 4 and Table 2 summarizes how much the DDD control dimension changes each baseline DID result. FRED-IP has an attenuation ratio of 1.016, so the control dimension leaves the production estimate essentially intact and slightly larger in magnitude relative to the baseline. BLS-PPI has an attenuation ratio of 0.694, indicating partial attenuation of the positive price estimate. Census-M3 has an attenuation ratio of -0.961, reflecting the near-zero DDD estimate on the opposite side of its small baseline estimate. This comparison helps explain why the strongest interpretation rests on FRED-IP rather than on a pooled reading across all outcomes.

The Bartik IV robustness check in Figure 6 reaches a similar conclusion for the primary FRED-IP specification. The OLS D1b estimate is $\beta = -0.7831$ with a 95% CI of $[-1.387, -0.176]$. The 2SLS Bartik IV estimate is $\beta = -0.8685$ with a 95% CI of $[-1.441, -0.316]$. The first-stage $F = 1,981.6$, which reduces concern that the robustness result is being driven by a weak first stage. The IV estimate is slightly larger in magnitude than the OLS estimate, but the two intervals point to the same substantive conclusion: higher tariff exposure is associated with lower real manufacturing production. This robustness check is therefore consistent with minimal attenuation bias, but it is not the headline design. It supports the DDD evidence rather than replacing it. The full first-stage, reduced-form, and 2SLS results underlying this comparison are reported in Appendix Table A6.




Figure 1: Tariff exposure by NAICS-2 manufacturing sector. The figure summarizes the distribution of tariff exposure across NAICS-2 manufacturing sectors.

Notes: The figure summarizes tariff exposure across broad NAICS-2 manufacturing sectors.




Figure 2: FRED-IP trends by exposure tercile. The figure compares pre/post industrial production trends across tariff-exposure terciles.

Notes: The figure compares industrial production trends across tariff-exposure terciles and is used as descriptive support for the main production-side analysis.




Figure 3: FRED-IP event study — baseline vs. D1b DDD.

Notes: Window truncated to $k = -18$ to $+19$ for legibility. Full ± 24 -month results, including the COVID-period tail ($k=20-22$, where bootstrap CIs widen substantially), are available in the underlying results table. The $k=-1$ point reflects the project's renormalization convention (subtracting the pre-period mean, with the omitted reference month included in the denominator) and is not, by construction, equal to zero.




Figure 4: Attenuation ratio comparison. Ratios are BLS-PPI ratio = 0.694, FRED-IP ratio = 1.016, and Census-M3 ratio = -0.961.

Notes: A negative Census-M3 attenuation ratio reflects a near-zero DDD estimate on the opposite side of the baseline estimate and should not be interpreted as a substantive sign reversal.




Figure 5: Back-of-the-envelope economic impact. The calculation implies annual loss = -51.9 billion dollars, cumulative loss = -311.4 billion dollars, and implied jobs = -168 thousand jobs.

Notes: The figure reports back-of-the-envelope implied output and employment effects based on the primary FRED-IP DDD estimate.




Figure 6: OLS vs. Bartik IV comparison. Point estimates are OLS = -0.7831, 95% CI [-1.387, -0.176], and 2SLS = -0.8685, 95% CI [-1.441, -0.316].

Notes: The figure compares the primary FRED-IP D1b estimate with the corresponding Bartik IV estimate.

Table 1A — Main DDD Results: BLS-PPI

Specification	$\hat{\beta}$	SE	Cluster- p	Boot- p	Boot 95% CI	N	Clusters
Panel A. BLS-PPI (log PPI, 77 NAICS-4 clusters)							
Baseline DID — Full sample	0.332**	(0.145)	0.022	0.031	[0.028, 0.650]	8,316	77
Baseline DID — Truncated (Feb 2020)	-0.017	(0.068)	0.807	0.818	[-0.154, 0.108]	4,774	77
DDD Primary — world_pen $\times$ month $\dagger$	0.231	(0.162)	0.155	0.301	[-0.165, 0.622]	8,316	77
DDD Secondary — world_pen $\times$ {Post, CovidPost}	0.226	(0.157)	0.152	0.147	[-0.071, 0.522]	8,316	77
DDD Falsification — Truncated	-0.060	(0.067)	0.373	0.717	[-0.368, 0.227]	4,774	77
DDD Stress test — china_pen $\times$ {Post, CovidPost}	0.179	(0.160)	0.262	0.474	[-0.281, 0.643]	8,316	77

Table 1B — Main DDD Results: FRED-IP

Specification	$\hat{\beta}$	SE	Cluster- p	Boot- p	Boot 95% CI	N	Clusters
Panel B. FRED-IP (log IP, 31 NAICS-4 clusters)							
Baseline DID — Full sample	-0.771**	(0.305)	0.012	0.011	[-1.375, -0.133]	3,348	31
Baseline DID — Truncated (Feb 2020)	-0.119	(0.218)	0.585	0.579	[-0.529, 0.286]	1,922	31
DDD Primary — world_pen $\times$ {Post, CovidPost}†	-0.783***	(0.264)	0.003	0.005	[-1.387, -0.176]	3,348	31
DDD Secondary — world_pen $\times$ month	-0.809***	(0.275)	0.003	0.041	[-1.594, -0.029]	3,348	31
DDD Falsification — Truncated	-0.060	(0.190)	0.752	0.756	[-0.414, 0.287]	1,922	31
DDD Stress test — china_pen $\times$ {Post, CovidPost}	-0.874***	(0.257)	0.001	0.003	[-1.495, -0.285]	3,348	31

Table 1C — Main DDD Results: Census-M3

Specification	$\hat{\beta}$	SE	Cluster- p	Boot- p	Boot 95% CI	N	Clusters
Panel C. Census-M3 (log nominal shipments, 85 NAICS-4 clusters)							
Baseline DID — Full sample	0.105	(0.146)	0.475	0.470	[-0.182, 0.387]	9,180	85
Baseline DID — Truncated (Feb 2020)	0.058	(0.058)	0.324	0.314	[-0.063, 0.172]	5,270	85
DDD Primary — world_pen $\times$ month $\dagger$	-0.101	(0.138)	0.466	0.730	[-0.682, 0.438]	9,180	85
DDD Secondary — world_pen $\times$ {Post, CovidPost}	-0.098	(0.134)	0.463	0.476	[-0.354, 0.166]	9,180	85
DDD Falsification — Truncated	-0.012	(0.056)	0.833	0.956	[-0.516, 0.477]	5,270	85
DDD Stress test — china_pen $\times$ {Post, CovidPost}	-0.124	(0.136)	0.360	0.779	[-0.950, 0.705]	9,180	85

Notes: Panel rows correspond to estimation cells C1, C4, D1b, D1, D2, D3 (see text). $\hat{\beta}$ is the coefficient on `exp_dtariff_tv` $\times$ Post (DID cells) or `exp_dtariff_tv` interacted with the DDD control (DDD cells). All specifications include NAICS-4 and month fixed effects. Standard errors clustered at the NAICS-4 level. Boot- p and Boot 95% CI from wild cluster bootstrap (999 replications, Rademacher weights, seed 301). Truncated sample ends February 2020 (pre-COVID). $\dagger$ Primary specification. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (cluster-robust).

Table 2 — DDD Robustness Panel: Baseline DID versus Primary DDD

Outcome	Baseline $\hat{\beta}$	Baseline Cluster- p	DDD $\hat{\beta}$	DDD Cluster- p	DDD Boot- p	DDD Boot CI	Attenuation ratio	Clus- ters
BLS-PPI (log PPI)	0.332**	0.022	0.231	0.155	0.301	[-0.165, 0.622]	0.694	77
FRED-IP (log IP) †	-0.771**	0.012	-0.783***	0.003	0.005	[-1.387, -0.176]	1.016	31
Census-M3 (log shipments)	0.105	0.475	-0.101	0.466	0.730	[-0.682, 0.438]	-0.961	85

Notes: Baseline DID uses cell C1. Primary DDD uses cell D1b for FRED-IP and cell D1 for BLS-PPI and Census-M3. † = primary anchor outcome. Attenuation ratio = DDD $\hat{\beta}$ / Baseline $\hat{\beta}$. Ratio > 1 indicates the DDD control does not absorb the treatment effect (i.e., the effect is not explained by the import-penetration time path). Ratio < 1 indicates partial absorption. Census-M3 shows a negative effect that is imprecisely estimated. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (cluster-robust).

Table 3 — Attenuation and Causal Readiness Summary

Outcome	$\text{corr}(E_i, \text{world_pen})$	$\text{corr}(E_i, \text{china_pen})$	Attenua- tion ratio	D2 falsifi- cation $\hat{\beta}$	D2 Cluster- p	D3 stress-test $\hat{\beta}$	D3 Cluster- p	Overall status
BLS-PPI	-0.420	-0.479	0.694	-0.060	0.373	0.179	0.262	Partial / inconclusive
FRED-IP	-0.315	-0.299	1.016	-0.060	0.752	-0.874***	0.001	Robust — DDD and stress test significant; falsification clean
Census- M3	-0.388	-0.460	-0.961	-0.012	0.833	-0.124	0.360	Null — DDD primary near zero and imprecisely estimated

Notes: $\text{corr}(E_i, \text{world_pen})$ is the cross-sectional correlation between eventual tariff exposure and world import penetration across NAICS-4 industries. Negative values confirm that more-exposed industries are not simply high-import-penetration industries, validating the DDD exclusion logic. D2 = truncated falsification (should be close to zero). D3 = china penetration stress test. Status labels: “Robust” = DDD primary significant and falsification null; “Partial / inconclusive” = DDD primary not significant at conventional levels; “Null” = DDD primary estimate not statistically distinguishable from zero and imprecisely estimated. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$ (cluster-robust).

9 Robustness Checks

The first robustness check is the truncated falsification design, D2, which asks whether the main pattern appears in the pre-COVID period where the identifying logic should not generate a large treatment effect. Across the three panels, the D2 estimates are small and statistically indistinguishable from zero. In Tables 1A-1C, BLS-PPI has D2 beta = -0.060, cluster-p = 0.373, and boot-p = 0.717. FRED-IP has D2 beta = -0.060, cluster-p = 0.752, and boot-p = 0.756. Census-M3 has D2 beta = -0.012, cluster-p = 0.833, and boot-p = 0.956. The uniform absence of significant D2 estimates is important because it addresses the main concern that high-exposure industries were already moving differently before the post-2018 treatment window. These estimates do not prove parallel trends in a mechanical sense, but they are consistent with the parallel-trends assumption holding in the pre-COVID period. They also strengthen the interpretation of the FRED-IP production result by showing that the negative primary estimate is not mirrored by a comparable falsification effect in the truncated sample.

The second robustness check is the China penetration stress test, D3, which asks whether the main findings survive an additional control structure tied more directly to China-specific penetration. The D3 pattern is consistent with the hierarchy of evidence in the main results. In Table 2, BLS-PPI has D3 beta = 0.179, cluster-p = 0.262, and boot-p = 0.474. This estimate remains positive but inconclusive, so it does not provide a strong price-side result. FRED-IP has D3 beta = -0.874, cluster-p = 0.001, and boot-p = 0.003. This estimate is statistically significant and directionally consistent with the D1b primary result, reinforcing the conclusion that tariff exposure is associated with lower real manufacturing production.

Census-M3 has $D3\text{ beta} = -0.124$, $\text{cluster-p} = 0.360$, and $\text{boot-p} = 0.779$. As in the main specification, the shipments panel remains null. Table 3 therefore helps summarize the same cross-outcome pattern: the production estimate is robust, the producer-price estimate is suggestive but not precise, and nominal shipments do not show a reliable response.

The third check concerns pre-treatment dynamics. Formal pre-trend tests have limited statistical power in the FRED-IP and BLS-PPI panels given the small number of clusters (31 and 77, respectively). Visual inspection of the event-study coefficients does not reveal systematic pre-treatment trends in the FRED-IP panel. The truncated falsification results (D2) provide a more powerful test of parallel trends and are near zero in all panels, consistent with the identifying assumption. For that reason, the pre-trend evidence should be read as supportive but not decisive on its own. The FRED-IP event-study overlay is most useful when combined with the D2 falsification estimates: the visual dynamics do not point to obvious pre-treatment divergence, while the truncated falsification coefficients are close to zero and statistically weak across BLS-PPI, FRED-IP, and Census-M3. This combination is the strongest robustness evidence for the paper’s identifying comparison. It does not remove all concerns about nonrandom tariff exposure, but it reduces the specific worry that the FRED-IP result is simply the continuation of a preexisting differential decline among treated industries. Appendix Tables A3, A4, and A5 report additional robustness checks for leave-one-out stability, alternative treatment timing, and NAICS-2 trend controls.

10 Implications

10.1 Mechanisms and Interpretation

Although the paper’s reduced-form design is not intended to fully decompose the channels through which tariff exposure lowers real production, two candidate mechanisms are consistent with the pattern of results and are worth making explicit. The first is an input-cost channel: because `exp_dtariff_tv` is constructed from tariff rates on the HS6 products an industry imports, weighted by each product’s share of the industry’s 2017 import bill (Appendix A8), industries with higher exposure are those that relied more heavily on now-tariffed imported inputs. Raising the cost of those inputs mechanically raises marginal cost for a given level of output, which can lead cost-minimizing firms to reduce production even absent any change in final demand. This channel is consistent with the FRED-IP result: more-exposed industries show lower real production after the shock, and the effect survives controlling for general, non-China import penetration, so it is not simply a story about industries that were already shrinking due to broader import competition.

A second candidate mechanism is foreign retaliation and lost export competitiveness. Several U.S. trading partners, including China, imposed retaliatory tariffs on U.S. exports following the Section 301 actions. To the extent that industries with higher Section 301 exposure also faced more retaliation on their own export markets, part of the estimated production decline could reflect reduced foreign demand for U.S. output rather than, or in addition to, higher input costs. The exposure measure used here, `exp_dtariff_tv`, is built entirely from U.S. import tariffs on inputs and does not incorporate foreign retaliatory tariffs on U.S. exports, so

this design cannot separate an input-cost channel from a retaliation-driven demand channel using the data assembled here. This is a genuine limitation of the identification strategy rather than a modeling choice: distinguishing the two channels would require industry- or product-level data on retaliatory tariff exposure specific to each affected industry’s export destinations, which is left for future work.

The joint pattern across the three outcome panels is informative about mechanism, even though it cannot fully identify one. If the dominant channel were a pure demand contraction with flexible prices, tariff exposure would be expected to lower both real production and producer prices together, since a negative demand shock ordinarily pushes price down alongside quantity. Instead, the BLS-PPI point estimate moves in the positive direction, albeit imprecisely (Table 1A), which is more consistent with an input-cost channel: costs rise, some of that increase is passed through to producer prices, and firms simultaneously produce less because the cost increase is not fully passed on to their own customers. Census-M3 shows a negative but imprecisely estimated effect on nominal shipments. On balance, the cross-outcome pattern is more supportive of a cost-side channel than of a pure export-demand explanation, but the paper’s design does not rule out some contribution from retaliation-driven demand effects, and the two channels are not mutually exclusive.

10.2 Economic Magnitude

The economic magnitude calculations translate the primary FRED-IP estimate into an interpretable scale. Figure 5 reports the back-of-the-envelope calculation, which applies the D1b coefficient to aggregate FRED-IP exposure. On that basis, the implied annual loss is

-\$51.9 billion, the cumulative loss over six years is -\$311.4 billion, and the implied employment effect is -168 thousand jobs. These values are useful for conveying scale, but they should not be interpreted as a full welfare accounting. The calculation does not account for general equilibrium effects, downstream linkages, or supply-chain adjustments. It also does not model changes in markups, sourcing networks, or firm entry and exit. The purpose is therefore narrower: to show that the estimated production response is economically meaningful if applied to aggregate exposed manufacturing activity.

10.3 Policy Implications

The policy implications follow directly from the cross-outcome pattern. Section 301 tariffs appear to have reduced domestic industrial production in exposed manufacturing industries, even controlling for global import-competitive pressures. That interpretation is strongest for FRED-IP and should remain cautious because tariff exposure was not randomly assigned. Still, the result is difficult to reconcile with a simple protectionist prediction in which tariffs unambiguously raise domestic output among exposed industries. Census-M3 shows a negative but imprecisely estimated effect on nominal shipments, so this margin does not provide clear evidence on its own. The Bartik IV result, $\beta = -0.8685$, is consistent with OLS, so the production effect is not primarily driven by endogenous tariff targeting of declining industries. This matters for policy evaluation because the relevant margin is not only whether imports become more costly, but whether domestic producers expand activity after the shock. Taken together, the evidence suggests that tariff protection designed to shield domestic producers may have unintended negative effects on output in exposed industries in the short-to-medium

run. The implication is not that tariffs never protect particular firms or product lines, but that broad exposure to tariffed inputs and trade-policy disruption can weaken manufacturing activity even when the policy goal is domestic industrial support.

11 Conclusions

This paper studies the causal effects of Section 301 China tariffs on U.S. manufacturing at the NAICS-4 industry level. The empirical contribution is to bring granular industry-level identification to a policy episode that is often discussed in aggregate terms. The core design is a triple-difference framework that compares industries with different tariff exposure before and after the policy shock while using world, non-China import penetration as the control dimension. That control dimension is central because tariff exposure was not randomly assigned across industries and may otherwise be confounded with broader import-competition dynamics. By separating China-specific tariff exposure from general import-penetration pressure, the design provides causal evidence on production-side effects rather than only documenting price incidence or aggregate trade-flow changes. The result is a more disciplined account of how a major tariff intervention translated into real manufacturing activity, producer prices, and nominal shipments across detailed industries. This framing also clarifies why production, prices, and shipments need to be evaluated jointly rather than treated as interchangeable measures of industrial performance.

The headline finding is that the strongest evidence appears in real production. In the primary FRED-IP D1b specification, $\beta = -0.783$ and $\text{boot-p} = 0.005$, with an attenuation ratio of 1.016. The attenuation ratio indicates that adding the DDD control dimension does

not absorb the production effect. The Bartik IV estimate, $\beta = -0.8685$, is similar in sign and magnitude, which supports the interpretation that the result is not primarily an artifact of endogenous tariff targeting. The secondary outcomes are more mixed and help bound the interpretation. BLS-PPI is suggestive but inconclusive, with $\text{boot-p} = 0.301$ and an attenuation ratio of 0.694. Census-M3 shows a negative effect on nominal shipments, with an attenuation ratio of -0.961, although this estimate is imprecisely estimated. The back-of-the-envelope economic magnitudes are large enough to be policy relevant: annual loss = -\$51.9 billion, cumulative loss = -\$311.4 billion, and implied jobs = -168 thousand. These scale calculations should be read as descriptive implications of the reduced-form estimate, not as a complete welfare analysis.

Several limitations define the scope of the findings without changing the central interpretation. FRED-IP has 31 clusters, so inference is credible but power-limited. The analysis is reduced-form and does not estimate general equilibrium effects or firm-level heterogeneity. The sample ends in 2023, so longer-run supply-chain adjustments are outside the window. These boundaries point to natural next steps. Future research could use firm/plant-level data to study heterogeneity by size, export orientation, and vertical integration. Future research could also use a structural model to decompose input-cost versus demand channels, especially where tariff exposure changes both production costs and output-market conditions. Finally, future research could extend the analysis with longer post-tariff panels for medium-run adjustment estimation. Those extensions would complement the evidence here by clarifying mechanisms, persistence, and distribution across industries. Within the present design, however, the main implication is clear: Section 301 tariff exposure appears to have reduced

real manufacturing production in more exposed industries, even after accounting for broader global import-competitive pressures.

12 Appendix

The appendix provides supplementary detail supporting the main-text descriptive, robustness, and identification claims. These tables document descriptive patterns and additional robustness checks, but they are not required to follow the main argument.

Table A1 — Panel-Level Summary Statistics

Variable	N	Mean	SD	Min	Max
Tariff exposure (exp_dtariff)	85	0.207	0.050	0.034	0.250
World import penetration (world_pen)	85	0.389	0.274	0.011	0.987
China import penetration (china_pen)	85	0.118	0.149	0.000	0.557
Log industrial production (FRED-IP panel)	31	4.572	0.064	4.423	4.707
Log producer price index (BLS-PPI panel)	77	5.087	0.363	4.031	6.678
Log nominal shipments (Census-M3 panel)	85	9.829	1.352	5.978	11.285
Treated (binary)	85	0.329	0.473	0	1

Notes: N differs across outcome variables because BLS-PPI, FRED-IP, and Census-M3 cover different subsets of the 85 NAICS-4 manufacturing industries in the underlying panel, consistent with the cluster counts reported in the Data section. Exposure, penetration, and treatment status are time-invariant industry-level values; outcome variables are full-sample monthly means.

Table A2 — Correlation matrix, FRED-IP panel

Variable	Tariff exposure	World penetration	China penetration	Eventual exposure
exp_dtariff	1.000	-0.558	-0.347	0.831
world_pen	-0.558	1.000	0.652	-0.315
china_pen	-0.347	0.652	1.000	-0.299
E_i^{eventual}	0.831	-0.315	-0.299	1.000

Notes: Pearson correlations are calculated across the 31 NAICS-4 industries covered by the FRED-IP anchor panel. E_i^{eventual} is the post-July-2018 mean of the time-varying tariff exposure measure.

Table A3: Leave-one-out stability check, FRED-IP

Excluded NAICS-2	$\hat{\beta}$	Cluster- p	Boot- p	CI lower	CI upper	Clusters remaining
None (full-sample reference)	-0.783	0.003	0.005	-1.387	-0.176	31
31	-0.332	0.017	0.082	-0.709	0.037	18
32	-1.033	0.062	0.061	-2.070	0.048	16
33	-0.949	0.000	0.003	-1.672	-0.240	28

Notes: The FRED-IP panel contains only NAICS-2 sectors 31, 32, and 33. Excluding sectors 31, 32, and 33 leaves 18, 16, and 28 NAICS-4 clusters, respectively. These small-cluster results are indicative stability checks, not strict robustness pass/fail inference. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A4: Alternative treatment timing robustness

Outcome	Timing	Primary spec	$\hat{\beta}$	Cluster- p	Boot- p	CI lower	CI upper	Clusters
BLS-PPI	July 2018	D1: world_pen x month	0.231	0.155	0.301	-0.165	0.622	77
BLS-PPI	August 2018	D1: world_pen x month	0.231	0.155	0.280	-0.178	0.631	77
BLS-PPI	September 2018	D1: world_pen x month	0.235	0.157	0.313	-0.172	0.659	77
BLS-PPI	January 2019	D1: world_pen x month	0.243	0.158	0.266	-0.168	0.645	77
FRED-IP	July 2018	D1b: world_pen x {Post, CovidPost}	-0.783	0.003	0.005	-1.387	-0.176	31
FRED-IP	August 2018	D1b: world_pen x {Post, CovidPost}	-0.784	0.003	0.004	-1.383	-0.158	31
FRED-IP	September 2018	D1b: world_pen x {Post, CovidPost}	-0.797	0.003	0.001	-1.406	-0.189	31
FRED-IP	January 2019	D1b: world_pen x {Post, CovidPost}	-0.853	0.001	0.002	-1.477	-0.243	31
Census-M3	July 2018	D1: world_pen x month	-0.101	0.466	0.730	-0.682	0.438	85
Census-M3	August 2018	D1: world_pen x month	-0.101	0.466	0.738	-0.687	0.489	85
Census-M3	September 2018	D1: world_pen x month	-0.107	0.447	0.733	-0.668	0.482	85
Census-M3	January 2019	D1: world_pen x month	-0.138	0.355	0.645	-0.688	0.443	85

Notes: July 2018 rows report the primary reference timing. Alternative exposure timing begins in August 2018, September 2018, or January 2019. For FRED-IP D1b, the world-penetration Post boundary moves with the alternative treatment date while CovidPost remains March 2020. Significance:

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

Table A5: NAICS-2 linear trends robustness

Specification	$\hat{\beta}$	Cluster- p	Boot- p	CI lower	CI upper	N	Clusters
FRED-IP D1b + NAICS-2 linear trends	-1.042	0.000	0.003	-1.733	-0.292	3,348	31

Notes: The model includes trend_31 and trend_32; NAICS-2 sector 33 is the omitted reference trend. This is a robustness check, not the primary specification. Significance: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A6 — Bartik IV First-Stage and 2SLS Results (FRED-IP)

Estimate	Coef.	SE	p	F	Diagnostic
Panel A. D1b IV estimates					
OLS benchmark	-0.7831	0.2580	0.0024		
First stage	0.9856	0.0221	0.0000	1,981.64	
Reduced form	-0.8560	0.2751	0.0019		
2SLS	-0.8685	0.2847	0.0023	1,981.64	
Panel B. Diagnostic gate					
corr(Z, exposure)	0.973636				passed

Notes: The endogenous regressor is the validated D1b `exp_dtariff_tv` column. The transformed Bartik instrument is $Z * (\text{exp_dtariff_tv} / \text{exp_dtariff})$, preserving the validated monthly timing structure while replacing only the cross-sectional peak exposure level with Bartik Z. Estimates use the D1b NAICS-4 and month fixed effects plus `world_pen_x_Post` and `world_pen_x_CovidPost` controls, residualized before OLS/first-stage/reduced-form/2SLS estimation. Bootstrap inference uses 999 Rademacher wild-cluster repetitions with centered p-values and confidence intervals.

12.1 Appendix A7: Back-of-Envelope Calculation Details

The calculation applies the primary FRED-IP D1b coefficient to the mean post-period exposure of the 11 treated NAICS-4 industries in the pinned FRED-IP panel. Mean exposure is 0.2487 and the D1b coefficient is -0.783, producing an exact proportional effect of

$$\exp(-0.783 \times 0.2487) - 1 = -0.1769. \quad (3)$$

The annual value-added effect uses the 2018 Census Annual Survey of Manufactures value-added base for the treated industries:

$$-0.1769 \times \$293,398,311,000 = -\$51,907,789,796, \quad (4)$$

which rounds to the reported annual loss of -\$51.9 billion. The cumulative figure applies this annual rate for six years, 2018-2023:

$$-\$51,907,789,796 \times 6 = -\$311,446,738,779, \quad (5)$$

which rounds to -\$311.4 billion. This extrapolation is a flat annual-rate scale calculation rather than a dynamic or general-equilibrium estimate.

For employment, the annual dollar effect is allocated across the 11 treated industries in proportion to each industry's 2018 Census ASM value added. Each allocated effect is divided by that industry's 2018 value added per worker, using private-ownership employment from the Quarterly Census of Employment and Wages. Summing the industry estimates gives an implied employment change of -168,391 workers, reported in the main text as -168 thousand. This calculation does not impose or estimate a labor-demand elasticity.

12.2 Variable Construction Appendix (A8)

Let h index HS6 products, j index NAICS-4 industries, and t index months. The construction code forms a fixed 2017 China-import share for each HS6 product within industry j ,

$$s_{hj}^{China} = \frac{w_{hj} M_{h,2017}^{China}}{\sum_{h' \in j} w_{h'j} M_{h',2017}^{China}}, \quad (6)$$

where w_{hj} is the product-to-industry mapping weight and $M_{h,2017}^{China}$ is average 2017 China imports. If τ_{ht} is the Section 301 tariff rate applying to product h in month t , the time-varying exposure measure is

$$\text{exp_dtariff_tv}_{jt} = \sum_{h \in j} s_{hj}^{China} \tau_{ht}. \quad (7)$$

The penetration measures use scaled imports and 2017 annualized domestic shipments. Let M_j^{World} and M_j^{China} denote the project's scaled world and China import measures, and let $Y_{j,2017}$ denote 2017 domestic shipments. The construction code defines absorption as

$$A_j = Y_{j,2017} + M_j^{World}, \quad (8)$$

with exports omitted by design for this cross-industry ranking proxy. The two penetration controls are

$$\text{world_pen}_j = \frac{M_j^{World}}{A_j} \quad \text{and} \quad \text{china_pen}_j = \frac{M_j^{China}}{A_j}. \quad (9)$$

The code applies the same import-unit scaling before forming both numerators and the absorption denominator. `world_pen` is the primary DDD control dimension; `china_pen` replaces it only in the D3 stress-test.

12.3 Data Construction Narrative

This subsection describes the panel construction pipeline in narrative form, complementing the formal variable definitions in Appendix A8. The goal is to make each step from raw source to estimation sample traceable, since the exposure and penetration measures combine six underlying data sources at different native levels of aggregation.

Step 1, tariff schedule assembly. Section 301 tariff rates and their effective tranche dates are pulled from USITC DataWeb at the HS6 product level. Each HS6 line is assigned the ad valorem rate that applied to it in each calendar month between July 2018 and September 2019, producing a product-by-month tariff-rate series, τ_{ht} , that is zero before a product’s tranche date and equal to the applicable statutory rate afterward.

Step 2, product-to-industry concordance. Each HS6 product is mapped to a NAICS-4 manufacturing industry, and a fixed 2017 China-import weight is computed for each product within its industry, giving each product a time-invariant share of that industry’s China-import bill. This fixed-weight approach, standard in the tariff-exposure literature, holds the industry composition of imports constant at its pre-shock (2017) level so that changes in the constructed exposure measure over time reflect only changes in tariff policy, not contemporaneous changes in import composition that could themselves be a response to the tariffs.

Step 3, exposure aggregation. Within each industry-month, the product-level tariff rates are aggregated using the fixed 2017 weights to produce the continuous, time-varying exposure measure, `exp_dtariff_tv`, following the formula in Appendix A8. Because the underlying product-level tariff rates step up discretely on their respective tranche dates, the resulting

industry-level series is a step function that rises across the four Section 301 waves, with the rate and timing of the increase differing by industry according to product mix.

Step 4, penetration controls. World (non-China) and China import values are pulled from WTO Tariff Data at the same HS6 level, aggregated to NAICS-4 using the same concordance as Step 2, and scaled consistently before being divided by an absorption denominator equal to 2017 domestic shipments (from the BEA API) plus scaled world imports. This produces the time-invariant `world_pen` and `china_pen` controls described in the Data section and Appendix A8.

Step 5, outcome panels and merge. The three outcome series, `log_ppi` (BLS API), `log_ip` (FRED API), and `log_shipments_nom` (Census API), are each pulled at their native NAICS-4 monthly frequency and merged onto the exposure and penetration panel by NAICS-4 code and calendar month. Because the three source series do not cover an identical set of NAICS-4 industries, the merge is conducted separately for each outcome, which is why the three estimation panels differ in cluster count (31, 77, and 85 for FRED-IP, BLS-PPI, and Census-M3, respectively) even though they draw on a common underlying exposure and penetration construction.

Step 6, back-of-envelope inputs. The economic-magnitude calculation in Appendix A7 draws on two additional cross-sectional sources not used elsewhere in the estimation: 2018 Census Annual Survey of Manufactures value added by industry, used to translate the percentage production effect into a dollar figure, and Quarterly Census of Employment and Wages private-ownership employment, used together with 2018 value added per worker to translate the dollar effect into an implied employment count. These two sources enter only the illustrative

scale calculation and play no role in identification.

13 References

- Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price. 2016. “Import Competition and the Great U.S. Employment Sag of the 2000s.” *Journal of Labor Economics* 34 (S1): S141–98. <https://doi.org/10.1086/682384>.
- Amiti, Mary, Stephen Redding, and David Weinstein. 2019. “The Impact of the 2018 Trade War on u.s. Prices and Welfare.” *Journal of Economic Perspectives* 33 (4): 187–210. <https://doi.org/10.1257/jep.33.4.187>.
- Autor, David, David Dorn, and Gordon Hanson. 2013. “The China Syndrome: Local Labor Market Effects of Import Competition in the United States.” *American Economic Review* 103 (6): 2121–68. <https://doi.org/10.1257/aer.103.6.2121>.
- Autor, David, David Dorn, and Gordon H. Hanson. 2016. “The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade.” *Annual Review of Economics* 8: 205–40. <https://doi.org/10.1146/annurev-economics-080315-015041>.
- Boehm, Christoph E., Aaron Flaaen, and Nitya Pandalai-Nayar. 2019. “Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 tōhoku Earthquake.” *Review of Economics and Statistics* 101 (1): 60–75. https://doi.org/10.1162/rest_a_00750.
- Bown, Chad P. 2021. “The US-China Trade War and Phase One Agreement.” *Journal of Policy Modeling* 43 (4): 805–43. <https://doi.org/10.1016/j.jpolmod.2021.02.009>.
- Callaway, Brantly, and Pedro H. C. Sant’Anna. 2021. “Difference-in-Differences with Multiple Time Periods.” *Journal of Econometrics* 225 (2): 200–230. <https://doi.org/10.1016/j.jeconom.2020.12.001>.
- Fajgelbaum, Pablo, Pinelopi Goldberg, Patrick Kennedy, and Amit Khandelwal. 2020. “The Return to Protectionism.” *Quarterly Journal of Economics* 135 (1): 1–55. <https://doi.org/10.1093/qje/qjz036>.
- Flaaen, Aaron, Ali Hortaçsu, and Felix Tintelnot. 2020. “The Production Relocation and Price Effects of U.S. Trade Policy: The Case of Washing Machines.” *American Economic Review* 110 (7): 2103–27. <https://doi.org/10.1257/aer.20190611>.
- Flaaen, Aaron, and Justin Pierce. 2019. *Disentangling the Effects of the 2018–2019 Tariffs*

- on a Globally Connected u.s. Manufacturing Sector*. Finance and Economics Discussion Series Nos. 2019-086. Board of Governors of the Federal Reserve System. <https://www.federalreserve.gov/econres/feds/files/2019086pap.pdf>.
- Francois, Joseph, and Laura M. Baughman. 2003. *The Unintended Consequences of U.S. Steel Import Tariffs: A Quantification of the Impact During 2002*. Trade Partnership Worldwide, LLC.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift. 2020. “Bartik Instruments: What, When, Why, and How.” *American Economic Review* 110 (8): 2586–624. <https://doi.org/10.1257/aer.20181047>.
- Handley, Kyle, Fariha Kamal, and Ryan Monarch. 2020. *Rising Import Tariffs, Falling Export Growth: When Modern Supply Chains Meet Old-Style Protectionism*. Working Paper No. 26611. National Bureau of Economic Research. <https://www.nber.org/papers/w26611>.
- Office of the United States Trade Representative. 2018. *Findings of the Investigation into China’s Acts, Policies, and Practices Related to Technology Transfer, Intellectual Property, and Innovation Under Section 301 of the Trade Act of 1974*. Office of the United States Trade Representative.
- Office of the United States Trade Representative. 2026. *The People’s Republic of China*. USTR Country/Region Trade Summary. <https://ustr.gov/countries-regions/china-mongolia-taiwan/peoples-republic-china>.
- Pierce, Justin R. 2011. “Plant-Level Responses to Antidumping Duties: Evidence from U.S. Manufacturers.” *Journal of International Economics* 85 (2): 222–33. <https://doi.org/10.1016/j.jinteco.2011.07.006>.
- Pierce, Justin, and Peter Schott. 2016. “The Surprisingly Swift Decline of u.s. Manufacturing Employment.” *American Economic Review* 106 (7): 1632–62. <https://doi.org/10.1257/aer.20131578>.
- Sun, Liyang, and Sarah Abraham. 2021. “Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects.” *Journal of Econometrics* 225 (2): 175–99. <https://doi.org/10.1016/j.jeconom.2020.09.006>.