Tariff Reversals and Import Adjustment: The Role of Prior Policy Exposure

Hazwan Haini[1]

Centre for Strategic and Policy Studies (CSPS), Brunei Darussalam

October 2026

Abstract

Do imports recover when tariffs are reduced? Answering this question requires separating exposure to a tariff reversal from exposure to the preceding increase. We examine the February 2026 U.S. tariff reset using public consumption-import data and a frozen statutory-tariff model. A fixed basket covers 44 origins and 4,818 products at the six-digit Harmonized System level. Poisson specifications absorb origin–product, origin–month, and product–month fixed effects. In the January 2025–June 2026 window, a 10-percentage-point larger net tariff cut is associated with 10.4% lower post-reset imports in the unadjusted specification. Conditioning on the preceding tariff increase changes this association to 1.4%, with a 95% sensitivity interval from −6.7% to 10.1%. Conditional estimates in shorter and extended windows are also statistically indistinguishable from zero. Apparent late-post improvements in unadjusted event studies do not survive freely varying controls for prior exposure. Joint pre-event tests nevertheless reject in both conditional follow-up windows, limiting causal interpretation. The findings document substantial specification sensitivity rather than identify irreversible supply-chain adjustment or a zero tariff response. They show why the policy history embedded in a tariff reversal matters for interpreting short-run trade evidence, and provide a transparent public-data framework for examining that distinction.

Keywords: Tariff reversals; import adjustment; trade policy uncertainty; global value chains; Poisson pseudo-maximum likelihood

JEL classification: F13; F14; C23

[1]Corresponding author. Centre for Strategic and Policy Studies (CSPS), Bandar Seri Begawan, Brunei Darussalam. Email: hazwan.haini@csps.org.bn.

1. Introduction

A tariff reduction can occur before the adjustment to the preceding increase is complete. Importers may still be drawing down inventories, renegotiating contracts, changing suppliers, or responding to uncertainty about the durability of the new regime. The observed response to a reversal consequently combines several objects: the incentive created by the current tariff change, the continuing consequences of earlier policy, and the composition of trade that remains available to respond. Distinguishing these objects matters for interpreting the economic performance of a trading system subject to repeated policy changes. A negative association between tariff cuts and subsequent imports need not mean that lower tariffs reduce trade. An apparent recovery need not isolate the reversal's contribution either.

The February 2026 U.S. tariff reset provides a setting in which this distinction is unusually relevant. The Supreme Court's decision in *Learning Resources, Inc. v. Trump* rejected the use of the International Emergency Economic Powers Act (IEEPA) to impose tariffs. Subsequent customs guidance ended collection under the affected provisions from February 24, while a new temporary surcharge under Section 122 of the Trade Act took effect at the same boundary (Supreme Court of the United States, 2026; U.S. Customs and Border Protection, 2026a, 2026b). The resulting change was a replacement of tariff components, with heterogeneous net consequences across origins and products. It was not a uniform movement from protection to free trade. Moreover, the exposures being reduced had been built up during the preceding policy cycle.

This paper asks how the measured import response to that reset changes when the preceding tariff increase is included in the empirical comparison. The question is narrower than whether the court decision raised aggregate real income or permanently reorganized global supply chains. It concerns the relationship between a measured policy reversal and subsequent trade, and the stability of that relationship across economically meaningful comparison periods. We use a fixed origin–product basket to compare the same trade cells over time. The central empirical exercise places the net reset exposure and the preceding tariff increase in the same conditional-mean specification, then allows both exposures to have freely varying monthly associations.

The analysis combines publicly available Census imports for consumption with a frozen release of The Budget Lab at Yale's Tariff Rate Tracker. The basket contains 44 origins and 4,818 six-digit Harmonized System (HS6) products. Monthly outcomes span January 2025 through August 2026; the primary tariff-history window ends in June 2026. The exposure is the difference between modeled total tariff rates on February 19 and February 24, aggregated with fixed 2024 product weights. The preceding increase is the difference between January 2026 and January 2025 mean modeled rates. February 2026 is excluded because it contains the transition. Poisson pseudo-maximum likelihood (PPML) specifications absorb origin–product, origin–month, and product–month fixed effects, with standard errors clustered by origin.

Three findings organize the evidence. First, a negative static association in the primary window is highly sensitive to conditioning on earlier policy. A 10-percentage-point larger net tariff cut is associated with 10.37% lower post-reset import value in the unadjusted specification. Once the preceding tariff increase is interacted with the post indicator, the reset association becomes 1.36%. Its 95% interval under the reported $t(G-1)$ sensitivity reference extends from −6.70% to 10.11%. The estimation sample is identical across these two specifications. The change therefore reflects the conditional comparison rather than the addition or removal of trade observations.

Second, neither a shorter pre-period nor a longer outcome window yields a precisely estimated conditional reset association. The October 2025–June 2026 specification produces a point estimate of 2.29%, with an interval from −6.40% to 11.79%. The January 2025–August 2026 specification produces 1.95%, with an interval from −6.48% to 11.13%. In unadjusted event studies, the reset-exposure association improves between March–April and May–June. These improvements become negative and statistically insignificant when the preceding increase receives its own unrestricted monthly interactions. The same sensitivity applies to the exploratory July–August comparison.

Third, the conditional event specifications retain differences before the reset. The joint pre-event test rejects in the short window at $p = 0.0349$ and in the extended window at approximately 9.03×10^{-9}. The short-window result is particularly informative because its unadjusted pre-event test does not reject. Conditioning on policy history thus changes both the post-event narrative and the diagnostic assessment of the pre-period. The evidence does not provide a stable causal estimate of the reset, nor does it establish that its true effect is zero. It shows that conclusions about persistence or recovery are sensitive to the policy history embedded in the exposure.

The first contribution is to evidence on the reversal stage of tariff adjustment. Studies of the 2018–2019 trade war document tariff incidence, trade contraction, production relocation, and reallocations across suppliers (Amiti et al., 2019; Fajgelbaum et al., 2020, 2024; Flaaen et al., 2020). More recent work examines the much broader 2025 tariff increases and their short-run price and trade consequences (Cavallo et al., 2025; Fajgelbaum & Khandelwal, 2026). These studies motivate the importance of adjustment after a change in protection. Our contribution is not a competing welfare estimate. It is evidence that the association measured when tariffs fall can change substantially once earlier exposure is allowed to continue affecting the comparison.

The second contribution is a measurement distinction between the size of a reversal and the policy path leading to it. The net reset and preceding increase have a correlation of 0.8603 after removing origin and product means on the complete, unweighted pair grid. This overlap makes a simple interpretation of reset exposure problematic even when that exposure is constructed from fixed, pre-period trade weights. We show the consequences directly by reporting specifications with and without the history measure on matching samples. The overlap statistic is a descriptive

diagnostic; it is not the PPML information matrix, an instrument-strength measure, or an identification argument.

The third contribution is to the interpretation of early evidence from repeated trade-policy changes. Static coefficients, event coefficients relative to the final pre-reset month, and differences between post-reset months answer different questions. We report them together with joint pre-event tests, import-value retention, origin exclusions, and covariance diagnostics. This approach follows the broader methodological concern that specification stability and pre-period evidence deserve attention alongside headline estimates (Roth, 2022; Roth et al., 2023). The contribution is a transparent accounting of what survives the conditional comparison, rather than a new estimator or a rule for selecting a favorable pre-test.

The economic stakes extend beyond bilateral import totals. Trade relationships can embody fixed costs, supplier knowledge, and complementary production arrangements. A temporary tariff increase may therefore have consequences after the measured rate falls. Yet the same observed persistence can arise from inventories, prices, or correlated exposure to an earlier policy shock. The present results distinguish the empirical sensitivity of these interpretations from proof of any particular mechanism. This distinction is relevant to the study of global value chains and resilience, while keeping claims about national income, manufacturing employment, and security outside what the outcome data can establish.

Section 2 describes the institutional setting and economic benchmarks. Section 3 explains the data, fixed basket, and exposure construction. Section 4 presents the conditional-mean specifications and scope of identification. Sections 5 and 6 report the estimates and sensitivity exercises. Section 7 develops their economic interpretation, and Section 8 concludes.

2. Institutional setting and economic benchmarks

2.1. A change in the composition of protection

The relevant institutional sequence separates a judicial decision from the collection rules facing an import entry. The Supreme Court decided *Learning Resources* on February 20, 2026, holding that IEEPA did not authorize the tariffs at issue (Supreme Court of the United States, 2026). Customs guidance subsequently specified that the corresponding additional duties would no longer be collected on goods entered for consumption, or withdrawn from warehouse for consumption, from February 24 (U.S. Customs and Border Protection, 2026a). This entry-based boundary is the timing reference for the empirical design. It does not establish the timing of an order, a production decision, or the arrival of a shipment.

The replacement regime is equally important. The February 20 proclamation imposed an additional 10% surcharge under Section 122 for a specified temporary period, with exemptions and interactions with other tariff authorities. Customs implementation guidance placed its start on February 24 (Executive Office of the President, 2026; U.S. Customs and Border Protection,

2026b). Removing an IEEPA component therefore need not imply an equal reduction in the total rate. The analysis uses the net modeled difference across the two comparison dates, including replacement and other modeled components. Positive exposure denotes a reduction in the total modeled rate; negative exposure denotes an increase.

Two features limit an interpretation of this sequence as a simple natural experiment. The common decision date does not randomize the cross-sectional size of the net rate change. That size depends on the protection previously imposed and the exemptions and components that remain. In addition, the replacement regime and subsequent tariff actions affect expectations and realized incentives during the outcome period. The institutional discontinuity motivates the event calendar, but neither its legal origin nor its common timing supplies an exclusion restriction for the magnitude of origin–product exposure.

2.2. Current incentives, inherited adjustment, and uncertainty

A conventional demand response provides a useful benchmark. Other things equal, a lower tariff reduces the wedge between the foreign supply price and the domestic tariff-inclusive price. Evidence from earlier U.S. tariff episodes shows why the distinction between these prices matters: the incidence at the border need not coincide with the change in retail prices (Amiti et al., 2020; Cavallo et al., 2021). Our outcome is customs import value excluding U.S. duties. Even a positive quantity response can be accompanied by changes in exporter prices and the within-HS6 composition of goods. The sign and magnitude of a value association are consequently not a direct estimate of tariff pass-through.

Inherited adjustment provides a second benchmark. In models with sunk entry costs, a firm that has left a destination may require a different incentive to return than a firm needs to continue serving it. Dixit (1989) and Baldwin and Krugman (1989) explain how large changes in trade conditions can leave persistent consequences. Roberts and Tybout (1997) provides evidence that earlier exporting experience matters for participation, and Alessandria and Choi (2014) embeds exporter dynamics in an analysis of liberalization. These mechanisms make it plausible that the response to a tariff reduction depends on the preceding increase. They do not imply that persistence in aggregate product trade identifies sunk costs: the same observation can result from several margins moving together.

Supplier relationships give this possibility a production-network interpretation. Global sourcing involves discrete choices about which countries and suppliers to use, alongside decisions about purchase quantities (Antràs et al., 2017). Network relationships can affect firm performance and the transmission of production shocks (Bernard et al., 2019; Boehm et al., 2019). Evidence on switching in U.S.–China trade also shows that relationship-specific costs can matter for sourcing decisions (Monarch, 2022). A tariff increase can therefore alter the state from which a later reduction operates. An origin–HS6 series, however, aggregates many buyers

and suppliers. A fall in that series cannot establish which relationships changed or whether a supplier network became less resilient.

Inventories provide a different explanation for delayed or apparently reversed responses. Firms can change the timing and size of orders when trade costs are volatile. Alessandria et al. (2010b) connects lumpy international purchases to inventory decisions, while Alessandria et al. (2010a) shows how inventory adjustment can shape aggregate trade dynamics. Front-loading before an anticipated increase can elevate a pre-period comparison; subsequent drawdowns can depress later imports even if final expenditure changes less. A reversal during this adjustment process could therefore coincide with weak customs entries without implying a persistent loss of productive capacity. Monthly entry values alone do not distinguish purchases for immediate use from the timing of inventory replenishment.

Uncertainty adds a forward-looking margin. Lowering today's tariff does not necessarily reduce the expected future cost of serving a market by the same amount. Handley and Limão (2015) and Handley and Limão (2017) show how uncertainty can affect trade and investment through decisions involving sunk costs. Handley and Limão (2022) surveys the broader evidence. The distinction between a tariff level and the perceived durability of that level is relevant when a legal ruling is followed by a temporary replacement instrument. Our data contain modeled statutory rates, not firms' subjective policy expectations. The design can assess the sensitivity of associations to prior rate exposure, but it cannot separately estimate an uncertainty channel.

2.3. Relation to tariff and trade-adjustment research

The closest empirical literature establishes the importance of measuring both policy changes and the margins through which they operate. Amiti et al. (2019) and Fajgelbaum et al. (2020) evaluate the 2018 tariffs using price, trade, and welfare frameworks. Flaaen et al. (2020) studies production relocation and prices in washing machines, and Fajgelbaum et al. (2024) examines reallocations involving countries outside the main bilateral dispute. The review by Fajgelbaum and Khandelwal (2022) places these responses in their wider economic context. These settings show why the response of one origin cannot be read as the response of total trade. They also provide a reason to include both origin–month and product–month fixed effects when examining differential adjustment within a large set of trade cells.

The newer tariff cycle makes the role of history especially salient. Fajgelbaum and Khandelwal (2026) studies the short-run consequences of the 2025 increases, while Cavallo et al. (2025) tracks retail-price responses using high-frequency microdata. Those papers investigate the escalation stage with outcomes and identification strategies suited to their questions. We examine the association around a later replacement of tariff components, with preceding exposure explicitly included. The distinction is substantive: a cell with a large net reduction can also be a cell that experienced a large earlier increase. We do not assume that a coefficient measured during escalation can simply be reversed to predict adjustment after the reset.

Related evidence on the removal of policy uncertainty and the dynamics of trade agreements reinforces this point. Pierce and Schott (2016) links a change in the risk of higher U.S. protection to manufacturing adjustment. Baier and Bergstrand (2007) emphasizes that trade-agreement effects develop over time and that policy selection matters. These studies differ from a short tariff-reset window, but they caution against treating a change in the current statutory rate as a sufficient description of the economic intervention. The empirical benchmarks here are accordingly a current-reset association, an association conditional on earlier protection, and a dynamic comparison that allows both to vary over time.

3. Data and measurement

3.1. Public trade data and the fixed product universe

The trade outcome is monthly imports for consumption from the Census international trade API at the origin–HS6 level (U.S. Census Bureau, n.d.-b). Imports for consumption cover merchandise entering consumption channels directly and merchandise withdrawn from customs warehouses or foreign trade zones for consumption. Customs values exclude U.S. import duties and associated freight and insurance charges (U.S. Census Bureau, n.d.-a). These conventions match the collection boundary used to define the reset, but they also mean that the outcome is neither tariff-inclusive expenditure nor a physical quantity. General imports and imports for consumption are distinct concepts and are not substituted for one another in the analysis.

Origins are selected before examining post-reset outcomes. The universe consists of the top 40 non-U.S. suppliers ranked by 2024 consumption-import value, augmented by eleven Southeast Asian economies, including Timor-Leste. Removing overlaps leaves 44 origins. This selection gives substantial trade coverage while retaining smaller Southeast Asian suppliers. It is a defined research sample rather than the complete set of U.S. trading partners. The analysis is not designed to estimate a separate Southeast Asian treatment effect, and the presence of smaller economies in the universe does not establish that the estimates transport to any particular one of them.

Product eligibility is determined by tariff measurement. The same positive-2024 HTS10 basket must have complete known free or ad valorem base-duty coverage throughout the primary model window and a valid reset comparison. This yields 4,818 HS6 products for every selected origin and 211,992 origin–product pairs. The original balanced panel contains 3,815,856 rows over January 2025–June 2026. Its basket represents 80.54% of reported ordinary-chapter import value from the selected origins during that period, including February. The coverage restriction excludes products with unmeasured specific or compound base duties and some classification or coverage changes. It should not be interpreted as a random sample of products.

The fixed grid preserves absent eligible consumption-import outcomes as zeros. The original panel contains 1,387,540 reported rows and adds 2,428,316 absent eligible outcomes. A further

358,575 source-reported rows explicitly contain zero consumption-import value. Independent country–month controls reconcile the underlying collection totals before balancing; import-value sums within the screened basket are unchanged by the added rows. An added zero means no separately reported HS6 consumption-import value in a reconciled country–month. It does not identify a firm's exit, the absence of every physical shipment, or a broken buyer–supplier relationship. Other financial variables on added rows are left missing.

Table 1 summarizes coverage and estimation-sample retention. The primary comparison excludes February and contains 17 months. The short comparison uses October 2025–June 2026, leaving eight included months. The extended comparison adds July–August outcomes on the unchanged pair universe, giving 20 calendar months and 19 included months after February is excluded. Effective PPML samples are smaller because all-zero pair effects and additional separated or singleton observations do not contribute to the fitted slope comparison. Retention in import value is high in each window, but high value retention does not remove the product-universe restriction.

Table 1. Sample coverage and estimation-sample retention

	Primary window	Short window	Extended window
Period	Jan. 2025–Jun. 2026	Oct. 2025–Jun. 2026	Jan. 2025–Aug. 2026
Included months	17	8	19
Origins in input	44	44	44
HS6 products	4,818	4,818	4,818
Eligible panel rows	3,603,864	1,695,936	4,027,848
Positive input rows	973,096	462,620	1,089,955
PPML observations	1,647,643	686,922	1,874,810
Positive retained rows	972,604	462,110	1,089,498
Origin clusters	44	43	44
Input value (USD billion)	3,538.615	1,669.916	4,030.287
Import value retained (%)	99.874	99.876	99.878

Notes: Origins are the top 40 non-U.S. 2024 consumption-import suppliers plus eleven Southeast Asian economies, with overlaps removed. Eligible rows refer to the fixed origin–HS6 grid after February 2026 is excluded. PPML retention is based on the corresponding unadjusted static fit; adding the prior-increase control preserves the estimation sample within each window. The short-window PPML sample drops Timor-Leste entirely. The same pre-estimation zero-pair screen used for the follow-up fits preserves all positive input outcomes and import value. The original January 2025–June 2026 basket, including February, covers 80.54% of selected-origin ordinary-chapter import value. Extended outcomes retain unverified later policy history.

3.2. Net reset exposure and the preceding increase

Tariff measures come from the published v1.2 release of The Budget Lab at Yale's Tariff Rate Tracker, vintage 2026-07-24-09 (The Budget Lab at Yale, 2026). The analysis freezes that release rather than silently replacing it with a subsequently revised series. Its archived snapshot

hashes identify the exact model inputs. The tracker assembles statutory components at the HTS10–country level; it remains a model of applicable rates, including assumptions about preference use and metal content. It is not a record of duty paid by each import transaction. The underlying tariff classification is documented by the U.S. International Trade Commission (U.S. International Trade Commission, n.d.).

Let c index origin, h an HTS10 line within HS6 product p, and d a date. Let τ_{chd} denote the modeled total tariff in fractional units. The common within-product weight w_{hp}^{2024} is the share of line h in all-origin U.S. consumption imports within product p in 2024. The HS6 model rate and reset exposure are

$$\tau_{cpd} = \sum_{h \in \mathcal{H}_p} w_{hp}^{2024} \tau_{chd}, \qquad D_{cp} = 100 \left(\tau_{cp,\text{Feb19}} - \tau_{cp,\text{Feb24}} \right). \tag{1}$$

Both comparison dates are in 2026. The weights sum to one within the eligible product basket and are common across origins. Holding them fixed prevents contemporaneous changes in the composition of imports from mechanically changing measured exposure. It does not make the underlying tariff schedule exogenous.

The clean reset comparison requires every positive 2024 weight to be represented, known free or ad valorem base duties, and unchanged modeled base rates between the two dates. The total modeled rate includes the replacement Section 122 regime and other tariff components. Exposure is therefore a net policy comparison, rather than a measure of IEEPA removal alone. Preference and content assumptions can cause modeled and transaction-specific rates to differ. Canada and Mexico receive explicit exclusion checks because these assumptions can be especially relevant in their preferential trade arrangements.

Monthly rates are calendar-day means of the eligible daily model states, respecting the source intervals' inclusive endpoints. Define $\overline{\tau}_{cp,m}$ as this monthly rate. The preceding increase, in percentage points, is

$$R_{cp} = 100 \left(\overline{\tau}_{cp,\text{Jan2026}} - \overline{\tau}_{cp,\text{Jan2025}} \right). \tag{2}$$

All inputs to this control end before the reset. It is a summary of earlier protection, not a complete distributed history of every tariff action. The regression variables are $x_{cp} = D_{cp}/10$ and $r_{cp} = R_{cp}/10$, so one unit is a 10-percentage-point larger net reduction or preceding increase. Their residual correlation after removing origin and product means is 0.8603 on the unweighted complete pair grid.

3.3. Timing, extension, and reproducibility

January 2026 is the event-study reference month; February is month zero and is excluded. March is the first full month after the collection-regime boundary. This convention avoids assigning a single regime to the transition month, but it does not eliminate anticipation, shipping lags, or inventory timing. Monthly entry values can reflect orders made under earlier expectations. The

short window reduces the span of the pre-period without changing the event date or product universe. It also reduces the information available for diagnosing earlier differential movements.

The July–August supplement extends outcomes while leaving initial reset exposure fixed. Its values reconcile to the source country–month controls, and it contains exactly the same origin–HS6 pairs. Later monthly tariff rates are not imputed. Crucially, a published model vintage can contain scheduled states without verifying what was subsequently implemented. The analysis does not independently reconcile the later policy history, so the extended specifications are exploratory outcome-window sensitivities. This distinction concerns the policy interpretation of the extension, not whether its observed import values were collected successfully.

The data and output records preserve separate checks for source-value reconciliation, unique keys, monthly coverage, exposure invariance, and estimation-sample loss. The prior-increase control has been reconciled to the monthly-rate inputs. The saved coefficient and covariance files make the reported transformations and linear contrasts reproducible without refitting the models. Appendix A provides variable definitions and the origin list; Appendix C describes numerical inference checks. This documentation identifies the data vintage and reporting conventions on which the substantive comparisons depend.

4. Empirical framework and identification

4.1. Static conditional-mean comparisons

Let Y_{cpt} denote consumption-import value in millions of current U.S. dollars. Define P_t as an indicator for March 2026 onward within the relevant estimation window. The history-adjusted static specification is

$$\mathrm{E}[Y_{cpt} \mid \mathcal{X}] = \exp\left\{\alpha_{cp} + \delta_{ct} + \gamma_{pt} + \beta x_{cp} P_t + \theta r_{cp} P_t\right\}, \tag{3}$$

where α_{cp} denotes origin–product fixed effects, δ_{ct} origin–month fixed effects, and γ_{pt} product–month fixed effects. The unadjusted benchmark imposes $\theta = 0$. Within each window the adjusted and unadjusted models retain the same estimation observations. A change in the reported reset coefficient therefore does not arise from a different realized PPML sample.

Origin–product effects absorb persistent differences in bilateral product trade. Origin–month effects absorb shocks that proportionately affect all covered products from an origin in a month. Product–month effects absorb shocks shared by covered origins within a product. The reset comparison is consequently identified statistically by remaining origin–product variation in exposure interacted with the event calendar. These effects do not absorb every country-specific product shock, nor do they make the residual exposure independent of the preceding policy path. They also absorb much of the variation through which aggregate equilibrium responses could operate, so the coefficient is not an estimate of the change in total U.S. imports.

PPML accommodates zero outcomes without adding an arbitrary constant and avoids requiring a correctly specified conditional variance (Santos Silva & Tenreyro, 2006). Estimation uses the high-dimensional implementation described by Correia et al. (2020). Import value is scaled by a common unit conversion; exposure and coefficient interpretations remain unchanged by that rescaling. Computational removal of all-zero pair effects is distinguished from selecting a sample of firms that survive. Positive outcomes and their associated zero months are retained unless the estimator removes additional separated or singleton observations. Table 1 reports the resulting sample and value retention.

The coefficient β describes a multiplicative conditional-mean association per 10-percentage-point larger cut. We report $100[\exp(\widehat{\beta}) - 1]$ for its percentage transformation. This is a comparison of exposure slopes across the pre/post periods under the maintained conditional-mean specification. It is not an aggregate growth rate, an import-demand elasticity with respect to $1 + \tau$, or a national welfare change. When the conditional mean is misspecified, PPML summarizes the associated score restrictions rather than automatically recovering an average causal response. The role of the history control is therefore assessed through specification sensitivity, not through an asserted causal decomposition of the unadjusted coefficient.

4.2. Monthly associations and post-period contrasts

Let $k = t - t_0$ index months relative to February 2026, and let $\mathcal{K}$ contain the included event months other than January's $k = -1$ reference. The history-adjusted event specification is

$$\mathrm{E}[Y_{cpt} \mid \mathcal{X}] = \exp\left\{\alpha_{cp} + \delta_{ct} + \gamma_{pt} + \sum_{k \in \mathcal{K}} \left(\beta_k x_{cp} + \theta_k r_{cp}\right) \mathbf{1}[t - t_0 = k]\right\}. \tag{4}$$

February is excluded from $\mathcal{K}$. The primary window has 12 pre-reference and four post-reset reset interactions. The short window has three pre-reference and four post-reset interactions. The extended window has 12 and six, respectively. History-adjusted models give r_{cp} the same unrestricted monthly flexibility as x_{cp}. The unadjusted event models omit the θ_k terms.

This specification asks whether exposure gradients in particular months differ from January 2026, after conditioning on the prior-increase gradient. It does not impose that the earlier policy association changes only once at the reset. Joint pre-event tests concern the β_k coefficients for months before January; joint post tests concern the included months after February. Both are tests of reset-exposure associations conditional on the other regressors. They are not tests that the history coefficients themselves are jointly zero.

Differences between later and earlier post-reset months summarize adjustment within the post period. For example, define

$$C_{MJ,MA} = \frac{\beta_3 + \beta_4}{2} - \frac{\beta_1 + \beta_2}{2}, \qquad \mathrm{se}(\widehat{C}) = \sqrt{a'\widehat{V}a}, \tag{5}$$

where March and April are $k = 1, 2$, May and June are $k = 3, 4$, and a is the corresponding contrast vector. The saved covariance matrix $\widehat{V}$ retains dependence across monthly coefficients. We also report mean post relative to January, June minus March, and, in the extension, July–August minus March–April. These summaries were selected after inspecting earlier results and are exploratory. Their exponentiated values describe ratios of geometric-mean exposure associations, not arithmetic means of monthly percentage changes.

4.3. Scope of identification and inference

A causal interpretation would require more than a common legal event. Conditional on the fixed effects and specified history controls, origin–product cells with different reset exposure would need suitable counterfactual paths, and no other exposure-correlated change could account for their relative outcomes. Exposure intensity must also support a meaningful comparison: common trends alone do not generally identify differences in responses across continuous doses without additional assumptions (Callaway et al., 2024). The present regressions do not implement a continuous-treatment causal estimator or establish those assumptions. The event coefficients provide diagnostics for the maintained comparisons.

Four threats are central. First, larger reductions are related to larger preceding increases, whose consequences may be nonlinear, delayed, or imperfectly summarized by a single January-to-January measure. Second, firms can anticipate policy changes and alter inventory and entry timing before the measured boundary. Third, the frozen model combines statutory information with preference and content assumptions; its retrospective vintage is not a measure of what each firm expected at the time. Fourth, replacement and subsequent tariff actions can be correlated with the initial reset exposure. Product-specific seasonality within origins and changes in the mix of HTS10 lines are further sources of residual movement. Rejected pre-event tests reinforce these concerns; a non-rejection in a shorter window would not eliminate them (Roth, 2022).

Standard errors are clustered by origin, allowing arbitrary dependence over products and time within an origin. The effective number of clusters is 44 in the primary and extended windows and 43 in the short window. This choice assumes that remaining cross-origin dependence is sufficiently limited for the reported inference. Product–month effects remove common product-level mean shifts but do not guarantee that assumption. Multiway clustering addresses a different dependence structure (Cameron et al., 2011); replacing it with origin clustering changes the maintained inference assumptions rather than repairing them automatically. Appendix C reports why an earlier rank-deficient multiway covariance is not used to support joint claims.

Individual coefficients and contrasts are accompanied by a $t(G-1)$ reference as a finite-cluster sensitivity calculation, while the joint tests use the estimator's asymptotic chi-squared reference. These calculations do not constitute validated small-cluster bootstrap inference (Cameron & Miller, 2015; MacKinnon et al., 2023). Numerical convergence and full-rank slope covariance are checked separately from identification. No analytical correction for possible finite-sample

bias in the high-dimensional PPML estimates is reported; such concerns can matter in related three-way settings (Weidner & Zylkin, 2021). The estimates are therefore presented as conditional associations, with sensitivity evidence used to delimit their interpretation.

5. Results

5.1. The static association changes with prior exposure

Table 2 presents the core comparison. Columns (1)–(2) use the primary window, columns (3)–(4) the short window, and columns (5)–(6) the extended window. The first column within each pair omits the preceding increase; the second includes it. Each pair uses the same fixed effects, outcome, exposure scaling, and estimation sample. The exercise therefore isolates how the fitted reset association changes when earlier tariff exposure enters the conditional comparison. The percentage transformations below the coefficient rows make the economic scale explicit.

Table 2. Tariff-reset associations with and without prior policy exposure

	(1) Primary	(2) Primary	(3) Short	(4) Short	(5) Extended	(6) Extended
Reset exposure × post	−0.1095***	0.0135	−0.0076	0.0226	−0.0922**	0.0193
	(0.0346)	(0.0411)	(0.0301)	(0.0440)	(0.0373)	(0.0428)
Prior increase × post	–	−0.0931***	–	−0.0227	–	−0.0842***
		(0.0193)		(0.0298)		(0.0140)
Reset association (%)	−10.37	1.36	−0.76	2.29	−8.81	1.95
$t(G-1)$ p-value	0.003	0.744	0.801	0.609	0.018	0.654
Prior-increase control	No	Yes	No	Yes	No	Yes
All three fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,647,643	1,647,643	686,922	686,922	1,874,810	1,874,810
Origin clusters	44	44	43	43	44	44
Pseudo R^2	0.9445	0.9445	0.9498	0.9498	0.9455	0.9455

Notes: PPML with origin–HS6, origin–month and HS6–month fixed effects. The outcome is consumption-import value in millions of current U.S. dollars. Exposures are scaled per 10 percentage points. February 2026 is excluded. Origin-clustered standard errors appear in parentheses. Stars use the $t(G-1)$ reference as a sensitivity calculation: *** $p<0.01$, ** $p<0.05$, * $p<0.10$. This is not bootstrap inference. Primary: January 2025–June 2026; short: October 2025–June 2026; extended: January 2025–August 2026. Post begins in March 2026. The prior increase is January 2026 minus January 2025 mean modeled tariffs. Percentage associations are $100[\exp(\widehat{\beta})-1]$. Pseudo R^2 values are the rounded values in the estimation logs; they are not evidence of causal validity. Extended results are exploratory.

Figure 1 displays the same six static associations with their intervals. The matching-window comparisons show both the movement of the point estimates and the uncertainty remaining after conditioning on prior exposure.

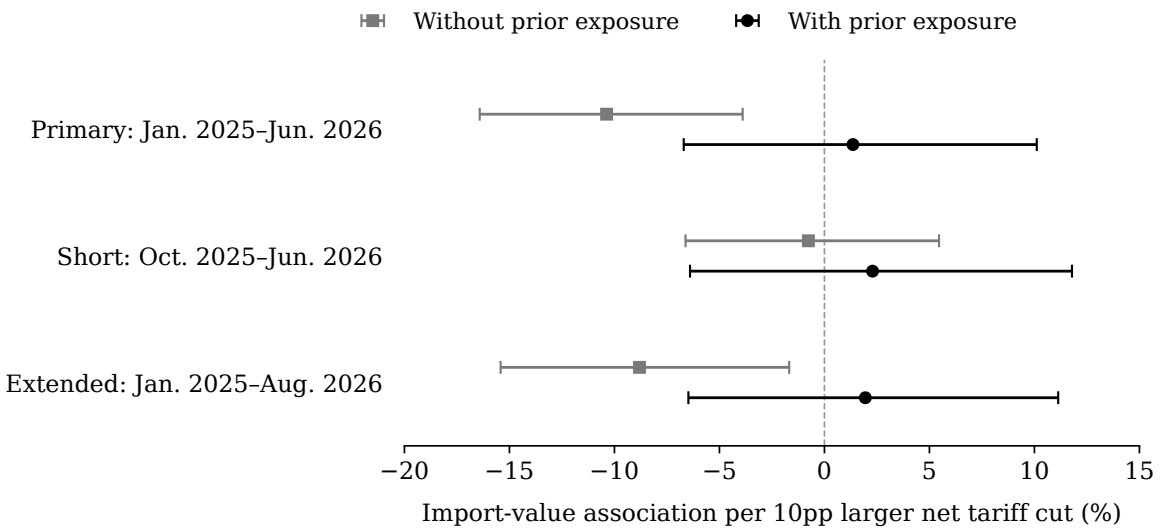


Figure 1. Static reset associations and prior tariff exposure

Notes: The six estimates in Table 2, transformed as $100[\exp(\widehat{\beta}) - 1]$. Bars are 95% origin-clustered $t(G-1)$ sensitivity intervals. Both specifications within each window retain the same observations. The extended window remains exploratory.

In the primary window, the unadjusted reset coefficient is −0.1095 with a clustered standard error of 0.0346. Its percentage association is −10.37%, and its $t(G-1)$ sensitivity p-value is 0.0029. Taken alone, this result could suggest that cells receiving larger net tariff reductions remained relatively weak after the reset. Such a description is incomplete, however, because the exposure is closely related to the preceding increase. Column (2) includes that increase interacted with the same post indicator and changes the reset coefficient to 0.0135, with standard error 0.0411 and $p = 0.744$. The corresponding association is 1.36%.

The prior-increase coefficient in column (2) is −0.0931, with standard error 0.0193. Its percentage transformation is −8.89% per 10-percentage-point larger preceding increase. This coefficient also remains a conditional association. It does not establish that the earlier increase caused the entire unadjusted reset relationship, nor does the movement between the two reset coefficients constitute a structural allocation of their contributions. The informative fact is that the negative reset association does not remain when the model permits earlier protection to have its own post-period gradient. The two specifications fit exactly the same 1,647,643 observations.

The interval is as important as the change in sign. Appendix Table B1 gives the transformed 95% intervals for all three history-adjusted static models. For the primary window the interval is [−6.70%, 10.11%]. Thus the data and maintained inference admit both a moderately negative and a materially positive reset association. The estimate cannot establish complete recovery, persistent damage, or an economically negligible response. In particular, a nonsignificant estimate is not an equivalence test around zero. The disappearance of the unadjusted negative estimate is evidence about specification sensitivity, not proof that tariff reductions have no consequences.

The short window changes the comparison period substantially. Without history controls, its point estimate is already close to zero: -0.0076, or -0.76%, with $p = 0.801$. Adding the preceding increase changes the reset association to 2.29%, with $p = 0.609$. Its 95% interval is $[-6.40\%, 11.79\%]$. The prior-increase coefficient is also imprecise in this window. This contrast with the primary window indicates that the longer pre-period contributes to the negative unadjusted static association. It does not establish that the shorter window is the correct counterfactual; that judgment requires examining its dynamic comparison as well.

In the extension, the unadjusted coefficient is -0.0922, or -8.81%, with $p = 0.0176$. Conditioning on the preceding increase changes it to 0.0193, or 1.95%, with $p = 0.654$. The 95% interval is $[-6.48\%, 11.13\%]$. The prior-increase coefficient remains negative, at -0.0842. Extending the outcome horizon therefore does not recover a stable negative reset association conditional on history. Because later policy is not independently reconciled, this finding is an exploratory consistency check on the same initial exposure rather than an estimate of the response to a maintained tariff regime through August.

5.2. Event profiles and the choice of reference period

The unadjusted primary event profile moves across the post-reset months. Relative to January 2026, the post coefficients move from a negative March association toward a positive June association. Yet the jointly tested pre-reference coefficients are far from jointly zero. Table 3 reports the pre-event rejection at $p \approx 3.07 \times 10^{-12}$. A visual account based only on the slope from March to June would miss the earlier exposure-related movements. Those earlier movements are also relevant to why the static comparison, which pools the full pre-period, differs from an event summary normalized to its final month.

The two summaries do not estimate the same object. The static coefficient imposes a common exposure gradient within its pooled pre and post periods. Event coefficients permit separate monthly gradients and subtract the January gradient through normalization. A nearly zero mean of post coefficients relative to January can therefore coexist with a negative pooled static coefficient if earlier pre-months differ from January. Moreover, nonlinear PPML weighting means that the static coefficient need not equal a simple arithmetic difference of averages of the unrestricted monthly coefficients. Neither summary should be substituted for the other when describing recovery.

Allowing the preceding increase its own monthly interactions weakens the unadjusted dynamic post-event evidence. In the primary window the joint post p-value increases from 0.0716 to 0.2904, while the joint pre test still rejects at approximately 1.40×10^{-8}. This is consistent with the static sensitivity: movements associated with the net reset are difficult to separate from the earlier exposure. The retained pre-event rejection also shows that a linear history control with unrestricted monthly coefficients does not make the remaining reset comparison behave like a validated counterfactual.

Table 3. Joint tests of reset-exposure event coefficients

Specification	Pre df	Pre p	Post df	Post p	Origins
Primary	12	< 0.001	4	0.072	44
Primary, history	12	< 0.001	4	0.290	44
Short	3	0.423	4	0.018	43
Short, history	3	0.035	4	0.344	43
Extended	12	< 0.001	6	0.001	44
Extended, history	12	< 0.001	6	0.272	44

Notes: PPML, the same three fixed effects and samples as Table 2, with freely varying monthly reset-exposure interactions. History specifications also interact the preceding tariff increase freely with every non-reference month. January 2026 is the reference, and February is excluded. Joint Wald p-values use the asymptotic chi-squared reference reported by Stata with origin-clustered covariance. All requested restrictions are retained, and each tested covariance block is full rank. Primary pre-test values are 3.07×10^{-12} without history and 1.40×10^{-8} with history. Extended pre-test values are 1.33×10^{-12} and 9.03×10^{-9}, respectively. These tests are diagnostic, and no specification is selected because its pre-test fails to reject.

Figure 2 shows the short-window event profile conditional on prior exposure. The conditional post coefficients are individually imprecise, and their joint p-value is 0.3443. More strikingly, the conditional pre test rejects at 0.0349, compared with 0.4230 in the unadjusted short event model. None of the three conditional lead coefficients is individually significant at 5% under the reported t reference. There is no contradiction: a joint test uses their covariance and a different joint reference distribution. The result illustrates why the collection of pointwise intervals cannot replace a joint diagnostic.

The unadjusted short-window pre non-rejection consequently does not provide an independently robust reason to prefer that specification. Shortening the window changes the baseline and the restrictions being tested, and adding the economically motivated second exposure changes the diagnostic again. We report all of these comparisons rather than using the pre-test to select a preferred result. A stronger causal analysis would have to justify the counterfactual through additional design information or a defensible sensitivity framework. The present changes in test outcomes cannot supply that justification by themselves.

Figure 3 presents the conditional extended profile. The joint post p-value is 0.2725, compared with 0.0015 without history controls. Its pre test rejects at approximately 9.03×10^{-9}. Thus the strong joint post variation in the unadjusted extension is not maintained in the conditional model, while differential pre-event movements remain. The extension adds observed outcome months but does not transform the original exposure into a complete description of the tariff incentives prevailing in those months. Its conditional profile is informative about sensitivity, with that policy-history limit retained.

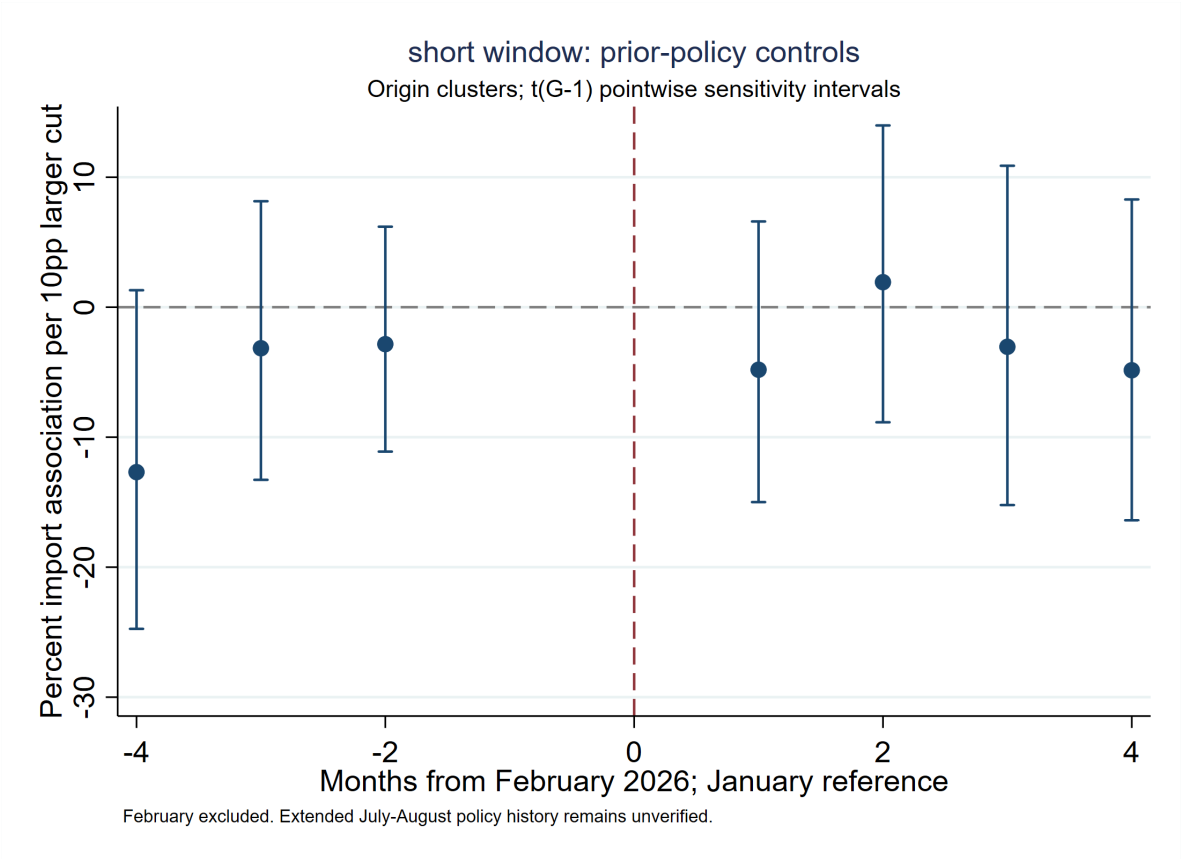


Figure 2. Short-window associations conditional on prior tariff exposure

Notes: Original supplied follow-up figure, unchanged. The sample is October 2025–June 2026, excluding February. Both reset exposure and the preceding increase receive unrestricted month interactions; January 2026 is the reference. All three fixed effects are included. Intervals are pointwise $t(G-1)$ sensitivities with 43 origin clusters.

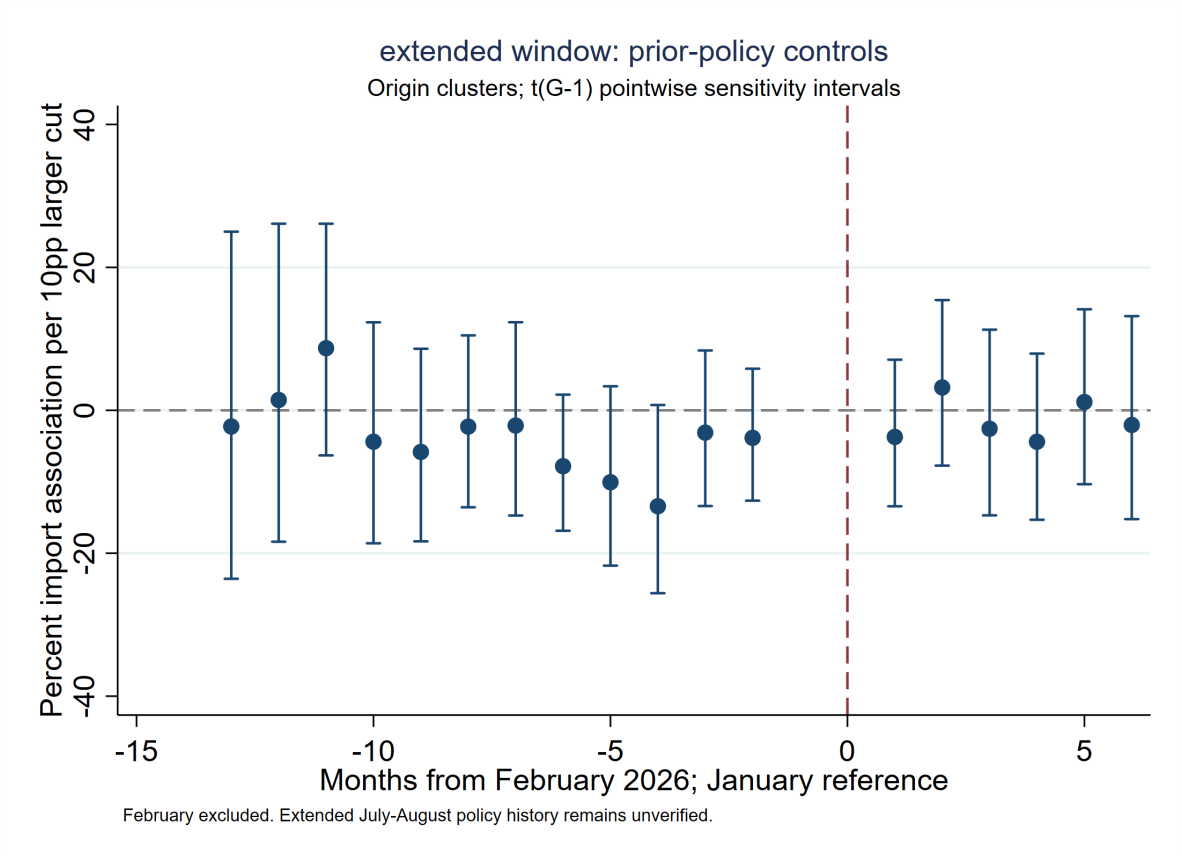


Figure 3. Extended-window associations conditional on prior tariff exposure

Notes: Original supplied follow-up figure, unchanged. The sample is January 2025–August 2026, excluding February. Reset and prior-increase exposures receive unrestricted monthly interactions; January 2026 is the reference. All three fixed effects are included. Intervals are pointwise $t(G-1)$ sensitivities with 44 origin clusters. July–August policy history remains independently unverified.

5.3. Apparent post-reset improvement is not stable to history controls

Table 4 examines the apparent improvement between early and later post-reset months using the full covariance of the event coefficients. In the unadjusted short window, the May–June association relative to March–April is 7.22%, with a 95% interval of $[1.37\%, 13.40\%]$ and $p = 0.016$. The corresponding history-adjusted comparison is -2.49%, with interval $[-7.25\%, 2.52\%]$ and $p = 0.315$. These estimates do not support describing the initial positive comparison as a history-independent recovery pattern. We do not report a formal cross-specification test of the difference between the two estimates.

Table 4. Changes in the reset-exposure association relative to March–April

Specification	Later period	Association (%)	95% interval	p-value
Short, unadjusted	May–Jun.	7.22	[1.37, 13.40]	0.016
Short, history	May–Jun.	−2.49	[−7.25, 2.52]	0.315
Extended, unadjusted	May–Jun.	6.89	[0.66, 13.51]	0.030
Extended, history	May–Jun.	−3.17	[−7.95, 1.85]	0.206
Extended, unadjusted	Jul.–Aug.	9.24	[0.01, 19.32]	0.050
Extended, history	Jul.–Aug.	−0.13	[−8.85, 9.42]	0.976

Notes: Contrasts are differences between averages of event coefficients, evaluated with the complete saved covariance matrix. For example, the May–June comparison is $(\beta_3 + \beta_4 - \beta_1 - \beta_2)/2$. The percentage column is $100[\exp(c'\widehat{\beta}) - 1]$, a ratio of geometric-mean exposure associations, not aggregate import growth or an arithmetic average of percentage effects. Intervals and p-values use the $t(G - 1)$ reference; 43 origins enter the short window and 44 the extended window. All specifications use PPML and the three fixed effects. These comparisons were selected after reviewing the earlier estimates and are exploratory, with unadjusted p-values.

The extended window gives a similar comparison. Without prior exposure, May–June relative to March–April is 6.89%, with $p = 0.030$. With prior exposure it is -3.17%, with $p = 0.206$. The unadjusted July–August comparison is 9.24%, with $p = 0.050$ after rounding; its unrounded value is 0.0498. The conditional comparison is -0.13%, with $p = 0.976$. The latter interval, $[-8.85\%, 9.42\%]$, admits sizable movements in either direction. Neither the unadjusted borderline result nor its conditional disappearance establishes a lasting recovery trajectory.

The remaining conditional contrasts lead to the same assessment. Appendix Table B2 reports all seven rather than only the later-versus-earlier averages. Mean post relative to January is -2.73% in the short window and -1.43% in the extension, with $p = 0.626$ and 0.787. June minus March is close to zero in both conditional models. Every conditional contrast is statistically insignificant under the reported reference. These results constrain a claim of a precisely measured conditional rebound, but they do not provide evidence that every economically relevant response has been ruled out.

The combined evidence is more informative than either a significant unadjusted comparison or an isolated insignificant conditional coefficient. On the same samples, both the negative pooled reset association and the apparent late-post improvement change when prior exposure enters. At the same time, the conditional event profiles retain pre-event differences. The supported finding is substantial dependence on the empirical treatment of policy history and baseline timing. A structural interpretation in terms of irreversibility, uncertainty, or restored sourcing would require information beyond those comparisons.

6. Sensitivity and inference

6.1. Origin exclusions and pre-reset activity

Table 5 evaluates whether the negative unadjusted primary association depends on particular origins or the inclusion of pairs with little pre-reset activity. Excluding China yields −7.75%, with $p = 0.015$. Excluding Brazil yields −12.75%, with $p = 0.001$. The sign therefore does not depend exclusively on either of these origins. Excluding Canada and Mexico yields −7.42%, with a wider interval and $p = 0.135$. These changes indicate that origin composition matters for precision and magnitude. They are separate exclusion exercises, not a complete leave-one-origin-out analysis or a test of structural heterogeneity between the excluded and retained groups.

Table 5. Origin exclusions and pre-reset activity restrictions

	Baseline	Exclude China	Exclude Canada/Mexico	Exclude Brazil	Q4-active pairs
Reset exposure × post	−0.1095***	−0.0807**	−0.0771	−0.1364***	−0.1136***
	(0.0346)	(0.0319)	(0.0505)	(0.0389)	(0.0354)
Association (%)	−10.37	−7.75	−7.42	−12.75	−10.73
$t(G-1)$ p-value	0.003	0.015	0.135	0.001	0.003
Prior-increase control	No	No	No	No	No
All three fixed effects	Yes	Yes	Yes	Yes	Yes
Observations	1,647,643	1,573,673	1,510,835	1,602,312	1,208,457
Origin clusters	44	43	42	43	43
Pseudo R^2	0.9445	0.9463	0.9376	0.9450	0.9420

Notes: PPML with origin–HS6, origin–month and HS6–month fixed effects. The outcome is consumption-import value in millions of current U.S. dollars. Exposures are scaled per 10 percentage points. February 2026 is excluded. Origin-clustered standard errors appear in parentheses. Stars use the $t(G-1)$ reference as a sensitivity calculation: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. This is not bootstrap inference. All columns use the primary window. Q4-active pairs have positive consumption-import value in at least one month of October–December 2025; this restriction uses only pre-reset outcomes. These checks evaluate the unadjusted benchmark. They do not condition on the preceding tariff increase and therefore do not resolve the sensitivity documented in Table 2.

The preferential-trade exclusion is especially relevant to exposure measurement. The model's treatment of preference eligibility and content can differ from realized transaction-level claims. Removing Canada and Mexico changes both that source of model dependence and a substantial part of the trading sample. Its result consequently cannot isolate measurement error alone. More generally, each exclusion allows the remaining observations to be weighted differently by PPML. Differences across these columns should be read as sensitivity to sample composition rather than estimates of how the reset affects one country relative to another.

The Q4-active restriction keeps pairs with positive import value in at least one month of October–December 2025. Unlike the computational screen for all-zero fixed effects, this activity definition uses only outcomes before the reset. It leaves 71,599 eligible pairs and produces a PPML association of −10.73%, with $p = 0.0026$. The resemblance to the baseline indicates that the unadjusted negative association is not confined to pairs that were inactive throughout that pre-period. It does not demonstrate persistence within firms, because the unit remains an aggregate origin–product cell.

The corresponding linear probability model is reported separately in Appendix Table B3. Its outcome is positive reported consumption-import value. A 10-percentage-point larger cut is associated with a 0.787-percentage-point lower probability of positive reported presence, with a 95% interval from −1.655 to 0.081 percentage points and $p = 0.074$. This estimate is not significant at 5%, and it is not conditioned on the preceding tariff increase. It cannot resolve the history sensitivity of the value result. The reporting interpretation of zeros also prevents translating it into a claim about firm exit or the number of destroyed trade relationships.

Together, the origin and activity exercises preserve part of the unadjusted pattern under alternative samples. They do not supply the omitted history comparison. The decisive sensitivity remains the change between otherwise matching columns in Table 2, together with the conditional event tests. Treating all checks as votes for or against a single coefficient would obscure their distinct purposes: some concern origin influence, some concern the universe of active cells, and others concern the policy exposure being compared. Each addresses a different potential explanation for the benchmark association.

6.2. Sample loss and numerical inference

The PPML estimation sample retains 99.874% of eligible primary-window import value, 99.876% in the short window, and 99.878% in the extension. Additional removal after the follow-up's all-zero pair screen leaves the same effective samples as their corresponding unadjusted models. This reconciliation is useful because a large loss of grid rows can otherwise be mistaken for a large loss of observed trade. Most grid exclusions concern all-zero pair effects. The remaining lost positive outcomes and values are separately recorded by origin. Neither the proportion of rows retained nor the high value-retention rate establishes that the retained sample is representative of excluded products or policy exposures.

All reported study fits converge, all requested slope coefficients are present, and the covariance blocks used for the reported joint tests retain their requested rank. Appendix Table C1 documents the event specifications. Coefficient standard errors, probability calculations, exponentiated intervals, and post-period contrasts have been independently reconciled to the saved covariance outputs. These checks establish internal consistency of reporting. They do not establish that the cluster reference distribution is accurate in finite samples or that the exposure comparison is causally valid. The number and dependence structure of origins remain material limitations despite the large number of product–month observations.

7. Economic interpretation and scope

7.1. What the sensitivity reveals about reversal-stage evidence

The results distinguish the size of a measured tariff reversal from the state of trade when the reversal arrives. A cell can have a large net reduction precisely because it accumulated substantial protection earlier. If imports were already adjusting along that earlier exposure gradient, the unadjusted reset interaction combines current and inherited differences. The observed correlation between the two measures and the movement of the fitted coefficients make this concern concrete. They do not identify which part of the earlier adjustment reflects prices, quantities, selection, or expectations. Prior exposure is informative as a conditioning variable while remaining an incomplete description of the economic state.

This distinction changes the interpretation of apparent persistence. The negative unadjusted static coefficient could be consistent with delayed adjustment following earlier protection, but it is also sensitive to how the earlier comparison months enter. Its disappearance after conditioning does not refute dynamic sourcing costs. It shows that this particular coefficient does not isolate them. Likewise, the unadjusted post-period improvement could be consistent with inventories normalizing or sourcing resuming, but that comparison is not stable when history is allowed a flexible monthly association. Evidence of mechanisms requires outcomes that discriminate among those explanations, rather than a preferred label for the same aggregate series.

The conditional estimates also leave economically meaningful uncertainty. Their intervals include positive and negative associations of magnitudes that could matter for particular import relationships. No equivalence bounds were specified and tested, and the analysis does not deliver a precise zero. The contribution is therefore not a finding that tariff cuts are ineffective. It is evidence that a simple reversal-stage comparison can provide an unstable answer when current relief and earlier protection are closely related. Reporting that instability is useful precisely because a stronger claim about recovery or irreversibility would otherwise be easy to attach to the benchmark coefficient.

7.2. Trade values, production networks, and aggregate performance

The conference's broader questions about economic performance require keeping the measured margin distinct from welfare. Structural analyses such as Caliendo and Parro (2015) and Arkolakis et al. (2012) connect trade changes to welfare through explicit assumptions about technology, substitution, expenditure, and equilibrium adjustment. Our conditional import-value coefficients do not supply those objects. They do not quantify consumer incidence, tariff revenue, domestic manufacturing protection, or the terms of trade. The fixed effects also remove common origin and product movements that may be important to aggregate outcomes. Aggregating the fitted exposure coefficient mechanically would therefore be inappropriate.

Production-network research nevertheless gives the question economic relevance. A buyer's response to a lower tariff can depend on suppliers available, contractual relationships already formed, and complementary domestic activities. Such dependence provides a reason for studying reversals separately from increases. The present data can show whether aggregate origin–product values display a stable exposure gradient after conditioning on a summary of prior policy. They cannot reveal whether the same U.S. buyer returned to the same foreign seller, whether a supplier retained productive capacity, or whether a domestic input bottleneck changed. Network resilience remains an economic motivation rather than an estimated outcome.

The same distinction applies to external validity. The inclusion of Southeast Asian origins makes their trade part of the pooled comparison, but it does not yield an ASEAN-wide causal response or a country-specific estimate. Smaller origins can contribute relatively little import value, and one is absent from the short-window PPML sample. Their product mix, preference treatment, and capacity to enter new relationships may differ from those of the largest suppliers. The pooled result is most directly informative about the specified screened basket and exposure comparison. Relevance to diversification in smaller economies depends on examining those additional margins rather than assuming transportability from the aggregate coefficient.

7.3. Evidence that would separate the remaining explanations

A more discriminating empirical account would combine policy and economic histories at a finer level. A transaction-consistent tariff history would distinguish the initial modeled net change from subsequent amendments, exclusions, preference claims, and content rules. Linked buyer–supplier or firm records could distinguish continued trade, switching, reentry, and exit. Compatible quantity and price measures could separate expenditure composition from shipment volume. Inventory information and order dates would help establish whether entry timing reflects a current incentive or an earlier purchasing decision. These are distinct information requirements; simply extending the outcome series does not satisfy them.

The counterfactual also needs explicit treatment. A fuller tariff-history control may absorb more inherited adjustment, but adding controls is not by itself an identification strategy. The parallel-trends sensitivity framework of Rambachan and Roth (2023) provides a way to articulate

departures from a counterfactual restriction in suitable designs. It has not been implemented or validated for the nonlinear, continuous-exposure comparison reported here. Similarly, a different cluster procedure would require demonstrating its applicability to this fixed-effect structure and the source of dependence. Neither an unreported correction nor a favorable alternative pre-window can be assumed to solve the present problem.

The immediate economic interpretation is thus bounded but substantive. Initial tariff relief, inherited protection, and later policy evolution are empirically distinct. Their overlap can change conclusions drawn from short-run import data even on a fixed product universe with extensive fixed effects. A temporary replacement regime makes that distinction especially relevant to the extension, because maintaining the initial exposure comparison through later months requires further policy accounting. This paper supplies evidence on the sensitivity and diagnostic limits of that comparison, clarifying what observed trade adjustment can establish in a changing policy regime.

8. Conclusion

This paper examines U.S. import adjustment around the February 2026 tariff reset using a fixed public-data basket and a frozen model of statutory tariff exposure. The central comparison separates the net reset from the preceding tariff increase. The unadjusted primary-window specification associates a 10-percentage-point larger cut with 10.37% lower post-reset import value. Conditioning on the preceding increase changes the association to 1.36%, with a wide interval that includes zero. Shorter and extended windows likewise provide imprecise conditional estimates. These changes occur on matching estimation samples, rather than through changes in retained trade observations.

Dynamic evidence reinforces the importance of that distinction. Apparent improvements between early and later post-reset months in unadjusted models do not survive freely varying monthly controls for prior exposure. At the same time, conditional event specifications retain pre-reset differences in both follow-up windows. The findings therefore do not establish a causal zero response, irreversible trade damage, or a history-independent recovery. They document that the apparent response to a tariff reversal is sensitive to the policy path and the baseline against which adjustment is measured.

This sensitivity has a clear implication for empirical interpretation. A reversal is not economically characterized by its current rate change alone when the size of relief is closely related to earlier protection. Fixed product weights and rich fixed effects address important measurement and comparison problems, but they do not remove that relationship. Static averages, event coefficients relative to the last pre-month, and changes between post-months retain different reference points. Presenting them together prevents one favorable summary from obscuring evidence that limits another.

The contribution is a transparent reversal-stage comparison with its scope made explicit. The outcome captures reported customs import value in a screened basket, and the extended months do not have an independently reconciled later tariff history. Distinguishing prices, quantities, inventory adjustment, and supplier relationships would require additional information. Within those limits, the paper shows why prior policy exposure belongs at the center of an empirical assessment of tariff reversals. It provides a reproducible basis for further work on trade adjustment without converting specification-sensitive associations into claims about national welfare or long-run supply-chain resilience.

Declarations

Funding. The author received no financial support for the research, authorship, or preparation of this manuscript.

Conflicts of interest/competing interests. The author declares no conflicts of interest or competing interests.

Institutional disclaimer. The views expressed in this manuscript are those of the author and do not necessarily represent the views of the Centre for Strategic and Policy Studies (CSPS), its Board of Directors, or CSPS Management.

Use of artificial intelligence tools. Artificial intelligence tools were used to support language refinement, structural editing, formatting, and presentation of the manuscript. The author retained full responsibility for the conceptual framing, data analysis, interpretation of results, policy discussion, and final content. No AI tool was used as a substitute for authorial judgment, and the author reviewed and approved the final manuscript before submission. ChatGPT also assisted with preparing data-construction and Stata code and checking exported numerical results.

Data availability. U.S. Census international trade data and The Budget Lab at Yale's Tariff Rate Tracker are publicly accessible from the sources cited in the text. The analysis uses the archived v1.2 tariff release, vintage 2026-07-24-09. Analysis files, preparation code, and saved estimation outputs are available from the author on reasonable request. Reproduction requires retaining the frozen source vintage identified in the measurement documentation.

References

Alessandria, G., & Choi, H. (2014). Establishment heterogeneity, exporter dynamics, and the effects of trade liberalization. *Journal of International Economics*, *94*(2), 207–223. https://doi.org/10.1016/j.jinteco.2014.08.006

Alessandria, G., Kaboski, J. P., & Midrigan, V. (2010a). The great trade collapse of 2008–09: An inventory adjustment? *IMF Economic Review*, *58*(2), 254–294. https://doi.org/10.1057/imfer.2010.10

Alessandria, G., Kaboski, J. P., & Midrigan, V. (2010b). Inventories, lumpy trade, and large devaluations. *American Economic Review*, *100*(5), 2304–2339. https://doi.org/10.1257/aer.100.5.2304

Amiti, M., Redding, S. J., & Weinstein, D. E. (2019). The impact of the 2018 tariffs on prices and welfare. *Journal of Economic Perspectives*, *33*(4), 187–210. https://doi.org/10.1257/jep.33.4.187

Amiti, M., Redding, S. J., & Weinstein, D. E. (2020). Who's paying for the US tariffs? A longer-term perspective. *AEA Papers and Proceedings*, *110*, 541–546. https://doi.org/10.1257/pandp.20201018

Antràs, P., Fort, T. C., & Tintelnot, F. (2017). The margins of global sourcing: Theory and evidence from US firms. *American Economic Review*, *107*(9), 2514–2564. https://doi.org/10.1257/aer.20141685

Arkolakis, C., Costinot, A., & Rodríguez-Clare, A. (2012). New trade models, same old gains? *American Economic Review*, *102*(1), 94–130. https://doi.org/10.1257/aer.102.1.94

Baier, S. L., & Bergstrand, J. H. (2007). Do free trade agreements actually increase members' international trade? *Journal of International Economics*, *71*(1), 72–95. https://doi.org/10.1016/j.jinteco.2006.02.005

Baldwin, R., & Krugman, P. (1989). Persistent trade effects of large exchange rate shocks. *The Quarterly Journal of Economics*, *104*(4), 635–654. https://doi.org/10.2307/2937860

Bernard, A. B., Moxnes, A., & Saito, Y. U. (2019). Production networks, geography, and firm performance. *Journal of Political Economy*, *127*(2), 639–688. https://doi.org/10.1086/700764

Boehm, C. E., Flaaen, A., & Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake. *The Review of Economics and Statistics*, *101*(1), 60–75. https://doi.org/10.1162/rest_a_00750

Caliendo, L., & Parro, F. (2015). Estimates of the trade and welfare effects of NAFTA. *The Review of Economic Studies*, *82*(1), 1–44. https://doi.org/10.1093/restud/rdu035

Callaway, B., Goodman-Bacon, A., & Sant'Anna, P. (2024). *Difference-in-differences with a continuous treatment* (Working Paper No. 32117). National Bureau of Economic Research. https://doi.org/10.3386/w32117

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2011). Robust inference with multiway clustering. *Journal of Business & Economic Statistics*, *29*(2), 238–249. https://doi.org/10.1198/jbes.2010.07136

Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, *50*(2), 317–372. https://doi.org/10.3368/jhr.50.2.317

Cavallo, A., Gopinath, G., Neiman, B., & Tang, J. (2021). Tariff pass-through at the border and at the store: Evidence from US trade policy. *American Economic Review: Insights*, *3*(1), 19–34. https://doi.org/10.1257/aeri.20190536

Cavallo, A., Llamas, P., & Vazquez, F. (2025). *Tracking the short-run price impact of U.S. tariffs* (Working Paper No. 34496). National Bureau of Economic Research. https://doi.org/10.3386/w34496

Correia, S., Guimarães, P., & Zylkin, T. (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal: Promoting communications on statistics and Stata*, *20*(1), 95–115. https://doi.org/10.1177/1536867x20909691

Dixit, A. (1989). Hysteresis, import penetration, and exchange rate pass-through. *The Quarterly Journal of Economics*, *104*(2), 205–228. https://doi.org/10.2307/2937845

Executive Office of the President. (2026, February 20). *Imposing a temporary import surcharge to address fundamental international payments problems: Proclamation 11012*. Retrieved October 23, 2026, from https://www.whitehouse.gov/presidential-actions/2026/02/imposing-a-temporary-import-surcharge-to-address-fundamental-international-payments-problems/

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., & Khandelwal, A. K. (2020). The return to protectionism. *The Quarterly Journal of Economics*, *135*(1), 1–55. https://doi.org/10.1093/qje/qjz036

Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., Khandelwal, A. K., & Taglioni, D. (2024). The US-China trade war and global reallocations. *American Economic Review: Insights*, *6*(2), 295–312. https://doi.org/10.1257/aeri.20230094

Fajgelbaum, P. D., & Khandelwal, A. K. (2022). The economic impacts of the US–China trade war. *Annual Review of Economics*, *14*(1), 205–228. https://doi.org/10.1146/annurev-economics-051420-110410

Fajgelbaum, P. D., & Khandelwal, A. K. (2026). *Tariffs in 2025: Short-run impacts on the U.S. economy* (Working Paper No. 35064). National Bureau of Economic Research. https://doi.org/10.3386/w35064

Flaaen, A., Hortaçsu, A., & Tintelnot, F. (2020). The production relocation and price effects of US trade policy: The case of washing machines. *American Economic Review*, *110*(7), 2103–2127. https://doi.org/10.1257/aer.20190611

Handley, K., & Limão, N. (2015). Trade and investment under policy uncertainty: Theory and firm evidence. *American Economic Journal: Economic Policy*, *7*(4), 189–222. https://doi.org/10.1257/pol.20140068

Handley, K., & Limão, N. (2017). Policy uncertainty, trade, and welfare: Theory and evidence for China and the United States. *American Economic Review*, *107*(9), 2731–2783. https://doi.org/10.1257/aer.20141419

Handley, K., & Limão, N. (2022). Trade policy uncertainty. *Annual Review of Economics*, *14*(1), 363–395. https://doi.org/10.1146/annurev-economics-021622-020416

MacKinnon, J. G., Nielsen, M. Ø., & Webb, M. D. (2023). Cluster-robust inference: A guide to empirical practice. *Journal of Econometrics*, *232*(2), 272–299. https://doi.org/10.1016/j.jeconom.2022.04.001

Monarch, R. (2022). "It's not you, it's me": Prices, quality, and switching in U.S.-China trade relationships. *The Review of Economics and Statistics*, *104*(5), 909–928. https://doi.org/10.1162/rest_a_01015

Pierce, J. R., & Schott, P. K. (2016). The surprisingly swift decline of US manufacturing employment. *American Economic Review*, *106*(7), 1632–1662. https://doi.org/10.1257/aer.20131578

Rambachan, A., & Roth, J. (2023). A more credible approach to parallel trends. *Review of Economic Studies*, *90*(5), 2555–2591. https://doi.org/10.1093/restud/rdad018

Roberts, M. J., & Tybout, J. R. (1997). The decision to export in Colombia: An empirical model of entry with sunk costs. *American Economic Review*, *87*(4), 545–564. https://pure.psu.edu/en/publications/the-decision-to-export-in-colombia-an-empirical-model-of-entry-wi/

Roth, J. (2022). Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights*, *4*(3), 305–322. https://doi.org/10.1257/aeri.20210236

Roth, J., Sant'Anna, P. H., Bilinski, A., & Poe, J. (2023). What's trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics*, *235*(2), 2218–2244. https://doi.org/10.1016/j.jeconom.2023.03.008

Santos Silva, J. M. C., & Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, *88*(4), 641–658. https://doi.org/10.1162/rest.88.4.641

Supreme Court of the United States. (2026, February 20). *Learning Resources, Inc. v. Trump, No. 24-1287*. Retrieved October 23, 2026, from https://www.supremecourt.gov/opinions/25pdf/24-1287_4gcj.pdf

The Budget Lab at Yale. (2026, July 24). *Tariff rate tracker: Published model release v1.2, vintage 2026-07-24-09*. Retrieved October 23, 2026, from https://github.com/Budget-Lab-Yale/tariff-rate-tracker/releases/tag/v1.2

U.S. Census Bureau. (n.d.-a). *International trade definitions*. Retrieved October 23, 2026, from https://www.census.gov/foreign-trade/reference/definitions/index.html

U.S. Census Bureau. (n.d.-b). *International trade: Monthly U.S. imports by Harmonized System commodity*. Retrieved October 23, 2026, from https://api.census.gov/data/timeseries/intltrade/imports/hs.html

U.S. Customs and Border Protection. (2026a). *CSMS #67834313: Ending collection of international emergency economic powers act duties*. Retrieved October 23, 2026, from https://content.govdelivery.com/accounts/USDHSCBP/bulletins/40b11c9

U.S. Customs and Border Protection. (2026b). *CSMS #67844987: Imposing temporary section 122 duties*. Retrieved October 23, 2026, from https://content.govdelivery.com/accounts/USDHSCBP/bulletins/40b3b7b

U.S. International Trade Commission. (n.d.). *Harmonized Tariff Schedule information*. Retrieved October 23, 2026, from https://www.usitc.gov/harmonized_tariff_information

Weidner, M., & Zylkin, T. (2021). Bias and consistency in three-way gravity models. *Journal of International Economics*, *132*, Article 103513. https://doi.org/10.1016/j.jinteco.2021.103513

Appendix A. Measurement, coverage, and origins

Table A1 records the variables used in the reported specifications. The full input grid has 44 origins, 4,818 products, and 211,992 pairs. The original 18-month basket and the outcome-only extension are kept distinct. Missing later tariff-history verification is not replaced with an assertion that all earlier monthly policy states have been independently legally audited.

Table A1. Variable definitions and measurement conventions

Variable	Definition and interpretation
Consumption-import value	Monthly Census imports for consumption at origin–HS6 level, in current U.S. dollars. Regressions divide by one million. Customs value excludes U.S. import duties; it is not a quantity measure.
Net reset exposure, D_{cp}	Modeled total rate on February 19 minus February 24, 2026, multiplied by 100. Positive values are net reductions; negative values are increases. The same fixed 2024 HTS10 weights enter both dates.
Preceding increase, R_{cp}	January 2026 minus January 2025 mean modeled tariff, in percentage points. All inputs end before the reset. This summary is not a complete distributed policy history.
Regression exposures, x_{cp}, r_{cp}	$D_{cp}/10$ and $R_{cp}/10$. A unit is a 10-percentage-point larger cut or preceding increase.
Within-HS6 weights	All-origin U.S. 2024 consumption-import shares of positive-weight HTS10 lines within each HS6 product. Weights are common across origins and fixed across model dates.
Clean tariff basket	Complete representation of positive 2024 weights, known free/ad valorem base duties, unchanged base rates across reset dates, and full daily coverage over the original January 2025–June 2026 monthly panel. Specific and compound base duties lacking comparable measurement are excluded.
Monthly modeled rate	Calendar-day mean of daily states, respecting inclusive interval endpoints. Clean monthly rates extend through June 2026; no later monthly rates are imputed.
Post indicator	One from March 2026 within the relevant window. February is excluded.
Event month	Months relative to February 2026. January is -1 and is the reference; March is 1.
Added outcome zero	An absent eligible consumption-import value within a reconciled origin–month. Other financial variables remain missing. This is not a firm-exit indicator.
Q4-active sample	Positive consumption-import value in at least one month of October–December 2025. Eligibility uses only pre-reset outcomes. There are 71,599 eligible pairs.
Reported presence	One for positive separately reported origin–HS6 consumption-import value; zero otherwise. LPM coefficients multiplied by 100 are percentage-point changes.
Computational active-pair screen	At least one positive outcome somewhere in the relevant estimation window, used only to remove all-zero pair effects before PPML. All zero months within active pairs remain. This is distinct from Q4-active sample eligibility.

Table A2 lists the frozen origin universe. Codes are Census origin codes. Ordinary product chapters are 01–97. Origin selection uses the 2024 ranking plus the Southeast Asian augmentation; the product screen is a separate restriction. The extended outcomes do not expand either universe.

Table A2. Frozen origin universe

Code	Origin	Code	Origin
1220	Canada	4890	Turkey
2010	Mexico	5081	Israel
2230	Costa Rica	5170	Saudi Arabia
3010	Colombia	5330	India
3310	Ecuador	5380	Bangladesh
3330	Peru	5460	Myanmar
3370	Chile	5490	Thailand
3510	Brazil	5520	Vietnam
4010	Sweden	5530	Laos
4099	Denmark (excl. Greenland)	5550	Cambodia
4120	United Kingdom	5570	Malaysia
4190	Ireland	5590	Singapore
4210	Netherlands	5600	Indonesia
4231	Belgium	5601	Timor-Leste
4279	France	5610	Brunei
4280	Germany	5650	Philippines
4330	Austria	5700	China
4370	Hungary	5800	South Korea
4419	Switzerland	5830	Taiwan
4550	Poland	5880	Japan
4700	Spain	6021	Australia
4759	Italy	7910	South Africa

Notes: Source: frozen Census origin selection based on 2024 consumption imports plus the Southeast Asian augmentation. Timor-Leste has no retained observations in the short-window PPML fits. Names are abbreviated for display; original Census names and rankings are retained with the replication records.

The original January 2025–June 2026 basket contains USD 3,719,869,130,636 of consumption-import value, compared with USD 4,618,430,957,640 across the selected origins' reported ordinary-chapter products over that period. Their ratio is 80.54%; both totals include February. These coverage totals differ from Table 1, which excludes February and reports window-specific eligible and retained values. The source collection reconciles country–month totals independently before construction of the restricted basket.

The frozen tariff release is v1.2, vintage 2026-07-24-09. Its source archive SHA-256 is 12c967386935cbdac13c78418811d66f8348c09b50f64c6eec6ad2c8d44e0df5. The release manifest records an upstream working tree with uncommitted changes, making the archived release and individual snapshot hashes important identifiers beyond a Git commit. This is a provenance characteristic of the published model, not evidence that each modeled rate equals a realized transaction duty. The supplied audit records preserve the baseline and extension dataset hashes and the full original model provenance.

Appendix B. Additional estimates and contrasts

Table B1 reports the transformed intervals for the three static history-adjusted comparisons. These are the same coefficients as the corresponding columns in Table 2. No additional estimation is involved.

Table B1. Intervals for static associations conditional on prior exposure

Window	Association (%)	95% interval	p-value
Primary, history	1.36	[−6.70, 10.11]	0.744
Short, history	2.29	[−6.40, 11.79]	0.609
Extended, history	1.95	[−6.48, 11.13]	0.654

Notes: Percent transformations of the reset coefficients in columns (2), (4) and (6) of Table 2. All intervals and p-values use the origin-clustered $t(G-1)$ sensitivity reference.

Table B2 provides every conditional post-period comparison saved by the follow-up analysis. In the short window, mean post averages March–June; in the extension it averages March–August. The June-minus-March and May–June-minus-March–April contrasts use identical calendar definitions in both windows. All comparisons are exploratory, and no multiple-testing adjustment has been applied. Their role is to summarize the observed profiles rather than supply independent confirmatory tests.

The reported-presence sensitivity in Table B3 uses a different outcome and estimator from the value specifications. Its pre-reset activity restriction fixes sample eligibility before the reset. It remains an unadjusted exposure comparison and does not identify individual firms or suppliers. The distinction between a zero customs value and the absence of a physical shipment is therefore retained in interpreting the probability coefficient.

Table B2. All conditional post-period contrasts

Window	Contrast	Coefficient	SE	p-value
Short	Mean post vs. January	−0.0277	0.0564	0.626
Short	June minus March	−0.0004	0.0438	0.992
Short	May–June minus March–April	−0.0252	0.0248	0.315
Extended	Mean post vs. January	−0.0144	0.0532	0.787
Extended	June minus March	−0.0070	0.0441	0.874
Extended	May–June minus March–April	−0.0323	0.0251	0.206
Extended	July–August minus March–April	−0.0013	0.0453	0.976

Notes: All seven saved conditional contrasts, including those not displayed in Table 4. The coefficient is a linear combination of log-mean exposure coefficients; standard errors use the full covariance matrix. Inference uses the same origin clusters and $t(G-1)$ sensitivity reference. Exploratory comparisons; p-values are unadjusted.

Table B3. Reported import presence among pre-reset active pairs

	Linear probability model
Reset exposure × post	−0.0079*
	(0.0043)
Association (percentage points)	−0.787
95% interval (percentage points)	[−1.655, 0.081]
$t(G-1)$ p-value	0.074
Observations	1,213,681
Origin clusters	43
R^2	0.5609

Notes: The dependent variable is one for positive separately reported origin–HS6 consumption-import value. The sample is Q4-active pairs in the primary window. All three fixed effects are included, without the prior-increase control. Origin-clustered standard errors are in parentheses. Coefficients are probability changes per 10pp larger cut; multiplying by 100 yields percentage points. The $t(G-1)$ sensitivity reference determines the interval and star: * $p < 0.10$. This outcome does not identify firm exit or physical shipment cessation.

Appendix C. Covariance diagnostics and reproducibility

Table C1 describes numerical checks for the event specifications used in the paper. The constant's normalization can make the full stored covariance have one fewer rank than its dimension even when the requested slope block is full rank. The table accordingly evaluates the slopes and the pre/post restriction blocks directly. Every reported joint test retains the stated number of restrictions. Small negative eigenvalues at floating-point tolerance in the full matrix are distinguished from a deficient requested slope block.

Table C1. Covariance checks for the reported event specifications

Model	Slopes	Slope rank	Pre block full rank	Post block full rank
origin_event	16	16	Yes	Yes
priorpath_event	32	32	Yes	Yes
short_event	7	7	Yes	Yes
short_prior_event	14	14	Yes	Yes
extended_event	18	18	Yes	Yes
extended_prior_event	36	36	Yes	Yes

Notes: Ranks are reconstructed from the exported covariance matrices. The normalization-dependent constant is excluded from the slope rank. Each pre/post test retains all restrictions listed in Table 3. Numerical rank is an inference diagnostic, not an identification test.

An earlier two-way-clustered event covariance in the source archive has dimension 17 and numerical rank 11, with incomplete support for the intended unrestricted event test. This paper does not use that covariance to support a joint inference claim. Origin-clustered specifications were estimated and audited separately. They maintain a different assumption about remaining dependence across origins, as discussed in Section 4; their usable numerical covariance does not establish that the assumption is satisfied. No validated wild-cluster bootstrap or full leave-one-origin-out series is claimed.

The reported post-period contrasts use $a'\widehat{V}a$ from each model's saved covariance, rather than summing marginal coefficient variances. Probability calculations and transformed intervals have been reconciled to those files. The paper's tables are assembled directly from the completed estimation outputs, with the original result files retained unchanged.

Reproduction has two distinct stages. Data preparation requires the frozen public Census collection and tariff release, the fixed basket and weights, and the stated construction rules. Estimation requires the supplied Stata preparation and analysis files, with their documented software versions and the same input datasets. The manuscript source package additionally includes the CSV outputs used for every table, a model inventory, and numerical reconciliation records. These permit checking the paper's reported estimates without rerunning a large PPML model. They do not replace a complete source-data reconstruction when the objective is independent replication from public inputs.