

Online Appendix:

“Tariff Front-Running” by Carreras-Valle and Lee

This online appendix contains the following items.

A. Additional Tables and Figures.

- Figure A1: the share of product-country partners and of firms affected by tariff increases over time.
- Figure A2: the inventory response by inventory type.
- Figure A3: the inventory response under alternative industry fixed effects.
- Figure A4: the robustness of the firm-level responses (Figure 3 of the main text) to the firm-level tariff-increase threshold of 0.01, 0.1, and 1 p.p.
- Figure A5: the robustness of the firm-level responses (Figure 3 of the main text) to defining the existing set of partners over the preceding two, four, and six quarters.

B. Data and Empirical Analyses.


- **B.1. Linking Bill-of-Lading Shipments to Each Firm.** The procedure that links the bill-of-lading shipments to Compustat firms.
- **B.2. Per-Wave Estimation.** Wave-by-wave estimates of the firm-level responses of imports, the number of partners, and inventories, and the weights with which the six waves aggregate to the pooled estimates of the main text.
- **B.3. Import Growth of Treated and Never-Treated Firms Before the Tariff Waves.** A specification check: cohort by cohort, the 2015–2016 import growth of the firms later treated versus the never-treated firms.
- **B.4. Magnitude of Import Adjustments.** Descriptive statistics behind the partner-level results of Section 3.3 of the main text: the import volumes of the affected products from the tariff-affected and the unaffected source countries in the year before and the year of the anticipatory adjustment, and the transitions in the number of source countries, against a never-treated benchmark.

C. Model.

- **C.1. Proofs of the Propositions.** The proofs of Propositions 1 and 3 of the main text, and the interval-wise formulation of the inventory manager’s problem implied by Proposition 2.
- **C.2. Solving the Transition Paths.** The algorithm that solves for the perfect-foresight transition path of aggregate output, prices, and wages.
- **C.3. Sensitivity Analyses for the Aggregate Dynamics.** The transition of aggregate output and prices under alternative values of the predetermined parameters.


A Additional Tables and Figures

Figure A1: Shares of product-country partners and firms affected by tariff increases




Note: Trade partners are defined at the product-country level, where products are classified at the HS6 level. The green dashed line (left axis) is the share of partners with a tariff increase of at least 1 p.p. The blue solid line (right axis) is the share of firms with a tariff increase, defined as an import-volume-weighted tariff increase of at least 0.1 p.p.

Figure A2: Inventory response by inventory type




Note: Each panel reports the estimated $\{\beta_k\}$ in Equation (2) of the main text for an inventory measure: (a) total inventories (the sum of raw materials, work-in-process, and finished goods), and its components (b) finished goods, (c) raw materials, and (d) work-in-process. The dark and light shaded areas are the ± 1 and ± 1.96 standard-error bands (68% and 95% confidence intervals), respectively. Standard errors are clustered at the firm level.

Figure A3: Inventory response under alternative industry fixed effects




Note: Each panel reports the estimated $\{\beta_k\}$ in Equation (2) of the main text for total inventories under alternative industry fixed effects: (a) the baseline NAICS 6-digit $\times$ event fixed effects, (b) NAICS 4-digit, (c) NAICS 2-digit, and (d) no industry fixed effects (2016 firm controls only). The panels share a common vertical axis. The dark and light shaded areas are the ± 1 and ± 1.96 standard-error bands (68% and 95% confidence intervals), respectively. Standard errors are clustered at the firm level.

Figure A4: Robustness of the firm-level results (Figure 3 of the main text) to the firm-level tariff-increase threshold



Note: Each panel reports the estimated $\{\beta_k\}$ in Equation (2) of the main text with the firm-level tariff-increase threshold set to (a) 0.01 p.p., (b) 0.1 p.p. (the baseline), and (c) 1 p.p., for (i) imports, (ii) the number of product-country partners, and (iii) inventories. At each threshold, cohorts and the excluded below-threshold group are re-classified under the same design. Panels in each row share a common vertical axis. The dark and light shaded areas are the ± 1 and ± 1.96 standard-error bands (68% and 95% confidence intervals), respectively. Standard errors are clustered at the firm level.

Figure A5: Robustness of the firm-level results (Figure 3 of the main text) to the length of the existing-set window



Note: Each panel reports the estimated $\{\beta_k\}$ in Equation (2) of the main text with the existing set of partners, $E_{i,t}$, and the existing import volumes, $v_{i,h,c,t}$, defined over the preceding (a) two, (b) four (the baseline), and (c) six quarters, for (i) imports, (ii) the number of product-country partners, and (iii) inventories. At each window length, the firm-level tariff increase, the cohorts, and the excluded below-threshold group are re-computed under the same design. Panels in each row share a common vertical axis. The dark and light shaded areas are the ± 1 and ± 1.96 standard-error bands (68% and 95% confidence intervals), respectively. Standard errors are clustered at the firm level.

B Data and Empirical Analyses

B.1 Linking Bill-of-Lading Shipments to Each Firm

We merge the BoL data with Compustat through text-analysis algorithms. The main challenge is that firm names vary both across shipments in the BoL data and between the BoL and Compustat datasets. We merge the BoL data with Compustat using a combination of firm identifiers and text-based name matching. We first assign shipments to firms using the identifier provided by PIERS, which links many shipments directly. For the remaining shipments, we use approximate name matching. Specifically, we standardize each firm name by stripping legal suffixes (e.g., Inc., Corp., Ltd.) and common industry terms (e.g., Shipping, Trading, Group); we call what remains the name’s “core.” We then compare cores using Levenshtein similarity and require a score of at least 90% for a match. To set this cutoff, we hand-labeled 500 randomly selected firm pairs and chose the threshold that maximized matching accuracy against these labels. This procedure increases the number of matched observations by more than 20% relative to using the PIERS identifiers alone. Finally, we link BoL firms to Compustat firms using the same name-standardization and matching procedure, again requiring a Levenshtein similarity score of at least 90%.

B.2 Per-Wave Estimation

Our baseline estimates in Equation (2) of the main text pool the six tariff waves. Here we report the estimates wave by wave. For each event quarter g and horizon k , we estimate, on the sample of cohort- g and never-treated firms,

$$\Delta \tilde{y}_{i,g+k} = \beta_{g,k} D_{i,g} + \Gamma'_{g,k} X_i^{2016} + \alpha_{j(i),g,k} + \epsilon_{i,g,k}, \quad (\text{B1})$$

which is Equation (2) with a wave-specific coefficient $\beta_{g,k}$. Figure B1 reports the per-wave responses of imports, the number of partners, and inventories, one wave per row. The anticipatory build-up appears across the waves; the per-wave paths are noisier than the pooled path, most visibly for the two 2019 cohorts (the List 3 escalation and List 4A), which have the fewest treated firms (128 and 75, and 60 and 30 in the inventory sample) and the shortest post-event horizons.

Because every term other than β_k in Equation (2) is event-specific, the pooled coefficient is exactly a convex combination of the per-wave coefficients. By the Frisch–Waugh–Lovell

theorem,

$$\beta_k = \sum_g w_{g,k} \beta_{g,k}, \quad w_{g,k} = \frac{\sum_{i \in S_{g,k}} \tilde{D}_{i,g,k}^2}{\sum_{g'} \sum_{i \in S_{g',k}} \tilde{D}_{i,g',k}^2}, \quad (\text{B2})$$


where $S_{g,k}$ is the wave- g estimation sample at horizon k and $\tilde{D}_{i,g,k}$ is the treatment indicator residualized on the controls and industry fixed effects within $S_{g,k}$. If there were no firm-level controls or industry-event-horizon fixed effects, only a wave-specific intercept, $\sum_i \tilde{D}^2$ would reduce to $n_{g,k} p_{g,k}(1 - p_{g,k})$, yielding the familiar group-size-times-treatment-variance weights of stacked designs (Cengiz et al. 2019). At $k = -1$, the weights ($w_{g,-1}$) across the six waves (2018Q1 through 2019Q3) are (0.173, 0.126, 0.292, 0.269, 0.089, 0.051) for imports and the number of partners, and (0.191, 0.103, 0.310, 0.265, 0.086, 0.045) for inventories, whose estimation sample is restricted to firms reporting inventory components. The weighted combinations in Equation (B2) reproduce the pooled estimates of Equation (2). Two features are worth noting: the two List 3-related waves and Lists 1–2 carry about two-thirds of the weight, and no single wave dominates the pooled path, the largest weight, on Lists 1–2, being below one-third.

Lastly, regarding the specification check mentioned in Section 3.2 of the main text, we highlight that among fifteen coefficients at $k = -6$ and across the baseline and four waves that yield the $k = -6$ estimates, only two were statistically significant at the 5% level.

B.3 Import Growth of Treated and Never-Treated Firms Before the Tariff Waves

This appendix conducts the specification check discussed at the end of Section 3.2 of the main text. We ask whether the firms that would later face tariff increases were already changing their imports before the higher tariff was in sight. For each firm, we compute the growth of its annual imports from 2015 to 2016 (in asinh differences) and regress it, cohort by cohort, on the treatment indicator over the corresponding estimation sample of Equation (2) (the cohort’s treated firms and the never-treated controls with non-missing 2016 characteristics). Both years precede the earliest trade-war investigation (April 2017), so any difference could reflect a pre-existing trend rather than anticipation. Table B1 shows that the treated–control differences are statistically indistinguishable from zero in every cohort. The last row replaces the treatment indicator with the continuous measure $D_{i,g}MS_i$ of Appendix B of the main text, pooled across cohorts, with the same result.

Figure B1: Per-wave responses of imports, the number of partners, and inventories



Note: Each panel reports the estimated $\{\beta_{g,k}\}$ in Equation (B1) for one event quarter (rows) and one outcome (columns): (a) current imports, (b) the number of current product-country partners, and (c) inventories. Each cohort is traced over its available horizons (the sample ends in 2019Q4); vertical axes differ across panels. The dark and light shaded areas are the ± 1 and ± 1.96 standard-error bands (68% and 95% confidence intervals), respectively. Standard errors are clustered at the firm level.

Table B1: Import growth before the anchor year: treated versus never-treated firms

Cohort	Treated firms	Coefficient	S.E.	<i>p</i> -value	Obs.
2018Q1	282	0.12	0.28	0.66	1,386
2018Q2	286	0.19	0.22	0.40	1,390
2018Q3	722	0.24	0.20	0.23	1,826
2018Q4	483	0.24	0.26	0.34	1,587
2019Q2	128	0.20	0.46	0.66	1,232
2019Q3	75	0.78	0.57	0.17	1,179
Continuous measure, pooled	1,976	1.78	1.28	0.16	3,030

Notes: Each row reports a separate regression of the firm’s 2015–2016 import growth, $\text{asinh}(\text{imports}_{i,2016}) - \text{asinh}(\text{imports}_{i,2015})$, on the cohort’s treatment indicator. The sample pools the cohort’s treated firms with the never-treated controls. Firms with no imports in 2015–2016 enter with zero growth. The last row pools all cohorts and replaces the indicator with the continuous measure $D_{i,g}MS_i$ of Appendix B of the main text. Standard errors are heteroskedasticity-robust, and *p*-values are two-sided.

B.4 Magnitude of Import Adjustments

In Section 3.3 of the main text, we uncovered that, in the anticipation of higher tariffs, firms adjust their imports along the within-product dimension: both the extensive- and intensive-margin responses are significant for the affected product–affected country and affected product–unaffected country partners (types (i) and (ii) in Figure 4 of the main text). In this section, we assess the magnitude of these adjustments: how much of an affected product a firm imports from the countries about to face the tariff and from the other countries in the year before the tariff, and how its number of source countries changes.

To address this question, we compare the import volumes of a firm-product across the two years preceding its tariff increase. Consider firm i that imports product h . Let g_{ih} denote the first event quarter at which some country c supplying product h to firm i (i.e., $(h, c) \in E_{i,g}$) experiences a tariff increase, $g_{ih} \equiv \min\{g \in \mathcal{G} : TI_{i,h,c,g} = 1 \text{ for some } c\}$, with $TI_{i,h,c,g}$ as defined in Section 3.3 of the main text. In our main empirical results (Figures 3 and 4 of the main text), the anticipatory adjustments build over the year preceding the event. Therefore, we compare the *adjustment year*, quarters $g_{ih} - 4$ through $g_{ih} - 1$, with the *pre-adjustment year*, $g_{ih} - 8$ through $g_{ih} - 5$, which precedes the anticipatory build-up. For each firm-product pair (i, h) , we split the source countries into *tariff-affected* countries—those for which the tariff on (h, c) increased at g_{ih} —and *unaffected* countries, and sum the firm-product’s imports from each type over each year.

On the intensive margin, Table B2 reports the quartiles of these import volumes across firm-products. Panel (a) covers the affected firm-products with positive imports of the product

in both years.¹ At the median, a firm-product imported \$77k of the product from the tariff-affected countries in the pre-adjustment year and \$107k in the adjustment year. Imports from unaffected countries are zero at the median in both years, as fewer than half of the affected firm-products source the product from an unaffected country in either year (46% and 42%), consistent with the prevalence of single-country sourcing (Table A1 of the main text). Total imports of the product rise from \$148k to \$176k at the median. Panel (b) restricts attention to the firm-products that source the product from both types of countries in both years, for which the two channels can be compared within the firm-product: the median imports from tariff-affected countries rise from \$284k to \$329k, and those from unaffected countries from \$239k to \$257k.

Table B2: Import volumes of the affected products by source-country type (USD thousand)

Source countries	Pre-adjustment year			Adjustment year		
	Q1	Median	Q3	Q1	Median	Q3
<i>(a) All firm-products (Nobs. 36,309)</i>						
Tariff-affected	9	77	412	23	107	487
Unaffected	0	0	80	0	0	77
All	31	148	728	37	176	853
<i>(b) Sourcing from both types of countries (Nobs. 9,079)</i>						
Tariff-affected	58	284	1,374	68	329	1,523
Unaffected	48	239	1,157	48	257	1,306
All	225	872	3,454	243	964	3,891

Notes: Each cell reports a quartile, across the affected firm-products of treated firms, of the firm-product's imports of the product from the given type of source countries over the year. Panel (a): firm-products with positive imports of the product in both years. Panel (b): the subset with positive imports from both tariff-affected and unaffected countries in both years. Imports from tariff-affected countries are positive in the adjustment year by construction, since that year defines existing-set membership at g_{ih} .

On the extensive margin, let $PA_{i,h}$ and $DA_{i,h}$ denote the sets of countries from which firm-product (i, h) imports during the pre-adjustment year and the adjustment year, respectively. For each affected firm-product of a treated firm, we count the number of source countries in $PA_{i,h}$ and in $DA_{i,h}$ and tabulate the transitions from $\{0, 1, 2, 3+\}$ pre-adjustment countries to $\{1, 2, 3+\}$ adjustment-year countries (Table B3, Panel (a)). Panel (b) provides the benchmark that mirrors the comparison in Equation (3) of the main text: the never-treated firms'

¹Import flows are lumpy: about half of the affected firm-products record no imports of the product in the pre-adjustment year, whereas every affected firm-product imports in the adjustment year by construction. Including these firm-products would make the comparison across the two years reflect whether a shipment happened to fall in the pre-adjustment year rather than how import volumes adjusted. We therefore restrict attention to the firm-products with positive imports in both years.

partners for the same products over the same event windows, subject to the same selection (products in the firm’s existing set at the event) and reweighted to the event composition of Panel (a). The modal affected product records no active source during the pre-adjustment year—import flows are lumpy—and a single source is the mode among the remainder. A firm-product exhibits at least one source country during the adjustment year by construction: that year coincides with the quarters defining existing-set membership at g_{ih} , so a count of zero cannot occur in either panel. Among affected products sourced from a single country in the pre-adjustment year, 22% expand to two or more countries during the adjustment year, against 10% of the benchmark products; among those with no active source, 10% resume with two or more countries, against 4%. The pattern holds in every row of the matrix.

Table B3: Transitions in the number of source countries (row shares, %)

Countries in $PA_{i,h}$	Nobs.	Countries in $DA_{i,h}$		
		1	2	3+
<i>(a) Affected products of treated firms</i>				
0	34,548	90.0	8.0	2.1
1	22,793	78.5	16.4	5.2
2	6,683	35.7	41.0	23.2
3+	6,885	8.3	17.3	74.3
<i>(b) Benchmark: never-treated firms</i>				
0	1,772	96.3	3.6	0.3
1	706	90.8	7.8	2.0
2	91	59.5	35.3	6.9
3+	60	23.7	25.7	57.3

Notes: Each row reports the distribution of the number of source countries during the adjustment year ($DA_{i,h}$) for firm-products with the given number of source countries in the pre-adjustment year ($PA_{i,h}$). Shaded cells denote the transitions that add source countries. There is no zero column: a count of zero source countries during the adjustment year cannot occur by construction. Panel (a): affected firm-products of treated firms at g_{ih} . Panel (b): never-treated firms’ partners for the same products over the same event windows, with row shares reweighted to the event composition of Panel (a).

C Model

C.1 Proofs of the Propositions

In this section, we provide the proofs of the propositions in the main text. For notational ease, let $\Xi_c \equiv T_c (\tau_c w_c)^{-\theta}$ denote the *sourcing potential* of country c . Given w , f , θ , and ρ ,

the import manager's problem in Equation (10) of the main text depends on the tariffs only through $\{\Xi_c\}_{c \in C}$, and the marginal cost of the foreign input in Equation (9) of the main text, $z_i(\mathcal{C}_i) = \gamma \left(\sum_{c \in \mathcal{C}_i} \Xi_c \right)^{-1/\theta}$, is decreasing in the sourcing potential of every country in the sourcing set.

Proposition 1. *Given the input demand n_i , the optimal sourcing set $\mathcal{C}_i(n_i)$ satisfies the following properties:*

- (i) *If $c \in \mathcal{C}_i(n_i)$ and $\Xi_c \leq \Xi_{c'}$, then $c' \in \mathcal{C}_i(n_i)$.*
- (ii) *If $n_{\text{low}} < n_{\text{high}}$, then $\mathcal{C}_i(n_{\text{low}}) \subset \mathcal{C}_i(n_{\text{high}})$, and $z_i(\mathcal{C}_i(n_{\text{low}})) \geq z_i(\mathcal{C}_i(n_{\text{high}}))$.*

Proof of Proposition 1. (i) Suppose not: $c \in \mathcal{C}_i(n_i)$, $\Xi_c \leq \Xi_{c'}$, and $c' \notin \mathcal{C}_i(n_i)$. Consider the alternative sourcing set $(\mathcal{C}_i(n_i) \setminus \{c\}) \cup \{c'\}$. It contains the same number of countries as $\mathcal{C}_i(n_i)$, hence incurs the same fixed cost, and it delivers a weakly lower marginal cost, since $\Xi_{c'} \geq \Xi_c$. Therefore, it fulfills the input demand n_i at a weakly lower total cost, which contradicts the optimality of $\mathcal{C}_i(n_i)$ (ties are broken in favor of the country with the higher sourcing potential).

(ii) Suppose that $\mathcal{C}_i(n_{\text{low}}) \not\supseteq \mathcal{C}_i(n_{\text{high}})$. (By part (i), both sets consist of the top-ranked countries in sourcing potential, so the only possible cases are $\mathcal{C}_i(n_{\text{low}}) \supsetneq \mathcal{C}_i(n_{\text{high}})$, $\mathcal{C}_i(n_{\text{low}}) = \mathcal{C}_i(n_{\text{high}})$, and $\mathcal{C}_i(n_{\text{low}}) \subsetneq \mathcal{C}_i(n_{\text{high}})$.) Since $\mathcal{C}_i(n_{\text{low}})$ is the optimal sourcing set for the input demand n_{low} ,

$$z_i(\mathcal{C}_i(n_{\text{low}})) n_{\text{low}} + |\mathcal{C}_i(n_{\text{low}})| wf \leq z_i(\mathcal{C}_i(n_{\text{high}})) n_{\text{low}} + |\mathcal{C}_i(n_{\text{high}})| wf. \quad (\text{C1})$$

Likewise, for n_{high} ,

$$z_i(\mathcal{C}_i(n_{\text{high}})) n_{\text{high}} + |\mathcal{C}_i(n_{\text{high}})| wf \leq z_i(\mathcal{C}_i(n_{\text{low}})) n_{\text{high}} + |\mathcal{C}_i(n_{\text{low}})| wf. \quad (\text{C2})$$

Adding Equations (C1) and (C2) yields

$$(z_i(\mathcal{C}_i(n_{\text{high}})) - z_i(\mathcal{C}_i(n_{\text{low}}))) (n_{\text{high}} - n_{\text{low}}) \leq 0. \quad (\text{C3})$$

Since $n_{\text{high}} > n_{\text{low}}$, Equation (C3) implies that $z_i(\mathcal{C}_i(n_{\text{high}})) \leq z_i(\mathcal{C}_i(n_{\text{low}}))$: the marginal cost under $\mathcal{C}_i(n_{\text{high}})$ is weakly lower than under $\mathcal{C}_i(n_{\text{low}})$. Since the fixed cost is strictly lower under $\mathcal{C}_i(n_{\text{high}})$ as well (as $|\mathcal{C}_i(n_{\text{low}})| > |\mathcal{C}_i(n_{\text{high}})|$), the total cost of fulfilling n_{low} is strictly lower with $\mathcal{C}_i(n_{\text{high}})$ than with $\mathcal{C}_i(n_{\text{low}})$. This contradicts the optimality of $\mathcal{C}_i(n_{\text{low}})$ for the input demand n_{low} . Hence, $\mathcal{C}_i(n_{\text{low}}) \subset \mathcal{C}_i(n_{\text{high}})$, and $z_i(\mathcal{C}_i(n_{\text{low}})) \geq z_i(\mathcal{C}_i(n_{\text{high}}))$ follows because the sum of the sourcing potentials is weakly larger over the larger set. $\square$

Proposition 3. Suppose that for country c , its tariff strictly increases from τ_c to τ'_c , and let $\boldsymbol{\tau}$ and $\boldsymbol{\tau}'$ denote the tariff vectors before and after this increase, respectively. Then, for all $n_i \geq 0$, $r_i(n_i, \boldsymbol{\tau}') \geq r_i(n_i, \boldsymbol{\tau})$.

Proof of Proposition 3. Before the tariff change, the import manager solves the cost-minimization problem in Equation (10) of the main text given the sourcing potentials $\{\Xi_1, \dots, \Xi_C\}$. After the tariff change, the sourcing potential of country c falls from Ξ_c to $\Xi'_c = T_c (\tau'_c w_c)^{-\theta} < \Xi_c$, while those of the other countries are unchanged, so the import manager solves the same problem given $\{\Xi_1, \dots, \Xi_{c-1}, \Xi'_c, \Xi_{c+1}, \dots, \Xi_C\}$. For any candidate sourcing set $\mathcal{C}_i$, the fixed cost is unchanged and the marginal cost $z_i(\mathcal{C}_i)$ is weakly higher under the new sourcing potentials (strictly higher if $c \in \mathcal{C}_i$). The minimized total cost, $r_i(n_i, \boldsymbol{\tau}') n_i$, is therefore weakly higher than $r_i(n_i, \boldsymbol{\tau}) n_i$ for every $n_i \geq 0$. $\square$

Last, Proposition 2 of the main text allows us to solve the inventory manager's dynamic problem, Equation (11) of the main text, interval by interval. Specifically, we solve C value functions, one for each interval of input demand, $[n_k, n_{k+1})$, over which the marginal cost of the foreign input is constant at z_k . The value function is the upper envelope of these interval-specific value functions,

$$V(s_i, \nu_i, \vec{\tau}) = \max_{k \in \{1, \dots, C\}} V_k(s_i, \nu_i, \vec{\tau}),$$

where, for each k ,

$$\begin{aligned} V_k(s_i, \nu_i, \vec{\tau}) &= \max_{p_i, y_i, s'_i, n_i} p_i y_i - r_i(n_i, \boldsymbol{\tau}) n_i + \beta \mathbb{E}_{\nu'_i} [V(s'_i, \nu'_i, \vec{\tau}')] \\ \text{s.t. } y_i &= \nu_i p_i^{-\epsilon} (1 - \mu)^{\frac{\epsilon}{\sigma}} Y^{\frac{\epsilon}{\sigma}} P^\epsilon y_M^{1 - \frac{\epsilon}{\sigma}}, \\ y_i &= q_i, \\ q_i &\leq s_i, \\ s'_i &= (1 - \delta)(s_i - q_i + n_i), \\ n_i &\in [n_k, n_{k+1}), \\ \frac{\partial(r_i(n_i, \boldsymbol{\tau}) n_i)}{\partial n_i} &= z_k. \end{aligned}$$

Within each interval, the marginal cost of the foreign input is a constant, which simplifies the value function iteration.

C.2 Solving the Transition Paths

In this section, we detail the algorithm used to solve for the perfect-foresight path of aggregate output, prices, and wages, $\{Y_t, P_t, w_t\}_{t=1}^T$, and the resulting distribution of firms over inventories. We index the periods of the transition by $t = 1, \dots, T$; in the timing of the main text, the news arrives at $t_{\text{news}} = -3$ and the tariffs rise at $t_\tau = 0$. The aggregates at each date are pinned down by three conditions in the three unknowns (Y, P, w) : labor market clearing, the aggregate output in Equation (4) of the main text, and the aggregate price in Equation (6) of the main text,

$$\text{labor market clearing: } y_d = \ell = L - f \int_0^1 |\mathcal{C}_i| di, \quad (\text{C4})$$

$$\text{aggregate output: } Y = (\mu^{1/\sigma} y_d^a + (1 - \mu)^{1/\sigma} y_M^a)^{1/a}, \quad a = \frac{\sigma-1}{\sigma}, \quad (\text{C5})$$

$$\text{aggregate price: } P = (\mu w^{1-\sigma} + (1 - \mu) p_M^{1-\sigma})^{1/(1-\sigma)}, \quad (\text{C6})$$

where $p_d = w$, and the final-good bundle and its price index aggregate across firms,

$$y_M = \left(\int_0^1 \nu_i^{1/\epsilon} y_i^{\frac{\epsilon-1}{\epsilon}} di \right)^{\frac{\epsilon}{\epsilon-1}}, \quad p_M = \left(\int_0^1 \nu_i p_i^{1-\epsilon} di \right)^{\frac{1}{1-\epsilon}}.$$

The demand for each final good in Equation (11) of the main text depends on the aggregates only through the demand shifter $\Lambda \equiv (1-\mu)^{\epsilon/\sigma} P^\epsilon Y^{\epsilon/\sigma} y_M^{1-\epsilon/\sigma}$, so that $y_i = \nu_i p_i^{-\epsilon} \Lambda$. The wage is recovered by inverting the consumer's demand for the domestic good, $w = (\mu P^\sigma Y / y_d)^{1/\sigma}$.

Step 0, initial guess. Let $\{Y_t^p, P_t^p, w_t^p\}_{t=1}^T$ be the current guess for the paths, initialized at the initial, low-tariff steady-state values for $t < t_\tau$ and at the final, high-tariff steady-state values for $t \geq t_\tau$. The tariffs change once, at t_τ .

Step 1, backward pass. Set $V_{T+1} = V^{\text{final}}$, the value function in the final, high-tariff steady state. For $t = T, T-1, \dots, t_{\text{news}} + 1$, evaluate the demand shifter implied by the guessed path, and apply the Bellman operator of Equation (11) of the main text, denoted by $\mathcal{B}$, *once*:

$$\Lambda_t = \Lambda(Y_t^p, P_t^p, w_t^p), \quad V_t = \mathcal{B}(V_{t+1}; \Lambda_t, w_t^p, \{z_k(\tau_t)\}_{k=1}^C),$$

where $\{z_k(\tau_t)\}_{k=1}^C$ are the marginal costs of Proposition 2 of the main text evaluated at the date- t tariffs, which take one of two precomputed values, before and after t_τ . Before the news arrives, the economy is in the initial, low-tariff steady state.

Step 2, forward sweep with within-date clearing. Start from the distribution of

inventories in the initial steady state, s_1 . For $t = 1, \dots, T$, solve

$$\mathcal{R}_t(Y, P, w) = \begin{pmatrix} \log \hat{Y}(Y, P, w) - \log Y \\ \log \hat{P}(Y, P, w) - \log P \\ \log \hat{w}(Y, P, w) - \log w \end{pmatrix} = 0,$$

where $(\hat{Y}, \hat{P}, \hat{w})$ are the values implied by Equations (C4) to (C6) after simulating every firm's decisions at (Y, P, w) , given the inherited inventories s_t and the date- t draw of demand shocks. We solve this three-dimensional root-finding problem for $(\log Y, \log P, \log w)$ with Broyden's method. Then, given the policies implied by the aggregates, we roll the distribution of inventories forward, $s_{t+1} = (1 - \delta)(s_t - q_t + n_t)$.

Step 3, outer update. Compare the new paths with the guessed ones in logs and take a damped step,

$$\log X_t^p \leftarrow \log X_t^p + \varrho (\log X_t - \log X_t^p), \quad X \in \{Y, P, w\},$$

with relaxation parameter $\varrho = 0.5$. We repeat Steps 1 through 3 until the paths converge, that is, until

$$\max_{t \leq T} \max \{ |\Delta \log Y_t|, |\Delta \log P_t|, |\Delta \log w_t| \} \leq \text{tol},$$


where Δ denotes the change between successive iterations.

C.3 Sensitivity Analyses for the Aggregate Dynamics

The effect of firms' anticipatory response on aggregate output and prices is robust to the different predetermined parameter values. Figure C1 shows the transition of aggregate output and price for different values of the elasticity of substitution across foreign input varieties, ρ ; the Fréchet parameter governing the dispersion of productivity draws, θ ; the elasticity of demand across final-good varieties, ϵ ; the storage cost of inventories, δ ; and the elasticity of substitution between the domestic good and the final-good bundle, σ .

A lower value of ρ decreases the import substitutability and increases the firm's incentives to anticipate the tariff. Similarly, a higher value of the Fréchet parameter, θ , reduces the volatility within the range of imports across countries, which lowers the benefit of adding an additional country to the sourcing set. Both limit the firm's ability to substitute away from the tariff, which creates further incentives to protect their production and have a smooth transition to the new steady state. Lowering the demand elasticity, ϵ , which increases the markup the firm sets, lowers the total response of output due to higher markups. A lower

Figure C1: Robustness exercises to different parameter values



Note: The figures show the index for the transition path of aggregate output and price for different values of the elasticity of substitution across foreign input varieties, ρ ; the Fréchet parameter governing the dispersion of productivity draws, θ ; the elasticity of demand across final-good varieties, ϵ ; the storage cost of inventories, δ ; and the elasticity of substitution between the domestic good and the final-good bundle, σ .

cost of inventory storage, δ , allows for a deeper anticipatory response that leads to a larger response of output and prices pre-tariff and a smoother transition afterward.

References

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. “The Effect of Minimum Wages on Low-Wage Jobs.” *The Quarterly Journal of Economics*, 134(3): 1405–1454.