

Tariffs, Rules of Origin, and the Reorganization of North American Supply Chains*

Ángel García-Lascurain Fernández[†]

Arizona State University

October 2026

[\[Click here for the latest version\]](#)

Abstract

This paper quantifies how tariffs and rules of origin drive the reorganization of North American supply chains around Mexico, and what that reorganization implies for welfare. Between 2018 and 2025 three shocks reshaped this landscape: the U.S.-China trade war, the 2020 NAFTA-to-USMCA transition, and a 2025 U.S. executive order sharply raising tariffs on Mexican imports outside USMCA preference. I develop a global sourcing model in which heterogeneous firms choose source countries and, product by product, whether to export under USMCA, thus having to comply with rules of origin, and calibrate it on matched census-customs microdata. Through the lens of the model, the transition to USMCA was roughly welfare-neutral for Mexico, while the shift in relative sourcing efficiency across countries raised welfare. The higher tariffs on non-compliant goods imposed under the second Trump administration, along with a Mexican tariff on source countries outside its free-trade network, both lower welfare by raising the effective cost of choosing not to comply with USMCA. Eliminating the agreement altogether would impose a substantial welfare loss on Mexico, one that grows sharply with the tariff on non-compliant goods.

JEL codes: F13, F14, F15, L23.

Keywords: rules of origin, USMCA, global sourcing, heterogeneous firms.

*I am grateful to Wyatt Brooks, Samuel Young, Victor Zuluaga and José Ramón Morán for the helpful comments and feedback on this project. I also thank Natalia Volkow and the Microdata Laboratory personnel at the Mexican National Institute of Statistics and Geography (INEGI) for facilitating the access to the data. All analysis of confidential data was carried out on-site in the INEGI Microdata Laboratory in Mexico City in accordance with Mexican laws. All errors are my own.

[†]Department of Economics, W.P. Carey School of Business, Arizona State University, (e-mail: agarc475@asu.edu)

1 Introduction

Over the last several years North American trade policy has become increasingly complex. For nearly three decades, trade in the region operated under the North American Free Trade Agreement (NAFTA), which guaranteed frictionless commerce among the United States, Canada and Mexico as long as a sufficient share of each product was sourced within the region. Firms in turn built their continental supply chains around this certainty. This is no longer the case. Since 2018, a bilateral tariff war between the United States and China, the replacement of NAFTA by the more restrictive United States-Mexico-Canada Agreement (USMCA), and a sequence of emergency U.S. tariffs on goods that fall short of USMCA’s regional content requirements have substantially altered what it costs to produce in Mexico and export to the other two nations. Preferential access can no longer be taken as given, and has instead become a changing policy margin that firms have to actively choose, comply with or forgo at every shipment.

Mexican manufacturing exports are at the center of the economy that these shocks are reorganizing. These exports amount to roughly a third of Mexican GDP, of which about 80% is destined for the United States, and are built on a long tradition of foreign investment in export-oriented plants, which today are embedded in the IMMEX program (the modern successor to the 1960s maquiladora regime). On the output side, firms face a choice on every shipment: they can either comply with USMCA’s rules of origin and ship tariff-free to the U.S., or forgo preference and pay the Most Favored Nation (MFN) tariff. Given that this choice shapes where firms source their inputs, where they sell and how much they earn, movement on the compliance margin ends up propagating through the broader economy.

The 2018 trade war raised product-specific tariffs on half of the U.S.–China trade universe. Rather than triggering reshoring, these tariffs diverted trade toward “bystander” economies, where the sectoral specialization that each country already had determined how much of this trade it was able to absorb (Fajgelbaum et al., 2024). Given its geographic proximity, deep supply-chain linkages and preferential access under NAFTA, Mexico was a natural candidate to absorb part of this trade. Between 2017 and 2025, U.S. imports from China fell by roughly half in real terms, whereas U.S. imports from Mexico grew 31%. At the same time, Mexico’s own imports from China grew 32%, with intermediate and capital goods accounting for roughly 60% of that growth.¹ The scale of this restructuring can also be observed on the ground, as foreign-headquartered firms alone added roughly half a million manufacturing workers in Mexico over the same period (Utar, 2026). What these volumes do not show, however, is what happened on the preferential margin. When USMCA replaced NAFTA

¹UN Comtrade bilateral flows at HS-6, mapped to end-use categories using the UN BEC concordance.

in July 2020 with tighter rules of origin, the share of Mexico-to-U.S. trade qualifying for preference did not increase. Instead, it fell from 56% in 2017 to 48% in 2024. It was only after March 2025, when an executive order under the International Emergency Economic Powers Act (IEEPA) imposed an additional 25% (additional to the existing MFN rate) tariff on Mexican imports while exempting USMCA-compliant goods, that this share jumped to about 85%. Firms had been rerouting trade through Mexico for years, but increasingly outside of the preferential regime, and it was not until the cost of staying outside of it became large enough that they formalized this trade under USMCA. While these tariffs were eventually struck down by the U.S. Supreme Court, the episode provides a glimpse at how firms react when not complying with a free trade agreement becomes prohibitively expensive. Following the ruling, the tariff on Mexican goods outside USMCA preference currently stands at 10%.

The goal of this paper is to study the trade-off that firms face between cost-minimizing global sourcing and free-trade-agreement compliance, and to quantify its welfare implications. The relevance of this margin goes beyond what has been observed historically. USMCA is now subject to an annual joint review by the three governments, meaning that rules of origin can change from year to year. This removes the sourcing certainty around which continental supply chains were built, and turns the compliance decision into an ongoing recalculation rather than a one-time choice. The framework developed in this paper allows one to understand what each of these policy scenarios may mean for the economy.

In order to do so, I develop a global sourcing model in which heterogeneous firms combine domestic labor with domestically-sourced and imported intermediates, choose the set of countries they source from out of a global menu, and decide, product by product, whether or not to comply with NAFTA/USMCA rules of origin. Compliance grants the firm tariff-free access to the U.S. and Canadian markets, but constrains its sourcing toward North America. Opting out, on the other hand, allows the firm to source globally without any constraints at the cost of paying the MFN tariff, which is the non-discriminatory rate that a WTO member applies to imports from any other WTO member. Given that domestic value added automatically counts toward regional-content requirements, the production structure plays a direct role in the compliance margin. Firms evaluate all of their feasible sourcing strategies and only enter production if the best one yields positive profits. Aggregate income and the domestic price level adjust in order to clear the Mexican goods market, whereas a competitive non-manufacturing sector pins down the wage and absorbs the reallocation of labor.

I calibrate the model at two different vintages using matched Mexican census-customs firm microdata: a pre-trade-war NAFTA equilibrium (2018 vintage) and a post-USMCA equilibrium (2023 vintage). The transition from NAFTA to USMCA rules shows to be close

to welfare-neutral, as the two-vintage decomposition attributes -0.06% of real income to it. When looking at the growth in real income between the two vintages, the largest components that can be computed on a common panel are the growth in foreign demand ($+7.1\%$) and the shift in relative sourcing conditions ($+2.1\%$), under which imports from China rise the most. Around the 2023 baseline, however, content rules are not welfare-neutral. Tightening the regional value content (RVC) threshold by 25 points lowers compliance by 14 points and real income by 0.6% , given that most firms keep complying and buy North American content that they would not otherwise buy. What makes the agreement valuable is the outside tariff. With MFN tariffs in place, removing USMCA costs Mexico 2.8% of its real income, whereas at the 25-point surcharge of the 2025 order the loss amounts to 17.6% . With the preferential route available, the same surcharge costs only 0.28% . Similarly, content rules cost more when non-compliance is expensive: a 20-point tightening costs about three times as much real income at the 2025 tariff level as at observed tariffs, which is more than the cost of the surcharge itself. Lastly, a 25% tariff on intermediates from China, a country with which Mexico holds no free trade agreement, cuts Chinese sourcing by 29% at no cost in real income, since the tariff rebate offsets the higher input costs. Asia excluding China is the region that gains the most in proportion ($+2.8\%$), while North American, European and domestic sourcing each rise by 1 to 1.3% .

This paper contributes to three different strands of the literature. The first is the quantitative global sourcing literature ([Antràs et al., 2017](#), [Eaton and Kortum, 2002](#), [Head and Mayer, 2019](#), [Head et al., 2024](#), [Morán and Cebrenros, 2025](#), [Tintelnot, 2017](#)), which typically operates at the firm-country level. To this literature, the paper adds a joint firm-level recovery of the capabilities to source from different countries and sell to different destinations, which is done using Mexican customs and census microdata. The second is the theoretical and empirical literature on rules of origin in preferential agreements ([Conconi et al., 2018](#), [Grossman, 1981](#), [Ju and Krishna, 2002](#), [Ornelas and Turner, 2024](#)). Here, the paper contributes a structural model that jointly represents different product-specific compliance regimes, which together cover the NAFTA and USMCA manufacturing universe, and takes it to economy-wide firm-level data rather than to a single sector. The closest antecedents on this margin are [Morán and Cebrenros \(2025\)](#), who extend the global sourcing model of [Antràs et al. \(2017\)](#) with an endogenous NAFTA-versus-WTO compliance choice using Mexican customs data and target bilateral trade flows and firm profits, and [Head et al. \(2024\)](#), who take a complementary structural approach for USMCA that is calibrated to the automotive sector and trace the Laffer curve of regional-content requirements on North American employment. Both of these studies, however, restrict their attention to a specific subset of the treaty’s product-level rules. The third is the literature on the reorganization of global supply

chains in an environment of rising trade conflicts (Alfaro and Chor, 2025, Fajgelbaum et al., 2024, 2020, Freund et al., 2024, Gopinath et al., 2025, Grossman et al., 2023, Utar, 2026, Utar et al., 2025), which has studied how trade and production reorganize across countries, including their diversion toward Mexico, but has not modelled the compliance friction that determines whether this diverted trade enters the United States under preference or not. To this literature, the paper adds a structural welfare quantification of Mexican consumer real income under the joint policy environment of trade conflict and the NAFTA-to-USMCA transition, which to the best of my knowledge no existing structural analysis of USMCA rules of origin has reported.

The remainder of this paper is organized as follows. Section 2 documents the empirical patterns in trade flows and USMCA compliance that motivate the model. Section 3 presents the structural framework. I follow by describing the calibration in Section 4 and the estimation results in Section 5. Section 6 presents the counterfactual analysis and Section 7 concludes.

2 Motivating Facts

Since 2018, North American manufacturing trade has reorganized substantially in favor of Mexico, while the share of MX→US trade claiming NAFTA/USMCA preference has moved through three distinct regimes. There are three policy events that plausibly drive this sequence. The first took place between 2018 and 2019, when, in response to the statements of USTR (2018) on forced technology transfer and other unfair trade practices by China, the United States invoked Section 301 of the Trade Act of 1974 in order to impose tariffs on roughly half of its manufactured imports from China. The second was the replacement of NAFTA by USMCA in July 2020, which raised the required North American content for about two fifths of MX→US trade, concentrated in the automotive core, while leaving most of the remainder unchanged or looser (Table 2). The third came in March 2025, when the second Trump administration’s IEEPA executive order imposed a 25% tariff on Mexican imports while exempting goods meeting USMCA rules of origin, which sharply raised the cost of failing to claim preference. Figure 1 traces four headline series across this window. The rest of this section organizes the evidence in two parts: the reorganization of trade flows (Section 2.1) and the evolution of USMCA preference utilization (Section 2.2).




Figure 1: **Trade flows and NAFTA/USMCA compliance, 2015–2025.** Manufacturing HS6 (16–97), deflated to 2023 USD. Monthly through Dec 2025.

2.1 The Reorganization of Trade Flows

The first place where Mexico’s response to the 2018 trade war can be observed is in trade volumes. Between 2017 and 2025, U.S. imports from China fell sharply (−54% in real terms), and Mexico picked up part of the displaced demand (MX→US: +31%) while at the same time deepening its own dependence on Chinese inputs (MX←CN: +32%). It is worth pointing out that roughly 60% of the MX←CN growth comes from intermediate and capital goods,² whereas only the remaining 40% corresponds to final-consumption goods.

A natural question that arises is whether the new MX→US trade took place in the product codes that Mexico was already exporting in 2017 or in newly-emerging product lines, that is, whether the trade response occurred on the intensive or on the extensive margin. The answer shows to be overwhelmingly intensive. Table 1 cross-tabulates MX→US manufacturing trade by 301 exposure and pre-shock exporter status, and one can observe that the reorganization concentrates along existing 301-exposed product lines, given that the (exposed, existing-exporter) cell alone holds 90% of 2017 MX→US trade. The two (new) rows are small. Furthermore, 98% of their 2025 value corresponds to subheadings created by the 2022 revision of the Harmonized System, which by construction have no 2017 trade and no Section 301 coding, meaning that they mostly record reclassification rather than new products.

The rightmost column of Table 1 reports the trade-weighted US MFN tariff that each cell

²BEC Rev. 5 intermediate and capital-goods categories.

faces, which is the rate that any Mexican exporter not claiming USMCA preference pays, applied to Mexican goods on the same WTO most-favored-nation terms as to any other origin. The exposed-and-existing cell, which holds 90% of MX→US trade, faces an average MFN of 3.3%, and the aggregate trade-weighted MFN across all four cells is 3.6%. Thus, the cost of skipping USMCA preference on a typical MX→US shipment is small in absolute terms. This is relevant for the second part of this section, as Section 2.2 shows that a meaningful share of USMCA-eligible trade was left unclaimed under this modest MFN baseline, and that this only changed once the March 2025 IEEPA carveout raised the cost of not claiming preference.

Table 1: MX→US manufacturing trade by 301 exposure and pre-shock exporter status, 2017–2025

301	Exporter status	Trade share (%)		Trade value (\$bn)		Avg MFN (%)
		2017	2025	2017	2025	
Exposed	Existing (2017 > 0)	90.3	85.1	325.6	403.7	3.3
Exposed	New (2017 = 0)	0.0	0.1	0.0	0.4	1.2
Not exposed	Existing (2017 > 0)	9.7	9.3	35.0	44.3	6.5
Not exposed	New (2017 = 0)	0.0	5.4	0.0	25.8	7.5
Total		100.0	100.0	360.6	474.2	3.6


Notes: MX→US manufacturing trade (HS 16–97), 2023 USD; annual (2017) and monthly (2025) Comtrade. 301 exposure: [Bown \(2021\)](#). MFN: USITC HTS, trade-weighted within cell; in the (not exposed, new) cell only 0.6% of 2025 value carries a rate.

2.2 USMCA Preference Utilization

USMCA inherits the four-regime structure that NAFTA used for its Product Specific Rules of Origin, and adjusts each of these regimes asymmetrically. Under *Change in Tariff Classification (CTC)*, a good qualifies if its non-North American inputs undergo a specified tariff shift through production. *Regional Value Content (RVC)* requires that a minimum share of the good’s value originate in North America. *Hybrid* regimes allow a good to qualify under either CTC or RVC, whichever is easier to satisfy. Lastly, *Yarn-Forward* applies to textiles and requires that upstream yarn originate in North America. This last regime is omitted throughout the paper, given that its multi-stage production requirements go beyond the scope of the quantitative model. The largest change takes place in the RVC regime, where the light-vehicle threshold rises from 62.5% to 75% (phased in over 2020–2023) and is supplemented by a new Labor Value Content rule (40–45% of value produced by workers earning at least \$16/hour) and a 70% North American steel and aluminum purchasing requirement, neither of which existed under NAFTA. CTC tightens in a more subtle way, through HS-specific

tariff-shift updates that generally require shifts from more granular subheadings. Hybrid in turn inherits both of these tightenings.

Figure 2: The USMCA tightening of required North American content is asymmetric



Notes: Ribbons bundle HS6 subheadings moving identically NAFTA→USMCA; width proportional to average 2018–2023 MX→US export value. Violins are trade-weighted marginals (2018 for NAFTA, 2023 for USMCA); labels mark the fifteen HS2 chapters with the most trade in subheadings whose requirement moved, with each chapter's average 2018–2023 share of MX→US manufacturing trade.

In order to compare the two agreements, Figure 2 places every manufacturing subheading on a single “required North American content” axis, where the width of each ribbon is proportional to the average 2018–2023 MX→US export value. For RVC subheadings, this axis corresponds to the explicit regional-value-content threshold written into the rule, that is, the minimum share of the final good’s value that must originate in North America. For CTC subheadings, where the rule requires non-North-American inputs to undergo a specified tariff-classification shift through production, I compute the implied floor on the NA content of the final good that is consistent with that shift given the subheading’s I-O input mix. Hybrid subheadings, which qualify under either rule, are placed on the axis whose parameter is defined under both agreements (the RVC threshold when both carry one, and the CTC floor otherwise), while subheadings that are Hybrid under both agreements are drawn on both axes. Given this common scale, the two violins represent the trade-weighted marginal distributions of required NA content on either side of the transition (NAFTA on the left and USMCA on the

right), whereas the connecting ribbons represent the distinct NAFTA→USMCA transitions in required content, each of which bundles the HS6 subheadings that move identically.

What one can observe at the product level is an asymmetric and concentrated tightening rather than a uniform rightward shift. The USMCA marginal (the right-hand violin) keeps most of its mass where the NAFTA distribution was, while at the same time developing a second and much higher mode near 0.75 that was essentially absent before. When restricting attention to the regimes whose binding axis is the explicit regional-value threshold (RVC and Hybrid), the share of MX→US trade with a required content of 70% or more rises from essentially 0% under NAFTA to about 30% under USMCA. This upper mode consists almost entirely of motor vehicles and parts (HS 87), which by themselves account for 29% of MX→US manufacturing trade and converge on the 75% RVC threshold, whereas machinery (HS 84, 19%) and electrical equipment (HS 85, 21%) account for most of the remaining trade-weighted increase. In the opposite direction, about 12% of trade loosens (the blue ribbons).

Table 2 tabulates the direction of the rule change within each USMCA regime, and shows that the three regimes have very different variation profiles. The tightening concentrates in the RVC regime, where 64.2% of RVC-regime trade in 2023 falls under an HS6 whose required NA content rose, whereas only 2.1% loosened, and another 30% was reclassified from CTC into RVC at the transition. CTC is the most balanced of the three, as roughly a quarter of its trade tightens, a quarter loosens and nearly half remains unchanged, meaning that the net movement within the regime is close to zero. Hybrid lies in between these two, with 37% tightening and only 6% loosening. In aggregate, weighting across all three regimes by 2023 trade, about 40% of manufacturing trade faces a stricter required NA content, which is more than three times the 12% that loosens, while the remainder is either unchanged (36%) or reclassified across the CTC↔RVC divide (10%).

Table 2: Direction of the NAFTA→USMCA rule change, by USMCA regime

USMCA regime	Within-regime direction of change (% of regime trade)			
	Tighter ($\Delta > 0$)	Unchanged ($\Delta = 0$)	Looser ($\Delta < 0$)	Regime switch (CTC↔RVC)
RVC	64.2	3.3	2.1	30.4
CTC	25.5	46.7	25.9	1.9
Hybrid	36.6	57.8	5.6	0.0
All three (aggregate)	40.5	36.1	11.8	10.0

Notes: Δ from Figure 2. Aggregate row weighted by 2023 MX→US exports (residual 1.5% Yarn-Forward). Regime switch counts pure CTC→pure RVC crossings only; Hybrid transitions fall in the tighter/unchanged/looser bins.

Table 3 traces the compliance response across the three main regimes. Between 2017 and

Table 3: USMCA compliance regimes: scope, trade share, and compliance change

Regime	Trade share (%)				Compliance rate (%)				Avg MFN
	2017	2024	2025	2025	2017	2024	2025	2025	(%)
			Jan–Feb	Mar–Dec			Jan–Feb	Mar–Dec	
RVC	31.6	31.4	28.3	27.5	94.5	80.8	78.2	84.4	7.6
CTC	37.3	35.9	36.3	33.0	30.8	31.6	30.3	84.3	2.6
Hybrid	29.5	31.9	34.5	38.7	46.4	34.2	29.1	85.9	1.4
Aggregate	98.4	99.1	99.1	99.2	55.9	48.0	43.6	85.0	3.7

Notes: The Feb/Mar cut isolates the IEEPA order (a 25% tariff on top of the MFN rate on Mexican imports outside USMCA preference, in force from March 4, 2025, with USMCA-compliant goods exempt from March 7). Value-weighted within regime. Shares are against the full four-regime manufacturing universe; residual is Yarn-Forward, which is omitted.

2024, the aggregate compliance rate fell by about 8 percentage points, and the magnitude of the decline in each regime depends on whether its binding constraint is based on value or on classification. The two regimes whose binding constraint depends on the *value* of inputs absorb almost all of the decline, as RVC falls 14 percentage points (from 94.5% to 80.8%) and Hybrid falls 12 percentage points (from 46.4% to 34.2%), whereas CTC barely moves. Given that the trade-weighted MFN on the three modeled regimes was 3.7% (Table 3),³ the option value of skipping USMCA on eligible trade was small.

This changed in March 2025, when the IEEPA executive order imposed a 25% tariff on Mexican imports on top of the existing MFN rate while exempting goods meeting USMCA rules of origin, which raised the cost of failing to claim preference overnight. The compliance response was both immediate and large. Aggregate value-weighted compliance stayed below its 2024 level in Jan–Feb 2025 ($\sim 44\%$) and then jumped to $\sim 85\%$ in Mar–Dec 2025 (Table 3), with CTC and Hybrid both nearly tripling and RVC rising by about six points. The size of the jump suggests that the bulk of the pre-2025 gap reflected the fact that several firms were willing to forgo compliance in favor of unconstrained sourcing and paying the MFN tariff. The sharp rise in compliance in 2025 reflects how responsive the compliance margin is to the cost of non-compliance once that cost is raised substantially.

The descriptive evidence in this section shows what happened, but it does not explain why it happened or what it implies for the real economy. In particular, it does not tell us what compliance costs the firms that comply, how much of the pre-2025 decline is due to USMCA as opposed to the trade war, or what would have happened under either policy in isolation. In order to answer these questions, one needs a model of the firm’s sourcing and compliance decision, which is what Section 3 develops.

³Table 1 reports 3.6% because it averages over all MX→US manufacturing trade, including Yarn-Forward and products without a coded rule, weighted by 2017 and 2025 trade, whereas Table 3 covers only the three modeled regimes, weighted by 2017, 2024 and March–December 2025 trade.

3 Model

The empirical evidence points to three features that the model must include. The first is that compliance changes are heterogeneous within each regime, meaning that preference claims cannot be a deterministic product attribute, given that not every shipment of the same product carries the same claim. The model therefore needs an endogenous compliance margin at the firm-product level. The second is that the changes are asymmetric across regimes in a way that tracks USMCA’s regime-specific rule changes, so the model needs a regime-specific rules-of-origin technology, in which the regional-value-content rule (a minimum North American share of the good’s value) binds differently from the change-in-tariff-classification rule (non-North American inputs must change tariff class through production). The third is that the trade reorganization concentrates on 301-exposed product lines while Mexico’s reliance on Chinese intermediates deepened alongside it, so the model needs a sourcing decision over countries that ties input origin to compliance feasibility. Together, these three features allow the model to attribute the observed compliance changes to each policy channel separately, and to report counterfactual estimates of what compliance would have done under the trade war alone, under USMCA alone, under the rise in tariffs on non-compliant goods alone, and under their interaction.

3.1 Environment

I develop a global sourcing model in the tradition of [Eaton and Kortum \(2002\)](#), [Antràs et al. \(2017\)](#), and [Tintelnot \(2017\)](#), which I extend with free trade agreement rules of origin and an endogenous compliance decision at the firm-product level. The central tension in the model is between minimizing input costs through global sourcing and restructuring the supply chain in order to qualify for tariff-free access to the U.S. and Canadian markets. Firms take the trade-policy parameters as given, and the model traces how these parameters map into sourcing, compliance and welfare.

The world is divided into two layers, one for sourcing and one for sales. On the sourcing side, intermediates are produced in six country blocks:

$$s \in \mathcal{S} = \{\text{U.S. \& Canada (USCA), Mexico (MX), Europe (EU),} \\ \text{China (CH), Asia excl. China (AX), Rest of World (RoW)}\}.$$

On the sales side, output is consumed in three destination blocks:

$$d \in \mathcal{D} = \{\text{U.S. \& Canada (US\&CA), Mexico (MX), Rest of World (RoW)}\}.$$

The exact composition of $\mathcal{S}$ and $\mathcal{D}$ does not drive the mechanisms of the model. The U.S. and Canada share a compliance regime and are therefore pooled on both the sourcing and the demand side, given that from Mexico's perspective they enter the model symmetrically as a single tariff-free North American partner. China is separated from the rest of Asia so that a non-FTA partner can be tariffed in isolation. Splitting destinations in turn allows the firm to choose compliance product by product against the share of its revenue that is US&CA-bound, rather than as an all-or-nothing plant-level decision.

3.2 Consumers

Consumers in each destination d split total expenditure Y_d between manufactures and a homogeneous non-manufacturing good $C_{o,d}$ through a Cobb-Douglas upper tier,

$$U_d = C_{M,d}^{\chi_d} C_{o,d}^{1-\chi_d}, \quad \chi_d \in (0, 1), \quad (1)$$

where $C_{M,d}$ is the manufacturing composite and the non-manufacturing good is the numeraire ($p_o \equiv 1$). A constant share of expenditure falls on manufactures,

$$E_d \equiv \chi_d Y_d, \quad (2)$$

with the remainder $(1 - \chi_d)Y_d$ spent on the non-manufacturing good. The manufacturing composite is itself a CES aggregate over a continuum of differentiated varieties $\omega \in \Omega_d$, with exogenous appeal weights $b_d(\omega) > 0$ scaling effective consumption:

$$C_{M,d} = \left(\int_{\omega \in \Omega_d} (b_d(\omega) q_d(\omega))^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}}, \quad (3)$$

with $\sigma > 1$ the elasticity of substitution. Utility maximization subject to the manufacturing budget $\int p_d(\omega) q_d(\omega) d\omega = E_d$ yields the quantity demanded

$$q_d(\omega) = E_d P_d^{\sigma-1} p_d(\omega)^{-\sigma} b_d(\omega)^{\sigma-1}, \quad (4)$$

where $P_d = \left(\int (p_d(\omega)/b_d(\omega))^{1-\sigma} d\omega \right)^{1/(1-\sigma)}$ is the CES price index of destination d . The only difference across destinations lies in what is endogenous. For the foreign markets $d \in \{\text{US\&CA, RoW}\}$, total expenditure Y_d is exogenous, so E_d is taken as given and the share χ_d is absorbed into its calibrated level. For Mexico, on the other hand, manufacturing non-wage income Π_M is endogenous and non-manufacturing income Π_o is exogenous (both are defined in Section 3.4), so E_M moves with Mexican income $w_m \bar{L} + \Pi_M + \Pi_o$ through (34), which is

what transmits the feedback from firm activity to domestic demand through Π_M .

3.3 Firms

Primitives. Each potential firm f in mass $\bar{M}$ is characterized by three objects: a product set $\mathcal{P}_f$ (empirically set at the HS6 level), a vector of country-specific sourcing capabilities $\{a_{f,s}\}_{s \in \mathcal{S}}$, and a vector of destination demand shifters $\{b_{f,d}\}_{d \in \mathcal{D}}$. The capability $a_{f,s}$ captures the firm's supplier networks and institutional relationships in source country s . For instance, U.S.-owned firms tend to have a high $a_{f,\text{USCA}}$, Chinese-owned firms tend to have a high $a_{f,\text{CH}}$, and so on. The shifter $b_{f,d}$ plays the analogous role on the demand side, scaling revenue at destination d in order to capture features such as market access, distribution networks, and brand recognition. Each product $p \in \mathcal{P}_f$ in turn determines the compliance regime and the MFN tariff that the firm faces.

Sourcing Strategy. Each firm sources a continuum of intermediate inputs $j \in [0, 1]$ from source countries in $\mathcal{S}$. To source from country s , the firm knows how well it can source from it through $a_{f,s}$ and pays a fixed cost $w_s \kappa_s \geq 0$ in units of Mexican labor. Let $\mathcal{S}_f \subseteq \mathcal{S}$ denote the set of countries from which the firm sources. I refer to $\mathcal{S}_f$ as the firm's *sourcing strategy*, and assume that it is shared across all of the firm's products and destinations.⁴

Conditional on paying the fixed cost, the firm draws input-level productivities $z_s(j) \sim \text{Fréchet}(T_s a_{f,s}^\theta, \theta)$ independently across inputs and source countries. These draws absorb all of the firm-, input- and country-specific factors (such as supplier match quality, shipping reliability and product fit) that shape the realized landed cost. $T_s > 0$ is the absolute sourcing advantage of country s , θ governs the dispersion of the productivity draws, and the firm-country capability $a_{f,s}$ shifts the location of the distribution upward, which raises the likelihood that country s offers a lower landed cost for a given input. The per-input landed cost is

$$c_{f,s}(j) = \frac{w_s \tau_s}{z_s(j)} \quad (5)$$

where w_s reflects the cost of inputs from each source country s and τ_s is the iceberg trade cost of sourcing from s . For each input j , the firm picks the cheapest source, $s^*(j) = \arg \min_{s \in \mathcal{S}_f} c_{f,s}(j)$. The intermediate cost index is then the integral of these per-input minima over the input continuum, whose closed form follows from standard Fréchet algebra:

⁴Customs data record input origin and output destination separately, but not which input made which output, so a shared $\mathcal{S}_f$ is what the data support.

$$\begin{aligned}
c_f^M(\mathcal{S}_f) &= \int_0^1 \min_{s \in \mathcal{S}_f} \frac{w_s \tau_s}{z_s(j)} dj \\
&= \Gamma\left(1 - \frac{1}{\theta}\right) \left(\sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-1/\theta}.
\end{aligned} \tag{6}$$

The expected share of intermediates sourced from country s is the integral over j of the indicator that s offers the minimum landed cost, which the same algebra evaluates as

$$\begin{aligned}
\mu_{f,s}(\mathcal{S}_f) &= \int_0^1 \mathbf{1}\left\{ \frac{w_s \tau_s}{z_s(j)} = \min_{s' \in \mathcal{S}_f} \frac{w_{s'} \tau_{s'}}{z_{s'}(j)} \right\} dj \\
&= \frac{T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta}}{\sum_{s' \in \mathcal{S}_f} T_{s'} a_{f,s'}^\theta (w_{s'} \tau_{s'})^{-\theta}}.
\end{aligned} \tag{7}$$

Production Technology. Each active firm produces output by combining Mexican labor ℓ_f with an intermediate composite M_f through a Cobb-Douglas technology:

$$q_f = \ell_f^\gamma M_f^{1-\gamma}, \tag{8}$$

where $\gamma \in (0, 1)$ is the value-added share of labor. Labor is hired locally at wage w_m , and its cost in turn counts as North American value added by construction. The intermediate composite M_f aggregates a continuum of inputs $j \in [0, 1]$, each sourced from the lowest-cost country in the firm's sourcing set. Note how this allows the possibility for a firm to comply with the trade agreement requirements not only by sourcing inputs from North America, but also through the value added by Mexican labor.

Combining (8) with (6), cost-minimization over labor and the intermediate composite yields the Cobb-Douglas unit cost:

$$c_f(\mathcal{S}_f) = B w_m^\gamma [c_f^M(\mathcal{S}_f)]^{1-\gamma}, \tag{9}$$

where $B \equiv \gamma^{-\gamma}(1-\gamma)^{-(1-\gamma)}$ is a constant derived from the cost-minimization problem.

Compliance: Three Regimes.⁵ Rules of origin bind only on US&CA-bound output. For each product $p \in \mathcal{P}_f$, the firm chooses a compliance indicator $d_{fp} \in \{0, 1\}$, where $d_{fp} = 1$ means that the firm accepts the applicable input restriction in exchange for tariff-free access

⁵Yarn-Forward is omitted given that textiles and apparel account for under 2% of MX→US manufacturing trade, and that their multi-tier yarn→fabric→garment origin chain goes beyond the model's single-tier sourcing.

to the US&CA market, whereas $d_{fp} = 0$ means that it pays the MFN tariff τ_p^{MFN} on US&CA-bound shipments and sources freely. The different regimes are quantified as follows:

- **Change in Tariff Classification (CTC).** A CTC rule for output p names, at a treaty-determined HS depth (chapter, heading or subheading), a specific set $\mathcal{P}(p)$ of input HS classes whose tariff classification must change *through production* in order for the input to qualify. A non-NA input from $\mathcal{P}(p)$ whose tariff class is unchanged through production fails the rule, whereas an NA-sourced input always passes, regardless of its classification. I take the I-O technology at output p as binding, meaning that any input in $\mathcal{P}(p)$ to which the I-O table assigns a positive expenditure share at p is interpreted as failing the classification-change test under p 's production technology, and must therefore be NA-sourced for compliance. The rule then reduces to a single product-specific scalar, where $\phi_p(j)$ is the I-O expenditure share of input class j in p 's intermediate bundle:

$$\alpha_p = \sum_{j \in \mathcal{P}(p)} \phi_p(j),$$

which is the I-O share of p 's intermediate bundle allocated to those restricted-and-unchanged classes. Through the lens of the model, the rule maps to a subset of the firm's continuum of input varieties $j \in [0, 1]$ of measure α_p that must be NA-sourced, while the residual measure $1 - \alpha_p$ is drawn from the unrestricted set $\mathcal{S}_f$. Given that input varieties are drawn iid Fréchet across countries, the cost of the restricted subset depends only on its measure and not on which varieties it contains, so the compliant intermediate cost is a Cobb-Douglas aggregation of the restricted (NA-only) and unrestricted sub-continua:

$$c_{f,p}^{M,CTC}(\mathcal{S}_f) = \Gamma\left(1 - \frac{1}{\theta}\right) \left(\sum_{s \in NA \cap \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-\alpha_p/\theta} \left(\sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-(1-\alpha_p)/\theta}. \quad (10)$$

Both trade agreements further allow a *de minimis* relaxation, under which CTC is treated as satisfied whenever non-conforming non-originating inputs account for at most $\delta_p = 10\%$ of the good's value under USMCA and 7% under NAFTA. Given that α_p is a share of intermediate spending and that intermediates make up a fraction $1 - \gamma$ of the good's cost, the tolerance trims the binding restricted measure to $\alpha_p^{\text{eff}} = \max(\alpha_p - \delta_p/(1 - \gamma), 0)$, where the good's value is measured at total cost, which is the same basis as the net-cost RVC thresholds below. This effectively allows firms to have a small margin of deviation from the CTC requirement without losing preference eligibility.

- **Regional Value Content (RVC).** The North American content share of production must reach a product-specific threshold μ_p^{RVC} . Given that Mexican labor counts as North American content, the content share under an unrestricted sourcing strategy is

$$\mu_f^{NA}(\mathcal{S}_f) \equiv \gamma + (1 - \gamma) \sum_{s \in NA \cap \mathcal{S}_f} \mu_{f,s}(\mathcal{S}_f), \quad (11)$$

where the first term is the automatic contribution of Mexican labor and the second is the NA share of the firm's intermediate sourcing. The rule is met when $\mu_f^{NA}(\mathcal{S}_f) \geq \mu_p^{RVC}$. A firm below the threshold can still qualify by switching part of its inputs to North American sources, at the cost that is derived in [Appendix A.4](#).

- **Hybrid.** The firm may qualify under either rule and takes the cheaper compliant cost:

$$c_{f,p}^{Hyb}(\mathcal{S}_f) = \min(c_{f,p}^{CTC}(\mathcal{S}_f), c_{f,p}^{RVC}(\mathcal{S}_f)).$$

The compliance choice is product- and destination-specific (only US&CA-bound shipments are exposed), whereas the sourcing strategy $\mathcal{S}_f$ is shared across all of the firm's products and destinations. This means that a firm can tailor its compliance vector $\{d_{fp}\}$ product by product, but that its menu of feasible source countries is fixed once it pays the sourcing fixed costs. The single-sourcing-strategy assumption keeps the discrete-choice problem tractable, and at the same time aligns the model with the customs data, which record firm-level import composition rather than product-specific intermediate flows.

Conditional on $\mathcal{S}_f$ and the per-product compliance indicator d_{fp} , the destination-specific landed unit cost is

$$c_{f,p,d}(\mathcal{S}_f, d_{fp}) = \begin{cases} c_{f,p}(\mathcal{S}_f) & \text{if } d = \text{US\&CA and } d_{fp} = 1, \\ c_f(\mathcal{S}_f) (1 + \tau_p^{MFN})^{\sigma/(\sigma-1)} & \text{if } d = \text{US\&CA and } d_{fp} = 0, \\ c_f(\mathcal{S}_f) & \text{if } d \in \{\text{MX, RoW}\}, \end{cases} \quad (12)$$

where $c_f(\mathcal{S}_f)$ is the unrestricted unit cost from [\(9\)](#) and $c_{f,p}(\mathcal{S}_f)$ is the compliant unit cost under the regime that governs product p , that is, one of $c_{f,p}^{CTC}(\mathcal{S}_f)$, $c_{f,p}^{RVC}(\mathcal{S}_f)$ and $c_{f,p}^{Hyb}(\mathcal{S}_f)$. Note that the choice $d_{fp} = 1$ requires at least one North American country in $\mathcal{S}_f$ under every regime, given that both the CTC restriction α_p and the RVC restriction $\alpha_{f,p}^*$ draw the restricted inputs from NA sources. Under compliance, the firm accepts the applicable input restriction and ships tariff-free at $c_{f,p}$, whereas under non-compliance it sources without restriction at c_f but pays the MFN tariff on its US&CA-bound revenue. The exponent

$\sigma/(\sigma - 1)$ turns an ad valorem duty on the fob value into the cost wedge that enters revenue. The US&CA buyer pays $p(1 + \tau_p^{MFN})$, the firm receives $p = \frac{\sigma}{\sigma-1}c_f$, and the duty $\tau_p^{MFN} p q$ accrues to the U.S. Treasury, meaning that the firm's fob revenue and profit on a non-compliant product scale as $(1 + \tau_p^{MFN})^{-\sigma}$, which is exactly what the wedge gives in (13). MX and RoW shipments are unaffected by the compliance choice.

Destination Markets. Each firm chooses a destination set $\mathcal{D}_f \subseteq \mathcal{D}$, and pays a fixed entry cost $w_m \kappa_d$ in units of Mexican labor for each $d \in \mathcal{D}_f$. Each product of an active firm is a differentiated variety at every destination that it serves. Given constant-markup pricing over the duty-free unit cost, the firm's fob revenue from product p at destination d is

$$R_{f,p,d}(\mathcal{S}_f, b_{f,d}) = \varpi_{f,p} E_d P_d^{\sigma-1} (c_{f,p,d}(\mathcal{S}_f))^{1-\sigma} b_{f,d}^{\sigma-1}, \quad (13)$$

where $c_{f,p,d}(\mathcal{S}_f)$ is the destination-specific landed unit cost from (12), the markup constant $(\frac{\sigma}{\sigma-1})^{1-\sigma}$ is absorbed into $b_{f,d}$, and $\varpi_{f,p}$ is product p 's share of the firm's observed exports times its number of products, so that the weights average to one within the firm.

Profits and the Firm Problem. The firm jointly chooses its sourcing set $\mathcal{S}_f$, its destination set $\mathcal{D}_f$, and a per-product compliance vector $\{d_{fp}\}_{p \in \mathcal{P}_f}$ in order to maximize expected profits:

$$\pi_f = \max_{\mathcal{S}_f \subseteq \mathcal{S}, \mathcal{D}_f \subseteq \mathcal{D}, \{d_{fp}\}} \frac{1}{\sigma} \sum_{d \in \mathcal{D}_f} \sum_{p \in \mathcal{P}_f} R_{f,p,d} - w_m \sum_{s \in \mathcal{S}_f} \kappa_s - w_m \sum_{d \in \mathcal{D}_f} \kappa_d, \quad (14)$$

subject to the applicable regime-specific compliance condition for each compliant product. As in Antràs et al. (2017), solving (14) is a combinatorial problem, in which the firm evaluates every non-empty source-country subset $\mathcal{S}_f \subseteq \mathcal{S}$, combined with every destination subset $\mathcal{D}_f \subseteq \mathcal{D}$ and the per-product compliance vector, and selects the strategy that yields the highest profit.

Endogenous Entry and Exit. Let $\pi_f(w_m)$ denote the maximized expected profit from (14), with the dependence on the Mexican wage made explicit. Firm f is *active* if and only if some non-empty $(\mathcal{S}_f, \mathcal{D}_f)$ yields positive profits:

$$f \in \mathcal{A}(w_m) \iff \pi_f(w_m) > 0. \quad (15)$$

The endogenous active set $\mathcal{A}(w_m) \subseteq [0, \bar{M}]$ is the subset of potential firms for whom some sourcing, destination and compliance strategy is profitable. Inactive firms do not produce, source inputs, or employ Mexican labor. As in Tintelnot (2017), the mass of potential firms $\bar{M}$ is taken as a primitive, so that there is no free-entry condition. Considering that firm heterogeneity is multidimensional (the full capability vector $\{a_{f,s}\}$ and demand vector $\{b_{f,d}\}$),

there is no single productivity cutoff separating active from inactive firms.

3.4 Market Clearing and Equilibrium

Mexico is endowed with a fixed labor force $\bar{L}$, which is mobile across the manufacturing sector modeled above and a homogeneous non-manufacturing sector. The non-manufacturing sector is perfectly competitive and produces the numeraire good with labor alone under constant returns, $Q_o = A_o L_o$, where A_o is its labor productivity. Free entry equates price to marginal cost, $p_o = w_m/A_o$, and given that the non-manufacturing good is the numeraire ($p_o \equiv 1$), this pins the Mexican wage to outside productivity,

$$w_m = A_o. \quad (16)$$

The non-manufacturing sector thus acts as a competitive outside option that is large enough to fix the wage at which manufacturing hires, meaning that labor reallocates across the two sectors at a constant w_m , which is the wage-pinning closure of [Antràs et al. \(2017\)](#).⁶

With the wage pinned, the labor market allocates $\bar{L}$ across the two sectors passively. Manufacturing employs whatever labor firms demand at $w_m = A_o$, and the non-manufacturing sector absorbs the residual $L_o = \bar{L} - L_m \geq 0$. Letting $VC_f \equiv \frac{\sigma-1}{\sigma} \sum_{d,p} R_{f,p,d}$ denote total variable cost, manufacturing labor demand aggregates over the active set $\mathcal{A} \equiv \mathcal{A}(w_m)$:

$$L_m = \int_{f \in \mathcal{A}} \left[(\gamma + (1 - \gamma)\mu_{f,MX}) \frac{VC_f}{w_m} + \sum_{s \in \mathcal{S}_f} \kappa_s + \sum_{d \in \mathcal{D}_f} \kappa_d \right] df, \quad (17)$$

where the bracket collects three sources of labor demand: production labor (the γ share of variable cost under Cobb-Douglas), Mexican-sourced intermediates (the $(1 - \gamma)\mu_{f,MX}$ share), and the fixed costs of sourcing and destination entry. Note that equation (17) is an accounting identity that reads off L_m from the firm-level solution, rather than a market-clearing condition that solves for a wage. The pinned-wage closure is interior whenever $L_m \leq \bar{L}$.

For the foreign destinations $d \in \{\text{US\&CA}, \text{RoW}\}$, manufacturing expenditure and prices (E_d, P_d) are taken as exogenous, considering that Mexico is a small open economy relative to those markets and its firms' decisions are not expected to materially move external incomes or price levels. The Mexican pair (E_M, P_M) , on the other hand, is endogenous ([Appendix A.7](#)).

The Mexican price index P_M in turn aggregates the prices of Mexican-firm varieties sold

⁶Non-manufacturing accounts for about 75% of Mexican employment (INEGI, ENOE), which supports the perfectly-elastic labor-supply pole.

at MX through the CES dual of (3):

$$(P_M)^{1-\sigma} = \int_{f \in \mathcal{A}: MX \in \mathcal{D}_f} \sum_{p \in \mathcal{P}_f} \left(\frac{p_{f,p,MX}}{b_{f,MX}} \right)^{1-\sigma} df. \quad (18)$$

The welfare metric that I report in the counterfactuals is Mexican consumer real income under the Cobb-Douglas upper tier. The ideal price index dual to $U = C_M^\chi C_o^{1-\chi}$ is the geometric mean $P \propto P_M^\chi P_o^{1-\chi}$, so deflating aggregate nominal Mexican income by P gives

$$W_M = \frac{w_m \bar{L} + \Pi_M + \Pi_o}{P_M^\chi P_o^{1-\chi}}, \quad (19)$$

where the three numerator components are, in order, the economy-wide wage bill (both sectors pay $w_m = A_o$ over the fixed labor supply $\bar{L}$), manufacturing non-wage income

$$\Pi_M = \int_{f \in \mathcal{A}} \pi_f df, \quad (20)$$

and non-manufacturing income Π_o . Given perfect competition and constant returns in that sector, free entry forces zero economic profits, so $\Pi_o \equiv 0$. Setting $P_o = 1$ as the numeraire, the denominator of (19) collapses to P_M^χ . Since $w_m = A_o$ and $\bar{L}$ are invariant under the pinned-wage closure, movements in the welfare numerator operate entirely through Π_M .

Definition (Equilibrium). Given preference (σ, χ) , technology (θ, γ, A_o) , firm $(\mathcal{P}_f, \{a_{f,s}\}, \{b_{f,d}\})$, policy $(\{\alpha_p, \mu_p^{RVC}, \tau_p^{MFN}\})$, foreign-market $(\{E_d, P_d\}_{d \neq MX}, \{T_s, w_s, \tau_s\})$, and fixed-cost $(\{\kappa_s, \kappa_d\})$ primitives, together with the Mexican labor force $\bar{L}$, an equilibrium is a tuple

$$(\mathcal{A}, \{\mathcal{S}_f, \mathcal{D}_f, (d_{fp})_{p \in \mathcal{P}_f}\}_{f \in \mathcal{A}})$$

such that, at the pinned wage $w_m = A_o$:

- (i) Consumers satisfy the Cobb-Douglas upper tier (1) and CES demand (4) at each destination.
- (ii) Each active firm $f \in \mathcal{A}$ solves (14) for $(\mathcal{S}_f, \mathcal{D}_f, \{d_{fp}\})$, taking the domestic pair (E_M, P_M) and the foreign pairs $\{(E_d, P_d)\}_{d \neq MX}$ as given.
- (iii) Entry satisfies (15): $\mathcal{A} = \{f : \pi_f(w_m) > 0\}$.
- (iv) Domestic demand is consistent: E_M satisfies (34) with Π_M from (20), and P_M satisfies (18), both at the firms' choices.

4 Calibration

4.1 Data

I take the model to firm-level Mexican data drawn from two sources: the *Profile of Export Manufacturing Enterprises* (PEME) customs microdata and the Mexican *Economic Census* (CE). PEME records the universe of Mexican customs transactions at the firm $\times$ HS8 $\times$ partner $\times$ year level, and from it I construct firm-level import-source shares $\mu_{f,s}$ for $s \neq \text{MX}$ and export-destination shares $\omega_{f,d}$ for $d \in \{\text{US\&CA}, \text{RoW}\}$. From the CE, which covers the population of Mexican manufacturing establishments in 2019 and 2024, I recover each firm’s domestic material purchases, domestic sales, employment and paid wage bill, ownership and industry. Establishments are aggregated to the firm level through the INEGI firm identifier, summing levels and taking industry and ownership from the largest establishment. This same INEGI firm identifier links CE firms to PEME firms, so that each matched firm carries its CE characteristics alongside the PEME-derived sourcing and destination shares. Given that the census asks about the firm’s operations in the year before it is taken, I match the 2019 and 2024 censuses to PEME 2018 and 2023 respectively, and express every peso magnitude of the 2018 vintage in constant 2023 pesos. Details about the results of these firm-level matches are available in [Appendix C](#).

A census firm enters the model as an individual row when it matches a PEME firm and trades on at least one side, that is, when it records foreign sales or imported materials. This includes import-only firms, which source abroad but do not export. Each import-only firm is assigned a single potential export product, namely the modal HS6 of its six-digit SCIAN class in the [Pierce and Schott \(2012\)](#) concordance weighted by Mexican exports to the U.S. and Canada, so that if the imputed demand shifters send it to the US&CA market in a counterfactual, it faces that product’s rule and MFN rate. All other census firms, both those without a customs record and those with one but no trade, are collapsed into bins by NAICS-3 industry rather than being assigned individual rows. These firms span a wide size range and account for a non-trivial share of manufacturing activity. In the Mexican price index, the non-traders enter as an additive term K in the CES sum, $P_M^{1-\sigma} = T + K$, where T aggregates the traders’ price terms and K is calibrated to the non-traders’ baseline share. Given that the non-traders are modeled as sourcing and selling only domestically at the pinned wage, their prices, and hence K , are invariant across counterfactuals. Nevertheless, since K enters the CES sum, it still dampens the response of P_M through the traders’ share. The sales of the non-traders move with Mexican demand, and their variable profits, which are a fraction $1/\sigma$ of those sales, enter manufacturing non-wage income Π_M .

The firm-level block runs on the traders in both vintages, which amount to 5,719 firms in

2018 and 5,133 in 2023. The calibration proceeds in two layers. The first is a set of externally calibrated parameters that are fixed from data or from the literature, and the second is a set of structural parameters that are estimated through model inversions, Heckman-corrected regressions and a moment-inversion estimator that matches aggregate participation shares. Table 4 summarizes both layers.

4.2 Externally Calibrated Parameters

Preferences (σ). Following Antràs et al. (2017), I pin the CES elasticity σ from the firm’s first-order condition under monopolistic competition, under which revenue equals a constant markup over variable cost, $R_f = \frac{\sigma}{\sigma-1} VC_f$, with $VC_f = wL_f + M_{t,f}$. I take the median of the firm-level markup R_f/VC_f across all CE manufacturing firms in each census wave, pool the two medians with weights proportional to each wave’s total revenue, and invert the pooled markup, so that the same σ enters both vintages. Pooling is a deliberate choice, as it imposes that σ , a deep preference parameter, is constant over the five-year window, meaning that the decomposition in Section 6 is attributable to the inputs that do vary. The median markup is 2.00 in the 2019 wave and 1.82 in the 2024 wave, and the pooled markup is 1.91, which inverts to $\sigma = 2.10$.

Technology (γ). Under Cobb-Douglas production, the labor share of variable cost equals the production-function exponent (Appendix A.2), which justifies calibrating γ as a cost share. Given that input intensity differs sharply across manufacturing branches, I compute the share at the NAICS-3 sector level as a ratio of sums, $\gamma_s = \sum wL / \sum (wL + M_t)$, where the paid wage bill wL and own-materials purchases M_t are aggregated across *all* CE manufacturing firms in sector s (and not only across the traders of the model block), and then assign each model firm its sector’s value. The sector shares are computed within each census wave and then frozen at their simple mean across the two waves. Appendix Figure 9 shows the per-wave and pooled values across the 21 NAICS-3 sectors. When weighted by the wage bill of the estimation panels, the pooled shares aggregate to 0.204 (2018) and 0.207 (2023).

Dispersion (θ). The Fréchet shape governing the dispersion of firm-level sourcing capabilities is set to $\theta = 1.79$ following Antràs et al. (2017).

Foreign-market block: country sourcing shifters (ψ_s). The bundle $\psi_s \equiv T_s(w_s\tau_s)^{-\theta}$ is identified from mean firm sourcing shares. Given that the underlying T_s , w_s and τ_s are not separately identified, I work with ψ_s directly and normalize $\psi_{\text{USCA}} \equiv 1$. I choose US–Canada as the normalizing coordinate because it has the highest source-block participation among the non-home blocks, which provides a stable non-zero denominator for the firm-level ratio $\mu_{f,s}/\mu_{f,\text{USCA}}$ that is used in the structural step below. Without loss of generality, setting

Table 4: Calibration: Parameter Values and Identification

Parameter	Object	Identified from
<i>A priori / data</i>		
σ	CES elasticity	Unweighted median firm markup $R_f/VC_f = \sigma/(\sigma - 1)$ per census wave; revenue-weighted mean of the two medians, inverted (CE)
γ_s	Cobb-Douglas VA share (NAICS-3)	Sector cost share $\sum wL / \sum (wL + M_t)$, mean of the two census waves (CE)
α_p, μ_p^{RVC}	Product RoO parameters	NAFTA/USMCA + INEGI I-O
τ_p^{MFN}	MFN tariff	USITC
$E_d, P_d, d \in \{\text{US\&CA, RoW}\}$	Foreign mfg. expenditure / price	World Bank WDI
$\psi_s \equiv T_s(w_s \tau_s)^{-\theta}$	Country sourcing shifter	Mean firm sourcing shares (CE and PEME)
$\bar{L}$	Mexican labor force	CE paid employment, manufacturing, trade and services
A_o	Non-mfg. productivity (pins w_m)	CE wage bill per paid worker, manufacturing, trade and services
$\chi \equiv \chi_{MX}$	Mfg. expenditure share (MX)	INEGI use table (consumption)
$\theta = 1.79$	Fréchet shape	Antràs et al. (2017)
$\delta = \sqrt{\log 3}$	Fixed-cost dispersion	Pinned, Morán and Cebreros (2025) ; estimated in Antràs et al. (2017)
<i>Structural</i>		
$a_{f,s}$ (per source country)	Firms' sourcing capabilities	Model based inversion and Heckman two-step over missing cells
$b_{f,d}$ (per destination)	Firms' demand shifters	Model based inversion and Heckman two-step over missing cells
κ_s (per source)	Source-entry fixed cost	Moment inversion on source participation shares
κ_d (per destination)	Destination-entry fixed cost	Moment inversion on destination participation shares
<i>Equilibrium</i>		
L_m	Manufacturing employment	Labor allocation (17) (non-mfg. residual)
E_M, Π_M, P_M	Equilibrium MX expenditure / income / price	Domestic-demand fixed point (34), (18) and (20)

$a_{f,s} \equiv 1$ collapses (7) to $\bar{\mu}_s/\bar{\mu}_{USCA} = \psi_s$, so that the remaining $\hat{\psi}_s$ are recovered directly from the sample means of the CE–PEME sourcing shares. The firm-specific $a_{f,s}$ are then restored in Section 4.3.

Foreign-market block: manufacturing expenditure and prices (E_d, P_d). The CES revenue equation $R_{f,p,d} = \varpi_{f,p} E_d P_d^{\sigma-1} c_{f,p,d}^{1-\sigma} b_{f,d}^{\sigma-1}$ takes the manufacturing expenditure E_d and the price P_d as exogenous on the foreign side, with the upper-tier share χ_d absorbed into E_d (2). I normalize $E_d \equiv P_d \equiv 1$ in the 2018 vintage for both $d \in \{\text{US\&CA}, \text{RoW}\}$. For the 2023 vintage, I take the 2018→2023 growth of nominal GDP and of the implicit GDP deflator from the World Bank WDI and set E_d to their ratio. The aggregation $\text{US\&CA} = \text{USA} + \text{CAN}$ and $\text{RoW} = \text{WLD} - \text{USA} - \text{CAN} - \text{MEX}$ is consistent with the model’s three-destination structure. The Mexican pair, on the other hand, is not anchored to data. (Π_M, P_M) is *read off* the master solve through the aggregation identities (18) and (20) at the pinned wage $w_m = A_o$, and in the counterfactual exercises of Section 6 (E_M, P_M) moves with income through (34).

Policy ($\alpha_p, \mu_p^{RVC}, \tau_p^{MFN}$). The product-level RoO parameters (α_p, μ_p^{RVC}) come from USMCA Annex 4-B⁷ and from Conconi et al.’s (2018) coding of NAFTA Annex 401, and are refined with the INEGI 2018 input-output table in order to capture each product’s input intensity. INEGI’s I-O table is available at a six-digit SCIAN disaggregation, so I used Pierce and Schott (2012)’s HS-NAICS concordance tables to match the specific input content requirements determined at the HS system to the NAICS input requirements. More information on the matching between these two classifications is available in Appendix B. MFN tariffs τ_p^{MFN} come from the USITC 2023 Harmonized Tariff Schedule. RVC thresholds enter in their *net-cost* formulation rather than in their *transaction value* formulation (which is roughly 10pp higher), given that the condition (11) is a cost share. All parameters are attached at the HS6 level, which is the granularity at which Annex 4-B writes the product-specific rules and at which the three regimes are cleanly separable. The parameter file spans the HS2012, HS2017 and HS2022 revisions, with codes linked across revisions through the UN Stats conversion tables, so that the customs codes of either vintage find their rule. It assigns a rule to 98% of Mexican manufacturing exports to the U.S. in both 2018 and 2023, with the remainder being Yarn-Forward.

α_p is computed from the INEGI input-output table conditional on the HS depth at which

⁷Annex 4-B is distributed only as legal text and not as a machine-readable table, so I built a parser that reads the annex and extracts each product-specific rule (CTC clauses, RVC thresholds and de minimis exceptions) into an HS6-level dataset. At each rule head, the parser regex-matches the text against a small library of standard treaty templates (“a change to X from any other Y”, “a regional value content of not less than Z percent”) in order to classify the regime and pull out the numerical thresholds. Rules written at a coarser HS depth (chapter or heading) are inherited onto their HS6 children where subheading-level specifics are absent.

the rule is written, so that the resulting HS6 parameter still reflects the binding restriction. As a clarifying example, consider HS 8471.30 (portable automatic data processing machines, i.e. laptops). USMCA Annex 4-B specifies the following CTC restriction:

A change to a good of subheading 8471.30 from any other good within that subheading or any other subheading, except from subheading 8471.41 through 8471.50.

The carve-out closes a substantial-assembly loophole. Without it, a Mexican firm could import a near-complete processing-unit module at HS 8471.50 from China, snap it into a chassis, and have the input technically satisfy the tariff shift given that 8471.50 differs from 8471.30. According to the input-output table, inputs from the 8471.41–50 range account for 14.6% of HS 8471.30’s intermediate spending, which yields $\alpha_{8471.30} = 0.146$.⁸

Table 5 summarizes the resulting HS6 parameter universe that is used as a model input, and Figure 3 below shows the underlying trade-weighted distributions of α_p and μ_p^{RVC} . On average across HS6 subheadings, pure-RVC coverage collapses to an automotive core at the 75% net-cost threshold, Hybrid coverage roughly doubles, and the CTC input restriction α_p tightens on both its pure and its Hybrid arm.

Closure ($\bar{L}$, A_o , χ). The labor-market closure adds a competitive non-manufacturing sector that pins the Mexican wage. Its productivity A_o sets the level of that wage and equals the census wage bill per paid worker across manufacturing, trade and services, which amounts to 165.4 thousand 2023 pesos in the 2019 census and 168.7 thousand in 2024. The labor force $\bar{L}$ is census paid employment in the same sectors, 16.3 and 20.6 million workers, where the manufacturing block employs the labor that its firms demand at that wage (17) and the competitive non-manufacturing sector employs the residual. The manufacturing expenditure share $\chi = 0.37$ is the share of Mexican final consumption that falls on manufactures, taken from the INEGI use table. It weights the manufacturing price index in the welfare deflator (19) and does not enter the demand loop (34).

Compliance rate. The reported compliance rate is the share of US&CA-bound revenue that is shipped under preference, $\sum_{f,p} d_{fp} R_{f,p,US\&CA} / \sum_{f,p} R_{f,p,US\&CA}$, where d_{fp} is the optimal choice in (14), which makes it directly comparable to the empirical benchmark in Table 3. Products with no coded rule, which account for 1.2% of US&CA-bound revenue in the 2018 panel and 1.8% in the 2023 panel, can never comply and pay no tariff. Given that a zero-MFN product whose rule does not bind complies at no cost, I also report a *strict* rate that counts as compliant only those products with a positive MFN rate. Its data counterpart, the value claimed on positive-MFN lines over all value, is 52.3% in 2018 and 46.5% in 2024.

⁸The parameter file used in the estimation codes this rule as a change of subheading, so it also restricts inputs classified in 8471.30 itself, which are another 0.3% of spending, and gives $\alpha_{8471.30} = 0.149$, whereas the clause “any other good within that subheading” exempts them.

Table 5: NAFTA vs USMCA Policy Parameters: Summary Stats

	NAFTA	USMCA	Δ
<i>CTC regime</i>			
HS6 subheadings	3,284	3,146	−138
α_p (mean / median)	14.4 / 4.9	19.0 / 7.7	+4.6 / +2.8
<i>RVC regime</i>			
HS6 subheadings	717	79	−638
μ_p^{RVC} (mean / median)	50.0 / 50.0	71.9 / 75.0	+21.9 / +25.0
<i>Hybrid regime</i>			
HS6 subheadings	818	1,575	+757
α_p (CTC arm, mean / median)	4.5 / 0.4	23.6 / 16.3	+19.1 / +15.9
μ_p^{RVC} (RVC arm, mean / median)	50.7 / 50.0	40.9 / 50.0	−9.8 / 0.0
<i>MFN tariff τ_p^{MFN} (common across policy regimes)</i>			
		<i>Equal-wtd</i>	<i>TW (2023)</i>
τ_p^{MFN} (mean / median)		3.3 / 2.4	3.3 / 1.4
Share with $\tau_p^{MFN} > 0$		58.8%	61.5%
Share with $\tau_p^{MFN} \geq 5$		24.2%	13.9%


Notes: Values in %; Δ in pp. Equal-weighted across the 5,077 HS6 subheadings of the parameter file in chapters 16–97, which include codes of the HS2012, HS2017 and HS2022 revisions. TW (2023): weighted by 2023 U.S. imports from Mexico (UN Comtrade). RVC on net-cost basis.

4.3 Structural Estimation

The remaining parameters are the firm-level capability and demand vectors $\mathbf{a}_f = (a_{f,s})_{s \neq \text{USCA}}$ and $\mathbf{b}_f = (b_{f,d})_{d \in \mathcal{D}}$, and the location parameters of the per-source and per-destination fixed-cost distributions $\{\mu_s, \mu_d\}$. The strategy consists in recovering $(\mathbf{a}_f, \mathbf{b}_f)$ for every firm by inverting the model on the cells where the firm is observed, and predicting them from a Heckman-corrected regression on the cells where it is not. I then estimate the fixed-cost location vector by inverting the model on aggregate participation shares, with one moment per parameter. The normalizing coordinate is $a_{f,\text{USCA}} \equiv 1$, given that US–Canada is the source block with the highest participation (Table 7), meaning that the ratio denominator $\mu_{f,\text{USCA}}$ is non-zero for the largest possible subsample of firms. No normalization is imposed on the demand side, as $b_{f,d}$ is inverted from (22) at all three destinations and its imputation regression is run on $\log b_{f,d}$ directly.

Firm primitives $(\mathbf{a}_f, \mathbf{b}_f)$: inversion on observed cells. On the cells where firm f actually sources from s or sells to d , the primitives can be backed out directly from the model. Log-inverting the Fréchet share equation (7) at $a_{f,\text{USCA}} \equiv 1$ gives the capability, while summing the CES revenue equation (13) across the firm’s products at destination d

Figure 3: Trade-weighted CDFs of the policy parameters α_p and μ_p^{RVC} under NAFTA and USMCA



Notes: 2018 weights for NAFTA, 2023 for USMCA.

and solving for the demand shifter gives:

$$\log a_{f,s} = \frac{1}{\theta} [\log(\mu_{f,s}/\mu_{f,USCA}) - \log(T_s(w_s\tau_s)^{-\theta})], \quad (21)$$

$$b_{f,d}^{\sigma-1} = \frac{\sum_p R_{f,p,d}}{E_d P_d^{\sigma-1} \sum_{p \in \mathcal{P}_f} \varpi_{f,p} c_{f,p,d}(\mathcal{S}_f)^{1-\sigma}}, \quad (22)$$

where (E_d, P_d) is the market bundle of destination d , Mexico included, set as in Section 4. These inversions are exact on $\{f : \mu_{f,s} > 0, \mu_{f,USCA} > 0\}$ and $\{f : Y_{f,d} > 0\}$, and are undefined elsewhere. Filling the unsourced cells with zeros would eliminate the extensive margin in the counterfactuals, given that no firm could ever flip into a new sourcing country or destination under a policy shock. I instead fill the missing cells with a Heckman-corrected regression prediction.

Firm primitives ($\mathbf{a}_f, \mathbf{b}_f$): Heckman two-step on missing cells. Fitting the mean functions in (21) and (22) by standard OLS on the selected subsample would be upward-

biased, given that the firms observed sourcing from s are precisely those with a high latent $a_{f,s}$, so the fitted coefficients would overstate the unconditional expectation. I therefore stack the firm–source pairs for the five non-USCA source blocks into a single long panel (US–Canada, being the normalizing coordinate, is not imputed) and fit one Heckman (1979) two-step regression for the capability latent, and the analogous regression stacking all three destinations for the demand latent. Each regression also pools the two vintages, meaning that the 2018 and 2023 panels are stacked with a vintage fixed effect in both the selection and the outcome stage, so that the two equilibria share one imputation surface up to a level shift, and neither vintage’s imputed cells rest on its own sample alone. The capability panel comprises the firms with $\mu_{f,\text{USCA}} > 0$ (4,939 in 2018 and 4,452 in 2023), for which the normalized log-ratio is defined, times the five blocks, which gives 46,955 firm–block rows in all, whereas the remaining 780 and 681 firms receive fitted values from the estimated equation. The demand panel stacks the exporters of both vintages, that is, the firms with US&CA or RoW sales (4,921 and 4,449), times the three destinations (28,110 rows), and import-only firms receive fitted values on the foreign coordinates. Note that fitting the demand side on every firm would let the import-only firms, which have no foreign sales by construction, enter the selection equation as firms that chose to stay out of both foreign markets. Helpman et al. (2008) apply the same correction to selection into export markets, which here is extended to selection into both source and destination markets.

$$\text{Stage 1 (selection): } \Pr(\mu_{f,s} > 0 \mid X_{f,s}, W_f) = \Phi(X'_{f,s}\delta_a + \delta_a^W W_f), \quad (23)$$

$$\text{Stage 2 (outcome): } \log(\mu_{f,s}/\mu_{f,\text{USCA}}) = X'_{f,s}\beta_a + \lambda_a \text{IMR}(X'_{f,s}\hat{\delta}_a + \hat{\delta}_a^W W_f) + u_{f,s}, \quad (24)$$

where the outcome-stage regressor vector $X_{f,s}$ stacks firm size ($\log L_f$), source-block fixed effects, ownership-country fixed effects, NAICS-3 sector fixed effects and the vintage fixed effect, and W_f is the *exclusion restriction*, which enters the selection stage only. This exclusion is the firm’s establishment count ($\log N_{\text{estab}}$) on the capability side and its product-range breadth ($\log N_{\text{HS6}}$) on the demand side, and Table 8 reports the first-stage coefficient on each. $\text{IMR}(z) = \phi(z)/\Phi(z)$ is the inverse Mills ratio. Source heterogeneity is identified within the single pooled regression, with the source fixed effects picking up the mean log-share differential across blocks.

The choice of exclusions mirrors the model’s own split between the extensive and the intensive margin. A multi-establishment firm spreads the fixed cost of a foreign-supplier relationship (procurement staff, customs registration, supplier search) over a larger operational

base, which makes entry into a source country more likely, whereas conditional on sourcing from s , the share equation (7) is scale-free and pinned down by $a_{f,s}$ alone. Symmetrically, a broad product portfolio raises the variable-profit return to paying a destination's entry cost once, but conditional on serving d , the inversion (22) nets the firm's product mix out of $b_{f,d}$, so breadth should not move the shifter itself. The identifying assumption is thus that establishment count and product breadth shift fixed-cost capacity, rather than proxy for latent capability or appeal conditional on the included controls. Under bivariate normality of $(\varepsilon_{f,s}, u_{f,s})$, the linear-IMR form is the closed-form correction with $\lambda_a = \rho_a \sigma_u$, and the fitted $\hat{\beta}_a$ then recovers the bias-corrected unconditional mean $X'_{f,s} \hat{\beta}_a = \mathbb{E}[\log(\mu_{f,s}/\mu_{f,\text{USCA}}) \mid X_{f,s}]$. Plugging this into (21) delivers the imputed $\hat{a}_{f,s}$ on the missing cells:

$$\log \hat{a}_{f,s} = \frac{1}{\theta} [X'_{f,s} \hat{\beta}_a - \log(T_s(w_s \tau_s)^{-\theta})] \quad \text{for } \mu_{f,s} = 0 \text{ or } \mu_{f,\text{USCA}} = 0,$$

while the closed-form inversion (21) is retained on the observed cells. The symmetric procedure on the pooled firm–destination panel gives $\hat{b}_{f,d}$. On both sides, an imputed cell is capped at the largest value inverted from data in the same block or destination. The resulting firm-by-source matrix $\hat{\mathbf{A}}$ and firm-by-destination matrix $\hat{\mathbf{B}}$ are the basis for the counterfactual extensive-margin responses, as under a policy shock a firm with an inactive cell whose imputed capability is high enough to clear the fixed cost may flip into a new sourcing or destination set.

Fixed costs (κ_s, κ_d): moment inversion on aggregate participation shares with lognormal draws. Firm-country and firm-destination fixed costs are drawn iid from log-normal distributions, following Antràs et al. (2017) and Morán and Cebrenros (2025):

$$\kappa_{f,s} \sim \text{LogNormal}(\mu_s, \delta), \quad \kappa_{f,d} \sim \text{LogNormal}(\mu_d, \delta), \quad (25)$$

with one draw per firm per source country and one per firm per destination. κ is the fixed cost that enters the firm's profit function (14), and μ is the log-space location parameter of its distribution, with $E[\kappa_c] = \exp(\mu_c + \delta^2/2)$. I pin the dispersion at $\delta = \sqrt{\log 3}$, following Morán and Cebrenros (2025).

Following Antràs et al. (2017), I pin $\kappa_s^{MX} \equiv 0$, given that every Mexican manufacturer sources domestically by construction. Even if a firm sources no material inputs domestically, every firm sources at least some labor domestically, so there is no meaningful *entry* fixed cost into the Mexican input market. The μ of the remaining eight blocks (five non-MX sources and three destinations) are estimated. The estimation recovers the eight location parameters

$$\boldsymbol{\mu} = \{\mu_s^{\text{USCA}}, \mu_s^{\text{AX}}, \mu_s^{\text{CH}}, \mu_s^{\text{EU}}, \mu_s^{\text{RoW}}, \mu_d^{\text{US\&CA}}, \mu_d^{\text{MX}}, \mu_d^{\text{RoW}}\},$$

one per estimated block. The results table reports $\hat{\kappa}_c = \exp(\hat{\mu}_c + \delta^2/2)$, and the rest of the paper refers to fixed costs as κ .

Solving the model at candidate $\boldsymbol{\mu}$. At every candidate $\boldsymbol{\mu}$, each firm receives its own draw $\kappa_{f,c} = \exp(\mu_c + \delta z_{f,c})$ at each of the nine blocks, with $\kappa_{f,s}^{MX} = 0$ by pin. The standard normals $z_{f,c}$ are drawn once and held fixed across every evaluation, so that the simulated moments use common random numbers. The firm's problem (14) is then solved once for each firm at its own draw, which gives the sourcing set $\mathcal{S}_f$ and the destination set $\mathcal{D}_f$.

Moment vector. Eight participation moments identify the eight location parameters, one for each. Five of them identify the non-MX source-side fixed costs through the extensive margin of sourcing:

$$m_s^{\text{source}} = \Pr(\mu_{f,s} > 0), \quad s \in \{\text{USCA}, \text{AX}, \text{CH}, \text{EU}, \text{RoW}\}.$$

The other three identify the destination-side fixed costs through the extensive margin of selling into each destination:

$$m_d^{\text{dest}} = \Pr(Y_{f,d} > 0), \quad d \in \{\text{US\&CA}, \text{MX}, \text{RoW}\}.$$

Further moments are computed but not targeted, and serve as checks in Section 5.1. Among them are the aggregate compliance rate (the share of MX→US import value claimed under preference, from USA Trade Online) and the share of firms sourcing from at least two foreign blocks, which contains information on the joint distribution of participation that the marginal shares do not.

Estimator. Let $\hat{\mathbf{m}}(\boldsymbol{\mu})$ collect the eight simulated participation rates at a candidate $\boldsymbol{\mu}$, let $\mathbf{m}$ collect their data counterparts, and stack the gaps as $\hat{\mathbf{g}}(\boldsymbol{\mu}) = \hat{\mathbf{m}}(\boldsymbol{\mu}) - \mathbf{m} \in \mathbb{R}^8$. The estimator is the root

$$\hat{\boldsymbol{\mu}} : \quad \hat{\mathbf{g}}(\hat{\boldsymbol{\mu}}) = \mathbf{0}. \tag{26}$$

Participation in block c is weakly decreasing in μ_c , given that raising μ_c raises $\kappa_{f,c}$ for every firm, which can only remove firms from the block, so each gap is monotone in its own parameter. With eight parameters against eight monotone moments, (26) is exactly identified, and I solve it by inversion rather than by minimizing a distance. This is the just-identified case of the simulated method of moments, in which the moment conditions (26) hold with equality up to the $1/N$ step of the simulated shares. Recovering parameters by exact inversion, with one moment per parameter, is standard practice in quantitative trade models. To state an example, Melitz and Redding (2015) choose the fixed exporting cost so as to reproduce the observed share of exporting firms, and, as Redding and Rossi-Hansberg (2017)

note for the class of exactly identified models that grew out of [Eaton and Kortum \(2002\)](#), observed outcomes map one-to-one into the unobserved primitives that rationalize them as an equilibrium.

5 Estimation Results

Section [5.1](#) reports the structural estimates and the fit of the model to targeted and untargeted moments in both vintages. Section [5.2](#) then examines the estimated firm primitives in the 2023 vintage, while Appendix [E](#) reports the same objects for 2018.

5.1 Parameter Estimates and Model Fit

Table 6: Structural parameter estimates, both vintages

Parameter	2018 vintage		2023 vintage	
	$\hat{\kappa}$	SE_{κ}	$\hat{\kappa}$	SE_{κ}
<i>Sourcing fixed costs κ_s</i>				
US–Canada	3.68	0.42	3.52	0.47
Asia excl. China	0.92	0.07	0.98	0.09
China	0.67	0.07	0.80	0.09
Europe	0.76	0.08	1.04	0.09
Rest of World	0.42	0.04	0.60	0.05
<i>Destination entry fixed costs κ_d</i>				
US&CA	15.40	1.26	19.90	1.65
Mexico	49.68	4.45	50.95	5.95
Rest of World	4.52	0.65	4.76	0.74
Firms	5,719		5,133	
firm×HS6 products	136,303		122,478	

Notes: $\hat{\kappa}$ is the mean of the fitted distribution, $\exp(\hat{\mu} + \delta^2/2)$. SE_{κ} is the bootstrap standard deviation of $\hat{\kappa}$ itself ($B = 100$ firm-cluster replicates), not of $\hat{\mu}$. $\kappa_s^{\text{MX}} \equiv 0$. Fixed costs are in units of Mexican labor (worker-years at the pinned wage).

There are two features of the estimates in Table [6](#) that are worth pointing out. The first is that, on the sourcing side, US–Canada is the most expensive block to enter in both vintages, at about 3.5 worker-years, whereas every other block costs about one worker-year or less. Between the two vintages, the US–Canada cost falls by 4% while the others rise, the least for Asia excluding China (+6%) and the most for Europe (+36%) and the Rest of World (+42%), with China in between (+19%). On the destination side, the domestic market has the largest cost, at about 50 worker-years in both vintages. The reason is that nearly every firm has a high inverted domestic demand shifter, so only a large entry cost can

reconcile those shifters with the 11 to 14% of firms that do not sell at home. The US&CA entry cost rises by 29% and the Rest-of-World cost by 5%. When measured against firm size, these costs are small, as the largest of them, domestic entry, is about 0.3% of mean firm revenue, and US–Canada sourcing is about 0.02%. The second feature is that all eight location parameters are interior, with bootstrap intervals well inside the inversion bracket, and that the inversion meets every participation target (Table 7).

Model fit. The inversion matches all eight targeted participation shares (Table 7). Among the untargeted moments, the share of firms sourcing from at least two foreign blocks is close to the data (88.7% against 86.6% in 2018, 90.5% against 87.4% in 2023), and so are the US&CA share of revenue (44% and 42%) and total sales, within 4% of the census. The model overpredicts the compliance rate (86.5% against 56% in 2018, 86.8% against 48% in 2023), mostly because zero-MFN products comply at no cost; on positive-MFN lines the gap narrows to 59.0% against 52.3% and 60.7% against 46.5%. What remains is a cost of claiming that the model omits: once the 2025 surcharge made non-compliance expensive, observed compliance rose from 44% to 85% (Table 3), so the counterfactuals speak to the response to policy rather than to the level of utilization at low tariffs.

Selection-corrected imputation. Table 8 reports the outcome stage of the two pooled Heckman regressions. On both sides, the inverse-Mills loading is positive and precisely estimated, at 3.76 (s.e. 0.36) for capabilities and 2.62 (s.e. 0.30) for demand. This means that the firms observed sourcing from a block, or selling in a market, have a higher latent capability or appeal than what their observables predict, so unadjusted OLS on the observed cells would overstate the latent means, and the unconditional fit that fills the missing cells lies below the observed cells. The capability surface, which is estimated on the log share ratio, that is, $\theta \log a_{f,s}$ up to a block constant (21), rises with firm size and shifts up by 0.19 in the 2023 vintage (0.11 in $\log a_{f,s}$).

Estimated firm primitives. Table 9 reports the cross-firm distribution of $\log a_{f,s}$ and $\log b_{f,d}$ implied by the inversions and the Heckman-corrected fits above, together with the share of firms for which each coordinate is inverted from data rather than imputed. Capabilities have wide spreads (standard deviations of 1.5 to 1.6 log units) and a clear ordering across sources, where Mexico has the highest mean, the AX, CH and EU blocks cluster near a common lower mean, and RoW lies below them. The demand coordinates are ordered as in the data, with Mexico first and the Rest of World, which is served by just over half of the firms, three to four log units below the other two.

Table 7: Targeted and untargeted moments: data vs. model, both vintages

Moment	2018 vintage			2023 vintage		
	Data	Model	Diff	Data	Model	Diff
<i>Source-side participation shares</i> $\Pr(\mu_{f,s} > 0)$ (targeted)						
$s = \text{US\&CA}$	0.864	0.864	+0.000	0.867	0.868	-0.000
$s = \text{AX}$	0.738	0.739	-0.000	0.762	0.762	-0.000
$s = \text{CH}$	0.784	0.784	+0.000	0.815	0.815	+0.000
$s = \text{EU}$	0.759	0.759	-0.000	0.770	0.770	+0.000
$s = \text{RoW}$	0.646	0.645	+0.000	0.650	0.650	+0.000
<i>Destination-side participation shares</i> $\Pr(Y_{f,d} > 0)$ (targeted)						
$d = \text{US\&CA}$	0.780	0.780	-0.000	0.799	0.799	-0.000
$d = \text{MX}$	0.861	0.861	+0.000	0.892	0.892	+0.000
$d = \text{RoW}$	0.527	0.527	-0.000	0.555	0.556	-0.000
<i>Untargeted moments</i>						
Compliance rate	0.559	0.865	-0.306	0.480	0.868	-0.388
Compliance rate, positive-MFN lines	0.523	0.590	-0.067	0.465	0.607	-0.142
Multi-source share (≥ 2 foreign)	0.866	0.887	-0.021	0.874	0.905	-0.031
Revenue share, US&CA	0.444	0.440	+0.004	0.416	0.415	+0.001
Revenue share, Mexico	0.494	0.499	-0.005	0.529	0.531	-0.001
Revenue share, RoW	0.062	0.061	+0.000	0.054	0.054	+0.000
Domestic share of materials	0.367	0.432	-0.065	0.328	0.466	-0.139

Notes: Diff = Data – Model. Targeted moments coincide up to the bisection tolerance. Compliance rate: share of US&CA-bound value claimed under the agreement in force, NAFTA rules for the 2018 vintage (equilibrium E_A of the decomposition) and USMCA rules for the 2023 vintage; data are the observed MX→US utilization in 2017 and 2024 (Table 3). Positive-MFN lines: value claimed on lines with a positive MFN rate over all value, data for 2018 and 2024 (Section 4.2). Multi-source moment counts foreign blocks only. Revenue shares and the domestic share of materials are means of firm shares weighted by census revenue and census materials, model at each vintage’s baseline solve. The 2018 vintage is inverted, and its baseline solved, under USMCA rules, here and in Table 6, Appendix D and Appendix E; only the two compliance-rate rows re-solve it under NAFTA rules, as equilibrium E_A .

5.2 Validating the Estimated Firm Primitives

This subsection reports the 2023 vintage throughout, while Appendix E reports each object for the 2018 vintage.

Ownership and the geography of the primitives. Tables 10a and 10b split the two blocks of Table 9 by the firm’s ownership country for the 2023 vintage, and on the capability side one can observe that the geography lines up with ownership. Whenever a group’s home region appears in the sourcing menu, it is that group’s best coordinate (the bold diagonal of Table 10a). Mexican-owned firms hold their highest capability in Mexico (+0.86), AX-owned firms in Asia excluding China (+1.52, the largest entry in the panel), Chinese-owned firms in China (+1.40) and EU-owned firms in Europe (+0.74), whereas US–CA-owned firms, whose home block is the normalizing one ($a_{f,\text{USCA}} \equiv 1$), lie below zero on every other block, meaning that their home block is their best coordinate as well. The AX- and CH-owned peaks are

Table 8: Heckman second stage: capability and demand imputation equations

	Capability $\log(\mu_{f,s}/\mu_{f,USCA})$	Demand $\log b_{f,d}$
$\log L_f$	0.329 (0.034)	0.400 (0.018)
Inverse Mills ratio ($\hat{\lambda}$)	3.756 (0.362)	2.620 (0.296)
Constant	-2.523 (0.456)	6.366 (0.152)
<i>First stage (selection probit)</i>		
Exclusion: $\log N_{estab}$	0.043 (0.011)	
Exclusion: $\log N_{HS6}$		0.086 (0.008)
Source FE	Yes	
Destination FE		Yes
Ownership FE	Yes	Yes
Sector (NAICS-3) FE	Yes	Yes
Year FE	Yes	Yes
N (panel)	46,955	28,110
N (selected)	36,370	22,466

Notes: Heckman (1979) two-step on the stacked 2018 and 2023 panels, pooled across blocks, with a vintage fixed effect in both stages; SEs in parentheses. Ownership base: Mexico (six source-block groups on the capability side; Mexico, US-CA and other foreign on the demand side). Each exclusion enters only its own side's selection probit. Capability panel: firms with $\mu_{f,USCA} > 0$ times five blocks; demand panel: exporters (US&CA or RoW sales) times three destinations.

Table 9: Estimated firm primitives: cross-firm distribution (2023 vintage)

Coordinate	mean	sd	observed
<i>Capabilities</i> $\log a_{f,s}$ ($USCA \equiv 0$)			
Mexico	0.377	1.480	0.55
Asia excl. China	-0.581	1.532	0.73
China	-0.502	1.575	0.75
Europe	-0.451	1.519	0.72
Rest of World	-0.923	1.597	0.62
<i>Demand shifters</i> $\log b_{f,d}$			
US&CA	8.583	2.185	0.80
Mexico	9.828	2.163	0.89
Rest of World	5.519	2.353	0.56

Notes: 5,133 traders. Direct inversion on observed cells; Heckman-fitted mean on missing cells, capped at the largest inverted value of the coordinate. “observed”: share of firms with the coordinate inverted from data; a capability coordinate counts as observed only for firms that also source from US-Canada, the normalizing block. $\log a_{f,USCA} \equiv 0$. Means and standard deviations aggregate the ownership-group statistics of the replication workbook by firm counts. 2018 vintage in Appendix E.

about two log points above the corresponding Mexican-owned means, and only the RoW-owned group, which is a residual composite of origins, has no peak in its own block. The demand side shows the same home-region pattern (Table 10b), as in each market the group with the highest mean shifter is the one from that region: Mexican-owned firms in Mexico (10.3), US–CA-owned firms in the US&CA market (9.5) and the remaining foreign owners in the Rest of World (6.2).

These level shifts have direct behavioral content. A higher $a_{f,s}$ raises both the probability that block s clears the firm’s fixed-cost hurdle and the expenditure share $\mu_{f,s}$ conditional on entry, while a higher $b_{f,d}$ scales revenue in market d , so that foreign affiliates source more from, and sell more into, their parents’ home regions. Note that this pattern comes from the inverted cells.

Table 10a: Estimated sourcing capabilities by ownership group (2023 vintage)

	Ownership group					
	Mexico	U.S. & CA	Asia excl. China	China	Europe	Rest of World
Firms	2,639	1,470	228	23	513	260
<i>Mean capability log $a_{f,s}$, by source block ($USCA \equiv 0$)</i>						
Mexico	0.86	−0.61	0.85	0.57	0.55	0.35
US–Canada	0.00	0.00	0.00	0.00	0.00	0.00
Asia excl. China	−0.74	−0.71	1.52	0.58	−0.45	−0.50
China	−0.50	−0.70	0.09	1.40	−0.25	−0.52
Europe	−0.36	−1.11	−0.52	−0.47	0.74	0.03
Rest of World	−0.82	−1.46	−0.94	−0.55	−0.15	−0.44

Notes: Means over all firms in each group, observed and Heckman-imputed cells together; per-coordinate observed shares in Table 9. Bold: each group’s home block. $\log a_{f,USCA} \equiv 0$ for every firm, so the US–Canada row is zero and each entry is a capability relative to US–Canada sourcing. Owners from Asia excl. China are those from Japan, Korea and Singapore, the Asian origins coded in both censuses; owners from other Asian countries are in the Rest of World group. 2018 vintage in Appendix E.

Table 10b: Estimated demand shifters by ownership group (2023 vintage)

	Ownership group		
	Mexico	U.S. & CA	Other foreign
Firms	2,639	1,470	1,024
<i>Mean demand shifter log $b_{f,d}$, by destination</i>			
Mexico	10.34	9.06	9.61
U.S. & CA	8.12	9.48	8.49
Rest of World	5.37	5.32	6.19

Notes: Means over all firms in each group, observed and Heckman-imputed cells together. Owners are grouped by destination block, so “Other foreign” pools the firms owned from Asia excl. China, China, Europe and the Rest of World in Table 10a. Bold: each group’s home market. 2018 vintage in Appendix E.

Variation explained by observables. Table 11 reports the share of the cross-firm variation in each $\log a_{f,s}$ and $\log b_{f,d}$ latent that is explained by ownership-country and industry (NAICS-3) fixed effects and their interaction. Let the primitive coordinates be indexed by k , which is a source block s on the capability side or a destination d on the demand side, and let $x_{f,k}$ be firm f 's estimated value of coordinate k , that is, $x_{f,k} = \log a_{f,s}$ for $k = s$ and $x_{f,k} = \log b_{f,d}$ for $k = d$, inverted or imputed as in Table 9. Let $o(f)$ be the firm's ownership group and $n(f)$ its NAICS-3 sector. For each coordinate, I estimate four regressions by OLS,

$$\begin{aligned} x_{f,k} &= \xi_{k,o(f)}^o + e_{f,k}^{\text{own}}, & x_{f,k} &= \xi_{k,n(f)}^n + e_{f,k}^{\text{ind}}, \\ x_{f,k} &= \xi_{k,o(f)}^o + \xi_{k,n(f)}^n + e_{f,k}^{\text{both}}, & x_{f,k} &= \xi_{k,o(f),n(f)}^{on} + e_{f,k}^{\times}, \end{aligned} \quad (27)$$

where $\xi_{k,o}^o$ is a fixed effect for ownership group o , $\xi_{k,n}^n$ is a fixed effect for sector n , $\xi_{k,o,n}^{on}$ is a fixed effect for each ownership-sector cell, all of them specific to coordinate k , and $e_{f,k}^m$ is the residual of specification m . The table reports the R^2 of each of these regressions. The exercise addresses a natural question that arises about the Heckman imputation, namely how much of the heterogeneity in $(\mathbf{a}_f, \mathbf{b}_f)$ is captured by the coarse observables (ownership and sector) that feed the mean function, and how much of it is idiosyncratic. Even when fully interacted, the observables explain at most 25% of the cross-firm variation in any primitive coordinate. Ownership is the stronger of the two observables for the MX, EU and RoW capability coordinates and ties with sector for AX, whereas sector is stronger for China and for all three demand coordinates. Given that the imputed cells are fitted values of a mean function that contains the same ownership and sector dummies, the R^2 are biased upward. Nevertheless, the conclusion that observables explain a minority of the variation survives this bias.

Cross-country correlation structure of the primitives. Figure 4 reports the sample correlation of $(\log a_f, \log b_f)$ across the 2023 firms. The capability block (Panel a) shows a coherent East-Asian cluster, as the AX-CH correlation of 0.54 is the largest pairwise correlation, above the 0.26 to 0.49 range of the other block pairs. Firms with an above-average China-sourcing capability tend to have an above-average AX capability as well, which is consistent with a latent “regional supplier-network access” factor. This correlation rationalizes part of the sourcing-reallocation pattern in Section 6.6, where sourcing shifts toward Asia excluding China about twice as much, in proportion, as toward any other block when China becomes more expensive. The demand block (Panel b) is weakly and positively correlated, at 0.29 for US&CA-RoW, 0.26 for Mexico-RoW and 0.13 for US&CA-Mexico, meaning that firms with a high appeal in one foreign market tend to have it in the other as well rather than specializing in one.

Table 11: Share of the cross-firm variation in $\log a_{f,s}$ and $\log b_{f,d}$ explained by ownership and sector (2023 vintage)

Latent	Coordinate	R^2 from fixed effects for			
		Own	Sector	Own + Sector	Own $\times$ Sector
$\log a$	MX	0.188	0.093	0.234	0.254
	AX	0.094	0.094	0.173	0.189
	CH	0.020	0.069	0.094	0.114
	EU	0.122	0.045	0.158	0.177
	RoW	0.063	0.045	0.099	0.119
$\log b$	US&CA	0.071	0.186	0.223	0.237
	MX	0.068	0.103	0.151	0.163
	RoW	0.021	0.111	0.138	0.151

Notes: 5,133 traders. R^2 from OLS on ownership dummies alone, NAICS-3 dummies alone, their union, and full interaction, equation (27). Ownership has six levels for the capability rows and three levels (Mexico, US-CA, other foreign) for the demand rows. 2018 vintage in Appendix E.

Figure 4: Correlation structure of estimated firm primitives (2023 vintage)

	MX	AX	CH	EU	RoW
MX	1.00				
AX	.26	1.00			
CH	.33	.54	1.00		
EU	.47	.35	.38	1.00	
RoW	.38	.30	.32	.49	1.00

(a) Sourcing capabilities $a_{f,s}$

	US&CA	MX	RoW
US&CA	1.00		
MX	.13	1.00	
RoW	.29	.26	1.00

(b) Demand shifters $b_{f,d}$

Notes: 5,133 traders. Panels show the lower triangle of the correlation matrix (upper triangle omitted by symmetry). Panel (a) omits USCA ($\log a_{f,USCA} \equiv 0$). 2018 vintage in Appendix E.

6 Counterfactual Analysis

This section begins with a decomposition of the observed 2018–2023 transition, and then sweeps individual policy instruments around the 2023 baseline: rules of origin (Section 6.2), the MFN tariff (Section 6.3), the agreement itself (Section 6.4), the interaction between content rules and the outside tariff (Section 6.5), and tariffs on Chinese inputs (Section 6.6). It is worth pointing out that, unlike the decomposition, each of these policy experiments changes the cost of the option that was the cheapest for a firm at baseline, given that a tighter rule of origin raises the cost of complying, whereas a higher tariff raises the cost of not complying. Each firm then re-optimizes, meaning that responses are heterogeneous: for some firms not complying remains the cheapest option, for others complying becomes the cheapest option, and some others keep complying at a higher cost. The aggregate patterns

reported below summarize these firm-level choices. In every counterfactual, the foreign pairs (E_d, P_d) are held fixed and the Mexican pair (E_M, P_M) moves with income through the demand loop (34) (Section 3.4).

6.1 Two-Equilibrium Decomposition

I calibrate the model at two different vintages. Equilibrium E_A uses CE 2019 and PEME 2018 with NAFTA rules of origin and pre-trade-war sourcing conditions $\{\psi_s\}$, whereas equilibrium E_B uses CE 2024 and PEME 2023 with USMCA rules and post-trade-war $\{\psi_s\}$. There are five inputs that move between the two: relative sourcing conditions $\{\psi_s\}$, rules of origin (α_p, μ_p^{RVC}) , the fixed costs of sourcing and destination participation (κ_s, κ_d) , the foreign expenditure and price pairs (E_d, P_d) , and the remaining fundamentals (the vintage level of the pooled imputation surface and the trader panel itself, given that the Fréchet shape θ is pinned at 1.79 and the sector cost shares γ_s are frozen at their two-wave mean throughout). In order to separate them, Table 12 constructs a ladder of intermediate equilibria in which each of the first four inputs is flipped in isolation to its later-vintage value. Equilibrium (5) flips all four jointly, and E_B additionally re-estimates the remaining objects at the later vintage. The ladder also includes the joint sourcing-and-rules step (1+2), which recovers the two-input version of this exercise but plays no role in the identity that follows.

Table 12: Two-equilibrium decomposition: which inputs change at each step.

		$\{\psi_s\}$	α_p, μ_p^{RVC}	κ_s, κ_d	E_d, P_d	Imp., panel
E_A	Baseline	Pre-TW	NAFTA	2018 est.	2018	2018
(1)	+ Sourcing cond.	Post-TW	NAFTA	2018 est.	2018	2018
(2)	+ RoO	Pre-TW	USMCA	2018 est.	2018	2018
(3)	+ Entry costs	Pre-TW	NAFTA	2023 est.	2018	2018
(4)	+ Foreign demand	Pre-TW	NAFTA	2018 est.	2023	2018
(1+2)	+ Sourcing & RoO	Post-TW	USMCA	2018 est.	2018	2018
(5)	+ All four	Post-TW	USMCA	2023 est.	2023	2018
$\bar{E}_B$	Full recalibration	Post-TW	USMCA	2023 est.	2023	2023

Notes: Bold marks an input set to its later-vintage value. $\theta = 1.79$ is pinned at every step. The last column collects the objects re-estimated only at E_B : the vintage level of the pooled Heckman imputation surface, the closure scalars ($w_m = A_o, \bar{L}$ and the domestic anchor), and the trader panel itself (5,719 firms at E_A through (5), 5,133 at E_B). The sector cost shares γ_s are frozen at their two-wave mean and are therefore common to every row. The 2018 primitives are inverted under USMCA rules; E_A re-solves that panel under NAFTA rules. Row (1+2) flips sourcing conditions and rules jointly; it is the two-input version of this exercise, reported for reference, and does not enter the six-component identity below.

Letting Y_k denote the value of an outcome Y at equilibrium $k \in \{A, 1, 2, 3, 4, 5, B\}$, with

$Y \in \{W_M, \Pi_M, P_M, M_{MX,s}\}$, the total change admits the exact decomposition

$$\begin{aligned} \underbrace{Y_B - Y_A}_{\text{total change}} = & \underbrace{Y_1 - Y_A}_{\substack{\text{sourcing conditions} \\ (\Delta_\psi)}} + \underbrace{Y_2 - Y_A}_{\substack{\text{USMCA rules} \\ (\Delta_{RoO})}} + \underbrace{Y_3 - Y_A}_{\substack{\text{entry costs} \\ (\Delta_\kappa)}} + \underbrace{Y_4 - Y_A}_{\substack{\text{foreign demand} \\ (\Delta_D)}} \\ & + \underbrace{Y_5 - Y_A - \Delta_\psi - \Delta_{RoO} - \Delta_\kappa - \Delta_D}_{\substack{\text{interaction} \\ (\Delta_\times)}} + \underbrace{Y_B - Y_5}_{\substack{\text{structural residual} \\ (\Delta_S)}}. \end{aligned}$$

Two of these components reflect the trade war. Δ_ψ is the shift in relative sourcing conditions. Given that the bundle $\psi_s \equiv T_s(w_s \tau_s)^{-\theta}$ is backed out from observed sourcing shares, its 2018–2023 movement is plausibly dominated by the trade war and the diversion that it set off, but it also embeds foreign productivity, wage and freight drift, and these components are not separately identified. Δ_κ is the entry-cost channel of the same shock, as the estimated source-side fixed cost falls for US–Canada and rises for every other block, while the destination-side costs rise, the most for the US&CA market (Table 6). Δ_{RoO} is the partial effect of the NAFTA to USMCA parameter flip, Δ_D isolates the real growth of foreign expenditure (P_d stays at one), and $\Delta_\times$ collects the non-additive interaction of the four. Δ_S , in turn, is a mixture. It moves the pooled imputation surface to its later-vintage level and re-inverts the primitives on the later panel (γ_s is frozen across vintages and therefore contributes nothing to it), and it also includes compositional turnover, given that equilibrium E_B is solved on the 2023 firm panel (5,133 firms) while E_A and the intermediate steps use the 2018 panel (5,719 firms) (Appendix C). Note that the identity is an exact accounting over calibrated inputs, and not a causal decomposition of shocks. Table 13 reports the results of the six-component decomposition.

The rules-of-origin flip shows to be close to welfare-neutral, as Δ_{RoO} lowers real income by 0.06% and leaves the price index unchanged. The welfare movement comes from elsewhere. Foreign demand growth contributes +7.1% and the shift in sourcing conditions +2.1%, with an interaction of +0.2%, whereas the entry-cost channel is small (−0.06%). The Δ_ψ component is led by China. At 2018 primitives, the 2023 sourcing conditions raise imports from China by 13%, the largest increase of any source, and from Europe and the Rest of World by 9 to 10%, while they lower imports from US–Canada by 2% and the price index by 1.8%. This sign should not be read as the trade war raising Mexican welfare, given that $\{\psi_s\}$ absorbs foreign productivity, wage and freight drift alongside the tariff channel. The structural residual on real income is large and positive (+14.8%), and it splits about evenly between the price-index term and the wage bill $w_m \bar{L}$, which is 29% higher at E_B because w_m and $\bar{L}$ are measured on the 2024 census. E_B aggregates the CES price index over a different panel,

Table 13: Two-equilibrium decomposition: results

Outcome	E_A	Components						E_B
		Δ_ψ	Δ_{RoO}	Δ_κ	Δ_D	$\Delta_\times$	Δ_S	
<i>Aggregates</i>								
Manufacturing profits Π_M	100	+1.93	-0.09	-0.08	+9.47	+0.18	+0.36	111.78
Price index P_M	100	-1.83	+0.00	+0.01	-0.00	-0.00	-14.68	83.49
Real income W_M	100	+2.14	-0.06	-0.06	+7.11	+0.19	+14.84	124.16
<i>Intermediate purchases by source $M_{MX,s}$</i>								
US-Canada	100	-2.07	-0.03	-0.03	+9.80	-0.20	-6.28	101.19
Mexico	100	+1.85	-0.08	-0.04	+8.95	+0.17	-5.89	104.96
Asia excl. China	100	+2.25	-0.10	-0.02	+10.53	+0.19	-20.95	91.91
China	100	+13.07	-0.04	-0.03	+10.63	+1.34	-19.34	105.63
Europe	100	+9.23	-0.46	-0.04	+10.12	+0.80	-20.21	99.44
Rest of World	100	+10.13	-0.21	-0.06	+9.18	+0.89	-9.57	110.37

Notes: $E_A = 100$, components in % of the E_A level, E_B as an index. E_A and steps (1) to (5) are solved on the 2018 panel (5,719 firms), E_B on the 2023 panel (5,133 firms). Peso magnitudes are in constant 2023 pesos at every step, since the 2018 census and customs values are deflated with the INPC before estimation (Section 4). Π_M sums the profits of the traders and the non-traders. The Δ_S column of the price index reflects the re-aggregation of the index over a different panel, not a price change; see text.


which includes a non-trader term that doubles with the non-traders' share of domestic sales (18% to 30%), meaning that the -14.7% move in P_M under Δ_S is a re-aggregation rather than a price change, and that the level of W_M is not comparable across the two panels. The first four components, all of which are computed on the 2018 panel, are the economically interpretable part of the ladder. The profit row, which does not depend on the price index, shows the same ordering of the four flipped inputs, with a residual of only +0.4%. Section 6.2 sweeps the same rule parameters around their USMCA baseline in isolation.

6.2 Rules-of-Origin Strictness


Figure 5 sweeps the two content parameters around their USMCA baselines while holding everything else fixed, with the CTC input restriction α_p moving over ± 25 pp and the RVC threshold μ_p^{RVC} over $[-75, +25]$ pp. The compliance responses are both large and regime-specific.

There are two features that stand out. The first is that the compliance responses are sizeable. Around the USMCA baseline, a ± 25 pp swing in α_p moves compliance by +3.7pp at the loosened end and -16.6pp at the tightened end, which is a 20pp range. This response is driven by the CTC regime, whose rate falls from 98% to 43% across the sweep, and by the Hybrid regime, which gives up 7pp relative to the baseline at the tight end even though it can fall back on its RVC arm. The μ_p^{RVC} sweep over $[-75, +25]$ pp gives a 22pp range (+7.5pp at


Figure 5: Counterfactual sweeps of rules-of-origin strictness




(a) Compliance (α_p sweep)




(b) Compliance (μ_p^{RVC} sweep)




(c) Sourcing composition (α_p sweep)



(d) Sourcing composition (μ_p^{RVC} sweep)



(e) Welfare and components (α_p sweep)



(f) Welfare and components (μ_p^{RVC} sweep)

Notes: 2023 vintage. Horizontal axis: change in pp from the USMCA value of α_p (left column) or μ_p^{RVC} (right column); the vertical line marks the baseline. Top row: “All products” is compliance over all US&CA-bound revenue (Section 4.2). Middle row: % change relative to the baseline in each origin’s share of Mexican sourcing, $\bar{\mu}_s = M_{MX,s} / \sum_{s'} M_{MX,s'}$. Bottom row: % change in real income, the traders’ profits (non-traders excluded) and the Mexican price index; the price index moves by less than 0.001% along these sweeps and its line coincides with zero.

the loose end and -14.2pp at the tight end). Given that RVC compliance is priced rather than binary (Appendix A.4), the response to a tighter threshold is smooth. Firms whose content falls short buy the missing North American share through the restriction α^* until its cost exceeds the tariff, so the RVC rate declines gradually from 81% at the baseline to 42% at $+25\text{pp}$ rather than collapsing at a cutoff, and reaches 99% once the threshold is loosened by 40pp or more. The response is regime-specific by construction, as only CTC and Hybrid move under the α_p sweep, whereas only RVC and Hybrid move under the μ_p^{RVC} sweep.

The second feature is that the welfare panels (e, f) are not flat. Tightening α_p by 25pp lowers W_M by 0.16% and the traders' profits by 0.24%, whereas tightening μ_p^{RVC} by 25pp lowers them by 0.61% and 0.90%, and loosening raises W_M by up to 0.05% and 0.18% respectively. The price index moves by less than a thousandth of a percent and sourcing shares move by at most 0.4% (panels c, d), meaning that the cost works through profits rather than through prices or sourcing composition. Most firms stay compliant when the rule tightens (compliance is still 70 to 73% of US&CA-bound revenue at the tight ends), and they do so by buying North American content that they would not otherwise buy. This cost is worth paying because preference is valuable, as at observed tariffs the agreement is worth 2.8% of real income (Section 6.4). Thus, a 25-point tightening of the RVC threshold costs about a fifth of what the agreement is worth. Sections 6.3 and 6.5 show how this cost grows with the price of non-compliance.

6.3 The Price of Non-Compliance: MFN Surcharge Sweeps

Figure 6 adds a uniform surcharge to every product's MFN tariff on US&CA-bound shipments, which is swept from -15 to $+100$ percentage points. The surcharge is floored at zero, so the left end removes the tariff on every product with an MFN rate at or below 15 points, which is essentially all US&CA-bound revenue, whereas $+25$ reproduces the design of the March 2025 order. Panel (a) shows that compliance rises steeply through the observed point and saturates early, as it is 78% with no tariff, 87% at the observed rates, 96% at $+5\text{pp}$, 97% at $+10\text{pp}$ and 98% at $+25\text{pp}$, which is the coverage ceiling set by the 1.8% of revenue with no coded rule (Section 4.2). With no tariff the rate stays high because zero-MFN products count as claiming, and the strict rate is 8% there and 61% at observed rates. The three regimes respond in a similar way. CTC goes from 88% at observed tariffs to 99.6% at $+5\text{pp}$, Hybrid from 99.0% to 99.8%, and RVC from 81% to 95% at $+5\text{pp}$, 98% at $+10\text{pp}$ and 99.5% at $+25\text{pp}$, meaning that buying content through α^* pays off once the tariff is a few points higher. Panel (c) shows that welfare falls monotonically but slowly, by 0.18% at $+5\text{pp}$, 0.28% at $+25\text{pp}$ and 0.31% at $+100\text{pp}$, with the traders' profits falling by 0.26, 0.42 and 0.45%. Marginal firms are pushed into compliance at a cost below the tariff that they avoid, and

the residual tariff bill shrinks as they switch to preference, from 0.09% of income at observed rates to 0.03% at +25pp.


This sweep qualifies the paper’s reading of the 2025 episode. In the model, firms had the capacity to comply under every rule type well before the surcharge arrived, given that 83% of the response of compliance to a +25pp surcharge is already obtained at +5pp. In the data, the jump from 44% to 85% between January–February and March–December 2025 (Table 3) came from the CTC and Hybrid regimes, whose rates nearly tripled, whereas RVC compliance rose by six points. The model matches the RVC level at observed tariffs (81% in both) but not its small 2025 response, as it reaches 99% at +25pp against the 84% observed. It places CTC and Hybrid at 88 and 99% before the surcharge, against 32 and 34% in the 2024 data, and at +25pp it puts aggregate compliance at 98%, against the 85% observed after the order. The gap at observed tariffs is the cost of claiming that the model omits, which is concentrated in the CTC and Hybrid regimes. For those regimes, whose compliance in the data rose to 84 and 86% once the tariff was in place, the 2025 response shows that this cost is small relative to a 25pp tariff.

6.4 Eliminating the Agreement

The sweeps so far vary the terms of the agreement, whereas this exercise removes it altogether. With no preferential route in place, every MX→US&CA shipment pays the product’s MFN tariff regardless of its content, and the compliance choice d_{fp} disappears. The experiment takes place on the export side only, as Mexican imports of U.S. and Canadian inputs keep their treatment, given that the model carries no Mexican tariff schedule on North American inputs. Table 14 reports the outcome on the 2023 vintage along the surcharge grid of Section 6.3.

At 2023 tariffs, removing the agreement lowers Mexican real income by 2.8%. Tariff payments to the United States rise from 0.09% to 1.6% of nominal income, and the loss is spread across the firm block, as manufacturing sales fall by 4.1% and intermediate purchases from every source, including Mexico’s own suppliers, fall by 3.2 to 5.5%, with U.S.–Canadian inputs falling the most because the firms that lose preferential access are the ones that rely on them the most. Sourcing composition barely changes, meaning that the loss is a scale effect on US&CA-bound sales rather than a reallocation. The loss is large relative to the content-rule experiments because much of the economy is exposed. US&CA-bound sales are 42% of the panel’s revenue (Section 5.1), nearly all of it carries a coded rule and 70% of it a positive MFN rate, and each point of duty lowers the fob revenue of a non-compliant product by about $\sigma = 2.1$ percent. Furthermore, manufacturing profits Π_M , of traders and non-traders, are about three quarters of the income measure, with the wage bill at the pinned wage being the rest, so the loss in profits passes almost one for one into real income.

Figure 6: MFN surcharge sweep



Notes: 2023 vintage. Horizontal axis: surcharge in pp added to each product's own MFN rate, floored at zero; 0 is the observed schedule (solid line) and +25 the March 2025 design (dotted line). Panel (a): "All products" is compliance over all US&CA-bound revenue (Section 4.2). Panel (b): % change in each origin's share of Mexican sourcing. Panel (c): the price index moves by less than 0.001% along the sweep and its line coincides with zero. All changes relative to the observed schedule.

Table 14: Eliminating the agreement

MFN variation (pp)	−5	Baseline	+5	+25	+100
<i>Values, $X^{\text{no agr.}}(s)/X^{\text{USMCA}}(s) - 1$ (%)</i>					
Income $w_m \bar{L} + \Pi_M$	−1.36	−2.80	−6.52	−17.59	−33.72
Manufacturing sales	−1.99	−4.10	−9.56	−25.80	−49.47
Price index P_M	+0.00	+0.00	+0.01	+0.02	+0.05
Real income W_M	−1.36	−2.80	−6.52	−17.60	−33.73
<i>Intermediate purchases by source $M_{MX,s}$ (%)</i>					
$s = \text{USCA}$	−3.08	−5.52	−11.03	−27.38	−51.27
$s = \text{MX}$	−1.85	−3.87	−8.93	−23.93	−45.73
$s = \text{AX}$	−1.41	−3.37	−9.40	−27.73	−54.59
$s = \text{CH}$	−1.72	−4.13	−10.31	−28.77	−55.83
$s = \text{EU}$	−1.46	−3.15	−7.98	−23.38	−46.25
$s = \text{RoW}$	−1.77	−3.59	−8.34	−22.82	−44.17


Notes: 2023 vintage (5,133 firms). Export side only: with no agreement every MX→US&CA shipment pays the product’s MFN rate whatever its content (zero for the 1.8% of revenue with no coded rule, Section 4.2), and compliance falls mechanically to zero. The surcharge s is in pp added to each product’s own MFN rate, 0 the observed 2023 schedule and +25 the March 2025 design. Both regimes are compared at a common s , so the tariff sweep itself is differenced out.

Thus, what the agreement is worth is a function of the tariff that it exempts firms from. The loss from eliminating it rises from 2.8% of real income at the observed schedule to 6.5% at +5pp, 17.6% at the +25pp IEEPA level and 33.7% at +100pp, and on the 2018 panel, solved under USMCA rules, the same figures are 2.8, 7.1, 20.1 and 38.9%. It is worth pointing out the asymmetry with Section 6.3, where with the preferential route available the same surcharges cost at most a third of a percent, given that firms respond by complying rather than by paying. In other words, the 2025 order raised the value of USMCA membership to Mexico about sixfold, and the compliance jump documented in Section 2 is the firm-level counterpart of that revaluation.

6.5 Content Rules Under a Binding Outside Option

Sections 6.2 and 6.3 vary one instrument at a time, whereas Figure 7 crosses them. Each content parameter is swept over ± 20 pp against a nine-point MFN surcharge grid that runs from −5pp (close to free trade) through 0 (the observed schedule) to +50pp, with +25pp being the 2025 design. Each panel shows the resulting real-income surface.

Figure 7: Rules-of-origin strictness at different levels of the outside tariff




Notes: 2023 vintage. Each cell is one (content parameter, MFN surcharge) pair, shaded by the % change in real income against the observed schedule at baseline rules. Colour runs dark for low to bright for high on a viridis ramp; each panel has its own scale, shown in its legend. The surcharge steps are unevenly spaced and drawn at equal width, so the horizontal axis is ordinal rather than to scale.

The two panels make the same point, which is that content rules matter more when non-compliance is expensive. At observed tariffs, tightening α_p or μ_p^{RVC} by 20 points costs 0.14% and 0.47% of real income, whereas at the 25-point surcharge of 2025 the same tightening costs 0.48% and 1.28%, given that firms stay compliant and buy North American content instead of paying the tariff.

6.6 Tariffs on Chinese Intermediates

The last experiment is a stylized version of the tariffs that Mexico imposed in January 2026, which raised duties to 20 to 35% (up to 50%) on 1,463 tariff lines from countries without a trade agreement, China included. I model it as a uniform tariff τ_{CH} on Chinese intermediates entering Mexico, which is swept from 0 to 100% (Figure 8). Given that $\psi_s \equiv T_s(w_s \tau_s)^{-\theta}$, raising the effective input price of Chinese intermediates $w_{CH} \tau_{CH}$ by the factor $(1 + \tau_{CH})$ is equivalent to scaling the sourcing shifter by $(1 + \tau_{CH})^{-\theta}$, so the counterfactual enters the model as $\psi_{CH} \rightarrow \psi_{CH} \cdot (1 + \tau_{CH})^{-\theta}$. Tariff revenue is rebated lump-sum to the representative Mexican consumer under a balanced-budget assumption, and is added as a fiscal transfer to the numerator of (19), which has no tariff-revenue term at baseline.

Figure 8: Tariffs on Chinese intermediates entering Mexico



Notes: 2023 vintage. Panel (a): % change in imports by origin (China on right axis), tariff-inclusive spending; dotted line at 25%.

Panel (a) traces the sourcing response. A 25% tariff on Chinese intermediates cuts Chinese sourcing by 29%, at 50% the cut deepens to 46%, and it takes a 100% tariff to reach 66%. These are tariff-inclusive purchases, whereas at customs value the cuts are 43, 64 and 83%. The displaced volume is small in aggregate, and the source that gains the most in proportion lies outside North America, as at 25% Asia excluding China rises by 2.8%, while Europe, the Rest of World and US–Canada rise by 1.3% and Mexican domestic sourcing by 1.0%. Asia excluding China is the block whose capabilities co-move the most with China’s (Figure 4a), meaning that the policy substitutes first toward geographies whose input cost and productivity draws are close to China’s, and that its ability to redirect trade to North America specifically is limited by this substitution pattern.

Panel (b) shows the welfare implications of the same policy. Small tariffs raise real income slightly, by 0.05% at 5% and 0.07% at 10%, given that the tariff rebate outweighs the fall in profits and the rise in the price index. At 25, 50 and 100%, the change in W_M is -0.01 , -0.31 and -0.93% , with the price index rising by 0.5, 0.8 and 1.1% and the traders’ profits falling by 0.7, 1.2 and 2.0%, net of a tariff rebate worth 0.5 to 0.6% of income. Aggregate compliance rises across the sweep (from 86.8% at zero tariff to 88.3% at 25% and 90.9% at 100%), as forcing firms off Chinese inputs raises their North American content and pushes a small mass of them over the compliance threshold. When compared to the other instruments, a 25% upstream tariff costs almost nothing in real income while cutting Chinese sourcing by far more than any content-rule sweep, which leaves sourcing shares within half a percent of baseline, whereas a 25-point tightening of the RVC threshold costs 0.6% of real income and eliminating the agreement costs 2.8% (Section 6.4). Thus, if the policy goal is to have less

Chinese content in North American supply chains, upstream input tariffs work at a margin where rules of origin are essentially inert. The trade-off is that they raise sourcing from the rest of Asia proportionally more than from North America.

7 Conclusion

In this paper I study how firms organize North American supply chains when tariff-free access to the U.S. market must be purchased with rules-of-origin compliance, and what this trade-off implies for Mexican welfare across the trade war and the NAFTA→USMCA transition. The descriptive evidence shows that trade reorganized toward Mexico while preference utilization fell, two facts that are consistent with several policy channels, and the model allows one to separate them.

Through the lens of the model, the rule transition itself shows to be close to welfare-neutral. Foreign-demand growth and the shift in sourcing conditions toward Chinese inputs account for the interpretable part of the aggregate income change across the two vintages. Compliance responds strongly to the cost of non-compliance under every rule type, and the welfare cost of rules of origin rises with that cost.

The policy implications come from the ranking of the instruments, which in turn depends on the tariff environment. At historical tariff levels, the agreement is worth 2.8% of Mexican real income, and tightening its content rules already costs real income, as a 25-point rise in the regional value content (RVC) threshold costs 0.6%, about a fifth of the agreement's value, given that firms keep complying and buy North American content rather than lose preference. Similarly, a 25-point rise in the change-in-tariff-classification (CTC) input restriction costs 0.16% of real income. At the 2025 tariff level, the agreement is worth about six times as much, and the cost of tighter content rules rises about threefold, above the cost of the tariff itself. Upstream tariffs on Chinese intermediates, on the other hand, force real sourcing reallocation at almost no welfare cost at a 25% rate, although they raise sourcing from the rest of Asia proportionally more than from North America. Rules of origin and the outside tariff are therefore complements, in the sense that the tariff decides whether firms comply, and in turn how much the rules cost. If the goal is to limit Chinese content in North American supply chains, a tariff on Chinese inputs works at a margin where rules of origin do not.

References

Alfaro, L. and Chor, D. (2025). An Anatomy of the Great Reallocation in US Supply Chain Trade. Nber working paper, National Bureau of Economic Research.

- Antràs, P., Fort, T. C., and Tintelnot, F. (2017). The Margins of Global Sourcing: Theory and Evidence from US Firms. *American Economic Review*, 107(9):2514–2564.
- Bown, C. P. (2021). The US–China trade war and Phase One agreement. *Journal of Policy Modeling*, 43(4):805–843.
- Conconi, P., García-Santana, M., Puccio, L., and Venturini, R. (2018). From final goods to inputs: the protectionist effect of rules of origin. *American Economic Review*, 108(8):2335–2365.
- Eaton, J. and Kortum, S. (2002). Technology, Geography, and Trade. *Econometrica*, 70(5):1741–1779.
- Fajgelbaum, P., Goldberg, P., Kennedy, P., Khandelwal, A., and Taglioni, D. (2024). The US-China Trade War and Global Reallocations. *American Economic Review: Insights*, 6(2):295–312.
- Fajgelbaum, P. D., Goldberg, P. K., Kennedy, P. J., and Khandelwal, A. K. (2020). The Return to Protectionism. *The Quarterly Journal of Economics*, 135(1):1–55.
- Freund, C., Mattoo, A., Mulabdic, A., and Ruta, M. (2024). Is US Trade Policy Reshaping Global Supply Chains? *Journal of International Economics*.
- Gopinath, G., Gourinchas, P.-O., Presbitero, A. F., and Topalova, P. (2025). Changing Global Linkages: A New Cold War? *Journal of International Economics*, 153:104042.
- Grossman, G. M. (1981). The Theory of Domestic Content Protection and Content Preference. *The Quarterly Journal of Economics*, 96(4):583–603.
- Grossman, G. M., Helpman, E., and Lhuillier, H. (2023). Supply Chain Resilience: Should Policy Promote International Diversification or Reshoring? *Journal of Political Economy*, 131(12):3462–3496.
- Head, K. and Mayer, T. (2019). Brands in Motion: How Frictions Shape Multinational Production. *American Economic Review*, 109(9):3073–3124.
- Head, K., Mayer, T., and Melitz, M. (2024). The Laffer Curve for Rules of Origin. *Journal of International Economics*.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.

- Helpman, E., Melitz, M., and Rubinstein, Y. (2008). Estimating trade flows: Trading partners and trading volumes. *The Quarterly Journal of Economics*, 123(2):441–487.
- Ju, J. and Krishna, K. (2002). Regulations, Regime Switches and Non-Monotonicity When Non-Compliance is an Option: An Application to Content Protection and Preference. *Economics Letters*, 77(3):315–321.
- Melitz, M. J. and Redding, S. J. (2015). New trade models, new welfare implications. *American Economic Review*, 105(3):1105–1146.
- Morán, J. R. and Cebrreros, A. (2025). Rules of Origin and the Use of NAFTA. Working paper, Banco de México.
- Ornelas, E. and Turner, J. L. (2024). The Costs and Benefits of Rules of Origin in Modern Free Trade Agreements. *Journal of International Economics*.
- Pierce, J. R. and Schott, P. K. (2012). A concordance between ten-digit US harmonized system codes and SIC/NAICS product classes and industries. *Journal of Economic and Social Measurement*, 37(1-2):61–96.
- Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative spatial economics. *Annual Review of Economics*, 9:21–58.
- Tintelnot, F. (2017). Global Production With Export Platforms. *The Quarterly Journal of Economics*, 132(1):157–209.
- USTR (2018). Findings of the investigation into China’s acts, policies, and practices related to technology transfer, intellectual property, and innovation under Section 301 of the Trade Act of 1974. Section 301 report, Office of the United States Trade Representative, Executive Office of the President.
- Utar, H. (2026). Elsewhere in North America: How U.S. Tariffs on China Boosted Mexico’s Manufacturing Employment and Output. Cesifo working paper, CESifo.
- Utar, H., Cebrreros, A., and Torres, L. (2025). The US-China Trade War and the Relocation of Global Value Chains to Mexico. *Review of Economics and Statistics*, pages 1–47.

A Model Derivations

This appendix collects the derivations of the closed-form expressions that are used in Section 3. Section A.1 presents the Fréchet closed forms for the intermediate cost index and the sourcing shares. Section A.2 derives the Cobb-Douglas unit cost from cost minimization. Section A.3 derives the compliant unit cost under the CTC regime, and Section A.4 derives the cheapest compliant restriction under RVC. Section A.5 derives the Mexican CES price index that is used in the Mexican block closure. Section A.6 decomposes manufacturing labor demand. Lastly, Section A.7 describes how the domestic-demand loop is solved.

A.1 Fréchet Closed Forms

Conditional on the firm's sourcing set $\mathcal{S}_f$, the cost of sourcing input j from country s is given by

$$c_{f,s}(j) = \frac{w_s \tau_s}{z_s(j)},$$

where $z_s(j) \sim \text{Fréchet}(T_s a_{f,s}^\theta, \theta)$ is an idiosyncratic productivity draw, independent across countries and inputs. The firm picks the lowest-cost source for each input,

$$s^*(j) = \arg \min_{s \in \mathcal{S}_f} c_{f,s}(j).$$

Standard Fréchet algebra (Eaton and Kortum, 2002) gives the probability that country s offers the minimum cost as

$$\mu_{f,s}(\mathcal{S}_f) = \int_0^1 \mathbb{1}\{s^*(j) = s\} dj = \frac{T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta}}{\sum_{s' \in \mathcal{S}_f} T_{s'} a_{f,s'}^\theta (w_{s'} \tau_{s'})^{-\theta}}, \quad (28)$$

which is equation (7) in the main text. Integrating the minimum cost over the Fréchet distribution in turn yields the intermediate cost index

$$c_f^M(\mathcal{S}_f) = \Gamma\left(1 - \frac{1}{\theta}\right) \left(\sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-1/\theta}, \quad (29)$$

which corresponds to equation (6). Countries with a higher technological capability (T_s), stronger firm-country compatibility ($a_{f,s}$), or lower effective cost ($w_s \tau_s$) account for a larger share of the firm's intermediate sourcing.

A.2 Cobb-Douglas Unit Cost

The firm's production function combines Mexican labor ℓ and an intermediate composite M through a constant-returns Cobb-Douglas:

$$q_f = \ell_f^\gamma M_f^{1-\gamma}, \quad \gamma \in (0, 1).$$

Given the Mexican wage w_m and the intermediate-composite price $c_f^M(\mathcal{S}_f)$ from (6), the firm chooses ℓ and M in order to minimize cost subject to producing q_f units:

$$\min_{\ell, M} w_m \ell + c_f^M M \quad \text{s.t.} \quad \ell^\gamma M^{1-\gamma} \geq q_f.$$

The Lagrangian is $\mathcal{L} = w_m \ell + c_f^M M - \lambda(\ell^\gamma M^{1-\gamma} - q_f)$, with first-order conditions

$$w_m = \lambda \gamma \ell^{\gamma-1} M^{1-\gamma}, \quad c_f^M = \lambda(1 - \gamma) \ell^\gamma M^{-\gamma}.$$

Dividing the two FOCs delivers the Cobb-Douglas cost-share property,

$$\frac{w_m \ell}{c_f^M M} = \frac{\gamma}{1 - \gamma}, \tag{30}$$

which states that labor's share in total variable cost equals the production-function exponent γ . This identity in turn justifies the calibration $\hat{\gamma} = wL/(wL + M_t)$ in Section 4.

Rearranging (30) gives labor as a function of the intermediate composite at the optimum,

$$\ell = \frac{\gamma}{1 - \gamma} \frac{c_f^M}{w_m} M.$$

Substituting this expression into the binding production constraint $\ell^\gamma M^{1-\gamma} = q_f$ in turn collapses the constraint to a single equation in M :

$$q_f = \left(\frac{\gamma}{1 - \gamma} \frac{c_f^M}{w_m} M \right)^\gamma M^{1-\gamma} = \left(\frac{\gamma}{1 - \gamma} \right)^\gamma \left(\frac{c_f^M}{w_m} \right)^\gamma M,$$

where the second equality uses $M^\gamma \cdot M^{1-\gamma} = M$. Solving for M and substituting back into the labor expression yields the conditional input demands

$$\ell_f^*(q_f) = q_f \left(\frac{\gamma}{1 - \gamma} \frac{c_f^M}{w_m} \right)^{1-\gamma}, \quad M_f^*(q_f) = q_f \left(\frac{1 - \gamma}{\gamma} \frac{w_m}{c_f^M} \right)^\gamma.$$

Plugging these into the cost objective $w_m \ell_f^* + c_f^M M_f^*$ gives total cost

$$C(q_f) = q_f w_m^\gamma (c_f^M)^{1-\gamma} \left[\left(\frac{\gamma}{1-\gamma} \right)^{1-\gamma} + \left(\frac{1-\gamma}{\gamma} \right)^\gamma \right] = q_f \cdot c_f.$$

The bracketed sum collapses to a constant: factoring out the common term $\gamma^{-\gamma}(1-\gamma)^{-(1-\gamma)}$, the first summand becomes γ and the second becomes $(1-\gamma)$, so the bracket reduces to $\gamma^{-\gamma}(1-\gamma)^{-(1-\gamma)}(\gamma + (1-\gamma)) = \gamma^{-\gamma}(1-\gamma)^{-(1-\gamma)} \equiv B$. The unit cost is therefore given by

$$c_f(\mathcal{S}_f) = B w_m^\gamma [c_f^M(\mathcal{S}_f)]^{1-\gamma}, \quad B \equiv \gamma^{-\gamma}(1-\gamma)^{-(1-\gamma)}, \quad (31)$$

which corresponds to equation (9) in the main text. Note that the constant B has no economic content, as it arises mechanically from the Cobb-Douglas algebra.

A.3 CTC Compliant Unit Cost

Under CTC compliance, a measure- α_p subset of the firm's continuum of intermediate input varieties $j \in [0, 1]$ must originate in North America. Given that input varieties are drawn iid Fréchet across countries, the law of large numbers over the continuum implies that the cost contribution of any measure- α_p subset depends only on its measure, and not on which specific varieties it contains. The restricted and unrestricted subsets can therefore be treated as two sub-continua, each with its own effective Fréchet scale parameter. For the restricted NA-only subset, the effective scale is $\sum_{s \in NA \cap \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta}$, whereas for the unrestricted subset it is $\sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta}$. The Cobb-Douglas aggregation of the two subsets in turn gives the compliant intermediate cost index:

$$c_{f,p}^{M,CTC}(\mathcal{S}_f) = \Gamma\left(1 - \frac{1}{\theta}\right) \left(\sum_{s \in NA \cap \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-\alpha_p/\theta} \left(\sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-(1-\alpha_p)/\theta}.$$

Substituting into the Cobb-Douglas production structure $c_f = B w_m^\gamma [c_f^M]^{1-\gamma}$ in turn yields the compliant unit cost:

$$\begin{aligned} c_{f,p}^{CTC}(\mathcal{S}_f) &= B w_m^\gamma \left[\Gamma\left(1 - \frac{1}{\theta}\right) \right]^{1-\gamma} \\ &\times \left(\sum_{s \in NA \cap \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-\alpha_p(1-\gamma)/\theta} \left(\sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} \right)^{-(1-\alpha_p)(1-\gamma)/\theta}. \end{aligned} \quad (32)$$

The intermediate-share exponents are scaled by $(1 - \gamma)$ considering that only the $(1 - \gamma)$ share of total cost is paid to intermediates under Cobb-Douglas production. Writing $\mu_f^{NA,M} \equiv \sum_{s \in NA \cap \mathcal{S}_f} \mu_{f,s}(\mathcal{S}_f)$ for the North American share of the unrestricted intermediate bundle, (32) collapses to the compact form used in the solver,

$$c_{f,p}^{CTC}(\mathcal{S}_f) = c_f(\mathcal{S}_f) (\mu_f^{NA,M})^{-\alpha_p(1-\gamma)/\theta},$$

since $\sum_{s \in NA \cap \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta} = \mu_f^{NA,M} \sum_{s \in \mathcal{S}_f} T_s a_{f,s}^\theta (w_s \tau_s)^{-\theta}$ by (7).

A.4 RVC Compliant Unit Cost

Under an RVC threshold the firm chooses the restricted measure itself. Restricting a measure α of the input continuum to North American sources leaves the unrestricted sub-continuum with NA share $\mu_f^{NA,M}$, so the NA share of the whole intermediate bundle becomes $\alpha + (1 - \alpha)\mu_f^{NA,M}$ and the content share of production becomes $\gamma + (1 - \gamma)[\alpha + (1 - \alpha)\mu_f^{NA,M}]$. Setting this equal to μ_p^{RVC} and solving for α gives

$$\alpha_{f,p}^*(\mathcal{S}_f) = \frac{\mu_p^{RVC} - \gamma - (1 - \gamma)\mu_f^{NA,M}}{(1 - \gamma)(1 - \mu_f^{NA,M})} = \frac{\mu_p^{RVC} - \mu_f^{NA}(\mathcal{S}_f)}{(1 - \gamma)(1 - \mu_f^{NA,M})},$$

which is clamped to $[0, 1]$, so that a firm that is already above the threshold restricts nothing. Given that the restricted cost is increasing in the restricted measure, α^* is the cheapest compliant restriction, and the RVC compliant cost is $c_{f,p}^{RVC}(\mathcal{S}_f) = c_f(\mathcal{S}_f) (\mu_f^{NA,M})^{-\alpha_{f,p}^*(1-\gamma)/\theta}$. A firm with no North American source has $\mu_f^{NA,M} = 0$ and therefore cannot comply whenever $\alpha^* > 0$.

A.5 Mexican CES Price Index

The CES price index P_M in equation (18) follows from the standard demand-side dual to the destination- d aggregator in (3). Starting from

$$\max_{\{q_d(\omega)\}} \left(\int_{\omega \in \Omega_d} (b_d(\omega) q_d(\omega))^{(\sigma-1)/\sigma} d\omega \right)^{\sigma/(\sigma-1)} \quad \text{s.t.} \quad \int p_d(\omega) q_d(\omega) d\omega = E_d,$$

the first-order condition at ω gives

$$b_d(\omega)^{(\sigma-1)/\sigma} q_d(\omega)^{-1/\sigma} = \lambda p_d(\omega),$$

which can be inverted for the demand function:

$$q_d(\omega) = \lambda^{-\sigma} b_d(\omega)^{\sigma-1} p_d(\omega)^{-\sigma}.$$

Plugging this back into the budget constraint pins down the multiplier:

$$E_d = \lambda^{-\sigma} \int b_d(\omega)^{\sigma-1} p_d(\omega)^{1-\sigma} d\omega \Rightarrow \lambda^{-\sigma} = \frac{E_d}{\int (p_d(\omega)/b_d(\omega))^{1-\sigma} d\omega}.$$

Defining the CES price index P_d implicitly through $\lambda^{-\sigma} = E_d P_d^{\sigma-1}$ in order to recover the demand expression (4), $q_d(\omega) = E_d P_d^{\sigma-1} p_d(\omega)^{-\sigma} b_d(\omega)^{\sigma-1}$, the price index satisfies

$$P_d^{1-\sigma} = \int_{\omega \in \Omega_d} \left(\frac{p_d(\omega)}{b_d(\omega)} \right)^{1-\sigma} d\omega. \quad (33)$$

In the model, varieties are firm-product pairs $\{(f, p)\}$ rather than a continuum. The set of active varieties sold at destination d in equilibrium is $\{(f, p) : f \in \mathcal{A}, d \in \mathcal{D}_f, p \in \mathcal{P}_f\}$, the demand shifter collapses to the firm-level value $b_{f,d}$ (shared across that firm's products at destination d), and the price is $p_{f,p,d}$. Replacing the integral in (33) with the corresponding sum and specializing to $d = MX$ yields equation (18):

$$(P_M)^{1-\sigma} = \int_{f \in \mathcal{A}: MX \in \mathcal{D}_f} \sum_{p \in \mathcal{P}_f} \left(\frac{p_{f,p,MX}}{b_{f,MX}} \right)^{1-\sigma} df.$$

Note that the $(1-\sigma)$ exponents on either side cannot be cancelled, given that the right-hand side is a sum of $(p/b)^{1-\sigma}$ terms. The equivalent rooted form $P_M = (\sum (p/b)^{1-\sigma})^{1/(1-\sigma)}$ is algebraically identical.

A.6 Manufacturing Labor Demand Decomposition

Total Mexican manufacturing labor compensation $w_m L_m$ collects three sources across the active set $\mathcal{A}$. Under Cobb-Douglas production, a share γ of each firm's variable cost $VC_f \equiv \frac{\sigma-1}{\sigma} \sum_{d \in \mathcal{D}_f} \sum_{p \in \mathcal{P}_f} R_{f,p,d}$ is paid directly to Mexican production labor, contributing the wage-bill component

$$E^{\text{prod}} = \int_{f \in \mathcal{A}} \gamma \cdot VC_f df.$$

A second share $(1-\gamma)\mu_{f,MX}$ of variable cost pays for Mexican-sourced intermediates. Treating upstream Mexican suppliers as having zero economic profit and labor as their sole factor, this intermediate spending equals an equivalent volume of Mexican labor compensation at

the upstream stage:

$$E^{\text{int}} = \int_{f \in \mathcal{A}} (1 - \gamma) \mu_{f, MX}(\mathcal{S}_f) \cdot VC_f df.$$

Third, each active firm pays sourcing fixed costs $w_m \sum_{s \in \mathcal{S}_f} \kappa_s$ and destination entry costs $w_m \sum_{d \in \mathcal{D}_f} \kappa_d$, both in units of Mexican labor by assumption, contributing

$$E^{\text{fixed}} = \int_{f \in \mathcal{A}} w_m \left(\sum_{s \in \mathcal{S}_f} \kappa_s + \sum_{d \in \mathcal{D}_f} \kappa_d \right) df.$$

In the executed calibration, the non-trading bins lie outside the firm block (they enter through the price-index term K and their domestic sales, Section 4), so the integral runs over the traders. If the bins were reinstated as firm rows, they would enter the same accounting with $\mathcal{S}_f = \{\text{MX}\}$, $\mathcal{D}_f \subseteq \{\text{MX}\}$ and $\mu_{f, MX} = 1$, with their variable bracket reducing to VC_f . Summing the three components gives total manufacturing labor compensation $w_m L_m$, and dividing by the pinned wage $w_m = A_o$ in turn yields the labor-demand equation (17) in the main text:

$$w_m L_m = \int_{f \in \mathcal{A}} \left[(\gamma + (1 - \gamma) \mu_{f, MX}) VC_f + w_m \sum_{s \in \mathcal{S}_f} \kappa_s + w_m \sum_{d \in \mathcal{D}_f} \kappa_d \right] df.$$

A.7 Solving the Domestic-Demand Loop

Under the Cobb-Douglas upper tier (2), Mexican manufacturing expenditure moves with Mexican income,

$$E_M = E_M^0 \frac{I}{I^0}, \quad I \equiv w_m \bar{L} + \Pi_M + \Pi_o, \quad (34)$$

where the superscript 0 marks master-equilibrium values. In this expression, Mexican manufacturing spending moves at the baseline propensity E_M^0/I^0 (0.87 in 2018 and 0.94 in 2023) rather than at χ . The level E_M^0 is the baseline domestic manufacturing sales of the model firms and the non-traders (Section 4). Anchoring it on these baseline sales rather than on χI^0 matters because domestic sales are gross of the intermediates that firms sell to each other. Each counterfactual iterates on (E_M, P_M) as follows. Given the pair, firms solve (14), their choices update Π_M through (20) and P_M through (18), and the loop stops when (34) and (18) hold at the firms' choices. The master equilibrium is a fixed point by construction.

B SCIAN6–HS6 Concordance and Match Quality

The CTC input-restriction parameter α_p is built from INEGI’s 2018 Mexican input–output (I-O) table at the six-digit SCIAN industry level, whereas the rules of origin and the model itself operate at HS6. In order to bridge the two, I use the [Pierce and Schott \(2012\)](#) NAICS6↔HS6 concordance, trade-weighted by MX→US and MX→CA export value, paired with a NAICS↔SCIAN crosswalk that drops to the deepest SCIAN level (6, 5, or 4 digits) at which INEGI reports. Pushing the SCIAN I-O flows through this concordance on both the output and the input side yields an HS6×HS6 input–output matrix, from which α_p is read off as the restricted-input share of each output HS6. Legal-list HS6 codes that lack direct I-O backing are imputed from the flow-weighted mean of their HS4 siblings, and then of their HS2 siblings, with a zero residual fallback. Codes of another HS revision take the value of the codes that they link to through the UN conversion tables, and a handful of codes whose rule restricts no inputs take zero from the rule itself.

The translation shows to be successful where trade actually concentrates. Table 15 reports match rates by treaty, both in counts and weighted by MX→US&CA export value. Under USMCA, α_p is computed directly from the I-O for 60% of HS6 codes but for 91–93% of trade value, given that unmatched codes form a low-trade long tail while vehicles, electronics and machinery map cleanly. The revision link accounts for 7% of codes and 7% of 2023 trade value, mostly codes that were renumbered at the HS2022 boundary. NAFTA’s direct-coding share is lower (49% of codes and 72% of trade value through [Conconi et al.’s](#) Annex 401 parse), with HS2-level imputation absorbing 13–21% of value. Zero fallbacks are negligible under either treaty (under 2% of codes and 0.2% of trade value), and the parameter file covers all manufacturing (HS2 16–97) MX→US&CA export value in both periods.

C CE–PEME Match Quality and Vintage Diagnostics

The trader panel is based on the census–customs match, in which CE manufacturing firms (collapsed to the firm identifier) are merged onto the PEME firm×HS6 customs panel. This appendix reports the match diagnostics that were produced inside the laboratory (Table 16), as well as two vintage-specific measurement findings that qualify the 2023 panel.

There are three findings worth pointing out. **Match rates and attrition.** Almost every customs trader in the manufacturing universe finds its census record in 2018 (97% of PEME firms and 99% of customs export value), and most of them do in 2023 (83% of firms and 86% of value). Seen from the census side, 78% and 66% of CE manufacturing exporters find a customs record, where the remainder are plausibly firms whose foreign revenue is recorded

Table 15: SCIAN6–HS6 match quality for the CTC input-restriction parameter α_p

Provenance of α_p	HS6 codes	% of HS6	% of modeled trade value	
			2018	2023
<i>Panel A. USMCA</i>				
Direct I-O HS6 match	3,032	59.7	93.1	91.0
HS4 match	875	17.2	0.2	0.3
HS2 match	497	9.8	4.7	0.5
HS revision link	348	6.9	0.4	6.6
Rule (no restricted inputs)	12	0.2	0.0	0.0
Yarn-Forward rail	277	5.5	1.5	1.5
Zero fallback	36	0.7	0.2	0.2
<i>Panel B. NAFTA</i>				
Direct I-O HS6 match	2,482	48.9	72.2	72.1
HS4 match	1,136	22.4	5.2	6.0
HS2 match	901	17.7	21.1	13.3
HS revision link	206	4.1	0.0	7.1
Yarn-Forward rail	258	5.1	1.3	1.3
Zero fallback	94	1.9	0.2	0.2

Notes: The 5,077 manufacturing HS6 (HS2 16–97) of the parameter file, which include codes of the HS2012, HS2017 and HS2022 revisions. Value columns weighted by 2018/2023 MX→US&CA exports (UN Comtrade, Mexico as reporter).

Table 16: CE–PEME match diagnostics by vintage

	CE 2019 $\times$ PEME 2018	CE 2024 $\times$ PEME 2023
PEME manufacturing firms	6,854	6,344
matched to a CE firm	6,676 (97.4%)	5,235 (82.5%)
PEME export value matched	99.0%	85.6%
PEME import value matched	98.9%	83.1%
CE manufacturing exporters ($rev_f > 0$)	5,224	4,781
matched to a PEME firm	4,065 (77.8%)	3,175 (66.4%)
CE foreign revenue matched	97.0%	86.2%
CE manufacturing importers ($M_f > 0$)	3,202	15
matched to a PEME firm	2,360 (73.7%)	8
Trader panel (firms)	5,719	5,133
<i>PEME firm-level trade (USD)</i>		
Exports: mean / median	\$60.4m / \$3.3m	\$84.8m / \$5.8m
Imports: mean / median	\$48.8m / \$3.8m	\$62.6m / \$5.3m
HS6 products per firm: mean / median	27.1 / 9	28.3 / 10
<i>CE foreign revenue / PEME exports, matched firms</i>		
p25	0.0085	0.0104
p50	0.0189	0.0178
p75	0.0252	0.0279
mean	39.8	1.09

Notes: Direct firm-identifier join within the manufacturing universe. Ratio: CE foreign revenue (thousand pesos) over PEME exports (USD); exchange-rate parity ≈ 0.019 (2018), 0.018 (2023).

by the census but whose shipments clear customs under a different legal entity. Keeping the matched firms with positive revenue, materials and wage bill in the census that trade on at least one side, aggregating establishments to the firm level, and dropping the exporters whose products are all Yarn-Forward takes the matched set from 6,676 and 5,235 firms to the 5,719 and 5,133 traders that are used in the paper. The attrition is uneven, 957 firms in 2018 against 102 in 2023, and comes mostly through the positive-wage-bill condition. Before the 2021 subcontracting reform, many manufacturers employed their workers through a separate staffing company and reported no payroll of their own, and the reform moved those workers onto the manufacturers' payrolls, which also raises $\bar{L}$, the paid wage bill and the per-wave γ_s between the censuses. The 2023 match is weaker on both sides than the 2018 one, and the unmatched customs value (14% of exports) is one candidate component of the structural residual of the decomposition.

Census–customs value discrepancies. At the median, the CE-to-PEME export ratio lies at exchange-rate parity (0.019 against a parity value of 0.0192 in 2018, and 0.018 against 0.0178 in 2023), which simultaneously validates the match and pins down the units (CE monetary variables are in thousands of pesos and PEME in USD). The interquartile range is


wide, from about half of parity at the 25th percentile to 1.3 to 1.6 times parity at the 75th, and the 2018 mean is dragged far above parity by a heavy right tail of firms that report census foreign revenue far in excess of their customs exports. These discrepancies plausibly mix consignment and maquila revenue-recognition conventions (census revenue records the processing fee while customs records the gross shipment value, which compresses the ratio for maquila-type firms), intra-year timing, and identifier-level aggregation differences. For the model, composition $(\mu_{f,s}, \omega_{f,d})$ comes from PEME and levels come from CE, so these wedges matter mainly for the levels-based stitch. Documenting their firm-level structure is part of the next laboratory request.

The CE 2024 imported-materials collapse. In the CE 2019 extraction, 3,202 manufacturing firms report positive imported materials ($M_f > 0$), whereas in the CE 2024 extraction only 15 do. The imported-materials item is thus effectively absent from the 2024 extraction, which is most plausibly an extraction or questionnaire-coverage issue to be confirmed with INEGI, rather than an economic fact. For a 2024 importer, the census own-materials total then already contains the imports, meaning that adding the customs imports on top of it would count them twice. For the non-maquila customs importers of the 2023 panel whose census materials show no imported part (3,131 firms), the customs imports are therefore subtracted from the census total in order to obtain domestic materials, floored at zero. At the median of these firms, customs imports are a third of census own materials, and for 462 of them (15%) they exceed it, so their domestic materials are set to zero. Maquila firms are left unchanged, given that their customs imports are consigned inputs that the census does not record as own materials. In all, 1,677 firms in 2018 and 1,698 in 2023 (29 and 33%) have zero domestic materials in the panel. Total materials M_t , which feeds γ_s and the markups behind σ , is the census total in both vintages and is therefore not affected. Any moment built on the census domestic/imported split remains unavailable in 2024, and the asymmetry between vintages enters the decomposition both through ψ_{MX} , which moves in Δ_ψ , and through the re-inverted capabilities in Δ_S .

D Discrete-Choice Fit by Block

Table 17 reports the sourcing- and destination-participation fit block by block, together with the share of firms whose entire sourcing set the model reproduces. Data and model shares coincide on the eight targeted margins by construction of the inversion, so the informative columns are the “correct %” ones, which score each firm’s in/out status cell by cell. On the sourcing side, the model gets 72 to 77% of firm-block cells right, which is above the 54 to 70% that a random assignment with the same shares would achieve in every block but

Figure 9: Sector cost shares γ_s across census vintages



Notes: γ_s by NAICS-3 over all CE manufacturing firms. Open circles: CE 2019; filled circles: CE 2024; diamonds: the pooled value used in both vintages.

US–Canada, where the high participation rate puts that benchmark at 76 to 77%. On the destination side, it gets 71 to 82% right, against 50 to 81% for the random benchmark.

E Estimated Firm Primitives: 2018 Vintage

This appendix reports for the 2018 vintage (5,719 traders) the objects that Section 5.2 reports for 2023.

E.1 Cross-Firm Distribution

Table 18 is the 2018 counterpart to Table 9. Capabilities again have standard deviations of 1.5 to 1.6 log units, with Mexico having the highest mean, the AX, CH and EU blocks near a common lower mean, and RoW below them. Between the two vintages, the mean rises for every capability coordinate, by 0.05 to 0.12 log units, the most for Mexico and the least for Europe. On the demand side, the mean imputed log b is 8.1 at US&CA, 8.9 at Mexico and 4.6 at RoW, against 8.7, 9.9 and 6.5 among the firms observed in each market.

Table 17: Discrete-choice fit: sourcing and destination participation

Margin	2018 vintage			2023 vintage		
	Data %	Model %	Correct %	Data %	Model %	Correct %
<i>Sourcing: share of firms sourcing from</i>						
US–Canada	86.4	86.4	74.5	86.7	86.8	75.2
Asia excl. China	73.8	73.9	72.7	76.2	76.2	73.3
China	78.4	78.4	74.0	81.5	81.5	77.1
Europe	75.9	75.9	74.7	77.0	77.0	75.6
Rest of World	64.6	64.5	72.8	65.0	65.0	72.5
<i>Destination: share of firms selling to</i>						
US&CA	78.0	78.0	73.1	79.9	79.9	75.7
Mexico	86.1	86.1	78.4	89.2	89.2	81.7
Rest of World	52.7	52.7	74.3	55.5	55.6	71.4
Exact sourcing set (% of firms)		37.3			38.5	

Notes: “Correct %”: share of firms whose in/out status matches the data. Data and model shares coincide by construction of the inversion. MX-source participation is 100% in the model by construction ($\kappa_s^{MX} \equiv 0$), against 71% (2018) and 67% (2023) of firms with positive domestic materials in the data; omitted.

Table 18: Estimated firm primitives: cross-firm distribution (2018 vintage)

Coordinate	mean	sd	observed
<i>Capabilities</i> $\log a_{f,s}$ ($USCA \equiv 0$)			
Mexico	0.259	1.518	0.58
Asia excl. China	−0.656	1.515	0.70
China	−0.595	1.507	0.73
Europe	−0.496	1.496	0.71
Rest of World	−1.006	1.551	0.61
<i>Demand shifters</i> $\log b_{f,d}$			
US&CA	8.567	2.318	0.78
Mexico	9.751	2.163	0.86
Rest of World	5.574	2.304	0.53

Notes: 5,719 traders. Construction as in Table 9.

E.2 Ownership Groups

Tables 19a and 19b are the 2018 counterparts to Tables 10a and 10b, and the home-region pattern is the same. Mexican-owned firms peak in Mexico (+0.71), AX-owned firms in Asia excluding China (+1.65), Chinese-owned firms in China (+1.49) and EU-owned firms in Europe (+0.87), whereas US–CA-owned firms lie below zero on every other block, and the RoW-owned group again has no peak in its own block. On the demand side, the highest mean shifter in each market again belongs to the group from that region: Mexican-owned firms in Mexico (10.3), US–CA-owned firms in the US&CA market (9.7) and the remaining foreign owners in the Rest of World (6.4).

Table 19a: Estimated sourcing capabilities by ownership group (2018 vintage)

	Ownership group					
	Mexico	U.S. & CA	Asia excl. China	China	Europe	Rest of World
Firms	3,241	1,546	244	20	488	180
<i>Mean capability log $a_{f,s}$, by source block ($USCA \equiv 0$)</i>						
Mexico	0.71	-0.83	0.85	-0.03	0.40	0.39
US-Canada	0.00	0.00	0.00	0.00	0.00	0.00
Asia excl. China	-0.85	-0.73	1.65	-0.03	-0.39	-0.53
China	-0.64	-0.72	0.05	1.49	-0.29	-0.65
Europe	-0.41	-1.14	-0.54	0.04	0.87	-0.19
Rest of World	-0.93	-1.48	-1.14	-0.66	-0.21	-0.32

Notes: Construction and groups as in Table 10a. Bold: each group's home block.

Table 19b: Estimated demand shifters by ownership group (2018 vintage)

	Ownership group		
	Mexico	U.S. & CA	Other foreign
Firms	3,241	1,546	932
<i>Mean demand shifter log $b_{f,d}$, by destination</i>			
Mexico	10.29	8.81	9.44
U.S. & CA	8.06	9.67	8.51
Rest of World	5.47	5.27	6.45

Notes: Construction and groups as in Table 10b. Bold: each group's home market.

E.3 Variation Explained by Observables

Table 20 is the 2018 counterpart to Table 11, and comes from the same regressions (27). The fully interacted observables explain at most 27% of the cross-firm variation in any coordinate (US&CA demand, 0.265), and ownership is again the stronger observable for the capability coordinates other than China, whereas sector is the stronger one for the three demand coordinates.

E.4 Correlation Structure

Figure 10 reports the 2018-vintage counterpart to Figure 4 in the main text. The correlation structure is stable across vintages, as the East-Asian bloc effect (the AX-CH capability correlation) is 0.53 in 2018 and 0.54 in 2023, and the US&CA-RoW destination correlation is 0.25 in 2018 and 0.29 in 2023. This stability supports the interpretation of these correlations as structural features of the trader panel rather than as artifacts of one vintage, although part of it is built in, given that the two vintages share one pooled imputation surface.

Table 20: Share of the cross-firm variation in $\log a_{f,s}$ and $\log b_{f,d}$ explained by ownership and sector (2018 vintage)

Latent	Coordinate	R^2 from fixed effects for			
		Own	Sector	Own + Sector	Own $\times$ Sector
$\log a$	MX	0.198	0.090	0.241	0.259
	AX	0.112	0.101	0.193	0.217
	CH	0.020	0.099	0.123	0.144
	EU	0.124	0.046	0.160	0.181
	RoW	0.055	0.047	0.097	0.120
$\log b$	US&CA	0.089	0.210	0.257	0.265
	MX	0.089	0.102	0.168	0.179
	RoW	0.029	0.098	0.134	0.142

Notes: 5,719 traders. Construction as in Table 11.

Figure 10: Correlation structure of estimated firm primitives (2018 vintage)

	MX	AX	CH	EU	RoW
MX	1.00				
AX	.24	1.00			
CH	.25	.53	1.00		
EU	.46	.34	.33	1.00	
RoW	.35	.29	.30	.48	1.00

(a) Sourcing capabilities $a_{f,s}$

	US&CA	MX	RoW
US&CA	1.00		
MX	.03	1.00	
RoW	.25	.22	1.00

(b) Demand shifters $b_{f,d}$

Notes: 5,719 traders, 2018 vintage. Panels show the lower triangle of the correlation matrix (upper triangle omitted by symmetry). Panel (a) omits USCA ($\log a_{f,USCA} \equiv 0$).