

The Laffer Curve for Rules of Origin*

Keith Head[†]

Thierry Mayer[‡]

Marc Melitz[§]

October 9, 2022

Abstract

We analyze how heterogeneous firms in a regional trade area (RTA) respond to rules of origin (RoO). Firms can source a continuum of inputs from both within and outside the RTA, and choose whether to comply with the RoO or pay a tariff penalty. We show how a Laffer curve for RoOs arises naturally in this setting: stricter content requirements initially expand regional part sourcing, but contract it when set at levels above a threshold. The parameters of the model are fit to data on regional part cost shares for all autos sold in North America. The calibrated model quantifies the impact of stricter RoOs imposed by the 2020 revision to NAFTA (USMCA). The stricter content requirement (62.5% to 75%) would raise employment by only 1.2%, while increasing auto prices assembled in the region by 0.3%. The higher requirement initially proposed by U.S. negotiators (85%) would lead to both higher prices and *lower* employment.

1 Introduction

The increasingly global structure of supply chains draws attention to the rules of origin (RoO) that govern whether firms can take advantage of trade liberalization. Regional trade agreements (RTAs) impose requirements on goods in order to qualify for the preferential tariff treatment. With localized supply structures, compliance is straightforward, but the parts making up modern goods often come from origins spanning the world. As RTAs proliferated at the same time as value chains globalized, RoOs are increasingly used as a protectionist device to shelter domestic suppliers. Historically there have been numerous shifts, both up and down, in the strictness of the rules. For example, the 1965 US-Canada Auto Pact used a 50% content rule. This was raised to 62.5% when the North American Free Trade Agreement (NAFTA) was enacted in 1994. The Transpacific Partnership (TPP) proposed to lower the rule to just 45% in 2016, only to have a new President

*We acknowledge support from Global Affairs Canada and the Social Studies and Humanities Research Council of Canada (435-2019-0381). We thank Anhua Chen and George Hu for superb research assistance, Emily Blanchard and Felix Tintelnot for helpful suggestions, and many seminar participants.

[†]Sauder School of Business, University of British Columbia, CEPR, keith.head@sauder.ubc.ca

[‡]Sciences Po, CEPR, and CEPR, thierry.mayer@sciencespo.fr

[§]Harvard University and NBER, mmelitz@harvard.edu

pull the US out of the TPP and propose instead a rise in the requirement to 85%. Canada and Mexico balked at such a high rate; the US negotiators finally settled on an increase to 75%, bolstered with additional binding requirements. In this capsule history we see no sign of a consensus on the ideal restrictiveness of rules of origin.

This paper points to the potential for stricter rules to be counterproductive—even if their goal is purely protectionist. Consider a simple, but representative, example. An assembler in Canada obtains engines within North America in order to comply with the old rules of origins but sources transmissions from Japan. Compliance with USMCA rules would require finding a North American source for the transmission. Suppose the firm deems that option to be too costly. Once car sales within North America are no longer compliant with the more restrictive rules, the producer no longer has any incentive to locate engine production within USMCA. They can then select their preferred location outside North America. The scope for perverse effects expands when the firm has the option to relocate assembly as well, which it may well do, given that non-compliance with the rule of origin prevents the use of tariff preferences. Once car assembly moves out of the region, the incentive to source any part from within the region declines since the imported car must pay the full tariff regardless of the source of its parts. This associated decrease in regional part demand is magnified by trade costs for parts—because the relocation of assembly then increases the delivered cost for regional parts. Those incentives for the relocation of part production and assembly will hold for vehicles intended for USMCA consumers—but will be particularly salient for vehicles currently assembled in the USMCA area for export to Asia or Europe.

We develop a theoretical model that highlights these competing incentives of tighter rules of origins on the relocation of production into and out-of the area of a regional trade agreement (RTA). The model shows how the negative impact on relocation out-of the RTA increasingly dominates as the rules of origins are tightened. We show that with a continuum of parts, the mechanisms described above naturally lead to what we call a rule of origin (RoO) Laffer curve. Initial increases in local content requirements shift parts production inside the region. A sufficiently strict rule depresses regional parts production. With the option to shift assembly, the final share of parts produced within the region can be lower than it would have been with no RoO at all. As with the Laffer curve for taxes, governments will generally want to avoid being on the wrong side of the Laffer curve's peak. These concerns motivate empirical work to determine at what point the relocation effects of stricter RoOs become dominant.

One rationale offered for rules of origin is to allow different tariff rates within the free trade area by preventing outside imports from circumventing high-tariff destinations through low-tariff entry points. This argument is unlikely to explain the RoOs we see in practice for two reasons. First, only a very low RoO would be required to prevent this kind of high-tariff circumvention.¹ Second, at least in North America, the pressure for higher RoOs mainly comes from the *low* MFN country, the US, who would stand to gain from facilitating evasion of the higher MFN tariffs in Canada and Mexico.

Two prior papers address our question of interest, which is whether even as protec-

¹Felbermayr et al. (2019) show that for 86% of bilateral product pairs, tariff circumvention is unprofitable due to small differences in MFN tariffs and non-negligible transportation costs.

tionist devices, strict rules of origin could fail to achieve their goals. The earliest work we know of in that vein is Grossman (1981). Grossman’s Proposition 3 states that small increases in local content requirements have ambiguous effects on industry value added, defined as the sum of value added in components and in final goods. Industry output is more likely to fall the less sensitive intermediate production is to its price and the more sensitive is final good production to the intermediate input price. In our setup we consider potentially large increases in content requirements. A key result in our setup is that value added in the components sector itself can decline. While our main result does not work via declines in final goods production, we show quantitatively how they exacerbate the negative effect on components.

Ju and Krishna (2002) continue the investigation of ambiguous effects of stricter rules launched by Grossman (1981). The key novelty of their setup is that *ex ante* identical firms potentially choose different equilibrium responses to the input requirements. There are regions in their model in which some firms comply, but others opt to just pay the MFN tariff. This can lead to non-monotonic impacts of rule of origin. In contrast to their model, the firms in our framework are fundamentally heterogeneous, with some having pre-determined tendencies towards offshore sourcing of parts. The advantages of assuming continuous firm type heterogeneity are 1) that it smoothens out the response to policy changes, and 2) realism (the data on content shares in the auto industry show wide variation). Ornelas and Turner (2022) adopt a different approach emphasizing relationship-specific investment under incomplete contracts. This gives rise to a role for stricter RoOs to solve the associated under-investment problem.

Rules of origins generate a very specific type of friction on the import of intermediate goods. A key mechanism in our model is that limiting the use of imported intermediates will tend to raise costs. This assumption finds support in a well-established empirical literature examining the broader impact of access to imported intermediate inputs on firm performance. These studies all find that improved access to intermediate inputs lead to sizable cost and productivity gains for the affected firms.²

Some recent studies have focused more specifically on the impact of frictions that arise specifically from feature of rules of origin. There too, the impact on firm performance is substantial. Demidova et al. (2012) find that Bangladeshi firms that are less affected by stringent rules of origins are 21% more productive than firms that are more affected by those rules. Sytsma (2022) estimates that rules of origin effectively cut the preferential margin faced by Bangladeshi apparel exporters by three-fourths. When the EU eased those rules, the number of exporters to the EU increased. Bombarda and Gamberoni (2013) similarly finds that when the EU allowed exporters to cumulate content across FTA partner countries, this stimulated both the extensive and intensive margins of exporting to and from the EU. Anson et al. (2005) find that the compliance costs associated with NAFTA rules of origins on Mexican firms amount to the equivalent of a 6% *ad valorem* tariff on the affected intermediate goods. That cost is high enough to negate the tariff advantage for many Mexican exporters to the U.S. and Canada relative to exporters from countries outside NAFTA. Conconi et al. (2018) empirically demonstrate the “cascade

²See Amiti and Konings (2007), Bombardini et al. (2021), Goldberg et al. (2010), Gopinath and Neiman (2014), and Halpern et al. (2015).

effect” whereby rules of origin shift protection from final goods to intermediate inputs. In particular, they find that imports of intermediates goods from third countries decline relative to NAFTA partners.

Methodologically, our model owes its greatest debt to the continuum model of sourcing introduced by Eaton and Kortum (2002), as well as the EK-within-the-firm models of Tintelnot (2017) and Antras et al. (2017). The model posits that final products are made from a continuum of parts. This turns out to be a very helpful assumption for thinking about the impact of rules of origin. We use it to formalize the intuition that content requirements could be set so high as to not only raise consumer costs, but actually hurt the domestic producers. The continuum model highlights the key mechanisms for thinking about how value chains adjust in response to stricter local content requirements. With a single part, we would not obtain our main results.

This paper is organized as follows. We begin with a parsimonious model that illustrates the basic trade-off associated with rules of origin. The model predicts a hump-shaped relationship between regional content requirements (RCR) and the realized regional content shares (RCS) that we refer to as the RoO Laffer curve. The base model takes assembly locations as fixed. The next section allows for assembly relocation. We first consider the option of relocating assembly within the region. We show how this gives trade negotiators in the largest country an incentive to push for stricter content requirements as this then induces relocation of assembly to that country (and away from smaller countries in the region). This fits the anecdotal evidence for the new USMCA, as the US negotiators pushed for substantially higher requirements than those desired by the Canadian and Mexican negotiators. Section 3 considers the option of relocating assembly outside the region and introduces differences in the trade costs to deliver parts to the two assembly locations. We show how such trade costs along with the assembly relocation choice amplify the negative impact of higher content requirements for regional part sourcing. High content requirements can then lead to a lower regional part share than in the absence of any content rules, as firms relocate assembly outside the region.

Section 4 then uses carline-level data on NAFTA cost shares to estimate the core parameters of our model when assembly location is fixed in the short-run. To fit the data more closely, we extend the basic model of section 2 to incorporate additional sources of heterogeneity. The calibrated model yields a Laffer curve disciplined by the data for the North American vehicle industry. The striking result from this calibration is that the increase in the RCR brought in by the USMCA appears to take North American content close to its maximum value. This calibration omits the negative impact on vehicle production—and derived demand for parts—in North America caused by increases in costs generated by stricter RoOs. Adding more structure for competition, demand, and assembly employment, Section 5 develops an exact hat algebra method in the tradition of Dekle et al. (2007) to predict changes in aggregates such as the price index and employment. The final exercise of the paper uses this method to plot the Laffer curve for employment in North America. We find that new USMCA’s rise in content requirements brings North America right to the top of the Laffer curve, so that further increases in content requirements would induce lower employment. The increase in employment induced by the USMCA content requirements is small: just 1.2%. And those employment gains come at the cost of a 0.3% higher price index for autos assembled in North America.

2 Sourcing and RoOs with a continuum of parts

Rules of origins generate competing incentives for the location of both parts and assembly within a regional trade area (RTA). Stricter rules intended to relocate the production of parts inside the RTA could potentially induce an opposite relocation effect away from the RTA. We develop a simple model focusing on the sourcing decision for parts in order to demonstrate how these opposing effects arise naturally in this setting and how the negative effects are likely to dominate when RoOs are tightened beyond a threshold. In order to focus on those opposing forces for part sourcing, we initially abstract from the associated assembly location decision. We then show how this additional choice interacts with the part sourcing decision and compliance with the RoO.

2.1 Model Setup

A firm uses a unit continuum of parts that can be sourced either domestically (within the RTA) or from a Foreign source.³ The cost of each part is drawn from a Weibull distribution with parameter $\theta \geq 1$. We normalize the mean cost of domestic parts (over the unit continuum) to one. The mean cost of the foreign sourced parts is $\delta > 0$.⁴ This parameter varies across firms. Firms with $\delta > 1$ have a lower domestic production cost for parts. For now, we ignore assembly costs and the assembly location choice, and focus on the compliance choice to a RoO.

2.1.1 Free Trade (No Policy Restrictions)

If a firm δ faces no restriction on part sourcing, it will source an unrestricted share of domestic parts

$$\chi_U(\delta) = (1 + \delta^{-\theta})^{-1}. \quad (1)$$

Given the Weibull cost draws, this sourcing decision equalizes the average cost of domestic-sourced parts with the average cost of foreign-sourced parts. This average cost is equal to $C_U(\delta) = \chi_U(\delta)^{1/\theta}$ and also captures the total cost of parts (given the unit continuum). Both $\chi_U(\delta)$ and $C_U(\delta)$ are increasing in δ : A firm with a bigger foreign cost advantage (lower δ) will choose a lower domestic part share and benefit from a lower total cost.

2.1.2 Rules of Origin

A RoO mandates that firms source a minimum share χ_R of their parts domestically (or alternatively a minimum domestic cost share), or face an MFN tariff rate on the final good exported within the RTA. We model this additional cost as an average tariff $\tau > 1$ incurred

³For now, we ignore any heterogeneity between countries in the RTA and only consider a single sourcing location for parts. Later on, we will introduce this heterogeneity when we endogenize the assembly location choice (as well as the part sourcing) between different countries in a RTA. Thus we refer to the regional part share as “domestic” —in contrast to the Foreign-sourced parts.

⁴In other words the distributions of the cost draws c sourced domestically ($i = D$) and in Foreign ($i = F$) are distributed $G_i(c) = 1 - \exp\left(-(c/\gamma_i)^\theta\right)$ with $\gamma_D \equiv \gamma$, $\gamma_F \equiv \delta\gamma$, $\gamma \equiv 1/\Gamma(1 + 1/\theta)$.

across all final good units produced. In the appendix, we show how this average tariff is scaled down from the MFN tariff rate as the share of within-RTA exports decreases when final good demand has a constant price elasticity (there is no scaling down of the MFN rate when all units are exported within the RTA, and final good demand is irrelevant). Since final good demand is only relevant for the determination of the average tariff τ , we do not introduce it explicitly.⁵ If a firm chooses to comply with the RoO and avoid the tariff, it sources progressively more expensive parts domestically (relative to foreign-sourced) until the minimum threshold is met. Those sourcing choices are identical to the ones a firm would make if a tariff $\rho > 1$ were imposed on foreign parts (with the tariff revenue subsequently rebated back to the firm). A tariff rate of ρ would induce a firm to increase its domestic share above $\chi_U(\delta)$ to

$$\chi_R = [1 + (\rho\delta)^{-\theta}]^{-1}. \quad (2)$$

If this firm were paying the tariff cost, its hypothetical total part cost would rise from $C_U(\delta)$ to $\chi_R^{1/\theta}$, which represents both the average cost of domestically sourced parts, as well as the average cost of foreign sourced parts *inclusive* of the hypothetical tariff.⁶

If a binding RoO $\chi_R > \chi_U(\delta)$ were mandated instead of a tariff, then the firm would make the same sourcing decisions, but its cost would not include the tariff ρ yielding a total part cost:

$$\begin{aligned} C(\chi_R, \delta) &= \chi_R \cdot \chi_R^{1/\theta} + (1 - \chi_R) \cdot \frac{\chi_R^{1/\theta}}{\rho} \\ &= \chi_R^{\frac{1+\theta}{\theta}} + (1 - \chi_R) \frac{1+\theta}{\theta} \delta, \end{aligned} \quad (3)$$

using $\rho = [\chi_R / (1 - \chi_R)]^{1/\theta} \delta^{-1}$ from (2). The first term captures the cost of domestic parts (same cost as under the tariff ρ) while the second one captures the cost of foreign parts rebated by the tariff ρ . The cost share associated with this RoO χ_R is:

$$\begin{aligned} \lambda(\chi_R, \delta) &= \frac{\chi_R^{\frac{1+\theta}{\theta}}}{C(\chi_R, \delta)} \\ &= \left[1 + \left(\frac{1 - \chi_R}{\chi_R} \right)^{\frac{1+\theta}{\theta}} \delta \right]^{-1}. \end{aligned} \quad (4)$$

Note that $\lambda(\chi_R, \delta)$ is monotonic in χ_R so we can think of a RoO as being imposed based on the share of parts χ_R or alternatively its cost share λ_R . Most free trade agreements specify

⁵With constant price elasticity demand, we show how this scaling down for the average tariff depends only on exogenous demand parameters (market demands across regions) and trade costs (transport costs and MFN tariffs). Thus, the decomposition of this average τ between the export share and the MFN tariff is inconsequential for a given final good. However, this means that variation in τ across products is possible even when goods face the same MFN tariff when there is variation in demand across products. We discuss the implications of variation in τ later on.

⁶Recall that the Weibull draws induce sourcing decisions that equalize the average cost parts by sourcing location.

the regional content rule of origin as cost shares though in some cases specific parts are mandated to be sourced with the region. This makes the rule look more like a part share rule. In the case where a cost-share λ_R is mandated, the required part share χ_R is just given by the inverse:

$$\chi_R = \left[1 + \left(\frac{\lambda_R^{-1} - 1}{\delta} \right)^{\frac{\theta}{\theta+1}} \right]^{-1}. \quad (5)$$

For the ensuing analysis, we assume that the rule is specified in terms of a part share χ_R , though this can be the outcome of the inversion above based on a cost-share rule λ_R .

A binding RoO $\chi_R > \chi_U(\delta)$ engenders an increase in the firm's total part cost relative to its unrestricted (lower bound) cost $C_U(\delta)$ given by

$$\tilde{C}(\chi_R, \delta) = \begin{cases} C(\chi_R, \delta)/C_U(\delta) > 1 & \chi_R > \chi_U(\delta), \\ 1 & \chi_R \leq \chi_U(\delta). \end{cases} \quad (6)$$

$\tilde{C}(\chi_R, \delta)$ is strictly decreasing in δ when the RoO is binding: firms with a greater comparative advantage in Foreign parts (lower δ) face a higher cost penalty of compliance for a given RoO χ_R . Absent the RoO, those firms would have chosen a lower part share $\chi_U(\delta) < \chi_R$, and it is therefore more costly to increase that share to χ_R . And $\tilde{C}(\chi_R, \delta)$ is strictly increasing in χ_R when the RoO is binding: a higher RoO induces higher cost penalties for all firms. However, a RoO need not be binding, as some firms with high δ may have an unrestricted part share *above* the RoO: $\chi_U(\delta) \geq \chi_R$. This will be the case for firms with $\delta \geq \delta^\circ$ such that $\chi_U(\delta^\circ) = \chi_R$, implying $\delta^\circ = (\chi_R^{-1} - 1)^{-1/\theta}$. Those firms can stick with their unrestricted part share $\chi_U(\delta)$ and still comply with the RoO with no cost penalty: $\tilde{C}(\chi_R, \delta) = 1$, $\forall \delta \geq \delta^\circ$.

2.1.3 Comparison with Eaton-Kortum Model of Trade in Goods


A rule of origin χ_R has the same welfare impact—the cost increase $\tilde{C}(\chi_R, \delta)$ imposed on firms—as an equivalent tariff ρ (which raises the domestic part share to χ_R) would in a two-country version of the Eaton and Kortum (2002) model, where $\delta^{-\theta}$ represents the foreign country's technological (absolute) advantage. But those welfare gains are no longer iso-elastic in the domestic share χ_R (with elasticity $1/\theta$) due to the revenue generated by the tariff. Just like a tariff, a rule of origin χ_R generates a distortion with no direct cost—unlike an iceberg trade cost. At the unrestricted equilibrium the elasticity of welfare with respect to either the rule χ_R or the tariff ρ is zero. However, that welfare elasticity then increases monotonically with either χ_R or ρ . The welfare penalty from a rule of origin—just like a tariff—is therefore more convex than a real cost that induces the same sourcing decisions. In the limit when $\chi_R = 1$, all compliant firms are forced into autarky sourcing with a part cost $C(1, \delta) = 1$ (the average cost of domestically sourced parts). Their cost disadvantage $\tilde{C}(1, \delta) = C_U(\delta)^{-1} = \chi_U(\delta)^{-1/\theta}$ is equal to the full gains from trade (autarky to free trade) in the E-K model representation.⁷

⁷ $\lim_{\chi_R \rightarrow 1} \delta^\circ = +\infty$, so a RoO $\chi_R = 1$ is binding for all firms.

2.2 Compliance

The top panel of Figure 1 shows the cost penalty $\tilde{C}(\chi_R, \delta)$ for three firms with different δ s as a function of the RoO χ_R . In this and all subsequent figures, we set $\theta = 4$ and a tariff level $\tau = 1.1$. As previously mentioned, this cost penalty increases with χ_R and is higher for firms with higher foreign-cost advantages (lower δ). We also see when a RoO χ_R is low enough ($\chi_R \leq \chi_U(\delta) \iff \delta \geq \delta^o$) to be non-binding, eliminating the cost penalty: $\tilde{C}(\chi_R, \delta) = 1$. As also anticipated, a given rule of origin χ_R is more likely to be binding for the firms with higher foreign-cost advantages as their unrestricted part share $\chi_U(\delta)$ is lower.

Figure 1: Compliance Cost and Sourcing Decision for 3 Firms



The policy maker cannot force firms to comply with a RoO χ_R . A firm δ can choose

to be non-compliant with the rule and pay the average tariff τ . It will do so whenever $\tilde{C}(\chi_R, \delta) \geq \tau$, and then revert to its unconstrained part sourcing with domestic share $\chi_U(\delta)$ and cost $C_U(\delta) = \chi_U(\delta)^{1/\theta}$. This leads to a cutoff rule for compliance: Only firms with $\delta > \delta^*$ such that $\tilde{C}(\chi_R, \delta^*) = \tau$ will comply with the RoO χ_R given a tariff punishment τ . δ^* increases monotonically with the rule of origin χ_R : A tougher RoO leads more firms to choose non-compliance. However, even in the autarky sourcing limit when $\chi_R = 1$, δ^* is finite: $\lim_{\chi_R \rightarrow 1} \delta^* = (\tau^\theta - 1)^{-1/\theta} \equiv \bar{\delta}$. Firms with δ above this threshold will comply with any RoO as their autarky cost disadvantage $\tilde{C}(1, \delta)$ is bounded and below the tariff cost τ . This is the case for the firm with $\delta = 1.25$ in the figure.


And conversely more firms with $\delta > \delta^*$ comply as the RoO becomes progressively looser and δ^* decreases. Some of those firms with $\delta \geq \delta^\circ > \delta^*$ do not face any compliance penalty, $\tilde{C}(\chi_R, \delta) = 1$, because their unrestricted part share $\chi_U(\delta)$ is above the RoO χ_R . Those firms comply with the RoO, but choose their unrestricted part share $\chi_U(\delta)$, whereas the remaining compliant firms face a cost penalty $\tilde{C}(\chi_R, \delta) > 1$ and set their part share at the level of the RoO χ_R . We label the latter firms compliant-constrained; and the former compliant-unconstrained. To summarize, only the compliant-constrained firms choose a part share at the level of the RoO χ_R . The remaining firms are either compliant-unconstrained or non-compliant, and they choose their unrestricted part share $\chi_U(\delta)$. Letting $\chi(\chi_R, \tau, \delta)$ denote the chosen part share, then:

$$\chi(\chi_R, \tau, \delta) = \begin{cases} \chi_R & \delta^* < \delta \leq \delta^\circ, \\ \chi_U(\delta) & \text{otherwise.} \end{cases} \quad (7)$$

The bottom panel of Figure 1 shows this chosen part share as a function of the RoO χ_R for the same three firms. When the RoO is low enough, all three firms are compliant-unconstrained and choose their respective part share $\chi_U(\delta)$. As the RoO increases, it starts binding for those firms and induces a cost penalty above 1. The firms then become compliant-constrained and set a part share at the level of the RoO χ_R . As the RoO increases even further, the cost penalty rises above the average tariff τ for the two firms with the lower δ s, and they then stop complying (non-compliant) and revert to their unconstrained part share $\chi_U(\delta)$.

Figure 2 plots the cost penalty $\tilde{C}(\chi_R, \delta)$ (top panel) and chosen part share $\chi(\chi_R, \tau, \delta)$ (bottom panel, in yellow) against the firm's δ for a given χ_R (equal to 0.7 in the figure). The bottom panel also adds the unrestricted part share $\chi_U(\delta)$ (in blue). This highlights the determination of the δ° cutoff at $\chi_U(\delta^\circ) = \chi_R$. The determination of the δ^* cutoff at $\tilde{C}(\chi_R, \delta^*) = \tau$ is shown in the top panel. As we previously mentioned, the cost penalty $\tilde{C}(\chi_R, \delta)$ is strictly decreasing in δ whenever the RoO χ_R is binding for the compliant-constrained firms with $\delta < \delta^\circ$, and then flat at one for the compliant-unconstrained firms with $\delta \geq \delta^\circ$. The bottom panel highlights how the compliant-constrained firms increase their domestic part share to satisfy the RoO χ_R and thus deviate from their unconstrained part share $\chi_U(\delta)$; whereas the compliant-unconstrained need not deviate from that unconstrained share in order to comply with the RoO.

Figure 2: Compliance Cost and Compliance Decision



2.3 Laffer Curve for Rules of Origin

The aggregate domestic part share is given by $X(\chi_R) = \int \chi(\chi_R, \tau, \delta) dF(\tau, \delta)$, where $F(\tau, \delta)$ is the joint distribution of firm-level δ s and τ s. For now, we assume a common average τ across firms and focus on the cross-firm variation in δ , which implies a univariate distribution $F(\delta)$ for the aggregate part share $X(\chi_R)$. The variation in that domestic part share at the firm-level as a function of χ_R was shown in the bottom panel of Figure 1 for three firms with different δ s. For firms with high $\delta \geq \bar{\delta}$ that always comply with any RoO (including $\chi_R = 1$), such as the firm with $\delta = 1.25$ in the Figure, a higher χ_R can never induce a lower part share choice $\chi(\chi_R, \tau, \delta)$: Either the firm is compliant-unconstrained and thus does not adjust its part share with the RoO χ_R , or it is compliant-constrained and increases its part share one-for-one with the RoO χ_R . However, the firms with $\delta < \bar{\delta}$ will respond to a higher RoO χ_R by adjusting their domestic part share non-monotonically: As χ_R increases from 0, it will initially induce a firm to increase its part share (when the rise in χ_R induces the firm to switch from compliant-unconstrained to compliant-constrained); but further increases in χ_R will then induce a sharp drop in the part share once the firm switches to non-compliance. This hump-shaped response is shown for the remaining two firms with $\delta = \{0.8, 1\}$ in the Figure (more accurately, a “triangle” shape for a single firm).

Aggregating over a distribution $F(\delta)$ for $\delta < \bar{\delta}$ smoothes out the firm-level responses into a smooth hump-shaped aggregate part share $X(\chi_R)$ curve. This curve starts (at $X(0)$) and stops (at $X(1)$) at the same level when all firms (again, with $\delta < \bar{\delta}$) choose their unconstrained part share $\chi_U(\delta)$:⁸

$$X(0) = X(1) = \frac{1}{F(\bar{\delta})} \int_0^{\bar{\delta}} \chi_U(\delta) dF(\delta).$$

We call this hump-shaped part share curve a Laffer curve for Rules of Origin due to its similarity with the hump-shaped Laffer curve for income tax as a function of the tax rate. Just like that original Laffer curve, a higher RoO χ_R is intended to increase the aggregate part share $X(\chi_R)$ by forcing firms to comply with a higher threshold—but can lead to decreases in $X(\chi_R)$ by inducing firms to switch to non-compliance.

The aggregate part share $X(\chi_R)$ for the set of always compliers (with $\delta \geq \bar{\delta}$) is non-decreasing (and must be increasing once χ_R rises above a threshold). If the distribution $F(\delta)$ is such that the set of always-compliers is dominant, then it is possible for the aggregate curve $X(\chi_R)$ to inherit that non-decreasing shape. Otherwise, the $X(\chi_R)$ curve will still be hump-shaped, but with $X(1) > X(0)$ due to the always-compliers.


We now parametrize $F(\delta)$ to quantitatively assess the Laffer curve for $X(\chi_R)$. We use a symmetric distribution for the percent cost advantage for domestic production $\log \delta$ (whenever negative, this represents the percent cost advantage in favor of Foreign part production). We use a Normal distribution for this cost advantage $\log \delta$ centered at zero so that there is no country-wide comparative advantage in favor of either country. And we set the standard deviation at 0.2. This implies that 10 percent of firms have a cost advantage multiplier greater than 1.39 (5 percent in favor of either domestic or Foreign);

⁸Recall that when $\chi_R = 0$, all firms are compliant-unconstrained; and when $\chi_R = 1$, all firms with $\delta < \bar{\delta}$ are non-compliant. In both cases, all firms choose the same part share $\chi_U(\delta)$.

20 percent have a cost advantage multiplier greater than 1.29 (10 percent in favor of either country); and lastly that 50 percent of firms have a cost advantage multiplier greater than 1.14 (25 percent in favor of either country). We continue with the same other parameter choices that we used in the earlier Figures: $\theta = 4$ and $\tau = 1.1$.

Figure 3 shows how the shares of non-compliant (blue), compliant-constrained (yellow), and compliant-unconstrained (green) firms vary as a function of the RoO χ_R (those three shares always add up to 1). As we previously described in the general case, the share of compliant-constrained firms (orange curve) initially increases from zero as the RoO is tightened (increasing χ_R). Those firms increase their domestic part share and generate an increase in the aggregate domestic part share $X(\chi_R)$. However, as the RoO is further tightened, some complying firms (blue curve, with the highest cost differentials in favor of Foreign, i.e. the lowest δ s) stop complying and revert to their lower unconstrained part share $\chi_U(\delta)$ (and pay the tariff τ). As we described for the general case, this share must monotonically increase with the rule of origin (which raises the compliance cutoff δ^*) regardless of the parametrization for δ . The remaining set of compliant-unconstrained firms (green curve) decreases monotonically from 1 to 0. This holds regardless of the parametrization choice for δ as the δ° cutoff increases with χ_R .


Figure 3: Compliance with Rule of Origin: Firm Shares



Given our parametrization, there are roughly 17% of firms with $\delta \geq \bar{\delta}$ that are always-compliers, as represented by the limit for the yellow curve as $\chi_R \rightarrow 1$. Among those firms, a tightening of the rule of origin χ_R must lead to an increase in the average domestic share. We previously argued that, excluding those firms, the aggregate domestic cost share must exhibit a Laffer curve, regardless of the parametrization choice for δ . We show the Laffer curves induced by our specific parametrization in Figure 4. The yellow curve excludes those 17% of firms. We see that the aggregate share initially increases and then returns to

its initial level as $\chi_R \rightarrow 1$. The blue curve shows the aggregate share for all firms. Because that excluded portion of firms is not too substantial, we see that the $X(\chi_R)$ curve remains hump shaped, though with $X(1) > X(0)$ as we previously described for the general case.

Figure 4: Laffer Curve for Domestic Part Production




In the appendix, we show how different parametrization choices affect the shape of the Laffer curve for $X(\chi_R)$. Those curves are displayed in Appendix Figure B.1 for high and low parameter settings for μ , σ , θ , and τ . While the $X(\chi_R)$ curve changes shape and position in intuitive ways, the basic hump-shaped Laffer curve is robust to those deviations from our benchmark parameter settings.

Our model based on multiple-part sourcing is key in generating this Laffer curve. In contrast, a similar model with a single part would generate a non-decreasing part share curve $X(\chi_R)$. This is illustrated in Figure 5, which replicates Figure 1 when there is a single-part decision made by the three firms. For the purpose of illustration, we assume that half of the firm's part purchases are exogenously sourced within the RTA (domestic) and the firm is considering the sourcing decision for a single part accounting for the other half of its part purchases. The choice of this part threshold is irrelevant for our argument, and could further be heterogeneous across firms.

When the RoO χ_R is below 50%, all firms are compliant-unconstrained, regardless of their δ . Once χ_R rises above 50%, then Foreign-sourcing this single part induces non-compliance with the RoO. Firms with $\delta < 0.8$ choose this option and pay the 10% average tariff penalty because that leads to a lower cost than domestic-sourcing of the part. On the other hand, firms with $0.8 \leq \delta < 1$ choose compliance (constrained) and switch from Foreign to domestic sourcing. And firms with $\delta \geq 1$ are always compliant-unconstrained: they prefer domestic-sourcing for that part regardless of the RoO. Thus, Figure 5 illustrates how, across the distribution of δ s, a higher RoO χ_R can only induce firms to increase

Figure 5: Compliance Cost and Sourcing Decision with a Single Part



their domestic part share (or maintain it at its current level): There is no non-monotonic response at the firm-level, and there hence cannot be a non-monotonic response when aggregating over firms.

3 Extension to Assembly Location Choice

In order to highlight the inherent forces generating a Laffer-curve effect for a RoO policy, we have focused on the part sourcing decision and abstracted from the associated decision regarding the location of assembly. We now extend our model to incorporate the choice of assembly: first in another country within the RTA, then in a (Foreign) country outside the RTA. The compliance decision is then linked with the location of assembly, yet a similar Laffer-curve effect for the rule of origin prevails. When we add trade costs in parts between the RTA and the Foreign location, we show how the negative impact of tighter RoOs for the aggregate regional part share is amplified. When those trade costs are high enough, a tighter RoO can then lower the part share below its initial level with no RoO. And for even higher trade costs, the hump in the Laffer curve can disappear and then tighter RoOs can only lead to decreases in the regional part share as firms relocate assembly to Foreign. For simplicity, we restrict the analysis to the choice of a single assembly location, precluding the splitting of assembly across locations. In essence, this imposes a restriction on the returns to scale in assembly.

3.1 Intra-RTA Assembly Relocation

We initially consider the assembly location choice between heterogeneous countries within the RTA, and further extend the model to analyze a Foreign assembly location. A RoO policy induces heterogeneous effects across countries within an RTA because the penalty for non-compliance is only applied to within-RTA exports. Thus, a country with a larger domestic market (such as the United States within NAFTA) faces a lower average tariff penalty for non-compliance: a greater share of output is sold domestically. If this were the only dimension of heterogeneity across countries in an RTA, then all assembly would relocate to the larger market with the lower average tariff penalty. We thus introduce another dimension of cost heterogeneity in assembly. First, we note that—when transport costs are low—cost heterogeneity in part production (such as our Weibull cost draws) would not affect the assembly location choice: firms can still make their unconstrained sourcing decisions for parts within the RTA independent of the assembly location (we introduce transport costs for parts later when we extend the assembly location to Foreign).

Along with a domestic assembly location (D), a firm can also choose an assembly location elsewhere in the RTA (R). We model relative differences in assembly costs between D and R as a multiplicative cost shifter δ_A^R . Just like we did for parts, we normalize the domestic assembly cost parameter at 1 (this involves a new normalization of costs, which now include assembly). Then, given a total part cost C , the total production cost—including assembly—is C in D and $\delta_A^R C$ in R . Without loss of generality (this is the only distinction between D and R), we assume that assembly production is more costly in R : $\delta_A^R > 1$. Also without loss of generality, we maintain our normalization of an average

part cost within the RTA to 1 (relative to the average of δ for Foreign parts). This subsumes adding separate Weibull cost draws for part production in both D and R .⁹ And lastly, due to differences in their domestic market sizes and MFN tariffs, the average cost penalty for non-compliance with the RoO differ across assembly location choices: $\tau^R > 1$ versus $\tau^D > 1$.

3.1.1 Part Shares and Cost

Regardless of assembly location (D vs R), firms use the same share of regional parts $\chi_U^D(\delta) = \chi_U^R(\delta) = \chi_U(\delta)$; and whenever $\delta \geq \delta^\circ$ (recall that $\chi_U(\delta^\circ) = \chi_R$), a firm can costlessly comply with the RoO at χ_R . It is then compliant-unconstrained. For assembly in D , the compliant-constrained, compliant-unconstrained, and non-compliant costs are the same as previously derived: respectively, $C^D(\chi_R, \delta) = C(\chi_R, \delta)$, $C_U^D(\delta) = C_U(\delta)$, and $\tau^D C_U^D(\delta)$. For assembly in R , the costs include the multiplicative assembly shifter δ_A^R , and the cost penalty for non-compliance is τ^R instead of τ^D . Thus, the compliant-constrained, compliant-unconstrained, and non-compliant costs are $C^R(\chi_R, \delta) = \delta_A^R C^D(\chi_R, \delta)$, $C_U^R(\delta) = \delta_A^R C_U^D(\delta)$, and $\tau^R C_U^R(\delta) = \tau^R \delta_A^R C_U^D(\delta)$. For compliant-constrained firms with $\delta < \delta^\circ$, the cost penalty associated with RoO compliance is the same for both assembly locations (and equal to our previously derived cost penalty): $\tilde{C}^D(\chi_R, \delta) = \tilde{C}^R(\chi_R, \delta) = \tilde{C}(\chi_R, \delta) = C(\chi_R, \delta)/C_U(\delta)$.

3.1.2 Compliance and Assembly Location

Given our assumption of an assembly cost advantage in D ($\delta_A^R > 1$), a compliant firm unaffected by any difference in tariff penalties will always choose assembly in D . If the average tariff advantage in favor of assembly in R does not outweigh the assembly cost advantage in D , then any non-compliant firm will also choose assembly in D . In this case, assembly in D always dominates assembly in R . We therefore focus on the remaining case when the tariff advantage in R outweighs the cost advantage in D : $\tau^R \delta_A^R < \tau^D$. In this case, a non-compliant firm prefers assembly in R (NCR), whereas a compliant firm prefers assembly in D (CD). Thus, the assembly location choice will be tied to the compliance decision with the RoO χ_R , and will be based on the lowest cost between $\tau^R C_U^R(\delta) = \tau^R \delta_A^R C_U^D(\delta)$ (NCR) and $C^D(\chi_R, \delta)$ (CD) so long as $\delta < \delta^\circ$ (otherwise, the firm is compliant-unconstrained and will assemble in D).

This assembly/compliance choice will be determined by a cutoff δ^{DR*} such that

$$\tilde{C}^D(\chi_R, \delta^{DR*}) = \tilde{C}(\chi_R, \delta^{DR*}) = \tau^R \delta_A^R.$$

Firms with δ below this cutoff will choose NCR, while firms above this cutoff will choose CD. As $\tau^R \delta_A^R > 1$ (both $\tau^R > 1$ and $\delta_A^R > 1$), the cutoff δ^{DR*} will be below δ° . Firms with δ

⁹If part production in D and R were subject to separate Weibull cost draws with mean δ^D and δ^R , then a firm would pick the lowest cost part leading to a Weibull distribution of cost (across both D and R) with mean $\left[(\delta^D)^{-\theta} + (\delta^R)^{-\theta} \right]^{-1/\theta}$. This is our new numeraire. We consider transport costs for parts later between the RTA and F along with the assembly location in F .

between those two cutoffs will be compliant-constrained, whereas firms with $\delta \geq \delta^\circ$ will be compliant-unconstrained—but all those compliant firms will assemble in D .

Comparative statics for the cutoff δ^{DR*} are identical to the ones we derived earlier for δ^* —since they are both based on the same cost penalty function $\tilde{C}(\chi_R, \delta)$. In particular, increases in the RoO χ_R will lead to a higher cutoff δ^{DR*} , and a relocation of assembly from D to R along with a switch from compliance to non-compliance with the RoO. Recall that our key assumption distinguishing D and R is: $\tau^R < \tau^R \delta_A^R < \tau^D$. In words, R is characterized by a lower tariff penalty on non-compliant intra-RTA exports, and this is offset by a relatively higher assembly cost in R (otherwise, assembly in R would dominate assembly in D regardless of the RoO χ_R). As we show later on for the case of NAFTA, the tariff penalty τ s are substantially lower for the United States, due to its large domestic market and associated lower levels of intra-NAFTA exports (see Figure 9). Thus, our model therefore explains why a country with the largest domestic market within an RTA (such as the United States for NAFTA) may prefer a higher RoO χ_R : It confers an assembly location advantage.¹⁰

3.2 Foreign Assembly Location

We now extend our assembly location choice to a foreign location, F , outside the RTA. For simplicity, we go back to a single assembly location (D) within the RTA. Extending our previous analysis to allow for another assembly location within the RTA is straightforward, but it requires the handling of more cases. We also incorporate trade costs for parts between the D and F locations, modeled as a symmetric *ad valorem* cost shock $\kappa \geq 1$. When the assembly location was restricted to be in D , such trade costs could just be subsumed in the cost difference δ . However, this isomorphism no longer works once multiple assembly locations are available.

In this scenario, the firm chooses whether to assemble in D or in a foreign country F . There are three important features of assembly in F . First, compliance is impossible because assembly within the region is a necessary condition for almost all rules of origin. This means that foreign assembly entails paying the MFN tariff for exports to the RTA. Those firms are in regime NCF (non-compliant foreign). Firms that produce in D then chose whether to comply with χ_R (regime C) or not (regime NCD). The second feature is that the symmetric $D - F$ trade cost for parts, κ , implies that foreign inputs are cheaper when assembling in country F and correspondingly domestic inputs are more expensive there. Whereas δ could be seen as incorporating the cost of transporting foreign inputs to the domestic market in the model without assembly location choice, δ now only includes the production cost differential associated with foreign parts. The third distinction of country F is that it might have lower assembly costs. Since a foreign assembly cost advantage is isomorphic to a lower trade cost penalty for non-compliance τ^F , this section absorbs those differentials into τ^F . The τ index applicable to domestic production is either 1 (C) or τ^D if noncompliant (NCD).

¹⁰Throughout, we abstract from a normative analysis that would weigh the employment and producer surplus gains associated with higher part and assembly production against the distortions induced by the RoO.

In order to derive equilibrium compliance and location choice decisions, we now need to consider separately the costs patterns for the cases of domestic and foreign assembly.

Starting with the *domestic assembly case*, we define new parts share and cost functions which are tied to the base model functions through the following mapping:

$$\chi_U^D(\delta) = \chi_U(\delta\kappa), \quad C_U^D(\delta) = C_U(\delta\kappa), \quad \text{and} \quad C^D(\chi_R, \delta) = C(\chi_R, \delta\kappa).$$

Note that $\chi_U^D(\delta) \geq \chi_U(\delta)$ because κ raises the domestic part share. Firms will be unconstrained by the content rule when $\chi_U^D(\delta) \geq \chi_R$. This will be the case when $\delta \geq \delta^{D^\circ}$, where $\chi_U^D(\delta^{D^\circ}) = \chi_R$. This cutoff is shifted down from the previous cutoff by κ : $\delta^{D^\circ} = \delta^\circ \kappa^{-1}$.

In the *case of foreign assembly*, the unconstrained domestic part share is

$$\chi_U^F(\delta) = \frac{\kappa^{-\theta}}{\kappa^{-\theta} + \delta^{-\theta}} = \chi_U(\delta/\kappa) \leq \chi_U(\delta). \quad (8)$$

When the firm switches the assembly location from D to F , there is a discrete drop in part share from $\chi_U^D(\delta)$ to $\chi_U^F(\delta)$. The unconstrained cost of parts when assembling in F is

$$C_U^F(\delta) = (\kappa^{-\theta} + \delta^{-\theta})^{-1/\theta} = \kappa C_U(\delta/\kappa). \quad (9)$$

Firm δ has the following cost function based on the regime:

$$C(\delta) = \begin{cases} C_U^D(\delta) & \text{compliant with a non-binding } \chi_R, \\ C^D(\chi_R, \delta) & \text{compliant with binding } \chi_R, \\ \tau^D C_U^D(\delta) & \text{non-compliant, assembly in } D, \\ \tau^F C_U^F(\delta) & \text{non-compliant, assembly in } F. \end{cases} \quad (10)$$

To define the low-cost outcome for a given $\{\chi_R, \delta\}$ tuple, it proves useful to calculate three sets of relative costs.

For assembly in D , define the relative cost of compliance as

$$\tilde{C}^D(\chi_R, \delta) = \begin{cases} \frac{C^D(\chi_R, \delta)}{C_U^D(\delta)} & \delta < \delta^{D^\circ}, \\ 1 & \delta \geq \delta^{D^\circ}. \end{cases}$$



Note that $\tilde{C}^D(\chi_R, \delta) = \tilde{C}(\chi_R, \delta\kappa)$. So it has exactly the same shape of old $\tilde{C}(\chi_R, \delta)$ with a lower bound of 1 when $\delta \geq \delta^{D^\circ}$. Define the relative unconstrained cost (based on assembly location) as

$$\tilde{C}_U(\delta) = \frac{\tau^F C_U^F(\delta)}{C_U^D(\delta)}$$

The cost $\tilde{C}_U(\delta)$ is increasing from $\tau^F \kappa^{-1}$ to $\tau^F \kappa$ as δ goes from 0 to ∞ , and $\tilde{C}_U(1) = \tau^F$.

Thus, the decision rule is given by the following inequalities: whenever $\tilde{C}^D(\chi_R, \delta) \geq \tau^D$, NCD is preferred to C (and vice-versa). When $\tilde{C}_U(\delta) \geq \tau^D$, NCD is preferred to

Figure 6: Cost minimizing compliance with F assembly option
 $\delta_1=0.75$ $\delta_2=1$



NCF (and vice-versa). Lastly, whenever $\tilde{C}^D(\chi_R, \delta) \geq \tilde{C}_U(\delta)$, NCF is preferred to C (and vice-versa). In other words, the choice of C, NCD, NCF is determined by

$$\min \left\{ \tilde{C}^D(\chi_R, \delta), \tau^D, \tilde{C}_U(\delta) \right\}.$$


Figure 6 shows the cost penalty of complying while assembling domestically for two firms, one with low cost of foreign parts (blue, $\delta_1 = 0.75$) and the second with relatively high cost of foreign parts (orange, $\delta_2 = 1$). The horizontal blue and orange lines correspond to the cost penalty of assembling in F and then exporting to the FTA region relative to non-compliant assembly in the domestic country. The reason this line is horizontal because the costs in the numerator and denominator are both unconstrained and hence not influenced by the content rule. For both firms, at a very low χ_R , compliance at $\chi_U^D(\delta) > \chi_R$ is the chosen option. As χ_R rises, the rule soon becomes binding for firm 1, but then becomes too costly to justify. Firm 1 therefore switches to NCF. In the lower panel of the figure, the compliant phase sees content rising one for one with the rule, until the sudden drop when requirement reaches about 55%. In contrast to previous figures, and the case of firm 2, when firm 1 stops complying its regional content falls below the level where it had started when χ_R was 0.3. This sheer drop occurs solely because of $\kappa > 1$. In contrast, the high-foreign-cost has a longer regime of compliance and then only drops back to the starting point (just over 60%). This is because it stays domestic even under non-compliance. In the diagram we see this because τ_D , the penalty for non-compliance when producing domestically, is lower than $\tilde{C}_U\delta_2$. No matter how strict the RoO is, non-compliance abroad is never attractive for this firm.

Figure 7 graphs the $X(\chi_R)$ Laffer curves under three different settings. The non-changing parameters are $\theta = 4$, $\tau^D = \tau^F = 1.1$, and $\sigma = 0.2$. For each case of κ , the mean (μ) of the log-normal draws for δ is shifted down by $\ln(1/\kappa)$. This means that when assembly in F is not an option, the effect of κ on all outcome variables is exactly offset by the shift in δ (because then only the product of κ and δ matters). So regardless of κ , the blue curve always shows the outcome when assembly in F is ruled out. This is the same Laffer curve used earlier in the paper. When κ is higher than 1.1 (the level of τ^F), then the minimum of the $\tilde{C}_U(\delta)$ is below 1, and there are some always NCF firms (even when $\chi_R = 0$). This is why the green line for $\kappa = 1.2$ starts below 0.5. In this case, when χ_R is low, there are no NCD firms and all compliant firms (CD) are compliant-unconstrained. So the initial increases in χ_R from zero have no effect: Both NCF and compliant-unconstrained firms are unaffected by an increase in χ_R . At some point, χ_R rises to a threshold where the CD firm with the lowest δ becomes compliant-constrained. Further increases in χ_R then have two opposite effects on X : some firms become CD constrained—and this raises X ; but other firms switch from CD to NCF—and this lowers X . With $\kappa = 1.15$ (purple curve), that balance is initially positive and then quickly turns negative. But with the green curve depicting $\kappa = 1.2$, that balance is always negative.

4 Calibration for North American auto supply chains

The existence of a hump-shaped relationship between the content requirement and the realized regional content share holds under a wide range of parameters. But as illustrated

Figure 7: ROO Laffer curve with option to relocate assembly



in Appendix figure B.1, the shape and peak of the curve varies considerably as we alter the parameters. In this section, we discipline the parameters with data from the North American automotive industry. In so doing we can offer a tentative answer to whether the rise in the RCR from 62.5% to 75% put the North American auto sector on the wrong side of the Laffer curve.

The quantification we carry out takes assembly locations as given. This can be justified as a medium-run policy analysis, in that firms can adjust the sources of parts but cannot change the location of their final assembly plants. The fixed location model has two fewer parameters as it does not need κ , the cost of shipping parts between regions, or τ^F , which implicitly incorporates relative assembly costs between the RTA and foreign location. In our final simulation, we incorporate market share changes that capture some aspects of assembly relocation. This is because the firms whose costs rise the most (because they stay compliant or start paying tariffs) suffer reductions in output.

In the form shown in section 2, the fixed-location model has just four parameters— μ , σ , θ , and τ —that need to be calibrated. In the model, the equilibrium distribution of regional content shares across firms reflects the distribution of δ , along with θ and τ . We therefore use such data to construct simulated method-of-moments estimates of the parameter values of the δ distribution (namely μ , σ), along with the other two parameters that are added to allow for additional heterogeneity to better fit the data.

The data for the calibration comes from the car industry. The appeal of this industry comes partly because it figured centrally in the negotiation of the new USMCA agreement. The other appealing feature of this sector is the availability of a rich data set on a

core variable in the model, λ , the share of parts costs sourced within the regional trade agreement. The *American Automobile Labelling Act* (AALA) provides annual reports showing the cost share for North American partners for all models of cars and pickup trucks sold in the US and Canada. We use a second source of data on cars and trucks, the IHS Markit automotive sales module, to obtain the market destinations of car models made in North America. As we discuss below, such data are needed to construct an index of τ that takes into account the fact that substantial shares of car output stay in the country where they are assembled and hence are not at risk of paying MFN tariffs. Combining these two data sets allows us to measure λ and τ at the level of a “carline,” a unit that always has the brand, model, and plant location, and often has additional detail on engine size or body type.

Table 1: Counts and medians of NAFTA parts cost shares

Assembled in:	Cars/MPVs		Pickup Trucks	
	Count	Median λ (%)	Count	Median λ (%)
Canada	159	70	2	91
Mexico	291	57	10	75
USA	904	62	119	75

Note: Counts are carline-year observations from 2011 to 2019

Table 1 provides a few statistics from the AALA data. It shows for each NAFTA assembly country the number and median parts cost share (λ in the model) for cars (including SUVs and MPVs) and light trucks (e.g. Ford 150). Two features are notable. Canada has the highest NAFTA content for both types of vehicles. Within countries, light trucks have higher NAFTA parts content and are more likely to be assembled in the US. Both of these comparisons support the model prediction that compliance rates increase with the τ index. Canadian assembly faces high τ because around 80% of production heads to the US. Meanwhile, trucks have high τ because of their 25% tariffs.

4.1 Initial comparison of model and data densities

We begin by considering the model as specified in the theory section, as this will help identify what extensions we need to make in order to achieve a reasonable fit to the observed distribution of NAFTA costs shares. The model predicts the aggregate regional parts share in equilibrium (X) as a function of the rule of origin. In the theory section, we express the rule as a required share of parts (χ_R). While this facilitates exposition, the primary NAFTA rule is a Regional Content Requirement (RCR), expressed in terms of the share of total cost incurred within the region.

We will match the moments of simulated parts cost shares (λ) to the observed AALA data between 2011 and 2019, a period when the actual RCR was 62.5%. Given an RCR, the parts costs requirement is given by

$$\lambda_R = \frac{\text{RCR} - \alpha}{1 - \alpha}, \quad (11)$$

where α denotes the share of final assembly in total costs. First consider the distribution of domestic parts costs shares for car models that are unconstrained by the RCR. Recalling that $F(\delta)$ is the CDF of the foreign cost advantage—with $\ln \delta \sim \mathcal{N}(\mu, \sigma)$ —the cumulative density of λ is given by (see Appendix section A.2 for detailed derivations):

$$G(\lambda) = \Phi \left(\frac{\ln(\lambda/(1-\lambda)) - \theta\mu}{\theta\sigma} \right). \quad (12)$$

Since the log odds ratio of λ is normally distributed for unconstrained firms, the maximum likelihood estimates of $\theta\mu$ and $\theta\sigma$ are given by the mean and standard deviation (respectively) of $\ln(\lambda/(1-\lambda))$. From this, we can see that the parameters μ and σ characterizing the heterogeneity of δ are not separately identifiable from θ when all firms are unconstrained. For this reason, we do not attempt to estimate θ and instead fix it at $\theta = 4$.¹¹

Figure 8: Density of λ : model vs data




Figure 8 plots the density of parts costs shares $g(\lambda)$, obtained by differentiating $G(\lambda)$ shown in equation (12), as an orange line. The parameters for $g(x)$ are the same as in the baseline case of our theory section ($\theta = 4, \mu = 0, \sigma = 0.2, \tau = 1.1$). For this figure we want the regional content requirement to equal the parts cost share requirement (λ_R) so we set $\alpha = 0$. With a binding rule of origin, the model predicts a distribution of cost shares, λ , depicted in green in Figure 8, that is very different from the unconstrained orange density. In this simple setup of the model, there is a single threshold δ^* (such that $\tilde{C}(\chi_R, \delta^*) = \tau$) beyond which firms begin to comply with the rule. This set of parameters implies that 71% of firms have a δ high enough to choose to comply exactly at $\lambda_R = \text{RCR} = 62.5\%$

¹¹Appendix C provides the rationale for choosing $\theta = 4$.

as shown on the spike in the figure. The green shaded density shows the firms that have high costs of domestic inputs (non-compliers) or such low costs that they choose NAFTA shares above the requirement.

The black line in figure 8 depicts the empirical distribution of regional shares for car models produced in Canada, Mexico or the USA, pooled over the year for which we have data and for which the RCR is 62.5%, that is 2011–2019. Two predictions of the simple model are at odds with what we observe in the AALA data: 1) the “hole” in the λ distribution between $\lambda_U(\delta^*) = 18.2\%$ and $\lambda_R = \text{RCR} = 62.5\%$, 2) the spike at λ_R .¹²

As in Eaton et al. (2011), the basic model predicts stronger partitioning of firms into decision regimes than what is observed in the data. In that paper, the problem is that firms do not comply with strict hierarchies across export destinations that the single-heterogeneity model predicts. The authors address this problem by introducing additional realistic dimensions of heterogeneity in the form of idiosyncratic entry cost and demand shocks. In that spirit, our calibration departs from the unrealistic assumption that all firms have the same τ and α . Furthermore, we take into account measurement error by simulating the observed λ to equal the model-generated λ with measurement error. These three generalizations of the model result in it no longer predicting the large empty region followed by bunching exactly the threshold for compliance. Of the three, the heterogeneous τ is the most complex, so we address that first in the next subsection.

4.2 Measurement of τ index heterogeneity

Firms that have plants located within NAFTA sell to destination countries d which we group in three sets, the assembly location ℓ , countries other than ℓ in the regional agreement, $\mathcal{R}$ and other foreign countries outside the agreement $\mathcal{F}$. For each firm, we can define τ as an index capturing the impact on profits of being hit with MFN tariffs ($\tau_d > 1$) in case of non-compliance. This index therefore varies with the dependence of each firm’s sales on different destinations d . We show in appendix A.1 that, for constant price elasticity of demand η , the τ index can be written as

$$\tau = \left(\frac{1 + \sum_{d \in \mathcal{R}} (A_d/K) \tau_d^{1-\eta}}{1 + \sum_{d \in \mathcal{R}} (A_d/K)} \right)^{1/(1-\eta)}, \quad (13)$$

where the A_d are demand shifters and $K \equiv A_\ell + \sum_{d \in \mathcal{F}} A_d \tau_d^{1-\eta}$, a term accounting for all sales outside the markets of RTA partners. K therefore does not depend on the individual compliance decision. A_d potentially incorporates carline-specific demand in each market, transport costs, and a d -specific price index (if demand is CES).

The next step is to calibrate the parameters A_d/K in terms of observables. Let r_d denote the ratio of destination $d \in \mathcal{R}$ sales relative to sales in all $d \in \mathcal{R}$ for a given firm/model. Under CES monopolistic competition, r_d is given by

$$r_d \equiv \frac{p_d q_d}{p_\ell q_\ell + \sum_{d \in \mathcal{F}} p_d q_d} = \frac{A_d \tau_d^{\mathbb{I}(1-\eta)}}{A_\ell + \sum_{d \in \mathcal{F}} A_d \tau_d^{1-\eta}} = \frac{A_d}{K} \tau_d^{\mathbb{I}(1-\eta)}, \quad (14)$$

¹²The latter spike is shown as a broken bar since its true height would be far above the modes of the other densities.

where $\mathbb{I}$ is a Boolean for firms that do not comply with the rule of origin and hence must pay the MFN tariff. We can use this formula to bracket the A_d/K calibration as $r_d \leq A_d/K \leq r_d \tau_d^{\eta-1}$. In practice, the gap between the lower and upper bound is quite small for NAFTA markets, so we simplify the empirical τ index to be

$$\tau = \left(\frac{1 + \sum_{d \in \mathcal{R}} r_d \tau_d^{1-\eta}}{1 + \sum_{d \in \mathcal{R}} r_d} \right)^{1/(1-\eta)}. \quad (15)$$

In order to measure r_d , we use IHS Markit sales data, the same source as Head and Mayer (2019), which provides, for every car model and year combination, the total number sold from a particular assembly plant to a particular final market. We can therefore consider car models produced in Canada, Mexico and USA and for all of them use the sales volume ratios, $q_d/(q_e + \sum_{d \in \mathcal{F}} q_d)$, as proxies for the destination-priced sales ratios, r_d . To match our use of AALA data from 2011 to 2019, we calculate the τ index starting in 2011 and continuing to the most recent year for which we have data, 2018, yielding 1069 observations for carlines assembled in the NAFTA countries.

We expect r_d to vary across models because of heterogeneity in demand as well as differences in transport costs from production location to the consumer. Similarly, τ_d varies because tariffs differ by narrow product code (The US imposes a 25% tariff on light trucks but only 2.5% on passenger cars) and destination (Mexico's MFN tariff is as high as 35%, Canada charges 6%).

Figure 9: Distribution of τ index across models assembled in NAFTA




Figure 9 plots the estimated densities of τ for three assembly locations: Canada (red), Mexico (green), and the USA (blue). The Mexican distribution is particularly interesting. Most of the density lies around a mode near 1.025, the US MFN on passenger cars. There is a second mode near 1.23 corresponding to pickup trucks and commercial vehicles mainly exported to the US where imports of this vehicle type are subject to the 25%

MFN duty. Note that US cars have the lowest τ index, despite the higher MFNs in Canada (6%) and Mexico (35%). This is because of the relatively low shares of US-assembled cars destined to the NAFTA trade partners. The black line pools observations for carlines assembled in all three countries. The distributions turn out to be very irregular and are hard to fit with simple parametric densities. For example, we show the very poor fit of a log-normal distribution for the pooled data.

4.3 Estimating the heterogeneity distribution parameters

There are four parameters to be estimated: the μ and σ from the log-normal δ distribution, and the concentration parameters for α heterogeneity and measurement error. We do this via simulated method of moments.

Each carline in the simulation is characterized by a δ , τ and α values. The number of draws takes the number of carlines in the AALA data, and multiplies it by 20 in order to have a set of draws large enough to resemble a continuum. δ is log-normal distribution with parameters μ and σ . τ is sampled from the distribution shown in figure 9. The carline-level share of assembly in total costs, α , is drawn from a Beta distribution (to ensure that this share stays within 0 and 1) with mean $\bar{\alpha} = 0.15$ and concentration parameter ν_α .¹³ Appendix C provides the data sources underlying the values selected for each non-estimated parameter.

With the draws of δ , α and τ for each simulated carline, we compute the vector of compliance decisions and optimally chosen regional share of parts costs λ as follows.

1. Compute the unconstrained share, $\lambda_U = \chi_U$ (the share of parts costs is equal to share of parts by Eaton and Kortum (2002) property B).
2. Compute the parts costs share in case of compliance to the RCR. For that, we need to transform it into a carline-level share of parts costs (λ_R) using equation (11).
3. Use equation (5) to map λ_R to χ_R . Since δ and α vary, χ_R is firm-specific.
4. With χ_R and χ_U , compute $\tilde{C} = C_R/C_U$, and compare it to τ in order to obtain the equilibrium compliance decision by each carline, yielding equilibrium λ for each firm:

$$\bar{\lambda}(\text{RCR}, \alpha, \delta, \tau) = \begin{cases} \lambda_R & \text{if } \tilde{C}(\lambda_R, \delta) < \tau \text{ and } \lambda_U(\delta) < \lambda_R, \\ \lambda_U(\delta) & \text{otherwise.} \end{cases} \quad (16)$$


5. Allow for measurement error by letting the final value of λ be determined by a draw from a Beta distribution centered around $\bar{\lambda}$ with concentration parameter ν_λ .¹⁴

¹³The Beta distribution is conventionally parameterized in terms of shape parameters a and b . A more useful version—at least for our purposes—has parameters $\bar{\alpha}$ and ν_α , where $\bar{\alpha}$ is the mean and ν_α is referred to as concentration. The latter relates to variance as $\nu = \mu(1 - \mu)/\text{var}(\alpha) - 1$. The correspondence back to the standard parameterization is $a = \bar{\alpha}\nu_\alpha$ and $b = (1 - \bar{\alpha})\nu_\alpha$.

¹⁴Among the sources of error are the AALA exemption for reporting Mexico content if it is below 15%. Additional measurement error comes from rounding which the law permits to the nearest 5%. We also intend for the error to capture deviations from the continuum assumption in the model. Since many parts have non-negligible cost shares, a firm that intends to “just comply” will in fact be observed to over-comply depending on the share of the last part.

The distribution of simulated parts costs shares (λ^{model}) is then compared to the distribution of North American parts cost shares that we observe in the AALA data (λ^{data}). We compute kernel densities of each of the two distributions, round those to the nearest percentile, and compute the L_2 norm between the 100 centiles. We select among our parameter to set to minimize the L_2 norm between the density of the data and the density of the (measurement error augmented) model.¹⁵

Figure 10: Fitted density and RoO Laffer curve



The blue line in figure 10(a) plots the density of λ^{model} . The black line in the figure depicts λ^{data} (the same kernel density shown in Figure 8). The high quality of the fit is revealed by the proximity of the two densities.

To provide external validation, we examined data on the fraction of intra-NAFTA trade that takes advantage of the zero tariffs applicable to imports that meet the rule of origin. We can then compare these preference utilization rates (PUR) for auto trade (HS 8703) to those that emerge from the simulation based on the calibration described above. The true PUR for US-made cars entering Canada is 97% in 2019 (before the change in the regional content requirement in 2020).¹⁶ The calibrated model obtains a 92% PUR. 24.3% of carlines comply constrained (CC) and 67.4% comply unconstrained (CU).

Figure 10(b) plots the calibrated Laffer curve, mapping from content requirements (RCR) to realized content shares (RCS). The content share of an individual carline comprises the share of costs coming from domestic assembly, α , plus the component share

¹⁵The L_2 norm is also referred to as the Euclidean distance and is given by $\sqrt{\sum_{i=1}^{100} (\lambda_i^{\text{model}} - \lambda_i^{\text{data}})^2}$, where i are centiles. The search iterates over a grid of 366,912 potential parameter values of $\{\mu, \sigma, \nu_\alpha, \nu_\lambda\}$.

¹⁶The very high PUR observed for autos in NAFTA is in keeping with the finding of Krishna et al. (2021) that larger firms are more able to overcome fixed documentation costs and learn over time how to comply with RoOs. In this case, the major auto makers have been complying with RoOs since the 1965 Auto Pact.

of total costs, $1 - \alpha$, multiplied by the share of parts costs that are procured within the region, λ . The aggregate RCS is the average value of $\alpha + \lambda(1 - \alpha)$ across all the simulated carlines.

The Laffer curve in figure 10(b) begins at about 67% when there is no rule of origin. Increases in the RCR have no effect until they reach about 40%. This is because of the combination of assembly costs and unconstrained preferences for locally-sourced parts. For the estimated parameters, the Laffer curve peaks at an RCR of 80%. Note that the realized content share would be much lower at just over 70%. Although the 62.5% rule, under the original NAFTA, was on the upward sloping part of the Laffer curve, the visual impression of a steep upward slope is the consequence of the narrow vertical range. Raising the RCR by 20% ($75/62.5 - 1$) increases the RCS by only 2.5% ($70.1/68.4 - 1$).

5 Simulating counterfactual market aggregates

In our theoretical derivations, we derived predictions for the unweighted share of parts sourced within the regional agreement (X). This is as far as we can go without imposing further structure on demand, market structure, and the relative size of assembly versus parts employment. However, by adding a small set of additional assumptions, we can construct all the key market aggregates: the price index, vehicle output, and employment. In particular, we now impose CES monopolistic competition for carlines with elasticity η .

We use the simulated aggregates to derive the Laffer curve for employment. This is of great interest because it incorporates two important mechanisms: 1) when RoOs make regionally-assembled cars more expensive, they will decline in market share and this will reduce the derived demand for regional parts and 2) the decline in market share causes further employment losses in final manufacturing (vehicle bodies and assembly).

The demand-side extension allows us to evaluate employment changes associated with any given change in a RoO policy. For any given RoO χ_R , we can simulate its impact on firm outcomes, namely their compliance choice and the resulting sourcing choice (share of regional parts) and associated part cost. In our calibration, we have relied on additional assumptions for an assembly cost share α in order to convert the part share χ predicted by our theoretical model to the RCR specified by the RoO policy (62.5% for NAFTA and 75% for the new USMCA); and for CES monopolistic competition for carlines with elasticity η in order to empirically measure the distribution of tariffs based on the observed market shares of the carlines in our data.

The simulated carlines are represented by N^D draws of the triplet $(\delta_j, \tau_j, \alpha_j)$ distributed according to the calibrated parameters. Throughout, we use the superscript D to denote the set of all carlines $j \in D$ assembled within the RTA. In addition, there is also a set of Foreign-assembled carlines (superscript F). For now, we assume that the location of assembly is fixed, and therefore that the sourcing choices of the Foreign carlines are independent of the RoO policy. For each draw j , a given RCR maps to a RoO part-share χ_R , and our model then predicts a compliance decision (Comply-Unconstrained, Comply-

Constrained, Non-Compliant), and an associated regional part share χ_j and part cost C_j :

$$\chi_j = \begin{cases} \chi_U(\delta_j) \\ \chi_R \\ \chi_U(\delta_j) \end{cases} \quad C_j = \begin{cases} C_U(\delta_j) & \text{Comply-Unconstrained} \\ C(\chi_R, \delta_j) & \text{Comply-Constrained} \\ \tau_j C_U(\delta_j) & \text{Non-Compliant} \end{cases}$$

5.1 Exact hat aggregation of simulation draws

For a given RoO, we evaluate the changes in χ_j and C_j relative to a benchmark policy with no rules of origins ($\chi_R^\circ = 0$):

$$\hat{\chi}_j = \frac{\chi_j}{\chi_j^\circ} = \frac{\chi_j}{\chi_U(\delta_j)}, \quad \hat{C}_j = \frac{C_j}{C_j^\circ} = \frac{C_j}{C_U(\delta_j)}.$$

We assume that rules of origin do not impact the share α_j of assembly costs. Given constant markups, the price for carline j relative to the $\chi_R^\circ = 0$ benchmark is then given by the impact of the part cost change $\hat{C}_j$ on total cost:

$$\hat{P}_j = \hat{C}_j^{1-\alpha_j}.$$

Our last assumption is that under this benchmark policy $\chi_R^\circ = 0$, market shares are independent of the $(\delta_j, \tau_j, \alpha_j)$ draws. In other words, those parameters confer a *comparative advantage* with respect to the location of part production—but *no absolute cost advantage*.¹⁷ Hence, there is no variation in market shares across draws j under $\chi_R^\circ = 0$: that distribution is degenerate. We can now use the “exact hat algebra” developed by Dekle et al. (2007) to aggregate the carline price changes $\hat{P}_j$ into the CES price index change for all carline draws $j \in D$:¹⁸

$$\hat{P}^D = \left(\sum_j \frac{1}{N^D} \hat{P}_j^{(1-\eta)} \right)^{\frac{1}{1-\eta}}. \quad (17)$$

The aggregate CES price index P incorporates $\hat{P}^D$ as well as the Foreign-assembled carlines F . Changes in the rules of origin do not affect either the part or assembly cost for those carlines. Thus, the CES price index P^F for that bundle of Foreign carlines is invariant to the RoO: $\hat{P}^F = 1$. Using this and applying the same exact hat algebra, the impact of the RoO on the overall price index is then given by the price index change for those two groups of carlines (D and F) along with their initial market shares:

$$\hat{P} = \left[(1 - s^{F^\circ}) \left(\hat{P}^D \right)^{1-\eta} + s^{F^\circ} \right]^{\frac{1}{1-\eta}}, \quad (18)$$

¹⁷Our theoretical model was developed with a version where low δ conferred an absolute cost advantage in part production. However, as we noted, this is only for analytical tractability and only affects the initial market shares of firms (which we did not use in any of our theoretical derivations). All the predictions related to changes in firm outcomes would have been identical under an assumption where δ does not confer any absolute advantage, as we assume here.

¹⁸Under $\chi_R^\circ = 0$, there is no variation in the market shares and they are thus equal to $1/N^D$ for all j .

where s^{F° is the share of Foreign-assembled carlines under $\chi_R^\circ = 0$. We describe below how we infer this share s^{F° based on the share of Foreign-assembled carlines under NAFTA, which we observe in our data. This price index change captures the overall impact of the RoO for consumers.

We can also use the impact of a RoO on the price index $\hat{P}$ to infer the market share changes for any simulated carline j , and the sets of all Domestic-assembled and Foreign-assembled carlines:

$$\hat{s}_j = \left(\frac{\hat{P}_j}{\hat{P}} \right)^{1-\eta}, \quad \hat{s}^D = \left(\frac{\hat{P}^D}{\hat{P}} \right)^{1-\eta}, \quad \hat{s}^F = \left(\frac{1}{\hat{P}} \right)^{1-\eta}. \quad (19)$$

Assuming a constant consumer expenditure on cars (unaffected by the RoO), the associated change in output (production) for any carline j is:

$$\hat{q}_j = \frac{\hat{s}_j}{\hat{P}_j}.$$

In turn, this carline production change will induce a change in regional part production given by

$$\hat{L}_j = \hat{\chi}_j \hat{q}_j = \hat{\chi}_j \frac{\hat{s}_j}{\hat{P}_j}.$$

We assume that employment for a part producer changes in proportion to its output and therefore denote this production change $\hat{L}_j$. Those employment changes can also be aggregated across all carlines $j \in D$ in a similar way:

$$\hat{L}^D = \hat{X}^D \frac{\hat{s}^D}{\hat{P}^D},$$

where $\hat{X}^D = X^D / X^{D^\circ}$ denotes the average domestic part share change. The average domestic part share X^D averages the carline part shares χ_j weighted by their output shares. As we did for the initial market shares under $\chi_R^\circ = 0$, we also assume that the initial output and employment shares are independent of the $(\delta_j, \tau_j, \alpha_j)$ draws. This implies that the initial ($\chi_R^\circ = 0$) average part share is unweighted and that the subsequent output weights are equal to the output changes $\hat{q}_j$:

$$X^{D^\circ} = \frac{1}{N^D} \sum_j \chi_U(\delta_j), \quad X^D = \frac{\sum_j \hat{q}_j \chi_j}{\sum_j \hat{q}_j}.$$

We can also calculate the assembly employment change related to the change in the final car output $\hat{q}_j$ for each carline.¹⁹ Just like we did for employment in the part sector, we assume that this employment is proportional to output, so the employment change for each carline is given by $\hat{L}_{Aj} = \hat{q}_j$. Aggregating across the set of all carlines, this output and employment change is given by $\hat{L}_A^D = \hat{s}^D / \hat{P}^D$.

¹⁹We think of assembly broadly as any production activity that is tied to domestic assembly of motor vehicles.

Lastly, we can also aggregate the changes in total employment $\hat{L}_T^D = \hat{L}_A^D + \hat{L}^D$ across parts and assembly so long as we can observe the ratio of assembly to parts employment $\zeta \equiv L_A^D/L^D$ for a given RoO. Let $\zeta^\circ \equiv L_A^{D^\circ}/L^{D^\circ}$ denote this ratio for all domestic-assembled carlines under $\chi_R^\circ = 0$. We describe below how we infer this ratio based on its observed level ζ under NAFTA. The total employment change for all domestic carlines is then given by:

$$\hat{L}_T^D = \frac{L_A^D + L^D}{L_A^{D^\circ} + L^{D^\circ}} = \frac{L_A^{D^\circ}}{L_A^{D^\circ} + L^{D^\circ}} \hat{L}_A^D + \frac{L^{D^\circ}}{L_A^{D^\circ} + L^{D^\circ}} \hat{L}^D = \frac{\zeta^\circ}{\zeta^\circ + 1} \hat{L}_A^D + \frac{1}{\zeta^\circ + 1} \hat{L}^D.$$

Inferring the initial market share of Foreign-assembled carlines Our derivation for the price index change in (18) relies on the initial share of Foreign-assembled carlines s^{F° when $\chi_R^\circ = 0$. We do not observe this share, but we do observe the subsequent share of Foreign-assembled carlines s^F under the NAFTA RCR of 62.5%. Given $\hat{P}^F = 1$ (no change in costs for those carlines), the change in market share from $\chi_R^\circ = 0$ to the NAFTA RCR is given by $\hat{s}^F = \hat{P}^F = \hat{P}^{\eta-1}$ (see 19): An increase in the price index induced by an increase in the cost of domestic-assembled carlines will in turn induce an increase in the market share of the Foreign carlines whose cost remains unchanged. Thus, the initial Foreign share s^{F° can be recovered from its share s^F under any RCR using

$$s^{F^\circ} = s^F \hat{P}^{1-\eta}. \quad (20)$$

Jointly, (18) and (20) allow us to recover both the price index change $\hat{P}$ and the initial Foreign share s^{F° using the simulation outcome (the set of price changes $\hat{P}_j$) under the NAFTA of 62.5% and the associated Foreign-assembled share s^F .²⁰ That initial share s^{F° can then be used to simulate the price index change $\hat{P}$ (and all the other welfare metrics) for any other RoO simulation.


Inferring the initial ratio of assembly to parts employment Our aggregation of employment changes in the assembly and part sectors relies on the initial ($\chi_R^\circ = 0$) ratio of assembly to parts employment for all domestic-assembled carlines ζ° . We do not observe this ratio, but we do observe it subsequent level ζ under NAFTA RCR of 62.5%. Since $\zeta = \zeta^\circ \hat{L}_A^D / \hat{L}^D$, we can recover ζ° based on that observed ratio ζ and our predicted employment changes $\hat{L}_A^D$ and $\hat{L}^D$ under NAFTA.

5.2 Laffer Curves for North American Rules of Origin

Figure 11 plots the percentage change in parts shares and employment relative to a non-binding rule of origin—any RCR from 0% to 31.5%, under our parameter values. The Laffer curves peak for at 81% for parts shares (X), 78% for parts employment (L), and just 75.4% for sectoral employment including assembly and bodies (L_T). Thus, the US proposal for an RCR of 85% is on the wrong side of all three curves. The USMCA (RCR

²⁰In our simulation results in the next section, we use $s^F = .25$: Under NAFTA, 25% of the carlines consumed in the United States and Canada are assembled outside of NAFTA.

Figure 11: Laffer curves for parts and employment



=75%) increase in parts employment relative to NAFTA (RCR=62.5%) is 2.3%, much less than industry associations predicted. The increase in sector employment for the 85% RCR is just 1.2%. Comparisons to actual data are problematic because of the Covid disruption in 2020 and the fact that we only have one complete year of employment data since the USMCA entered into force. There is very preliminary evidence that employment has fallen with assembly falling more than parts (as predicted). Preference utilization has also fallen.

The exact hat algebra also generates counterfactual changes in the price index, as shown in equation (17), a key component of welfare. The rise in the RCR from 62.5% to 75% raises the price index of regionally assembled carlines (P^D) by 0.3%. Further increase from 75% to the US proposal of 85% would raise the price index by 0.4%. A major reason these increases are so modest is that only 13% of carlines are compliant-constrained before and after the stricter RoO is imposed.²¹

6 Conclusion

We analyzed how heterogeneous firms respond to rules of origin when they face the choice of within-RTA or outside-RTA part sourcing for a large number of parts. Firms can choose to comply with the RoO, which raises their part cost when their unrestricted sourcing choices would not comply with the RoO. Or firms can choose not to comply and pay the tariff penalty on their intra-RTA exports. When the RoO is below a threshold, the former compliance effect dominates and the average regional part sourcing increases with the RoO—at the expense of higher costs, and hence higher prices for consumers.

²¹Our companion paper, Head et al. (2022), breaks down changes in prices and employment according to the compliance status of each carline before and after various changes in the RoO.

But when the RoO rises above a threshold, the latter non-compliance effect dominates and a tighter RoO then leads to both *lower* regional part sourcing and higher costs and consumer prices. We call this the RoO Laffer curve.

We do not analyze the policy planner's optimal choice of a RoO, which would depend on the weights placed on producer surplus and employment in both the part and final good sectors, as well as consumer surplus for the final good. However, there is no specification for those weights which would lead to an optimal RoO on the backward bending portion of the Laffer curve. Instead, we quantitatively assess the shape of the Laffer curve for North American autos, using a version of our model calibrated to fit data on regional part cost shares for all autos sold in North America. We obtain the Laffer curve by simulating and aggregating the predictions of our model across the range of counterfactual RoOs. We also derive the Laffer curve for North American auto employment combining both part production and assembly, which takes into consideration the negative impact of higher auto prices on reduced demand for autos assembled in North America (with a substitution effect towards auto imports into North America that are not affected by the RoO). This Laffer curve reveals that the 2020 revision of NAFTA to the new USMCA, which raised the auto content requirements from 62.5% to 75%, will raise employment by only 1.2% while increasing auto prices assembled in the region by 0.3%. And the higher requirement initially proposed by U.S. negotiators (85%) is on the backward bending portion of the Laffer—so that it would induce *reductions* in auto employment along with higher auto prices. We also show that assembly plants in Canada and Mexico are disproportionately affected by increases in the RoO because the penalty for non-compliance is much more severe given their reliance on intra-NAFTA exports to the large US market. This is notwithstanding the U.S.'s relatively lower MFN tariff on passenger vehicles at 2.5%: US-assembled autos are predominantly sold within the U.S. and are therefore unaffected by the RoO.

Several aspects of our quantitative predictions tend to understate the negative impacts of the stricter RoO brought in by the USMCA. First, the rise in the RCR to 75% was only one of several new requirements, including core parts requirements and wage requirements that are binding in Mexico. Second, our predictions are based on the version of our theoretical model when assembly location is fixed. Theoretically, we show that this assembly relocation choice along with differences in trade costs to deliver parts across the assembly locations amplify the negative impact of higher RoOs. High content requirements can then lead to a lower regional part share than in the absence of any content rules, as firms relocate assembly outside the region. This generates an even more severe impact for total auto employment given the associated losses in the assembly sector. Those incentives for the relocation of part production and assembly will hold for vehicles intended for USMCA consumers—but will be particularly salient for vehicles currently assembled in the USMCA area for export to Asia or Europe.

References

Amiti, M. and J. Konings (2007). Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *The American Economic Review* 97, 1611–1638.

- Anson, J., O. Cadot, A. Estevadeordal, J. d. Melo, A. Suwa-Eisenmann, and B. Turchudur (2005). Rules of Origin in North–South Preferential Trading Arrangements with an Application to NAFTA. *Review of International Economics* 13(3), 501–517.
- Antras, P., T. C. Fort, and F. Tintelnot (2017). The margins of global sourcing: theory and evidence from US firms. *American Economic Review* 107(9), 2514–64.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica* 63(4), 841–890.
- Bombarda, P. and E. Gamberoni (2013). Firm Heterogeneity, Rules of Origin, and Rules of Cumulation. *International Economic Review* 54(1), 307–328.
- Bombardini, M., C. Keith Head, M. D. Tito, and R. Wang (2021). How the breadth and depth of import relationships affect the performance of canadian manufacturers. *Canadian Journal of Economics/Revue canadienne d'économie* 54(4), 1525–1561.
- Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2018). What drives home market advantage? *Journal of International Economics* 110, 135–150.
- Conconi, P., M. García-Santana, L. Puccio, and R. Venturini (2018). From Final Goods to Inputs: The Protectionist Effect of Rules of Origin. *American Economic Review* 108(8), 2335–2365.
- Conlon, C. and J. Gortmaker (2020). Best practices for differentiated products demand estimation with PyBLP. *The RAND Journal of Economics* 51(4), 1108–1161.
- DeGroot, M. H. and M. J. Schervish (2002). *Probability and statistics*. Addison-Wesley.
- Dekle, R., J. Eaton, and S. Kortum (2007). Unbalanced trade. *American Economic Review* 97(2), 351–355.
- Demidova, S., H. L. Kee, and K. Krishna (2012, July). Do trade policy differences induce sorting? Theory and evidence from Bangladeshi apparel exporters. *Journal of International Economics* 87(2), 247–261.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Eaton, J. and S. Kortum (2012). Putting Ricardo to work. *The Journal of Economic Perspectives* 26(2), 65–89.
- Eaton, J., S. Kortum, and F. Kramarz (2011). An anatomy of international trade: evidence from French firms. *Econometrica* 79(5), 1453–1498.
- Felbermayr, G., F. Teti, and E. Yalcin (2019). Rules of origin and the profitability of trade deflection. *Journal of International Economics* 121, 1032–48.
- Fontagné, L., H. Guimbard, and G. Orefice (2022). Tariff-based product-level trade elasticities. *Journal of International Economics* 137, 103593.

- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010). Imported Intermediate Inputs and Domestic Product Growth: Evidence from India. *The Quarterly Journal of Economics* 125(4), 1727–1767.
- Goldberg, P. K. and F. Verboven (2001). The evolution of price dispersion in the European car market. *The Review of Economic Studies* 68(4), 811–848.
- Gopinath, G. and B. Neiman (2014). Trade Adjustment and Productivity in Large Crises. *American Economic Review* 104(3), 793–831.
- Grossman, G. M. (1981). The theory of domestic content protection and content preference. *The Quarterly Journal of Economics* 96(4), 583–603.
- Halpern, L., M. Koren, and A. Szeidl (2015, December). Imported Inputs and Productivity. *American Economic Review* 105(12), 3660–3703.
- Head, K. and T. Mayer (2019). Brands in motion: How frictions shape multinational production. *American Economic Review* 109(9), 3073–3124.
- Head, K., T. Mayer, and M. Melitz (2022). The unintended consequences of high regional content requirements. Manuscript.
- Ju, J. and K. Krishna (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.
- Krishna, K., C. Salamanca, Y. Suzuki, and C. Volpe-Martincus (2021). Learning to use trade agreements. Working Paper 29319, National Bureau of Economic Research.
- Ornelas, E. and J. L. Turner (2022). The costs and benefits of rules of origin in modern free trade agreements. Discussion Paper 1867, Centre for Economic Performance.
- Sytsma, T. (2022, February). Improving preferential market access through rules of origin: Firm-level evidence from Bangladesh. *American Economic Journal: Economic Policy* 14(1), 440–72.
- Tintelnot, F. (2017). Global production with export platforms. *The Quarterly Journal of Economics* 132(1), 157–209.

A Derivations

A.1 τ index

Proposition. Consider a profit-maximizing firm under monopolistic competition. Marginal costs are a constant c . The firm sells to destination markets $i = 1, 2, \dots, N$ with constant elasticity demand $q_i = A_i p_i^{-\eta}$ ($\eta > 1$) and exogenous market demand $A_i > 0$. The firm is indifferent

between increases in delivered cost $\tau_i \geq 1$ per market and a common delivered cost increase $\tau \geq 1$ common across markets such that:

$$\tau = \left(\sum_{i=1}^N \frac{A_i}{\sum_{i=1}^N A_i} \tau_i^{1-\eta} \right)^{1/(1-\eta)}.$$

Proof. Consider the ratio of operating profits (gross of a common fixed cost) from the per market τ_i and a common cost τ :

$$\frac{\sum_{i=1}^N A_i (\tau_i c)^{1-\eta}}{\left(\sum_{i=1}^N A_i \right) (\tau c)^{1-\eta}} = \frac{\sum_{i=1}^N \frac{A_i}{\sum_{i=1}^N A_i} \tau_i^{1-\eta}}{\tau^{1-\eta}} = 1.$$

□

Corollary. Consider two destinations with market demands A_1 and A_2 and delivered costs $\tau_1 \leq \tau_2$. Then $\tau_1 \leq \tau \leq \tau_2$ such that:

- τ increases from τ_1 to τ_2 as A_2 goes from 0 to ∞
- τ decreases from τ_2 to τ_1 as A_1 goes from 0 to ∞
- τ decreases from $\exp\left(\frac{A_1}{A_1+A_2} \ln \tau_1 + \frac{A_2}{A_1+A_2} \ln \tau_2\right)$ (the geometric mean of τ_1 and τ_2) to τ_1 as η increases from 1 to ∞

A.2 Distribution of λ_U

(DeGroot and Schervish, 2002, p. 160) show that if X has CDF $F(x)$ and $Y = r(X)$ with inverse function $X = s(Y)$, then if $r' > 0$, $G(y) = F(s(y))$ and if $r' < 0$, $G(y) = 1 - F(s(y))$.

Recall that $\lambda_U = \chi_U = (1 + \delta^{-\theta})^{-1}$, so $r(\cdot)$ is monotonically increasing. Inverting, $s(Y) = (Y/(1 - Y))^{1/\theta}$. Applying the theorem to δ with CDF $F(\delta)$, we have a CDF for λ of


$$G(\lambda) = F\left[\left(\frac{\lambda}{1-\lambda}\right)^{1/\theta}\right] = \Phi\left(\frac{\ln(\lambda/(1-\lambda)) - \theta\mu}{\theta\sigma}\right),$$

where the second equality comes from inserting the log-normal functional form assumption, with $\Phi(\cdot)$ being the CDF of the standard normal. The PDF, $g(\lambda)$ is just the derivative of $G(\lambda)$ with respect to λ , yielding


$$g(\lambda) = f\left[\left(\frac{\lambda}{1-\lambda}\right)^{1/\theta}\right] \frac{1}{\theta\lambda^2} \left(\frac{\lambda}{1-\lambda}\right)^{1+1/\theta} = \phi\left[\frac{\ln o(\lambda) - \theta\mu}{\theta\sigma}\right] \frac{o(\lambda)^{1+1/\theta}}{\theta\lambda^2}, \quad (\text{A.1})$$

where $o(\lambda) \equiv \lambda/(1 - \lambda)$ is the odds ratio of the λ .


Figure B.1: The RoO Laffer Curve varying μ , σ , and θ




(a) δ means (μ) differ



(b) δ dispersions (σ) differ



(c) θ differences



(d) τ differences

B The RoO Laffer curve: alternative parameter settings

Figure B.1 shows the sensitivity, or lack thereof in some cases, of shape of the RoO Laffer curve to changes in the key parameters. Panels (a) and (b) consider different distributional parameters for $F(\delta)$ (the distribution of the home cost advantage). Increasing μ , the mean of $\ln \delta$ increases regional content for any χ_R ; it also shifts the peak of the curve to the right and increases the fraction of always compliers. Raising σ , the standard deviation of $\ln \sigma$ leaves the χ_R that generates peak local content approximately unchanged. However, the curve flattens and the fraction of always compliers increases. Panel (c) varies θ , the parameter that determines the strength of comparative advantage within the firm affects the shape of the curve and the location of the peak. However, even large changes (doubling it to $\theta = 8$, halving it to $\theta = 2$) do not lead to curves that are strikingly different from each other. Finally, in panel (d), we show that the curve looks very different when we consider more extreme tariff values. With a 25% tariff, a high χ_R can generate 80% compliance. With $\tau = 1.025$ the actual MFN tariff the US charges for motor vehicles, average regional content shares peak at less than 55% for χ_R around the old NAFTA rule of 62.5%.

C Sources of non-calibrated parameters

We conduct a grid search to determine the best-fit parameters characterizing the δ distribution (μ and σ), as well as the concentration parameters for the α distribution and measurement/specification error of the model. There are four important parameters that we obtain either from the literature or from aggregate data: $\bar{\alpha} = 0.15$, $\theta = 4$, $\eta = 4$, and $\zeta = 0.63$. We provide the basis for these parameter settings below.

For the mean of the α distribution we rely on aggregated industry sources for three different car producing countries (which happen to be the homes of the three authors). Canadian Annual Survey of Manufacturing (ASM) data (<https://www.ic.gc.ca/app/scr/app/cis/manufacturing-fabrication/33611>) for NAICS 33611 (which includes pick-up trucks) records materials (M) of \$47.7bn and shipments (S) of \$56.4bn (both in CAD). Applying the formula $\alpha = 1 - M/S = 0.15$. Applying the same formula to the US ASM data (<https://data.census.gov/cedsci/table?q=AM1831BASIC&n=336111>) yields $\alpha = 0.18$. In similar French manufacturing data, we calculate $\alpha = 0.16$. These calculations attribute all costs that are not materials to assembly. This means that they are probably upper bounds. However, one extra bit of anecdotal evidence favors $\alpha = 0.15$: Stellantis CEO Carlos Tavares stated that “suppliers represent 85% of the production costs of a vehicle” (Detroit News May 16, 2022).

The parameter θ in our model captures heterogeneity across parts in the relative cost of foreign and regional inputs. It also acts as the trade elasticity for parts sourcing. We therefore draw on the literature that has estimated trade elasticities. We use $\theta = 4$ because it is the preferred value in Eaton and Kortum (2012). They note that $\theta = 4$ finds support in price data and is “in line with several earlier studies based on other evidence.” We have also carried out the quantification using $\theta = 6$, which has the justification of being close to the absolute value the tariff elasticity estimate of -6.3 obtained by Fontagné et al. (2022) for the harmonized system code 8708 (parts of motor vehi-

cles). Estimates for all HS codes are posted at <https://sites.google.com/view/product-level-trade-elasticity>. The fit to the AALA data is the same out to three decimals for $\theta = 4$ and $\theta = 6$.

To set the CES elasticity η , we reviewed several sources of average own price elasticities for cars. Conlon and Gortmaker (2020) Table 8 replicates the estimation of Berry et al. (1995) and obtains an average own-price elasticity of -3.93 . Also in Table 8 they implement a set best practices and recover an average own-price elasticity of -3.46 . For European countries in 1990, Goldberg and Verboven (2001) find average own price elasticities ranging from -4.09 to -6.21 . Using more recent data, Coşar et al. (2018) report in their online appendix (table B.5) own-price elasticities for five car models ranging from -14.1 to -15.7 . Using 2000–2018 data from IHS Markit, and an entirely different methodology, Head and Mayer (2019) estimate $\eta = 3.87$. The similarity of this estimate to the BLP replication motivates us to set $\eta = 4$ in our simulations.

The parameter ζ is defined as the ratio of assembly employment to parts employment (L_A/L in our notation). Our source is the Employment, Hours, and Earnings from the Current Employment Statistics survey (United States national data). For L_A we sum the NAICS of 3361 and 3362 (Motor vehicles, Motor vehicle bodies and trailers). We measure L as NAICS 3363 (Motor vehicle parts). The Series Id for the three variables are CEU3133610001, CEU3133620001, and CEU3133630001, available here https://www.bls.gov/iag/tgs/iagauto.htm#emp_national. Averaging over the available data prior to 2020 (which is shocked by massive Covid employment reductions just prior to the USMCA entering into force), that is 2012–2019, we obtain $\zeta = 0.63$.