The Unintended Consequences of High Regional Content Requirements[1]

Keith Head*, Thierry Mayer† and Marc Melitz‡

1. Introduction

When the Trump administration launched its revision of the treaty governing trade between the US and its neighbors, the US negotiators emphasized the need for stricter rules of origin. The US Trade Representative, Robert Lighthizer, reportedly asked his counterparts to raise the Regional Content Requirement (RCR) to 85%, a large increase from the level set in 1993 (62.5%).[2] Canada and Mexico balked at such a high rate, and the three parties finally settled on an increase to 75%, bolstered with additional binding requirements. The political appeal of stricter origin rules lies in the hope that they will increase domestic employment in the parts industry. Lighthizer (2020) acknowledged this intent, writing "The USMCA rebalances the NAFTA to promote increased production in the United States and North America."

From an economic standpoint it is hard to justify onerous restrictions on sourcing. If the goal is merely to limit imports of parts, then tariffs on parts would be a more efficient tool. While trade agreements that lack a common external tariff need some rule of origin to prevent back-door entry to the high-tariff market via the low-tariff country, this issue was not relevant in the USMCA negotiation for two reasons. First, because the actual differences in tariffs were small, so much smaller content restrictions would be sufficient to prevent this tariff-hopping.[3] Second, it was the lower-tariff member, the US, that was asking for the stricter rules.

Going back to the work of Grossman (1981), economists have investigated whether, even as protectionist devices, strict rules of origin could fail to achieve their goals. Grossman's Proposition 3 states that small increases in local content requirements

[1]We thank participants at seminars, and particularly Emily Blanchard and Felix Tintelnot, for helpful suggestions. Constanza Abuin, Anhua Chen, and George Hu provided superb research assistance.
*Sauder School of Business, University of British Columbia and CEPR (keith.head@sauder.ubc.ca)
†Sciences-Po, CEPII and CEPR, (thierry.mayer@sciencespo.fr)
‡Harvard University, CEPR and NBER, (mmelitz@harvard.edu)
[2]Husisian et al. (2018) note the 85% proposal in their overview of the USMCA.
[3]Felbermayr et al. (2019) present evidence that this argument applies to most rules of origin.

have ambiguous effects on industry value added, defined as the sum of value added in components and in final goods.

> "Whereas the content protection policy causes an increase in the output of domestic components, it will normally result in a concomitant contraction of final good production. Which effect will dominate depends on how sensitive intermediate good production is to changes in its output price, and how sensitive final good production is to changes in the price of its intermediate input."

In this chapter, we extend the Grossman approach to take into account the very large number of diverse parts that go into modern manufactured goods such as automobiles. For each part, the firm decides whether to source it from inside the region (where there is a free trade agreement) or from outside countries. The core tradeoff the firm faces is that within-region sourcing helps it comply with Rules of Origin (RoO), but necessitates forgoing opportunities to obtain cheaper parts elsewhere. In section 4, we give an overview of the theoretical model developed in Head et al. (2022) that analyzes these tradeoffs. We show that RoOs generate competing incentives for part sourcing within an RTA. Even though the rules are intended to relocate production of parts within the RTA, they can have the opposite effect when they are overly restrictive. This main result does not work via declines in final goods production, as in Grossman (1981). However, we also quantify the negative impact of higher costs induced by the RoOs for part production. This quantification exercise predicts how *any* given RoO would affect market share changes and the associated production and employment changes across all vehicle plans selling in the region. Drawing on the attractive aggregation properties of our model, we derive average price, market share, production, and part employment changes across groups of carlines—including the group of all carlines assembled within the region.

This paper is organized as follows. Section 2 provides an overview of recent changes in rules of origin that impacted the auto industry in North America and Europe. The following section presents empirical patterns of sourcing in North America that inform the model and the way we quantify it. Section 4 summarizes the key mechanisms of the model developed in Head et al. (2022). We then estimate the model to fit the pattern of sourcing observed at the level of individual car models prior to the 2020 changes in RCRs. Section 5 describes how we use that fitted model to evaluate the impact of counterfactual RoOs. Section 6 reports the effects of changing those rules for both NAFTA and the EU-UK trade agreement.

2. Changing rules of origin in North America and Europe

Rules of origin in the auto industry were first introduced in the 1965 Auto Pact between Canada and the United States. To avoid non-US companies setting up sales enterprises in Canada to serve the US market, it was agreed that only cars with 50% content from the US and Canada would benefit from the new tariff-free regime.[4] In the negotiation of the North American Free Trade Agreement in 1991 the American side sought a more restrictive rule. Irwin (2017) describes the initial negotiating positions and how they reached the peculiar regional content requirement of 62.5%:

> Rules of origin were particularly important in the case of automobiles. The US auto industry wanted high North American content rules to ensure that Mexico did not become an export platform for Japanese or other foreign producers who would simply send parts to Mexico for assembly and then ship the vehicles into the United States... For NAFTA, the United Auto Workers pushed for an 80 percent rule, Ford and Chrysler 70 percent, and General Motors 60 percent. Mexico and Canada wanted to keep the 50 percent requirement in the US-Canada FTA, but reluctantly accepted 60 percent. US negotiators had promised auto producers a number higher than 60 percent to prevent their opposition. While they were able to persuade Mexico to go to 65 percent, Canada remained firm at 60 percent and so the negotiators split the difference and arrived at a 62.5 percent rule.

Irwin goes on to describe how the US compromise led to an apoplectic call to the US trade negotiator from Ford's CEO, who felt betrayed by the failure to obtain the promised 65%. The case points to the central importance assigned to rules of origin, as well as the presumption that US producers would benefit from a stricter rule of origin than the one the US had settled on for NAFTA.

When President Trump's negotiators set out to replace NAFTA, one of their focal points was stricter rules of origin for the Auto industry. Eventually Canada, the US, and Mexico agreed in 2019 to replace the 1994–2020 NAFTA with a new agreement called the USMCA (in the United States). Lighthizer (2020) offered the following justification for stricter rules of origin:

> The USMCA rebalances the NAFTA to promote increased production in the United States and North America and to ensure that non-parties do not gain unwarranted benefits through the agreement. The USMCA features

[4] Anastakis (2005) provides a book-length treatment of this pioneering regional agreement.

> innovative rules of origin for automobiles and automobile parts that, once fully implemented, will create strong incentives to invest and manufacture in the United States and North America.

The new agreement devoted 39 pages in an appendix to the new rules, so we cannot do full justice to their complexity here. The following were the main ways in which the requirements for qualifying for tariff-free treatment became more difficult for the auto sector:

1. The minimum North American regional content requirement (RCR) was increased to 75% (from 62.5%).
2. A new Labor Value Content (LVC) rule requires that 40–45% of auto content be made by workers earning at least $16 per hour.
3. 70% percent of both the steel and the aluminum going into each car must originate in North America.
4. Six "super-core" parts—including engines and transmissions—must themselves comply with the 75% RCR.

The new requirements are clearly intended to discourage firms from sourcing parts from outside North America: if the vehicles currently assembled in the USMCA area with non-USMCA parts do not satisfy the new higher requirements, they will no longer qualify for duty-free imports within the USMCA area. The $16 hourly wage minimum also tilts sourcing preferences against Mexico in favor of Canada or the US. This is because either factory wages must quadruple from about $4 per hour, or the cars made with Mexican parts become non-compliant and have to pay tariffs. While this Mexico-specific feature of the USMCA RoO is important, it does not fit well within our modelling structure, so we leave further quantification of its consequences to future work. However, our model does say something about the qualitative effects of the labor-value requirement. The policy appears to be designed to lower the attractiveness of Mexico as a supplier. However, a less competitive Mexican supply sector also raises the expected costs of cars assembled elsewhere in North America. Thus, it could bring additional unintended consequences, beyond those that we quantify in this chapter.

Table 1 provides some early evidence on how the car industry is responding to the phasing in of the USMCA's stricter rules of origin. We see that in 2019 compliance with the agreement was very high, at least for those cars and light trucks shipped across the borders within North America. By 2021, the RCR had risen to 69%. The striking outcome is large drops in preference utilization for cars shipped from Mexico into the US and even larger drop for imports into Canada. RoO compliance

Table 1 – Use of preferential tariffs by US and Canada

Year:	2019	2020	2021	2019	2020	2021
RCR:	62.5%	66%	69%	62.5%	66%	69%
Importer: USA	Origin:					
Product	Canada			Mexico		
8703 (Cars)	99.2	97.9	97.7	99.4	95.2	86.8
8704 (Trucks)	97.8	93.7	94.2	100.0	99.8	99.8
Importer: Canada	Origin:					
Product	USA			Mexico		
8703 (Cars)	97.3	97.6	86.3	99.2	96.5	81.5
8704 (Trucks)	96.8	97.7	96.7	99.1	98.5	98.9

for exported Mexican trucks remains higher, in line with the much higher penalty for non-compliant trucks imported into the US: a 25% tariff.

The other major regional trade agreement, the European Union, had no need for rules of origin since it is a customs union with a common external tariff. This came to an end in 2020 with the conclusion of the negotiations creating the European Union and United Kingdom Trade and Cooperation Act (TCA). While the status of fisheries and Northern Ireland garnered more press attention, debates over rules of origin again proved to be a sticking point. A "Swiss-style" agreement would have retained better access to the EU market, but the UK government demanded that its negotiators "Give us Canada."[5]

The inevitable consequence of a Canada-style deal would be rules of origin. Predictably, based on Canada's history of negotiations with its larger trade partner, the EU wanted stricter rules than the UK. Michel Barnier, the chief EU negotiator, gave a speech in the summer of 2020 arguing "Do we really want to take a risk with rules of origin that would allow the UK to become a manufacturing hub for the EU, by allowing it to assemble materials and goods sourced all over the world, and export them to the single market as British goods: tariff-and-quota-free?"[6] The final version of the TCA specified that motor vehicles would satisfy the RoO provided that the Maximum value of Non-Originating Materials (MaxNOM) was kept below 45%. The minimum RCR for regionally-sourced parts is therefore $100-45=55\%$,

[5]"Inside the Brexit deal: the agreement and the aftermath" George Parker, Peter Foster, Sam Fleming and Jim Brunsden, *Financial Times* January 21, 2021.
[6]"What's driving the EU on rules of origin?" Jim Brunsden, *Financial Times* October 29, 2020.

more lenient than NAFTA—even before the 2020 rule changes.

Section 6 quantifies the consequences for consumers and producers of these recent changes in RoOs in North America and Europe. We also consider counterfactuals of stricter RoOs that might have been enacted. Before those numerical exercises, we need to introduce our model. To ground the model, we first describe data on sourcing of automotive parts in North America.

3. Regional parts use in NAFTA: key patterns

We use two data sources on regional parts use in North America. The first is extremely detailed data on sourcing of engines and transmissions, two of the highest value components of internal combustion engine (ICE) vehicles. The second source is data from the American Automobile labelling act which examines sourcing of all components aggregated together.

3.1. Sourcing of engines and transmissions (IHS data)

Figure 1 displays the 2018 production shares of all the main powertrain sourcing configurations, which we define as a pair of countries where the first provides the engine and the second supplies the transmission. The source of the data is the automotive consultancy IHS Market. They provide the number of units manufactured in each plant for all firms, detailed by engine and transmission source. The fill color of squares shows where the engine was produced whereas circles do the same for transmissions. Even with all non-NAFTA source countries aggregated into a single rest-of-the-world (RoW) group, there are a large number of possibilities. To keep the figure readable, we only show configurations that account for at least one percent of local production.

The main takeaway from figure 1 is the heterogeneity in sourcing patterns, even when considering just two components. The most common configurations differ across the three countries. When assembly takes place in Canada, vehicles with both engines and transmissions from the US are the most common configuration, accounting over 30% of the cars assembled there. US factories use domestic engines and transmissions for over 40% of vehicles. In Mexico, USA-USA accounts for about 10% of assembly.[7] Canadian parts are often included in the powertrain for cars assembled in Canada, but much less so in the US. Outside those two countries, Canadian engines and transmissions have negligible use.

[7] By contrast in the main manufacturing countries outside North America—Japan, Korea, and Germany—the USA-USA pairing is used for just one percent of cars.

Figure 1 – Heterogeneity in engine and transmission sourcing configurations for North American vehicles

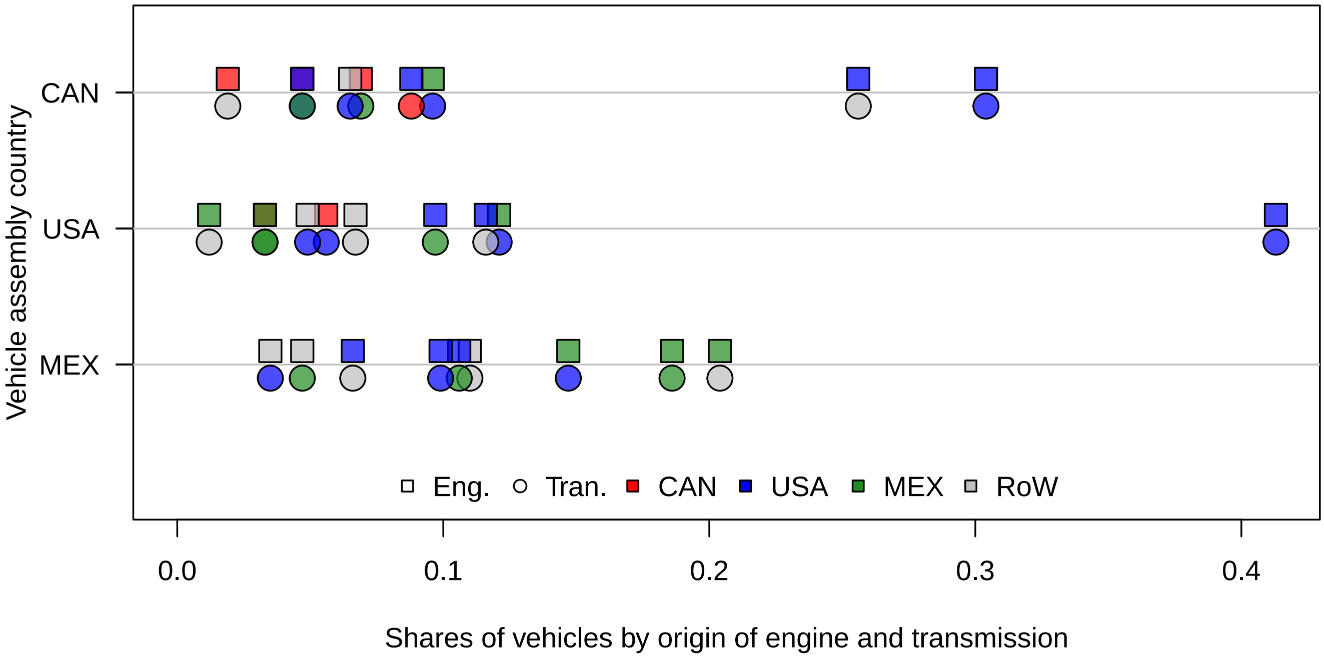

Note: The horizontal axis shows, for each assembly country (or group) the share of cars assembled in that area using various engine (squares) and transmission (circles) sourcing configurations. The assembly countries depicted on the vertical axis are Canada (CAN), the United States (USA), and Mexico (MEX). Configurations are included if they account for 1% of cars in each country.

The diversity of configurations observed for just two parts establishes the importance of allowing for heterogeneity *within* countries. This features prominently in the model described in the next section. One of the key ideas in the model is that some parts are likely to be sourced domestically even with rather lenient rules of origin. Firms would be more reluctant to bring sourcing of other parts into the region, and would do so only when compelled by a stricter RoO. One factor underlying this asymmetry could be differences in the part-specific cost of remote sourcing.

Figure 2, also based on the IHS Markit data, provides compelling evidence that remote sourcing of engines is relatively rare throughout the global vehicle industry. On the other hand, long-distance sourcing seems less costly for transmissions. Thus, in the context of our model, engines are examples of parts that firms source locally even without pressure from RoOs, whereas transmissions are the marginal part that would be added only to avoid incurring tariffs when rules are strict.

Figure 2 plots the cumulative distribution functions (CDF) of distances for engines and transmissions.[8] For every distance between an engine or transmission factory and an assembly factory, we calculate the share of all vehicles made from engines or transmissions transported less than that distance. The thinner blue and red lines in figure 2 depict these CDFs for engines and transmissions, respectively. In 2000, we see that over half of all cars are built using engines that travelled less than 238 kilometers (347 km for transmissions). Nearly 20 years later, the median distance that a transmission was shipped had almost doubled to 682 km. In contrast, the median engine was transported an even shorter distance than before.

To what extent do these observed distances simply reflect geographic clustering of plants? To answer this, we compute a benchmark CDF based on plant locations under a null hypothesis of random sourcing. That is, in this hypothetical data generating process for distances, each engine is equally likely to end up in every car. Thus, the fraction of engines from plant A travelling *d* km to plant B would be equal to plant B's share of world vehicle production. The thicker lines in panels (a) and (b) graph these CDFs in 2000 and 2018. We see that median distances under the null are vastly larger—about 7,900 km in 2000 and 7,640 km in 2018. In other words, if distance did not matter, we should see much higher shares of engines and

[8]Figure 2 applies the great circle formula to calculate the distance between engine (or transmission) factories and the final vehicle assembly factory. Since engines and transmissions are too heavy and bulky for air shipment, so road, rail, or sea distances would be more accurate. Past work finds high correlations between great circle and actual road distances within countries. For intercontinental trade, air routes diverge in a more severe way from sea routes. Thus we should expect that any measurement error is larger for long distances, but we see relatively little trade at distances over 2000km.

Figure 2 – The distribution of sourcing distances

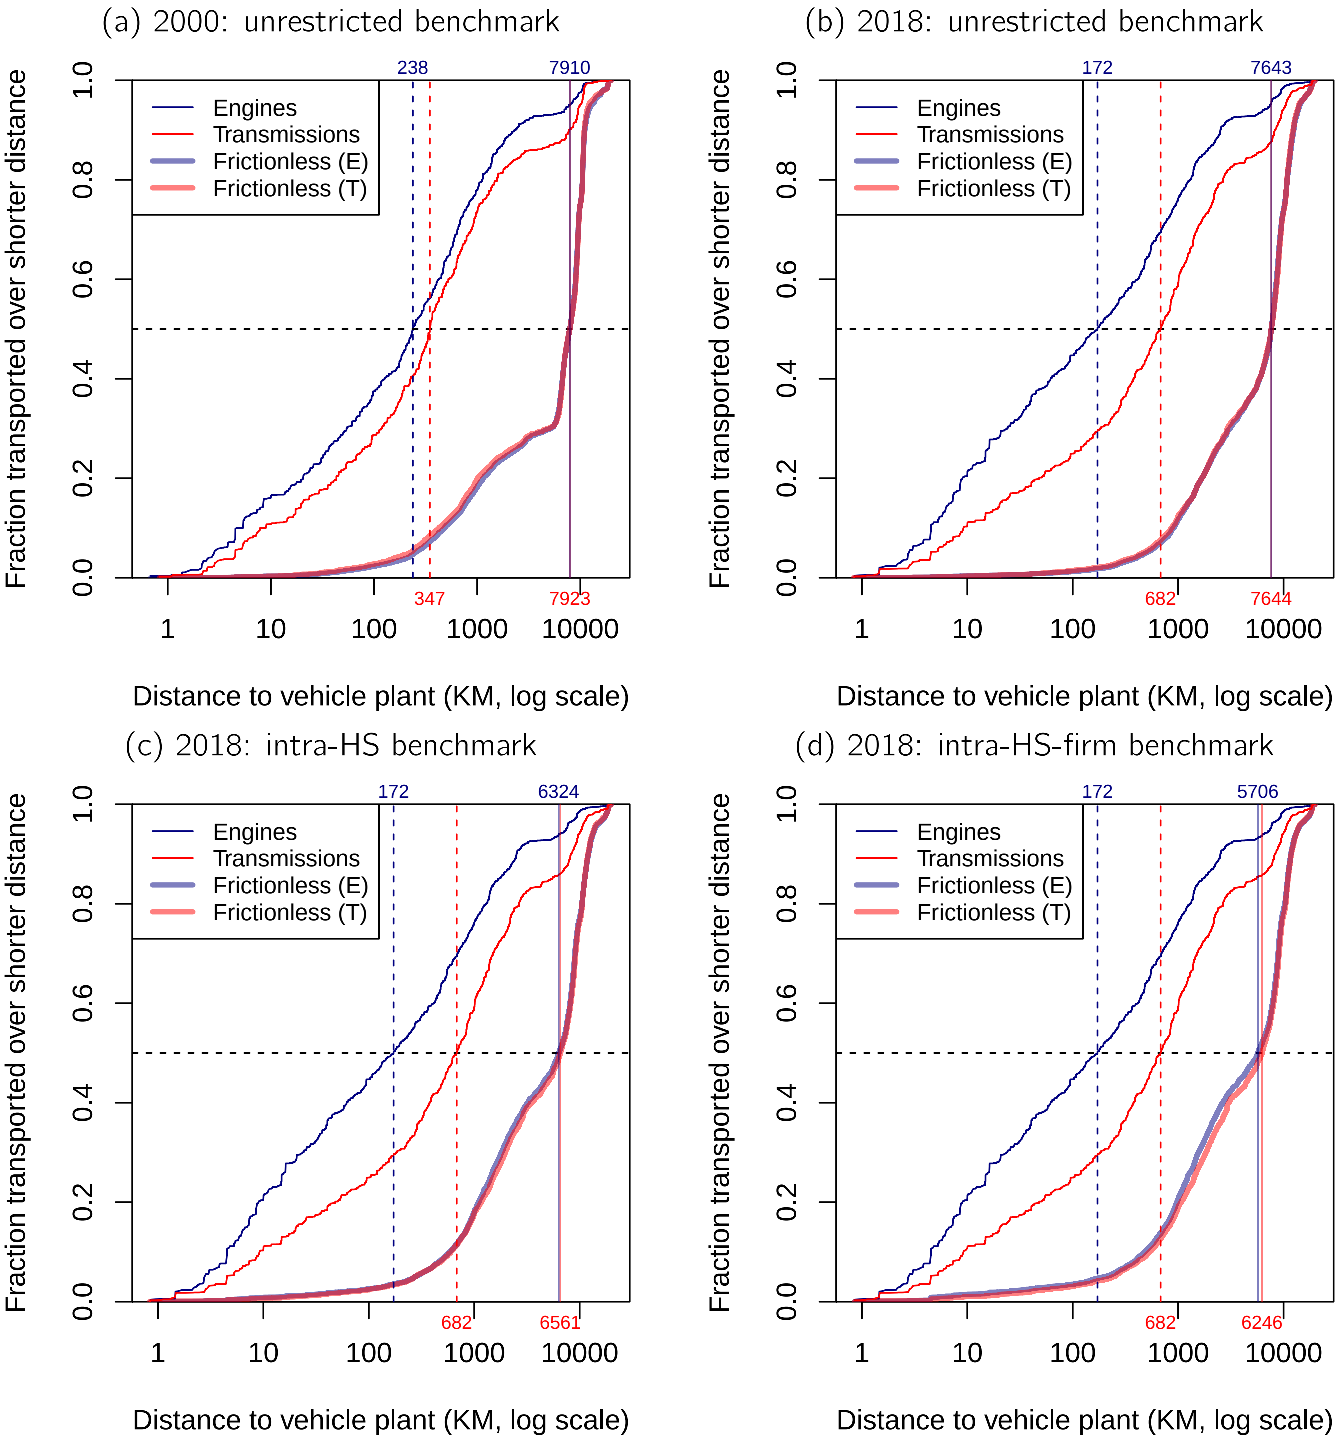

Note: Each lines graphs the fraction of engines (blue) or transmissions (red) transported by less than or equal distance from their point of manufacture to the final assembly location. The thick lines are benchmarks expected under frictionless (random) sourcing. The benchmarks in figures (a) and (b) treat engines and transmissions as homogeneously usable by any vehicle. Figures (c) and (d) respect distinctions in the HS code of each vehicle and the vertical relationships between the core parts and assemblers.

transmissions crossing oceans.

The null benchmark of panels (a) and (b) ignore some simple constraints. Automatic transmissions made in Japan will not be transported to factories in Europe to equip manual transmission cars. The relatively high displacement engines made for pickup trucks in North America will not end up in cars assembled in Japan. Panel (c) takes into account these product-compatibility constraints by recalculating the benchmark CDFs. This lowers the median benchmark distance by about 1,000 km, but obviously cannot explain the much shorter actual distances. Panel (d) constructs a benchmark that obeys additional data constraints. It takes into account that if a factory builds an engine that in reality goes to a Mazda factory, then even in the random benchmark it must still end up in Mazda factory (albeit not the same one). This rules out, among other things that it ends up in India, where Mazda has no factories. This additional element of realism in the benchmark only drives down the median by an additional 300 km (transmissions) or 600 km (engines). Evidently, the bias towards proximate sources is not something that can be eliminated by simple benchmark corrections.

Shipping heavy car parts and coordinating with distant assembly plants is costly. This implies that many parts would be sourced regionally even in the absence of rules of origin. The unconstrained regional sourcing is an important part of our model. The point to note is that these benefits of local sourcing differ, even within components of the powertrain. A more extreme contrast between parts would be between car seats and electronics: The former are almost always assembled locally while the latter almost all come from Asia. We now turn to broader evidence on the sourcing of all car parts going into cars sold in Canada and the US.

3.2. North American input cost shares (AALA data)

Our source of data regarding variation in regional cost shares is based on annual reports mandated by The American Automobile Labeling Act (AALA) of 1992. The law requires that "A label with the US/Canada content percentage and related additional information must be displayed on these vehicles up to the time of first retail sale." According to AALA, each new passenger motor vehicle must be labeled with the following information:

1. The percentage of US/Canadian equipment (parts) content
2. The name and percentage content for any countries other than the US and Canada that individually contributes 15 percent or more of the equipment content (with a maximum of two countries)

3. The countries of final assembly, engine manufacture, and transmission manufacture.

The data are available in PDF form on the AALA website.[9]. Information on component suppliers other than the US and Canada begins in 2011. The cost share data is reported by AALA at the carline level, which usually corresponds to a brand-model assembled at a specific factory. AALA often provides more detail for carlines, with information such as engine size.

Figure 3 – NAFTA regional cost share by location of production (CDFs)

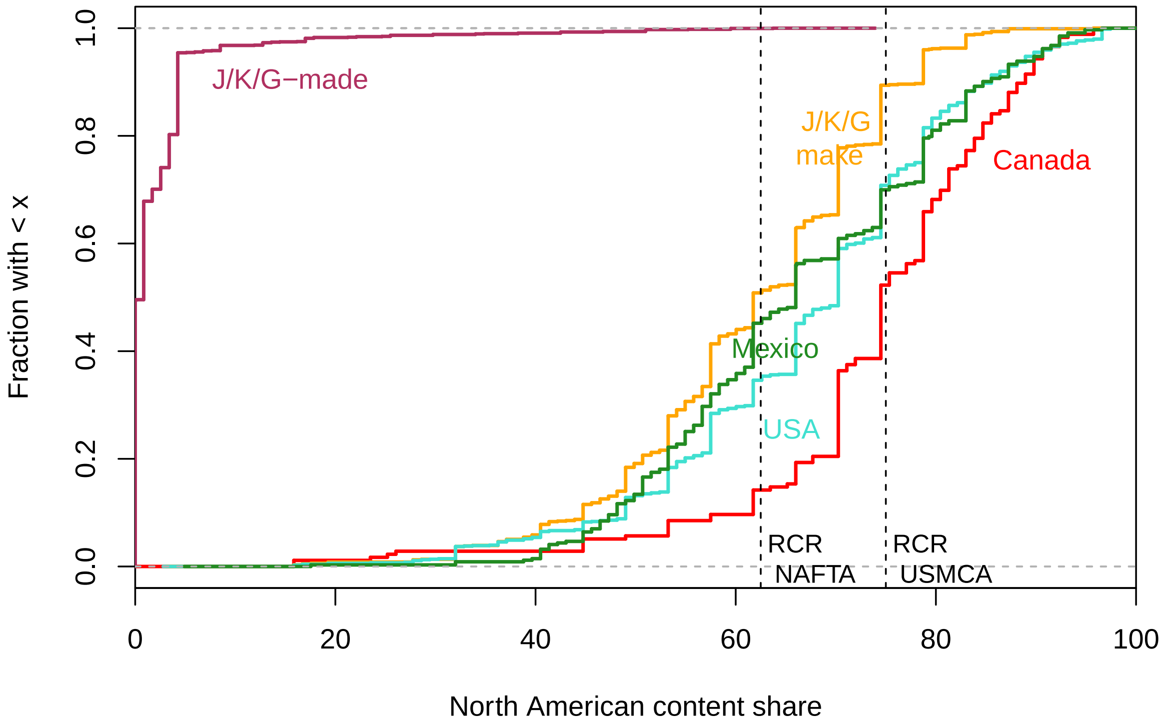

We represent the model-level AALA data as a collection of cumulative densities in figure 3. These are plotted with the original data pooled over the 2011–2020 period. We plot the CDFs separately for the cars that are the most potentially affected by the RoO, i.e. those produced in Canada, Mexico and the US. We also present separate densities for the Japanese, Korean, and German brands that are produced in NAFTA (J/K/G *make*). Finally, we also plot a density for the models sold in the US but assembled in Japan, Korea, or Germany (J/K/G *made*).

The AALA reports give estimates of the share of parts costs, not accounting for

[9] https://www.nhtsa.gov/part-583-american-automobile-labeling-act-reports

assembly costs. In order to compare those numbers to the RCR, we therefore need to add on the regional costs attributable to assembly. Figure 3 computes the overall regional cost share under the assumption that final assembly amounts to 15% of the total production cost of each regionally made car.

Four main findings emerge. The majority of carlines in each NAFTA country have cost shares that indicate compliance with the 62.5% RCR prescribed by the original NAFTA. Second, compliance is highest in Canada, lowest in Mexico, and intermediate in the US. Car brands headquartered in the three major car-producing countries outside NAFTA have lower NAFTA inputs shares even when producing in NAFTA. Finally, North American cost shares for cars assembled outside North America tend to be very small.

4. A theoretical model of parts sourcing

As we previously discussed, rules of origins (RoO) can generate competing incentives for the location of part production within a regional trade area (RTA). Those rules are intended to relocate the production of parts within the RTA; but when they are overly restrictive, the impact on regional sourcing will be reversed and part sourcing will be relocated outside the region. We now sketch a simple model based on our companion paper Head et al. (2022) that illustrates why RoOs will induce such a hump-shaped response for that regional part share. In order to focus on the sourcing decision for parts and the intuition for this hump-shaped response—which we call the Laffer curve for RoOs—we keep the location of assembly fixed. Our companion paper shows how RoOs will also impact that assembly location choice, and how overly restrictive RoOs will not only lead to lower regional part sourcing, but also induce final good producers to relocate assembly outside the region.

4.1. Model Structure

The potential for the downward sloping segment of the RoO Laffer curve, where stricter RoOs lead to reductions in the regional part share, arises when final good firms (a carline producer in our data) make sourcing decisions for many parts. Although we would technically only need a minimum of two parts to highlight this effect, we develop a model with a continuum of parts due to its analytical tractability. And it also fits well with our empirical application where car producers make sourcing decisions on a very large number of parts.

Each car part can be sourced from either within the region at one cost or outside the region, denoted Foreign, at a different cost. Each part cost for regional and

Foreign production is modeled as a stochastic draw from a Weibull distribution with parameter $\theta \geq 1$.[10] We normalize the mean cost for regional production to one. The mean cost of the Foreign-sourced parts is $\delta > 0$. This parameter varies across firms. Firms with $\delta > 1$ have a lower regional production cost for parts *on average*. As we mentioned earlier, we ignore the assembly location choice in order to focus on the part sourcing decision (regional or Foreign); and we therefore do not model the associated assembly costs until the quantification in section 6.

Free Trade (No Rules or Origin) When there are no RoOs, a firm δ decides whether to source each part from either within or outside the region based on whichever cost is lower. This is the firm's unrestricted part-sourcing choice, which we denote with a subscript U. The resulting share of regionally-sourced parts is given by the probability that the regional cost for a given part is lower than the Foreign cost. Given our distributional assumptions for the Weibull cost draws, that probability and resulting share is:

$$\chi_U(\delta) = \left(1 + \delta^{-\theta}\right)^{-1}. \tag{1}$$

Firms with higher δs have a comparative advantage in regional part production and hence source a higher share of their parts domestically. This sourcing decision then leads to a total parts cost (aggregating over both the regional and Foreign parts) of $C_U(\delta) = \chi_U(\delta)^{1/\theta}$. As we will see below, these cost differences will be inconsequential for a firm's response to a RoO, because that will only depend on how a RoO *increases* the firm's cost above this benchmark $C_U(\delta)$.

Rules of Origin A RoO mandates that firms source a minimum fraction of their parts χ_R regionally, or else it will face a Most Favored Nation (MFN) tariff rate on the final good exported within the RTA. We model this additional cost as an average tariff $\tau > 1$ incurred across all final good units produced. In the quantification in section 5, we will construct this average tariff rate based on the share of a carline's within-RTA exports relative to all its other sales. If a firm chooses to comply with the RoO and avoid the tariff, it sources progressively more expensive parts regionally (relative to foreign-sourced) until the minimum threshold is met. In our companion paper, we show how the sourcing choices to comply with a RoO χ_R are equivalent to the ones the firm would make if a tariff were imposed on foreign

[10] The parameter θ governs the variance of the cost draws. As θ increases, the variance decreases. In the limit as θ goes to infinity, the variance goes to zero and there is no variation in the cost draws around their mean.

parts (with the tariff revenue subsequently rebated back to the firm). We also describe the connections between a RoO specified as a regional part χ_R and a RoO specified as a regional cost share λ_R: a mandated minimum cost-share for regionally-produced parts. Both types of RoOs have qualitatively identical effects on regional part-sourcing because there is a monotonic relationship between χ_R and λ_R. This connection is important in the quantification because RoOs for cars in NAFTA and the EU-UK TCA are specified as cost shares.

When a binding RoO $\chi_R > \chi_U(\delta)$ is mandated, the firm's total part cost increases from $C_U(\delta)$ to:

$$C(\chi_R, \delta) \;=\; \chi_R^{\frac{1+\theta}{\theta}} + (1-\chi_R)^{\frac{1+\theta}{\theta}}\,\delta. \tag{2}$$

This represents an increase in the firm's total part cost relative to its unrestricted (lower bound) cost $C_U(\delta)$ given by the ratio

$$\tilde{C}(\chi_R, \delta) = C(\chi_R, \delta)/C_U(\delta) > 1.$$

This cost ratio captures the compliance cost penalty associated with the RoO χ_R. It is represented in the top panel of Figure 4 as a function of the RoO χ_R for three different firms. Anticipating our empirical application, we use our fitted distribution for δs across NAFTA-assembled carlines. Firm 2 has $\delta_2 = 0.12$, which is the median δ (representing a 12% average cost advantage for NAFTA-produced parts).[11] We then show two other firms (δ_1 and δ_3) that are, respectively, at the 5% and 95% percentile for that empirical distribution. For any given firm—a given δ—there is a range where its unrestricted sourcing choice $\chi_U(\delta)$ is above χ_R and therefore complies with the RoO. There is no cost associated with compliance, so $\tilde{C}(\chi_R, \delta)$ is at its lower-bound of one. We denote this case compliant-unconstrained. As the RoO χ_R rises above $\chi_U(\delta)$, compliance with the RoO entails a cost compliance penalty $\tilde{C}(\chi_R, \delta) > 1$. As anticipated, this cost penalty then increases monotonically with the RoO χ_R: compliance becomes increasingly costly as the RoO becomes more restrictive. Looking across firms, we see that, as expected, the compliance cost with a given RoO χ_R is always higher for firms with lower δ whenever they are not unconstrained: those firms have a comparative advantage in Foreign-sourced parts, so complying with a given RoO is more expensive.

4.2. Compliance

As we mentioned, a firm δ can choose not to satisfy the RoO χ_R and instead pay the average tariff τ. It will do so whenever the compliance cost is greater

[11] We also set $\theta = 4$.

Figure 4 – Compliance Cost and Sourcing Decision for 3 Firms

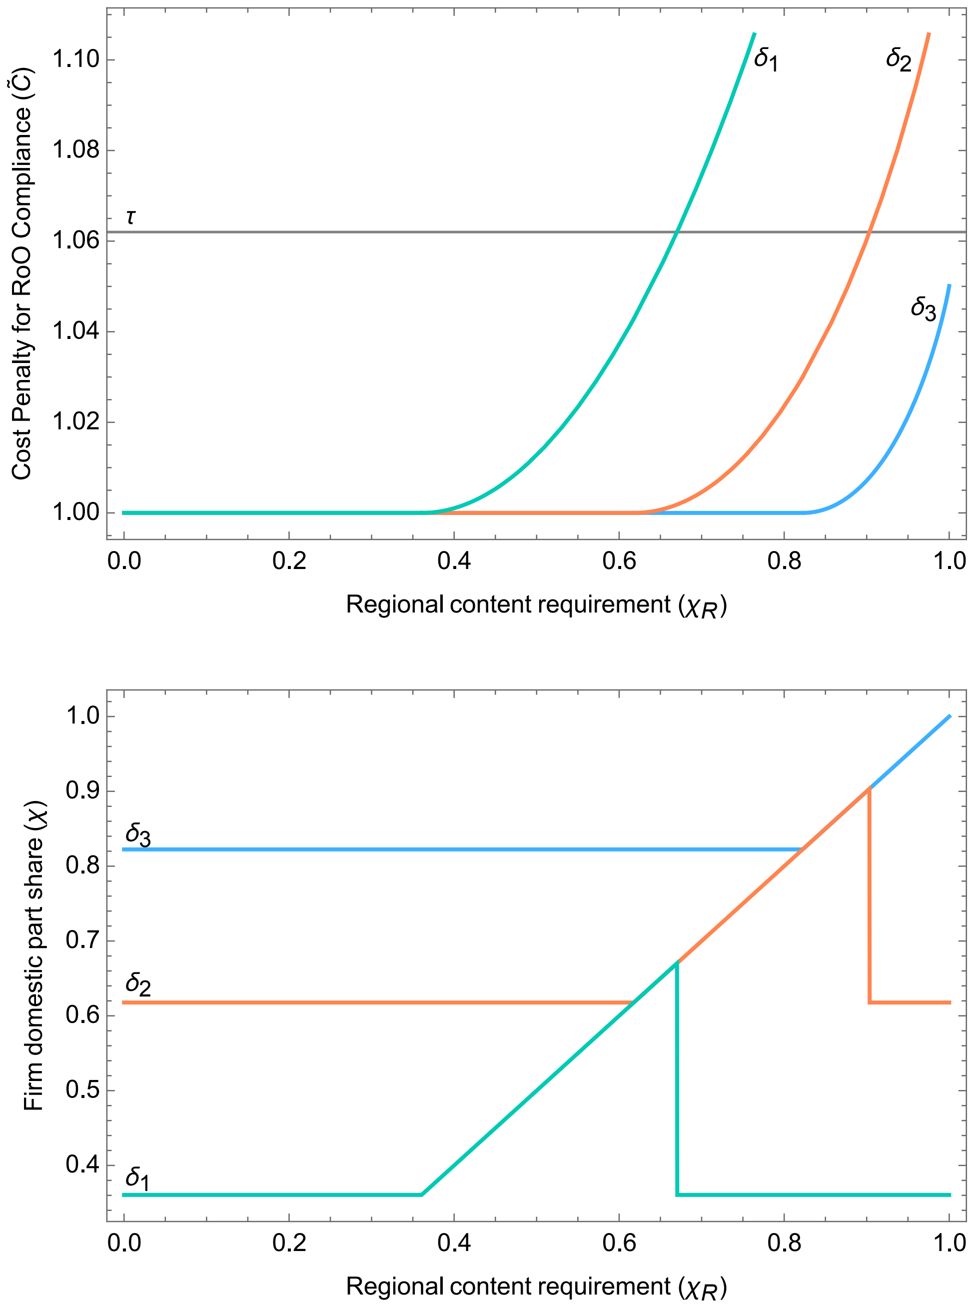

than the tariff penalty: $\tilde{C}(\chi_R, \delta) \geq \tau$. In this case, we label the firm as *non-compliant*, and it then reverts to its unconstrained part sourcing with regional share $\chi_U(\delta)$ and associated cost $C_U(\delta) = \chi_U(\delta)^{1/\theta}$. The horizontal line in the top panel of Figure 4 shows the example of a 6.2% tariff penalty. Continuing with our anticipated empirical application, this represents the non-compliance tariff that would be paid on average across all vehicles assembled in Mexico based on the empirical proportion of Mexican-assembled vehicles that are exported to its NAFTA partners, the United States and Canada and their associated MFN tariffs.

The bottom panel of Figure 4 shows the regional part share chosen by the three firms, given their compliance decision. When the RoO χ_R is low enough, all three firms are *compliant-unconstrained* and choose their unrestricted part share $\chi_U(\delta)$. This corresponds to the case of no compliance cost penalty, $\tilde{C}(\chi_R, \delta) = 1$, in the top panel. As the RoO χ_R increases, firm 1, followed by firm 2 and then firm 3 become *compliant-constrained*: The compliance cost penalty $\tilde{C}(\chi_R, \delta)$ rises above one, but remains below the tariff penalty τ. In this case, the firms choose the regional share χ_R to comply with the RoO. This is captured by the 45-degree increasing line in the bottom panel: a chosen regional share equal to the RoO. As the RoO χ_R further increases, firm 1 and then firm 2 choose non-compliance: the cost penalty is higher than the tariff penalty. In those cases, their chosen regional part-shares drop back to their initial unrestricted levels $\chi_U(\delta)$. Note that firm 3 will never choose to be non-compliant: Complying with even the most restrictive RoO of 100% is still less costly than the tariff penalty. We label firms of this type as *always-compliers*.

4.3. Laffer Curve for Rules of Origin

Setting aside those firms that are always-compliers, we see in Figure 4 that increasing a RoO from 0% to 100% will initially induce firms to increase their regional part-share—when they are compliant-constrained—but will then induce those firms to sharply reduce their part-share once the RoO rises above a threshold where the firms choose non-compliance. In our companion paper we show that this non-monotonic response, in this individual firm case an inverted-V, requires a firm-sourcing decision over multiple parts. When there is a single part, that non-monotonic sourcing response disappears: Increasing the RoO can never induce a firm to *reduce* its regional part-share. And we also show that as we smooth that inverted-V sourcing response at the firm-level over a set of firms with heterogeneous δ, then the average regional sourcing share becomes a smooth inverted-U Laffer curve. So long as we exclude the always-compliers, then the average regional part-share returns to its initial ($\chi_R = 0$) level as the RoO increases to its 100% upward bound. When we consider the full set of firms including always-compliers, then the average regional part-share remains above its initial level as the RoO increases to its upward bound.[12]

5. Simulating policy changes in the model

The model delineated in the previous section provides key qualitative insights. Most importantly, it demonstrates the unintended consequences of an overly strict set of rules of origin. When the cost of compliance is higher than the penalty for non-compliance, firms will opt into non-compliance, cutting regional input use down to their unconstrained levels. The key unanswered questions are whether recent policy changes put North America into this range of counter-productive rules. Answering this question requires us to calibrate several different dimensions of heterogeneity. We do this by finding parameter values that induce the best fit between our simulated data and the observed data for the pre-USCMCA period when the RCR was 62.5%.

When taking the model to the data, we have to take a stand on the level at which the content decision is made. While the model refers to "firms," the AALA reports show that different carlines owned by the same firm use very different shares of North American inputs. For example, the made-in-Mexico Ford Fiesta uses 80% North American parts, whereas the US-assembled Ford Mustang has 46% of its parts originating in North America. The Volkswagen Golf R, made in Germany, has only 1% of North American parts but the Golf GTI assembled in Mexico has 42%. The US-assembled VW Passat has 61% for the version with a 2.0 liter engine (made in Mexico) and just 30% for the 3.6 liter version (engine imported from Germany).[13] Thus, the data suggest that the content decision is taken in response to variation in relative costs (δ in the model) at the level of specific carlines. The actual decision-maker could be a plant manager or global headquarters. In the model it, does not matter whether the decision is centralized, because profit maximization implies that costs should be minimized for each carline. There is a single compliance decision for all the vehicles that come out of the same production line, regardless of their final destination. This assumption comes from observation in the IHS Markit data that it is extremely rare for the same carline to source a given engine or transmission from more than one country. Also, the AALA data provide single NAFTA shares for each carline.

[12]Hypothetically, if the distribution of δs is such that it is dominated by always-compliers, then it is possible for the average regional part-share to monotonically increase with the RoO. However, we show that this is not the case for NAFTA.

[13]All these percentages are cost shares from the 2019 AALA report.

It is important to simulate the model at the carline-level because the tariff penalty for non-compliance (τ in the model) varies greatly across carlines because of their different sales destinations. For example, the Ford Mustang has 2018 sales of 76,000 units in the US. These cars will not pay any tariff penalties for non-compliance with USMCA rules, nor will the roughly 12,000 units headed to Australia and China.[14] Only the 7,600 Mustangs sold in Canada and the 1,900 sold in Mexico will face MFN tariffs as a penalty for non-compliance with the USMCA RoO. The situation of the Ford Fiesta made in Mexico is very different. It sends the lion's share of its total production (66,000 cars) to its USMCA partners: to the US (52,000 cars) and Canada (1,200 cars). Meanwhile, only 4,500 Fiestas stay in Mexico. The overwhelming dominance of export sales to NAFTA partners gives the Fiesta plant very strong compliance incentives, as compared to the Mustang. We capture this important source of heterogeneity by using the IHS Markit data to compute tariff penalties for every carline.

The tariff penalty tends to be much lower than the MFN tariffs because large shares of output in the regional plants of a carmaker tend to stay within the country of production or go to markets outside the region (as in the Mustang example). Table 2 provides more granular information for the twenty largest tariff penalties.

We use a simulation of our model to estimate the underlying heterogeneity parameters. The idea is that carlines receive their comparative advantage "draws" according to a particular "guess" for the mean and standard deviation of δ. At the same time, they draw a parameter determining the importance of assembly costs for that carline. Then the simulated carlines each decide whether to comply with a content requirement of 62.5%. Depending on the assembly cost share, this RCR converts to a particular parts costs share (λ_R in the model), which in turn converts to an implied share of regional parts (χ_R in the model). If compliance is too costly relative to the tariff penalty, then the carline selects its unconstrained cost minimizing North American parts share. The result is a vector of parts costs shares emerging from the simulated model. Recognizing that the model is an approximation, and the data reporting in AALA is far from perfect, the simulation builds in random measurement error.[15] The result is a simulation-based distribution of North American parts shares, which we compare to the actual distribution from

[14] The tariffs China imposes on US exports do not depend on their North American content.

[15] Among 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.

Table 2 – Top tariff penalties for USMCA carlines in 2018

Brand	Model	Assembly country	Tariff penalty	sh. rest of RTA
Chevrolet	Silverado	Mexico	1.23	0.96
Toyota	Tacoma	Mexico	1.22	0.97
Nissan	NV200	Mexico	1.22	0.99
Ram	2500/3500	Mexico	1.21	0.94
Ram	ProMaster	Mexico	1.20	0.92
GMC	Sierra	Mexico	1.20	0.99
Ram	1500	Mexico	1.19	0.96
GMC	Sierra	Canada	1.18	0.80
Mercedes-Benz	Sprinter	United States	1.09	0.73
Chevrolet	Silverado	Canada	1.05	0.29
Volkswagen	Golf SportWagen	Mexico	1.03	0.96
Chevrolet	Cruze	Mexico	1.03	0.93
Nissan	Note	Mexico	1.03	0.86
Volkswagen	Golf	Mexico	1.03	0.81
GMC	Terrain	Mexico	1.03	0.98
Toyota	Corolla	Canada	1.03	0.85
Infiniti	QX50	Mexico	1.03	0.93
Buick	Regal	Canada	1.03	0.97
Dodge	Journey	Mexico	1.03	0.94
Dodge	Charger	Canada	1.03	0.89

Note: Head et al. (2022) provides the formula used to compute the carline-level tariff penalty in a way that takes into account market share changes in response to tariff changes.

the AALA reports. We quantify the discrepancy in terms of the sum of squared deviations between model and data. The algorithm then repeats the procedure for a large grid of different guesses for the parameters, selecting the ones that achieve the best fit between simulation and observation. Head et al. (2022) provides a more formal description of this procedure for estimating the model parameters.

The estimated parameters allow the distribution of the simulated carlines to tightly fit the distribution of North American content reported by AALA. To provide external validation for the quantified version of the model, we follow the common practice of considering a feature in the data that was not part of the original moment-matching exercise. For this purpose, we compare the implied RoO compliance rates (also referred to as preference utilization rates) for auto trade (HS 8703) to those that emerge from the simulation based on the calibration described above. As shown in table 1, the true rate of preference utilization for US-made cars entering Canada is 97% in 2019 (before the change in the regional content requirement in 2020). The calibrated model obtains a rate of 92%. Thus, our model is able to closely mirror the distribution of North American content rates at the carline level and also match reasonably well the RoO satisfaction rates observed for aggregate trade flows within North America.

After obtaining the best-fit values, we can solve the model for *any* potential RCR. This requires computing how each individual carline will respond to a stricter RCR. Depending on their parameter draws, they might increase regional parts shares just enough to match the new requirement, or they might opt into non-compliance. Based on this decision, the change in costs (from increasing regional content in response to a stricter rule) or the tariff penalties (from opting not to comply with a stricter rule) will reallocate market share towards foreign carlines, as well as those domestic carlines that were not complying before the stricter rule. In computing the changes in this step, we take advantage of the aggregation properties of the constant elasticity of substitution (CES) demand system. This provides an exact aggregation for the resulting changes in the price index and employment in the next section.

6. Quantification of the impact of RoO changes

In this section we use our model, with parameters chosen to fit the distribution of regional content by North American carlines, to quantify the effects of two recent changes in RoOs. The first is the tightening of RoOs for North American vehicle trade, which was one of the most salient features of the USMCA. The second is the application of rules of origin to UK-EU trade, required by Britain's exit from the customs union in the final Brexit deal.

We evaluate changes in the strictness of the RoO, as measured by changes in the RCR for the enacted policies. We also consider alternative RCR levels that *might* have been chosen. For each policy change, we report outcomes for groups of carlines based on their compliance decisions before and after the RoO changes. For example, the first group in each table is the one for carlines that comply exactly with the old RoO but then decline to comply with the new RoO. The first numerical column shows the share of carlines in each group (in percent). The last four columns report the simulated changes induced by the change in the RCR. These outcome variables comprise the percentage changes in the price index, the group's market share, the weighted average regional parts share, and employment.

6.1. USMCA

Table 3 – Increase in RCR from NAFTA (62.5%) to USMCA (75%)

Compliance status under:		Share of	Percent changes in			
			Price	Mkt.	Parts	Parts
NAFTA (RCR=62.5%)	USMCA (RCR=75%)	carlines		share	share	Emp.
Comply-constrained	Non-compliant	16.90	0.57	-1.05	-10.40	-11.85
Comply-unconstrained	Non-compliant	7.10	0.27	-0.16	0.02	-0.40
Comply-constrained	Comply-constrained	7.30	1.32	-3.23	20.97	15.53
Comply-unconstrained	Comply-constrained	34.50	0.21	0.00	8.26	8.03
Non-compliant	Non-compliant	8.30	0.00	0.65	0.00	0.65
Comply-unconstrained	Comply-unconstrained	25.80	0.00	0.65	0.00	0.65
All	All	100.00	0.28	-0.20	2.80	2.30

Notes: "Share of carlines" refers to the percentage of all domestic carlines in the corresponding status tuple. "Parts share" is a quantity-weighted average of the shares of parts from NAFTA origins across regionally assembled carlines. "Parts Emp." is employment in parts manufacture for domestically assembled vehicles.

Table 3 describes the simulated outcomes for the USMCA increase in the RCR from 62.5% to 75%. According to the calibrated model, just over a third of carlines switch from complying unconstrained to complying at the minimum required level of 75%. These carlines will increase their regional parts shares by about 8%. The increase in average costs for the group is just one fifth of a percent. There is no discernible reduction in market share for this group and its employment rises by almost the same amount as its average parts shares. Greater employment gains are recorded by the 7.3% of carlines that were just complying at 62.5% and raise their regional content up to 75%. These carlines increase their parts shares (X) by 21%, slightly more than the overall cost change of $0.75/0.625 - 1 = 20\%$. The implied rise in employment is just under 16%. The dampening comes from the three percent market share reduction for this group, which itself follows from their 1.32% rise in their average price.

The increase in employment for the constrained compliers is mostly offset by a reduction in employment by carlines that stop complying, once faced with the 75% RCR. The overall employment gain is just 2.3%, much lower than the naive expectation of 20% ($0.75/0.625 = 1.2$) that would follow from assuming that all carlines mechanically comply with the RoO. While the employment gains are modest, so are the price increases faced by consumers: the price index for regionally assembled cars rises by just 0.28%. As predicted by the convex cost curves shown in figure 4, there will be a higher cost of further rises in the RCR.

Table 4 – Increase in RCR from USMCA (75%) to US negotiating point (85%)

Compliance status under:			Percent changes in			
USMCA (RCR = 75%)	US ask (RCR = 85%)	Share of carlines	Price	Mkt. share	Parts share	Parts Emp.
Comply-constrained	Non-compliant	28.90	0.54	-0.75	-10.69	-11.84
Comply-unconstrained	Non-compliant	4.30	0.24	0.13	0.01	-0.10
Comply-constrained	Comply-constrained	12.90	1.38	-3.19	14.50	9.34
Comply-unconstrained	Comply-constrained	18.00	0.25	0.13	7.15	7.02
Non-compliant	Non-compliant	32.40	0.00	0.86	0.00	0.86
Comply-unconstrained	Comply-unconstrained	3.40	0.00	0.86	0.00	0.86
All	All	100.00	0.39	-0.29	0.07	-0.60

Notes: “Share of carlines” refers to the percentage of all domestic carlines in the corresponding status tuple. “Parts share" is a quantity-weighted average of the shares of parts from NAFTA origins across regionally assembled carlines. “Parts Emp.” is employment in parts manufacture for domestically assembled vehicles.

Table 4 reports the results of a counterfactual rise in the RCR from 75% to 85% (the original US ask during the USMCA negotiations). The last row of the right-most column gives an interesting message for policy. It shows that had the US succeeded in negotiating an 85% RCR, this would have *reduced* employment in the parts industry. 85% is on the wrong side of the Laffer curve for employment, although it is approximately the peak for the regional parts share. Compared to the move from 62.5% to 75%, the further ten percentage point (ppt) increase in the RCR causes the share of carlines dropping out of compliance to rise to 29%. Those carlines, whose average tariff penalty is just 1.2%, reduce their regional parts by nearly 11%. By contrast, only 13% of carlines decide to remain compliant with the 85% RCR. Those mainly consist of light trucks (as we see in table 2) who face a much larger average tariff penalty of 7.8%.

The negative result of the 85% RCR for employment in the parts sector, as opposed to the slight positive change for the parts share, comes from the demand side. The carlines that are constrained compliant with the RoO at both levels see their market shares fall by 3.2%. This means that even though their sourcing pattern is using 14.5% more regional parts, substitution away from the more expensive compliant cars limits employment gains to just 9.3%. Consumer price increases from the stricter RoOs remain modest at 0.4%.

Table 5 – Increase in RCR from 75% to 100% regional content

Compliance status under:		Share of	Price	Percent changes in Mkt.	Parts	Parts
USMCA (RCR = 75%)	RCR = 100%	carlines		share	share	Emp.
Comply-constrained	Non-compliant	38.90	1.06	-0.97	-9.80	-11.61
Comply-unconstrained	Non-compliant	24.40	0.90	-0.47	-0.01	-1.37
Comply-constrained	Comply-constrained	3.00	11.09	-25.44	45.55	-2.31
Comply-unconstrained	Comply-constrained	1.40	5.79	-13.65	32.24	7.94
Non-compliant	Non-compliant	32.40	0.00	2.22	0.00	2.22
All	All	100.00	0.98	-0.72	-3.04	-4.67

Notes: "Share of carlines" refers to the percentage of all domestic carlines in the corresponding status tuple. "Parts share" is a quantity-weighted average of the shares of parts from NAFTA origins across regionally assembled carlines. "Parts Emp." is employment in parts manufacture for domestically assembled vehicles.

The Inflation Reduction Act (IRA) of 2022 included a $7500 subsidy to consumers who purchase electric vehicles (EV). It also required that by 2029, in order to receive the subsidy, the EV would need a battery whose components were 100% made in North America (or other trade agreement partners). It was reported that no EV currently on the market uses batteries that comply with that requirement.[16]

[16] *The Verge*, August 8, 2022

This extreme content rule motivated us to consider an equally extreme revision to NAFTA: going to a 100% RCR (from the current USMCA level). Our parameter estimates imply that only 4.4% of carlines would comply with this policy and their prices would rise by 11% if already constrained at the 75%, or by 6% if newly constrained. For the other 95% of carlines that would stop complying with the RoO, prices would rise by about 1% (except for one third who were already non-compliant). The bottom line number is that a policy feature, ostensibly designed to be pro-employment, would actually reduce employment by almost five percent in the parts industry.

The phase-in of the 100% content rule for batteries is a feature of the IRA that we highlight because it relates to our model. The Senate actually voted down a motion (by a Republican who opposed the overall legislation) to implement the 100% rule immediately rather than start in 2024 with a 40% requirement. This suggests that a goal of the policy is to induce relocation of the production of battery inputs to North America over the next five years. Currently, China's share of world refining for minerals used in batteries is 59% (lithium) and 75% (cobalt).[17] Our model does not consider plant location decisions by components suppliers. In principle, this might bolster the case for stricter RoOs. However, opening up the possibility of plant relocation can also dramatically worsen the employment effects of stricter RoOs, as we show in a model extension developed in Head et al. (2022). Knowing that they will not comply with the RoO erodes the firm's rationale for local assembly. Firms that decide to relocate outside the region not only reduce assembly jobs; due to high trade costs on intermediate inputs, they sharply reduce their use of inputs from the region they exited. Recall that figure 3 shows that Japanese and German makers use far lower shares of North American inputs in their cars assembled outside North America.

Table 6 compares the old NAFTA RoO to a hypothetical situation without any RoO. This could be interpreted as a North American customs union. This case presents the most straightforward set of outcomes since all carlines are initially unconstrained. Roughly two thirds remain unconstrained with the 62.5% RoO, reflecting the inherent desirability of local sourcing (to avoid transport costs). Moving from no content requirement to 62.5% leads to a 2% increase in employment whereas prices and market shares hardly change. The rise in employment comes almost entirely from a quarter of the carlines moving from unconstrained choices to using higher North American content as a result of the 62.5% rule becoming binding. Those carlines collectively increase production (and hence jobs) by 10% with very little in terms of offsetting effects since only 8% of carlines begin to pay

[17] *Business Insider* August 10, 2022

Table 6 – Changes due to imposing the NAFTA content requirement

Compliance status under:			Percent changes in			
No RoO (RCR = 0%)	NAFTA (RCR = 62.5%)	Share of carlines	Price	Mkt. share	Parts share	Parts Emp.
Comply-unconstrained	Non-compliant	8.30	0.51	-1.27	0.15	-1.62
Comply-unconstrained	Comply-constrained	24.30	0.25	-0.52	10.62	9.76
Comply-unconstrained	Comply-unconstrained	67.40	0.00	0.23	0.00	0.23
All	All	100.00	0.10	-0.08	2.21	2.03

Notes: "Share of carlines" refers to the percentage of all domestic carlines in the corresponding status tuple. "Parts share" is a quantity-weighted average of the shares of parts from NAFTA origins across regionally assembled carlines. "Parts Emp." is employment in parts manufacture for domestically assembled vehicles.

tariffs.

6.2. Brexit and the UK-EU TCA

We now apply the parameters estimated for the North American data to consider the impact of the new rules of origin brought in by the post-Brexit trading arrangement between the UK and the remaining 27 EU members. The reason we do not re-estimate the parameters is that the AALA data contains only those cars sold in the US and therefore omits many of the mass-market cars in Europe.[18] Also, the coverage of country-level costs outside Canada and the US has many omissions due to the 15% reporting threshold. The parameters estimated for North American carlines are still relevant for counterfactuals in Europe. This is because the mean δ reflects high transport costs for parts seen worldwide (see figure 2). Moreover, the standard deviation of δ reflects cost heterogeneity across carlines based on access to regional or third-country parts suppliers. Thus, in North America there are substantial differences in the geographic structure of supply chains for the "Big 3" US and the Japanese producers on one hand—who have developed their North American supply chains over decades—and the German producers on the other hand who have only recently entered the North American market. Similar differences are at work in Europe.

There are two notable differences between our post-Brexit simulations and those we conduct for the USMCA. First, the TCA imposes a RoO of just 55%, compared to the USMCA's 75%. Going in the other direction, the EU and UK tariffs on

[18] Renault, Peugeot, Seat, and Skoda are examples of popular brands in Europe that are not offered in the US.

Figure 5 – The distribution of the tariff penalty for the UK/EU TCA

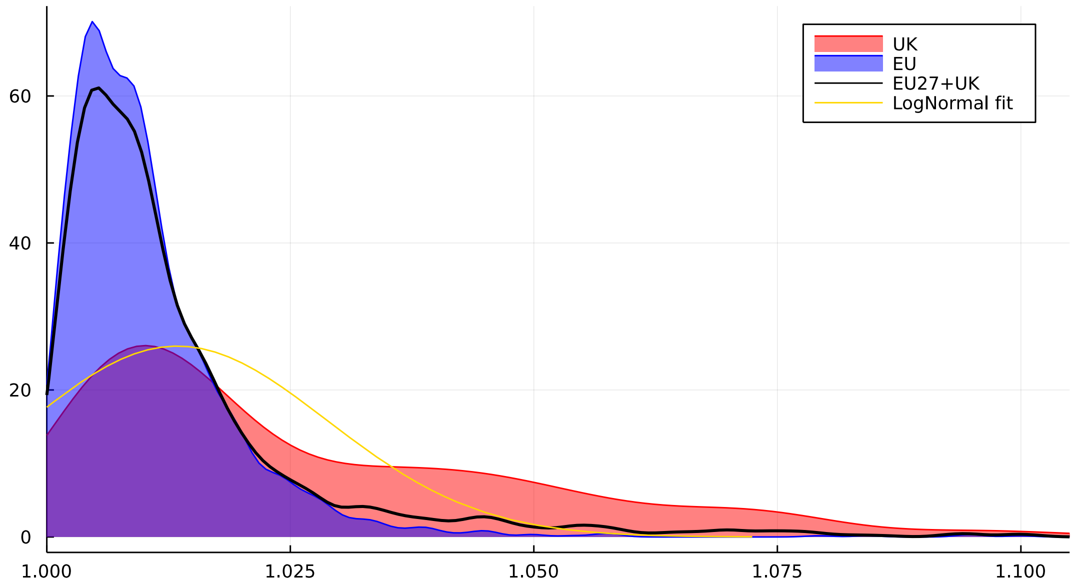

non-compliant cars are 10%, four times the 2.5% charged in the US. However, the tariff penalty does not just depend on the MFN tariff, but also on the destination of export sales. As seen in figure 5, the shipment-weighted tariff penalty has a mode that is *much* lower than 10%. Reflecting its smaller market size, cars assembled in the UK face a longer, thicker tail of high tariff penalties than cars assembled in the EU27. Eight of the top 10 tariff penalties shown in table 7 are for UK-made carlines with two Toyota models so strongly oriented towards the continent that their effective tariff penalties of 8% are very close to the MFN tariff.

Table 8 considers the impact of moving from a customs union with regional content requirement to a free trade agreement with an RCR of 55%.[19] The first striking point is that 85% of carlines in the UK and EU27 remain unconstrained under the RCR of 55%. This is because the fraction of comply-unconstrained depends only on the RCR and the carline-specific parameters, which are drawn from the same distribution for both economies. What differs in the UK/EU simulation are the tariff penalties, but they only influence the decision of whether to just comply (constrained) or not comply. Recall that the tariff penalty is generally higher in the UK. Hence, we see that the EU27 assembles more than twice the UK fraction of non-compliant carlines. In both countries, the cost increases from non-compliance are high enough that they more than offset the increase in parts share and thus

[19]There are some complexities in the UK-EU TCA as regards electric vehicles.

Table 7 – Top tariff penalty indexes for UK/EU TCA in 2018

Brand	Model	Assembly country	tariff penalty	sh. rest of TCA
Toyota	Avensis	United Kingdom	1.08	0.86
Toyota	Auris	United Kingdom	1.08	0.81
Opel	Astra	United Kingdom	1.08	0.81
Opel	Vivaro	United Kingdom	1.08	0.63
Nissan	Qashqai	United Kingdom	1.06	0.61
Honda	CR-V	United Kingdom	1.05	0.59
Volkswagen	Scirocco	Portugal	1.05	0.56
Nissan	Juke	United Kingdom	1.05	0.56
Nissan	Leaf	United Kingdom	1.05	0.54
Audi	A1	Spain	1.05	0.51
Opel	Mokka	Spain	1.04	0.49
Mini	Clubman	United Kingdom	1.04	0.41
Mini	Mini	United Kingdom	1.03	0.38
Land Rover	Range Rover Evoque	United Kingdom	1.03	0.36
Nissan	Navara	Spain	1.03	0.26
Ford	Fiesta	Germany	1.03	0.33
Opel	Corsa	Germany	1.03	0.32
Audi	TT	Hungary	1.03	0.29
Jaguar	F-Type	United Kingdom	1.02	0.27
Jaguar	E-PACE	Austria	1.02	0.27

Notes: Head et al. (2022) provides the formula used to compute the carline-level tariff penalty. The share rest of TCA is EU 27 sales divided by total sales for UK-assembled cars and UK sales divided by total sales for EU-assembled cars.

Table 8 – UK/EU TCA adopts a 55% RCR, replacing customs union

Compliance status under: No RoO (RCR = 0%)	TCA (RCR = 55%)	Share of carlines	Price	Percent changes in Mkt. share	Parts share	Parts Emp.
United Kingdom:						
Comply-unconstrained	Non-compliant	0.90	0.77	-2.17	0.30	-2.63
Comply-unconstrained	Comply-constrained	14.10	0.28	-0.73	12.70	11.56
Comply-unconstrained	Comply-unconstrained	85.00	0.00	0.10	0.00	0.10
All	All	100.00	0.05	-0.03	1.30	1.22
European Union at 27:						
Comply-unconstrained	Non-compliant	2.40	0.55	-1.56	0.17	-1.93
Comply-unconstrained	Comply-constrained	12.60	0.20	-0.52	10.72	9.92
Comply-unconstrained	Comply-unconstrained	85.00	0.00	0.09	-0.00	0.09
All	All	100.00	0.04	-0.03	1.01	0.94

Notes: Same parameters as NAFTA counterfactuals but different distribution of the tariff penalty.

lead to falling employment. Nevertheless, employment gains among the 13 (EU27) or 14 (UK) percent of carlines that comply at the 55% level are large enough to produce a one percent increase in parts employment.

The results in Table 9 indicate that further employment gains were available if that had been the object of the TCA negotiators. Although roughly 10% of the carlines that were constrained at 55% would opt into paying MFN duties at an RCR of 75%, their employment losses would not be severe enough to offset the rising employment of carlines that become or stay exactly compliant. With its larger tariff penalty, the UK sees the biggest gains (7%) while the EU27 has gains of just over 3%. It is worth emphasizing that the naive calculation based upon the ratio of RCRs (0.75/0.55) would imply a 36% increase.

7. Policy implications and discussion

The USMCA was welcomed by the chief lobbyist for Canadian auto parts manufacturers, Flavio Volpe. In an interview that agreement as: "That deal [USMCA]... is the best single positive hit for supplier business across North America in the history of the auto business. We think there's going to be 25% more in absolute volume

Table 9 – Changes due to UK/EU TCA moving to a USMCA 75% RCR

Compliance status under: TCA (RCR = 55%)	Alt. TCA (RCR = 75%)	Share of carlines	Percent changes in: Price	Mkt. share	Parts share	Parts Emp.
United Kingdom:						
Comply-constrained	Non-compliant	8.40	1.47	-3.06	-13.29	-17.16
Comply-unconstrained	Non-compliant	6.40	0.64	-0.62	0.08	-1.17
Comply-constrained	Comply-constrained	5.80	2.73	-6.58	38.14	25.62
Comply-unconstrained	Comply-constrained	52.90	0.49	-0.18	13.18	12.43
Non-compliant	Non-compliant	0.90	0.00	1.29	0.00	1.29
Comply-unconstrained	Comply-unconstrained	25.70	0.00	1.29	-0.00	1.29
All	All	100.00	0.57	-0.42	7.91	6.84
European Union at 27:						
Comply-constrained	Non-compliant	12.00	0.87	-1.78	-9.89	-12.26
Comply-unconstrained	Non-compliant	17.70	0.53	-0.77	0.06	-1.23
Comply-constrained	Comply-constrained	0.70	2.38	-6.04	36.52	25.29
Comply-unconstrained	Comply-constrained	41.70	0.35	-0.23	10.57	9.94
Non-compliant	Non-compliant	2.40	0.00	0.81	0.00	0.81
Comply-unconstrained	Comply-unconstrained	25.60	0.00	0.81	0.00	0.81
All	All	100.00	0.36	-0.26	3.89	3.26

Notes: Same parameters as NAFTA counterfactuals but different distribution of the tariff penalty.

bought from local suppliers." The head of the Mexican auto parts industry association predicted a ten percent increase in production in Mexico's part sector.[20] In contrast, the calibrated version of our model implies a much smaller effect of 2.3% (Table 3, bottom row).

What is it about our model that implies much lower employment gains from RoO increases than naive calculations? The key point is that complying with a strict rule of origin is a choice. The benefit is preferential tariff access to the other North American markets. However, so long as the US maintains its 2.5% MFN tariff on finished cars, this is not a huge penalty. Moreover, some German factories in the US may care far more about their sales in other markets—such as China, for example—than they do about losing sales in Mexico or Canada. If bringing transmission sourcing to North America will add to the costs and make the vehicle non-competitive in China, the firm might prefer not to comply on sales to Mexico or Canada and then source engines from Europe as well if the only reason they had only sourced locally was to comply with the old NAFTA rules.

The results from our quantification suggest that the old NAFTA rule and the current TCA rule are both under the parts employment-maximizing levels. However, the original Trump administration demand of 85% would have been counterproductive even from a purely protectionist standpoint. Our results also suggest the 100% content requirements for batteries for EVs are likely to lower employment while significantly raising the costs of EV adoption.

References

Anastakis, D. (2005). *Auto Pact: Creating a Borderless North American auto industry 1960–1971*. University of Toronto Press.

Felbermayr, G., Teti, F., and Yalcin, E. (2019). Rules of origin and the profitability of trade deflection. *Journal of International Economics*, 121:1032–48.

Grossman, G. M. (1981). The theory of domestic content protection and content preference. *The Quarterly Journal of Economics*, 96(4):583–603.

Head, K., Mayer, T., and Melitz, M. (2022). The Laffer curve for rules of origin. Manuscript.

Husisian, G., Gomez-Strozzi, A., and Alvarez, A. (2018). International trade: A new dawn for North American trade. Technical report, Foley LLP.

Irwin, D. A. (2017). *Clashing Over Commerce*. University of Chicago Press.

[20] Reuters, October 1, 2018

Lighthizer, R. E. (2020). *2020 Trade Policy Agenda and 2019 Annual Report*. Office of the US Trade Representative, Washington, DC.