



HAL
open science

Brands in motion: How frictions shape multinational production

Keith Head, Thierry Mayer

► **To cite this version:**

Keith Head, Thierry Mayer. Brands in motion: How frictions shape multinational production. American Economic Review, 2019, 109 (9), pp.3073 - 3124. 10.1257/aer.20161345 . hal-03391876

HAL Id: hal-03391876

<https://sciencespo.hal.science/hal-03391876v1>

Submitted on 21 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Brands in motion: How frictions shape multinational production*

Keith Head[†]

Thierry Mayer[‡]

January 10, 2019

Revision 2

Abstract

Following the 2016 Leave vote in the referendum on UK membership in the EU and the election of Donald Trump, trade agreements have entered a period of great instability. To predict the impact of possible disruptions to existing arrangements requires counterfactual analysis that takes into account the complex set of factors influencing the production and marketing strategies of multinational corporations. We estimate a model of multinational decision-making in the car industry. This model predicts the production reallocation and consumer surplus consequences of changes in tariffs and non-tariff barriers induced by US-led protectionism, Brexit, Trans-Pacific and Trans-Atlantic integration agreements.

*This research has received funding from the European Research Council (ERC) under the Grant Agreement No. 313522. Participants at presentations at Banque de France, Berkeley, Bilkent, Bocconi, CREI (UPF), EITI, ERWIT, Harvard-MIT trade seminar, HKU, HKUST, Kiel, Princeton, Society of Economic Dynamics, UBC, University of Toronto, and Warwick provided insightful comments. Suggestions by Kerem Coşar, Swati Dhingra, Gene Grossman, Beata Javorcik, Jan de Loecker, Esteban Rossi-Hansberg, and David Weinstein led to significant modifications to earlier versions of the paper. Comments by the editor and three anonymous referees contributed to major revisions. Jules Hugot, David Rodrigues, Eve Sihra, Xiaonan Sun, and Chenying Yang performed valuable research assistance.

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

[‡]Sciences Po, Banque de France, CEPII, and CEPR, thierry.mayer@sciencespo.fr

1 Introduction

After decades in which free trade agreements proliferated and deepened in scope, 2016 appeared to mark a major turning point. The Leave vote in the UK referendum on EU membership and the election of Donald Trump brought long-standing integration arrangements to the brink of disruption. Pure trade models are ill-equipped for predicting the outcomes of regional dis-integration because they omit an increasingly important feature of the world economy: multinational production (MP). The foreign affiliate structures of multinational corporations (MNCs) complicate matters because they introduce new sets of bilateral relationships. In addition to the origin-destination goods flows of standard trade models, MP models feature interactions between headquarters and subsidiary locations. MNCs must decide which of their network of production facilities will serve each market. Furthermore, because MNCs are typically multiproduct firms, they face decisions over which subset of varieties to offer in each of the markets where they choose to operate distribution facilities. Each of these decisions is likely to be influenced by distinct bilateral frictions.

Data limitations present a major challenge in estimating an economy-wide model of MP that encompasses these decisions and the corresponding frictions. We therefore study a single industry, cars, where multinational production is prevalent: Multinational brands, those produced in more than one country, account for 99.5% of cars sold in the OECD. It is also a sector where firm-level trade flows are available for all the main producing and consuming nations. This allows us to estimate the impacts of trade integration based on variation in tariffs on final cars and parts as well as the presence of integration agreements that go well beyond tariff cuts. The estimated model predicts the consequences for producers and consumers of the shocks to trade policy that are currently being debated: US threats to its NAFTA partners and other major producers to massively raise tariffs under Section 232 national security provisions, Brexit, Trans-Pacific and Trans-Atlantic integration agreements. This paper offers the first quantitative assessment of such proposals that takes into account the micro-structure of multinational production.

To meet our goals of estimating structural parameters and solving for responses to counterfactual policies, we need a tractable MP framework. A salient feature of MP in the car industry is the prevalent use of export platforms: 50% of cars sold in OECD markets are assembled in locations that are neither the headquarter nor the consuming country. The recently developed quantitative framework—that we refer to as the double CES (constant elasticity of substitution) MP model—tractably incorporates export platforms. It combines a CES heterogeneous-firm product market structure as in Melitz (2003) with a constant-elasticity sourcing decision adapted from Eaton and Kortum (2002). Important contributors to the development of this framework include Ramondo (2014), Ramondo and Rodríguez-Clare (2013), Irarrazabal et al. (2013), Arkolakis et al. (2018), Antras et al. (2017), and (closest to our setup) Tintelnot (2017). Comparative statics in these papers generally hinge on two parameters: the first governs *substitutability between products* from the view of consumers, whereas the second describes the *interchangeability of potential production locations* from the firm’s perspective. The double CES framework extends the gravity equation to a setting where it is costly for headquarters to coordinate foreign assembly and distribution

affiliates. Gravity has proven to be a powerful tool for understanding international trade flows; its most attractive features being tractable estimation, good fit to the data, and the ability to conduct counterfactuals with minimal data requirements.¹ Our implementation of the double CES framework here maintains those three desirable features.

With the goal of making plausible predictions for the consequences of trade policy changes, we extend the double-CES MP model in three ways:

1. Multinational firms are also multiproduct firms who offer different sets of varieties in each market. In the car industry, the typical brand makes ten models, of which they only offer one quarter in a typical market. They frequently opt out of serving markets at all (the typical brand serves only a third of the market-year combinations). These facts motivate the inclusion of two extensive margins of adjustment to policy shocks: which markets to serve and what subset of varieties to offer in each.²
2. While frictions between production and consumption locations are standard in pure trade models, the possibility that the firm is headquartered in a third country calls for the consideration of two additional frictions. Costs of producing far from headquarters are a well-established feature of the MP framework (going back to Ramondo (2014)). An innovation of our paper is to incorporate a third friction between headquarters and the market. We interpret this third friction as marketing costs, with both variable and fixed components. The new friction will play a crucial role in explaining our first extension of the MP framework, the market entry margins. A key motivation for incorporating the new frictions is that modern “deep” integration agreements contain whole chapters that do not operate on the origin-destination path traversed by goods. Rather, topics such as harmonization of standards, protection of investments, and facilitation of temporary movement of professionals, mainly affect the flows of *headquarters* services to production and distribution affiliates. Because our data tracks the three countries where a brand is headquartered, produces and sells its products, we are able to econometrically identify the new frictions separately from traditional trade costs.
3. The model incorporates *external* increasing returns to scale (IRS) by specifying the marginal costs of each plant in a given country as a function of total car production in that country. This contrasts with the internal scale economies modeled by Goldberg and Verboven (2001) and Antras et al. (2017) in related contexts. There are two benefits of introducing external IRS into the double-CES MP framework. First, there is ample evidence suggesting that external (Marshallian) scale economies are important in practice.³ One well-supported mechanism

¹Dekle et al. (2007) were the first to use the CES structure of gravity to implement what Costinot and Rodriguez-Clare (2014) call the exact hat algebra approach to counterfactuals.

²Prior work on multinational production omits the two market entry margins, as Arkolakis et al. (2018) assume single-product firms, Antras et al. (2017) consider a single market, and Tintelnot (2017) assumes that firms offer a unit mass of varieties in every market. The variety-entry extensive margin for multiple-product firms was incorporated in a pure trade model by Bernard et al. (2011).

³Spatial concentration of car production has been an important feature of the industry since its founding, as seen

comes from forward and backward linkages between assemblers and parts suppliers.⁴ Second, building external IRS into the cost function allows for important interdependencies across markets without sacrificing tractability, one of the chief attractions of the double-CES MP model. As in Tintelnot (2017) and Antras et al. (2017), firm-level decisions of where to assemble cars remain independent from each other, facilitating estimation. Aggregate-level production decisions nevertheless become *interdependent*, with a change in frictions between two countries affecting outcomes in third countries. Tintelnot (2017) and Antras et al. (2017) build in interdependencies through the mechanism of firms paying a fixed cost to add countries to the choice set. This approach would not be computationally feasible in our context because of the above-mentioned extensive margins that are essential features of the industry.

The data we use come from an automotive industry consultancy that tracks production at the level of *brands* (Acura, BMW, Chevrolet) and *models* (RDX, X5, Corvette). Because our paper uses car data, it invites comparison to a series of papers that have considered trade and competition in this industry. Goldberg (1995), Verboven (1996), and Berry et al. (1999) investigate quantitative restrictions on imports of cars into the US and EU markets. More recently, an independent and contemporaneous paper by Coşar et al. (2018) combines a demand side from Berry et al. (1995) with the MP model of Tintelnot (2017). These papers feature multi-product oligopoly and use either nested or random coefficients differentiated products demand systems. The advantage of these approaches is that they allow for variable markups and yield more realistic substitution patterns than the monopolistic competition with symmetric varieties demand assumed in the double CES model. This method has two disadvantages in our context. First, it severs the connection to the gravity equation from trade. Second, to implement the rich substitution models, the researcher needs to know the prices and continuous characteristics of all the models. Such data are only available for a drastically reduced set of brands, models, and markets.⁵ This would make it impossible for us to consider the global production reallocations associated with the mega-regional agreements.

The chief concern about CES for the purposes of this paper is that it might exaggerate the degree of substitution between models in very different segments of the car market. This could lead to erroneously large responses to trade policy changes (e.g. a Brexit-induced tariff on Polish-made Fiat 500s would be unlikely to trigger much substitution towards UK-made Land Rovers). We mitigate this concern, while maintaining all the computational advantages of CES, by also estimating and simulating a version of the demand side that nests varieties within market segments.⁶ This follows the research line of Goldberg (1995) and Verboven (1996), with two important mod-

in the production clusters around Detroit and Paris. More recently, the examples of plant agglomerations in cities of Slovakia (4 plants), Central Mexico (11 plants), Northern France (6 plants) and the I-75 corridor in the USA (about 10 plants) point to the persistent importance of Marshallian economies.

⁴See Smith and Florida (1994), Klier and McMillen (2008) and Schmitt and Biesebroeck (2013).

⁵The Coşar et al. (2018) data set has 9 markets and 60 brands compared to the 76 markets and 138 brands in our estimating sample.

⁶Grigolon and Verboven (2014) show that nested logit can match fairly closely the cross-price elasticities of a random coefficients model.

ifications. As in Björnerstedt and Verboven (2016), substitution within each nest takes the CES form (albeit with quantity shares). Second, to characterize the maximal extent of divergence from symmetric CES, our formulation restricts all substitution to occur *within segments*. The unified and segmented markets versions of the model therefore bracket the extent of substitution between car models in different segments. The interval between these extreme approaches turns out to be fairly narrow in terms of the outcomes of our counterfactual scenarios.

We recover the structural parameters of the MP model through the sequential estimation of four equations corresponding to four key decisions made by multinational firms: (1) from which of the firm’s factories to source each variety, (2) the quantity supplied to each market (3) which varieties to offer in each market the brand is distributed in, and (4) where to distribute the brand. The two first equations deliver credible estimates of the two pivotal elasticities of the double CES framework: Identifying from variation in car tariffs, we estimate a sourcing elasticity of 7.7. The sourcing equation delivers a brand’s cost index for supplying models to a given destination. Variation in this index identifies the demand-side CES and estimates it as 3.87. We find that regional integration has substantial effects on all three dimensions of frictions. This is fully in line with the observation that export platforms are organized on a regional basis: 85% of export platform for OECD markets occurs within regional trade agreements. Combining all four equations, the double-CES framework performs well when applied to the global car industry data. The bilateral trade flows predicted by the model match the data with a correlation of 0.74. The new features that we incorporate into the MP framework—the market entry margins for models and brands and the marketing costs—prove to be quantitatively important. The median *ad valorem* equivalents of the combined variable and fixed components of marketing costs (68%) are larger than the trade costs (24%) and frictions between headquarters and assembly locations (31%) already standard in the MP literature.

The results from counterfactual trade policy changes improve our understanding of the impacts of the creation and dissolution of regional integration agreements. We predict substantial reallocations of production in response to a set of policies that have recently been proposed or implemented. Perhaps the most striking outcomes are the dramatic output reductions that the Canadian and Mexican car industries would have suffered if the US had followed through on threats to abrogate NAFTA and impose 25% Section 232 (national security) tariffs on its neighbors. Our simulations show that even when the two countries impose retaliatory tariffs, there would be a 40% decline in Mexican production and a even larger 67% cut in Canada. Meanwhile, consumer surplus in all three countries would decline (by as much as 6% in Canada). Through the exemptions to Section 232 tariffs that Canada and Mexico obtained in NAFTA renegotiations, the two countries stand to increase production substantially if the US imposes such tariffs on the rest of the world. Under this trade war scenario, US-based plants also increase production sharply, mainly at the expense of Korea, Japan and Germany, which collectively lose between 2.4 and 2.8 million cars. The other major trade disruption we consider, Brexit, causes relatively minor production losses in the UK (about 2/3 of a typical size plant), but the country’s consumers pay up to

8% more for cars.

A second set of counterfactuals predicts the consequences of new integration agreements that the US opted not to join. We predict that joining the Trans-Pacific Partnership (TPP) or forming a deep trade agreement with the EU (TTIP) would have lowered production in the US. On the other hand, plants in Canada stand to increase production by going ahead with both trans-oceanic agreements. Membership in the 11-member CPTPP boosts Canadian output by one third, primarily because our estimates imply a more than six percent cost reduction for Japanese-owned plants in Canada. By far the greatest consumer benefits in all the policy scenarios go to the Vietnamese whose car price index falls by over 25% with Trans-Pacific integration. Most of this comes from eliminating 44% tariffs on imported Japanese cars, however. The reductions in marketing costs from Trans-Pacific integration are predicted to generate quite large reactions for the entry margins introduced in our paper: for instance, the CPTPP (excluding the US) expands model entry by Japanese brands in Canada by around 20% and reduces the probability that Chevrolet enters the Vietnamese market by more than 30 percentage points.

The paper continues in five main sections. We first discuss and display some of the important empirical features of the global car industry, using the nearly exhaustive firm-level information on where each variety is designed, assembled and sold. Drawing on these facts, the next section generalizes the existing models to include marketing frictions and market-entry decisions at the model and brand level. We then show how the structural parameters of the MP model can be recovered from four estimating equations corresponding to four key decisions made by multinationals. Following estimation, we present the key methodological aspects of our counterfactual exercises. Finally, we use those methods to project the outcomes of (1) trade wars provoked by US imposition of national-security tariffs on cars, (2) soft and hard versions of Brexit, (3) Trans-Pacific and (4) Trans-Atlantic integration agreements.

2 Data and model-relevant facts

Recent work on multinational production uses data sets that cover all manufacturing or even the universe of multinational activities (including services). The drawback of such data sets is the absence of complete micro-level flows. This forces the theory to do more of the work in the estimation process. We concentrate on a single activity within a single sector—the assembly of passenger cars. As this focus raises the issue of the external validity of our results, we think it worthwhile to emphasize compensating advantages of studying the car industry.

The first and foremost advantage of the car industry is the extraordinary richness of the data compiled by IHS Markit.⁷ IHS uses new car registration information (and probably other sources of information) to obtain annual flows at the level of individual models identifying the assembly plant and country of sale. From it we extract origin-destination flows for 4791 car models sold by

⁷Other attractive aspects of the car industry include its size (passenger cars alone constitute 4% of global trade and the broader industry accounts for 5 to 6% of employment in the US and EU) and prominence in public debate.

138 brands over the 2000–2016 period.⁸

What we refer to as a “model” is a combination of three variables in the original dataset: 1) “sales nameplate” which IHS defines as the “Name under which the vehicle is sold in the respective country”;⁹ 2) the “bodytype” defined as “Vehicle silhouette without doors designation”;¹⁰ 3) the “program,” which IHS defines as the “code used by OEMs to identify vehicle throughout design lifecycle.” Programs constitute redesigns, or new generations of a model.¹¹

The empirical analysis in the main text maps the theoretical concept of varieties to models and the concept of firms to brands. Models appear to be the natural counterpart to the concept of varieties. As implied by our theory for individual varieties, we show that models sold in a particular market are almost always sourced from a single assembly location. There are several reasons we employ brands, rather than parent corporations, to correspond to the theoretical concept of the firm. First, the brand is the common identity across models that is promoted to buyers via advertising and dealership networks. This suggests that the brand’s home is the one relevant for marketing frictions. Second, most of the brands under common ownership were originally independent firms (e.g. Chevrolet and Opel (GM), Ferrari and Chrysler (Fiat), Volvo (Geely), Mini (BMW)). Partly for historical reasons, brand headquarters often correspond to the location where models are designed. For example, while Jaguar is owned by Tata Motors, based in India, Jaguar’s cars are designed at the brand’s headquarters in Coventry in the UK. We think of the brand’s headquarters as a principal source of tangible (e.g. engines) and intangible (e.g. designs, managerial oversight) inputs used by the assembly plants.

There are two potential sources of concern when using the brand/model concepts. The first is that headquarter inputs may originate mainly from a higher level than the brand headquarters. A second worry comes from the industry practice of re-badging: different brand/model combinations might cover what is essentially the same underlying car. The richness of the IHS data enables us to replicate all our analysis using an alternative approach that deals with those concerns. The alternative specifies varieties as particular car designs using the identifiers for the “platform” (the underbody of the car) to which we add the above-defined program and bodytype. The concept of firm is the “Design Parent”, the corporation that has managerial control over the design of the platform used by each variety. We discuss the results from implementing this approach, shown in full in Appendix F.2, where relevant as we report stylized facts and regression results.

We identify the brand headquarters (*i*) as the country in which each brand was founded. In the case of spin-off brands like Acura, we use the headquarters of the firm that established the brand (Japan in this case). Unlike the few available government-provided data sets used in the literature, we are not restricted to parent firms or affiliates based in a single reporting country. Rather, our data set is a nearly exhaustive account of global car headquarters, assembly and sales


⁸Due to entry and exit, there are fewer brands and models in each year. For instance, 2128 models were offered by 120 brands in 2016. Appendix E explains the cuts we applied to the original IHS dataset.

⁹Our sample includes 2377 distinct nameplates such as the 500 (Fiat), Twingo (Renault), 3 (Mazda).

¹⁰The bodytypes are sedan, SUV, hatchback, MPV, wagon, coupe, convertible, and roadster.

¹¹The Renault Twingo, for instance, has had three generations to date: X06 (1993–2012), X44 (2007–2014), and X07 (2014–).

Figure 1: Example: Fiat 500 & Punto production organization



locations. Our estimating sample comprises the shipments of cars assembled in 52 countries by brands headquartered in 21 countries and sold in 76 national markets.

Figure 1 illustrates some of the important aspects of our data using the case of the brand Fiat in 2013 for two of its main models, the Punto and the 500, and seven markets. Fiat makes the 199 program of the Punto in Italy, selling the cars to domestic and EU consumers (the dashed line reflects the 66 cars sold in Mexico). A budget version of the Punto (code 310) is made in Brazil for several markets in South America. The Fiat 500, mainly made in Poland (for EU markets) and Mexico (for the Americas), exemplifies the importance of regional export platforms in the car industry. A striking feature of the Fiat example is that no market is assigned to more than one assembly location for a given model. This pattern of single sourcing generalizes very broadly as we show in Fact 1 below. The absence of the Punto in the US market provides an example of selective model-market entry. Fiat does not distribute any of its models in 11 of the 76 markets (mainly in Asia). We show in Fact 3 below that most models and brands are offered in only a minority of the potential markets.

We now turn to describing three empirical facts that bear on the specific features of the model we estimate. The first two relate to key tractability assumptions of the existing model whereas the third represents a feature that we argue should be added to the standard model.

2.1 Fact 1: Almost all models are single-sourced

At the level of detail at which trade data is collected (6 digit HS), most large countries import from multiple source countries. This is part of the reason why the Armington assumption that products are differentiated by country of origin became so commonplace in quantitative models of trade.

In the car industry we have finer detail because specific models of a car are more disaggregated than tariff classifications. At the level of models, for a specific market, firms almost always source from a *single* origin country. This is not because all models are produced at single locations. In 2016, about a fifth of all models are produced in more than one country and we observe four that are produced in ten or more countries. Rather, it is because firms match assembly sites to markets in a one-to-many mapping.

Table 1: Numbers of sources for each market-model-year

# Sources	All model-markets			Brands with 10+ locations		
	Count	Col %	Cum %	Count	Col %	Cum %
1	320,069	97.7	97.7	194,901	97.1	97.1
2	7,386	2.3	99.9	5,651	2.8	99.9
3	243	.1	100	208	.1	100
4	5	0	100	5	0	100

Table 1 shows that 98% of the model-market-year observations feature sourcing from a single assembly country. Sourcing from up to four countries happens occasionally but it is very rare. This is true for models produced by brands that have ten or more *potential* production countries, where potential sites are measured by the number of countries where the brand conducts assembly (of any model). In 97% of the cases, these models are still single-sourced.

2.2 Fact 2: Most markets are not highly concentrated

Firms in the car industry are not, of course, “massless” as assumed in the monopolistic competition model. The pertinent question is whether the monopolistic competition provides a useful approximation for answering the questions considered in this paper. The serious drawback of assuming oligopolistic price setting as in Atkeson and Burstein (2008) is that we would no longer be able to express flows as a closed-form multiplicative solution in terms of frictions. This would lose the connection to gravity and therefore also make it impossible to use the simple and direct estimation methods derived in the next sections.

One defense of the use of monopolistic competition is that, in some respects, the industry is not as concentrated as one might imagine. Table 2 shows that most markets feature many competitors. For three quarters of market-years we consider, more than 191 models are available. Even at the highest level of aggregation, the sales parent,¹² three quarters of the markets have at

¹²“Sales Parent” is defined by IHS as “The company who owns the brand at the current point in time.”. For example,

Table 2: Market share concentration in car sales, 2000–2016

Level	Inter-Quartile-Range market shares				Concentration % market-years		
	count	median	CR5	top	low	moderate	high
model	191–372	.05–.1	17–37	5–12	94	5	1
brand	33–49	.32–.98	49–74	14–31	51	33	16
parent	18–24	1–2.53	69–83	22–36	6	66	28
—— MPV	9–13	3.01–7.21	80–96	26–49	0	44	56
—— SUV	13–19	2.16–4.53	67–85	21–38	8	61	31
—— bigCar	13–16	2.16–4.79	74–88	25–40	1	56	43
—— smallCar	14–18	1.34–3.88	74–89	24–40	1	59	40
—— sportlux	7–14	2.48–8.2	84–96	32–54	0	29	71

All figures are calculated over all market-year combinations (76 countries, 2000 to 2016). CR5 is the combined share of the top 5. Markets classified as low ($H < 1000$), medium ($1000 \leq H \leq 2000$), and high ($H > 2000$) concentration based on EU Commission thresholds. The first 3 rows calculate shares of whole passenger car market; the last 5 rows use parent firm shares within market segments.

least 18 competing firms. Column (2) shows that median market shares are small (mainly under 5%), implying that oligopoly markups for the majority of firms would be close to those implied by monopolistic competition. Column (3) shows the concentration ratio for the top five actors at each level of aggregation. In three quarters of the market-years, the top 5 brands account for less than 74% of the market. The last three columns show that at the highest levels of ownership (parent), EU merger guidelines would be “unlikely to identify horizontal competition concerns” for 72% of the market-year combinations.¹³ Even within segments, the majority of markets are moderately concentrated except in the case of MPVs and sport and luxury cars.

To be clear, we are not arguing that oligopoly is irrelevant in the industry. The largest firms are big enough to have endogenous markups that significantly exceed those implied by monopolistic competition. Nevertheless, even under a data generating processes that matches the level of concentration observed in the industry, an estimated CES monopolistic competition model can deliver surprisingly accurate predictions for trade policy counterfactuals. Head and Mayer (2018) simulate data from a BLP framework featuring oligopoly, rich substitution in demand, and multi-product firms that internalize cannibalization effects. The CES-MC model is capable of closely approximating the aggregated counterfactuals for BLP-generated data under settings that replicate data moments for parent firms in Table 2 (5-firm concentration ratios of 70–80% and an average of 10 models per firm¹⁴). There is no theorem guaranteeing the close fit we have found in these simulations generalizes to all situations. However, the simulations establish that the mere fact that CES-MC omits many theoretically desirable features does not systematically prevent it from being

Volkswagen is the Sales Parent of Audi, Bentley, Bugatti, Lamborghini, SEAT, Skoda and Volkswagen. There is a many-to-one mapping between brands and their sales parent.

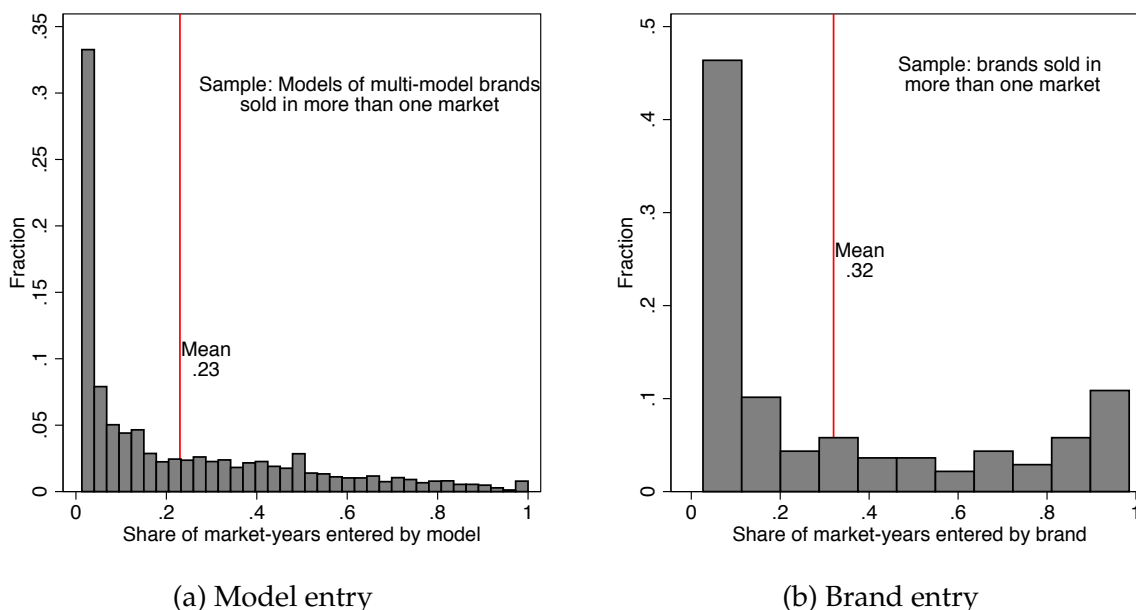
¹³We use the EU threshold because it is intermediate between the corresponding Herfindahl thresholds used by the US Department of Justice (1800) and the Federal Trade Commission (2500).

¹⁴On average, brands offer 6.8 models and parent firms offer 13.4 models.

a useful tool for counterfactual policy exercises. The success of the CES-MC framework in these simulations reinforces its appeal for our purposes, given its tractability, low data requirements, and connection to the gravity equation.

2.3 Fact 3: Most brands and models are offered in a minority of markets

Figure 2: Entry rates for models and brands



In the MP model presented in the next section, the firm decides where to establish distribution networks and which of its varieties to offer in each of those markets. Here we show that the model-level entry margin is very important for multi-model brands in the car industry. Panel (a) of figure 2 depicts the histogram of $\bar{\mathbb{I}}_{mn}$, the model-level mean of the binary variable \mathbb{I}_{mnt} indicating model m is offered in market n in year t . The sample comprises model-market-years where the brand is available, the model is offered in more than one market, and the brand makes more than one model. We observe that brands almost never serve a market with *all* their models and only 15% of models are offered in the majority of the markets where the brand is available. With the average entry rate being just 23%, it seems clear that the standard MP framework should be augmented to include the extensive margin of model-level entry. A potential concern with these figures is that we may be underestimating entry due to the re-badging phenomenon. For example, Mazda sells the car design specified by platform “C1” and program “J68C” as the “Axela” in Japan but as the “3” everywhere else. We thus treat the Axela as being offered in just 1.5% of the market-years. Using the firm-variety methodology described in Appendix F.2 we see that the hatchback version of C1-J68C has an 78% entry rate. However, looking across all varieties the average entry rate is just 24%, slightly larger than the average across all models. Figure F.1 in this appendix

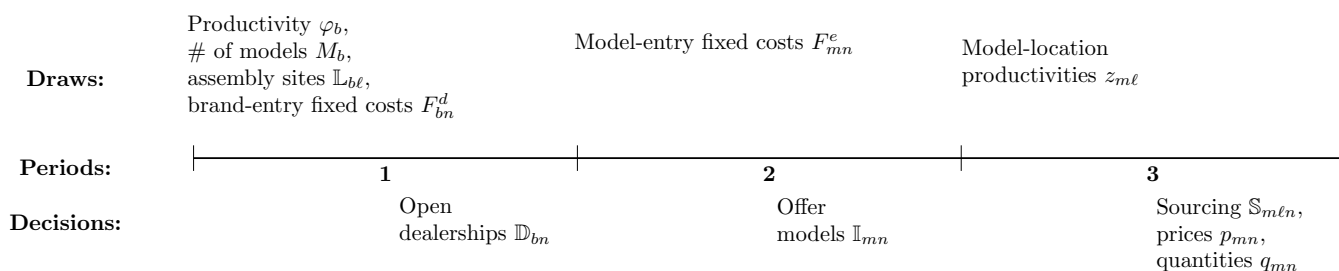
shows that the whole distribution of entry rates is visually unchanged after removing the re-badging issue. The takeaway is that whether we define varieties as the consumer sees them or based on firm-level design distinctions, they tend to be offered in about one quarter of the places where they might be offered.

Panel (b) of figure 2 shows the distribution of entry rates by brands over markets. The key distinguishing feature with the model entry histogram is the existence of a local peak of high entry rates for a few brands. Ten brands enter more than 95% of the available market-years. At the other extreme, 43 brands enter 5% or fewer. On average, a brand enters less than a third of countries it could sell in.

3 The double CES model of multinational production

There are B brands, M symmetrically differentiated car models, and N countries. Brand b is endowed with a headquarter country $i(b)$; a location-independent productivity φ_b ; a portfolio of $M_b = \sum_m \mathbb{M}_{bm}$ models, where $\mathbb{M}_{bm} = 1$ if b owns m ; and a set of production facilities, with $\mathbb{L}_{b\ell} = 1$ for countries where b can manufacture any of its models.

Figure 3: Timeline of when shocks are realized and decisions are made



The sequence of decisions is depicted in Figure 3. At the beginning of period 1, brands are endowed with their defining attributes and learn F_{bn}^d , the fixed costs of creating a dealership network in each potential market. Brands then decide where to establish distribution facilities, corresponding to $\mathbb{D}_{bn} = 1$. For markets the brand has entered, it learns the model-entry fixed costs, F_{mn}^e , in period 2 and decides which models to offer in each market ($\mathbb{I}_{mn} = 1$). In the third period, brands learn the model-location productivity shocks, $z_{m\ell}$, and select the source ℓ , that minimizes the delivered cost to market n for model m (denoted $c_{m\ell n}$), subject to $\mathbb{L}_{b\ell} = 1$. Selected sources have $\mathbb{S}_{m\ell n} = 1$. Based on realized costs, firms set prices, p_{mn} , and read quantities, q_{mn} , off their demand curves. In order to be able to estimate the model sequentially, we assume that, conditional on the observables (e.g. distance from market n to the headquarters of brand b), shocks to market entry costs (F_{bn}^d and F_{mn}^e) are uncorrelated with each other and with $z_{m\ell}$, the production cost shock.¹⁵

¹⁵Ciliberto et al. (2018) consider more general patterns of correlations between cost and demand shocks in a framework that also involves oligopoly pricing and multiple equilibria in the entry game. These features lead to heavy computational burden and challenging parameter identification that go beyond the scope of our paper.

Brand-level profits are given by aggregating the model-level gross profits, π_{mn} , and netting out all fixed costs:

$$\Pi_b = \sum_{m=1}^M \mathbb{M}_{bm} \left[\sum_{n=1}^N \mathbb{I}_{mn} (\pi_{mn} - F_{mn}^e) \right] - \sum_{n=1}^N \mathbb{D}_{bn} F_{bn}^d \quad (1)$$

where
$$\pi_{mn} = \left(p_{mn} q_{mn} - \sum_{\ell=1}^N \mathbb{S}_{m\ell n} c_{m\ell n} q_{m\ell n} \right)$$

In summing over models above and throughout the paper, we follow our data and our programming in treating models as discrete entities. However, firm behavior is modelled as being monopolistically competitive: brands make pricing and entry decisions as if both brands and models were massless.

Profit maximization is constrained by $\mathbb{I}_{mn} \leq \mathbb{D}_{bn}$, and $\mathbb{S}_{m\ell n} \leq \mathbb{L}_{b\ell}$. That is, a brand distribution network in market n is necessary if any of the brands' models are to be offered there and the brand must have a plant in location ℓ if it is to be used as a source for any model.

This paper considers location choices as pre-determined variables that constrain subsequent choices. This makes sense for analysis of the medium-run consequences of policy changes. As we have already mentioned, one reason we take this approach is to avoid the computational challenges of modelling plant location choice when the profit function is neither globally submodular nor supermodular. A second reason for taking locations as given is the strong persistence observed in the set of locations where each brand operates: three quarters of all cars produced in 2016 (the year for which we run counterfactuals) were assembled in brand-country combinations that were already active in 2000. This fraction rises to 88% for OECD countries and to 94% if we draw the line in 2007, a decade before our counterfactuals.

3.1 Consumer preferences and demand

In our data we observe only quantities, not expenditures, and therefore wish to use a specification in which firm-level sales volumes are expressed as a share of total quantity demanded. As in the recent work of Fajgelbaum et al. (2011), we derive demand from the discrete choices across models by logistically distributed consumers. In contrast to that paper, however, our formulation retains the constant elasticity of substitution. Following Hanemann (1984), under conditions detailed in appendix A, households denoted h choose m to minimize $p_{mn(h)}/\psi_{mh}$, where $p_{mn(h)}$ is the price of model m in the market n where household h is located and ψ_{mh} is the quality that household perceives. With ψ_{mh} distributed Fréchet with shape parameter η (an inverse measure of customer heterogeneity), quantity demanded for model m in market n is given by

$$q_{mn} = \left(\frac{p_{mn}}{P_n} \right)^{-\eta} Q_n \quad \text{where} \quad P_n \equiv \left(\sum_j \mathbb{I}_{jn} p_{jn}^{-\eta} \right)^{-1/\eta}, \quad (2)$$

where Q_n denotes aggregate new car purchases in the market. The Fréchet shape parameter η is the first of the two elasticities of substitution that drive outcomes in this framework. As with σ in the Dixit-Stiglitz framework, η is the own price elasticity of demand and also determines the sales to profit ratio. The key difference from the Dixit-Stiglitz setup is that here market shares q_{mn}/Q_n are expressed in terms of *quantities* rather than expenditures.

The delivered price of model m in n under monopolistic competition is a constant markup $\eta/(\eta - 1)$ of marginal cost. Substituting this price into the demand curve, sales are

$$q_{mn} = \left(\frac{\eta}{\eta - 1} c_{mn} \right)^{-\eta} Q_n P_n^\eta, \quad (3)$$

where c_{mn} is the marginal cost of model- m cars delivered to market n . Delivered costs depend on the sourcing decision, which in turn depends on a comparison of assembly and trade costs across candidate supply countries.

3.2 Costs (including frictions)

With marginal costs taken as given, each firm looks for the best site, conditional on its set of potential locations: $c_{mn} = \min\{c_{m\ell n}, \forall \ell \text{ such that } \mathbb{L}_{b\ell} = 1\}$. The marginal cost of assembling model m in country ℓ depends on the costs of inputs as well as three productivity determinants, which we elaborate on in successive paragraphs.

Plants combine inputs obtained locally with inputs imported from the headquarter country with Cobb-Douglas technology. Let w_ℓ and w_i denote the costs of a composite factor in the assembly and headquarter countries respectively. On top of worker wages and efficiency, w captures the price and variety of parts available in each country. Parameter $1 - \alpha$ denotes the cost share of headquarter-country inputs.¹⁶ Frictions applicable to inputs sourced from the headquarter country are captured in $\tau_{i\ell}^H \geq 1$. The $\tau_{i\ell}^H$ term includes tariffs country ℓ imposes on key inputs (engines, transmissions) from country i . We use $\gamma_{i\ell} \equiv (\tau_{i\ell}^H)^{1-\alpha}$ as the composite measure of the costs of separating assembly from headquarters to emphasize the similarity to the corresponding friction in Arkolakis et al. (2018). The difference is that our γ reflects input transfers from headquarters whereas their γ is a penalty in terms of lost productivity associated with the transfer of operational methods from HQ to the assembly country. The combined cost of inputs can therefore be expressed as $w_\ell^\alpha w_i^{1-\alpha} \gamma_{i\ell}$.

Turning to productivity determinants, the first we consider is a brand-level shifter, φ_b , familiar from models of firm-level heterogeneity. The second productivity shifter, $z_{m\ell}$, is a model-location shock capturing how well-suited country ℓ is for assembly of model m . It is also a familiar component of the multinational production framework growing out of the product-country heterogeneity term in Eaton and Kortum (2002). Following the literature, $z_{m\ell}$ is distributed Fréchet

¹⁶ The notion that plants rely heavily on inputs from their HQ country is consistent with de Gortari's (2017) observation that exports from German-owned plants in Mexico contain much higher German content than US-owned plants exports to the US.

with shape parameter θ . The final determinant of productivity is external economies of scale, a novel element in multinational production models, but recently incorporated in trade models by Kucheryavyi et al. (2016). As with that paper and the related empirical estimation carried out in Bartelme et al. (2018), productivity is specified as a power function of industry size in country ℓ . Parameter ς is the elasticity of costs with respect to the amount of production in the assembly country, q_ℓ . External economies of scale correspond to $\varsigma < 0$. While many micro-foundations could underlie the external IRS posited here, the Marshallian mechanism of downstream production attracting a denser network of parts suppliers is very plausible in this industry. Combining the input cost terms and the three productivity shifters, marginal costs are given by

$$c_{m\ell} = \frac{w_\ell^\alpha w_i^{1-\alpha} \gamma_{i\ell} q_\ell^\varsigma}{\varphi_b z_{m\ell}}. \quad (4)$$

Compared to other papers in the multinational production literature, the main innovation in equation (4) is that it allows for returns to scale that are external to the plant. Based on the large absolute size of car assembly plants and their tendency to cluster spatially, increasing returns that are both internal and external to the plant are likely to be important in the industry. Incorporating internal increasing returns would be computationally challenging. One critical problem in our context is that internal IRS would violate the conditions underlying the tractable equations for the multinational production model that we employ in the estimation and the solution of the model. This is especially true with respect to the case where marginal costs are decreasing in plant scale. Then the outcome of each sourcing decision affects the relative attractiveness of the sources in every other decision. This means there are vast sets of possible output configurations to evaluate. An alternative way to obtain internal IRS is to model it in the form of constant variable costs combined with a fixed cost per plant as in Tintelnot (2017) and Antras et al. (2017). This retains the tractability of the sourcing equation but imposes a complex computation of optimal plant location. While those two papers have solved this problem within their frameworks, our extensive margins on the model and brand distribution side prevent us from applying their solutions.¹⁷

Fortunately, many of the key aggregate implications of internal IRS also carry over to the external IRS set-up. For example, suppose the US imposes a tariff on cars from Mexico. This will tend to lower production in Mexico, causing the plants there to contract, with each moving up its downward-sloping marginal cost curve. The rise in marginal cost will further reduce demand from the US (an amplification effect) and also lower demand from Canada (an interdependency effect). Suppose instead that marginal costs are constant, but a fixed cost must be paid to keep each plant open. Then the least profitable plants will close and drop out of the brand's sourcing set for sales in both the US and Canada. We would therefore see declines in realized sourcing from Mexico in both markets. These are the same qualitative industry-country-level predictions entailed by external increasing returns. As the assumption of external IRS allows the model to

¹⁷The algorithm employed by Antras et al. (2017) to solve for the optimal choice set requires a super-modular objective function. Although Arkolakis and Eckert (2017) generalizes the algorithm to handle sub-modular problems, our objective function has both super and sub-modular regions.

incorporate interdependencies and amplification effects in response to trade policy changes at the industry-country level, while still being able to exploit the attractive functional forms implied by the assumption of constant returns at the brand level, it seems like an attractive compromise. The outcomes under alternative formulations of IRS will vary across firms, of course, and the aggregate magnitudes need not be the same. Finding a way to handle IRS internal to the firm in MP models with extensive margins on the distribution side remains an interesting topic for future research.

The delivered marginal cost of model m from assembly country ℓ to market n is

$$c_{m\ell n} = c_{m\ell} \tau_{\ell n} \delta_{in}, \quad (5)$$

where $\tau_{\ell n} \geq 1$ represents conventional trade costs such as tariffs and freight, and $\delta_{in} \geq 1$ captures variable distribution and marketing costs. The δ_{in} friction includes the added cost of operating dealership networks abroad, as they may be easier to manage over shorter distances, and with RTA visas (or free movement of labor in the case of economic unions) facilitating visits from head office managers. Increases in variable costs brought about by foreign regulatory requirements would also be reflected in δ_{in} .¹⁸

3.3 Sourcing decision

Brands choose the optimal production locations for each model they intend to sell in a market from the set of countries where the brand has assembly facilities, i.e. $\mathbb{L}_{b\ell} = 1$. The firm's optimal strategy is to single-source for each model-market combination from the country offering the minimum $c_{m\ell n}$. The probability that ℓ is selected is the probability that $c_{m\ell n}$ is lower than the alternatives:

$$\begin{aligned} \text{Prob}(\mathbb{S}_{m\ell n} = 1 \mid \mathbb{L}_{b\ell} = 1) &= \text{Prob}(c_{m\ell n} \leq c_{mk n}, \forall k \text{ with } \mathbb{L}_{bk} = 1) \\ &= \text{Prob}(\ln z_{m\ell n} - \alpha \ln w_{\ell} - \ln \gamma_{i\ell} - \ln \tau_{\ell n} - \varsigma \ln(q_{\ell}) \\ &> \ln z_{mk n} - \alpha \ln w_k - \ln \gamma_{ik} - \ln \tau_{kn} - \varsigma \ln(q_k), \forall k \text{ with } \mathbb{L}_{bk} = 1). \end{aligned}$$

Firm level productivity, φ_b , the friction δ_{in} , and the HQ cost factor w_i cancel out of this probability since they affect all ℓ locations the same way. The probability of selecting origin ℓ from the set of locations where the brand has a plant ($\mathbb{L}_{b\ell} = 1$) as the source of model m in market n is the same

¹⁸For example, foreign car makers complained about the additional costs of daytime running lamps when Canada mandated them for new cars in 1990. Another telling example comes from the 2018 renegotiation of the Korea-US RTA (<http://money.cnn.com/2018/03/27/news/economy/us-south-korea-trade-deal/index.html>). The revised deal allows US car makers to export up to 50,000 vehicles per year to South Korea that do not comply with South Korean safety rules (up from 25,000).

for all models of a given brand:

$$\text{Prob}(\mathbb{S}_{m\ell n} = 1 \mid \mathbb{L}_{b\ell} = 1) = \left(\frac{w_\ell^\alpha \gamma_{i\ell} \tau_{\ell n} q_\ell^\zeta}{C_{bn}} \right)^{-\theta}, \text{ with } C_{bn} \equiv \left(\sum_k \mathbb{L}_{bk} (w_k^\alpha \gamma_{ik} \tau_{kn} q_k^\zeta)^{-\theta} \right)^{-1/\theta}. \quad (6)$$

θ is the second CES in this framework, playing the same role as in Eaton and Kortum (2002). C_{bn} is the multinational production cost index, summarizing the firm's costs of serving market n . Versions of this equation appear in Arkolakis et al. (2018) as equation (6), Tintelnot (2017) as equation (9), and Antras et al. (2017) equation (7).¹⁹ We are the first to estimate this equation directly with both $\tau_{\ell n}$ and γ_{in} frictions, because such estimation requires variety-level data on sourcing for multiple markets and for firms with many different headquarters.

3.4 Brand-level market shares

All models are ex-ante symmetric. Taking expectations over the z shocks implicit in c_{mn} we can use equation (3) to derive expected model-level sales in market n as

$$\mathbb{E}[q_{mn}] = \mathbb{I}_{mn} \left(\frac{\eta}{\eta - 1} \right)^{-\eta} P_n^\eta Q_n \mathbb{E}[c_{m\ell n}^{-\eta} \mid \mathbb{S}_{m\ell n} = 1]. \quad (7)$$

Expected $c_{m\ell n}^{-\eta}$ is multiplicative in the expectation of $z_{m\ell n}^\eta$ conditional on ℓ being the lowest cost location for mn . Adapting a result from Hanemann (1984), this expectation is

$$\mathbb{E}[z_{m\ell n}^\eta \mid \mathbb{S}_{m\ell n} = 1] = [\text{Prob}(\mathbb{S}_{m\ell n} = 1 \mid \mathbb{L}_{b\ell} = 1)]^{-\frac{\eta}{\theta}} \Gamma\left(1 - \frac{\eta}{\theta}\right),$$

where $\Gamma(\cdot)$ denotes the Gamma function. Combining this result with the cost function equations (4) and (5), the ℓn and $i\ell$ cost factors cancel with their counterparts in $\text{Prob}(\mathbb{S}_{m\ell n} = 1 \mid \mathbb{L}_{b\ell} = 1)$. Substituting back into (7) leads to a simple multiplicative expression for expected market share:

$$\mathbb{E}[q_{mn}/Q_n] = \mathbb{I}_{mn} \kappa_1 \left(\frac{\varphi_b P_n}{w_i^{1-\alpha} \delta_{in}} \right)^\eta C_{bn}^{-\eta}, \quad (8)$$

where $\kappa_1 \equiv \left(\frac{\eta}{\eta-1} \right)^{-\eta} \Gamma\left(1 - \frac{\eta}{\theta}\right)$.

Summing over the models that b sells in n , the expected market share of brand b in market n (conditional on having a distribution network in n and offering M_{bn} models) is

$$\mathbb{E}[q_{bn}/Q_n \mid \mathbb{D}_{bn} = 1, M_{bn}] = \sum_m \mathbb{M}_{bm} \mathbb{E}[q_{mn}/Q_n] = \kappa_1 M_{bn} (\varphi_b/w_i^{1-\alpha})^\eta \delta_{in}^{-\eta} P_n^\eta C_{bn}^{-\eta}, \quad (9)$$

¹⁹Like Tintelnot (2017), we assume independent productivity shocks whereas the Arkolakis et al. (2018) formulation allows for them to be correlated.

where $M_{bn} = \sum_m \mathbb{M}_{bm} \mathbb{I}_{mn}$ and the price index is re-expressed as

$$P_n = \kappa_1^{-1/\eta} \left(\sum_b M_{bn} (\varphi_b / w_i^{1-\alpha})^\eta \delta_{in}^{-\eta} C_{bn}^{-\eta} \right)^{-1/\eta} \quad (10)$$

The number of models a brand offers in a market, M_{bn} , is endogenous but it can be moved to the left-hand side of (9) to obtain an expression for the brand's average market share in market n :

$$\mathbb{E} \left[\frac{q_{bn}}{M_{bn} Q_n} \mid \mathbb{D}_{bn} = 1 \right] = \exp(\ln \kappa_1 - \eta \ln \delta_{in} - \eta \ln C_{bn} + \eta \ln(\varphi_b / w_i^{1-\alpha}) + \eta \ln P_n). \quad (11)$$

This equation for expected average market shares is a generalized linear model of the determinants of $\ln \delta_{in}$. The γ or τ frictions enter via the multinational production cost index C_{bn} .

The coefficient on $\ln C_{bn}$ identifies the demand elasticity η . Our identification exploits the CES-monopolistic competition implication of complete passthrough. This allows us to identify the demand elasticity by substituting in a delivered cost index (C_{bn}) in place of the expected price. This index varies across markets and periods for a given brand mainly because of tariff and RTA variation. Therefore our key identifying assumption is that trade policy variation is orthogonal to variation in the quality of a brands' varieties. Because the cost index C_{bn} also depends on local productivity-adjusted wages in the set of countries where the brand assembles, we also need to assume that brands do not systematically source high-demand varieties from high-wage countries. We return to this issue when interpreting our elasticity estimates.

3.5 Model-market entry decision

The incentive to enter a market depends on expected profitability. To explain why all models of a given brand do not always enter (or stay out of) a given market, we introduce mn heterogeneity in the form of fixed market-entry costs. Entry costs increase proportionately to a new set of frictions denoted δ_{in}^e , the fixed cost counterpart of δ_{in} , representing systematic increases in fixed costs associated with separation between the headquarters country and the market. For example, regulations are often claimed to mandate product specifications that the home-based firms have already adopted. Redesigning a model to conform with foreign product regulations, and promoting the model to make consumers aware of it are two examples of costs that enter δ_{in}^e . With HQ and destination-market inputs used for marketing in proportions ζ and $1 - \zeta$, the fixed costs are given by $F_{mn}^e = w_i^\zeta w_n^{1-\zeta} \delta_{in}^e \epsilon_{mn}^e$. The fixed costs shock, ϵ_{mn}^e , is log-normal with parameters $\mu_n^e + \beta_b^e$ and σ_e . Country characteristics such as size and costs of registering a new product are captured in μ_n^e whereas brand-specific determinants of entry costs are in β_b^e .²⁰

The probability that entry occurs, $\mathbb{I}_{mn} = 1$, is the probability that model-level expected profits

²⁰We implicitly assume β_b^e is linearly decreasing with $\ln \varphi_b$ so as to make the model invariant to the scale of φ_b . This allows to normalize one brand's productivity to be one when we extract the structural parameters.

net of fixed costs are positive:²¹

$$\text{Prob}(\mathbb{I}_{mn} = 1) = \text{Prob}(\mathbb{E}[\pi_{mn}] > F_{mn}^e). \quad (12)$$

With a constant demand elasticity, η , variable profits are given by

$$\mathbb{E}[\pi_{mn}] = \mathbb{E}[p_{mn}q_{mn}]/\eta = \mathbb{E}[p_{mn}^{1-\eta}]P_n^\eta Q_n/\eta. \quad (13)$$

The brand foresees that it will choose the optimal assembly location after learning the realizations of the model-location productivity shocks, z_{ml} . Applying the moment generating equation from Hanemann (1984),

$$\mathbb{E}[p_{mn}^{1-\eta}] = \kappa_2 \left(\frac{w_i^{1-\alpha} \delta_{in}}{\varphi_b} \right)^{1-\eta} C_{bn}^{1-\eta},$$

where $\kappa_2 \equiv \left(\frac{\eta}{\eta-1} \right)^{1-\eta} \Gamma \left(1 + \frac{1-\eta}{\theta} \right)$. Substituting this expression into (13),

$$\mathbb{E}[\pi_{mn}] = \frac{\kappa_2}{\eta} \left(\frac{w_i^{1-\alpha} \delta_{in}}{\varphi_b} \right)^{1-\eta} C_{bn}^{1-\eta} Q_n P_n^\eta. \quad (14)$$

Taking logs on both sides of the inequality in (12), substituting in expected profits, and incorporating the distributional assumptions for F_{mn}^e , the expected share of models offered is

$$\begin{aligned} \mathbb{E}[M_{bn}/M_b] = \text{Prob}(\mathbb{I}_{mn} = 1) = & \Phi[(\ln \kappa_2 - \ln \eta + (\eta - 1) \ln(\varphi_b/w_i^{1-\alpha}) - (\eta - 1) \ln C_{bn} \\ & - (\eta - 1) \ln \delta_{in} - \beta_b^e - \ln w_i^\zeta - \ln \delta_{in}^e + \ln Q_n + \eta \ln P_n - \ln w_n^{1-\zeta} - \mu_n^e)/\sigma_e]. \end{aligned} \quad (15)$$

The terms of this equation indexed b or i will be captured collectively with a brand fixed effect. The last four terms on the second line go into a destination fixed effect. This entry equation produces the sensible prediction that the share of models offered in a market increases with its size, quality and efficiency of the brand, and declines with frictions, fixed costs and local competition (a low P_n). The entry decision also depends negatively on C_{bn} , the expected cost of serving n , which is lower for a brand if its plants are located in countries that have either low assembly costs or low transport costs to the market, both costs being part of C_{bn} . The probability of entry is invariant to an increase of wages everywhere by the same proportion.²²

3.6 Brand entry (distribution networks)

Individual models differ in terms of comparative advantage and marketing fixed costs but each model has the same expected value of profits, $\mathbb{E}[\pi_{mn}]$. The brand's total profit in market n gross of

²¹A consequence of monopolistic competition is that concerns over cannibalization are absent from the model entry decision.

²²Multiplying the composite factor cost w by λ lowers profits by $-[1 + (\eta - 1)(1 - \alpha)] \ln \lambda$ directly. There is also the $-(\eta - 1)\alpha \ln \lambda$ effect via C_{bn} and the $\eta \ln \lambda$ effect via the price index. These terms cancel each other.

costs of a brand-level distribution network is

$$\mathbb{E}[\pi_{bn}] = M_b \times (\mathbb{E}[\pi_{mn}] - \mathbb{E}[F_{mn}^e | \mathbb{I}_{mn}]) \times \text{Prob}(\mathbb{I}_{mn} = 1), \quad (16)$$

where $\mathbb{E}[\pi_{mn}]$ comes from equation (14) and $\text{Prob}(\mathbb{I}_{mn} = 1)$ comes from equation (15). Expected fixed costs conditional on the model being profitable enough to offer is

$$\mathbb{E}[F_{mn}^e | F_{mn}^e < \mathbb{E}[\pi_{mn}]] = \exp([\mu_n^e + \beta_b^e + \ln(w_i^\zeta w_n^{1-\zeta}) + \ln \delta_{in}^e] + 0.5\sigma_e^2) \frac{\Phi(z_{bn} - \sigma_e)}{\Phi(z_{bn})}, \quad (17)$$

where $z_{bn} \equiv (\ln \mathbb{E}[\pi_{mn}] - [\mu_n^e + \beta_b^e + \ln(w_i^\zeta w_n^{1-\zeta}) + \ln \delta_{in}^e])/\sigma_e$ is the “standardized” expected net profitability of an individual model—which is the same for all models of a given brand.²³ With this notation, $\text{Prob}(\mathbb{I}_{mn} = 1) = \Phi(z_{bn})$. Using this equality and plugging expected fixed costs (17) into (16) yields expected brand profit gross of setting up distribution facilities in country n as

$$\mathbb{E}[\pi_{bn}] = M_b \left[\mathbb{E}[\pi_{mn}] \Phi(z_{bn}) - w_i^\zeta w_n^{1-\zeta} \delta_{in}^e \exp(\mu_n^e + \beta_b^e + 0.5\sigma_e^2) \Phi(z_{bn} - \sigma_e) \right]. \quad (18)$$

The probability of brand entry is the probability that expected profits of b in n exceed fixed costs of establishing distribution facilities for the brand. As with model-level entry costs, we assume brand-level distribution costs are the product of headquarter wages, frictions, and a shock term: $F_{bn}^d = w_i^\zeta w_n^{1-\zeta} \delta_{in}^d \epsilon_{bn}^d$. The brand-destination shock to fixed costs of *brand* entry, ϵ_{bn}^d , is log-normal with parameters $\mu_n^d + \beta_b^d$ and σ_d . Country-level and brand-level determinants of the fixed costs associated with setting up a new business are captured by μ_n^d and β_b^d respectively.

Taking logs of the brand entry condition $\mathbb{E}[\pi_{bn}] > F_{bn}^d$ leads to the following probability of brand entry :

$$\text{Prob}(\mathbb{D}_{bn} = 1) = \Phi \left(\frac{\ln \mathbb{E}[\pi_{bn}] - [\mu_n^d + \beta_b^d + \ln(w_i^\zeta w_n^{1-\zeta}) + \ln \delta_{in}^d]}{\sigma_d} \right). \quad (19)$$

4 Empirical implementation

Equations (6), (11), (15) and (19) collectively describe firms’ behavior in the model. We now consider the empirical implementation of those four equations.

4.1 Friction determinants

We start by specifying the empirical content of frictions. The frictions governing trade costs (τ), HQ input transfer costs (γ), variable marketing costs (δ), and fixed model-entry (δ^e) and brand-entry (δ^d) costs, are exponential functions of the observable determinants (some of which vary

²³There is homogeneity of degree zero in z_{bn} with respect to wages for the same reason as equation (15).

over time, hence the new subscript t) denoted $\mathbf{X}_{\ell nt}$, $\mathbf{X}_{i\ell t}$ and \mathbf{X}_{int} :

$$\begin{aligned}\tau_{\ell nt} &= \exp(\mathbf{X}'_{\ell nt}\boldsymbol{\rho}), & \gamma_{i\ell t} &= \exp(\mathbf{X}'_{i\ell t}\mathbf{g}), & \delta_{int} &= \exp(\mathbf{X}'_{int}\mathbf{d}), \\ \delta_{int}^e &= \exp(\mathbf{X}'_{int}\mathbf{f}^e), & \delta_{int}^d &= \exp(\mathbf{X}'_{int}\mathbf{f}^d),\end{aligned}\tag{20}$$

where $\boldsymbol{\rho}$, \mathbf{g} , \mathbf{d} , \mathbf{f}^e and \mathbf{f}^d are vectors of the primitive friction cost parameters.

The \mathbf{X} vectors include the standard explanatory variables used in gravity equations: home, distance, and common language. These variables have already been shown to matter for trade flows and affiliate sales. The differences in subscripts are of critical importance to the estimation. Thus $\text{home}_{\ell n}$ indicates that the assembly plant is in the same country as where the car is bought, whereas $\text{home}_{i\ell}$ equals one when the plant is located in the headquarters country, and finally home_{in} turns on when consumer and brand share the same home country. Distance is the average number of kilometers on great-circle route between the main cities in the corresponding countries. Language indicates that the countries share an official language.

In keeping with our focus on the role of trade policies in determining the pattern of multinational production, the \mathbf{X} vectors include additional determinants that are novel to our study. First, in $\mathbf{X}_{\ell nt}$ we have the log of one plus the tariff each country n imposes on ℓ -origin passenger cars in year t . We also include in $\mathbf{X}_{\ell nt}$ an indicator for a “deep” regional trading agreement between ℓ and n in year t , set equal to one if the agreement includes customs-related procedures or services.

In $\mathbf{X}_{i\ell t}$ we include tariffs on imported inputs (major components only) from the headquarters country. As with tariffs on assembled cars, the input tariffs enter with the functional form $\ln(1 + \text{tariff})$. As with the determinants of τ , we allow γ to depend on the existence of a deep integration agreement. In the $i\ell$ dimension, depth is obtained via an investment chapter, or if the RTA includes a services agreement or customs-related procedures. The last of these is likely to be important if the assembly factor relies on the headquarters country for car parts.

The frictions in the in dimension, δ_{in} , δ_{in}^e and δ_{in}^d , differ from the previous \mathbf{X} vectors in two important respects. First, there is no analogue to tariffs in this dimension. To capture the idea that LDCs may be more protective in their regulations of domestic brands, we interact home_{in} with LDC_n , an indicator that the country in question is not a member of the OECD. Our distinctive indicator of depth for RTAs in the in dimension is the inclusion of a chapter on technical barriers to trade (TBTs), which often include provisions for mutual recognition of standards. As in the other dimensions, a sufficient condition to qualify as a deep agreement (in all dimensions) is the inclusion of services. The rationale here is that the operation of car dealerships is a service activity.

Appendix E provides more detail on measurement of the friction determinants, in particular the sources and procedures used for the tariffs and the deep RTA indicators.

4.2 Estimating equations

We now express the four equations that identify the structural parameters in an estimable way in terms of observed variables with associated coefficients and fixed effects (denoted $\text{FE}^{(j)}$ where

$j = 1, 2, 3, 4$).

(1) Sourcing

We transform the sourcing equation into its estimable version by substituting the τ and γ frictions from equation (20) into (6) and setting $\theta\alpha(\ln w_{\ell t} - \ln w_{\ell T}) = \mathbf{W}'_{\ell t}\mathbf{v}_1$, where $\mathbf{W}_{\ell t}$ comprises two proxies for changes in production cost: per capita income and the price level of GDP, and \mathbf{v}_1 is the set of associated coefficients.²⁴ Both proxies are expressed as logs of indices that take values of 1 in $T = 2016$.

In our setup, the probability that brand b sources model m from country ℓ to serve consumers in n in year t is the same across all b 's models. We can therefore aggregate the binary decisions into a count variable, summing over the number of models owned by b and sourced from ℓ : $S_{b\ell nt} \equiv \sum_m \mathbb{M}_{bm} S_{m\ell nt}$. Expected sourcing counts are

$$\mathbb{E}[S_{b\ell nt} \mid \mathbb{L}_{b\ell t} = 1] = \exp[\text{FE}_{\ell}^{(1)} - \mathbf{W}'_{\ell t}\mathbf{v}_1 - \theta\zeta \ln q_{\ell t} - \theta\mathbf{X}'_{\ell nt}\boldsymbol{\rho} - \theta\mathbf{X}'_{i\ell t}\mathbf{g} + \text{FE}_{nt}^{(1)}]. \quad (21)$$

The destination-time fixed effect is $\text{FE}_{nt}^{(1)} = -\ln(\sum_k \mathbb{L}_{bkt} \exp[\text{FE}_k^{(1)} - \mathbf{W}'_{kt}\mathbf{v}_1 - \theta\zeta \ln q_{kt} - \theta\mathbf{X}'_{knt}\boldsymbol{\rho} - \theta\mathbf{X}'_{ikt}\mathbf{g}])$, and the assembly-country fixed effects are interpreted as $\text{FE}_{\ell}^{(1)} = -\theta\alpha \ln w_{\ell T}$.

Brands select sources from the set of countries in which they currently have plants ($\mathbb{L}_{b\ell t} = 1$).²⁵ Equation (21) can be consistently estimated using Poisson PMLE.²⁶ Substituting the estimated coefficients and fixed effects into (6) yields C_{bn} which we need in the next two estimation steps, market share and model entry.

(2) Brand-level market shares

The second key equation to be estimated is the intensive margin of brand-level sales in each market n , year t . Including the measurable version of our δ_{int} frictions into (11), we obtain the following estimable equation of the brand's average market share over its models:

$$\mathbb{E}\left[\frac{q_{bnt}}{M_{bnt}Q_{nt}} \mid \mathbb{D}_{bnt} = 1\right] = \exp\left[\text{FE}_b^{(2)} - \mathbf{W}'_{i(b)t}\mathbf{v}_2 + \text{FE}_{nt}^{(2)} - \eta\mathbf{X}'_{int}\mathbf{d} - \eta \ln C_{bnt}\right], \quad (22)$$

where $\eta(1 - \alpha)(\ln w_{i(b)t} - \ln w_{i(b)T}) = \mathbf{W}'_{i(b)t}\mathbf{v}_2$ captures the evolution of HQ-related costs through changes in income per capita and GDP price. Notation $i(b)$ designates the HQ country of brand

²⁴The sign of per capita income is ambiguous since it reflects productivity (cost-lowering) and wages (cost-raising). On the other hand, price level of GDP should have a negative influence on sourcing since it captures exchange rate over-valuation.

²⁵Our choice set assumption differs from Coşar et al. (2018) who estimate a cost function that assumes that only the countries currently producing a model enter the set of alternative sourcing locations. For example in the Coşar et al. (2018) approach the choice set for the Renault Twingo would be France and Colombia in 2006, whereas in 2008 the choice set would switch to Colombia and Slovenia (because Renault relocated all its Twingo production for Europe from France to Slovenia in 2007). In our approach, all the countries where Renault is active in a given year are included in the choice. Thus, France, Slovenia, and Colombia (and Turkey etc.) are sourcing options in every year. The distinction between these approaches could be seen as one of short and medium runs (in the long run, brands can expand the set of countries where they have factories).

²⁶The fact that a multinomial discrete choice model can be estimated using Poisson with fixed effects on counts, and yielding identical results to conditional logit was discovered by Guimaraes et al. (2003).

b and $T = 2016$. The structural interpretation of the fixed effects becomes $FE_b^{(2)} = \eta \ln \varphi_b - \eta(1 - \alpha) \ln w_{i(b)T}$ and $FE_{nt}^{(2)} = \ln \kappa_1 + \eta \ln P_{nt}$. The C_{bnt} included as the last control comes from the sourcing probability results from equation (21) where $C_{bnt} = (\sum_k \mathbb{L}_{bkt} (w_{kt}^\alpha \tau_{knt} \gamma_{ikt} q_k^\zeta)^{-\theta})^{-1/\theta}$. This regression allows us to estimate the δ_{int} determinants and provides our estimate of η . The natural way to estimate the moment condition shown in equation (22) is Poisson PML because it does not require an additional homoskedastic log-normality assumption for the error term.²⁷

(3) Model entry decision

As with the sourcing decision, we use the fact that our model predicts the entry probability of models to be constant for a given brand to specify the regression as a fractional probit with left-hand side variable being the share of models offered by b in market n and year t .

Substituting $\delta_{int} = \exp(\mathbf{X}'_{int} \mathbf{d})$ and $\delta_{int}^e = \exp(\mathbf{X}'_{int} \mathbf{f}^e)$ into equation (15) and introducing fixed effects, we obtain the estimable version of the model-market entry equation,

$$\mathbb{E} \left[\frac{M_{bnt}}{M_{bt}} \mid \mathbb{D}_{bnt} = 1 \right] = \Phi \left[\text{CST}^{(3)} + \mathbf{X}'_{int} \mathbf{e} - (\eta - 1) \ln C_{bnt} + FE_b^{(3)} - \mathbf{W}'_{i(b)t} \mathbf{v}_3 + FE_{nt}^{(3)} \right], \quad (23)$$

where the constant, $\text{CST}^{(3)}$ is given by $(\ln \kappa_2 - \ln \eta) / \sigma^e$. The coefficients on the gravity determinants in \mathbf{X}_{int} have structural interpretations given by $\mathbf{e} = -[(\eta - 1)\mathbf{d} + \mathbf{f}^e] / \sigma^e$. Thus, the coefficients on the in friction determinants combine the δ_{in} variable marketing cost effects with the δ_{in}^e fixed marketing costs. Changes in HQ-related costs also involve both determinants: $\frac{(1-\alpha)(\eta-1)+\zeta}{\sigma^e} (\ln w_{i(b)t} - \ln w_{i(b)T}) = \mathbf{W}'_{i(b)t} \mathbf{v}_3$.

All the γ and τ geography effects are captured in the $\ln C_{bnt}$ term, the (inverse) index of how well-positioned brand b 's assembly plants are to serve market n in t . Structural interpretation of fixed effects are $FE_b^{(3)} = \left[(\eta - 1) \ln \varphi_b - \frac{(1-\alpha)(\eta-1)+\zeta}{\sigma^e} \ln w_{i(b)T} - \beta_b^e \right] / \sigma^e$ and $FE_{nt}^{(3)} = [\ln Q_{nt} + \eta \ln P_{nt} - (1 - \zeta) \ln w_{nt} - \mu_n^e] / \sigma^e$.

(4) Brand entry decision

The empirical version of brand entry is obtained inserting $\delta_{int}^d = \exp(\mathbf{X}'_{int} \mathbf{f}^d)$ into (19), with the headquarter inputs needed for brand entry fixed costs specified as $\frac{\zeta}{\sigma_d} (\ln w_{i(b)t} - \ln w_{i(b)T}) = \mathbf{W}'_{i(b)t} \mathbf{v}_4$:

$$\text{Prob}(\mathbb{D}_{bnt} = 1) = \Phi \left(\frac{1}{\sigma_d} \ln \mathbb{E}[\pi_{bnt}] - \mathbf{X}'_{int} \mathbf{f}^d / \sigma_d + FE_b^{(4)} - \mathbf{W}'_{i(b)t} \mathbf{v}_4 + FE_{nt}^{(4)} \right). \quad (24)$$

Inverting the coefficient of our calculated profitability of brand b in market n gives a direct estimate of the standard deviation of log fixed costs, σ_d . The structural interpretations of the brand and destination-time fixed effects are $FE_b^{(4)} = -(\beta_b^d + \zeta \ln w_{i(b)T}) / \sigma_d$ and $FE_{nt}^{(4)} = -(\mu_n^d + \ln w_{nt}^{1-\zeta}) / \sigma_d$. We estimate equation (24) as a binary probit with the constructed $\ln \mathbb{E}[\pi_{bn}]$ on the right-hand side, together with the friction determinants, brand and destination fixed effects.

²⁷Santos Silva and Tenreyro (2006) elaborate on this advantage in the context of gravity equations but it is equally applicable to the estimation of any constant elasticity relationship.

4.3 Identification of structural parameters

Equations (21), (22), (23) and (24) estimated sequentially, yield all the parameters needed to solve the model. Our model is specified such that there is only one estimate for each parameter of interest.

Sourcing: Coefficients from the sourcing equation (21) have structural interpretations $-\theta\rho$, $-\theta\mathbf{g}$, and $-\theta\varsigma$. Thus we can calculate τ and γ friction parameters, as well as the scale elasticity, by dividing our estimates by $-\theta$, the coefficient on car tariffs. The fixed effects on origin countries combined with θ allows us to recover $w_\ell^\alpha = \exp(-\text{FE}_\ell^{(1)}/\theta)$. We recover the share of headquarters' country input in the total costs of production by using the ratio of our two direct price shifters in γ and τ . Recalling that $\ln \gamma_{ilt} = (1 - \alpha) \ln \tau_{ilt}^H$, we estimate $1 - \alpha$ by dividing the coefficient on $\ln(1 + \text{parts tariff}_{ilt})$ by the coefficient on $\ln(1 + \text{car tariff}_{\ell nt})$.

Market share: The coefficients on the friction determinants correspond to $-\eta\mathbf{d}$. Dividing by $-\eta$, the coefficient on $\ln C_{bn}$ in equation (22), yields the vector of δ friction parameters \mathbf{d} . The price indices, P_n , are proportional to $\exp(\text{FE}_{nt}^{(2)}/\eta)$, the exponentiated destination fixed effects divided by the consumer elasticity. A combination of brand-related parameters involving physical productivity and the factor costs in the headquarters of brand b is given by $\varphi_b/w_{i(b)}^{1-\alpha} = \exp(\text{FE}_b^{(2)}/\eta)$.

Model Entry: We obtain σ_e as $1 - \eta$ divided by the coefficient on $\ln C_{bn}$ in equation (23). Model entry fixed costs depend on $\ln w_{nt}^{1-\zeta} + \mu_{nt}^e + \ln \delta_{in}^e$. Inverting the definition of the destination fixed effects, $\ln w_{nt}^{1-\zeta} + \mu_{nt}^e = \ln Q_{nt} + \eta \ln P_{nt} - \text{FE}_{nt}^{(3)}/\sigma_e$. The friction parameters for model entry (needed to compute δ_{in}^e) are obtained from the coefficients on friction determinants, \mathbf{e} , combined with variable version of δ costs obtained from the market share equation, using the formula $\mathbf{f}^e = -\mathbf{e}\sigma_e - \mathbf{d}(\eta - 1)$. The remaining components of model entry fixed costs F_{mnt}^e , namely $\beta_b^e + \ln w_{i(b)}^\zeta$, require more involved manipulations of $\text{FE}_b^{(2)}$ and $\text{FE}_b^{(3)}$, that we relegate to appendix B. There we also show how the fixed effects from market share and model entry regressions can be used to reconstruct $\mathbb{E}[\pi_{mnt}]$, and then $\mathbb{E}[\pi_{bnt}]$, which is needed in the brand entry equation.

Brand entry: Equation (24) estimates σ_d as the inverse of the coefficient on $\ln \mathbb{E}[\pi_{bnt}]$. Destination and brand fixed effects yield estimates of $\mu_n^d + \ln w_{nt}^{1-\zeta}$ and $\beta_b^d + \ln w_{i(b)t}^\zeta$ respectively. Multiplying the coefficients on the friction determinants by σ_d yields the vector \mathbf{f}^d .

The market share and model entry equations both depend on $\ln C_{bnt}$, a variable generated using the estimated coefficients of the sourcing equation. The key determinant of brand entry, $\ln \mathbb{E}[\pi_{bnt}]$, also depends on C_{bnt} as well as estimates of fixed costs and the price index extracted from the market share and model entry equations. Since the presence of these generated regressors has the potential to bias the standard errors, we report bootstrap standard errors for all equations. As brands make repeated decisions (sourcing, entry, etc.), it is important to make standard errors

robust to possible correlation in the errors by clustering. In the bootstrap setting this is achieved by drawing *all* the observations for a given brand cluster in each of the four equations.²⁸

5 Results

5.1 Baseline estimates

Table 3 reports the coefficients for each of the four estimating equations.

Sourcing estimates

Column (1) reports our sourcing results. The estimates reveal the importance of trade costs in selecting sources. Home effects are large: the implied increase in the odds of choosing a location is obtained by exponentiating the coefficient. Plants located in the market being served have odds of being chosen that are 2.6 times higher. Distance from the market also significantly reduces the share of models sourced from an assembly country.

The coefficient of -7.7 on the log of one plus the car tariff implies $\theta = 7.7$ as the critical elasticity of substitution between sources.²⁹ Deep regional trade agreements augment the odds of being chosen by 28%, even after accounting for the tariffs applied by the destination market to the different possible origins of the car. Both tariffs and deep RTA effects will be important for our counterfactuals where we experiment with scenarios involving different combinations of RTA and tariff changes.

The estimates of the γ frictions are much less precise, with standard errors several times those estimated for trade frictions. Two of the effects, distance and language, do not even enter with the expected sign, although neither is significantly different from zero. The significant effect is that assembly locations in the brand's home country are $\exp(2.248) \approx 9.5$ times more likely to be selected. The elasticity on the car parts tariff can be used to infer the share of assembly costs attributable to components from the headquarters country, $(1 - \alpha)$ in the cost equation, which is about 37% ($2.87/7.7$). While the precise value of this ratio should be taken with caution, we now have direct evidence of the importance of intermediate inputs from the headquarters country. This feature of the MP model has major qualitative and quantitative implications for the impact of trade liberalization, as we shall see in the counterfactuals. Deep RTAs between assembly and

²⁸ There are 45 brands that cannot enter the sourcing stage estimation since they only have a single country to source models from. We therefore conduct two separate draws in each repetition of the bootstrap. First, we draw (with replacement) among the 93 brands that can source from at least two countries in a given year. The estimated parameters allow us to construct $\ln C_{bn}$ for all brand-destination combinations relevant in the next three steps. We then draw (with replacement) from the full set of 138 car brands and use this bootstrap sample for brand market share, model entry, and brand entry decisions. This completes the procedure needed for one replication. In order to choose the number of replications, we follow procedures described in Andrews and Buchinsky (2000). This involves 1) running a number of initial bootstrap repetitions (500 in our case), 2) calculating an excess kurtosis statistic for every parameter, 3) running 500 additional bootstrap repetitions. The number of added repetitions is sufficient to set the percentage deviation bound (pdb) for all parameters of interest below 5% with a confidence level of 95%, taking into account the deviations from normality implied by the step 2) calculation (see their equation 3.4).

²⁹ The estimate of θ when using all the locations of the parent firm as options for sourcing also rounds to 7.7 as can be seen in Table F.2. The firm-variety approach also shows similar coefficients for the other determinants of trade costs.

Table 3: Baseline results

Decision:	Sourcing	Market share	Model entry	Brand entry
Dep. Var:	$S_{b\ell nt}$	$\frac{q_{bnt}}{M_{bnt}Q_{nt}}$	$\frac{M_{bnt}}{M_{bt}}$	\mathbb{D}_{bnt}
Method:	PPML	PPML	frac. probit	probit
	(1)	(2)	(3)	(4)
home $_{\ell n}$	0.973 (0.142)			
ln dist $_{\ell n}$	-0.323 (0.06)			
language $_{\ell n}$	-0.042 (0.068)			
ln (1+ car tariff $_{\ell n}$)	-7.696 (0.398)			
Deep RTA $_{\ell n}$	0.246 (0.085)			
home $_{i\ell}$	2.248 (0.49)			
ln dist $_{i\ell}$	0.166 (0.136)			
language $_{i\ell}$	-0.218 (0.308)			
ln (1+ parts tariff $_{i\ell}$)	-2.872 (3.197)			
Deep RTA $_{i\ell}$	0.495 (0.301)			
ln q_{ℓ}	0.27 (0.076)			
home $_{in}$		0.816 (0.249)	0.26 (0.068)	0.718 (0.422)
home $_{in} \times \text{LDC}_n$		-0.028 (0.36)	0.829 (0.118)	1.328 (0.576)
ln dist $_{in}$		-0.339 (0.104)	-0.059 (0.017)	0.009 (0.069)
language $_{in}$		0.289 (0.14)	0.068 (0.036)	0.017 (0.122)
Deep RTA $_{in}$		-0.04 (0.14)	0.121 (0.042)	0.165 (0.117)
ln C_{bn}		-3.874 (0.97)	-0.512 (0.221)	
ln $\mathbb{E}[\pi_{bn}]$				0.595 (0.067)
Observations	347542	46299	46300	128589
R^2	0.824	0.601	0.715	0.618
Fixed effects:	ℓ, bnt	b, nt	b, nt	b, nt

Standard errors bootstrapped with brand b clusters over 1000 replications (see footnote 28). R^2 is the squared correlation of fitted and true dependent variables except in specification (4) where the pseudo- R^2 is reported. Each regression controls for log per-capita income and price level of the assembly country.

headquarter countries are estimated to have a larger effect on sourcing than deep RTAs between assembly and consumer countries, but the standard error is also larger.³⁰

Our method estimates external economies of scale based on the magnitude of the revealed preferences of brands for assembly locations with high aggregate output ($q_{\ell t}$). We estimate $-\theta\varsigma = 0.27$ which, given our $\theta = 7.7$ estimate, implies $\varsigma = -0.035$. As pointed out by Goldberg and Verboven (2001), unobserved factors can both make a location attractive and increase its aggregate production. This would lead to upward bias in our estimate of $-\theta\varsigma$. Our estimation mitigates this through the inclusion of country-specific fixed effects (identifying the degree of external returns in the within dimension). However, our estimate of the external scale elasticity should still be seen as an upper bound since time-varying cost shocks could be correlated with changes in $q_{\ell t}$.

Two simple and common ways to mitigate the endogeneity bias that can be applied to our case are: (1) lagging $q_{\ell t}$, (2) constructing a Bartik-style prediction for $q_{\ell t}$ to be used in a control function approach. The Bartik prediction applies changes in total demand Q_{nt} to brand-origin-destination market shares fixed at their 2002 levels. We also report an even more demanding Bartik specification where the shares to be shifted are brand-origin market shares. In regressions reported in appendix F.1, we followed the two approaches. Lagging lowers the magnitude of increasing returns to $\varsigma = -0.032$, whereas the Bartik approach reduces it further to $\varsigma = -0.024$ (for the most demanding one, $\varsigma = -0.028$ for the other one). Neither approach is perfect, but both values support our use of $\varsigma = -0.035$ as the upper bound of the parameter governing the strength of interdependencies.

The sole estimate of the external returns elasticity in the motor vehicle industry that we know of is from a recent version of Bartelme et al. (2018). Using origin-destination trade flows, they estimate a compound parameter analogous to our ς as 0.15. This is four times larger than our non-instrumented “upper bound” estimate even though they employ a demand-based instrument. There are two quantitatively important reasons to expect their study to yield larger estimates. First, the trade elasticity entering their estimate is 5.7, which is 26% smaller than our $\theta = 7.7$. Second, our sourcing method neutralizes variation in the number, quality, and productivity of brands and models produced in each country. Since the Bartelme et al. (2018) method is identified by variation from more aggregated trade flows, their scale elasticity is more encompassing as it also captures any additional mechanisms linking the national scale of production to *firm*-level performance variables. As we hold such variables constant, $\varsigma = -0.035$ should be interpreted as the output scale elasticity on the costs of different countries in assembling a given variety.

The literature also provides some estimates of *internal* returns to scale, estimated from data on prices or costs. These papers omit the multinational production dimension that causes the combinatorial computational challenge. Our implied $\varsigma = -0.035$ lies within the interval of those

³⁰Appendix F.3 presents a set of estimates of γ frictions from an alternative moment conditions that are consistent with the double-CES MP model. The main takeaway from Table F.3 is that the coefficients on Deep RTA $_{i\ell}$ and on tariffs on car parts are stronger and more significant than in our baseline results. However, since the estimates of θ are also larger, the AVE of deep RTA remain very similar to the baseline. The ratio of coefficients between car and parts tariffs also provides comparable alternative estimates of $1 - \alpha$ ranging between 29% and 50%.

estimates of internal returns to scale for the car industry. They are bigger than the range of values reported by Goldberg and Verboven (2001), -0.006 to -0.03 , but smaller than the -0.11 and -0.07 provided by Verboven (1996) and Fuss and Waverman (1990).

Brand-level market share estimates

Determinants of a brand's market share are estimated in Column (2) of Table 3. The estimate of $\eta = 3.87$ (from the coefficient on $\ln C_{bnt}$) is substantially smaller than the θ obtained in the sourcing decision. It implies that there is considerably more heterogeneity in consumer evaluations of brands than in car maker evaluations of assembly locations. One concern with our estimate of η is that it could be biased towards zero if brands with high unobserved demand shocks systematically locate their assembly in high-wage countries. However, the markup implied by our estimate ($\eta/(\eta - 1) = 35\%$) lies within the highly dispersed set of results found in the Industrial Organization literature on cars. The three pioneering papers, Goldberg (1995), Berry et al. (1995) and Feenstra and Levinsohn (1995) report average markups of 38%, 24%, and 18%, respectively. Verboven (1996) and Berry et al. (1999) show markups of specific models that range from 8 to 36% in the former paper and 24% to 42% in the latter. Most recently, Coşar et al. (2018) report in their Table 11 average firm-market markups ranging from 6% (Peugeot in Brazil) to 12.4% (Peugeot in France).³¹

Among the determinants of marketing frictions, consumers are more than twice as likely to select a home brand, corroborating the large home bias found by Coşar et al. (2018). In addition, we estimate that increasing consumer distance from headquarters sharply lowers market shares, even controlling for distance from the consumer to the assembly location, which is captured by C_{bn} in the same regression. Sharing a common language reduces variable marketing costs, increasing the average market share of a brand in those markets by around a third compared to destinations where consumers speak a different language.

The effect of deep regional agreements is perversely negative in this regression, but its magnitude is small, and is very imprecisely estimated. Deep RTA status operates very strongly on the extensive margins: sourcing in column (1), model entry in column (3), and brand entry in column (4).

Model entry estimates

Column (3) of Table 3 shows that all the marketing cost determinants have the expected signs and are highly significant. More models are offered in the home country of the brand, especially when this country is a developing one. Spatial proximity promote entry as well. Deep RTAs between the headquarter country (i) and the market (n) increase the fraction of models offered by 14% (calculated as the average semi-elasticity). As it seems unlikely that RTAs change preferences, we see the deep RTA_{int} effects as supporting the cost-shifter interpretation. Under this approach,

³¹The firm-variety estimate of $\eta = 1.9$ (standard error of 1) shown in Appendix F.2 implies markups over 100%, well outside this range. We attribute this low, noisy estimate to measurement error in C_{bn} caused by the firm-variety approach.

our δ^e frictions include various types of marketing efforts, in particular managing dealership networks. This may be facilitated by the freer movement of skilled workers that is a commonly included provision of RTAs (e.g. NAFTA, EU). The RTA_{int} effect may also capture the greater ease of compliance with regulatory standards if the head office lies within the region and is therefore more able to exert influence on specific requirements in harmonized rules. Note also that the significance of this fixed cost dimension of RTA effects contrasts with the weak impact of the same variable on brand-level sales. This suggests that deep RTAs reduce the fixed costs of model entry between HQ and destination (δ_{int}^e), rather than the variable marketing costs (δ_{int}), that affect brand sales as well.

The overall cost of serving n for brand b (C_{bn} , constructed from sourcing estimates of the first column) strongly reduces the share of models offered in a market as expected. Dividing $(1 - \eta)$ by that coefficient provides our estimate of $\sigma^e = 5.61$. The fixed cost of introducing new models in a market therefore exhibits very large variation.

Brand entry estimates

The last column of our baseline table shows that a number of determinants for model entry are also relevant for whether the brand is present altogether in a market. Domestic entry is naturally a dominant feature of the data (among the exceptions are Acura, Lexus, Infiniti, Isuzu, and Scion—which are not sold in Japan—and Hummer which continued to be sold in Japan and Taiwan after ending sales in the USA). Deep RTAs also reduce the fixed costs of establishing distribution networks, resulting in a 27% larger probability of brand entry. The inverse of the coefficient obtained on $\ln \mathbb{E}[\pi_{bn}]$ yields our estimate of $\sigma^d = 1.68$.

5.2 Interpreting the structural parameters

Table 4: Friction parameters

	Variable costs			Fixed costs	
Friction:	τ	γ	δ	δ^e	δ^d
Estimate:	ρ	\mathbf{g}	\mathbf{d}	\mathbf{f}^e	\mathbf{f}^d
home	-0.126	-0.292	-0.211	-0.856	-1.206
home \times LDC			0.007	-4.674	-2.231
ln distance	0.042	-0.022	0.088	0.078	-0.016
common language	0.005	0.028	-0.075	-0.165	-0.029
RTA (deep)	-0.032	-0.064	0.010	-0.708	-0.278

Elasticities used to obtain frictions: $\theta = 7.7$, $\eta = 3.87$, $\sigma^e = 5.61$ and $\sigma^d = 1.68$. Calculation of those frictions are described in section 4.3 and use coefficients from Table 3.

Using the procedures detailed in section 4.3, we now proceed to report and interpret the struc-

tural parameters underlying our estimates. The three sets of parameters relevant for variable costs—trade costs τ , multinational production costs γ and marketing costs δ —are reported in the first three columns of Table 4. For each, the coefficient on the k th element of \mathbf{X} maps to proportional increases in price (an *ad-valorem* equivalent) of $\exp(\rho^{(k)} \Delta X^{(k)}) - 1$.

In the case of the τ frictions, we can relate our estimates to what is known from direct measurement of the frictions. The elasticity of τ with respect to distance is of particular interest to us since it has been estimated on its own using various types of data in the literature, including the effect of physical distance on freight costs. The ℓn distance cost elasticity in column (1) of Table 4 is $\rho^{\text{distance}} = 0.042$. Coşar et al. (2018) report a somewhat smaller value of $\rho^{\text{distance}} = 0.016$ (Table 12, column IV). Both estimates of ρ^{distance} , fit in the “reasonable range” of 0.01 to 0.07 in the literature summarized by Head and Mayer (2013). Our results imply that the distance effects on trade flows can be fully explained without reference to the “dark matter” invoked by Head and Mayer (2013) to explain aggregate distance elasticities of -1 or higher. This is not surprising since the main candidate explanations for dark matter—poor information and differences in preferences—should be accounted for in the δ_{in} marketing cost parameters.

The estimates from the market share equation imply large variable costs in the headquarter-market dimension. The distance elasticity is 0.088, more than double the corresponding transport cost elasticity 0.042. Our elasticity is also larger than the Wang (2017) estimate of 0.044 based on export sales of foreign-owned manufacturers in China. By contrast, the home bias in marketing costs for cars is 0.21, quite comparable to Wang’s (2017) estimate of 0.24 for US-headquartered firms, but lower than his 0.95 for Japan and Korea affiliates.

How should we interpret the parameters shown in Table 4? The first thing to note is that the three chief variable cost frictions in the model, $\tau_{\ell n}$, $\gamma_{i\ell}$ and δ_{in} all require a normalization to be meaningful. Put another way, any of these frictions could be scaled up by a constant without changing any of the endogenous variables. The normalization we use is the internal friction within the United States.³² An estimated $\delta_{in} = 1.3$, for example means that firms headquartered in i inflate their delivered prices to consumers in n by 30% more than firms headquartered in the US inflate costs for their home-country consumers.

The fixed cost parameters δ_{in}^e and δ_{in}^d are also defined relative to a reference dyad. Thus again if we estimate $\delta_{in}^e = 1.3$ it means that fixed costs of adding another model are 30% higher for firms from i offering models in n than US firms adding a model in the US. This interpretation also holds for δ_{in}^d . To facilitate comparisons with the variable cost frictions, we want to convert δ_{in}^e and δ_{in}^d into their *ad valorem* equivalents (AVE). This can be accomplished using the following thought experiment: Let gross profits be a given ratio of fixed costs. Suppose we shock fixed costs by δ_{in}^e . Then, in order to keep the previous ratio (and thus model entry probability) unchanged, gross profits must rise by the same proportion. Using (15), we can find a δ'_{in} that would achieve this proportional increase, i.e. $(\delta'_{in})^{-(\eta-1)} = \delta_{in}^e$. Inverting we obtain $\delta'_{in} = (\delta_{in}^e)^{-1/(\eta-1)} < 1$. We define the $\text{AVE}(\delta_{in}^e) = 1 - (\delta_{in}^e)^{-1/(\eta-1)}$. Determining the AVE for δ_{in}^d is more complex because

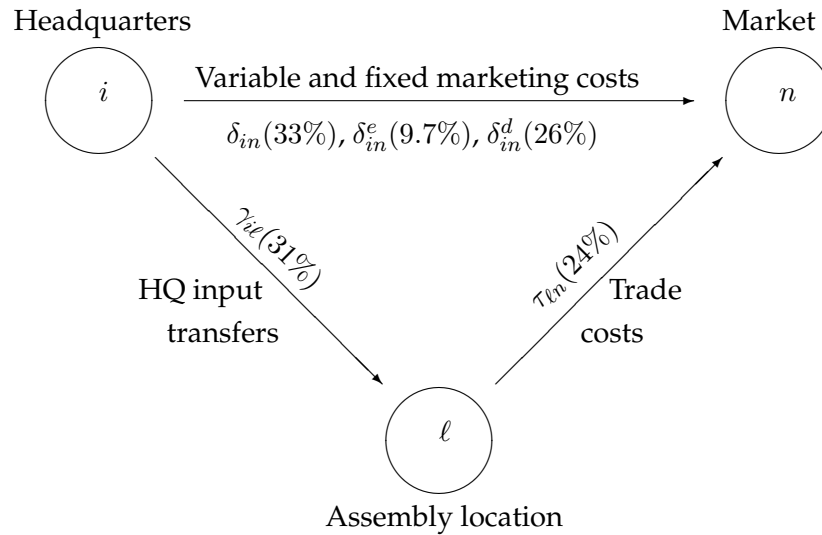
³²This means countries with smaller internal distances than the US can have $\tau_{nn} < 1$.

δ'_{in} affects $\mathbb{E}[\pi_{bn}]$ through multiple non-separable channels. We can still conduct an analogous thought experiment. Define δ'_{in} as the variable marketing cost that would magnify expected gross profits of the brand by the same factor as δ_{in}^d :

$$\frac{\mathbb{E}[\pi_{bn}](\delta'_{in})}{\mathbb{E}[\pi_{bn}](\delta_{in} = 1)} = \delta_{in}^d$$

The δ'_{in} can be found by defining a function $g(\delta'_{in}) = \mathbb{E}[\pi_{bn}](\delta'_{in}) - \delta_{in}^d \mathbb{E}[\pi_{bn}](\delta_{in} = 1)$, and then solving for a root.

Figure 4: Frictions impeding multinational flows



Note: Friction *ad valorem* equivalents shown in (). See text for explanation.

Results of those calculations are reported in figure 4 adapted from Arkolakis et al. (2018) with the addition of an edge corresponding to variable and fixed marketing costs δ_{in} , δ_{in}^e , and δ_{in}^d . On each edge of the triangle we report the relevant frictions, which is the median value calculated in our sample for the year 2016. The variable frictions are $\tau_{ln} - 1 = 24\%$, $\gamma_{il} - 1 = 31\%$, and $\delta_{in} - 1 = 33\%$. The sales of a model produced in ℓ , headquartered in i and sold in n would therefore face a total cost-increasing friction of 116% ($\tau_{ln}\gamma_{il}\delta_{in} - 1$). This is not out of line with the figures provided in Table 7, column (1) of Anderson and Van Wincoop (2004), ranging from 91 to 174%. Being the largest of the three frictions, variable marketing costs are quantitatively important enough to warrant inclusion in the multinational production framework. Set on top of that, the extra burden of fixed marketing costs have AVEs of 9.7% and 26% for model and brand entry, reinforcing the finding that the new dimension of frictions we added to the MP model is quantitatively very important.³³ The only frictions that change in our counterfactual experiments

³³Wang (2017) estimates the first and third components of the marketing costs, using foreign affiliate trade data from

are tariffs and RTAs. Tariffs are *ad valorem* already. The AVE of RTAs for the three dimensions of our friction triangle are 3.2% for τ , 6.4% for γ and a total of 29% for the combined effects of δ , δ^e , and δ^d .

The friction estimates shown in Figure 4 do not distinguish cost-based interpretations of $\tau_{\ell n}$, $\gamma_{i\ell}$, and δ_{in} from preference-based interpretations. For example, a desire by consumers to “buy local” to support workers has the same effect on sourcing as an increase in $\tau_{\ell n}$. Similarly, if Japanese workers had a reputation for quality control, then Toyota’s assembly facilities outside Japan would have their sourcing shares reduced in a way that would be isomorphic to an increase in $\gamma_{i\ell}$. Finally, spatially correlated taste differences (e.g. for fuel economy, safety, or shape) could be equivalent in their effects on market shares to a rise in δ_{in} due to higher distribution costs in remote markets. Allowing for such preference effects in the utility function would just add more parameters that could not be identified separately from the ones existing in our specifications. To estimate separately the cost and demand-side effects would require a different estimation strategy that uses price information. Such a data requirement would severely limit the geographic scope of the study. For the purposes of our counterfactuals on how integration affects production and trade, we do not need to disentangle cost mechanisms from preference mechanisms.³⁴

The index of local assembly costs in each country, $c_\ell \equiv w_\ell^\alpha q_\ell^\zeta$, is a key parameter of the model because it tells us where production would gravitate in the absence of frictions. We obtain the 2016 levels as $c_\ell = \exp(-FE_\ell^{(1)}/\theta)q_\ell^\zeta$ from the estimates in the sourcing equation. The c_ℓ can only be identified up to a scalar so we express them all as cost advantages with respect to the United States, i.e. $100 \times (c_{USA} - c_\ell)/c_\ell$.

Figure 5 graphs the cost advantage of the 47 assembly countries we use in estimation. The clear “winner” for the car industry is South Korea with Japan as runner up. Egypt is the outlier in the other direction. The implied differences in unit assembly costs are quite small for the main European brand headquarters. France, the UK, and Germany are within a few percentage points from each other. Canada is also very similar to its southern neighbor. The similarity in costs between these countries suggests that friction changes have the potential to cause substantial reallocations in production.


5.3 Segment-level estimates

Up until this point we have treated the market for cars as one in which all car models substitute symmetrically for each other. In reality, we think the Camry and Accord mid-sized sedans are closer substitutes for each other than either would be for a Sienna minivan (all Honda models). To allow for more realistic substitution patterns, we follow the tradition initiated by Goldberg (1995) and Verboven (1996) which groups models according to their primary function. In this extension,

China. Since Wang’s sample includes all manufacturing industries, the large magnitudes he obtains suggest that these frictions are important beyond the car industry.

³⁴Goldberg and Verboven (2001) separate out the home bias into demand and cost components. Coşar et al. (2018) estimate cost-based ($\gamma_{i\ell}$) frictions of distance from a brand’s home. They also have a home-brand effect in preferences that would operate as a δ_{in} effect in our model.

Figure 5: Cost advantage inferred from sourcing decisions



households receive idiosyncratic utility shocks that determine the choice of the segment from which they then select a model. This approach leads to a modified version of equation (22), in which market shares are measured *within segments* denoted s :

$$\mathbb{E} \left[\frac{q_{bnst}}{M_{bnst}Q_{nst}} \mid \mathbb{D}_{bnt} = 1 \right] = \exp \left[\text{FE}_{bs}^{(2)} - \mathbf{W}'_{i(b)t} \mathbf{v}_{2s} + \text{FE}_{nst}^{(2)} - \eta_s \mathbf{X}'_{int} \mathbf{d}_s - \eta_s \ln C_{bnt} \right]. \quad (25)$$

All of the demand parameters are segment-specific. While estimating an upper branch decision to determine Q_{nst} would be possible, the effect would simply be to bring back inter-segment substitution. Instead, we take Q_{nst} as exogenous in our counterfactual simulations, a modelling decision that restricts all substitution to take place *within segments*.³⁵ Bracketing the range of substitution is important for our counterfactuals since the unified market assumption has the potential to exaggerate the response to changes in frictions.

Allowing for segment-specific demand elasticities and nested substitution also affects the model-level expected profits: Equation (14) features the size of the market and the price index, which are now both segment-specific. In addition, the size of η_s determines the response to the cost index C_{bn} . Model entry should therefore also be run for each segment. The dependent variable changes from M_{bnt}/M_{bt} to M_{bnst}/M_{bst} . Once parameters from market share and model entry regressions

³⁵Fixed Q_{nst} can be rationalized as the limiting value of a nested model in which the “segment shock” has very large variance.

have been collected, $\mathbb{E}(\pi_{bnt}) = \sum_s \mathbb{E}(\pi_{bnst})$ is computed to yield an estimation of brand entry choices, that takes into account the underlying segmentation of the car market.³⁶

Table 5 provides segment-level estimates for the market share equation (upper panel) and model and brand entry (lower panel). We consider five segments: small cars, big cars, multi-purpose vehicles (MPVs, which includes minivans), sport utility vehicles (SUVs), and a combination of sport and luxury cars.³⁷ The most important set of coefficients obtained is the one on $\ln C_{bn}$, revealing η_s . The demand side elasticity for small cars, at -4.4 is reasonably close from the one obtained in the baseline regression (-3.87). The price response is larger for MPVs, big cars, and SUVs, all around -6 . The sport&lux η of just 0.458 implies an infinite markup over costs, which makes it impossible to include this segment in aggregate profit calculations or counterfactuals. The average η_s , excluding sport&lux, is -5.6 . This higher elasticity of market shares with respect to changes in C_{bn} will tend to increase responsiveness to tariff changes in the segment version of the model, offsetting to some extent the elimination of cross-segment substitution.

The segment specification exhibits a similar pattern of friction estimates to the ones obtained in Table 3. Market shares of home brands are higher in every segment; physical distance has a robust negative effect. Deep RTAs again lack significant effects on the intensive margin. However, the positive impact of deep RTAs on model and brand entry show up in the segment specification. The model entry effects are generally smaller in the segment specification whereas the brand entry effects of RTAs are estimated to be larger.

The model entry regressions have the expected negative signs on $\ln C_{bn}$ for each segment other than sport&lux. These coefficients are used to identify the dispersion parameters on fixed costs distributions for model entry (σ_s^e). In brand entry, column (6), it is the coefficient on $\ln \mathbb{E}(\pi_{bn})$ which plays this role for estimating σ^d . Both sets of dispersion parameters are reported in Table 6. The σ_s^e are larger for several segments but σ^d hardly changes when moving to segments. This table also calculates the structural parameters associated with all the frictions. A comparison with the corresponding frictions in the last three columns of Table 4 shows lower tariff-equivalents for home bias (ranging from 3–12% excluding the sport&lux outlier). There are higher effects for deep RTAs on brand entry fixed costs, and heterogeneous deep RTA effects on model entry across the segments, with SUVs and big cars having very large fixed costs of adding models.

6 Counterfactual methods

The counterfactuals investigate a set of trade policy scenarios involving shocks to tariffs and the deep RTA indicators. They treat as exogenous country-level new car purchases (Q_n), each brand's total number of models (M_b), and each brand's set of production locations, $\mathbb{L}_{b\ell}$. Since the data used

³⁶The precise steps to construct $\mathbb{E}(\pi_{bnst})$ are the same as in appendix B, with the appropriate s subscript when the segment dimension is relevant.

³⁷We based the categorization on a combination of information on the "Global sales sub-segment" (a functional categorization) and the "Global sales segment" (a size categorization) of the model specified in IHS original data. Our segments therefore represent roughly similar-sized sets of models, grouped by categories suggested by the industry consultancy from which we bought the data.

Table 5: Segment-level market share and market entry estimates

Dep. Var:	Average market share					
Method:	PPML					
Segment:	(1)	(2)	(3)	(4)	(5)	
	small cars	big cars	MPV	SUV	sport& lux	
home_{in}	0.545 (0.379)	0.383 (0.338)	0.191 (0.362)	0.553 (0.297)	1.21 (0.174)	
$\text{home}_{in} \times \text{LDC}_n$	0.198 (0.442)	-0.467 (0.749)	-0.366 (0.853)	-0.399 (0.31)	-1.143 (0.698)	
$\ln \text{dist}_{in}$	-0.382 (0.117)	-0.269 (0.123)	-0.259 (0.102)	-0.121 (0.104)	-0.155 (0.081)	
language_{in}	0.216 (0.179)	0.327 (0.147)	0.285 (0.216)	0.518 (0.19)	0.219 (0.105)	
Deep RTA_{in}	0.042 (0.193)	-0.127 (0.147)	-0.083 (0.182)	-0.171 (0.132)	0.165 (0.169)	
$\ln C_{bn}$	-4.442 (1.195)	-6.046 (1.045)	-6.257 (1.665)	-5.558 (0.795)	-0.458 (0.692)	
Observations	29564	26692	16415	29379	17225	
R^2	0.476	0.494	0.388	0.434	0.542	
Dep. Var:	Model entry (fraction)					Brand entry
Method:	Fractional probit					Binary probit
Segment:	(1)	(2)	(3)	(4)	(5)	(6)
	small cars	big cars	MPV	SUV	sport&lux	all
home_{in}	0.029 (0.079)	0.218 (0.09)	0.222 (0.131)	0.152 (0.084)	0.224 (0.074)	0.845 (0.471)
$\text{home}_{in} \times \text{LDC}_n$	0.779 (0.132)	0.076 (0.226)	-0.094 (0.257)	0.541 (0.195)	-0.563 (0.162)	1.088 (0.614)
$\ln \text{dist}_{in}$	-0.064 (0.022)	-0.07 (0.03)	-0.154 (0.033)	-0.085 (0.03)	-0.045 (0.041)	-0.065 (0.073)
language_{in}	0.088 (0.052)	0.04 (0.038)	-0.022 (0.041)	0.049 (0.041)	0.08 (0.049)	-0.2 (0.149)
Deep RTA_{in}	0.047 (0.04)	0.155 (0.057)	0.045 (0.092)	0.089 (0.046)	0.052 (0.063)	0.365 (0.118)
$\ln C_{bn}$	-0.953 (0.311)	-0.529 (0.403)	-0.733 (0.497)	-0.148 (0.333)	0.178 (0.499)	
$\ln \mathbb{E}[\pi_{bn}]$						0.554 (0.082)
Observations	27871	24864	15838	25794	15916	113207
R^2	0.747	0.681	0.757	0.711	0.464	0.650

Standard errors bootstrapped with brand b clusters over 1000 replications (see footnote 28). R^2 is the squared correlation of fitted and true dependent variables. Each regression controls for log per-capita income and price level of the assembly country.

Table 6: Friction parameters (segment-level)

Friction:	δ	δ^e	δ^d
Estimate:	\mathbf{f}	\mathbf{f}^e	\mathbf{f}^d
Small cars ($\eta = 4.44, \sigma^e = 3.61$)			
home	-0.123	0.318	-1.524
home \times LDC	-0.045	-2.660	-1.964
ln distance	0.086	-0.066	0.118
common language	-0.049	-0.149	0.361
RTA (deep)	-0.010	-0.138	-0.659
Big cars ($\eta = 6.04, \sigma^e = 9.54$)			
home	-0.063	-1.757	-1.524
home \times LDC	0.077	-1.119	-1.964
ln distance	0.044	0.443	0.118
common language	-0.054	-0.112	0.361
RTA (deep)	0.021	-1.584	-0.659
Multi-purpose vehicles ($\eta = 6.26, \sigma^e = 7.17$)			
home	-0.031	-1.427	-1.524
home \times LDC	0.058	0.368	-1.964
ln distance	0.041	0.885	0.118
common language	-0.046	0.397	0.361
RTA (deep)	0.013	-0.395	-0.659
Sport utility vehicles ($\eta = 5.56, \sigma^e = 30.73$)			
home	-0.100	-4.225	-1.524
home \times LDC	0.072	-16.950	-1.964
ln distance	0.022	2.517	0.118
common language	-0.093	-1.074	0.361
RTA (deep)	0.031	-2.872	-0.659
Sport and luxury cars ($\eta = 0.46, \sigma^e = 3.05$)			
home	-2.641	-2.115	-1.524
home \times LDC	2.496	3.071	-1.964
ln distance	0.339	0.321	0.118
common language	-0.478	-0.502	0.361
RTA (deep)	-0.359	-0.355	-0.659

Elasticities used to obtain frictions: (1) the segment-specific ones are given in the table, (2) for brand entry, $\sigma^d = 1.80$. Calculation of those frictions are described in the text and use coefficients from Tables 3 and 5.

for each variable comes from a single year (2016, the last available in our sample) we suppress the time subscripts in this section. We start by highlighting the features of our counterfactuals that require further elaboration.

6.1 Three methodological considerations

The first such feature is our treatment of country-level external returns to scale. Our counterfactuals solves for the equilibrium for a given q_ℓ and then updates q_ℓ to deliver new unit costs and a new equilibrium. This iteration continues until a fixed point is reached. There is concern over existence and uniqueness of equilibria with increasing returns. However, our simulations suggest that for our parameter values, there is a unique fixed point. Kucheryavyy et al. (2016) find that a sufficient condition for uniqueness is that the trade elasticity multiplied by the scale elasticity ($-\zeta\theta$ in our notation) should be less than one. Our estimates $\theta = 7.7$ and $\zeta = -0.035$ imply $-\zeta\theta = 0.27 \ll 1$, suggesting a unique equilibrium. As we show in Appendix D, $-\zeta\theta < 1$ is a sufficient solution to avoid explosive outcomes, which is necessary for existence and uniqueness of an internal equilibrium. To isolate the impact of increasing returns and quantify the importance of interdependencies, we also conduct a non-IRS version of the counterfactuals. That setting treats q_ℓ as another proxy for local assembly costs (w_ℓ^α), along with GDP per capita and the exchange rate over-valuation index. Thus it is held constant at the observed level even as the policy variables are changed. We focus on the IRS results but comment on some interesting differences relative to the non-IRS case.

The second methodological aspects of our counterfactuals to be detailed regards the treatment of segments. We approach the issue of market segments with two boundary assumptions in order to ensure that how we handle demand substitution patterns does not exaggerate the response to policy changes. The unified car market assumption makes every car model a symmetric substitute for every other model. In contrast, the segmented market assumption shuts down between-segment substitution by fixing the Q_{ns} at the 2016 levels. To see how this could matter in counterfactuals, consider the response of Smart production in France to the ending of NAFTA preferences. The C'_{bn} for Smart's rivals who produce in North America will rise whereas Smart's C_{bn} will be unchanged. Therefore Smart's sales are expected to rise. The difference is that, under segmented markets, Smart achieves a higher market share in small cars, a relatively unimportant segment in the US, whereas in a unified market Smart also gains at the expense of North American SUV production. This will have aggregate effects on French car production since the brands that produce in France and also have a distribution presence in the US all make small cars.

A third important aspect of the counterfactuals is the method of solving for changes relative to the factual policies. The method we report in the main text solves the full model under the current set of frictions, computing expected values of sales, sourcing shares, model-level and brand-level entry decisions. The same set of calculations is repeated under the counterfactual set of frictions to obtain a new set of expected values. We refer to the comparison of predictions under current and alternative sets of frictions as Difference in Expected Values (DEV). DEV requires estimates of

662 parameters.³⁸ An alternative approach, described in detail in Appendix C, is called Exact Hat Algebra (EHA). The EHA approach requires just 10 parameters ($\eta, \theta, \sigma^e, \sigma^d, \alpha, \varsigma$, and the 4 friction parameters for deep RTAs). EHA allows observed levels of the endogenous variables to “stand in” for parameter estimates as well as unobservables. Thus, by definition, EHA replicates the actual data, whereas DEV sometimes errs by large amounts in predicting the factual levels of production in each country (q_ℓ).³⁹ EHA has two important disadvantages. The first is that brand entry cannot be handled by this method because it is a binary variable. The second concern in using EHA is that it does not allow a brand to start sourcing from an assembly country that was not used prior to the shock. Any zero remains zero, no matter how large the change in frictions. Appendix I provides a full set of results for this alternative approach to counterfactual simulations, together with a section discussing when and why EHA and DEV results differ.

6.2 The solution algorithm

Our DEV approach solves the model in levels twice: once at the factual level of frictions, and once for the same frictions evaluated under the counterfactual scenario. The equations summarizing the equilibrium are the sourcing decision (6), the market share (9), the price index (10), and the two entry equations (15) and (19). The identification of structural parameters needed for those equations is detailed in section 4.3 and in appendix B. Solving the model involves nested fixed point iterations with an inner, a middle, and an outer loop.

1. The inner loop solves a system of two non-linear equations obtained from the price index (10) and the model entry probability (15). It takes as given the vectors of brand entry (\mathbb{D}_{bn}) and the multinational cost index, C_{bn} , which is determined by the set of frictions and national production, q_ℓ . Fixed point iteration yields equilibrium values of P_n and M_{bn} .
2. The middle loop takes these two variables and feeds them into expected market shares, equation (9). Combined with sourcing probabilities, $\text{Prob}(\mathbb{S}_{b\ell n} = 1)$, from equation (6) the vector of equilibrium flows is given by

$$\mathbb{E}[q_{b\ell n}] = \mathbb{E} \left[\frac{q_{bn}}{Q_n} \mid \mathbb{D}_{bn} = 1, M_{bn} \right] \times \text{Prob}(\mathbb{S}_{b\ell n} = 1) \times Q_n. \quad (26)$$

Next, we sum over $\mathbb{E}[q_{b\ell n}]$ for all b and n to obtain the expected value of q_ℓ , then used to update C_{bn} and $\text{Prob}(\mathbb{S}_{b\ell n} = 1)$. The inner loop (step 1) is re-run with these new inputs to output a new vector of quantities. The process iterates until the vector of q_ℓ stops changing.

³⁸There are 10 elasticities, 23 structural friction parameters, 47 production country FEs ($\text{FE}_\ell^{(1)}$), three times the 74 market FEs ($\text{FE}_n^{(2)}$, $\text{FE}_n^{(3)}$ and $\text{FE}_n^{(4)}$) and three times the 120 brand FEs ($\text{FE}_b^{(2)}$, $\text{FE}_b^{(3)}$ and $\text{FE}_b^{(4)}$). DEV is feasible here because all the parameters in our model are identified.

³⁹EHA was developed by Dekle et al. (2007) and Arkolakis et al. (2012). Eaton et al. (2013) is an example of a paper using DEV even though two of the authors helped to originate the EHA approach. The reason was that the later paper abandoned the continuum assumption to consider granular data of the type that is important here as well.

3. The brand entry vector is then updated in an outer loop using the rule that entry occurs when expected profits exceed fixed entry costs, $\mathbb{E}[\pi_{bn}] > F_{bn}^d$, with expected profits and fixed costs calculations being detailed in Appendix B. Since the inner and middle loops depend upon this vector of entry, the algorithm iterates over the three loops until the set of profitable brand-market combinations stabilizes.

Handling the decisions of brands to enter or not in counterfactuals is far from straightforward. The theory specifies the distribution of brand entry fixed costs as log-normal. If we drew fixed costs from the complete distribution, there would be a large number of instances of brands that are in fact available in a country but would be absent even in the *factual* version of the solution of the model. Conversely, there would also be many false entrants in the simulation. If a major brand were falsely absent or present it would severely endanger the realism of the counterfactual results. We therefore follow König et al. (2017) in drawing from a distribution that is truncated such that fixed costs of factual entrants are *lower* than our predicted value of their gross profits. Brands that are absent from a particular market (Renault in the US for example), take their fixed cost draws from the portion of the distribution where fixed costs *exceed* gross profits. The counterfactual policies maintain the exact draw of fixed costs for each brand-destination but recomputes expected gross profits. This allows for some brands to be drawn into or out of its factual entry decision. To obtain expectations we replicate the solution of the model with 1000 draws of fixed costs.⁴⁰

The discrete nature of brand entry has to be taken into account when computing the equilibrium. In contrast to entry at the model level, where the only relevant object is the share of models the brand decides to offer, the identity of which brand enters matters for the outcomes of the counterfactual. This is because each brand has its own network of potential assembly locations ($L_{b\ell}$) and its own mass of models (M_b) with associated productivity φ_b . This means we have to keep track of which particular brands have entered or exited as a result of a policy change. The algorithm iterates until a fixed point in the brand-entry vector is reached. At each iteration, the computation of expected profits takes the price index for each market (calculated in the inner loop where brand presences are held fixed) as a given. In practice, entrants do affect the price index. In the small car markets (e.g. Bulgaria, Ukraine) this decline in the price index can be large enough that it induces exit in the following iteration. This leads to a rise in the price index which can attract firms back into the market. This process repeats itself in an oscillatory pattern, at which point we terminate the iteration. Note that this oscillation only occurs because brand-level entry has to be considered as an integer problem. The expected fraction of models offered in a market ($\mathbb{E}[M_{bn}/M_b]$) is a continuous variable, so there is no integer issue in the inner loop.

The segmented market version of DEV has a few important differences. In step 1 (the inner loop), price indices, the mass of models offered, and all three marketing costs (δ , δ^e and δ^d) need to

⁴⁰ A subtle aspect of this approach is that it does not actually guarantee that entry choices in the simulation of factual policies match the real decisions 100% of the time. This is because the truncation is based on the gross profits calculated using price indexes extracted from the fixed effects in the market share regression. The P_n obtained in the solved model will differ, occasionally by enough to alter the entry decision of individual brands that were near the entry/exit threshold.

be defined at the segment level (the sourcing probabilities, as well as τ and γ frictions do not have a segment dimension). The inner loop therefore solves for equilibrium P_{ns} and M_{bns} . The middle loop (run as a second step) updates national output by summing up the segment-level sales of the brand that are expected to be sourced from different origin countries:

$$\mathbb{E}[q_\ell] = \sum_b \sum_n \mathbb{E}[q_{b\ell n}] = \sum_b \sum_n \text{Prob}(\mathbb{S}_{b\ell n} = 1) \sum_s \mathbb{E} \left[\frac{q_{bns}}{Q_{ns}} \mid \mathbb{D}_{bn} = 1, M_{bns} \right] \times Q_{ns}. \quad (27)$$

Lastly, the decision of the brand to enter a market depends on the sum of the profits to be earned in each segment where the brand has models. Therefore, the outer loop updates the vector of brand entry decisions which are transformed as $\mathbb{E}[\pi_{bn}] = \sum_s \mathbb{E}[\pi_{bns}] > F_{bn}^d$.

Before turning to solutions of counterfactual policies, it is important to demonstrate that the endogenous variables solved for under factual trade policies do not depart too much from their data counterparts. Starting with the decision where brands are offered, the entry rate in fact is 36.2%, slightly lower than the 36.6% average in the simulation with a maximal difference of 0.82%. Figures G.1 and G.2 in Appendix G show the fit of one run of the DEV simulation (for unified and segmented markets respectively) under the factual set of policies. The correlations between simulated and actual are high for all four variables examined. In the unified markets with IRS case, we obtain the following correlations: 0.98 for the price index, 0.63 for brand-origin-destination sales, 0.74 for origin-destination flows and 0.86 for aggregate origin-level output. The correlations are even higher for the non-IRS and segmented versions of the model.

7 Counterfactual results

The main motivation for estimating the model of this paper is the investigation of counterfactual trade policy changes. We report results on eight different scenarios: two different versions for each of four types of policy experiments.

1. Trumpit/Section 232:

- (a) The USA imposes a 25% tariff on cars and parts from Canada and Mexico and ends NAFTA (a deep RTA). Canada and Mexico retaliate with equal tariffs.
- (b) The USA imposes the same 25% tariffs on all origins *except* Canada and Mexico. The targeted countries reciprocate.

2. UK exit from the European Union:

- (a) Soft Brexit: a free trade agreement retains tariff-free trade between the UK and the EU27 but rescinds all deeper integration measures.
- (b) Hard Brexit: in addition to rescinding the deep integration measures, the EU27 and UK impose the current EU MFN tariffs.

3. Trans-Pacific integration with and without the US:

- (a) TPP: a deep integration agreement between Australia, Brunei, Canada, Chile, Japan, Malaysia, Mexico, New Zealand, Peru, Singapore, the United States, and Vietnam.
- (b) Comprehensive and Progressive Transpacific Partnership (CPTPP): a revised TPP agreement omitting the US (in force since December 2018).

4. Trans-Atlantic Integration:

- (a) Comprehensive Economic Trade Agreement (CETA): deep integration between the EU28 and Canada (provisionally applied since September 2017).
- (b) CETA + Transatlantic Trade and Investment Partnership (TTIP): EU28 deep agreements with Canada *and* the US.

Figures 6, 7, 8, and 9 display predictions of the model for each of the four policy experiments. Each figure contains two panels showing changes in output on the left and percent changes in consumer surplus on the right. For each of the ten most affected countries, we plot outcomes for unified (circles) and segmented markets (squares) versions of the model. We contrast the two policy variants by showing one in red and the other in blue. The counterfactual underlying these figures solves the model to calculate the difference in expected values (DEV). We also present results using the EHA method for counterfactuals in appendix I. All numbers used for graphical displays of both DEV and EHA versions of counterfactual scenarios can be found from the detailed Tables in appendix J.

7.1 Trumpit/Section 232

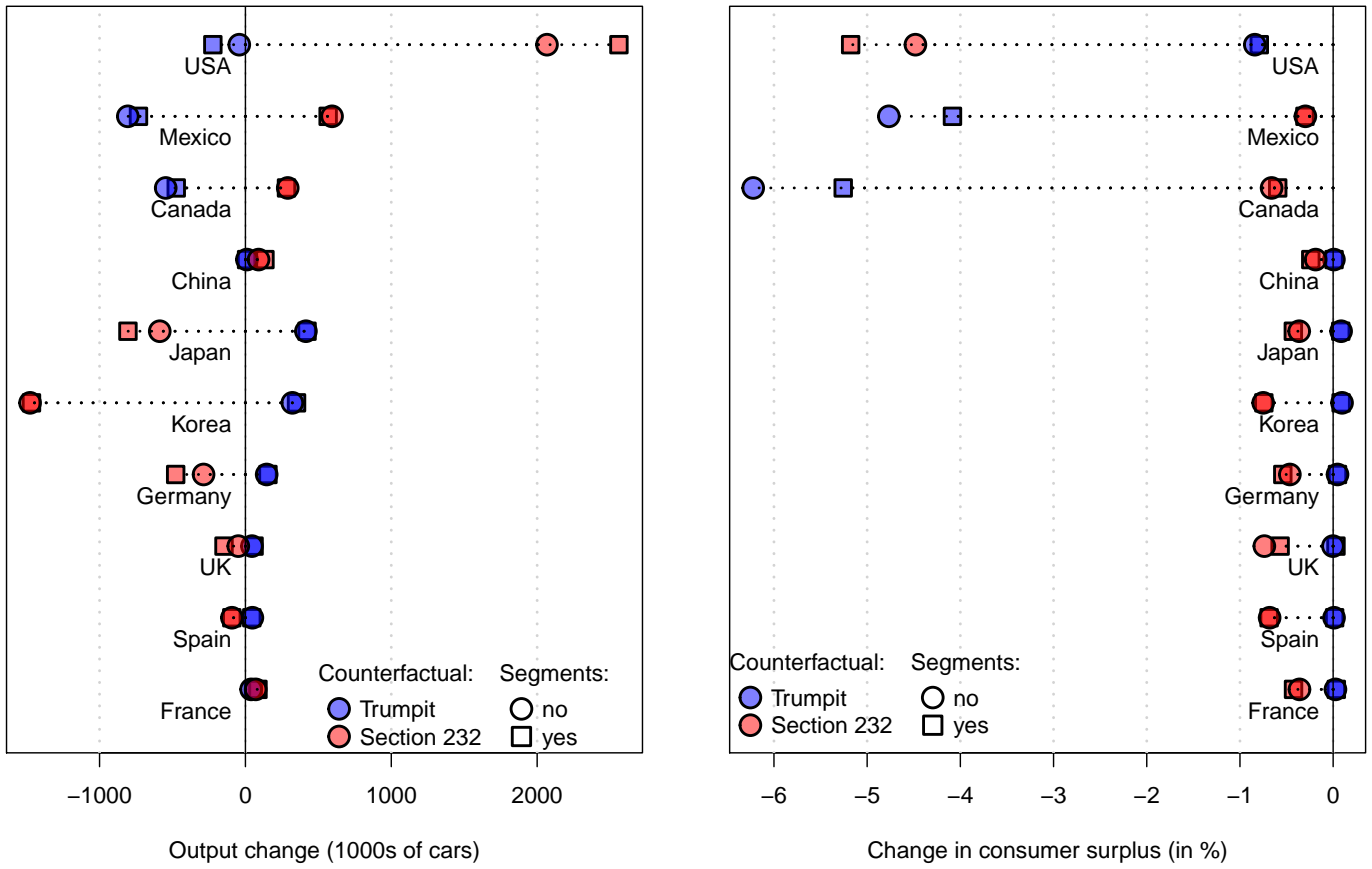
Renegotiation of the North American Free Trade Agreement began shortly after the inauguration of Donald Trump. In the final stages of the negotiation, the US president threatened to leave NAFTA and impose 25% tariffs in the auto sector under Section 232 national security provisions. We refer to this combined threat as “Trumpit.”

Figure 6 depicts in blue the imposition of 25% tariffs on both assembled cars and parts from Canada and Mexico. The latter are assumed to retaliate with identical tariffs on the US (while maintaining full NAFTA preferences with each other). Our simulations point to disastrous outcomes for Canada and Mexico, whose combined losses sum to 1.2–1.4 million cars.⁴¹ These losses would be equivalent to shutting two median-sized plants in Canada (producing 243 thousand cars per year) and five of the smaller (150 thousand) plants operating in Mexico. Production also declines in the US but the loss of 42–223 thousand cars represents less than one percent of the initial level. Canada stands to lose 67–68% of its car industry, while Mexico loses 38–41% of its production.⁴²

⁴¹As the true amount of substitution between segments lies between the segmented and unified polar cases, we report ranges of outcomes except when they round to the same value.

⁴²The percentage changes in production are all based on the initial levels predicted by the model.

Figure 6: Trumpit / Section 232



Why do Canada and Mexico fare so badly under Trumpit? The first factor is that the shock is very large, including both the final and upstream tariffs and also the sizeable tariff equivalents of lost deep integration. In total, τ frictions rise to 28.2% while γ frictions rise to 31.4%. The shock to τ applies to large shares of Canadian and Mexican production. In 2016 the US purchased 83% of Canadian-made cars and 55% of Mexican cars. In contrast, those countries' combined purchases amounted to only 8.4% of the cars assembled in the US. The shock to the γ friction also applies broadly in Canada and Mexico. Our estimates imply that parts from the US account for 37% of the assembly costs of the US brands who account for 55% of Canadian and 36% of Mexican production in the 2016 data. The final key fact explaining our simulation predictions is that almost all the brands made in Canada (11 of 12) and Mexico (10 of 14) are also made in the US. The large estimated sourcing elasticity ($\theta = 7.7$) implies these brands will heavily shift production for the US market to US assembly sites when Trumpit raises the production and delivery costs of using sites in Canada and Mexico.

The asymmetric damage to the factories of US brands in Canada and Mexico from higher γ frictions is absent in a pure trade model. The Canadian and Mexican governments would likely realize that putting tariffs on imported parts from the US is a defective way to retaliate. To inflict the greatest damage on US production with the lowest losses to Canadian production, it would make more sense to limit retaliation to final goods. To investigate what part of the losses seen in the Trumpit scenario come from avoidable cost increases, we ran a version including this modification. Dropping the retaliatory tariffs on car parts does indeed shrink the losses to Mexican and Canadian production by 104,000 and 76,000 cars, respectively in the segmented case.⁴³

The blue symbols in the right panel of Figure 6 show that car buyers in all three countries lose from Trumpit. Losses in the US are minor (about 1%) but Canada loses 5–6% of consumer surplus while Mexico loses 4–5%. Consumer losses in both countries can be partially mitigated by “smart retaliation”: exempting parts lowers the losses by about 0.4% for both countries.

The Trumpit counterfactual hurts all three participants, both in terms of reduced production (cumulating to 1.4 million fewer cars made in North America) and lower consumer surplus. The disproportionate losses by Canada and Mexico help to explain why they were willing to agree to NAFTA revisions that were widely perceived as unfavorable. While the members of NAFTA lose from Trumpit, the seven other countries shown in Figure 6 gain both in terms of output and consumer surplus. This is specially true for Japan, Korea, and Germany, who collectively increase production by around 900 thousand cars. Since our estimates rank those three countries among the four lowest cost assemblers, the reallocation of sourcing decisions for the US market heavily favors them.

The changes in production and consumer surplus displayed in Figure 6 take into account adjustments in sourcing, market shares, model entry, and brand entry. The last of these is of particular interest because it is a binary outcome that conditions the three subsequent decisions in any given replication. The simulations calculate the probability of opening a distribution network over

⁴³With unified markets, Mexican and Canadian production losses contract by 119,000 and 87,000 cars.

1000 repetitions with and without the policy intervention. We find that Trumpit does not change entry probabilities for the major brands. However, US brands that serve few foreign markets substantially lower their likelihoods of entering Canada and Mexico. The most extreme case is Buick, which enters Mexico in every replication without Trumpit but only in 54% of the Trumpit replications. In contrast, a number of second-tier EU brands increase their entry propensity in both markets. For example, Citroën’s chance of entering Mexico rises from 11% to 27% due to Trumpit. Skoda, Dacia, and Opel more than triple their brand entry probabilities in Mexico. In practice, we find that brand entry issues do not have much quantitative impact because the brands that are on the margin of entering and staying account for very little output and also have limited sourcing alternatives.

The NAFTA revision (agreed upon in October 2018) exempted shipments originating in Canada and Mexico of up to 2.6 million cars each from any subsequent Section 232 tariffs. At the time of writing, this threat remains in effect for other countries, which motivates our second variant (labeled Section 232). The red symbols in Figure 6 apply to this drastic scenario, in which the US imposes 25% tariffs on cars and parts imported from all non-North American partners, who retaliate with equivalent tariffs.⁴⁴

US imposition of Section 232 tariffs on non-North American car imports would cut production in Japan and Korea by large amounts, with German production also sharply reduced. Their combined losses sum to 2.4–2.8 million cars. The beneficiaries would be North American plants, with those in the US experiencing the largest increase in production (up to 2.6 million cars added in the segmented case). The large rises in predicted exports of Mexico and Canada to the US, 416–431 ths. and 228–230 ths. respectively, may explain the US insistence on limiting the number of car imports to be exempted from Section 232. Consumers lose surplus in all countries, with the price index faced by US buyers rising by 4.5–5.2%. The losses for Mexican and Canadian consumers—who do not impose any tariffs in this scenario—stem mainly from rising production costs at the factories of Asian and EU brands in the US, where car parts rise in cost by 25%.

Increasing returns to scale play an important role in the counterfactual predictions displayed in Figure 6. The estimated magnitude of scale economies is modest: a doubling of national output reduces costs by just 2% ($2^{-0.035} = 0.976$). However, the large tariff response elasticities lead to substantial *amplification* of production changes. To quantify the IRS-induced magnification effects, we re-solve the counterfactual holding domestic production (q_{lt}) inside the cost function constant at the factual levels. This is equivalent to treating country scale as an omitted cost term that does not vary in response to policy changes. IRS generates much larger predicted output losses for Mexico and Canada in the Trumpit scenario (8 and 11 percentage points larger, respectively, for unified markets). The amplification effects turn out to be strikingly systematic. Regressing the Trumpit country-level change in log output in the IRS case against the change in log output for the non-IRS case, we obtain a coefficient of 1.36 (for segments and unified markets). The fit (R^2) of this

⁴⁴A few countries (India, Iran, and Vietnam, for instance) have tariffs larger than 25% in 2016. The simulation leaves those tariffs unchanged.

simple linear regression is 0.99. As explained in Appendix D, a simplified version of our model predicts the amplification coefficient to be $1/(1 + \varsigma\theta) = 1.37$. Similar values of the amplification effect show up in all our counterfactuals.

The presence of increasing returns exerts a second, more subtle, effect on the conduct of the counterfactual responses: it generates interdependencies across markets. China's production changes in response to Section 232 tariffs provide a useful illustration. China's increase in shipments to the domestic market under IRS (122–168 ths. cars) is more than double its increase in the constant returns simulation. This occurs in part because nearby suppliers Japan and Korea experience rising production costs under IRS because of their diminished scale (shown in red in Figure 6). Why do China's domestic shipments rise (by 56–74 ths. cars) even in the absence of IRS? Chinese car tariffs do not change because our experiment imposes 25% retaliatory duties across all countries and China's tariffs were already at 25%. The production increase in the home market stems from a second interdependency caused by γ effects. For example, BMW exports SUVs from the US to China. These have substantial German content (37% according to our estimates) which implies that higher Section 232 tariffs on car parts raise the cost of German cars exported to China. Thus, even without higher Chinese tariffs on such cars, non-US brands assembled in the US become less competitive in China compared to locally assembled cars. All in all, such brands export 45% fewer cars to China under the Section 232 tariffs (without IRS effects). In sum, the two interdependencies, one from IRS and the other from HQ-sourced parts, lead to the interesting prediction that US section 232 tariffs imposed on non-North American producers would actually expand the number of cars assembled in China.

Figure 6 shows that, contrary to our initial expectations, the segment version of the model does not systematically dampen the impact of changes in trade policies. This is because switching from unified markets to segments has three effects. First, it limits substitution within segments, which acts as a dampener. Working in the opposite direction, the demand elasticities are larger for segments than for unified car markets. This implies a greater market share response to policy changes. Another factor that increases responses is that segmenting tends to reduce the number of competing brands. The net effect is larger absolute impacts in the segmented version for the US and Japan under Trumpit and Section 232 but slightly smaller effects for Canada and Mexico. As the segmented and unified magnitudes tend to be similar, we comment only on the segmented results in the discussions of the rest of the policy experiments.

7.2 Brexit

Since the 2016 Leave vote, debate revolved around whether Brexit will be “soft” or “hard.” The soft Brexit case captures the scenario in which Britain retains tariff-free access to the EU but loses the deep integration aspects of the RTA such as free mobility of professionals and the ability to influence EU regulations on car standards. We also simulate a hard Brexit scenario where UK exports face the EU's 10% MFN tariffs while the UK reciprocates at the same rates. In both scenarios the UK cannot “roll over” existing EU trade agreements (with South Korea, Mexico and Turkey

notably), so trade with all those countries reverts to MFN tariffs. The two post-Brexit scenarios set all the deep RTA dummies to zero if they correspond to (1) dyads involving the UK and EU, (2) a dyad involving the UK and one of the countries having a preferential deep RTA with the EU.

Figure 7: Brexit

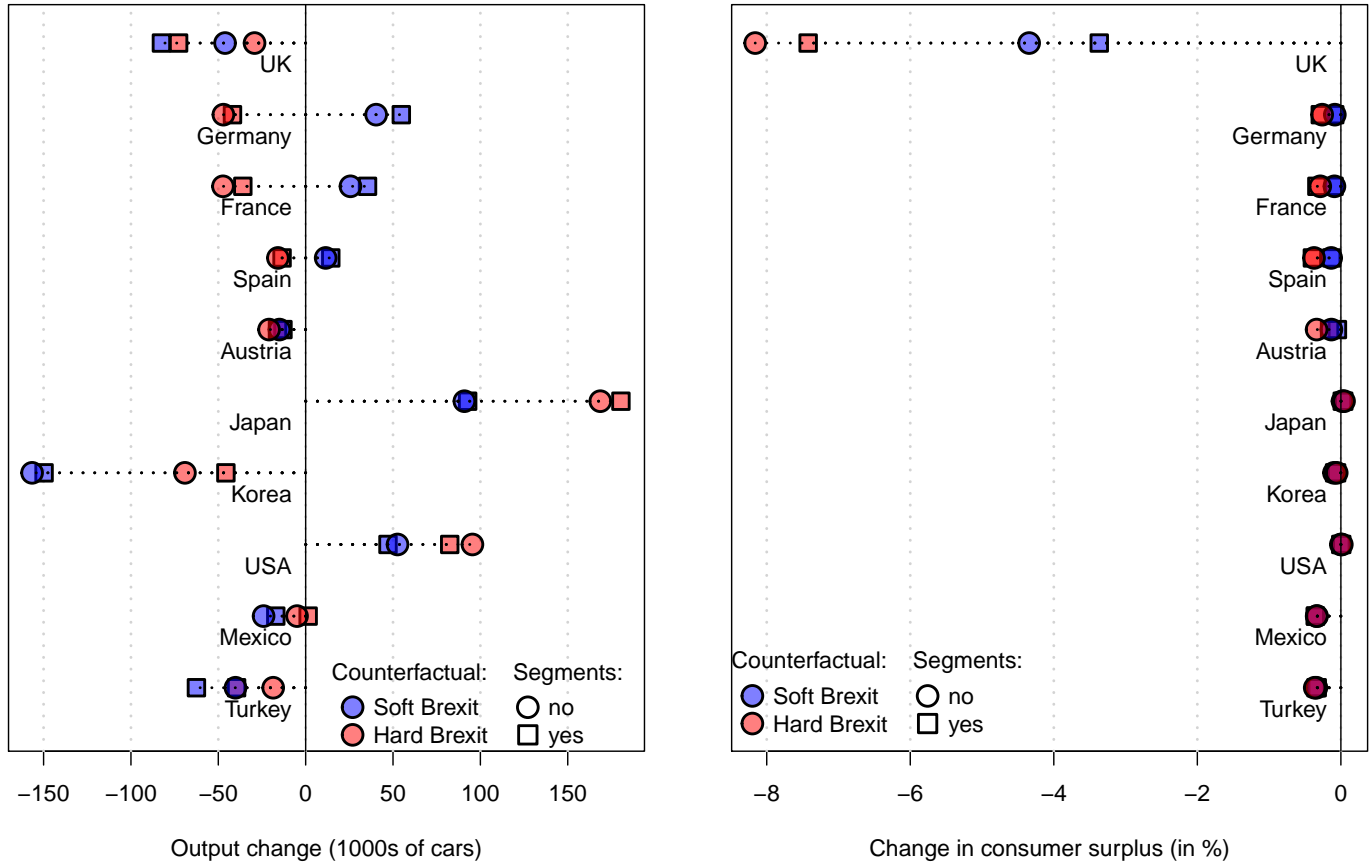


Figure 7 points to poor outcomes for both the UK car industry and the British car buyer. Workers in UK plants lose slightly more (83 ths. cars) under Soft Brexit than with Hard Brexit (73 ths. cars). To put those numbers in perspective, they amount to losses of between 2/3 and 3/4 of Honda’s Swindon factory, which in 2016 assembled 110 thousand cars employing around 3100 workers.⁴⁵ The relatively small net production losses in the UK mask large changes in the “gross flows.” Under Hard Brexit, UK factories increase production for the local market by 290 ths. cars, while their sales on the continent fall by 259 thousand. Sales outside the EU are predicted to fall by 103 ths. cars because UK-made cars lose their tariff-free status in Turkey, Mexico, and other countries where the EU signed FTAs.

The protective effect of 10% tariffs more than offsets the lost exports to the Continent under Hard Brexit. This prediction contrasts with the Trumpit prediction that Canada loses sales in both the US and at Home. Two factors underlie this difference. First, US brands make the majority of

⁴⁵*Financial Times*, March 13, 2015. Swindon is an interesting plant because it is about average in size and accounts for the majority of Honda sales in the EU.

cars in Canada and their plants experience production cost increases (γ effects) under Trumpit. Falling US-brand production even harms the Japanese makers, through Canada's loss of external scale economies. Meanwhile in Britain, domestic brands accounted for 45% of production, whereas EU brands account for just 8% of production (2016 data). A second factor is that the EU27 runs a substantial trade surplus with the UK in cars. This means that as tariffs in both directions rise to 10%, the UK has less to lose.

While these simulation results do not predict disastrous consequences to auto workers from a "no deal" Brexit, the cost to consumers is much more severe. Consumers lose 3.4% of their surplus with Soft Brexit but their losses more than double to 7.4% when 10% tariffs are imposed on EU-made cars. The tariffs imposed on UK cars exported to the EU have very little effect on consumer surplus there. One problem the UK faces is the fact that major multinationals like Ford and Nissan have factories in Spain they can easily switch production to in order to continue to serve the EU27 tariff-free. On the other hand, none of the major EU brands operate factories in the UK, implying that British buyers will have to pay more for their cars.

7.3 Trans-Atlantic Integration

While recent reversals on trade integration have dominated attention, some new deep integration agreements have recently been implemented. The Comprehensive Economic and Trade Agreement (CETA) between Canada and the EU28 still awaits ratification by different European national parliaments, but it has been applied provisionally since September 21st, 2017. By contrast, negotiations on the Transatlantic Trade and Investment Partnership (TTIP), an integration agreement between the EU and the US, were put to an indefinite halt following the 2016 US elections. We therefore consider the first scenario to be the full implementation of CETA: the abolition of tariffs between the EU and Canada, combined with deep integration. On top of those, the second scenario applies the same policy changes to the EU28-USA country pairs.

Figure 8 shows that most gains in production under CETA accrue to Canada (56 ths. cars, 8% production increase), followed by Germany (43 ths. cars) and Britain (25ths. cars). The countries losing sales are the US, Korea, Japan, and Mexico (in that order). Korea and Mexico face erosion of their existing preferences granted by the EU and Canada as part of pre-existing FTAs. The only country where consumers have notable gains is Canada, where surplus rises by 1.8%. Canadian production represents too small a share of sales in Europe even after CETA to make much difference for EU customers.

Adding the US to the trans-Atlantic liberalization of car trade would lead to potentially large output gains for Germany (7%), Spain (16%), and Poland (74%).⁴⁶ These countries increase production by 795 ths. cars. Interestingly, their losses at home from lowering tariffs on US-made cars are negligible: Germany loses sales of just 31 ths. cars at home and Poland and Spain actually *increase* production for their home market. These gains come from a mix of increasing returns and

⁴⁶Appendix H provides evidence that such large production increases are feasible within the medium run contemplated in our policy counterfactuals.

γ effects (US-owned factories are important in both countries). Meanwhile production falls in the US as they cede more sales in their home market than they gain in Europe. This asymmetry is initially surprising since it is the EU that lowers its tariffs more under trans-Atlantic free trade (10% vs 2.5% MFN tariffs). The US losses can be understood in terms of the tendency of trade liberalization to reallocate output to the low-cost producers. In the prediction of our model—and in the data—the EU collectively has a cost advantage, exporting about six times more to the US than what it imports.⁴⁷

Consumers gain widely, but modestly, from trans-Atlantic integration. Gains in Germany, the US, Spain and Britain all round to 1%. Canadian buyers actually gain more (2.3%) *with the US* in the agreement. This is because the German brands VW, Mercedes, and BMW all have manufacturing plants in the US. Deep integration implies a 6% reduction in the production costs of those plants (γ effects), which are also cheaper to distribute in Canada as a result of the agreement (δ effects).

7.4 Trans-Pacific Integration

Figure 9 displays the predicted impact of the Transpacific Partnership, showing outcomes for the original agreement involving the US and its successor, the CPTPP, which came into force December 30, 2018. The TPP is mainly of interest as a “might-have-been” policy since the US exited the agreement in January 2017. Despite claims that TPP was more of a regulatory agreement than a trade agreement, the TPP and CPTPP include substantial tariff cuts for some country pairs. In 2016 Japanese exporters faced a 44% tariff when exporting to Vietnam, and 6–7% to Canada and New-Zealand. US exports faced a 55% duty in Vietnam and 23% in Malaysia.⁴⁸

Transpacific integration shows how policy impacts on the location of production and the fortunes of individual brands are intertwined. In terms of production, the most affected country is Canada. It is predicted to increase output by 33% under TPP and by 42% under CPTPP. In terms of brand nationalities, Japan is the primary beneficiary of both TPP and CPTPP. Both policies increase the share of the Canadian market served by Japanese brands by about ten percentage points. Under TPP, Japanese brand market shares rise by six percentage points in the US and they even rise a small amount (one percentage point) under CPTPP.

Several forces at work underlie these predicted reallocations in production and market shares. They represent aspects of our framework that do not feature in traditional models. Japanese plants, which account for 45% of Canadian production in 2016, are predicted to obtain a 6.4% cost reduction as a result of the deep γ_{il} aspect of the agreement. Elimination of the 2.7% Canadian tariff on auto parts (an average across the main auto parts HS codes) will further reduce costs at Toyota and Honda’s Ontario plants. Under TPP, similar gains would also accrue to the Japanese plants in the US, although US parts tariffs are only 1.4% on average.

⁴⁷The model predicts an export/import ratio of 6.2, very close to the 5.6 ratio in the 2016 data.

⁴⁸Japan signed a free trade agreement with Malaysia that entered into force in July of 2006, and stipulated a 9-year phase-out period for cars.

Figure 8: CETA (EU-Canada deep RTA)


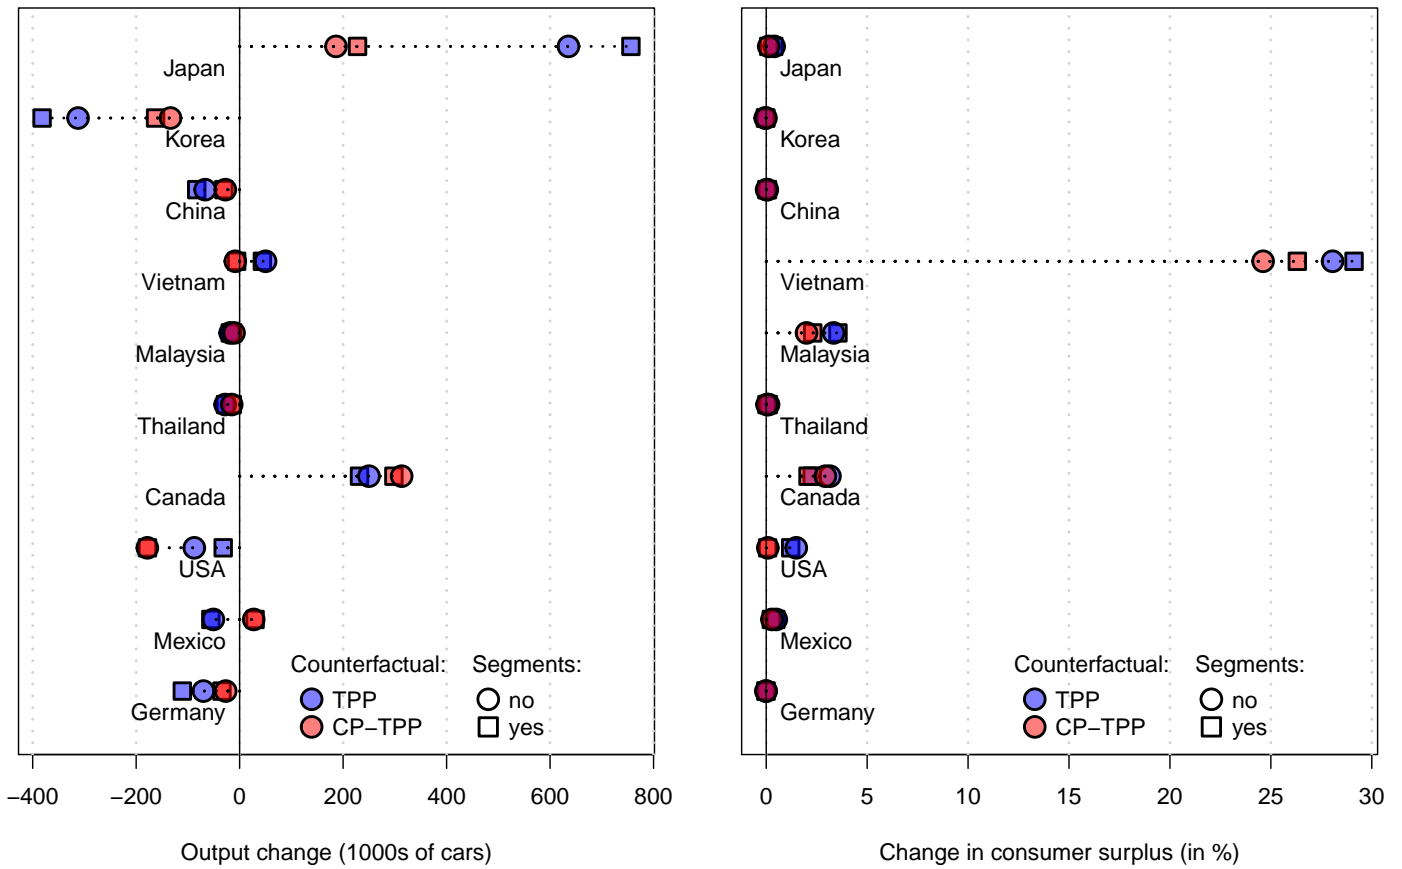


Figure 9: Transpacific Partnership



Even though the γ effects are mainly balanced, under TPP the reduction in marketing costs from Japan leads to increased market shares for Japanese brands in the US with much of the additional demand sourced from Canada. Under CPTPP, the situation is different. Although Japanese brand shares do not rise much in the US, they increase the fraction of models sourced from their Canadian plants. Meanwhile, because Japan-origin cars avoid Canada's 6% car tariff, Japan's sourcing share rises in the Canadian market. The total number of Japanese cars sold in Canada rises enough to offset this negative effect.

Lower costs of marketing Japanese cars are predicted to stimulate entry of nearly 20% more Japanese car varieties in Canada under the CPTPP, and would stimulate a similar increase in the US had the TPP been implemented. Since a significant fraction of these new models will be sourced from the Toyota and Honda plants in Canada, model entry stimulates production gains. There are also Japanese brands that would have been unlikely to enter Canada but are more likely to do so under CPTPP: Acura, Scion and Infiniti increase their entry probability in Canada by a substantial range (19%, 18%, 8%).

The γ gains our model predicts for Canadian production are not present in the conventional trade model used by Global Affairs Canada to predict the consequences of CPTPP.⁴⁹ This study finds close to negligible effects of CPTPP on Canadian output. Another study conducted in 2012 by Van Biesebroeck et al. (2012)⁵⁰ looked at the impact of several RTAs for automobile production in Canada. While their use of detailed data on car characteristics allowed the authors to use the whole apparatus of BLP-type demand estimation, they do not account for the γ -type frictions we include here. This is critical in terms of predicted outcomes. When we run our CPTPP counterfactual without any γ liberalization, we find a very small overall impact on Canadian output, confirming the results of the two mentioned studies.

The other country strongly affected by TPP and CPTPP is Vietnam. Vietnam increases production with the TPP scenario. This is because of improved access to the US market and more efficient operations at Ford and Chevrolet plants which are expected to stimulate a big increase in exports to the US. However, the Japanese brands currently assembling in Vietnam will radically increase their sourcing from Japan. If the US is not in the agreement, the net effect on production becomes negative for Vietnam. The brand entry margin is also very active in Vietnam under CPTPP. Because of the increased competitiveness of Japanese brands, many competitors are predicted to exit. Land-Rover exits Vietnam in 46% of the replications, Chevrolet also faces a large exit probability at 32%. Mercedes-Benz, BMW and Hyundai have a probability of exiting the country between 14 and 19%. By contrast, Subaru and Lexus (currently not sold in Vietnam) enter in 17% and 11% of the replications. The bright side of CPTPP for Vietnam is that the reduction of the 55% tariffs on assembled cars leads to an over 26% increase in the surplus of Vietnamese car buyers.

⁴⁹<http://international.gc.ca/trade-commerce/trade-agreements-accords-commerciaux/agr-acc/cptpp-ptpgp/impact-repercussions.aspx?lang=eng>

⁵⁰<http://www.international.gc.ca/economist-economiste/analysis-analyse/studies-etudes/studies-etudes-01.aspx?lang=eng>

8 Conclusion

We deploy extremely detailed data from the car industry to estimate the structural parameters of an extended version of the double-CES MP model. We use this framework to predict the medium-run consequences of numerous policy proposals circulating in the post-2016 context (after the victories of Leave in the Brexit referendum, and of Donald Trump in the US elections). Our counterfactuals simulate adjustments in firm-level decisions of (1) which markets to enter, (2) the fraction of varieties to offer, (3) the quantity of cars to supply in each of those markets, and (4) which location to source from.

Several insights emerge from our eight counterfactual exercises that have broad applicability. First, the stakes for consumers and producers in the outcomes of trade policy are often sizeable. Canada could lose two thirds of its auto sector in one scenario (Trumpit). The Vietnamese price index of cars falls by a quarter in another (TPP). A common factor underlying the large effects predicted here are the magnitudes of the two key elasticities: 7.7 for substitution between assembly sites (θ) and 3.87 for substitution between varieties (η). Moreover, the tariff changes proposed in the car industry are far from negligible. Second, the structure of multinational production matters a great deal. The much larger output reductions in Canada when losing duty-free access to the US market as compared to the analogous losses for the UK upon a Hard Brexit illustrates how the origins and networks of plants for each brand shape counterfactual outcomes. One reason why Canada suffers so much from the end of NAFTA is because 11 of the 12 brands that produce in Canada have plants in the US which they can switch production to. Third, some striking results, like the boom of Canadian car production under the CPTPP hinge importantly on the γ effects of RTAs. The reduced cost of operating Japanese plants in Canada impacts the American market (lowering output) even though the US retains the same trade policies. A second important source of interdependencies associated with trade policy changes is external returns to scale. Policies that substantially contract or expand output of important producers have spillover effects on other markets as we saw in the prediction that Chinese car output actually expands in response to Section 232 tariffs.

We view the counterfactual outcomes as the medium-run response because they hold the set of production locations constant. In this time frame, each brand's network of sourcing alternatives strongly shapes the responses to policy changes. The long-run decision of opening and closing production operations in different countries is of course very interesting. Our focus on the medium-run follows from the desire to keep the estimation tractable and the scope of this paper finite. The medium-run is also important in its own right because production networks are strongly persistent. Even over a 17-year period, covering a major disruptive crisis for this industry, 88% of OECD countries' car production still takes place within brand-country combinations that existed in 2000.

The large estimated benefits of producing, designing, and selling within a country, within an RTA, or with nearby countries all motivate future research to identify the mechanisms that underlie these frictions. This is particularly the case for the δ effects we have added to the framework.

While we specify them as marketing costs, preference-based mechanisms may play important roles. A third topic (along with plant location choices) calling for more research is the decision of where to source the components to be assembled in each car plant. Due to data limitations, we focused on the role of tariffs in raising costs for parts originating from headquarters, but actual sourcing problems are further complicated by rules of origin. These and other aspects of multinational expansion strategies provide a full agenda for future research.

References

- Anderson, J. E. and E. Van Wincoop (2004). Trade costs. *Journal of Economic literature* 42(3), 691–751.
- Anderson, S., A. De Palma, and J. Thisse (1992). *Discrete choice theory of product differentiation*. MIT Press.
- Andrews, D. W. and M. Buchinsky (2000). A three-step method for choosing the number of bootstrap repetitions. *Econometrica* 68(1), 23–51.
- 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.
- Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Arkolakis, C. and F. Eckert (2017). Combinatorial discrete choice. mimeo.
- Arkolakis, C., N. Ramondo, A. Rodríguez-Clare, and S. Yeaple (2018). Innovation and production in the global economy. *American Economic Review*.
- Atkeson, A. and A. Burstein (2008). Pricing-to-market, trade costs, and international relative prices. *The American Economic Review* 98(5), 1998–2031.
- Bartelme, D., A. Costinot, and A. Rodríguez-Clare (2018). External economies of scale and industrial policy: A view from trade.
- Bernard, A. B., S. J. Redding, and P. K. Schott (2011). Multiproduct firms and trade liberalization. *The Quarterly Journal of Economics* 126(3), 1271–1318.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.
- Berry, S., J. Levinsohn, and A. Pakes (1999). Voluntary export restraints on automobiles: evaluating a trade policy. *American Economic Review* 89(3), 400–430.
- Björnerstedt, J. and F. Verboven (2016). Does merger simulation work? Evidence from the Swedish analgesics market. *American Economic Journal: Applied Economics* 8(3), 125–64.

- Bresnahan, T. F. and V. A. Ramey (1993). Segment shifts and capacity utilization in the US automobile industry. *The American Economic Review* 83(2), 213–218.
- Bresnahan, T. F. and V. A. Ramey (1994). Output fluctuations at the plant level. *The Quarterly Journal of Economics* 109(3), 593–624.
- Ciliberto, F., C. Murry, and E. T. Tamer (2018). Market structure and competition in airline markets. Discussion Paper 13346, CEPR.
- Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2018). What drives home market advantage? *Journal of International Economics* 110, 135–150.
- Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In E. Helpman, G. Gopinath, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, Chapter 4, pp. 197–261. Elsevier.
- de Gortari, A. (2017). Disentangling global value chains. Technical report, Mimeo, Harvard University.
- Dekle, R., J. Eaton, and S. Kortum (2007). Unbalanced trade. *American Economic Review* 97(2), 351–355.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Eaton, J., S. Kortum, and S. Sotelo (2013). International trade: Linking micro and macro. In D. Acemoglu, M. Arellano, and E. Dekel (Eds.), *Advances in Economics and Econometrics Tenth World Congress*, Volume II: Applied Economics. Cambridge University Press.
- Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011). Income distribution, product quality, and international trade. *Journal of Political Economy* 119(4), 721–765.
- Feenstra, R. C. and J. A. Levinsohn (1995). Estimating markups and market conduct with multidimensional product attributes. *Review of Economic Studies* 62(1), 19–52.
- Fuss, M. and L. Waverman (1990). The extent and sources of cost and efficiency differences between US and Japanese motor vehicle producers. *Journal of the Japanese and International Economies* 4(3), 219–256.
- Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the US automobile industry. *Econometrica*, 891–951.
- 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.
- Grigolon, L. and F. Verboven (2014). Nested logit or random coefficients logit? A comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics* 96(5), 916–935.

- Guimaraes, P., O. Figueirido, and D. Woodward (2003). A tractable approach to the firm location decision problem. *The Review of Economics and Statistics* 85(1), 201–204.
- Hanemann, W. M. (1984). Discrete/continuous models of consumer demand. *Econometrica*, 541–561.
- Head, K. and T. Mayer (2013). What separates us? Sources of resistance to globalization. *Canadian Journal of Economics* 46(4), 1196–1231.
- Head, K. and T. Mayer (2018). Poor substitutes? Counterfactual methods in IO and trade compared. mimeo.
- Irrazabal, A., A. Moxnes, and L. D. Opmolla (2013). The margins of multinational production and the role of intrafirm trade. *Journal of Political Economy* 121(1), 74–126.
- Klier, T. and D. P. McMillen (2008). Evolving agglomeration in the U.S. auto supplier industry. *Journal of Regional Science* 48(1), 245–267.
- König, M. D., D. Rohner, M. Thoenig, and F. Zilibotti (2017). Networks in conflict: Theory and evidence from the great war of africa. *Econometrica* 85(4), 1093–1132.
- Kucheryavy, K., G. Lyn, and A. Rodríguez-Clare (2016). Grounded by gravity: A well-behaved trade model with industry-level economies of scale. Technical Report 22484, National Bureau of Economic Research.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725.
- Muffatto, M. (1999). Introducing a platform strategy in product development. *International Journal of Production Economics* 60, 145–153.
- Ramondo, N. (2014). A quantitative approach to multinational production. *Journal of International Economics* 93(1), 108–122.
- Ramondo, N. and A. Rodríguez-Clare (2013). Trade, multinational production, and the gains from openness. *Journal of Political Economy* 121(2), 273–322.
- Santos Silva, J. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics* 88(4), 641–658.
- Schmitt, A. and J. V. Biesebroeck (2013). Proximity strategies in outsourcing relations: The role of geographical, cultural and relational proximity in the European automotive industry. *Journal of International Business Studies* 44(5), 475–503.
- Smith, D. and R. Florida (1994). Agglomeration and industrial location: An econometric analysis of Japanese-affiliated manufacturing establishments in automotive-related industries. *Journal of Urban Economics* 36(1), 23–41.

- Tintelnot, F. (2017). Global production with export platforms. *The Quarterly Journal of Economics* 132(1), 157–209.
- Van Biesebroeck, J., H. Gao, and F. Verboven (2012). Impact of FTAs on Canadian auto industry. *DFAIT Canada* 31.
- Verboven, F. (1996). International price discrimination in the European car market. *The RAND Journal of Economics*, 240–268.
- Wang, Z. (2017). Headquarters gravity. Technical report, Mimeo, Shanghai University.

For Online Publication

A Constant elasticity of substitution discrete choice

Following Hanemann (1984)'s equation (3.5), let utility of household h be given by

$$U_h = u \left(\sum_m \psi_{mh} c_{mh}, z_h \right),$$

with z the outside good. The model-household parameters ψ_{mh} convert car use into equivalent units of psychological car services.⁵¹

Unlike the more familiar RUM with unitary demand, we model the c_{mh} as continuous choice variables. There are two interpretations for cars. One involves households with multiple members who share some number of cars. For example with two adults and one teenager in the household $c_h = 1$ if each member has their own car, but would be $c_h = 1/3$ if the three household members shared a single car. Obviously, unless households are very large (car-sharing groups might be an illustration), the continuity assumption is violated by integer issues.

A second interpretation involves endogenous use of a durable good. Suppose that each new car delivers 1 unit of lifetime services. Then $\sum_t c_{ht} = 1$. By driving sparingly or maintaining intensively in a given year, c_{ht} can be reduced, prolonging the duration of use. In this case $c_{ht} = 0.2$ would correspond to using 1/5 of the car's operating life each year. Assuming a steady state and aggregating over all households, the annual demand for new cars of model m in market n is given by $q_{mn} = \sum_h c_{mh}$. Summing across all models, the household's annual consumption is $c_h \equiv \sum_m c_{mh}$. Summing across all households and models, we have $\sum_h \sum_m c_{mh} = Q_n$, where Q_n denotes aggregate number of new cars sold in country n . We have implicitly assumed that in our steady state car replacements are spread evenly over periods, to avoid all consumers buying new cars in the fifth year and no sales at all in between.

Consumers choose c_{mh} for each model of the set of models available in market n and spend the remainder of their income, y_h , on outside good z with price normalized to one. Thus they maximize U_h subject to $\sum_m p_m c_{mh} + z_h = y_h$. Denoting the Lagrange multiplier as λ , and the partial derivatives with respect to $\sum_m \psi_{mh} c_{mh}$ and z_h as u_1 and u_2 , the first order conditions are

$$u_1 \psi_{mh} = \lambda p_m \quad \forall m \text{ with } c_{mh} > 0; \quad \text{and} \quad u_2 = \lambda.$$

Combining we have

$$\frac{u_1}{u_2} = \frac{p_m}{\psi_{mh}} \quad \forall m \text{ with } c_{mh} > 0$$

This equation implies a relationship between $\sum_m \psi_{mh} c_{mh}$ and p_m/ψ_{mh} that can only hold for $c_{mh} > 0$ and $c_{mh'} > 0$ under the measure 0 event that $\frac{p_m}{\psi_{mh}} = \frac{p_{m'}}{\psi_{m'h}}$ for $m \neq m'$. Otherwise each

⁵¹For example, ψ_{mh} could be the number of driving kilometers expected by the buyer over the lifetime of the model.

household h will select its preferred model m_h^* and consume c_h units while consuming $c_{m'h} = 0$ on all $m' \neq m_h^*$. In other words, the indifference curves between any pair of varieties m and m' , holding z constant, are linear, implying a corner solution. Thus c_h is given by

$$\frac{u_1(\psi_{mh}c_h, y - p_m c_h)}{u_2(\psi_{mh}c_h, y - p_m c_h)} = \frac{p_m}{\psi_{mh}} \quad \text{for } m = m_h^*$$

The preferred choice, m^* , is given by the argmin of p_m/ψ_{mh} (Hanemann, 1984, p. 548). Since a monotonic transformation of p_m/ψ_{mh} preserves the ranking, this is equivalent to maximizing $\ln \psi_{mh} - \ln p_m$. Parameterizing $\psi_{mh} = \beta_m \exp(\epsilon_{mh})$, the probability a given household chooses model m is

$$\text{Prob}(p_m/\psi_{mh} < p_j/\psi_{hj}) = \text{Prob}(\epsilon_{mh} + \ln \beta_m > \epsilon_{jh} + \ln \beta_j + \ln p_m - \ln p_j), \forall j \neq m.$$

With ϵ distributed according to the CDF $\exp(-\exp(-\eta\epsilon))$ (Gumbel with scale parameter $1/\eta$), the resulting choice probabilities at the level of market n are

$$\mathbb{P}_{mn} = \frac{\beta_m^\eta (p_{mn})^{-\eta}}{\Phi_n}, \quad \text{where } \Phi_n \equiv \sum_{j \in \mathcal{M}_n} \beta_j^\eta (p_{jn})^{-\eta}.$$

The above equation can be re-expressed in the standard conditional logit form by taking logs and then taking the exponential of each term in the numerator and denominator.

Aggregate expected sales of model m in n are

$$\mathbb{E}[q_{mn}] = \sum_h \mathbb{P}_{mn} c_h = \mathbb{P}_{mn} \sum_h c_h = \mathbb{P}_{mn} Q_n.$$

The elasticity of demand with respect to the price of model m is $-\eta(1 - \mathbb{P}_{mn})$, which goes to $-\eta$ as $\mathbb{P}_{mn} \rightarrow 0$. Intuitively, demand becomes more responsive to price as η increases because η is *inversely* related to the amount of heterogeneity in consumer preferences.

Expected sales of any model are proportional to the aggregate size of the market expressed in volumes, regardless of $u(\cdot)$. Furthermore, *income does not affect the choice between models* but, depending on the form of $u(\cdot)$, the consumption of cars can have any income expansion path. For example, under the Cobb-Douglas case, explored by Anderson et al. (1992), the optimal consumption of the chosen car is $c_{mh} = (\alpha y_h)/p_m$, for $m = m_h^*$. Non-homothetic demand will be obtained from all other assumed $u(\cdot)$. The quasi-linear case where $U_h = (\sum_m \psi_{mh} c_{mh})^\alpha + z_h$, yields $c_{mh} = \left(\frac{p_m}{\alpha \psi_{mh}^\alpha}\right)^{1/(\alpha-1)}$. The share of expenditure spent on cars will therefore fall with income. An opposite conclusion can be obtained with $U_h = \sum_m \psi_{mh} c_{mh} + z_h^\alpha$, which gives the demand for the chosen car model $c_{mh} = \frac{y_h - \left(\frac{\psi_{mh}}{\alpha p_m}\right)^{1/(\alpha-1)}}{p_m}$. In this case, car expenditure as a share of income is increasing in income.

B Constructing expected profits from estimates

The brand entry estimation requires to empirically measure $\mathbb{E}[\pi_{bnt}]$. Equation (18) shows that we need a number of intermediate estimates for that. In particular, we need to measure $\mathbb{E}[\pi_{mnt}]$ and the model entry fixed costs parameters $\mu_{nt}^e + \beta_b^e + \ln(w_{it}^\zeta w_{nt}^{1-\zeta})$.

The country (nt) fixed effects in equation (22) have structural interpretations given by $\text{FE}_{nt}^{(2)} = \ln \kappa_1 + \eta \ln P_{nt}$. In the model entry equation (23), which involves a constant, the country fixed effects are interpreted as

$$\hat{\sigma}_e \text{FE}_{nt}^{(3)} = \ln Q_{nt} + \eta \ln P_{nt} - \ln(w_{nt}^{1-\zeta}) - \mu_{nt}^e - (\ln Q_{1T} + \eta \ln P_{1T} - \ln(w_{1T}^{1-\zeta}) - \mu_{1T}^e). \quad (\text{B.1})$$

with the model-entry constant given by

$$\hat{\sigma}_e \text{CST}^{(3)} = \ln \kappa_2 - \ln \eta + (\eta - 1)(\ln \varphi_1 - (1 - \alpha) \ln w_{i(1)T}) - \beta_1^e + (\ln Q_{1T} + \eta \ln P_{1T} - \ln(w_{1T}^{1-\zeta}) - \mu_{1T}^e). \quad (\text{B.2})$$

Only relative levels of productivity (φ_b) and headquarter wages ($w_{i(b)}$) matter for market shares. Therefore, we can normalize $\varphi_1 = w_{i(1)T} = 1$, implying

$$(\ln Q_{1T} + \eta \ln P_{1T} - \ln(w_{1T}^{1-\zeta}) - \mu_{1T}^e) = \hat{\sigma}_e \text{CST}^{(3)} - (\ln \kappa_2 - \ln \eta) + \beta_1^e. \quad (\text{B.3})$$

Substituting (B.3) into (B.1), replacing $\eta \ln P_{nt}$ with $\text{FE}_{nt}^{(2)} - \ln \kappa_1$, and isolating the unknown parameters, we obtain

$$\ln(w_{nt}^{1-\zeta}) + \mu_{nt}^e + \beta_1^e = \ln Q_{nt} + (\text{FE}_{nt}^{(2)} - \ln \kappa_1) - \hat{\sigma}_e (\text{FE}_{nt}^{(3)} + \text{CST}^{(3)}) + (\ln \kappa_2 - \ln \eta). \quad (\text{B.4})$$

We then use the fixed effect of brand b in the entry and market share equations to add the missing β_b^e :

$$\begin{aligned} \hat{\sigma}_e \text{FE}_b^{(3)} &= -(\beta_b^e - \beta_1^e) + (\eta - 1) \ln \varphi_b - (\eta - 1)(1 - \alpha) \ln w_{i(b)T}, \quad \text{and} \\ \text{FE}_b^{(2)} &= \eta [\ln \varphi_b - (1 - \alpha) \ln w_{i(b)T}]. \end{aligned} \quad (\text{B.5})$$

Multiplying the second line by $\frac{\eta-1}{\eta}$, we isolate β_1^e as:

$$\beta_1^e = \beta_b^e - \left(\frac{\eta-1}{\eta} \text{FE}_b^{(2)} - \hat{\sigma}_e \text{FE}_b^{(3)} \right), \quad (\text{B.6})$$

and rewrite equation (B.4) as

$$\begin{aligned} \ln(w_{nt}^{1-\zeta}) + \mu_{nt}^e + \beta_b^e &= \ln Q_{nt} + (\text{FE}_{nt}^{(2)} - \ln \kappa_1) - \hat{\sigma}_e (\text{FE}_{nt}^{(3)} + \text{CST}^{(3)}) + (\ln \kappa_2 - \ln \eta) \\ &\quad + \frac{\eta-1}{\eta} \text{FE}_b^{(2)} - \hat{\sigma}_e \text{FE}_b^{(3)} \end{aligned} \quad (\text{B.7})$$

The model entry fixed cost central parameter is therefore obtained by adding $\ln(w_{i(b)t}^\zeta)$:

$$\begin{aligned} \ln(w_{i(b)t}^\zeta w_{nt}^{1-\zeta}) + \mu_{nt}^e + \beta_b^e = & \ln Q_{nt} + (\text{FE}_{nt}^{(2)} - \ln \kappa_1) - \hat{\sigma}_e(\text{FE}_{nt}^{(3)} + \text{CST}^{(3)}) + (\ln \kappa_2 - \ln \eta) \\ & + \frac{\eta - 1}{\eta} \text{FE}_b^{(2)} - \hat{\sigma}_e \text{FE}_b^{(3)} + \left(\frac{\eta - 1}{\eta} \mathbf{v}_2 - \hat{\sigma}_e \mathbf{v}_3 \right) \mathbf{W}'_{i(b)t}, \end{aligned} \quad (\text{B.8})$$

$\mathbf{W}_{i(b)t}$ including the two proxies for $\ln w_{i(b)t}$.

The last step needed for reconstructing $\mathbb{E}[\pi_{bnt}]$ is a measurement of $\mathbb{E}[\pi_{mnt}]$. We obtain it from estimates of the market share and entry equations:

$$\begin{aligned} \mathbb{E}[\pi_{mnt}] = & \frac{\kappa_2}{\eta} Q_n \\ & \exp \left(\frac{\eta - 1}{\eta} (\text{FE}_b^{(2)} - \mathbf{W}'_{i(b)t} \mathbf{v}_2) - (\eta - 1) \mathbf{X}'_{int} \mathbf{d} - (\eta - 1) \ln C_{bnt} + (\text{FE}_{nt}^{(2)} - \ln \kappa_1) \right) \end{aligned} \quad (\text{B.9})$$

In DEV counterfactuals, we also need to reconstruct the price index and brand entry fixed costs parameters. The price index is reconstructed as

$$P_n = \left(\sum_b \kappa_1 M_{bn} \exp \left(\text{FE}_b^{(2)} - \mathbf{W}'_{i(b)t} \mathbf{v}_2 - \eta \mathbf{X}'_{int} \mathbf{d} - \eta \ln C_{bnt} \right) \right)^{-1/\eta}, \quad (\text{B.10})$$

The central parameter of the distribution costs is retrieved as

$$\mu_n^d + \beta_b^d + \ln(w_{i(b)t}^\zeta w_{nt}^{1-\zeta}) = -\hat{\sigma}_d (\text{FE}_{nt}^{(4)} + \text{FE}_b^{(4)} - \mathbf{W}'_{i(b)t} \mathbf{v}_4). \quad (\text{B.11})$$

C Exact Hat Algebra (EHA) with the double-CES MP model

EHA solves for an equilibrium in the proportional changes of all variables of interest following a change in frictions. Variables taking a hat symbol are defined as ratios of $\hat{x} = x'/x$, where x is the initial level, and x' is the level attained after the change. The main advantages are that (1) it computes predicted (exact) percentage changes from the actual data, (2) related, it allows for unobservables in the actual decisions (as long as they are unaffected by the counterfactual), (3) it minimizes the data and parameter requirements.

The EHA method used here includes two non-standard features: First, we allow for proportional changes in the fraction of models offered in a market. Second, we incorporate external returns to scale in a multinational production setting. The change in total output located in country k , denoted \hat{q}_k , affects outcomes through \hat{C}_{bn} . Therefore, the problem can be decomposed in an inner and outer loop, analogous to the structure used in the pure trade model of Kucheryavyi et al. (2016). The first subsection of this appendix presents the details of the inner loop. The second subsection presents the overall solution algorithm.

C.1 EHA with a continuous model-entry margin: the inner loop

Starting with the sourcing decision, equation (6), algebraic manipulations of CES shares yield

$$\hat{C}_{bn} = \left(\sum_k \mathbb{L}_{bk} \text{Prob}(\mathbb{S}_{bkn} = 1) (\hat{\gamma}_{ik} \hat{\tau}_{kn} \hat{q}_k^\zeta)^{-\theta} \right)^{-1/\theta}. \quad (\text{C.12})$$

Since all models are the same in expectation, the $\text{Prob}(\mathbb{S}_{bkn} = 1) = s_{bkn} \equiv q_{bkn}/q_{bn}$, i.e. the share of cars brand b sells in n that it sources from country k . This quantity share has the same expected value as the sourcing count share, S_{bkn}/M_{bn} , but it allows for internal consistency in the counterfactuals. The updating function for C_{bn} depends on this sourcing share and on the changes in frictions, both of which we observe.

The price index updating function is

$$\hat{P}_n = \left(\sum_b \mathbb{D}_{bn} \frac{q_{bn}}{Q_n} \hat{M}_{bn} (\hat{\delta}_{in} \hat{C}_{bn})^{-\eta} \right)^{-1/\eta}. \quad (\text{C.13})$$

We observe the initial market share of b in n , q_{bn}/Q_n , and \hat{C}_{bn} is obtained from (C.12). To determine \hat{M}_{bn} , we need to investigate how the number of models offered in each market changes in the counterfactual.

One of the novel aspects of our EHA approach to counterfactuals is to account for changes in model entry, \hat{M}_{bn} . The condition for model entry (12), combined with the definition of fixed model entry costs, F_{mn}^e , determines the probability of model entry. In the EHA method we replace the initial probability of model entry with the initial share of models offered by brand b in market n .

$$\frac{M_{bn}}{M_b} = \text{Prob}(\mathbb{I}_{mn} = 1) = \Phi \left[\frac{\ln \mathbb{E}[\pi_{mn}] - (\mu_n^e + \beta_b^e + \zeta \ln w_i + (1 - \zeta) \ln w_n + \ln \delta_{in}^e)}{\sigma_e} \right]. \quad (\text{C.14})$$

In the counterfactual, we have the following entry shares:

$$\frac{M'_{bn}}{M_b} = \Phi \left[\frac{\ln \widehat{\mathbb{E}[\pi_{mn}]} + \ln \mathbb{E}[\pi_{mn}] - (\mu_n^e + \beta_b^e + \zeta \ln w_i + (1 - \zeta) \ln w_n + \ln \hat{\delta}_{in}^e + \ln \delta_{in}^e)}{\sigma_e} \right]. \quad (\text{C.15})$$

Collecting terms, one can rewrite

$$\frac{M'_{bn}}{M_b} = \Phi \left[\frac{\ln \widehat{\mathbb{E}[\pi_{mn}]} - \ln \hat{\delta}_{in}^e}{\sigma_e} + \frac{\ln \mathbb{E}[\pi_{mn}] - (\mu_n^e + \beta_b^e + \zeta \ln w_i + (1 - \zeta) \ln w_n + \ln \delta_{in}^e)}{\sigma_e} \right]. \quad (\text{C.16})$$

Substituting the second term in brackets with equation (C.14)

$$\frac{M'_{bn}}{M_b} = \Phi \left[\frac{\ln \widehat{\mathbb{E}[\pi_{mn}]} - \ln \hat{\delta}_{in}^e}{\sigma_e} + \Phi^{-1} \left(\frac{M_{bn}}{M_b} \right) \right]. \quad (\text{C.17})$$

Finally, the percent change in number of models offered is

$$\hat{M}_{bn} = \Phi \left[\frac{\ln \widehat{\mathbb{E}[\pi_{mn}]} - \ln \hat{\delta}_{in}^e}{\sigma_e} + \Phi^{-1} \left(\frac{M_{bn}}{M_b} \right) \right] \frac{M_b}{M_{bn}}. \quad (\text{C.18})$$

Equation (C.18) has many known components (M_b and M_{bn} are observed, σ_e is estimated, and $\hat{\delta}_{in}^e$ is part of the counterfactual experiment). The last needed part is to update the expected profits from entry of model m ($\widehat{\mathbb{E}[\pi_{mn}]}$). Using (14), we obtain the last element as $\widehat{\mathbb{E}[\pi_{mn}]} = \hat{\delta}_{in}^{1-\eta} \hat{C}_{bn}^{1-\eta} \hat{P}_n^\eta$, and therefore

$$\hat{M}_{bn} = \Phi \left[\frac{\ln(\hat{\delta}_{in}^{1-\eta} \hat{C}_{bn}^{1-\eta} \hat{P}_n^\eta) - \ln \hat{\delta}_{in}^e}{\sigma_e} + \Phi^{-1} \left(\frac{M_{bn}}{M_b} \right) \right] \frac{M_b}{M_{bn}}. \quad (\text{C.19})$$

C.2 The algorithm

The inner/outer loop procedure works as follows:

1. For a given level of \hat{q}_k , we have a system of three equations (details in Appendix C.1) determining the three endogenous objects that determine outcomes in the counterfactual: updates of the cost index, the price index, and the number of varieties offered:

$$\hat{C}_{bn} = \left(\sum_k \mathbb{L}_{bk} s_{bkn} (\hat{\gamma}_{ik} \hat{\tau}_{kn} \hat{q}_k^\zeta)^{-\theta} \right)^{-1/\theta}, \quad (\text{C.20})$$

$$\hat{P}_n = \left(\sum_b \mathbb{D}_{bn} \frac{q_{bn}}{Q_n} \hat{M}_{bn} (\hat{\delta}_{in} \hat{C}_{bn})^{-\eta} \right)^{-1/\eta}, \quad (\text{C.21})$$

$$\hat{M}_{bn} = \Phi \left[\frac{\ln(\hat{\delta}_{in}^{1-\eta} \hat{C}_{bn}^{1-\eta} \hat{P}_n^\eta) - \ln \hat{\delta}_{in}^e}{\sigma_e} + \Phi^{-1} \left(\frac{M_{bn}}{M_b} \right) \right] \frac{M_b}{M_{bn}}, \quad (\text{C.22})$$

where s_{bkn} is the share of sales in market n sourced from country k . A fixed point iteration with dampening solves for the equilibrium values of the system (C.20), (C.21) and (C.22). In our counterfactual, our main variable of interest is the percentage change in quantities

$$\hat{q}_{bln} = \hat{q}_{bn} \times \hat{s}_{bln}.$$

Again the EHA approach is very useful here, since it can be used to show that the changes in sourcing probability and brand market share are only functions of changes in frictions (known) and of the three endogenous variables \hat{C}_{bn} , \hat{P}_n , and \hat{M}_{bn} solved by the fixed point iteration:

$$\hat{s}_{bln} = \left(\frac{\hat{\gamma}_{il} \hat{\tau}_{ln} \hat{q}_l^\zeta}{\hat{C}_{bn}} \right)^{-\theta}, \text{ and } \hat{q}_{bn} = \hat{M}_{bn} \left(\frac{\hat{\delta}_{in} \hat{C}_{bn}}{\hat{P}_n} \right)^{-\eta} \quad (\text{C.23})$$

2. With the \hat{q}_{bln} generated in the inner loop, country-level output is updated using

$$q'_\ell = \sum_n \sum_b \hat{q}_{bln} q_{bln}, \quad \text{and therefore} \quad \hat{q}_\ell = q'_\ell / q_\ell.$$

Since \hat{C}_{bn} contains $\hat{q}_\ell^{-\theta_s}$, it needs to be updated. The inner loop is then run again, giving a new vector of \hat{q}_ℓ . This outer loop is run until we reach a fixed point in the vector of country-level output change.

In the segmented market version of our model, each segment can be considered in isolation when updating the market share equation. Start with the identity decomposing the sales of brand b from a plant in ℓ when serving n :

$$q_{bln} = \sum_s q_{blns} = s_{bln} \times q_{bn} = s_{bln} \times \sum_s q_{bns}.$$

In changes

$$\hat{q}_{bln} = \frac{q'_{bln}}{q_{bln}} = \hat{s}_{bln} \times \frac{\sum_s q'_{bns}}{\sum_s q_{bns}}. \quad (\text{C.24})$$

The expression describing sourcing share \hat{s}_{bln} in equation (C.23) is unchanged, since it is not affected by any segment-level determinant. However, the new level of production at the segment level is

$$q'_{bns} = q_{bns} \hat{q}_{bns} = q_{bns} \hat{M}_{bns} \left(\frac{\hat{\delta}_{ins} \hat{C}_{bn}}{\hat{P}_{ns}} \right)^{-\eta_s}, \quad (\text{C.25})$$

with changes in the price index and number of offered models given by

$$\begin{aligned} \hat{P}_{ns} &= \left(\sum_b \mathbb{D}_{bns} \frac{q_{bns}}{Q_{ns}} \hat{M}_{bns} (\hat{\delta}_{ins} \hat{C}_{bn})^{-\eta_s} \right)^{-1/\eta_s}, \\ \hat{M}_{bns} &= \Phi \left[\frac{\ln(\hat{\delta}_{ins}^{1-\eta_s} \hat{C}_{bn}^{1-\eta_s} \hat{P}_{ns}^{\eta_s}) - \ln \hat{\delta}_{ins}^e}{\sigma_e^s} + \Phi^{-1} \left(\frac{M_{bns}}{M_{bs}} \right) \right] \frac{M_{bs}}{M_{bns}}. \end{aligned} \quad (\text{C.26})$$

The algorithm is very similar to the unified markets case. The inner loop solves for \hat{C}_{bn} , \hat{P}_{ns} and \hat{M}_{bns} , using (C.20) and (C.26) which gives \hat{q}_{bln} from (C.25) and (C.24). The outer loop then sums over the new shipments to obtain total output in each country, which enters back \hat{C}_{bn} in the next iteration. The process is repeated until no further change is detected in any of those endogenous variables.

D External returns to scale and amplification

In this appendix, we discuss how our estimates of external IRS generate amplification in the output response to trade policy scenarios. Total output in country ℓ is simply the sum over brands and destinations of expected sales of b in n sourced from ℓ ($q_\ell = \sum_b \sum_n q_{bln}$). Multiplying (9) by

(6) and dropping expectations for simplicity, we obtain

$$q_{b\ell n} = \kappa_1(\gamma_{i\ell}\tau_{\ell n})^{-\theta}(w_\ell^\alpha q_\ell^\zeta)^{-\theta}(\varphi_b/w_i^{1-\alpha})^\eta Q_n P_n^\eta M_{bn} \delta_{in}^{-\eta} C_{bn}^{-\eta-1}. \quad (\text{D.27})$$

Let $K_{b\ell n}$ denote all the shifters of $q_{b\ell n}$ that do not directly depend on q_ℓ (ignoring for now the indirect effects via M_{bn} , C_{bn} and P_n):

$$K_{b\ell n} \equiv \kappa_1(\gamma_{i\ell}\tau_{\ell n})^{-\theta}(w_\ell^\alpha)^{-\theta}(\varphi_b/w_i^{1-\alpha})^\eta Q_n P_n^\eta M_{bn} \delta_{in}^{-\eta} C_{bn}^{-\eta-1}.$$

Therefore we can write

$$q_{b\ell n} = q_\ell^{-\varsigma\theta} K_{b\ell n}.$$

Aggregate output is therefore $q_\ell = \sum_b \sum_n q_{b\ell n} = q_\ell^{-\varsigma\theta} \sum_b \sum_n K_{b\ell n}$, and

$$q_\ell = \left(\sum_b \sum_n K_{b\ell n} \right)^{\frac{1}{1+\varsigma\theta}}.$$

Let us note at this stage that $K_{b\ell n}$ is also the value taken by $q_{b\ell n}$ under constant returns to scale, that is when $\varsigma = 0$. We can therefore define $q_\ell^{\text{CRS}} = \sum_b \sum_n K_{b\ell n}$ and

$$q_\ell^{\text{IRS}} = \left(q_\ell^{\text{CRS}} \right)^{\frac{1}{1+\varsigma\theta}}.$$

Consider a tariff or a wage shock that shifts national output by a factor $\hat{q}_\ell^{\text{CRS}}$ (the immediate effect). We can therefore express the full proportional change in output under IRS (the sum of immediate and amplification effects) as

$$\hat{q}_\ell^{\text{IRS}} = \left(\hat{q}_\ell^{\text{CRS}} \right)^{\frac{1}{1+\varsigma\theta}}. \quad (\text{D.28})$$

We refer to the amplification effect as the situation when the elasticity of the IRS output to CRS output ($1/(1+\varsigma\theta)$) is greater than one. With $\varsigma < 0$, we have $\varsigma\theta < 0$ and therefore amplification through external IRS. As $\varsigma\theta$ approaches -1 , there will be an infinitely large response to any small initial shock to output in ℓ . Hence, we need $-\varsigma\theta < 1$ which is analogous to the condition establishing uniqueness of equilibrium with external returns in the pure trade model of Kucheryavyy et al. (2016), since θ is the trade elasticity (the coefficient of log sales on log trade costs being $-\theta$ in equation (D.27)) and ς is the scale elasticity.

Taking logs of (D.28) yields a regression that can be used to quantify the amplification effect implied by our estimates of external returns to scale.

$$\ln \hat{q}_\ell^{\text{IRS}} = \frac{1}{1+\varsigma\theta} \ln \hat{q}_\ell^{\text{CRS}} + \epsilon_\ell.$$

The error term ϵ_ℓ is intended as a way to allow for the misspecification due to the fact that in reality our model has three endogenous variables that indirectly depend on q_ℓ . Our estimate of

$-\zeta\theta = 0.27$ implies that we should obtain a coefficient of $1/(1-0.27) = 1.37$. We run the regression on all of our 8 counterfactuals taking the sample to be all 47 assembly countries. Depending on the policy experiment, the coefficient on changes in logs falls between 1.32 and 1.37. The fit is extremely tight, with $R^2 = 0.99$ in every case. The slight under-estimates of amplification in our regressions are not surprising given that we took M_{bn} , C_{bn} and P_n as exogenous in our thought experiment involving K_{bn} . The change in q_ℓ must adjust P_n so as to hold the total number of cars consumed constant in each country. Since we hold world output the same, it seems intuitive that the simple calculations above will tend to *overstate* the amount of IRS output expansion.

E Data Appendix

E.1 Exclusions from the raw IHS data

- In order to restrict attention to vehicles with comparable substitution patterns, we eliminated light commercial vehicles as a car type, to work only with passenger cars. We also dropped pick-up trucks and vans because over 90% of their sales are registered as commercial vehicles.
- We delete shipments of unknown brand or assembly country. There were 62 countries in the IHS data where assembly location was unavailable for all sales and all years (mainly Caribbean and African countries). We also required that at least 90% of the total car sales in a country must come from identified brands, leading us to drop 6 more countries for recent years (Algeria, Bolivia and Peru before 2008, Chile before 2002, Kazakhstan and Belarus before 2005). The remaining 76 markets constituted 98.8% of world automotive sales in the 2016 IHS data. The market-years we lose are also dropped as production sites based on the fact that in most case, their output is essentially meant for domestic consumption.
- Norway is only an option for Think and in those cases it is the only option; therefore a Norway fixed effect cannot be estimated.
- We drop De Tomaso because it is only sold in one market (Kuwait) for two years and the estimations of equation (22) and (23) cannot identify its brand fixed effect. The same is true for Troller, which only sells in Brazil.
- AIL and Pyeonghwa Motors are dropped because the IHS data does not show their production in the headquarters countries (respectively, Israel and North Korea) even though other information reports they do assemble car in those locations during the time frame of our data.
- FSO and TVR are only present in 2000 in our dataset, Moskvich sells in Ukraine until 2001, they are dropped since we consider years starting in 2002 for estimation.
- The Vauxhall brand name is only used in the UK for cars that are elsewhere sold as Opel. Because we want to consider potential relocation from UK to Germany and vice-versa in particular for the Brexit counterfactual, Vauxhall is renamed Opel.
- We eliminated the observations where a brand's total production in a given origin was less than 10 cars a year. Those mostly involved extinct models being sold out of left-over inventories (Mazda selling to Switzerland one unit of the 121 model from a closed factory in the UK several years after production was stopped for instance).
- We drop 43 brands that never had more than one model. They cannot be included in the estimation of the model-entry equation because their brand dummy is a perfect predictor.

Such firms are typically very small, having (collectively) a median share of a market-year of just 0.002%, with the maximum market share of 1.4% in China in 2004.

- We drop two countries from the counterfactuals, that have so few brands that the computation of the equilibrium sometimes failed because of zero brand entry: Pakistan (4 brands in 2016) and Venezuela (6). These markets are retained in the estimation, however.

E.2 Other data sources: gravity variables, tariffs and RTAs

The time-invariant determinants of frictions (distance, home, contiguity, common language) and GDP per cap variables come from the CEPII gravity database.⁵² Tariff information for both assembled cars and parts comes from the WITS database managed by the World Bank. WITS compiles individual country declarations of their applied MFN and preferential *ad valorem* tariffs, as well *ad valorem* equivalents (AVE) of any specific tariffs. The car tariff is the simple average of the tariffs in HS heading 8703. The car parts tariff is the simple average of the three 4-digit HS headings associated with major components (8706, 8707, and 8708), together with the relevant HS6 categories for engines and associated parts (840733, 840734, 840820, 840991, and 840999). There are many holes in the data which we fill via linear interpolation. When the data is missing for the most recent years, we use the last available year. When a preferential rate exists, we use it. For the rest of dyad-years, we use the MFN tariff inclusive of the AVE of specific tariffs. We also corrected a number of issues in available WITS data regarding recent RTAs that are important for our purposes. Korea signed a number of recent agreements (with the EU and USA in 2012, Canada in 2015, Peru in 2012, Turkey in 2013, Australia in 2015, New-Zealand, Vietnam and China in 2016), for which WITS tariff data is either not accurate or not available. Japan signed RTAs with Peru (entered into force in 2012) and Australia (entered into force in 2015), for which the preferential rates were not mentioned in WITS. Colombia also lacked preferential rates for agreements with the USA and Canada (entry into force in 2012) and the EU (entry into force in 2013). For all those cases (and a few other for which it turned out that cars and parts were mostly exempted), we went to individual text and tariff schedules of those agreements to compute the tariff rate relevant for the bilateral pair in the relevant years. This involves in particular to take into account the “Staging” variable specified (usually giving the number of equal cuts in years) applied to the “Base rate” (the MFN rate at the date of entry into force). Sometimes, even that level of detail is not enough. For instance, the Korea-USA agreement finally decided to postpone the negotiated phasing in by five years (<https://www.uskoreacouncil.org/wp-content/uploads/2014/12/Automotive-Provisions.pdf>). We took those into account when mentioned on the countries’ relevant website. The correction was done both for assembled cars and parts. Note that the staging of tariff liberalization can be very different across bilateral pairs.⁵³

⁵²Updated for the purpose of the paper, and available at http://www.cepii.fr/cepii/en/bdd_modele/presentation.asp?id=8 or <https://sites.google.com/site/hiegravity/data-sources>.

⁵³For example, Colombia’s 10% MFN car engine tariff went immediately to zero for its FTA with the US but took 8 years to expire in the FTA with the EU. Korea reduced its tariffs on assembled cars from the EU and the USA to 0 in a

Those changes are particularly important for obtaining realistic numbers for our counterfactuals, since Korea and Japan are estimated to be the two top places to produce cars in terms of productive efficiency (lowest estimated costs of production).

The RTA database maintained by the WTO provides the dates, membership and topics covered for each trade agreement.

F Robustness

F.1 Robustness on external returns to scale

Table F.1 reports results of four robustness checks for the magnitude of external returns to scale. Column (1) simply lags the total output variable by one year. Columns (3) and (4) proceed to purge total output from potential simultaneity bias, using two different versions of a Bartik instrument (using a control function approach since our estimation procedure is PPML rather than OLS).

Table F.1: Sourcing decision with alternative output variables

	(1)	(2)	(3)	(4)
$\ln(1 + \text{car tariff}_{\ell n})$	-7.688 (0.890)	-7.817 (0.916)	-7.820 (0.914)	-7.815 (0.914)
log lagged country output	0.243 (0.057)			
log current country output		0.278 (0.082)	0.217 (0.107)	0.189 (0.110)
residual term of Bartik first-stage			0.121 (0.127)	0.137 (0.104)
Observations	346322	338722	338722	338722
ς	0.032	0.036	0.028	0.024

Standard errors clustered by production country. Fixed effects for brand-destination and production country are included in all columns. Column (1) replaces current output with lagged version. Column (2) replicates our benchmark results for the same sample as the control function in columns (3) and (4), which are two versions of the control function Bartik approach, described in this appendix.

The procedure for constructing the instrument is the following: Start by computing the share of each possible origin country in the sales of each brand-destination combination in a base year (2002 in our case). Then compute the predicted level of sales multiplying this initial share by the level of demand faced by the brand in that destination each subsequent year. This keeps the sourcing shares (our endogenous variable in the final regression) unchanged. Then sum those predicted sales for each origin-year, yielding predicted output of the country-year. This prediction is used in a regression explaining the true level of output ($q_{\ell t}$). Finally use the residual of that

few years, but cars were exempted from any liberalization in the RTA with China (http://fta.mofcom.gov.cn/korea/annex/fujian2_A_hfgsjr_en.pdf).

regression in the sourcing decision as a control function, the idea being that the coefficient of q_{lt} now reflects the change in output that *comes purely from shifts in demand in destinations where the origin was selling in the base year*. The approach is the same in column (4) but with a more demanding specification, where the shares to be shifted are the market shares of each brand-origin. This means that the only reason why total output is allowed to change is the *overall* demand of each destination country, leaving the market shares and the sourcing decisions of all brands at their base level in 2002.

Our baseline ς is 3.5%. Lagging origin output reduces this value to 3.2%. Using the Bartik instrumentation reduces it more substantially to 2.8 and 2.4% when using the less and more demanding specifications respectively. The Bartik instrumentation loses a moderate number of observations for origin-years that were not assembling cars in 2002 (Bulgaria, Algeria and Morocco are 3 examples of countries that start producing cars later in our sample). Column (2) replicates our baseline estimation on this reduced sample. The coefficient on car tariffs (which is used with the one on q_{lt} to reveal ς) is very stable. The coefficients on q_{lt} stay positive and statistically significant, but are reduced in an expected way: as mentioned in the main text, our estimated parameter for external returns to scale should be seen as an upper bound, which will tend to give the largest scope for interdependencies across markets in our counterfactuals.

F.2 Estimates using the firm-variety approach

Variety v corresponds to an “underneath the hood” concept of product differentiation—in contrast to models which were “re-badged” versions of cars that were physically very similar. We define distinct varieties using three variables in the IHS database:

Platform “All-new ground up redesign would constitute a new Platform designation.” Muffatto (1999) points out that companies vary in terms of how many aspects of the design go into the platform designation. At a minimum, platforms include a common underbody and suspension. Broader definitions include engines, transmissions, and exhaust systems.

Program “Code is used by OEMs to identify Vehicle throughout design lifecycle.” We think of programs as constituting more minor redesigns, or new generations within a given platform.

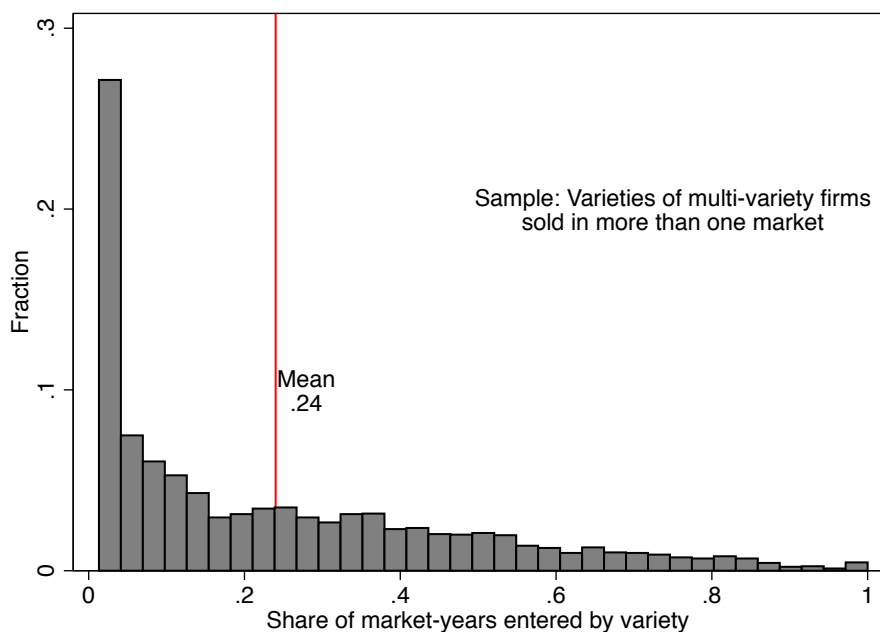
Body type Distinguishes between sedans, hatchbacks, wagons, etc.

Firm f here corresponds to the IHS variable “Design Parent: The company/OEM responsible for the design of the vehicle platform.” Except for a small number of cases that we manually corrected, platforms map many to one to Design Parents. We think of this as the engineering/design approach. While it does not provide a clear ownership criteria, IHS allows for firms to be designated as “parents” even if ownership is less than 50%. For example Kia has Hyundai as a parent even though Hyundai owned about 34% of Kia stock in December 2015.

The biggest problem with the Design Parents (DP) is that IHS only reports it as of 2016. Thus, going back in time, it gives incorrect DPs. For example, it makes no sense to think of Tata as the DP

for Jaguar cars before 2008 when the brand was owned by Ford. We are able to track ownership changes for brands over time, as the latter often correspond to distinct, stock-selling corporations (e.g. Audi, Nissan). However, it is more difficult to track ownership of platforms. Brands map many-to-many to Design Parents (they map many-to-one to Sales Parents). The reason is that brands market (and even manufacture) platforms designed by other firms. The IHS Engineering Group identifier is very helpful in a few cases (Chrysler-Fiat, Mazda-Ford). For others the brand-platform mapping seems clean enough.

Figure F.1: Market coverage by multi-variety firms



There were two main concerns about the brand-model approach employed in the main text of the paper. The first concern is that much of the low entry rates observed at the brand-model level could be an artifact of re-badging strategies. Thus while Honda seems to sell the Legend in Japan only, the same variety is in fact available in many markets as the Acura RL.⁵⁴ Figure F.1 reproduces Figure 2 using the firm-variety concept. The mean entry rate rises, as expected, but only from 0.23 to 0.24. The whole distributions of entry rates are visually very similar.

The second concern with the brand-model approach is that parent-firm headquarters might be making the critical management and parts supply decisions and that the brand headquarters might be less relevant from the point of view of γ_{il} frictions. For example, the top management of Renault-Nissan in Paris might provide all the brands of the group (Renault, Nissan, Dacia and Lada) with designs and production technologies. Using France, Japan, Romania, and Russia,

⁵⁴ Another form of re-badging holds the model name constant while changing the brand name. For example the platform B0, program H79 is sold in roughly equal amounts as a “Duster” under the brand Renault and as a Dacia (a Romanian brand acquired by Renault).

respectively, as the brand headquarters might therefore incorrectly specify the relevant frictions.

Table F.2 re-estimates the baseline specification from Table 3. The sample size in the sourcing equation (column 1) falls by 20%. There is a greater number of possible assembly locations when taking account of all the parent firms' production facilities but there are far fewer sourcing shares when grouping by firm instead of brand. The estimates for the τ_{ln} determinants are very similar to those reported in Table 3: The key elasticity of sourcing response, θ , rounds to 7.7 in both regressions. The imprecision of the estimates of the γ_{il} determinants in the sourcing equations persists with the new set of headquarter i locations. Distance continues to have the wrong sign. Parts tariffs now enter with a highly significant coefficient but it implies a cost share of HQ-provided intermediates that exceeds one. In sum, using firm headquarters does not markedly improve the γ estimates. In column (2), the firm average market share equation estimates an η of 1.9, considerably lower than the 3.87 obtained for brands. This η implies markups above 100% that are drastically higher than other estimates in the literature. The implausibly low η estimates suggest that the firm-level market share equation suffers from measurement error in the calculation of C_{bn} . Firms aggregate an often highly heterogeneous set of plants producing very distinct sets of cars. Geely-Volvo and Tata-Jaguar are examples of cases where plants from one brand are essentially irrelevant for the other brand's production. Columns (3) and (4) show that deep RTAs promote variety and firm entry but with smaller coefficients and higher standard errors. This corroborates our view that the brand/model concept is more appropriate.

Table F.2: Results with the firm-variety approach

Decision:	Sourcing	Market share	Model entry	Brand entry
Dep. Var:	$S_{b\ell nt}$	$\frac{q_{bnt}}{M_{bnt}Q_{nt}}$	$\frac{M_{bnt}}{M_{bt}}$	\mathbb{D}_{bnt}
Method:	PPML	PPML	frac. probit	probit
	(1)	(2)	(3)	(4)
home $_{\ell n}$	0.891 (0.285)			
ln dist $_{\ell n}$	-0.251 (0.073)			
language $_{\ell n}$	-0.066 (0.112)			
ln (1+ car tariff $_{\ell n}$)	-7.725 (0.905)			
Deep RTA $_{\ell n}$	0.184 (0.131)			
home $_{i\ell}$	1.447 (0.400)			
ln dist $_{i\ell}$	0.135 (0.106)			
language $_{i\ell}$	-0.064 (0.314)			
ln (1+ parts tariff $_{i\ell}$)	-11.109 (3.024)			
Deep RTA $_{i\ell}$	-0.577 (0.303)			
ln q_{ℓ}	0.227 (0.063)			
home $_{in}$		0.578 (0.283)	0.249 (0.082)	0.597 (0.520)
home $_{in} \times \text{LDC}_n$		0.882 (0.461)	1.144 (0.113)	3.689 (0.609)
ln dist $_{in}$		-0.371 (0.095)	-0.079 (0.022)	-0.027 (0.130)
language $_{in}$		0.253 (0.198)	0.085 (0.055)	0.006 (0.198)
Deep RTA $_{in}$		0.064 (0.113)	0.079 (0.034)	0.095 (0.155)
ln C_{bn}		-1.913 (0.995)	-0.279 (0.244)	
ln $\mathbb{E}[\pi_{bn}]$				0.932 (0.127)
Observations	281583	21932	21933	62472
R^2	0.781	0.594	0.761	0.707
Fixed effects:	ℓ, bnt	b, nt	b, nt	b, nt
S.E. cluster:	ℓ	f	f	f

Standard errors in parentheses. r^2 is squared correlation of fitted and true dependent variables except in specification (4) where the pseudo- r^2 is reported. Each regression controls for log per-capita income and price level of the assembly country.

F.3 Sourcing parameter estimates from quantities

We can express the expected sales of a brand to a given market from any of the country ℓ where it is producing by simply multiplying expected sales of b in n (9) by the probability it sources its models from ℓ (6).

$$\mathbb{E}[q_{b\ell n} | \mathbb{D}_{bn} = 1, \mathbb{L}_{b\ell} = 1] = \kappa_1 (\gamma_{i\ell} \tau_{\ell n})^{-\theta} \times (w_\ell^\alpha q_\ell^\zeta)^{-\theta} \times (\varphi_b / w_i^{1-\alpha})^\eta Q_n P_n^\eta M_{bn} \delta_{in}^{-\eta} C_{bn}^{-\eta-1}. \quad (\text{F.1})$$

Equation (F.1) can be used to obtain additional sets of estimates of γ frictions, with the difference that they combine the extensive margin of the sourcing equation with the intensive value of sales in each market the brand serves. The regression includes two sets of fixed effects, one for the country of origin, and one for the brand-destination-time combination, which takes into account the third term in (F.1). There are three ways to specify the LHS of the regression. As when estimating total sales of b in n , we can use the unitary coefficient prediction to divide $q_{b\ell n}$ by $M_{bn} Q_n$. Alternatively, we can let fixed effects absorb M_{bn} and Q_n without imposing the constraint, or have an intermediate approach where the dependent variable is market share $q_{b\ell n} / Q_n$.

We also evaluate the robustness of estimates regarding whether (first 3 columns) or not (last 3 columns) external economies of scale are considered. The main takeaway from Table F.3 is that the coefficients on Deep RTA $_{i\ell}$ and on tariffs on car parts are stronger and more significant than in our baseline results. However, since the estimates of θ are also larger, the AVE of deep RTA remain very similar to the baseline. The ratio of coefficients between car and parts tariffs also provides comparable alternative estimates of $1 - \alpha$ ranging between 29% and 50%.

G Fit of the DEV simulation


Figure G.1 shows the fit of one run of the DEV simulation under the factual set of tariffs and RTA policies. The equilibrium price index is a key element of the model, (inversely) summarizing the degree of competition on each destination, once all actors have solved for the optimal sourcing, model entry and brand entry choices. Its fit with estimated one shown in panel (a) is quite remarkable: Japan, Korea, Germany, China and the US are among the most competitive markets, while Iran, Vietnam and Egypt are at the other end of the spectrum. Producers present in those latter markets are still protected by very high tariffs, which lowers entry and overall competition for local consumers. Panel (b) graphs true brand-origin-destination sales against simulation-predicted sales with both expressed on a log scale. The data cluster around the 45-degree line, obtaining a correlation (in logs) of 0.63. Panel (c) aggregates flows at the country-pair level, and the fit is even more impressive, at 0.74. Part of the high explanatory power stems from the presence of Q_n in the prediction in both graphs. Nevertheless, the figure does show that the estimated model captures the main variation in the data, whereas failure to do so would have raised concerns about its suitability for conducting counterfactuals. Panel (d) further aggregates and shows the equilibrium output of each producing country against the true one in 2016. Black hollow circles simply

Table F.3: Bilateral brand sales regression provide alternative estimates of γ

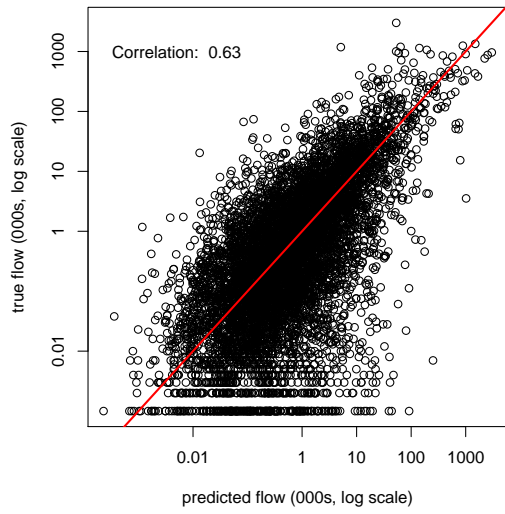
Dep. Var:	$\frac{q_{bln}}{M_{bn}Q_n}$	q_{bln}	$\frac{q_{bln}}{Q_n}$	$\frac{q_{bln}}{M_{bn}Q_n}$	q_{bln}	$\frac{q_{bln}}{Q_n}$
	(1)	(2)	(3)	(4)	(5)	(6)
Trade costs						
home $_{\ell n}$	1.497 (0.277)	1.738 (0.309)	1.392 (0.264)	1.506 (0.276)	1.746 (0.315)	1.397 (0.264)
ln dist $_{\ell n}$	-0.742 (0.079)	-0.609 (0.105)	-0.687 (0.074)	-0.746 (0.079)	-0.606 (0.106)	-0.688 (0.074)
language $_{\ell n}$	0.031 (0.180)	0.145 (0.175)	-0.044 (0.146)	0.025 (0.181)	0.155 (0.176)	-0.047 (0.146)
ln (1+ car tariff $_{\ell n}$)	-10.878 (0.803)	-12.999 (1.821)	-11.722 (0.960)	-10.882 (0.798)	-12.943 (1.839)	-11.730 (0.965)
Deep RTA $_{\ell n}$	0.535 (0.158)	1.039 (0.211)	0.523 (0.168)	0.535 (0.157)	1.038 (0.213)	0.520 (0.167)
MP frictions						
home $_{i\ell}$	2.530 (0.526)	1.399 (0.409)	2.194 (0.566)	2.540 (0.532)	1.373 (0.425)	2.206 (0.570)
ln dist $_{i\ell}$	0.229 (0.139)	-0.088 (0.101)	0.210 (0.126)	0.219 (0.139)	-0.109 (0.104)	0.201 (0.126)
language $_{i\ell}$	0.034 (0.330)	0.156 (0.267)	-0.040 (0.282)	0.021 (0.328)	0.132 (0.274)	-0.049 (0.279)
ln (1+ parts tariff $_{i\ell}$)	-4.225 (2.153)	-6.548 (2.219)	-4.898 (2.364)	-3.242 (2.146)	-5.082 (2.291)	-3.913 (2.349)
Deep RTA $_{i\ell}$	0.652 (0.291)	0.604 (0.286)	0.641 (0.335)	0.667 (0.294)	0.609 (0.295)	0.656 (0.338)
log current country output				0.268 (0.078)	0.652 (0.054)	0.292 (0.076)
Observations	375473	375473	375473	375473	375473	375473
rsq	0.927	0.943	0.832	0.926	0.944	0.833

Estimation with PPML. Standard errors in parentheses, clustered by origin country. All regressions have origin and brand-market-year fixed effects. r^2 is squared correlation of fitted and true dependent variables except in specification (4) where the pseudo- r^2 is reported. Each regression controls for log per-capita income and price level of the assembly country.


Figure G.1: Fit of the DEV simulation: predictions vs data



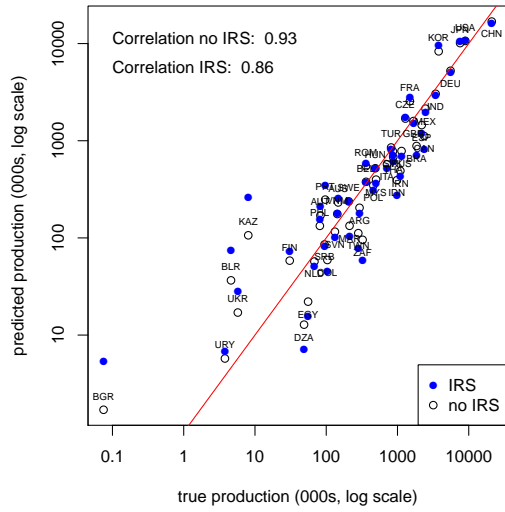
(a) Price indices (P_n)



(b) Brand flows (q_{bn})



(c) Trade flows (q_{ln})



(d) National output (q_{ℓ})

re-iterate the global ability of our model to explain global patterns of the data. The blue ones show equilibrium in the IRS situation, where the output of a country feeds back into lower production costs, and requiring an outer loop to solve for the vector of production. The difference between the two scenarios is very clear: when the model predicts that a country should produce more than its actual production (which serves as an initializing guess), this difference is amplified by the scenario with endogenous external economies of scale. Conversely, initial negative deviations are worsened.

Figure G.2: Fit of the DEV simulation: predictions vs data

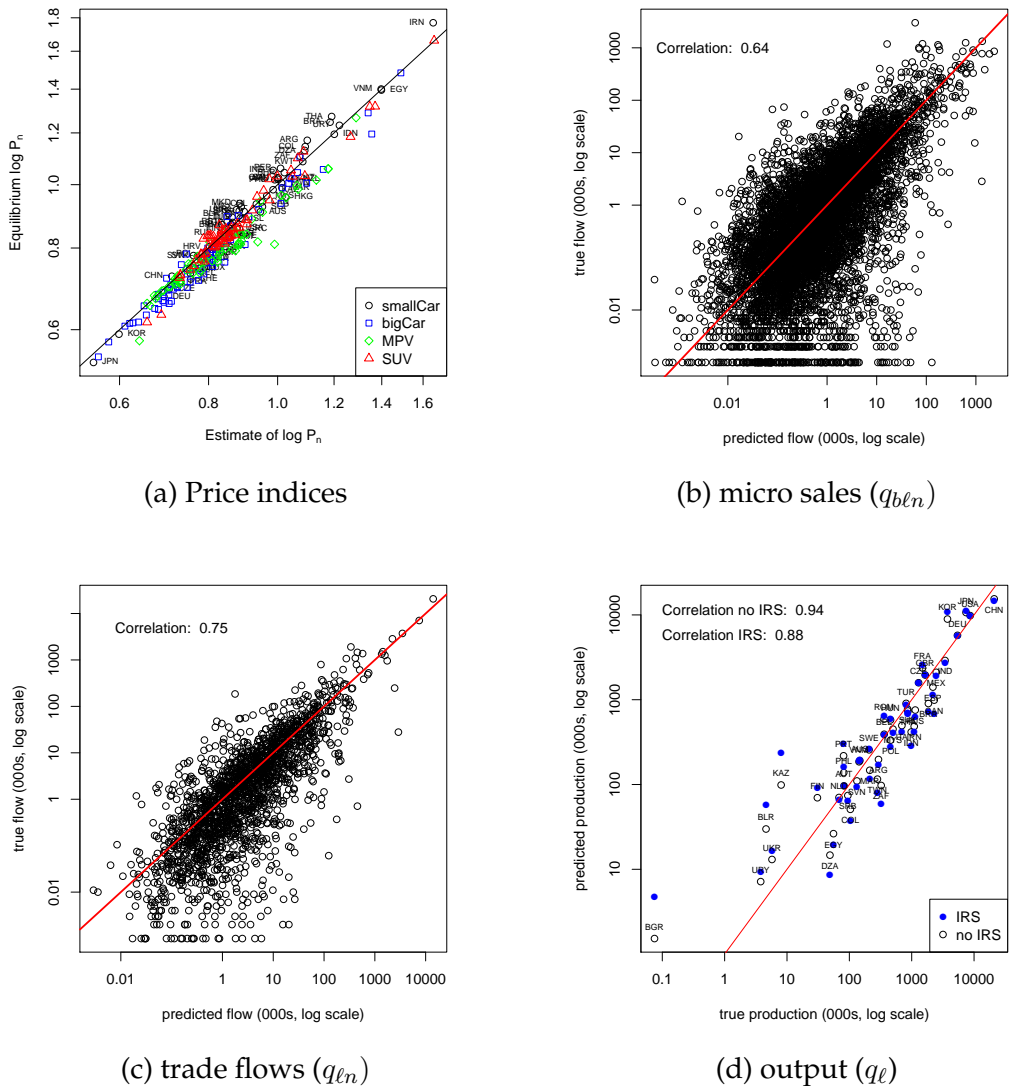


Figure G.2 shows the fit of the segmented market version of DEV. The correlations between the true flows and the flows predicted by the model are very similar and even slightly higher than in the unified market case.

H Evidence on capacity constraints in the medium-run

Our model does not have rising marginal costs due to plant-level capacity constraints. This assumption underlies our discrete choice formulation of the sourcing decision. If rising marginal costs were indeed a key feature of the data, the sourcing decision would require equating marginal costs across plants and the firm would often use multiple plants to serve the same market. In our data we observe 98% single-country sourcing. Thus it is very rare for firms to source the same car from two different countries. This fact is in line with the assumption of constant (or decreasing) marginal costs.

Our model allows for non-constant returns to scale that are external to the firm. Our estimates imply *increasing* returns, which we interpret as arising from Marshallian effects such as labor market pooling and, especially endogenous numbers of input suppliers. These effects would work to offset the short-run tendency of marginal costs at the plant-level to rise following a demand shock. Whether the industry-level increasing returns or plant-level decreasing returns dominate in the aggregate should depend on the time frame of the counterfactual.

Our notion of the “medium run” is the period in which the brand can adjust all four of the margins we estimate (The short run would involve only intensive margin choices and the long run would allow for adding or dropping countries in the production choice set). Medium-run does not correspond to a specific amount of calendar time. However, we have observed that most changes in the sourcing decision tend to happen when the brand introduces a new model generation. The modal duration of a program is six years and only one third of the models last longer than that. We therefore think of the medium run as approximately six years.

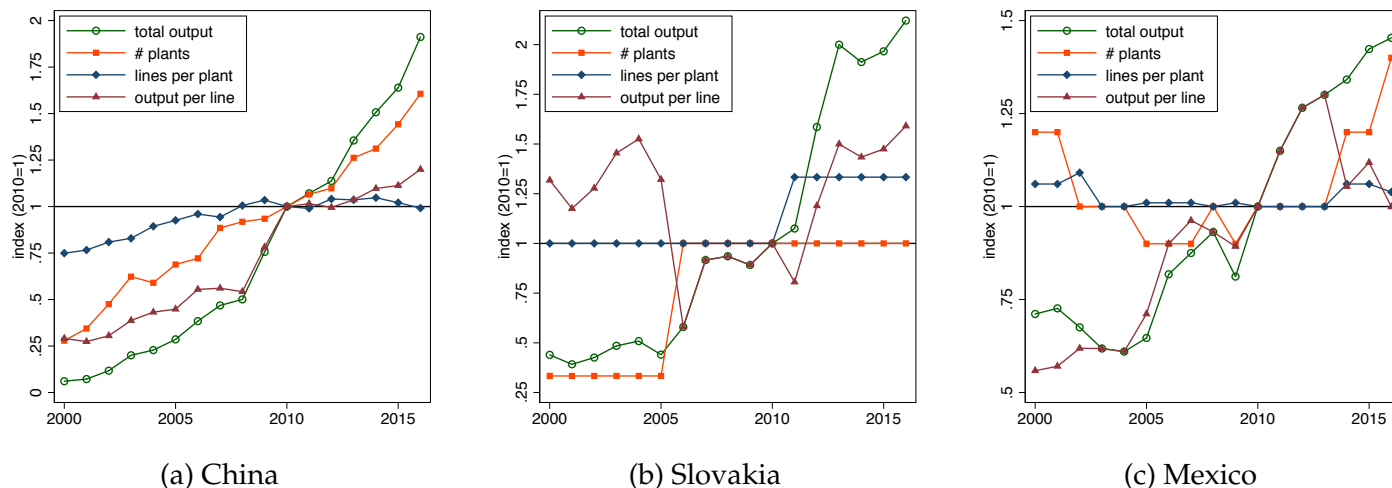
We do not have direct evidence on the shape of the marginal cost curve over this medium-run scenario. However, to the extent that capacity constraints are binding and marginal costs sharply increasing in output, we would expect not to see any large increases in national car output over 5–7 year time frames.

There are two relevant cases for evaluating the potential for large output increases. The first is where the country’s factories already have substantial excess capacity. While we do not observe capacity in our data, we use the maximum past production level as a proxy. This definition was inspired by Bresnahan and Ramey (1993) which used the maximum historical number of production shifts. Their approach does not allow for increases in line speed which Bresnahan and Ramey (1994) observe to “have a sizable impact on the variance of output at quarterly frequencies.” At annual frequencies we would expect even greater scope for line speed increases since Bresnahan and Ramey (1994) report that line speed increases are mainly obtained by adding workers to the line.

Using the past-peak measure of capacity, we find that many countries in our study have substantial excess capacity in 2016. Poland’s 2016 production was just 58% of its 2010 maximum and Belgium in 2016 had fallen to 41% of its 2007 peak. France and Italy are in a similar situation at with current quantities 44% and 51% of past peaks. Hence, for these countries we would not anticipate any significant limits to responding to higher demand from either Trans-Atlantic in-

tegration or Brexit. Another important country in our counterfactuals with current production under capacity is Japan (84%).

Figure H.1: Output increases on 3 margins in China, Slovakia, and Mexico



The second case of interest are countries producing quantities near their historical maxima. Important examples include China, Korea, Germany, Slovakia, the US, Canada, and Mexico. For these countries to expand output by large amounts they would have to construct new plants, add production lines to existing plants, or increase production per line (either by adding workers to increase line speed or adding extra shifts). To judge how feasible this might be, Figure H.1 decomposes output growth for three countries that have dramatically expanded their car industries in the last 17 years: China, Mexico, and Slovakia. These countries exhibit substantial growth from each of these sources, depending on the time range. In the case of China, rising incomes led to an astounding 32-fold increase in car production between 2000 and 2016, an annual rate of increase of over 23%. While the rate of growth has abated, Chinese car output still managed to grow by 91% from 2010 to 2016, roughly corresponding to our medium-run scenario. It did so mainly by increasing the number of plants (37 were opened in the 7-year period) but output per production line also grew by 20%. In the early 2000s growth was more evenly divided between new plants, new production lines, and expand output per line.

China is undoubtedly an extreme case, but even countries serving mature markets such as Slovakia and Mexico can be used to illustrate the potential of countries to increase production rapidly in the medium run. Slovakia experienced a boom in new car investments following accession to the European Union in 2004. From 2004 to 2010, output doubled. By 2016, car production had doubled again. Over the longer time frame, all the production increase can be attributed to new plants and production lines as output per line did not increase. However, since 2011 production has increased entirely on the intensive margin with a 97% increase in output per production line. The Mexican car industry traces its success back to NAFTA in 1993 but it increased production by a factor of 2.4 since 2004. For the first 9 years, all the increases were on the intensive margin but

since 2013 four new plants were opened in Mexico along with one extra production line.

To summarize, there are three aspects of the data suggesting that capacity issues are unlikely to pose binding constraints in the medium-run (≈ 6 year) time frame relevant for our policy counterfactuals. First, single-country sourcing is almost universal. Second, our estimates support downward sloping industry cost curves. Third, large output increases, featuring increases in plants, production lines, and output per line, have been observed in multiple countries over 6-year time frames.

I EHA version of counterfactual scenario results

In this section of the appendix, we report results using the Exact Hat Algebra (EHA) approach described in the text, and in more details in appendix C. We also highlight the differences between EHA and DEV approaches to counterfactual analysis.


A key pattern is that the source of quantitative differences between EHA and DEV almost always arises from the different ways the two methods deal with a zero realized flow. EHA is rightly praised for its lower informational requirements. It also allows the counterfactual to implicitly hold constant deviations from expected values that might arise from unobservable frictions. To obtain bilateral zero flows, EHA implicitly assumes infinite frictions. As a result, any country ℓ that fails to be chosen by brand b to supply any models to market n , in the actual data will also remain a zero flow in under any counterfactual policy. In DEV, on the other hand, the computed flows (both factual and counterfactual) are expected values (and therefore greater than zero so long as the brand enters the market).

I.1 Trumpit and Section 232

The Trumpit counterfactual is an interesting case to study the role of zero flows and how the two methods deal with them. For example, Ford does not actually send cars from its Spanish or Polish factories to the North American market. The large British plants of Toyota and Nissan did not send a single car in 2016 to any of the three NAFTA members. In EHA, this is interpreted as a huge and persistent friction. As a consequence, the US market carries out zero substitution from Nissan’s Mexico factory to its UK factory in response to Trumpit. However, Ford, Toyota and Nissan (among many others) have positive *expected* flows from their European plants to North America in DEV. Trumpit leads them to source substantially more from those EU-based operations to serve North-American markets. For Poland, Spain, and the UK, the EHA counterfactual shows a much weaker response compared to DEV as expected, reflecting the fact that the latter method accounts for the possibility of UK-made Nissans to be sold in the US. The most extreme case is Poland for which DEV increases exceeds those of EHA by up to seven percentage points.

India’s Trumpit outcomes reveals the way EHA implicitly takes into account model “residuals,” i.e. deviations from expected values. Since India is a high-cost producer and faces near 34% MFN tariffs on exports to Mexico, we would not expect it to be much of an exporter to Mexico.

Figure I.1: Trumpit and Section 232 (Exact Hat Algebra method)



But in fact Chevrolet, Ford, Volkswagen, Hyundai, and Suzuki all export large amounts. Trumpit implies that US-made cars face 25% tariffs on its exports to Mexico. This creates an expansion in Indian auto manufacturing of around 4% (74 th. cars). DEV does not build this unexplained ability of India to export cars to Mexico into the counterfactual (or factual). As a consequence, it only predicts an increase of 0.65% for India. The case for Spain is simply the reverse of India. Spain is a low-cost maker with zero-tariff access to Mexico. It would be expected to be a major exporter to Mexico before NAFTA but in fact it only exports about 25 thousand cars (implying a bad “residual” under EHA). Hence the average increase is almost ten times as large under DEV (+3.84%) than under EHA (+0.3%).

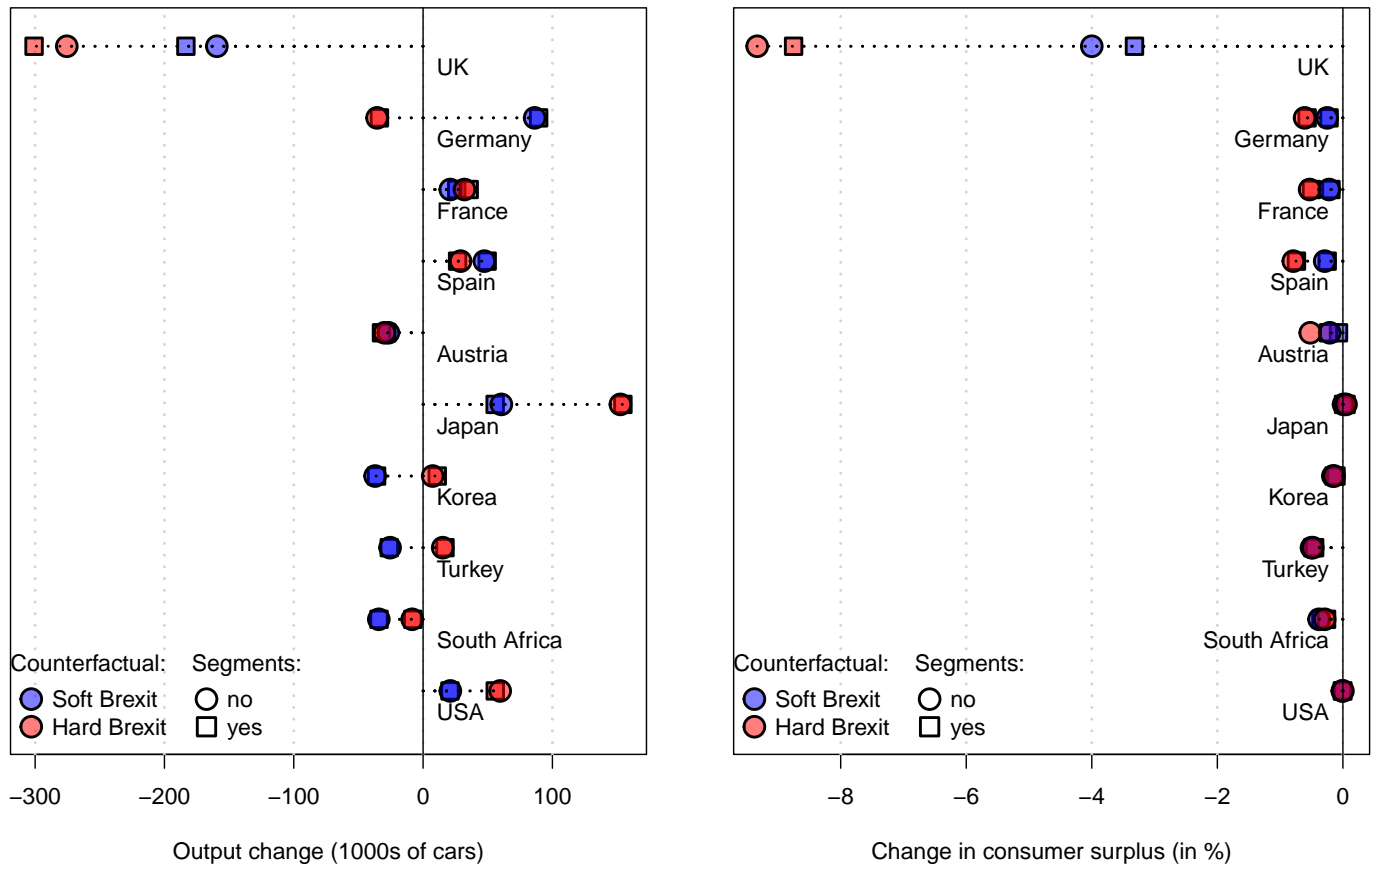
The net effect of cases like that of Spain is adding to the US output gains under EHA. This is because EHA rules out re-sourcing to a number of countries in Europe where US-owned plants do not currently send positive amounts of cars to the US. Therefore, under EHA, US factories do not have to share the gains from reduced access to Canadian and Mexican plants with their EU counterparts. Under DEV, the EU plants gains are large enough to generate a small net loss in US production under Trumpit. For the same reason, DEV dampens losses to US car buyers, who benefit from the option of importing EU-made cars. It seems plausible to us that the DEV outcomes for the US are the more likely ones to prevail in the medium-run as multinationals increase sourcing from their European factories to reflect the loss of preferential access of Mexican and Canadian plants to the US market.

The principal winners and losers in terms of output changes in the Section 232 scenario are the same under EHA and DEV. Mexico and Canada gain even more under EHA, 57% and 77% (as compared to 30% and 40% with DEV). On the flip side, Japanese, Korea, and Germany all lose more. The exception to the rule is the US, where car production is predicted to *fall* by 3.5% under EHA, whereas DEV had predicted a 26% increase. The first-order explanation for these differences is that DEV understates the shares of the US market served by Mexican and Canadian plants. In 2016 data those shares are 10% and 14%, much higher than DEV predictions of 4.3% and 2.9%, respectively. The emphasis EHA puts on initial flows as proxies for persistent unobserved frictions accounts for why factories in Canada and Mexico expand so much when Section 232 raises the costs of serving the US from other locations. While this might seem like a point in favor of EHA, these large expansions are not as plausible given features of the revised NAFTA that our simulations do not consider, namely tighter rules of origin and quotas on the amount of cars exported to the US that would escape Section 232 tariffs.

I.2 United Kingdom exits the European Union (Brexit)

Figure I.2 points to outcomes for both the UK car industry and the British car buyer that are much worse under EHA than under DEV. The industry contracts by 18% under hard Brexit, and still by 11% for Soft Brexit. In absolute terms, losses for the UK would be 183 thousand cars under Soft Brexit and more than 300 thousands under Hard Brexit. Consumer losses also rise slightly under EHA.

Figure I.2: Brexit



The structure of multinational production and British consumer tastes both contribute to make Brexit even more detrimental to the UK production in the EHA case than under DEV. Major German brands such as VW, BMW and Mercedes-Benz lack UK factories. This limits the scope for re-orienting the supply for UK demand to UK production locations (BMW and Mercedes-Benz for instance re-optimize their sourcing strategy to increase the share of models imported from their US plants). Because of the positive residual in the German brands' markets shares in the UK, this effect is magnified under EHA.

Aside from the UK, the big losers from Soft Brexit are South Korea, Turkey, and South Africa. All three countries suffer from the loss of tariff-free access to the UK that they obtained through trade agreements with the EU. South Africa is especially hard hit in the EHA setting (around -11%) because it has a surprisingly high sourcing share in the data (18 times higher than South Africa's exports to France for example). South Africa's production losses under EHA are large enough to trigger add-on costs from lower scale economies. Turkey's situation reverses the EHA/DEV differences (Turkey's share in UK purchases are only a third of their share in France).

The Brexit counter-factual is also valuable to illustrate how scale economies lead to market interdependencies. In a constant returns world, Honda's sourcing decision for sales in the US would be unaffected by Brexit. This is because there would be no tariff changes and, as a non-EU brand, no γ changes. With increasing returns, the smaller scale of the UK car industry raises Honda's UK plant's relative costs. The simulation predicts that US will lower its UK sourcing probability for the Honda Civic by 5%. Bigger effects arise when friction changes are combined with scale economies. The US probability of sourcing Mini hatchbacks and convertibles from the Netherlands falls by 32% when Brexit raises the costs of assembling the Mini in a no longer deeply integrated RTA by an estimated 6%. Increasing returns would magnify the reduction to 37%.

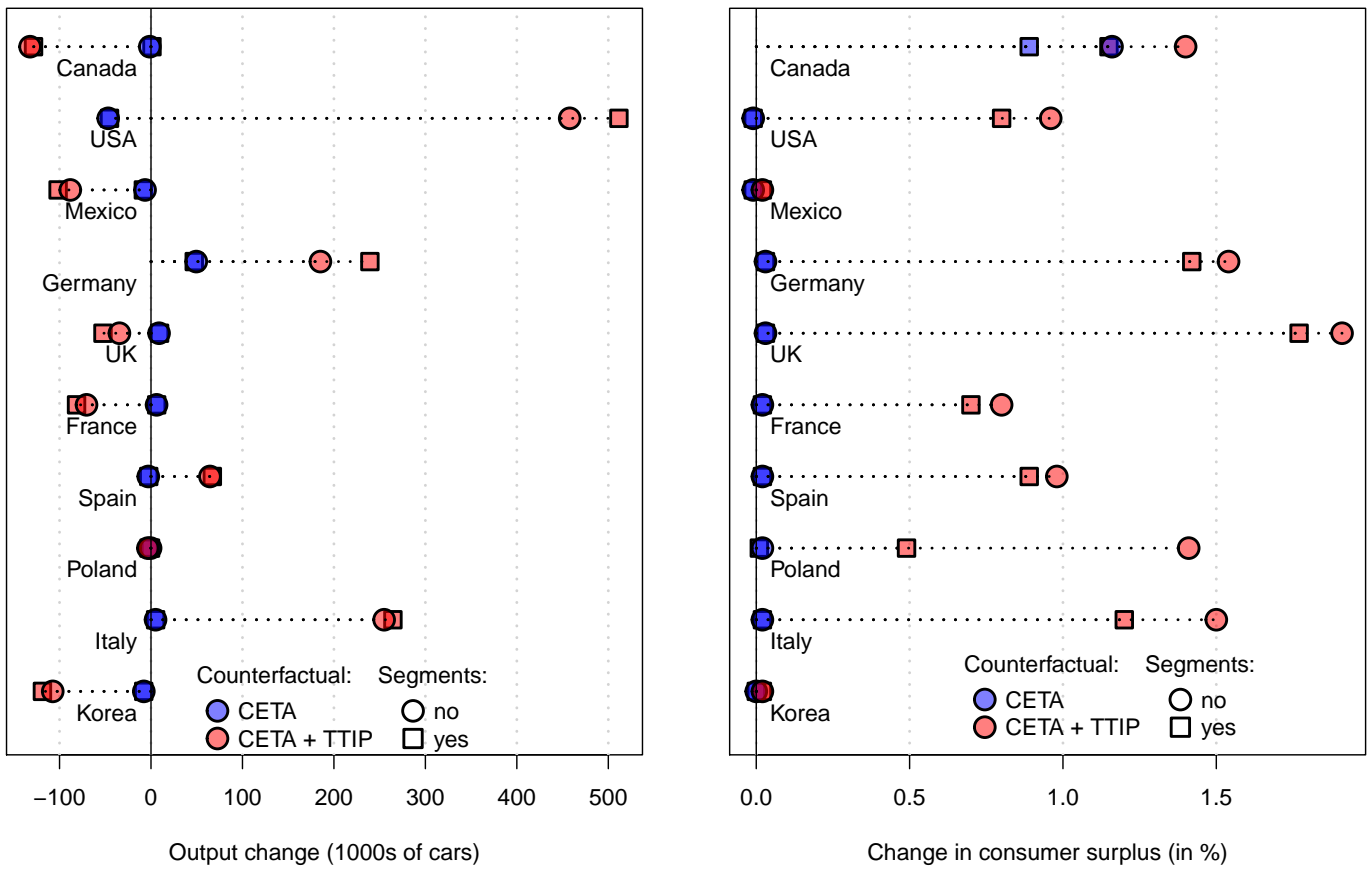
I.3 Trans-Atlantic Integration

Figure I.3 shows the effect of a deep trade agreement between the European Union (all current 28 members) and Canada (CETA) as well as an extension including the US (CETA+TTIP). Our counterfactuals point to very different outcomes from CETA for Canada under EHA and DEV. The DEV method predicts an 8% production increase, while under EHA, Canadian production is essentially unchanged. The reason EHA predicts such small effects is that Canada has negligible exports to the EU. Our estimates predict it should export about one percent of cars to the UK, for example. The share in actual data is one tenth that. Similar ratios apply to other EU destinations. Canada does not import large shares from the EU either, so under EHA only Germany and Italy experience sizable output changes.

Not surprisingly, including the US in trans-Atlantic integration leads to much bigger impacts, just as it did in the DEV method reported in the main text. As with DEV, consumers on both sides of the Atlantic benefit by about one percent, on average. Also in common with DEV, the Asian exporters Korea and Japan (not shown) lose from the changes in preferential market access.

The most dramatic difference between DEV and EHA is that in the latter, auto production in

Figure I.3: Trans-Atlantic Integration (Exact Hat Algebra method)



the US *increases* when it is included in trans-Atlantic integration. Once again, the explanation comes from the fact that in the 2016 data many brands that manufacture in Europe have zero exports to United States. EHA implicitly assumes prohibitively large frictions and does not allow these brands to respond to the export opportunities created by TTIP. Without this added competition, there is negligible offsetting of the roughly 400 ths. cars in exports to Europe which the US gains under both EHA and DEV.

Under CETA+TTIP, Italy stands out as a major gainer (45%) for EHA and this comes entirely from one interesting fact about 2016 car flows. Fiat factories in Italy successfully export Jeeps to the US. EHA allows this to expand massively. In contrast, Ford's factories in Poland *should* be exporting already to the US and therefore rapidly expand through a mix of τ and γ effects. In practice, Ford uses this plant exclusively to serve EU car buyers.

I.4 Trans-Pacific Integration

Figure I.4: Trans-Pacific Integration (Exact Hat Algebra method)

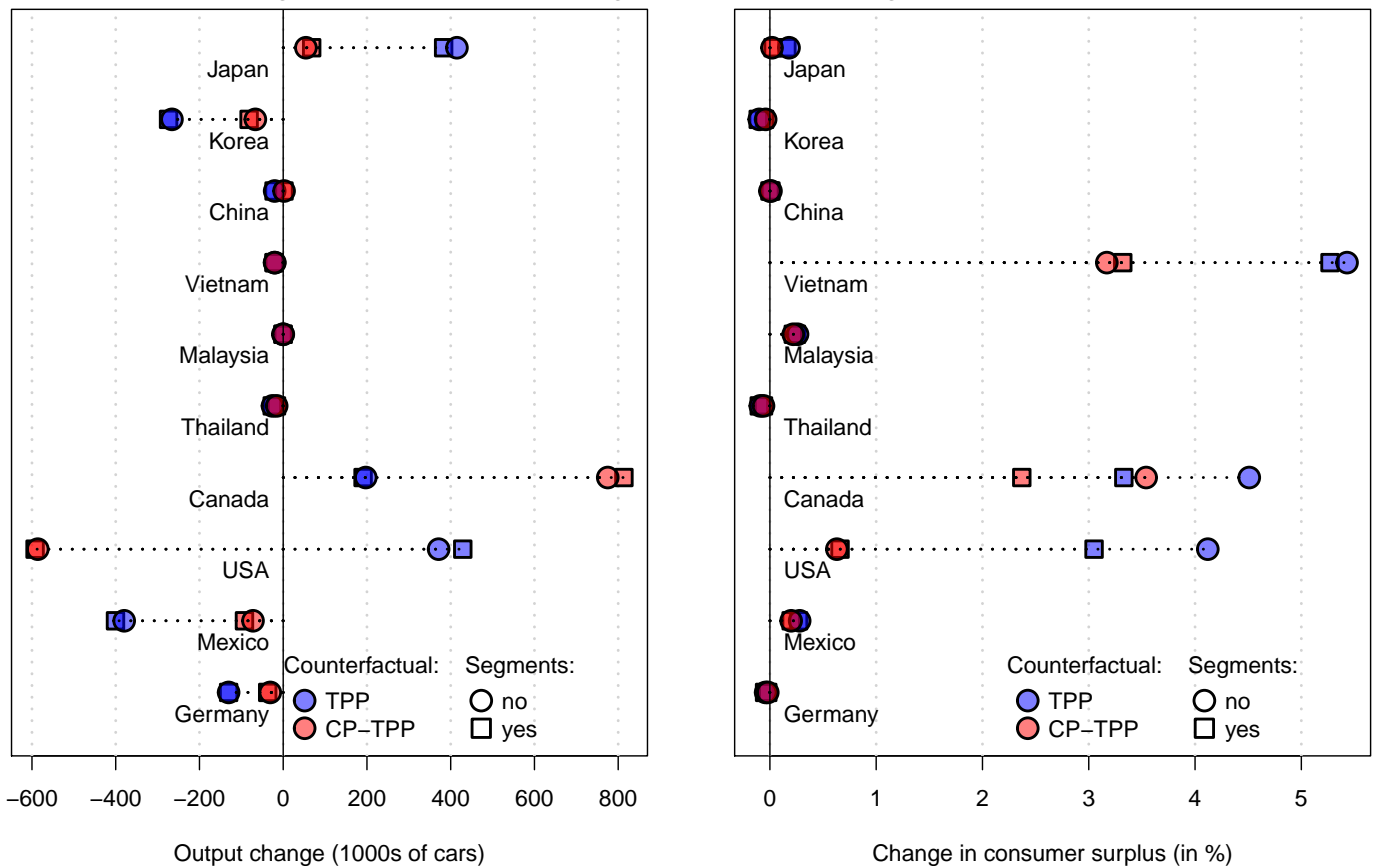


Figure I.4 shows the EHA outcomes for trans-Pacific integration, again with and without the US on board. The most striking outcome of either a TPP or CPTPP is the large predicted gain in Canadian car production. It would be tempting to attribute those gains to the elimination

of high tariffs protecting Vietnamese and Malaysian markets and non-tariff barriers on the large Japanese market. Under CPTPP, Canadian car plants would certainly gain preferential access to those markets relative to the US, although in practice these plants would have to find a way to attain the 45% CPTPP rule of origin when denied the ability to count US inputs as part of the regional content. These considerations turn out to be irrelevant under the EHA method because Canada's actual exports to Vietnam and Malaysia are zero and Canada's entire exports to Japan amount to just 288 cars (a mix of Cadillacs and Chryslers).

Vietnam under TPP or CPTPP is a showcase for how different EHA and DEV solutions can be. Under DEV, Vietnam increases production with the TPP scenario, through a combination of improved access to the US market and γ effects boosting the expected sales of US brands with operations in Vietnam. In reality, Vietnam production is 100% oriented towards the local market. This implies no gains in exports to the US under EHA. However, the Japanese makers present in Vietnam will radically increase their sourcing from Japan, leading to the production losses we see for Vietnam under EHA.

I.5 Summary of EHA differences from DEV

An important overall conclusion of the EHA result is that the choice between solution methods is not innocuous. It often has serious quantitative effects and sometimes changes the sign of the output changes. Fortunately, the direction of the consumer surplus changes caused by our policy experiments are the same whenever they are non-negligible (over 0.5% in absolute value in either method). Even their magnitudes tend to be fairly robust across methods: the median absolute difference among the non-negligible changes is 0.6%. The output differences come in large part from the way brand-market zero flows are handled under EHA. We see these discrepancies as motivating the search for empirical evidence on whether EHA or DEV counterfactuals are more accurate in practice.

J Tables detailing counterfactual scenario results

In this section, we provide four tables in which the counterfactual results are detailed for a wider range of countries than in Figures 6 to 9 in the text for DEV and Figures I.1 to I.4 for EHA in the appendix. For each set of counterfactual scenarios (Trumpit/Brexit/Transatlantic Integration/Transpacific Integration), we give the unified and segmented versions for each of the DEV and EHA methods.

Table J.1: Trumpit DEV (unified)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
25% duties and loss of deep RTA						
MEX	-10	-569	-228	-806	-40.65	-4.77
CAN	-47	-412	-87	-546	-67.74	-6.22
JPN	11	343	60	415	3.93	0.09
KOR	2	268	52	322	3.32	0.10
DEU	5	117	24	146	2.88	0.05
ESP	2	36	10	48	4.01	0.01
GBR	3	35	6	45	2.95	-0.00
BRA	28	12	4	44	6.41	-0.55
USA	568	-634	24	-42	-0.39	-0.84
FRA	1	34	5	39	1.40	0.02
CZE	0	29	7	36	2.04	0.03
BEL	0	25	8	33	8.90	0.03
POL	0	21	7	28	9.06	0.01
TUR	0	17	5	22	2.74	0.04
IND	5	15	2	21	0.72	-0.00
25% tariffs applied on major countries exc. Canada & Mexico						
USA	2707	-11	-629	2067	19.34	-4.48
KOR	14	-1320	-170	-1476	-15.23	-0.75
MEX	23	431	140	594	29.95	-0.30
JPN	37	-744	120	-587	-5.57	-0.36
CAN	21	228	42	290	35.98	-0.66
DEU	14	-325	24	-287	-5.66	-0.46
BEL	-3	-80	-40	-122	-32.83	-0.60
POL	-2	-65	-31	-98	-31.95	-0.64
ESP	0	-85	-8	-92	-7.77	-0.68
CHN	122	-50	18	90	0.55	-0.19
AUS	-2	-66	-17	-85	-34.46	-1.39
FRA	21	-26	76	70	2.51	-0.36
GBR	17	-87	22	-49	-3.21	-0.74
BRA	-9	-41	4	-46	-6.59	-0.35
KAZ	0	-33	-5	-38	-14.56	-1.10

The RTA column sums USA, Canada, and Mexico (excluding domestic shipments).

Table J.2: Trumpit DEV (segmented)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
25% duties and loss of deep RTA						
MEX	-8	-512	-212	-733	-38.43	-4.08
CAN	-49	-343	-82	-474	-67.41	-5.26
JPN	11	350	62	423	3.72	0.08
KOR	2	286	62	350	3.24	0.09
USA	415	-644	6	-223	-2.31	-0.79
DEU	5	127	22	155	2.76	0.05
GBR	3	49	8	59	3.06	0.03
FRA	2	45	8	55	2.13	0.03
ESP	1	34	9	44	3.84	0.02
BRA	26	12	4	42	5.96	-0.51
CZE	0	31	8	39	2.39	0.03
BEL	0	20	5	26	6.67	0.05
TUR	0	18	5	23	2.64	0.03
SVK	0	16	4	20	2.83	0.02
POL	0	15	4	20	6.99	0.01
25% tariffs applied on major countries exc. Canada & Mexico						
USA	3205	-5	-639	2561	26.48	-5.18
KOR	17	-1312	-172	-1467	-13.55	-0.74
JPN	38	-965	122	-805	-7.09	-0.43
MEX	23	416	128	567	29.73	-0.30
DEU	8	-501	14	-478	-8.53	-0.54
CAN	18	230	33	282	40.02	-0.59
GBR	11	-181	23	-147	-7.56	-0.57
CHN	168	-54	20	134	0.92	-0.24
BEL	-3	-82	-44	-129	-33.70	-0.62
ESP	-0	-83	-10	-93	-8.20	-0.69
POL	-2	-54	-34	-89	-31.76	-0.14
FRA	24	-21	83	86	3.35	-0.43
AUS	-2	-43	-15	-60	-30.86	-1.14
TUR	2	-68	15	-52	-5.98	-0.43
BRA	-8	-44	3	-49	-6.90	-0.37

The RTA column sums USA, Canada, and Mexico (excluding domestic shipments).

Table J.3: Trumpit EHA (unified)

Country	Change in shipments (ths. cars) to:			% Chg.	CS % Chg.	
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>			<i>Total</i>
25% duties and loss of deep RTA						
CAN	6	-1836	-34	-1864	-78.80	-9.00
USA	2304	-498	45	1851	20.90	-2.71
MEX	-80	-1170	-101	-1351	-54.60	-4.29
JPN	2	521	23	546	7.20	0.22
KOR	2	322	21	346	9.20	0.26
DEU	4	154	8	166	3.00	0.05
IND	-0	76	3	79	4.50	0.14
BRA	15	16	10	42	2.30	-0.06
GBR	0	32	2	33	2.00	0.07
ITA	2	26	1	28	4.80	0.07
THA	-0	20	-1	19	2.20	0.07
HUN	0	16	1	17	3.40	0.04
SVK	-0	16	0	17	1.90	0.05
CHN	9	4	1	15	0.10	-0.01
COL	7	5	0	12	11.40	-0.76
25% tariffs applied on major countries exc. Canada & Mexico						
CAN	76	1601	30	1707	72.10	-1.48
JPN	4	-1533	66	-1463	-19.40	-0.75
MEX	83	1219	86	1388	56.10	-0.22
KOR	21	-962	0	-941	-25.10	-2.32
DEU	14	-532	-13	-531	-9.60	-1.01
CHN	192	-20	6	178	0.90	-1.11
USA	855	-190	-825	-159	-1.80	-6.95
ITA	-14	-120	-14	-147	-25.00	-1.07
CZE	4	1	66	71	5.50	-0.73
GBR	23	-137	46	-68	-4.10	-1.33
FRA	15	-15	52	52	3.50	-0.51
SWE	-1	-44	-1	-46	-22.30	-0.73
IND	0	-44	0	-43	-2.50	-0.72
TUR	6	-1	37	43	5.20	-0.46
ESP	1	-18	-20	-37	-1.70	-0.70

The RTA column sums USA, Canada, and Mexico (excluding domestic shipments).

Table J.4: Trumpit EHA (segmented)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
25% duties and loss of deep RTA						
CAN	-1	-1780	-35	-1816	-79.10	-8.53
USA	2174	-499	39	1714	19.90	-2.71
MEX	-86	-1181	-105	-1371	-55.40	-3.96
JPN	3	542	26	571	7.70	0.24
KOR	4	359	26	389	10.50	0.29
DEU	4	172	10	186	3.40	0.06
IND	-0	72	3	74	4.20	0.14
BRA	15	18	10	43	2.30	-0.05
GBR	0	37	2	39	2.30	0.08
ITA	2	28	1	32	5.40	0.07
HUN	0	19	2	21	4.60	0.05
THA	-0	21	-1	19	2.20	0.08
SVK	0	18	1	19	2.20	0.05
CHN	11	5	1	17	0.10	-0.01
COL	7	6	0	13	12.80	-0.73
25% tariffs applied on major countries exc. Canada & Mexico						
CAN	81	1659	23	1763	76.80	-1.48
MEX	85	1241	82	1408	56.90	-0.22
JPN	5	-1484	82	-1397	-18.80	-0.73
KOR	22	-939	1	-917	-24.70	-2.19
DEU	5	-510	-73	-578	-10.70	-0.97
USA	683	-192	-789	-297	-3.50	-6.59
CHN	187	-21	5	171	0.90	-1.05
ITA	-18	-120	-24	-162	-27.40	-1.01
CZE	5	1	79	84	6.50	-0.67
FRA	18	-15	64	67	4.50	-0.47
IND	-0	-45	-8	-53	-3.00	-0.73
TUR	8	-1	45	52	6.40	-0.38
GBR	26	-137	60	-51	-3.10	-1.29
ESP	1	-19	-30	-48	-2.20	-0.67
SWE	-1	-46	0	-47	-22.60	-0.43

The RTA column sums USA, Canada, and Mexico (excluding domestic shipments).

Table J.5: Brexit DEV (unified)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
Soft Brexit: shallow FTA (0 tariffs, no deep integration measures)						
KOR	2	-143	-16	-157	-1.61	-0.10
JPN	2	54	35	91	0.86	0.02
USA	9	31	12	53	0.49	-0.00
GBR	116	-78	-85	-46	-3.04	-4.34
DEU	11	19	10	40	0.80	-0.08
TUR	1	-32	-9	-40	-4.96	-0.34
FRA	6	13	7	26	0.91	-0.09
MEX	5	-26	-3	-24	-1.21	-0.33
CZE	1	11	6	17	0.99	-0.07
AUT	-0	-8	-6	-15	-7.17	-0.13
ESP	2	6	4	11	0.96	-0.13
ROM	0	7	3	10	1.74	-0.08
NLD	-0	-5	-4	-10	-18.51	-0.20
MAR	0	-8	-0	-8	-7.89	-0.28
CAN	0	4	3	7	0.86	-0.01
Hard ("no deal") Brexit: 10% tariffs in both directions						
JPN	2	125	41	169	1.60	0.04
USA	9	72	14	95	0.89	0.01
KOR	2	-66	-5	-69	-0.71	-0.06
FRA	10	-60	2	-47	-1.68	-0.29
DEU	21	-73	4	-47	-0.93	-0.26
GBR	255	-201	-84	-29	-1.93	-8.16
CZE	1	-30	2	-27	-1.55	-0.25
AUT	-0	-13	-8	-21	-10.01	-0.34
TUR	1	-14	-6	-19	-2.29	-0.36
ESP	4	-21	1	-16	-1.33	-0.37
RUS	1	11	2	14	2.06	-0.00
IND	1	11	2	13	0.46	0.01
NLD	-0	-7	-5	-13	-24.24	-0.52
CAN	0	9	3	12	1.54	0.01
SVK	0	-12	0	-12	-1.66	-0.33

The RTA column sums the UK and the EU27 (excluding domestic shipments).

Table J.6: Brexit DEV (segmented)

Country	Change in shipments (ths. cars) to:			% Chg.	CS % Chg.	
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>			<i>Total</i>
Soft Brexit: shallow FTA (0 tariffs, no deep integration measures)						
KOR	4	-139	-14	-150	-1.38	-0.09
JPN	2	51	40	93	0.82	0.02
GBR	112	-92	-103	-83	-4.25	-3.37
TUR	-0	-37	-25	-63	-7.20	-0.33
DEU	12	28	14	55	0.97	-0.09
USA	13	18	16	47	0.50	-0.00
FRA	7	20	9	35	1.38	-0.09
CZE	1	12	8	20	1.23	-0.09
MEX	6	-22	-1	-17	-0.90	-0.35
ROM	0	11	4	15	2.52	-0.09
ESP	2	8	5	14	1.26	-0.13
AUT	-0	-6	-6	-13	-13.47	-0.05
NLD	-1	-5	-5	-11	-16.07	-0.15
CHN	7	1	2	10	0.10	-0.00
SVK	0	5	3	8	1.18	-0.10
Hard ("no deal") Brexit: 10% tariffs in both directions						
JPN	3	129	49	180	1.59	0.04
USA	15	51	17	82	0.85	0.01
GBR	290	-259	-103	-73	-3.75	-7.42
KOR	4	-50	-0	-46	-0.42	-0.06
DEU	26	-74	5	-42	-0.75	-0.29
TUR	1	-18	-22	-40	-4.55	-0.35
FRA	14	-53	4	-36	-1.39	-0.33
CZE	1	-27	3	-23	-1.41	-0.33
AUT	-0	-8	-8	-16	-16.75	-0.16
RUS	1	11	2	15	2.35	-0.01
IND	1	11	2	14	0.50	0.00
NLD	-0	-7	-6	-14	-20.99	-0.46
ESP	4	-19	2	-14	-1.18	-0.40
CHN	1	8	2	11	0.10	-0.00
CAN	0	6	3	10	1.40	-0.01

The RTA column sums the UK and the EU27 (excluding domestic shipments).

Table J.7: Brexit EHA (unified)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
Soft Brexit: shallow FTA (0 tariffs, no deep integration measures)						
GBR	68	-162	-65	-159	-9.60	-4.00
DEU	21	50	16	86	1.60	-0.25
JPN	2	46	12	61	0.80	0.01
ESP	5	36	6	47	2.20	-0.29
KOR	8	-45	-0	-37	-1.00	-0.16
ZAF	-0	-31	-3	-34	-10.60	-0.38
NLD	-0	-17	-11	-28	-40.60	-0.29
AUT	-0	-14	-13	-27	-32.40	-0.21
TUR	6	-32	1	-26	-3.10	-0.50
USA	5	13	3	21	0.20	-0.01
FRA	7	12	3	21	1.40	-0.22
CZE	1	13	3	17	1.30	-0.14
IND	0	6	5	11	0.60	0.02
POL	1	7	1	9	2.00	-0.33
ROM	0	5	2	8	2.10	-0.15
Hard ("no deal") Brexit: 10% tariffs in both directions						
GBR	211	-416	-71	-276	-16.50	-9.33
JPN	4	130	18	153	2.00	0.05
USA	7	49	4	60	0.70	0.00
IND	0	36	6	43	2.40	0.08
DEU	39	-80	6	-36	-0.60	-0.61
FRA	16	13	3	32	2.10	-0.53
AUT	-0	-15	-15	-30	-36.00	-0.52
ESP	11	12	6	29	1.30	-0.79
NLD	-0	-16	-12	-29	-42.00	-0.69
TUR	7	6	2	15	1.90	-0.48
THA	0	10	2	13	1.40	0.04
ZAF	1	-9	-0	-8	-2.60	-0.29
SVK	0	-9	0	-8	-1.00	-0.49
KOR	9	-5	3	7	0.20	-0.14
MEX	2	4	2	7	0.30	-0.05

The RTA column sums the UK and the EU27 (excluding domestic shipments).

Table J.8: Brexit EHA (segmented)

Country	Change in shipments (ths. cars) to:			% Chg.	CS % Chg.	
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>			<i>Total</i>
Soft Brexit: shallow FTA (0 tariffs, no deep integration measures)						
GBR	51	-164	-70	-183	-11.10	-3.32
DEU	21	53	16	89	1.60	-0.22
JPN	3	40	13	56	0.70	0.01
ESP	5	39	6	50	2.30	-0.25
KOR	8	-45	1	-36	-1.00	-0.14
ZAF	-1	-31	-3	-34	-10.50	-0.35
NLD	-1	-18	-11	-29	-43.00	-0.23
AUT	-0	-14	-14	-28	-34.30	-0.07
TUR	5	-31	1	-26	-3.20	-0.48
FRA	7	16	3	26	1.70	-0.19
CZE	1	18	3	22	1.70	-0.12
USA	6	11	3	21	0.20	-0.01
SVK	0	12	2	14	1.60	-0.16
POL	1	8	1	10	2.30	-0.16
ROM	0	7	3	10	2.60	-0.10
Hard ("no deal") Brexit: 10% tariffs in both directions						
GBR	210	-432	-78	-301	-18.30	-8.75
JPN	5	129	21	155	2.10	0.05
USA	8	44	4	56	0.60	0.00
IND	0	31	6	37	2.10	0.07
FRA	17	16	3	36	2.40	-0.50
DEU	40	-79	5	-34	-0.60	-0.57
AUT	-0	-16	-16	-32	-39.20	-0.23
NLD	-0	-19	-13	-32	-46.90	-0.57
ESP	11	11	5	27	1.20	-0.74
TUR	6	9	2	17	2.10	-0.45
THA	0	12	2	15	1.70	0.05
ROM	1	8	3	11	3.10	-0.36
KOR	10	-3	5	11	0.30	-0.11
CHN	7	0	0	8	0.00	0.00
ZAF	1	-8	-0	-8	-2.40	-0.26

The RTA column sums the UK and the EU27 (excluding domestic shipments).

Table J.9: CETA (unified)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
CETA						
USA	-10	-56	-2	-68	-0.64	-0.01
CAN	-12	63	13	65	8.03	1.62
DEU	-2	43	3	43	0.85	0.06
KOR	-0	-32	-3	-35	-0.36	-0.01
JPN	-0	-29	-4	-33	-0.32	-0.01
MEX	-1	-14	-3	-18	-0.89	0.00
GBR	-2	16	1	15	1.01	0.08
ESP	-1	9	1	9	0.77	0.09
BEL	-0	7	1	8	2.03	0.07
POL	-0	6	1	7	2.14	0.08
ITA	-1	5	0	5	0.99	0.08
TUR	-0	-4	-1	-5	-0.63	0.01
CZE	-0	4	1	5	0.28	0.06
SVK	-0	4	1	5	0.65	0.07
PRT	-0	4	0	4	1.31	0.11
CETA + TTIP						
DEU	-37	407	69	438	8.65	0.89
KOR	-3	-256	-45	-303	-3.13	-0.08
BEL	5	213	53	271	72.67	1.11
POL	3	178	45	226	73.68	1.21
JPN	-3	-171	-40	-214	-2.03	-0.05
ESP	0	155	40	195	16.35	1.31
USA	-618	416	15	-188	-1.76	1.13
MEX	-11	-102	-22	-136	-6.84	0.37
FRA	-31	-42	-12	-85	-3.02	0.66
TUR	-2	-31	-5	-38	-4.69	0.29
ROM	-1	23	11	33	5.67	0.89
CZE	-4	-20	-8	-32	-1.84	0.71
ITA	-8	32	6	29	5.46	1.04
CHN	-12	-13	-4	-29	-0.18	0.01
RUS	-4	-15	-3	-23	-3.30	0.15
GBR	-33	54	-1	19	1.26	1.33
CAN	-21	19	-5	-7	-0.89	2.13

RTA is EU28 plus Canada in top panel (CETA) and also includes US in the lower panel.

Table J.10: CETA (segmented)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
CETA						
USA	-11	-61	-3	-75	-0.77	-0.01
CAN	-12	56	12	56	7.89	1.77
DEU	-3	41	4	43	0.76	0.06
KOR	-0	-36	-6	-42	-0.39	-0.01
JPN	-0	-29	-3	-32	-0.29	-0.01
GBR	-1	23	3	25	1.27	0.07
MEX	-1	-13	-2	-16	-0.82	0.00
ESP	-1	10	1	11	0.92	0.08
CZE	-0	7	1	8	0.49	0.07
BEL	-0	7	1	8	2.03	0.06
POL	-0	5	1	6	2.11	0.02
TUR	-0	-5	-1	-5	-0.61	0.00
SVK	-0	5	1	5	0.76	0.06
FRA	-2	7	0	5	0.19	0.05
HUN	-0	3	0	3	0.54	0.06
CETA + TTIP						
DEU	-31	361	74	404	7.21	0.81
KOR	-6	-285	-62	-353	-3.26	-0.06
JPN	-5	-220	-63	-288	-2.54	-0.05
BEL	6	188	53	246	64.22	1.01
POL	3	158	48	210	74.56	0.29
ESP	2	140	40	182	15.95	1.14
MEX	-11	-103	-21	-135	-7.10	0.44
FRA	-32	-29	-13	-73	-2.83	0.72
GBR	-33	85	-1	50	2.61	1.06
USA	-499	405	44	-50	-0.52	1.28
TUR	-2	-38	-7	-47	-5.45	0.22
CHN	-14	-16	-4	-34	-0.22	0.03
ITA	-4	30	7	32	7.74	1.00
ROM	-1	14	10	23	3.87	0.68
RUS	-4	-15	-4	-23	-3.60	0.22
CAN	-20	17	-3	-6	-0.80	2.32

RTA is EU28 plus Canada in top panel (CETA) and also includes US in the lower panel.

Table J.11: CETA EHA (unified)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
CETA						
DEU	-1	47	3	49	0.90	0.03
USA	-4	-41	-1	-47	-0.50	-0.01
JPN	-0	-10	0	-10	-0.10	0.00
GBR	-0	8	1	9	0.50	0.03
KOR	-0	-8	-0	-8	-0.20	0.00
MEX	-0	-6	-0	-6	-0.30	-0.01
FRA	0	6	0	6	0.40	0.02
ITA	0	4	0	5	0.80	0.02
HUN	0	4	0	5	1.00	0.01
ESP	-0	-2	-0	-3	-0.10	0.02
SVK	0	3	0	3	0.30	0.01
SWE	0	2	0	2	1.20	0.03
CZE	-0	-1	-0	-2	-0.10	0.01
CAN	-17	14	1	-1	-0.10	1.16
BEL	-0	1	0	1	0.30	0.02
CETA + TTIP						
USA	-155	439	174	458	5.20	0.96
ITA	19	226	11	255	43.20	1.50
DEU	-68	244	9	185	3.30	1.54
JPN	-6	-128	-17	-150	-2.00	-0.04
CAN	-24	-107	-2	-132	-5.60	1.40
KOR	-7	-90	-10	-107	-2.90	0.02
CHN	-90	-2	-1	-93	-0.50	0.05
MEX	-6	-80	-3	-88	-3.60	0.02
CZE	-6	-73	-5	-84	-6.50	1.19
FRA	-21	-46	-3	-71	-4.70	0.80
ESP	-0	49	16	64	2.90	0.98
TUR	-6	-39	-2	-48	-5.80	0.41
GBR	-34	4	-5	-35	-2.10	1.91
SVK	-0	-32	-2	-34	-4.00	1.15
ZAF	-6	-24	-4	-33	-10.30	0.58

RTA is EU28 plus Canada in top panel (CETA) and also includes US in the lower panel.

Table J.12: CETA EHA (segmented)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
CETA						
DEU	-1	44	4	47	0.90	0.03
USA	-5	-40	-1	-45	-0.50	-0.01
JPN	-0	-10	0	-10	-0.10	0.00
GBR	-0	9	1	9	0.60	0.03
MEX	-0	-8	-0	-8	-0.30	-0.01
KOR	-0	-8	-0	-8	-0.20	0.00
FRA	-0	6	0	6	0.40	0.02
ITA	0	4	0	5	0.90	0.02
HUN	0	4	0	5	1.00	0.01
SVK	0	3	0	3	0.30	0.01
ESP	-0	-2	-0	-2	-0.10	0.02
SWE	0	2	0	2	1.10	0.01
CZE	-0	-2	-0	-2	-0.20	0.01
CAN	-15	14	2	1	0.00	0.89
BEL	-0	1	0	1	0.30	0.02
CETA + TTIP						
USA	-125	449	187	512	5.90	0.80
ITA	18	235	12	265	44.80	1.20
DEU	-59	272	27	239	4.40	1.42
JPN	-7	-137	-24	-169	-2.30	-0.04
CAN	-23	-104	-2	-128	-5.60	1.15
KOR	-10	-96	-13	-119	-3.20	0.02
MEX	-6	-92	-4	-102	-4.10	0.02
CHN	-97	-2	-1	-100	-0.50	0.05
CZE	-7	-83	-6	-97	-7.50	1.13
FRA	-24	-52	-5	-81	-5.50	0.70
ESP	-0	50	17	67	3.10	0.89
TUR	-8	-44	-3	-55	-6.80	0.39
GBR	-35	-10	-7	-53	-3.20	1.77
SVK	-0	-41	-3	-45	-5.20	0.89
ZAF	-6	-25	-4	-35	-10.70	0.60

RTA is EU28 plus Canada in top panel (CETA) and also includes US in the lower panel.

Table J.13: TPP (unified)

Country	Change in shipments (ths. cars) to:			% Chg.	CS % Chg.	
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>			<i>Total</i>
TPP with USA						
JPN	-52	700	-12	636	6.03	0.39
KOR	-5	-253	-54	-312	-3.22	-0.05
CAN	19	136	95	250	31.02	3.16
USA	-303	80	136	-88	-0.82	1.49
DEU	-2	-56	-12	-70	-1.38	-0.01
CHN	-50	-14	-3	-67	-0.41	0.05
MEX	-10	-33	-9	-51	-2.56	0.50
VNM	-32	41	41	50	27.66	28.06
GBR	-4	-26	-12	-42	-2.79	0.00
THA	-4	-16	-8	-28	-4.53	0.11
TUR	-1	-16	-9	-26	-3.16	0.04
FRA	-1	-14	-8	-23	-0.83	-0.00
CZE	-0	-12	-9	-21	-1.21	-0.00
IND	-2	-14	-4	-20	-0.67	-0.01
ESP	-1	-14	-5	-19	-1.62	0.03
CP-TPP (without USA)						
CAN	25	68	221	313	38.87	2.93
JPN	-20	235	-29	186	1.77	0.12
USA	-74	-96	-9	-178	-1.67	0.08
KOR	-2	-99	-33	-134	-1.38	-0.01
CHN	-20	-7	-2	-28	-0.17	0.02
MEX	-3	24	6	27	1.37	0.28
DEU	-1	-19	-6	-27	-0.53	-0.01
GBR	-2	-10	-8	-20	-1.31	0.00
THA	-2	-9	-5	-15	-2.48	0.04
FRA	-1	-7	-5	-13	-0.47	-0.00
TUR	-0	-6	-6	-12	-1.53	0.01
CZE	-0	-6	-5	-12	-0.65	-0.00
MYS	-16	7	-2	-11	-2.96	2.00
IND	-1	-6	-3	-9	-0.30	-0.01
AUS	-4	9	3	9	3.63	0.77
VNM	-32	8	16	-8	-4.61	24.62

The RTA column includes TPP12 (including US) in top panel and excludes US in lower panel.

Table J.14: TPP (segmented)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
TPP with USA						
JPN	-41	773	24	756	6.66	0.37
KOR	-8	-292	-82	-382	-3.52	-0.04
CAN	14	122	95	232	32.93	2.28
DEU	-4	-85	-21	-111	-1.98	-0.02
CHN	-65	-15	-3	-83	-0.57	0.07
GBR	-4	-43	-13	-60	-3.08	-0.01
MEX	-10	-38	-7	-56	-2.91	0.47
VNM	-32	34	42	44	23.09	29.12
USA	-252	68	152	-32	-0.33	1.21
THA	-4	-16	-7	-27	-4.04	0.10
TUR	-1	-17	-8	-26	-3.04	0.03
ESP	-1	-16	-6	-23	-2.04	0.03
FRA	-1	-14	-8	-23	-0.87	-0.00
CZE	-0	-12	-7	-19	-1.19	-0.00
IND	1	-15	-4	-18	-0.64	-0.03
CPTPP (without USA)						
CAN	20	60	218	298	42.36	2.07
JPN	-16	250	-6	228	2.01	0.12
USA	-77	-92	-9	-178	-1.85	0.07
KOR	-3	-110	-49	-162	-1.50	-0.02
DEU	-1	-21	-12	-34	-0.61	-0.01
CHN	-22	-6	-2	-30	-0.20	0.03
MEX	-3	23	10	30	1.57	0.25
GBR	-2	-12	-8	-21	-1.11	0.00
THA	-2	-8	-4	-14	-2.09	0.04
MYS	-18	8	-3	-13	-3.20	2.31
FRA	-1	-6	-4	-11	-0.43	-0.00
TUR	-0	-5	-5	-11	-1.22	0.01
CZE	-0	-5	-5	-10	-0.61	0.00
AUS	-2	8	3	9	4.72	0.63
ESP	-0	-4	-3	-7	-0.65	0.01
VNM	-33	8	19	-6	-2.98	26.31

The RTA column includes TPP12 (including US) in top panel and excludes US in lower panel.

Table J.15: TPP EHA (unified)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
TPP with USA						
JPN	-0	411	4	415	5.50	0.18
MEX	-30	-340	-10	-380	-15.30	0.28
USA	259	48	64	371	4.20	4.12
KOR	-9	-239	-18	-266	-7.10	-0.10
CAN	36	159	3	197	8.30	4.51
DEU	-3	-117	-10	-130	-2.40	-0.04
GBR	-1	-33	-4	-38	-2.30	-0.05
ITA	-2	-23	-2	-26	-4.40	-0.05
THA	-0	-23	-3	-26	-2.90	-0.09
CHN	-13	-5	-1	-20	-0.10	0.01
VNM	-20	0	0	-20	-14.00	5.43
HUN	-0	-13	-2	-15	-3.10	-0.03
SWE	-0	-10	-1	-11	-5.30	-0.05
ZAF	-0	-8	-1	-9	-2.70	-0.05
IND	0	-6	-1	-7	-0.40	-0.01
CP-TPP (without USA)						
CAN	83	20	673	775	32.80	3.54
USA	-461	-113	-12	-586	-6.60	0.63
MEX	-6	-16	-50	-72	-2.90	0.20
KOR	-1	-32	-34	-67	-1.80	-0.04
JPN	1	144	-91	54	0.70	0.02
DEU	0	-17	-14	-31	-0.60	-0.02
VNM	-21	0	0	-21	-14.40	3.17
THA	-0	-14	-2	-16	-1.80	-0.06
GBR	-0	-6	-6	-11	-0.70	-0.02
HUN	-0	-3	-2	-5	-1.00	-0.01
IND	0	-4	0	-4	-0.20	-0.01
ITA	-0	-1	-3	-4	-0.60	-0.01
FRA	0	-2	-1	-3	-0.20	-0.01
SWE	-0	-1	-2	-2	-1.20	-0.02
CHN	4	-1	-0	2	0.00	0.00

The RTA column includes TPP12 (including US) in top panel and excludes US in lower panel.

Table J.16: TPP EHA (segmented)

Country	Change in shipments (ths. cars) to:				% Chg.	CS % Chg.
	<i>Domestic</i>	<i>RTA</i>	<i>ROW</i>	<i>Total</i>		
TPP with USA						
USA	306	53	69	429	5.00	3.05
MEX	-31	-359	-11	-401	-16.20	0.27
JPN	2	375	6	383	5.20	0.16
KOR	-12	-240	-22	-274	-7.40	-0.11
CAN	35	153	3	191	8.30	3.33
DEU	-3	-115	-12	-130	-2.40	-0.05
GBR	-1	-29	-3	-33	-2.00	-0.05
THA	-0	-23	-3	-26	-3.00	-0.10
CHN	-15	-5	-1	-22	-0.10	0.01
VNM	-21	0	0	-21	-15.00	5.27
ITA	-1	-16	-1	-18	-3.10	-0.04
HUN	-0	-13	-2	-15	-3.30	-0.03
ZAF	-0	-9	-1	-10	-2.90	-0.06
SWE	-0	-8	-1	-9	-4.50	-0.03
IND	0	-7	-1	-8	-0.40	-0.01
CP-TPP (without USA)						
CAN	82	20	711	814	35.50	2.37
USA	-471	-109	-12	-593	-6.90	0.66
MEX	-7	-18	-67	-92	-3.70	0.20
KOR	-2	-32	-48	-81	-2.20	-0.05
JPN	1	145	-78	68	0.90	0.03
DEU	-0	-16	-21	-37	-0.70	-0.02
VNM	-22	0	0	-22	-15.30	3.32
THA	-0	-14	-2	-16	-1.80	-0.06
GBR	-0	-5	-6	-12	-0.70	-0.02
HUN	-0	-3	-3	-5	-1.20	-0.01
IND	0	-5	0	-5	-0.30	-0.01
ITA	-0	-1	-4	-4	-0.70	-0.01
FRA	0	-2	-1	-3	-0.20	-0.01
ZAF	0	-1	-2	-3	-0.80	-0.03
SWE	-0	-1	-2	-3	-1.30	-0.01

The RTA column includes TPP12 (including US) in top panel and excludes US in lower panel.