Brands in motion: How frictions shape multinational production*

Keith Head† Thierry Mayer‡

November 5, 2016

Abstract

Using disaggregated data on car assembly and trade, we estimate a model of multinational production. Decisions of which markets to enter, how much to sell in each, and which assembly locations to select for each market depend on three types of friction. In addition to the trade and multinational production costs emphasized in past work, we incorporate a third friction: regardless of production origin, selling costs in a market rise with separation from the brand's headquarters. The estimation transparently recovers all the structural parameters. We then simulate the consequences of controversial trade policy changes: TPP, TTIP, Brexit, and NAFTA abrogation.

*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, EITI, ERWIT, Harvard-MIT trade seminar, HKU, HKUST, Kiel, Princeton, Society of Economic Dynamics, UBC, University of Toronto, and Warwick provided valuable 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. Jules Hugot, David Rodrigues, Eve Sihra, Xiaonon 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

In 2016 the fate of an unprecedented set of far-reaching integration agreements awaited political outcomes. Proposed transpacific and transatlantic agreements (covering 40% and 45% of the world economy) faced vigorous opposition, while unilateral withdrawals from existing agreements had been approved by British voters and threatened by a US presidential candidate. Pure trade models are ill-equipped for predicting the outcomes of such proposals 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 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 also face decisions over which subset of varieties to offer in each of the markets where they 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 and, most vitally, data on production and bilateral 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. We use the estimated model to predict the consequences for producers and consumers of the shocks to trade policy that are currently being debated. To avoid opaque estimation and black-box counterfactuals, we employ strong functional forms that make estimation tractable and modularize the decision making of the MNC. This involves some compromises in terms of the amount of "industry realism" that can be accommodated. We argue quantitatively that the functional forms we use can still produce reliable predictions for aggregates even if they are just approximations for a more complex underlying model.

The unified quantitative framework utilized here could be termed the double CES (constant elasticity of substitution) model of multinational production. 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) via Tintelnot (2016). Important contributors to the development of the double CES framework for MP include Ramondo (2014), Ramondo and Rodríguez-Clare (2013), Irarrazabal et al. (2013), and Arkolakis et al.

(2013). While the specifics vary across these papers, comparative statics in these models 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. This is the first MP model implementation to estimate two extensive margins (the sourcing of each variety for each market and variety-market entry) and two intensive margins (variety- and firm-level sales on each origin-destination path).

The first contribution of this paper is to use the extensive and intensive margin equations to estimate the friction parameters relevant to the MP model based on micro data. In contrast to Arkolakis et al. (2013), Tintelnot (2016), and Coşar et al. (2016) we have nearly exhaustive firm-variety-source-market level flows. This rich data set permits estimation in which all cost parameters are identified transparently through discrete choice and gravity-style flow regressions implied by the model. The sourcing and intensive margin equations deliver credible estimates of the two pivotal elasticities of the double CES framework. A second contribution of the paper is to extend the basic MP model to include a new friction, between HQ and market, and a new decision, which varieties to offer where. By way of contrast, Tintelnot (2016) assumes a unit mass of varieties for each firm, the entirety of which are offered in every market. From a policy perspective, the most important contribution comes from using the fully estimated model to predict how changes in regional integration would reshape the activities of multinational corporations. This paper offers the first quantitative assessment of proposals to integrate and dis-integrate regional economies (TPP, TTIP, Brexit, NAFTA) that takes into account the microeconomic structure of multinational production.

We utilize data provided by an automotive industry consultant which tracks production at the level of *brands* (Acura, BMW, Chevrolet) and *models* (RDX, X5, Corvette). We view brands as the appropriate counterpart of firms in the theory, as they have more continuity over time and similarity in product offerings than the parent corporations (for example, Tata's $1,600 Nano model has very little to do with the Jaguar-brand cars that came under Tata ownership in 2008). Car models correspond to the natural understanding of varieties in monopolistic competition. We organize the estimating framework around the brand-level decisions over which countries to offer each model and which countries to source assembly from for each model-market pair.

The "brands in motion" in the paper title refers to two types of metaphorical movement. The first is the transfer of brand-specific inputs from the headquarters to plants in other countries. This friction is already emphasized in the previously cited literature on multinational production. The second sense of mobility is one that has not yet figured explicitly in prior work: To what extent can a brand transfer its success in the home market into foreign markets? Since the impediments to moving technology to the assembly location are called multinational production (MP) frictions, we term the impediments to moving market success abroad multinational sales (MS) frictions. MS frictions are the cost disadvantage incurred when the market is distinct and distant from the headquarters country—regardless of the location of production. A key motivation for incorporating MP and MS 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.

The MP and MS frictions combine with the familiar trade frictions associated with separation between production and consumption locations to shape firms' decisions between exporting from home and producing abroad to serve host, home, and third markets. To distinguish those new frictions from traditional trade costs, we show that one needs data tracking the three countries where a brand is headquartered, produces and sells its products. The idea is therefore to use the simplest modeling structure that permits transparent identification of these new frictions without committing to sets of assumptions that are context-specific. The double CES framework yields such a structure and can be seen as an extension of the gravity equation to a setup where coordination of foreign assembly and distribution affiliates by headquarters is costly. Gravity has proven to be a powerful tool for understanding international trade flows; its most attractive features being tractability, straight-forward estimation, and good fit to the data. The gravity equation—extended to incorporate MP—again performs strongly in our application to the car industry.

Because our paper utilizes 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. (2016) combines a demand side from Berry et al. (1995) with the MP model of Tintelnot (2016). These papers feature 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 richer and more realistic substitution patterns than the monopolistic competition with symmetric varieties demand assumed in the double CES model.

Although nested and random-coefficient logit demand can capture the compelling idea that some car models substitute more readily for each other than they do for models with very different attributes, such richness comes with two major costs. 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 attributes of all the models. Such data are only available for a drastically reduced set of brands, models, and markets.[1] This would make it impossible for us to consider the global production reallocations associated with the mega-regional agreements.

The chief advantage of the CES assumption is that it leads to a linear-in-parameters specification where implementation is straightforward and identification is transparent. Since the counterfactuals are not embedded in a setting containing many industry-specific assumptions, our results illustrate general features of the MP model that would be expected to apply in other industries that share the same broad features. We are not complacent about the strong restrictions imposed in the CES monopolistic competition set-up. However, we take some comfort from the fact that the own-price elasticities we estimate and the implied markups lie well within the range of estimates obtained using richer demand structures. This offers reassurance that the symmetry assumption of CES does not do too much violence to the central moments of the data.[2] Appendix B rigorously validates the usefulness of the CES model for public policy prediction even under market structures considered more realistic for the car industry: CES counterfactuals closely approximate data generated using Berry et al. (1995) assumptions in Monte Carlo simulations.

The results obtained in this paper offer insights to the design of models of the allocation of multinational production across countries. Both trade costs and multinational sales frictions are strongly significant in all specifications. Tariffs on final cars have major effects on the sourcing decision with an elasticity of -8.05 and the allocation of market shares (elasticity of -3.77). With regard to the former, we find the double CES framework performs well when applied to the global car industry data. The core parameters we obtain are internally consistent across the different estimating equations. They also make sense when compared to estimates from independent sources. In terms of a simple measure of fit, the flows predicted by the model match the data with a correlation of 70%. The new features that we incorporate into the MP framework—the variety-market entry margin and the multinational sales friction—prove to be quantitatively important.

[1] The Coşar et al. (2016) data set has 9 markets and 60 brands compared to the 73 markets and 145 brands in our estimating sample.

[2] Adao et al. (2015) estimate a "mixed CES" model of bilateral trade and demonstrate that, while the cross-price elasticities are affected, the central elasticity of trade with respect to trade costs hardly changes when the authors incorporate coefficient heterogeneity to allow the model to exhibit richer patterns of substitution.

We think it is probably true more broadly that firms do not export all their varieties to every market. Including this margin does not over-complicate estimation or simulation. Combining the model-level intensive margin with the entry margin, we estimate sizeable and fairly robust home, distance, and RTA effects associated with headquarter-market separation. The ad valorem equivalents of the combined variable and fixed components of MS frictions are several times larger than corresponding trade costs and MP frictions.

The results from counterfactual trade policy changes improve our understanding of the impacts of regional integration agreements. As one would expect in a pure trade model, liberalization has negative third-country effects via the path of erosion of trade preferences. For example, the United States' NAFTA partners lose production when the US integrates more closely with the EU in TTIP. A qualitatively different third-country effect comes from reduction in MP frictions associated with RTAs. They raise the competitiveness of multinational subsidiaries in the new integration area, boosting exports to the rest of the world. Another distinctive feature of our results as compared to a pure trade model is that "deep integration" sometimes magnifies and sometimes reverses the effects of tariffs-only agreements. As an example of the former, the deep aspects of TTIP multiply US consumer gains by a factor of ten. As an example of the latter, Canadian production declines under a tariffs-only version of the TPP whereas a deep integration version of the agreement is predicted to raise production.

The paper continues in four main sections. We first discuss and display some of the important empirical features of multinational production and trade in our dataset featuring 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 multinational sales frictions and a model-market entry decision. We then show how the structural parameters of the MP model can be recovered from estimating equations that capitalize on the disaggregated nature of our data. Finally, we evaluate the effects of two proposed "mega-regional" integration agreements (TPP and TTIP), as well as two regional "dis-integration" policies (the UK exit from the EU, and a US exit from NAFTA) using a counter-factual solution of the model.

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 the IHS Automotive consultancy (formerly Polk).[3] From it we extract origin-destination flows for 2361 car models sold by 145 brands.[4] What we refer to as a "model", IHS calls the "sales nameplate" and defines as the "Name under which the vehicle is sold in the respective country."[5] 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 2000 to 2013.

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 the 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 multinational sales 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. For instance, 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, respectively, as the brand headquarters might therefore incorrectly specify the relevant frictions. A second

[3]Other 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.

[4]Appendix C lists and explains the deletions we applied to the original IHS dataset.

[5]Examples of models with the brand shown in parentheses are the 500 (Fiat), Twingo (Renault), 3 (Mazda).

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), the "program" (a distinction between minor redesigns), and the body type (hatchbacks, sedans, etc.) The concept of firm is the "Design Parent", the corporation that has managerial control over the design of the platform used by each variety. The results from implementing this approach, shown in Appendix D, are not systematically weaker or stronger, and in some cases are strikingly similar. We make comparisons 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 locations. Our estimating sample comprises the shipments of cars assembled in 49 countries by brands headquartered in 22 countries and sold in 73 national markets.

Figure 1 illustrates the different types of production actually done by a large brand, Fiat in 2013, for two of its main models and seven markets. Fiat sells the Punto to domestic and EU consumers from its home plant in Italy. Italian imports of the Fiat 500 from its Polish plant is an example of vertical MP. Horizontal MP occurs in each assembly location: sales in Mexico of the Mexican-made 500, and the local sales of the 500 from the Polish plant, as well as the Brazilian sales of the Punto assembled there. There are also many examples of export platform flows, which are mainly organized along regional lines, a feature that our regressions will reveal is of key importance. 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 below. The fact that the US does not import the Punto from any source provides an example of selective model-market entry. We show below that this phenomenon is more the norm than the exception.

Figure 2 displays in panel (a) the relative shares of different forms of multinational production. We see that in 2000 home production was still prevalent, accounting for about two thirds of total production. By 2013, foreign production—mainly oriented towards consumers in the same country as the overseas assembly plant—had taken the lead.

Figure 1: Example: Fiat 500 & Punto production organization

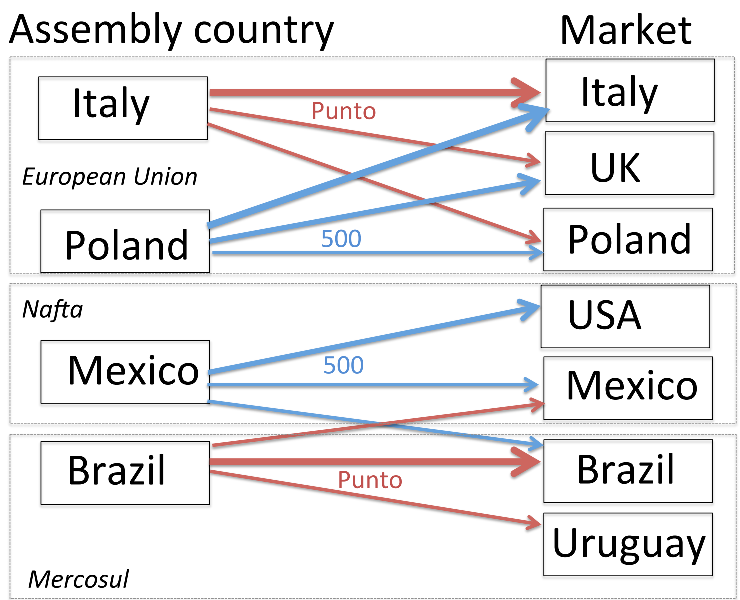

Restricting attention to the traditional major markets for cars in panel (b) of Figure 2, the picture is substantially altered in one respect: most of the rise in MP for local sales disappears. This change reflects the massive importance China has assumed as host of MP. In OECD markets, export platform MP is much larger than MP for the host market, and the gap (between dashed and solid blue lines) increases over time. This underscores the empirical relevance of incorporating export platform MP as in Tintelnot (2016). Furthermore, the share of export platform occurring inside RTA borders is extremely high over the whole period (never going below 75%, and even higher for OECD markets) confirming the anecdotal evidence of the Fiat example presented in figure 1.

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

Figure 2: Some key ratios of multinational production types

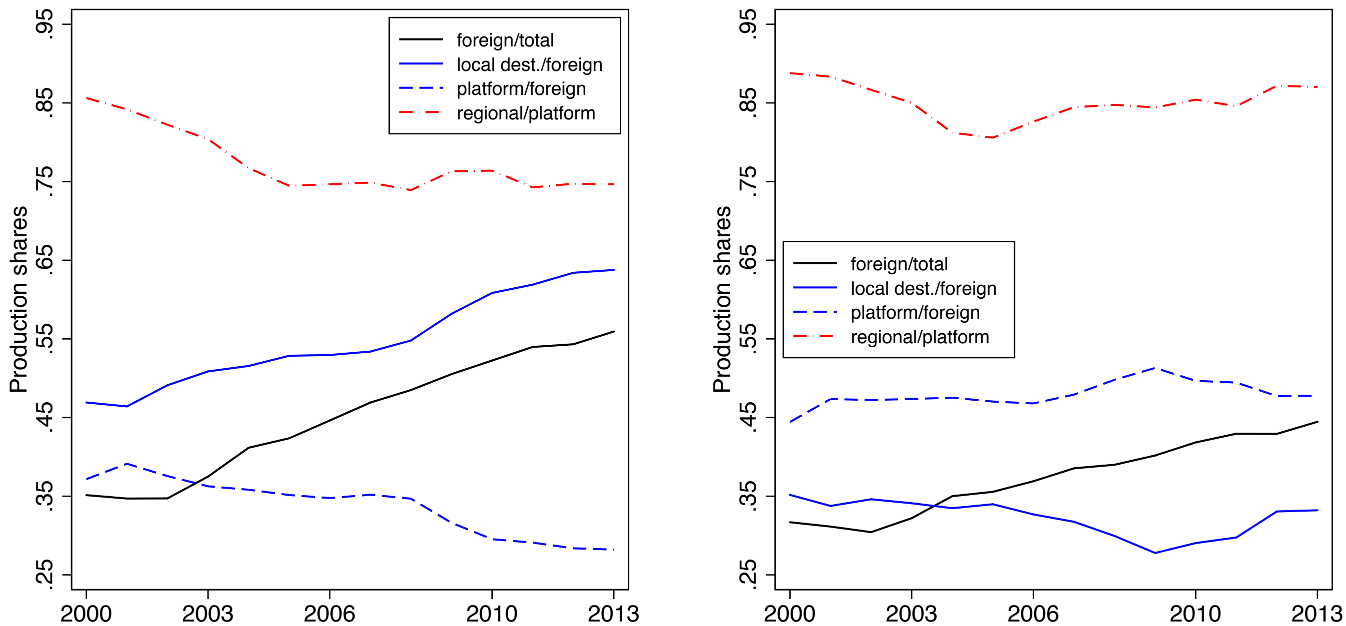

(a) all markets (b) OECD markets only

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. About a quarter of all models are produced in more than one country and we observe six 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

	All model-markets			Brands with 10+ locations		
# Sources	Count	Col %	Cum %	Count	Col %	Cum %
1	196,741	95.3	**95.3**	114,624	94.0	94.0
2	8,383	4.1	99.4	6,216	5.1	99.1
3	1,177	0.6	100.0	995	0.8	100.0
4	52	0.0	100.0	47	0.0	100.0
5	6	0.0	100.0	6	0.0	100.0

Table 1 shows that 95% of the model-market-year observations feature sourcing from a single assembly country. Sourcing from up to five 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 94% of the cases, these models are still single-sourced. The prevalence of single sourcing is even more marked when we define varieties as platform-program-bodytypes, as in Appendix D. In that case we find 98% of the variety-market-years are supplied from a single source country. This suggests some of the dual sourcing in Table 1 may be attributable to changing assembly locations when new generations of a model ("programs") are introduced.

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 section.

Two lines of argument support the usefulness of monopolistic competition in this context. First, in several important respects, the industry is less concentrated than might be imagined. Second, we show that even under data generating processes that approximate the level of concentration observed in the industry, an estimated CES monopolistic competition model can deliver accurate predictions for trade policy counterfactuals.

Table 2 shows in the first two columns that most markets feature many competitors and market shares are typically small. Consequently, with symmetric differentiation between all firms, oligopoly markups for most firms would be close to those implied by monopolistic competition.

The IQR of the number of models offered is quite high, with three quarters of markets having more than 164 models. The top model market share never exceeds 13%, while the CR5 has a maximum of 40%. With aggregation up to the brand or firm level, as shown in the two following rows, concentration measures naturally rise. Even at the highest level of ownership (the design parent), US merger guidelines would classify close to half the market-years as unconcentrated and just 15% as highly concentrated. Three quarters of

Table 2: Market share concentration in car sales, 2000–2013

		Inter-Quartile-Range market shares			Concentration % market-years		
Level	count	median	CR5	top	low	mod.	high
model	164–284	.05–.13	21–40	6–13	97	3	0
brand	33–49	.33–.97	49–75	14–32	70	20	10
firm	17–22	1.1–2.78	70–83	21–37	48	37	15

All figures are calculated over all market-year combinations (73 countries, 2000 to 2013). CR5 is the combined share of the top 5. Concentration shows the share of market-years that US merger guidelines classify as unconcentrated (H<1500) , moderately (1500≤ H≤ 2500) or highly concentrated (H>2500).

the market shares attained by the largest firm are less than 37%.

To be clear, we are not arguing that oligopoly is irrelevant in the industry. The largest firms are big enough to have endogenous markups even under CES. In Appendix B we simulate data from a BLP framework featuring oligopoly, rich substitution in demand, and multiproduct firms that internalize cannibalization effects. We find the CES model tracks BLP-generated data well under settings that replicate data moments of Table 2 (5-firm concentration ratios of 70–80% and an average of 10 models per firm[6]). 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 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 models are offered in a minority of markets

In the MP model presented in the next section, the firm decides which of its varieties to offer in each of the markets where it has distribution facilities. The extensive margin was already incorporated into a model of multiple-product, multiple destination firms by Bernard et al. (2011). It has not been incorporated in prior work on multinational production since Arkolakis et al. (2013) assume single-product firms and Tintelnot (2016) assumes a unit mass of varieties that the firm offers in every market. Here we show that the model-level entry margin is very important for multi-model brands in the car industry. Figure 3 depicts the histogram of $\bar{\mathbb{I}}_{mn}$, the model-level mean of the binary variable

[6]In the average market, the average brand offers 5.4 models and the average firm offers 13.5 varieties.

Figure 3: Market coverage by multi-model brands

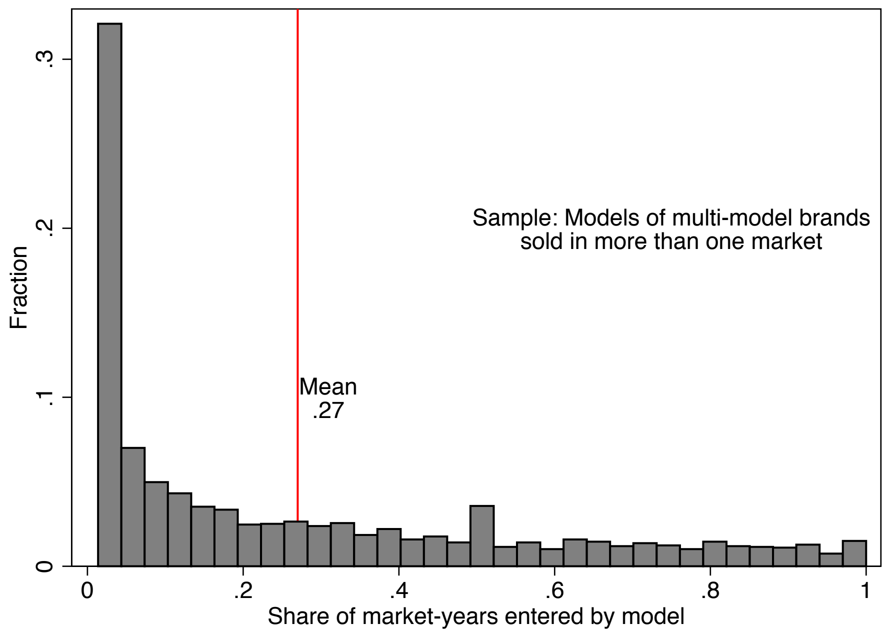

$\mathbb{I}_{mnt}$ indicating model m is offered in market n in year t. The sample comprises model-market-years where the brand is available, the model is offered in more than one market, and the brand makes more than one model. We observe that brands almost never serve a market with *all* their models and only 21% of models are available in the majority of the markets where the brand is available. With the average entry rate being just 27%, 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 1.6% of the market-years. Using the firm-variety methodology described in Appendix D we see that the hatchback version of C1-J68C has an 80% entry rate. However, looking across all varieties the average entry rate is just 24%, slightly lower than the average across all models. The reason the average declines seems to be that firms have operations in more countries than brands, which increases the set of places where entry does not occur. Also by distinguishing programs and body types, this approach actually has more total varieties. 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.

3 The double CES model of multinational production

The structure underlying our estimating equations is one in firms make a binary decision over whether to offer a particular variety in a given market. Then they make a multinomial decision over the assembly location and finally they set the quantity of cars to be delivered from each assembly country to each market. A firm would ideally site all assembly in the country offering the lowest input costs. However, it also wants to produce close to consumers (to avoid trade costs) and close to headquarters (to avoid MP frictions). The geographic distribution of buyers depends on the aggregate demand for cars in each country. The market shares obtained by each car brand depend on the fixed and variable costs of offering models in destinations separated from headquarters (MS frictions).

Figure 4: Frictions impeding multinational flows

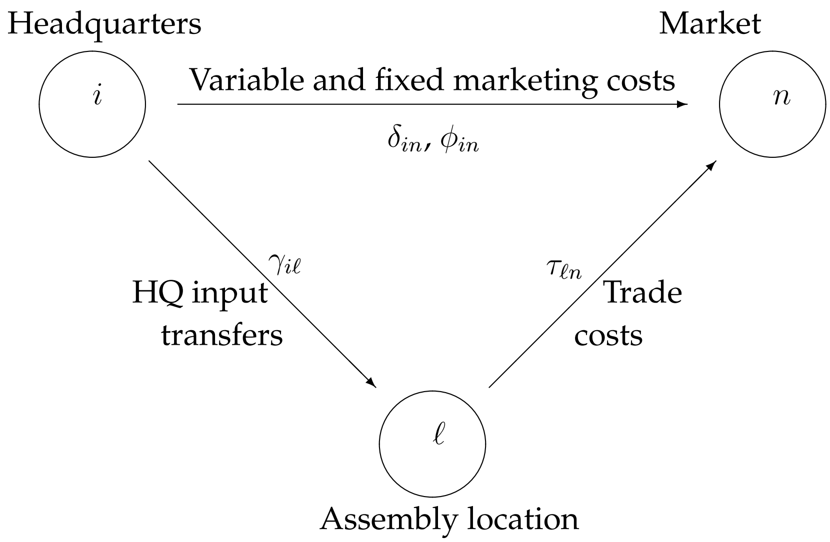

Figure 4, adapted with one major change from Arkolakis et al. (2013), depicts the three frictions schematically. The first friction, conventionally denoted $\tau_{\ell n}$, is the multiplicative increase in costs associated with shipping goods from assembly location ℓ to destination market n. A second friction, denoted $\gamma_{i\ell}$ following Arkolakis et al. (2013), is the increase in production costs incurred when a car is assembled remotely from the headquarters. The novel MS frictions we introduce in Figure 4 are δ_{in} and ϕ_{in}. The former is a rise in delivered marginal costs due to separation between market and headquarters—regardless of production location. The latter is a fixed cost of making each additional variety available in market n, given that headquarter services such as design, reputation, and dealership management are coming from headquarters.

The next subsections consider (in reverse order) three decisions to be made for each model m: whether to offer it in a given market n, which assembly location ℓ to source it from, and the amount to ship from each source to each market. In making each model-level decision, we take the brand's presence as an assembler or distributor in the country as given.

3.1 Consumer preferences and demand

In our data we observe only quantities, not expenditures, and therefore need 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. We parameterize ψ_{mh} in terms of a common reputation and a household-level idiosyncratic shock: $\psi_{mh} = \beta_m \exp(\epsilon_{mh})$, where ϵ is Gumbel with scale parameter $1/\eta$. The probability household h chooses model m from the set $\mathcal{M}_n$ of models available in n is given by

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

To facilitate aggregation, we set $\beta_m = \beta_b$ for all models of a given brand.[7] Quantity demanded for model m in market n is therefore given by

$$q_{mn} = \mathbb{P}_{mn} Q_n = \frac{\beta_b^\eta p_{mn}^{-\eta}}{\Phi_n} Q_n.$$

η is the first of the two CES that drive outcomes in this framework. It plays the same role as σ in the Dixit-Stiglitz framework. The key difference is that demand is expressed in terms of *quantity* shares, $\mathbb{P}_{mn}$, and aggregate quantities (Q_n), rather than value shares and aggregate expenditures.

[7] This assumption is relaxed in one regression specification.

3.2 Quantities conditional on sourcing location

Equilibrium price p_{mn} depends on the delivered unit cost of model m from assembly country ℓ to market n. Assembly costs per unit depend on four inputs, labor costs in ℓ, the costs of intangible inputs provided by the headquarters (such as managerial oversight), the costs of bundles of intermediates from the host and headquarters (HQ) countries. The production function is Cobb-Douglas with cost shares α_1, α_2, α_3, and α_4. Inputs from the HQ country i are subject to trade costs $\tau_{i\ell}^H$ for HQ services (e.g. managerial control and knowledge transfer) and $\tau_{i\ell}^I$ for tangible intermediate inputs (car parts). Costs also depend inversely on a TFP term combining a Melitzian brand-level shifter and a term representing the idiosyncratic match between model m and country ℓ. Letting W denote wages, S a skill adjustment, P a price index for car parts, φ_b brand b's productivity, and $\zeta_{m\ell}$ the idiosyncratic productivity term, assembly costs are given by

$$C_{m\ell} = \frac{(W_\ell/S_\ell)^{\alpha_1}[(W_i/S_i)\tau_{i\ell}^H]^{\alpha_2} P_\ell^{\alpha_3}(P_i\tau_{i\ell}^I)^{\alpha_4}}{\varphi_b \exp(\zeta_{m\ell})}.$$

Including intermediates is important since they account for three quarters of the value of motor vehicles shipments (source: STAN for 2007). Thus $\alpha_3 + \alpha_4 \approx 0.75$.

The notation can be made more compact and brought closer to what our empirics can identify by defining a single cost index w that combines cost determinants from each location, specifically $w_\ell = (W_\ell/S_\ell)^{\alpha_1} P_\ell^{\alpha_3}$ and $w_i = (W_i/S_i)^{\alpha_2} P_i^{\alpha_4}$. The trade costs applicable to intangible inputs and parts from the headquarters country are combined to form $\gamma_{i\ell} = (\tau_{i\ell}^H)^{\alpha_2}(\tau_{i\ell}^I)^{\alpha_4}$. The $\gamma_{i\ell}$ performs the same role as MP costs, also denoted γ, in Arkolakis et al. (2013). The difference is merely one of interpretation, with our γ reflecting input costs from headquarters and theirs being a friction (expressed as a penalty in terms of lost productivity) associated with transfer of operational methods from HQ to assembly country.[8] Combining these simplifications, assembly costs simplify to

$$C_{m\ell} = \frac{w_\ell w_i \gamma_{i\ell}}{\varphi_b \exp(\zeta_{m\ell})}.$$

Delivery of assembled goods is itself subject to a pair of frictions based on transportation and tariffs on final cars ($\tau_{\ell n}$) and variable marketing costs (δ_{in}). The δ_{in} capture 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 in the case of customs unions) facilitating visits from head office managers. Increases in variable costs brought

[8]Javorcik and Poelhekke (2016) provide support for the hypothesis that foreign affiliates are more productive due to continuous injections of HQ services.

about by foreign regulatory requirements would also be reflected in δ_{in}.[9]

Delivered costs are therefore given by

$$c_{m\ell n} = C_{m\ell}\tau_{\ell n}\delta_{in} = \frac{w_\ell w_i \gamma_{i\ell}\tau_{\ell n}\delta_{in}}{\varphi_b \exp(\zeta_{m\ell})},$$

The delivered price of model m in n is related to marginal costs via the constant markup of CES monopolistic competition, $\frac{\eta}{\eta-1}$. Substituting price into the demand curve, the equilibrium quantity of model m made in ℓ, delivered to n is

$$q_{m\ell n} = \begin{cases} \beta_b^\eta \left(\frac{\eta}{\eta-1}\frac{w_\ell w_i \tau_{\ell n}\delta_{in}\gamma_{i\ell}}{\varphi_b}\right)^{-\eta} Q_n \Phi_n^{-1}\exp(\eta\zeta_{m\ell}) & \text{if } \ell = \ell^*_{mn} \\ 0 & \text{otherwise} \end{cases}$$

where ℓ^*_{mn} is the optimal location, for which $c_{m\ell n}$ is minimized.

As our empirical implementation of the MP models considers flows $q_{m\ell n}$ as a function of friction determinants, it does 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 flows 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 sales 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 flows 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 three more parameters that could not be identified separately from the existing three 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. Instead, our priority is to use the near-exhaustive coverage of markets and models found in the quantity data. We leave to other work the decomposition of frictions into cost and preferences. In that vein, Coşar et al. (2016) restrict the number of markets they study so that they can use price data and 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.

[9] For example, foreign car makers complained about the additional costs of daytime running lamps when Canada mandated them for new cars in 1990.

Expected q depends upon the expected $\exp(\eta\zeta_{m\ell})$. Assuming the $\zeta_{m\ell}$ are distributed Gumbel with scale parameter $1/\theta$, Hanemann (1984) shows that the expected $\exp(\eta\zeta_{m\ell})$, conditional on ℓ being the lowest cost location for m is

$$\mathbb{E}[e^{\eta\zeta_{m\ell}} \mid \ell = \ell^*_{mn}] = \mathbb{P}_{\ell|bn}^{-\frac{\eta}{\theta}} \Gamma\left(1 - \frac{\eta}{\theta}\right),$$

with $\mathbb{P}_{\ell|bn}$ the probability of selecting origin ℓ as source of model m for brand b, and $\Gamma()$ the Gamma function. Therefore expected sales are multiplicative in the determinants of market, origin, brand, frictions and the probability of choosing ℓ.

$$\mathbb{E}[q_{m\ell n} \mid \ell = \ell^*_{mn}] = \kappa_1 \frac{Q_n}{\Phi_n} \left(\frac{w_\ell w_i \tau_{\ell n} \gamma_{i\ell} \delta_{in}}{\beta_b \varphi_b}\right)^{-\eta} \mathbb{P}_{\ell|bn}^{-\frac{\eta}{\theta}}, \tag{1}$$

where $\kappa_1 \equiv \left(\frac{\eta}{\eta-1}\right)^{-\eta} \Gamma\left(1 - \frac{\eta}{\theta}\right)$. As these flows depend on the optimal location for a model-market combination, we now turn to that choice.

3.3 Sourcing decision

Brands choose the optimal source for each model they intend to sell in a market from the set of countries where the brand has assembly facilities, denoted $\mathcal{L}_b$. The probability that $\ell \in \mathcal{L}_b$ is selected is the probability that $c_{m\ell n}$ is lower than the brand's alternatives:

$$\text{Prob}(\ell = \ell^*_{mn}) = \text{Prob}(c_{m\ell n} \leq c_{mkn}, \forall k \in \mathcal{L}_b)$$

$$= \text{Prob}(\zeta_{m\ell} - \ln w_\ell - \ln \gamma_{i\ell} - \ln \tau_{\ell n} > \zeta_{mk} - \ln w_k - \ln \gamma_{ik} - \ln \tau_{kn})$$

The MS 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 as the source of model m in market n is the same for all models of a given brand.

$$\mathbb{P}_{\ell|bn} = \frac{(w_\ell \gamma_{i\ell} \tau_{\ell n})^{-\theta}}{D_{bn}}, \text{ with } D_{bn} \equiv \sum_{k \in \mathcal{L}_b} (w_k \gamma_{ik} \tau_{kn})^{-\theta}. \tag{2}$$

θ is the second CES in this framework, playing the same role as in Eaton and Kortum (2002). Versions of this equation appear in Arkolakis et al. (2013) as equation (6) and Tintelnot (2016) as equation (9), who use it as a building block in their models.[10] In contrast, we estimate the equation directly. So far as we know, no previous study has been able to

[10]Like Tintelnot (2016), we assume independent productivity shocks whereas the Arkolakis et al. (2013) formulation allows for them to be correlated.

do so, mainly because variety-level sourcing data is so hard to find.

3.4 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 model-market entry costs, F_{mn}. One way to imagine this is that each model receives a draw of the necessary amount of marketing costs that would be required to allow it to compete symmetrically with other models in a given market.

We assume that entry decisions are made prior to learning the realizations of the model-location productivity shocks, $\zeta_{m\ell}$. Therefore, entry decisions are made assuming that optimal assembly locations will be chosen. Expected profit net of entry costs for model m in market n can be expressed in terms of the expected price:

$$\mathbb{E}[\pi_{mn}] = \mathbb{E}[p_{mn}q_{mn}]/\eta - F_{mn} = \mathbb{E}[p_{mn}^{1-\eta}]\beta_b^{\eta}K_n - F_{mn}, \tag{3}$$

where $K_n \equiv Q_n\Phi_n^{-1}/\eta$. The probability that entry, denoted $\mathbb{I}_{mn} = 1$, occurs is the probability that expected profits (net of fixed costs) are positive:

$$\text{Prob}(\mathbb{I}_{mn} = 1) = \text{Prob}(\mathbb{E}[\pi_{mn}] > 0) = \text{Prob}(F_{mn} < \mathbb{E}[p_{mn}^{1-\eta}]\beta_b^{\eta}K_n)$$

Taking logs on both sides of the inequality,

$$\text{Prob}(\mathbb{I}_{mn} = 1) = \text{Prob}(\ln F_{mn} < \ln \mathbb{E}[p_{mn}^{1-\eta}] + \eta \ln \beta_b + \ln K_n).$$

The fixed costs of model entry are distributed logistically with location parameter $J_n\phi_{in}$ and scale parameter $1/\lambda$. Country-characteristics such as size and costs of registering a new business are captured in J_n whereas ϕ_{in} is 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. Costs of redesigning a model to comply with foreign product regulations would enter ϕ_{in}.

The entry probability for a model is given by

$$\text{Prob}(\mathbb{I}_{mn} = 1) = \Lambda[\lambda \ln \mathbb{E}[p_{mn}^{1-\eta}] + \eta\lambda \ln \beta_b + \lambda(\ln K_n - \ln J_n - \ln \phi_{in})].$$

We now need to take into account how the firm forms expectations for prices. Using the moment generating function, we obtain

$$\mathbb{E}[p_{mn}^{1-\eta}] = \kappa_2 \varphi^{\eta-1} \delta^{1-\eta} D_{bn}^{(\eta-1)/\theta},$$

where $\kappa_2 \equiv \left(\frac{\eta}{\eta-1}\right)^{1-\eta} \Gamma\left(1 + \frac{1-\eta}{\theta}\right)$. Hence, after substitution of the components of K_n and of the expected price, the probability of entering is

$$\begin{aligned} \text{Prob}(\mathbb{I}_{mn} = 1) =& \Lambda[\lambda(\ln \kappa_2 - \ln \eta) - \lambda(\eta - 1) \ln \delta_{in} - \lambda \ln \phi_{in} + \overbrace{\frac{\lambda(\eta-1)}{\theta} \ln D_{bn}}^{\text{brand-market}} \\ &+ \underbrace{\lambda(\eta - 1)(\ln \varphi_b - \ln w_i) + \lambda\eta \ln \beta_b}_{\text{brand}} + \underbrace{\lambda(\ln Q_n - \ln \Phi_n - \ln J_n)}_{\text{market}}]. \end{aligned} \quad (4)$$

This entry equation produces the sensible prediction that the likelihood of entering a market increases with its size, quality and efficiency of the brand, and declines with frictions, fixed costs and local competition (Φ_n). The entry decision also depends positively on the denominator term from the sourcing decision (equation 2). The reason is that the expected cost of serving a given market will be lower for a brand if its plants are located in countries that are low cost suppliers to market n, because they have either low assembly costs or low transport costs to the market, since both costs are contained in D_{bn}.

3.5 Aggregation to brand-level quantities

Summing over the set $\mathcal{M}_{bn}$ of models that b sells in n, brand-level flows are denoted $q_{b\ell n}$. The realized flow depends on all the $\zeta_{m\ell}$ shocks that determine the sourcing decisions for each model. It also depends on the set of models that brand b decides to offer in market n. The *expected* sales of brand b to market n, conditional on the set of models offered in each market and ℓ being chosen as the low-cost assembly location, is given by

$$\mathbb{E}[q_{b\ell n}] = \sum_{m \in \mathcal{M}_{bn}} \mathbb{E}[q_{m\ell n} \mid \ell = \ell^*_{mn}] \times \mathbb{P}_{\ell|bn}.$$

Substituting equation (2) into (1) and simplifying, we re-express expected model-level flows from ℓ to n as

$$\mathbb{E}[q_{m\ell n} \mid \ell = \ell^*_{mn}] = \kappa_1 \left(\frac{\varphi_b \beta_b}{w_i \delta_{in}}\right)^{\eta} \frac{D_{bn}^{\eta/\theta}}{\Phi_n}.$$

Multiplying this value by the formula for $\mathbb{P}_{\ell|bn}$ from equation (2) and summing across the M_{bn} models that brand b offers in n, we obtain

$$\mathbb{E}[q_{b\ell n}] = \kappa_1(\gamma_{i\ell}\tau_{\ell n})^{-\theta}\underbrace{w_\ell^{-\theta}}_{\text{origin}}\underbrace{(\beta_b\varphi_b/w_i)^\eta}_{\text{brand}}\underbrace{\frac{Q_n}{\Phi_n}}_{\text{market}}\underbrace{M_{bn}\delta_{in}^{-\eta}D_{bn}^{\frac{\eta}{\theta}-1}}_{\text{brand-market}}. \tag{5}$$

This is the equivalent of equation (10) of Tintelnot (2016), except for the discrete choice CES demand (where aggregate quantity demanded replaces aggregate expenditure), non-unit masses of models, and the presence of MS frictions. The key result is that aggregation of models changes the parameter governing the responses of trade flows to the $\gamma_{i\ell}$ and $\tau_{\ell n}$ frictions. Whereas it was η in the model-level equation (1), it is θ here in the brand-level equation (5). The former captures the homogeneity of tastes of consumers over models, whereas the latter characterizes the homogeneity in productivity across locations that might assemble a given model. An interesting difference arises with responses to δ_{in}, which persist in being governed by the demand-side elasticity, η. This is because the MS friction characterizes the HQ-destination pair of countries, and therefore does not include any determinant related to the cost of where the car is actually produced. We use equation (5) as to obtain another set of estimates of τ and γ, 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.

A final equation from the model sums across production locations for each brand to obtain the aggregate expected brand sales by market:

$$\mathbb{E}[q_{bn}] = \sum_\ell \mathbb{E}[q_{b\ell n}] = \kappa_1 M_{bn}(\beta_b\varphi_b/w_i)^\eta\delta_{in}^{-\eta}\frac{Q_n}{\Phi_n}D_{bn}^{\frac{\eta}{\theta}}. \tag{6}$$

This equation cannot be used to estimate γ or τ since they enter only via D_{bn}. However, the equation is well-suited for estimating δ and will provide our estimates for the counterfactual exercises. The coefficient on $\ln D_{bn}$ is useful in two respects. First we can use it, in combination with an η estimate, to back out λ from the entry equation. Second, we can compare the estimate of $\widehat{\eta/\theta}$ with the ratios of individual estimates of η and θ to be obtained from tariff variation.

4 Results

We now consider the empirical implementation of the five equations describing firms' behavior in the model. We start by expressing each equation in an estimable way in terms of fixed effects and observed variables with associated coefficients. For each of the equations, we have vectors of observable fictions denoted $\mathbf{X}_{\ell nt}$, $\mathbf{X}_{i\ell t}$ and $\mathbf{X}_{int}$.

$$\tau_{\ell nt} = \exp(\mathbf{X}'_{\ell nt}\boldsymbol{\rho}), \quad \gamma_{i\ell t} = \exp(\mathbf{X}'_{i\ell t}\mathbf{g}), \quad \delta_{int} = \exp(\mathbf{X}'_{int}\mathbf{d}), \quad \phi_{int} = \exp(\mathbf{X}'_{int}\mathbf{f}), \tag{7}$$

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

The $\mathbf{X}$ vectors include the standard explanatory variables used in gravity equations: home, distance, contiguity, 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 home$_{\ell n}$ indicates that the assembly plant is in the same country as where the car is bought, whereas 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. Contiguity and language indicate that the countries share a land border or 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. This functional form ensures that the coefficient provides a direct estimate of θ (or η in the model-level estimating equation). We also include in $\mathbf{X}_{\ell nt}$ an indicator for a "deep" regional trading agreement between ℓ and n in year t. Depth in this dimension can be obtained via inclusion of customs-related procedures, or the inclusion of services in the agreement.

In $\mathbf{X}_{i\ell t}$ we include tariffs on imported inputs (engine parts and other major components only) from the headquarters country. As with tariffs on assembled cars, the input tariffs enter with the functional form $\ln(1 + \text{tariff})$. This implies a structural interpretation of either $-\eta\alpha_4$ or $-\theta\alpha_4$ where α_4 is the headquarters country input share of the total costs of production. We can estimate α_4 by dividing the coefficient on $\ln(1 + \text{parts tariff}_{i\ell t})$ by the coefficient on $\ln(1 + \text{car tariff}_{\ell nt})$ in any of the equations including these covariates. 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} and ϕ_{in}, 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 before, 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 C provides more detail on measurement of the friction determinants, in particular the sources and procedures used for the tariffs and the deep RTA indicators.

4.1 Estimating the sourcing equation

We transform the sourcing equation into its estimable version by substituting equation (7) into (2) and specifying assembly costs as $\ln w_{\ell t} = \upsilon_{0\ell} + \upsilon_1 \ln y_{\ell t} + \upsilon_2 \ln P^y_{\ell t}$, where $y_{\ell t}$ is per capita income, and $P^y_{\ell t}$ is the price level of GDP.[11] Both are expressed as indices, taking values of 1 in 2013, such that $\ln w_{\ell 2013} = \upsilon_{0\ell}$. The probability brand b sources model m from country $\ell \in \mathcal{L}_{bt}$ to serve consumers in n in year t is

$$\mathbb{P}_{\ell|bnt} = \frac{\exp[\mathrm{FE}_\ell - \theta\upsilon_1 \ln y_{\ell t} - \theta\upsilon_2 \ln P^y_{\ell t} - \theta \mathbf{X}'_{\ell nt}\boldsymbol{\rho} - \theta \mathbf{X}'_{i\ell t}\mathbf{g}]}{\sum_{k\in\mathcal{L}_{bt}} \exp[\mathrm{FE}_k + \upsilon_1 \ln y_{kt} + \upsilon_2 \ln P^y_{kt} - \theta \mathbf{X}'_{knt}\boldsymbol{\rho} - \theta \mathbf{X}'_{ikt}\mathbf{g}]}. \tag{8}$$

The assembly-country fixed effects are structurally interpreted as $\mathrm{FE}_\ell = -\theta\upsilon_{0\ell}$.

The model implies that we should estimate a conditional logit where each brand-destination combination is faced with as many choices as the number of countries in which it has plants, the set denoted $\mathcal{L}_{bt}$. This approach differs from Coşar et al. (2016) 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. (2016) 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).

Column (1) of Table 3 reports our sourcing results. The estimates reveal the impor-

[11] The sign of υ_1 is ambiguous since y reflects productivity (cost-lowering) and wages (cost-raising). On the other hand, υ_2 should be positive since P^y captures exchange rate over-valuation.

Table 3: Baseline results

Dep. Var:	ℓ^*_{mnt}	$\ln q_{m\ell n}$		$\frac{q_{b\ell n}}{M_{bn}Q_n}$	$\frac{q_{bn}}{M_{bn}Q_n}$	$\ln\left(\frac{q_{bn}}{M_{bn}Q_n}\right)$	$\mathbb{I}_{mnt}$
Method:	cond. logit	OLS	OLS	Poisson	Poisson	OLS	logit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Trade costs							
home$_{\ell n}$	1.035^a	0.774^a	0.962^a	1.523^a			
	(0.395)	(0.242)	(0.250)	(0.335)			
ln dist$_{\ell n}$	-0.294^b	-0.179^b	-0.151^c	-0.699^a			
	(0.114)	(0.077)	(0.084)	(0.105)			
contig$_{\ell n}$	0.139	0.117	0.170^c	0.198^c			
	(0.135)	(0.091)	(0.096)	(0.114)			
language$_{\ell n}$	-0.096	0.144^c	0.094	-0.045			
	(0.103)	(0.080)	(0.070)	(0.137)			
ln (1+ car tariff$_{\ell n}$)	-8.052^a	-3.501^a	-4.034^a	-10.698^a			
	(1.697)	(0.974)	(0.998)	(0.823)			
Deep RTA$_{\ell n}$	0.254^c	0.140	0.248^c	0.527^a			
	(0.154)	(0.129)	(0.127)	(0.145)			
MP frictions							
home$_{i\ell}$	2.356^b	-0.047	0.511	2.394^a			
	(1.079)	(0.466)	(0.333)	(0.694)			
ln dist$_{i\ell}$	0.170	0.160	0.150	0.218			
	(0.312)	(0.132)	(0.110)	(0.223)			
contig$_{i\ell}$	0.004	-0.087	0.353	-0.182			
	(0.433)	(0.299)	(0.236)	(0.373)			
language$_{i\ell}$	-0.099	-0.019	-0.698^a	0.226			
	(0.462)	(0.294)	(0.241)	(0.362)			
ln (1+ parts tariff$_{i\ell}$)	-2.039	-0.516	-0.106	-4.014^a			
	(1.605)	(1.170)	(1.011)	(1.412)			
Deep RTA$_{i\ell}$	0.421	-0.002	0.199	0.554^c			
	(0.484)	(0.221)	(0.133)	(0.308)			
MS frictions							
home$_{in}$		0.843^a	0.865^a		0.720^a	0.984^a	0.657^a
		(0.301)	(0.283)		(0.250)	(0.340)	(0.131)
home$_{in}$ $\times$ LDC$_n$		0.522	0.718		0.220	-0.474	0.937^a
		(0.834)	(0.804)		(0.235)	(0.844)	(0.280)
ln dist$_{in}$		-0.121	-0.178		-0.381^a	-0.115	-0.110^a
		(0.096)	(0.108)		(0.104)	(0.118)	(0.041)
contig$_{in}$		0.132	0.107		0.001	0.341^b	0.205^a
		(0.146)	(0.147)		(0.112)	(0.149)	(0.074)
language$_{in}$		0.212^b	0.313^a		0.291^b	0.350^a	-0.063
		(0.080)	(0.078)		(0.129)	(0.106)	(0.071)
Deep RTA$_{in}$		0.356^b	0.257^c		-0.074	0.201	0.151^a
		(0.138)	(0.142)		(0.131)	(0.185)	(0.047)
$\ln \hat{\mathbb{P}}_{b\ell n}$		0.057	-0.018				
		(0.085)	(0.096)				
$\ln \hat{D}_{bn}$					0.520^a	0.751^a	0.129^b
					(0.079)	(0.100)	(0.050)
Observations	2214350	227995	227995	292987	39759	39759	760841
r2	0.506	0.495	0.647	0.111	0.653	0.662	0.143
S.E. cluster:	ℓ	ℓ, n	ℓ, n	ℓ	b	b, d	b

Standard errors in parentheses. Significance: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$. r^2 is squared correlation of fitted and true dependent variables except in specifications (1) and (7) where pseudo-r^2 is reported. Each regression controls for log per-capita income and price level of the assembly country.

tance 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 almost triple the odds of being chosen. Distance from the market also significantly reduces the probability of being selected.

The coefficient on the log of one plus the car tariff rate estimates $-\theta$ in this equation. The conditional logit estimate is 8.05 and is our preferred estimate for θ. The brand-level sales equation will confirm this high level of substitutability between the alternative plants available to brands. Deep regional trade agreements augment the odds of being chosen by 29%, even after accounting for the tariffs applied by the destination market to the different possible origins of the car.

The estimates of the MP 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 $\exp(2.356) \approx 11$ 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, α_4 in the cost equation, which is about 25% (2.04/8.05). Deep RTAs between assembly and headquarter countries are estimated to have a larger effect on sourcing than deep RTAs between assembly and consumer countries, but the standard error is so large as to prevent any precise inferences from column (1).

4.2 Model-level intensive margin for sales

Having estimated the determinants of the sourcing of model m, we turn to the analysis of the micro-level sales equation (1). Using the matrix notation for frictions, this equation is transformed into the following estimating form:

$$\begin{aligned}\ln q_{\ell mnt} \quad = \quad & \mathrm{FE}_\ell - \eta\upsilon_1 \ln y_{\ell t} - \eta\upsilon_2 \ln P^y_{\ell t} + \mathrm{FE}_b - \eta\upsilon_1 \ln y_{it} - \eta\upsilon_2 \ln P^y_{it} + \mathrm{FE}_{nt} \\ & -\eta\mathbf{X}'_{\ell nt}\boldsymbol{\rho} - \eta\mathbf{X}'_{i\ell t}\mathbf{g} - \eta\mathbf{X}'_{int}\mathbf{d} - (\eta/\theta)\ln\hat{\mathbb{P}}_{\ell|bnt} + \nu_{\ell mnt}. \qquad (9)\end{aligned}$$

The structural terms underlying the origin, brand and destination effects are $\mathrm{FE}_\ell = -\eta\upsilon_{0\ell}$, $\mathrm{FE}_b = \eta(\ln\beta_{b(m)} + \ln\varphi_{b(m)} - \upsilon_{0i})$ and $\mathrm{FE}_{nt} = \ln Q_{nt} - \ln\Phi_{nt}$. An important question relates to the presence of the $\ln\hat{\mathbb{P}}_{\ell|bnt}$ term on the RHS of the regression equation. We observe $\ln q_{m\ell nt}$ only for the locations actually chosen by the brand as the lowest cost sources for model m deliveries to market n. Locations with low w_ℓ, $\tau_{\ell n}$ and/or $\gamma_{i\ell}$ are attractive locations and can be chosen even if $\zeta_{m\ell}$ (the random part of productivity that is specific to that model and plant) is low. Therefore, the error term is negatively correlated with variables that increase attractiveness, potentially leading to biased estimates. Hanemann (1984)'s results suggest a Heckman-like two stage procedure. In the first step, one estimates equation (8), the conditional logit sourcing decision. From the results, we then calculate $\ln \hat{\mathbb{P}}_{\ell|bnt}$ and add it to RHS of the $\ln q_{m\ell nt}$ equation. The error term, $\nu_{m\ell nt}$ includes the ζ productivity shock as well as any errors that arise from mis-measuring frictions or mis-specification.

Column (2) of Table 3 estimates equation (9) while column (3) replaces brand effects (FE_b) with model fixed effects. In keeping with the estimates from the sourcing decision, we find the quantity conditional on being selected to respond strongly and robustly to trade frictions.

The model-level sales equation estimates η based of the coefficient on ln (1+ car tariff$_{\ell n}$). It is 3.50 in column (2) and 4.03 in column (3). We take the average of those two figures as our estimated $\hat{\eta} = 3.77$[12]

This estimate of η is substantially smaller than the ones for θ obtained in the sourcing decision or the brand-level sales. It implies that there is considerably more heterogeneity in consumer evaluations of brands than in car maker evaluations of assembly locations.

Our estimate of $\eta = 3.77$ implies a constant markup of 36% for all models in all markets. Past research has produced a wide range of markup estimates for the auto industry. Three early 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. (2016) find model-market markups ranging from 4.6% (Renault Clio in Brazil) to 12.3% (Clio in France). Our implied markup lies in the upper region of the highly dispersed set of results found in this literature.

As with sourcing, individual determinants of MP frictions are mainly insignificant and often take the incorrect sign. As a group, they are not significant at standard levels with brand fixed effects, and with model fixed effects the significance derives from a perversely signed language effect (Table 4 provides joint significance tests of the three sets of frictions for all our regressions). Trade and MS costs are collectively highly significant in every specification. When included with other categories of friction, the MS frictions obtain the highest joint significance.

Among the determinants of multinational sales (MS) frictions, consumers are more than twice as likely to select a home brand, corroborating the significant home bias found

[12]This lies between two leading estimates of average own price elasticities for car models of 3.28 and 5.00 obtained by Goldberg (1995) and Berry et al. (1995).

by Coşar et al. (2016). Sharing a common language or being members of a deep regional trade agreement (RTA) also reduce the MS frictions between the HQ and the destination country. The latter effect, which raises model sales by 43% (with brand fixed effects) and 29% (with model fixed effects), will be important for our counterfactuals where we experiment with different scenarios of RTA changes.

4.3 Brand-level intensive margin for sales

Brand-level exports are predicted in equation (5). This equation includes M_{bn}, the number of models that a brand chooses to sell in n, on the right-hand side. As it is an endogenous variable that enters with a unitary elasticity, we pass it to the left-hand side and re-express the dependent variable as average market share per model. Adding the time dimension, we obtain:

$$\mathbb{E}\left[\frac{q_{b\ell nt}}{M_{bnt}Q_n}\right] = \exp[-\theta \mathbf{X}'_{\ell nt}\boldsymbol{\rho} - \theta \mathbf{X}'_{i\ell t}\mathbf{g} + \mathrm{FE}_{\ell} + \upsilon_1 \ln y_{\ell t} + \upsilon_2 \ln P^{y}_{\ell t} + \mathrm{FE}_{bnt}]. \tag{10}$$

Our brand-level equation uses brand-market-year (bnt) fixed effects and assembly country (ℓ) fixed effects that capture the cost index of producing each model: The structural parameters underlying the fixed effects are $\mathrm{FE}_{bnt} = \eta(\ln(\varphi_b\beta_b) - \upsilon_{0\ell} - \ln\delta_{int}) + \ln\left[(Q_{nt}/\Phi_{nt})D_{bnt}^{\frac{\eta}{\theta}-1}\right]$ and $\mathrm{FE}_{\ell} = -\theta\upsilon_{0\ell}$. Since brands have only one HQ i, δ_{int} is not identified in this estimation. The primary use of equation (10) is to provide another set of estimates for θ, $\alpha_4\theta$, and the $\gamma_{i\ell t}$ and $\tau_{\ell nt}$ parameters.

Equation (10) is estimated using Poisson PML, primarily because this method keeps all the observations with $q_{b\ell nt} = 0$. Santos Silva and Tenreyro (2006) emphasize that a second advantage of this method is its robustness to deviations from homoskedasticity that lead to bias in linear-in-logs regressions. Column (4) estimates the same parameters as the column (1) specification, but with a key difference that it includes variation in $q_{b\ell nt}$ when flows are positive. As in the multinomial PML advanced by Eaton et al. (2013), the dependent variable is divided by the size of the destination market, Q_{nt}. This specification should be just as robust as Poisson PML on the non-normalized flows but it puts less weight on the larger trade values since its objective function is focused on minimizing deviations in shares.[13]

The first key finding of column (4) in Table 3 is that the coefficients on the brand-level τ and γ frictions determinants tend to be larger than their model-level counterparts. For example, the model-level τ distance elasticity in column (2) is -0.179, whereas the brand-

[13] Head and Mayer (2014) elaborate on these issues.

level distance elasticities are about four times as large: -0.699. This is exactly what one should expect if $\theta > \eta$, a relationship that also finds some support in the comparison of the coefficient on tariffs applied between the origin and destination of the trade flow.

The ratio of the parts to cars tariffs provides a second estimate of $\alpha_4 = 4.0/10.7 = 0.37$, higher, but in the same ballpark as the 0.25 estimate from column (1). Given the imprecision of the numerator and denominator estimates, the confidence intervals of the ratio are wide. Still, 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.

A second use of brand sales is to sum up the cars shipped from different origins to a given destination. Including the measurable version of our frictions into (6), we obtain the following estimable equation of the brand's average market share, *regardless* of assembly location

$$\mathbb{E}\left[\frac{q_{bnt}}{M_{bnt}Q_n}\right] = \exp\left[-\eta \mathbf{X}'_{int}\mathbf{d} + \mathrm{FE}_b - \eta\upsilon_1 \ln y_{it} - \eta\upsilon_2 \ln P^y_{it} + \mathrm{FE}_{nt} + \left(\frac{\eta}{\theta}\right)\ln \hat{D}_{bnt}\right]. \quad (11)$$

The structural interpretation of the fixed effects becomes $\mathrm{FE}_b = \eta(\ln(\varphi_b\beta_b) - \upsilon_{0i})$ and $\mathrm{FE}_{nt} = \ln(Q_{nt}/\Phi_{nt})$. The $\ln \hat{D}_{bnt}$ included as the last control comes from the sourcing probability results from equation (8) where $D_{bnt} = \sum_{k\in\mathcal{L}_{bt}}(\gamma_{ikt}w_{kt}\tau_{knt})^{-\theta}$ is the denominator of the choice probability. This regression has the advantage of permitting estimation of the δ_{int} terms because it employs brand and destination-year effects but does not require brand-destination-year effects.

We first estimate equation (11) with Poisson PML and report the results in column (5). In this case there are no zeros since the destination-years in the sample are limited to those where the brand has distribution activity. Since we can estimate the linear-in-logs version of equation (11) without selection bias, we do so for comparison purposes and report the results in column (6).

The key finding from brand-level regressions aggregating over sources is the confirmation of the importance of δ_{in} frictions already found in the model-level sales from columns (2) and (3). There is a very large premium of brand sales in their home country. In addition, we estimate that increasing consumer distance from headquarters lowers market shares, even controlling for distance from the consumer to the assembly location. Language effects maintain their negative impact on MS frictions at approximately the same level as observed in the model-level equations (as expected since the structural interpretation of the coefficient is the same). The effect of deep regional agreements is not estimated very precisely at the brand level, as compared to the model-level estimates from columns (2) and (3). One of the important result for our counterfactual is that this variable sees its most solid impact on the extensive margin of model entry, to which we turn next.

The column (5) and (6) results also corroborate the findings from columns (1)–(4) that $\theta > \eta$. This comes from the coefficient on $\ln \hat{D}_{bnt}$, which has a theoretical value of η/θ. Averaging the $\hat{\eta}$ from column (2) and (3) and dividing by the $\hat{\theta}$ averaged across columns (1) and (4) delivers an implied η/θ equal to 0.40, very similar to the 0.52 from column (5).

4.4 Market entry decision

Substituting $\delta_{int} = \exp(\mathbf{X}'_{int}\mathbf{d})$ and $\phi_{int} = \exp(\mathbf{X}'_{int}\mathbf{f})$ into equation (4) and introducing fixed effects, we obtain the estimable version of the model-market entry equation,

$$\text{Prob}(\mathbb{I}_{mnt} = 1) = \Lambda\left[\mathbf{X}'_{int}\mathbf{e} + \lambda(\eta - 1)\left(\frac{1}{\theta}\ln \hat{D}_{bnt} - \upsilon_1 \ln y_{it} - \upsilon_2 \ln P^y_{it}\right) + \text{FE}_b + \text{FE}_{nt}\right], \tag{12}$$

where the coefficients on the gravity determinants in $\mathbf{X}_{int}$ have structural interpretations given by $\mathbf{e} = -\lambda(\eta - 1)\mathbf{d} - \lambda\mathbf{f}$. Thus, the coefficients on the in friction determinants combine the δ_{in} variable marketing cost effects with the ϕ_{in} fixed marketing costs. All the γ and τ geography effects are captured in the $\ln D_{bnt}$ term. It can be seen as an index of how well-positioned brand b's assembly plants are to serve market n in t.

Column (7) of Table 3 shows that, with the exception of language, all the MS frictions 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 and contiguity promote entry as well. Deep RTAs between the headquarter country (i) and the market (n) raise the odds of model entry by 16%. As it seems unlikely that RTAs change preferences, we see the deep RTA$_{int}$ effects as supporting the cost-shifter interpretation. Under this approach, our MS 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 MS friction effect of RTA 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 (ϕ_{int}), rather than the variable marketing costs (δ_{int}), that affect brand sales as well.

Table 4: Statistical significance of friction categories in baseline results

	Regression specification from Table 3						
Dep. Var:	ℓ^*_{mnt}	$\ln q_{m\ell n}$		$\frac{q_{b\ell n}}{M_{bn}Q_n}$	$\frac{q_{bn}}{M_{bn}Q_n}$	$\ln\left(\frac{q_{bn}}{M_{bn}Q_n}\right)$	$\mathbb{I}_{mnt}$
Method:	cond. logit	OLS	OLS	Poisson	Poisson	OLS	logit
Test Statistic (d.f.)	χ^2 (6)	F(6, 48)	F(6, 48)	χ^2 (6)	χ^2(6)	F(6, 72)	χ^2(6)
Trade costs: $\tau_{\ell nt}$:	263.85	8.03	10.35	1500.81			
MP costs $\gamma_{i\ell t}$:	95.73	1.64†	2.25†	136.33			
MS costs δ_{int}, ϕ_{int}:		8.81	11.15		114.03	7.22	65.91

†: All tests have p-values near zero except 1.64 (0.16) and 2.25 (0.05).

Table 4 provides a summary of the statistical significance for each of the different frictions. Trade costs collectively have much higher relevance than MP frictions, as measured by F and χ^2 statistics. One might be worried that weak γ effects come from using brand headquarters instead of firm headquarters, which would add measurement error to all the frictions. Assuaging this concern, the finding of much a higher τ than γ significance is maintained when we use the firm-variety approach in Table 11. The significance of the δ effects can only be compared to that of τ and γ in the model-level estimates from the second and third columns. We see here that MS costs are the most significant of the three frictions. This is not the case using the parent-platform approach in Table 11 but the δ effects there continue to be highly significant and more significant than the γ determinants.

4.5 Backing out the frictions and assembly costs

The coefficients on tariffs in Table 3 identify the cost response elasticities η and θ. Combining them with the coefficients on the other friction variables, we now calculate the underlying friction parameters. These are the parameters that convert, for example, distance differences into cost differences. We also infer the relative costs (before frictions) of each assembly country (w_ℓ).

The estimates of the sourcing equation (2) reported in column (1) of Table 3 provide estimates of $-\widehat{\theta\boldsymbol{\rho}}$ and $-\widehat{\theta \mathrm{g}}$. Dividing by $-\hat{\theta} = -8.05$ from the sourcing equation, we obtain $\hat{\boldsymbol{\rho}}$ and $\hat{\mathrm{g}}$ and report them as the first two columns of Table 5.

In the case of the τ frictions shown in the first column of Table 5, 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. Our preferred estimate of the ℓn distance effect in column (1) of Table 5

Table 5: Friction parameters (preferred estimates)

Friction:	τ	γ	δ	ϕ
Estimate:	$\hat{\boldsymbol{\rho}}$	$\hat{\mathbf{g}}$	$\hat{\mathbf{d}}$	$\hat{\mathbf{f}}$
home	-0.129	-0.293	-0.226	-1.754
home $\times$ LDC			0.034	-3.487
ln distance	0.037	-0.021	0.066	0.217
contiguity	-0.017	-0.001	-0.045	-0.618
common language	0.012	0.012	-0.085	0.464
RTA (deep)	-0.031	-0.052	-0.017	-0.499

Elasticities used to obtain frictions: $\hat{\theta} = 8.05$, $\hat{\eta} = 3.77$, $\hat{\lambda} = 0.28$. Calculation of those frictions are described in the text and use coefficients from Table 3.

is $\hat{\rho} = 0.037$. Coşar et al. (2016) report a somewhat smaller value of $\hat{\rho} = 0.016$ (Table 12, column IV). Both estimates of $\hat{\rho}$, the delivery cost of 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. In a way this is not surprising in this context. Information is clearly not a problem in the sourcing equation since car firms presumably know their own costs. Moreover taste differences and trust issues (other candidates for dark matter) should show up mainly in the MS frictions.

Estimates of the variable multinational sales friction parameters ($\mathbf{d}$) are obtained from the brand-level results in columns (5) and (6). Dividing these coefficients by $-\hat{\eta} = -3.77$ delivers two sets of estimates of $\widehat{\mathbf{d}}$. We average them and report the result in the third column of Table 5.

The fixed cost parameters are obtained from the entry equation (4). Using its estimable form shown in equation (12), we calculate an estimate of $\mathbf{f}$ using the $\hat{\mathbf{d}}$ we just obtained:

$$\hat{\mathbf{f}} = -\hat{\mathbf{e}}/\hat{\lambda} - (\hat{\eta} - 1)\hat{\mathbf{d}}. \tag{13}$$

The remaining unknown, $\hat{\lambda}$, is inferred from the coefficient on $\ln D_{bn}$ in this entry equation combined with the one from brand-level equation (11). Taking the ratio of entry (columns 7) to brand (averaged columns 5 and 6) coefficients and multiplying by $\hat{\eta}/(\hat{\eta} - 1)$ yields $\hat{\lambda} = 0.28$. Plugging this into equation (13) completes the set of preferred parameters for each of the frictions.

Alternative estimates of the same underlying frictions could be obtained from specifications in Table 3: Columns (2), (3) and (4) can each be used to back out $\hat{\boldsymbol{\rho}}$ and $\hat{\mathrm{g}}$ while columns (2) and (3) can also deliver $\hat{\mathbf{d}}$. We show in appendix E that the preferred and alternative estimates of the frictions correlate strongly. The key finding for our purposes concerns the RTA effects on the τ and γ frictions. While they are imprecisely estimated in the sourcing equation, the point estimates of the frictions are almost the same as the corresponding estimates based on the brand-level sales equations where the degree of significance for both was considerably higher.

How should we interpret the parameters shown in Table 5? For τ, γ, and δ the coefficient on the kth element of **X** maps to proportional increases in price (an ad-valorem equivalent) of $\exp(\rho^{(k)}\Delta X^{(k)})-1$. For the parameters determining fixed costs the mapping is less obvious. One natural approach is to consider the decrease in marginal costs (and hence price) that would be needed to keep a hypothetical variety indifferent between entry and staying out. Referring back to equation (3), we see that multiplying any of the variable cost shocks by $1-s$ raises expected profit gross of fixed costs by $(1-s)^{1-\eta}$. The proportional change in fixed costs from doubling a continuous variable cost determinant with attached friction f would be 2^f. Solving for the *ad valorem* equivalent (AVE), we have $s = 1-(2^f)^{1/(1-\eta)}$. For any of the friction-*reducing* binary variables, the corresponding formula is $s = 1-\exp(-f)^{1/(1-\eta)}$. Applying these two formulas respectively to a doubling of distance or the end of a deep RTA, we obtain AVEs of 5.3% and 16.5%. The corresponding variable cost (δ) effects on MS frictions are $2^{0.066}-1 = 4.7\%$ and $\exp(0.017)-1 = 1.7\%$. Combining δ and ϕ AVEs we obtain 10% for distance doubling and 18.2% for RTAs. These are substantially larger than their τ (2.6% and 3.1%) and γ (-1.4% and 5.3%) counterparts, indicating that marketing costs are quantitatively important enough to warrant inclusion in the multinational production framework. The tariff equivalent of the variable cost disadvantage for foreign-headquartered cars (in an OECD market) is $(\exp(0.226) = 25\%$, almost twice as high as the $\exp(0.129) = 14\%$ penalty for foreign-assembled cars. Set on top of that, the extra burden of fixed costs for foreign headquarters has an AVE of 47%.

The index of local assembly costs in each country, w_ℓ, is a key parameter of the model because it tells us where production would gravitate in the absence of frictions. We obtain the 2013 levels as $\hat{w}_\ell = \exp(-\widehat{\mathrm{FE}}_\ell/\hat{\theta})$ from the estimates in the sourcing equation. The w_ℓ can only be identified up to a scalar so we express them all as cost advantages with respect to Italy, i.e. $-(\hat{w}_\ell - \hat{w}_{\mathrm{ITA}})/\hat{w}_{\mathrm{ITA}}$.

Figure 5 (a) graphs the cost advantage of the twenty four lowest cost assembly countries. The clear "winner" for the car industry is South Korea with Japan as runner up. Egypt is the outlier in the other direction in Figure 5 (b), which depicts the 23 highest

Figure 5: Cost advantage inferred from sourcing decisions

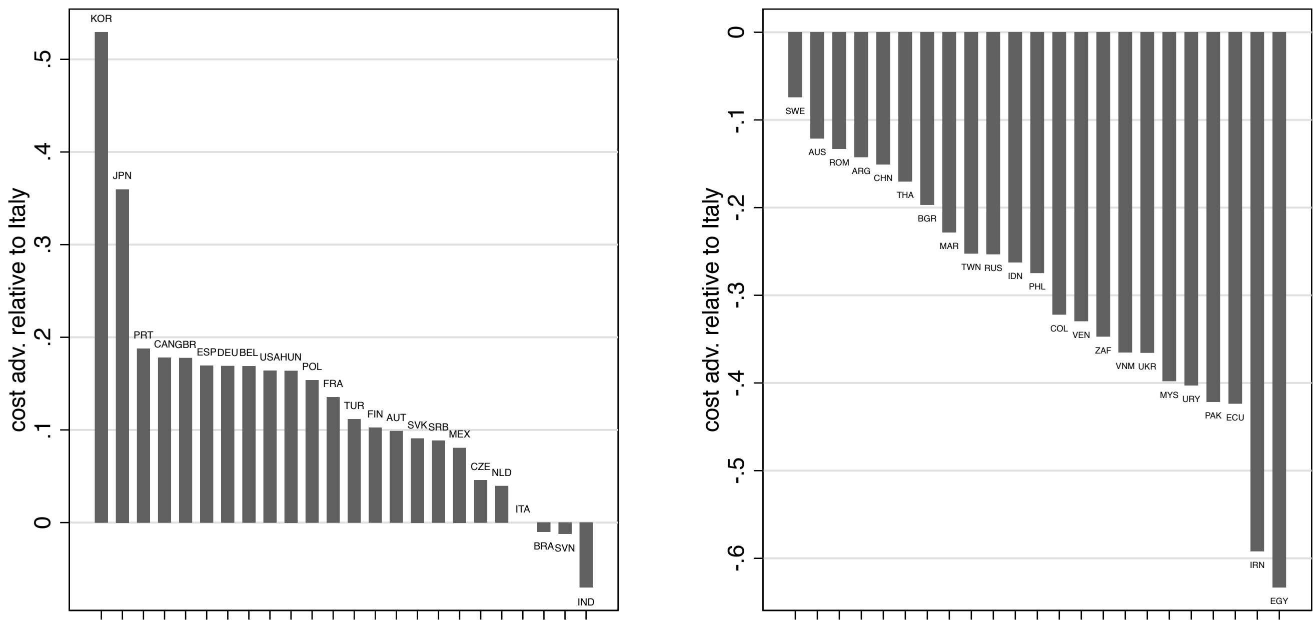

(a) Top 24

(b) Bottom 23

cost countries. 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 Counterfactuals

We motivated the paper with the issue of how regional integration agreements reshape the spatial allocation of multinational production and trade. Having estimated the equations implied by the model, we conduct counterfactual policy changes to investigate the impact of preferential integration on the location of production (sourcing), allocations of models (entry), shipments across markets (brand-level sales) and consumer surplus.

Four prospective RTA changes are the subject of public debate at the time of writing:

1. Dismantling NAFTA: the proposals of presidential candidate Donald Trump to raise tariffs on imports from Mexico to 35%, and/or exit from NAFTA.

2. Brexit: the exit of the United Kingdom from the EU.
3. Enactment of the Trans-Pacific Partnership (TPP): Australia, Brunei (not in IHS data), Canada, Chile, Japan, Malaysia, Mexico, New Zealand, Peru, Singapore, the United States, and Vietnam.
4. Enactment of the Transatlantic Trade and Investment Partnership (TTIP), an integration agreement between the EU and the US.

In the next two subsections, we provide some details on the data and parameters that are the inputs into the counterfactual as well as the algorithm that generates the outputs: changes in how much is produced where and changes in consumer surplus.

5.1 Summary of exogenous variables and parameters

The exogenous variables in the counterfactuals are the $\mathbf{X}$ matrix of friction determinants, country-level new car purchases (Q_n), the brand's total number of models (M_b), each brand's set of production locations, $\mathcal{L}_b$, and its set of countries with brand dealerships. The counterfactuals shock two of the friction determinants: tariffs and the deep RTA dummies. The counterfactuals are all carried out using data from 2013 (the last available in our sample) so we suppress the time subscripts in this section.

Section 4.5 describes how we back out the friction estimates from our regressions. There are two remaining sets of brand-specific and country-specific variables that play important parts in the analysis:

Brand effects are obtained by combining estimates derived from the OLS and Poisson estimates of equation (11). Using $\widehat{\frac{\beta_b\varphi_b}{w_i}} = \exp(\widehat{\text{FE}}_b/\hat{\eta})$, we divide the estimated brand fixed effects from each regression by $\hat{\eta} = 3.77$, exponentiate, and then average.

Market-specific fixed costs of entry : We solve for the implied market-specific fixed costs of adding a model, J_n, utilizing the market fixed effects in equation (12) and our estimate of $\hat{\lambda} = 0.28$.

$$J_n = \widehat{\text{FE}}_n/\hat{\lambda} + \ln Q_n - \ln \widehat{\Phi}_n,$$

where $\widehat{\Phi}_n = \hat{\kappa}_1 \sum_b M_{bn} \left(\widehat{\frac{\beta_b\varphi_b}{w_i}}\right)^{\hat{\eta}} \hat{\delta}_{in}^{-\hat{\eta}} \widehat{D}_{bn}^{\frac{\hat{\eta}}{\hat{\theta}}}$, our estimate of Φ_n that uses actual data for M_{bn} (rather than solving for endogenous entry levels as we do in the actual simulation stage). At this stage, $\widehat{D}_{bn} = \sum_{k\in\mathcal{L}_b}(\hat{w}_\ell\hat{\gamma}_{ik}\hat{\tau}_{kn})^{-\hat{\theta}}$ can be completely calculated from the set of parameters at hand.

5.2 Algorithm solving for endogenous variables

The endogenous variables in our model are M_{bn} (entry counts), Φ_n (needed for consumer surplus), and $q_{b\ell n}$ (which determines the impact of changes on firms and workers). The counterfactual exercises require an algorithm for dealing with the simultaneity between the model entry decision and the overall index of competition in the market, Φ_n.

The goal is to solve for expected values of brand-origin-destination flows under factual, $\bar{q}_{b\ell n}$, and counterfactual, $\widetilde{q}_{b\ell n}$, settings. It is useful to express the brand-level equation (5) for factual and counterfactual sales as being multiplicatively separable between two probabilities:

$$\bar{q}_{b\ell n} = Q_n \bar{\mathbb{P}}_{bn} \bar{\mathbb{P}}_{\ell|bn} \qquad \text{and} \qquad \widetilde{q}_{b\ell n} = Q_n \widetilde{\mathbb{P}}_{bn} \widetilde{\mathbb{P}}_{\ell|bn}$$

The first probability, $\mathbb{P}_{bn}$ gives the expected share of sales in n going to brand b. The second probability, $\mathbb{P}_{\ell|bn}$, governs the sourcing decision. It is straightforward to calculate $\widetilde{\mathbb{P}}_{\ell|bn}$, since this only involves an update of the frictions in the numerator and denominator, $\widetilde{D}_{bn} = \sum_{k \in \mathcal{L}_b} (\hat{w}_\ell \tilde{\gamma}_{ik} \tilde{\tau}_{kn})^{-\hat{\theta}}$, of the probability formula. Calculating $\widetilde{\mathbb{P}}_{bn}$ is trickier since model choice depends on the availability and prices of models in each market, captured in $\widetilde{\Phi}_{bn}$, but entry itself depends on the same index. We therefore need to solve simultaneously for equilibrium levels of M_{bn} and Φ_{bn}, first as expected values, the factual, and then in the counterfactual scenario.

The expected price index in n, $\bar{\Phi}_n$, is the same as $\hat{\Phi}_n$ except that the actual number of models sold by brand b in market n is replaced with its expected value ($\bar{M}_{bn}$):

$$\bar{\Phi}_n = \hat{\kappa}_1 \sum_b \bar{M}_{bn} \left(\frac{\widehat{\beta_b \varphi_b}}{w_i} \right)^{\hat{\eta}} \hat{\delta}_{in}^{-\hat{\eta}} \widehat{D}_{bn}^{\frac{\hat{\eta}}{\hat{\theta}}}. \tag{14}$$

Using the logit estimates of the brand effects ($\widehat{\mathrm{FE}}_b$), together with other estimated parameters, the expected number of entrants is M_b times the probability of models from brand b entering market n (equation 4):

$$\begin{aligned}\bar{M}_{bn} = & M_b \times \Lambda \Big[\hat{\lambda}(\ln \hat{\kappa}_2 - \ln \hat{\eta}) - \hat{\lambda}(\hat{\eta} - 1) \ln \hat{\delta}_{in} - \hat{\lambda} \ln \hat{\phi}_{in} \\ & + \frac{\hat{\lambda}(\hat{\eta} - 1)}{\hat{\theta}} \ln \widehat{D}_{bn} + \widehat{\mathrm{FE}}_b + \hat{\lambda}(\ln Q_n - \ln \bar{\Phi}_n - \widehat{J}_n) \Big]. \end{aligned} \tag{15}$$

The algorithm solves this set of nonlinear equations through an iteration process using (14) and (15). We begin with a guess of $\bar{\Phi}_n$ in equation (14), where $\bar{M}_{bn}$ is *initialized as the* *count of realized entrants*. This permits calculation of the implied model entry flows, $\bar{M}_{bn}$ from equation (15), leading to a new value of $\bar{\Phi}_n$ in (14). This *tâtonnement* process is not a contraction mapping, so a dampening factor (set equal to 0.3) is used to reach the fixed points for each market. Since there is no feedback to $\widehat{\mathbb{P}}_{\ell|bn}$ we can use the $\widehat{D}_{bn}$ obtained at realized friction values for all iterations.

Substituting in those factual expected values $\bar{M}_{bn}$ and $\bar{\Phi}_n$, the expected brand-market shares are given by

$$\bar{\mathbb{P}}_{bn} = \frac{\hat{\kappa}_1 \bar{M}_{bn} \left(\widehat{\frac{\beta_b \varphi_b}{w_i}}\right)^{\hat{\eta}} \hat{\delta}_{in}^{-\hat{\eta}} \widehat{D}_{bn}^{\frac{\hat{\eta}}{\hat{\theta}}}}{\bar{\Phi}_n}. \tag{16}$$

To obtain expected brand-level shipments from ℓ to n, one then just needs to plug in values to $\bar{q}_{b\ell n} = Q_n \bar{\mathbb{P}}_{bn} \bar{\mathbb{P}}_{\ell|bn}$.[14]

Figure 6 graphs true brand-origin-destination sales ($q_{b\ell n}$) against simulation-predicted sales ($\bar{q}_{b\ell n}$) with both expressed on a log scale. The data cluster around the 45-degree line, obtaining a correlation (in logs) of 0.7. Part of the high explanatory power stems from the presence of Q_n in the prediction. 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.

The counterfactual solution of the endogenous variables is then obtained after changing the frictions τ, γ, δ, and ϕ from their "hat" settings to their counterfactual "tilde" settings by changing the level of tariffs and turning on or off the corresponding deep RTA indicator. The iteration described above provides $\widetilde{M}_{bn}$ and $\widetilde{\Phi}_{bn}$, giving counterfactual market shares $\widetilde{\mathbb{P}}_{bn} = \left[\hat{\kappa}_1 \widetilde{M}_{bn} \left(\widehat{\frac{\beta_b \varphi_b}{w_i}}\right)^{\hat{\eta}} \tilde{\delta}_{in}^{-\hat{\eta}} \widetilde{D}_{bn}^{\frac{\hat{\eta}}{\hat{\theta}}}\right] / \widetilde{\Phi}_n$. Combining with new sourcing probabilities, one can calculate the counterfactual flows $\widetilde{q}_{b\ell n} = Q_n \widetilde{\mathbb{P}}_{bn} \widetilde{\mathbb{P}}_{\ell|bn}$.

5.3 Results from counterfactual experiments

Tables 6, 7, 8, and 9 depict changes in production destined to three aggregated markets: the domestic market of the country listed, members of the RTA other than that country, and the rest of the world (ROW). We also show the level of *expected* production in the baseline (factual) situation. Figures 7, 8, 9, and 10 show graphically the most affected countries, in terms of changes in the country's share of world production (panel a).

Consumer surplus changes are reported in percent in the final column. They are calculated as $\left(\widetilde{\Phi}_n / \bar{\Phi}_n\right)^{\frac{1}{\hat{\eta}}} - 1$ (this formula is the logit equivalent of the change in the price

[14] Again, since $\widehat{D}_{bn}$ does not involve either $\bar{M}_{bn}$ or $\bar{\Phi}_n$, $\widehat{\mathbb{P}}_{\ell|bn}$ is not affected by the iterative looping procedure, and we can use $\bar{\mathbb{P}}_{\ell|bn} = \widehat{\mathbb{P}}_{\ell|bn}$.

Figure 6: Fit of $q_{b\ell n}$ data to expected values of solved model

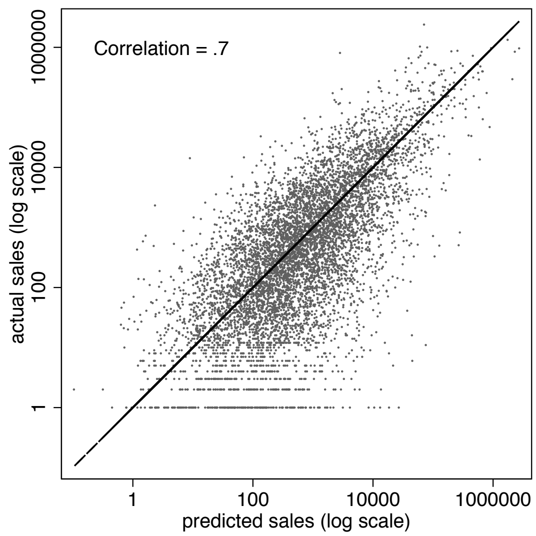

index in the CES demand system, see Anderson et al. (1992) for details). Panel (b) of each relevant figure also illustrates changes in consumer surplus. We leave exhaustive welfare analysis (including changes in profits, impact on workers, and tariff revenues) for future work.

5.3.1 Dismantling the North American Free Trade Agreement (NAFTA)

The first counterfactual considers three potential ways that trade preferences between the US, Mexico, and Canada might be altered. While dissolving NAFTA looked like an unlikely scenario before the summer of 2016, it has gained topical interest with the proposals of candidate Trump in the US presidential campaign. It is also of independent interest because it bears on the long-standing question of whether free trade with Mexico has been bad for US manufacturing. Finally, it provides a useful setting to distinguish quantitatively between changes in the trade costs of final goods versus the whole ensemble of frictions incorporated into our framework.

Our first scenario takes seriously the proposal by Donald Trump to impose 35% tariffs on imports from Mexico. Anticipating an equal retaliation, this counterfactual also resets the Mexican tariff on US cars and parts to 35%. Paralleling the Brexit abbreviation, we refer to this case as the "Trumpit" counterfactual. The second scenario envisions the dissolution of NAFTA, where all three members end deep integration and revert to imposing

MFN tariffs on each other.[15]

The first panel of Table 6 lists the most impacted nations in the Trumpit scenario. Mexican production and, presumably, employment are predicted to decline by a devastating 41%. The Mexican plants not only lose sales in the US market due to the application of the Trump tariff to cars, the Chevrolet and Ford plants also lose thousands in sales to Canada and other markets. One major reason is that the 35% retaliatory duties significantly raise these plants parts costs, since our estimates imply that one quarter of the costs of assembly are attributable to parts imported from the HQ country. The loss of a deep RTA integration further increases assembly costs by about 5%. The combined cost shock to US-owned plants in Mexico is so large that they actually lose small amounts of sales in the Mexican market. Only the non-US-owned plants expand production in Mexico, which helps to explain why the increase in production for the domestic market is just 7,090 cars.

The main beneficiaries from the collapse of Mexican exports to the USA are the US-based producers, followed by plants located in Japan and Canada (which keep their preferential access to the US market). It is notable that most EU-based producers also gain because their relative access has been made better by the spike in the US-Mexico tariff. Not surprisingly, Trumpit is very bad for Mexican car buyers who experience consumer surplus losses of almost 6%.

The outcomes in the Trumpit scenario are powerful illustrations of the role played by MP and MS frictions. This can be seen in the three rows in the middle of Table 6. There we limit the friction changes to the tariffs on final cars and the loss of deep RTA in the ℓn (assembly to market) dimension. In this thought experiment the rise in production in Mexico for the Mexican market is eight times as high. Now *all* brands increase production for the host-country market and there are no cuts in sales to third-country markets. In terms of aggregate effects, the decline in Mexican production is 1.6 times higher and the loss to Mexican consumers is 1.4 times higher in the counterfactual involving the full set of frictions.

In the end-of-NAFTA scenario, depicted in orange in Figure 7, all three members reduce production. The gain in US production for its home market is now dominated by the lost sales to Canada and Mexico, resulting in an overall fall of 0.9%. This is much smaller than the drops in the two partners' production, with Canada losing 24%, and Mexico 17% due to the loss of their favored status in the US market. This joint loss represents nearly one percent of world production, roughly equally shared between Canada and Mexico,

[15]This implicitly assumes that pre-NAFTA agreements, such as the 1965 Canada-US Auto Pact, are not reinstated.

Table 6: Undoing NAFTA

Country	Change in production by market (# cars)				% Chg.	Base (mn)	CS % Chg.
	Domestic	*RTA*	*ROW*	*Total*			
Trump tariffs: reciprocal 35% duties and loss of deep RTA for USA-MEX pair							
MEX	7090	-439073	-133738	-565721	-41.1	1.376	-5.57
JPN	704	132542	10780	144026	1.1	12.552	0
CAN	3324	53827	3624	60775	4.6	1.312	-.16
KOR	177	38082	13698	51957	.7	6.999	0
DEU	1167	43458	6339	50964	1.1	4.623	-.02
USA	239441	-212549	22415	49307	.5	9.378	-.4
BRA	26485	5717	2442	34644	2	1.743	-.26
ESP	300	17398	3103	20801	1.7	1.205	-.03
FRA	600	16800	3315	20715	.8	2.71	-.02
GBR	701	13704	1266	15671	1.2	1.315	-.04
ITA	278	9884	1214	11376	2.4	.482	-.03
BEL	82	8817	2256	11155	3.1	.361	-.02
POL	54	9128	1939	11121	3.6	.306	-.03
ARG	1550	2429	5587	9566	2.1	.453	-.16
IND	3134	4036	1362	8532	.3	2.604	-.02
Trump tariffs: $\tau_{\ell n}$ effects only							
MEX	57266	-408506	0	-351240	-25.5	1.376	-4.06
CAN	0	48154	0	48154	3.7	1.312	0
USA	230971	-222930	0	8041	.1	9.378	-.39
End of NAFTA							
CAN	15201	-300129	-34188	-319116	-24.3	1.312	-4.38
MEX	43747	-213320	-68591	-238164	-17.3	1.376	-6.15
JPN	796	218626	9253	228675	1.8	12.552	0
KOR	233	78608	11620	90461	1.3	6.999	0
USA	283982	-384658	18813	-81863	-.9	9.378	-.44
DEU	967	61595	5308	67870	1.5	4.623	-.01
BRA	15411	9593	1863	26867	1.5	1.743	-.15
ESP	254	22634	2625	25513	2.1	1.205	-.03
FRA	477	21966	2724	25167	.9	2.71	-.02
GBR	566	21438	1037	23041	1.8	1.315	-.03
BEL	65	12534	1921	14520	4	.361	-.02
POL	45	11920	1751	13716	4.5	.306	-.02
ITA	228	10341	1070	11639	2.4	.482	-.02
CHN	4811	4953	1206	10970	.1	10.879	-.01
IND	1717	7489	1194	10400	.4	2.604	-.01

Elasticity parameter relevant for the Consumer Surplus calculation is $\eta = 3.77$.

Figure 7: Dismantling NAFTA

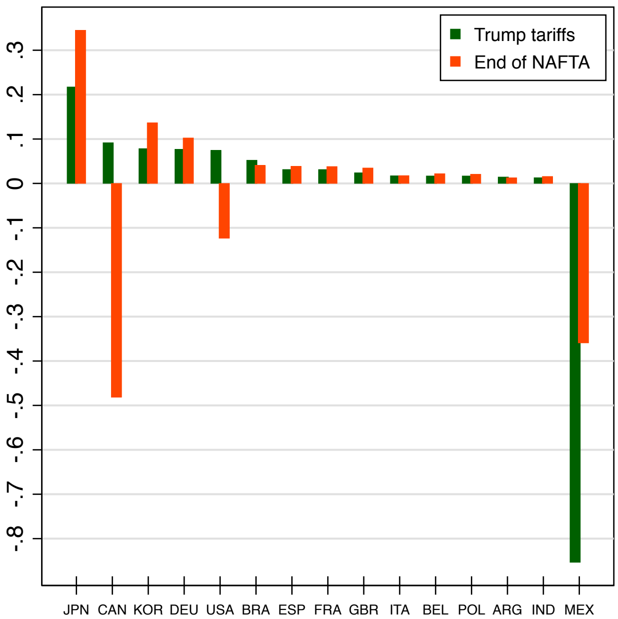

(a) Change in % share of world output

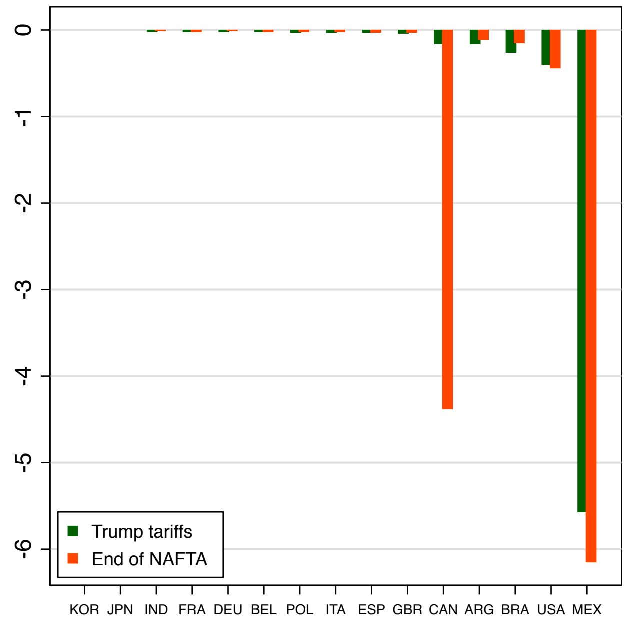

(b) % Change in consumer surplus

whereas a similar total decrease is entirely borne by Mexico in the Trumpit scenario.

The Canadian and Mexican outcomes further worsen due to increases in delivered costs of tangible and intangible (γ) inputs suffered by US-owned plants. The two partners concede sales in the US market, in their respective domestic markets, and most importantly, in ROW markets. As a consequence, production in Korea, Japan, Germany, and Belgium rise not only because of higher sales in NAFTA, but also because of higher market shares in ROW markets.

An example of the richness of the MP effects comes from comparing how Korean and Japanese production benefit from NAFTA dissolution. As the costs of serving the Canadian market from the US rise, Toyota's probability of sourcing from the US falls by 5.3 percentage points (from 10.9% to 5.6%). Toyota's Japanese plants are the biggest beneficiaries, gaining 2.6 points, but its Canadian plants grab 1.6%. Hyundai's US plant also sees a sharp fall in its sourcing probability (8 points). However, as Hyundai lacks a plant in Canada, it reallocates nearly the entire probability (6.8 points) to its Korean plants.

A final point of interest is that the increases in variable and fixed marketing costs (δ and ϕ) have asymmetric effects on consumers. Because there are no Canadian or Mexican-based brands, the US customer is unaffected by the rise in those two costs. On the contrary, Canadian and Mexican consumers witness a steep increase in the marketing costs of varieties produced by US brands (a large share of their consumption basket). US brands collectively drop 3 models in Canada and 8 in Mexico. Combining all mechanisms, consumer surplus losses come to 4% in Canada and 6% in Mexico.

5.3.2 United Kingdom exits the European Union (Brexit)

Figure 8: The effects of Brexit

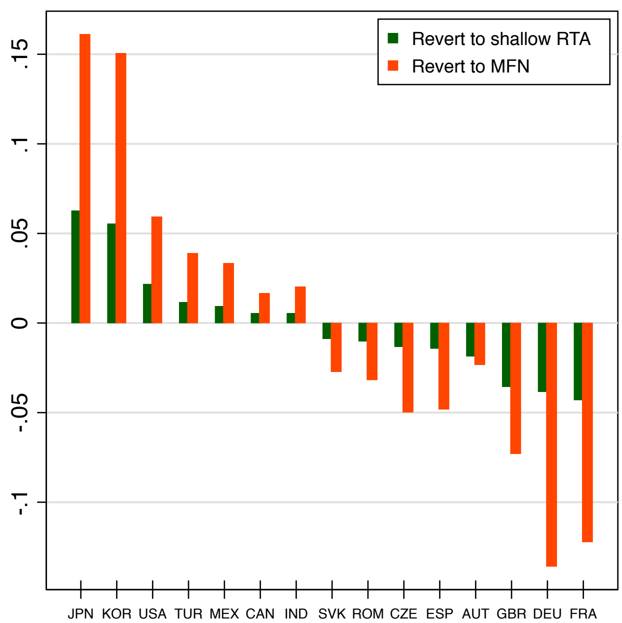

(a) Change in % share of world output

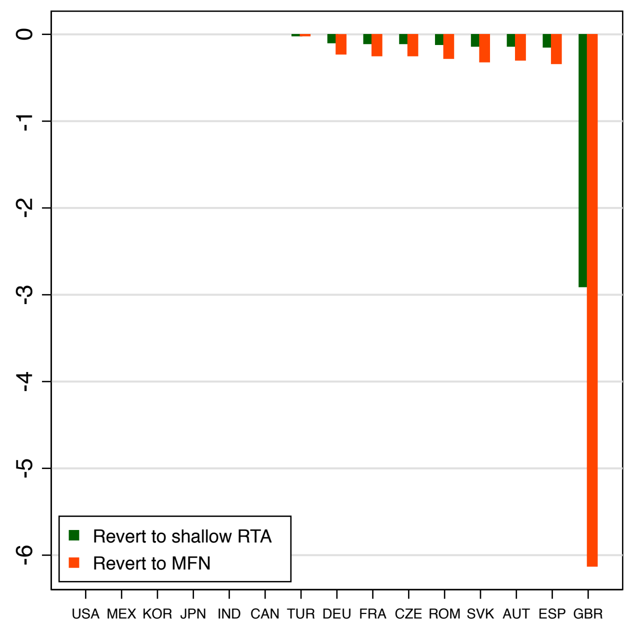

(b) % Change in consumer surplus

Following the June 23, 2016 vote for the UK to leave the EU, it was unclear what trade arrangement would be negotiated between the two parties. Here we consider two potential post-membership relationships between the UK and EU. The shallow RTA 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 then simulate the scenario where UK exports face the EU's MFN tariffs while the EU reciprocates at the same rates. Both cases hold constant all the RTA relationships the UK currently enjoys through its EU membership (e.g. in particular the customs union with Turkey and the EU-Mexico FTA).

Table 7: Brexit: UK exit from the EU

Country	Change in production by market (# cars)				% Chg.	Base (mn)	CS % Chg.
	Domestic	*RTA*	*ROW*	*Total*			
Revert to shallow RTA							
JPN	35	40386	1140	41561	.3	12.552	0
KOR	18	35223	1458	36699	.5	6.999	0
FRA	3694	-32434	256	-28484	-1.1	2.71	-.11
DEU	7836	-34583	1381	-25366	-.5	4.623	-.1
GBR	47472	-68231	-2782	-23541	-1.8	1.315	-2.91
USA	538	13412	445	14395	.2	9.378	0
AUT	-225	-6261	-5788	-12274	-6.6	.187	-.14
ESP	984	-10737	324	-9429	-.8	1.205	-.15
CZE	227	-9069	142	-8700	-.7	1.188	-.11
TUR	45	7566	68	7679	1.1	.683	-.02
ROM	82	-6796	49	-6665	-1.5	.444	-.12
MEX	14	5985	119	6118	.4	1.376	0
SVK	40	-5900	76	-5784	-.9	.621	-.14
CAN	24	3427	122	3573	.3	1.312	0
IND	9	3392	168	3569	.1	2.604	0
Revert to MFN							
JPN	39	105514	1276	106829	.9	12.552	0
KOR	20	98117	1633	99770	1.4	6.999	0
DEU	16428	-108059	1547	-90084	-1.9	4.623	-.23
FRA	8833	-90144	287	-81024	-3	2.71	-.25
GBR	124238	-169474	-3108	-48344	-3.7	1.315	-6.13
USA	603	38170	498	39271	.4	9.378	0
CZE	536	-33663	159	-32968	-2.8	1.188	-.25
ESP	2315	-34613	363	-31935	-2.7	1.205	-.34
TUR	51	25681	76	25808	3.8	.683	-.02
MEX	16	21915	133	22064	1.6	1.376	0
ROM	191	-21297	55	-21051	-4.7	.444	-.28
SVK	93	-18206	85	-18028	-2.9	.621	-.32
AUT	-71	-8858	-6488	-15417	-8.3	.187	-.3
IND	10	13208	188	13406	.5	2.604	0
ITA	1861	10484	52	12397	2.6	.482	-.3

Elasticity parameter relevant for the Consumer Surplus calculation is $\eta = 3.77$.

The post-Brexit shallow RTA scenario sets all the deep RTA dummies to zero if they correspond to dyads involving the UK and EU. As indicated in section 4.5, this involves raising τ by 3%, γ by 5%, δ by 2% and ϕ by an AVE of 16%. The first panel of Table 7 shows a 24 thousand reduction in car production in the UK. Assuming roughly one worker per 50 cars (a rough average among the major plants in the UK), this works out to a loss of about 500 assembly jobs. Although production for the home market rises, it does not come close to offsetting losses in exports to the EU and to rest-of-world (ROW). The latter comes from a rise in γ for Opel, the only EU-headquartered brand still assembling cars in the UK in 2013. The idea is that absent deep integration, supplying inputs from Opel's headquarters in Germany becomes more costly. Consumer surplus in the UK falls by about 3% due mainly to higher import prices. The number of models sold in the UK only falls slightly for most brands, with the largest impact on the entry margin being a predicted cut of two Volkswagen models.

The 7% reduction in Austria's production occurs because it hosts three UK brands (and no home-based brands). Expressed in terms of world production in panel (a) of figure 8, UK and France incur the greatest losses. Consumers losses are minimal in countries other than the UK as appears clearly in panel (b) of the same figure.

The scenario in which the UK fails to maintain tariff-free access to the EU market dramatically increases the losses for UK production and consumer surplus. Imposing the current 10% MFN duty on EU imports doubles consumer losses to 6% while UK car production shrinks by 48 thousand cars—or about 1,000 workers—compared to what the model predicts under full integration with the EU. The large rises in production in Turkey (3.8%) and Mexico (1.6%) are consequences of the assumption that the UK stays in trade agreements with those countries that previously signed RTAs with the EU.

5.3.3 Trans-Pacific Partnership (TPP)

Table 8 displays the predicted impact of TPP on the fifteen countries with the largest (absolute) changes in output. The first panel of the table treats the TPP as if it were a standard free trade agreement. While the trade liberalization is modest for some rich countries like the USA or Japan that already impose almost no tariff on their future trade partners, some tariff cuts will be drastic. Japanese exporters face a 38% tariff when exporting to Vietnam, 20% to Malaysia, and 6–7% to Canada, Mexico, New-Zealand or Peru. US exports to Vietnam even face a record 58% duty in the last year of our sample (and 22% to Malaysia). As one would expect, each of the major car-producing members reduces output for its home market. In the case of the US, the reduction amounts to 63 thousand cars. Japan's tiny reduction in home market sales is dwarfed by nearly half a million increase in car

Table 8: Implementing TPP

Country	Change in production by market (# cars)				% Chg.	Base (mn)	CS % Chg.
	Domestic	*RTA*	*ROW*	*Total*			
Shallow RTA							
JPN	-2019	476445	-9399	465027	3.7	12.552	.01
USA	-62687	-31512	4684	-89515	-1	9.378	.19
KOR	-157	-78682	-2436	-81275	-1.2	6.999	0
MYS	-52380	2863	1315	-48202	-12	.403	9.12
CAN	-20251	-10194	-450	-30895	-2.4	1.312	1.26
DEU	-146	-26692	-758	-27596	-.6	4.623	0
THA	-3552	-17289	-522	-21363	-2.7	.797	.09
FRA	-100	-19807	-687	-20594	-.8	2.71	0
GBR	-148	-18338	-588	-19074	-1.5	1.315	0
VNM	-13037	7253	23739	17955	8.5	.211	26.16
CHN	-5224	-11196	-208	-16628	-.2	10.879	.01
IDN	-510	-14068	-656	-15234	-3	.513	.04
MEX	-14108	-616	288	-14436	-1	1.376	2.06
IND	-1175	-10320	-411	-11906	-.5	2.604	.01
ESP	-33	-10661	-351	-11045	-.9	1.205	0
Deep RTA							
JPN	-69561	1000621	-105040	826020	6.6	12.552	.24
KOR	-2053	-222809	-20280	-245142	-3.5	6.999	.04
CAN	-9509	88235	98588	177314	13.5	1.312	4.32
USA	-319002	15355	148199	-155448	-1.7	9.378	2.16
DEU	-2429	-75855	-10536	-88820	-1.9	4.623	.04
CHN	-38721	-19622	-2766	-61109	-.6	10.879	.08
MEX	-16715	-33786	-5638	-56139	-4.1	1.376	2.48
MYS	-56490	5615	-188	-51063	-12.7	.403	10.1
GBR	-2649	-31436	-8925	-43010	-3.3	1.315	.07
FRA	-1748	-28776	-10939	-41463	-1.5	2.71	.04
VNM	-13169	21504	31036	39371	18.6	.211	27.77
BRA	-11551	-18244	-4112	-33907	-1.9	1.743	.11
THA	-5618	-22843	-3715	-32176	-4	.797	.14
ESP	-513	-23698	-4188	-28399	-2.4	1.205	.09
AUS	-11093	19634	18274	26815	11.3	.238	5.55

Elasticity parameter relevant for the Consumer Surplus calculation is $\eta = 3.77$.

Figure 9: The effects of TPP

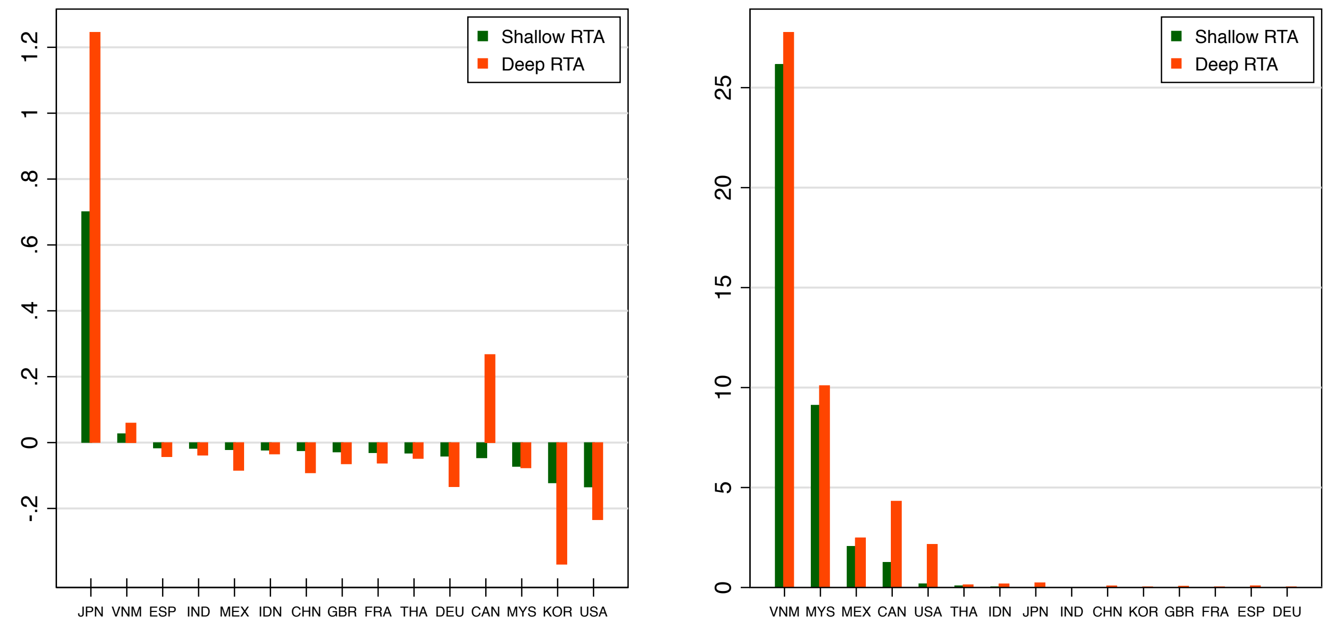

(a) Change in % share of world output

(b) % Change in consumer surplus

sales to the other TPP countries. Overall, Japan expands production by 3.7% while the US industry contracts by 1%. An explanation for the big surge in Japanese production is that it is estimated to have a substantial cost advantage over the other TPP producers as seen in Figure 5.

From the point of view of Japanese workers, the deeper integration results shown in the lower panel would be appealing, as a deep RTA doubles the gains in exports to TPP partners to pass one million. However, TPP also lowers costs of Japanese plants in the US, Australia, Canada, Mexico, Malaysia, and Vietnam by 5%, leading to an overall gain of 826 thousand cars (6.6% of its 2013 production). This γ effect is big enough to convert Canada from a net production loss of 31 thousand to a production gain of 178 thousand. TPP constitutes a major reshuffling even when expressed in terms of world output of cars. Figure 9 shows that Japan gains more than 1.2% of world car production, mostly at the expense of Korea and of the United States.

While the TPP looks bad for US auto workers, there are large predicted benefits for US consumers. The Japanese brands add six models in both the Canadian and US markets. The US price index for cars falls by 2.2% when all frictions are removed. Canadian consumers gain nearly twice as much. Naturally, the consumer surplus gains are much larger in Vietnam and Malaysia that reduce tariffs from an initially very high level. An interesting difference between these two countries regards the change in total production. While Figure 5 tells us that the cost disadvantage of both producers is comparable, their tariff pattern is not. While tariffs on car parts are generally smaller than on finished cars, Vietnam still taxes parts imported from the US and Japan at respective rates of 18 and 10%, while the rate is 1% in Malaysia. The Toyota and Nissan factories hosted in Malaysia are therefore predicted to see only a modest growth in their sales to ROW, while the corresponding increases in Japanese and US-owned plants located in Vietnam are predicted to be very large.

5.3.4 Transatlantic Trade and Investment Partnership (TTIP)

Table 9: Implementing TTIP

Country	Change in production by market (# cars)				% Chg.	Base (mn)	CS % Chg.
	Domestic	*RTA*	*ROW*	*Total*			
Shallow RTA							
USA	-62527	265188	-1308	201353	2.1	9.378	.13
KOR	-83	-45141	-1088	-46312	-.7	6.999	0
JPN	-179	-43489	-2150	-45818	-.4	12.552	0
DEU	-20880	-3884	2104	-22660	-.5	4.623	.25
FRA	-8126	-13601	-417	-22144	-.8	2.71	.26
MEX	-264	-13838	-743	-14845	-1.1	1.376	.02
CAN	-379	-11002	-573	-11954	-.9	1.312	.02
BEL	-359	8270	3335	11246	3.1	.361	.32
POL	-289	7692	3553	10956	3.6	.306	.38
ITA	-2883	-4909	-281	-8073	-1.7	.482	.36
Deep RTA							
USA	-340549	487197	37970	184618	2	9.378	1.04
JPN	-2108	-156677	-25422	-184207	-1.5	12.552	.01
KOR	-1001	-121050	-13371	-135422	-1.9	6.999	.02
BEL	1928	97603	29725	129256	35.8	.361	.92
POL	1009	89535	31592	122136	40	.306	1.08
ESP	-2929	86455	26100	109626	9.1	1.205	1.17
DEU	-47811	107911	4483	64583	1.4	4.623	.71
FRA	-21035	-29634	-4783	-55452	-2	2.71	.7
CAN	-4507	-42039	-6025	-52571	-4	1.312	.23
MEX	-2953	-33944	-8330	-45227	-3.3	1.376	.29

Elasticity parameter relevant for the Consumer Surplus calculation is $\eta = 3.77$.

Figure 10: The effects of TTIP

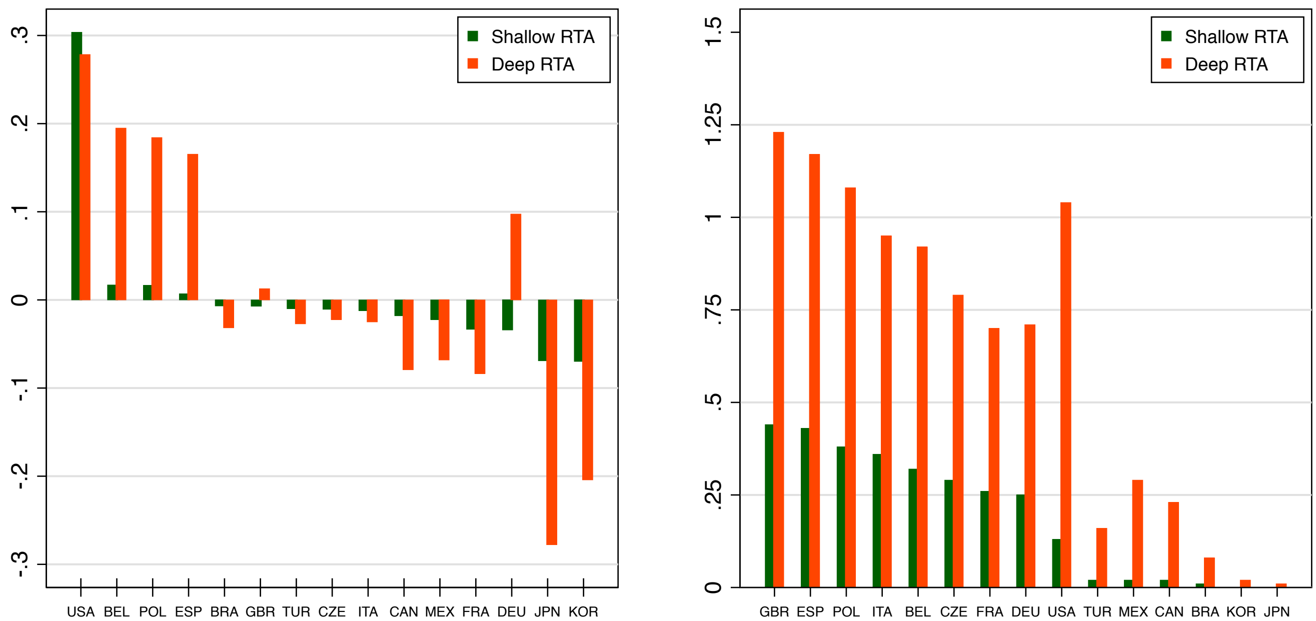

(a) Change in % share of world output

(b) % Change in consumer surplus

Let us now turn to another prospective agreement, TTIP, which would liberalize trade and MP between the US and the EU. The first scenario turns all tariffs between our sample's members of the EU and the US to zero. The deep integration scenario adds the effect of deep RTA dummies to those country pairs. Table 9 shows that Korea and Japan experience the greatest production losses, collectively 1.1% with a shallow RTA and 3.4% with a deep RTA. A preferential trade agreement between the EU and the US would lower Canadian and Mexican auto production by about 1% each. Canada's losses rise to 4% under deep integration. This shows that erosion of preferences can be a major concern and helps to explain why Canada negotiated its own integration agreement with the EU in 2014 (known as CETA, this pact has yet to be ratified). US losses from domestic sales (because of increased competition in the US market) are dwarfed by increased sales to EU countries in the shallow RTA scenario. This comes from the fact that the trade liberalization is asymmetric. In 2013, the average MFN tariff on passengers cars of the EU was 10%, while the US one was only 1.25%. The lower trade liberalization in the EU to US direction translates into smaller gains for the US consumer, compared to the Italian, French or German one. Note that in the second scenario (where liberalization is symmetric), the gains to the US consumer are an order of magnitude larger, due to cheaper access to EU-branded cars.

The cases of Belgium and Poland are quite interesting. With TTIP, Ford's Belgian factory approximately doubles its probability of being selected as the low-cost source for shipments to the US. Deep integration also doubles the probability of sourcing from Poland for Chrysler and Ford (relative to the status quo). In the case of France, with no US-owned plant, the lost market share in the EU clearly is much larger than the gains in the US market of the Toyota and Smart French plants.

6 Conclusion

This paper shows how the double CES structure of the modern multinational production framework can be exploited to achieve tractable estimation of all the structural parameters. Our methods are applicable to datasets where the researcher observes variety-level shipments to each market, tracking the production location and headquarters pertaining to each flow. A major contribution of this paper is to formulate the empirical predictions of the double CES model as four linear-in-parameters "workhorse" equations. Two of them describe choices at the extensive margins: whether to offer each variety in each market and from which factory to source each variety for the markets in which they are offered. The variety entry decision, which is new to our framework, proves to be the key channel through which deep integration agreements affect multinational sales. Two other decisions, the sales of each model and brand total sales, are intensive margins that resemble gravity equations. We deploy extremely detailed data from the car industry to estimate all the decisions just as the model dictates. To keep the scope of this paper finite, we have not attempted to estimate the firm-level decisions to establish assembly or distribution facilities in a country, nor have we considered the question of how many varieties to offer in all. These and other aspects of multinational expansion strategies provide a full agenda for future research.

A clear takeaway from our results is that multinational firms operate on a regional basis. One reason is that regional agreements eliminate MFN tariffs, which we show to generate large responses by multinational firms: 8.05 for substitution between assembly sites (θ) and 3.77 for substitution between varieties (η). Going beyond tariff reductions, the *deep integration* aspects of regional agreements appear to lower frictions in three distinct dimensions. They have a tariff equivalent of 3% on the assembly-to-market path and 5% on the headquarters-to-assembly path. On the HQ-to-market path, deep integration (in the form of harmonized standards and/or reduced impediments to trade in services) lowers variable marketing costs by 2% and fixed entry costs by an amount that would equate to a 16% tariff. These large estimates invite further research to elucidate the mechanisms that underlie multinational sales frictions.

Our counterfactuals provide policy-relevant outcomes for production and consumer surplus for four controversial changes in the structure of regional agreements. For example, we predict a complete Brexit would cut British car output by 4% and car buyers' surplus by 6%. The simulations also exhibit the rich patterns of adjustment that characterize the double CES framework of MP. One illustration comes from comparing the boosts Korean and Japanese production receive from NAFTA dissolution. As the costs of serving the Canadian market from the US rise, Toyota's probability of sourcing from the US falls by 5.3 percentage points (from 10.9% to 5.6%). Toyota's Japanese plants are the biggest beneficiaries, gaining 2.6 points, but its Canadian plants grab 1.6%. Hyundai's US plant also sees a sharp fall in its sourcing probability (8 points). However, as Hyundai lacks a plant in Canada, it reallocates nearly the entire probability (6.8 points) to its Korean plants. A second illustration comes from the prediction that a US-Mexico tariff war would lead to a 41% collapse in Mexican production but have a negligible impact on US production. These different outcomes arise in part because of a natural asymmetry in the MP framework: US brands are assembled in Mexico but there are no Mexican brands assembled in the US. Given the importance our estimates attribute to parts imported from headquarters, Mexican assembly costs rise when the US and Mexico impose tariffs on each other, while US assembly costs are unaffected. These two examples drive home the point that a full understanding of national outcomes from policy changes requires consideration of the geographic structure of MNCs' production networks.

References

Adao, R., A. Costinot, and D. Donaldson (2015). Nonparametric counterfactual predictions in neoclassical models of international trade. Working Paper 21401, National Bureau of Economic Research.

Anderson, S., A. De Palma, and J. Thisse (1992). *Discrete choice theory of product differentiation*. MIT Press.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

Arkolakis, C., N. Ramondo, A. Rodríguez-Clare, and S. Yeaple (2013). Innovation and production in the global economy. Working Paper 18972, National Bureau of Economic Research.

Atkeson, A. and A. Burstein (2008). Pricing-to-market, trade costs, and international relative prices. *The American Economic Review 98*(5), 1998–2031.

Bas, M., T. Mayer, and M. Thoenig (2015). From micro to macro: demand, supply, and heterogeneity in the trade elasticity. Discussion Paper 10637, CEPR.

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.

Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2016). Taste heterogeneity, trade costs, and global market outcomes in the automobile industry. mimeo.

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.

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 Congres*, 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.

Fernandes, A., P. Klenow, S. Meleshchuk, M. D. Pierola, and U. A. Rodriguez-Clare (2015). The intensive margin in trade: Moving beyond pareto. mimeo.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the us automobile industry. *Econometrica*, 891–951.

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 (2014). Gravity equations: Workhorse, toolkit, and cookbook. In E. Helpman, G. Gopinath, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier.

Head, K., T. Mayer, and M. Thoenig (2014). Welfare and trade without pareto. *American Economic Review 104*(5), 310–16.

Irarrazabal, A., A. Moxnes, and L. D. Opromolla (2013). The margins of multinational production and the role of intrafirm trade. *Journal of Political Economy 121*(1), 74–126.

Javorcik, B. and S. Poelhekke (2016). Former foreign affiliates: cast out and outperformed? Forthcoming, *Journal of the European Economic Association*.

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.

Tintelnot, F. (2016). Global production with export platforms. Forthcoming, *Quarterly Journal of Economics*.

Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics*, 240–268.

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.[16]

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

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

order conditions are

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

Combining we have

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

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 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

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

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 $\mathbb{P}_{mn} \to 0$. 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 $u()$. Furthermore, *income does not affect the choice between models* but, depending on the form of $u()$, 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 $u()$. 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 CES counterfactuals approximate BLP

Here we investigate the implications of estimating CES-based equations when the true data generating process (DGP) is the random coefficients logit model with Bertrand-Nash oligopolistic multiproduct price setting. As this is a mouthful, we will refer to this DGP as "BLP." For a wide value of parameter settings, counterfactuals based on CES estimates yield results that give good guidance on the "true" quantitative response of the aggregate share of domestic firms. For some plausible settings, the declines in the domestic share associated with tariff reductions are accurate out to the third decimal point.

B.1 The random-coefficients multiproduct oligopoly model (BLP)

There are M models subscripted with m and N households subscripted with h. The (indirect) utility of household h is given by

$$U_{mh} = \beta_h x_m - \alpha_h p_m + \xi_m + \varepsilon_{mh},$$

with x_m being a characteristic of the model and p_m being its price. We think of x_m as an observed component of quality, whereas the unobserved component is captured in ξ_m. Assuming that the individual random term of households for specific models, ε_{mh}, is distributed Gumbel, the choice probability of household h for model m takes the usual logit form:

$$\mathbb{P}_{mh} = \frac{\exp(\beta_h x_m - \alpha_h p_m + \xi_m)}{\sum_i \exp(\beta_h x_i - \alpha_h p_i + \xi_i)}. \tag{17}$$

We specify β_h and α_h such that

$$\beta_h \sim \mathcal{N}(\mu_\beta, \sigma_\beta) \quad \text{and} \quad \ln \alpha_h \sim \mathcal{N}(\mu_\alpha, \sigma_\alpha).$$

On the supply side, the model-specific primitives, x_m and ξ_m, are also assumed to be normally distributed. Firms are allowed to own several models, and unobserved quality ξ has both a model-level and a firm-level component (common to all varieties manufactured by firm f). For marginal costs, we follow the literature in assuming $\ln c_m = \gamma_0 + \gamma_1 x_m + \gamma_2 \xi_m + \nu_m$. We set ν_m to be normal as well since the linear combination of the three normal shocks will be normal and therefore marginal costs will be log-normal.[17]

With all individuals buying one car, the market share of model m simply averages individual probabilities from equation (17) over the N consumers:

$$s_m = \frac{\sum_h \mathbb{P}_{mh}}{N} = \frac{1}{N} \sum_h \frac{\exp(\beta_h x_m - \alpha_h p_m + \xi_m)}{\sum_i \exp(\beta_h x_i - \alpha_h p_i + \xi_i)}. \tag{18}$$

Profits for firm f are given by model-level profits $\pi_m = N s_m (p_m - c_m)$ for the models owned by the firm. Define the symmetric M-by-M co-ownership matrix as $\mathbf{\Omega}_{jm} = 1$ if models j and m have a common owner and zero otherwise.

The FOC for model m's price in the Bertrand-Nash equilibrium is

$$s_m + (p_m - c_m) \frac{\partial s_m}{\partial p_m} = - \sum_{j \neq m} \Omega_{jm} (p_j - c_j) \frac{\partial s_j}{\partial p_m}.$$

Isolating p_m on the left-hand side, while noting p_m is implicit in the s_m on the right-hand side, we have

$$p_m = c_m - \frac{s_m + \sum_{j \neq m} \Omega_{jm} (p_j - c_j) \frac{\partial s_j}{\partial p_m}}{\frac{\partial s_m}{\partial p_m}}$$

Since matrix formulations can drastically improve computation time, we express this equation in terms of the "optimal" response vector.

$$\mathbf{p}^\star = \mathbf{c} + \frac{\mathbf{s} + \mathbf{r}}{\mathbf{d}}, \tag{19}$$

where $\mathbf{r}$ is the effect on profits earned by the rest (r is a mnemonic for rest) of the firm's models caused by raising model m's price and $\mathbf{d}$ is minus the derivative of market share

[17] This distribution is a natural choice since it ensures positive costs, and under CES demand would lead to log-normal sales distributions, a dominant feature of the micro-level data in many countries (see Head et al. (2014) or Fernandes et al. (2015) for instance).

with respect to own price (elements of $\mathbf{d}$ are given by $-\frac{\partial s_m}{\partial p_m}$). Define an M by M matrix of cross-price derivatives on market share as $\mathbf{D}$. Elements of $\mathbf{D}$ are given by $D_{jm} = \Omega_{jm}\frac{\partial s_j}{\partial p_m}$ for $j \neq m$ and zero otherwise. With the aid of $\mathbf{D}$ we can express the cross-variety profit impact of price rises compactly as

$$\mathbf{r} = \boldsymbol{D}(\mathbf{p} - \mathbf{c}).$$

Numerical experimentation indicates that this system is not a contraction mapping but it can be solved via fixed point iteration on the equation

$$\mathbf{p}^{i+1} = \omega\mathbf{p}^{\star} + (1-\omega)\mathbf{p}^{i},$$

where ω is the weight accorded to the new best response and $1-\omega$ is the weight on the previous vector of iterations.

B.2 The CES approximation

We take logs of the market shares generated in equation (18) and use it as the dependent variable in the following regression equation:

$$\ln s_m = k + \beta x_m - \eta \ln c_m + \varphi_f + \epsilon_m, \tag{20}$$

where φ_f are firm-level fixed effects, which account for differences in mean unobserved quality (ξ_m in the BLP model) of each model across manufacturer/owner firm f. This estimation provides an estimate of the cost elasticity, denoted $\hat{\eta}$, which is also the price elasticity when CES demand is combined with monopolistic competition. Thus, the CES-MC model predicts prices

$$\hat{p}_m = \frac{\hat{\eta}}{\hat{\eta}-1}c_m.$$

Using this price and regression coefficients yields CES-predicted market shares, to be compared with the BLP-based market shares from equation (18):

$$\hat{s}_m = \frac{\exp(\hat{\beta}x_m + \hat{\varphi}_f)\hat{p}_m^{-\hat{\eta}}}{\sum_i \exp(\hat{\beta}x_i + \hat{\varphi}_f)\hat{p}_i^{-\hat{\eta}}}. \tag{21}$$

There are two fundamental differences between "true" s_m and CES-MC predicted $\hat{s}_m$. First, the two functional forms are radically different with s_m constructed from an *inner* logit (which is not CES) and an *outer* summation over heterogeneous-coefficient consumers. Second, the prices determining market shares are different. The s_m are based on the equilibrium oligopoly prices that take into account cannibalization effects, whereas $\hat{s}_m$ predicts prices by applying a constant markup (based on the estimated cost elasticity) to costs. With multiproduct oligopolists, such markups would not be optimal even under CES demand.

The question we explore next is whether the $\hat{s}_m$, despite all their differences, might nevertheless succeed in capturing the cross-sectional variation in s_m and more crucially for our purposes, the changes in aggregate domestic market shares.

B.3 Montecarlo results for tariff reductions

Based on the model described in the previous section, we simulate a tariff reduction scenario using the following data generating process.

1. For each of 100 models, draw quality and marginal costs. For each of 1000 households draw preferences.

2. Allocate 50 models to each of two different countries and impose a 20% tariff (specified as a higher marginal cost of delivering the product) on foreign varieties.

3. Run the fixed point iteration to find the BLP equilibrium prices and market shares.

4. Run equation (20), i.e. a regression of log share on log cost (controlling for firm dummies and model quality x) to estimate $\hat{\eta}$, the constant elasticity of market share with respect to cost, which also determines the markup under CES-MC.

5. Substitute $\hat{\eta}$ into equation (21) to obtain CES analytical prices and market shares.

6. Reduce the tariff by 10%, and redo steps 3 and 5.[18]

After completing these steps, we calculate the difference between the BLP generated market shares and the CES predictions at both the model and country level. We average the results of 1000 replications.

This simulation is carried out under four different sets of BLP parameter values. In each setting we adjust the mean value of price responsiveness, μ_α, such that the average CES cost elasticity estimate ($-\hat{\eta}$) rounds to -5.0.[19] The rationale for fixing the same

[18]In calculating counterfactuals, we are therefore letting initial market shares in the CES approximation differ from the BLP data, rather than employing the "exact hat algebra" method to the initial "true" market shares, as detailed in Costinot and Rodriguez-Clare (2014).

[19]This corresponds to the average firm-level elasticity of trade with respect to tariffs found by Bas et al. (2015), pooling over 6-digit industries.

cost elasticity across settings is that we consider $\hat{\eta} = 5$ as data in the spirit of Arkolakis et al. (2012). Under all four settings, the average sensitivity of consumers to quality is normalized to $\mu_\beta = 1$.

The settings differ in the heterogeneity of the consumers' α and β parameters affecting sensitivity to price and to quality. The first setting imposes homogeneous coefficients in consumer indirect utility. With $\sigma_\alpha = \sigma_\beta = 0$, demand takes the logit form with an absolute price elasticity that is increasing in price. The second setting introduces heterogeneous valuations of quality, raising σ_β to one. This opens the possibility of richer substitution parameters than setting 1 because now high quality goods are closer substitutes for other high quality goods.

The two last settings allow for heterogeneous price sensitivity. Setting 3 calibrates σ_α to obtain the same pass-through rate as implied by CES-MC with $\eta = 5$, namely a derivative of price with respect to costs of $5/(5-1) = 1.25$.[20] Following the practice in exchange rate passthrough studies, we also show the elasticity of passthrough. When the rate is 1.25, the elasticity is close to one (the value expected for CES-MC). Setting 4 raises σ_α to 1.1, the average of the values implied by the CGLT estimates combined with Gini indices of income inequality for their nine markets.[21]

Our first implementation, shown as panel (a) of Table 10, retains a market structure approximating monopolistic competition. We also rule out unobserved quality differences, setting $\sigma_\xi = 0$. These restrictions allow us to focus on the effects of changing the functional form to random coefficients logit. We assume all 100 varieties are independently owned. This eliminate cannibalization effects. Without multiple models per firm, we cannot estimate firm fixed effects so we just include a dummy for domestic models. Setting 1 of panel (a) eliminates rich substitution, so the main departure from CES-MC is the logit shape of the demand form. The other departure is that firms are not atomistic. Column (5) of Table 10 shows that the collective share of the top 5 variety-firms is 27% rather than the 5% that would be expected with 100 symmetric firms. This does not seem to affect the results much as we see a passthrough rate of one, exactly what Anderson et al. (1992) obtain analytically for the case of symmetric firms.

The very strong correlations between market share levels (0.96) and changes (0.94) show that CES is capable of closely approximating logit market shares. High correlations at the model level are neither a sufficient nor a necessary condition for good fit of aggre-

[20] Passthrough rates are obtained numerically by averaging across all foreign firms the decrease in their price divided by the decrease in costs (from the tariff reduction).

[21] The variance of σ_α sums variance from income and other sources of heterogeneity, with $\sigma_\alpha^2 = a_2^2\sigma_y^2 + \sigma_a^2$. CGLT (2015 NBER WP) estimate $a_2 = -0.709$ and $\sigma_a = 1.003$. We calculate σ_y (the variance of log incomes) for each country using data on its Gini index (denoted G) and the formula $\sigma_y = \sqrt{2}\Phi^{-1}[(1+G)/2]$.

Table 10: Monte Carlo simulation of the BLP data generating process

	Setting				CR5	Passthru			Corr		Agg. %Δ		Agg. bias	
#	σ_α	σ_β	μ_α	$-\hat{\eta}$	(%)	rate	elas.	$\frac{\partial \ln s_m}{\partial \ln p_m}$	s_m	Δs_m	BLP	CES	Avg.	S.D.
	Panel (a): 100 firms with 1 model each ("monopolistic competition")													
1	0.0	0.0	0.00	-5.0	27	1.00	0.84	-6.5	0.96	0.94	-15.9	-16.6	-0.8	1.4
2	0.0	1.0	0.00	-5.0	27	0.99	0.83	-6.5	0.81	0.85	-15.7	-16.8	-1.1	1.9
3	0.6	1.0	-0.10	-5.0	32	1.25	0.98	-4.7	0.93	0.95	-16.9	-17.1	-0.1	1.2
4	1.1	1.0	-0.15	-5.0	40	1.67	1.15	-3.3	0.91	0.88	-17.7	-16.8	0.9	3.2
	Panel (b): 10 firms with 10 models each (multiproduct oligopoly)													
1	0.0	0.0	0.10	-5.0	71	0.96	0.81	-7.1	0.82	0.84	-15.8	-16.8	-1.0	2.0
2	0.0	1.0	0.11	-5.0	71	0.96	0.81	-7.2	0.78	0.81	-15.6	-16.9	-1.3	2.1
3	0.7	1.0	0.05	-5.0	74	1.25	0.98	-4.9	0.89	0.91	-17.1	-16.9	0.2	1.8
4	1.1	1.0	0.16	-5.0	78	1.52	1.09	-3.9	0.85	0.82	-18.2	-16.8	1.3	3.8

Note: CR5 is the 5-firm concentration ratio. "Corr" show Pearson correlations between levels and changes of BLP and CES market shares. "Agg. Δ" sums Δs_m (BLP) or $\Delta \hat{s}_m$ (CES) for all models produced in the home country. "Agg. Bias" shows the mean and standard deviations of the BLP aggregate change subtracted from the CES prediction.

gates. On the one hand, there may be outliers that drive a substantial macro response without substantially weakening the correlation, and on the other hand, offsetting deviations could mute the aggregate response. Summing changes at the national level, we find a less than one percentage point deviation between the BLP change and the change predicted by the CES model.

In setting 2 we see that the richer substitution patterns brought about by $\sigma_\beta = 1$ lower the model-level correlations and raise the aggregate bias. However the gap between a 15.7% change and 16.8% is still quite small—less than the standard deviation across replications. Setting 3 exhibits negligible bias and we suspect that this arises because both models have identical pass-through rates and nearly the same price-elasticities of demand (-5.0 in CES vs -4.7 in BLP). Setting 4 shows that even with the heterogeneity of price sensitivity increased to the level implied by the CGLT estimates, there is not much deterioration in the performance of the CES approximation. As in settings 1 and 2 the average deviation is about one percentage point in absolute value, although the sign of the bias is now reversed. The high σ_α leads to extraordinarily high passthrough rates—a \$1 increase in costs leads to a \$1.67 increase in prices—of tariffs into higher import prices. However, the restriction that cost elasticities of market share remain at -5 leads to a much lower value of μ_α which shrinks the absolute price elasticity to 3.3. Thus, foreign firms raise their price more, but lower substitutability between varieties dampens the response of domestic market shares.

Panel (a) shows that neither the logit form nor the rich substitution allowed by random coefficients prevents the CES approximation from matching counterfactuals. The next step, considered in panel (b), is to bring in oligopoly with multiple varieties sold by each firm. In this case we also follow the standard practice in BLP of allowing for unobserved quality ($\sigma_\xi > 0$).[22] While we are not trying to fit the simulations to an individual car market, we want the settings to correspond roughly to the number of firms and models observed in actual markets. Having 10 firms with 10 models each generates a realistic amount of oligopoly since this leads to concentration ratios for the top 5 firms (CR5) between 70 and 80%, in line with the CR5s reported by Coşar et al. (2016) for 9 markets in their Table 1. Those figures also match the ones we report in our main text Table 2, i.e. an average number of 10 models per firm, and and an interquartile range of 70 to 83 % for the CR5 over the 73 markets and 14 years of our sample.

Figure 11: Initial model market shares: CES vs BLP

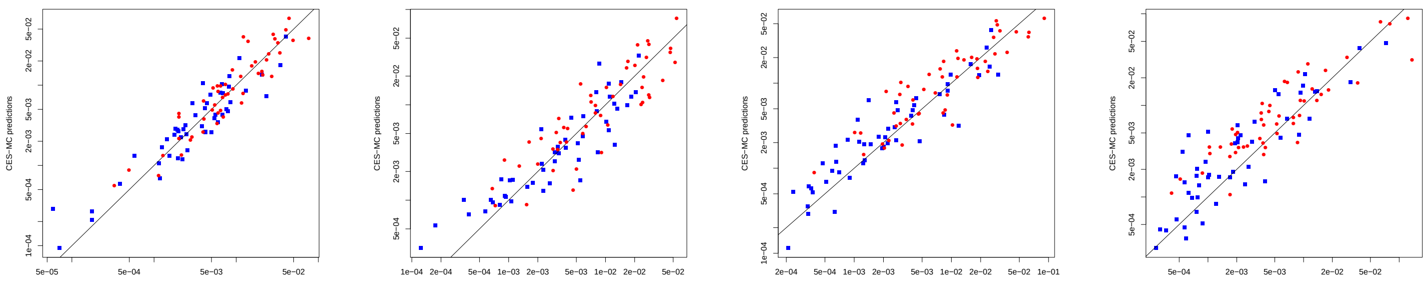

(a) Setting 1 (b) Setting 2 (c) Setting 3 (d) Setting 4

Setting 1 now differs from CES-MC in three new dimensions that were excluded from this setting in panel (a): oligopoly, cannibalization, and unobserved quality. As results of these changes, correlations at the variety level fall (from 0.96 to 0.82 in levels and from 0.94 to 0.84 in changes). Figures 11(a) and 12(a) display the performance of the CES prediction at the micro level, with the BLP-generated data on the horizontal axis and the CES prediction on the vertical axis. The foreign models are represented with blue squares and the domestic models with red dots. Both figures show that most models are close to the 45-degree line where CES predictions match the BLP generated data.

[22]Specifically, we add $\xi \sim \mathcal{N}(0, \sqrt{5})$ while lowering the standard deviation of observed quality to $\sigma_x = \sqrt{5}$ such that the combined quality variance remains one. We assume that half the variance in σ_ξ is firm level and half is model level. The CES estimation now includes firm dummies.

Figure 12: Changes in model-level market shares: CES vs BLP

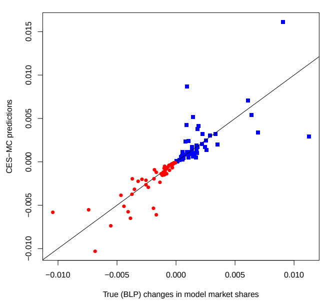

(a) Setting 1

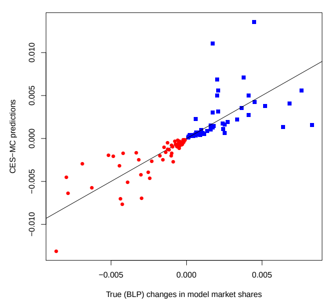

(b) Setting 2

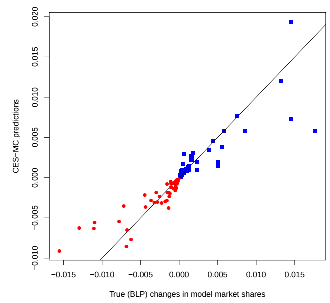

(c) Setting 3

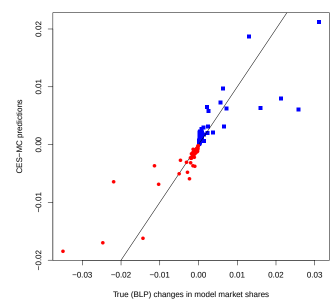

(d) Setting 4

Figure 13: Percent changes in domestic market share: CES vs BLP

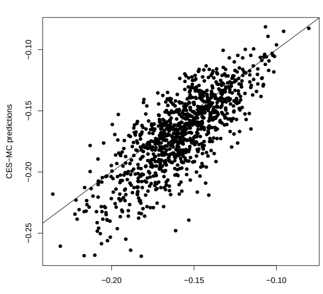

(a) Setting 1

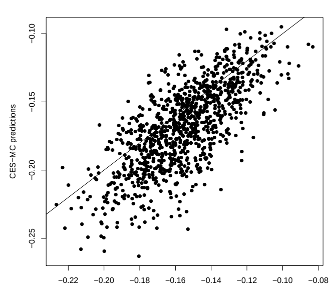

(b) Setting 2

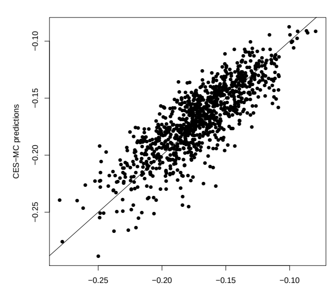

(c) Setting 3

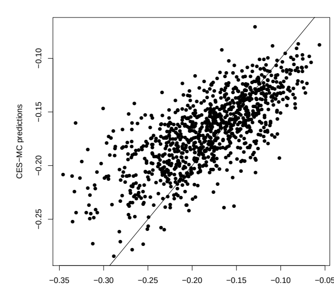

(d) Setting 4

Aggregating across models, we graph the change in the shares of all domestically produced models in response to the 10% tariff reduction in Figure 13. There is a small downward bias in the CES prediction but it is only slightly worse than what was obtained under monopolistic competition. Adding richer substitution patterns in settings 2 lowers correlations marginally and hardly changes the macro level fit. Setting 3, with passthrough rates corresponding to CES-MC, also continues to fit the BLP data very well. The prediction and the data both round to a 17% domestic output reduction. Figure 13(c) illustrates the remarkable predictive performance of the CES approximation across 1000 repetitions, with the points clustering close to the 45 degree line. This fit is especially striking in light of the individual models that have large deviations in Figure 12(c). This illustrates the idea that rich substitution can induce the CES approximation to fail badly for individual models, while nevertheless maintaining strong predictive power at the macro level.

Setting 4 of panel (b) shows the largest case of CES bias in Table 10. While the CES prediction does not change relative to panel (a), the BLP domestic output reduction rises by a half a percentage point. Figure 13(d) shows that in repetitions featuring large aggregate reductions under BLP, CES tends to under-predict the absolute change (the opposite of the pattern exhibited in figure 13(b)). The average bias, however, is about one third of the standard deviation across replications.

The CES-MC model performs very well in panel (b) despite the myopic pricing policy which omits adjustments to avoid cannibalization between the varieties that a firm offers. Equation (19) allows us to decompose the final price into three components: cost, a markup term given by $\mathbf{s}/\mathbf{d}$, and the cannibalization adjustment given by $\mathbf{r}/\mathbf{d}$. The simulation results show that on average just 1.6% to 1.8% of the price is attributable to the cannibalization adjustment. Under CES-MC, the markup share obtained by dividing $\mathbf{s}/\mathbf{d}$ by the price vector is just a constant given $1/\eta = 0.2$. In setting 3 the simulations show this markup term averages 0.2, with an interquartile range across models in a given replication of 0.19 to 0.21. This close similarity may account for why CES is so successful in this setting. In setting 4, where we observed deteriorating performance of CES, the markup term is 0.26 on average, with much more dispersion: the IQR is 0.22–0.30.

C Data Appendix

C.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 22 countries in the IHS data where assembly location was unavailable for all sales. We also required that at least 90% of the total car sales in a country must come from identified brands, leading us to drop Algeria and Cuba as well. The remaining 73 markets constituted 98% of world automotive sales in the 2013 IHS data.
- Norway is only an option for Think and in those cases it is the only option; therefore a NOR 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 (6) and (4) cannot identify its brand fixed effect.
- 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.
- 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).
- We drop 36 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 0.003%, with the maximum market share of 1.76% in China in 2013.

C.2 Other data sources

The time-invariant determinants of frictions (distance, home, contiguity, common language) come from the CEPII gravity database. 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. 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. 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).

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

D Estimates using the firm-variety approach

Table 12 reports results from the same set of specifications shown in Table 3 re-estimated using a different empirical mapping between firms and the varieties they produce. 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 (Polk) 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 2013.

The biggest problem with the Design Parents is that IHS only reports Design Parent (DP) as of 2013. 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 14: Market coverage by multi-variety firms

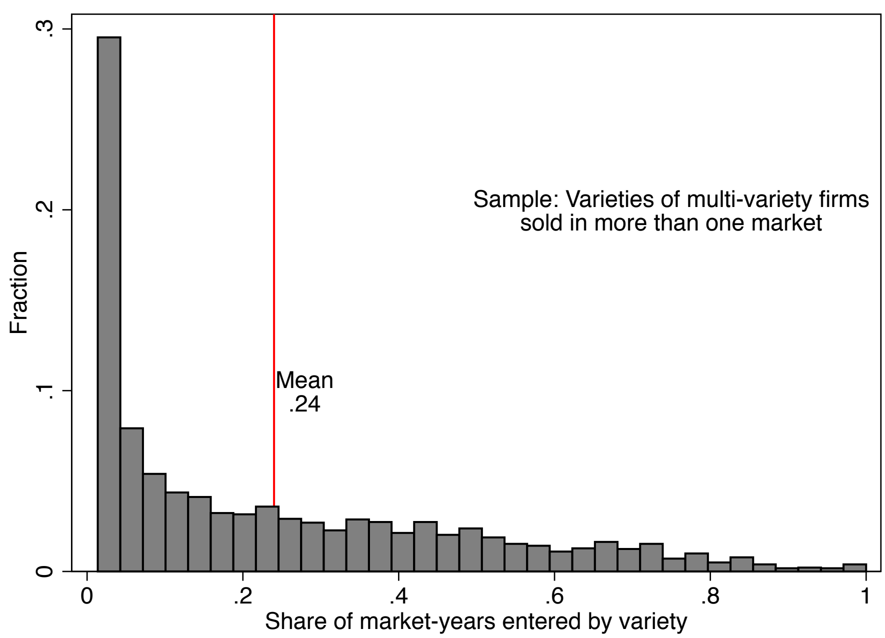

There were two main concerns about the brand-model approach. The first is that parent-firm headquarters might be making the critical management and parts supply decisions and that the brand headquarters might be less relevant from the point of view of $\gamma_{i\ell}$ (MP) frictions. Thus, SEAT assembly plants in Spain should perhaps be considered in a different country from their headquarters if the VW headquarters in Germany is supplying key parts and managerial oversight. The second 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, a nearly identical car is in fact available in many markets as the Acura RL.[23]

[23] Another form of re-badging holds the model name constant while changing the brand name. For exam-

Table 11: Statistical significance of friction categories in baseline results

	Regression specification from Table 3						
Equation:	sourcing	variety:		firm:	firm		entry
Method:	cond. logit	OLS	OLS	Poisson	Poisson	OLS	logit
Test Statistic (d.f.)	χ^2 (6)	F(6, 48)	F(6, 48)	χ^2 (6)	χ^2(6)	F(6, 72)	χ^2(6)
Trade costs: $\tau_{\ell n}$:	174.67	19.23	35.39	1198.61			
MP costs $\gamma_{i\ell}$:	81.05	2.88†	3.57†	106.21			
MS costs δ_{in}, ϕ_{in}:		4.23	4.88		36.78	10.25	124.10

†: All tests have p-values near zero except 2.88 (0.02) and 3.57 (0.01).

Table 12 reports results that alleviate both of these concerns. As detailed in Appendix D, we have reconstructed the data set using the parent company that designed the platform of the car in place of the brand and identifiers of the actual design in place of the model. We then re-estimated the baseline specification from 3. The sample size in the sourcing equation (column 1) doubles, reflecting a greater number of possible assembly locations when taking account of all the parent firms' production facilities. We are struck by the similarity in the estimates of the two key elasticities: $\hat{\theta}$ is 8.0 with brand/model and 8.3 with firm/variety whereas the averages of the specification (2) and (3) estimates $\hat{\eta}$ are 3.8 (brand/model) and 3.9 (firm/variety). The imprecision of the estimates of the $\gamma_{i\ell}$ determinants persists with the new set of headquarter i locations. Furthermore, the significant role of entry cost frictions (our δ_{in} and ϕ_{in}) comes through strongly even after redefining varieties in a way that rules out re-badging concerns. Only one coefficient out of seven loses statistical significance (and that one, contiguity, remains positive but one third smaller) in column (7) of Table 12.

E Internal consistency of friction estimates

As discussed in the main text, our regressions generate alternative estimates for the preferred parameters displayed in Table 5 and used in the counterfactuals. One approach would be to estimate a system in which a given parameter is constrained to take the same value in each equation where it appears. This would be more efficient if the equations are all correctly specified. However, our approach allows us to avoid mis-specification from one equation contaminating estimates from a correctly specified equation. We then can compare estimates from different equations to determine how much robustness there is in the estimates each equation offers.

ple 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).

Table 12: Results with the firm-variety approach

Dep. Var:	ℓ^*_{mnt}	$\ln q_{m\ell n}$		$\frac{q_{b\ell n}}{M_{bn}Q_n}$	$\frac{q_{bn}}{M_{bn}Q_n}$	$\ln\left(\frac{q_{bn}}{M_{bn}Q_n}\right)$	$\mathbb{I}_{mnt}$
Method:	cond. logit	OLS	OLS	Poisson	Poisson	OLS	logit
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Trade costs							
$\text{home}_{\ell n}$	0.967^a	1.033^a	1.295^a	1.828^a			
	(0.346)	(0.172)	(0.155)	(0.311)			
$\ln \text{dist}_{\ell n}$	-0.190^c	-0.312^a	-0.307^a	-0.752^a			
	(0.113)	(0.066)	(0.058)	(0.103)			
$\text{contig}_{\ell n}$	0.178	0.129	0.211^b	0.218^c			
	(0.118)	(0.100)	(0.093)	(0.125)			
$\text{language}_{\ell n}$	-0.136	0.141	0.134^b	0.036			
	(0.139)	(0.098)	(0.060)	(0.184)			
ln (1+ car $\text{tariff}_{\ell n}$)	-8.289^a	-2.921^b	-4.977^a	-9.468^a			
	(2.014)	(1.262)	(0.842)	(1.221)			
Deep $\text{RTA}_{\ell n}$	0.231	0.442^a	0.436^a	0.513^a			
	(0.171)	(0.092)	(0.091)	(0.133)			
MP frictions							
$\text{home}_{i\ell}$	1.852^a	-0.429	0.344	1.699^a			
	(0.610)	(0.513)	(0.245)	(0.482)			
$\ln \text{dist}_{i\ell}$	0.203	0.048	-0.070	0.164			
	(0.204)	(0.123)	(0.083)	(0.114)			
$\text{contig}_{i\ell}$	-0.008	-0.254	-0.319^c	0.097			
	(0.273)	(0.297)	(0.182)	(0.418)			
$\text{language}_{i\ell}$	0.232	-0.049	-0.284	-0.250			
	(0.324)	(0.367)	(0.186)	(0.265)			
ln (1+ parts $\text{tariff}_{i\ell}$)	-7.880^c	-1.156	-2.395^a	-7.418^a			
	(4.073)	(1.545)	(0.789)	(1.929)			
Deep $\text{RTA}_{i\ell}$	-0.259	-0.010	0.104	0.119			
	(0.312)	(0.208)	(0.084)	(0.252)			
MS frictions							
home_{in}		0.554^b	0.436^b		0.310	0.666	0.405^a
		(0.247)	(0.178)		(0.290)	(0.437)	(0.147)
$\text{home}_{in} \times \text{LDC}_n$		0.743	0.259		0.977^b	0.278	1.655^a
		(0.787)	(0.694)		(0.461)	(0.881)	(0.225)
$\ln \text{dist}_{in}$		-0.038	0.009		-0.446^a	-0.340^b	-0.080^a
		(0.072)	(0.061)		(0.113)	(0.144)	(0.031)
contig_{in}		0.063	0.069		-0.207	0.059	0.139
		(0.110)	(0.113)		(0.139)	(0.199)	(0.085)
language_{in}		0.196^a	0.249^a		0.036	0.190^c	-0.077
		(0.070)	(0.055)		(0.139)	(0.105)	(0.072)
Deep RTA_{in}		0.044	-0.026		0.007	0.050	0.100^a
		(0.137)	(0.111)		(0.125)	(0.184)	(0.035)
$\ln \hat{\mathbb{P}}_{b\ell n}$		0.026	-0.183^b				
		(0.122)	(0.069)				
$\ln \hat{D}_{bn}$					0.265^b	0.615^a	0.212^a
					(0.124)	(0.180)	(0.058)
Observations	4826626	265449	265449	252555	19031	19031	1050816
r2	0.437	0.439	0.642	0.119	0.753	0.636	0.102
S.E. cluster:	ℓ	ℓ, n	ℓ, n	ℓ	f	f, d	f

Standard errors in parentheses. Significance: ^{c} $p < 0.1$, ^{b} $p < 0.05$, ^{a} $p < 0.01$. r^2 is squared correlation of fitted and true dependent variables except in specifications (1) and (7) where pseudo-r^2 is reported.

Figure 15: Friction parameters across equations

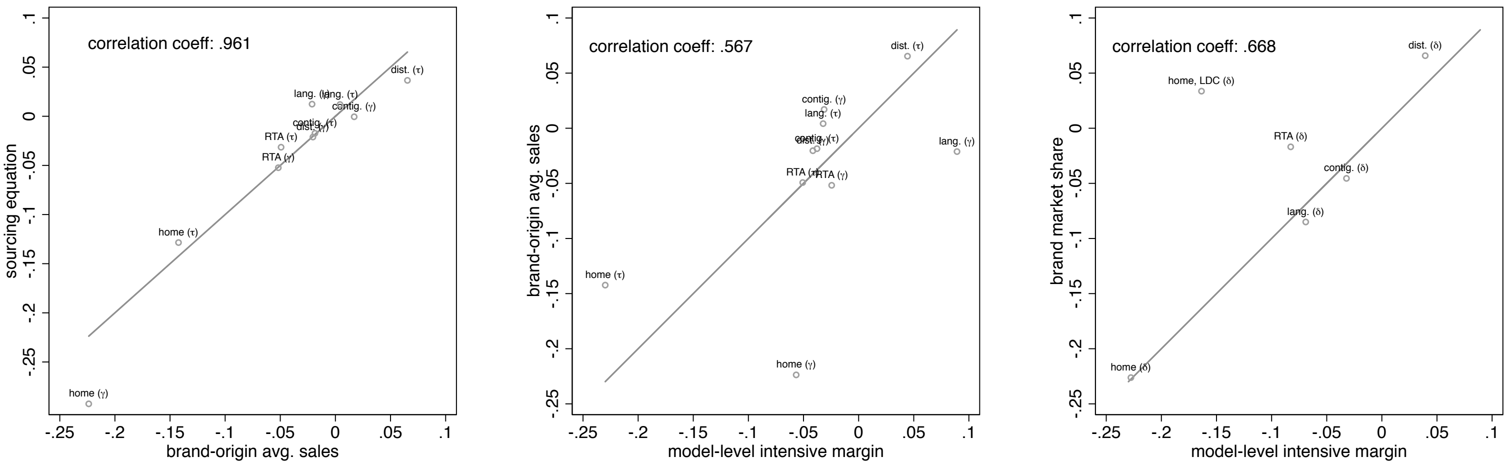

(a) τ and γ

(b) τ and γ

(c) δ

Figure 15 shows that the alternative estimates of structural parameters are mainly very similar to each other. Parameters τ and γ can be obtained from both the sourcing and the brand-origin-destination equations, columns (1) and (4) of Table 3. Panel (a) shows an impressive 0.96 correlation . A comparable consistency is displayed in panel (b), for the same parameters estimated with sourcing and model-level sales (obtained by averaging columns (2)and (3) of Table 3). Panel (c) compares estimates for δ frictions obtained from the model entry (column 7) and brand-destination sales equations (averaged over columns 5 and 6). Again, we find a strong correlation. This congruence of structural parameters estimated from quite different firm-level decisions and econometric models gives an added degree of confidence in the robustness of our structural parameters estimates.