

CEP Discussion Paper No 1551

June 2018

**Brands in Motion: How Frictions Shape
Multinational Production**

**Keith Head
Thierry Mayer**

Abstract

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

Key words: trade agreements, Brexit, tariffs, non-tariff barriers

This paper was produced as part of the Centre's Trade Programme. The Centre for Economic Performance is financed by the Economic and Social Research Council.

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

Keith Head, University of British Columbia, CEPR and visitor (2015) at Centre for Economic Performance, London School of Economics. Thierry Mayer, Sciences Po, Banque de France, CEPII and CEPR.

Published by
Centre for Economic Performance
London School of Economics and Political Science
Houghton Street
London WC2A 2AE

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means without the prior permission in writing of the publisher nor be issued to the public or circulated in any form other than that in which it is published.

Requests for permission to reproduce any article or part of the Working Paper should be sent to the editor at the above address.

1 Introduction

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

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

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

its most attractive features being tractable estimation, good fit to the data, and the ability to conduct counterfactuals with minimal data requirements.¹ Our implementation of the double CES framework here maintains those three desirable features.

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

1. Multinational firms are also multiproduct firms who offer different sets of varieties in each market. In the car industry, the typical brand makes ten models, of which they only offer one quarter in a typical market. They frequently opt out of serving markets at all (the typical brand serves only a third of the market-year combinations). These facts motivate the inclusion of two extensive margins of adjustment to policy shocks: which markets to serve and what subset of varieties to offer in each.²
2. While frictions between production and consumption locations are standard in pure trade models, the possibility that the firm is headquartered in a third country calls for the consideration of two additional frictions. Costs of producing far from headquarters are a well-established feature of the MP framework (going back to Ramondo (2014)). An innovation of our paper is to incorporate a third friction between headquarters and the market. We interpret this third friction as marketing costs, with both variable and fixed components. The new friction will play a crucial role in explaining our first extension of the MP framework, the market entry margins. A key motivation for incorporating the new frictions is that modern “deep” integration agreements contain whole chapters that do not operate on the origin-destination path traversed by goods. Rather, topics such as harmonization of standards, protection of investments, and facilitation of temporary movement of professionals, mainly affect the flows of *headquarters* services to production and distribution affiliates. Because our data tracks the three countries where a brand is headquartered, produces and sells its products, we are able to econometrically identify the new frictions separately from traditional trade costs.
3. Inspired by the observed prominence of spatial clustering in the auto industry and by recent theoretical advances,³ we allow for external returns to scale in the double-CES MP framework. An important benefit from this extension is to allow for interdependencies in how production and consumer surplus are impacted by shocks to frictions. A central element for tractable estimation of sourcing decisions both in our setup and in the versions of the double CES MP model developed by Tintelnot (2017) and Antras et al. (2017) is the CES functional

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

²Prior work on multinational production omits the two market entry margins, as Arkolakis et al. (2013) assume single-product firms, Antras et al. (2017) consider a single market, and Tintelnot (2017) assumes that firms offer a unit mass of varieties in every market. The variety-entry extensive margin for multiple-product firms was incorporated in a pure trade model by Bernard et al. (2011).

³See Grossman and Rossi-Hansberg (2010) and Kucheryavyi et al. (2016).

form, which relies on independent cost-minimization, conditional on the choice set. Tintelnot (2017) and Antras et al. (2017) build in interdependencies through the mechanism of firms paying a fixed cost to add countries to the choice set. This approach would not be computationally feasible in our context because of the above-mentioned extensive margins that are essential features of the industry.⁴ We instead introduce interdependencies in our setup via the channel of external returns to scale.

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

The chief concern about CES for the purposes of this paper is that it might exaggerate the degree of substitution between models in very different segments of the car market. This could lead to erroneously large responses to trade policy changes (e.g. a Brexit-induced tariff on Polish-made Fiat 500s would be unlikely to trigger much substitution towards UK-made Land Rovers). We mitigate this concern, while maintaining all the computational advantages of CES, by also estimating and simulating a version of the demand side that nests varieties within market segments.⁶ This follows the research line of Goldberg (1995) and Verboven (1996), with two important modifications. As in Björnerstedt and Verboven (2016), substitution within each nest takes the CES form (albeit with quantity shares). Second, to characterize the maximal extent of divergence from symmetric CES, our formulation restricts all substitution to occur *within segments*. The unified and segmented markets versions of the model therefore bracket the extent of substitution between car models in different segments. The interval between these extreme approaches turns out to be

⁴The algorithm employed by Antras et al. (2017) to solve for the optimal choice set requires a super-modular objective function. Although Arkolakis and Eckert (2017) generalizes the algorithm to handle sub-modular problems, our objective function has both super and sub-modular regions.

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

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

fairly narrow in terms of the outcomes of our counterfactual scenarios.

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

The results from counterfactual trade policy changes improve our understanding of the impacts of the creation and dissolution of regional integration agreements. As one would expect in a pure trade model, tariff changes have third-country effects via the path of erosion of trade preferences. For example, the end of NAFTA would stimulate exports to North America from Japan, Korea and the EU. 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. Canada's integration with remaining members of the TPP, leads to a major boost in exports to the United States (by Japanese-owned plants) which is not predicted by conventional trade models. The presence of returns to scale in our model amplifies by 20–30% the reallocation of production following policy changes. In addition, IRS introduces market interdependencies: As an example, the simulation predicts Brexit to lower the shipments of UK-made Hondas to the US, despite the absence of any change in UK-US trade costs. Reductions in marketing costs can generate quite large reactions for the entry margins we introduced: for instance, Trans-Pacific integration excluding the US is expected to stimulate model entry by Japanese brands in Canada by around 22% and to reduce the probability that Chevrolet enters the Vietnamese market by 30 percentage points.

The paper continues in five main sections. We first discuss and display some of the important empirical features of the global car industry, using the nearly exhaustive firm-level information on where each variety is designed, assembled and sold. Drawing on these facts, the next section generalizes the existing models to include marketing frictions and market-entry decisions at the

model and brand level. We then show how the structural parameters of the MP model can be recovered from four estimating equations corresponding to four key decisions made by multinationals. Following estimation, we present the key methodological aspects of our counterfactual exercises. Finally, we use those methods to project the outcomes of NAFTA termination, soft and hard version of Brexit, and Trans-Pacific and Trans-Atlantic integration.

2 Data and model-relevant facts

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

The first and foremost advantage of the car industry is the extraordinary richness of the data compiled by IHS Markit.⁷ IHS uses new car registration information (and probably other sources of information) to obtain annual flows at the level of individual models identifying the assembly plant and country of sale. From it we extract origin-destination flows for 4791 car models sold by 138 brands over the 2000–2016 period.⁸

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

The empirical analysis in the main text maps the theoretical concept of varieties to models and the concept of firms to brands. Models appear to be the natural counterpart to the concept of varieties. As implied by our theory for individual varieties, we show that models sold in a particular market are almost always sourced from a single assembly location. There are several reasons we employ brands, rather than parent corporations, to correspond to the theoretical concept of the firm. First, the brand is the common identity across models that is promoted to buyers via advertising and dealership networks. This suggests that the brand’s home is the one relevant for marketing frictions. Second, most of the brands under common ownership were originally independent firms (e.g. Chevrolet and Opel (GM), Ferrari and Chrysler (Fiat), Volvo (Geely), Mini

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

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

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

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

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

(BMW)). Partly for historical reasons, brand headquarters often correspond to the location where models are designed. For example, while Jaguar is owned by Tata Motors, based in India, Jaguar's cars are designed at the brand's headquarters in Coventry in the UK. We think of the brand's headquarters as a principal source of tangible (e.g. engines) and intangible (e.g. designs, managerial oversight) inputs used by the assembly plants.


There are two potential sources of concern when using the brand/model concepts. The first is that headquarter inputs may originate mainly from a higher level than the brand headquarters. A second worry comes from the industry practice of re-badging: different brand/model combinations might cover what is essentially the same underlying car. The richness of the IHS data enables us to replicate all our analysis using an alternative approach that deals with those concerns. The alternative specifies varieties as particular car designs using the identifiers for the "platform" (the underbody of the car) to which we add the above-defined program and bodytype. The concept of firm is the "Design Parent", the corporation that has managerial control over the design of the platform used by each variety. We discuss the results from implementing this approach, shown in full in Appendix E, 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 52 countries by brands headquartered in 21 countries and sold in 76 national markets.

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

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

Figure 1: Example: Fiat 500 & Punto production organization



2.1 Fact 1: Almost all models are single-sourced

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

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

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

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

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

2.2 Fact 2: Most markets are not highly concentrated

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

One defense of the use of monopolistic competition is that, in some respects, the industry is not as concentrated as one might imagine. Table 2 shows that most markets feature many competitors. For three quarters of market-years we consider, more than 191 models are available. Even at the highest level of aggregation, the sales parent,¹² the three quarters of the markets have at least 18 competing firms. Column (2) shows that median market shares are small (mainly under 5%), implying that oligopoly markups for the majority of firms would be close to those implied by monopolistic competition. Column (3) shows the concentration ratio for the top five actors at each level of aggregation. In three quarters of the market-years, the top 5 brands account for less than 74% of the market. The last three columns show that at the highest levels of ownership (parent), EU merger guidelines would be “unlikely to identify horizontal competition concerns” for 72% of the market-year combinations.¹³ Even within segments, the majority of markets are moderately concentrated except in the case of MPVs and sport and luxury cars.

To be clear, we are not arguing that oligopoly is irrelevant in the industry. The largest firms are big enough to have endogenous markups that significantly exceed those implied by monopolistic competition. Nevertheless, even under a data generating processes that matches the level of concentration observed in the industry, an estimated CES monopolistic competition model can deliver surprisingly accurate predictions for trade policy counterfactuals. Head and Mayer (2018) simulate data from a BLP framework featuring oligopoly, rich substitution in demand, and multi-product firms that internalize cannibalization effects. The CES-MC model is capable of closely approximating the aggregated counterfactuals for BLP-generated data under settings that replicate data moments for parent firms in Table 2 (5-firm concentration ratios of 70–80% and an average

¹²“Sales Parent” is defined by IHS as “The company who owns the brand at the current point in time.”. For example, Volkswagen is the Sales Parent of Audi, Bentley, Bugatti, Lamborghini, SEAT, Skoda and Volkswagen. There is a many-to-one mapping between brands and their sales parent.

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

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

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

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

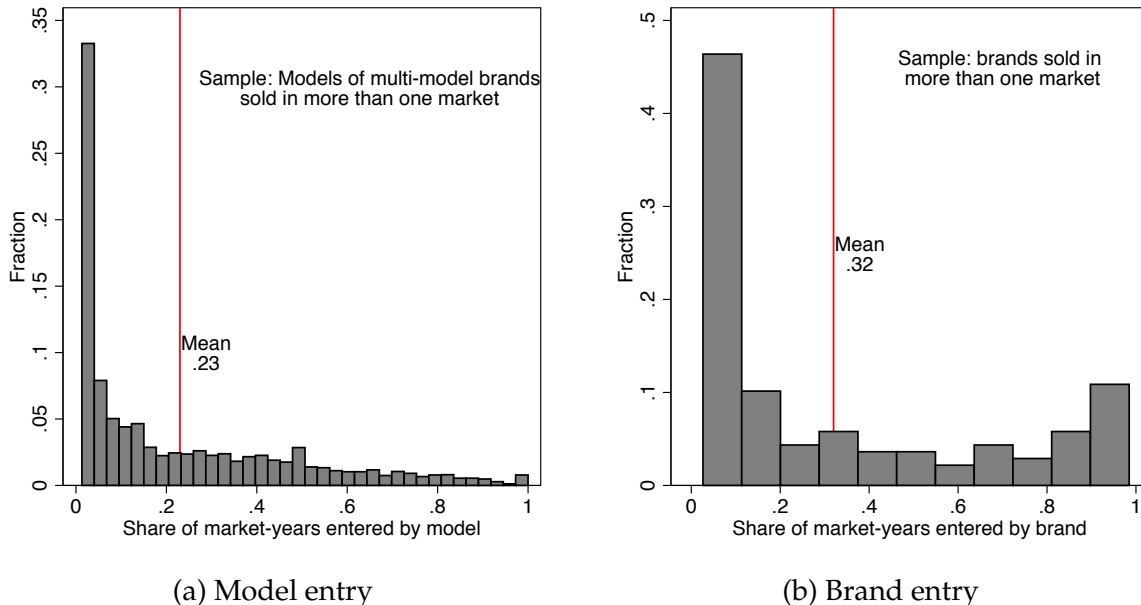
of 10 models per firm¹⁴). There is no theorem guaranteeing the close fit we have found in these simulations generalizes to all situations. However, the simulations establish that the mere fact that CES-MC omits many theoretically desirable features does not systematically prevent it from being a useful tool for counterfactual policy exercises. The success of the CES-MC framework in these simulations reinforces its appeal for our purposes, given its tractability, low data requirements, and connection to the gravity equation.

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

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

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

Figure 2: Entry rates for models and brands



hatchback version of C1-J68C has an 78% entry rate. However, looking across all varieties the average entry rate is just 24%, slightly larger than the average across all models. Figure E.1 shows the whole distribution of entry rates is visually unchanged after removing the re-badging issue. The takeaway is that whether we define varieties as the consumer sees them or based on firm-level design distinctions, they tend to be offered in about one quarter of the places where they might be offered.

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

3 The double CES model of multinational production

There are B brands, M symmetrically differentiated car models, and N countries. Brand b is endowed with a headquarter country $i(b)$; a location-independent productivity (φ_b); a portfolio of $M_b = \sum_m \mathbb{M}_{bm}$ models, where $\mathbb{M}_{bm} = 1$ if b owns m ; and a set of production facilities, with $\mathbb{L}_{bl} = 1$ for countries where b can manufacture any of its models. We see the decision to take location choices as exogenous as appropriate for the analysis of the medium-run consequences of policy changes because of the strong persistence that we observe in the set of locations where each brand operates: three quarters of all cars produced in 2016 (the year for which we run counterfactuals) were assembled in brand-country combinations that were already active in 2000. This fraction

risers to 88% for OECD countries and to 94% if we draw the line in 2007, a decade before our counterfactuals.

The sequence of decisions is as follows.

1. **Brand entry:** Brands learn F_{bn}^d , the fixed costs of creating a dealership network and then decide where to establish distribution facilities, corresponding to $\mathbb{D}_{bn} = 1$.
2. **Model entry:** For markets the brand has entered, it then learns the model-entry fixed costs, F_{mn}^e and decides which models to offer in each market ($\mathbb{I}_{mn} = 1$).
3. **Sourcing:** Brands decide the source ℓ , that minimizes the delivered cost to market n for model m (denoted $c_{m\ell n}$), subject to $\mathbb{L}_{b\ell} = 1$. Selected sources have $\mathbb{S}_{m\ell n} = 1$.
4. **Market shares:** Firms set p_{mn} and read q_{mn} off their demand curves.

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

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

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

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

3.1 Consumer preferences and demand

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

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

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

The monopolistic competition assumption implies the delivered price of model m in n is a constant markup, $\frac{\eta}{\eta-1}$ of marginal cost. Substituting price into the demand curve, sales are

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

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

3.2 Costs (including frictions)

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

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

where $z_{m\ell}$ is a productivity shock distributed Fréchet with shape parameter θ . In the numerator, w_ℓ and w_i are the costs of a composite factor (including efficiency-adjusted labor, capital and locally-sourced intermediates) in the assembly and headquarter countries. Parameter $1 - \alpha$ denotes the cost share of inputs obtained from the headquarters. Frictions applicable to inputs sourced from the headquarter country are captured in $\tau_{i\ell}^H \geq 1$. We use $\gamma_{i\ell} \equiv (\tau_{i\ell}^H)^{1-\alpha}$ as the notation for costs involved in separating assembly from headquarters to emphasize the similarity to the corresponding friction in Arkolakis et al. (2013). The difference is mainly one of interpretation, with our γ reflecting input costs from headquarters and theirs being a penalty in terms of lost productivity associated with transfer of operational methods from HQ to assembly country. Relevant for counterfactual policy changes, our $\gamma_{i\ell}$ is a function of tariffs on key inputs (engines, transmissions) obtained from the HQ country. Finally, ς is the elasticity of costs to the scale of production in the assembly country. External economies of scale correspond to $\varsigma < 0$.

Parts are important in the auto industry. We do not incorporate them explicitly in the cost function, due to lack of appropriate data. However equation (4) implicitly features three important roles for parts. First, w_ℓ should be thought of as capturing not only worker wages and productivity but also the price and variety of parts available in country ℓ . Second, tariffs paid on parts imported from the brand's headquarters are included in the friction $\gamma_{i\ell}$.¹⁵ Finally, the Marshallian

¹⁵ The notion that plants rely heavily on inputs from their HQ country is consistent with de Gortari's (2017) observation that exports from German-owned plants in Mexico contain much higher German content than US-owned plants

mechanism of downstream production attracting a denser network of parts suppliers is a likely explanation of the external returns to scale captured by q_ℓ^ς .

The high level of spatial clustering of the auto industries that we observe in many countries provides supportive evidence for the empirical relevance of such Marshallian economies.¹⁶ Allowing for external IRS is also important in our analysis because it generates interdependencies in how country-level production reacts to changes in frictions: With $\varsigma < 0$, the decision to source cars to be sold in one market from a particular country boosts the incentive to use that same assembly site to serve other markets.¹⁷

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

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

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

3.3 Sourcing decision

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

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

exports to the US.

¹⁶Spatial concentration of car production has been an important feature of the industry since its founding, as seen in the production clusters around Detroit and Paris. More recently, the examples of plant agglomerations in cities of Slovakia (4 plants), Central Mexico (11 plants), Northern France (6 plants) and I-75 corridor in the USA (about 10 plants) point to the persistent importance of Marshallian economies. Smith and Florida (1994), Klier and McMillen (2008) and Schmitt and Biesebroeck (2013) are examples of papers that find support for the empirical relevance of Marshallian economies in the car industry using firm-level location choices.

¹⁷Antras et al. (2017) emphasize the empirical importance of interdependencies in their global sourcing model. They generate those interdependencies through the use of source country-specific fixed costs.

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

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

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

θ is the second CES in this framework, playing the same role as in Eaton and Kortum (2002). C_{bn} is the multinational production cost index, summarizing the firm's costs of serving market n . Versions of this equation appear in Arkolakis et al. (2013) as equation (6), Tintelnot (2017) as equation (9), and Antras et al. (2017) equation (7).¹⁹ We are the first to estimate this equation directly with both $\tau_{\ell n}$ and γ_{in} frictions, because such estimation requires variety-level data on sourcing for multiple markets and for firms with many different headquarters.

3.4 From model-level to brand-level sales

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

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

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

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

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

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

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

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

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

¹⁹Like Tintelnot (2017), we assume independent productivity shocks whereas the Arkolakis et al. (2013) formulation allows for them to be correlated.

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

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

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

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

This equation for expected average market shares is linear in the determinants of $\ln \delta_{in}$. Note that the γ or τ frictions are included in the multinational production cost index C_{bn} . The coefficient on $\ln C_{bn}$ identifies the demand elasticity η .

3.5 Model-market entry decision

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

The probability that entry occurs, $\mathbb{I}_{mn} = 1$, is the probability that model-level expected profits net of fixed costs are positive:²¹

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

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

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

The brand foresees that it will choose the optimal assembly location after learning the realizations

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

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

of the model-location productivity shocks, $z_{m\ell}$. Applying the moment generating equation from Hanemann (1984),

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

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

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

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

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

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

3.6 Brand entry (distribution networks)

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

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

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

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

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

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

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

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

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

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

4 Empirical implementation

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

4.1 Friction determinants

We start by specifying the empirical content of frictions. The frictions governing trade costs (τ), HQ input transfer costs (γ) variable marketing costs (δ), and fixed model-entry (δ^e) and brand-entry (δ^d) costs, are exponential functions of the observable determinants (some of which vary over time, hence the new subscript t) denoted $\mathbf{X}_{\ell nt}$, \mathbf{X}_{ilt} and \mathbf{X}_{int} :

$$\begin{aligned} \tau_{\ell nt} &= \exp(\mathbf{X}'_{\ell nt} \boldsymbol{\rho}), & \gamma_{ilt} &= \exp(\mathbf{X}'_{ilt} \mathbf{g}), & \delta_{int} &= \exp(\mathbf{X}'_{int} \mathbf{d}), \\ \delta_{int}^e &= \exp(\mathbf{X}'_{int} \mathbf{f}^e), & \delta_{int}^d &= \exp(\mathbf{X}'_{int} \mathbf{f}^d), \end{aligned} \quad (20)$$

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

The \mathbf{X} vectors include the standard explanatory variables used in gravity equations: home, distance, and common language. These variables have already been shown to matter for trade flows and affiliate sales. The differences in subscripts are of critical importance to the estimation. Thus $\text{home}_{\ell n}$ indicates that the assembly plant is in the same country as where the car is bought, whereas $\text{home}_{i\ell}$ equals one when the plant is located in the headquarters country, and finally

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

home_{in} turns on when consumer and brand share the same home country. Distance is the average number of kilometers on great-circle route between the main cities in the corresponding countries. Language indicates that the countries share an official language.

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

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

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

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

4.2 Estimating equations

We now express the four equations that identify the structural parameters in an estimable way in terms of observed variables with associated coefficients and fixed effects (denoted $\text{FE}^{(j)}$ where $j = 1, 2, 3, 4$).

(1) Sourcing

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

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

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

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

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

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

Our method estimates external economies of scale based on the magnitude of the revealed preferences of brands for assembly locations with high aggregate output ($q_{\ell t}$). As pointed out by Goldberg and Verboven (2001), unobserved factors can both make a location attractive and increase its aggregate production. This would lead to upward bias in our estimate of $-\theta \varsigma$. Our estimates should therefore be seen as an upper bound of the degree of external increasing returns. We are not aware of structural estimates of the external returns elasticity in the car industry with which to compare our estimates of ς . The literature does provide some estimates of *internal* returns to scale, estimated from data on prices or costs. These papers omit the multinational production dimension that causes the combinatorial computational challenge. Our implied $\varsigma = -0.035$ is within the interval of those estimates of internal returns to scale for the car industry. They are bigger than the range of values reported by Goldberg and Verboven (2001), -0.006 to -0.03 , but smaller than the -0.11 and -0.07 provided by Verboven (1996) and Fuss and Waverman (1990). Two simple and common ways to mitigate the endogeneity bias that can be applied to our case are: i) lagging $q_{\ell t}$, ii) constructing a Bartik-style prediction for $q_{\ell t}$ to be used in a control function approach. In unreported regressions, we followed the two approaches. The Bartik prediction applies changes in total demand Q_{nt} to brand-origin-destination market shares fixed at their 2002 levels. Lagging lowers the magnitude of increasing returns to $\varsigma = -0.032$, whereas the Bartik approach reduces it further to $\varsigma = -0.024$. Neither approach is perfect, but both values support our use of $\varsigma = -0.035$ as the upper bound of the parameter governing the strength of interdependencies.

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

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

(2) Brand-level market shares

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

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

where $\eta(1 - \alpha)(\ln w_{i(b)t} - \ln w_{i(b)T}) = \mathbf{W}'_{i(b)t} \mathbf{v}_2$ captures the evolution of HQ-related costs through changes in income per capita and GDP price. Notation $i(b)$ designates the HQ country of brand b and $T = 2016$. The structural interpretation of the fixed effects becomes $\text{FE}_b^{(2)} = \eta \ln \varphi_b - \eta(1 - \alpha) \ln w_{i(b)T}$ and $\text{FE}_{nt}^{(2)} = \ln \kappa_1 + \eta \ln P_{nt}$. The C_{bnt} included as the last control comes from the sourcing probability results from equation (21) where $C_{bnt} = \left(\sum_k \mathbb{L}_{bkt} (w_{kt}^\alpha \tau_{knt} \gamma_{ikt} q_k^\zeta) \right)^{-1/\theta}$. This regression allows us to estimate the δ_{int} determinants and provides our estimate of η . The natural way to estimate the moment condition shown in equation (22) is Poisson PML because it does not require an additional homoskedastic log-normality assumption for the error term.²⁷

(3) Model entry decision

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

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

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

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

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

(4) Brand entry decision

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

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

$\mathbf{W}'_{i(b)t}\mathbf{v}_4$:

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

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

4.3 Identification of structural parameters

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

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

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

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

Brand entry: Equation (24) estimates σ_d as the inverse of the coefficient on $\ln \mathbb{E}[\pi_{bnt}]$. Destination

and brand fixed effects yield estimates of $\mu_n^d + \ln w_{nt}^{1-\zeta}$ and $\beta_b^d + \ln w_{i(b)t}^\zeta$ respectively. Multiplying the coefficients on the friction determinants by σ_d yields the vector \mathbf{f}^d .

5 Results

5.1 Baseline estimates

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

Sourcing estimates

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

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

The estimates of the γ frictions are much less precise, with standard errors several times those estimated for trade frictions. Two of the effects, distance and language, do not even enter with the expected sign, although neither is significantly different from zero. The significant effect is that assembly locations in the brand's home country are $\exp(2.248) \approx 9.5$ times more likely to be selected. The elasticity on the car parts tariff can be used to infer the share of assembly costs attributable to components from the headquarters country, $(1 - \alpha)$ in the cost equation, which is about 37% ($2.87/7.7$). While the precise value of this ratio should be taken with caution, we now have direct evidence of the importance of intermediate inputs from the headquarters country. This feature of the MP model has major qualitative and quantitative implications for the impact of trade liberalization, as we shall see in the counterfactuals. Deep RTAs between assembly and headquarter countries are estimated to have a larger effect on sourcing than deep RTAs between assembly and consumer countries, but the standard error is also larger.²⁹

Brand-level market share estimates

²⁸The estimate of θ when using all the locations of the parent firm as options for sourcing also rounds to 7.7 as can be seen in Table E.1. The firm-variety approach also shows similar coefficients for the other determinants of trade costs.

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

Table 3: Baseline results

Decision:	Sourcing	Market share	Model entry	Brand entry
Dep. Var:	$S_{b\ell nt}$	$\frac{q_{bnt}}{M_{bnt}Q_{nt}}$	$\frac{M_{bnt}}{M_{bt}}$	\mathbb{D}_{bnt}
Method:	PPML	PPML	frac. probit	probit
	(1)	(2)	(3)	(4)
home $_{\ell n}$	0.973 (0.318)			
ln dist $_{\ell n}$	-0.323 (0.078)			
language $_{\ell n}$	-0.042 (0.114)			
ln (1+ car tariff $_{\ell n}$)	-7.696 (0.893)			
Deep RTA $_{\ell n}$	0.246 (0.137)			
home $_{i\ell}$	2.248 (0.616)			
ln dist $_{i\ell}$	0.166 (0.117)			
language $_{i\ell}$	-0.218 (0.366)			
ln (1+ parts tariff $_{i\ell}$)	-2.872 (3.118)			
Deep RTA $_{i\ell}$	0.495 (0.445)			
ln q_{ℓ}	0.270 (0.071)			
home $_{in}$		0.816 (0.186)	0.260 (0.065)	0.718 (0.411)
home $_{in} \times \text{LDC}_n$		-0.028 (0.238)	0.829 (0.113)	1.328 (0.498)
ln dist $_{in}$		-0.339 (0.094)	-0.059 (0.016)	0.009 (0.064)
language $_{in}$		0.289 (0.130)	0.068 (0.034)	0.017 (0.123)
Deep RTA $_{in}$		-0.040 (0.116)	0.121 (0.042)	0.165 (0.098)
ln C_{bn}		-3.874 (0.699)	-0.512 (0.211)	
ln $\mathbb{E}[\pi_{bn}]$				0.595 (0.060)
Observations	347542	46299	46300	128589
R^2	0.824	0.601	0.715	0.618
S.E. cluster:	ℓ	b	b	b

Standard errors clustered at the level of countries (ℓ) or brands (b). R^2 is squared correlation of fitted and true dependent variables except in specification (4) where the pseudo- R^2 is reported. Each regression controls for log per-capita income and price level of the assembly country.

Determinants of a brand's market share are estimated in Column (2) of Table 3. The estimate of $\eta = 3.87$ (from the coefficient on $\ln C_{bn}$) is substantially smaller than the θ obtained in the sourcing decision. It implies that there is considerably more heterogeneity in consumer evaluations of brands than in car maker evaluations of assembly locations. An own-price elasticity of 3.87 implies a markup of 35% under monopolistic competition. 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. (2018) report in Table 11 average firm-market markups ranging from 6% (Peugeot in Brazil) to 12.4% (Peugeot in France). Our implied markup lies in the upper region of the highly dispersed set of results found in this literature.³⁰

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

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

Model entry estimates

Column (3) of Table 3 shows that all the marketing cost determinants have the expected signs and are highly significant. More models are offered in the home country of the brand, especially when this country is a developing one. Spatial proximity promote entry as well. Deep RTAs between the headquarter country (i) and the market (n) increase the fraction of models offered by 14% (calculated as the average semi-elasticity). As it seems unlikely that RTAs change preferences, we see the deep RTA_{int} effects as supporting the cost-shifter interpretation. Under this approach, our δ^e frictions include various types of marketing efforts, in particular managing dealership networks. This may be facilitated by the freer movement of skilled workers that is a commonly included provision of RTAs (e.g. NAFTA, EU). The RTA_{int} effect may also capture the greater ease of compliance with regulatory standards if the head office lies within the region and is therefore more able to exert influence on specific requirements in harmonized rules. Note also that the significance of this fixed cost dimension of RTA effects contrasts with the weak impact of the same variable on brand-level sales. This suggests that deep RTAs reduce the fixed costs of model entry

³⁰The firm-variety estimate of $\eta = 1.5$ (standard error of 1) shown in Appendix E implies markups of 200%, well outside this range. We attribute this low, noisy estimate to measurement error in C_{bn} caused by the firm-variety approach.

between HQ and destination (δ_{int}^e), rather than the variable marketing costs (δ_{int}), that affect brand sales as well.

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

Brand entry estimates

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

5.2 Interpreting the structural parameters

Table 4: Friction parameters

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

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

Using the procedures detailed in section 4.3, we now proceed to report and interpret the structural parameters underlying our estimates. The three sets of parameters relevant for variable costs—trade costs τ , multinational production costs γ and marketing costs δ —are reported in the first three columns of Table 4. For each, the coefficient on the k th element of \mathbf{X} maps to proportional increases in price (an *ad-valorem* equivalent) of $\exp(\rho^{(k)} \Delta X^{(k)}) - 1$.

In the case of the τ frictions, we can relate our estimates to what is known from direct measurement of the frictions. The elasticity of τ with respect to distance is of particular interest to us since it has been estimated on its own using various types of data in the literature, including the

effect of physical distance on freight costs. The ℓn distance cost elasticity in column (1) of Table 4 is $\rho^{\text{distance}} = 0.042$. Coşar et al. (2018) report a somewhat smaller value of $\rho^{\text{distance}} = 0.016$ (Table 12, column IV). Both estimates of ρ^{distance} , fit in the “reasonable range” of 0.01 to 0.07 in the literature summarized by Head and Mayer (2013). Our results imply that the distance effects on trade flows can be fully explained without reference to the “dark matter” invoked by Head and Mayer (2013) to explain aggregate distance elasticities of -1 or higher. This is not surprising since the main candidate explanations for dark matter—poor information and differences in preferences—should be accounted for in the δ_{in} marketing cost parameters.

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

How should we interpret the parameters shown in Table 4? The first thing to note is that the three chief variable cost frictions in the model, $\tau_{\ell n}$, $\gamma_{i\ell}$ and δ_{in} all require a normalization to be meaningful. Put another way, any of these frictions could be scaled up by a constant without changing any of the endogenous variables. The normalization we use is the internal friction within the United States.³¹ An estimated $\delta_{in} = 1.3$, for example means that firms headquartered in i inflate their delivered prices to consumers in n by 30% more than firms headquartered in the US inflate costs for their home-country consumers.

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

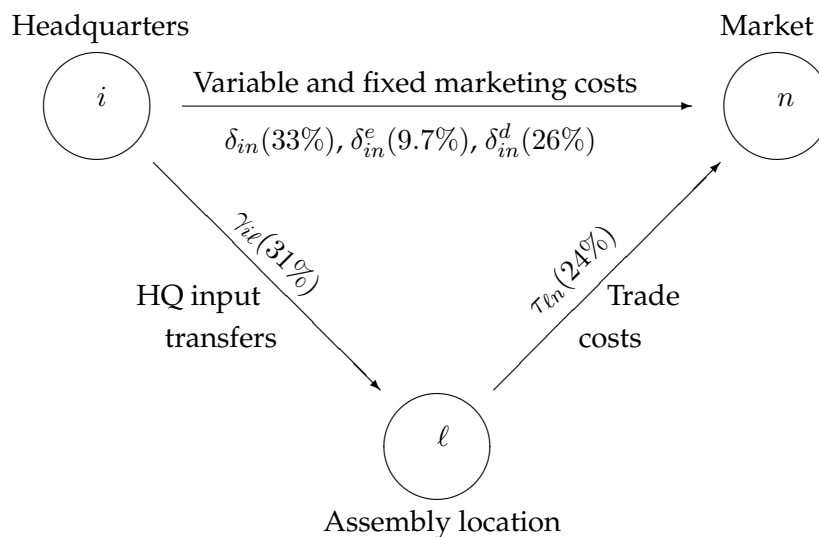
$$\frac{\mathbb{E}[\pi_{bn}](\delta'_{in})}{\mathbb{E}[\pi_{bn}](\delta_{in} = 1)} = \delta_{in}^d$$

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

³¹This means countries with smaller internal distances than the US can have $\tau_{nn} < 1$.

solving for a root.

Figure 3: Frictions impeding multinational flows



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

Results of those calculations are reported in figure 3 adapted from Arkolakis et al. (2013) with the addition of an edge corresponding to variable and fixed marketing costs δ_{in} , δ_{in}^e , and δ_{in}^d . On each edge of the triangle we report the relevant frictions, which is the median value calculated in our sample for the year 2016. The variable frictions are $\tau_{ln} - 1 = 24\%$, $\gamma_{il} - 1 = 31\%$, and $\delta_{in} - 1 = 33\%$. The sales of a model produced in ℓ , headquartered in i and sold in n would therefore face a total cost-increasing friction of 116% ($\tau_{ln}\gamma_{il}\delta_{in} - 1$). This is not out of line with the figures provided in Table 7, column (1) of Anderson and Van Wincoop (2004), ranging from 91 to 174%. Being the largest of the three frictions, variable marketing costs are quantitatively important enough to warrant inclusion in the multinational production framework. Set on top of that, the extra burden of fixed marketing costs have AVEs of 9.7% and 26% for model and brand entry, reinforcing the finding that the new dimension of frictions we added to the MP model is quantitatively very important.³² The only frictions that change in our counterfactual experiments are tariffs and RTAs. Tariffs are *ad valorem* already. The AVE of RTAs for the three dimensions of our friction triangle are 3.2% for τ , 6.4% for γ and a total of 29% for the combined effects of δ , δ^e , and δ^d .

The friction estimates shown in Figure 3 do not distinguish cost-based interpretations of τ_{ln} , γ_{il} , and δ_{in} from preference-based interpretations. For example, a desire by consumers to “buy local” to support workers has the same effect on sourcing as an increase in τ_{ln} . Similarly, if

³²Wang (2017) estimates the first and third components of the marketing costs, using foreign affiliate trade data from China. Since Wang’s sample includes all manufacturing industries, the large magnitudes he obtains suggest that these frictions are important beyond the car industry.

Japanese workers had a reputation for quality control, then Toyota’s assembly facilities outside Japan would have their sourcing shares reduced in a way that would be isomorphic to an increase in $\gamma_{i\ell}$. Finally, spatially correlated taste differences (e.g. for fuel economy, safety, or shape) could be equivalent in their effects on market shares to a rise in δ_{in} due to higher distribution costs in remote markets. Allowing for such preference effects in the utility function would just add more parameters that could not be identified separately from the ones existing in our specifications. To estimate separately the cost and demand-side effects would require a different estimation strategy that uses price information. Such a data requirement would severely limit the geographic scope of the study. For the purposes of our counterfactuals on how integration affects production and trade, we do not need to disentangle cost mechanisms from preference mechanisms.³³

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

Figure 4: Cost advantage inferred from sourcing decisions




Figure 4 graphs the cost advantage of the 47 assembly countries we use in estimation. The clear

³³Goldberg and Verboven (2001) separate out the home bias into demand and cost components. Coşar et al. (2018) estimate cost-based ($\gamma_{i\ell}$) frictions of distance from a brand’s home. They also have a home-brand effect in preferences that would operate as a δ_{in} effect in our model.

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

5.3 Segment-level estimates

Up until this point we have treated the market for cars as one in which all car models substitute symmetrically for each other. In reality, we think the Camry and Accord mid-sized sedans are closer substitutes for each other than either would be for a Sienna minivan (all Honda models). To allow for more realistic substitution patterns, we follow the tradition initiated by Goldberg (1995) and Verboven (1996) which groups models according to their primary function. In this extension, households receive idiosyncratic utility shocks that determine the choice of the segment from which they then select a model. This approach leads to a modified version of equation (22), in which market shares are measured *within segments* denoted s :

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

All of the demand parameters are segment-specific. While estimating an upper branch decision to determine Q_{nst} would be possible, the effect would simply be to bring back inter-segment substitution. Instead, we take Q_{nst} as exogenous in our counterfactual simulations, a modelling decision that restricts all substitution to take place *within segments*.³⁴ Bracketing the range of substitution is important for our counterfactuals since the unified market assumption has the potential to exaggerate the response to changes in frictions.

Allowing for segment-specific demand elasticities and nested substitution also affects the model-level expected profits: Equation (14) features the size of the market and the price index, which are now both segment-specific. In addition, the size of η_s determines the response to the cost index C_{bn} . Model entry should therefore also be run for each segment. The dependent variable changes from M_{bnt}/M_{bt} to M_{bnst}/M_{bst} . Once parameters from market share and model entry regressions have been collected, $\mathbb{E}(\pi_{bnt}) = \sum_s \mathbb{E}(\pi_{bnst})$ is computed to yield an estimation of brand entry choices, that takes into account the underlying segmentation of the car market.³⁵

Table 5 provides segment-level estimates for the market share equation (upper panel) and model and brandy entry (lower panel). We consider five segments: small cars, big cars, multi-purpose vehicles (MPVs, which includes minivans), sport utility vehicles (SUVs), and a combi-

³⁴Fixed Q_{nst} can be rationalized as the limiting value of a nested model in which the “segment shock” has very large variance.

³⁵The precise steps to construct $\mathbb{E}(\pi_{bnst})$ are the same as in appendix B, with the appropriate s subscript when the segment dimension is relevant.

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

Dep. Var:	Average market share					
Method:	PPML					
Segment:	(1)	(2)	(3)	(4)	(5)	
	small cars	big cars	MPV	SUV	sport& lux	
$home_{in}$	0.545	0.383	0.191	0.553	1.210	
	(0.303)	(0.285)	(0.293)	(0.236)	(0.134)	
$home_{in} \times LDC_n$	0.198	-0.467	-0.366	-0.399	-1.143	
	(0.313)	(0.417)	(0.696)	(0.233)	(0.585)	
$\ln dist_{in}$	-0.382	-0.269	-0.259	-0.121	-0.155	
	(0.102)	(0.096)	(0.078)	(0.088)	(0.063)	
$language_{in}$	0.216	0.327	0.285	0.518	0.219	
	(0.152)	(0.114)	(0.177)	(0.163)	(0.095)	
Deep RTA_{in}	0.042	-0.127	-0.083	-0.171	0.165	
	(0.159)	(0.125)	(0.162)	(0.113)	(0.130)	
$\ln C_{bn}$	-4.442	-6.046	-6.257	-5.558	-0.458	
	(0.932)	(0.722)	(1.388)	(0.622)	(0.554)	
Observations	29564	26692	16415	29379	17225	
rsq	0.476	0.494	0.388	0.434	0.542	
Dep. Var:	Model entry (fraction)					Brand entry
Method:	Fractional probit					Binary probit
Segment:	(1)	(2)	(3)	(4)	(5)	(6)
	small cars	big cars	MPV	SUV	sport&lux	all
$home_{in}$	0.029	0.218	0.222	0.152	0.224	0.835
	(0.069)	(0.085)	(0.123)	(0.077)	(0.061)	(0.462)
$home_{in} \times LDC_n$	0.779	0.076	-0.094	0.541	-0.563	1.076
	(0.128)	(0.191)	(0.209)	(0.187)	(0.131)	(0.514)
$\ln dist_{in}$	-0.064	-0.070	-0.154	-0.085	-0.045	-0.059
	(0.019)	(0.027)	(0.029)	(0.030)	(0.037)	(0.068)
$language_{in}$	0.088	0.040	-0.022	0.049	0.080	-0.203
	(0.051)	(0.032)	(0.036)	(0.039)	(0.046)	(0.128)
Deep RTA_{in}	0.047	0.155	0.045	0.089	0.052	0.362
	(0.037)	(0.055)	(0.092)	(0.043)	(0.056)	(0.086)
$\ln C_{bn}$	-0.953	-0.529	-0.733	-0.148	0.178	
	(0.282)	(0.344)	(0.373)	(0.317)	(0.411)	
$\ln \mathbb{E}[\pi_{bn}]$						0.566
						(0.055)
Observations	27871	24864	15838	25794	15916	113207
rsq	0.747	0.681	0.757	0.711	0.464	0.651

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. Each regression controls for log per-capita income and price level of the assembly country.

nation of sport and luxury cars.³⁶ The most important set of coefficients obtained is the one on $\ln C_{bn}$, revealing η_s . The demand side elasticity for small cars, at -4.4 is reasonably close from the one obtained in the baseline regression (-3.9). The price response is larger for MPVs, big cars, and SUVs, all around -6 . The sport&lux η of just 0.458 implies an infinite markup over costs, which makes it impossible to include this segment in aggregate profit calculations or counterfactuals. The average η_s , excluding sport&lux, is -5.6 . This higher elasticity of market shares with respect to changes in C_{bn} will tend to increase responsiveness to tariff changes in the segment version of the model, offsetting to some extent the elimination of cross-segment substitution.

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

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

6 Counterfactual methods

The counterfactuals investigate a set of trade policy scenarios involving shocks to tariffs and the deep RTA indicators. They treat as exogenous country-level new car purchases (Q_n), each brand's total number of models (M_b), each brand's set of production locations, \mathbb{L}_{bl} . Since the data used for each variable comes from a single year (2016, the last available in our sample) we suppress the time subscripts in this section.

The plausibility of counterfactual exercises depends in large part on trust in the underlying assumptions and estimated parameters. On three important issues we investigate more than one setting. In particular we have IRS and non-IRS settings, segmented versus unified car markets, and most importantly two different ways to solve the model (proportional changes or levels). We

³⁶We based the categorization on a combination of information on the “Global sales sub-segment” (a functional categorization) and the “Global sales segment” (a size categorization) of the model specified in IHS original data. Our segments therefore represent roughly similar-sized sets of models, grouped by categories suggested by the industry consultancy from which we bought the data.

Table 6: Friction parameters (segment-level)

Friction:	δ	δ^e	δ^d
Estimate:	\mathbf{f}	\mathbf{f}^e	\mathbf{f}^d
Small cars ($\eta = 4.44, \sigma^e = 3.61$)			
home	-0.123	0.318	-1.476
home \times LDC	-0.045	-2.660	-1.901
ln distance	0.086	-0.066	0.104
common language	-0.049	-0.149	0.359
RTA (deep)	-0.010	-0.138	-0.640
Big cars ($\eta = 6.04, \sigma^e = 9.54$)			
home	-0.063	-1.757	-1.476
home \times LDC	0.077	-1.119	-1.901
ln distance	0.044	0.443	0.104
common language	-0.054	-0.112	0.359
RTA (deep)	0.021	-1.584	-0.640
Multi-purpose vehicles ($\eta = 6.26, \sigma^e = 7.17$)			
home	-0.031	-1.427	-1.476
home \times LDC	0.058	0.368	-1.901
ln distance	0.041	0.885	0.104
common language	-0.046	0.397	0.359
RTA (deep)	0.013	-0.395	-0.640
Sport utility vehicles ($\eta = 5.56, \sigma^e = 30.73$)			
home	-0.100	-4.225	-1.476
home \times LDC	0.072	-16.950	-1.901
ln distance	0.022	2.517	0.104
common language	-0.093	-1.074	0.359
RTA (deep)	0.031	-2.872	-0.640
Sport and luxury cars ($\eta = 0.46, \sigma^e = 3.05$)			
home	-2.641	-2.115	-1.476
home \times LDC	2.496	3.071	-1.901
ln distance	0.339	0.321	0.104
common language	-0.478	-0.502	0.359
RTA (deep)	-0.359	-0.355	-0.640

Elasticities used to obtain frictions: $\sigma^d = 1.77$. Calculation of those frictions are described in the text and use coefficients from Tables 3 and 5.

elaborate on each of these methods below.

The IRS version of our counterfactuals solves for the equilibrium for a given q_ℓ and then updates q_ℓ to deliver new unit costs and a new equilibrium. This iteration continues until a fixed point is reached. There is concern over existence and uniqueness of equilibria with increasing returns. However, our simulations suggest that for our parameter values, there is a unique fixed point. Kucheryavy et al. (2016) find that a sufficient condition for uniqueness is that the trade elasticity multiplied by the scale elasticity (θ_ς in our notation) should be less than one. Our estimates $\theta = 7.7$ and $\varsigma = 0.035$ imply $\theta_\varsigma = 0.27 \ll 1$, suggesting a unique equilibrium. To isolate the impact of increasing returns and quantify the importance of interdependencies, we also conduct a non-IRS version of the counterfactuals. That setting treats $q_{\ell t}$ as another proxy for local assembly costs ($w_{\ell t}^\alpha$), along with GDP per capita and the exchange rate over-valuation index. Thus it is held constant at the observed level even as the policy variables are changed.

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

The third setting option in the counterfactuals is the method of solving for changes relative to the factual policies. One method follows the path of Dekle et al. (2007) and Arkolakis et al. (2012) in using Exact Hat Algebra (EHA). This takes advantage of the CES equations for sourcing shares and market shares. The main advantages are that i) it computes predicted (exact) percentage changes from the actual data, ii) related, it allows for unobservables in the actual decisions (as long as they are unaffected by the counterfactual), iii) it minimizes the data and parameter requirements.

EHA has two important disadvantages. The first is that brand entry cannot be handled by this method because it is a binary variable. The second concern in using EHA is that it does not allow a brand to start sourcing from an assembly country that was not used prior to the shock. Any zero remains zero, no matter how large the change in frictions. This motivates us to complement the EHA method with a second approach that solves the full model under the current set of frictions, computing expected values of sales, sourcing shares, model-level and brand-level entry decisions. This can be redone under the counterfactual set of frictions and one can then simply take the Difference in Expected Values (DEV) as the outcome of the counterfactual.³⁷ One drawback of DEV

³⁷Eaton et al. (2013) also follow this approach to run counterfactuals although two of the authors where key pre-

is that it requires estimates of 662 parameters.³⁸ By sharp contrast, the EHA approach requires just 10 parameters ($\eta, \theta, \sigma^e, \sigma^d, \alpha, \varsigma$, and the 4 friction parameters for deep RTAs). EHA allows observed levels of the endogenous variables to “stand in” for parameter estimates as well as unobservables. Thus, by definition, EHA replicates the actual data, whereas DEV sometimes errs by large amounts in predicting the factual levels of production in each country (q_ℓ).

6.1 Exact Hat Algebra (EHA)

EHA solves for an equilibrium in the proportional changes of all variables of interest following a change in frictions. Variables taking a hat symbol are defined as ratios of $\hat{x} = x'/x$, where x is the initial level, and x' is the level attained after the change. The EHA method used here includes two non-standard features: First, we allow for proportional changes in the fraction of models offered in a market. Second, we incorporate external returns to scale in a multinational production setting. The change in total output located in country k , denoted \hat{q}_k , affects outcomes through \hat{C}_{bn} . Therefore, the problem can be decomposed in an inner and outer loop, analogous to the structure used in the pure trade model of Kucheryavyy et al. (2016).

The inner/outer loop procedure works as follows:

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

$$\hat{C}_{bn} = \left(\sum_k \mathbb{L}_{bk} s_{bkn} (\hat{\gamma}_{ik} \hat{\tau}_{kn} \hat{q}_k)^{-\theta} \right)^{-1/\theta}, \quad (26)$$

$$\hat{P}_n = \left(\sum_b \mathbb{D}_{bn} \frac{q_{bn}}{Q_n} \hat{M}_{bn} (\hat{\delta}_{in} \hat{C}_{bn})^{-\eta} \right)^{-1/\eta}, \quad (27)$$

$$\hat{M}_{bn} = \Phi \left[\frac{\ln(\hat{\delta}_{in}^{1-\eta} \hat{C}_{bn}^{1-\eta} \hat{P}_n^\eta) - \ln \hat{\delta}_{in}^e}{\sigma_e} + \Phi^{-1} \left(\frac{M_{bn}}{M_b} \right) \right] \frac{M_b}{M_{bn}}, \quad (28)$$

where s_{bkn} is the share of sales in market n sourced from country k . A fixed point iteration with dampening solves for the equilibrium values of these variables. In our counterfactual, our main variable of interest is the percentage change in quantities

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

Again the EHA approach is very useful here, since it can be used to show that the changes

cursors in developing the EHA approach (Dekle et al. (2007)). The model used in Eaton et al. (2013), by dropping the continuum of firms assumption, also has to give up the “macro restrictions” that Arkolakis et al. (2012) show are needed to be able to use EHA.

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

in sourcing probability and brand market share are only functions of changes in frictions (known) and of the three endogenous variables \hat{C}_{bn} , \hat{P}_n , and \hat{M}_{bn} solved by the fixed point iteration:

$$\hat{s}_{bln} = \left(\frac{\hat{\gamma}_{il} \hat{\tau}_{ln} \hat{q}_\ell^s}{\hat{C}_{bn}} \right)^{-\theta}, \text{ and } \hat{q}_{bn} = \hat{M}_{bn} \left(\frac{\hat{\delta}_{in} \hat{C}_{bn}}{\hat{P}_n} \right)^{-\eta} \quad (29)$$

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

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

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

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

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

In changes

$$\hat{q}_{bln} = \frac{q'_{bln}}{q_{bln}} = \hat{s}_{bln} \times \frac{\sum_s q'_{bns}}{\sum_s q_{bns}}. \quad (30)$$

The expression describing sourcing share \hat{s}_{bln} in equation (29) is unchanged, since it is not affected by any segment-level determinant. However, the new level of production at the segment level is

$$q'_{bns} = q_{bns} \hat{q}_{bns} = q_{bns} \hat{M}_{bns} \left(\frac{\hat{\delta}_{ins} \hat{C}_{bn}}{\hat{P}_{ns}} \right)^{-\eta_s}, \quad (31)$$

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

$$\begin{aligned} \hat{P}_{ns} &= \left(\sum_b \mathbb{D}_{bns} \frac{q_{bns}}{Q_{ns}} \hat{M}_{bns} (\hat{\delta}_{ins} \hat{C}_{bn})^{-\eta_s} \right)^{-1/\eta_s}, \\ \hat{M}_{bns} &= \Phi \left[\frac{\ln(\hat{\delta}_{ins}^{1-\eta_s} \hat{C}_{bn}^{1-\eta_s} \hat{P}_{ns}^{\eta_s}) - \ln \hat{\delta}_{ins}^e}{\sigma_e^s} + \Phi^{-1} \left(\frac{M_{bns}}{M_{bs}} \right) \right] \frac{M_{bs}}{M_{bns}}. \end{aligned} \quad (32)$$

The algorithm is very similar to the unified markets case. The inner loop solves for \hat{C}_{bn} , \hat{P}_{ns} and \hat{M}_{bns} , using (26) and (32) which gives \hat{q}_{bln} from (31) and (30). The outer loop then sums over the new shipments to obtain total output in each country, which enters back \hat{C}_{bn} in the next iteration. The process is repeated until no further change is detected in any of those endogenous variables.

6.2 Difference in Expected Values (DEV)

While EHA simulations calculate percentage changes in the endogenous variables directly, the DEV approach solves the model in levels twice: once at the factual level of frictions, and once for the same frictions evaluated under the counterfactual scenario. The equations summarizing the equilibrium are the sourcing decision (6), the market share (9), the price index (10), and the two entry equations (15) and (19). The identification of structural parameters needed for those equations is detailed in section 4.3 and in appendix B. As in EHA, solving the model involves nested fixed point iterations, but with an additional loop to accommodate brand entry. The computation of the equilibrium entails an inner, middle, and outer loop.

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

$$\mathbb{E}[q_{b\ell n}] = \mathbb{E} \left[\frac{q_{bn}}{Q_n} \mid \mathbb{D}_{bn} = 1, M_{bn} \right] \times \text{Prob}(\mathbb{S}_{b\ell n} = 1) \times Q_n. \quad (33)$$

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

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

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

policies maintain the exact draw of fixed costs for each brand-destination but recomputes expected gross profits. This allows for some brands to be drawn into or out of its factual entry decision. To obtain expectations we replicate the solution of the model with 1000 draws of fixed costs.³⁹

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

The segmented market version of DEV has a few important differences. In step 1, price indices, the mass of models offered, and all three marketing costs (δ , δ^e and δ^d) need to be defined at the segment level (the sourcing probabilities, as well as τ and γ frictions do not have a segment dimension). The inner loop therefore solves for equilibrium P_{ns} and M_{bns} . The middle loop (run as a second step) updates national output by summing up the segment-level sales of the brand that are expected to be sourced from different origin countries:

$$\mathbb{E}[q_\ell] = \sum_b \sum_n \mathbb{E}[q_{b\ell n}] = \sum_b \sum_n \text{Prob}(S_{b\ell n} = 1) \sum_s \mathbb{E} \left[\frac{q_{bns}}{Q_{ns}} \mid \mathbb{D}_{bn} = 1, M_{bns} \right] \times Q_{ns}. \quad (34)$$

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

Before turning to solutions of counterfactual policies, it is important to demonstrate that the endogenous variables solved for under factual trade policies do not depart too much from their data counterparts. Starting with the decision where brands are offered, the entry rate in fact is

³⁹ A subtle aspect of this approach is that it does not actually guarantee that entry choices in the simulation of factual policies match the real decisions 100% of the time. This is because the truncation is based on the gross profits calculated using price indexes extracted from the fixed effects in the market share regression. The P_n obtained in the solved model will differ, occasionally by enough to alter the entry decision of individual brands that were near the entry/exit threshold.

36.2%, slightly lower than the 36.6% average in the simulation with a maximal difference of 0.82%. Figures G.1 and G.2 in Appendix G show the fit of one run of the DEV simulation (for unified and segmented markets respectively) under the factual set of policies. The correlations between simulated and actual are high for all four variables examined. In the non-segment, IRS case, we obtain the following correlation coefficients: 0.98 for the price index, 0.63 for brand-origin-destination sales, 0.74 for country-pair flows and 0.86 for aggregate country-level output. The correlations are even higher for the non-IRS and segmented versions of the model.

7 Counterfactual results


The main motivations for estimating the model of this paper is the investigation of counterfactual trade policy changes. We report results on seven different scenarios.

1. Terminating NAFTA: end of North American tariff preferences and deeper integration.
2. UK exit from the European Union:
 - (a) Soft Brexit: a free trade agreement retains tariff-free trade between the UK and the EU27 but all deeper integration measures are rescinded.
 - (b) Hard Brexit: in addition to the cessation of deep integration measures, the EU27 and UK impose the current EU MFN tariffs.
3. Trans-Pacific integration with and without the US:
 - (a) TPP: a deep integration agreement between Australia, Brunei, Canada, Chile, Japan, Malaysia, Mexico, New Zealand, Peru, Singapore, the United States, and Vietnam.
 - (b) Comprehensive and Progressive Transpacific Partnership (CPTPP): a revised TPP agreement (signed March 8, 2018), which omits the US.
4. Trans-Atlantic Integration:
 - (a) Comprehensive Economic Trade Agreement (CETA): deep integration between the EU28 and Canada that has been provisionally applied since September 2017.
 - (b) CETA + Transatlantic Trade and Investment Partnership (TTIP): EU28 deep agreements with Canada *and* the US.

7.1 Termination of NAFTA

Renegotiation of the North American Free Trade Agreement began shortly after the inauguration of Donald Trump who repeatedly threatened to pull the US out of the agreement if his goals are not met. The policy experiment we consider is the dissolution of NAFTA. All three members end deep integration and revert to imposing MFN tariffs on each other. Car import tariffs rise to the

Figure 5: Termination of NAFTA



MFN levels of 6% in Canada, 2.5% in the United States, and 34% in Mexico. Tariffs on parts will rise as well but those tariffs are low in all three countries.⁴⁰ A more likely scenario at the time of writing would be an Article 2205 withdrawal by the US which would leave the agreement in place between Canada and Mexico. However, the experiment of ending NAFTA entirely has the advantage of serving as a way to quantify the benefits created by the agreement as a whole for each member.

Figure 5 illustrates the impact of ending NAFTA under eight different settings in the estimation and simulation methods. These symbols use shape, size and shading to represent the 2^3 combinations for three different settings. Shape identifies the simulation method, with triangles depicting EHA, and circles depicting DEV. Larger symbols indicate the presence of increasing returns. Blue symbols use the segmented market assumption, red is unified, and purple corresponds to cases where the results overlap. The left panel shows percentage change in car production in each country while the right panel shows changes in consumer surplus (CS).

The results shown in Figure 5 point to potentially disastrous outcomes for Canada and Mexico. Canada stands to lose 44–47% of its car industry under increasing returns settings. Mexico could lose up to 29% of its production. These losses would be equivalent to shutting two to four median-sized plants in Canada (producing 243 thousand cars per year) and two to five of the smaller (150 thousand) plants operating in Mexico.

Four issues are important for understanding these large losses. First, the shares of Canadian and Mexican car production headed to US in 2016 was very high: 83% for Canada and 55% for Mexico. Second, although the US tariff on Canadian and Mexican cars would only rise to 2.5%, the τ frictions add another 3.2% for all shipments into the US. Third, the γ friction rises because of the higher input tariffs (1.35% in Mexico, 2.74% in Canada) combined with the loss of deep integration, which adds 6.4% to costs of US brands assembled in Canada (55% in 2016) and Mexico (36%). Fourth, almost all the brands made in Canada (11 of 12) and Mexico (10 of 14) are also made in the US. Since the sourcing elasticity, θ is so large (7.7), this implies these brands can and will heavily shift production for the US market to US assembly sites when NAFTA raises costs of using sites in Canada and Mexico.

The outcome for US car production is method-dependent. Whereas US output contracts by 2.4% or less under DEV, it expands by up to 11% under EHA. Although production in the US rises in the EHA simulation of NAFTA, the combined market share of US brands in the US actually declines from 36% to 34%. In Canada and Mexico the rise in C_{bn} caused by the end of NAFTA is much more severe: 31% to 24% in Canada and 29% to 23% in Mexico. The damage to firms caused by impeding their use of offshore factories is something a pure trade model would not pick up. Car buyers in all three countries lose as well, with the worst case scenario being Canada's 5.4% loss under unified markets, IRS, and EHA. Even under the best-case scenarios, Canadian and Mexican car buyers lose about three percent of their consumer surplus. US car buyer losses are

⁴⁰The NAFTA termination scenario assumes that pre-NAFTA agreements, such as the 1965 Canada-US Auto Pact, are not reinstated.

under 2%.

Allowing for increasing returns generates much larger predicted losses for Canada as can be seen in Figure 5. The model actually exhibits surprisingly systematic magnification effects from external returns to scale. When regressing the country-level percentage change in output of the IRS case against the change in the non-IRS case, we obtain a coefficient of 1.28 for DEV and 1.24 for EHA (segmented markets in both cases), with near perfect fit ($R^2 > 0.99$ in both cases).

In comparison to the predictable magnification that IRS brings to the counterfactuals, it is harder to characterize the changes brought about by segmenting the car market into functional categories. Earlier we suggested that Smart's gains from NAFTA termination would be limited because it can only take market share from other small cars. Figure 5 shows that, contrary to our initial expectations, the segment version of the model does not systematically dampen the impact of changes in trade policies. This is because switching to segments has two other effects that can amplify policy changes. First, η for small cars (4.4) is larger than the η for for the unified car market (3.9). This implies a great market share response to policy changes. A less obvious difference is that segmenting tends to reduce the number of competitors, which also magnifies reallocations between brands. Thus, Smart gains more market share *within the small car segment* when its rivals' costs of serving the US market rise. The net effect of all the changes introduced by segmenting is that terminating NAFTA raises Smart's sales in the USA from its French plant by 6.7% with segments, compared to 5.6% without them. Aggregating up to the country level, Figure 5 shows the segments setting has very little net impact. The right panel of Figure 5 shows that under segments Canadian and Mexican car buyers have smaller losses of surplus. This can be attributed to the larger η , implying greater willingness to substitute to car models whose prices were not increased by NAFTA termination.

Figure 5 reveals that for the US, Korea, India, Spain and Brazil the most important setting decision is whether to solve for changes via the EHA or DEV method. The most obvious difference between EHA and DEV is that we can only incorporate brand entry and exit in the latter method. With the end of NAFTA, one would expect brands to revisit their entry strategy, at least in the medium run. The DEV counterfactual results calculate the change in the probability of opening a distribution network over 1000 repetitions. The simulations predict that the US brands with low propensities to enter foreign markets will tend to drop out of Canada and especially Mexico. The most extreme case is Buick which enters Mexico in every replication under the factual simulation but only 40% of the counterfactual replications. On the other hand, a number of second-tier EU brands increase their entry propensity in both markets. For example Citroën has only an 10% chance of entering Mexico with NAFTA but a 27% chance without NAFTA. Skoda, Dacia, and Opel more than triple their brand entry probabilities in Mexico. In Canada, the biggest increases in entry probability are for the Japanese brands, Acura, Scion, and Infiniti. Smart, both designed and produced in Europe also sees a large increase in entry probability. In practice, we find that brand entry issues do not have much quantitative impact because the brands that are on the margin of entering and staying account for very little output and also have limited sourcing alternatives.

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

For example, Ford does not actually send cars from its Spanish or Polish factories to the North American market. The large British plants of Toyota and Nissan did not send a single car in 2016 to any of the three NAFTA members. In EHA, this is interpreted as a huge and persistent friction. As a consequence, the US market carries out zero substitution from Nissan's Mexico factory to its UK factory in response to NAFTA termination. However, Ford, Toyota and Nissan (among many others) have positive *expected* flows from their European plants to North America in DEV. Terminating NAFTA leads them to source substantially more from those EU-based operations to serve North-American markets. For Poland, Spain, and the UK, the EHA counterfactual shows a much weaker response compared to DEV as expected, reflecting the fact that the latter method accounts for the possibility of UK-made Nissans to be sold in the US. The most extreme case is Poland for which DEV increases exceeds those of EHA by up to 6 percentage points.

India's end-of-NAFTA outcomes reveals the way EHA implicitly takes into account model "residuals," i.e. deviations from expected values. Since India is a high-cost producer and faces near 34% MFN tariffs on exports to Mexico, we would not expect it to be much of an exporter to Mexico. But in fact Chevrolet, Ford, Volkswagen, Hyundai, and Suzuki all export large amounts. Termination of NAFTA implies that US-made cars face the same 34% tariffs to export to Mexico. This creates an expansion in Indian auto manufacturing ranging from 2.2 to 3% (depending on IRS and market segmentation). DEV does not build this unexplained ability of India to export cars to Mexico into the counterfactual (or factual). As a consequence, it only predicts a maximal increase of 0.45% for India. The case for Spain is simply the reverse of India. Spain is a low-cost maker with zero-tariff access to Mexico. It would be expected to be a solid exporter to Mexico before NAFTA but in fact it only exports about 25 thousand cars (implying a bad "residual" under EHA). Hence the average increase is almost ten times as large under DEV (+2.54%) than under EHA (+0.275%).

The net effect of cases like that of Spain is that ending NAFTA is good for the US under EHA. This is because EHA rules out re-sourcing to a number of countries in Europe where US-owned plants do not currently send positive amounts of cars to the US. Therefore, under EHA, US factories do not have to share the gains from reduced access to Canadian and Mexican plants with their EU counterparts. Under DEV, the EU plants gains are large enough to generate a small net loss in US production in a hypothetical NAFTA termination. For the same reason, DEV dampens losses to US car buyers, who benefit from the option of importing EU-made cars. It seems plausible to

us that the DEV outcomes for the US are the more likely ones to prevail in the medium-run as multinationals increase sourcing from their European factories to reflect the loss of preferential access of Mexican and Canadian plants to the US market.

7.2 United Kingdom exits the European Union (Brexit)

In the aftermath of the 2016 Leave vote, debate revolved around whether Brexit will be “soft” or “hard.” We consider two potential post-membership relationships between the UK and EU along those lines. The soft Brexit case captures the scenario in which Britain retains tariff-free access to the EU but loses the deep integration aspects of the RTA such as free mobility of professionals and the ability to influence EU regulations on car standards. We then simulate the hard Brexit scenario where UK exports face the EU’s 10% MFN tariffs while the UK reciprocates at the same rates. In both scenarios we assume that the UK will be unable to “roll over” existing EU trade agreements, thus reverting to MFN duties with Turkey, South Korea, South Africa, among others. The two post-Brexit scenarios set all the deep RTA dummies to zero if they correspond to i) dyads involving the UK and EU, ii) a dyad involving the UK and one of the countries having a preferential deep RTA with the EU.

Figures 6 and 7 point to poor outcomes for the both the UK car industry and the British car buyer. The worst-case scenario for workers in UK plants would be the Hard Brexit (EHA, segments, IRS) where the industry contracts by 18.3%. With Soft Brexit, the model still predicts roughly 10% losses under EHA and 4% under DEV. In absolute terms, the Soft Brexit losses for the UK would range from 46 to 160 thousand cars. To put those numbers in perspective, Honda’s Swindon factory assembles 110 thousand cars in 2016 with approximately 3100 workers in 2015.⁴¹ The drop in production in the Hard Brexit case would rise up to 275 thousand cars, up to two and a half Swindon-equivalents (around 7,700 jobs lost).⁴² Consumers lose 3–4% of their surplus with Soft Brexit but their losses roughly double to 8–9% in the event of a Hard Brexit.

Since the British buyers favour continental brands, their losses from 10% tariffs and even larger marketing and trade frictions are easy to comprehend. The cause of the production losses was not obvious until we dug into the demand and sourcing patterns. The basic problem the UK faces is the fact that major multinationals like Ford and Nissan have factories in Spain they can easily switch production to in order to continue to serve the EU27 tariff-free. On the other hand, major German brands such as VW and Mercedes-Benz lack UK factories. This limits the scope for re-orienting the supply for UK demand to UK production locations. This effect is more severe under EHA which takes into account the positive residual in the German brands’ markets shares in the UK.

Aside from the UK, the big losers from Soft Brexit are Turkey and South Africa. Both countries suffer from the loss of tariff-free access to the UK that they obtained through trade agreements

⁴¹Financial Times, March 13, 2015.

⁴²Swindon is an interesting case since it is i) very close in size to the average plant operating in the UK, ii) the only plant producing Hondas in the EU in 2016, iii) producing essentially all Civic and CR-V models sold in the EU, more than half of all Honda sales in the region.

Figure 6: Soft Brexit (shallow FTA)





Figure 7: Hard Brexit (MFN tariffs)



with the EU. South Africa is especially hard hit in the EHA setting (around -9%) because it has a surprisingly high sourcing share in the data (18 times higher than South Africa's exports to France for example). South Africa's production losses under EHA are large enough to trigger add-on costs from lower scale economies. Turkey's situation reverses the EHA/DEV differences (Turkey's share in UK purchases are only a third of their share in France).

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

7.3 Trans-Pacific Integration

Figure 8 displays the predicted impact of TPP on the fifteen countries with the largest sales in the zone in 2016. This figure refers to the initially planned set of 12 participants including the United States. TPP removes all tariffs and sets the deep RTA dummies to one between participants. This policy is mainly of interest as "might-have-been" since the US exited the agreement in January 2017. Figure 9 projects the outcome of the CPTPP, the successor agreement involving the 11 remaining members which was concluded in January 2018.

While there have been some claims that TPP was more of a regulatory agreement than a trade agreement, this view neglects the fact that TPP and CPTPP include substantial tariff cuts for some country pairs. Japanese exporters face a 44% tariff when exporting to Vietnam, and 6–7% to Canada and New-Zealand. US exports face a 55% duty in Vietnam and 23% in Malaysia.⁴³

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

There are several forces at work underlying the gains our simulations predict and they represent aspects of our framework that do not feature in traditional models. Japanese plants, which

⁴³Japan signed a free trade agreement with Malaysia that entered into force in July of 2006, and stipulated a 9-year phase-out period for cars.

Figure 8: Transpacific Partnership (TPP, including US)





Figure 9: CPTPP (TPP without US)



account for 45% of Canadian production, are predicted to obtain a 6.4% cost reduction as a result of the deep γ_{il} aspect of the agreement. Elimination of the 2.7% Canadian tariff on auto parts (an average across the main auto parts HS codes) will further reduce costs at Toyota and Honda's Ontario plants. Under TPP, similar gains would also accrue to the Japanese plants in the US, although US parts tariffs are only 1.4% on average. Even though the γ effects are mainly balanced, under TPP the reduction in marketing costs from Japan, leads to increased market shares for Japanese brands in the US with much of the additional demand sourced from Canada. Under CPTPP, the situation is different. Japanese brand shares stay about the same in the US but increase the fraction of models sourced from their Canadian plants. Meanwhile, because Japan-origin cars avoid Canada's 6% car tariff, Japan's sourcing share rises in the Canadian market. The total number of Japanese cars sold in Canada rises enough to offset this negative effect.

Lower costs of marketing Japanese cars are predicted to stimulate entry of about 22% more Japanese car varieties in Canada under the CPTPP, and would stimulate about the same increase in the US had the TPP been implemented. Since a significant fraction of these new models will be sourced from the Toyota and Honda plants in Canada, model entry stimulates production gains. There are also Japanese brands that would have been unlikely to enter Canada but are more likely to do so under CPTPP: Acura, Scion and Infiniti increase their entry probability in Canada by a substantial range (between 10 and 22%). Daihatsu and Suzuki, two brands not currently sold in Canada are predicted to enter the market in around 4% of the CPTPP Monte Carlo repetitions. As these brands are not made in Canada, their entry takes market share from brands who do produce locally and therefore tends to lower production.

The γ gains our model predicts for Canadian production are not present in the conventional trade model used by Global Affairs Canada to predict the consequences of CPTPP.⁴⁴ This study finds close to negligible effects of CPTPP on Canadian output. Another study conducted in 2012 by Van Biesebroeck et al. (2012)⁴⁵ looked at the impact of several RTAs for automobile production in Canada. While their use of detailed data on car characteristics allowed the authors to use the whole apparatus of BLP-type demand estimation, they do not account for the γ -type frictions we include here. This is critical in terms of predicted outcomes. When we run our CPTPP counterfactual without any γ liberalization, we find a very small overall impact on Canadian output, confirming the results of the two mentioned studies.

The other country strongly affected by TPP and CPTPP is Vietnam. Here, more than ever, we observe how different EHA and DEV solutions can be. Under DEV, Vietnam increases production with the TPP scenario. This is because of improved access to the US market and more efficient operations at Ford and Chevrolet plants which are expected to stimulate a big increase in exports to the US. In reality, Vietnam production is 100% oriented towards the local market. This implies no gains in exports to the US under EHA. However, the Japanese makers present in Vietnam

⁴⁴<http://international.gc.ca/trade-commerce/trade-agreements-accords-commerciaux/agr-acc/cptpp-ptpgp/impact-repercussions.aspx?lang=eng>

⁴⁵<http://www.international.gc.ca/economist-economiste/analysis-analyse/studies-etudes/studies-etudes-01.aspx?lang=eng>

will radically increase their sourcing from Japan, leading to the production losses we see for Vietnam under EHA. Without the US in the agreement, Vietnam loses car production even with DEV. Another distinctive aspect of CPTPP for Vietnam is that the brand entry margin is very active. Because of the increased competitiveness of Japanese brands, many competitors are predicted to exit. Land-Rover is predicted to exit Vietnam in 42% of the replications, Chevrolet also faces a large exit probability at 30%. Mercedes-Benz, BMW and Hyundai have a probability of exiting the country between 12 and 16%. By contrast, Subaru and Lexus (currently not sold in Vietnam) are predicted to enter in 14 and 11% respectively. The bright side of CPTPP is that the reduction of the 55% tariffs on assembled cars leads to major benefits for Vietnamese car buyers who the DEV version predicts to receive 25% lower price indexes.

7.4 Trans-Atlantic Integration

The last set of counterfactual scenarios we are considering relate to the couple of agreements that span the Atlantic. The first one is the Comprehensive Economic and Trade Agreement (CETA) between Canada and the EU. Although the agreement still awaits ratification by different European national parliaments, it has been applied provisionally since September 21st of 2017. By contrast, negotiations on the Transatlantic Trade and Investment Partnership (TTIP), an integration agreement between the EU and the US, were put to an indefinite halt following the 2016 US elections. We therefore consider the first scenario to be the full implementation of CETA: the abolition of tariffs between the EU and Canada, combined with deep integration. On top of those, the second scenario applies the same policy changes to the EU28-USA country pairs.

Our counterfactuals point to three possible outcomes from CETA for Canada. With increasing returns, the DEV method predicts an 8% production increase. This falls to 6% when the IRS magnification effect is eliminated. Under EHA, however Canadian production is essentially unchanged. The reason EHA predicts such small effects is that Canada has negligible exports to the EU. Our estimates predict it should export about one percent of cars to the UK for example. The share in actual data is one tenth that. Similar ratios apply to other EU destinations. Canada does not import large shares from the EU either so under EHA only Germany and Italy experience sizable output changes. Canadian car buyers are expected to gain 1–2% more surplus but the CETA impact on European buyers is predicted to be negligible.

Adding the US to the trans-Atlantic liberalization of car trade would lead to potentially large output gains for Spain (up to 16%), Italy (up to 45%) and Poland (up to 74%).⁴⁶ Only Italy gains much under EHA and this comes entirely from one interesting data fact. Jeeps made in Fiat factories in Italy are successfully exported to the US. EHA allows this to expand massively. The reverse case is Poland. Ford's factories there *should* be exporting already to the US and therefore rapidly expand through a mix of τ and γ effects. In practice, Ford uses this plant exclusively to serve EU car buyers. Japan and Korea will both see modest contractions in output in either form of trans-

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

Figure 10: CETA (EU-Canada deep RTA)





Figure 11: CETA + TTIP



Atlantic agreement. Consumers in all 30 countries would gain in the range of 1–2% from CETA plus TTIP. The main reason for these gains is the efficiency improvements our estimates predict for the EU factories of US-origin brands.

Several insights emerge from our seven counterfactual exercises that we anticipate to be broadly applicable. First and foremost, the choice between solution methods is not innocuous. It often has serious quantitative effects and occasionally changes the sign of the output changes (though the sign of consumer surplus changes are much more robust). Second, the structure of multinational production matters a great deal. It was not obvious that Brexit should reduce British car production. If there had been VW and Renault plants in the UK and no Toyota and Nissan plants on the continent, UK car production might have expanded. Third, even modest scale economies (a doubling of national scale only reduces costs by 2% according to our estimates) can result in very sizable amplification of the losses or gains from changes in frictions. Fourth, some startling results, like the boom of Canadian car production under the CPTPP hinge importantly on the γ effects of RTAs. Under CPTPP, the Canadian plants of Toyota and Honda experience a 6% reduction in costs due to deep integration and 2.7% cheaper headquarter inputs because of phased-out tariffs on parts imported from Japan.

8 Conclusion

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

We refer to those decisions collectively as the medium-run response, because they hold the set of production locations constant. In this time frame, each brand's network of sourcing alternatives strongly shapes the responses to policy changes. For instance, one reason why Canada suffers so much from the end of NAFTA is because 11 of the 12 brands that produce in Canada have plants in the US which they can switch production to. The long-run decision of opening and closing production operations in different countries is of course very interesting. Our focus on the medium-run follows from the desire to keep the estimation tractable and the scope of this paper finite. The medium-run is also important in its own right because production networks are strongly persistent. Even over a 17-year period, covering a major disruptive crisis for this industry, 88% of OECD countries' car production still takes place within brand-country combinations that existed in 2000.

The policies we investigate are predicted to have large impacts on all the margins. Production in Canada and Mexico fall by up to 47% and 29% in the event of NAFTA termination respectively.

The recently signed CPTPP increases Japanese brands market share in Canada from 42% to 51%. Because the model estimates that the Ontario factories of Toyota and Honda will become more efficient due to deep integration between Canada and Japan, the simulations predict Japanese market shares will even rise from 43% to 45% in the US, even though the US is no longer participating in the agreement. Sourcing of Minis from the Netherlands for sales to US customers falls by 37% under Brexit because of rising frictions exacerbated by lost scale economies. A common factor underlying the large effects predicted here are the magnitudes of the two key elasticities: 7.7 for substitution between assembly sites (θ) and 3.87 for substitution between varieties (η). Moreover, tariffs in the car industry are far from negligible (6% in Canada, 10% in the EU to give OECD examples). Finally we estimate that deep integration has substantially lowered multinational frictions between the headquarters and the locations of production and sales. The loss of such integration benefits by discarding long-established trade relationships is projected to harm car buyers through higher prices and fewer varieties, with the worst case being the 9% rise in the price index in the UK in the event of a Hard Brexit.

These large estimated benefits of producing, designing, and selling within a country, within an RTA, or with nearby countries all motivate future research to identify the mechanisms that underlie these frictions. This is particularly the case for the δ effects we have added to the framework. While we specify them as marketing costs, preference-based mechanisms may play important roles. A third topic (along with plant location choices) calling for more research is the decision of where to source the components to be assembled in each car plant. Due to data limitations, we focused on the role of tariffs in raising costs for parts originating from headquarters, but actual sourcing problems are further complicated by rules of origin. These and other aspects of multinational expansion strategies provide a full agenda for future research.

References

- Anderson, J. E. and E. Van Wincoop (2004). Trade costs. *Journal of Economic literature* 42(3), 691–751.
- Anderson, S., A. De Palma, and J. Thisse (1992). *Discrete choice theory of product differentiation*. MIT Press.
- Antras, P., T. C. Fort, and F. Tintelnot (2017). The margins of global sourcing: theory and evidence from us firms. *American Economic Review* 107(9), 2514–64.
- Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). New trade models, same old gains? *American Economic Review* 102(1), 94–130.
- Arkolakis, C. and F. Eckert (2017). Combinatorial discrete choice. mimeo.
- Arkolakis, C., N. Ramondo, A. Rodríguez-Clare, and S. Yeaple (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.
- Bernard, A. B., S. J. Redding, and P. K. Schott (2011). Multiproduct firms and trade liberalization. *The Quarterly Journal of Economics* 126(3), 1271–1318.
- Berry, S., J. Levinsohn, and A. Pakes (1995). Automobile prices in market equilibrium. *Econometrica*, 841–890.
- Berry, S., J. Levinsohn, and A. Pakes (1999). Voluntary export restraints on automobiles: evaluating a trade policy. *American Economic Review* 89(3), 400–430.
- Björnerstedt, J. and F. Verboven (2016). Does merger simulation work? Evidence from the Swedish analgesics market. *American Economic Journal: Applied Economics* 8(3), 125–64.
- Bresnahan, T. F. and V. A. Ramey (1993). Segment shifts and capacity utilization in the US automobile industry. *The American Economic Review* 83(2), 213–218.
- Bresnahan, T. F. and V. A. Ramey (1994). Output fluctuations at the plant level. *The Quarterly Journal of Economics* 109(3), 593–624.
- Coşar, K., P. Grieco, S. Li, and F. Tintelnot (2018). What drives home market advantage? *Journal of International Economics* 110, 135–150.
- Costinot, A. and A. Rodriguez-Clare (2014). Trade theory with numbers: Quantifying the consequences of globalization. In E. Helpman, G. Gopinath, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, Chapter 4, pp. 197–261. Elsevier.
- de Gortari, A. (2017). Disentangling global value chains. Technical report, Mimeo, Harvard University.
- Dekle, R., J. Eaton, and S. Kortum (2007). Unbalanced trade. *American Economic Review* 97(2), 351–355.
- Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica* 70(5), 1741–1779.
- Eaton, J., S. Kortum, and S. Sotelo (2013). International trade: Linking micro and macro. In D. Acemoglu, M. Arellano, and E. Dekel (Eds.), *Advances in Economics and Econometrics Tenth World Congress*, Volume II: Applied Economics. Cambridge University Press.
- Fajgelbaum, P., G. M. Grossman, and E. Helpman (2011). Income distribution, product quality, and international trade. *Journal of Political Economy* 119(4), 721–765.
- Feenstra, R. C. and J. A. Levinsohn (1995). Estimating markups and market conduct with multidimensional product attributes. *Review of Economic Studies* 62(1), 19–52.

- Fuss, M. and L. Waverman (1990). The extent and sources of cost and efficiency differences between us and Japanese motor vehicle producers. *Journal of the Japanese and International Economies* 4(3), 219–256.
- Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets: The case of the US automobile industry. *Econometrica*, 891–951.
- Goldberg, P. K. and F. Verboven (2001). The evolution of price dispersion in the European car market. *The Review of Economic Studies* 68(4), 811–848.
- Grigolon, L. and F. Verboven (2014). Nested logit or random coefficients logit? A comparison of alternative discrete choice models of product differentiation. *Review of Economics and Statistics* 96(5), 916–935.
- Grossman, G. and E. Rossi-Hansberg (2010). External economies and international trade redux. *The Quarterly Journal of Economics* 125(2), 829–858.
- Guimaraes, P., O. Figueirido, and D. Woodward (2003). A tractable approach to the firm location decision problem. *The Review of Economics and Statistics* 85(1), 201–204.
- Hanemann, W. M. (1984). Discrete/continuous models of consumer demand. *Econometrica*, 541–561.
- Head, K. and T. Mayer (2013). What separates us? Sources of resistance to globalization. *Canadian Journal of Economics* 46(4), 1196–1231.
- Head, K. and T. Mayer (2018). Poor substitutes? Counterfactual methods in IO and trade compared. mimeo.
- Irrazabal, A., A. Moxnes, and L. D. Oromolla (2013). The margins of multinational production and the role of intrafirm trade. *Journal of Political Economy* 121(1), 74–126.
- Klier, T. and D. P. McMillen (2008). Evolving agglomeration in the U.S. auto supplier industry. *Journal of Regional Science* 48(1), 245–267.
- König, M. D., D. Rohner, M. Thoenig, and F. Zilibotti (2017). Networks in conflict: Theory and evidence from the Great War of Africa. *Econometrica* 85(4), 1093–1132.
- Kucheryavyi, K., G. Lyn, and A. Rodríguez-Clare (2016). Grounded by gravity: A well-behaved trade model with industry-level economies of scale. Technical Report 22484, National Bureau of Economic Research.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725.
- Muffatto, M. (1999). Introducing a platform strategy in product development. *International Journal of Production Economics* 60, 145–153.

- Ramondo, N. (2014). A quantitative approach to multinational production. *Journal of International Economics* 93(1), 108–122.
- Ramondo, N. and A. Rodríguez-Clare (2013). Trade, multinational production, and the gains from openness. *Journal of Political Economy* 121(2), 273–322.
- Santos Silva, J. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics* 88(4), 641–658.
- Schmitt, A. and J. V. Biesebroeck (2013). Proximity strategies in outsourcing relations: The role of geographical, cultural and relational proximity in the european automotive industry. *Journal of International Business Studies* 44(5), 475–503.
- Smith, D. and R. Florida (1994). Agglomeration and industrial location: An econometric analysis of japanese-affiliated manufacturing establishments in automotive-related industries. *Journal of Urban Economics* 36(1), 23–41.
- Tintelnot, F. (2017). Global production with export platforms. *The Quarterly Journal of Economics* 132(1), 157.
- Van Biesebroeck, J., H. Gao, and F. Verboven (2012). Impact of FTAs on canadian auto industry. *DFAIT Canada* 31.
- Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics*, 240–268.
- Wang, Z. (2017). Headquarters gravity. Technical report, Mimeo, Shanghai University.

For Online Publication

A Constant elasticity of substitution discrete choice

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

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

with z the outside good. The model-household parameters ψ_{mh} convert car use into equivalent units of psychological car services.⁴⁷

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

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

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

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

Combining we have

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

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

⁴⁷For example, ψ_{mh} could be the number of driving kilometers expected by the buyer over the lifetime of the model.

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

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

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

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

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

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

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

Aggregate expected sales of model m in n are

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

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

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

B Constructing expected profits from estimates

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

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

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

with the model-entry constant given by

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

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

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

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

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

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

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

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

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

and rewrite equation (B.4) as

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

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

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

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

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

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

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

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

The central parameter of the distribution costs is retrieved as

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

C EHA with a continuous model entry margin

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

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

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

The price index updating function is

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

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

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

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

In the counterfactual, we have the following entry shares:

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

Collecting terms, one can rewrite

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

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

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

Finally, the percent change in number of models offered is

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

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

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

D Data Appendix

D.1 Exclusions from the raw IHS data

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

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

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

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

The time-invariant determinants of frictions (distance, home, contiguity, common language) and GDP per cap variables come from the CEPII gravity database.⁴⁸ Tariff information for both assembled cars and parts comes from the WITS database managed by the World Bank. WITS compiles individual country declarations of their applied MFN and preferential *ad valorem* tariffs, as well *ad valorem* equivalents (AVE) of any specific tariffs. The car tariff is the simple average of the tariffs in HS heading 8703. The car parts tariff is the simple average of the three 4-digit HS headings associated with major components (8706, 8707, and 8708), together with the relevant HS6 categories for engines and associated parts (840733, 840734, 840820, 840991, and 840999). There are many holes in the data which we fill via linear interpolation. When the data is missing for the most recent years, we use the last available year. When a preferential rate exists, we use it. For the rest of dyad-years, we use the MFN tariff inclusive of the AVE of specific tariffs. We also corrected a number of issues in available WITS data regarding recent RTAs that are important for our purposes. Korea signed a number of recent agreements (with the EU and USA in 2012, Canada in 2015, Peru in 2012, Turkey in 2013, Australia in 2015, New-Zealand, Vietnam and China in 2016), for which WITS tariff data is either not accurate or not available. Japan signed RTAs with Peru (entered into force in 2012) and Australia (entered into force in 2015), for which the preferential rates were not mentioned in WITS. Colombia also lacked preferential rates for agreements with the USA and Canada (entry into force in 2012) and the EU (entry into force in 2013). For all those cases (and a few other for which it turned out that cars and parts were mostly exempted), we went to individual text and tariff schedules of those agreements to compute the tariff rate relevant for the bilateral pair in the relevant years. This involves in particular to take into account the “Staging” variable specified (usually giving the number of equal cuts in years) applied to the “Base rate” (the MFN rate at the date of entry into force). Sometimes, even that level of detail is not enough. For instance, the Korea-USA agreement finally decided to postpone the negotiated phasing in by five years (<https://www.uskoreacouncil.org/wp-content/uploads/2014/12/Automotive-Provisions.pdf>). We took those into account when mentioned on the countries’ relevant website. The correction was done both for assembled cars and parts. Note that the staging of tariff liberalization can be very different across bilateral pairs.⁴⁹

⁴⁸Updated for the purpose of the paper, and available at http://www.cepii.fr/cepii/en/bdd_modele/presentation.asp?id=8 or <https://sites.google.com/site/hiegravity/data-sources>.

⁴⁹For example, Colombia’s 10% MFN car engine tariff went immediately to zero for its FTA with the US but took 8 years to expire in the FTA with the EU. Korea reduced its tariffs on assembled cars from the EU and the USA to 0 in a

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

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

E Estimates using the firm-variety approach

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

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

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

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


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

The biggest problem with the Design Parents (DP) is that IHS only reports it as of 2016. Thus, going back in time, it gives incorrect DPs. For example, it makes no sense to think of Tata as the DP for Jaguar cars before 2008 when the brand was owned by Ford. We are able to track ownership changes for brands over time, as the latter often correspond to distinct, stock-selling corporations (e.g. Audi, Nissan). However, it is more difficult to track ownership of platforms. Brands map many-to-many to Design Parents (they map many-to-one to Sales Parents). The reason is that brands market (and even manufacture) platforms designed by other firms. The IHS Engineering Group identifier is very helpful in a few cases (Chrysler-Fiat, Mazda-Ford). For others the brand-platform mapping seems clean enough.

There were two main concerns about the brand-model approach employed in the main text of the paper. The first concern is that much of the low entry rates observed at the brand-model

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

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



level could be an artifact of re-badging strategies. Thus while Honda seems to sell the Legend in Japan only, the same variety is in fact available in many markets as the Acura RL.⁵⁰ Figure E.1 reproduces Figure 2 using the firm-variety concept. The mean entry rate rises, as expected, but only from 0.23 to 0.24. The whole distributions of entry rates are visually very similar.

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

Table E.1 re-estimates the baseline specification from Table 3. The sample size in the sourcing equation (column 1) falls by 20%. There a greater number of possible assembly locations when taking account of all the parent firms' production facilities but there are far fewer sourcing shares when grouping by firm instead of brand. The estimates for the τ_{ln} determinants are very similar to those reported in Table 3: The key elasticity of sourcing response, θ , rounds to 7.7 in both regressions. The imprecision of the estimates of the γ_{il} determinants in the sourcing equations persists with the new set of headquarter i locations. Distance continues to have the wrong sign. Parts tariffs now enter with a highly significant coefficient but it implies a cost share of HQ-provided

⁵⁰Another form of re-badging holds the model name constant while changing the brand name. For example the platform B0, program H79 is sold in roughly equal amounts as a "Duster" under the brand Renault and as a Dacia (a Romanian brand acquired by Renault).

intermediates that exceeds one. In sum, using firm headquarters does not markedly improve the γ estimates. In column (2), the firm average market share equation estimates an η of 1.5, considerably lower than the 3.87 obtained for brands. This η implies markups of 200% that are drastically higher than other estimates in the literature. The implausibly low η estimates suggest that the firm-level market share equation suffers from measurement error in the calculation of C_{bn} . Firms aggregate an often highly heterogeneous set of plants producing very distinct sets of cars. Chevrolet-Opel, Geely-Volvo and Tata-Jaguar are examples of cases where plants from one brand are essentially irrelevant for the other brand's production. Columns (3) and (4) show that deep RTAs promote variety and firm entry but with smaller coefficients and higher standard errors. This corroborates our view that the brand/model concept is more appropriate.

Table E.1: Results with the firm-variety approach

Dep. Var:	$S_{b\ell nt}$	$\frac{q_{bnt}}{M_{bnt}Q_{nt}}$	\mathbb{I}_{mnt}	\mathbb{D}_{bnt}
Method:	PPML	PPML	frac. probit	probit
	(1)	(2)	(3)	(4)
home $_{\ell n}$	0.932 (0.288)			
ln dist $_{\ell n}$	-0.238 (0.075)			
language $_{\ell n}$	-0.069 (0.114)			
ln (1+ car tariff $_{\ell n}$)	-7.741 (0.914)			
Deep RTA $_{\ell n}$	0.220 (0.127)			
home $_{i\ell}$	1.656 (0.397)			
ln dist $_{i\ell}$	0.202 (0.107)			
language $_{i\ell}$	-0.034 (0.321)			
ln (1+ parts tariff $_{i\ell}$)	-11.862 (3.256)			
Deep RTA $_{i\ell}$	-0.369 (0.265)			
ln q_{ℓ}	0.239 (0.063)			
home $_{in}$		0.611 (0.288)	0.251 (0.083)	0.663 (0.520)
home $_{in} \times \text{LDC}_n$		0.967 (0.461)	1.154 (0.115)	3.952 (0.622)
ln dist $_{in}$		-0.381 (0.095)	-0.081 (0.022)	-0.082 (0.129)
language $_{in}$		0.250 (0.200)	0.084 (0.056)	0.032 (0.196)
Deep RTA $_{in}$		0.063 (0.117)	0.079 (0.035)	0.113 (0.154)
ln C_{bn}		-1.509 (1.019)	-0.231 (0.270)	
ln $\mathbb{E}[\pi_{bn}]$				0.929 (0.145)
Observations	281583	21932	21933	62472
rsq	0.783	0.595	0.761	0.705
S.E. cluster:	ℓ	f	f	f

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

F Sourcing parameter estimates from quantities

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

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

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

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

G Fit of the DEV simulation


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

Table F.1: Bilateral brand sales regression provide alternative estimates of γ


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

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


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




(a) Price indices (P_n)



(b) Brand flows (q_{bln})



(c) Trade flows (q_{ln})



(d) National output (q_ℓ)

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

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

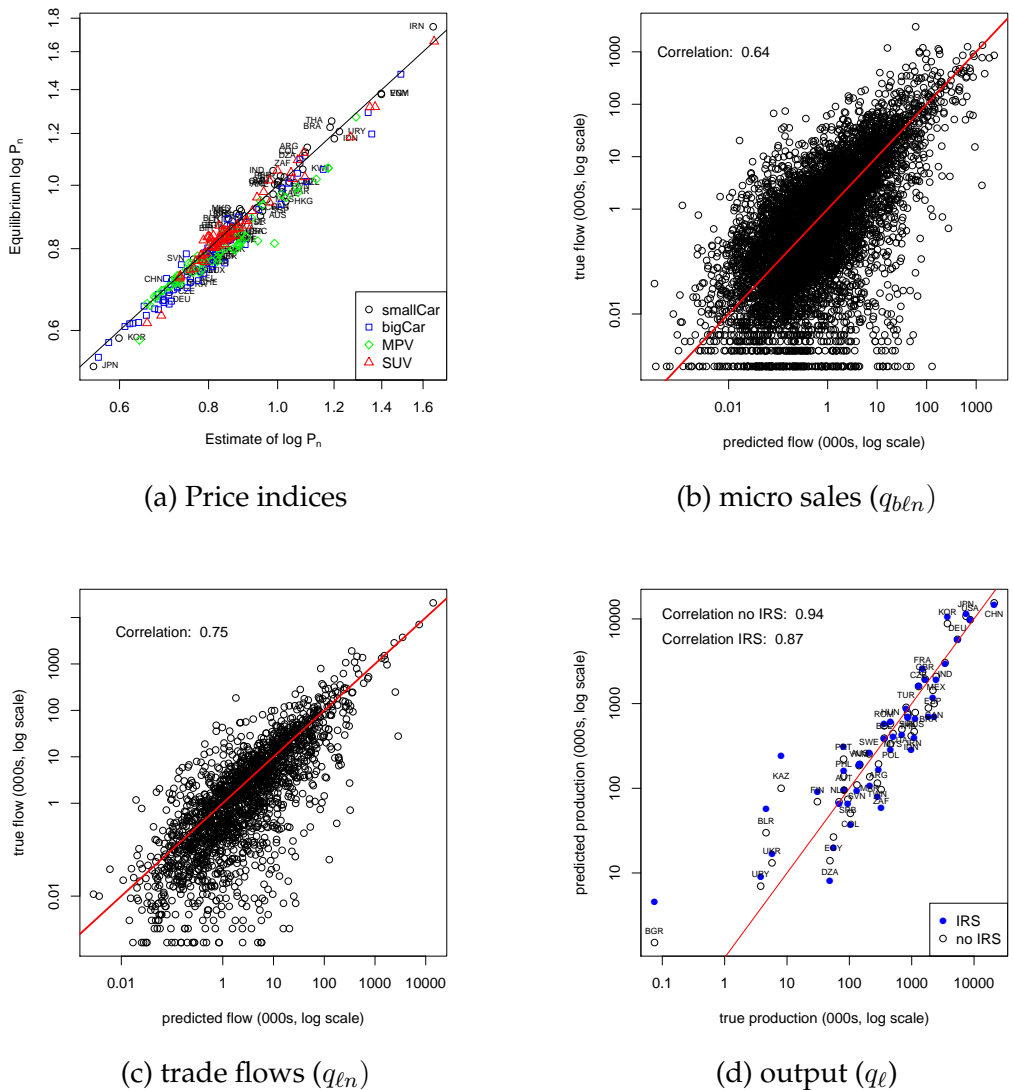


Figure G.2 shows the fit of the segmented market version of DEV. The correlations between the true flows and the flows predicted by the model are very similar and even slightly higher than in the unified market case.

H Evidence on capacity constraints in the medium-run

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

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

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


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

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

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

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

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



The second case of interest are countries producing quantities near their historical maxima. Important examples include China, Korea, Germany, Slovakia, the US, Canada, and Mexico. For these countries to expand output by large amounts they would have to construct new plants, add production lines to existing plants, or increase production per line (either by adding workers to increase line speed or adding extra shifts). To judge how feasible this might be, Figure H.3 decomposes output growth for three countries that have dramatically expanded their car industries in the last 17 years: China, Mexico, and Slovakia. These countries exhibit substantial growth from each of these sources, depending on the time range. In the case of China, rising incomes led to an astounding 32-fold increase in car production between 2000 and 2016, an annual rate of increase of over 23%. While the rate of growth has abated, Chinese car output still managed to grow by 91% from 2010 to 2016, roughly corresponding to our medium-run scenario. It did so mainly by increasing the number of plants (37 were opened in the 7-year period) but output per production line also grew by 20%. In the early 2000s growth was more evenly divided between new plants, new production lines, and expand output per line.

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

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

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

CENTRE FOR ECONOMIC PERFORMANCE
Recent Discussion Papers

1550	Natalie Chen Dennis Novy	Currency Unions, Trade and Heterogeneity
1549	Felipe Carozzi	The Role of Demand in Land Re-Development
1548	Nick Bloom Brian Lucking John Van Reenen	Have R&D Spillovers Changed?
1547	Laura Alfaro Nick Bloom Paola Conconi Harald Fadinger Patrick Legros Andrew F. Newman Raffaella Sadun John Van Reenen	Come Together: Firm Boundaries and Delegation
1546	Jörg Kalbfuß Reto Odermatt Alois Stutzer	Medical Marijuana Laws and Mental Health in the United States
1545	Maria Molina-Domene	Specialization Matters in the Firm Size-Wage Gap
1544	Jonathan Colmer	Weather, Labor Reallocation and Industrial Production: Evidence from India
1543	Andrew B. Bernard Esther Ann Bøler Swati Dhingra	Firm-to-Firm Connections in Colombian Imports
1542	Gerard H. Dericks Hans R.A. Koster	The Billion Pound Drop: The Blitz and Agglomeration Economics in London
1541	Andrew B. Bernard Andreas Moxnes	Networks and Trade

1540	Does Gentrification Displace Poor Households? An ‘Identification-Via-Interaction’ Approach	Sevrin Waights
1539	Zhuan Pei Jörn-Steffen Pischke Hannes Schwandt	Poorly Measured Confounders Are More Useful On the Left than On the Right
1538	Konstantin Büchel Stephan Kyburz	Fast Track to Growth? Railway Access, Population Growth and Local Displacement in 19 th Century Switzerland
1537	Stephen Gibbons Maria Sánchez-Vidal Olmo Silva	The Bedroom Tax
1536	Pawel Adrjan Brian Bell	Pension Shocks and Wages
1535	Tom Kirchmaier Stephen Machin Matteo Sandi Robert Witt	Prices, Policing and Policy: The Dynamics of Crime Booms and Busts
1534	Laurent Bouton Paola Conconi Francisco Pino Maurizio Zanardi	Guns, Environment and Abortion: How Single-Minded Voters Shape Politicians Decisions
1533	Giulia Giupponi Stephen Machin	Changing the Structure of Minimum Wages: Firm Adjustment and Wage Spillovers
1532	Swati Dhingra Rebecca Freeman Eleonora Mavroeidi	Beyond Tariff Reductions: What Extra Boost From Trade Agreement Provisions?

The Centre for Economic Performance Publications Unit
Tel: +44 (0)20 7955 7673 Email info@cep.lse.ac.uk
Website: <http://cep.lse.ac.uk> Twitter: @CEP_LSE