

What Drives Home Market Advantage?*

A. Kerem Coşar^a, Paul L. E. Grieco^b, Shengyu Li^c, and Felix Tintelnot^d

^aUniversity of Virginia, Department of Economics, CEPR and CESifo

^bPennsylvania State University, Department of Economics

^cDurham University, Business School

^dUniversity of Chicago, Department of Economics, and NBER

January 30, 2017

Abstract

In the automobile industry, as in many tradable goods markets, firms usually earn their highest market share within their domestic market. The goal of this paper is to disentangle the supply- and demand-driven sources of the home market advantage. While trade costs, foreign production costs, and taste heterogeneity all matter for market outcomes, we find that a preference for home brands is the single most important driver of home market advantage - even after controlling for brand histories and dealer networks. Furthermore, we also find that consumers favor domestically producing brands even if these brands originated from a foreign country. Therefore, our results suggest a novel demand effect of FDI: Establishing local production increases demand for the brand even in the absence of any cost savings.

*We are grateful to Arnaud Costinot, Michael Dinerstein, Jonathan Eaton, Teresa Fort, Charlie Murry, and Ralph Ossa for insightful conversations, to our discussants Treb Allen, Penny Goldberg, Keith Head, Ferdinando Monte and Mahmut Yasar for valuable comments and suggestions, and to seminar participants at the Princeton IES summer workshop, NBER summer institute and various presentations for comments and suggestions. This work was completed in part with resources provided by the University of Chicago Research Computing Center. Kathie Chuang, Jacob Dorn, Adrian Geilen, Zhida Gui, John Lennox, and Ivy Malao provided excellent research assistance. The authors gratefully acknowledge the support of the National Science Foundation (under grants SES-1459950 and SES-1459905). Coşar wishes to thank the University of Chicago Booth School of Business' Initiative on Global Markets and Stockholm School of Economics, and Tintelnot wishes to thank the Andrew and Betsy Rosenfield Program in Economics, Public Policy, and Law at the Becker Friedman Institute for financial support. Correspondence: kerem.cosar@gmail.com, paul.grieco@psu.edu, shengyu.li@durham.ac.uk, tintelnot@uchicago.edu.

1 Introduction

In tradable goods industries, it is typical for firms to earn their highest market shares in their domestic market. This home market advantage persists despite substantial integration of international markets over the past several decades. There is no shortage of explanations (e.g., trade costs, investment frictions, home preference, taste heterogeneity for characteristics) for this empirical regularity, but different explanations have substantially different policy implications. The goal of this paper is to quantify the sources of home market advantage, and to understand their implications for international trade and investment. In particular, what are the roles of tariff and non-tariff barriers, transportation costs, and foreign production costs in explaining global market outcomes? How important are consumer preferences, either for particular characteristics or simply for national brands? The automobile industry provides an interesting case for analyzing these questions. The industry accounts for over 10 percent of world trade in manufactured goods (WTO, 2013) and bears the features of many oligopolistic industries, producing differentiated and tradable goods, while domestic producers command a dominant share in their home markets.

Apart from the auto industry’s importance in world trade and manufacturing employment, the availability of data also makes this industry suitable for our analysis. We have compiled a rich and unique dataset of global sales and supply. The sales data inform us about prices and quantities (as opposed to revenue only) by automobile model, as well as several characteristics (e.g., horsepower, size, weight and fuel efficiency) in nine countries across three continents. On the supply side, we have worldwide data on the assembly plant locations of each model. We propose a structural model that exploits two features of the data to separate preference-based incentives to purchase local products from supply-side frictions such as trade and investment costs. First, the availability of price data allows us to measure the willingness to pay for models. Second, the prevalence of foreign direct investment (FDI) provides variation between national brand identity and assembly location across models, helping to disentangle their demand and supply effects. Moreover, we are able to separately identify country-level preferences for observed characteristics (such as heterogeneous taste for fuel efficiency across countries) from home preference—a systemic preference for purchasing local products. Such a preference could arise from several sources including consumer information asymmetries between local and non-local products, nationalistic feelings, or a stronger ability on

the part of local brands to respond to localized tastes for unobserved characteristics (e.g., body styling or interior features such as cup holders).¹

Traditionally, models of international trade have relied on relatively restrictive demand systems (e.g. constant elasticity of substitution in Krugman 1980, Eaton and Kortum 2002, Melitz 2003, and Anderson and van Wincoop 2003) to analyze market outcomes.² While these approaches represent tractable means of analysis, they may be limited in their ability to capture rich substitution patterns that are a feature of horizontally differentiated oligopolistic industries such as cars. Quantitative applications have also been limited by the availability of only revenue data without credible price and quantity information. As a result, they may lead to biased estimates of trade costs and an underappreciation of preference differences across national markets. We incorporate a random coefficients approach to estimate demand, allowing for both within- and across-market heterogeneity in consumer preferences. This more flexible approach enables us to consistently estimate demand- and supply-driven mechanisms behind market segmentation.³ The model yields estimates of product-specific markups, trade costs, and production costs from detailed industry level data. This extends, in an industry equilibrium, the analysis of recent quantitative trade models with multinational production (Ramondo and Rodriguez-Clare 2013; Arkolakis et al. 2013; and Tintelnot 2016), which were also limited by the availability of only revenue data on multinationals' foreign affiliate sales for the aggregate manufacturing sector.

Separating the underlying drivers of home market advantage is important for understanding how globalization affects market outcomes and welfare. In the absence of preference-based drivers, one would expect the home market advantage to vanish if all trade barriers were removed. Despite large reductions in formal barriers to trade and foreign investment, as well as in transport and communication costs, significant home market advantages persist. If this is to a large extent accounted for by demand-related factors, the implications for policy and firm strategies are very

¹In a world with trade costs and increasing returns to scale, local producers may obtain larger national market shares for reasons distinct from home preference if they happen to supply the goods that are in high relative demand in the home country (the home market effect in Krugman 1980). Auer (2014) applies this idea to the automotive industry in order to rationalize the sluggish response of trade flows when trade costs fall.

²With the notable exception of non-homothetic preferences (Fieler 2011; Fajgelbaum, Grossman and Helpman 2011; Fajgelbaum and Khandelwal 2014) which are used to rationalize a certain pattern in the trade data, namely the prevalence of north-north trade. However, in that literature consumer preferences are identical across countries, and it is income levels that vary.

³In previous work (Coşar, Grieco and Tintelnot 2015), we estimated supply-side border frictions in wind turbine trade using detailed geographic data on firm sales.

different from models that only consider supply-side frictions. We find that demand-side differences are significant drivers of home market advantage relative to supply side policy frictions. This limits the extent to which trade liberalization can foster market integration; ignoring these limitations would lead policy makers to over-emphasize the benefits of trade liberalization.

Our model builds on the random coefficients demand model developed by Berry, Levinsohn and Pakes (1995). While we follow the general approach in estimating demand and cost side in that workhorse model in the industrial organization literature, we make several departures from the standard analysis: On the demand side, we recover a brand-market-specific shifter to consumers' utilities that we decompose into a rich set of observables such as brand nationality, dealer density, and brand entry date to each market. On the cost side, we use variation in assembly and headquarter locations to estimate trade and foreign production costs, while accounting for firms' endogenous sourcing decision across the set of available plants. This allows us to directly recover the cost elasticities of various trade related observables such as assembly-to-market distance and headquarter-to-assembly distance. Finally, we develop a procedure to unpack the contributions of tariffs, trade/FDI costs, home preference, and taste heterogeneity in explaining market shares.

To quantify home market advantage, we calculate the difference in market share when a model is sold at home versus abroad, controlling for model and market fixed effects. Using our structural estimates, we then evaluate the contributions of potential drivers of home market advantage by computing counterfactual prices and shares after removing various demand and supply components of our model and re-computing the home market advantage statistic. We find that home market advantage is most sensitive to the removal of home preference for domestic brands, declining by about 47 percent. In contrast, when we remove all tariffs, trade, and foreign production costs, the home market advantage declines by only 19 percent. Importantly, our results are obtained by controlling for the impact of a brand's entry date into a market, dealer density, and whether the brand has any local assembly. Since, for historical reasons, these factors are correlated with home status, treating them as unobserved would lead to even larger estimates of home preference. While the trade literature has often used trade costs and home bias in preference interchangeably, one of

the key contributions of this paper is to disentangle the role of these two forces.⁴

In addition to the importance of home preference for brands, we also uncover a novel demand impact of foreign production. Since brands typically assemble models in a variety of locations, we are able to separate how local production of *any* model affects local taste for the brand from the usual cost effect of local assembly for a particular model. We find this preference for brands with domestic assembly plants, regardless brands' country of origin, to be significant and roughly two-thirds of the home preference effect. As such, the domestic assembly demand effect represents an intriguing rationale for FDI beyond reducing trade costs.

A number of papers have looked at the car industry to study questions in international trade. Among them, Feenstra (1988), Goldberg (1995), and Berry, Levinsohn and Pakes (1999) analyze the effects of Japanese voluntary export restraints on the American auto market; Goldberg and Verboven (2004) use panel data from the car industry to demonstrate a strong positive effect of the Euro on price convergence; Brambilla (2007) investigates firms' responses to trade policy in South America; McCalman and Spearot (2013) study firms' offshoring strategies using data on North American light truck production locations; and Miravete, Moral and Thurk (2015) analyze the role of emission policy as an entry barrier using Spanish car market data.

Closer to our work, Goldberg and Verboven (2001) study price dispersion in the European car market and also find evidence that consumers favor national brands. Their paper, however, does not quantify the importance of demand versus cost drivers for market segmentation. By analyzing data beyond the European market and accounting for dealer networks and brand-market histories, we bring in important variation, such as larger differences in gas prices, supply barriers, and potentially also in preferences across the Atlantic.

In contemporaneous work, Head and Mayer (2015) estimate a model of international trade and foreign production using sourcing data from the automotive industry, finding that foreign sales are impeded significantly by trade, foreign production, and multinational sales frictions. Similar to our paper, the latter captures a general disadvantage in selling outside the home market. Our

⁴See for example the following quote from the gravity model review article by Anderson (2011): "It is usual to impose identical preferences across countries. [...] Henceforth trade cost is used without qualification but is understood to potentially reflect demand-side home bias. Declines in trade costs can be understood as reflecting homogenization of tastes."

contribution is to disentangle the demand- and supply-related components of this disadvantage.⁵

Our findings are related to Atkin (2013) who demonstrates that neglecting taste heterogeneity for food leads to biased estimates of gains from trade. They also complement an extensive literature in marketing illustrating the importance of brand preferences (Bronnenberg, Dhar and Dube 2009; Bronnenberg, Dube and Gentzkow 2012). Relative to these papers, we contribute by jointly analyzing cost and preference differences across markets and quantifying their impact in an international context. While previous work has documented that each of these channels is important, our work is the first to directly compare the impact of tastes and costs on home market advantage.

The next section describes the data and presents the stylized facts motivating our analysis. Section 3 formulates a model of international competition in the automobile market. We estimate the model in Section 4 and evaluate the drivers of home market advantage in Section 5.

2 Data and Descriptive Evidence

Our dataset covers the market for passenger cars in 6 EU countries (Belgium, Germany, France, Italy, Spain, Great Britain), Brazil, Canada and the US for the period 2007-2011.⁶ For each available market-year, we observe model-level sales (i.e. number of new cars sold), prices (MSRP) and product characteristics (such as length, width, weight, and fuel efficiency). The characteristics of a model are the sales-weighted average trim-level characteristics and vary across markets and years. We also constructed a dataset of assembly countries for each model-year in the demand data. Complementing these datasets are market-specific variables such as gas prices, import tariffs on cars, sales taxes, the level and dispersion of household income, as well as brand-market-specific variables such as a brand’s entry date into a market and the number of car dealers in the market. Appendix A describes the construction of the dataset.

Some manufacturing groups own multiple brands. In what follows, we distinguish firms (manufacturing groups such as General Motors (GM) and Volkswagen (VW)), brands (such as Chevrolet, Vauxhall, and Opel owned by GM; Audi and Seat owned by VW) and models (such as Chevrolet

⁵While their dataset reports model level flows between all assembly locations and many more destination markets, our dataset has more detailed information about the product—most notably, prices and characteristics—in fewer yet important markets. Having information on prices and quantities separately is key in identifying demand- from supply-related factors.

⁶Brazilian market data is missing for 2007, and Canadian data is available for 2008-2009 only. Total sales cover more than 90% of total new passenger car sales in the European markets and 80% of sales in the American markets.

Table 1: Market concentration

	Sales	Firms	Top 5	Brands	Top 5	Models	Top 5
BEL	496,171	20	0.68	39	0.44	314	0.13
BRA	2,635,547	17	0.82	24	0.81	101	0.36
CAN	1,137,750	16	0.65	34	0.50	207	0.22
DEU	3,011,988	20	0.71	38	0.54	323	0.18
ESP	1,082,868	21	0.72	39	0.44	290	0.16
FRA	2,046,077	20	0.81	38	0.65	271	0.25
GBR	2,026,497	22	0.63	39	0.47	311	0.21
ITA	2,016,114	22	0.70	41	0.51	283	0.26
USA	10,390,308	19	0.68	40	0.53	291	0.14

Notes: Average number of passenger cars sold annually in each country over the data period. Market shares by top manufacturing group (firms), brands and models are revenue-based.

Malibu and Opel Corsa). In cases where a firm owns foreign brands, we distinguish the headquarter country from a brand’s nationality. For example, GM is a US firm, but Vauxhall and Opel are British and German brands, respectively. In other words, a brand’s nationality is defined as the country from which it historically originates. Accordingly, a brand enjoys home status in its country of origin (such as Opel in Germany). Across all years and markets, the dataset encompasses 28 firms, 60 brands and 598 models. Firms are headquartered in 12 different countries, and brands are associated with 15 different countries.

We now document a number of facts that guide our modeling choices. The oligopolistic nature of the car industry is well known. While measures of concentration vary across markets (Table 1), the top 5 firms account for an average of 55% of total revenues across all market-years. Similarly, the market share of the top 5 brands is 36%.

Table 2 presents market shares by brand nationality. The diagonal in bold highlights the dominance of home brands (Belgium, Brazil and Canada do not have a national brand in our dataset). Spanish and British brands have marginal sales outside of their markets. Similarly, Italian brands have low sales in other European markets but a stronger presence in Brazil due to FDI by Fiat in South America. The most striking difference is between Germany and France: in both markets, home brands account for more than half of the sales, whereas German brands’ market share in France is only 19%, which is relatively higher than the French market share of 9% in Germany.

Table 2: Market shares by brand nationality

	Market share of brands from						
	DEU	ESP	FRA	GBR	ITA	USA	Other
BEL	0.34	0.02	0.26	0.02	0.04	0.09	0.23
BRA	0.23	-	0.11	0.00	0.22	0.31	0.13
CAN	0.07	-	-	0.01	-	0.34	0.58
DEU	0.55	0.02	0.09	0.01	0.03	0.08	0.21
ESP	0.26	0.09	0.26	0.01	0.03	0.11	0.22
FRA	0.19	0.02	0.52	0.01	0.04	0.07	0.16
GBR	0.23	0.02	0.13	0.18	0.02	0.16	0.25
ITA	0.24	0.01	0.15	0.02	0.30	0.12	0.17
USA	0.08	-	-	0.01	0.00	0.40	0.52

Notes: Each row presents the revenue-based market share of brands originating from countries listed in the columns, adding up to one subject to rounding error. - means that brands from the origin country are not sold in the market, and 0.00 implies a market share of less than one percent. Other includes Japan, Korea, China, India, Sweden, Malaysia, Czech Republic, Romania and Russia.

Brands' differential market shares across countries are driven by an extensive margin of model offerings as well as an intensive margin of sales per model. In order to decompose these two margins, we follow Bernard et al. (2009) and start with the identity

$$s_{bmt} = \bar{s}_{bmt} \cdot N_{bmt},$$

where s_{bmt} is the share of brand b in total market m sales in year t , and N_{bmt} is the number of models offered. We then separately project $\ln(N_{bmt})$ and $\ln(\bar{s}_{bmt})$ on $\ln(s_{bmt})$. Table 3 reports the results. The intensive margin accounts for 58 to 62 percent of the overall variation, depending on whether the market share is in revenues or units sold. Variation in the popularity of brands across countries is not simply due to the number of products offered.

To gauge the extent of the home market advantage, we project market shares on a dummy variable that takes the value one if a model is at home and zero otherwise. Table 4 presents the results. The first two columns are the brand-market-year (bmt) level, while the last column is at the model-market-year (jmt) level. Given the importance of the extensive margin documented above, column II also controls for the (log) number of models that the brand offers. Fixed effects control for brands' and models' global popularity and market-year specific conditions. We find a large and

Table 3: Market share decomposition

	I	II	III	IV
	$\ln(\bar{s}_{bmt})$	$\ln(N_{bmt})$	$\ln(\bar{s}_{bmt})$	$\ln(N_{bmt})$
$\ln(s_{bmt})$	0.618 (0.006)	0.382 (0.006)	0.577 (0.007)	0.423 (0.007)
Observations	1473	1473	1473	1473
R^2	0.810	0.619	0.780	0.654
Share	Units	Units	Revenue	Revenue
Margin	Intensive	Extensive	Intensive	Extensive

Notes: Standard errors in parentheses. This table shows the share of variation in s_{bm} coming from the intensive margin (average model market share of the brand) and the extensive margin (number of models offered by the brand). Accordingly, columns 1 and 2 as well as columns 3 and 4 add up to one. All regressions are estimated with year fixed effects.

significant home market effect: being a home brand increases market share at the model level by 238 percent.⁷ We label this effect “home market advantage.”

The data also exhibits significant cross-country variation in product offerings and sales. Table 5 presents average prices and characteristics (weighted by sales) of cars sold in each market. The average car in the North American market is larger in horsepower and size (columns 2 and 3) and less fuel efficient (column 4) than the typical car sold in Europe and Brazil. Differences in gas prices, however, affect the average cost of a mile (last column) that consumers face in each market. While some of this variation is due to the extensive margin, characteristics also differ within models. Controlling for model fixed effects, we see less powerful engines and much higher fuel economy in European countries and Brazil compared to the North American market (see Table B.2 in Appendix B). These systematic differences highlight the importance of controlling for product characteristics when estimating consumer demand.

On the production side, there are 50 countries that assemble cars. 43% of the models (255 out of 598) are assembled in more than one country, accounting for 64% of total revenue. The market share of models assembled in 5 or more countries is 30% (see Figure B.1 in the appendix). Our data inform us about the set of countries in which each model is assembled, but we do not know exactly which assembly plant was used to serve which market. Column 1 of Table 6 presents the average

⁷We calculate home market advantage using $100 \cdot (\exp(\hat{c}) - 1)$, where \hat{c} is the coefficient on the dummy for being a home brand. If we follow Kennedy (1981) and van Garderen and Shah (2002) and calculate home market advantage using $100 \cdot (\exp(\hat{c} - \frac{1}{2}V(\hat{c})) - 1)$, where $V(\hat{c})$ is the estimate of its variance, we obtain nearly identical results, since the coefficient is very precisely estimated.

Table 4: Home market advantage

	I	II	III
	$\ln(s_{bmt})$	$\ln(s_{bmt})$	$\ln(s_{jmt})$
Home brand	1.679 (0.082)	1.066 (0.061)	1.219 (0.032)
$\ln(N_{bmt})$		1.533 (0.042)	
Observations	1473	1473	8841
R^2	0.793	0.895	0.720
Market-year FE	Yes	Yes	Yes
Brand FE	Yes	Yes	
Model FE			Yes

Notes: Standard errors in parentheses. All regressions are estimated with market-year (mt) fixed effects. The first two columns are at the brand-market-year (bmt) level and use brand fixed effects. The last column is at the model-market-year (jmt) and uses model fixed effects.

number of countries in which models consumed in a particular market are assembled (weighted by sales). For instance, there are 3.8 countries in which models sold in Brazil are assembled, while models sold in Canada are assembled in 5 countries. Geography, units costs and trade policy are important determinants of these potential supply locations. Brazil is the most protected country in our dataset, with an MFN import tariff on cars equal to 35%, and the US is the most open, with an MFN import tariff of 2.5%.⁸ (column 4). The resulting tariff-jumping FDI leads to a higher market share for the models that are assembled domestically (column 3) in countries with high tariffs. For example, the market share of models with domestic assembly plants is 86 percent in Brazil. There is also substantial variation across markets in the distance to the nearest assembly plant.⁹

To gain some initial insight into how model characteristics and supply locations correlate with prices, Table 7 presents the results of a hedonic price regression with model and market-year fixed effects. While this regression is unable to distinguish demand and supply effects from each other, it does give an indication of how equilibrium prices move with model characteristics. As

⁸The so-called US “chicken tax” of 25% applies on light trucks, which we exclude from our analysis. SUVs are imported as passenger cars.

⁹The distances in Table 6 include internal distance within a country. We use bilateral and internal distances from the CEPII data (Mayer and Zignago 2011) calculated as population-weighted distances between the biggest cities of two countries. Internal distances range from 66 km in Belgium to 1,853 km in the U.S.

Table 5: Prices and characteristics

	Price	HP/Wt	Size	MPG	Gas price	MPD
BEL	27,078	58.4	7.6	34.4	5.2	6.6
BRA	13,808	62.1	6.8	30.1	2.1	14.4
CAN	24,301	91.9	8.3	22.2	2.9	7.7
DEU	30,981	66.8	7.6	29.3	5.9	4.9
ESP	26,143	60.8	7.6	32.6	4.7	7
FRA	25,161	57.2	7.3	35.5	5.3	6.7
GBR	28,859	65.5	7.5	30.4	5.5	5.6
ITA	24,252	57.6	7	33.4	5.3	6.4
USA	26,363	97.9	8.7	20.8	2.5	8.3

Notes: All variables are averages across models weighted by market share over the data period. Prices are deflated to their 2005 local currency levels and converted to USD at that year's exchange rate. HP/Wt denotes horsepower per weight (kg) times 1,000. Size is meter length times meter width. MPG is miles per gallon. Gas prices are per gallon in 2005 USD. MPD is miles per dollar (MPG/price). The Appendix describes data sources and tabulation procedures.

Table 6: Supply locations

	Supply Locations	Average Distance	Domestic Share (%)	MFN Tariff (%)
BEL	4.4	1095	9	10
BRA	3.8	1897	86	35
CAN	5	3669	26	6.1
DEU	4.7	1077	51	10
ESP	4.5	1627	33	10
FRA	4.2	1027	41	10
GBR	4.5	1440	17	10
ITA	4	1333	20	10
USA	4.2	3625	54	2.5

Notes: All figures weighted by market share. Supply location is the average number of countries in which models sold in a market are assembled, weighted by models' market share. Average distance is the average distance across models to the closest supply location including internal distance, see footnote 9. Domestic share is the market share of models which have an assembly plant in the country. Implied internal distances capture differences in land area across countries. MFN (most favored nation) is the non-discriminatory tariff rate applied to WTO members that are not in a free trade agreement with the country.

expected, horsepower per weight, size and fuel efficiency are all associated with higher prices. The coefficient on distance to nearest assembly is also positive and significant. Under the assumption that consumers do not care about distance to assembly after controlling for other factors, this result implies that assembly distance has some power as an instrument for demand. Most interestingly, a model being assembled at home is negatively correlated with price, while a model being associated with a home brand is positively correlated with price. These two features can be separated due to FDI: some home brands' models are produced abroad and imported while some foreign brands' models are assembled within the market. This finding supports our assumption in the structural model below that brand nationality affects the demand for cars, while car models' assembly location affects their production costs.¹⁰

Table 7: Price regression

	$\ln(\text{price}_{jmt})$
$\ln(\text{hppwt}_{jmt})$	0.267 (0.011)
$\ln(\text{size}_{jmt})$	0.544 (0.039)
$\ln(\text{mpd}_{jmt})$	0.029 (0.010)
$\ln(\text{dist}_{jmt})$	0.019 (0.002)
Domestic assembly by model	-0.013 (0.004)
Domestic assembly by brand	-0.007 (0.003)
Home brand	0.023 (0.004)
Observations	8841
R^2	0.985
Market-year FE	Yes
Model FE	Yes

Notes: See table 5 for the description of parameters. Home brand is one if a model belongs to a national brand and zero otherwise. Local assembly by model is one if there is an assembly plant for a model in the country and zero otherwise. Domestic assembly by brand is one if there is at least one model owned by the same brand and assembled in the country, and zero otherwise. Regression controls for market-year and model fixed effects.

¹⁰As is well known, the MSRP can differ from the actual transaction price. In Appendix B.1, we analyze online car sales from three countries. We document that while brands differ in the extent to which they give discounts on MSRP, they do not give systematically larger discounts at home than they do abroad.

3 Model

We model the national market for cars in a given year. We first give a brief overview of the model describing the assumptions on the timing of actions. We then discuss demand and supply in more detail in the subsections below.

Firms (e.g. Toyota) are endowed with a set of brands (e.g. Toyota, Lexus) which in turn are endowed with a set of models (e.g. Toyota Corolla, Toyota Camry) to sell within the market. We take as given the headquarter location of a firm and the country of origin of a brand. Each model is endowed with a set of characteristics (e.g., size, fuel efficiency) and a set of assembly locations where the model can be produced. At the start of the year, all firms observe a set of demand and supply shocks for each model that are uncorrelated with model or assembly location characteristics. This assumption implies that a firm chooses to offer a car in a location before observing the model-market demand or supply shock. This assumption—which is common to many random coefficients demand estimations (e.g., Berry, Levinsohn and Pakes, 1995)—is reasonable in our context because while it is relatively easy to adjust a car’s price in reaction to local market conditions, the decision to release a model in a country generally involves a significant period of design and marketing research prior to entry. Similarly, moving the assembly of a certain model to a plant requires a planning and retooling time. Having observed their own and competitors’ demand and supply shocks, firms simultaneously choose prices at the model level according to a Nash-Bertrand equilibrium – taking into account cannibalization across brands within the firm. We follow the literature on automobile pricing in assuming that prices are set at the model level and consumers face a single price. Consumers then observe these prices and make purchases. Finally, automakers select the assembly location from which to source ordered cars. We allow for heterogeneity in production costs at the car-assembly location level so that a firm may choose to source cars from multiple assembly locations to supply the same model to a market. In addition to allowing for potential gains to variety in assembly locations, cost heterogeneity makes the model much more computationally tractable to estimate.

3.1 Demand

The utility to consumer i in market m from purchasing model j is,¹¹

$$u_{jmi} = \bar{u}(x_{jm}, p_{jm}, \beta_{mi}, \alpha_{mi}) + \xi_{jm} + \varepsilon_{jmi} \quad (1)$$

where x_{jm} represents the model characteristics—e.g. horsepower per weight, size, fuel efficiency—and p_{jm} represents the price. Importantly, brand-country fixed effects are also included in x_{jm} . These brand-country fixed effects capture everything that varies at the brand-country level such as dealer networks, the year of a brand’s first entry in a market, or whether the brand first originated in the country. We assume that consumers are indifferent between the assembly locations for a particular model, hence the assembly location used to serve model j to market m does not enter x_{jm} .¹² Consumers are, however, allowed to have different preferences for brands that have local production (for example, consumers may prefer to support local manufacturing employment). The terms β_{mi} and α_{mi} represent tastes for characteristics and price sensitivity, respectively. Differences in β_{mi} across individuals and countries may arise due to innate preferences or differences in the prices of complementary goods such as parking space and gasoline. Each model receives a market-year-specific demand shock, ξ_{jm} , which is common to all consumers within a market. Finally, each consumer receives an idiosyncratic utility shock for each model, ε_{jmi} , which is distributed according to the Type-I extreme value distribution.

Consumers in each market observe the set of available products and choose the model that maximizes their utility from all available models and a no-purchase option. We normalize the utility of the no-purchase option to $u_{0mi} = \varepsilon_{0mi}$ and let C_m be the set of cars consumers can choose from within market m . Each consumer chooses the option that maximizes her utility,

$$d_{mi} = \operatorname{argmax}_{j \in C_m \cup 0} u_{jmi}.$$

Integrating out the idiosyncratic consumer taste shock, the probability that a consumer with tastes

¹¹For readability we omit the time subscript, t , from the model section.

¹²This assumption is reasonable because if consumers indeed cared about the assembly location of a particular model, manufacturers should offer different MSRPs for the same model produced in different locations, which we don’t observe in the markets that we study.

$(\beta_{mi}, \alpha_{mi})$ buys model j is,

$$Pr(d_{mi} = j | \beta_{mi}, \alpha_{mi}) = \frac{e^{\bar{u}_{jmi} + \xi_{jm}}}{1 + \sum_{k \in C_m} e^{\bar{u}_{kmi} + \xi_{km}}}.$$

Market shares for model j then can be calculated by integrating these individual-specific probabilities over the distribution of consumer tastes in the market:

$$s_{jm} = \int Pr(d_{mi} = j | \beta_{mi}, \alpha_{mi}) dF_m(\beta_{mi}, \alpha_{mi} | \theta^d). \quad (2)$$

The demand parameters, θ^d , govern the distribution of tastes.

3.2 Supply

Firms supply models to consumers by sourcing them from available assembly locations, which were determined prior to demand and cost shocks being revealed to the firms. The marginal cost of sourcing a car i of model type j for market m from location ℓ is,

$$c_{jmli} = c_1(h_{jm}, \kappa) c_2(g_{jm\ell}, \delta) e^{\omega_{jm} - \nu_{jmli}} \quad (3)$$

where $c_1(\cdot)$ represents model- and market-specific costs of selling model j in market m , which are determined by a vector of observable model and market characteristics h_{jm} (such as global production costs of the vehicle and local distribution costs) and a vector of parameters κ . Similarly, $c_2(\cdot)$ represents the effect of costs due to sourcing model j from an assembly plant in location ℓ to be sold in market m . It depends on a vector of known market-assembly-model characteristics $g_{jm\ell}$ (such as distance to the sourcing country from the market and the firm's headquarter location and productivity in the assembly location) and a vector of parameters δ . The structural error term ω_{jm} represents a shock to the marginal costs of selling model j in a given market m . Finally, costs at the car level are affected by an idiosyncratic shock, ν_{jmli} . This final cost is revealed to the manufacturer at the time a car is ordered, after prices for models are set. Producers have full knowledge of ω_{jm} and all other cost shifters besides ν_{jmli} when setting prices. As we show below, the idiosyncratic error ν_{jmli} introduces the possibility of “gains from diversification” in assembly locations and rationalizes the possibility that some models are sourced from multiple assembly

locations.

The manufacturer minimizes costs by sourcing cars from the lowest cost location from its set of available assembly locations, L_j ,

$$c_{jmi} = \min_{\ell \in L_j} c_{jm\ell i}.$$

However, the firm must set prices prior to the $\nu_{jm\ell i}$ shock being revealed; therefore it must set prices according to its expected cost of supplying a model by integrating over $\nu_{jm\ell i}$. We assume $\nu_{jm\ell i}$ is distributed Type-I extreme value with scale parameter σ_ν .¹³ Given this assumption, the probability of sourcing a car from location ℓ is,

$$Pr(i \text{ is sourced from } \ell) = \frac{c_2(g_{jm\ell}, \delta)^{-1/\sigma_\nu}}{\sum_{k \in L(j)} c_2(g_{jmk}, \delta)^{-1/\sigma_\nu}}, \quad (4)$$

where we exploit the fact that minimizing cost is equivalent to maximizing the negative of the logarithm of cost. Therefore, the logarithm of the average marginal cost to sell a car of model j is,¹⁴

$$\log c_{jm} = \log c_1(h_{jm}, \kappa) - \sigma_\nu \log \left(\sum_{k \in L(j)} \exp \left(\frac{-\log c_2(g_{jmk}, \delta)}{\sigma_\nu} \right) \right) + \omega_{jm}. \quad (5)$$

Or equivalently, the average marginal cost for model j in market m is,

$$c_{jm} = c_1(h_{jm}, \kappa) \left(\sum_{k \in L(j)} c_2(g_{jmk}, \delta)^{-1/\sigma_\nu} \right)^{-\sigma_\nu} \exp(\omega_{jm}). \quad (6)$$

In these expressions, the second term captures the fact that manufacturers endogenously choose to source cars from the lowest cost locations. The intuition behind this formula is straightforward. Lower cost locations are more likely to be used as sources, which is reflected in the fact that they contribute the most to the sum over locations. Moreover, as more locations are added, this sum increases, further reducing costs. The expression for costs is computationally tractable, since it is differentiable in all parameters. The value of σ_ν captures “gains from variety” in the sense that

¹³We could relax the assumption that $\nu_{jm\ell i}$ is independent across i at the cost of additional notation. For example, we could divide the year into a large number of discrete time sub-periods and let each consumer who purchases a car within a sub-period receive the same draw of $\nu_{jm\ell i}$. This would be consistent with the shock reflecting unanticipated backlogs or shocks to assembly location productivity during the year.

¹⁴A constant from integrating the Type-1 extreme value distribution is absorbed in $\log c_1(h_{jm}, \kappa)$.

the value of an additional assembly location is scaled by σ_ν . Furthermore, as $\sigma_\nu \rightarrow 0$, firms always source from the single location that has the lowest average cost, and (5) becomes,

$$\lim_{\sigma_\nu \rightarrow 0} \log c_{jm} = \log c_1(h_{jm}, \kappa) + \min_{k \in L(j)} \left\{ \log c_2(g_{jmk}, \delta) \right\} + \omega_{jm}.$$

So as $\sigma_\nu \rightarrow 0$, only variation in g_{jmk} at the minimum cost location affects the marginal cost of a model. However, in the limit as $\sigma_\nu \rightarrow 0$, average costs are no longer differentiable with respect to the source-destination-specific cost parameters.¹⁵ The supply-side parameters to estimate are $\theta^s = (\delta, \kappa, \sigma_\nu)$.

3.3 Pricing Equilibrium

Firms choose prices to maximize profits given demand and the average marginal cost of a model c_{jmt} , which is determined by the cost minimization across available assembly locations as described above. Since a mass of consumers purchases cars, c_{jmt} is exactly known to firms when they set prices, even though they do not know $\nu_{jml\hat{i}}$ until consumer i purchases a car. For the same reason, firms know from (2) exactly what the shares will be given a vector of prices within the market p_m . Therefore, firm f 's profit maximization problem is to choose prices for its portfolio of models within a market $J_m(f)$ to maximize profits,¹⁶

$$\max_{\{p_{jm}\}_{j \in J_m(f)}} \sum_{j \in J_m(f)} [p_{jm} - c_{jm}] \cdot N_m \cdot s_{jm}(p_{jm}; p_m^{-j}), \quad (7)$$

where N_m is the exogenous number of potential buyers and p_m^{-j} is the vector of prices for models other than j . A Nash-Bertrand equilibrium strategy profile is a vector p_m such that $s_{jm} = s_{jm}(p_m)$ and all firms are maximizing profits. Therefore, prices satisfy the system of first order conditions for every price, p_{jm} .

$$s_{jm}(p_m) + \sum_{k \in J_m(f)} [p_{km} - c_{km}] \frac{\partial s_{km}(p_m)}{\partial p_{jm}} = 0. \quad (8)$$

¹⁵In fact, in our application we find that σ_ν is close to zero, so this error primarily plays a technical role in convexifying the firms' sourcing problem

¹⁶We control for differences in tax regimes (e.g., value added versus retail sales tax) across markets using country market dummies in the specification of costs. The model could be extended to explicitly account for differing tax regimes given stronger assumptions about how the base amount for these regimes is determined.

3.4 Identification

Before turning to the estimation section, we discuss how the parameters of our model are identified, with a particular focus on the identification of the cost-side parameters. As we show below, these parameters are identified even though we do not directly observe the source-destination-specific trade flows of cars. Identification of the demand-side parameters is relatively standard: the demand parameters are identified via moment condition assumptions on the model-market demand shocks ξ_{jm} . As shown by Berry (1994), there is a one-to-one mapping between the demand shocks and observed market shares given demand parameters and observed prices. So, given a demand parameter θ^d , we can numerically recover the complete vector of demand shocks within a market,

$$\xi_{jm} = s_j^{-1}(s_m, p_m; \theta^d).$$

We then identify the model using a vector of instruments z_{jm} such that $E[\xi_{jm}z_{jm}] = 0$.¹⁷ The model precludes price from being used as an instrument since it is endogenously determined.

With the demand parameters identified, we are able to recover marginal cost for each model by inverting the firms' first order conditions at observed prices and shares as in Nevo (2001). For clarity, we suppress the market subscripts and focus on a single market. Given demand parameters and observed prices and shares, all the terms in (8) are known with the exception of the vector of marginal costs c_m . Note that firms internalize their cross-price effect on other models that they sell, but not on competitor models. If we define Ω such that,

$$\Omega_{jk} = -\frac{\partial s_k(p_m)}{\partial p_j} \cdot \mathbf{1}[j, k \text{ jointly owned}],$$

then we can write (8) in vector notation, $s(p_m) - \Omega(p_m - c_m) = 0$, and we can easily solve this for the vector of marginal costs,

$$c_m = [p_m - \Omega^{-1}s(p_m)].$$

Once costs are recovered, we can identify the cost-side parameter θ^s from (5) and the assumption that $E[\omega_{jm}|(h_{jm}, g_{jml})] = 0$. While identification of model- and market-specific costs, κ , is straightforward given regularity conditions that will be satisfied by our parameterization, the

¹⁷We discuss the the demand instruments we use in estimation in Section 4.1.

contribution of trade- and location-specific production costs, parameterized by δ and σ_ν , is more subtle.¹⁸ First, consider identification of σ_ν , the variance of the idiosyncratic car cost shock. Suppose that all assembly locations were identical and geography was symmetric, that is, for a given model, $c_2(g_{jml}, \delta) = \bar{c}_2$. In this case, the only reason to source from a particular location would be due to the extreme value error, ν_{jml} . There would be a cost advantage to operating multiple assembly locations in that the manufacturer would get a new draw of this idiosyncratic cost shock for each location. Therefore, the extent to which marginal costs decline as we vary the number of production locations identifies σ_ν . In the extreme, suppose $\sigma_\nu = 0$. Then, an additional assembly location will not reduce marginal costs at all. With σ_ν identified, we can identify the parameters on assembly location characteristics from the variation in these characteristics. This variation will affect average costs in two ways. First, it will change the cost associated with that assembly location conditional on it being used, and second, it will change the probability that the plant is used to source cars. Again, consider the extreme case when $\sigma_\nu = 0$. Then, *only* variation in the lowest cost assembly location's characteristics affect average costs. Therefore, variation in an assembly characteristic across locations first identifies which is the lowest cost and then identifies the parameter for that characteristic based on the size of the change in cost. If σ_ν is positive, then variation in characteristics of all assembly locations affect costs, but its impact is weighted by the quantity of cars each location provides. In summary, each element of δ is identified as long as it affects $c_2(g_{jml}, \delta)$ for some model j where trade flows are positive between market m and assembly location ℓ . This is the case even though we do not directly observe trade flows because we can use variation in model costs c_{jm} and g_{jml} to infer the effect of δ .

4 Estimation

The model is estimated in two stages. We first estimate the demand side, and use firms' profit maximization conditions and the estimated demand parameters to recover the marginal cost of supplying each model to each country. We then use these recovered costs to estimate the supply side.¹⁹ Since we use data from multiple years, we introduce a time subscript, t , below.

¹⁸We are assuming a location normalization in $c_2(g_{jmk}, \delta)$, as is common in discrete choice models, without loss of generality. A scale normalization on σ_ν is not necessary as we explain below.

¹⁹In principle, demand and supply can be estimated jointly, which would improve efficiency at the cost of computational tractability.

4.1 Demand Parameterization and Estimation

We start by parameterizing the utility function to be quasi-linear in price and car characteristics, and quadratic in tastes for size:

$$\begin{aligned} \bar{u}(x_{jmt}, p_{jmt}, \beta_{mi}, \alpha_{mti}) = & \beta_{mi}^{hp} \text{hppwt}_{jmt} + \beta_{mi}^{md} \text{mpd}_{jmt} + \beta_{mi}^{sz} \text{size}_{jmt} + \beta_m^{sz2} \text{size}_{jmt}^2 \quad (9) \\ & - \alpha_{mti} p_{jmt} + \iota_{mti} + \psi_{mb(j)}, \end{aligned}$$

where hppwt_{jmt} is the horsepower of the car divided by its weight (a measure of acceleration capability), size_{jmt} is the size of the car (length times width in meters), and mpd_{jmt} is miles per dollar at market price for gas (according to city fuel efficiency rating). This specification allows for consumers' marginal taste for size to increase or decrease in the amount of the characteristic provided. We would expect the marginal utility of size to decrease as a car gets larger.

We allow tastes for characteristics to vary by country. First, we assume that the tastes for hppwt_{jmt} and size_{jmt} are normally distributed with country-specific means and common variance,²⁰

$$\begin{bmatrix} \beta_{mi}^{hp} \\ \beta_{mi}^{sz} \end{bmatrix} \sim N \left(\begin{bmatrix} \bar{\beta}_m^{hp} \\ \bar{\beta}_m^{sz} \end{bmatrix}, \begin{bmatrix} \sigma_{hp}^2 & 0 \\ 0 & \sigma_{sz}^2 \end{bmatrix} \right),$$

while the quadratic parameter for size is market-specific but constant across consumers within a market.²¹

Second, we allow consumers preferences for miles – a proxy for cost of ownership – to be income dependent.²² Specifically, consumers tastes for miles per dollar are also normally distributed according to,

²⁰In principle we could allow the variances to vary by market, however because they enter the objective function in a nonlinear way, doing so would greatly increase the computational complexity of estimation. Moreover, since we have only 4-5 years of data from each market it is unclear that these parameters could be precisely estimated at the market level.

²¹This specification leads to the intuitive interpretation that a model with size_{jmt} provides a marginal utility for size to the median consumer within market m as given by

$$\text{med} \left(\frac{\partial u_{jmti}}{\partial \text{size}_{jmt}} \right) = \bar{\beta}_m^{sz} + 2\beta_m^{sz2} \text{size}_{jmt},$$

while other consumers' marginal utility for size is normally distributed around this level with variance σ_{sz}^2 .

²²Because we do not observe individual consumers, we simulate their income, inc_{mti} , from a log normal distribution fitted to mean household income and the Gini ratio for each market.

$$\beta_{mti}^{md} \sim N(\bar{\beta}_m^{md} + \pi_{md} \log \text{inc}_{mti}, \sigma_{md}^2).$$

Third, consumers' price-sensitivity, α_{mti} , is distributed log-normally conditional on consumer i 's income according to,

$$\log \alpha_{mti} \sim N(\bar{\alpha} + \pi_\alpha \log \text{inc}_{mti}, \sigma_\alpha^2).$$

Letting price sensitivity vary with income allows for non-homotheticity of preferences across models.²³

Fourth, we allow for consumers' tastes for purchasing automobiles to be heterogeneous across consumers with a mean that depends on the country and year to capture shocks to automobile demand,

$$\iota_{mti} \sim N(\bar{\iota}_{mt}, \sigma_\iota^2).$$

The final term in (9) is a brand-country fixed effect. The brand-country fixed effects, $\psi_{mb(j)}$, capture revealed preference for different brands within each country.²⁴ For each model j , $b(j)$ represents its brand. The same firm may operate multiple brands. That is, the Toyota Corolla is of brand 'Toyota' while the Lexus RX 450 is of brand 'Lexus' even though they are offered by the same firm (Toyota). Separating brands within firms is important since firms frequently use branding as a method of accentuating product differentiation. To the extent that consumers exhibit a preference for their home brands, this preference is absorbed into these brand-country fixed effects.

Under this parameterization of the demand model, $\theta^d = (\bar{\beta}_m^x, \beta_m^{sz2}, \sigma_x, \pi_{md}, \bar{\alpha}, \pi_\alpha, \sigma_\alpha, \sigma_\iota, \bar{\iota}_{mt}, \psi_{mb})$ represents the parameters to estimate, where $x \in \{hp, sz, md\}$. As discussed above, given θ^d and the observed market shares, there is a one-to-one mapping to the vector of demand shocks $\xi(\theta^d)$. We approximate the market shares using Halton sequences to integrate out the distribution of consumer taste, and we solve this mapping numerically to recover the demand shocks implied by the data for a given parameter values. We estimate the model by minimizing the generalized method

²³We have experimented with allowing the taste for cars relative to the outside good to vary with income, but this leads to imprecise estimates of that additional parameter and π_α . Note that we include country-year fixed effects which account for cross-country non-homotheticity.

²⁴As a robustness check, we have also estimated a version of the model with model fixed effects that are constant across countries.

of moments objective function,

$$\hat{\theta}^d = \underset{\theta^d}{\operatorname{argmin}} \xi(\theta^d)' Z \hat{W} Z' \xi(\theta^d),$$

where Z is a matrix of instruments and \hat{W} is a consistent estimate of the optimal weight matrix obtained from a first stage estimate.

To identify the price and other nonlinear parameters (eight in all), we employ three types of instruments. First, we make use of the sums of the characteristics of other products in the market—the so-called “BLP” instruments used in Berry, Levinsohn and Pakes (1995) which include the number of rival products in the market and the sum of the characteristics of rival products (horsepower per weight, size, and miles per dollar), yielding four instruments total. Second, we make use of the “differentiation IVs” advocated recently by Gandhi and Houde (2016). For each characteristic, we use two measures of the number of rival products in the vicinity a product, providing six additional instruments.²⁵ Finally, we leverage the assembly information in our data to use cost shifters as instruments. We use the distance to and the tariff rate on the model’s nearest assembly location, the model’s total number of assembly locations, and a dummy for domestic assembly location for the model.²⁶ This provides us with four cost side instruments, giving us 14 instruments total. Hence the model is over-identified. We have run several alternative specifications and have found that the results—in particular our measures of home preference and relative elasticities and markups—are qualitatively robust.

4.2 Demand Estimates

The demand estimates are presented in Table 8. Estimates of the tastes for characteristics are listed by country across columns at the top of table, with estimates of the standard deviation of the linear coefficients and income effects in the last two columns. Estimates of price sensitivity are provided at the bottom of the table. The price sensitivity parameters are all strongly significant and indicate that price sensitivity decreases in income ($\pi_\alpha < 0$). We also find that the taste for

²⁵Specifically, for each model j we count the number of competing models k such that difference in characteristics, $|x_j - x_k|$, is within 0.1 and 1 standard deviations of the unconditional distribution of characteristic differences. This statistic provides a measure of “local” competition in product space. See Gandhi and Houde (2016) for more details.

²⁶Note that domestic assembly at the model level is not perfectly correlated with domestic assembly at the brand level, and is more likely to affect the cost of a particular model than the brand reputation of the model.

Table 8: Parameter estimates for the demand model

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA	Random Coef. Std	Income Interactions
HP per Weight	0.225 (0.205)	0.487 (0.234)	-0.011 (0.189)	0.801 (0.310)	0.331 (0.175)	0.338 (0.162)	0.279 (0.173)	0.383 (0.196)	-0.051 (0.157)	0.390 (0.149)	
MPDCITY	0.128 (0.129)	0.436 (0.241)	0.281 (0.284)	0.437 (0.296)	-0.474 (0.183)	0.298 (0.248)	0.392 (0.264)	0.036 (0.186)	0.208 (0.263)	0.008 (0.015)	-0.533 (0.287)
Size	8.400 (10.570)	0.048 (2.779)	7.553 (2.374)	-3.669 (3.176)	5.620 (2.822)	6.597 (1.729)	5.485 (2.237)	1.613 (2.099)	8.390 (2.157)	0.041 (0.370)	
Size Square	-0.165 (0.658)	0.290 (0.217)	-0.261 (0.139)	0.547 (0.250)	-0.074 (0.203)	-0.200 (0.125)	-0.101 (0.163)	0.169 (0.170)	-0.344 (0.124)		

Random coef. automobile tastes

σ_ν
3.182
(4.490)

Price Sensitivity Parameters

$\bar{\alpha}$ π_α σ_α
2.524 -0.997 0.053
(0.190) (0.296) (0.087)

Notes: The units for HP per weight, size, and price are horse power per 100 kg, m^2 , and 10 thousand dollars, respectively. MPDCITY is miles per dollar in city driving. See table 5 for the description of parameters. This specification uses brand-country dummies. Weighted bootstrap standard errors in parenthesis.

miles per dollars decreases with income (π_{md}), consistent with low income consumers being more concerned with the cost of ownership of the vehicle.

Considering the tastes for characteristics, we find evidence that the tastes for product characteristics differ across countries, conditional on income. The estimated mean taste coefficients for horse power per weight, size, and miles per dollar are remarkably similar across Canada and the US, and quite different from the other countries. Since product characteristics are correlated with each other and they are subject to technological constraints, the estimated taste coefficients for characteristics are a bit difficult to interpret. Most countries tend to have positive taste for horse power per weight, while the taste for this characteristic is negative in the US and Canada. American and Canadian consumers tend to have lower taste for fuel efficiency than consumers in Germany and Great Britain. The taste for size is concave for six out of nine countries, with a significant coefficient on the quadratic term in four countries.

The brand-country effects estimated in this step, ψ_{mb} , capture revealed tastes for particular brands within country m . These will include consumers' valuation of home brand status, as well as the unobserved quality of the brand (such as reliability), its marketing cachet, the availability of dealerships and repair shops for the brand, et cetera. To assess the strength of home preference, we project the brand-market fixed effects on home status, a series of brand-market level controls, as well as brand and market fixed effects. Specifically, we follow Chamberlain (1982) and (Nevo, 2001, 322-323) to estimate,

$$\hat{\psi}_{mb} = \rho_1 \cdot \mathbf{1}[b \text{ is a home brand in } m] + \rho_2 \cdot \mathbf{1}[b \text{ assembles in } m] + \eta X_{mb} + \mu_b + \mu_m + u_{mb}, \quad (10)$$

where ρ_1 represents the preference for brands that originated at home and ρ_2 represents the preference for brands that have an assembly plant in the country, regardless of their origin. In our data, almost all home brands have domestic assembly while only some multinationals assemble in foreign markets, such as Toyota in the US.²⁷ The controls, X_{mb} , account for brand-market characteristics that are likely to be correlated with home status. We include in X_{mb} the number of years that brands have been selling in each country prior to the start of our data, and the number of

²⁷The single exception in our data to the rule that home brands have a domestic assembly plant is Smart. It is a German brand but is assembled completely outside Germany.

dealerships per household.²⁸ These variables capture various aspects of market penetration such as cumulative advertising, number of experienced mechanics, and the convenience for the consumer. Finally, we include a full set of market and brand fixed effects. The brand fixed effects absorb the global appeal of a brand. We refer to the first parameter of interest, ρ_1 , as home preference; it is the revealed preference for buying home brands, after controlling for observable characteristics of offerings, overall brand quality, and the brand history and dealer network in a market. This preference may have several sources, including the ability of home brands to provide unobserved characteristics that better fit their home market, consumers’ nationalistic feelings for domestic brands, or consumers’ stronger familiarity with domestic brands. We refer to the second parameter of interest, ρ_2 , as domestic assembly preference. It captures consumers’ potential support for local jobs and their perception of a brands’ long term commitment to supply their market.

The parameters in (10) can be consistently estimated if u_{mb} is uncorrelated with the right hand side variables. There are two major concerns one might have about this exogeneity assumption.²⁹ First, whether or not to have assembly plants and dealer networks were decisions endogenously made by the brand at some point in time. It is difficult to find valid instruments to control for this endogeneity explicitly. Many brands, however, first started producing in foreign countries and established dealer distribution networks a number of decades ago. Therefore, the current brand-country-specific demand shock, u_{mb} , is not likely to be strongly correlated with these past decisions. For instance, Toyota has been selling in the US market since 1958 and established production in 1986 in response to Japanese voluntary export restraints; similarly, Ford has been exporting to Germany since 1907 and established production in 1926 in response to tariff barriers. More systematically, among all 113 brand-country combinations with local production in our data, the median year of establishing local production is 1963 and the 75th percentile is 1990. Nevertheless, this may lead our estimates of the taste for domestic assembly ρ_2 and the taste for dealer density to be overstated, if those activities are exercised more when the brand-country-specific demand shock, u_{mb} , is high. By the same logic, this would bias home preference, ρ_1 , downward. Despite this possible downward bias, we find—as we discuss below—that the taste for home origin is large. Second, brands that

²⁸While the number of dealerships changes slightly from year to year, it is highly persistent, so we fix it over time in our data and rely on cross-brand variation. See the data description in Appendix A for details.

²⁹Note that even if this assumption does not hold, the estimates of brand preferences, elasticities and markups will still be consistent under our timing assumptions since they rely on moment conditions in ξ_{jmt} rather than u_{mb} . Furthermore, the estimation of the cost parameters below will be unaffected.

Table 9: Estimates for home brand and domestic assembly statuses

	I	II	III	IV	V
	$\hat{\psi}_{mb}$	$\hat{\psi}_{mb}$	$\hat{\psi}_{mb}$	$\hat{\psi}_{mb}$	$\hat{\psi}_{mb}$
Home brand, ρ_1	1.138 (0.074)	0.776 (0.093)	0.662 (0.095)	0.598 (0.097)	0.513 (0.099)
Domestic assembly, ρ_2		0.465 (0.062)	0.477 (0.062)	0.382 (0.064)	0.396 (0.064)
Years in market			0.004 (0.001)		0.003 (0.001)
Dealer density				0.121 (0.020)	0.116 (0.020)
Brand FE	Yes	Yes	Yes	Yes	Yes
Market FE	Yes	Yes	Yes	Yes	Yes

Notes: Dependent variable $\hat{\psi}_{mb}$ is the estimated brand-market fixed effects. Home brand is one if brand b originates from market m and zero otherwise. Domestic assembly is one if brand b has an assembly plant in market m and zero otherwise. Years in market is the number of years that brand b has been selling in market m as of 2011. Dealer density is number of dealers per household as of 2015. Home preference assumed to be homogeneous across countries. Weighted bootstrap standard errors in parenthesis.

originated at home or that have domestic assembly may also have higher advertising expenditures. We have tried to obtain such data, but it was impossible to get this information for all brands in all markets. However, advertising expenditure is presumably highly correlated with dealer networks and years in market, which we include in the specification as controls.

Table 9 presents our estimates of the drivers of brand preference varying the set of controls. Across all specifications, the home preference estimate is substantial and highly significant. The impact of history and dealer density attenuates the estimate of home preference by up to 35 percent, suggesting that the head start of national brands plays an important role in their dominance.³⁰ For our preferred specification with all controls in column V, the median consumer's willingness to pay for a home brand over a foreign one with domestic assembly is about \$600 in the US (if all other brand and model characteristics were identical). Relative to a foreign brand with no domestic assembly, the median US consumer is willing to pay about \$1,050 more for a home brand. This corresponds to roughly 4 percent of the average car price in the US. The willingness to pay numbers vary slightly across countries due to differences in the income distribution.

³⁰This result is robust to alternative specification of market-brand characteristics, such as rank of entry (rather than years in market) and including controls in logarithms (as opposed to levels).

Table 10: Country-specific estimates

	Home Brand (ρ_1)	Domestic Assembly (ρ_2)
BRA		-0.186 (0.572)
BEL		0.220 (0.146)
CAN		0.199 (0.194)
DEU	0.441 (0.263)	-0.194 (0.291)
ESP	1.032 (0.462)	0.360 (0.168)
FRA	0.853 (0.226)	0.364 (0.189)
GBR	0.250 (0.338)	0.463 (0.174)
ITA	0.945 (0.351)	0.320 (0.271)
USA	0.316 (0.181)	0.701 (0.163)

Notes: This table presents results for a specification in which home preference, ρ_1 , and the taste for brands with domestic assembly, ρ_2 , are heterogeneous across countries. Weighted bootstrap standard errors in parentheses.

To investigate how home preference varies across countries, we also estimate the model allowing home preference, ρ_1 , and taste for domestic assembly, ρ_2 , to be interacted with market country. The results, presented in Table 10, show that the home preference is not driven by outlier countries. Because most countries have relatively few home brands, these estimates are less precise, although most estimates remain statistically significant. However, there is some heterogeneity in home preference across countries. We find it to be highest in Spain, France, and Italy. These results are consistent with the findings of Goldberg and Verboven (2001) who find a strong preference for domestic brands in European car markets between 1980-1993. The taste for domestic assembly is strongest in the United States and Great Britain, and insignificant in a number of countries.

We finish this subsection by discussing possible drivers of home preference and of preference for locally producing brands. Consumers may have an innate preference for purchasing local brands arising out of nationalistic feelings. For example, Germans may prefer Volkswagen because they

view it as a German brand, while Italians might derive extra utility from purchasing a Fiat. Consumers may prefer locally producing brands due to a desire to support local employment or because they show a larger commitment to supply products and components also in the future in the local market. Home preference may also be capturing systematic differences in *unobserved* product characteristics (such as cup holders) for which tastes across countries differ. If home brands are better in supplying these characteristics, our estimate of home preference will reflect those as well. One may view the ability of local brands to quickly adjust to innovations in local preferences as a potentially important aspect of home preference. That said, the results documented in Table B.2 of Appendix B (see Section 2 for details) demonstrate that firms substantially customize their *observed* product characteristics across markets - even within the same model. Hence if cup holders are in higher demand in the US than in Europe, it is in the best interest of foreign brands to react and supply them to the US market. To the extent that they adjust these unobserved features in the same direction as home brands—as they seem to be doing for observed characteristics—it is unlikely that long-run differences in unobserved product characteristics will be the major driver of our estimate of ρ_1 .

4.3 Markups

The demand parameters directly imply elasticities and markups for each model. We provide the implied own and cross elasticities for selected models in Appendix Table B.3. In general, the model predicts plausible own and cross elasticity patterns, with models of similar classes competing against each other more strongly. We use the the estimates of elasticities to compute markups for each model according to the firm’s first order condition (8). Appendix Table B.4 presents the median (across years) of the implied markups for a selection of models. We find that markups range between

Table 11: Weighted average markups of firms across markets (percent)

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	7.6	8.5		9.2	7.2	9.5	8.4	11.8	7.4
Ford	6.3	8.5	9.0	8.4	7.5	9.2	9.4	8.5	10.0
GM	6.6	8.9	9.5	8.6	7.6	9.4	9.3	8.7	11.3
PSA	6.0	9.7		8.4	8.2	12.4	8.7	8.7	
Toyota	6.1	8.2	9.9	8.1	7.0	8.9	8.1	8.2	9.6
VW	7.5	9.7	8.5	10.8	8.6	9.8	9.5	8.6	9.2

roughly 6 percent to 13, with the United States having relatively high markups, and Brazil having relatively low markups. As a basis for comparison, GM reported that it achieved an average margin of 10.3 percent for North American operations in 2015, and states that part of its strategic plan is to maintain margins “between 9 and 10 percent” (General Motors, 2016).³¹

However, our interest is not so much on the average level of markups as the relationship between markups of home versus and foreign brands. To illustrate the impact of home preference on markups, Table 11 displays the share-weighted average markup of the six major manufacturing firms across all countries in our sales data. Across the table, we see that firms tend to charge their highest markups in their home countries. Volkswagen (VW) in Germany, General Motors (GM) in the United States, and most strikingly Peugeot (PSA) in France and Fiat in Italy. The pattern of home country markups relative to markups of the same firm in other countries is consistent with demand drivers playing an important role for home market advantage. Overall, the model produces intuitive estimates of markups—and hence marginal costs—across countries and models.

4.4 Supply Parameterization and Estimation

In the second stage of our estimation procedure, we use the costs implied by the demand model to estimate the supply side using nonlinear least squares. To do so, we parameterize $c_1(h_{jmt}, \kappa)$ and $c_2(g_{jmlt}, \delta)$, which determine the costs associated with selling model j in market m and the costs associated with sourcing model j from assembly location ℓ , respectively.

For country-model-specific costs, we assume,

$$\begin{aligned} \log c_1(h_{jmt}, \kappa) &= \kappa^{hp} \log hp_{jmt} + \kappa^{wt} \log wt_{jmt} + \kappa^{sz} \log size_{jmt} \\ &+ \kappa^{mg} \log mpg_{jmt} + \kappa_{mt} + \kappa_j. \end{aligned} \tag{11}$$

As opposed to the demand side, we allow costs to be determined by horsepower and weight separately, rather than by their ratio. This is intuitive because we would expect both to increase the cost of a car, whereas on the demand side we use their ratio as a measure of acceleration while accounting for size separately. We also include miles per gallon (mpg)—rather than miles per

³¹Of course, comparing accounting-based margins to our estimates of margins based on pricing and estimated elasticities is only a rough check on our estimation procedure. Our estimates are slightly lower than others in the literature for US automobiles for earlier time periods (e.g., Berry, Levinsohn and Pakes, 1995).

dollar—on the cost side. This is because the price of gas in m should affect demand for fuel-efficient vehicles, but not the cost of producing fuel-efficient vehicles. Finally, the supply side includes market country-year and model fixed effects. In contrast, the demand side includes brand-country effects. We prefer this specification because it allows us to control for the substantial variation in unobserved costs of models within brands on the supply side while being flexible about how tastes for brands vary across countries on the demand side.³²

The final element of the supply side is the assembly-location-specific cost function $c_2(g_{jmlt}, \delta)$,

$$\begin{aligned} \log c_2(g_{jmlt}, \delta) &= \delta^{mdist} \log \text{dist}_{m\ell} + \delta^{dom} \mathbf{1}[\ell = m] + \delta^{hqdist} \log \text{dist}_{h(j)\ell} \\ &+ \log(1 + \delta^{trf} \cdot \text{tariff}_{m\ell t}) + \delta^{xr} \log \text{fxrate}_{\ell t} + \delta_{\ell}. \end{aligned} \quad (12)$$

The first two terms capture the effect of trade costs, including a direct effect of distance as well as a dummy to control for domestic trade, in a traditional iceberg-like fashion.³³ The next term, $\delta^{hqdist} \log \text{dist}_{h(j)\ell}$, accounts for the impact of distance between a firm’s headquarters and the assembly location. Costs may be larger for distant plants due to monitoring or communication costs between a headquarters and its plants, or due to shipment of intermediate inputs.³⁴ The parameter δ^{trf} captures the proportion of the model’s cost subject to import tariffs.³⁵ Import tariffs are *ad valorem* based on the reported port cost of the car, which is likely to be lower than the marginal cost of the car implied by profit maximization, since the latter includes internal shipment and marketing costs. Below, we estimate the model both holding δ^{trf} fixed at one (the case where tariffs are paid on the full marginal cost) and allowing it to be estimated. We would expect only a portion of the marginal cost to be applied to the tariff, if a significant fraction of the marginal cost were incurred within the market country, due to distribution / retailing costs. The second to last term, $\delta^{xr} \log \text{fxrate}_{\ell t}$ captures the effect of exchange rate variation in source locations to costs. $\text{fxrate}_{\ell t}$

³²We have also estimated several alternative specifications, including model fixed effects on both demand and supply sides and country-brand effects on both sides. The results are qualitatively similar. Note that the effect of characteristics on cost are still identified when model fixed effects are included on the supply side, due to variation in the characteristics of a model both across countries and across years.

³³Recall that we use internal distance when the assembly and market countries are the same ($m = \ell$), see footnote 9 in Section 2 for details.

³⁴See Giroud (2013) and Tintelnot (2016) for a discussion and evidence for such frictions.

³⁵While identification comes mainly from the cross-section, there is also time variation in tariffs due to several events during the data period: some assembly countries become a member of the World Trade Organization (Ukraine’s entry in 2008), the EU and US reclassify countries in their Generalized System of Preferences and, finally, free trade agreements come into force (EU-Korea FTA in 2011).

is local currency per USD, normalized to one for the base year 2007. A depreciation of assembly country currencies would decrease dollar costs if $\delta^{xr} < 0$. Finally, we control for productivity differences across assembly locations with a location fixed effect, δ_ℓ , which is common to all plants within a country. This term absorbs both productivity difference across assembly countries and measurement error of internal distances within the assembly country.

The vector of supply parameters to estimate is $\theta^s = (\kappa, \delta, \sigma_\nu)$. The estimator for the supply side is the minimizer of the nonlinear least squares objective function,

$$\hat{\theta}^s = \underset{\theta^s}{\operatorname{argmin}} \sum_{m=1}^M \sum_{t=1}^{T_m} \sum_{j=1}^{J_{mt}} \omega_{jmt}(\theta^s)^2,$$

where,

$$\omega_{jmt}(\theta^s) = \log \hat{c}_{jmt} - \log c_1(h_{jmt}, \kappa) + \sigma_\nu \log \left(\sum_{k \in L_t(j)} \exp \left(\frac{-\log c_2(g_{jmkt}, \delta)}{\sigma_\nu} \right) \right).$$

To account for the fact that marginal costs \hat{c}_{jmt} are functions of demand estimates, standard errors are calculated using a weighted bootstrap across both demand and supply stages.

In practice, we find that the objective function is minimized at low values of $\hat{\sigma}_\nu$, indicating that the gains to variety are very small. For such values of σ_ν , the probability of sourcing from any particular location converges to either one or zero depending on whether that location is the minimum cost supplier or not. At low values of σ_ν , model predictions about marginal costs become insensitive to variation in σ_ν . As a result, the objective function becomes almost flat in this region, although other parameters are precisely estimated. Therefore, we fix $\bar{\sigma}_\nu = 0.01$ and estimate the remaining parameters of the model.³⁶ A low estimate of σ_ν is consistent with the findings of Head and Mayer (2015), who report that at the model level, firms almost always source a specific market from a single origin country.

Table 12: Marginal cost estimates

	I	II	III	IV
Horsepower, κ^{hp}	0.300 (0.045)	0.300 (0.045)	0.317 (0.044)	0.316 (0.044)
Weight, κ^{wt}	0.151 (0.036)	0.151 (0.037)	0.154 (0.036)	0.153 (0.036)
Size, κ^{sz}	0.363 (0.015)	0.363 (0.015)	0.367 (0.015)	0.367 (0.014)
Miles per gallon, κ^{mg}	0.055 (0.013)	0.056 (0.013)	0.054 (0.013)	0.055 (0.013)
Assembly-to-market distance, δ^{mdist}	0.004 (0.004)	0.004 (0.004)	0.016 (0.003)	0.016 (0.004)
Domestic assembly, δ^{dom}	-0.008 (0.005)	-0.008 (0.005)	-0.001 (0.002)	-0.001 (0.004)
Assembly-to-HQ distance, δ^{hqdist}	-0.001 (0.008)	-0.000 (0.009)	0.003 (0.007)	0.003 (0.007)
Tariff, δ^{trf}			0.686 (0.060)	0.682 (0.063)
FX rate, δ^{xr}		-0.013 (0.015)		-0.019 (0.015)
Fixed σ_ν	0.01	0.01	0.01	0.01

Notes: Estimates of the supply function specified in Section 4.4. Car costs, distance measures, tariff, and car characteristics are in logarithm. The coefficient on tariff, δ^{trf} , is fixed to one in columns I and II. Weighted bootstrap standard errors in parentheses.

4.5 Supply Estimates

The estimates of the supply side are presented in Table 12. Considering the effect of characteristics on the cost side (top panel), horsepower, size, weight and fuel efficiency all have the expected sign and are statistically significant. Turning to sourcing related costs (bottom panel), we find that for our preferred specification (column IV), production at a domestic (δ^{dom}) plant decreases costs. The cost elasticity of distance, δ^{mdist} , is comparable to the estimate of Head and Mayer (2015)—0.036—and within the range of estimates summarized by Head and Mayer (2013). It is slightly lower than the estimate from the reduced form price regression in Table 7, which could be interpreted as a cost elasticity under the assumption of perfect competition where firms always source from the nearest location. Assembly-to-HQ distances (δ^{hqdist}) increase marginal costs, consistent with non-trivial

³⁶The estimates and counterfactual results are robust to fixing σ_ν within a range of [0.001,0.2]. Moreover, an LR-type test rejects $\sigma_\nu > 0.2$ at the 95% confidence level.


monitoring and management costs related to remote production. However, this effect is smaller than the assembly to market cost elasticity and is not statistically significant.

We also find that properly controlling for tariffs matters for the magnitude of these effects. In Columns I-II, the tariff incidence parameter, δ^{trf} , is fixed at one—so the tariff applies to the full marginal cost of the car—and in Columns III-IV, we allow δ^{trf} to be estimated. When estimating δ^{trf} , we find that it is significantly below one, implying that the tariff is applied to less than the full marginal cost of the car, consistent with the presence of a portion of marginal costs being due to destination-specific internal delivery and marketing. Estimating δ^{trf} also has a substantial effect on the estimates of trade costs. In particular, the cost of assembly-to-market distance (δ^{mdist}) is four times larger when tariffs are controlled for (from column I to III). At the same time, the benefits of being domestic decrease substantially. This is intuitive because tariff rates are positively correlated with distance. Moreover, tariffs are naturally zero when the assembly plant is domestic, and they tend to be low between contiguous countries due to regional trade agreements. As a consequence, fixing the tariff incidence parameter above its estimated value induces downward bias on the impact of distance and an upward bias on the benefits of domesticity. The exchange rate parameter δ^{xr} in columns II and IV is negative as expected, though small in magnitude (a 10 percent depreciation of the assembly country currency would decrease total production costs – which includes costs incurred in the assembly, market, and possibly other supplier countries – by about .19 percent). From column III to IV, its inclusion does not significantly alter other estimates.

Our model delivers implied trade flows between countries at the model level through (3) and (4). We use this to conduct an out-of-sample test of our cost estimates by investigating how well they match aggregate trade flows. Specifically, we aggregate model-level flows up to the country-pair level and compare them to trade flows reported in the WITS database of the World Bank.³⁷ Figure 1 presents the scatter plot comparing our implied trade flows (in logs) with those in the trade data together with the best linear predictor of the data given our model flows. If our model perfectly replicated the aggregate data, the estimated slope of this regression would be exactly one, and R^2 would be one. In fact, the regression estimates a slope of 0.69, and the R^2 of this regression is 0.44. There are many reasons why we fail to match the aggregated trade flows perfectly. Our costs

³⁷This data is trade flows reported by importers in HS6 product categories associated with assembled cars. See Appendix A.4 for details.

Figure 1: Predicted trade flows and data



are not intended to represent accounting costs reported to customs upon import, but are instead the marginal costs the firm uses for setting prices—including costs incurred internal to the market country. Moreover, there is likely measurement error in both the aggregated trade flow data and in our data on market shares and prices used to estimate our model.³⁸ Finally, some mis-specification of our parametric functional forms used in estimation is inevitable. Overall, we believe the fact that the implied flow data matches the aggregate data as well as it does provides some degree of confidence that the model is capturing the essential drivers of market outcomes.

As an additional external validity check, we compare model-generated import shares in total sales with their empirical counterparts in Table 13. For each market, the ratios are total imported units divided by total units sold over the data period. Except for Belgium, the model captures the level of import penetration quite well: the correlation between data and model-generated ratios is 0.91.³⁹

³⁸Moreover, our model prices, shares and characteristics are themselves aggregations of finer trim-level data on new automobile sales.

³⁹Being a gateway to Europe, Belgium has a high level of re-exports, resulting in high gross imports which are not sold domestically—hence a import/sales ratio greater than one. The correlation between empirical and model-generated import penetration is 0.65 if we include Belgium.

Table 13: Import share in sales

	Model	Data		Model	Data
BRA	0.136	0.180	FRA	0.799	0.789
CAN	0.740	0.658	GBR	0.886	0.978
DEU	0.573	0.813	ITA	0.864	0.812
ESP	0.761	0.815	USA	0.525	0.453

Notes: Imported units divided by total units sold. See Appendix A.4 for details. For Belgium, model prediction is 0.949 versus the empirical value of 2.658, which is greater than one due to high level of re-exports.

5 What Drives Home Market Advantage?

In Section 2, we illustrated that firms tend to have substantially larger market shares in their home market. In Sections 3 and 4, we proposed and estimated a structural model that accounts for various demand and supply channels that could generate this home market advantage. These include tariffs, trade and remote production costs, cross-country heterogeneity in tastes for characteristics, and a preference for domestic brands. In this section, we use these estimates to assess the role of each in contributing to the home market advantage.

To do so, we re-estimate the same regression across a series of counterfactuals:

$$\log(s_{jmt}) = \lambda \cdot \mathbf{1}[b(j) \text{ is a home brand in } m] + \alpha_j + \gamma_{mt} + \varepsilon_{jmt}, \quad (13)$$

where the parameter λ measures the extent of “home market advantage.” We do not interpret λ causally, but instead use it as a metric of how home brands correlate with higher market shares in the data. Prices and market shares exactly match the data at the structural estimates, so the results from this exercise under the baseline are equivalent to our preliminary analysis in Section 2 (Table 4, Column 3), which implies that brands have on average 238 percent higher market share (i.e., more than triple) at home versus abroad. For each counterfactual, we re-calculate equilibrium costs, prices, and market shares, holding fixed model offerings within each market and the set of available assembly locations for each model.⁴⁰ However, we allow firms to re-optimize their sourcing decisions from this set. Therefore, these counterfactuals should be interpreted as “medium run” in the sense that firms can adjust sourcing and prices but neither the entry/exit of models into markets nor the construction/closure of assembly plants. Using these counterfactual market shares, we re-estimate (13) to determine the change in the home market advantage. Clearly, some of these counterfactuals do not represent changes that are achievable via policy. However, our goal is to use them as thought experiments to illustrate the drivers of home market advantage.

Table 14 displays the results. The first column reports the estimated coefficients under each counterfactual. The second column reports the implied change in the market share difference

⁴⁰As is well known, discrete choice demand models with consumer heterogeneity in tastes for characteristics and price could have multiple equilibria in the pricing game. We have not found such multiple equilibria for our estimates, but we also cannot rule out that they occur. We use the iterated best response algorithm starting from the initial equilibrium in order to compute the new equilibrium.

Table 14: Home market advantage under counterfactual scenarios

	Coefficient λ	Home Market Advantage (% Chg)
Baseline	1.22	
<i>Supply:</i>		
All tariffs eliminated	1.18	-4.9
No international trade frictions	1.09	-17.6
No multinational production frictions	1.22	-0.4
No tariffs, trade or multinational production frictions	1.08	-18.5
<i>Market access and prices of complementary goods:</i>		
Equalized dealer networks	1.04	-22.8
All countries have German gas prices	1.23	1.5
<i>Demand:</i>		
No home preference	0.82	-46.7
No local assembly preference	0.98	-30.7
No local assembly and no home preference	0.57	-67.6

Notes: Column 2 reports changes in the home market advantage statistic ($100 \times (\exp(\lambda) - 1)$).

between the baseline and the counterfactual estimate of λ . We now discuss each of the scenarios in turn.⁴¹ The baseline parameters which we modify below are the estimates in Tables 8, 9 (column V), and 12 (column IV).

Supply We begin by examining supply-side explanations for the home market advantage. In general, we would expect trade and foreign investment frictions to lead to cost differences that favor home brands. However, we do account for the fact that brands offshore the production of some models sold in their home market. Nonetheless, the home market advantage declines in all of our scenarios where we remove supply-side frictions.

We first consider the removal of all tariffs on automobile trade. This results in a slight decline in home market advantage of 5 percent. One reason for this is that many popular models of foreign brands are either produced domestically or in countries against which tariffs are already low—if not zero—due to regional free trade agreements. As a result, eliminating tariffs has only a minor effect on costs, which feeds through to a small decline in home market advantage. The effect increases when we remove the impact of both tariffs and shipping costs from the model. The impact of

⁴¹We present brand-country level market shares under each counterfactual in Appendix B. Price and profit outcomes are available from the authors by request.

external shipping costs is removed by reducing shipping costs to their domestic level regardless of assembly plant location. Again we allow firms to re-optimize their pricing and sourcing decisions. In the new equilibrium, we find a 17.6 percent decline in the home market advantage relative to the baseline. Next, we consider the impact of remote production frictions on the model. Eliminating costs associated with the distance between assembly plants and headquarters location, has only a small effect on home market advantage, which declines by 0.4 percent. Finally, we consider a reduction in all cost-side trade and foreign investment frictions in concert. Overall, when we remove all cost-side frictions, the home market advantage statistic declines by 18.5 percent. This suggests that cost-side drivers are an important but far from complete driver of home market advantage.

Market access and prices of complementary goods We next investigate how much of home market advantage is driven by differences in sales networks and prices of complementary goods. Foreign firms tend to have fewer dealers and service operations, lowering the utility of potential buyers of these brands. When we equalize the dealer density across brands (we set it equal to the median dealer density observed in the data in each country), we find that the home market advantage statistic would fall by 22.8 percent. This suggests that differences in sunk investment in retailing (and possibly advertising) are important sources of the observed home market advantage.

We next consider the impact of gas price differences across countries on home market advantage. Taking the current product offerings as given, home market advantage would *increase* slightly by 1.5 percent if all countries faced German gas prices, which are the highest gas prices across Europe. As discussed in Section 2, car makers customize their product offerings to the local car market. Therefore it is not surprising that differences in the price of gas across countries cannot explain the observed home market advantage since most foreign brands offer products that comport with local tastes—at in terms of fuel economy.

Demand Finally in Table 14, we remove consumers’ direct preference for home brands by eliminating the structural home preference (i.e., $\rho_1 = 0$). We find that eliminating home preference alone, while maintaining local assembly preferences that are almost universally enjoyed by local brands, has a dramatic effect on the home market advantage, which falls by 46.7 percent. Note that this calculation controls for the fact that home brands tend to have more dealerships and a longer history in their own home country. Home brands, as well as selected foreign brands, also have

local assembly. If we eliminate the preference for local assembly ($\rho_2 = 0$), home market advantage falls by 30.7 percent. In the case of eliminating both home preference and taste for local assembly ($\rho_1 = \rho_2 = 0$), home market advantage falls by 67.6 percent. This leads us to conclude that demand-side effects, and home brand preference in particular, are the key channel that gives rise to home market advantage in the automobile industry in the medium run, while cost-side elements play a substantive but secondary role.

The value of a domestic brand An interesting feature of the automobile industry is that there have been several mergers where an international firm owns a domestic brand but maintains its “domestic” image in marketing campaigns (e.g., Volkswagen’s ownership of SEAT in Spain, GM’s ownership of Vauxhall and Opel in the UK and Germany, and Fiat’s recent purchase of Chrysler in the United States). Our results suggest that one benefit of operating “domestic” brands for foreign firms is due to consumers’ preferences for local brands. Hence home preference can provide a motive for foreign direct investment (via acquiring local brands), analogous to jumping tariffs by establishing foreign production. To explore this idea in more detail, we use our model to calculate the importance of brand-nationality to the profitability of home brands. Specifically, we remove the home preference from only the brand under consideration and re-calculate the equilibrium to see the impact of home preference on that brand. For this exercise, we use country-specific home preference estimates from Table 10 to highlight cross-country differences. Table 15 reports how removing home preference affects local brands’ prices, sales and profitability in their home country. The upper panel considers foreign-owned brands, while the lower panel lists the largest home brand in the country.⁴² We find that home preference is extremely important to domestic brands’ operating profits. Although firms do react to the loss of home preference by lowering prices slightly,⁴³ the reduction in profits is largely due to a dramatic reduction in sales when the home preference is removed. The effect varies substantially by brand. The largest profit loss occurs at the Spanish brand SEAT (62 percent), which is owned by Volkswagen but has a relatively small presence worldwide. Despite being foreign-owned, SEAT appears to actively cultivate a Spanish identity to the extent of offering models named the Leon, Toledo, and Alhambra. In contrast,

⁴²SEAT and Vauxhall are both foreign-owned and the largest brand in their home country, so we do not repeat them.

⁴³Recall that markups are typically below 10 percent (see Table 11), so a one percent change in price represents roughly a 10 percent change in the markup since costs are unchanged in this counterfactual.

Table 15: Value of domestic brand status for selected brands

	Percent Change in		
	Price	Quantity	Profit
Seat in Spain (VW)	-0.4	-59.6	-61.9
Vauxhall in UK (GM)	-0.2	-17.1	-19.4
Chrysler in US (Fiat)	-0.1	-23.8	-24.3
Opel in Germany (GM)	-0.2	-27.8	-29.4
VW in Germany	-0.4	-25.2	-28.0
Renault in France	-1.0	-45.6	-50.7
Fiat in Italy	-1.3	-46.6	-53.4
Chevrolet in US	-0.3	-21.8	-23.2

Volkswagen itself—a strong worldwide brand—loses a comparably modest 28 percent of its local operating profits when its home preference is eliminated in Germany. In summary, the impact of home preference does appear large enough to suggest that buying a local brand may be an attractive mode of entry for foreign firms due to consumers’ innate preference for local brands.

The demand and cost effects of FDI Instead of acquiring a local brand in a foreign market, firms often simply establish foreign assembly operations. Our estimates suggest that such foreign investment has both a cost reducing effect (due to closer proximity to the market and possibly cheaper production costs) and a demand effect (due to consumers’ valuation of local assembly). Consider for example the case of Toyota in the US. Suppose Toyota removed all its US operations. How would its (variable) profits in the US market be affected?⁴⁴ We disentangle the total effect by the following thought experiment. Suppose Toyota’s US assembly plants were removed, but Toyota would still be conceived of consumers as a brand with local assembly. This counterfactual scenario yields what we term the “cost effect of FDI”.⁴⁵ As shown in Table 16, we find that Toyota’s US profits would fall by 17 percent under such circumstances. If instead Toyota is no longer valued as a local producer, but maintained its assembly locations on the supply side, this gives us the “demand effect of FDI.” We find that Toyota’s profits in the US would fall by 29 percent without the demand effect of FDI. Combining the effects leads to 41 percent lower profits

⁴⁴Since we have not estimated the fixed costs of foreign plants, we can only analyze changes in Toyota’s variable profit.

⁴⁵Whenever a model has multiple assembly locations, we drop the US assembly location in the counterfactual. Whenever the only assembly location for a model is in the US, we move the assembly location in the counterfactual to Japan.

Table 16: Demand and cost effects of FDI for selected brands

	Percent Change in		
	Price	Quantity	Profit
<i>Removing Toyota plants in US</i>			
No local assembly preference	-0.4	-24.4	-28.5
Only cost effect	1.9	-14.2	-17.1
Combined	0.9	-35.7	-41.1
<i>Removing Ford plants in Germany</i>			
No local assembly preference	-0.2	-29.2	-30.8
Only cost effect	1.0	-11.9	-10.6
Combined	1.5	-38.0	-38.3

by Toyota in the US. We have similarly analyzed the profits of Ford in Germany (also shown in Table 16). The effects are qualitatively similar, with the demand effect of FDI exceeding the cost effect. We think that the demand effect of FDI deserves further investigation as a novel channel unaccounted for by the standard theories on multinational firms.

6 Conclusion

The automobile industry exhibits significant home market advantage in market shares. This paper proposes and estimates a structural model to disentangle the contribution of various demand- and cost-side elements to market outcomes. The estimates clearly establish the existence of both demand factors and supply frictions behind the empirical regularity of home market advantage.

On the demand side, consumers exhibit preference for their national brands relative to how these brands are viewed in the rest of the world, even after controlling for car and brand characteristics. They also show preference for brands with a local production presence. Moreover, there are distinct differences in tastes for characteristics across countries. On the cost side, tariffs, trade costs, and remote production costs all play a role in segmenting markets.

To establish the relative importance of these channels, we conduct a series of counterfactual experiments isolating each feature of the model. Removing all tariffs, trade and multinational frictions, i.e., supply factors, explains around 18.5 percent of home market advantage. Another aspect of market penetration, the density of the dealers, has an equally important role. Equalizing dealer networks reduces home market advantage by 22.8 percent. It turns out that, however, home

preference in demand is the major driver of the home market advantage, with an effect roughly twice as large as removing all cost-side frictions. Together with preference for local assembly, these demand side factors explain around two-thirds of home market advantage. Quantitative models that ignore such strong demand side effects are prone to overestimating trade and multinational costs.

While our work separates the importance of trade costs and preference differences across countries, we are not able to disentangle further the sources of home preference and of preference for locally producing brands. It could be that these tastes arise out of nationalistic feelings, the desire to support local employment, or for other reasons like the greater ability by home brands to cater to local tastes for unobserved product characteristics. Note that our analysis allows for heterogeneous taste for observable product characteristics. Importantly, whatever the sources of home preference are, they are clearly distinct from trade costs.

References

- Anderson, James E.** 2011. “The Gravity Model.” *Annual Review of Economics*, 3(1): 133–160.
- Anderson, James E., and Eric van Wincoop.** 2003. “Gravity with Gravitas: A Solution to the Border Puzzle.” *American Economic Review*, 93(1): 170–192.
- Arkolakis, Costas, Natalia Ramondo, Andres Rodriguez-Clare, and Stephen Yeaple.** 2013. “Innovation and Production in the Global Economy.” National Bureau of Economic Research, Inc NBER Working Papers 18972.
- Atkin, David.** 2013. “Trade, Tastes, and Nutrition in India.” *American Economic Review*, 103(5): 1629–63.
- Auer, Raphael.** 2014. “Product Heterogeneity, Cross-Country Taste Differences, and the Growth of World Trade.” Swiss National Bank.
- Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott.** 2009. “The Margins of US Trade.” *American Economic Review*, 99(2): 487–93.
- Berry, Steven.** 1994. “Estimating Discrete-Choice Models of Product Differentiation.” *RAND Journal of Economics*, 25(2): 242–262.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1995. “Automobile Prices in Market Equilibrium.” *Econometrica*, 63(4): 841–890.
- Berry, Steven, James Levinsohn, and Ariel Pakes.** 1999. “Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy.” *American Economic Review*, 89(3): 400–430.

- Brambilla, Irene.** 2007. "A Customs Union with Multinational Firms: The Automobile Market in Argentina and Brazil." mimeo.
- Bronnenberg, Bart J., Jean-Pierre H. Dube, and Matthew Gentzkow.** 2012. "The Evolution of Brand Preferences: Evidence from Consumer Migration." *American Economic Review*, 102(6): 2472–2508.
- Bronnenberg, Bart J., Sanjay K. Dhar, and Jean-Pierre H. Dube.** 2009. "Brand History, Geography, and the Persistence of Brand Shares." *Journal of Political Economy*, 117(1): 87–115.
- Chamberlain, Gary.** 1982. "Multi Variate Regression Models for Panel Data." *Journal of Econometrics*, 18: 5–46.
- Coşar, A. Kerem, Paul L.E. Grieco, and Felix Tintelnot.** 2015. "Borders Geography, and Oligopoly: Evidence from the Wind Turbine Industry." *Review of Economics and Statistics*, 97(3): 623–637.
- Eaton, Jonathan, and Samuel Kortum.** 2002. "Technology, Geography, and Trade." *Econometrica*, 70(5): 1741–1779.
- Fajgelbaum, Pablo D, and Amit K Khandelwal.** 2014. "Measuring the unequal gains from trade." National Bureau of Economic Research.
- Fajgelbaum, Pablo, Gene M. Grossman, and Elhanan Helpman.** 2011. "Income Distribution, Product Quality, and International Trade." *Journal of Political Economy*, 119(4): 721 – 765.
- Feenstra, Robert C.** 1988. "Quality Change under Trade Restraints in Japanese Autos." *The Quarterly Journal of Economics*, 103(1): 131–46.
- Fieler, Ana Cecilia.** 2011. "Nonhomotheticity and bilateral trade: Evidence and a quantitative explanation." *Econometrica*, 79(4): 1069–1101.
- Gandhi, Amit, and Jean-Francois Houde.** 2016. "Measuring Substitution Patterns in Differentiated Products Industries." University of Wisconsin-Madison and Wharton School.
- General Motors.** 2016. "Form 10-K, 2015." Retrieved from https://www.gm.com/content/dam/gm/en_us/english/Group4/InvestorsPDFDocuments/10-K.pdfhttps://www.gm.com/content/dam/gm/en_us/english/Group4/InvestorsPDFDocuments/10-K.pdf.
- Giroud, Xavier.** 2013. "Proximity and Investment: Evidence from Plant-Level Data." *The Quarterly Journal of Economics*, 128(2): 861–915.
- Goldberg, Pinelopi Koujianou.** 1995. "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry." *Econometrica*, 63(4): 891–951.
- Goldberg, Pinelopi Koujianou, and Frank Verboven.** 2001. "The Evolution of Price Dispersion in the European Car Market." *Review of Economic Studies*, 68(4): 811–48.
- Goldberg, Pinelopi Koujianou, and Frank Verboven.** 2004. "Cross-country price dispersion in the euro era: a case study of the European car market." *Economic Policy*, 19(40): 483–521.
- Head, Keith, and Thierry Mayer.** 2013. "What Separates Us? Sources of Resistance to Globalization." *Canadian Journal of Economics*, 46(4): 1196–1231.

- Head, Keith, and Thierry Mayer.** 2015. “Brands in Motion: How frictions shape multinational production.” mimeo.
- Kennedy, Peter E.** 1981. “Estimation with correctly interpreted dummy variables in semilogarithmic equations.” *American Economic Review*, 71(4).
- Krugman, Paul.** 1980. “Scale economies, product differentiation, and the pattern of trade.” *The American Economic Review*, 950–959.
- McCalman, Phillip, and Alan Spearot.** 2013. “Why trucks jump: Offshoring and product characteristics.” *Journal of International Economics*, 91(1): 82 – 95.
- Melitz, Marc J.** 2003. “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity.” *Econometrica*, 71(6): 1695–1725.
- Miravete, Eugenio J., Maria J. Moral, and Jeff Thurk.** 2015. “Innovation, Emissions Policy, and Competitive Advantage in the Diffusion of European Diesel Automobiles.” CEPR Discussion Paper No. DP10783.
- Nevo, Aviv.** 2001. “Measuring Market Power in the Ready-to-Eat Cereal Industry.” *Econometrica*, 69(2): 307–342.
- Ramondo, Natalia, and Andres Rodriguez-Clare.** 2013. “Trade, Multinational Production, and the Gains from Openness.” *Journal of Political Economy*, 121(2): 273 – 322.
- Tintelnot, Felix.** 2016. “Global Production with Export Platforms.” Forthcoming, *Quarterly Journal of Economics*.
- van Garderen, Kees Jan, and Chandra Shah.** 2002. “Exact interpretation of dummy variables in semilogarithmic equations.” *The Econometrics Journal*, 5(1): 149–159.
- WTO.** 2013. “International Trade Statistics 2013.” World Trade Organization.

Appendices

Appendix A Data

A.1 Demand Data

We purchased a dataset of trim-level unit sales, prices (MSRP) and characteristics (size, weight, horsepower, fuel type, transmission, wheel base) for nine markets (Belgium, Brazil, Canada, France, Germany, Italy, Spain, UK, US) and five years (2007-2011) from R.L.Polk & Company, a market research firm that got acquired by IHS Inc. in 2013. Data for the years 2007 and 2008 are missing for Brazil and Canada, respectively. Following the common practice in the literature, we aggregated sales to the model level since very small market shares at the trim level create numerical challenges for the BLP inversion. The aggregation used trim-level sales as weights to calculate average model prices and characteristics. We fill in the few cases of missing characteristics (most notably in Brazil), with the characteristics of the same models from the North American market. Prices at local currencies were translated into USD using the average annual exchange rate. In countries with a retail sales tax, we augment this price with the retail sales tax so it approximates the effective price to the consumer. This procedure generated 9,498 observations. We dropped pickup trucks since they constitute a somewhat unique segment in the US. We also dropped observations for 2010-2011 in Canada since information on SUV models sold there in these years was missing. This leaves us with 8,841 observations. Additional data come from OECD (sales tax data), Penn World Tables Version 8 (income per capita in 2005 PPP, price levels of GDP, exchange rates), and the World Bank (Gini coefficients)

All nominal variables in the data, MSRP and gas prices in each market-year, have been deflated to 2005 US dollars to be consistent with the PPP income per capita variable denominated in 2005 US dollars. All variables are first deflated to their 2005 levels using price level time series from the Penn World Tables, and then converted to USD using the nominal exchange rate of the year.

A.2 Supply Data

To locate the production locations of unique model-year combinations in the demand data, we purchased data on assembly plants by manufacturer groups and models between 2007-2011 from Ward's Communications. Assembly countries for model-years present in the demand data but missing in the purchased supply data were collected by research assistants from the Internet. The complete supply data encompasses 52 assembly countries. The models produced in Uruguay belong to the Chinese brand Geely for which fuel efficiency measures are missing. As a result, we drop Uruguay as an assembly location. Also, data for Kenya and Bangladesh overlap in that Toyota Land Cruiser is the only model produced in these countries. Since this leads to multicollinearity in estimating model and production location fixed effects, we drop Bangladesh. This leaves us with 50 countries from which the models in the demand data could be supplied. The countries in which manufacturer groups are headquartered constitute another dimension of the data, which is more easily accessible from online sources. There are 12 headquarter countries associated with the 28 manufacturing groups: China, Germany, Spain, France, the UK, Italy, Japan, Korea, Malaysia,

Russia, Sweden and the US.⁴⁶ The CEPII dataset (Head and Mayer 2013) provides us with the distances between headquarter and assembly countries, as well as the distance and the contiguity of our nine markets to assembly countries. Bilateral tariff data come from TARIC (EU Integrated Tariff Database), Canada Border Services Agency, USITC and WITS databases. Most of the bilateral tariffs were constant throughout the data period with two exceptions. The entry of Ukraine to the WTO led to a reduction of US tariffs from 10% to the MFN level of 2.5%. EU tariffs to S. Korea decreased from 10% to 3% in 2011 when a free trade agreement became effective. We ignore rules of origin requirements related to the regional value-added content in FTAs: for instance, according to NAFTA rules, a car can be imported from Mexico to the US tariff-free only if the regional value-added content is above 62.5%. The rest is subject to tariff. Unfortunately, systematic model-level data on location-specific value-added is not available. In our cost estimation (subsection 4.5), we make an attempt to account for the fact that only a fraction of an imported car’s cost is subject to import tariffs.

In order to investigate the sources of the brand-country fixed effects from the demand estimation, we supplement the demand data with information on brands’ years of entry into and the number of their dealers in each of our 9 markets. The year-of-entry data was collected by consulting various sources including the Internet, business history books and companies’ public relations agents. Data regarding the number of dealers was collected from Google Maps. There are 331 brand-country observations.

The number of car dealers for a manufacturer brand within a country is collected using Google Places API (<https://developers.google.com/places/webservice/search?hl=en>). This API provides a function called radar search that returns the search query given the key words, place types, center coordinates and radius of the area of interest. The query has detailed information including place id that can uniquely identify a place, coordinates and description. There is a limit to the number of results returned per search (200) and also the radius (50km). We set the keyword to be the name of the manufacturer brand and the place type to be “car dealer”.⁴⁷ Then we iterated over areas to cover the entire country by choosing different coordinate centers and set the radius to 50km. The area may cover places outside of the country, in which case we removed these results based on their coordinates. There may also be overlapped area search in the search iteration, and we removed the repeated results using place id. To avoid counting dealers of used cars, we did a radar search using used car as the keyword, and deleted a place if its id is found in the used car list.

A.3 Gas Prices

We collected data on unit price at the pump for gasoline, diesel and ethanol in all our markets. Data sources are the websites of the US Energy Information Administration, Natural Resources Canada, European Commission, and the Brazilian National Agency of Petroleum, Natural Gas and Biofuels (ANP). Fuel market shares in each country and year are used as weights in calculating average gas prices.

⁴⁶For two manufacturing groups, Chrysler-Fiat and Renault-Nissan, we assigned each firm a separate headquarter country: Chrysler in the US and Fiat in Italy, Renault in France and Nissan in Japan. During this period, key managerial decisions at Chrysler were still made in Detroit and the merger of the two companies wasn’t legally complete until 2014. Similarly, while the Renault-Nissan alliance coordinates global procurement, production and marketing, they still keep their separate management structures and brand identities.

⁴⁷Because Opel is a common location name in some countries, we use “Opel dealer” instead of “Opel” in the search, where “dealer” is translated into the local language.

A.4 Trade flows and Import Shares in Sales

Figure 1 and table 13 use trade flows in USD and in units downloaded from `wits.worldbank.org`. In the figure, we use USD import value from each source country. The table reports import penetration, i.e., imported units over units sold. We take total sales by market-year from `Marklines.com`. Since our data captures most of cars sold in each market, using sales from our baseline data slightly increases the empirical import/sales ratio in table 13.

In the external validity checks, we use values and quantities reported by importers in HS6 product categories associated with assembled cars. This spans all the subheadings under 2007 HS 4-digit heading 8703 (Motor cars and other motor vehicles principally designed for the transport of persons) excluding 870310 (Vehicles specially designed for traveling on snow; golf carts and similar vehicles). These HS product codes are: 870321, 870322, 870323, 870324, 870331, 870332, 870333, and 870390. This data includes many small flows due to personal imports of automobiles, so we exclude pairs with less than \$5 million in reported flows, which amounts to roughly 200 units. The reported results are robust to adjusting this cutoff.

Appendix B Additional Figures and Tables

B.1 Analysis of Price Discounts

As is well known, transaction prices can differ from the manufacturer’s suggested retail price, which introduces some measurement error into our price variable. Also, at least anecdotally, brands are known to differ in their discount policies. In this Appendix section, we analyze whether home brands differ systematically from foreign brands in their discount behavior. We downloaded online car prices for US, Germany, and France during the period December 2015 to February 2016.⁴⁸ We analyze discount behavior across the three countries. The left hand side variable is the price discount ($\frac{\text{online price}}{\text{MSRP}} - 1$). On the right hand side, we interact a dummy variable whether a brand is at home with a country dummy and include country fixed effects in the regression. We find that German brands tend to give smaller discounts (note that discount is a negative number so a positive coefficient implies smaller discounts), while French and US brands tend to give larger discounts. Importantly, though, brands do not systematically give larger discounts in their home market than in foreign markets as illustrated by the insignificant results for the home brand variables in the results in columns II and III when global model or brand fixed effects are included. Hence the global brand fixed effects in our demand estimation will absorb differences in the brands’ discount behavior.

Table B.1: Online price discounts

	<i>Separate by country</i>	<i>Global model fixed effect</i>	<i>Global brand fixed effect</i>
	I	II	III
	discount	discount	discount
Home brand in USA	-0.0173 (0.00846)	-0.000689 (0.0205)	0.00633 (0.0157)
Home brand in Germany	0.0199 (0.00841)	-0.00449 (0.00970)	0.00820 (0.00948)
Home brand in France	-0.0390 (0.0108)	0.00968 (0.0147)	-0.0168 (0.0145)
Observations	803	803	803
Country Fixed Effect	Yes	Yes	Yes
Model Fixed Effect	No	Yes	No
Brand Fixed Effect	No	No	Yes

Notes: Standard errors in parentheses.

⁴⁸Aggregate pricing data for the US was taken from *www.truecar.com* on December 22, 2015 in 5 cities (Philadelphia, Los Angeles, Washington D.C., and Dallas) and reflect the base car model. Aggregate pricing data for Germany was taken from *www.autohaus24.de* on January 23, 2016 and reflect the base car model. Pricing data for France was taken from *www.promoneuve.fr* on February 8, 2016 in every region and include car styles and possible optional equipment, where authorized dealers post prices for new cars. Pricing discounts for each country are calculated by first calculating the average price of the car model and MSRP for unique models and then calculating the discount.

B.2 Descriptive Evidence

Figure B.1: Models and market shares by number of assembly countries




Table B.2: Characteristics by market

	I	II	III
	$\ln(hppwt_{jmt})$	$\ln(size_{jmt})$	$\ln(mpg_{jmt})$
BEL	-0.276 (0.006)	-0.009 (0.001)	0.253 (0.006)
BRA	-0.043 (0.0111)	0.00308 (0.004)	0.188 (0.008)
CAN	-0.0001 (0.006)	0.0008 (0.001)	0.0156 (0.005)
DEU	-0.195 (0.006)	-0.006 (0.001)	0.157 (0.006)
ESP	-0.228 (0.006)	-0.007 (0.001)	0.228 (0.006)
FRA	-0.239 (0.006)	-0.006 (0.001)	0.267 (0.006)
GBR	-0.210 (0.006)	-0.008 (0.001)	0.189 (0.006)
ITA	-0.235 (0.006)	-0.008 (0.001)	0.229 (0.006)
Observations	8841	8841	8841
R^2	0.953	0.985	0.928
Year FE	Yes	Yes	Yes
Model FE	Yes	Yes	Yes

Standard errors in parentheses. US is the omitted dummy, so all coefficients showcase differences in country means against the US.

B.3 Model predictions

B.3.1 Elasticities and Markups

Table B.3 presents the elasticities and cross-elasticities for selected models in the subset of markets where these models compete. Looking at the own-price elasticities, we see they vary mildly across models. When we consider the cross-elasticities, the table illustrates that the model is able to capture the expected competitive patterns. The two luxury models, the Audi A6 and Mercedes E350, compete most strongly with each other. Similarly, Renault Clio, Toyota Corolla, and Ford Focus compete strongly with each other but not with the luxury vehicles.

Table B.3: Median own and cross-price elasticities for select models

	Audi A6	Ford Focus	Mercedes E350	Renault Clio	Toyota Corolla
Audi A6	-14.687	0.034	0.433	0.007	0.021
Ford Focus	0.064	-14.473	0.029	0.261	0.471
Mercedes E 350	0.230	0.006	-15.727	0.003	0.001
Renault Clio	0.008	0.311	0.001	-14.090	0.028
Toyota Corolla	0.003	0.483	0.001	0.269	-14.962

Notes: This table shows the substitution elasticity of models in the row with respect to the prices of models in the column. Each entry represents the median of elasticities across country-years.

Table B.4: Median markups of select models across years (percent)

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Audi A4	5.8	9.1	8.0	9.8	8.5	9.9	9.3	8.3	9.3
Audi A6		9.0	8.2	8.7	8.2	9.6	9.4	7.7	11.4
BMW 530		7.8		6.7	6.7	8.7	8.8	7.4	10.6
BMW X3		8.1	8.8	7.0	7.1	8.8	8.8	7.3	11.4
Chrysler 300		6.8	9.2	5.6	6.5	8.1	8.3	6.3	10.9
Ford Fiesta	5.4	8.8		9.8	8.2	9.3	8.9	8.8	7.9
Ford Focus	6.0	8.2	8.8	7.7	7.8	8.8	8.9	8.0	8.1
Honda Accord	6.4	7.2	9.1	6.6	6.9	8.3	7.7	6.7	9.8
Honda CR-V	7.8	7.4	9.6	6.6	7.1	8.7	7.9	7.0	9.9
Jaguar XF		6.9	8.0	5.5	6.1	7.7	8.8	6.0	11.9
Jeep Grand Cherokee	6.8	6.8	10.0	5.8	6.4	8.0	8.3	6.3	13.0
Lexus RX 450		5.8	6.9	4.8	4.9	6.8	6.9	5.2	8.5
Mercedes E 350		7.1	7.7	6.8	6.5	8.4	8.2	6.4	12.3
Mini New Mini	6.7	7.9	7.9	7.8	7.4	8.4	7.6	7.4	7.5
Renault Clio	4.6	9.5		9.7	8.6	12.3	8.6	8.6	
Toyota Corolla	5.7	8.0	10.3	8.0	7.5	8.5	8.7	7.9	8.9
Toyota RAV-4	6.7	7.5	9.3	6.6	7.0	8.7	7.8	7.2	10.2
VW Golf	6.8	9.4	9.0	11.4	8.8	9.6	9.1	8.6	7.6
VW Passat	6.0	9.1	8.2	10.2	8.8	9.5	8.9	8.6	9.2
VW Tiguan	8.3	9.6	8.8	10.5	9.1	10.3	9.6	8.8	10.0

Table B.4 presents the median (across years) of the implied markups for a selection of models in

all countries where those models appear. Intuitively, markups are lowest in Brazil, which is by far the lowest income country in our dataset. Several smaller models such as the VW Golf, Mini, and Ford Fiesta tend to have smaller markups in the United States than they do in European countries. The overall pattern of brands tending to have higher markups in their home country is also apparent in this table, although the popularity of model characteristics and product competition also clearly affects markups.

B.3.2 Trade and Foreign Production Frictions

To get a sense for the magnitude of the estimated trade frictions, we conduct two exercises which calculate the proportion of automobile costs that are due to external shipping and remote production, and showcase how these quantities vary across brands and countries. Note that this analysis computes costs actually paid in overcoming production frictions. It is not capturing the impact of production frictions that firms endogenously avoid (e.g., sourcing locally to avoid high tariffs and shipping costs). In the Section 5, we conduct a series of counterfactuals which allow firms to re-optimize production decisions when frictions are removed.

Table B.5: Weighted average external shipping cost (assembly to market cost, including domestic assembly effects) as percent of marginal cost

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	0.0	5.0		2.5	2.8	2.4	4.0	1.4	2.5
Ford	0.2	3.3	0.6	1.4	2.7	2.1	3.7	2.8	0.2
GM	0.2	4.1	0.6	1.3	2.1	1.9	1.7	2.1	0.2
PSA	0.4	4.4		2.2	1.4	1.3	3.2	2.2	
Toyota	0.3	5.6	0.9	3.2	3.5	2.4	3.7	2.8	1.1
VW	0.2	3.9	2.5	1.2	1.5	1.9	3.1	2.0	1.6

Table B.5 reflects the percentage of the total cost that is directly related to external shipping and domesticity (but not including tariffs). In this exercise, we use the estimates from column IV of Table 12 to calculate the proportion of costs paid that are due to shipping from an international location. That is, we calculate the change in total costs when we set the domestic dummy equal to one and the distance between the assembly countries and the destination market equal to the internal distance of the destination market. This calculation keeps tariffs, sourcing locations, and the distance between assembly and headquarters the same. We see some variation by firm and country, ranging between 0 percent (Fiat in Brazil) to 5.6 percent (Toyota in Belgium) of the marginal cost. As we would expect, these costs tend to be relatively low in the firm’s home country, despite the fact that even home firms import at least some proportion of their cars from abroad, generating positive external shipping costs. However, in Brazil, where many firms have local plants targeting South American markets, average shipping costs are actually lower as a share of costs than in home countries.⁴⁹

Table B.6 carries out a similar exercise by computing the proportion of costs due to sourcing from assembly locations outside the firm’s headquarter country. In this case, we compute the proportion of additional costs from assembling cars outside of the home headquarter country as a proportion of the overall cost. Not surprisingly, these costs tend to be smallest in the firm’s

⁴⁹It is also interesting to note that General Motors has its lowest average shipping costs to European markets in Germany, where its Opel subsidiary is based.

home country, although they are not zero since, again, firms source some models in home markets from abroad. These costs range from about 0 to about 1 percent of marginal costs of supplying a model to a market. As with shipping costs, the case of Brazil is especially interesting since remote assembly costs tend to be highest there. This is the flip side of the low shipping costs for the Brazilian market observed in Table B.5. Firms are endogenously choosing to locate assembly locations in Brazil, incurring remote production costs instead of paying higher shipping costs and high import tariffs to access the Brazilian market.

Table B.6: Weighted average remote production cost as percent of marginal cost

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	0.8	0.3		0.3	0.3	0.3	0.4	0.2	0.6
Ford	0.4	0.3	0.0	0.3	0.3	0.3	0.3	0.3	0.0
GM	0.4	0.4	0.0	0.4	0.4	0.4	0.4	0.4	0.0
PSA	0.8	0.2		0.3	0.2	0.2	0.3	0.3	
Toyota	1.0	0.5	0.8	0.6	0.5	0.6	0.5	0.6	0.6
VW	0.9	0.2	0.6	0.2	0.4	0.3	0.3	0.3	0.6

B.4 Counterfactuals

The following tables present market shares from the baseline model (which exactly matches share and price data at the model level) and the counterfactual scenarios we consider. Data is aggregated according to brand nationality.

Table B.7: Data: average area-level market shares of brands across markets (%)

	Data/Baseline								
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	31.0	8.9	34.1	8.3	11.3	6.6	15.9	11.6	39.6
EU brands	56.6	75.7	8.5	77.0	70.6	82.6	62.8	74.0	9.5
JPN brands	8.3	11.3	48.3	10.9	13.1	8.8	16.9	11.5	42.9
Other brands	4.2	4.0	9.1	3.7	5.1	2.0	4.3	2.9	7.9
Home brands				55.4	9.0	52.3	18.2	30.3	39.6

Table B.8: Supply counterfactuals: percentage point changes of average area-level market share of brands across markets

All tariffs eliminated									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-8.5	0.7	-5.3	1.1	1.8	0.3	2.3	2.0	-2.7
EU brands	-18.7	-4.6	2.9	-5.4	-7.2	-2.1	-7.4	-5.7	2.1
JPN brands	-3.5	3.0	0.4	3.6	4.0	1.3	4.2	2.9	-0.4
Other brands	30.7	1.0	2.1	0.7	1.3	0.4	0.8	0.7	1.0
Home brands				-5.3	-4.0	0.3	-6.0	-3.6	-2.7
No international trade frictions									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-0.5	-0.1	-2.9	0.6	1.0	0.2	1.2	1.1	-3.1
EU brands	-1.2	-2.8	1.5	-4.3	-5.0	-1.6	-4.7	-3.3	1.4
JPN brands	-0.5	2.3	0.3	3.1	3.0	1.1	2.9	1.8	1.0
Other brands	2.1	0.6	1.2	0.6	0.9	0.3	0.6	0.4	0.7
Home brands				-4.9	-3.1	-0.4	-4.4	-2.3	-3.1
No multinational production frictions									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	0.0	0.4	-0.3	0.4	0.4	0.3	0.5	0.5	-0.4
EU brands	0.0	-0.5	0.0	-0.6	-0.5	-0.4	-0.5	-0.6	0.1
JPN brands	0.3	0.1	0.3	0.1	0.1	0.1	0.0	0.1	0.5
Other brands	-0.3	0.1	-0.1	0.1	-0.0	0.0	0.0	0.0	-0.2
Home brands				-0.6	-0.1	-0.2	0.3	-0.5	-0.4
No tariffs, trade or multinational production frictions									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-11.0	0.9	-8.5	1.6	2.4	0.5	3.3	3.1	-6.2
EU brands	-24.6	-5.6	4.6	-7.0	-9.1	-2.8	-9.8	-7.9	3.9
JPN brands	-3.1	3.8	0.9	4.8	5.4	1.9	5.6	4.0	0.6
Other brands	38.7	1.0	3.0	0.7	1.3	0.5	0.9	0.8	1.6
Home brands				-6.1	-6.3	0.5	-9.6	-5.8	-6.2

Table B.9: Market access and prices of complementary goods counterfactuals: percentage point changes of average area-level market share of brands across markets

Equalize dealer networks									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-0.8	-0.0	-2.4	-0.8	-0.1	1.1	-0.7	-0.0	-1.4
EU brands	0.8	-0.9	3.3	-1.9	-0.5	-3.3	0.2	-1.0	0.9
JPN brands	-0.0	0.4	-1.0	2.1	0.3	1.8	0.3	0.7	0.2
Other brands	0.1	0.5	0.1	0.6	0.3	0.4	0.2	0.3	0.4
Home brands				-4.4	-0.8	-9.7	-0.5	-2.8	-1.4
All countries have German gas prices									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-9.0	-0.0	-0.1	0.0	-0.0	-0.0	0.0	0.1	-2.0
EU brands	5.8	0.2	0.1	0.0	1.0	0.1	0.2	-0.0	0.0
JPN brands	2.7	-0.1	0.5	0.0	-0.8	-0.0	-0.1	-0.0	2.2
Other brands	0.6	-0.0	-0.5	0.0	-0.2	-0.0	-0.1	-0.0	-0.2
Home brands				0.0	0.6	0.1	0.0	0.0	-2.0

Table B.10: Demand counterfactuals: percentage point changes of average area-level market share of brands across markets

No home preference									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	0.0	0.0	0.0	1.8	0.4	1.4	1.1	1.2	-9.6
EU brands	0.0	0.0	0.0	-4.9	-0.9	-4.0	-2.6	-2.8	1.8
JPN brands	0.0	0.0	0.0	2.4	0.4	2.1	1.2	1.3	6.6
Other brands	0.0	0.0	0.0	0.7	0.2	0.4	0.3	0.3	1.1
Home brands				-9.8	-3.2	-10.4	-5.7	-7.0	-9.6
No local assembly preference									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-0.7	-1.5	-2.8	-1.0	-0.6	1.4	1.5	-1.7	-1.2
EU brands	-0.7	-0.2	1.8	-1.8	-2.1	-2.4	0.5	0.0	1.9
JPN brands	0.0	1.3	-0.7	2.2	1.5	0.5	-2.4	1.3	-0.3
Other brands	1.3	0.5	1.6	0.6	1.2	0.4	0.4	0.4	-0.4
Home brands				-6.2	-1.2	-6.7	-3.9	-4.8	-1.2
No local assembly and no home preference									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US brands	-0.7	-1.5	-2.8	0.4	-0.3	2.9	2.5	-0.7	-10.7
EU brands	-0.7	-0.2	1.8	-6.4	-3.0	-6.3	-1.6	-2.4	4.0
JPN brands	0.0	1.3	-0.7	4.7	1.9	2.5	-1.5	2.5	6.1
Other brands	1.3	0.5	1.6	1.3	1.4	0.9	0.6	0.7	0.6
Home brands				-15.9	-4.0	-16.8	-8.6	-11.3	-10.7