Taste Heterogeneity, Trade Costs, and Global Market Outcomes in the Automobile Industry*

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

[a]University of Chicago, Booth School of Business
[b]Pennsylvania State University, Department of Economics
[c]Durham University, Business School
[d]University of Chicago, Department of Economics

April 8, 2015

Abstract

In the automobile industry, as in many tradable goods markets, firms earn their highest market share within their domestic market. This home market advantage persists despite substantial integration of international markets during the past several decades. The goal of this paper is to quantify the supply- and demand-driven sources of the home market advantage and to understand their implications for international trade and investment. Building on the random coefficients demand model developed by Berry, Levinsohn, and Pakes (1995), we estimate demand and supply in the automobile industry for nine countries in three continents, allowing for unobserved taste and cost variation at the model and market levels. The estimated model helps to examine the contributions of tariffs, trade and FDI costs, home preference, and taste heterogeneity to domestic firms' home market advantage as well as conduct policy analysis.

*Preliminary and incomplete; comments welcome. We thank Arnaud Costinot, Jonathan Eaton, Keith Head, and seminar participants at various institutions for valuable comments and suggestions. This work was completed in part with resources provided by the University of Chicago Research Computing Center. Coşar wishes to thank the University of Chicago Booth School of Business' Initiative on Global Markets for financial support. Correspondence: `kerem.cosar@chicagobooth.edu`, `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 during 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 the home market advantage, and to understand the implications of these frictions for international trade and investment. In particular, what is the role 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, durable and tradable goods, while domestic producers command a dominant share in their home markets (see Table 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. 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 estimating demand allowing for both within and across market heterogeneity in demand. This more flexible approach enables us to consistently estimate demand- and supply-driven mechanisms behind market segmentation.[1] Our approach also enables the estimation of costs of foreign production from detailed industry level data and extends the analysis of recent quantitative trade models with multinational production (Ramondo and Rodriguez-Clare 2013, Arkolakis, Ramondo, Rodriguez-Clare, and Yeaple 2013, and Tintelnot 2014), which were also limited by the availability of only revenue data on multinationals' foreign affiliate sales for the aggregate manufacturing sector.

Apart from its importance in world trade and manufacturing employment, the availability of data also makes the auto industry suitable for our analysis. We have compiled a rich and unique dataset of global demand and supply. The demand data informs us about prices and quantities (as opposed to sales revenue only) by model as well as several characteristics such as horsepower, size, weight and fuel efficiency in nine countries across three continents. On the supply side, worldwide data on the assembly plant locations of each model enables us to solve the optimal sourcing problem of producers. Through the use of a structural model, we rely on price data to separate "home preference"—demand-driven preferences for local products—from trade and investment costs. Moreover, we are able to separately identify country level preferences for characteristics (such as fuel efficiency) from home preference. This is important since in a world with trade costs, local firms endogenously produce products which are more amenable to local tastes and therefore obtain larger national market share for reasons distinct from national preferences (Krugman, 1980). As pointed out by Auer (2014), if endogenous specialization is a strong feature of the economy, trade will respond sluggishly after a trade liberalization.

A number of papers have looked at the car industry to study questions in international

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

trade. Among these papers, 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 (2001b) study price dispersion in the European car market and also find evidence for consumers favoring national brands; Goldberg and Verboven (2004) use panel data from the car industry to demonstrate a strong positive effect of the Euro on price convergence. McCalman and Spearot (2013) study firms' offshoring strategies using data on North American light truck production locations. More recently in contemporaneous work, Head and Mayer (2015) extend the workhorse multinational production models by adding a demand friction and quantify their model using trade flow data from the car industry.[2] Similarly to our paper, they find that foreign sales are impeded significantly by trade frictions, foreign production frictions, and a general difficulty to sell outside the home market that may be due to preferences.

We build on the random coefficients demand model developed by Berry, Levinsohn, and Pakes (1995), which study the U.S. automobile market, and Nevo (2001), who introduces product-fixed effects in this demand estimation framework. This framework uses a flexible structural approach to recover firm markups in a differentiated products market with multi-product firms.[3] Due to the knowledge of assembly locations, we can include a cost shifter (distance of assembly location to market) as an instrument for price in the demand estimation. After recovering demand elasticities, we use optimal pricing strategies to back out costs. Using variation in assembly and headquarter locations, we estimate trade and foreign production costs, as well as the cost of supplying of each characteristic. We also allow for the fact that the source location was endogenously chosen from the set of all available plants. Overall, we estimate both demand and supply in a consistent framework

[2] Their paper uses trade flow data between many more countries than ours, but does not include price or characteristics information and therefore requires stronger assumptions to estimate preferences.

[3] In an alternative approach, De Loecker and Warzynski (2012) use plant level production data to recover markups in a trade setting.

that allows for unobserved demand and cost heterogeneity at the model and market levels.

Next, we use the estimated model to examine the contributions of tariffs, trade and FDI costs, home bias, and taste heterogeneity to domestic firms' home market advantage. We evaluate the contribution of each of these items to firms' home market advantage by a reduced form regression of counterfactual market shares from the structural model on a home market dummy. The home market dummy captures, all else equal, how much larger firm's market shares are in their home country. We find that this home dummy in the reduced form regression falls by far the most if the structural home preference parameter is removed from the utility function, which indicates a key missing element of the existing trade and multinational production literature. In ongoing work, we are exploring how multinational firms are jumping home preferences (analogous to jumping tariffs) by acquiring foreign brands.

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 data set 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.[4] 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, weight and fuel efficiency). These characteristics vary across markets and years. We also constructed a

[4]Brazilian 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.

data set of assembly locations informing us about the countries in which the models in our demand data were assembled any given year. While we do not know exact sourcing locations, i.e. which of these assembly locations supplied a certain model to a market, we solve firms' cost minimizing sourcing decision in our estimation. We complement this data set with market-specific variables such as gas prices, import tariffs on cars, sales taxes, the level and dispersion of household income.

Some manufacturing groups own multiple brands in the global car market. In what follows, we distinguish firms (manufacturing groups such as GM and VW), brands (such as Vauxhall and Opel owned by GM, Audi and Seat owned by VW) and models (such as Vauxhall Corsa 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 fixed as the country from which it historically originates. Across all years and markets, the data set encompasses 27 firms, 60 brands and 596 models. Firms are headquartered in 12 different countries and brands are associated with 15 different countries.

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 35%.

Associating brands with their country of origin, 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 data set). 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. The most striking

Table 1: **Market Concentration**

	Sales	Firms	Top 5	Brands	Top 5	Models	Top 5
BEL	494,126	19	68 %	38	45 %	312	13 %
BRA	2,514,116	16	83 %	23	83 %	97	36 %
CAN	1,134,317	16	65 %	34	50 %	203	22 %
DEU	2,993,401	19	71 %	38	54 %	323	18 %
ESP	1,074,607	20	73 %	39	44 %	290	16 %
FRA	2,032,459	19	81 %	38	65 %	271	25 %
GBR	2,016,893	21	63 %	39	48 %	309	21 %
ITA	2,005,482	21	70 %	41	51 %	283	26 %
USA	10,361,189	18	69 %	40	53 %	291	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.

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.

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, Jensen, Redding, and Schott (2009) and start with

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	-	0.10	-	0.23	0.32	0.13
CAN	0.07	-	-	0.01	-	0.34	0.58
DEU	**0.56**	0.02	0.09	0.01	0.03	0.08	0.21
ESP	0.26	**0.09**	0.27	0.01	0.03	0.11	0.22
FRA	0.19	0.02	**0.53**	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	-	**0.40**	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. Other includes Japan, Korea, China, India, Sweden, Malaysia, Czech Republic, Romania and Russia. The bottom panel excludes these "other" countries and presents market shares within the brand-owning producers in our dataset.

Table 3: **Market Share Decomposition**

	(1) $\ln(\bar{s}_{bm})$	(2) $\ln(N_{bm})$	(3) $\ln(\bar{s}_{bm})$	(4) $\ln(N_{bm})$
$\ln(s_{bm})$	0.587^{***}	0.413^{***}	0.550^{***}	0.450^{***}
	(0.0078)	(0.0077)	(0.0080)	(0.0079)
Observations	1415	1415	1415	1415
R^2	0.803	0.668	0.773	0.695
Share	Units	Units	Revenue	Revenue
Margin	Intensive	Extensive	Intensive	Extensive

Notes: Standard errors in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$)

the identity

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

where s_{bm} is the share of brand b in total market m sales and N_{bm} is the number of models offered. All variables are averages over the entire data period. We then separately project $ln(N_{bm})$ and $ln(\bar{s}_{bm})$ on $ln(s_{bm})$. Table 3 reports the results. The intensive margin accounts 55 to 59 percent of the overall variation, depending on whether the market share is in revenues or units sold.

Table 4: **Home Market Advantage**

	(1) $\ln(s_{bm})$	(2) $\ln(s_{bm})$	(3) $\ln(s_{bm})$
Home brand	1.646^{***}		0.961^{***}
	(0.0818)		(0.0558)
$\ln(N_{bm})$		1.713^{***}	1.531^{***}
		(0.0385)	(0.0364)
Obs.	1415	1415	1415
R^2	0.785	0.888	0.908
Market-year FE	Yes	Yes	Yes
Brand FE	Yes	Yes	Yes

Notes: Standard errors in parentheses. * ($p < 0.10$), ** ($p < 0.05$), *** ($p < 0.01$)

To further gauge how these margins affect the home market advantage, we project $\ln(s_{bm})$, the log share of a brand in a market averaged over the data period, on a dummy variable that takes the value one if the brand is at home and zero otherwise, as well as on the (log) number of models that the brand offers there. Fixed effects control for brands' global popularity and market-year specific conditions. Table 4 confirms our assertion that the extensive margin of model offerings matters for explaining market share differences. Even then, however, there is a large and significant home market effect in the intensive margin. Being a home brand increases market share by 161 percent.[5]

In the production side, there are 49 countries that assemble cars. More than half of models (316 out of 596) are assembled in more than one country, accounting for 91% of total revenue. The market share of models assembled in 8 or more countries is 51%. While we do not know the exact sources of supply, we can analyze each market in terms of potential sources of supply. Column 1 of Table 5 presents the average number of countries in which models consumed in a particular market are assembled. For instance, there are 3.8 countries in which models sold in Brazil are assembled (weighted by models' market shares in Brazil) while models sold in Canada are assembled in 5 countries. Nearest of these plants are on average 1883 km away from Brazilian consumers while models sold in Canada and the US are assembled in more distant locations (column 2).[6] Geography, units costs and trade policy are important determinants of these potential supply locations. Brazil is the most protected country in our data set, with an MFN import tariff on cars equal to 35 % (column 4). The resulting tariff-jumping FDI leads to a higher market share for the models that are assembled domestically (column 3).

Table 6 present average prices and characteristics (weighted by sales) of cars sold in

[5]Using the home effect coefficient from column 4 of table 4, we calculate $\exp(0.961) - 1 = 1.61$

[6]These figures include countries' internal distances. We use the CEPII distance data and their weighted distance measure; the weighted internal distance in Brazil is 1,157km, 1,853km in the U.S., and 1,485 km in Canada.

Table 5: **Supply Locations**

	Supply Locations	**Average Distance**	**Domestic Share**	**MFN Tariff**
BEL	4.4	1,084	76 %	10 %
BRA	3.8	1,883	86 %	35 %
CAN	5.0	3,662	60 %	6.1 %
DEU	4.7	1,069	75 %	10 %
ESP	4.5	1,607	79 %	10 %
FRA	4.2	1,013	77 %	10 %
GBR	4.5	1,429	77 %	10 %
ITA	4.0	1,319	69 %	10 %
USA	4.2	3,618	67 %	2.5 %

Notes: Supply locations 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, weighted by models' market share. The distance measure is based on bilateral distances between the biggest cities of two countries, with inter-city distances being weighted by the population share of the city (Mayer and Zignago 2011). Domestic share is the market share of domestically assembled models. 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.

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.

Table 6: **Prices and Characteristics**

	Price	**HP/Wt**	**Size**	**MPG**	**Gas price**	**MP$**
BEL	32,459	58.2	7.6	34.4	7.3	4.7
BRA	23,689	62.4	6.8	30.1	5.6	5.4
CAN	30,568	91.7	8.3	22.3	2.9	7.6
DEU	35,729	66.6	7.6	29.4	7.3	4.1
ESP	31,633	60.7	7.6	32.7	5.4	6.1
FRA	29,537	57.1	7.3	35.6	7.0	5.1
GBR	31,253	65.3	7.5	30.5	7.0	4.3
ITA	27,514	57.5	7.0	33.5	7.2	4.7
USA	28,792	97.8	8.7	21.0	3.1	6.7

Notes: All variables are averages over models weighted by market share over the data period. Prices are in USD, converted from local currency using mean yearly exchange rates and averaged over the data period. 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 USD. MP$ is miles per dollar (MPG/price).

We finish this section with a reduced form regression of model-market-year specific log prices in table 7. To model characteristics, we add a home brand dummy, the distance to the nearest assembly location and a domestic assembly dummy that takes the value one if the nearest assembly location is the market itself. The home effect is statistically and economically significant, suggesting a price premium of two percent by national brands. Prices are also significantly correlated with the distance to the nearest assembly plant, suggesting that exploiting geographic variation in the assembly locations leads to cost shifters which can be used as instruments for prices in the demand estimation described below.

Table 7: **Price Regression**

	$\ln(price_{jmt})$
$\ln(hp.wt_{jmt})$	0.249***
	(0.011)
$\ln(size_{jmt})$	0.564***
	(0.040)
$\ln(mpg_{jmt})$	0.0191*
	(0.010)
$\ln(dist_{jmt})$	0.0170***
	(0.002)
Domestic assembly	-0.014***
	(0.004)
Home brand	0.019***
	(0.003)
Obs.	7929
R^2	0.985
Market-year FE	Yes
Model FE	Yes

Notes: See table 6 for the description of parameters. Home brand is one if a model belongs to a national brand and zero otherwise. Regression controls for market-year and model fixed effects.

3 Model

Our goal is to separately identify demand and supply factors in driving market outcomes in the global car industry. To do this we propose an equilibrium model of auto pricing which allows us to consistently estimate demand and use optimal firm pricing decisions to recover markups and costs. We then examine how proximity to a market affects both demand for automobiles and the cost of supplying automobiles.

We model the national market for cars in a given calendar year. Manufacturers are endowed with a set of models (i.e., Hyundai Sonata, Hyundai Elantra, etc.) to sell within the market. Each model is endowed with a set of characteristics which include characteristics of the car itself and the set of assembly plants which produce that model. At the start of the year, all manufacturers observe a set of demand and supply shocks for each model which are uncorrelated with model or assembly location characteristics. This assumption implies that a manufacturer chooses to offer a car in a location before observing the model-market demand or supply shock. This is reasonable 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 design period 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, manufacturers simultaneously choose prices at the model level according to a Nash-Bertrand equilibrium. We follow the literature on automobile pricing in assuming that prices are set at the model level and consumers face a single price, so we abstract away from haggling and trim-level pricing (e.g., offering multiple engine types in a particular model). The characteristics of a model are assumed to be the sales-weighted average trim-level characteristics. Consumers then observe these prices and make purchases. Finally, automakers select the assembly location from which to source ordered cars. We allow for heterogeneity at the car level so that a manufacturer may choose to source cars from multiple different assembly locations to supply the same market.

3.1 Demand

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

$$u_{jmti} = \bar{u}(x_{jmt}, p_{jmt}, \beta_{mi}, \alpha_{mti}) + \xi_{jmt} + \varepsilon_{jmti} \tag{1}$$

where x_{jmt} represents the model characteristics—e.g. horsepower per weight, size, fuel efficiency, or brand fixed effects—and p_{jmt} represents the price. Consumers within a market have heterogeneous tastes for characteristics and price, which may be distributed differently in different markets as represented by $(\beta_{mi}, \alpha_{mti}) \sim F_{mt}(\,\cdot\,|\theta^d)$. We assume this distribution is known to the researcher up to a vector of demand parameters, θ^d. Each models receives a market-year specific demand shock, ξ_{jmt} that is common to all consumers within a market. Finally, each consumer receives an idiosyncratic utility shock for each model, ε_{jmti} which is distributed according to the Type-I extreme value distribution.

Consumers in each market in a given year 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 $\bar{u}_{0mti} = \varepsilon_{0mti}$.[7] Let C_{mt} be the set of cars consumers can choose from in market m in year t as well as the no-purchase option. Each consumer chooses the option that maximizes her utility,

$$choice_{mti} = \operatorname*{argmax}_{j \in C_{mt}} u_{jmti}.$$

Integrating out the idiosyncratic consumer taste shock, we have the probability that each

[7]Our specification will include year and market dummy variables in $\bar{u}$ so this normalization does not preclude differences in the value of the outside good across markets.

consumer buys a car given their tastes $(\beta_{mi}, \alpha_{mi})$,

$$Pr(choice_{mti} = j|\beta_{mi}, \alpha_{mi}) = \frac{e^{\bar{u}_{jmti}+\xi_{jmt}}}{\sum_{k\in C_{mt}} e^{\bar{u}_{kmti}+\xi_{kmt}}}.$$

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

$$s_{jmt} = \int Pr(d_{mti} = j|\beta_{mi}, \alpha_{mi}) dF_{mt}(\beta_{mi}, \alpha_{mi}|\theta^d). \tag{2}$$

The demand parameters to estimate are the distribution of tastes for characteristics, θ^d. Following Berry, Levinsohn, and Pakes (1995), we will estimate these parameters by inverting the share equation to recover ξ_{jmt} and construct a set of moment conditions using exogeneity restrictions.

3.2 Supply

Manufacturers 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 cost of sourcing a car i of model type j for market m from location ℓ is,

$$c_{jm\ell ti} = \kappa(h_{jmt}, \theta_1^s)\delta(d_{jm\ell t}, \theta_2^s)e^{\omega_{jmt}-\nu_{jm\ell ti}} \tag{3}$$

where $\kappa_{jmt}(\cdot)$ represents model and market-specific costs of selling model j in market m at time t, which is determined by a vector of observable model and market characteristics h_{jmt} and a vector of parameters θ_1^s. Similarly, $\delta(\cdot)$ represents the effect of trade costs due to sourcing model j from an assembly plant in location ℓ to be sold in market m and depends on a vector of known market-assembly-model characteristics $d_{jm\ell}$ and a vector of parameters θ_2^s. The structural error term ω_{jmt} represents a shock to the marginal costs of selling model j in a given market m at time t. Finally, costs at the car level are affected by an idiosyncratic cost shock, $\nu_{jm\ell ti}$ that accounts for cost differences between assembly locations in their costs, this final cost is revealed to the manufacturer at the time a car is ordered, but after prices for models are set. Producers have full knowledge of ω_{jmt} and all other cost shifters besides $\nu_{jm\ell ti}$ when setting prices. As we show below, the idiosyncratic error $\nu_{jm\ell ti}$ introduces the possibility of "gains from diversification" in assembly locations and rationalizes the possibility that some models are sourced from multiple assembly locations.

Given this model for costs, the manufacturer minimizes costs by sourcing cars from the lowest cost location from its set of available assembly locations at time t, $L_t(j)$,

$$c_{jmti} = \min_{\ell \in L_t(j)} c_{jm\ell ti}.$$

However, the firm must set prices prior to the $\nu_{jm\ell t}$ shock being revealed, therefore it must set prices according to its expected cost of supplying a model by integrating over $\nu_{jm\ell ti}$. We assume $\nu_{jm\ell ti}$ is distributed Type-I extreme value with scale parameter σ_ν.[8] Given this assumption, the probability of sourcing a car from location ℓ is,

$$Pr(i \text{ is sourced from } \ell) = \frac{\exp\left(\frac{-\log\delta(d_{jm\ell},\theta_2^s)}{\sigma_\nu}\right)}{\sum_{k\in L(j)} \exp\left(\frac{-\log\delta(d_{jmk},\theta_2^s)}{\sigma_\nu}\right)}, \tag{4}$$

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,[9]

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

[9]A constant from intergrating the Type-1 extreme value distribution is absorbed in κ_{jmt}.

$$\log c_{jmt} = \log \kappa(h_{jmt}, \theta_1^s) - \sigma_\nu \log \left(\sum_{k \in L(j)} \exp \left(\frac{-\log \delta(d_{jmk}, \theta_2^s)}{\sigma_\nu} \right) \right) + \omega_{jmt}. \tag{5}$$

In this expression, 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 value of σ_{nu} captures "gains from variety" in the sense that the value of an additional assembly location is scaled by σ_ν. Furthermore, as $\sigma_\nu \to 0$, firms source always source from the single location that has the lowest average cost and (5) becomes,

$$\lim_{\sigma_\nu \to 0} \log c_{jmt} = \log \kappa(h_{jmt}, \theta_1^s) + \min_{k \in L(j)} \left\{ \log \delta(d_{jmk}, \theta_2^s) \right\} + \omega_{jmt}.$$

So as $\sigma_\nu \to 0$, only variation in d_{jmk} at the minimum cost location affects the marginal cost of a model.

The supply side parameters to estimate are $\theta^s = (\theta_1^s, \theta_2^s, \sigma_\nu)$.

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. Since a mass of consumers purchase cars, c_{jmt} is exactly known to manufacturers when they set prices, even though they do not know $\nu_{jm\ell ti}$ 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 $\mathbf{p}_m$. Therefore, firm f's profit maximization problem is to choose prices for its portfolio of models within a market J_{mt}^f to maximize profits,[10]

$$\max_{\{p_{jmt}\}_{j\in J_{mt}^f}} \sum_{j\in J_{mt}^f} \left[p_{jmt} - c_{jmt}\right] \cdot N_{mt} \cdot s_{jmt}(p_{jmt}; \mathbf{p}_m^{-j}), \tag{6}$$

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

$$s_{jmt}(\mathbf{p}_m) + \sum_{k\in J_{mt}(f)} \left[p_{kmt} - c_{kmt}\right] \frac{\partial s_{kmt}(\mathbf{p}_m)}{\partial p_{jmt}} = 0. \tag{7}$$

3.4 Identification

The demand parameters are identified via moment condition assumptions on the model-market demand shocks ξ_{jmt}. 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,

$$\xi_{mt} = s^{-1}(s_{mt}, p_{mt}; \theta^d).$$

We then identify the model using a vector of instruments z_{jmt} such that $E[\xi_{jmt} z_{jmt}] = 0$. Note that the model precludes model price from being used as in instrument since it is endogenously determined. Our model delivers three different classes of available instruments.

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

First, model characteristics (e.g., x_{mjt}) which are determined before demand shocks are realized are uncorrelated with demand shocks, though they will clearly be correlated with price. Second, as discussed in Berry, Levinsohn, and Pakes (1995), characteristics of other models are similarly available as instruments, since they effect prices though the markup term in (7). Finally, functions of observed variables that affect costs ($h_{jmt}, d_{jmt\ell}$) may be used as instruments since they are uncorrelated with ξ_{jmt} given the timing assumption but affect prices through the cost term in (7). In our case, these variables are the drivers of trade costs, such as the distance to assembly locations. We have experimented with all three types of instruments. In practice, we have found that model characteristics and cost shifters perform well, whereas functions of rival characteristics (the so-called BLP instruments) tend to be weak.[11] Therefore, we use characteristics and cost shifters as instruments in our preferred specification.

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-time subscripts and focus on a single market.Given demand parameters and observed prices and shares, all the terms in (7) are known with the exception of c_j. 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,

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

then we can write (7) in vector notation, $s(p) - \Omega(p - c) = 0$, and can easily solve this for the vector of marginal costs,

$$c = [p - \Omega^{-1} s(p)].$$

[11]This finding is consistent with the predictions of Armstrong (2014) relating to the viability of markup instruments when there are a large number of firms. These predictions are also corroborated by Monte Carlo simulations performed by Conlon (2013).

Once costs are recovered, we can identify the cost side parameter θ^s from (5) and the assumption that $E[\omega_{jmt}|(h_{jmt}, d_{jmt\ell})] = 0$. While identification of $\kappa(h_{jmt}, \theta_1^s)$ is straightforward given regularity conditions that will be satisfied by our parameterization, the contribution of individual assembly locations $\delta(d_{jmk}, \theta_2^s)$ and σ_ν are more subtle.[12] First, consider identification of σ_ν, the variance of the idiosyncratic car cost shock. Suppose that all assembly locations were identical, that is for a given model, $\delta(d_{jmk}, \theta_2^s) = \delta$. In this case, the only reason to source from a particular location would be due to the extreme value error, $\nu_{jmt\ell i}$. There would be a cost advantage to operating multiple assembly locations in that you would get a new draw of this idiosyncratic cost shock for each location. Therefore, the extent to which average costs decline as we vary the number of production locations identifies σ_ν. In the extreme, suppose $\sigma_\nu = 0$, then additional assembly location will not reduce average costs at all. For this reason, we associate the parameter σ_ν with "gains from variety" in assembly locations. 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 plant is used to source cars. Again, consider the extreme case when $\sigma_\nu = 0$. Then, variation in an assembly locations characteristics *only* affect average costs if that location is the low-cost location, so variation in an assembly characteristic across locations first identifies which is low cost (since only this one affects average costs), and then identifies the parameter for that characteristic based on the size of the change in costs. In summary, each element of θ_2^s is identified as long as it affects $\delta(d_{jm\ell}, \theta_2^s)$ 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 and

[12]We are of course assuming a location normalization in $\delta(d_{jmk}, \theta_2^s)$, as is common in discrete choice models, without loss of generality. A scale normalization on σ_ν is not necessary as we explain below.

$d_{jm\ell}$ to infer the effect of θ_2^s.

4 Estimation

The model is estimated in two stages. We first estimate the demand side. We then use firms' profit maximization conditions and the estimated demand parameters to recover the marginal cost of supplying each model to each country. Finally, we use these recovered costs to estimate the supply side.[13]

4.1 Demand Parameterization and Estimation

We start by parameterizing the utility function to be quasilinear in price, and quadratic in tastes for characteristics:

$$\begin{aligned}\bar{u}(x_{jmt}, p_{jmt}, \beta_{mi}, \alpha_{mti}) =& \beta_{mi}^{hp}\text{hppwt}_{jmt} + \beta_{m}^{hp2}\text{hppwt}_{jmt}^2 \\ &+ \beta_{mi}^{sz}\text{size}_{jmt} + \beta_{m}^{sz2}\text{size}_{jmt}^2 \\ &+ \beta_{mi}^{md}\text{mpd}_{jmt} + \beta_{m}^{md3}\text{mpd}_{jmt}^2 \\ &- \alpha_{mti}p_{jmt} + \iota_t + \psi_{mb(j)},\end{aligned} \tag{8}$$

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 traveled per dollar consumed based on market prices for gas (according to city fuel efficiency rating). This specification allows for the fact that consumers' marginal taste for a characteristic may be increasing or decreasing in the amount of the characteristic provided. For example, we would expect the marginal utility of size to decrease as the size of the car

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

gets larger. We further assume that linear tastes are normally distributed around a market specific mean with common variance across markets,[14]

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

while the quadratic parameters are market specific but are constant across consumers within a market. This 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.

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

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

Because we do not observe individual consumers, we simulate consumers' income from a log normal distribution determined by fitting mean household income and the Gini ratio for each market.[15]

[14]In 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.

[15]While we have annual household income data for each country, the Gini ratio is constant across years for a given country. Again, greater flexibility across countries could in principle be allowed at the cost of computational tractability due to the introduction of additional nonlinear parameters.

The final two terms in (8) are a set of time and brand-country fixed effects Time fixed effects capture global shocks to automobile demand. The brand-country fixed effect, $\psi_{mb(j)}$, captures revealed preference for different brands within each country. For each model j, $b(j)$ represents its marketing 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).[16] This is important since firms frequently use branding as a method of accentuating vertical differentiation. Moreover, to the extent that consumers exhibit a home preference for their domestic brands, this preference is absorbed in these brand country fixed effects. Below, we follow Nevo (2001) to recover home preferences using an auxiliary regression.

Under this parameterization of the demand model, $\theta^d = (\beta_m^{\bar{x}}, \beta_m^{x2}, \sigma_x, \bar{\alpha}, \pi_\alpha, \sigma_\alpha, \iota_t, \psi_{mb})$ 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 estimate the model by minimizing the generalized method 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. This estimator is asymptotically normal with variance covariance matrix provided in Berry, Levinsohn, and Pakes (1995) and Nevo (2000). For the instrument set we use model characteristics including time and brand dummies, a dummy for whether the model can be produced at a domestic assembly location, tariff rate of the closest available assembly location, the number of production locations interacted with a market dummy, and the minimum distance to an assembly location

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

interacted with a market dummy.[17]

4.2 Demand Estimates

The demand estimates are presented in Table 8. The estimates of the distribution of price sensitivity are statistically significant and indicate that price sensitivity is decreasing in income (π_p) with substantial dispersion conditional on income (σ_α). The other random coefficients are insignificant and smaller in magnitude, which may indicate that most heterogeneity in tastes is captured either through price sensitivities or in mean differences across countries. Considering the tastes for characteristics, the patterns are largely intuitive. All countries have a positive marginal utility for size, which is decreasing in size for most countries. Taste for horsepower is typically positive and appears close to linear in most countries. Taste for miles per dollar is most heterogeneous across countries. Looking across countries, Canada and the United States are remarkably similar in their tastes, while there are substantial differences between other countries.

The brand country dummy estimates, ψ_{mb}, capture revealed preference tastes for a particular brand within country m. These will include the unobserved quality of the brand (such as engineering expertise), its marketing cachet, the availability of dealerships and repair shops for the brand, et cetera. One key characteristic of a brand that is captured by the brand-country fixed effect is whether or not it is considered a "home brand" by consumers. For example, one might expect that Germans prefer Volkswagen because they view it as a German brand, while Italians might derive extra utility from purchasing Fiat. To asses the strength of home preference in brands, we follow Nevo (2001) and Chamberlain

[17] Results are robust to alternative specifications of the instruments. Instruments are standardized to ensure that the weight matrix is well-conditioned.

Table 8: Parameter Estimates for the Demand Model

Variable	Estimate									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA	R.C. Std
HP per Weight	-0.042	0.212	0.356	0.295	-0.480	0.804	0.487	0.161	0.321	0.047
	(0.681)	(0.505)	(0.413)	(0.480)	(0.426)	(0.486)	(0.384)	(0.477)	(0.559)	(0.929)
HP per Weight^2	0.011	0.001	-0.003	0.006	0.069	-0.029	-0.009	0.004	-0.005	
	(0.042)	(0.037)	(0.020)	(0.029)	(0.027)	(0.034)	(0.024)	(0.033)	(0.026)	
Size	2.633	6.325	6.855	3.040	8.273	7.465	5.742	5.606	5.429	0.218
	(1.814)	(1.684)	(1.697)	(1.630)	(1.303)	(1.152)	(1.056)	(0.974)	(1.471)	(0.605)
Size^2	0.099	-0.220	-0.281	0.001	-0.364	-0.337	-0.218	-0.219	-0.242	
	(0.143)	(0.113)	(0.097)	(0.107)	(0.077)	(0.068)	(0.071)	(0.074)	(0.081)	
MPDCITY	0.728	-1.996	-0.058	1.501	1.015	-0.232	-0.893	-2.700	-1.615	0.029
	(0.658)	(0.804)	(0.888)	(0.952)	(0.497)	(0.968)	(0.786)	(0.820)	(0.724)	(1.528)
MPDCITY^2	-0.037	0.167	0.007	-0.105	-0.093	0.070	0.097	0.286	0.072	
	(0.043)	(0.075)	(0.043)	(0.094)	(0.040)	(0.086)	(0.061)	(0.078)	(0.045)	

Other Nonlinear Parameters

Constant R.C.	$\overline{\alpha}$	π_p	σ_p
1.846	9.436	-0.730	0.782
(5.796)	(1.347)	(0.130)	(0.173)

The units for HP per weight, size, MPD, and price are horse power per 100 kg, m^2, mileage per dollar, and 10 thousand dollars, respectively. This specification uses brand-country dummies.

Table 9: Structural Home Preference Parameters, $\hat{\rho}$

	Country-Spec.	Homog.
Home	-	0.991
	-	(0.058)
DEU	0.800	-
	(0.145)	-
ESP	1.682	-
	(0.517)	-
FRA	1.221	-
	(0.153)	-
GBR	1.094	-
	(0.155)	-
ITA	1.509	-
	(0.169)	-
USA	0.629	-
	(0.107)	-

(1982) to estimate the home market effect from,

$$\psi_{mb} = \rho \mathbf{1}[b \text{ is a domestic brand in } m] + \eta_b + \mu_m + u_{mb}, \tag{9}$$

where ρ is the home market preference after controlling for country invariant brand quality and cross-country differences in preference for automobiles. This parameter, which we refer to as the structural home preference parameter, is the revealed preference for buying home brands. This preference may arise due to nationalistic feelings among consumers or through other underlying channels, such as an accumulation of consumer capital from having a long history within a country or a wide dealership network within the country which are common with native brands. If we had data on these channels—such as dealer network size or cumulative advertising expenditure—it could be added to (9) and estimated separately from the home preference. However, our estimates of θ^d and the implied costs from our model would not be affected by the addition of this data since we are already controlling for brand-country effects in (8).[18]

Table 9 presents our estimates of the home preference parameter ρ. Column I assumes that the parameter is constant across countries while Column II allows the preference to vary across countries where a national brand is sold. Overall, we find a substantial and highly significant preference for native brands, even after controlling for local preferences for characteristics, the quality of brands, and prices. In addition, Column II indicates that there is substantial heterogeneity in home preference across countries. We find that home preference is highest in Spain and Italy, where the home preference is estimated to be over 50 percent larger than the homogeneous estimate.[19] On the other hand, the countries with the smallest home preferences are Germany and the United States. These results are consistent with the findings of Goldberg and Verboven (2001a) when studying European car sales from 1980-1993, who find the strongest home preference in Italy, and the weakest in Germany.

4.3 Elasticities and Markups

While the demand parameters themselves are difficult to interpret on their own, they directly imply elasticities and markups for each model. In Table 10 we present the elasticities and cross-elasticities for a subset of models in the subset of markets where these models compete. Looking at the own-price elasticities, we see they vary in an intuitive way. The luxury models are are the least elastic while the three compact-to-midrange models are more elastic. The Toyota Corolla, which is never a home-brand in our market, has the highest median own-price elasticity. When we consider the cross-elasticities, the table il-

[18] These estimates would however give us the ability to run a counterfactual removing home preference while holding dealership networks fixed, as we will discuss below.

[19] The primary national brand in Spain is SEAT, which is wholly owned by Volkswagen, but still projects a Spanish identity to the extent of offering models named the Leon, Toledo, and Alhambra.

Table 10: Median own and cross-price elasticities for select models.

	Audi A6	Ford Focus	Mercedes E 350	Renault Clio	Toyota Corolla
Audi A6	-7.449	0.013	0.118	0.004	0.010
Ford Focus	0.030	-9.938	0.016	0.148	0.111
Mercedes E 350	0.068	0.004	-7.203	0.002	0.001
Renault Clio	0.005	0.219	0.001	-9.603	0.020
Toyota Corolla	0.003	0.274	0.001	0.213	-10.651

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.

lustrates 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, although the cross elasticities in the luxury class are generally smaller than for more quotidian models, suggesting that price competition is less fierce in in the luxury portion of the market. Interestingly, the cross-price elasticities for the Toyota Corolla with its competitor models—the Renault Clio and Ford Focus—are quite asymmetric, this is due to the fact that the median price-to-share ratio is higher for the Corolla (which never enjoys a home market advantage in our markets) than for these other models.

Our estimates of elasticities directly predict markups according to the firms first order condition (7). Table 11 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 data set. The highest markups tend to be found in the United States, although several smaller model such as the VW Golf, Mini, and Ford Fiesta tend to have smaller markups in the United States than they do in European countries. Looking across brands, luxury cars, such as the BMW 530 and Mercedes E 350, tend have the highest markups. SUVs such as the VW Tiguan command high markups in Europe and lower markups in the United States and Canada,

Table 11: Median markups of select models across years (percent).

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Audi A4	9.6	16.7	15.5	19.6	18.6	19.9	18.2	18.8	16.0
Audi A6		18.2	17.5	19.7	19.4	21.2	18.7	19.4	20.3
BMW 530		17.9		16.3	17.7	20.8	18.2	18.6	19.5
BMW X3		16.4	15.9	15.7	16.5	19.3	17.6	17.6	18.6
Chrysler 300		14.4	15.6	13.0	12.4	17.1	15.2	15.8	16.7
Ford Fiesta	9.0	10.4		14.1	13.7	13.2	13.5	13.1	12.9
Ford Focus	10.4	10.1	11.7	11.6	13.2	13.7	13.6	14.1	12.0
Honda Accord		11.9	14.0	11.6	14.5	16.2	14.4	13.7	14.0
Honda CR-V	11.7	11.8	14.9	11.9	14.6	16.3	14.3	14.7	13.8
Jaguar XF		16.0	17.2	13.7	16.8	18.7	18.1	15.9	20.5
Jeep Grand Cherokee		15.0	17.1	13.3	15.2	16.9	15.7	16.3	18.0
Lexus RX 450		16.5	17.6	14.7	16.5	19.1	16.5	16.0	21.0
Mercedes E 350		17.3	17.3	16.9	17.4	19.9	17.4	17.0	20.8
Mini New Mini	10.3	10.2	11.7	12.1	12.8	13.7	12.2	14.0	11.6
Renault Clio	8.1	11.3		14.0	14.8	17.5	12.5	12.6	
Toyota Corolla	10.2	9.9	13.5	11.9	11.8	12.5	10.2	11.4	13.0
Toyota RAV-4	11.1	12.1	13.9	11.9	14.3	16.1	14.4	14.9	14.2
VW Golf	10.8	12.9	11.3	18.0	17.0	15.9	15.1	16.0	11.7
VW Passat	10.7	14.9	13.1	19.3	18.1	18.6	16.2	18.2	14.3
VW Tiguan	11.1	15.6	13.4	18.9	18.8	19.0	17.3	18.4	14.0

where they face significantly more competition. Overall, the model produces intuitive estimates of markups—and hence marginal costs—across countries and models.

To illustrate the importance of home preference on markups, Table 12 presents weighted average markups aggregated to the firm-country level. Across the table, we see that firms tend to charge their highest markups in their home countries. Volkswagen in Germany, Ford and General Motors in the United States, and most strikingly PSA (Peugeot Citroen) 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 the findings of the home preference parameters in Table 9. Were we not accounting for home preference explicitly in the demand side—while simultaneously controlling for price and taste for characteristics, these

Table 12: Weighted average markups of manufacturers across markets (percent).

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	10.9	10.6	13.6	13.3	11.5	13.8	12.4	18.4	13.8
Ford	9.4	11.1	13.5	12.8	12.1	13.9	14.0	13.5	14.4
General Motors	10.4	11.1	14.0	13.1	12.4	13.8	13.7	13.5	16.1
PSA	9.5	12.3		12.9	13.5	19.5	12.8	13.7	
Toyota	10.1	10.9	14.1	12.8	12.1	14.2	12.8	13.0	14.9
Volkswagen	10.7	13.8	12.7	18.7	15.3	16.4	15.5	15.9	14.1

implied markups would be lower which would result in higher implied cost for home-branded models, biasing our estimates of costs.

4.4 Supply Parameterization and Estimation

In the second stage of our estimation procedure, we use the costs implied by the model to estimate our supply side using nonlinear least squares. To do so, we parameterize $\kappa(h_{jmt}, \theta_1^s)$ and $\delta(d_{jm\ell t}, \theta_2^s)$ 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 computational simplicity, we assume both follow a Cobb-Douglas-like formulation so that they are linear in logarithms.

For market-model costs, we assume,

$$\begin{aligned}\log \kappa(h_{jmt}, \theta_1^s) = &\kappa^{hp} \log \mathrm{hp}_{jmt} + \kappa^{wt} \log \mathrm{wt}_{jmt} + \kappa^{sz} \log \mathrm{size}_{jmt} \\ &+ \kappa^{mg} \log \mathrm{mpg}_{jmt} + \kappa^{m} + \kappa^{j} + \kappa^{t}.\end{aligned} \tag{10}$$

As opposed to the demand side, we allow costs to be determined by horsepower and weight separately, rather than their ratio. This is intuitive because we would expect both to increase the cost of the car, whereas on the demand side we were using their ratio as a measure of acceleration while leaving size to proxy for the size of the car. We also include miles per gallon—rather than miles per dollar—on the cost side, eliminating the effect of local gas prices. 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 model, market country, and time fixed effects. In contrast, the demand side includes brand-country and time effects on the demand side. We prefer this specification because it allows costs 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.[20] Even when we include a set of model fixed effects on the supply side, the effect of characteristics on cost are still identified due to variation in the characteristics of a model both across countries and across years.

The final element of the supply side is the assembly location specific cost function $\delta(d_{jm\ell}, \theta_2^s)$,

$$\begin{aligned} \log \delta(d_{jm\ell t}, \theta_2^s) = & \delta^{mdist} \log \mathrm{dist}_{m\ell} + \delta^{dom} \mathbf{1}[\ell = m] + \delta^{ctg} \mathbf{1}[\ell \text{ is contiguous to } m] \\ & + \zeta \log(1 + \mathrm{tariff}_{m\ell t}) + \delta^{hqdist} \log \mathrm{dist}_{h(j)\ell} + \phi_\ell. \end{aligned} \tag{11}$$

The first three terms capture the effect of trade costs, including a direct effect of distance as well as dummies to control for domestic and contiguous trade, in a traditional iceberg-like fashion.

The fourth term, $\zeta \log(1 + \mathrm{tariff}_{m\ell t})$, captures the effect of import tariffs on costs.[21] We estimate the parameter ζ to account for the fact that import tariffs are ad valorem based on the reported port cost of the car, which is likely to be lower than the marginal cost

[20] We 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.

[21] The time variation in tariffs is due to several types of events during the data period: some assembly countries become a member to 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).

of the car implied by profit maximization, since the later includes internal shipment and marketing costs. Below, estimate the model both holding ζ fixed at one (the case where tariffs are paid on the full marginal cost) and allowing it to be estimated.[22]

The second to last term, $\delta^{hqdist} \log \text{dist}_{h(j)\ell}$, accounts for the impact of distance between a firms headquarters and the assembly location. Costs may be larger for far away plants due to monitoring or communication costs between a headquarters and plants, or due to shipment of intermediate inputs.[23]

Finally, we control for productivity differences across assembly locations with a location fixed effect, ϕ_ℓ that 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.

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

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

where,

$$\omega_{jmt}(\theta^s) = \log c_{jmt} - \log \kappa(h_{jmt}, \theta_1^s) + \sigma_\nu \log \left(\sum_{k \in L_t(j)} \exp \left(\frac{-\log \delta(d_{jmkt}, \theta_2^s)}{\sigma_\nu} \right) \right).$$

In practice, we make two adjustments to this estimator. First, we find that the objective function is minimized at a very low value of σ_ν, indicating that the gains to variety are very small. We have found that $\hat{\sigma}_\nu = 0.003$. A low estimate of σ_ν is consistent with the findings

[22]It might be better to specify the tariff term as $\log(1 + \zeta \text{tariff}_{m\ell t})$, in which case ζ has the direct interpretation as the proportion of the model's cost subject to tariffs. This formulation complicates the model somewhat by introducing an additional nonlinearity, but we plan to experiment with it in a future version.

[23]See Giroud (2013) and Tintelnot (2014) for a discussion and evidence for such frictions.

of Head and Mayer (2015), who report that European data on the direct sourcing flows of cars at the model level indicates that vast the majority of models are sourced from a single location. In their data, only five percent or models are actually sourced from multiple locations to the same market. For such low values of σ_ν, the probability of sourcing from any particular location converges converges to either 1 (if that location is the minimum cost supplier) or 0 (if it is not) for most models, and the model predictions about marginal costs become insensitive to the precise value of σ_ν. As a result, the objective function becomes almost flat in this region, although estimates of $\log \kappa(h_{jmt}, \theta_1^s)$ and $(\log \delta(d_{jm\ell t}, \theta_2^s)/\sigma_\nu)$ are precisely estimated. Therefore, we fix σ_ν at 0.01 and estimate the remaining parameters of the model.[24]

A second adjustment is needed due to the presence of assembly country fixed effects, ϕ_ℓ. Our data on assembly locations comes from a worldwide matching car models to assembly locations, but we do not know the actual trade flows from an assembly location to markets at the model level.[25] On the other hand, our market data comes from 9 countries in Europe, North America, and South America only. Therefore, some assembly locations in our data are never used as sources for our market locations. In this case, the model correctly identifies the sourcing probability to be almost zero, but the assembly location dummy is not identified, as ϕ_ℓ in this case may be arbitrarily high. However, this loss of identification for a subset of "remote" assembly locations has no effect on identification of the remaining parameters, and we deal with it by treating the ϕ_ℓ as profiled nuisance parameters, effectively dropping the unidentified parameters from the estimation.[26] After

[24] In the future, we will verify that estimates and counterfactual results are insensitive to fixing σ_ν within a range around its minimizer.

[25] If we had this data, our supply side estimation could easily be extended to take advantage of it, which would provide much more precise information on σ_ν.

[26] This assumption has a mild implication for the counterfactual analysis, we assume that any change in cost is not so large that some model starts being sourced from an assembly location from which *no* model from *any* market in our data has been sourced. We do however allow models to start being sourced form assembly locations which had not previously sourced that particular model.

making these two adjustments, the estimator we use in this version of the paper is, for $\theta^s = (\kappa, \delta, \zeta)$

$$(\hat{\kappa}, \hat{\delta}, \hat{\zeta}) = \underset{(\kappa,\delta,\zeta)}{\text{argmin}} \left[\min_{\phi} \sum_{m=1}^{M} \sum_{t=1}^{T_m} \sum_{j=1}^{J_{mt}} \omega_{jmt}(\theta^s)^2 \right],$$

where,

$$\omega_{jmt}(\theta^s) = \log c_{jmt} - \log \kappa(h_{jmt}, \theta_1^s) + \bar{\sigma}_\nu \log \left(\sum_{k \in L_t(j)} \exp \left(\frac{-\log \delta(d_{jmkt}, \theta_2^s)}{\sigma_\nu} \right) \right),$$

and $\bar{\sigma}_\nu = 0.01$. This estimator is consistent and asymptotically normal. While the standard errors presented in this version do not correct for the fact that c_{jmt} is a function of estimated demand parameters, we will add this correction (or jointly estimate the demand and supply models) in a future version of the paper.

4.5 Supply Estimates

The estimates of the supply side are presented in Table 13. We present two parameterizations. In Column I the tariff incidence parameter, ζ is fixed at one—so the tariff applies to the full marginal cost of the car—and in Column II we allow ζ to be estimated. When estimating ζ, we find that it is roughly 0.9, 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 to do internal delivery and marketing. While we do not reject the null hypothesis that $\zeta = 1$, we do find that estimating ζ has a dramatic effect on the estimates of trade costs. In particular, the cost of distance rises, and the benefits of being domestic and contiguous both decrease (i.e., the coefficients δ^{dom} and δ^{ctg} become less negative). This is intuitive because tariff rates are positively correlated with distance and negatively correlated contiguity due to the presence of regional free trade areas.[27] As a consequence,

[27] Of course, tariffs are zero when the assembly plant is domestic.

Table 13: Cost Estimates.

Variable	I	II
Assembly to Market, δ^{mdist}	0.0011	0.0061
	(0.0037)	(0.0038)
Domestic Location, δ^{dom}	-0.0196	-0.0146
	(0.0059)	(0.0060)
Contiguous Location, δ^{ctg}	-0.0115	-0.0080
	(0.0037)	(0.0038)
Distance to HQ, δ^{hqdist}	0.0186	0.0174
	(0.0059)	(0.0061)
Tariff, ζ		0.8978
		(0.0854)
Horsepower, κ^{hp}	0.2681	0.2718
	(0.0509)	(0.0515)
Weight, κ^{wt}	0.1626	0.1662
	(0.0365)	(0.0367)
Size, κ^{sz}	0.3247	0.3249
	(0.0164)	(0.0164)
Miles per Gallon, κ^{mg}	-0.0143	-0.0152
	(0.0113)	(0.0112)
Fixed σ_ν	0.01	0.01

Car cost, distance measures, tariff, and car characteristics are in logarithm.

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 and contiguity.

Focusing on Column II of Table 13, we find that, after controlling for tariffs, distance and contiguity, the impact of distance distance between assembly location and market is positive, but relatively small in magnitude. Doubling the assembly to market distance only increases costs by 0.6 percent. However, it is important keep in mind there is extremely wide variation in distance to market in the automobile industry, with some cars being produced domestically and others being shipped across oceans. Moreover, there are strong cost benefits to producing cars domestically or in contiguous countries, doing so reduces

Table 14: Weighted average external shipping cost (assembly to market, including domestic and contiguity effects) in overall cost for firms across markets (percent of marginal cost).

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	0.0	3.2	0.6	1.8	2.4	2.0	2.8	1.3	0.3
Ford	0.2	2.5	0.9	1.7	2.8	2.0	3.1	2.8	0.3
General Motors	0.2	2.7	0.8	1.1	2.3	1.8	1.5	2.2	0.3
PSA	0.6	2.7		1.8	1.1	0.9	2.6	1.9	
Toyota	0.2	3.9	0.8	2.9	3.0	2.9	3.0	3.0	1.1
Volkswagen	0.1	2.2	2.4	0.7	1.3	1.4	2.4	2.1	2.3

marginal costs by 1.5 and 0.8 percent respectively. We also find a significant effect of distance between an assembly location and headquarters. This suggests that there are non-trivial monitoring and management costs related to producing remotely. Considering the effect of characteristics on the cost side, the impact of horsepower, size and weight all have the expected sign and are statistically significant. While the estimate of the cost to provide miles per gallon is slightly negative, it is not significant.

To get a sense of the importance of external "shipping costs," we compute the additional costs firms pay for each model relative to what they would pay if the model were produced in the sales market. To be clear, Table 14 isolates the the effect of distance, domesticity, and contiguity on costs relative to domestic production as a share of the total marginal cost of the model.[28] While there is some variation by firm and country, the shipping costs typically account for between one and three percent of the marginal cost. On a thirty thousand dollar car (roughly the median in our data) this amounts to between three hundred to nine hundred dollars.[29] As we would expect, these costs tend to be relatively

[28]These figures do not reflect differences in tariffs or the impact of moving production on remote production costs due to distance between assembly and headquarters. They do account for the fact that firms are endogenously choosing to source their models from the cost minimizing assembly location.

[29]Of course, this is the average without conditioning on where the car is produced, so it includes domestically produced cars and the average cost conditional on the car being produced abroad is higher.

low in the firm's home country, however even home firms tend to import at least some proportion of their cars from abroad, generating positive external shipping costs. However, in Brazil, where many firms choose to locate foreign assembly facilities aimed at producing for South American markets, average shipping costs are actually lower as a share of costs than in home countries.[30]

Table 15 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 home country, although they are not zero since, again, firms source at least some models in domestic markets from abroad. The costs tend to grow roughly in line with distance between the market and home country. 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 14. 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.

We conclude this section with an out of sample test of our cost estimates. Our model delivers implied trade flows between countries at the model level through (3) and (4). We aggregate these flows up to the country-pair levels and compare them to trade flows in the automobile sector reported in the WITS database of the World Bank.[31] Figure 1 presents

[30]It 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.

[31]This data is trade flows reported by importers in HS6 product categories associated with assembled cars. 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.

Table 15: Weighted average remote production cost (assembly to HQ) in overall cost for firms across markets (percent of marginal cost).

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	5.2	1.1	4.7	1.5	1.7	1.4	1.9	1.1	4.4
Ford	2.6	1.8	0.1	1.7	1.4	1.9	1.5	1.6	0.2
General Motors	2.5	2.5	0.3	2.5	2.4	2.5	2.4	2.5	0.2
PSA	5.2	1.4		1.5	1.3	1.0	1.8	1.7	
Toyota	6.4	2.9	4.6	3.0	2.5	3.1	2.7	3.0	3.6
Volkswagen	5.8	1.1	3.7	0.8	1.7	1.2	1.2	1.1	4.1

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 1. In fact, the regression estimates a slope of 0.77 and the R^2 of this regression is 0.35. There are many reasons why we fail to match the aggregated trade flows perfectly. Our costs are not intended to represent the costs at importing, 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.[32] 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 aggregated data as well as it does provides some degree of confidence that the model is producing capturing the essential drivers of market outcomes.

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

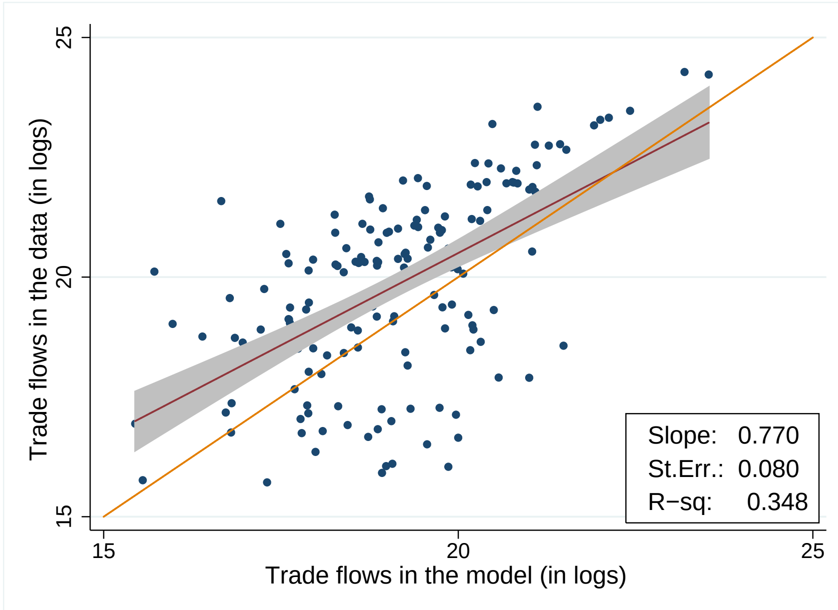

Figure 1: Predicted trade flows and data

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 have proposed and estimated a structural model that accounts for various demand and supply channels that could generate this home market effect. These include a preference for domestic brands, cross-country heterogeneity in tastes for characteristics, trade costs, tariffs and remote production costs. In this section, we use these estimates to asses the role of each in contributing to home market advantage.

To do so, we conduct a series of counterfactuals and re-run the same reduced form regression to reveal the extent of home market advantage. To begin with, we use our structural estimates as a baseline. At this baseline, the prices and market shares for each model exactly match those observed in the data. To establish the degree of home market advantage in the data, we run the following reduced form regression,

$$\log(s_{jmt}) = \beta \mathbf{1}[b(j) \text{ is a home brand in } m] + \alpha_j + \gamma_m + \varepsilon_{jmt} \tag{12}$$

and define the parameter β as the extent of "home market advantage." We do not interpret β causally, but instead use it as a useful metric of how home brands correlate with higher market shares in the data. We then propose a series of counterfactuals where we remove demand and cost channels which could lead to home market advantage from the structural model, and use the revised parameter estimates to recalculate equilibrium costs, prices, and shares under the model. We do compute the new equilibrium prices via iterated best response, starting from the observed equilibrium prices in the baseline.[33] When costs channels change, we allow firms to re-optimize their sourcing locations from the set of available assembly locations from each model. However, we hold the set of assembly locations, and the choice set within each country fixed. Therefore, these counterfactuals should be interpreted as "medium run" in the sense that firms can adjust sourcing and prices, but not entry or exit of models into markets or the construction or closure of assembly plants. Using these counterfactual market shares, we re-estimate (12) to determine the change in the home market advantage. Clearly, the majority of these counterfactuals do not represent changes that would be achievable via policy. However, our goal is use them as thought experiments to illustrate the drivers of home market advantage.

Table 16 displays the results from the reduced form regression of log model market shares on a home dummy as well as brand and country-year fixed effects. In the baseline, where the model prediction of market shares at the model level fits the data perfectly, the

[33] As 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 16: Home market advantage under counterfactual scenarios

Scenario	Normalized coefficient
Baseline	1.00
No Home Preference, homogeneous	0.05
No Home Preference, country-specific	0.07
All countries have French tastes for characteristics	0.92
All Tariffs Eliminated	0.97
No Shipping Frictions	0.90
No Remote Production Frictions	0.96
No Tariffs, Shipping or Remote Production Frictions	0.87

estimated coefficient implies that domestic models enjoy 130 percent larger market share at home than abroad. We normalize this coefficient to one to highlight how this result changes under various counterfactuals.

First, we remove consumers' direct preference for home brands by eliminating the structural home preference estimated in Table 9. We can do this in two ways, treating the home preference as homogenous across countries or as country specific. Either way, we find that eliminating home brand preference has a dramatic effect on the home market advantage, which falls to only 5 to 7 percent of its former level. This is by far the largest impact on home market advantage of any of the counterfactuals we attempt.

Another possibility is that home models designed to better match local tastes for characteristics. To consider this possibility, we ask how home market advantage would be affected if the distribution of tastes for characteristics in all markets was identical to those in France. While this does result in a reduction in the home dummy estimate, it still remains at 92 percent of its total value. One possible reason for this is because domestic brands tend to have a large number of models throughout characteristic space in their home countries. As a result, they are well positioned to respond to dramatic shifts in local tastes relative to foreign brands.

Next, we turn to supply side explanations for the home market advantage. We first consider the removal of all tariffs on automobile trade. This results in a decline in home market advantage, but it remains at 97 percent of its original level. Digging deeper into the data, one reason for this seems to be that many of the most popular models of foreign brands are either produced domestically or in countries where tariffs are already low—if not zero due to regional free trade agreements. As a result, eliminating tariffs has only minor effect on costs (Table A.2, in Appendix A) which feeds through to a small decline in home market advantage. We get a somewhat stronger effect when we remove shipping frictions—including the impact of distance as well as the domestic and contiguity effects—from the model. This results in a home market advantage that is 90 percent of the baseline, the largest individual effect we find on the cost side. Removing the remote production friction from the model has only a small effect home market advantage, which remains 96 percent of its baseline level.

Given the large difference between the impact of removing home preference on the demand side and the cost side effects, we consider one final counterfactual removing all cost-side frictions associated with trade from the model. The resulting home preference parameter is 87 percent of the baseline. This leads us to conclude that demand side effects, and home brand preference in particular, are the key channel which gives rise to home market advantage in the automobile industry in the medium run, while cost side elements play a substantive but supporting role.

It is important to keep in mind that, in the long run, firms endogenously determine production locations, their portfolio of models to offer in each country, and also fixed investments such as marketing and dealer network density which play a role in determining brand popularity within a country. While our counterfactuals hold these elements of the market fixed, they are able to shed light on their importance. One particularly interesting case is the counterfactual where we change local tastes in all markets to match those of France. Relative to the United States our estimates indicate that French consumers' tastes for size and horsepower per weight decrease much more quickly with the size and power of the car. As a result, French consumers tend to prefer smaller, less powerful cars relative to US consumers. Also, neither of the leading French manufacturers, Peugeot S.A. (PSA) or Renault, offer models in the United States.[34] When we run the counterfactual ascribing French tastes for characteristics to US consumers, market shares change dramatically (Table A.1). The primary beneficiaries of the shift are Japanese firms (primarily Toyota, Honda, and Nissan), who sell a wide variety of smaller models in the United States. In fact, their share in the United States rises from 43 to 74 percent, despite the fact that the increase their markups from 14 to 30 percent (Table A.4). At the same time, the market share of European brands collapses from 9.4 to 1.5 percent. European offerings in the United States are dominated by German brands Volkswagen, BMW, and Mercedes who tend to offer relatively more high-end powerful models and SUVs rather than the broader spectrum of models they sell in Europe.

We can contrast the outcomes from this experiment in the United States with those in Germany, where Peugeot and Renault are substantial players. In Germany, the Japanese share rises only slightly, from 10.9 to 13.6 percent, however the German share falls substantially, from 55.7 to 37.5 percent. In Germany, the primary beneficiary of a shift towards French tastes is the French brands themselves which is why the decline in EU shares is much smaller than the decline in German share, from 77.2 to 69.6 percent.

This exercise illustrates the importance of the fact that firms themselves are endogenously determining the choice set of consumers in reaction to their preferences. In the long run, we would expect firms to react to such a dramatic change in tastes by adjusting

[34] Renault acquired a 43 percent stake in the Japanese firm Nissan—which does sell in the United States—in 1999, but we regard Nissan as a Japanese brand.

their product offerings and also their constellation of assembly locations.[35] While our exercises hold model offerings and assembly locations fixed, we hope they also provide some indication for how our model could be used to shed light these incentives.

6 Conclusion

The automobile industry exhibits significant home market advantage in market shares. This paper proposes and estimates a structural model to estimate the relationship of various demand-side and cost-side elements to market outcomes. The estimates clearly establish the existence of both demand and cost side frictions. On the demand side, consumers exhibit strong preference for their domestic brands relative to how these brands are viewed in the rest of the world. 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 where we see how a common measure of home market advantage is effected by removing a particular feature of the model. We find that brand preference is a major driver of the home market advantage, with an effect more than seven times larger than removing all cost side frictions. This however does not mean that other features are not important. In particular, our counterfactual analysis focuses on the medium run while other factors could play a large role in determining both how models are introduced into markets and where assembly locations are located.

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

[35] In the—admittedly fantastic—case of Francophilia taking hold in the United States, one imagines this might extend to entry on the part of Peugeot and Renault.

marketing campaigns (e.g., Volkswagen 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 jumping home preferences can be another motive for acquiring foreign firms analogous to jumping tariffs by establishing foreign production.

References

Anderson, J. E., and E. van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93(1), 170–192.

Arkolakis, C., N. Ramondo, A. Rodriguez-Clare, and S. Yeaple (2013): "Innovation and Production in the Global Economy," NBER Working Papers 18972, National Bureau of Economic Research, Inc.

Armstrong, T. (2014): "Large Market Asymptotics for Differentiated Product Demand Estimators with Economic Models of Supply," Yale University.

Auer, R. (2014): "Product Heterogeneity, Cross-Country Taste Differences, and the Growth of World Trade," Swiss National Bank.

Bernard, A. B., J. B. Jensen, S. J. Redding, and P. K. Schott (2009): "The Margins of US Trade," *American Economic Review*, 99(2), 487–93.

Berry, S. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2), 242–262.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841–890.

——— (1999): "Voluntary Export Restraints on Automobiles: Evaluating a Trade Policy," *American Economic Review*, 89(3), 400–430.

Chamberlain, G. (1982): "Multi Variate Regression Models for Panel Data," *Journal of Econometrics*, 18, 5–46.

Coşar, A. K., P. L. E. Grieco, and F. Tintelnot (forthcoming): "Borders Geography, and Oligopoly: Evidence from the Wind Turbine Industry," *Review of Economics and Statistics*.

Conlon, C. (2013): "The Empirical Likelihood MPEC Approach to Demand Estimation," Columbia University.

De Loecker, J., and F. Warzynski (2012): "Markups and Firm-Level Export Status," *American Economic Review*, 102(6), 2437–2471.

Eaton, J., and S. Kortum (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.

Feenstra, R. C. (1988): "Quality Change under Trade Restraints in Japanese Autos," *The Quarterly Journal of Economics*, 103(1), 131–46.

Giroud, X. (2013): "Proximity and Investment: Evidence from Plant-Level Data," *The Quarterly Journal of Economics*, 128(2), 861–915.

Goldberg, P. K. (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 63(4), 891–951.

Goldberg, P. K., and F. Verboven (2001a): "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, 68(4), 811–848.

——— (2001b): "The Evolution of Price Dispersion in the European Car Market," *Review of Economic Studies*, 68(4), 811–48.

——— (2004): "Cross-country price dispersion in the euro era: a case study of the European car market," *Economic Policy*, 19(40), 483–521.

Head, K., and T. Mayer (2015): "Brands in Motion: Estimating multinational frictions in the car industry," mimeo.

Krugman, P. (1980): "Scale economies, product differentiation, and the pattern of trade," *The American Economic Review*, pp. 950–959.

McCalman, P., and A. Spearot (2013): "Why trucks jump: Offshoring and product characteristics," *Journal of International Economics*, 91(1), 82 – 95.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

Nevo, A. (2000): "A Practicitoner's Guide to Estimation of Random Coefficients Logit Models of Demand," *Journal of Economics & Management Strategy*, 9(4), 513–548.

——— (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry," *Econometrica*, 69(2), 307–342.

Ramondo, N., and A. Rodriguez-Clare (2013): "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 121(2), 273 – 322.

Tintelnot, F. (2014): "Global Production with Export Platforms," mimeo Princeton University.

WTO (2013): "International Trade Statistics 2013," Discussion paper, World Trade Organization.

Appendices

Appendix A Additional Counterfactual Tables

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

Table A.1: Average area-level market share of brands across markets (in percentage)

Data	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	31.7	8.9	34.2	8.3	11.2	6.5	15.9	11.6	39.3
EU Brands	55.8	75.8	8.4	77.2	70.7	82.7	62.8	74.1	9.4
JPN Brands	8.4	11.3	48.2	10.9	13.1	8.8	16.9	11.3	43.1
Other Brands	4.1	4.0	9.2	3.7	5.1	2.0	4.4	3.0	8.1
Home Brands				55.7	9.2	52.6	18.2	30.3	39.3
No Home Preference (Homog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	31.7	8.9	34.2	11.7	11.7	9.1	17.7	13.8	21.9
EU Brands	55.8	75.8	8.4	68.0	69.2	75.1	58.3	69.0	12.8
JPN Brands	8.4	11.3	48.2	15.4	13.8	12.9	19.0	13.6	55.1
Other Brands	4.1	4.0	9.2	4.9	5.3	2.8	4.9	3.5	10.2
Home Brands				37.1	3.9	33.2	8.5	17.7	21.9
No Home Preference (Heterog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	31.7	8.9	34.2	11.0	11.9	9.7	17.9	14.8	27.7
EU Brands	55.8	75.8	8.4	69.7	68.6	73.5	58.0	66.9	11.7
JPN Brands	8.4	11.3	48.2	14.6	14.0	13.8	19.2	14.5	51.1
Other Brands	4.1	4.0	9.2	4.7	5.4	3.0	4.9	3.8	9.5
Home Brands				40.6	2.0	29.1	7.8	12.4	27.7
All France Taste	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	34.3	7.9	8.6	9.8	6.9	6.5	14.5	11.9	10.6
EU Brands	60.0	68.0	7.1	69.9	68.1	82.7	56.0	69.1	1.5
JPN Brands	4.1	19.1	72.9	13.6	21.0	8.8	19.4	13.8	74.3
Other Brands	1.5	4.9	11.4	6.7	4.0	2.0	10.1	5.2	13.6
Home Brands				37.5	8.9	52.6	13.1	32.7	10.6
No Shipping Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	31.4	9.0	32.6	8.7	11.8	6.8	16.6	12.2	37.3
EU Brands	55.2	74.0	9.4	74.8	68.4	81.4	60.7	72.5	10.5
JPN Brands	8.3	12.6	48.1	12.4	14.3	9.7	17.9	12.1	43.8
Other Brands	5.1	4.4	9.9	4.0	5.5	2.2	4.7	3.1	8.4
Home Brands				52.7	7.5	51.7	16.2	28.8	37.3
No Remote Production Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	31.0	10.6	33.8	9.9	12.4	7.6	17.4	13.3	37.7
EU Brands	55.6	73.2	8.4	74.6	68.6	80.9	60.7	71.6	9.7
JPN Brands	10.0	11.9	49.8	11.5	13.8	9.4	17.3	11.9	44.9
Other Brands	3.4	4.3	8.0	4.0	5.2	2.1	4.6	3.2	7.7
Home Brands				53.5	8.8	52.1	19.6	27.8	37.7
No Tariff	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	24.6	9.6	29.0	9.3	12.6	6.8	17.8	12.6	37.2
EU Brands	41.3	69.8	10.7	71.4	64.0	79.8	54.9	69.6	11.1
JPN Brands	9.1	15.0	49.4	14.5	16.8	10.8	21.6	14.1	43.3
Other Brands	25.1	5.6	10.9	4.8	6.6	2.6	5.7	3.8	8.4
Home Brands				51.2	6.0	51.5	13.2	27.7	37.2
No Trade Friction	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	25.0	10.3	28.2	10.3	13.5	7.4	19.3	14.4	34.5
EU Brands	40.5	70.0	11.5	71.0	63.8	79.4	54.1	67.7	12.3
JPN Brands	10.7	14.3	50.0	14.2	16.6	10.7	21.2	14.1	44.9
Other Brands	23.8	5.4	10.3	4.6	6.1	2.5	5.4	3.8	8.3
Home Brands				51.0	4.2	52.5	12.1	23.5	34.5

Table A.2: Weighted average area-level cost (thousand dollars)

Data	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	19.51	26.12	28.51	25.77	24.32	23.62	23.77	21.66	25.23
EU Brands	18.84	29.87	36.82	32.67	28.29	25.34	29.80	24.42	35.47
JPN Brands	37.31	24.97	24.94	25.22	28.90	24.87	24.32	22.71	23.28
Other Brands	39.36	24.56	20.22	20.94	25.39	23.80	18.61	20.42	19.03
Home Brands				36.27	22.92	23.42	29.78	18.75	25.23
No Home Preference (Homog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	19.51	26.12	28.51	25.77	24.32	23.62	23.77	21.66	25.23
EU Brands	18.84	29.87	36.82	32.67	28.29	25.34	29.80	24.42	35.47
JPN Brands	37.31	24.97	24.94	25.22	28.90	24.87	24.32	22.71	23.28
Other Brands	39.36	24.56	20.22	20.94	25.39	23.80	18.61	20.42	19.03
Home Brands				36.27	22.92	23.42	29.78	18.75	25.23
No Home Preference (Heterog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	19.51	26.12	28.51	25.77	24.32	23.62	23.77	21.66	25.23
EU Brands	18.84	29.87	36.82	32.67	28.29	25.34	29.80	24.42	35.47
JPN Brands	37.31	24.97	24.94	25.22	28.90	24.87	24.32	22.71	23.28
Other Brands	39.36	24.56	20.22	20.94	25.39	23.80	18.61	20.42	19.03
Home Brands				36.27	22.92	23.42	29.78	18.75	25.23
All France Taste	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	19.51	26.12	28.51	25.77	24.32	23.62	23.77	21.66	25.23
EU Brands	18.84	29.87	36.82	32.67	28.29	25.34	29.80	24.42	35.47
JPN Brands	37.31	24.97	24.94	25.22	28.90	24.87	24.32	22.71	23.28
Other Brands	39.36	24.56	20.22	20.94	25.39	23.80	18.61	20.42	19.03
Home Brands				36.27	22.92	23.42	29.78	18.75	25.23
No Shipping Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	19.16	25.47	27.85	25.05	23.42	22.91	22.92	20.87	24.71
EU Brands	18.51	29.19	35.37	32.19	27.52	24.75	29.04	23.76	34.13
JPN Brands	36.51	24.01	24.23	24.26	27.74	23.91	23.41	21.83	22.61
Other Brands	37.75	23.63	19.41	20.20	24.38	22.90	17.89	19.65	18.45
Home Brands				35.83	22.71	22.88	29.35	18.36	24.71
No Remote Production Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	18.08	24.35	26.86	24.03	22.74	21.99	22.23	20.20	23.84
EU Brands	17.43	28.65	34.78	31.56	27.03	24.21	28.53	23.39	33.45
JPN Brands	34.29	23.62	23.33	23.91	27.33	23.42	22.98	21.46	21.80
Other Brands	37.48	23.13	19.28	19.67	24.06	22.43	17.53	19.21	17.98
Home Brands				35.11	21.84	22.30	28.19	18.05	23.84
No Tariff	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	19.21	25.07	28.49	24.62	23.01	22.75	22.50	20.63	25.21
EU Brands	18.58	29.07	35.20	31.97	27.45	24.60	29.09	23.74	34.65
JPN Brands	33.90	23.37	24.46	23.58	26.94	23.36	22.67	21.25	23.10
Other Brands	30.88	22.81	19.40	19.54	23.57	22.25	17.30	19.08	18.82
Home Brands				35.53	22.92	22.68	29.59	18.38	25.21
No Trade Friction	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	17.47	22.66	26.22	22.21	20.66	20.44	20.20	18.47	23.33
EU Brands	16.89	26.64	32.05	29.44	25.14	22.53	26.78	21.88	31.54
JPN Brands	30.50	21.27	22.26	21.52	24.47	21.17	20.63	19.32	21.05
Other Brands	28.24	20.75	17.76	17.82	21.56	20.24	15.74	17.32	17.19
Home Brands				32.73	21.63	20.74	27.60	17.28	23.33

Table A.3: Weighted average area-level markup (thousand dollars)

Data	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.98	2.97	4.05	3.28	3.00	3.34	3.34	3.03	3.99
EU Brands	1.97	4.00	5.56	5.33	4.00	4.56	4.50	3.96	6.32
JPN Brands	3.93	2.78	3.40	3.07	3.65	3.80	3.09	3.05	3.40
Other Brands	4.56	2.64	2.51	2.66	3.04	3.32	2.37	2.70	2.35
Home Brands				6.08	3.32	4.44	4.39	3.41	3.99
No Home Preference (Homog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.98	2.97	4.05	3.53	3.01	3.44	3.42	3.04	3.69
EU Brands	1.97	4.00	5.56	5.26	3.95	4.27	4.49	3.73	6.39
JPN Brands	3.93	2.78	3.40	3.25	3.65	3.74	3.13	3.05	3.58
Other Brands	4.56	2.64	2.51	2.80	3.04	3.34	2.40	2.70	2.47
Home Brands				5.95	3.14	3.94	4.08	2.85	3.69
No Home Preference (Heterog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.98	2.97	4.05	3.49	3.02	3.46	3.42	3.05	3.79
EU Brands	1.97	4.00	5.56	5.27	3.93	4.22	4.49	3.66	6.37
JPN Brands	3.93	2.78	3.40	3.22	3.65	3.73	3.13	3.04	3.52
Other Brands	4.56	2.64	2.51	2.77	3.03	3.35	2.40	2.70	2.43
Home Brands				5.98	3.08	3.84	4.06	2.66	3.79
All France Taste	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.93	3.30	2.90	3.19	3.29	3.34	3.19	2.70	2.54
EU Brands	1.85	4.59	4.54	5.56	4.41	4.56	4.61	3.65	5.15
JPN Brands	4.48	3.12	5.82	3.02	4.34	3.80	3.04	2.76	7.70
Other Brands	4.81	2.89	2.01	2.40	3.46	3.32	2.19	2.31	1.76
Home Brands				6.46	3.63	4.44	4.36	2.84	2.54
No Shipping Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.94	2.94	3.97	3.23	2.92	3.27	3.26	2.95	3.92
EU Brands	1.93	3.95	5.38	5.25	3.89	4.48	4.42	3.85	6.12
JPN Brands	3.81	2.71	3.32	2.99	3.52	3.69	3.00	2.95	3.32
Other Brands	4.41	2.57	2.46	2.60	2.93	3.21	2.29	2.62	2.30
Home Brands				6.01	3.25	4.35	4.33	3.29	3.92
No Remote Production Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.86	2.86	3.87	3.17	2.88	3.18	3.22	2.91	3.83
EU Brands	1.85	3.92	5.37	5.16	3.87	4.43	4.41	3.81	6.13
JPN Brands	3.75	2.70	3.26	2.96	3.52	3.66	2.98	2.93	3.25
Other Brands	4.39	2.54	2.38	2.56	2.94	3.18	2.26	2.58	2.24
Home Brands				5.90	3.12	4.30	4.29	3.22	3.83
No Tariff	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.73	2.91	4.04	3.21	2.90	3.27	3.25	2.92	3.97
EU Brands	1.71	3.93	5.33	5.17	3.87	4.48	4.42	3.83	6.16
JPN Brands	3.17	2.66	3.35	2.93	3.44	3.64	2.93	2.89	3.37
Other Brands	4.05	2.50	2.50	2.56	2.86	3.14	2.25	2.57	2.34
Home Brands				5.91	3.23	4.35	4.33	3.26	3.97
No Trade Friction	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	1.59	2.71	3.79	2.96	2.66	3.02	3.01	2.69	3.74
EU Brands	1.57	3.73	5.00	4.84	3.63	4.24	4.20	3.56	5.79
JPN Brands	2.92	2.48	3.12	2.71	3.21	3.40	2.72	2.68	3.14
Other Brands	3.65	2.32	2.32	2.36	2.66	2.92	2.07	2.37	2.17
Home Brands				5.53	3.05	4.11	4.15	2.96	3.74

Table A.4: Weighted average area-level markup (percent)

Data	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	10.1	11.2	13.7	13.1	12.3	14.0	14.0	13.7	15.0
EU Brands	10.5	12.8	14.0	16.1	13.9	17.9	14.4	16.1	16.6
JPN Brands	10.6	11.0	13.2	12.4	12.5	15.1	12.5	13.1	14.0
Other Brands	11.4	10.6	12.4	13.2	12.0	13.9	13.1	13.1	12.3
Home Brands				16.5	14.5	18.9	14.0	18.4	15.0
No Home Preference (Homog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	10.1	11.2	13.7	14.0	12.4	14.5	14.3	13.9	13.9
EU Brands	10.5	12.8	14.0	15.9	13.7	16.7	14.4	14.9	16.8
JPN Brands	10.6	11.0	13.2	13.1	12.5	14.9	12.7	13.2	14.8
Other Brands	11.4	10.6	12.4	13.8	12.0	14.1	13.3	13.1	13.0
Home Brands				16.2	13.7	16.9	13.0	15.3	13.9
No Home Preference (Heterog)	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	10.1	11.2	13.7	13.8	12.4	14.6	14.3	14.0	14.3
EU Brands	10.5	12.8	14.0	15.9	13.6	16.5	14.4	14.5	16.7
JPN Brands	10.6	11.0	13.2	13.0	12.5	14.9	12.7	13.2	14.5
Other Brands	11.4	10.6	12.4	13.7	12.0	14.1	13.3	13.1	12.8
Home Brands				16.2	13.4	16.5	13.0	14.2	14.3
All France Taste	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	9.4	12.0	9.6	12.0	13.4	14.0	13.0	11.5	9.5
EU Brands	9.5	14.2	10.9	16.0	15.3	17.9	14.3	13.9	12.8
JPN Brands	12.2	11.8	21.2	11.6	14.7	15.1	11.8	11.3	30.0
Other Brands	12.1	11.1	9.9	11.1	13.6	13.9	11.6	10.2	9.2
Home Brands				17.1	15.8	18.9	13.6	14.8	9.5
No Shipping Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	10.0	11.4	13.8	13.3	12.5	14.1	14.2	13.8	15.0
EU Brands	10.5	12.9	14.2	16.1	13.9	17.9	14.5	16.1	16.6
JPN Brands	10.5	11.2	13.3	12.6	12.6	15.2	12.6	13.2	14.1
Other Brands	11.4	10.8	12.7	13.4	12.1	14.0	13.2	13.2	12.5
Home Brands				16.5	14.3	19.0	14.0	18.1	15.0
No Remote Production Cost	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	10.1	11.6	13.9	13.6	12.7	14.3	14.4	14.1	15.2
EU Brands	10.7	13.0	14.3	16.1	14.0	18.2	14.7	16.1	17.0
JPN Brands	11.0	11.3	13.6	12.7	12.7	15.4	12.7	13.4	14.3
Other Brands	11.5	10.9	12.3	13.5	12.2	14.1	13.2	13.3	12.4
Home Brands				16.5	14.3	19.2	14.5	18.0	15.2
No Tariff	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	8.9	11.4	13.7	13.4	12.6	14.2	14.4	13.9	15.0
EU Brands	9.2	12.9	14.1	15.9	13.8	18.1	14.4	16.0	16.5
JPN Brands	9.3	11.3	13.3	12.7	12.7	15.4	12.7	13.4	14.0
Other Brands	12.8	10.9	12.9	13.7	12.2	14.1	13.5	13.3	12.4
Home Brands				16.3	14.1	19.1	13.8	17.9	15.0
No Trade Friction	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
US Brands	9.0	11.8	13.9	13.8	12.9	14.6	14.8	14.4	15.2
EU Brands	9.3	13.3	14.5	16.1	14.1	18.7	14.8	15.9	17.0
JPN Brands	9.5	11.6	13.6	12.9	13.0	15.9	12.9	13.6	14.3
Other Brands	12.6	11.1	13.1	13.9	12.4	14.4	13.6	13.6	12.6
Home Brands				16.6	14.0	19.8	14.2	17.3	15.2