Trade Costs or Taste Differences? Evidence from the Global Car Industry

K. Cosar, P. Grieco, S. Li, F. Tintelnot

August 2014

Preliminary: please do not cite

1 Introduction

The global automobile industry is characterized by a small number of multinational firms. In 2012 five firms—Toyota, General Motors, Volkswagen, Hyundai, and Ford—accounted for 49 percent of units produced, while the fifteen largest firms accounted for over 80 percent of world production. These firms use a combination of international trade and foreign direct investment (FDI) to access markets. The industry plays a significant role in overall trade flows: automobile products account for over 10 percent of world trade in manufactured goods. Thus, incorporating the oligopolistic nature of the industry is important to understanding a sizable fraction of world manufacturing trade. This paper aims to consistently estimate an oligopolistic model of automobile supply and demand incorporating data from three different continents in order to analyze the roles of trade costs and demand heterogeneity in firms' marketing, FDI, and export decisions.

Traditionally, models of global 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 international trade. 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 autos. They have also been limited by the availability of only revenue data without credible price and quantiy information. As a result, they may result in biased estimates of trade costs and an under appreciation 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 will be able to consistently estimate demand and supply driven mechanisms for understanding trade flows. 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, Arkolakis e.a. (2013), and Tintelnot (2014)), which were limited by the availability of only revenue data on multinationals' foreign affiliate sales for the aggregate manufacturing sector.

This paper builds off of our earlier work on the analysis of trade costs in an oligopolistic industry with detailed spatial production and sales data. In Cosar, Grieco, and Tintelnot (forthcoming), we show how to estimate distance and border related trade costs in an oligopolistic setting where demand preferences do not vary systematically across countries. We apply this model to analyze trade flows of wind turbines between Germany and Denmark—a business to business industry where the payoffs of a turbine design are tightly tied to its ability to produce electricity, and there is little scope for national preference for home produced goods. While similarly spatial production and sales data is available for the automobile industry, it is much more complicated on other dimensions. For example, preferences for automobile characteristics may change systematically across countries, home bias is considered an important feature of industry preferences, and a single firm sells several different products and operates a variety of assembly locations in multiple countries. Accounting for these features of the industry is important for providing consistent estimates of trade costs. In addition, they open the door to a more sophisticated analysis of trade flows in an oligopolistic environment.

A key advantage of the auto industry for our analysis is the availability of data from a variety of markets. We have compiled a dataset of car sales by model for nine countries across three continents, and are hopeful that we will be able to add additional countries in the future. In contrast to our earlier work, we are able to use data on prices and sales (as opposed to sales only) as well as several characteristics such as horsepower, size, weight and fuel efficiency. We have linked this sales data to worldwide data on the assembly plant locations of each model which enables us to uncover the role of trade costs in the marginal cost of supplying cars to the market. Through the use of a structural model, the availability of price data is useful in separating "home bias"—demand driven preferences for local products—from trade costs and policies (e.g. tariffs and environmental regulations that change the cost of driving a high horsepower car differentially across countries). Moreover, we are able to identify country level preferences for characteristics (such as fuel efficiency) from traditional home bias. 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. As pointed out by Auer (2014) in a dynamic extension of the home market effect model by Krugman (1980), if endogenous specialization is a strong feature of the economy, trade will respond sluggishly after a trade liberalization.

A well-known issue in demand estimation such as this is the endogeneity of prices; two features of our data enable us to address these concerns. First, because automakers frequently release the same model in multiple countries, we are able to allow for model fixed effects to account for unobserved model quality. Second, the assembly location data gives us a cost side instrument—distance to market—to control for endogeneity of prices. Our preliminary work suggests that this approach is successful in that the estimated price-elasticities imply price-cost margins which are consistent with other industry studies.

Going forward, our results can be used to understand the impact of trade costs and foreign direct investment in the auto industry. We will compare our estimates of trade costs to those from simpler demand systems to gauge the importance of cross-country demand differences and home bias in measuring trade frictions. Moreover, we can use our model to conduct counterfactual analysis to understand how the presence of a foreign assembly plant affects competition, trade flows, and consumer welfare across the region. Ultimately, we hope that our work will lead to a deeper understanding of the endogenous decisions of firms of where to locate assembly plants and which automobile models should be marketed in which countries.

The paper will shed light on several interesting issues about international trade and the auto industry. Within the auto industry, firms use both FDI and exports to supply markets while governments enact policies which affect sourcing decisions (such as preferential taxation or loan guarantees). However, there is relatively little research on why a firm might serve some markets via FDI, while simply exporting to others. Our study aims to fill this gap by examining the tradeoffs to these alternative methods of serving foreign markets. Moreover, the role of home bias is a common feature in autos which makes it difficult to assess the size of trade costs, which are often estimated to be large even after the dramatic liberalization of manufacturing trade. An accurate measure of the role of home bias in preferences versus trade costs is important to understanding how trade flows react to policy changes. While we study the auto industry specifically, our results should provide insights on the relative tradeoffs in other manufacturing industries as well. The key features of the auto industry we study are differentiated competition and sourcing of multiple markets via large investments in assembly plants. Two features common to many manufacturing industries, such as machinery and electronics.

2 Relation to the Literature

This paper adds to a large body of empirical work on the car industry. The distinctive feature of our work is that it first takes a global view on the automobile industry by studying outcomes in nine large automobile markets on three continents; second it incorporates a richer supply side by solving the firm's problem which market to serve from which assembly location. Both data and conceptual improvements result in a richer cost side estimation, which allows us to disentangle the effects of preferences, home bias, and trade and foreign production costs on market outcomes.

On the methodological side we follow the random coefficients demand model developed by Berry, Levinson, and Pakes (1995), which study the U.S. automobile market, and Nevo (2001), who introduces product-fixed effects in this demand estimation framework. Due to the knowledge of assembly locations, we can include a cost shifter (distance of assembly location to market) as instrument for price in the demand estimation. We also allow for the fact that the source location was endogenously chosen from the set of all available plants.

A number of papers have used the detailed data provided by the car industry to study questions in international trade. Among these papers, Goldberg (1995) analyzes the effects of U.S. import quotas on cars produced in Japan on market shares and consumer prices, while Goldberg and Verboven (2001) 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 and McCalman and Spearot (2013) study firms' offshoring strategies using data on North American light truck production locations.

3 Data and Descriptive Evidence

On the demand side, we purchased data from Polk, a well-known industry data vendor, informing us about annual registration of new passenger vehicles in nine countries between 2007-2011. The markets are Belgium, Brazil, Canada, France, Germany, Italy, Spain, United Kingdom and United States. Manufacturer's suggested retail price, quantity sold and key characteristics are observed at the model-trim level in each market-year. An example for an observation is the 2008 US sales and price of Toyota Corolla with a certain configuration of horsepower, weight, dimensions and fuel consumption.

On the supply side, we started building a dataset of assembly locations. For each model-year pair in our vehicle sales data, we would like to observe the countries in which the model is assembled that given year. Through an ongoing data collection effort, we are complementing and expanding a dataset purchased from WardsAuto. So far, our supply-side data contains information about *potential* sources of supply for each model-market pair. Unfortunately, we do not observe exact sourcing decisions of producers, i.e. we know the countries in which Toyota Corolla was being assembled but we do not know which of these plants a given market was supplied from. Currently, we are investigating the feasibility of obtaining data on sourcing and plan to allocate some of our budget to that purpose.

We complement this demand and supply data with market specific measures that matter for estimation, such as tariff rates, sales taxes and measures of within-country income inequality.

We now describe the key features of the data used in our preliminary analysis. Starting with demand, the left panel of Table 1 presents sales weighted average characteristics in the nine markets across years. Evidently, the "typical car" differs systematically across countries. The average horsepower in the US is more than twice its level in Brazil. Average price in Germany is 30 percent higher than in the US. On the supply side, the middle panel of table 1 shows that markets also differ by the average number of models offered by competing firms. In the intensive margin, the right panel of table 1 shows the share of sales in each market (rows) by manufacturing groups headquartered in various countries (columns). This reveals a disproportionately large share by national producers: German and French producers command a 50 and 59 percent domestic market share, respectively. This is despite the absence of policy-related trade barriers within the EU and plausibly similar tastes, suggesting bias for national brands.

To further document the importance of incorporating the supply side to an estimation of global competition in this industry, Figure 1 plots the fraction of models and their total market share in our sales data by the number of countries in which these models are assembled. Slightly more than half of the 596 unique models in our data are assembled in a single country, typically where the firm is also headquartered, yet their market share is only 30 percent. A total of 15 percent of models have assembly plants in 6 or more countries with a total market share around 33 percent. This variation shows the importance of incorporating multinational production to our analysis. Taking the choice of assembly locations as exogenous, we exploit the variation in distances and tariffs between assembly countries and the nine markets under study as cost shifters, identifying the effect of trade costs in market shares. Together with shipping distances and tariffs, variation in productivity across assembly countries help us to instrument for prices in our demand estimation.

Finally, we finish with a simple hedonic pricing analysis reflecting our key points. In Table 2, we present the results from regressing model-level log prices on a set of market and model dummies together with the log distance between the market and the nearest assembly country (mindist), and characteristics of interest. The results suggest that distance to supply sources is a significant determinant of costs, and thus prices, lending credibility to our instrumentation strategy. Unsurprisingly, horsepower and size are positively correlated with prices, suggesting that both the mean and variance of income in a market plays a role in shaping aggregate demand for car characteristics. Finally, the home dummy is large and significant, suggesting that there is a willingness to pay for national producers. We now describe a model of oligopolistic competition to estimate demand and supply-related factors in this industry.

Table 1: Average Characteristics, Models and Market Shares by Country

	Characteristic				Models	Market share of firms from				
Market	price	hppwt	size	mpgc	per firm	DEU	FRA	ITA	USA	Other
BEL	32.49	58.92	7.60	34.53	16.4	31	30	4	18	17
BRA	23.71	63.12	6.80	30.11	6.1	23	10	24	31	12
CAN	27.26	92.64	8.30	22.27	12.7	8	7	11	22	52
DEU	35.67	67.93	7.60	29.38	17	**50**	12	5	18	15
ESP	31.57	61.36	7.60	32.39	14.5	31	33	2	19	15
FRA	29.54	57.88	7.30	35.53	14.2	17	**59**	4	11	9
GBR	31.41	66.39	7.50	30.36	14.7	28	16	2	30	24
ITA	27.49	58.62	7.00	33.35	13.5	18	18	**32**	19	13
USA	27.21	98.57	8.70	21.00	16.2	8	8	8	**29**	47

Table 2: Hedonic price regression

	log price
log mindist	0.0225***
log HP per Weight	0.263***
log Size	0.586***
log MPCCITY	0.00799
Home dummy	0.0139***
Country dummies	YES
Model dummies	YES
N	7929

Figure 1: Models and Market Shares by Number of Assembly Countries

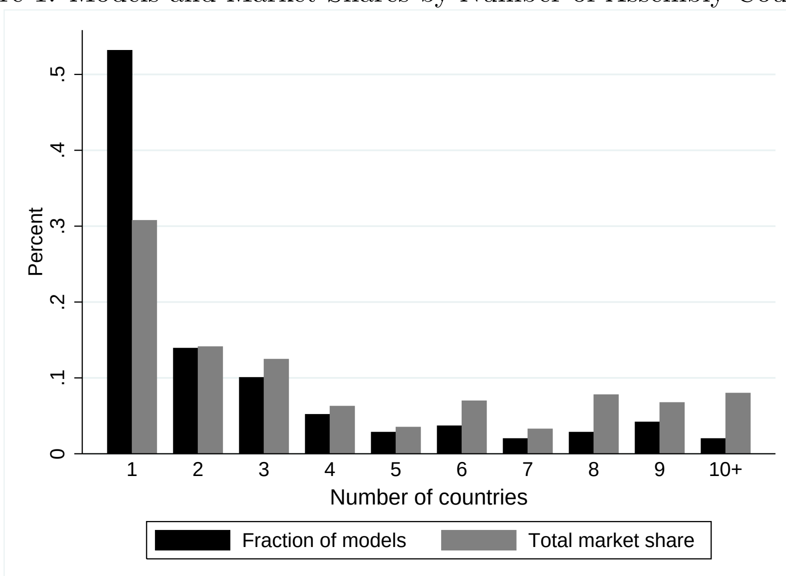

4 Model

The goal of our model is to measure the relative importance of demand and supply factors in trade flows. To do this, we propose an equilibrium model of price setting according to a random coefficients discrete choice framework (Berry et al 1995). The model is estimated in two parts. We will first consider the demand side, which will recover demand parameters such as home bias. We then will 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, and trade costs in particular.

4.1 Demand

Consumer i in country m gets the following utility from purchasing model j,

$$u_{jmi} = X_{jm}\beta_i + \eta \mathbf{1}\left[j \text{ is foreign model}\right] - \alpha_i p_{jm} + \psi_m + \xi_{jm} + \epsilon_{jmi} \tag{1}$$

where X_{jm} represent the model characteristics—horsepower, fuel efficiency, size, and weight. Note that these characteristics typically vary slightly across countries, which are valued by consumer i with preference parameters β_i, η represents consumers disutility of consuming foreign models, p_{jc} is the price of model j while α_i represents the consumers price sensitivity, and finally ψ_j is the unobserved utility of model j which is constant across countries, and ξ_{jm} is a country-model demand shock which accounts for the match quality of country m and model j. Finally, ϵ_{jmi} is a type-I extreme value draw which represents the consumers unobserved utility for model j.

We allow for heterogeneity in consumer's taste both within and across countries. For the tastes for characteristics, we assume tastes for individual i according to,

$$\beta_i = \bar{\beta}_m + \Sigma\zeta_i$$

Where $\bar{\beta}_m$ is the mean taste for the characteristic for country m, Σ is a diagonal matrix which determines the variance of the distribution of consumers tastes within the country,[1] and ζ_i is a draw from the standard normal distribution. Consumers also differ in their price-sensitivity, α_i, which we assume is log-normally distributed,

$$\log \alpha_i = \bar{\alpha} + \Pi D_{im} + \sigma^{\alpha}\zeta_i^{\alpha}.$$

Here, $\bar{\alpha}$ represents the (log) mean price sensitivity, D_{ic} is a vector of individual characteristics which are drawn from the country-level population distribution while Π is a vector of parameters that govern the relationship between these characteristics and price sensitivity. We have included only income in D_{ic} in our preliminary results, however we have also experimented with income squared.[2]

Following the literature, we assume that firms decide on which models to sell in which countries prior to the observation of the country-model error term, ξ_{jm}. Upon observing this shock, firms simultaneously set prices according to equilibrium conditions, which induces endogeneity between prices and the error term creating the need to instrument for prices. Let S_m represent the set of car's available in country c, plus an outside option of no purchase which we normalize as $u_{0mi} = \epsilon_{0mi}$, where ϵ_{0mi} is distributed type-I extreme value. Then the consumer's purchase problem is to simply choose the option from the choice set that maximizes her utility, $\max_{j \in S_m} u_{jmi}$. Integrating over all consumers, we arrive at a formula for each model's market share, which we observe in the data. To do so, it is convenient to re-write utility as the sum of δ_{jm}, which is a function of the linear parameters and does not vary across consumers within a market, and μ_{ici} which is a function of the non-linear parameters and contains all within market consumer heterogeneity.

$$u_{jmi} = \delta_{jm} + \mu_{jmi} + \epsilon_{jmi}$$

Where

$$\delta_{jm} = X_{jm}\bar{\beta}_m + \eta \mathbf{1}\left[j \text{ is foreign model}\right] - \bar{\alpha} p_{jm} + \psi_j + \xi_{jm}$$

and

$$\mu_{jmi} = X_{jm}\Sigma\zeta_i - \exp(\bar{\alpha} + \Pi D_{im} + \sigma^{\alpha}\zeta_i^{\alpha})p_{jm}$$

Then, we can integrate over the population characteristics, $(D_{I'm}, \zeta_i)$ to arrive at predicted market shares,

$$s_{jm}(\delta_m, \Pi, \Sigma) = \int\int \frac{\exp(\delta_{jm} + \mu_{jmi})}{1 + \sum_{k \in C(m)} \exp(\delta_{km} + \mu_{kmi})} dF(D_i) dF(\zeta_i) \qquad (2)$$

[1] This variance is assumed to be constant across countries for reasons of tractability.

[2] We are also working on including income effects in the distribution of β_i.

Berry et al (1995) show that (2) implies a one-to-one mapping between the country-level mean utility and market shares for a given set of non-linear parameters which are in our case, $(\Sigma, \bar{\alpha}, \Pi, \sigma^{\alpha})$. Thus we can numerically invert this relationship to find the implied mean utilities of each model, and the error term ξ_{jm}. Then, we can estimate the model according to the moment conditions $E[\xi_{jm}|Z_{jm}] = 0$ where Z_{jm} represents a vector of instruments which is correlated with price but uncorrelated with the model-country demand shock. In our setting, we use the model characteristics, functions of the characteristics of other models, and distance to nearest assembly plant producing model j as instruments. The first two are commonly used in empirical studies of demand, while the last one is likely to be correlated with costs due to the role of distance in trade costs.

4.2 Recovering Marginal Costs

Once the demand estimates are recovered we adapt Nevo (2001) to recover marginal costs. First, recall that in each year, firms choose prices taking other model characteristics and the set of models in the market as given. Therefore, we can write market shares and their derivatives as a function of the price vector p. These functions are known because we have estimated the demand parameters. In each country-year (subscript suppressed) each firm chooses prices for its models to maximize

$$\pi_f = \sum_{j \in J(f)} (p_j - c_j) M s_j(p) - C_f$$

where $J(f)$ is the set of models sold by f in this this market, M is market size, and C_j is a fixed cost to operate in this market which is irrelevant to what we are doing now, but here for future use.

The first order condition for p_j implies,

$$s_j(p) + \sum_{r \in J(f)} (p_r - c_r) \frac{\partial s_r(p)}{\partial p_j} = 0 \tag{3}$$

With the demand parameters estimated, and prices observed, all the terms in (3) 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,

$$\Omega_{jr} = -\frac{\partial s_r(p)}{\partial p_j} \cdot \mathbf{1}\left[j, r \text{ jointly owned}\right],$$

Then we can write (3) in vector notation, $s(p) - \Omega(p - c) = 0$, and can easily solve this for the vector of marginal costs,

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

4.3 Estimating Trade Costs

With the recovered costs in hand, the next step is to estimate the impact of trade costs on total marginal cost of supplying cars. There are several features of trade costs which we are interested in. First, and most straightforwardly, we want to understand how the distance from assembly location to market affects costs. Second, we want to see whether plants that are distant from company headquarters tend to be less productive than plants that are nearby. Finally, we want to explore the possibility that there may be gains from variety and specialization from producing the same model in multiple locations, or if the only reason to have multiple factories producing the same model is for the sake of market access.

For each model, we know the set of assembly locations where the model is produced, but we do not have precise data on the flows of models from assembly locations to each of our markets. Therefore, to estimate trade costs while allowing for the fact that there may be gains from variety or specialization leading to the same model being supplied from multiple locations, we adopt the following multiplicative form for the marginal cost of supplying a model j automobile to individual i in market m from assembly location ℓ,

$$c_{jm\ell i} = \Upsilon_m (1 + t_{m\ell}) c^p_{jm\ell} e^{-\nu_{jm\ell i}} \tag{4}$$

Where Υ_m represents local marketing costs (for example, land and labor costs for dealerships), $c^p_{jm\ell}$ is the average port cost of the model when it is sourced from location ℓ and $t_{m\ell}$ is the ad valorem tariff rate. Finally $\nu_{jm\ell i}$ is a car-specific idiosyncratic error term related to the individual characteristics of consumer i's order (color, trim packages, etc.) which may make the car easier to produce at some locations rather than others if there is a high degree of specialization across plants. On the other hand, if plants more or less duplicate their productive capabilities, the variance of this idiosyncratic term will be small. We assume $\nu_{jm\ell i}$ is drawn from a type-I extreme value distribution with scale parameter σ_ν. We further parameterize the port cost as,

$$\log c^p_{jm\ell} = X_{jm}\delta^X + \kappa_j + d_{m\ell}\tau + d_{o\ell}\gamma + \phi_\ell + \omega_{jm} \tag{5}$$

Here, X_{jm} represent the model's characteristics as it is marketed in country m, κ_j is a model level fixed effect,[3] $d_{m\ell}$ is the log distance between the market and the assembly location representing trade costs, $d_{o\ell}$ is the log distance between company headquarters and the assembly location which we take as a proxy for foreign production costs that grow larger as plants become more difficult to monitor from headquarters, ϕ_ℓ is a fixed effect reflecting the productivity differences between assembly locations, ω_{jm} is an error term reflecting the cost of marketing model j in country m.

We assume that firms' source each car from its lowest cost location. Integrating out the idiosyncratic cost term, the expected cost of a model given the set of available assembly locations $L(j)$ is,

$$\log c_{jm} = X_{jm}\delta^X + \upsilon_m + \kappa_j - \sigma_\nu \log\left(\sum_{k\in L(j)} \exp\left(\frac{-d_{mk}\tau - \log(1+t_{m\ell}) - d_{ok}\gamma - \phi_k}{\sigma_\nu}\right)\right) + \omega_{jm} \tag{6}$$

Where $\upsilon_m = \log \Upsilon_m$ is a market dummy which absorbs differences in marketing costs across countries. Since c_{jm} are recovered from the demand side above, we are able to estimate (6) using non-linear least squares.

This equation makes clear how σ_ν, which embodies gains from variety in assembly locations, is identified from data on costs. Suppose there are two identical models being sold in the country A: one is produced exclusively in country A, while the other is produced in both A and B. Trade costs are increasing in distance and there are no productivity differences between the two countries. If there are no gains from variety, than the average cost of supplying both cars to country A be equal, since trade costs are avoided when sourcing from A, it will never be profitable to ship a model from B to A. However, if there are gains from variety, then the average cost of the car produced in two locations will be lower, since the idiosyncratic cost component will sometimes mean it is profitable to import the car from B even though it faces higher trade costs. Notice that as $\sigma_\nu \to 0$, all cars from a given market are sourced from the same assembly plant and (6) simplifies to,

$$\log c_{jm} = X_{jm}\delta^X + \upsilon_m + \kappa_j + \min_{k\in L(j)}\{d_{mk}\tau + \log(1+t_{m\ell}) + d_{ok}\gamma + \phi_k\} + \omega_{jc}.$$

[3]In practice, model characteristics do vary across markets, so the cost of characteristics are separately identified.

This limit represents the case where there are no gains from variety and the sole reason to build multiple assembly plants is market access. In practice, we have found that σ_ν is in fact very low, and so we have also considered variants of this model which is a numerically tractable approximation when σ_ν is small.

5 Preliminary Results

Our baseline estimates of the demand system are presented in Table 3. This specification includes model and year dummies and allows for random coefficients on horsepower per 1000 kilogram (a proxy for power), fuel efficiency measured in miles per gallon, and size (length x width). As discussed above, price sensitivity is assumed to vary by income, and the distribution of income is lognormal with mean and variance parameters based on each countries household income distribution.

The first thing to note is that the taste for characteristics is very heterogenous between consumers within a country. All of the random coefficients on characteristics are statistically significant and they tend to be larger than the country level means. This should be interpreted to mean that while some consumers prefer powerful cars, others do not, and are more willing to tradeoff some characteristics for others. This heterogeneity results in a horizontally segmented market where models with similar characteristics compete primarily with each other, while there is less competition across automobile classes.

A second finding from Table 3 is that cross-country variation in mean tastes for characteristics is relatively small. With the exception of Brazilians, who on average seem to prefer larger more powerful cars and are less interested in fuel economy, differences between the other countries are typically not statistically or economically significant. Many national stereotypes, such as Americans' preference for large cars, fuel inefficient cars, are thus not born out by the model, despite the fact that these characteristics are over-represented in US car sales, as discussed above in our data section. It may be that these differences are driven more by supply-side factors or differences in regulation rather than cross-country heterogeneity in preferences. This finding appears to validate the common assumption in trade work that preferences for product characteristics vary little across countries.

Finally, a key estimate from our demand specification is η, the home bias dummy, which captures the advantage a domestic brand has over foreign competitors in its home market.[4] This coefficient is substantial and statistically significant. To directly interpret this coefficient, "willingness to pay" for a domestic brand of the median consumer in each country. Not surprisingly, this number is smallest in Brazil, where median income is lowest, but is still a substantial \$218. For other countries in the sample, the quantity ranges from \$700 in Spain to \$1081 in Belgium, or roughly 2.5-3.5 percent of the total price of a car.

We use the demand estimates to calculate the implied markups for each model in each country. Table 4 shows median markups for a sample of firms (including Toyota, GM, VW, Ford, PSA and Fiat). Not surprisingly, markups are notably smaller in Brazil, where consumers have lower income and therefore higher price sensitivity. For most producers, the markup in Brazil is around 10 percent, although Toyota is a notable exception, although Toyota offers only four models in Brazil. For the remaining countries, markups range from 13 to 15 percent, which seems plausible. Notably, they tend to be higher in home countries. For example Fiat's markups are highest in Italy, and the firm with the highest markups in Italy is Fiat. The same is true for VW and Germany.

[4] We apply the home dummy on the basis of the brands historic national identity even when it has been purchased by a foreign multinational. For example, Jaguar is considered a UK brand, and Seat is considered a Spanish brand, even though they have both been purchased by foreign multinationals.

Table 3: Estimated parameters of the demand equation

Variable	Estimate									
	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA	R.C. Std
Constant	-4.311	-8.109	1.034	-7.422	-6.903	-8.799	-9.347	-4.848	-2.995	6.470
	(2.976)	(3.952)	(4.982)	(4.893)	(4.938)	(4.603)	(5.067)	(3.949)	(5.501)	(3.261)
HP per Weight	0.062	-0.030	-0.042	-0.031	-0.013	-0.015	-0.012	-0.004	-0.044	0.078
	(0.017)	(0.033)	(0.036)	(0.032)	(0.028)	(0.035)	(0.033)	(0.032)	(0.040)	(0.019)
Size	1.923	0.386	-0.525	0.347	0.643	0.326	0.334	0.106	-0.339	0.821
	(0.397)	(0.450)	(0.546)	(0.540)	(0.522)	(0.501)	(0.557)	(0.473)	(0.568)	(0.307)
MPGCITY	-0.167	0.018	-0.036	0.004	-0.050	0.031	0.054	-0.022	0.012	0.163
	(0.050)	(0.042)	(0.048)	(0.042)	(0.038)	(0.041)	(0.039)	(0.038)	(0.047)	(0.053)
Price		$\overline{\alpha}$				π_p			σ_p	
		11.077				-0.819			1.186	
		(0.619)				(0.052)			(0.093)	
Home Dummy, η	1.131									
	(0.041)									

[1] The units for HP per weight, size, MPG, and price are horse power per 1000 kg, m^2, mileage per gallon, and 10 thousand dollars, respectively.

Table 4: Median markups of brands across markets – in percentage

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat	11.2	13.8	14.7	13.6	13.0	14.2	13.6	15.3	15.1
Ford	10.2	14.0	13.3	14.0	13.5	14.4	15.0	14.4	13.8
GM	10.6	15.3	13.8	14.3	13.5	14.6	15.3	15.1	14.7
PSA	10.5	15.3	-	14.2	14.3	16.0	14.7	14.7	-
Toyota	6.6	13.9	13.6	13.7	13.1	14.1	14.0	14.6	14.6
VW	10.9	15.5	13.1	15.4	14.5	14.8	14.7	14.9	13.6

[1] Percentage markups are defined as $p/c - 1$.

Table 5: Estimated parameters of the cost equation

Variable	Estimate
Assembly to Market (τ)	0.009
	(0.003)
Assembly to HQ (γ)	0.018
	(0.006)
HP per Weight	0.003
	(0.0002)
Size	0.056
	(0.008)
MPGCITY	-0.001
	(0.0004)
Fixed σ	0.01

The units for HP per weight, size, MPG, cost, and distance are horse power per 1000 kg, m^2, mileage per gallon, and 10 thousand dollars, miles, respectively.

This illustrates that home bias endows some market power on domestic automakers, however it does not explain the mechanism through which this home bias arises. It may have been established through a stronger sunk investment in the brand (i.e., longer advertising history or larger dealer networks) or on the other hand could be related to the patriotic feelings of consumers. In either case, recovering these differences in markups is an important so that they do not contaminate our estimates of trade costs.

Once markups and marginal costs are calculated, completing the estimation of trade costs is a relatively straightforward application of nonlinear least squares on (6). However, we have found that estimates of σ_ν are very close to zero. The finding that σ_ν is small is economically relevant. It means that gains from diversification are essentially irrelevant to building additional assembly plants, and instead multiple plants allow firms to extend their market reach by lowering trade costs. Unfortunately, a small value for σ_ν, generates numerical instability when estimating τ and γ due to the high value of the term inside the exponential function. To address this, we fix σ_ν at 0.01

and estimate the remaining parameters.[5]

On the cost side estimation, Table 5 present the results of our preliminary specification, which also includes dummies to control for model and year fixed effects. In addition, we control for productivity differences across assembly locations by including assembly country dummies, ϕ_ℓ, as a result the low cost assembly location is not necessarily the closest. Therefore, one potential reason to incur foreign the foreign production costs we estimate is to take advantage of a highly productive assembly location or favorable tariff rates. The results on the cost of power (horsepower per weight) and size are intuitive, in that more powerful and larger cars are more costly to build. The coefficient on fuel efficiency is small and not statistically significant. However, the most interesting results are the effect of log distance to the assembly locations on cost. We find evidence of both traditional trade costs and foreign production costs related to the distance between a manufacturers headquarters and the assembly location. The estimates imply that if the distance between the assembling location to the market increases by 10%, then the cost increases by \$27 (for a car with cost of \$30,000); if the assembling location to the headquarter increases by 10%, then the cost increases by \$54 (for a car with cost of \$30,000). While these numbers may appear small, one should keep in mind that they are per car supplied and are therefore substantial given the volume of trade and foreign production. Note that our estimate of trade costs is separate from the impact of tariffs, which are controlled for separately in our specification. When tariff rates are not controlled for, the trade cost coefficient more than doubles, which provides some indication that tariffs remain a substantial source of trade frictions in the global auto market. Somewhat surprisingly, the foreign production cost coefficient is higher than the trade costs coefficient (although the difference is not statistically significant). While they are less concrete than trade costs related to the physical shipping of a product, it appears that foreign production costs of managing production at a distance do impose significant costs on multinational firms.

To investigate the role of trade and foreign production costs in more detail, we calculate the impact of these costs on the overall cost of automobiles. First, we consider the magnitude of trade costs paid relative to what they would be if the car were sourced from the market country as a percentage of total costs (i.e., if firms chose to minimize trade costs). Table 6 provides the median trade cost for major firms in each country. Note that these costs are illustrative of the costs associated with trade, and may be offset by productivity differences across countries. Overall, trade costs measured in this way range from zero to 3 percent of the total cost of the car. Trade costs are particularly low in Brazil, since most cars sold in Brazil are sourced there or in another South American country. Geography and plant constellations are also apparent when examining trade costs of Toyota, which has a strong presence in North America, but pays substantial trade costs in Europe. Moreover, each of the European countries pay the lowest trade costs in their home country, although they still produce some cars elsewhere to be imported home. Also interestingly, while Volkswagen pays significant trade costs to produce in the US of 1.4 percent, Toyota pays much less, reflecting its reliance on North American plants.

We carry out a similar exercise with regard to foreign production costs in Table 7. Here, we calculate foreign production costs relative to the minimum cost available by producing the model in the headquarters country. We find that foreign production costs can be as high as 5 percent for some firms. In particular, there are significant costs in Brazil, again, this is because most Brazilian models are produced in South America, far away from the firms' headquarters. European firms pay significant foreign production costs for models marketed in North America which are typically produced in North American plants.

[5] We have also experimented with setting σ_ν at .0001 and .1 and have found similar results. Moreover, we have verified that we are able to accurately recover trade costs when σ_ν is low using this approach in a Monte Carlo experiment.

Table 6: Trade cost (due to market-assembly distance) in overall cost for brands across markets, percentage

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat Group Automobiles SpA	0.0	2.4	0.2	1.0	1.2	0.9	1.8	0.4	0.1
Ford Group	0.1	3.0	0.3	1.7	2.0	1.5	2.4	1.8	0.1
General Motors Group	0.1	2.6	0.3	1.1	1.6	1.4	0.9	1.4	0.1
PSA Group	0.4	2.3	-	1.0	1.0	0.7	1.6	1.0	-
Toyota Group	0.1	3.5	0.4	2.1	2.2	2.1	2.4	2.2	0.6
Volkswagen Group	0.1	2.2	1.6	0.5	0.8	1.0	1.5	1.0	1.4

Of course, the costs assessed in Tables 6 and 7 are a function of firms location decisions, so they are equilibrium outcomes. However, they illustrate some of the tradeoffs involved in firms' sourcing decisions. In particular, foreign production costs paid are not trivial, which implies that firms are able to gain significant cost advantages (either by taking advantage of productivity gains or in the reduction of trade costs in nearby markets) to offset these management costs. At the same time, some firms (particularly in Brazil) have taken actions that are qualitatively similar to attempting to minimize trade costs, resulting in relatively low trade costs paid. Finally, firms appear to be more willing to incur trade costs in Europe, where a single firm can conceivably supply several different markets.

Table 7: Foreign production cost (due to HQ-assembly distance) in overall cost for brands across markets, percentage

	BRA	BEL	CAN	DEU	ESP	FRA	GBR	ITA	USA
Fiat Group Automobiles SpA	5.3	1.0	5.1	1.4	1.5	1.3	1.7	1.0	5.1
Ford Group	2.6	2.0	0.1	1.8	1.6	2.2	1.8	2.0	0.2
General Motors Group	2.5	2.6	0.3	2.6	2.7	2.7	2.4	2.8	0.2
PSA Group	5.3	1.8	-	1.6	1.8	1.6	1.8	2.1	-
Toyota Group	6.5	2.9	4.7	3.0	2.6	3.2	2.8	3.1	3.7
Volkswagen Group	5.9	1.6	5.1	1.3	2.1	1.9	1.7	1.8	4.8

6 Further Applications

We described an ambitious database building effort and an innovative methodology to estimate trade costs and preferences in a key industry populated by oligopolistic multinational firms. We finish by describing a number of interesting questions that can be addressed by applying this machinery:

1. Effect of further trade and investment liberalizations: our results indicate that tariffs and

foreign production costs are important in determining costs and market shares. A counterfactual exercise, feasible with the data and results we have so far, is to reduce those tariffs and foreign production costs and simulate the effect on total automotive sales, market shares by producers and consumer surplus.

2. Importance of regulations: apart from tariffs, there is a wide range of regulations that affect the demand for cars, such as environmental standards and differential tax regimes based on engine power or vehicle size. The issue of harmonizing these regulations lies at the core of ongoing negotiations for the Transatlantic Trade and Investment Treaty between the US and the EU. It is widely claimed that such regulations present de-facto barriers to trade. Our framework could be used to analyze this claim and the quantitative effect of reducing such barriers. In our current estimates, regulatory differences across countries show up as a national preference for characteristics, but in as much as they can be expressed in dollar or ad-valorem terms we can vary them in counterfactual analysis. We are planning to collect additional data to facilitate this exercise.

3. The sources of home bias: so far, the home bias term is a residual. We are planning to shed further light on its determinants, such as the extent of dealer networks and consumer capital through reputation (crudely approximated by the initial entry date of a brand into a market and past advertising). Additional data collection effort will enable us to answer this question.

4. Endogenous production locations: in our current analysis we take the location of assembly locations as given. Endogenizing the assembly location decisions would enable to extend the counterfactuals to allow for firms re-optimizing their location choices. Tintelnot (2014) shows how to solve a model with fixed costs of establishing a plant location and the possibility of export platform sales in a model with monopolistic competition. Jia (2008) solves a strategic location game of two firms. We are exploring whether these approaches can be combined with our framework to solve the complex location decision.