Downstream Competition and Exclusive Dealing*

Franco Esteban Cattaneo†

November 15, 2018

Abstract

I empirically investigate the role of downstream competition in the use of exclusive dealing and quantify its effects in the formation of the car retailing networks. I estimate a model where exclusivity impacts both supply and demand, and dealers choose which brand to offer in a strategic manner. These choices frame product and brand availability in the market and the retail networks for manufacturers. In my model, dealers have incentives to add more brands in order to sell a wider set of products, but their interest to differentiate from local rivals limits this option. Moreover, manufacturers could raise costs anticompetitively to deter dealers from selling products of rival brands. I analyze the potential for this foreclosing channel by estimating fixed cost differences between exclusive and non-exclusive stores using moment inequalities. I find that multi-dealing has an average cost advantage between -€10,000 and €620,000. These numbers indicate that downstream competition, instead of anticompetitive motives, explains a more substantial part of the prevalence of exclusive dealing in the market.

Keywords: Exclusive dealing, retail networks, geographic competition, partial identification, moment inequalities, fixed costs

JEL Codes: L10, L42, L62, L81

---

*I am deeply grateful to my supervisor Michelle Sovinsky for her guidance and support. I thank Alessandra Allocca, Robert Aue, Thiago Cacicedo, Francesco Paolo Conteduca, Niklas Garnadt, Laura Grigolon, Phil Haile, Ruben Hipp, Nicolas Schutz, Hidenori Takahashi, Emanuele Tarantino, Frank Verboven and Yihan Yan for their useful suggestions. I also thank the attendants to the ENTER Jamboree 2018, and XXXIII Jornadas de Economía Industrial for their useful comments. The usual disclaimer applies.

†Center for Doctoral Studies in Economics, University of Mannheim, Germany (franco.esteban.cattaneo@gess.uni-mannheim.de).

1 Introduction

Despite attempts by the competition authorities to encourage car retail outlets to sell more than one brand, multi-brand dealerships are a rarity in Europe[1]. Manufacturers often force their retailers to sell other brands under dedicated corporate identities using separate showrooms and with different personnel. These requirements increase the costs of multi-brand dealing and deter retailers from taking such options. However, multi-dealing reduces the incentives for manufacturers to invest in their retail stores, especially with respect to marketing and brand image. This is an important consideration, as investments in promotional efforts can be substantial: for example, in the case of car retailing, in Spain alone, they account for 13.4% of total sales revenue.

In this paper, I study how the presence of exclusive contracts shapes the competition between retailers from an empirical perspective. On the one hand, it can be in the interest of the retailer and the manufacturer to sign exclusive dealing agreements in order to preserve the returns from investments in brand promotional activities (Besanko and Perry, 1993), and differentiate from competing retailers (Besanko and Perry, 1994). On the other hand, the manufacturer may use exclusivity to soften competition by raising the retailers' costs to offer other brands[2]. In particular, I assess the question of whether the prevalence of exclusive dealing between manufacturers and retailers emerges as a result of competition shaping the market or whether there is a scope for manufacturers to indirectly deter intra-dealer competition by increasing the costs for dealerships to sell for other brands.

The existing empirical literature has studied these two forces in isolation. More specifically, Asker (2016) looks at the possibility that exclusive dealing gives rise to foreclosure of competing brands, while Nurski and Verboven (2016) study the role of exclusivity as a means to achieve higher demand for the retailer. However, only a combined analysis can help competition authorities to quantify the merits of exclusive contracts.

[1] MEMO 10/217 by the European Commission: "The old rules have had little impact on favouring multi-dealerships, which continue to be determined by the size of the dealers and their geographical location – multi dealerships are more likely to happen in remote areas and within large dealer groups that have buyer power."

[2] Anticompetitive motives for exclusive dealing have been addressed extensively. E.g., Aghion and Bolton (1987), Rasmussen et al. (1991), Segal and Whinston (1996), Bernheim and Whinston (1998), or Calzolari and Denicolò (2015)

I estimate a structural model of demand and supply of the car retail market that quantifies the diverse effects of exclusivity. Including both sides of the empirical model to analyze exclusivity helps to construct a full picture for regulation of this kind of vertical restrictions. On the one hand, only modeling demand misses the potential effects of exclusive contracts endogenously shaping market structure for distribution. On the other hand, focusing on capturing supply differences in retail between exclusive and non-exclusive dealers does not allow for other motivations for exclusive dealing other than excluding rivals. My results indicate that downstream competition, instead of anticompetitive motives, explains a more substantial part of the prevalence of exclusivity in the car distribution market.

There are several challenges associated with this exercise. First, in order to estimate a model where exclusive dealing impacts both supply and demand, I require data on (i) car sales registrations, (ii) dealer locations, (iii) brands sold at each dealer, as well as (iv) demographics of consumers. I collected this information by combining existing sources with comprehensive self-collected data. The main difficulty in constructing these data was to distinguish and classify multi-dealerships and exclusive dealers since many appeared to be disguised under separate showrooms and names. To address this ambiguity, I define a firm as a multi-dealer when it has adjacent showrooms which belong to the same owner. These classification efforts resulted in a novel dataset, whose features I discuss in detail shortly.

Second, I model retailers' choices of which brands to offer, making my empirical framework the first one to include such a feature. I use these firm choices to uncover the fixed costs of operating a dealership. Specifically, following a similar principle to that of Asker (2016), I compare the estimated costs of exclusive and non-exclusive dealerships to estimate the difference in the costs of signing such arrangements for the retailers. These costs arise when it is more costly to establish a dealer selling more than one brand than a dealer selling each of these brands separately.

Similarly to entry games, local competition across dealers leads to a multiplicity of equilibria of the simultaneous move game in which retailers decide what brands to deal for, making maximum likelihood methods unfeasible for estimation. I estimate bounds to fixed costs using moment inequalities defined by equilibrium play (Pakes, 2010; Pakes, Porter, Ho, and Ishii, 2015) and methods developed to draw inference in these cases (Andrews and Soares, 2010). This approach enables me to overcome the problem of multiplicity at the cost of losing point identification of the parameters.

Another challenge related to the use of moment inequalities and equilibrium play for estimation is selection on unobservables. This problem arises when one makes inference based on observed choices to estimate fixed costs. In particular, when estimating fixed costs of selling for a particular brand, the upper bounds are identified by those dealers selling for that brand and the lower bounds by those who do not. Using this principle alone yields biased estimates since the dealers selling a particular brand most likely have better (lower) fixed costs than the dealers not selling it. The resulting (biased) estimates present themselves in the form of upper bounds for the parameters that are too low and lower bounds that are too high. Therefore, the estimated sets may not contain the true parameters.

A methodological contribution of this paper, then, is to introduce a strategy to deal with the selection problem that arises in these situations. The approach I propose is grounded on the observation that, conditional on observables, equilibrium play by dealers reflects their need to differentiate from neighboring rivals, and on the fact that these equilibrium choices are likely to replicate permuting brand offerings across dealers. The intuition is that the differentiation between dealers within a local market is a stable equilibrium prediction, and not which dealer sells what brand. Based on this idea, I propose conditions under which I can create new inequalities using multiple perturbations from equilibrium play. These inequalities allow me to derive moment conditions that are not dependent on choice and hence are free from selection.

In my empirical model, exclusive dealing comes into play in three ways. First, it enters the demand as a parameter in the consumers' utility of purchasing a car. This parameter captures enhanced consumer experience owing to improved customer service or better promotion. Second, multi-dealing enters the fixed costs paid by dealers as a cost shifter that I estimate. This parameter captures potential additional costs (or costs savings) related to selling for more than one brand. Finally, the choices of whether to deal exclusively are endogenous to the model, meaning that these decisions take the effects above into account, as well as the competitive environment in the market.

The demand framework has similar characteristics to the one in Nurski and Verboven (2016), where dealers differentiate from each other spatially, and exclusive contracts enter demand as a product characteristic. This demand shifter represents a taste for exclusivity, due to premium service or additional promotional and retailing efforts.

As mentioned above, I complete the model by allowing retailers to endogenously choose their brand offerings in a simultaneous move game setup. On the one hand, retailers want to differentiate from each other by offering different products to dealers geographically close. On the other hand, they want to sell popular products. Moreover, in the case of manufacturers raising costs of multi-dealing, exclusive dealing may appeal to downstream competitors because of its lower fixed costs.

In environments without intense competition, exclusive dealing has the effect of limiting the variety of products offered and narrowing demand for the retailer. Nevertheless, in the presence of fierce competition, single branding permits differentiation across smaller dealers. This interaction between spatial and product differentiation downstream is internalized by manufacturers, who set product prices in accordance with their distribution networks. Exclusive dealing eliminates competition among products of different brands within a retailer.

There is a longstanding debate about exclusive dealing contracts in competition policy because of their potential foreclosing effects. This controversy had its start in the literature with Posner (1976) and Bork (1978), whose work concluded that exclusive contracts could not deter entry from a more efficient competitor. My paper relates to the vast and rich theoretical literature that developed trying to refute this view. The main takeaway of this literature is that, although contracts of this kind might have exclusionary effects, their existence can be beneficial, by boosting investment and retailing efforts[3].

There is also growing literature on the empirics of exclusive dealing. My work links most directly to two papers in this literature and complements them. Asker (2016) develops a foreclosure test for the beer market in Chicago. He uses demand estimates and prices to infer distribution costs for brewers. He compares these costs between areas where Miller and Anheuser-Busch use exclusive contracts and areas where they do not and finds no statistical evidence of foreclosure. My approach shares similarities with Asker (2016) because I also use demand estimates to infer costs downstream and compare them between exclusive and non-exclusive dealers.

[3] Apart from the previously mentioned papers, Fumagalli and Motta (2006), and Simpson and Wickelgreen (2007) introduce the role of competition among firms in the downstream market as a force affecting the incentives to sign exclusive contracts and their potential for exclusion. Besanko and Perry (1994) explore the role of spatial differentiation across retailers. Sass (2005) provides a comprehensive overview of the main mechanisms used in the literature to rationalize the use of exclusive dealing.

However, I additionally include demand-side effects of exclusive dealing, and I focus on differences in fixed costs and allow for endogenous market structure, while he focuses on variable costs and keeps the market structure fixed.

The industry and demand modeling links my paper to Nurski and Verboven (2016). They estimate a model of spatial demand and perform counterfactuals that assess the collective incentives for incumbent manufacturers to maintain these agreements. While Nurski and Verboven (2016) make an extensive analysis of demand and manufacturers' incentives for exclusive contracts, I model the distribution network and estimate the fixed costs borne by these retailers. My model contains the channels for exclusive dealing of Nurski and Verboven (2016), where it shifts utility and it lowers product availability for rival brands. In addition, my model incorporates supply side motives for exclusive dealing, where retailers might deal with only one brand because it is cheaper for them to do so. To the best of my knowledge, this is the first paper that explores jointly supply and demand side mechanisms for exclusive dealing.

Other work in this area includes Ater (2015) who finds that exclusive contracts between fast-food restaurants and shopping malls impact competition negatively by lowering the number of restaurants, increasing prices and limiting total sales. Eizenberg et al. (2017) focus on the dynamic effects that exclusive contracts between Intel and PC makers had on the development of its competitor AMD. Chen (2014) analyzes the entry of specialty beers and does not find any foreclosing motives behind exclusive contracts by incumbent breweries.

This article is also related to the stream of literature on endogenous product offerings. Examples include Fan (2013) in the newspaper market, Draganska et al. (2009) on the variety of vanilla ice cream, or Eizenberg (2014) in the PC market. In my model, the strategic considerations that determine dealers' endogenous choice of brands are analogous to those shaping firms' decision to introduce a product in these papers. Dealerships decide with their brand offerings what bundles of goods to offer and with them determine (endogenous) product availability in the market.

Finally, the estimation of fixed costs relates to the literature that uses moment inequalities to overcome the problem of multiple equilibria (Ciliberto and Tamer, 2009; Pakes, 2010; Pakes, Porter, Ho, and Ishii, 2015). This approach has been used recently in a number of empirical applications in industrial organization and trade (e.g. Holmes, 2011; Morales, Sheu, and Zahler, 2015; Houde, Newberry, and Seim, 2017; Wollmann, 2018). One main difference across these papers is how they deal with the potential selection issues induced by structural disturbances in the fixed costs. I contribute to this literature by introducing a new way to circumvent this issue.

In summary, my contribution is manifold. First, I estimate a model that combines supply and demand to quantify the effects of exclusive dealing. Second, I determine retail brands within the model. Third, I introduce another strategy to deal with the selection problem common in the literature using choice data. Finally, I construct a novel dataset containing data on car sales and retail points in Spain.

The paper proceeds as follows. In the next section I describe the data. The model is presented in section 3. Sections 4 and 5 describe estimation and results respectively. Section 6 concludes.

2 Data

My dataset concerns the market for cars in Spain from July 2016 until August 2017. First, I use data on car registrations from the traffic registry. Second, I supplemented these car registry data with additional characteristics of cars that I collected from specialized magazines. Third, I collected data on car dealerships locations from web searches. Finally, I use information on population demographics and locations of consumers from Governmental Offices. I discuss each of these data in turn.

2.1 Car sales and characteristics data

I obtained data on car sales from the Spanish Directorate-General of Traffic (DGT)[4]. These data consist of daily information on all cars registered in the Spanish territory starting in December 2014. The data include a written description of the car model and its Vehicle Identification Number (VIN). They comprise a number of car characteristics such as engine displacement, horsepower, type of bodywork, number of seats, energetic propulsion, and the postal code and municipality where the car was registered.

I use observations corresponding to new cars, 4x4s or small pickups used for

[4]The data can be observed at `https://sedeapl.dgt.gob.es/WEB_IEST_CONSULTA/microdatos.faces`, and its documentation (in Spanish) at `https://sedeapl.dgt.gob.es/IEST_INTER/pdfs/disenoRegistro/vehiculos/matriculaciones/MATRICULACIONES_MATRABA.pdf`

non-commercial purposes. In total, in the period between July 2016 and August 2017, there are 1,091,932 registrations from 7,712 different municipalities. I chose this time window because it approximately coincides with the period on which I was able to collect the data on dealerships.

Table 1 shows market shares for car makes and models. It is notable that no make has a market share close to the 30%, which the General Vertical Block Exemption considers to be worrisome for legal vertical agreements[5]. In particular, no car make appears to dominate the market: all market shares are below 10% and the market leaders change across time periods and geographic regions[6].

Table 1: Market shares of best selling car makes and models

Make	Share	Sales	Model	Share	Sales
Peugeot	8.44%	70,805	Leon	2.51%	21,066
Renault	7.68%	64,416	Qashqai	2.44%	20,464
Volkswagen	7.22%	60,578	Sandero	2.22%	18,652
Seat	6.53%	54,790	Golf	2.19%	18,351
Ford	6.27%	52,570	Ibiza	2.16%	18,164
Opel	5.99%	50,266	Clio	2.01%	16,894
Citroen	5.90%	49,535	308	1.86%	15,594
Toyota	5.43%	45,577	Megane	1.82%	15,257
Nissan	5.15%	43,192	Corsa	1.79%	15,042
Kia	4.75%	39,817	Tucson	1.74%	14,613
Total	100%	839,086	Total	100%	839,086

I collected data on car characteristics from a series of specialized magazines (primarily autobild.es and autopista.es). These characteristics include list prices, measures of fuel consumption, car dimensions and weight. The data are detailed at the model (e.g. Ford Fiesta), version (e.g. Ford Fiesta 3P), and trim (e.g. Ford Fiesta 3P 2008 1.25 Duratec 82CV Trend) level.

I constructed a baseline model by merging the two datasets. First, I classified the models from the registry's string descriptions using automatized text analysis. Subsequently, I used information on bodywork type, measures, number of doors and horsepower to determine the car's version. Finally, I matched each registry entry to the car trim with the closest identifying characteristics. I define a baseline model

[5] OJ L 102, 23.4.2010, p. 1–7

[6] See Figure **??**, where I present market leaders by year and administrative region for the years previous to the sample.

as the mean of all merged trims. This linkage approach preserves a larger part of price variation in the data and controls for the fact that, especially in higher-end cars, within-model price dispersion plays a sizable role.

I excluded car models absent at more than 30 provinces[7] and car categories that are not in direct competition with passenger cars (e.g. big vans, luxury sports cars). I aggregated the data at the province level. This market definition preserves the geographic disaggregation of the data without having markets with market share of zero for products with low probability of being chosen. I dropped provinces outside the Iberian peninsula (i.e. Canary Islands, Balearic Islands, Ceuta and Melilla) because they are geographically apart from the rest of the country.

Table 2 shows some descriptive statistics for different car characteristics after matching them with the registry data. The data comprise 43 out of a possible 52 provinces[8] and 234 car models with significant variation in their characteristics. The average price is around €34,300, and the average horsepower around 144 CV, but they both have a large dispersion.

Table 2: Descriptive statistics of car characteristics

	Mean	Std. Dev	Min	Max	Obs
Model					
Horsepower	144.34	60.23	60	422	234
Weight (100 Kg.)	14.55	3.35	8.05	24.65	234
Size (m^2)	8.03	1.06	4.48	10.36	234
Fuel Cons. (l/km)	5.08	1.19	3.3	10.61	234
Price (€10,000)	3.43	2.16	1.02	14.86	234
Markets					
Provinces					43

2.2 Dealer data

I next require data on locations of dealerships as well as which brands are for sale at each dealership. This last part is particularly crucial as it drives the classification of a dealership as exclusive. Unfortunately, these data were not available in Spain so I collected them manually. First, I gathered the data on locations from online appli-

[7]Spain is divided into 23 Autonomous Communities that are subdivided into 52 provinces.

[8]These 43 provinces include 40.3 million inhabitants out of the total of 46.5 million in Spain.

cations that each manufacturer has available on their websites. These applications are normally used for manufacturers to inform about their available points of sale. From them, I obtained a list of dealerships with their location and services for each car brand. Second, I manually combined the observations that were points of sale for more than one brand.

I consider two observations from different brands to be available at the same dealership (making it a multi-dealer) if they are (i) located adjacent to each other geographically and (ii) have the same owner. This definition is based on the observed patterns for multi-dealers, where normally different car makes have separated showrooms and different names even if they are operated by the same owner. Since construction and geographic distribution varies across urban and rural areas, I also consider observations separated by a street intersection as contiguous, but not those separated by another building or dealership.

Using this definition, the data show 44% of dealerships are multi-dealers, i.e., offer more than one brand (which constitutes 66% of the total points of sale). There are some brands that have their dealership networks integrated. For example, Citroen and DS, or Renault and Dacia share all their points of sale. However, if I count brands with integrated networks as exclusive dealing, the percentage of multi-dealers declines substantially to 22% (41% of points of sale)

Figure 1: Shared dealership networks

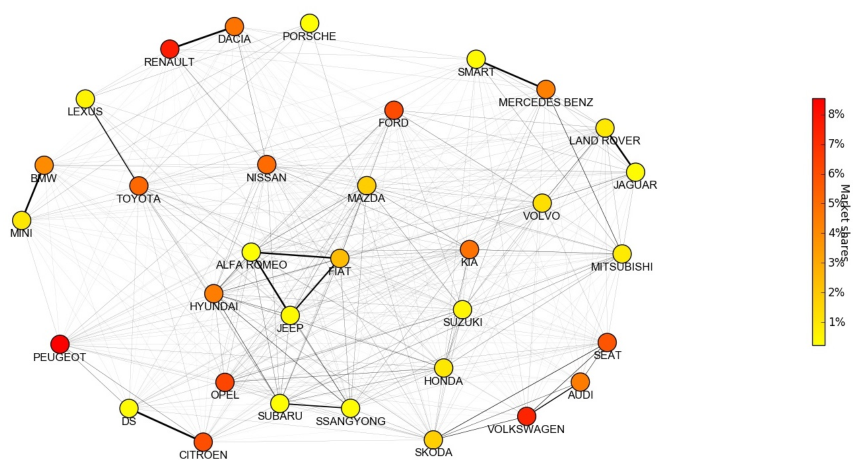

Figure 1 summarizes general patterns in distribution networks. Node colors represent the market share of the manufacturer in the market, whereas the thickness of node links are the percentage of shared dealerships between the two auto makes. This percentage is measured as the total number of shared dealerships over the number of dealerships for the smallest of the two brands. Particularly bold links between nodes of the same holding (e.g. VW Group, PSA, FCA) indicate that common dealerships are substantially more likely between makes belonging to the same company - extreme cases are Renault/Dacia and Citroen/DS where distribution is completely integrated.

The second observable pattern is the higher *degree*, in the sense of more shared dealerships, of some brands with relatively low market shares. This phenomenon is more prominent among Asian car makes (Honda, Mazda, Hyundai, Subaru), most of which do not belong to any particular holding group and share more dealerships with more different brands than market leaders like Volkswagen or Renault.

Figure 2: Number of dealerships per brand

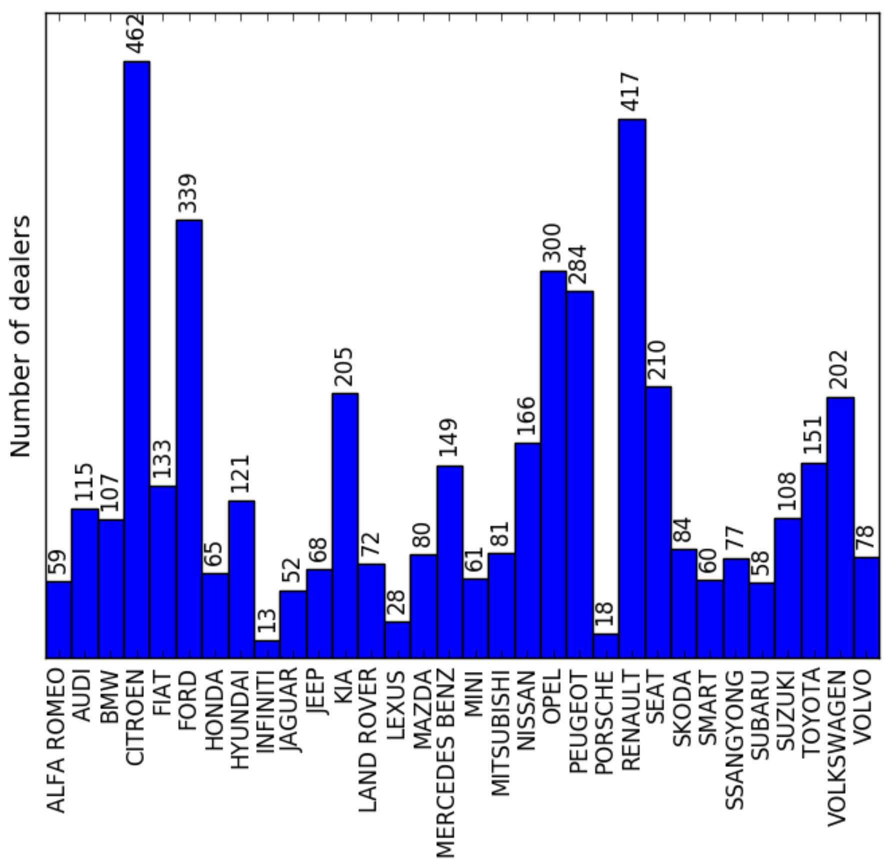

Figure 2 shows the number of dealerships per brand. There are clear differences in terms of dealer density across car makes and reasons to think that these differences are not solely driven by demand concerns. For example, Renault/Dacia have as many as 417 dealers in the whole territory whereas Volkswagen has 202 and their differences in sales are only 0.46 percentage points. Differences in dealer density are more pronounced in scarcely populated areas, where Renault/Dacia, Citroen/DS, and Peugeot are spread across provinces, whereas the rest of manufacturers are only present in urban areas.

These differences in geographic coverage are also reflected in Figure 3, where the points of sales for Honda, Volkswagen and Renault are plotted. Red points denote exclusive dealerships, while blue points are multi-dealerships. Another empirical regularity observable from Figure 3 is a larger tendency towards multi-dealerships in rural areas as compared to more densely populated areas. This tendency is normally attributed to the higher buyer power of the dealers present in these places due to the lower competition that they face.

Figure 3: Exclusive and Non-Exclusive dealerships for Honda, VW and Renault

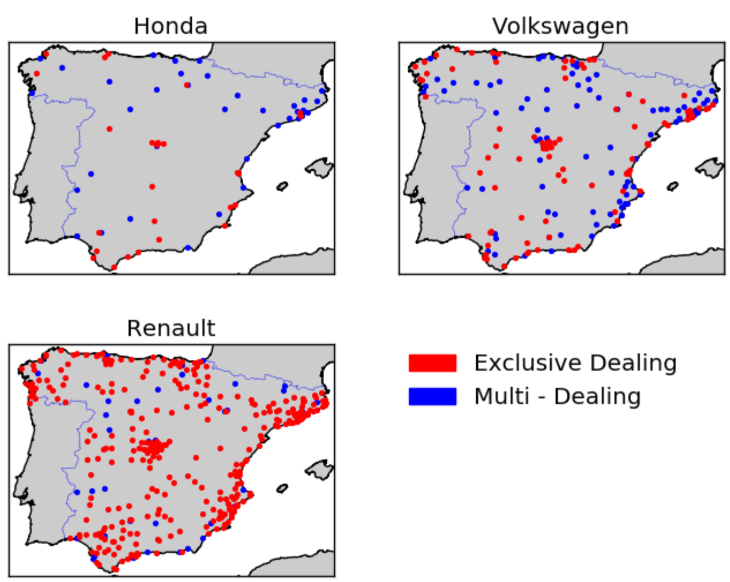

2.3 Geographic locations

One important consideration is how far consumers are willing to travel to purchase their cars. In order to incorporate this I use data on geographic positions at the municipality level from the National Geographic Institute (IGN). This dataset provides geocoordinates of the boundaries of each city. Using these boundaries, I drew a large number of random locations in each province. I weighted the draws by population size of each municipality within every market, so that the geographic distribution of consumers is consistent with the actual one within a province.

Table 3: Descriptive statistics of distance to closest dealer (in km)

	Renault	**Volkswagen**	**Mercedes**	**Mitsubishi**	**Infiniti**
Madrid	5.74	6.39	6.94	9.00	11.08
Barcelona	2.70	3.32	3.66	6.12	12.32
Murcia	8.51	12.19	12.86	15.21	23.93
La Coruña	7.28	8.05	12.95	14.92	23.87
Cáceres	15.39	21.54	22.85	35.43	148.44
Cuenca	23.34	25.70	24.75	46.25	98.75

Table 3 shows the average simulated distance to the closest dealers for a number of brands in a number of provinces. The brands in the table have different levels of dealer density and are placed in the columns from most dense (Renault) to less dense (Infiniti). The provinces on the rows consist of the two most densely populated markets (Barcelona and Madrid), two middle sized provinces (Murcia and La Coruña), and two very sparsely populated provinces (Cáceres and Cuenca).

It is visible that highly populated areas have greater dealer supply, and thus transport distances are significantly shorter. This pattern is consistent across brands. However, the change in distances is more than proportional when departing from more to less populated areas. This fact confirms the previous observation that differences in dealer density are more pronounced in areas with less urban development.

3 Model

The model consists of four stages. In the first stage, dealerships draw costs and decide what brands to offer taking into strategic consideration their local competitors. After dealer configurations are determined in the first stage, manufacturers determine their wholesale prices in the second stage. In the third stage, manufacturers set cars' list prices. Finally, consumers purchase their cars.

The primitives of the model are the utility parameters, the marginal costs for each car model and the fixed costs of establishing a dealership. In what follows, I introduce the model starting from the demand side.

3.1 Demand

I model demand using a random-coefficient-logit specification (Berry, Levinsohn, and Pakes, 1995). Individual i chooses what car $j \in J$ to buy. The indirect utility for an individual in market m from buying product j at dealership d is given by

$$u_{ijdm} = \delta_{jm} + \mu_{ijm} + \gamma_{idm} + \epsilon_{ijdm},$$

where $\delta_{jm} = x'_{jm}\beta + \alpha p_j + \xi_{jm}$ is the base utility for product j in market m. This term contains observable car characteristics x_{jm} and p_j that include fuel consumption, size, engine power, and price. The characteristics also include dummy variables for province and country of origin of the car. I use these fixed effects to proxy for unobserved effects that are market and brand specific. The ξ_{jm} includes car attributes that are observed by the consumer, but unobserved to the econometrician.

The heterogeneity in consumers is captured by $\mu_{ijm} + \gamma_{idm} + \epsilon_{ijdm}$, which consists of the ϵ_{ijd}, idiosyncratic consumer disturbances that are assumed to be distributed according to a type I extreme value distribution. Heterogeneity in consumers' sensitivity to prices are captured in the interaction term $\mu_{ijm} = \sigma y_{im} p_j$, where y_{im} represents income of consumer i.

The term $\gamma_{idm} = \gamma_1 \text{ED}_d + \gamma_2 \text{dist}_{id}$ explains the heterogeneity in consumers' access to dealerships in their choice sets. I follow Nurski and Verboven (2016) and capture the impact of these characteristics by using two attributes: distance to the dealer and a dummy variable equal to 1 if the dealer is exclusive (ED_d). The exclusive dealing dummy contains any possible demand effects that exclusive retailing can induce, e.g., being more prestigious, delivering a better service or enjoying additional promotional efforts.

Geographical distance to dealerships is also included in Albuquerque and Bronnenberg (2012) and adds a spatial dimension to the model. Its coefficient explains the impact that traveling distance to a point of sale has on the utility of consumers. I anticipate that coefficient is negative, as in Nurski and Verboven (2016) and Albuquerque and Bronnenberg (2012), meaning that consumers value proximity to dealers.

While I observe sales at a very local level, the exact point of sale where transactions take place as well as consumers' residences are unknown to me, so I need to simulate them. I simulate random consumer locations in each market from a distribution that draws with higher probability locations from municipalities that are more populated.[9] I compute for each of these simulated locations the distance to every dealer to get $\text{dist}(i, d)$, and I set ED_d equal to 1 if a dealer is exclusive and 0 otherwise.

An issue that arises with modeling purchases as a combination of product and dealer is that it expands the choice set for consumers exponentially, posing a substantial computational burden. I take a similar approach to Nurski and Verboven (2016) and assume the choice set to contain all possible car models from their closest available dealer. This assumption reduces the choice set in each market to a maximum of 238 products.

Reducing the dealerships in the choice set to the nearest simplifies computation, but imposes restrictions. This assumption eliminates a large part of dealership competition within car brands because no consumer can take two dealerships selling the same car into consideration. Competition among two retailers dealing for the same brand boils down to the spatial dimension, i.e., which of the two is closer to a consumer[10].

I believe this restriction is less relevant in the context of this paper since the focus is on downstream incentives to engage in exclusive contracts. Moreover, I limit its effects using a large number of simulations. In this manner, an area that has many dealers will have different closest dealers, whereas areas with fewer dealers will have the same closest dealers in every simulation.

I complete the discrete choice model of demand by introducing an "outside" option, which includes not purchasing a car, purchasing a car outside of the 238 models considered, or purchasing a car from a dealer-product combination outside of the ones allowed by the model. I assume this outside product to have a base

[9] Consumers are assumed to be distributed uniformly within a municipality. This distributional assumption does not seem to be restrictive given that Spain is characterized to have numerous small municipalities.

[10] The coordination of competition within a distribution network is the focus of some theoretical papers (e.g. Lin, 1990; O'Brien and Shaffer, 1993) and it is a rationale for exclusive dealing akin to that of exclusive territories (Rey and Stiglitz, 1988, 1995).

utility that is normalized to zero, i.e., $u_{i0m} = \epsilon_{i0m}$.

Following Nevo (2001), I group parameters into $\theta_1 = (\alpha, \beta)$, and $\theta_2 = (\sigma, \gamma_1, \gamma_2)$. Assuming that consumers purchase their most preferred car, the distribution of unobservables y_i, $\text{dist}(i, d)$, ED_d, and ϵ_{ijdm} define the simulated individual choice probability. Let x_m, p, and δ_m denote the vectors containing x_{jm}, p_j, and δ_{jm} for every car j and market m, then

$$A_{jm}(x_m, p, \delta_m; \theta_2) = \{(y_i, \epsilon_{ijdm}, \text{dist}(i, d), \text{ED}_d) \mid u_{ijdm} \geq u_{ikdm} \ \forall k = 0, 1, ..., J_m\}$$

defines the set of unobservables for which product j is chosen. Given parameters, the market shares for each product are defined as

$$s_{jdm} = \int\limits_{A_{jm}} \frac{\delta_{jm} + \mu_{ijm} + \gamma_{idm}}{1 + \sum_{k \in J} (\delta_{km} + \mu_{ikm} + \gamma_{idm})} dG_y(y) dG_d(\text{dist}, \text{ED}),$$

where the fraction term denotes the individual choice probability s_{ijdm}. Its formula comes from the assumed distribution function for ϵ_{ijdm}. The $G(\cdot)$ functions are the distribution functions for each unobservable. For simplicity, they are assumed to be independent of each other.

3.2 Price Competition

This part of the model follows the empirical literature that uses market data to infer marginal costs and upstream wholesale prices, (e.g. Sudhir, 2001; Brenkers and Verboven, 2006; Berto Villas-Boas, 2007). I assume that price setting takes place in two stages and that manufacturers set both. First, I assume that manufacturers set wholesale prices in order to maximize their profits. These wholesale prices are realized and observed. Subsequently, manufacturers set list prices such that retailers can extract a margin that is consistent with profit maximization.

Data limitations drive this assumption. I do not observe transaction prices, and therefore I use list prices to approximate them. Since there is one list price for each product across markets, manufacturers maximize the profits of the whole network of retailers so that, on average, retailers have incentives to comply with manufacturers' pricing instructions.

I explain the pricing of the model in an inverse order and start with the list prices. Let each dealership d have a profit function of the form

$$\pi_d = \sum_{m \in M} \sum_{b \in a_d} \sum_{j \in b} \left(p_j - p_j^w\right) \mathcal{M}_m s_{jdm}(\theta, p, a) - F_d(a_d),$$

where $\mathcal{M}_m$ is the size of market m in the set of all markets M. $q_{jdm} = \mathcal{M}_m s_{jdm}$ are the quantities of product j sold by dealer d in market m predicted by the demand model. I denote p_j and p_j^w as list and wholesale prices respectively.

I denote by $b \in B$ a brand in the set of all car brands. The set of brands that dealership d sells for is denoted by a_d, which is chosen from A_d - a subset of the power set of brands $\mathcal{P}(B)$. Finally, $a = (a_1, ...a_D)$ displays all brand offerings for all the dealers in the market, and $F_d(a_d)$ are the fixed costs of opening a dealership selling a_d. More detail on this will follow in subsections 3.3 and 4.

The profit function above yields list-pricing first-order conditions

$$\frac{\partial \pi_d}{\partial p_j} = \sum_{m \in M} \left(q_{jdm} + \sum_{b \in a_d} \sum_{k \in b} \left(p_j - p_j^w\right) \frac{\partial q_{kdm}(\theta, p, a)}{\partial p_j} \right) = 0 \text{ for all } j \in b \text{ and } b \in a_d.$$

These first order conditions are used by manufacturers $b \in B$ to set their list prices so as to maximize the profits of its joint network.

$$\sum_{d \in D} \mathbb{I}\{b \in a_d\} \cdot \frac{\partial \pi_d}{\partial p_j} = 0 \text{ for all } j \in b, \tag{1}$$

where $\mathbb{I}\{b \in a_d\}$ is an indicator variable that is equal to 1 if dealer d offers brand b (i.e., $b \in a_d$). From equation (1), one can rearrange its terms in matrix notation to get

$$q + \left(\sum_{d \in D} \Delta_d \right) (p - p^w) = 0, \tag{2}$$

where Δ_d is a $J \times J$ matrix where an element is equal to $\frac{\partial q_{jd}(a,\theta)}{\partial p_k} = \sum_{m \in M} \frac{\partial q_{jdm}(a,\theta)}{\partial p_k}$ if product j and k are sold by dealership d.

Notice that this expression is similar to the standard multi-product firms' pricing equations in many papers estimating demand (e.g. Berry, Levinsohn, and Pakes, 1995), except for the term including $\sum_{d \in D} \Delta_d$. In my model, this term emphasizes the role of dealer networks internalizing manufacturers' incentives. This conceptual difference can be best explained with an example. For two brands with integrated dealership networks (e.g., Renault and Dacia), the entries in this term are always going to be different from zero, and the sum is going to entail the same derivative as in a standard ownership matrix. In this case, integrated points of sale make compatible downstream pricing with common ownership. In general, common ownership is more prevalent in the pricing equation the more dealers the brands share.

Despite being restrictive, this modeling of list prices is sensible and captures a series of mechanisms that are important when studying exclusive dealing in a spatial market. First, notice that equation (1) is equivalent to the equilibrium condition arising from a model where all retailers set prices independently: p_j would still be the average retail price for product j. Second, list prices derived from these first order conditions capture the spatial dimension for product competition through the derivatives at the dealer level.

Finally, having solved for list prices, one can go back to the stage where wholesale prices are decided. Manufacturers' profit maximization for a firm f selling a series of brands b is

$$\max_{\{p_j^w\}} \Pi_f(p, p^w) = \sum_{b \in f} \sum_{j \in b} (p_j^w - c_j) q_j.$$

3.3 Entry

A dealer d in the pool of potential entrants E is located at l_d and chooses what brands to offer (a_d) from the set $A_d \subset \mathcal{P}(B)$. It can choose to offer one brand, e.g. $a_d = \{\text{Peugeot}\}$, or many brands, e.g. $a_d = \{\text{Peugeot, Suzuki, Subaru}\}$, or none, i.e. $a_d = \emptyset$. The last option corresponds to the case in which the entrant decides to stay out of the market. The set of entrants is $D \subseteq E$. Dealerships take their entry decisions a_d to maximize their expected profits given the choices of competing rivals ($\pi_d(a_d, a_{-d})$) and their information set $\mathcal{I}_d$:

$$\max_{a_d \in A_d} \mathbb{E}[\pi_d(a_d, a_{-d}) | \mathcal{I}_d] = \underbrace{\mathbb{E}\left[\sum_{m \in M} \sum_{j \in J_d} q_{jdm}(\theta, a)(p_j - c_j)\right]}_{\mathbb{E}[\text{VP}(a)]} - F_d(a_d). \tag{3}$$

where $F_d(a_d)$ are the fixed costs of establishing a dealer of type a_d. These costs are a function of the brands in a_d, and whether the dealer is exclusive or not. I assume the fixed costs to have a simple function of the form

$$F_d(a_d) = \sum_{b \in a_d} (F_b + \nu_d^b) + \mathbb{I}\{|a_d| > 1\} \cdot C_{MD} + \nu_d^l, \tag{4}$$

where, as explained in the previous subsection, F_b is the cost dealership d faces when offering brand b and C_{MD} are the potential additional costs that can occur when dealing with more than one brand. They also include structural disturbances ν_d^b and ν_d^l which represent unobserved idiosyncratic cost components that dealer d observes, but that I do not. ν_d^b are unobservable shocks to fixed costs that depend on the brand choices of dealer d, while ν_d^l are unobservable components to dealers' locations. I assume that these costs are such that $\mathbb{E}\left[\nu_d^b\right] = \mathbb{E}\left[\nu_d^l\right] = 0$. This functional form is very simple, but the parameter C_{MD} accounts for potential jumps in the cost function when transitioning from exclusive dealing to multi-dealing.

The term $\mathbb{E}[\mathrm{VP}(a)]$ denotes the expected variable profits of dealer d, and it entails two assumptions that facilitate estimation and are commonly shared in most applications (e.g. Holmes, 2011; Eizenberg, 2014; Houde, Newberry, and Seim, 2017). First, it implies that the expectations of the dealers are correct[11]. Second, it assumes that the dealers' information set does not contain any additional unobservable knowledge about its expected variable profits.

This profit function captures downstream competition in the model. Exclusive dealing enters both variable profits (through market shares) and fixed costs. Market shares bring strategic interactions between geographically close competitors into account. The estimated magnitude of the distance parameter in the demand determines, in turn, the relevant market for a dealership and the intensity of competition. If consumers are averse to driving far to buy a car, dealerships compete with each other locally, and more locally the higher this aversion is.

Multi-dealerships are more profitable in markets with higher isolation between points of sale as it allows the dealer to offer a more extensive selection of products and occupy a more substantial part of demand. In markets with a dense dealership

[11] Pakes (2010) and Pakes, Porter, Ho, and Ishii (2015) point out that a weaker condition on agents' expectations can also work. It is enough to assume that dealers do not have any systematic bias or deviation in their expectations. In other words, they can have wrong expectations, as long as they are not consistently wrong.

structure, exclusive dealing is favorable in that (i) it reduces costs, allowing competitors to stay in the market even with smaller sales, and (ii) it differentiates dealerships from each other by offering different sets of products, relaxing competition (Besanko and Perry, 1994).

4 Estimation

I estimate the model in three steps. First, I estimate the demand parameters $\theta_1 = (\alpha, \beta)$, and $\theta_2 = (\sigma, \gamma_1, \gamma_2)$. Using these estimates, I back out product unobservable characteristics $\hat{\xi}_{jm}(\hat{\theta}) = \delta(\hat{\theta}_2) - X_{jm}\beta + \alpha p_j$ and manufacturers' wholesale prices. Finally, I use all previous estimates together with equilibrium condition to estimate bounds on fixed costs and its parameters $F = (F_1, ..., F_B, C_{MD})$.

4.1 Estimation of demand parameters $\theta = (\theta_1, \theta_2)$

I estimate the demand model following the methods proposed in Berry (1994) and Berry, Levinsohn, and Pakes (1995, BLP). These estimation methods are based on equating predicted and observed market shares for every product and market, so as to then back out the value of average utility δ and minimize the difference $\xi(\theta) = \delta(\theta_2) - X\beta - \alpha p$. The model is estimated by General Method of Moments (Henceforth GMM, Hansen, 1982) using the moment condition

$$\mathbb{E}[Z'\xi(\theta)] = 0,$$

where Z is a matrix of instruments, and ξ is the vector of unobserved product characteristics. The estimates $\hat{\theta}$ are given by

$$\hat{\theta} = \arg\min_{\theta} \xi(\theta)' Z W^{-1} Z' \xi(\theta),$$

where W is an estimate of $\mathbb{E}(Z'\xi\xi'Z)$.

In order to control for potential correlation between unobservables ξ_{jm} and prices p_j, I use the set of instruments proposed in Berry, Levinsohn, and Pakes (1995). These BLP instruments include own car characteristics, sums of characteristics from the same manufacturer, and sums of car characteristics for rival products. I classified all car models into their market segments and performed these operations within segments for additional variation.

In addition to the BLP instruments, I incorporate as instruments neighboring demographics and rival dealer characteristics similar to Fan (2013) and Gentzkow and Shapiro (2010). These "Waldfogel" instruments (Waldfogel, 2003; Berry and Haile, 2009) make use of the geographic nature of dealer competition and modeling. For any dealership, its rival points of sale are those for which there exists at least one simulation draw with both of them in its choice set. Given this definition, I use as an instrument for a product in a dealership the demographics of simulation draws that have some of its rivals in their choice set, but not the original dealership.

The intuition of these instruments can be best described with an example. For dealership A, income in some neighboring area might not affect it directly because it does not receive any demand from it. It can, though, affect directly the demand of some rival retailer B that is closer to that area. In this manner, since endogenous variables for rival dealer B are affected by the income of this area, then it also affects through competition, the ones of dealership A. Similarly, since dealerships are locally competing, the distance to rival points of sale determines to a great extent whether a given location is considered to be far by consumers or not.

I require an additional assumption on the choice set of consumers to use the BLP instruments for the reason that their identification hinges on changes in the characteristics of rival products (Berry and Haile, 2014). I let the choice set of consumers be all models available at less than 80 kilometers of distance[12]. This assumption is sensible in the light of the very few cars that are bought from brands that are located far away and the low number entries that are registered in a province other than the one of purchase. It is also in line with empirical evidence, Murry and Zhou (2017) observe that less than 5% of car purchases take place at a distance further than 48 kilometers.

Finally, it is important to note that, as in most of the literature on endogenous product characteristics, estimating demand on observed dealerships might suffer from selection. This issue arises because choosing a brand for which to retail might be correlated with unobservable characteristics of the products offered by it. In this case, the timing of the model alleviates these concerns. When choosing which brands to deal, retailers are assumed not to know the realizations of unobserved product

[12] I performed robustness checks with 60 and 70 kilometer and they did not present any strong difference.

characteristics (ξ) and can only condition their choices on variables that are also observable to the econometrician. This argument is used in Eizenberg (2014), where it is also formalized.

4.2 Estimation of wholesale prices and unobserved product characteristics (ξ)

I recover product characteristics that are unobserved to the econometrician as the residual $\xi(\hat{\theta})$ product of $\delta(\hat{\theta}_2) - X\hat{\beta} + \hat{\alpha}p$ from the demand estimation parameters. The collection of all residuals ξ for a given product over all markets (i.e., the set $\{\xi_{jm}(\hat{\theta})\}_{m\in M}$), defines the empirical distribution function for the unobservables of that product that I use later for the simulation of expected variable profits.

Marginal costs are backed out using demand parameters and the distribution of consumer locations. Unfortunately, I do not observe transaction prices in the different stores, which limits the possibility of inferring wholesale prices at the store level. Following section 3.2, I use the equality in equation (2) to solve for the vector of wholesale prices p^w given the vector of list prices (p), the realized profile of brand offerings (a), demand parameters ($\hat{\theta}$), and the inferred unobservables ($\hat{\xi}_{jm}$)

$$p^w = p + \frac{q}{\sum_{d\in D} \Delta_d}.$$

I compute the product derivatives over prices at the dealer level approximating them via Monte Carlo integration. The derivatives are given by

$$\frac{\partial q_{jdm}(a,\hat{\theta})}{\partial p_k} = \mathcal{M}_m NS^{-1} \sum_{i}^{NS} \mathbb{I}\{d \in J_i\} \frac{\partial s_{ijdm}(a,\hat{\theta})}{\partial p_k},$$

where aside from previously introduced notation, I denote $\mathbb{I}\{d \in J_i\}$ an indicator function equal to 1 when dealer d is in the choice set of consumer i (characterized by J_i).

4.3 Estimation of fixed costs

I follow the literature using profit inequalities to estimate fixed costs (Ciliberto and Tamer, 2009; Pakes, 2010; Pakes, Porter, Ho, and Ishii, 2015). This approach flexibly accommodates multiplicity of equilibria and large action spaces. However, this comes at the expense of partial identification of fixed costs. In what follows, I describe the assumptions necessary to estimate the parameters in (4).

Assumption 1 (Best Response Condition) *If a_d is observed to be the strategy played by dealership d, then it must be the case that*

$$\max_{a_d \in A_d} \mathbb{E}[\pi_d(a_d, a_{-d})|\mathcal{I}_d] \geq \mathbb{E}[\pi_d(a'_d, a_{-d})|\mathcal{I}_d] \text{ for every } a'_d \in A_d \text{ and } d \in D.$$

Assumption 1 describes the common equilibrium assumption for this subgame. It says that if a vector of dealership choices is observed in the data, these actions are profit maximizing and hence no unilateral deviation could make them better off. This assumption is common to the literature. This best response condition delineates the principle on which the moment inequality conditions of this estimation strategy are built. That is, I add (and subtract) brands to the observed offerings in order to estimate bounds on the parameters.

The presence of structural disturbances ν^l_d, ν^b_d reconciles differences between the model predictions and observed actions. However, a problem of selection arises in that structural disturbances are not mean zero conditional on observed choices, even if they are unconditionally so. Pakes (2010) details several strategies to overcome this issue.

Location unobserved components are easy to control for given the separable functional form. Their disturbances are differenced out since I construct my moments by changing brands choices and keeping locations fixed. I introduce Assumption 2 in order to construct a way to circumvent the selection problem occurring with ν^b_d. In essence, I create counterfactual inequalities that hold no matter what decision retailers make in order to be able to use the unconditional expectation of ν^b_d. Since $\mathbb{E}[\nu^b_d] = 0$, these unconditional moments eliminate the selection effect. Let $a^{b-}_d = a_d \backslash \{b\}$, and $a^{b+}_d = a_d \cup \{b\}$.

Assumption 2 (Eventual (Un)Profitability) *Let $d, \tilde{d}$ be two observed dealerships with a_d and $a_{\tilde{d}}$ respectively, and suppose $b \in a_{\tilde{d}}$. Then, if dist$(d, \tilde{d}) < L$ there exists at least one $i_d \in \{0,1\}^{|-d|}$ with $a'_{-d} = i_d \cdot a_{-d} + (1 - i_d) \cdot a^{b-}_{-d}$ such that*

$$\mathbb{E}[\pi_d(a^{b+}_d, a'_{-d})|\mathcal{I}_d] \geq \mathbb{E}[\pi_d(a_d, a'_{-d})|\mathcal{I}_d].$$

Conversely, let $b \in a_d$, then there exists at least one $i_d \in \{0,1\}^{|-d|}$ with $a'_{-d} =$

$i_d \cdot a_{-d} + (1 - i_d) \cdot a_{-d}^{b_+}$ *such that*

$$\mathbb{E}[\pi_d(a_d^{b_-}, a'_{-d})|\mathcal{I}_d] \geq \mathbb{E}[\pi_d(a_d, a'_{-d})|\mathcal{I}_d].$$

There are two points to note about assumption 2. First, it specifies the area close to a dealer offering a specific brand (henceforth neighboring area). In these neighboring areas, the demand for a product of this brand is very similar for the original dealer as it would be for any neighboring dealer, should they also offer it. Logically, if a dealership is observed to be in a location, it means that there is enough demand around that area to sustain that dealership.

The spatial structure of the model implies that, in equilibrium, neighboring dealers normally tend to choose brands that their rivals are not choosing. This aspect is implied by the fact that, if two points of sale cannot relax their competition in the model through location choices, they are going to do so by their product offerings. This anti-coordination motive between dealers best responses could likely lead to a multiplicity of equilibria where the brand offerings are kept fixed, but what dealer offers which brand can be permuted.

It is useful to consider an illustrative example. In a completely isolated market with two dealerships and two brands (e.g., Renault and Seat), assume that the first dealership offers Renault and the second offers Seat. It is likely that, if the dealerships are similar in their observables, there is also another candidate equilibrium where the first dealership offers Seat and the second Renault. A dealership might not find it profitable to offer a brand (e.g. Seat) because there is another Seat dealer neighboring, but it might find it profitable were that competitor not offering Seat.

The second part of assumption 2 assures that there exist alternative profiles for which these profitable deviations can exist regardless of the unobservable ν_d^b within these neighboring areas. Assumption 2 basically states that, in these areas, any dealership could potentially deal for this brand profitably (unprofitably) if intrabrand retail competition is sufficiently relaxed (tightened). The radius of maximum geographic distance (L) for neighboring retailers is chosen small in order for this condition only to apply in areas where it is observed that there is enough demand for a dealership to offer the products of this brand.

Following the previous example, the assumption states that the neighboring dealership will surely find profitable to offer Seat if the rival dealer did not offer Seat and he was the unique dealer for that brand in a large area ($a_d^{b_-}$).

I use the two assumptions to construct moment conditions that do not suffer from selection. The strategy I employ consists of first creating a function that subtracts a given brand b from each dealership whenever it is possible, i.e., when that brand is offered. For dealers for which the brand is not offered, I create a multilateral deviation for which they could potentially offer b. This deviation is a perturbation of the equilibrium play because it subtracts brand b from all local competitors within a radius L[13] in order to make the choice of b attractive to these dealers. I finally select those observations that offer b profitably with a weight function and average over dealers to construct moment conditions. I repeat the same procedure in the opposite direction, i.e., adding a brand if possible and creating multilateral deviations if not, and for all brands.

Let $\Delta x\left(a_d, a'_d; a_{-d}\right)$ be defined as $x\left(a_d, a_{-d}\right) - x\left(a'_d, a_{-d}\right)$ for a function x and any $a_d, a'_d \in A_d$. I construct the function $\Delta r^u_b(a_d, a^{b-}_d, a_{-d})$ as

$$\Delta r^u_b(a_d, a^{b-}_d, a_{-d}) = \begin{cases} \mathbb{E}\left[\Delta \mathrm{VP}_d(a_d, a^{b-}_d; a_{-d})\right] - F_b - \mathbb{I}\left\{|a_d| \neq 2\right\} \cdot C_{MD}, & \text{if } b \in a_d, \\ \mathbb{E}\left[\Delta \mathrm{VP}_d(a^{b+}_d, a_d; a'_{-d})\right] - F_b - \mathbb{I}\{|a^{b+}_d| \neq 2\} \cdot C_{MD}, & \text{if } b \notin a_d, \end{cases} \tag{5}$$

where $\mathbb{I}\left\{|a_d| \neq 2\right\}$ is an indicator function equal to 1 if dealer d offers a quantity of brands (denoted as $|a_d|$) different from 2. These indicators multiply the multi-dealing costs, which are only relevant in the fixed costs function (4) whenever a dealership transitions from selling for one brand to selling for two brands, or vice versa.

For a dealership d, I define it to be in the neighborhood of b if there is another dealer d' selling b within a distance L from d. This definition is formalized by

$$\mathcal{N}^L_b = \left\{d \in D \mid \mathrm{dist}(d, d') < L \text{ for some } d' \in D \text{ such that } b \in a_{d'}\right\}.$$

With these neighborhoods, I construct the weight functions $g^1_d(b, a_d, a_{-d})$ and $g^2_d(b, a_d, a_{-d})$ with which I define two sets of B moment conditions. In particular, let g^1_b and g^2_b be

$$\begin{aligned} g^1_d(b, a_d, a_{-d}) = {} & \mathbb{I}\{b \in a_d\} \cdot \mathbb{I}\{|a_d| \neq 2\} + \mathbb{I}\{b \notin a_d\} \cdot \mathbb{I}\{d \in \mathcal{N}^L_b\} \cdot \mathbb{I}\{|a^{b+}_d| \neq 2\}, \quad \text{and} \\ g^2_d(b, a_d, a_{-d}) = {} & \mathbb{I}\{b \in a_d\} \cdot \mathbb{I}\{|a_d| = 2\} + \mathbb{I}\{b \notin a_d\} \cdot \mathbb{I}\{d \in \mathcal{N}^L_b\} \cdot \mathbb{I}\{|a^{b+}_d| = 2\}. \end{aligned} \tag{6}$$

[13] The results I report use L=15 kilometers, but they are robust to radius of 1, 5, 10, 20 and 30 kilometers.

The weight functions in (6) basically selects observations that lie within a $\mathcal{N}_b^L$ and divides them into two groups. The first one has all observations that do not transition from 2 to 1 brands in equation 5, and thus only carry information about the component F_b in the fixed costs. The second one contains the observations that transition from 2 to 1 brands when a brand is taken. These ones contain information about F_b and C_{MD}. Using (5) and (6), I construct the moment conditions

$$\begin{aligned} m_b^1 = & \quad |D|^{-1} \sum_{d \in D} g_d^1(b, a_d, a_{-d}) \Delta r_b^u \left(a_d, a_d^{b-}, a_{-d} \right) \geq 0, \quad \text{and} \\ m_b^2 = & \quad |D|^{-1} \sum_{d \in D} g_d^2(b, a_d, a_{-d}) \Delta r_b^u \left(a_d, a_d^{b-}, a_{-d} \right) \geq 0, \end{aligned} \tag{7}$$

for i.i.d. disturbances provided that the terms $|D|^{-1} \sum_{d \in D} g_d^1(b, a_d, a_{-d}) \cdot \nu_d^b$ and $|D|^{-1} \sum_{d \in D} g_d^2(b, a_d, a_{-d}) \cdot \nu_d^b$ vanish to 0 following the law of large numbers.

Similarly, I also define moments that determine the lower bounds for the parameters in a similar but opposite manner by defining function $\Delta r_b^l(a_d, a_d^{b+}, a_{-d})$, and weights $g_d^3(b, a_d, a_{-d})$ and $g_d^4(b, a_d, a_{-d})$ as

$$\Delta r_b^l(a_d, a_d^{b+}, a_{-d}) = \begin{cases} \mathbb{E}\left[\Delta \mathrm{VP}_d(a_d^{b-}, a_d; a'_{-d})\right] - F_b - \mathbb{I}\{|a_d^{b-}| > 1\} \cdot C_{MD}, & \text{if } b \in a_d, \\ \mathbb{E}\left[\Delta \mathrm{VP}_d(a_d, a_d^{b+}; a_{-d})\right] - F_b - \mathbb{I}\{|a_d| > 1\} \cdot C_{MD}, & \text{if } b \notin a_d, \end{cases} \tag{8}$$

$$\begin{aligned} g_d^3(b, a_d, a_{-d}) = & \quad \mathbb{I}\{b \in a_d\} \cdot \mathbb{I}\{d \notin \mathcal{N}_b^L\} \cdot \mathbb{I}\{|a_d^{b-}| > 1\} + \mathbb{I}\{b \notin a_d\} \cdot \mathbb{I}\{|a_d| > 1\}, \quad \text{and} \\ g_d^4(b, a_d, a_{-d}) = & \quad \mathbb{I}\{b \in a_d\} \cdot \mathbb{I}\{d \notin \mathcal{N}_b^L\} \cdot \mathbb{I}\{|a_d^{b-}| = 1\} + \mathbb{I}\{b \notin a_d\} \cdot \mathbb{I}\{|a_d| = 1\}. \end{aligned} \tag{9}$$

In this case, the critical inequality to identify the potential cost of multi-dealing is when adding a brand to an exclusive dealer. Converse to (6), the weight functions in (9) select all the observations that do not have $b \in a_d$ and add to them those that are unprofitable using Assumption 2. Using (8) and (9), I form the B moment inequalities

$$\begin{aligned} m_b^3 = & \quad |D|^{-1} \sum_{d \in D} g_d^3(b, a_d, a_{-d}) \Delta r_b^l(a_d, a_d^{b-}, a_{-d}) \geq 0, \quad \text{and} \\ m_b^4 = & \quad |D|^{-1} \sum_{d \in D} g_d^4(b, a_d, a_{-d}) \Delta r_b^l(a_d, a_d^{b-}, a_{-d}) \geq 0. \end{aligned} \tag{10}$$

These inequalities also hold provided that the terms $|D|^{-1} \sum_{d \in D} g_d^3(b, a_d, a_{-d}) \cdot \nu_d^b$ and $|D|^{-1} \sum_{d \in D} g_d^4(b, a_d, a_{-d}) \cdot \nu_d^b$ converge to 0.

The way moment inequalities 7 and 10 are constructed resembles Eizenberg (2014) on one side, and Pakes (2010) on the other. Eizenberg (2014) proposes an estimate that overcomes selection by replacing the missing values with conservative estimates of them. In his case, he uses the maximum and minimum of the differences in expected profits for the observed cases as an estimate to the missing upper and lower bounds. My approach shares some similarities in that it also approximates the value of unobserved choices and it does so conservatively. However, the big action set that dealers face in my game and the geographic component of supply and demand in the model do not allow this approach to yield any kind of informative result. This problem might be better described with an example. Following Eizenberg (2014), the fixed cost for a Nissan dealer will have 166 observed upper bounds and around 3,200 that are estimated to be the upper bound of these 166. Furthermore, this upper bound might come from a Nissan dealer in Barcelona or Madrid, which does not correspond to the possible expected revenues in less densely populated areas. The empirical content of such an estimate might come from below 5% of the observations.

Pakes (2010) discusses several ways to overcome selection. One these strategies uses unconditional averages due to inequalities that hold no matter what decision the agent has made[14]. My assumption is also formulated independent of agents own choices, but instead uses that of neighbors, which should not be inducing selection if these errors are independently distributed. Using this assumption, I can be conservative on which dealerships can eventually profitably offer products of a brand and still account for selection. While it is easy to justify that a Volkswagen dealer could eventually be profitable in some local geographic position in Barcelona, it is difficult to justify the same for an Infiniti dealer in Cáceres. Table 3 shows that the average consumer in that province needs to travel 148.44 kilometers to its closest point of sale of that brand. These traveling times imply that, so far, no dealership found offering Infiniti profitable there, even having no local rival.

In addition to these moments, I employ other moment inequalities based on Ho and Pakes (2014). I pair couples of observations d_1 and d_2 where d_1 subtracts brand b while d_2 adds it in order to form additional moments to identify C_{MD}. In order

[14]see assumption PC4b in Pakes (2010)

to avoid selection in these moments, I pair couples of equations in 5 and 8 to define

$$\Delta w(d_1, d_2, b) = \Delta r_b^u(a_{d_1}, a_{d_1}^{b-}, a_{-d_1}) + \Delta r_b^l(a_{d_2}, a_{d_2}^{b+}, a_{-d_2}), \tag{11}$$

which I combine with assumptions 1 and 2 and the weight functions defined before to form

$$\begin{aligned} m^5 = \quad & NM^{-1} \sum_{b\in B} \sum_{d_1\in D} \sum_{d_2\neq d_1} g_{d_1}^2(b, a_{d_1}, a_{-d_1}) g_{d_2}^3(b, a_{d_2}, a_{-d_2}) \Delta w(d_1, d_2, b) \geq 0, \quad \text{and} \\ m^6 = \quad & NM^{-1} \sum_{b\in B} \sum_{d_1\in D} \sum_{d_2\neq d_1} g_{d_1}^1(b, a_{d_1}, a_{-d_1}) g_{d_2}^4(b, a_{d_2}, a_{-d_2}) \Delta w(d_1, d_2, b) \geq 0, \end{aligned} \tag{12}$$

where NM denotes the total number of matches formed. Intuitively speaking, moments in (12) select observations that transition from multi- to exclusive dealer dropping one brand with some multi-dealer that adds that brand. The average of these pairings should be bigger than or equal to zero if we take into account equations 7 and 10. However, these new inequalities only depend on C_{MD} and add a further restriction by matching independent observations.

For the estimation of parameters $F = (F_1, ..., F_B, C_{MD})$, I define $m(F)$ to be the vector containing moments $\left[m_1^1, ..., m_B^1, ..., m^6\right]$. The identified set of parameters must satisfy each of these inequality or, equivalently, be a part of the space of parameters minimizing the objective function

$$\left[m(F)_{\_}\right]' \Sigma(F)^{-1} \left[m(F)_{\_}\right], \tag{13}$$

where Equation 13 is similar to the objective function in Chernozhukov et al. (2007). $m(F)_{\_}$ is a loss function that is different from zero whenever $m(F)$ is below 0, and it is 0 otherwise. $\Sigma(F)$ is the variance covariance matrix for the moments [15].

I use the method developed in Andrews and Soares (2010) to construct confidence sets that contain the true F_o in 95% of the cases. In practical terms, this methodology includes in the confidence set all vectors of parameters whose test statistic cannot reject the null hypothesis that these vectors are equivalent to the true parameter F_o. The acceptance and rejection regions are delimited by the 95% percentile of the distribution of test statistics from a large number of bootstrapped subsamples. In my application, I set the bootstrap subsamples to be one fourth of

[15] In this application, I assume all off-diagonal entries to be zero. i.e. I do not take correlation across moments into account.

the total sample (subsample size of around 836 observations), and the number of bootstrap repetitions at 10,000.

I perform a search for parameters in the confidence sets as a problem of constrained optimization where, starting from a value within the set, I look for the minimum and maximum values for each parameter independently subjected to the vector not being rejected by the test. Since the set of parameters is large, I cannot search for all parameters at the same time and I tackle this issue in two steps. In a first step, I pair moments including pairs of parameters (F_0, C_{MD}), (F_1, C_{MD}), ..., (F_B, C_{MD}) and perform a search for parameters independently. From the first step optimization, I collect the sets $\underline{C} = \{\underline{C}_0, ..., \underline{C}_B\}$ and $\overline{C} = \{\overline{C}_0, ..., \overline{C}_B\}$. Clearly, pairs of F_b and C_{MD} such that

$$C_{MD} \in \left[\sup \underline{C}, \inf \overline{C}\right] \tag{14}$$

were accepted for all optimization in the first step. In the second step, I search for parameters in the same pairs as in the first step, but with the bounds defined in (14) as additional constraints on the parameters.

4.4 Computing expected variable profits

In subsection 4.3, I laid out the estimation strategy for fixed costs. This procedure requires knowledge of the expected variable profits for each dealership. I simulated these profits for both the observed equilibrium and the perturbed strategy profiles using inferred margins, demand estimates and demographic data.

In the first step, I reorganize the dealerships according to the profile to be simulated. In the case of $\mathbb{E}[\mathrm{VP}(a)]$ no change is needed, while for unilateral perturbations $\mathbb{E}[\mathrm{VP}(a_d^{b-}, a_{-d})]$ or $\mathbb{E}[\mathrm{VP}(a_d^{b+}, a_{-d})]$ the change is simply subtracting or adding a point of sale[16] respectively. When simulating multilateral perturbations $\mathbb{E}[\mathrm{VP}(a_d^{b+}, a_{-d}^{b-})]$, I subtract brand b for all dealers within a circle of 15 kilometers around dealer d and add the brand to it.

After reorganizing dealers' offerings, I recalculate the distances from the simulated consumer locations to dealerships for each brand. Two additional assumptions are used in order to simulate individuals and their purchase decisions: (i) I assume

[16] These changes may also entail a change in exclusivity status along the process, which is considered.

that consumers are allocated uniformly within a municipality given a large set of simulated locations, and (ii) I use the recovered empirical distribution function of ξ, where these ξ are jointly distributed for all products across markets.

The simulation process for consumer purchases is as follows. First, I draw ξ_{jm} for each product in each market. Second, I draw locations and other demographics for each individual in every municipality within each market. For every simulated individual, I also draw disturbances ϵ_{ijmd} for each product from a Type 1 Extreme Value distribution. With these draws I assemble consumers' utility and compute the car purchases for each individual as the utility maximizing product that yields a utility higher than 0. This procedure is repeated a total of NS times.

Since all of these simulations are done in order to get the expected variable profits for some particular dealership, I can make some simplifications that reduce the computational expenses from these operations. First, I do not have to compute sales for municipalities where dealer d is not in any choice set. Furthermore, I do not need to simulate the purchases for products not sold in this dealership. It suffices to find one product (or outside option) yielding a higher utility than any of the products sold by this dealer in order to finish the simulation for an individual that does not buy from d. This practical shortcut is substantially less computationally intensive than finding the utility maximizer of the products sold by other dealers.

Finally, I calculate expected variable profits by multiplying inferred margins by the sales for each product, and average it over the number of simulations.

5 Results

5.1 Demand estimates and inferred margins

Table 4 presents demand estimates for different specifications. Column (1) is the baseline Logit without any random coefficients. Columns (2) and (3) add dealership characteristics. Column (4) includes all dealer variables (exclusivity and distance) and interactions of price with income. It is the specification used for the supply side estimates.

In line with intuition, the coefficients for distance and price are negative and significant across the different specifications. Consumers dislike paying more for their cars and traveling longer distances. The positive sign of the interaction term of income with price means that demand becomes less sensitive to price as income

Table 4: Estimates for the demand model

	(1) Logit	(2) RC Logit	(3) RC Logit	(4) RC Logit
Price	-2.232	-1.130	-1.163	-2.291
	(0.220)	(0.118)	(0.115)	(0.618)
Fuel Cons.	-0.344	-0.342	-0.332	-0.214
	(0.041)	(0.048)	(0.048)	(0.067)
HP / Weight	-0.113	-0.028	-0.020	0.070
	(0.037)	(0.039)	(0.039)	(0.074)
Size	1.276	1.291	1.325	2.066
	(0.118)	(0.132)	(0.130)	(0.429)
Cons.	-15.200	-13.548	-13.914	-17.706
	(0.978)	(1.074)	(1.051)	(2.461)
Distance		-0.556	-0.546	-0.353
		(0.060)	(0.060)	(0.110)
ED			0.200	-0.021
			(0.100)	(0.154)
Price $\times$ Income				0.073
				(0.023)
Origin f.e.	Yes	Yes	Yes	Yes
Province f.e.	Yes	Yes	Yes	Yes

increases. According to the estimates, for a consumer with an annual income of €20,000, a price increase of €1,000 for a car has a comparable effect to 6.07 kilometers of additional travel distance to the point of sale for this car. Exclusivity enters positively in specification (3), but loses significance when income random coefficients are added in column (4).

The rest of parameters have signs in line with what is expected: fuel consumption reduces the utility for the car while size has a positive sign. Horsepower over weight has a changing sign and it appears to be not significant in many of the specifications.

Figure 4 represents the distribution of own price elasticities for the different products and markets, while Table 5 shows which car models are pricing at the highest and lowest elasticity segments. Aside from luxurious cars, most of the local elasticities oscillate between -15% and -2.26%, with an average of -7.73% (median 6.30%) decrease in demand for a 1% price increase. Whereas demand estimates show that distance has a sizable effect on utility, most elasticity differences across products are driven by prices.

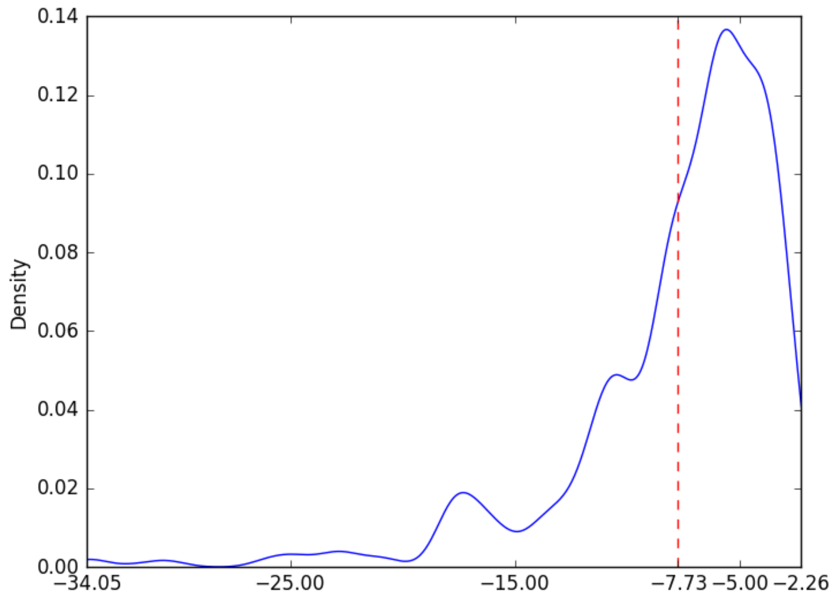

Figure 4: Distribution of elasticities at local markets

Table 5: Top 5 Highest and Lowest Elasticities

Brand	Model	Elasticity	Brand	Model	Elasticity
Highest 5 Elasticities			**Lowest 5 Elasticities**		
Land Rover	Range Rover	-33.78	Dacia	Dokker	-2.56
Porsche	Panamera	-30.70	Ford	Ka	-2.48
BMW	Serie 6	-25.75	Dacia	Logan	-2.33
Mercedes	Clase S	-24.62	Dacia	Sandero	-2.30
BMW	Serie 7	-23.05	Skoda	Citigo	-2.30
Median					
Volkswagen	Beetle	-6.39			
Mini	Paceman	-6.30			

Figure 5 shows the inferred marginal costs and margins for all cars. On the horizontal axis there are all car models ordered by price. The blue bars are all the inferred marginal costs, whereas the red area on top of them are the margins. The relatively small magnitude of the utility estimates for all car characteristics except price imply a relatively constant markup in absolute terms, or in other words, a markup that proportionately reduces as car list prices increase. The average markup is around €3,130.

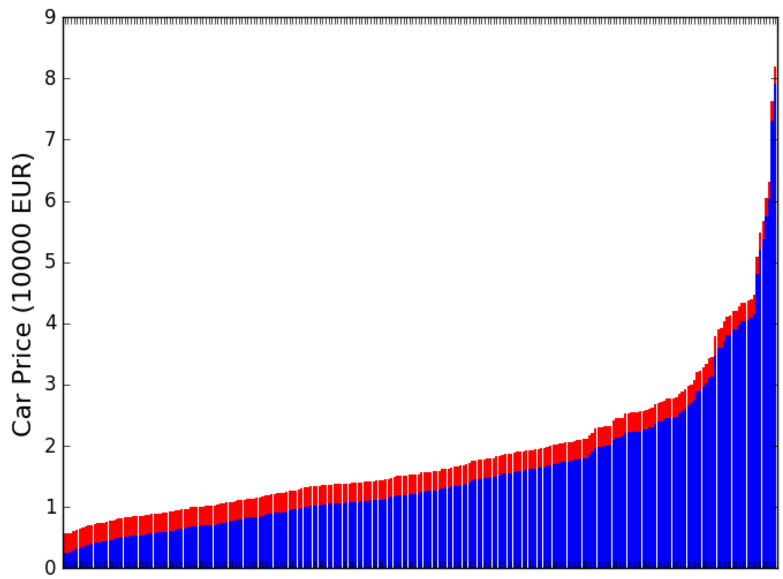

Figure 5: Distribution of dealer expected margin by car

5.2 Fixed Costs' estimates

Table 6 reports the fixed costs' estimates for the model. Columns Upper and Lower report the bounds corresponding to the 95% confidence set for these parameters. The confidence sets for the different brand related cost components show a very large difference across brands. In all cases, these bounds are relatively large, but in most cases they are bounded away from zero. Brand cost can be separated into three groups: brands like Alfa Romeo, Ssangyong, Subaru and Suzuki, whose costs are rather small (below €1 million) and their intervals not very wide; middle brands, with wider parameter intervals and higher upper bounds (e.g. Ford, Hyundai, Opel, Nissan, below €7 million). Finally, the third group is composed by popular brands and higher class manufacturers (e.g. Audi, BMW, Peugeot, Volkswagen, Mercedes) and their fixed costs can exceed the €7 million.

The parameter for the cost of multi-dealing is negative, but not significantly different from zero. This result rejects the hypothesis that manufacturers use exclusive contracts to deter other brands from their points of sale by raising their costs to offer products from other manufacturers. It is dealerships who choose to deal exclusively in their trade-off between differentiating from rival dealerships and offering more products.

Table 6: Fixed costs' estimates (in €10,000)

	Lower	Upper		Lower	Upper
Alfa Romeo	10.3750	112.3477	Mini	31.0048	233.3680
Audi	248.9550	967.3268	Mitsubishi	57.6487	253.7285
BMW	240.7112	1427.9737	Nissan	105.1268	453.3960
Citroen	158.1234	409.9916	Opel	112.3373	473.0761
Fiat	31.8118	199.4893	Peugeot	232.2222	982.2118
Ford	109.0567	382.4699	Porsche	1049.9087	5929.1565
Honda	35.2027	221.0466	Renault	314.3647	950.5661
Hyundai	110.0356	515.8113	Seat	243.6576	840.3873
Infiniti	43.2850	299.4948	Skoda	63.6640	367.2605
Jaguar	90.1613	601.0930	Smart	-9.9428	60.6519
Jeep	37.5273	266.4081	SsangYong	17.8420	118.3771
KIA	114.6440	580.4779	Subaru	1.9487	78.3876
Land Rover	116.1576	623.2758	Suzuki	15.3060	98.6510
Lexus	20.3505	35.8169	Toyota	117.6452	443.0447
Mazda	86.8278	440.2257	Volkswagen	213.7001	746.6376
Mercedes	248.3775	906.8245	Volvo	92.0989	499.0274
			Multi-Dealing	-62.5631	1.1340

The results in Table 6 are robust to different radius for the neighboring dealers and they also do not seem to vary substantially under alternative algorithms. A natural robustness check is to perform in the first step the parameter search jointly for several brands, instead of doing these searches parallely, in order to account for potential correlations remaining unaccounted by the estimates. I tested several of these combinations without finding qualitative differences.

Another potential concern with these estimates arises from the side of the specification. It is plausible to think that some manufacturers might have more power or more interest in raising costs than others, and therefore the parameter C_{MD} captures an average cost of multi-dealing across brands. Accommodating this heterogeneity across groups of brands is possible, although it would entail a reformulation of the moment conditions in (7) and (10) so as not to count the C_{MD} twice.

6 Concluding Remarks

This paper has added to the debate on the potential anticompetitive effects of exclusive dealing from an empirical perspective. I estimated a structural model that combined demand and supply, and incorporated the manifold effects of exclusivity. Exclusive contracts can boost demand through increased promotional effort and better dealer service, but it can also be used to raise the fixed costs of distributing for rival brands, and deter competition from other brands within dealers. Moreover, exclusive dealing can be used by dealers as a means to differentiate from their local competitors. In order to capture this differentiation motives, my model extended the literature allowing for dealers to choose the brands that they sell endogenously.

For estimating this model, I assembled a novel dataset that combines information on car sales and car retailers from Spain. These data contain sales for a large number of car models in the market, and the specific location and brand offerings for all the dealers for the 32 most popular car brands in the country.

Furthermore, my estimation of the supply side contributed to the recent literature using moment inequalities to estimate fixed costs. In particular, I proposed a way how to circumvent the potential selection on unobservables that might happen when using equilibrium choices to estimate parameters. My approach is well-suited to spatial markets and it can be used in problems where the agents have large action spaces. It is based on using counterfactual dealer offerings for those observations that are selected out in order to account for them.

The results of my estimation suggest that (i) there are no particular effects in utility derived from exclusive dealing, (ii) there are no sizable costs additionally when multi-dealing. These two findings lead to conclude that (iii) spatial competition among downstream competitors creates the conditions for which they take up exclusive dealing in order to differentiate from their rivals in the products they offer. These conclusions are in line with the Chicago school view that exclusive contracts are not anticompetitive because they are not used to exclude rivals, and have large implications for regulatory policy in retail markets.

Two questions for future research are whether there actually exist difficulties for smaller manufacturers to access to retailing points, even if driven by competitive forces, and, if so, what is the impact of this exclusion on welfare. The model and estimation of this paper laid the foundations to answering questions of this kind, and more generally to analyze the effects of vertical restraints on market structure and product variety.

References

Philippe Aghion and Patrick Bolton. Contracts as a barrier to entry. *American Economic Review*, 77(3):388–401, 1987.

Paulo Albuquerque and Bart J. Bronnenberg. Measuring the impact of negative demand shocks on car dealer networks. *Marketing Science*, 31(1):4–23, 2012.

Donald W. K. Andrews and Gustavo Soares. Inference for parameters defined by moment inequalities using generalized moment selection. *Econometrica*, 78(1): 119–157, 2010.

John Asker. Diagnosing foreclosure due to exclusive dealing. *Journal of Industrial Economics*, 64(3):375–410, 2016.

Itai Ater. Vertical foreclosure using exclusivity clauses: Evidence from shopping malls. *Journal of Economics & Management Strategy*, 24(3):620–642, 2015.

B. Douglas Bernheim and Michael D. Whinston. Exclusive dealing. *Journal of Political Economy*, 106(1):64–103, 1998.

Steven Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, pages 242–262, 1994.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica: Journal of the Econometric Society*, pages 841–890, 1995.

Steven T. Berry and Philip Haile. Nonparametric identification of multinomial choice demand models with heterogeneous consumers. *National Bureau of Economic Research*, w15276, 2009.

Steven T. Berry and Philip Haile. Identification in differentiated products markets using market level data. *Econometrica*, 82(5):1749–1797, 2014.

Sofia Berto Villas-Boas. Vertical relationships between manufacturers and retailers: Inference with limited data. *Review of Economic Studies*, 74(2):625–652, 2007.

David Besanko and Martin K. Perry. Equilibrium incentives for exclusive dealing in a differentiated products oligopoly. *RAND Journal of Economics*, pages 646–667, 1993.

David Besanko and Martin K. Perry. Exclusive dealing in a spatial model of retail competition. *International Journal of Industrial Organization*, 12, 1994.

Robert H Bork. *The Antitrust Paradox: A Policy at War with Itself*. 1978.

Randy Brenkers and Frank Verboven. Liberalizing a distribution system: The european car market. *Journal of the European Economic Association*, 4(1):216–251, 2006.

Giacomo Calzolari and Vincenzo Denicolò. Exclusive contracts and market dominance. *American Economic Review*, 105(11):3321–3351, 2015.

Chia-Wen Chen. Estimating the foreclosure effect of exclusive dealing: Evidence from the entry of specialty beer producers. *International Journal of Industrial Organization*, 37:47–64, 2014.

Victor Chernozhukov, Han Hong, and Elie Tamer. Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284, 2007.

Federico Ciliberto and Elie Tamer. Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6):1791 – 1828, 2009.

Michaela Draganska, Michael Mazzeo, and Katja Seim. Beyond plain vanilla: Modeling joint product assortment and pricing decisions. *Quantitative Marketing and Economics*, 7(2):105 – 146, 2009.

Alon Eizenberg. Upstream innovation and product variety in the us home pc market. *Review of Economic Studies*, 81(3):1003–1045, 2014.

Alon Eizenberg, Andras Pechy, and Michelle Sovinsky. The dynamics of technology adoption and vertical restraints: An empirical analysis. 2017.

Ying Fan. Ownership consolidation and product characteristics: A study of the us daily newspaper market. *American Economic Review*, 103(5):1598 – 1628, 2013.

Chiara Fumagalli and Massimo Motta. Exclusive dealing and entry, when buyers compete. *The American Economic Review*, 96(3):785–795, 2006.

Matthew Gentzkow and Jesse M. Shapiro. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71, 2010.

Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982.

Katherine Ho and Ariel Pakes. Hospital choices, hospital prices, and financial incentives to physicians. *American Economic Review*, 104(12):3841–3884, 2014.

Thomas J. Holmes. The diffusion of wal-mart and economies of density. *Econometrica*, 79(1):253 – 302, 2011.

Jean-François Houde, Peter Newberry, and Katja Seim. Economies of density in e-commerce: A study of amazon's fulfillment center network. *National Bureau of Economic Research*, 2017.

Y. Joseph Lin. The dampening-of-competition effect of exclusive dealing. *Journal of Industrial Economics*, 39:209–223, 1990.

Eduardo Morales, Gloria Sheu, and Andrés Zahler. Extended gravity. *NBER Working Paper*, 21351, 2015.

Charles Murry and Yiyi Zhou. Consumer search and automobile dealer co-location. 2017.

Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342, 2001.

Laura Nurski and Frank Verboven. Exclusive dealing as a barrier to entry? evidence from automobiles. *The Review of Economic Studies*, 83(3):1156–1188, 2016.

Daniel P. O'Brien and Greg Shaffer. On the dampening-of-competition effect of exclusive dealing. *Journal of Industrial Economics*, 41:215–221, 1993.

Ariel Pakes. Alternative models for moments inequalities. *Econometrica*, 78(6): 1783–1822, 2010.

Ariel Pakes, Jack Porter, Katherine Ho, and Joy Ishii. Moment inequalities and their application. *Econometrica*, 83(1):315–334, 2015.

Richard A Posner. *Antitrust law.* University of Chicago Press, 1976.

Eric B. Rasmussen, J. Mark Ramseyer, and John S. Wiley. Naked exclusion. *American Economic Review*, 81(5):1137–1145, 1991.

Patrick Rey and Joseph E. Stiglitz. Vertical restraints and producers' competition. *European Economic Review*, 32:561–568, 1988.

Patrick Rey and Joseph E. Stiglitz. The role of exclusive territories in producers' competition. *RAND Journal of Economics*, 26:431 – 451, 1995.

Tim R. Sass. The competitive effects of exclusive dealing: Evidence from the us beer industry. *International Journal of Industrial Organization*, 23(3-4):203 – 225, 2005.

Ilya R. Segal and Michael D. Whinston. Naked exclusion and buyer coordination. *Harvard University, Institute of Economic Research Working Papers: No. 1780*, 1996.

John Simpson and Abraham L. Wickelgreen. Naked exclusion, efficient breach, and downstream competition. *American Economic Review*, 97(4):1305–1320, 2007.

Karunakaran Sudhir. Structural analysis of manufacturer pricing in the presence of a strategic retailer. *Marketing Science*, 20(3):244–264, 2001.

Joel Waldfogel. Preference externalities: An empirical study of who benefits whom in differentiated-product markets. *RAND Journal of Economics*, 34(3):557–568, 2003.

Thomas G. Wollmann. Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, 208(6):1364 – 1406, 2018.