

Spatial Rents, Garage Location, and Competition in the London Bus Market

Marleen Marra* Florian Oswald*

November 14, 2023

Abstract

This paper estimates a structural model of the market for public bus transportation services in London, which links the value of a bus garage to expected profits from bus route procurement auctions. The model is used to derive the spatial rents associated with owning a garage with a certain centrality in the network of garages and routes, i.e., the benefits of having low transportation costs, being far removed from competitors, and having multiple garages clustered together. We exploit the unique features of the London bus industry to recover these rents using a standard discrete choice estimator, reflecting the highly complex location choice problem of heterogeneous multi-plant firms competing in an economy with spatial rents. Even a parsimonious specification is shown to capture remarkably well the observed changes in the garage-operator network since the privatization of this industry in 1994, as constructed from archival data and bus spotter sites. Counterfactual simulations reveal an efficiency loss of £18m to £27m per year, of 6.5%-9.8% of the total procurement cost of providing public bus transportation in London, resulting from operators holding out the sale of their garages. The degree to which profitable garage transactions are “blocked” is lower in periods when more firms enter the market, consistent with anti-competitive market-sharing behavior.

JEL codes: D44, L41, L92, R41

Keywords: Firm Location Choice, Procurement Auctions, Public Transport, Urban

*Department of Economics, Sciences Po, Paris, France. We thank Martin Bretschneider, Xinyu Dai and Ian Hsieh for excellent research assistance. We also thank participants at the TSE and UCL Econometrics of Games, Matching and Networks conference, EARIE, as well as Estelle Cantillon, Gabriel Ahlfeldt, Milena Almagro, Pierre-Philippe Combes, Francesco Decarolis, Pierre Dubois, Michele Fioretti, and Emeric Henry for comments. We are very grateful to Les Stitson from the London Omnibus Traction Society for compiling a dataset for this study based on their archives and for generously answering numerous questions about the history of garage ownership in London. Thanks also to Robert Munster, the owner of LondonBusRoutes.net, for answering questions and sharing data on bus route characteristics with us. We can be reached via e-mail at marleen.marra@sciencespo.fr and florian.oswald@sciencespo.fr.

1 Introduction

In many industries, the decision of where to locate matters for a firm's profitability. Minimizing the spatial frictions between firms and consumers is also important in an urban context, where the congestion and pollution externalities of transporting goods and services across space are especially large. Important questions in this sphere involve the extent of spatial competition. For instance: What if firms can protect their local monopoly rents by mutually agreeing to stay out of each others' vicinity, thereby avoiding competition and affecting the transport needs of cities? As the firm location choice problem quickly becomes intractable, especially when economies of density generate spillovers across locations (see, e.g., [Oberfield et al. \(Forthcoming\)](#)), the literature provides scant empirical evidence to address such questions.

This paper investigates the role of spatial rents within the market for public bus transportation services in London. Firms participating in this market maintain their own bus fleets from proprietary bus garages, and bid for contracts to operate bus routes in the Greater London area. Given the costs associated with driving empty buses between garages and bus routes, the firms face a tradeoff similar to those in classical location choice models (e.g., [Hotelling \(1929\)](#), [Salop \(1979\)](#)). We will show that in this market, firms seek to optimally locate their bus garages close to bus routes, but far away from competitors, who would reduce their local rents in future procurement auctions.

This industry is particularly suitable for the analysis of spatial rents in a procurement setting. The winning bidder receives a lump sum payment (i.e. their bid) to provide the bus service, whereas all bus fares are collected by the local authority Transport for London (TfL). Moreover, the service is tightly regulated — TfL specifies the exact route, frequency of operation, type and make of vehicle with precise specifications, maintenance and cleanliness requirements, and duration of the contract. Hence, the nature of this market severely restricts the set of choice variables for participating firms. One important margin which bus operators can compete on, however, is the location of their proprietary bus garages, which determine their efficiency in providing the tightly regulated service.

Specifically, operators are less competitive on routes far away from their garages, due to what the industry refers to as *dead mile costs* of transporting empty vehicles between the startpoints of the route and the garage. This generates local monopoly rents as the equilibrium mark-up is also higher when the bus garages of competing firms are further away from the route. In our companion paper ([Marra and Oswald \(2023\)](#)) we document that dead miles are substantial in this market and that strategic incentives contribute to the problem.¹ In this paper, we further

¹We find that supplying 100 contractable bus route minutes requires, on average, about 13 minutes of dead mile driving time and we confirm that this affects procurement outcomes, also consistent with results in [Cantillon and Pesendorfer \(2006, 2007\)](#). It turns out that stationing buses only in the operator's own garages — which we refer to as garage ownership frictions — are partly responsible for the dead mile bus movements in the city. Removing these frictions by letting a systems operator assign buses to garages, which corresponds to unbundling of the garage

assess the role of the network of garages in the procurement of London bus routes, using a structural model that also accounts for economies of density and endogenous location choice.

To do so, we build on the dataset of the garage ownership history in the London bus market since privatization in 1994 introduced in [Marra and Oswald \(2023\)](#). For each garage, we not only know ownership over time, but also its location within the global route network, as well as within the constantly evolving network of garages for each operator at each point in time. A key observation is that the garage-route network reveals substantial clustering of bus garages by operators in certain parts of the city. We explore whether, in addition to cost-based explanations, firms holding out on the sale of their garages is partly responsible for this observed clustering. As such, we use this particular setting to quantify a prevalent economy-wide problem of assets not reallocating to their most efficient uses due to their owners enjoying some degree of market power from them (see [Posner and Weyl \(2017\)](#)).

A two-stage structural model links the value of a garage to revenues from bus service contracts that operators could service from it. Spatial characteristics of expected revenues, including economies of density and local monopoly rents, are derived from the second (procurement auction) stage. The first stage uses those characteristics in a model of garage location choice. We show that in this setting, garage transaction prices do not need to be observed in order to estimate the value of the garage to an operator, and how this value depends on the spatial location of the garage in the network of bus garages and routes.

Furthermore, we propose an estimation strategy that in an innovative way casts the matching of operators to garages into a standard discrete choice problem. Our approach embeds a set of highly nonlinear measures of drive-time distances of garages to several points of interest, like for instance other garages (own or competitor's), route start points, route end points and route intersection points. We find that our parsimonious model with a low-dimensional representations of network centrality measures and estimated on only 192 garage ownership changes fits the data remarkably well.

Our estimates shed light on the relative importance of economies of density – or the *agglomeration benefits* of having multiple garages clustered closely together – and local monopoly rents. For instance, we find that an operator would be indifferent between having an additional own garage 4 minutes away (representing agglomeration benefits) or having one fewer competing operator's garage within a 10-minute drive (local monopoly rents). We also capture persistence in the garage-operator network by including the number of years the incumbent has been in the garage before selling it. The estimated incumbency benefit is substantial from the onset and decreases fast over time, which reflects that firms own garages for about 3 years, on average. Interpreting the incumbency benefit relative to the local monopoly rents, the incumbent operator is estimated to have a benefit over the other operators that would only be offset with four (additional) competing garages within a 10 minute drive. After one year, the incumbency network from the competitive part of the market, reduces dead miles by 14 percent.

benefit is already offset by only 1.5 competing garages.

With the help of our structural estimates we investigate whether garages transact sufficiently, or whether operators hold out the sale of their garages to some extent. We estimate that in 11.3 percent of cases, another operator has a higher latent utility than the incumbent. This is true conditional on a sizeable incumbency benefit, estimated economies of density, and other estimated cost factors including variation in the distance of the garage to the other garages of the different operators. Operator profits in this market are directly linked to revenues from route auctions, allowing us to compute the total efficiency loss from the private ownership of garages and the unrealized profitable ownership changes identified with the counterfactual simulations. The efficiency loss, or approximate loss to the UK taxpayer, is estimated to be about £17.7 - £26.6 million per year when firms include a mark-up of 10% to 15% in their route auction bid (as estimated in [Cantillon and Pesendorfer \(2007\)](#)). This is a sizeable efficiency loss equal to 6.5% – 9.8% of the total costs to procure bus routes in London.

We use the model finally to assess potentially non competitive behaviour in the London bus market. This line of inquiry is motivated by the striking patterns of spatial segregation among operators in the London Bus market. Moreover, the UK Competition Commission fined several bus operators for collusive market sharing behavior outside of London in 2011 and explicitly expressed concern about the London bus market.² Our approach is based on a simple but powerful idea: a necessary condition for collusion is the triggering of a punishment when a firm deviates to the competitive equilibrium. Hence, collusion must result in otherwise profitable garage entries being blocked due to the expected punishment being larger than the expected gains from entry.³ Our structural estimates of the garage-operator utility function are needed to separate a collusive interpretation of the observed spatial segregation from clustering due to agglomeration benefits in operating routes from multiple near-by garages. We assess the connection between the estimated punishments and factors that increase the likelihood of collusion within a market according to [Ivaldi et al. \(2007\)](#).

To do so, we explore all cases when the incumbent of a garage has a lower latent utility for that garage than another operator. As explained above, the difference in the utilities of each incumbent-operator pair in that set can be understood to capture a minimum expected “punishment” to the operator if he would obtain the garage. We find that the implied punishment value correlates strongly and inversely with entry. In terms of magnitudes, each new operator

²Specifically, the UK Competition Commission uncovered in 2011 that bus operators in various local markets engaged in a practice of avoiding competition with each other in designated “core territories,” resulting in a division of the market along geographical lines. This behavior was estimated to impose an annual cost ranging from £115 million to £305 million on UK consumers and taxpayers. While the investigation excluded London and Northern Ireland, the report expressed concerns about the possibility of broader geographic market segregation beyond the cases identified.

³The specific form of punishment is not specified but should be understood as a means of enforcing compliance, potentially through future route auctions or through interactions in other markets, infringements on the perpetrator’s own core territory, or other measures resembling a “price war”, as in e.g. [Green and Porter \(1984\)](#) and [Rotemberg and Saloner \(1986\)](#).

competing for bus routes in London decreases the probability of the highest-utility operator not being in the relevant garage by about 6 percentage points. While not conclusive, these results are consistent with the notion that markets are more collusion-prone in stable periods without entry and when the proceeds from collusion can be split among fewer firms. These results highlight a useful application of the estimated garage-utility model with spatial rents, especially since market-sharing collusion is not necessarily identifiable in bid data.⁴

The rest of the paper proceeds as follows. The next section places the paper in the literature. Section 3 describes the London bus market and garage-route network dataset used in this study, and section 4 provides descriptive evidence about the importance of the garage network. A structural model of garage location choice is provided in section 5, which also describes the empirical strategy to recover model parameters. Estimation results are presented in section 6. Section 7 uses the structural parameters to evaluate the efficiency loss of private ownership of garages in this market, and interprets the results through the lens of a collusive market-sharing equilibrium. Section 8 concludes.

2 Relation to the literature

This paper can be placed at the intersection of the literatures on firm entry (e.g., [Bresnahan and Reiss \(1990, 1991\)](#) and [Berry \(1992\)](#)) and collusion detection (e.g., [Porter and Zona \(1993, 1999\)](#), [Bajari and Ye \(2003\)](#), and [Kawai and Nakabayashi \(2022\)](#)) in IO, location choice (e.g., [Head and Mayer \(2004\)](#), [Combes and Gobillon \(2015\)](#)) and optimal transport networks (e.g., [Fajgelbaum and Schaal \(2020\)](#)) in trade, and urban transportation (e.g., [Balboni \(2019\)](#), [Gaduh et al. \(2022\)](#), and [De Palma et al. \(2011\)](#)). Some additional details are provided below, to highlight the contributions of the paper to the existing body of research.

Within the literature on firm entry, models with (product or geographical) location choice are especially related (e.g., [Mazzeo \(2002\)](#), [Seim \(2006\)](#), [Jia \(2008\)](#), [Holmes \(2011\)](#), [Ellickson, Houghton and Timmins \(2013\)](#), and [Oberfield et al. \(Forthcoming\)](#)). As in [Seim \(2006\)](#), the observed location choices of firms are used to recover structural elements of the location-firm utility function. However, a distinguishing feature of the bus garage setting is that only one firm can be at a location at each point in time. Our model is therefore built to capture the question *which firm has the highest utility for each location* rather than *where does each firm locate*. We refer to this inverted choice problem as *Garage chooses Operator* and describe under which assumptions the inversion can be made in section 5.2.3. Combined with having data on garage ownership changes over time, our approach allows for an off-the-shelf discrete choice estimation strategy.

⁴Collusion is not identifiable in bid data if operators bid competitively given the realised garage network. Indeed, [Waterson and Xie \(2019\)](#) investigate procurement data from the London bus market between 2003 and 2015 and find little evidence for collusive bidding, despite the fact that some market features make it prone to collusion (see 7.2.1).

Another difference with the firm entry literature is that our reduced form utility function is derived from a structural procurement auction model, based on the close link between the value of a garage and the expected profits derived from routes operated from these garages.⁵ Oberfield et al. (Forthcoming) propose a tractable limit case of the heterogeneous firm location choice problem with spillovers, and show that its predictions are consistent with data of firm location choice across industries in the US.

Our assessment of anti-competitive behavior relies on the implications of market-sharing for firm location choice, rather than bids themselves.⁶ The method relies on the absence of profitable actions being taken and does not require a *smoking gun* to separate firms into likely-colluders and likely-competitors. While our counterfactual results are based on necessary but not sufficient conditions for collusion following Ivaldi et al. (2007), we derive implied punishment values and find that they relate to factors known to increase the likelihood of collusion within a market. Similarly, Porter (1983), Ellison (1994), and Ishii (2009) examine patterns of price wars corresponding to predictions from repeated games, based on theory in Green and Porter (1984) and Rotemberg and Saloner (1986). Looking into a form of non-price collusion, this aspect of the paper relates to Sullivan (2017) who studies collusive (product) positioning. Our theoretical framework is consistent with the collusive network formation model in Belleflamme and Bloch (2004).

Hence, we consider a combination of spatial factors considered relevant in these fields; *agglomeration benefits* or *economies of density* as for instance in Head and Mayer (2004) (foreign investment) and Holmes (2011) (retail), spatial competition with *local monopoly rents* as in Miller and Osborne (2014) (cement) and Houde (2012) (gasoline), and *transportation costs* in procurement as in Olivares et al. (2012) (school meals) and Cantillon and Pesendorfer (2006, 2007) (bus transport).⁷ Collectively, we refer to these factors as *spatial rents* and confirm the importance of the firm's location in all three respects for the London bus market.⁸

⁵As such, we quantify economies of density resulting from garages clustering near each other, similar to clustering of retail stores as in e.g. Jia (2008) and Holmes (2011) (who also employs a continuous geography rather than counting stores in lumpy markets), but without having to solve for the market equilibrium.

⁶Chassang and Ortner (2019) also departs from the more standard analysis of bid patterns, instead exploiting the cartel's ability to implement effective punishments.

⁷As in Miller and Osborne (2014), we study a setting with geographic price discrimination and local monopoly rents, but in our setting it is the firm who pays the transportation costs and firms compete for contracts in procurement auctions. Cantillon and Pesendorfer (2006, 2007) study combinatorial auctions using London bus route procurement data. While they focus on the bidding stage of the model, our analysis centres on the garage network formation stage to recover the parameters of the operator's profit function based on observed garage network changes.

⁸We estimate these spatial rents based on firm location choices, holding the transport network itself fixed, rather than seeking to find an optimal network based on estimated consumer preferences as in the (optimal) transportation literature (e.g., Gaduh et al. (2022)).

3 Industry and data description

London Bus Services Ltd (London Buses) is part of Surface Transport within *Transport for London (TfL)*, in charge of delivering the Mayor’s Transport Strategy. This involves transporting over six million passengers on 7,700 scheduled buses and 675 different routes stretching over 490 million kilometers of road. Procuring bus operation services from private operators costs TfL on average £273m annually over the period 2003–2019, paid for by UK taxpayers. The market has been organized in this way since it was privatized along geographic lines in 1994. Today, six large (international) transportation firms and a few smaller operators compete for tenders.

For a comprehensive description of this market and our data we refer to our companion paper [Marra and Oswald \(2023\)](#). That paper also documents that the distance of an operator’s garage to the bus route in question has an economically meaningful effect on firms’ bidding behavior. This is partly so because of the tight regulation by TfL, who specify virtually all remaining operational details of the service, including the frequency, exact route and stops, vehicle make, type and other requirements, and number of buses needed for each route. While the costs associated with these specifications are therefore the same across operators, it remains at their discretion how to get the required bus fleet to the route in the most efficient way. Operators own their own garages, and the location of these garages *vis-a-vis* the route is crucial.⁹

For our analysis, we construct a comprehensive dataset of the operator and bus garage network in London between 1994–2019. The London Omnibus Traction Society provided us with a historic dataset based on their archive of bus schedules, matching the majority of routes tendered since 2005 to the garage that it was operated from, and containing information about openings and closings of garages. We complement this data with garage location information from the privately-run London Bus Routes website, all bus stops of each route and their locations from the TfL Open Data Initiative, and manual searches of garage addresses using the Google Maps API.

Our structural estimates are based on *movers*, e.g. changes to the ownership of garages observed over time. The estimation sample is based on garage-operator network changes taking place after 1994, thus excluding the first year when the market was being privatized and when the garage network was being formed. We start with 114 garages and exclude five garages with incomplete information about the ownership history between 1995–2019, seven garages that were shared at some point in time, and one temporary garage. Finally, 10 of the remaining 101 garages never changed ownership in our data, which gives us 91 unique garages for estimation.

In total, there are 192 instances in this sample where a garage changes ownership. Table 1 lists the type of ownership change that took place. In 27 percent of the times that the garage-

⁹In what follows, we refer to the firm operating buses from the garage as the garage *owner* —regardless of whether the firm has a freehold or long leasehold contract— and the data excludes the few cases where a garage is shared between two operators.

Table 1: Summary of all garage network changes 1995-2019

Event	Count	Share
New entry	91	0.474
Onwership change	51	0.266
Temporarily vacant	17	0.088
Closed down	33	0.172
Total	192	1

operator network changes, a garage is acquired from another operator. Most of the times (91) the garage network changes because an operator occupies a garage for the first time. In the remaining cases the network changes because a garage is closed down (31 times) or temporarily not used to operate bus routes (17 times). We interpret these latter cases as the garage being bought from or sold to *the market*, as described further in section 5.2.3.

Furthermore, the analysis presented in this paper is based on a unit of time that is a so-called *change date*, which we define as a date at which we observe any change in the garage-operator ownership structure. There are 99 change dates in our data from 1995 to 2019, with on average 1.94 ownership changes per change date. In 80% of cases there is a single operator who is active on a given change date (see figure 1a). The vast majority of network changes also involves just one garage per time (see figure 1b).

To best capture spatial rents associated with the geographic location of garages relative to routes and other garages, distances are measured as the exact drive times on the London road network.¹⁰ The average drive time to the two endpoints of a route is selected as primary dead miles variable. This measure reflects that a route needs to be serviced from both directions, so being near only one endpoint is not necessarily optimal.

4 The importance of the garage network: descriptive evidence

This section provides new insights into the garage-operator network and its relation to bidding in route auctions.

4.1 Spatial clustering of garages

Figure 2 shows a snapshot of our dataset in 2019, which links routes to the garage where they were operated from. A key observation in our paper is the spatial clustering of bus garages.

¹⁰Simple straight-line distance between points would miss a great deal of nonlinearities arising from geography and, ultimately, true transportation costs. We therefore use the [Open Source Routing Machine \(OSRM\)](#) to compute optimal drive times between all garages and all other garages and all 55,578 bus stops in London. Besides the distance to the endpoints of the routes, we also consider the distance to the closest stopover point, where at least 5 routes intersect. See [Marra and Oswald \(2023\)](#) for further details.




Figure 1: Garage network changes in sample

Notes: Histograms show the number of operators entering a garage (a) and the number of garages entered by an operator (b) on the same date. Based on all times when operators opened a new garage (the 91 “New entry” events in Table 1) or purchased an existing garage from another operator (the 51 “Ownership change” events in Table 1) between 1995-2019.

Specifically, the figure shows that garages of a certain operator group are concentrated in certain areas of London, rather than being uniformly spread out across the city. The structural model and counterfactual simulations are designed to understand the factors driving this observed spatial clustering. The model is based on changes in the garage ownership network over time. Figure 3 summarizes the number of times that a garage has changed owners in the data. Garages change ownership between 1-7 times, and no obvious spatial pattern can be detected as to which garages change ownership more frequently.

4.2 Link between garage network and revenue from route operation

In our companion paper (Marra and Oswald (2023)), we use our linked route-garage network dataset to document that the location of a garage vis-a-vis the route is economically important. Specifically, we regress the cost per mile of the winning operator on route characteristics, the number of bidders, and our dead miles metric. An additional minute of (start-stop) dead mile driving time increases the winning bid by £19,000 and £88,000, depending on which auction-

Bus Routes and Garages

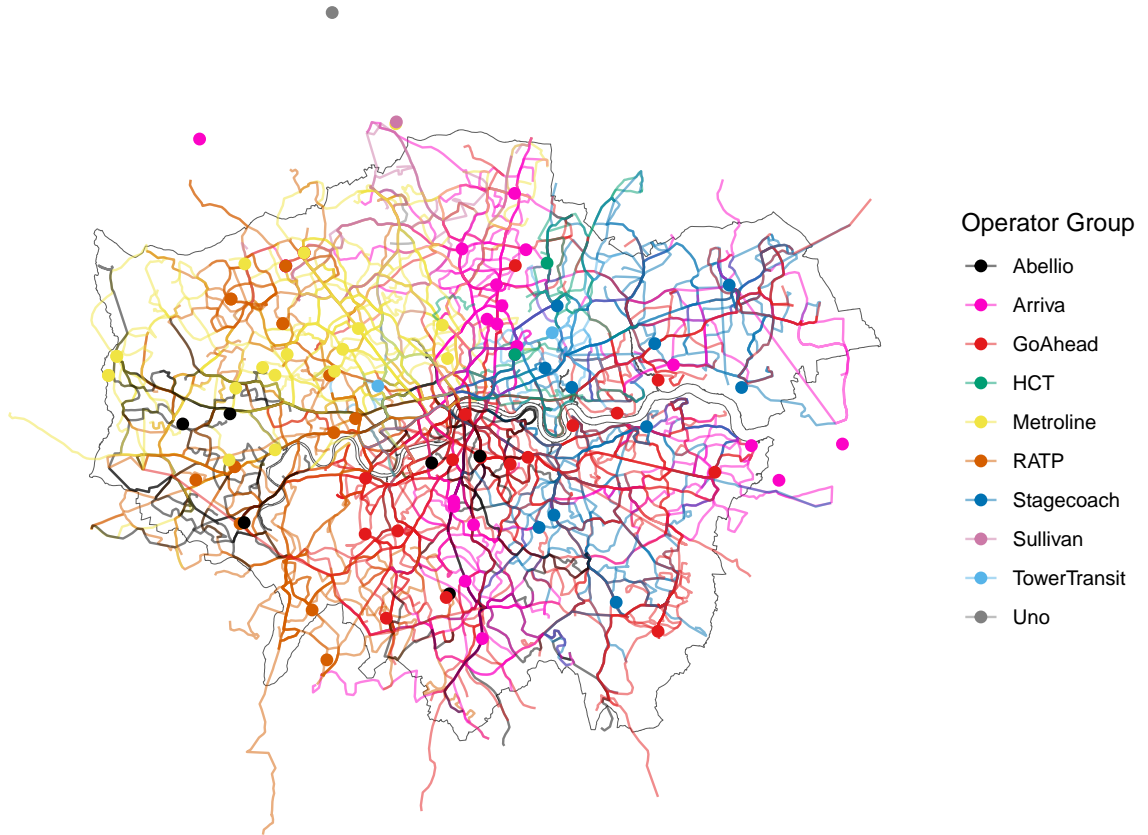


Figure 2: Route network and garage locations by operator group.

Notes: Garage-route network at the end of the sample (2019). Points denote locations of garages, line segments indicate bus routes, and colors indicate which operator owns the garage or operates the route.

level observables are included, as in columns (1)-(4) of table 3.¹¹

This paper focuses on the implications of dead miles costs for operators' location (garage) choice behavior, and the costs and benefits from being near competing or own garages. Hence, we complement the regression analysis presented in our companion paper to capture potential agglomeration benefits and competitive considerations. Specifically, we use the tender data — downloadable from the Transport for London website — to regress the winning bid on: i) auction characteristics, ii) our computed measures of the dead miles driving time between the route and the garage where the route is operated from, iii) the driving time between the garage where the route is operated from and the closest garage by the same operator (*closest own garage* to capture agglomeration benefits), and iv) the dead miles of the closest garage of a competitor of the winning bidder (to capture local monopoly rents).

¹¹As there may be selection on dead miles by the auctioneer, [Marra and Oswald \(2023\)](#) also confirms these conclusions using individual bid data provided by two firms in this market.

Table 2: Summary statistics of garage-route network variables

	Mean	SD	Min	Max
Accepted Bid (Million Pounds)	2.80	2.28	0.00	18.17
Dead Miles Start-Stop (Minutes)	13.23	5.18	2.55	41.45
Dead Miles Start-Stop (km)	7.82	3.55	1.31	26.94
Route Length	7.88	2.56	2.00	24.00
Number of Bidders	2.79	1.10	1.00	9.00
Peak Vehicle Requirement (PVR)	12.45	7.76	0.00	53.00
Closest Own Garage (mins)	12.20	10.65	0.00	48.80
Dead Miles of Closest Competitor (mins)	13.17	4.75	3.30	35.60
Total Number of:				
Auctions	1907.00			
Auctions w/ observed dead miles	1126.00			
Unique Routes	698.00			
Routes w/ observed operating garage	592.00			

Notes: Summarizing the variables used in the regressions in table 3. We label *Dead Miles Start-Stop* the average of time (distance) from garage to start or end point of a given bus route.

Garage Ownership 1994–2020



Figure 3: Number of ownership changes by garage

Notes: This map illustrates that different garages experience a different number of owners over time, in part depending on their location within the route network. Darker road segments imply a greater number of routes.

Table 3: Explaining the size of accepted bids for London bus routes (Start-Stop Distance)

	Accepted Bid (in Million Pounds Sterling)				
	(1)	(2)	(3)	(4)	(5)
Dead Miles Start-Stop (Minutes)	0.088*** (0.011)	0.095*** (0.011)	0.099*** (0.011)	0.019** (0.006)	0.005 (0.007)
Route Length	0.072*** (0.022)	0.077*** (0.021)	0.067** (0.021)	-0.013 (0.012)	-0.022+ (0.012)
Number of Bidders				-0.149*** (0.029)	-0.126*** (0.029)
Peak Vehicle Requirement (PVR)				0.237*** (0.004)	0.225*** (0.005)
Constant	1.330*** (0.209)	0.509+ (0.272)	0.172 (0.332)	-0.786*** (0.207)	-0.875*** (0.223)
Closest Own Garage (mins)					-0.006+ (0.004)
Dead Miles of Closest Competitor (mins)					0.043*** (0.008)
Year FE	-	✓	✓	✓	✓
Winner FE	-	-	✓	✓	✓
Num.Obs.	1457	1457	1457	1457	1457
R2	0.059	0.146	0.196	0.748	0.754
RMSE	2.08	1.98	1.92	1.08	1.06

Notes: This table shows results from regressions run at the winning bid level data. For each winning bid, we record the dead miles measure – if available (in minutes drivetime), route length (in km), the number of bidders in the relevant auction, the route’s Peak Vehicle Requirement, as well as a measure for agglomeration benefits (closeness to next own garage) and competitive pressure in the auction (dead miles of the closest competitor garage). We observe the identity of the operating garage for 592 out of 698 total routes. Variable summary statistics are reported in table 2. Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

These garage network variables capture additional variation in winning bids across route auctions but eliminate the dead miles effect found previously, at least when defining dead miles as the average number of minutes between the garage and to the two endpoints of the route (Start-Stop minutes). The estimates are given in column (5) of table 3. The effect of the distance between the garage and the route’s closest stopover point (Stopover minutes) on accepted bids is small but positive when including other network variables (see column (5) of table 4). For both distance measures, the results also point to a negative correlation between the accepted bid and the winning bidder having garages spatially clustered. This is unexpected given the observed spatial clustering of garages as documented above. However, the interpretation of the variables Dead Miles and Closest Own Garage are complicated: their coefficients reflect both cost effects and adjustments of the equilibrium mark-up.

Table 4: Explaining the size of accepted bids for London bus routes (Stopover Distance)

	Accepted Bid (in Million Pounds Sterling)				
	(1)	(2)	(3)	(4)	(5)
Dead Miles Stopover (Minutes)	0.042*** (0.010)	0.049*** (0.009)	0.050*** (0.009)	0.015** (0.005)	0.009+ (0.005)
Route Length	0.078*** (0.024)	0.084*** (0.023)	0.078*** (0.023)	-0.013 (0.013)	-0.024+ (0.013)
Number of Bidders				-0.151*** (0.032)	-0.126*** (0.032)
Peak Vehicle Requirement (PVR)				0.239*** (0.005)	0.225*** (0.005)
Constant	2.137*** (0.206)	1.213*** (0.285)	1.136** (0.350)	-0.749*** (0.221)	-0.913*** (0.246)
Closest Own Garage (mins)					-0.007+ (0.004)
Dead Miles of Closest Competitor (mins)					0.045*** (0.008)
Year FE	-	✓	✓	✓	✓
Winner FE	-	-	✓	✓	✓
Num.Obs.	1333	1333	1333	1333	1307
R2	0.028	0.118	0.166	0.740	0.748
RMSE	2.15	2.05	1.99	1.11	1.09

Notes: This table shows results from regressions run at the winning bid level data. For each winning bid, we record the dead miles measure – if available (in minutes drivetime), route length (in km), the number of bidders in the relevant auction, the route’s Peak Vehicle Requirement, as well as a measure for agglomeration benefits (closeness to next own garage) and competitive pressure in the auction (dead miles of the closest competitor garage). We observe the identity of the operating garage for 592 out of 698 total routes. Variable summary statistics are reported in table 2. Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

By contrast, the fact that operators benefit from being isolated from their competitors is crystal clear. The winning bid is £43,000-45,000 pounds higher for each minute more drive time between the garage where the route is operated from and a garage of a competing bidder (column (5) in tables 3-4), and the interpretation of this effect is straightforward as the driving time to competing garages does not go into an operator’s cost function. As having competitors located further away can only benefit an operator through reduced competition in the auction (abstracting from endogenous network formation explanations), the strongly positive effect of the Dead Miles of Closest Competitor confirms the importance of the spatial layout of the garage-operator-route network.

In the next section we propose a structural model that captures these salient market features, with local monopoly rents associated with garages that are located close to routes but far from

competitors, and with a cost function that allows for cost savings when operating routes from nearby garages, and we estimate their relative importance.

5 Empirical model

This section develops a static two-stage model that formally links the value of a firm's location to the spatial rents the firm can obtain, in order to estimate the determinants of these spatial rents from observed location choices. Each period t is divided into three stages, i.e. stages 0, 1 and 2. Upon entering stage 0, operators (i) observe the current state of the garage-operator network and draw idiosyncratic cost shocks to operate a given route r , denoted $\{v_{irt}\}_{r=1}^R$. Then, in stage 1 the garage-operator network is potentially changed, if any operator obtains a higher expected utility from a garage than the current owner. In stage 2 the operators bid on bus route contracts given this network, where it bears remembering that the location of garages is of first order importance for placing competitive bids. Firms understand this staging structure, and indeed infer the value of a garage in stage 1 from knowledge about expected profits from route auctions in stage 2 as will be shown below, but they do not reason beyond the current period t . In other words, we model operators as myopic decision makers, not as players involved in a dynamic game.

A salient feature of the London bus garage market is that garage use is exclusive, so the incumbent needs to be willing to vacate a garage in order for another firm being able to enter. This feature is exploited to simplify the garage entry problem into one where garages choose operators. We also show that in this setting, garage transaction prices do not need to be observed in order to estimate the the determinants of spatial rents. Next, we describe the bidding stage (stage 2) and the garage network formation stage (stage 1) in more detail.

5.1 Stage 2: Bidding for routes given a garage network

The route procurement auction is a first-price auction with package bidding. Motivated by the results in [Cantillon and Pesendorfer \(2007\)](#) and to simplify the exposition we proceed under the assumption that there are no cost synergies in route operation.¹²

First, consider what determines the marginal cost for operator i to operate route r at time t (c_{irt}). Route-specific features such as the number of bus stops on the route, the length of the route, the frequency of service, and the type of bus that should be run on it are important and fixed by TfL in their call for tenders. These are captured by the variable $\text{RouteSpecificCosts}_{r,t}$

¹²[Cantillon and Pesendorfer \(2007\)](#) focus on identification and estimation of cost in combinatorial auctions and find limited evidence for the cost synergies that justify a combinatorial auction format in the London bus route market. The median combination of two (three) routes is estimated to cost at least 11 (24) percent *more* to operate than the individual routes separately. One explanation for this finding could be that garage capacity constraints are more important than any cost savings from operating at a larger scale.

that are the same for all operators. More interesting are the elements of the cost function that differ across operators. The variable $DeadMiles_{irt}$ captures the shortest driving time between the endpoints of route r and the garage where operator i at time t operates the route from (typically, that is the closest garage). It can also be expected that operating routes from multiple garages clustered together is likely to lower the marginal cost, which is captured by the variable $DensityGarages_{it}$. Combined, the cost function contains

$$\begin{aligned} c_{irt} &= RouteSpecificCosts_{rt} + DeadMiles_{irt} + DensityGarages_{it} + v_{irt} \\ &\equiv \delta_{irt} + v_{irt} \end{aligned} \quad (1)$$

where δ_{irt} reflects the deterministic cost component that is publicly observable. Operators have private information about their idiosyncratic cost component v_{irt} , which we assume $v \sim^{i.i.d.} F_v$. They bid (b_{ir}) to maximize their expected profit from operating the route

$$(b_{irt} - c_{irt}) \Pr(\text{win}|b_{irt}), \quad (2)$$

where $\Pr(\text{win}|b_{irt})$ indicates the probability of winning the auction when submitting bid b_{irt} . Crucially, the optimal bidding strategy involves a trade-off between increasing the probability of winning and reducing the revenue conditional on winning. We do not impose that firms bid competitively. Instead, we adopt a weaker notion that firms' expected procurement profits are increasing in the probability that they win a contract for a given bid and decreasing in their route operation costs. This holds for the unique Bayes Nash equilibrium in the case without collusion, when the setting reduces to a standard Independent Private Value first price procurement auction. It also holds for a cartel with side-payments (or without side-payments when there are enough contracts auctioned), which are (close to) efficient (Pesendorfer (2000)). For simplicity of exposition, and given prior evidence towards this end by Waterson and Xie (2019), we illustrate the link between auction profits and garage location in the competitive scenario.

Under some regularity conditions, the unique symmetric Bayes Nash equilibrium $\sigma^*(c_{irt})$ solves (see e.g. Guerre, Perrigne and Vuong (2000))

$$b_{irt} = \sigma^*(c_{irt}) = c_{irt} + \frac{\Pr(\text{win}|\sigma^*(c_{irt}))}{\Pr'(\text{win}|\sigma^*(c_{irt}))}, \quad (3)$$

where $\Pr'(\text{win}|\sigma^*(c_{irt}))$ denotes the derivative of $\Pr(\text{win}|b_{irt})$ with respect to the bid. Linking these route operation costs to the value of a garage, it is clear that more centrally located garages are more valuable. This is because having more route endpoints in the vicinity of the garage is beneficial to all operators through lowering their c_{irt} (assuming a positive coefficient on $DeadMiles_{irt}$ in (1)). Moreover, if there are agglomeration benefits in the operation of routes

from multiple nearby garages, reflected in a positive coefficient on $\text{DensityGarages}_{it}$ in (1), garages that are more central in an operator's own network of garages are more valuable, too.

The location of the garages of competing bidders does not affect c_{irt} but it does affect expected profits from the auction. This can be seen by unpacking the mark-up term in the equilibrium bid function (3). For bidder i , the probability of winning is equal to one minus the probability that any competitor has a lower cost than $\sigma^{-1}(b_{irt})$ (given equilibrium bidding), i.e.

$$\begin{aligned} \Pr(\text{win}|b_{irt}) &= 1 - \prod_{h \neq i \in \{1, \dots, N_t\}} \Pr\left(v_{hrt} + \delta_{hrt} \leq \sigma^{-1}(b_{irt})\right) \\ &= 1 - \prod_{h \neq i \in \{1, \dots, N_t\}} F_v\left(\sigma^{-1}(b_{irt}) - \delta_{hrt}\right), \end{aligned} \quad (4)$$

with N_t denoting the number of bidders. Conditional on b_{irt} , the probability of winning increases in the distance between competitors' garages and the route (higher DeadMiles_{irt} and therefore δ_{hrt} for competitor h). Garages that are further away from the closest garage of any competing operator can therefore extract a higher mark-up over cost in the route procurement auction (or, equivalently, higher cartel side-payments). In other words, dead miles generate local monopoly rents to operator i for garages that are more isolated in the network of its competitors.

5.2 Stage 1: Garage Entry

Now that we have established how the location of garages affects expected rents from TfL route auctions, we turn to what this implies for the garage-operator match function in the first stage.

5.2.1 Link to second stage

The value of garage j to operator i in stage 1 of period t (π_{ijt}) can be specified as the sum over all routes of expected revenues generated through route procurements:

$$\pi_{ijt} = \sum_r \left[(\sigma^*(c_{jirt}) - c_{jirt}) \Pr(\text{win}|\sigma^*(c_{jirt})) \right]. \quad (5)$$

Relative to the specification of marginal costs in (1), j -subscripts are added to indicate that the value function is garage-specific due to the dead miles that need to be covered between the endpoints of route r and garage j (while in the route auction stage, bidder i already choose which garage to operate the route from). The probability of winning is decreasing in distance to the route, and the garage value function naturally captures that operators only bid on a subset of routes that are not too far from their garages if one imposes a minimum win probability above

which operators consider participating (as in [Cantillon and Pesendorfer \(2007\)](#)).¹³ Tying the garage value function to the expected auction revenues from the second stage, and to elements of the network of routes and garages, allows us to employ an empirical strategy that recovers elements of the garage value function from operators' garage entry decisions.

5.2.2 Reduced form

Let \mathcal{N}_t (with cardinality N_t) denote the set of operators in period t , each of which is indicated with subscript i . \mathcal{J} denotes the set of garages that is fixed over time and I_{jt} the identity of the operator owning garage j at time t . The payoff for operator i of owning garage j at time t is denoted by π_{ijt} . Given the strong link with the expected profits that can be generated in the procurement auction it is particularly important that π_{ijt} captures the spatial location of the garage in the network.¹⁴ Specifically, let π_{ijt} be

$$\pi_{ijt} = \pi \left(\Gamma_j^X, \Gamma_{ijt}^C, \Gamma_{ijt}^O, X_{jt} \right) + \epsilon_{ijt}, \quad (6)$$

with $\epsilon_{ijt} \perp \left(\Gamma_j^X, \Gamma_{ijt}^C, \Gamma_{ijt}^O, X_{jt} \right)$, and $\Gamma_j^X, \Gamma_{ijt}^C, \Gamma_{ijt}^O$ representing the potentially high-dimensional distance matrices capturing respectively the centrality of garage j in the exogenous route network, the centrality of garage j in the garages of competitors of i , and the centrality of garage j in the rest of the network of garages of operator i at that time. In its most extensive form, matrix Γ_j^X contains distances between garage j and all bus stops in the city.¹⁵ Bounded by a low number of observations, the dimensionality of $\Gamma_j^X, \Gamma_{ijt}^C$, and Γ_{ijt}^O is reduced in our empirical analysis. For instance, Γ_{ijt}^C is reduced to include only the distance to the closest garage from any competitor of operator i to garage j . This is reasonable if only the number and not the identity of competitors affect expected profits from the route procurement auction, which is justified in a model with symmetric bidders.

As we describe in detail in [Marra and Oswald \(2023\)](#), we use the actual driving time as our preferred measure of distance to fully account for the spatial nature of the problem, including geographic obstacles such as the River Thames that can only be crossed at certain points. Other elements of the garage value function potentially include observed garage characteristics, X_{jt} , and an additive idiosyncratic cost specific to each route and operator in each period, ϵ_{ijt} .

¹³Garages do not have unlimited capacity, so if the dead miles constraint is not binding before, the true garage value should also be subject to a constraint limiting the total expected Peak Vehicle Requirement of that route to the number of buses that can be parked in the garage.

¹⁴The garage value in (6) can be thought of as the reduced form of the structural garage value derived from expected procurement profits, defined in (5) in the appendix.

¹⁵The full matrix captures that some intermediate bus stops are used as so-called stopover points where drivers are replaced to take a (meal) break or switch routes, so as buses are driven empty between the garage and such stopover points they also contribute to the route's dead miles.

5.2.3 Garage chooses operator

We estimate the relationship between garage value and observables, $\pi(\cdot)$ in (6), assuming that the event of observing an ownership change is independent of $\pi(\cdot)$. The validity of this approach relies on the extent to which the empirical analysis captures all (first order) factors of relevance to how different operators value the different garages. A natural way to proceed would be to construct an estimator of $\pi(\cdot)$ based on the multinomial choice problem of operators choosing a garage out of the set of available garages. We take a different approach, based on the salient features of the London bus garage market where garages are never shared between operators and an operator needs to leave the premises before another firm can enter. Primarily because of these features, the problem is re-interpreted as garages choosing an operator out of the set of operators \mathcal{N}_t that are active in period t — owning at least one garage.

All operators are free to make an offer to purchase a garage at any time, and garage use is exclusive. We assume that operators can always match the price offered by other firms, so the garage price determination process can be described by a second price auction.¹⁶ Specifically, the equilibrium price for garage j at time t can be described as a function of all operators' realizations of π_{ijt} , as:

$$p_{jt}^* = Y_{jt}^{N_t-1:N_t}, \quad (7)$$

where $Y_{jt}^{N_t-1:N_t}$ denotes the second-highest value (π) for garage j at time t among all N_t market participants. Note that also the firm currently owning the garage is part of the order statistic $Y_{jt}^{N_t-1:N_t}$, because the highest-value bidder can only buy the garage at a price exceeding the incumbent's value.¹⁷ Interpreting transaction prices as resulting from an auction process in (7) makes clear that all firms face the same equilibrium price to enter into a garage; p_{jt}^* . That also holds for the incumbent, for whom p_{jt}^* is the opportunity cost of staying in the garage, and for the market in the form of a property developer or investor. We furthermore assume that the price affects the indirect value of a garage in the same way for all market participants, which applies for instance when the indirect garage utility equals $\pi_{ijt} - \alpha p_{jt}^*$. It is obvious that under this assumption, which rules out a random coefficient on the price variable, that p_{jt}^* drops out when choosing among $i = 0, \dots, N_t - 1$ on the basis of maximizing $\pi_{ijt} - p_{jt}^*$.¹⁸

Some garages are not yet developed at the start of the sample in 1995, some remain vacant before being bought by another operator, and a few are shut down permanently before the end

¹⁶The approach is proposed by [Allen, Clark and Houde \(2019\)](#) for empirical analysis of negotiated price markets. [Slattery \(2021\)](#) also models firm relocation subsidies offered by states and cities as a second price auction.

¹⁷This is like an auction with a secret reserve price, where the seller can also be modelled as an additional bidder.

¹⁸As such, the garage value function presented in (6) does not contain the price of buying the garage, and differs therefore from the *indirect* utility function that underlies a typical discrete choice model in IO. In fact, the prices paid for most transactions in our data are unobserved despite our best effort to recover them. When observing transaction prices, one can also directly estimate the garage value function from the equilibrium price in (7) by applying an estimation method from the auction literature. While supporting our results with estimates obtained with a vastly different estimation method and based on transaction prices rather than the identity of entrants remains worth pursuing, we have to date not managed to collect more than a handful of garage transaction prices.

of the sample in 2019. To avoid modelling the development of additional garages, and to reflect the highly congested London property market, the set of garages \mathcal{J} is constant over time and contains all garages that are ever in use over the course of our sample period. When a garage is not (yet) owned by a bus operator, it is vacant from the perspective of this study, and assigned to *the market* which is the outside option, $i = 0$, and is always included in the choice set \mathcal{N}_t . The mean utility of choosing the outside option is normalised to 0: $\pi(\Gamma_j^X, \Gamma_{0jt}^C, \Gamma_{0jt}^O, X_{jt}) = 0 \forall j, t$.

5.2.4 Timing assumptions

We exploit the time dimension of our data and assume that garages choose operators sequentially. This reflects the fact that in the vast majority of cases, only one garage changes ownership at a given date (see figure 1). As such, the observed garage ownership vector can in each period be interpreted as a Nash equilibrium of the garage–choose–operator game in the sense that there are no profitable unilateral deviations where a different operator would generate more utility in any of the garages. Combined with the abstraction from dynamics, this rules out multiple equilibria. When garage j chooses an operator at time t , it conditions on the observed state of the network of own- and competing garages at that point in time (Γ_{ijt}^O and $\Gamma_{ijt}^C \forall i$).¹⁹

Framing the problem as sequential choice also avoids complexities associated with a bundled choice problem. This is especially true as the garages are not mutual substitutes for all firms in the presence of agglomeration effects, relating to the canonical result in Milgrom (2000, Theorem 4, paraphrasing) that when auctioning multiple objects that are not mutual substitutes for all bidders, for some payoff functions no competitive equilibrium price vector (\mathbf{p}^*) exists.²⁰ Another way to justify the assumption of sequential choice, abstracting from bundles, is that the agglomeration benefit is never large enough relative to the garage purchase price to justify buying more than one garage at once. Again, this could explain why budget-constrained operators in reality change only one garage at a time in over 80 percent of cases where the

¹⁹Essentially, the sequential garage–choose–operator game is a type of myopic *alternating selection mechanism*, proposed by McAfee (1992) as a simple solution to divide assets among firms or individuals. However, the division of garages across operators is not final and transfers can be made, which removes the benefit of selecting first as in McAfee (1992).

²⁰To illustrate, consider garage A and garage B choosing among two operators simultaneously (the example is taken from Milgrom (2000)). Let their utilities be given by:

Garage:	A	B	A & B
Operator 1	a	b	a + b + c
Operator 2	a+d	b+d	a+b+d

$c > 0, c/2 < d < c$.

For operator 1, with $c > 0$ the garages have agglomeration benefits so they are not perfect substitutes even when ignoring their location in the bus route network and the relation to competing operators' garages. For operator 2, the garages are substitutes. Given the parameterization of $c/2 < d < c$, the value-maximizing allocation is for both garages to choose A. However, with $d > 0$, operator 2 will only not be chosen by garage A and garage B if their respective transaction prices are at least $a + d$ and $b + d$, and at those prices also operator 1 does not want to enter the combination of A and B. In our setting an additional complication comes from the utility of operator 1 for garage A potentially depending on whether or not operator 2 enters in garage B, and so on.

garage network changes (see figure 1b).

To account for persistence effects, a firm's tenure at a garage is included in some empirical specifications, as detailed in section 6. It is important to recognize that if the closeness of a competing garage is especially important, abstracting from dynamics in the garage entry decision might affect our empirical analysis.²¹

5.2.5 Closed form operator choice probabilities

The error term ϵ_{ijt} (in the reduced form garage value model in (6)) varies randomly across operators, garages, and over time. It contains a mixture of idiosyncratic route operation costs v_{irt} (in the route operation cost model in (1)), for all routes in the proximity of garage j , and measurement error associated with the low-dimensional representations of the network variables $(\Sigma_j^X, \Sigma_{ijt}^C, \Sigma_{ijt}^O)$. Assuming that $\epsilon_{ijt} \sim i.i.d. F_\epsilon$, where F_ϵ is Gumbell, the problem reduces to a conditional logit model with familiar closed form choice probabilities (McFadden (1984)). The estimation results based on the model outlined in this section therefore rely on the identities of the entrants at each date that the garage ownership vector changes, and are obtained by maximum likelihood estimation.

6 Estimation results

This section presents the estimation results of the reduced form model introduced in section 5.2.2 where *garages choose operators*, assuming that when a garage changes ownership, the entrant has the highest utility among all active operators at that time.

6.1 Estimation Sample

The unit of observation in the multinomial choice model is a garage ownership change, so that the estimates are based on all 192 ownership changes observed during our sample between 1995 and 2019 (i.e., all network changes reported in table 1). In this reversed choice setup, the number of choice alternatives equals the number of operators that have at least one garage at the date when the garage ownership change takes place. Our environment is static, but it is interesting to note that the number of alternatives has been increasing from about 6 in the start of our sample to about 9 at the end, with a peak of 12 active operators during two *change-dates* in 2003 and 2004 (see figure 7a).

²¹For example, it is conceivable that a firm enters a location even when his period payoff is negative as long as this makes all nearby garages less attractive to his competitors so that in the future he can expand there at lower cost. This describes a type of preemptive *capacity expansion* strategy, not too different from preemptive expansion by airlines facing a threat of entry of Southwest Airlines into their routes (Goolsbee and Syverson (2008)).

Table 5: Summary statistics, location choice model

	Mean	SD	Min	Max
Choice entry	0.09	0.28	0.00	1.00
No. bus route starts within 10 min. drive time (Γ_j^X)	14.52	7.07	0.00	34.00
Min. drive time to any own garage (Γ_j^O)	28.08	17.06	0.00	86.40
Avg. drive time to any own garage (Γ_j^O)	38.75	14.48	0.00	86.40
Avg. drive time to any comp. garage (Γ_j^C)	38.87	5.98	28.32	57.89
No. comp. garages within 10 min. drive time (Γ_j^C)	1.72	1.55	0.00	7.00
No. garages owned	6.52	5.11	1.00	19.00
Years owned by previous operator (T_{jt})	3.24	4.24	0.00	23.00
Incumbency benefit ($\exp(-T_{jt})$)	0.46	0.47	0.00	1.00
Garage is not vacant	0.89	0.31	0.00	1.00

Notes: Summary of the variables used for the estimation of parameters in the reduced form of garage-operator utility (6) in table 6. $\Gamma_j^O = 0$ is obtained only if operator j owns exactly one garage at the change date. Following that observation is the fact that $\min(\Gamma_j^O) = \mu(\Gamma_j^O) = 0$ is caused by operator-changedate pairs of this type. Excluding such operator-changedate pairs (1% of data), we observe $\min(\Gamma_j^O) = 0.70$ and $\mu(\Gamma_j^O) = 1.90$ instead.

We start by giving a descriptive overview of the variables used in estimation in table 5. The first row displays our outcome variable, which takes a value of one if a certain operator is chosen by a garage, zero otherwise, in any given choice situation (*garage-change-date*). Given that the number of operators involved in each situation varies, we cannot give the displayed number an intuitive interpretation. We simply state that 9% of the records of our sample data mark an entry event.

Each garage has an exogenous characteristic Γ_j^X – independent of ownership – which describes the location of the garage within the route network. We operationalize this measure by counting the number of bus route start points within 10 minutes drive time of any given garage, and obtain an average of 15 routes. In other words, within a relatively modest driving time of 10 minutes, the average garage could potentially be used to operate 15 routes. For the remainder of the table, it is useful to remember that at each point in time, we can characterise the network of garages in to *own* (O) and *competitors'* (C) garage networks. With regards to the network properties of an operator's own garages, denoted by Γ_j^O , we include the minimal drive time to any of an operator's own garages, which is on average 28 minutes – meaning that on average it takes 28 minutes to drive a bus from any of operator j 's garages to their closest next garage. We also report the time it takes on average to drive to all of an operator's other garages (39 minutes to the average garage). Turning to the network properties of competitor garages, we implement the measures for Γ_j^C by computing the average over drive times to all competitor garages (39 minutes). Our second measure for Γ_j^C is a count of the number of competitor garages within

Table 6: Estimation results for the reduced form of garage-operator utility (6)

	(1)	(2)	(3)	(4)	(5)
No. bus route starts within 10 min. drive time	0.029 (0.025)	0.106*** (0.021)	-0.023 (0.080)	0.066* (0.033)	0.030 (0.039)
No. bus route starts within 10 min. drive time (squared)			0.002 (0.003)		
Min. drive time to any own garage	-0.179*** (0.042)		-0.210*** (0.048)	-0.194*** (0.043)	-0.197*** (0.046)
Avg. drive time to any own garage		-0.077*** (0.023)			
Avg. drive time to any comp. garage			0.040* (0.017)		
No. comp. garages within 10 min. drive time	-0.636*** (0.189)	-0.547*** (0.160)	-0.596** (0.192)	-0.730*** (0.195)	-0.654*** (0.198)
No. garages owned			-0.152** (0.054)		
Incumbency benefit [$\exp(-T_{jt})$]				2.077*** (0.526)	2.625*** (0.667)
Garage is not vacant					1.795+ (0.923)
Inclusive value	2.551*** (0.568)	1.416** (0.447)	2.722*** (0.629)	3.250*** (0.638)	2.737*** (0.614)
Num.Obs.	5184	5184	5184	5184	5184
Pr(highest-utility operator correctly predicted)	0.37	0.344	0.354	0.385	0.38
Mean AUC	0.912	0.914	0.917	0.915	0.916
Own garage proximity (mins) offsets extra comp. garage	3.56	7.067	2.835	3.771	3.325
McFadden's pseudo- R^2 (ρ^2)	0.33	0.314	0.342	0.353	0.358
McFadden's pseudo- R^2 w.r.t. Null Model (ρ_0^2)	0.468	0.456	0.479	0.487	0.491
Number of Choice Situations	192	192	192	192	192

Notes: Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Maximum likelihood estimates from a nested logit discrete choice model where the level of observation is an ownership change, which can be an entry of a new operator in a garage or the vacating of the current operator without entry of another one. There are two nests: all operators are in one and the vacant choice is in another. Results based on the same specifications but with a logit model are reported in table 8 in the appendix. The data contains 192 ownership changes and in total 5,184 rows of the number of alternatives (operators) \times number of ownership changes. The own garage proximity statistic shows how much closer, in minutes drive time, the closest own garage of an operator must be to garage j to offset having one additional garage of a competing operator within a 10 minute drive of garage j . For four of the five columns, this number falls between 2.5 and 4. The statistic is obtained by dividing the coefficients of *No. comp. garages within 10 min. drive time* by the one of *Min./Avg. drive time to any own garage*. Incumbency benefit [$\exp(-T_{jt})$] is the exponential survival function of the number of years (T_{jt}) the incumbent of garage j owned the garage before the ownership change took place. AUC refers to Area Under the (ROC) curve, and we report the frequency-weighted averages over all curves as in [Provost and Domingos \(2003\)](#) – see section 6.3.

10 minutes drive time, which is 1.7 on average. This description makes salient the fact that on average, own garages cluster together spatially, whereas the distance to competitor garages is significantly larger. Finally in table 5, we see that on average an operator owns 6.5 garages, for a duration of on average 3.2 years, and that 89% of garages are non vacant at any date in our estimation sample.

To summarize, time-varying spatial distance variables are constructed to account for competitive effects, agglomeration benefits, and centrality in the route network.

6.2 Estimates of Spatial Rents

The main results of the paper are shown in table 6. These are based on a nesting structure that places all operators in one nest and the garage going vacant (the outside option) in another nest. The inclusive value is always statistically significant, rejecting non-nested logit demand.²² This is not surprising in the light of our parsimonious empirical specifications; it is more likely that two active operators have similar unobserved idiosyncratic utilities for adding a garage to their garage networks than the “the market” and another operator do. In addition, the estimates reveal a higher mean utility among all operators relative to the market, as can be explained by the operators being able to make direct use of the garage to service profitable TfL bus route contracts (see column 5).

Moreover, these structural estimates confirm the importance of spatial rents in this market. Garages that are closer to the first bus stop of routes in the route network procured by TfL are more attractive to all operators. This underscores that dead miles are important in the garage entry decision, and resonates with the idea that local monopoly rents are generated from having to drive fewer empty miles to service a route.²³ The fact that garages are only valuable to the extent that the firms can use them to submit more competitive bids to operate nearby bus routes is also picked up by the negative coefficient on having a garage owned by a competitor closer by. Moreover, the value of a garage is higher when it is closer to the nearest garage by the same operator, supporting the presence of agglomeration benefits in this market. Detecting agglomeration benefits in the garage entry decision is particularly relevant here as this will, at least to some extent, explain the observed spatial clustering of operators in certain parts of London. Omitting such an effect from the analysis would favour collusive market sharing as an alternative explanation for the observed clustering, as we assess in the counterfactuals. This effect is highly significant and consistent across specifications.

²²A likelihood ratio test between nested un-nested models confirms this at the 5% confidence level. The results for the simple logit model are highly similar and presented in table 8 in the appendix.

²³However, the statistical significance of the dead miles coefficient disappears when also accounting for an incumbency benefit and a dummy for non-vacant garages in column (5) or when omitting this pair in column (1). When using as agglomeration benefit measure the *average* drive time to the closest garage in an operator’s network in column (2), the dead miles measure is particularly strong.

To interpret the magnitudes of these spatial rents we compare the relative importance of agglomeration benefits and local monopoly rents. We find that, on average, the closest garage of an operator (to some garage j) needs to be about 3-4 minutes closer by road to offset the reduction in utility from having one additional garage by a competing operator within a 10 minute drive radius (from garage j).²⁴

We also consider two model extensions. The first extension allows for economies in operating multiple garages irrespective of their location. The results in column 3 suggest that having more garages makes an operator less likely to enter. However, an alternative interpretation relates to the dynamics of operators adjusting their garage network over time to achieve an ideal size. As the model estimates are based on such entry decisions, including the number of garages in $\pi(X_{ijt}; \theta)$ risks over-predicting the utility of entering a garage for small operators in the counterfactual simulations based on all garage-change-dates where no change of ownership occurred, we do not consider this extension in the counterfactual simulations.

The second extension addresses persistence in the garage ownership network, for instance due to fixed costs or having local route expertise. For example, if there are fixed costs involved with developing the garage to meet the needs of the new firm, these are sunk at the time of entry and are not, like p_{jt}^* pocketed by the previous owner.²⁵ Any benefits accrued from the initial investments will likely decrease over time (i.e., assets depreciate), so we include as explanatory variable *Incumbency benefit* = $\exp(-T_{jt})$ which equals the exponential survival function of the number of years T_{jt} that the incumbent has been in garage j at time t when the ownership changes.²⁶ With this specification it takes roughly five years for any benefits from having tailored the garage to one's needs to dissolve to zero. The large and positive coefficient adds additional persistence to the empirical model. We can interpret the magnitude of the incumbency benefit over time in relation to the benefit of being central in the garage network. Upon entering a garage, the incumbent operator has a benefit over the other operators that would only be offset with four (additional) competing garages within a 10 minute drive. After one year, the incumbency benefit is already offset by about 1.5 nearby competing garages. We consider this extension to be important, as we aim to set up our simulation analysis against interpreting non-moves as collusive.

²⁴These results refers to the row labelled *Own garage proximity (mins) offsets extra comp. garage* in table 6.

²⁵Recall that the competitive purchase price p_{jt}^* for garage j at time t is not relevant in our setting to determine the identity of the highest-utility operator, as all operators i including the incumbent will have to pay the same equilibrium (opportunity) cost p_{jt}^* to enter (or stay) in the garage.

²⁶The estimation results are similar in a specification based on the linear T_{jt} variable, but the in-sample fit is better with the exponential (see tables 9-10 in the appendix).




Figure 4: ROC curves for selected model (corresponding to column 5 of table 6).

6.3 Goodness-of-Fit and Model Selection

We want to rely on the models estimated in the previous subsection to inform us about the latent garage-operator values. Based on the economic grounds set out in the previous section, we favor the model specification in column 5 of table 6. Here we provide an evaluation of model performance.

Multinomial choice is a multiclass classification problem, for which no simple summary statistic of goodness-of-fit along the lines of the well-known R^2 measure exists. Instead, we rely on insights from the machine (and statistical) learning literature. We give a brief overview of so-called ROC (Receiver Operating Characteristic) curves, which are our preferred measure, in Appendix B for the binary classification case. In our multi-class classification setting, we rely on the *one-vs-all* classification metric, i.e. we investigate model performance in correctly predicting choice j out of a total set of options \mathcal{J} . That implies that we have one ROC curve for each operator. It describes how well the model performs in correctly identifying an operator-garage pair as having the highest utility (true positive rate) versus incorrectly doing so (false positive rate). We obtain for each potential choice of operator a so-called Area Under the Curve (AUC) measure, which is contained in $[0, 1]$, and informs about model fit. We follow [Provost and Domingos \(2003\)](#) and compute frequency-weighted averages over all ROC curves, to arrive at a mean AUC measure, which is reported for all specifications (the row “Mean AUC” in tables 6 and 8).

Figure 4 shows the ROC curves for all operators separately for our preferred specification. The in-sample performance is generally good, with operators Arriva, First and GoAhead standing out as the ones with poorer ROC measures. However, the average performance measured with the Mean AUC remains high (0.92), even when considering that this measure gives more weight to ROC curves from operators that have at least one garage for more change-dates in the sample. That our garage-utility model does well in terms of selecting the right operator on average (over all garage-change-dates) is shown by comparing the predicted and observed choice frequency (figure 8). The model also performs well in terms of predicting the share of times that each operator has the highest π_{ijt} . The worst in-sample fit is for the operators Arriva, Go-Ahead, and TransDev (and Vacant). It is interesting to note here that the logit model does better in fitting the choice frequency for Vacant than the nested logit model. Combined, this confirms that the model does well in describing which operator enters into a garage, but that it should not be used to determine when a new garage gets built or closed down.

In fact, there are minor differences in the various in-sample model fit statistics across specifications (in tables 6 (for nested logit) and 8 (for logit)). The AUC is about 0.9 for all specifications. The Aikaike Information Criterion is between 656-694. It decreases slightly when introducing the Incumbency benefit variable (column 4 versus 1), as the AIC penalizes the number of variables and the fit does not improve much. However, we prefer to include this variable to make sure that the counterfactuals are set up against mis-interpreting an incumbent-specific benefit from tenure at the garage as punishment. We finally prefer the nested logit specification, including a different intercept for non-vacant operators (column 5) even though the AIC is slightly lower than in the specification without it (column 4), as these estimates also imply stronger agglomeration benefits.

Ultimately, the garage entry process is more complex than how it is represented in this model. We abstract from dynamics, bundled choice, and —due observing only 192 choice situations— have to rely on simple representations of the network variables. However, the two McFadden’s pseudo- R^2 (or *likelihood ratio index*) measures show that our stylized model with only five variables does remarkably well in approximating the data, relative to either a random draw or a model with operator intercepts only.²⁷ With reliable estimates of the garage-operator utility function in hand, we now turn to the issue of estimating counterfactual utility for all garage-change-dates where no ownership change took place and how this relates to potential collusive behavior.

²⁷McFadden pseudo- R^2 values of 0.2-0.4 are typically considered to have an excellent fit (much better than an R^2 of the same magnitude in OLS regressions), see also [Domencich and McFadden \(1975\)](#).

7 Counterfactuals

In this section, we use the estimated garage-operator utility model to understand to what extent profitable garage ownership changes did not take place, and to provide plausible explanations for these lost opportunities.

7.1 Application 1: Efficiency Loss of Hold Out

We first apply our structural estimates to study the efficiency loss accruing from the private ownership of bus garages in London. The garage owner may refrain from selling the garage to another firm with a higher utility for it, waiting to fetch a higher price or in our setting not wanting to give competing operators an advantage in nearby bus route auctions. The issue of privately owned assets not reallocating to their most efficient uses is a widespread problem — motivating clever incentive-based mechanisms to address it (see [Posner and Weyl \(2017\)](#)) — and our structural estimates can be used to quantify the so-called “hold out problem” for the London bus garage market.

Our simulations are based on the following procedure. First, take the set of garages where no ownership change took place during a changedate t . Denote this set by \mathcal{J}_t^N , and the collection for all change-dates by $\mathcal{J}^N = \{\mathcal{J}_t^N\}_t$. We compute the deterministic part of the garage-operator utility ($\hat{\pi}_{ijt}$) based on our model estimates $\forall(t, j \in \mathcal{J}_t^N, i \in \mathcal{N}_t)$, which gives the expected utility for all operator-garage-changedate triples where no ownership change took place. We define $\hat{\psi}_{ijt}$ as the utils difference between the utility of operator i for garage j at time t and utility of the incumbent at that time (I_{jt}), if positive:

$$\hat{\psi}_{ijt} = \max(0, \hat{\pi}_{ijt} - \hat{\pi}_{I_{jt}}) \quad (8)$$

After excluding combinations where $i = 0$ in order not to ascribe behavioral attributes to “the market”, and after excluding $i = I_{jt}$ as choice alternative, a total of 71,775 operator-date-garage triples remain in \mathcal{J}^N . We find that in 11.3 percent of these cases, another operator would benefit more from having that garage in their own network than the incumbent does ($\hat{\psi}_{ijt} > 0$). Figure 9 compares the choice frequency (the share of all $|\mathcal{J}^N|$ times that an operator has the highest π_{ijt}) as predicted by the model and as observed in the data. A take-away from this figure is that when assigning the garage to the operator with the highest predicted π_{ijt} , there are no massive changes in terms of how many garages operators own. While there will still be changes in the configuration of the garage ownership network, the results reported below are not driven by, for instance, a fringe operator that the model predicts should be a large player instead.

The associated efficiency loss can also be presented in monetary terms, because the value of a garage is equal to the sums of expected profits that the garage generates through running

bus route contracts from it.²⁸ Specifically, as a monetary measure of efficiency loss we compute the average utility difference (ψ_{ijt}) as a share of the incumbent’s utility ($\pi_{I_{jt}}$) across all instances where the incumbent stayed in the garage, and multiply this by the average profits from the bus route procurement contracts. This is based on the notion that when an operator does not sell his garage to someone with a higher value for it, this must be motivated by expecting additional gains from (future) auction rents.²⁹

To compute this measure, we start from the set \mathcal{J}^N defined above, and we further exclude all cases where $\hat{\psi}_{ijt}$ exceeds its 99th percentile as well as negative estimated incumbent values. These measures are taken to minimize the dependence of our efficiency measure on outliers, and both go in the direction of underestimating the effect on the efficiency loss. Then, if there are multiple operators with a higher $\hat{\pi}_{ijt}$ than the incumbent, we consider only the potential entrant with the highest utility for the garage as this would be the firm that is predicted to enter in an efficient allocation. After making these changes, we are left with 6,401 instances where we have the estimated utility of the incumbent for a garage at a change-date and the estimated utility of the highest-utility potential entrant (i), and we denote this set by \mathcal{J}^{Ngt} . Next, we compute ζ as the average lost utility as a share of the incumbent’s utility across all observations in this set:

$$\zeta \equiv \frac{1}{|\mathcal{J}^{Ngt}|} \sum_{(i,j,t) \in \mathcal{J}^{Ngt}} \left(\frac{\hat{\psi}_{ijt}}{\hat{\pi}_{I_{jt}}} \right) \quad (9)$$

The estimated share equals $\zeta = 0.65$ between 1995 and 2019. This number is stable over time. When restricting the time period to 2004-2019 —matching the tender data— the share equals $\zeta = 0.66$. Due to the tight link between the value of the garage and revenues from the route auctions, additional rents to the operator are losses to the taxpayer. Annual revenues of operators from TfL route auctions in our data amount to £273m per year and [Cantillon and Pesendorfer \(2006\)](#) estimate that mark-ups equal 10-15 percent of the submitted bid in this market (based on structural estimates using data spanning 1995-2001). Hence, the hold out problem in the London bus market is estimated to cost the taxpayer approximately £17.7m to £26.6m per year ($\zeta \times 273 \times [0.1, 0.15]$), or about 6.5% to 9.8% of the total procurement cost of providing public bus transportation in London.

²⁸Note that if all garage transaction prices would be observed, the the garage transaction price would be included in the utility function (6) and this could lend a monetary interpretation to the estimated utility loss.

²⁹Simply put, when a competing operator is estimated to have a 20 percent higher utility than the incumbent for that garage ($\frac{\psi_{ijt}}{\pi_{I_{jt}}} = 1.2$) the associated efficiency loss would be 20 percent of the expected procurement profits for this (ijt)-combination, as the incumbent must expect to gain at least 20 percent more from winning tenders when not selling the garage.

7.2 Application 2: Interpretation as Market-Sharing

We now explore whether unrealized profitable transactions, as identified above, can be interpreted as resulting from collusive market-sharing agreements between operators. This line of inquiry is motivated by the fact that the firms in this sample have a history of anti-competitive behavior in the UK. Specifically, in the UK market for bus route services, the UK Competition Commission finds evidence for “[...] operator conduct by which operators avoid competing with other operators in ‘Core Territories’ (certain parts of an operator’s network which it regards as its ‘own’ territory) leading to geographic market segregation.” Their 2011 report estimated an annual cost for UK consumers and taxpayers between 115 and 305 million pounds from adverse effects on competition and voiced the concern “[...] that geographic market segregation might be a more widespread feature than we have identified”. The investigation excluded London and Northern Ireland.

Hence, we consider (tacit) market-sharing agreements that designate certain spatial areas of influence to operators.³⁰ We generally follow [Pesendorfer \(2000\)](#) in considering a static game between operators, with individual firms having incentives to deviate from the collusive agreement, and with compliance potentially coming from an infinitely repeated version of the game, or potentially from repeated interactions in other markets outside of London.³¹

A key idea that we rely on in this application is that incentive-compatibility is guaranteed by the expected punishment inflicted upon the firm deviating from a collusive agreement. On the other hand, all instances when a firm did enter into a garage must have been originated in a setting without (fear of) punishment — as otherwise we are looking at realizations off the (potentially collusive) path.

To formalize these ideas, we first assume that any garage network changes that we do observe could not have been blocked by a potential market-sharing agreement. This rules out that there was some (ineffective) punishment inflicted upon an entrant that is not high enough to deter the firm from entering. As such, our structural estimates of the garage-operator utility function in (5) reflect true values. We furthermore maintain that, without market-sharing agreements covering the garage j at time t , the incumbent of that garage would sell it at a price exceeding its valuation. Together with the assumption that all operators have the chance to meet this bid (so that the equilibrium price p_{jt}^* is independent of the identity of the bidder), we have that the highest-value firm enters the garage absent punishment.³²

³⁰We remain agnostic about the exact form that such an agreement takes. It might be a reciprocal agreement between two or more firms of comparable size (i.e. “I don’t enter your area if you don’t enter mine”) or it might be a unilateral threat (i.e. “Don’t enter my area or else”). For example, this type of collusion describes the territorial allocation of districts by the Cincinnati school milk cartel studied in [Porter and Zona \(1999\)](#).

³¹As in [Pesendorfer \(2000\)](#), compliance with the cartel mechanism implicitly requires a sufficiently large discount factor in the infinitely repeated version of the static game.

³²Hence, we consider that market-sharing agreements “block” some individually profitable transactions from occurring due to the threat of punishment (in whatever form). This can be considered a weak minimum condition for collusion of this type as reciprocity is not imposed. The necessary condition is consistent with the model of collusive network formation by [Belleflamme and Bloch \(2004\)](#). [Belleflamme and Bloch \(2004\)](#) model firms that

Table 7: Regressions of punishment on the Number of Alternatives

	$\psi : \psi > 0$	$\psi : \psi > 0$	$1(\psi > 0)$	$1(\psi > 0)$
(Intercept)	2.108*** (0.068)	2.967*** (0.748)	-1.377*** (0.044)	-1.068* (0.453)
No. alternatives	-0.002 (0.006)	-0.041 (0.035)	-0.058*** (0.004)	-0.059** (0.021)
Num.Obs.	8136	8136	71 775	71 775
R2	0.000	0.003		
R2 Adj.	0.000	0.000		
AIC	32 194.4	32 217.1	50 469.7	50 300.6
BIC	32 215.4	32 406.3	50 488.1	50 539.3
Log.Lik.	-16 094.175	-16 081.575	-25 232.855	-25 124.292
RMSE	1.75	1.75	0.32	0.32
Year-fixed effects?		✓		✓

Notes: Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

7.2.1 Relation to theory


Above, we report that in 11.3 percent of cases our model predicts that another operator has a higher utility from being in a garage than the incumbent ($\hat{\psi}_{ijt} > 0$). It must be recognized that this can be partly driven by noise or model misspecification. For example, the garage might not have been sold to an operator with a high level of expected utility if the incumbent drew a particularly high value of ϵ_{ijt} or for another unmodeled factor. We therefore relate the estimated “punishment” value ($\hat{\psi}_{ijt}$) against the characteristics that are known to make an industry more collusion-prone in theory, as summarized in [Marc Ivaldi, Bruno Jullien, Patrick Rey, Paul Seabright and Jean Tirole \(2007\)](#) (henceforth: IJRST).³³

The first thing to check is whether ψ relates to the number of firms in the market. It is expected that more firms make it more difficult to coordinate, and that deviations are more tempting as firms share the collusive profit (IJRST). We assess this by regressing either the amount of $\hat{\psi}_{ijt}$ when positive or a dummy variable indicating that $\hat{\psi}_{ijt}$ is positive on the number of firms with at least one garage in market (London) at that point in time (see table 7). The negative coefficient on “No. alternatives” confirms that periods with more firms are associated

are originally specialized on one market — e.g., the “Core Territory” of our bus operators — and that compete in (auction) markets afterwards or decide to stay out of each other’s territories by signing bilateral agreements. They show that agreements between multiple parties are more difficult to sustain than bilateral agreements. Moreover, in the presence of asymmetries such as incumbency benefit in the home market, market sharing agreements are more attractive as the loss from agreeing not to enter each other’s territory is lower than without incumbency benefits.

³³Several of these characteristics reflect the market in ways that are fixed over time and across space and so it cannot be verified whether they relate statistically to the estimated ψ ’s. However, the majority of those characteristics arguably apply to the market in question, such as the relatively symmetric firms interacting frequently in tenders for similarly sized and highly standardized contracts

Figure 5: Punishment and relation to firm entry



Notes: Displaying yearly averages of the estimated punishment when positive (plot a), and the punishment probability (plot b), across all operator-garage-changedate triples where no ownership change took place (e.g., all entries of \mathcal{J}^N as defined in section 7.1)). Red bars denote years with *significant entry*, defined as the entry by at least one operator who remained active for at least five years. Time series between 1995-2014 to accommodate identification of *significant entries* up to the end of the sample (2019).

with a lower probability of observing a positive ψ , in line with interpreting ψ as the expected punishment for deviating from a collusive agreement. According to this interpretation of the results, conditional on getting punished the punishment amount does not decrease in the number of firms. It should be kept in mind however that only a lower bound on the expected

punishment is identified, and slackness of this bound may be case-specific. The number of active firms varies over time, but these results are consistent when including year fixed effects.

Relatedly, entry of firms and entry barriers are also important. Without entry barriers, collusive profits would quickly erode through entry of new firms and the prospect of future entry reduces the scope for retaliation (IJRST). To explore this, we relate time series of ψ to the entry of operators. Figure 5 plots yearly averages of ψ when positive and of the probability of a positive ψ . Years that experienced entry of an operator that stayed in the market for at least five years are indicated by the red vertical bars. There are two high-entry periods; between 1997-2000 and in the period 2009 and 2011-2013, and two periods without entry (of a firm that is likely to exert substantial competitive pressure); between 2001-2008 and between 2014-2019. In line with the theoretical prediction when ψ is interpreted as punishment, both the probability of punishment and the average amount are low in the initial high-entry period, increase over time between 2001 and 2008, and then decrease again in the second high-entry period. What happens at the end of the 2010s is less clear. A steadily increasing average punishment value corresponds to the idea that the market entered a period of stability but the punishment probability is only increasing between 2013 and 2015 and decreasing afterwards.

Studying the spatial pattern of ψ 's also reveals interesting facts.³⁴ For garages that are more isolated from competing garages, the incumbent is more likely to punish and also to punish harsher. This corresponds to the idea that there are local monopoly rents associated with more isolated garages. In particular, the incumbent has higher expected route auction revenues when his closest competitor is further away and has therefore a greater incentive to protect this position. While such a spatial aspect of collusion is not particularly referenced in IJRST, it reflects the presence of dead miles as an entry barrier to competing in the route auction.

8 Conclusions

This paper studies spatial rents and firm conduct in the London bus market, building on a new database of bus garages used by private firms for the operation of bus routes in the Greater London Area. In particular, the dataset contains details about the location and ownership of bus garages over time, starting from 1994 when the London bus market was privatized via a split into 12 geographical areas. The data is constructed by combining archival data from the London Omnibus Traction Society and hobbyist (*bus spotters*) websites and is described in detail in [Marra and Oswald \(2023\)](#). The London bus procurement setting is particularly suitable for the analysis of spatial rents in relation to competitive forces. With the demand side being tightly regulated, we show that the location of a garage in the garage-route network is key determinant of a firm's profitability in this market.

³⁴The results discussed in this paragraph are presented in table 11 in the appendix.

We consider spatial rents in a broad sense, including the location of the garage within an operator's own network of garages (capturing *economies of density*), the location of the garage within the network of garages of competing operators (capturing *local monopoly rents*), and the location of the garage relative to the network of bus routes (capturing *dead miles* transportation costs). To study these factors, we develop a structural model to capture observed changes in garage ownership. Estimates reveal substantial local monopoly rents, economies of density in operating routes from multiple nearby garages, and capture persistence in the garage network with an incumbency benefit.

In terms of fit, a parsimonious model with low-dimensional representations of spatial centrality measures does remarkably well in approximating the data in terms of which operator enters which garage. To illustrate, our preferred specification has a *likelihood ratio index* (or McFadden pseudo- R^2) far exceeding or at the upper end of the [0.2-0.4] range that is considered to reflect an "excellent fit" by Domencich and McFadden (1975), both in comparison to a null model ($\rho_0^2 = 0.49$) with intercept only and in comparison to a model with operator fixed effects ($\rho^2 = 0.36$).

The model estimates are furthermore used to compute an estimate efficiency loss from operators holding out garages which competitors could use more efficiently in bidding. We find this loss to lie in between 6.5% and 9.8% of total procurement costs for London buses. We also assess whether collusive behavior can partly explain the large degree of observed spatial clustering of operators in certain parts of the city, by estimating the implied punishments a cartel would need to inflict on deviations. We find that these punishments are related to factors known to make a market more collusion-prone. As such, our results suggest that the concerns voiced by the UK Competition Commission about potential market sharing agreements in the London Bus Market are well-founded.

References

- Allen, Jason, Robert Clark, and Jean-François Houde.** 2019. "Search frictions and market power in negotiated-price markets." *Journal of Political Economy*, 127(4): 1550–1598.
- Bajari, Patrick, and Lixin Ye.** 2003. "Deciding between competition and collusion." *Review of Economics and statistics*, 85(4): 971–989.
- Balboni, Clare Alexandra.** 2019. "In harm's way? infrastructure investments and the persistence of coastal cities." PhD diss. London School of Economics and Political Science.
- Belleflamme, Paul, and Francis Bloch.** 2004. "Market sharing agreements and collusive networks." *International Economic Review*, 45(2): 387–411.
- Berry, Steven T.** 1992. "Estimation of a model of entry in the airline industry." *Econometrica*, 60(4): 889–917.
- Bresnahan, Timothy F., and Peter C. Reiss.** 1990. "Entry in monopoly markets." *The Review of Economic Studies*, 57(4): 531–553.
- Bresnahan, Timothy F, and Peter C Reiss.** 1991. "Entry and competition in concentrated markets." *Journal of Political Economy*, 99(5): 977–1009.
- Cantillon, Estelle, and Martin Pesendorfer.** 2006. "Auctioning bus routes: The London experience." In *Combinatorial Auctions.*, ed. Peter Cramton, Yoav Shoham and Richard Steinberg, Chapter 22, 573–592. Cambridge, Massachusetts: The MIT Press.
- Cantillon, Estelle, and Martin Pesendorfer.** 2007. "Combination bidding in multi-unit auctions." *CEPR Discussion Paper No. 6083*.
- Chassang, Sylvain, and Juan Ortner.** 2019. "Collusion in auctions with constrained bids: Theory and evidence from public procurement." *Journal of Political Economy*, 127(5): 2269–2300.
- Combes, Pierre-Philippe, and Laurent Gobillon.** 2015. "The empirics of agglomeration economies." In *Handbook of regional and urban economics*. Vol. 5, 247–348. Elsevier.
- De Palma, André, Robin Lindsey, Emile Quinet, and Roger Vickerman.** 2011. *A handbook of transport economics*. Edward Elgar Publishing.
- Domencich, Thomas A, and Daniel McFadden.** 1975. "Urban travel demand-a behavioral analysis."
- Ellickson, Paul B, Stephanie Houghton, and Christopher Timmins.** 2013. "Estimating network economies in retail chains: a revealed preference approach." *The RAND Journal of Economics*, 44(2): 169–193.

- Ellison, Glenn.** 1994. "Cooperation in the prisoner's dilemma with anonymous random matching." *The Review of Economic Studies*, 61(3): 567–588.
- Fajgelbaum, Pablo D, and Edouard Schaal.** 2020. "Optimal transport networks in spatial equilibrium." *Econometrica*, 88(4): 1411–1452.
- Fawcett, Tom.** 2004. "ROC graphs: Notes and practical considerations for researchers." *Machine learning*, 31(1): 1–38.
- Gaduh, Arya, Tilman Graff, Rema Hanna, Gabriel Kreindler, and Benjamin A Olken.** 2022. "Optimal Public Transportation Networks: Evidence from the World's Largest Bus Rapid Transit System in Jakarta."
- Goolsbee, Austan, and Chad Syverson.** 2008. "How do incumbents respond to the threat of entry? Evidence from the major airlines." *The Quarterly journal of economics*, 123(4): 1611–1633.
- Green, Edward J, and Robert H Porter.** 1984. "Noncooperative collusion under imperfect price information." *Econometrica: Journal of the Econometric Society*, 87–100.
- Guerre, Emmanuel, Isabelle Perrigne, and Quang Vuong.** 2000. "Optimal nonparametric estimation of first-price auctions." *Econometrica*, 68(3): 525–574.
- Head, Keith, and Thierry Mayer.** 2004. "Market potential and the location of Japanese investment in the European Union." *Review of Economic Studies*, 86(November): 959–972.
- Holmes, Thomas J.** 2011. "The Diffusion of Wal-Mart and Economies of Density." *Econometrica*, 79(1): 253–302.
- Hotelling, H.** 1929. "Stability in competition." *Economic Journal*, 39: 41–57.
- Houde, Jean-François.** 2012. "Spatial differentiation and vertical mergers in retail markets for gasoline." *American Economic Review*, 102(5): 2147–2182.
- Ishii, Rieko.** 2009. "Favor exchange in collusion: Empirical study of repeated procurement auctions in Japan." *International Journal of Industrial Organization*, 27(2): 137–144.
- Ivaldi, Marc, Bruno Jullien, Patrick Rey, Paul Seabright, and Jean Tirole.** 2007. "The economics of tacit collusion: implications for merger control." *Contributions to Economic Analysis*, 282: 217–239.
- Jia, Panle.** 2008. "What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry." *Econometrica*, 76(6): 1263–1316.
- Kawai, Kei, and Jun Nakabayashi.** 2022. "Detecting large-scale collusion in procurement auctions." *Journal of Political Economy*, 130(5): 1364–1411.

- Marra, Marleen, and Florian Oswald.** 2023. "Ownership Frictions in a Procurement Market: Evidence from London Buses." working paper, <https://floswald.github.io/pdf/london-buses.pdf>.
- Mazzeo, Michael J.** 2002. "Product choice and oligopoly market structure." *The RAND Journal of Economics*, 33(2): 221.
- McAfee, R Preston.** 1992. "Amicable divorce: Dissolving a partnership with simple mechanisms." *Journal of Economic Theory*, 56(2): 266–293.
- McFadden, Daniel L.** 1984. "Econometric analysis of qualitative response models." *Handbook of econometrics*, 2: 1395–1457.
- Milgrom, Paul.** 2000. "Putting auction theory to work: The simultaneous ascending auction." *Journal of Political Economy*, 108(2): 245–272.
- Miller, Nathan H., and Matthew Osborne.** 2014. "Spatial differentiation and price discrimination in the cement industry: Evidence from a structural model." *RAND Journal of Economics*, 45(2): 221–247.
- Oberfield, Ezra, Esteban Rossi-Hansberg, Pierre-Daniel Sarte, and Nicholas Trachter.** Forthcoming. "Plants in space." *Journal of Political Economy*.
- Olivares, Marcelo, Gabriel Y Weintraub, Rafael Epstein, and Daniel Yung.** 2012. "Combinatorial auctions for procurement: An empirical study of the Chilean school meals auction." *Management Science*, 58(8): 1458–1481.
- Pesendorfer, Martin.** 2000. "A study of collusion in first-price auctions." *The Review of Economic Studies*, 67(3): 381–411.
- Porter, Robert H.** 1983. "A study of cartel stability: the Joint Executive Committee, 1880-1886." *The Bell Journal of Economics*, 301–314.
- Porter, Robert H, and J Douglas Zona.** 1993. "Detection of bid rigging in procurement auctions." *Journal of Political Economy*, 101(3).
- Porter, Robert H, and J Douglas Zona.** 1999. "Ohio school milk markets: An analysis of bidding." *The RAND Journal of Economics*, 30(2): 263–288.
- Posner, Eric A, and E Glen Weyl.** 2017. "Property is only another name for monopoly." *Journal of Legal Analysis*, 9(1): 51–123.
- Provost, Foster, and Pedro Domingos.** 2003. "Tree induction for probability-based ranking." *Machine learning*, 52(3): 199–215.

- Rotemberg, Julio J, and Garth Saloner.** 1986. "A supergame-theoretic model of price wars during booms." *The American economic review*, 76(3): 390–407.
- Salop, Steven C.** 1979. "Monopolistic competition with outside goods." *The Bell Journal of Economics*, 141–156.
- Seim, Katja.** 2006. "An empirical model of firm entry with endogenous product-type choices." *RAND Journal of Economics*, 37(3): 619–640.
- Slattery, Cailin.** 2021. "Bidding for firms: Subsidy competition in the U.S." *Unpublished manuscript*.
- Sullivan, Christopher.** 2017. "The ice cream split: Empirically distinguishing price and product space collusion." *Available at SSRN 3321948*.
- Waterson, Michael, and Jian Xie.** 2019. "Testing for collusion in bus contracting in London." *Warwick Economics Research Papers No: 1196*.




Figure 6: Time series of Ownership Changes in garage data

Notes: We call *empty* garages those which are vacant, abandoned or not yet built at a certain point in time. *New* refers to a change of operator in a certain garage.


A Additional results

Table 8: Estimation results garage-operator utility (logit model)

	(1)	(2)	(3)	(4)	(5)
No. bus route starts within 10 min. drive time	0.055*** (0.016)	0.107*** (0.019)	0.046 (0.070)	0.111*** (0.023)	0.043 (0.030)
No. bus route starts within 10 min. drive time (squared)			0.000 (0.002)		
Min. drive time to any own garage	-0.070*** (0.008)		-0.079*** (0.009)	-0.063*** (0.008)	-0.074*** (0.009)
Avg. drive time to any own garage		-0.058*** (0.006)			
Avg. drive time to any comp. garage			0.020+ (0.011)		
No. comp. garages within 10 min. drive time	-0.284** (0.108)	-0.450*** (0.107)	-0.285** (0.110)	-0.312** (0.111)	-0.259* (0.111)
No. garages owned			-0.063** (0.020)		
Incumbency benefit [$\exp(-T_{jt})$]				1.487*** (0.379)	2.378*** (0.531)
Garage is not vacant					1.853** (0.585)
Num.Obs.	5184	5184	5184	5184	5184
Pr(highest-utility operator correctly predicted)	0.38	0.375	0.38	0.417	0.385
Mean AUC	0.911	0.916	0.915	0.912	0.915
Own garage proximity (mins) offsets extra comp. garage	4.056	7.764	3.614	4.923	3.499
McFadden's pseudo- R^2 (ρ^2)	0.32	0.314	0.332	0.337	0.349
McFadden's pseudo- R^2 w.r.t. Null Model (ρ_0^2)	0.461	0.456	0.47	0.474	0.484
Number of Choice Situations	192	192	192	192	192

Notes: Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Maximum likelihood estimates from a logit discrete choice model where the level of observation is an ownership change, which can be an entry of a new operator in a garage or the vacating of the current operator without entry of another one. Results based on the same specifications but with a nested logit model, with all operators in one nest and the vacant choice in another one, are reported in table 6.

(a) When owner changed



(b) When owner stayed




Figure 7: Number of active operators that a garage can choose from

Table 9: Estimation results with linear incumbency tenure (logit model)

	(1)	(2)	(3)	(4)
No. bus route starts within 10 min. drive time	0.067*** (0.017)	0.029 (0.028)	0.111*** (0.023)	0.043 (0.030)
Min. drivetime to any own garage	-0.069*** (0.008)	-0.075*** (0.009)	-0.063*** (0.008)	-0.074*** (0.009)
No. comp. garages within 10 min. drive time	-0.280* (0.109)	-0.250* (0.109)	-0.312** (0.111)	-0.259* (0.111)
Linear incumbency benefit (T_{jt})	-0.070+ (0.036)	-0.091* (0.039)		
Incumbency benefit [$\exp(-T_{jt})$]			1.487*** (0.379)	2.378*** (0.531)
Garage is not vacant		0.737+ (0.444)		1.853** (0.585)
Num.Obs.	5184	5184	5184	5184
AIC	686.9	686.1	673.1	663.3
Share of correct predictions	0.38	0.391	0.417	0.385
Mean AUC	0.91	0.912	0.912	0.915

Notes: Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.




Figure 8: In sample fit: observed vs predicted entry frequency by operator

Notes: Showing the in sample fit for the selected model (column 5 of table 6). Data choice frequency records the number of choice situations in data where an operator (including the Vacant operator) is the entrant, divided by the total number of choice situations (192). Model choice frequency records the number of choice situations where the operator is assigned the highest predicted utility by the selected model, divided by the total number of choice situations.




Figure 9: Out-of-sample observed vs predicted entry frequency by operator

Notes: Showing the out-of-sample fit for the selected model (column 5 of table 6). Data choice frequency records the number of choice situations in data where an operator (including the Vacant operator) is the incumbent, divided by the total number of choice situations (192). Model choice frequency records the number of choice situations where the operator is assigned the highest predicted utility by the selected model, divided by the total number of choice situations.

Table 10: Estimation results with linear incumbency tenure (nested logit model)

	(1)	(2)	(3)	(4)
No. bus route starts within 10 min. drive time	0.023 (0.031)	0.010 (0.038)	0.066* (0.033)	0.030 (0.039)
Min. drive time to any own garage	-0.212*** (0.046)	-0.215*** (0.048)	-0.194*** (0.043)	-0.197*** (0.046)
No. comp. garages within 10 min. drive time	-0.683*** (0.196)	-0.661*** (0.200)	-0.730*** (0.195)	-0.654*** (0.198)
Linear incumbency benefit (T_{jt})	-0.130** (0.050)	-0.137** (0.052)		
Incumbency benefit [$\exp(-T_{jt})$]			2.077*** (0.526)	2.625*** (0.667)
Garage is not vacant		0.545 (0.754)		1.795+ (0.923)
Inclusive value	3.209*** (0.689)	3.041*** (0.687)	3.250*** (0.638)	2.737*** (0.614)
Num.Obs.	5184	5184	5184	5184
AIC	674.1	675.4	659.6	656.1
Share of correct predictions	0.37	0.359	0.385	0.38
Mean AUC	0.914	0.914	0.915	0.916

Notes: Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Figure 10: Winning auction bid, normalized by Peak Vehicle Requirement

Time series of Average AcceptedBid/PVR (Million Pound Sterlings)
Including observations with missing dead miles




Table 11: Regressions of punishment on two measures of distance to competitors' garages

	$\psi : \psi > 0$	$\psi : \psi > 0$	$1(\psi > 0)$	$1(\psi > 0)$	$\psi : \psi > 0$	$\psi : \psi > 0$	$1(\psi > 0)$	$1(\psi > 0)$
(Intercept)	1.628*** (0.036)	1.592*** (0.093)	-2.907*** (0.024)	-3.266*** (0.058)	0.893*** (0.093)	0.873*** (0.127)	-4.728*** (0.060)	-4.926*** (0.079)
Min. dist. to competitors (km)	0.083*** (0.005)	0.085*** (0.005)	0.181*** (0.004)	0.193*** (0.004)				
Avg. dist. to competitors (km)					0.042*** (0.003)	0.044*** (0.003)	0.101*** (0.002)	0.102*** (0.002)
Num.Obs.	8136	8136	71 775	71 775	8136	8136	71 775	71 775
R2	0.028	0.031			0.021	0.024		
R2 Adj.	0.028	0.028			0.021	0.021		
AIC	31 963.4	31 984.6	48 714.5	48 045.6	32 024.9	32 041.2	48 467.4	47 999.6
BIC	31 984.4	32 173.8	48 732.8	48 284.4	32 045.9	32 230.3	48 485.8	48 238.3
Log.Lik.	-15 978.715	-15 965.324	-24 355.226	-23 996.824	-16 009.432	-15 993.582	-24 231.702	-23 973.802
RMSE	1.72	1.72	0.31	0.31	1.73	1.73	0.31	0.31
Year-fixed effects?		✓		✓		✓		✓

Notes: Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

B Goodness of fit measures: ROC curves

In this section, we illustrate the interpretation of ROC curves as goodness of fit measures. While our application is multinomial (or *multiclass*), the illustration is based on a binary classification exercise (i.e. with potential outcomes **positive** and **negative**). A classification algorithm C can achieve one of four outcomes when classifying a certain observation i , highlighted in the schematic *confusion matrix* table 12, leading to several performance measures: The *true positive rate* is defined by $\frac{TP}{P}$, the *false positive rate* is $\frac{FP}{N}$, *precision* is $\frac{TP}{TP+FP}$ and *accuracy* is $\frac{TP+TN}{P+N}$. Those quantities are related to test *sensitivity* (equal to the true positive rate), and *specificity* (one minus the false positive rate). If the classifier is probabilistic in nature (e.g. a Naive Bayes classifier or indeed our multinomial logit (i.e. softmax) model), the analyst has to choose a cutoff value c , whereby model predictions greater than c are labelled **positive**, and vice versa. Notice that each value for $c \in [0, 1]$ applied to a given classifier C (same data), will give rise a different confusion matrix, hence different performance measures. It is straightforward to choose the *best* cutoff level and show the model in its most favourable light, which is not desirable. The ROC curve is a two-dimensional graph displaying true positive versus false positive rates *for all possible cutoff values*, see [Fawcett \(2004\)](#) for an accessible introduction. For additional insight, Figure 11 displays an ROC curve based on our estimates, but only for a single operator.

Table 12: Confusion Matrix in a binary classification setting.

	Data: positive	Data: negative
Model: positive	True Positive (TP)	False Positive (FP)
Model: negative	False Negative (FN)	True Negative (TN)
Column Totals	P	N




Figure 11: ROC curve choice of GoAhead vs all other alternatives.

Notes: The plot decreases the cutoff level from 1 to 0 as one traverses the figure from right to left. For example, in the top right corner, and a cutoff level equal to 0, the model achieves a 100% true positive rate, because each prediction will be greater than zero. Subsequent larger cutoff values imply different performance along those lines. The 45 degree line is equivalent to randomly choosing an alternative, so should be seen as the minimum achievable performance. ROCs in the top left corner are generally indicating good performance. See [Fawcett \(2004\)](#) for more details.