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

Marleen Marra* Florian Oswald*

September 7, 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. local monopoly rents, transportation costs, and economies of density. 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 confirm that collusive behavior is partly responsible for the observed degree of spatial clustering of firms, beyond the estimated economies of density from operating routes from nearby garages. We estimate an approximate loss to the taxpayer of at least £17m per year, or about 6.2% of the total procurement cost of providing public bus transportation in London.

JEL codes: D44, L41, L92, R41

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

*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, Milena Almagro, Pierre Dubois, and colleagues at Sciences Po for helpful 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 techniques to address such questions.

This paper investigates the role of spatial rents and firm conduct 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 an agreed-upon 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, vehicle specifications, maintenance and cleanliness requirements, and duration of the contract. Hence, an important margin on which bus operators compete is the location of their proprietary bus garages in relation to the route, which to large extent determines their efficiency in providing the requested 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.

While studying this issue, the paper makes three main contributions. First, we collect and assemble a unique dataset on the garage ownership history in the London bus market since privatization in 1994. In doing so, we add significant new elements to what the existing literature, most notably Cantillon and Pesendorfer (2006, 2007), has discovered about this particular market. 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.

Second, in a two-stage model we link the value of a garage to revenues from bus service contracts that operators could service. 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.

Third, 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 garage clustering benefits and local monopoly rents. We find that, on average, the closest garage of an operator needs to be about 3 to 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 that garage.

With the help of our structural estimates we can generate intuition as to whether anti-competitive market sharing agreements exist between firms in the London bus route procurement market. 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. This line of inquiry is motivated by striking patterns of spatial segregation among operators in the London Bus market observed in the data. In line with concerns raised by the UK Competition Commission in 2011, this pattern might suggest the presence of uncompetitive practices associated with market-sharing.[1] Our structural estimates of the garage-operator utility function are needed to separate such an explanation from clustering due to agglomeration benefits in operating routes from multiple near-by garages. To this end, we propose a theory of operator conduct that allows for implicit agreements on market sharing among participants. Specifically, we formalize that a necessary condition for the existence of (implicit) market-sharing agreements is that operators abstain from entering garages located within the "core territories" of other operators due to

[1]The UK Competition Commission uncovered in 2011 that bus operators in various local markets engaged in a practice of avoiding competition with each others 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.

the apprehension of potential punishment.[2] The theory also posits that any changes in garage ownership that do take place must not violate any agreements, if they exist, as that would imply that the punishment is ineffective. With these fundamental concepts in mind, and having estimated the utility derived by operators from owning a garage, we proceed to estimate the counterfactual utility for all garage ownership changes that did *not* occur.

We estimate that in about 11 percent of cases, another operator has a higher latent utility than the incumbent. Subsequently, we assess the connection between the estimated punishments and factors known to increase the likelihood of collusion within a market.[3] When new competitors enter the London bus market, our method assigns low levels of implied punishment: collusion breaks down during such periods. In terms of magnitudes, each new operator 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. Given that the gain from collusion must be at least as high as the forgone profits when deviating, and because profits in this market are directly linked to revenues from route auctions, we are able to compute an estimate for the total efficiency loss from the private ownership of garages and (potentially collusive) unrealized profitable ownership changes. The approximate loss to the UK taxpayer is estimated to be at least £16.4 and up to £24.7 million per year when firms include a mark-up of 10% to 15% in their route auction bid (as in Cantillon and Pesendorfer (2007)). While conservatively based on our estimate of a lower bound on the punishment value, these are sizeable amounts equal to 6.3% – 9.5% of the total bus route procurement cost.

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 the dataset compiled for 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]The specific form of punishment is not specified but should be understood as a means of enforcing compliance, potentially through future route auctions, 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).

[3]Market-sharing collusion is not neccesarily identifiable in data from the bidding stage. 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. In section 7.2.2 we discuss the features of the London bus market that make it prone to collusive behavior, according to collusion-prone market features identified in theory models as summarized in Ivaldi et al. (2007).

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. 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.[4] 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.[5] 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

[4] As 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.

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

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).[6] 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.[7]

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, each weekday.[8] From our data on bus route procurement auctions, we compute an average expenditure on procuring bus operations of £273m annually over the period 2003–2019 (see table 9), paid for by UK taxpayers.[9] Before we introduce the constructed estimation sample with garage ownership over time, we give a brief outline of the history of London's bus system and the route procurement process.

3.1 Brief history of London's public bus system

Historically, many different independent companies owned and operated public transport in London and the UK, providing potentially wasteful duplicate services while competing for

[6] As 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.

[7] We 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)).

[8] Over 120 routes are 24h per day, seven days a week operations. See TfL's London's Bus Contracting and Tendering Process for more details.

[9] About one third of TfL funding comes from ticket fares, Congestion Charge, Ultra Low Emission Zone etc, with the remainder being split into grants funded by business taxation and borrowing (https://tfl.gov.uk/corporate/about-tfl/how-we-work/how-we-are-funded). Assuming for illustrative purposes that 2/3 of the £273m bus operating cost is funded only via business taxation, one finds a contribution to bus services of roughly 262 Pounds per London business, per year (there were 1.04 million businesses registered in London in 2022, see https://researchbriefings.files.parliament.uk/documents/SN06152/SN06152.pdf).

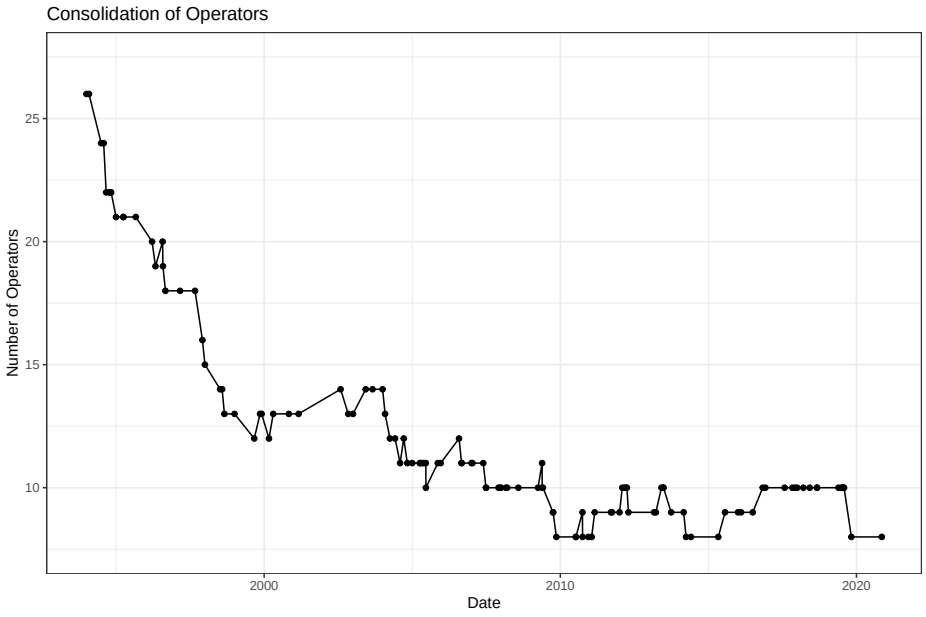


Figure 1: Time series of Ownership consolidation. Over time, through merger and acquisition as well as exit from the market, the number of active operators falls.

passengers.[10] From 1933, the state took an active role in the coordination of services with its successive transport authorities. The National Bus Company was created in 1969 in a wave of nationalizations. It owned a number of locally managed subsidiary companies who were in charge of bus operation. The London Regional Transport Act of 1984 started deregulation of the transport sector, calling for public tendering of bus routes and opening the market for independent operators. London's state-run entities were first divested along geographical lines into 12 companies before being fully privatized in 1994/1995.

Today, the market is dominated by six large (international) transport firms and a few smaller operators. The geographical split of the 12 divisions of London Regional Transport is especially relevant for the analysis in this paper, as it likely causes the current operators to gravitate towards a part of the city. Table 8 documents the current ownership structure of the 12 divisions. The tendering of bus routes allowed a great number of smaller operators to enter the market in or after 1994, which was indeed a design objective of the auction setup, as explained in detail in Cantillon and Pesendorfer (2007). Subsequent consolidation in the market lead to today's situation where a handful of large operator groups provide the bulk of bus services in London, which is illustrated in figure 1.

[10]Incidentally, one company (London General Omnibus Company) painted its buses bright red to stand out and later became the largest operator in London, which is the foundation for the current-day cherry red London bus.

3.2 Route auction data

TfL runs a continuous programme of tendering bus routes, offering contracts for about 15% of the system's routes year year. In practice, an invitation to tender is sent out every 2–4 weeks to a set of pre–approved operators. Winning firms receive their bids as financial compensation, while all ticket fares go to TfL. Contracts typically start 8 to 10 months after the award date, so that winning bidders can reorganise, update their bus fleet if needed, and (in rare cases) obtain the necessary garage space. The tendered contracts fully specify routes, schedules, minimum performance standards, and bus specifications. In practice, this means that bus operators have only few margins of adjustment, and it is likely that none of them are systematically related to quality of service, as all relevant aspects of the operation are contractually fixed. As such, the distance of an operator's garage to the bus route in question can be considered an important determinant of bidding behaviour.[11]

To study the link between the garage network and rents from the route auctions more formally, we obtain tender data from the TfL website.[12] The outcomes of all tenders since 2003 are available, and we gather information about the winning bidder, the contract value, cost per mile of route, and the number of bidders for all 1907 routes tendered tendered between March 13th 2003 and October 19th 2019. Routes are typically tendered for a period of 5 years, with a 2 year extension conditional on meeting certain performance standards, and on average a route shows up 2.5 times in the full sample. In addition, we obtain additional route characteristics from the London Bus Routes website, including each route's peak vehicle requirements (PVR, the number of buses that are needed to drive the schedule specified by TfL), length in miles, and frequency of service. Any missing route information is complemented using data web-scraped from a site for bus hobbyists.[13] We give an overview of 2003–2019 route procurement data in table 9.

3.3 Garage-operator network data

Bus garages can be considered part of the productive assets needed to fulfil bus route contracts (together with the fleet and the drivers). The garages are private property of the bus operator, meaning that the operator is either the freehold owner or long leaseholder (e.g., the government remains the primary owner of the land).[14] We refer to the firm operating buses from the garage

[11] Cantillon and Pesendorfer (2007) estimate that a one percent increase in closest distance between a route's endpoint and the operator's garage increases the bid by 0.68 percent. Their analysis abstracts from possible agglomeration benefits from having multiple garages clustered together as well as competitive effects from the operating garage being isolated in the garage-operator network. We explore these elements explicitly in section 4, while accounting for the geography of the market by using driving time on the road as distance measures.

[12] See https://tfl.gov.uk/forms/13796.aspx.

[13] See https://bus-routes-in-london.fandom.com/wiki/Bus_Routes_in_London_Wiki.

[14] TfL remains the primary owner of the land for the garages Ash Grove, Brixton Tramshed, Edgware, Fulwell, Twickenham, Walworth, West Ham, and Uxbridge, as documented in this Freedom of Information Request

as the garage *owner* —regardless of the freehold or leasehold contract status— and exclude the few cases where a garage is shared between two operators.[15]

Our empirical analysis relies on a comprehensive dataset of the operator and bus garage network in London between 1994-2019, which we compile from various sources. Information about which routes were run from which garage comes from the London Omnibus Traction Society, the UK's largest bus enthusiasts' organization. They provided valuable input by constructing historic data for this study based on their archive of bus schedules. The data that they provide matches the majority of routes tendered since 2005 to the garage that it was operated from.[16] The data also shows when a new garage was put into use or an old garage closed down. Matching this data to the tender data furthermore tells us when a garage changed ownership. Garage locations are obtained from the London Bus Routes website.[17] The website contains a file with – for the majority of garages – full information about the operator currently owning it, it's address, and TfL garage code. Missing addresses were completed manually, and corresponding longitude-latitude information is collected from the Google Maps API. We also obtained geo-coded locations of all bus stops and their sequences for each route from TfL's Open Data Initiative platform.[18]

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-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 and 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.

Our structural estimates are 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.

from 2017: https://tfl.gov.uk/corporate/transparency/freedom-of-information/foi-request-detail?referenceId=FOI-2567-1617.

[15]Our estimation strategy does not rely on the transaction price paid for acquiring the freehold or leasehold rights, nor on the duration of the contract, as explained in section 5.2.3. We failed in our attempts to obtain prices for all transactions since 1994 but we can report a median price of £1.4m for the sale contracts that we found through the UK Land Registry portal.

[16]The matching covers 85% of all routes, see table 2. See http://www.lots.org.uk/ for Society memberships. Data received on October 22nd 2019.

[17]See http://www.londonbusroutes.net/index.htm.

[18]See https://api-portal.tfl.gov.uk/docs.

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

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 2a). The vast majority of network changes also involves just one garage per time (see figure 2b).

3.4 Measurement of Drive Times

A crucial input to our empirical exercise is a precise measure of distance between points of interest in the garage-route network. In particular, simple straight-line distance between points would miss a great deal of nonlinearities arising from geography and, ultimately, the shape of the London road network. We therefore compute optimal drive times between all bus garages and between all garages and all 55,578 bus stops on the road network, relying on the Open Source Routing Machine (OSRM).[19] The resulting drive time measures are used in later sections to capture in a flexible way the location of a garage in relation to bus routes and other garages. A visual illustration of the distance a bus (or car) can cover on London's route network in a given amount of time is provided via so-called *isochrones* in figure 9.

4 The importance of the garage network: descriptive evidence

This section provides some initial insight into the garage-operator network and its relation to bidding in route auctions.

4.1 Spatial clustering of garages

Figure 3 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.

[19] In terms of functionality, this is identical to the google maps routes API, with the exception of information of congestion of roads at different times of the day. The benefit of OSRM is that it can be run for free on a local machine, so that computing a great number of optimal routes is feasible. The underlying maps are from OpenStreetMap. We compute 6.5 million optimal routes, which for the basic request on the google maps routes API (USD 5 per 1000 Queries – one route is one query) would cost USD 32,000.

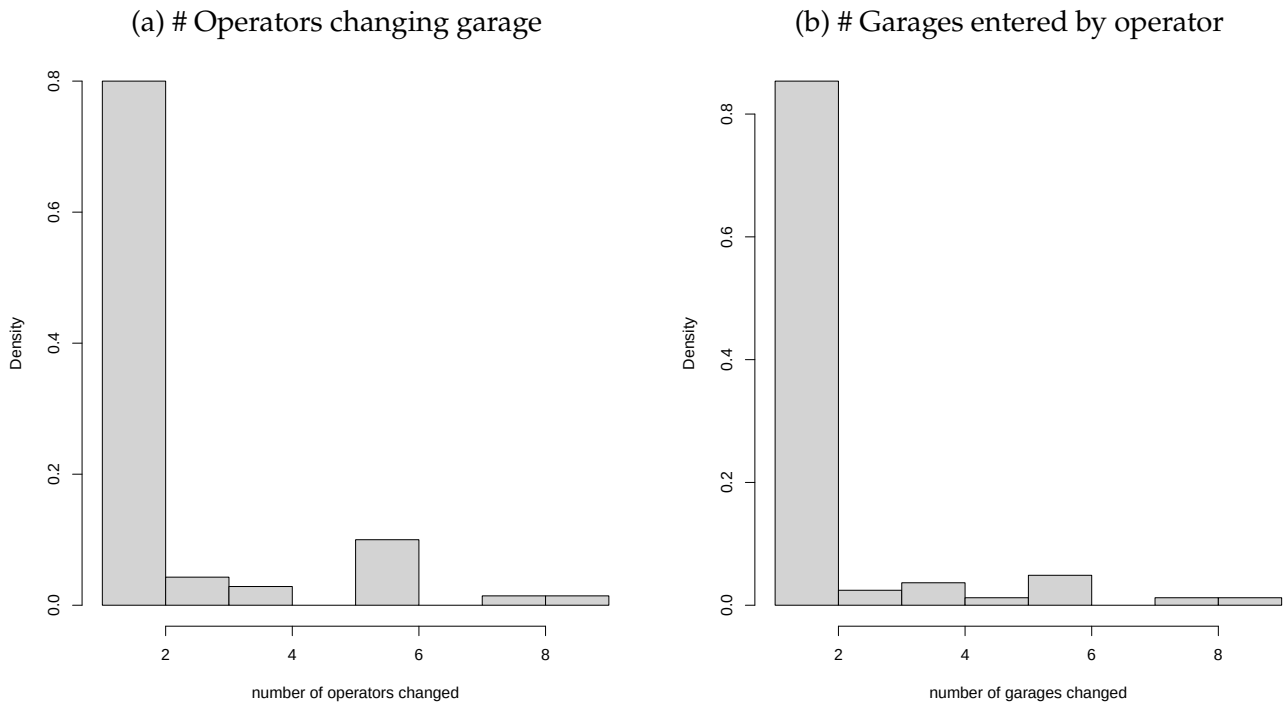


Figure 2: Garage network changes. Number of operators entering a garage (a) and number of garages entered by operator (b) on the same date. Based on all times when operators opened a new garage (the 91 "New entry" events in Table 1) or purchased an existing garage from another operator (the 51 "Ownership 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. To understand what drives this clustering, we rely on a structural model that recovers the relative importance of clustering synergies coming from lowering route operation costs, the value of being isolated from competitors, and the benefit of being close to routes. Furthermore, a counterfactual simulation exercise is designed to assess whether anti-competitive market sharing strategies are a potential driver of the observed clustering.

4.2 Dead miles

The second empirical fact we want to highlight is that operators in general want to minimize the distance between route and garage. This reflects the notion that it is costly to operate a route from a far-away garage, because the bus has to travel empty until revenue-generating operation can start. Conforming to industry practice, we use the term *dead miles* to describe the associated costs.[20]

Our preferred measure of dead miles is the actual driving time on London roads between a

[20] According to https://en.wikipedia.org/wiki/Dead_mileage: *"Dead mileage, dead running, light running, empty cars or deadheading in public transport [...] is when a revenue-gaining vehicle operates without carrying or accepting passengers, such as when coming from a garage to begin its first trip of the day."*

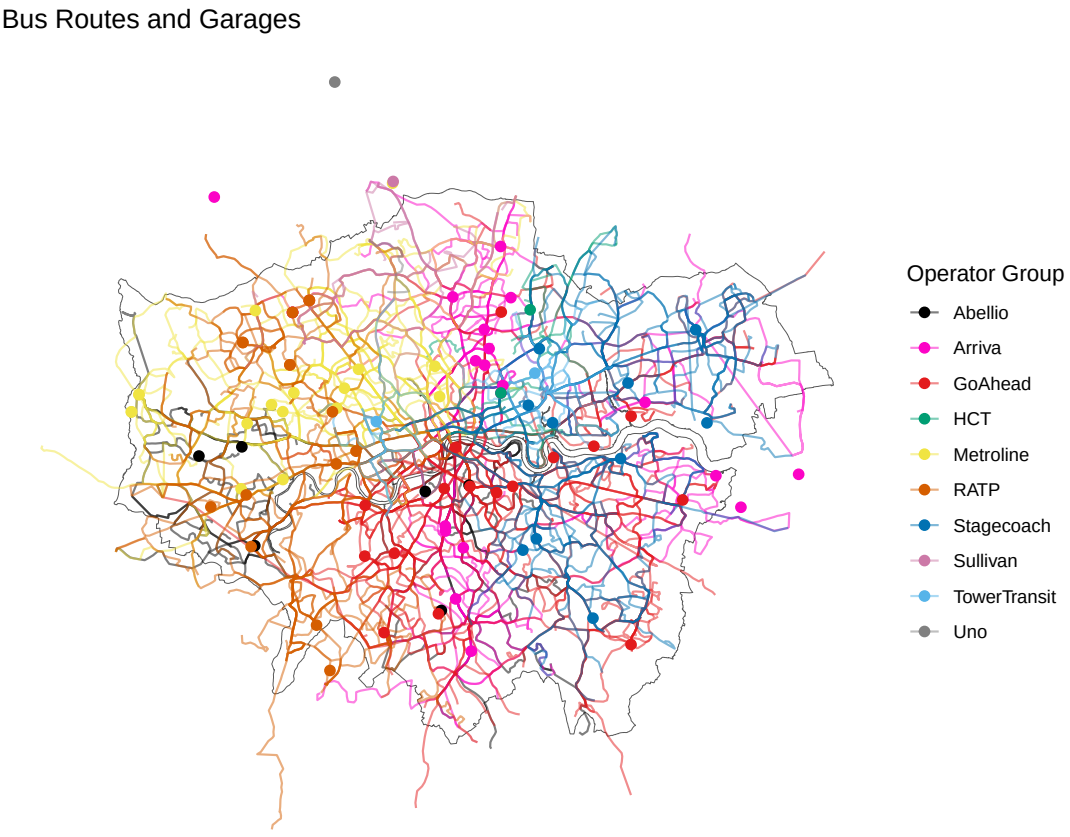


Figure 3: Route network and garage locations by operator group at end of the sample (2019). Each point denotes the location of a garage, an a line segment in a certain color indicates route ownership.

garage and the closest endpoint of the route (see section 3.4 on the construction of these drive time measures). The distribution of this dead miles measure for all garages to all routes in the auction data, regardless of whether the route is operated from that garage, is given in figure 10a. The average is around 40 minutes. Clearly, the realized dead miles that need to be covered —once a route is won and assigned to a garage of the winning operator— are much lower. On average, about 13 minutes (8 km) are actually driven empty, each time a bus is moved between the garage and the closest start or endpoint of the route, as shown in table 2. On average, we measure one dead mile for each mile of bus route. The total dead mile costs per route accumulate because multiple vehicles are employed on each route (to adhere to TfL's schedule requirements in both directions) and buses also return to the garage to facilitate meal breaks and shift changes of the drivers.[21]

[21]The average PVR (or *Peak Vehicle Requirement*, measuring the number of buses that is needed to operate the route at the busiest time - this is a contractual obligation for each route) for the routes in our auction data is 12. We also employ other measures of dead miles in our analysis, including the drive time between the garage and the closest big bus stop (so-called "stopover point", which we define as having at least 5 intersecting routes) and the average drive time to both endpoints of the route.

Table 2: Summary statistics of variables used in regression 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.

	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)	14.75	9.98	1.70	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			

4.3 Link between garage network and revenue from route operation

We use our linked route-garage network dataset to explore whether garage location affects the operator's expected route revenues. Specifically, we regress the cost per mile of the winning operator on route characteristics, the number of bidders, and our dead miles metric. The results are reported in table 3. It shows that the relationship between the size of the winning bid for a tendered route, and the location of the garage where the route is subsequently operated from. The outcome variable is the accepted bid measured in million pound sterling, and we control by default for route length, because it obviously determines operating costs on the route. The average bid conditional on dead miles and route length is £1.36m (Constant in column (1)). Depending on whether we additionally control for year and winning operator fixed effects, number of bidders and PVR, the estimated impact of an additional minute of dead mile driving time from operating garage to start or end of the route lies between £72,000 (column (1)) and £17,000 (column (4)), with goodness of fit increasing continuously from (1) through (4) to yield a final R^2 of 0.76. In other words, after we control for exogenous route (length, PVR) and market characteristics (number of bidders), the location of the garage – and thus the time it takes to drive on London's road network to the starting point of the tendered route – seems to have an economically important impact.[22]

Column (5) assesses the presence of potential agglomeration benefits and competitive consid-

[22]We repeat the same analysis for different measures of dead miles in the appendix. In table 10 we look at the impact of dead mile drivetime as an average to start point of the route, while in table 11 we investigate the relationship with respect to the average distance to *stopover bus stops*, i.e. bus stops where routes intersect, which can be an operational advantage to exchange drivers on the route. The dead mile effect is strongest for the average of start and stop points, otherwise the results are fairly similar.

	Accepted Bid (in Million Pounds Sterling)				
	(1)	(2)	(3)	(4)	(5)
Dead Miles Start-Stop (Minutes)	0.072***	0.079***	0.086***	0.017**	0.003
	(0.012)	(0.011)	(0.011)	(0.006)	(0.007)
Route Length	0.080***	0.080***	0.067**	-0.025*	-0.023+
	(0.023)	(0.022)	(0.022)	(0.012)	(0.013)
Number of Bidders				-0.140***	-0.114***
				(0.030)	(0.033)
Peak Vehicle Requirement (PVR)				0.233***	0.223***
				(0.005)	(0.005)
Constant	1.362***	0.670*	0.400	-0.568**	-0.790**
	(0.218)	(0.287)	(0.364)	(0.219)	(0.255)
Closest Own Garage (mins)					-0.008*
					(0.004)
Dead Miles of Closest Competitor (mins)					0.041***
					(0.008)
Year FE	-	✓	✓	✓	✓
Winner FE	-	-	✓	✓	✓
Num.Obs.	1126	1126	1126	1126	940
R2	0.053	0.145	0.191	0.762	0.789
RMSE	1.98	1.88	1.83	0.99	0.95

Table 3: Explaining the size of accepted bids for London bus routes.

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, which explains why the sample contains 1126 instead of all 1907 auctions, as shown in table 2. Statistical significance: + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

erations. We include 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), as well as the dead miles of the closest garage of a competitor of the winning bidder (competitive pressure). These variables capture additional variation (the R^2 increases to 0.79) but eliminate the dead miles effect found above. The estimates also point to a small and negative effect of the winning bidder having garages spatially clustered. However, the interpretation of these two variables is complicated as it might reflect both cost savings and adjustments of the equilibrium mark-up. By contrast, the fact that operators benefit from being isolated from their competitors is crystal clear. The winning bid is £41,000 pounds higher for each minute more drive time between the garage where the route is operated from and a garage of a competing bidder, and 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 derive with it, in order to estimate the determinants of these spatial rents from observed location choices. Specifically, in our empirical context the garage network is formed in the first stage and in the second stage operators bid on bus route contracts given this network. Firms are myopic and learn their own idiosyncratic route operation costs at the beginning of the period (in stage 1). 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, akin to the myopic *alternating selection mechanism* of dividing assets that McAfee (1992) show can result in an efficient allocation. 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.

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.[23]

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 $\text{RouteSpecificCosts}_{rt}$ 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

[23] 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.

garages clustered together is likely to lower the marginal cost, which is captured by the variable $\text{DensityGarages}_{it}$. Combined, the cost function contains

$$\begin{aligned} c_{irt} &= \text{RouteSpecificCosts}_{rt} + \text{DeadMiles}_{irt} + \text{DensityGarages}_{it} + \nu_{irt} \\ &\equiv \delta_{irt} + \nu_{irt} \end{aligned} \tag{1}$$

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

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

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 is 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))

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

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.

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

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 it's 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 at time t (π_{ijt}) can be specified as the sum over all routes of expected revenues generated through route procurements:

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

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)).[24] Tying the garage value function to the expected auction revenues from the second stage, and to elements

[24]Garages 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.

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.[25] Specifically, let π_{ijt} be

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

with $\epsilon_{ijt} \perp \left(\Gamma_j^X, \Gamma_{ijt}^C, \Gamma_{ijt}^O, X_{jt}\right)$, and Γ_j^X, Γ_{ijt}^C, Γ_{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.[26] Bounded by a low number of observations, the dimensionality of Γ_j^X, Γ_{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 mentioned in section 3.4, 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}.

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

[25] The 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.

[26] The 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.

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.[27] Specifically, the equilibrium price for garage j at time t can be described as a function of all operators' realizations of π_{ijt}, as:

$$p^*_{jt} = Y_{jt}^{N_t-1:N_t}, \tag{7}$$

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.[28] 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 $i = 0, \ldots, N_t - 1$ on the basis of maximizing $\pi_{ijt} - p^*_{jt}$.[29]

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 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

[27] The 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.

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

[29] As 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.

mean utility of choosing the outside option is normalised to 0: $\pi(\Gamma_j^X, \Gamma_{0jt}^C, \Gamma_{0jt}^O, X_{jt}) = 0 \; \forall j, t$.

Further, 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 2). 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. 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 $\Gamma_{ijt}^C \; \forall i$). Essentially, 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.[30]

Framing the problem as sequential choice, permitted in our case as we observe garage ownership changes over time, 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.[31] 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 garage network changes (see figure 2b).

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.[32]

[30] 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).

[31] To 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.

[32] For 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 4: Summary of variables used for the estimation of parameters in the reduced form of garage-operator utility (6). Estimation results in table 5.

	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: $\Gamma_j^O = 0$ 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.

When also assuming that the idiosyncratic value component $\epsilon_{ijt} \sim^{i.i.d.} F_\epsilon$, where F_ϵ is Weibull, 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*

Table 5: Estimation results for the reduced form of garage-operator utility in equation (6).

	(1)	(2)	(3)	(4)	(5)
No. bus route starts within 10 min. drive time	0.029	0.106***	−0.023	0.066*	0.030
	(0.025)	(0.021)	(0.080)	(0.033)	(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***		−0.210***	−0.194***	−0.197***
	(0.042)		(0.048)	(0.043)	(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***	−0.547***	−0.596**	−0.730***	−0.654***
	(0.189)	(0.160)	(0.192)	(0.195)	(0.198)
No. garages owned			−0.152**		
			(0.054)		
Incumbency benefit $[exp(-T_{jt})]$				2.077***	2.625***
				(0.526)	(0.667)
Garage is not vacant					1.795+
					(0.923)
Inclusive value	2.551***	1.416**	2.722***	3.250***	2.737***
	(0.568)	(0.447)	(0.629)	(0.638)	(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 12 in the appendix. The data contains 192 ownership changes and in total 5,184 rows of the number of alternatives (operators) × 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 cofficients 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.

in 2003 and 2004 (see figure 11a).

We start by giving a descriptive overview of the variables used in estimation in table 4. 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.[33] 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 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 4, 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 5. 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.[34] 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

[33] We illustrate the full distribution of drivetimes for two example garages in figure 8.

[34] A 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 12 in the appendix.

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.[35] 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.

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).[36]

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 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

[35]However, 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.

[36]These results refers to the row labelled *Own garage proximity (mins) offsets extra comp. garage* in table 5.

of entry and are not, like p^*_{jt} pocketed by the previous owner.[37] 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.[38] 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.

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 5. 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 5 and 12).

[37] Recall 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.

[38] The 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 table 14 in the appendix).

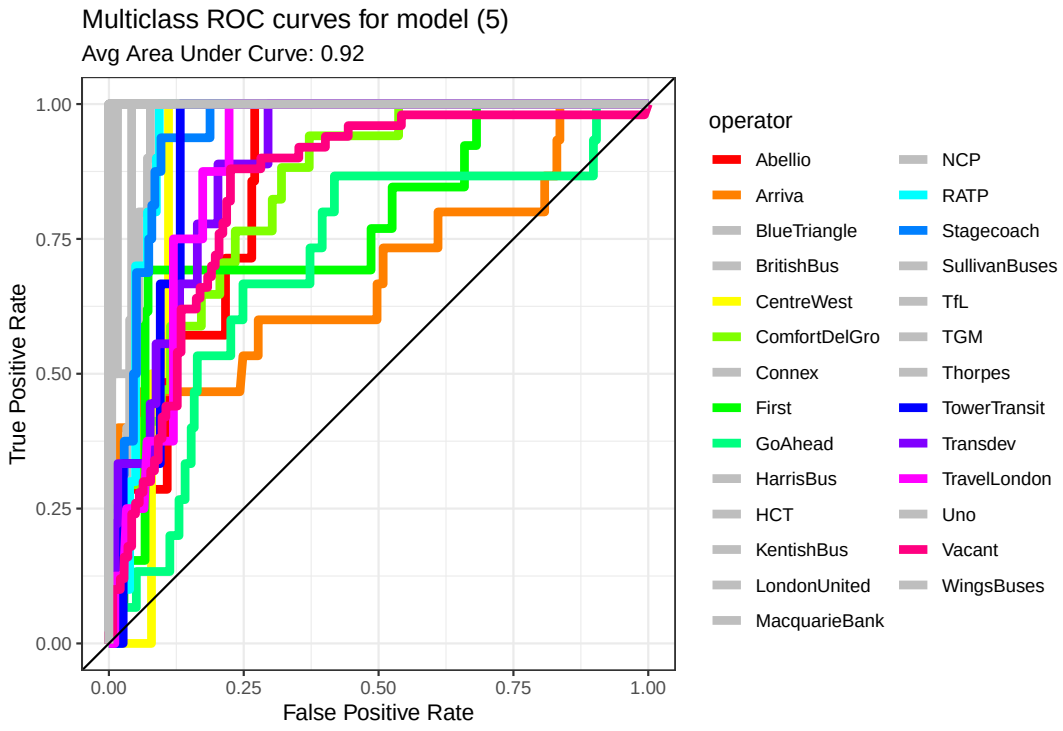


Figure 4: ROC curves for selected model (corresponding to column 5 of table 5).

Figure 4 shows the ROC curves for all operators separately for our preferred specification. The in-sample performance is generally better for the smaller operators compared to the Abellio, RATP, GoAhead, and other large operators. 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 12). 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 5 (for nested logit) and 12 (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, includ-

ing 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.[39] 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.

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 Overview of results

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 π_{ijt} based on our model estimates $\forall(t, j \in \mathcal{J}^N, i \in \mathcal{N}_t)$, which gives the expected utility for all operator-garage-changedate triples where no ownership change took place. The utils difference between the utility of the highest-utility operator (i) and the incumbent (I_{jt}) is taken to be the lost value from an unrealized transaction, when positive. After also taking out 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 about 10 percent of these cases, another operator would benefit more from having that garage in their own network than the incumbent does. Figure 13 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

[39] McFadden 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).

results reported below are not driven by, for instance, a fringe operator that the model predicts should be a large player instead.

7.2 Interpretation as collusion

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.[40] 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 form an infinitely repeated version of the game. However, we focus on necessary but not sufficient conditions for collusion. A key idea that we rely on is that incentive-compatibility of such firms is guaranteed by the punishment inflicted upon the firm deviating from it. 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. After formalizing the concepts of punishment and collusion in this setting, we analyze to what extent the estimated punishments relate to theoretical predictions about when collusion is more likely.

7.2.1 Conceptual framework

The following assumptions characterise necessary conditions for collusion in this setting, where ψ_{ijt} denotes the punishment inflicted upon firm i by I_{jt}, the incumbent of garage j at time t, when i deviates from the agreement.

Assumption 1 (Punishment is effective) *If firm $i \neq I_{jt}$ enters garage j at t, $\psi_{ijt} = 0$*

Assumption 2 (Free entry absent punishment) *Firm i' enters garage j at t iff $\{\pi_{i'jt} > \pi_{ijt}\ \forall i \in \mathcal{N}_t$ and $\psi_{i'jt} = 0\}$*

[40] We 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).

Assumption 1 rules out that there is some (ineffective) punishment inflicted upon an entrant that is not high enough to deter the firm from entering.[41] It implies that when we observe a firm entering into a garage, the utility of entering that garage for that firm is fully described by the (competitive) π_{ijt}. Assumption 2 is based on the idea that, without market-sharing agreement between the incumbent $i = I_{jt}$ and firm i', i would sell garage j at t to i' as long as $i's$ bid exceeds the incumbent's value. 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. These assumptions are made to interpret the behavior of operators entering into locations. An incumbent is always free to sell the garage if he chooses so. In this setting, collusion determines whether or not other firms purchase the garage when it would be in their best interest (absent punishment). This is formalized below.

Definition 1 (Collusion (necessary condition)) *A necessary condition for a collusive market-sharing agreement between firm i and I_{jt} that involves garage j at time t is that, despite $\pi_{ijt} > \pi_{i'jt}\ \forall i' \neq i \in \mathcal{N}_t$, firm i does not enter into garage j at t.*

Definition 1 spells out that in the presence of market-sharing agreements, some individually profitable transactions do not occur due to the threat of punishment. 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).[42] Moreover, the value of an unrealized profitable entry bounds the punishment value associated with violating an implicit or explicit market-sharing agreement, as formalized in Definition 2.

Definition 2 (Punishment (lower bound)) *If firms i and I_{jt} collude at t, a lower bound on the punishment value for violating the agrement equals*

$$\begin{aligned}\psi_{ijt} &= \max(0, (\{\pi_{ijt}\}_{i \neq I_{jt}} - p^*_{jt}) - (\{\pi_{ijt}\}_{i \neq I_{jt}} - p^*_{jt})) \\ &= \max(0, \pi_{i \neq I_{jt}, jt} - \pi_{i = I_{jt}, jt}). \end{aligned} \tag{8}$$

If firms i and I_{jt} do not collude at t, $\psi_{ijt} = 0$. Also, $\psi_{ijt} = 0$ for $i = I_{jt}$.

[41] In our setting it is natural to assume away imperfect information about the strategies played by firms, which abstracts from the main reason why there are punishments on the equilibrium path in Green and Porter (1984).

[42] Belleflamme and Bloch (2004) model firms that 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.

Table 6: Regressions of punishment on the Number of Alternatives

	$\psi : \psi > 0$	$\psi : \psi > 0$	$1(\psi > 0)$	$1(\psi > 0)$
(Intercept)	2.108***	2.967***	−1.377***	−1.068*
	(0.068)	(0.748)	(0.044)	(0.453)
No. alternatives	−0.002	−0.041	−0.058***	−0.059**
	(0.006)	(0.035)	(0.004)	(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.

7.2.2 Relation to theory

We next use the estimates of π_{ijt} $\forall ijt \in \mathcal{J}^N$ to construct estimates of the implied punishment values, according to the coceptual framework described in the previous section and as consistent with the collusive network formation model of Belleflamme and Bloch (2004). As such, ψ_{ijt} is defined by the piecewise function in (8) as the lower bound on the expected punishment to operator i by I_{jt}. It must be recognized that this estimated "punishment" 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.

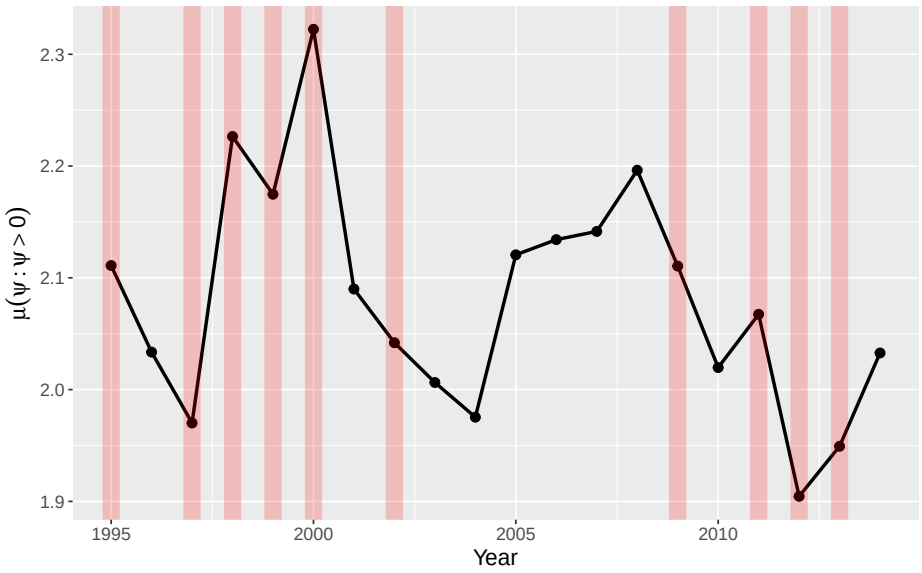
We evaluate our counterfactuals against the characteristics that are known to make an industry more collusion-prone in theory, as summarized in Marc Ivaldi, Bruno Jullien, Patrick Rey, Paul Seabright and Jean Tirole (2007) (henceforth: IJRST). Several of these characteristics reflect the market in ways that are fixed over time and across space and so it cannot be verified whether they relate statstically to the estimated ψ's. However, the majority of those characteristics arguably apply to the market in question.[43]

[43]In what follows, the text in quotes paraphrases other collusion-prone characteristics highlighted in IJRST. TfL has held procurement auctions about every two weeks since 1995 and provides detailed information about when which routes are being auctioned off and what the route characteristics are. Furthermore, it is clearly public information who operates which routes and who owns which garage, and TfL also publishes details about the winning and losing bids on its website. This already ticks off "frequent interaction" and "market transparency". As mentioned in the introduction, the call for tenders specifies in great detail what the route operation service should consist of. This is important as holding the route operation details fixed there is limited scope for cost asymmetries in this market, and it is known that "cost asymmetries hinder collusion". Moreover, with the route (and bus-) characteristics specified by TfL, so the firms are essentially providing a homogeneous service. To the extent that there are no perceived quality differences by TfL, the additional transparency that the "lack of product

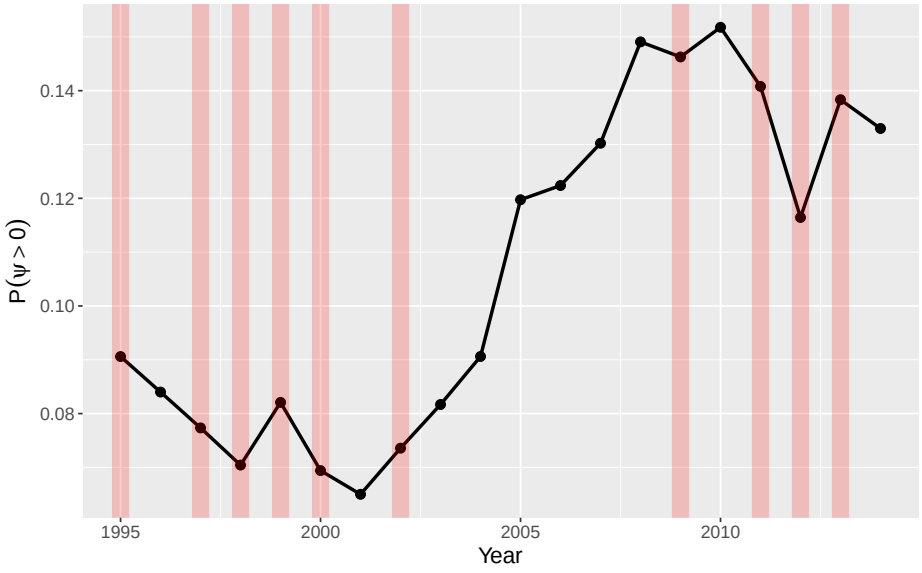
Figure 5: Punishment and relation to firm entry

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).

(a) mean $\psi(\psi > 0)$



(b) $1(\psi > 0)$



differentiation" and the "lack of drastic innovation" provides should facilitate collusion. There are also "no demand fluctuations" as the route network is fixed over time and will require service also in the future, and as revenues are determined solely by the final procurement price (all bus fares go to TfL). Not all identified characteristics apply unambiguously. We have no way of knowing whether "firms put sufficient weight on future profits". Without estimating a structural auction model, we also cannot identify whether profits increase over time, which is relevant as collusion is easier to sustain when "today's profits are small compared to tomorrow's ones". One can observe an increasing trend since 2006 in the winning bid normalized by the number of buses needed to operate the route — e.g. the peak vehicle requirement (PVR), as displayed in figure 14, although this measure is clearly affected by the firms' conduct. Also, while the larger operators compete for routes in other markets too, this "multi-market

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 6). The negative coefficient on "No. alternatives" confirms that periods with more firms are associated 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. 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.

contact" does not apply to the smaller firms. There are asymmetries in capacity constraints in terms of the number of garages that operators own, which is a market feature that hinders collusion by curtailing the power of capacity-constrained firms to retaliate. Finally, we find that higher punishment (measured in terms of the share of garages owned by an operator) also has a weakly positive association with more concentration in our data (see table 15 in the appendix), but the interpretation of this relationship is unclear. As noted by IJRST the relative size of firms is largely endogenous when they collude.

7.3 Efficiency loss

Theory also tells us that the gains from collusion must outweigh the benefits from deviating and triggering a punishment (e.g., IJRST, Igami and Sugaya (2022)). In our setting gains from collusion are losses to the taxpayer because the value of a garage is equal to the sums of profits the garage generates through running TfL route contracts from it. Note 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 damages. The loss in consumer surplus would be obtained by diving the estimated punishment by the coefficient on the transaction price when included in π_{ijt}, and summing total monetary punishments over time.

Following the same logic we can approximate the loss in surplus despite not observing garage transaction prices, as the estimated (minimum) punishment ($\hat{\psi}_{ijt}$) and incumbent value ($\hat{\pi}_{I_{jt}jt}$) are in the same units and route auction revenues are known. Specifically, we compute the average of the punishment share ($\frac{\hat{\psi}_{ijt}}{\hat{\pi}_{I_{jt}jt}}$) across ijt combinations in the entire sample between 1995 and 2019, excluding vacant garages and cases where "the market" is predicted to have the highest utility. We excluded cases where $\frac{\hat{\psi}_{ijt}}{\hat{\pi}_{I_{jt}jt}}$ exceeds the 95th percentile of this range, and negative estimated incumbent values. These measures are taken so that the overall average punishment share does not rely on outliers, and all go in the direction of underestimating the effect on consumer surplus. The punishment share is based on the estimated $\hat{\psi}_{ijt}$ for the potential entrant with the highest utility for a given garage j and *change-date* t, as this would be the firm that would have entered absent punishment.

The average punishment share is 0.63 between 1995 and 2019 in the sample where the incumbent stayed in the garage, and 0.62 is the unconditional average punishment adjusted by the probability that the incumbent stays. This number is stable over time. When restricting the time period to 2004-2019 —matching the tender data— the unconditional punishment share equals 0.63. Due to the tight link between the value of the garage and revenues from the route auctions, the additional rents extracted by the operators from the route auctions must therefore be at least 63 percent higher than without market-sharing. 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). This amounts to an approximate loss to the taxpayer of at least £17m and up to £25.4m per year, or about 6.2% to 9.3% of the total procurement cost of providing public bus transportation in London.

Table 7: 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***	1.592***	−2.907***	−3.266***	0.893***	0.873***	−4.728***	−4.926***
	(0.036)	(0.093)	(0.024)	(0.058)	(0.093)	(0.127)	(0.060)	(0.079)
Min. dist. to competitors (km)	0.083***	0.085***	0.181***	0.193***				
	(0.005)	(0.005)	(0.004)	(0.004)				
Avg. dist. to competitors (km)					0.042***	0.044***	0.101***	0.102***
					(0.003)	(0.003)	(0.002)	(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.

8 Conclusions

This paper studies spatial rents and firm conduct in the London bus market. To do so, we assemble what is the first consolidated 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.

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). As such, some garage locations are more desirable than others, and the degree of desirability differs for each operator, because it depends on where their other garages are located at any given point in time. In a nutshell, we use observed changes in garage ownership in this context to estimate a model of operator garage choice, which we use to predict a match surplus for each operator-garage pair in each moment of time.

We show that 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 assess whether collusive behavior can partly explain the large degree of observed spatial clustering of operators in certain parts of the city, above and beyond the estimated economies of density, and perhaps promoted by the privatization of the market along geographical lines in 1994. We estimate implied punishment values consistent with necessary conditions for collusive market-sharing agreements between operators. Positive punishment values are obtained when observing operators in garages that are not predicted to have the highest utility for the garage by the model. 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, Massachussets: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.

Igami, Mitsuru, and Takuo Sugaya. 2022. "Measuring the incentive to collude: the vitamin cartels, 1990–99." *The Review of Economic Studies*, 89(3): 1460–1494.

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.

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.

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*.

A Additional tables and figures

Table 8: Privatization of 12 divisions London Regional Transport

Division	Area	Original buyer	Current owner
CentreWest	West	management	Metroline & Tower Transit
East London	East	Stagecoach	Stagecoach
Leaside	River Lea	Cowie Group	Arriva
London Central	South Central	Go-Ahead	Go-Ahead
London General	South West	management	Go-Ahead
London Northern	North	MTL	Metroline
London United	South West	management	RATP
Metroline	North West	management	Metroline
Selkent	South East	Stagecoach	Stagecoach
South London	South	Cowie Group	Arriva
Westlink	Kingston	employees	RATP
London Coaches	Central	management	Arriva

Source for original buyer and current (2019) owner: this Wikipedia page, and the area denomination obtained from: this Fandom page.

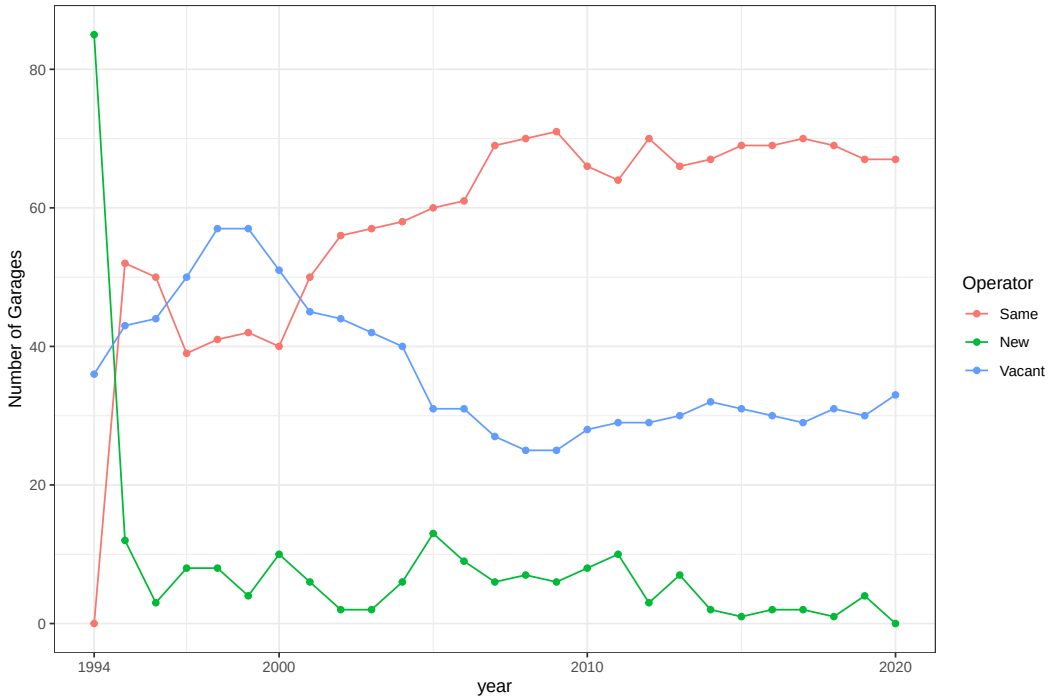


Figure 6: Time series of Ownership Changes in garage data. We call *empty* garages those which are vacant, abandoned or not yet built at a certain point in time. *New* refers to a change of operator in a certain garage.

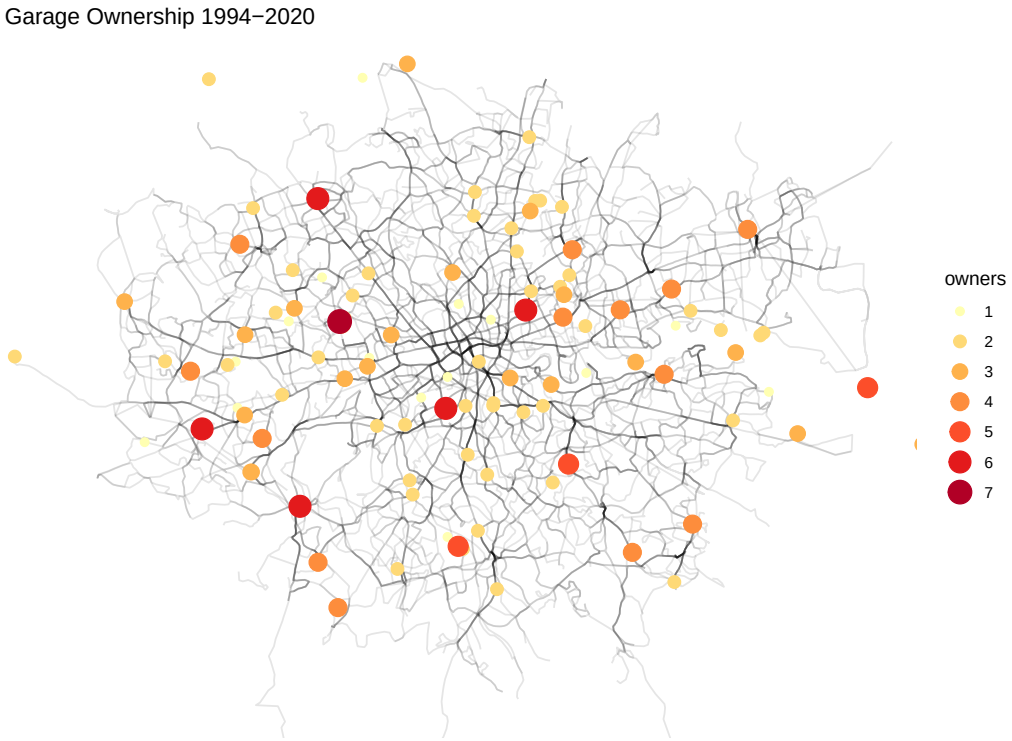


Figure 7: This map illustrates that different garages experience a different number of owners over time, in part depending on their location within the route network. Darker road segments imply a greater number or routes.

Table 9: Winning operator groups and auction characteristics from bus route procurement auctions data 2003–2019. Monetary values are given in million £

Winning Group	Total Earnings 2003-2019 (m£)	Routes Awared 2003-2019
GoAhead	1,205.7	441
Arriva	1,069.7	338
Stagecoach	854.0	318
Metroline	765.9	237
RATP	449.8	191
Abellio	403.3	138
First	382.0	149
TowerTransit	124.0	25
HCT	49.8	30
NCP	12.8	7
SullivanBuses	12.7	26
TGM	6.4	5
Connex	2.9	1
Uno	0.5	1
Routes auctioned / year	112.2	
Number of Bidders / year	312.6	
Bidders per Auction / year	2.9	
Median Revenue / year (m£)	2.3	
Total Revenue / year (m£)	273.0	
Total Revenue 2003-2019 (m£)	5,339.5	

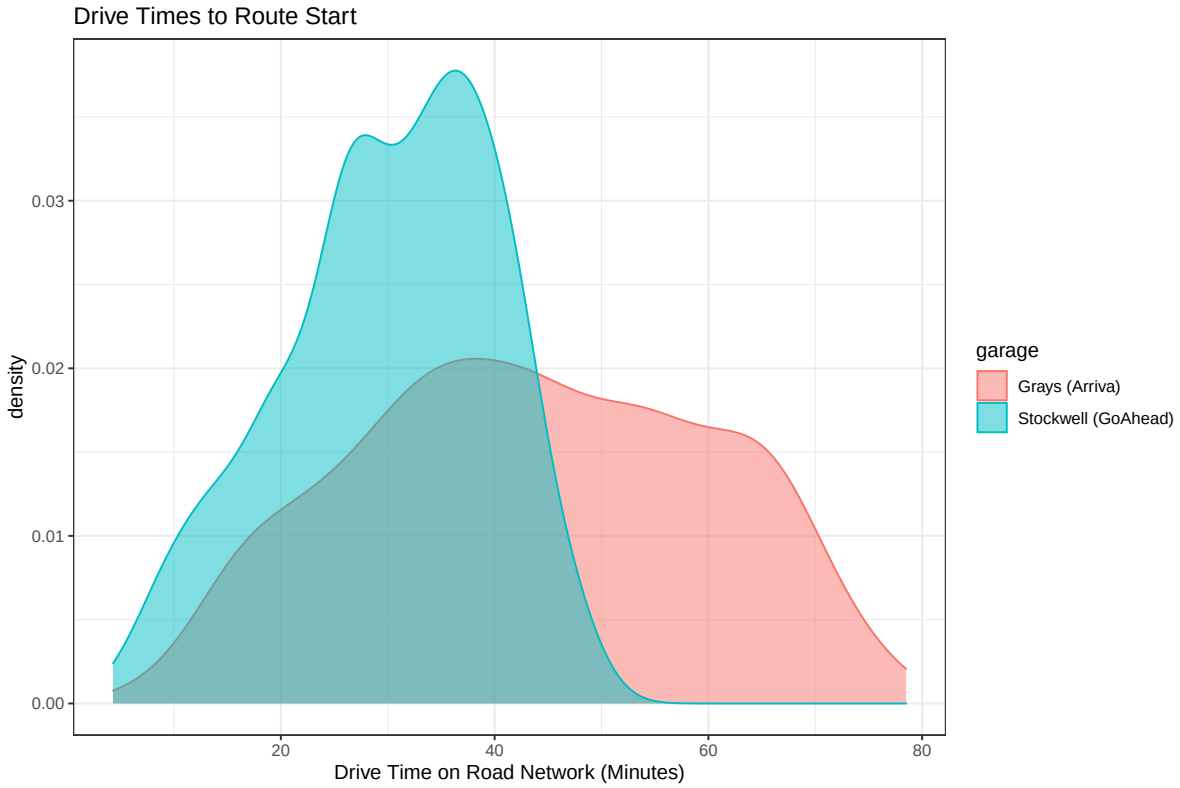
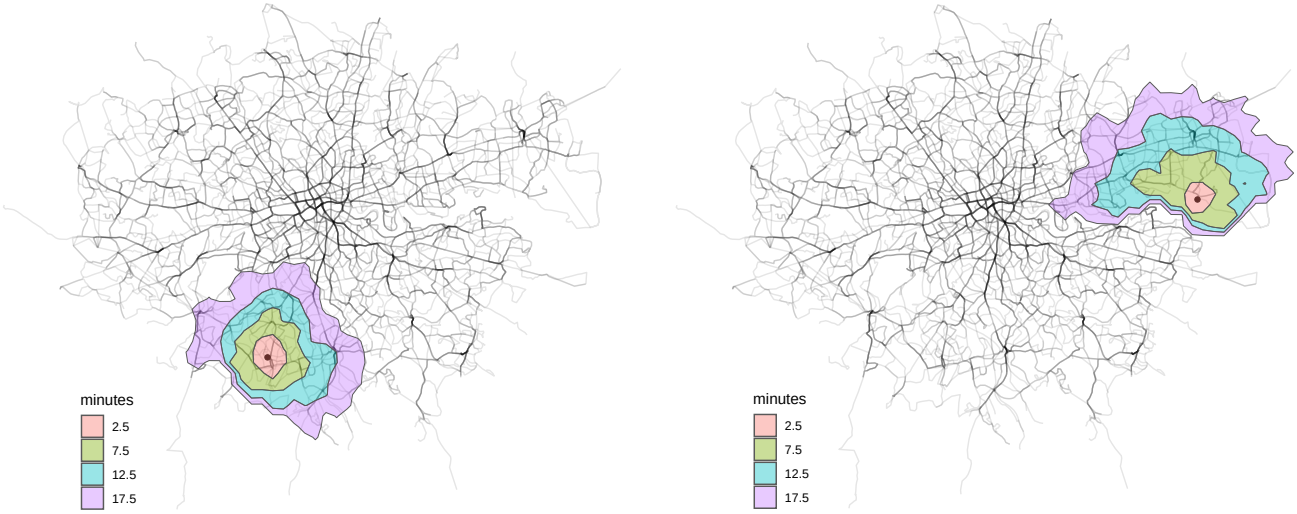


Figure 8: Illustration of our measure of driving time on London's road network from two garages to all possible bus route starting points. The Stockwell garage is centrally located, hence has shorter driving times to starting points of bus routes, whereas the Grays garage is the most easterly located garage in the network, hence exhibits considerably longer driving times to route starting points. The figure shows *all potential* drive times, not realized ones in the bidding process.

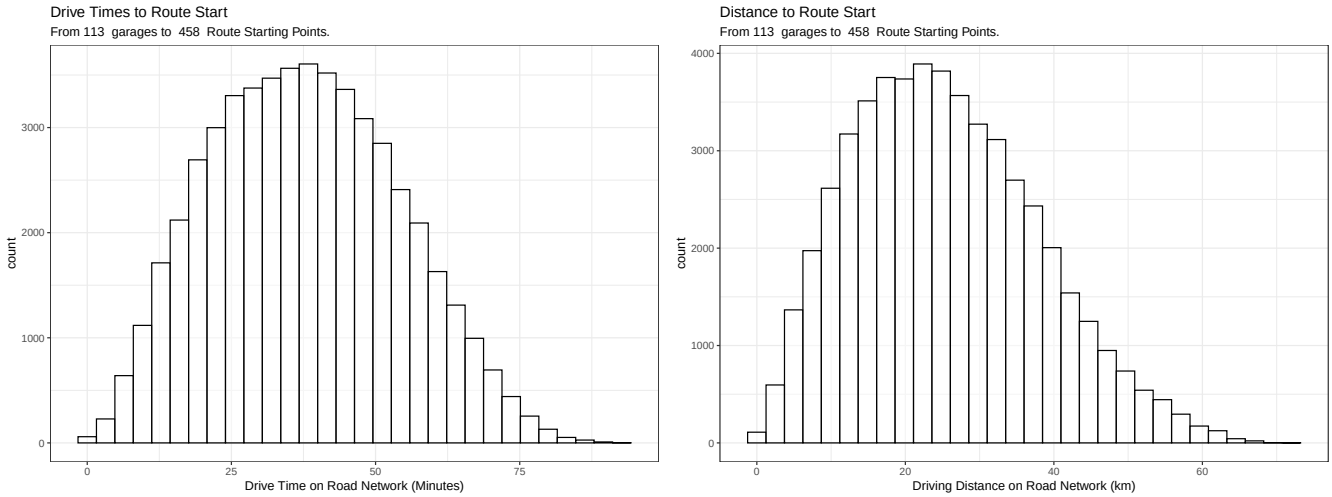
Figure 9: Isochrones for two example garages. An Isochrone connects all points reachable within a given time threshold on London's route network by bus/car. The solid black point represents the location of the garage.



(a) Isochrones for garage *Sutton*.

(b) Isochrones for garage *Rainham*.

Figure 10: Distribution of Driving Time and Distance from all garages to all potential route starting points



(a) Distribution of drive.

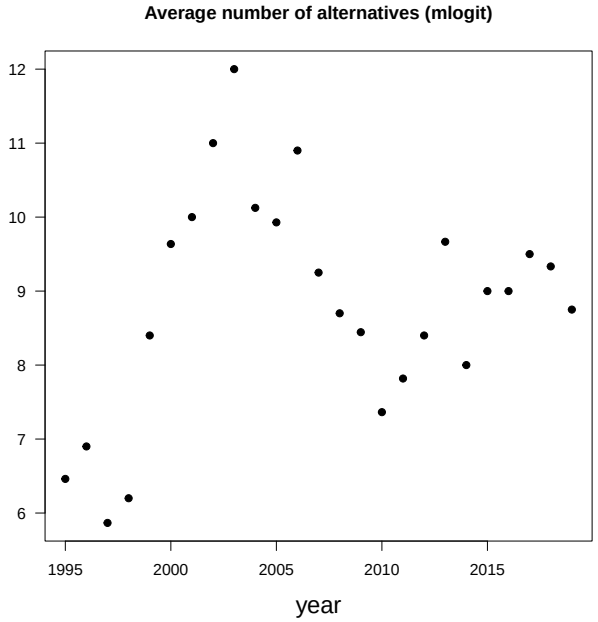
(b) Similar but for driving *distance* instead of time.

Table 10: Dead Mile Regressions based on Start (Minutes)

	Accepted Bid (in Million Pounds Sterling)				
	(1)	(2)	(3)	(4)	(5)
Dead Miles Start (Minutes)	0.009	0.014+	0.018*	0.006	0.0006
	(0.007)	(0.007)	(0.007)	(0.004)	(0.005)
Route Length	0.108***	0.115***	0.107***	-0.006	-0.023+
	(0.022)	(0.021)	(0.021)	(0.012)	(0.013)
Number of Bidders				-0.145***	-0.114***
				(0.029)	(0.033)
Peak Vehicle Requirement (PVR)				0.240***	0.223***
				(0.004)	(0.005)
Constant	2.119***	1.313***	1.098***	-0.715***	-0.786**
	(0.195)	(0.265)	(0.324)	(0.202)	(0.256)
Closest Own Garage (mins)					-0.008*
					(0.004)
Dead Miles of Closest Competitor (mins)					0.042***
					(0.008)
Year FE	-	✓	✓	✓	✓
Winner FE	-	-	✓	✓	✓
Num.Obs.	1476	1476	1476	1476	940
R2	0.018	0.103	0.152	0.746	0.789
RMSE	2.12	2.03	1.97	1.08	0.95

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

(a) When owner changed



(b) When owner stayed

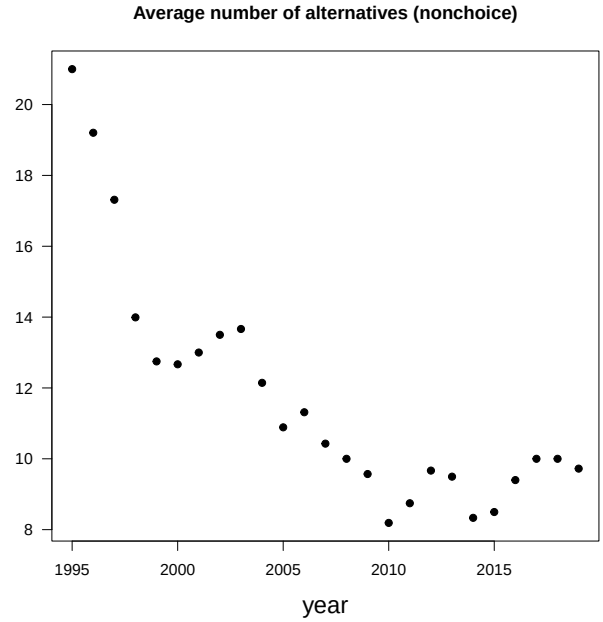


Figure 11: Active operators garage can choose from in *changedate*

Table 11: Dead Mile Regressions based on Stopover (Minutes)

	Accepted Bid (in Million Pounds Sterling)				
	(1)	(2)	(3)	(4)	(5)
Dead Miles Stopover (Minutes)	0.038***	0.046***	0.048***	0.014**	0.008
	(0.010)	(0.010)	(0.010)	(0.005)	(0.006)
Route Length	0.080***	0.085***	0.080***	-0.012	-0.028*
	(0.024)	(0.023)	(0.023)	(0.013)	(0.014)
Number of Bidders				-0.150***	-0.111**
				(0.032)	(0.037)
Peak Vehicle Requirement (PVR)				0.239***	0.223***
				(0.005)	(0.006)
Constant	2.170***	1.240***	1.154**	-0.745***	-0.854**
	(0.206)	(0.286)	(0.351)	(0.222)	(0.280)
Closest Own Garage (mins)					-0.007+
					(0.004)
Dead Miles of Closest Competitor (mins)					0.042***
					(0.008)
Year FE	-	✓	✓	✓	✓
Winner FE	-	-	✓	✓	✓
Num.Obs.	1333	1333	1333	1333	840
R2	0.025	0.115	0.164	0.740	0.785
RMSE	2.15	2.05	1.99	1.11	0.98

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

Table 12: Estimation results garage-operator utility (logit model)

	(1)	(2)	(3)	(4)	(5)
No. bus route starts within 10 min. drive time	0.055***	0.107***	0.046	0.111***	0.043
	(0.016)	(0.019)	(0.070)	(0.023)	(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***		−0.079***	−0.063***	−0.074***
	(0.008)		(0.009)	(0.008)	(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**	−0.450***	−0.285**	−0.312**	−0.259*
	(0.108)	(0.107)	(0.110)	(0.111)	(0.111)
No. garages owned			−0.063**		
			(0.020)		
Incumbency benefit $[exp(-T_{jt})]$				1.487***	2.378***
				(0.379)	(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 5.

Table 13: Estimation results with linear incumbency tenure (logit model)

	(1)	(2)	(3)	(4)
No. bus route starts within 10 min. drive time	0.067***	0.029	0.111***	0.043
	(0.017)	(0.028)	(0.023)	(0.030)
Min. drivetime to any own garage	−0.069***	−0.075***	−0.063***	−0.074***
	(0.008)	(0.009)	(0.008)	(0.009)
No. comp. garages within 10 min. drive time	−0.280*	−0.250*	−0.312**	−0.259*
	(0.109)	(0.109)	(0.111)	(0.111)
Linear incumbency benefit (T_{jt})	−0.070+	−0.091*		
	(0.036)	(0.039)		
Incumbency benefit [$exp(-T_{jt})$]			1.487***	2.378***
			(0.379)	(0.531)
Garage is not vacant		0.737+		1.853**
		(0.444)		(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.

Table 14: Estimation results with linear incumbency tenure (nested logit model)

	(1)	(2)	(3)	(4)
No. bus route starts within 10 min. drive time	0.023	0.010	0.066*	0.030
	(0.031)	(0.038)	(0.033)	(0.039)
Min. drive time to any own garage	−0.212***	−0.215***	−0.194***	−0.197***
	(0.046)	(0.048)	(0.043)	(0.046)
No. comp. garages within 10 min. drive time	−0.683***	−0.661***	−0.730***	−0.654***
	(0.196)	(0.200)	(0.195)	(0.198)
Linear incumbency benefit (T_{jt})	−0.130**	−0.137**		
	(0.050)	(0.052)		
Incumbency benefit [$exp(-T_{jt})$]			2.077***	2.625***
			(0.526)	(0.667)
Garage is not vacant		0.545		1.795+
		(0.754)		(0.923)
Inclusive value	3.209***	3.041***	3.250***	2.737***
	(0.689)	(0.687)	(0.638)	(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.

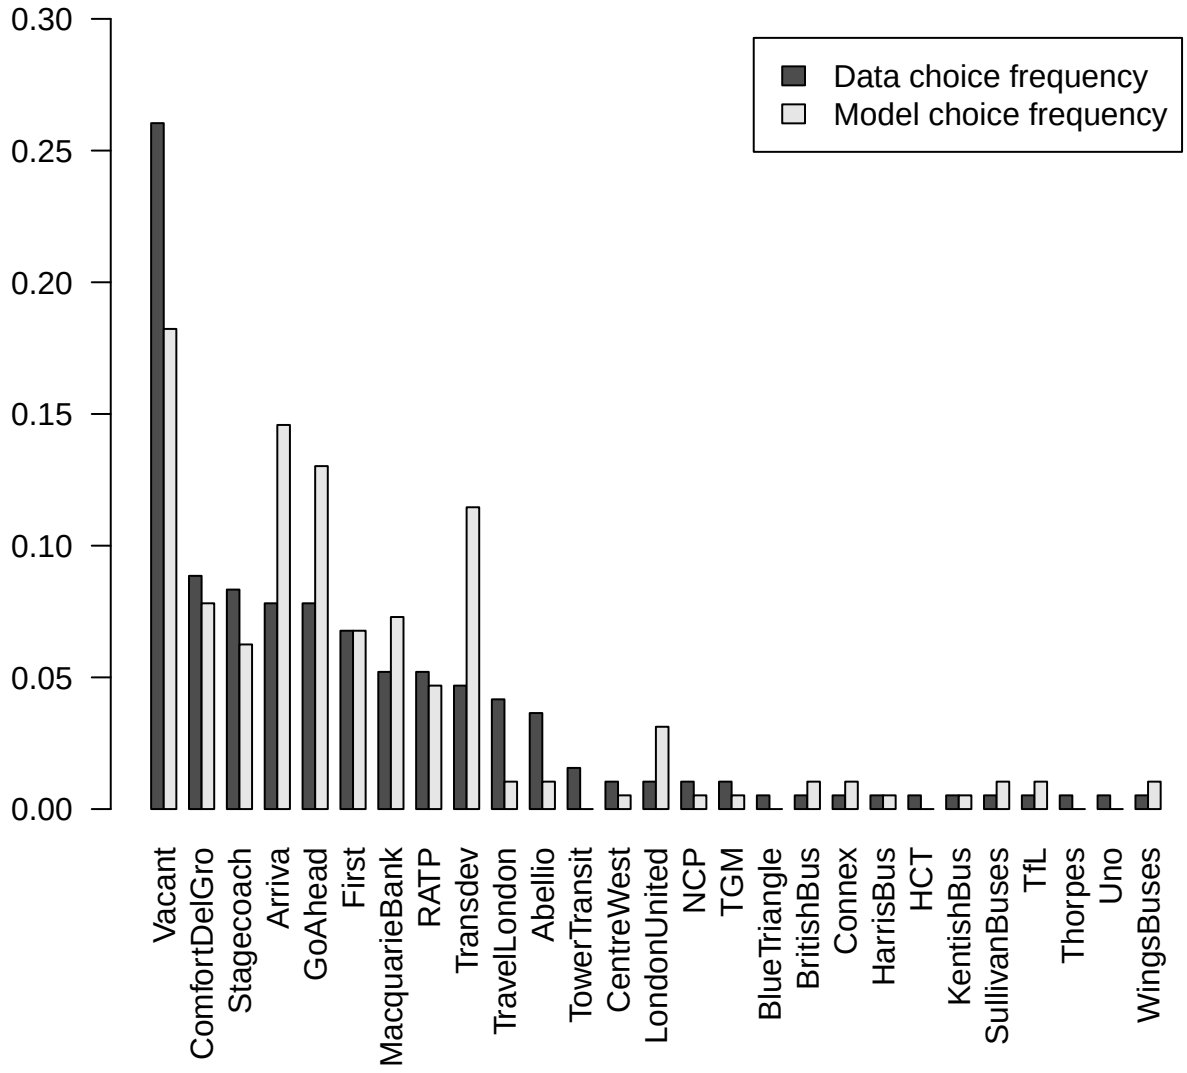


Figure 12: In sample fit: observed vs predicted entry frequency by operator for selected model (column 5 of table 5). Data choice frequency records the number of choice situations in data where an operator (including the Vacant operator) is the entrant, divided by the total number of choice situations (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.

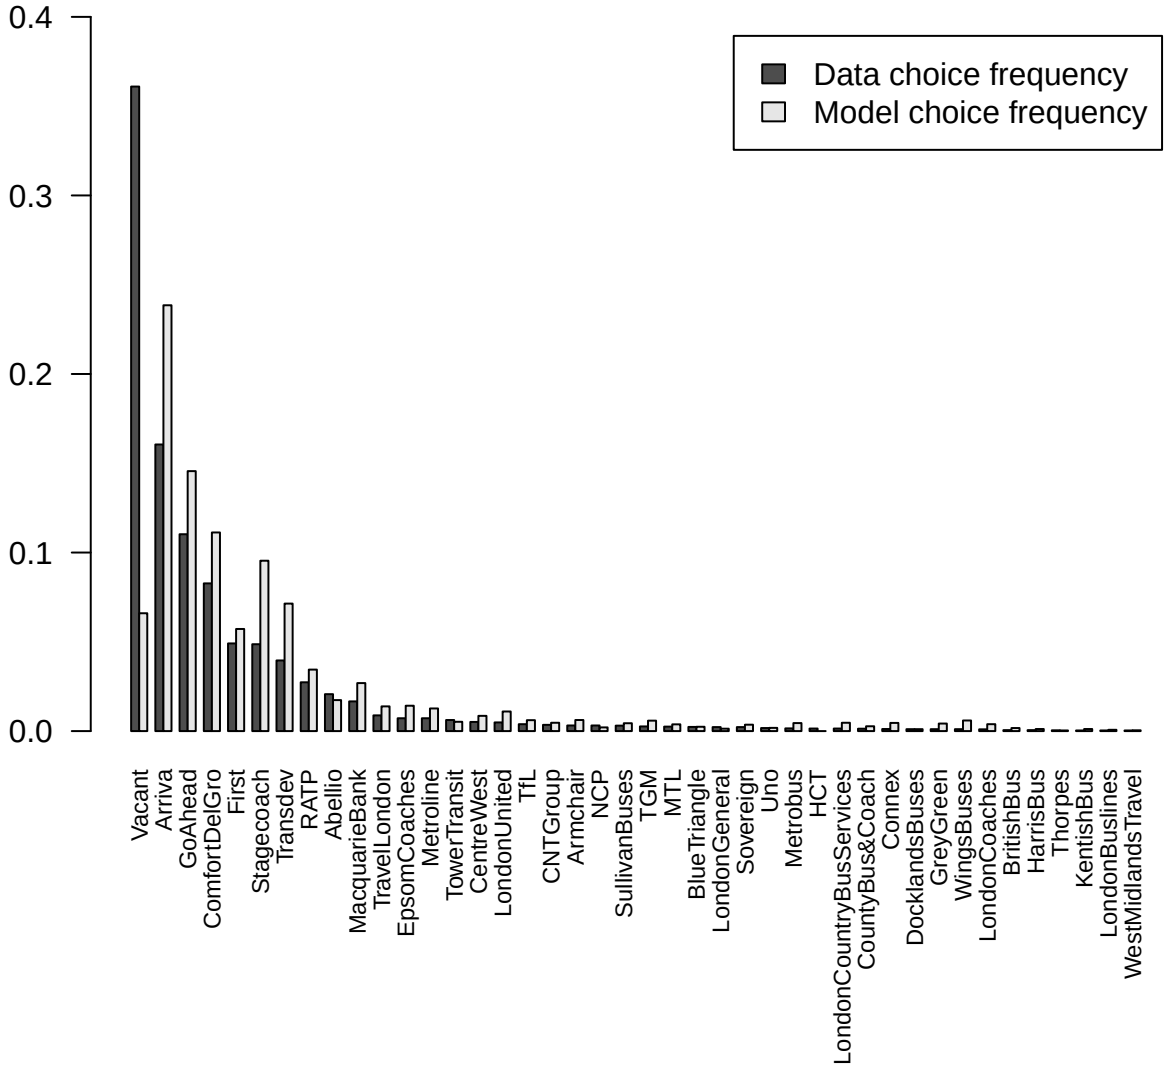


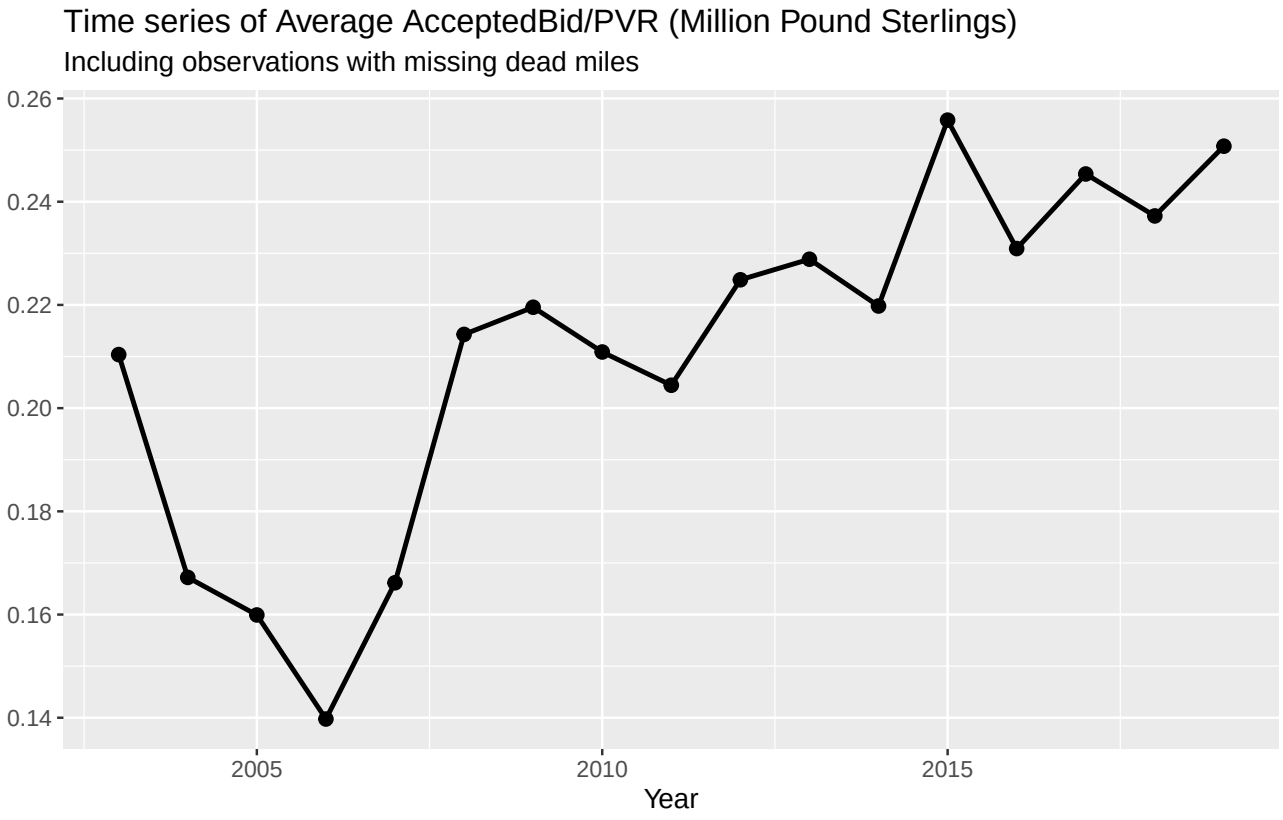
Figure 13: Out-of-sample observed vs predicted entry frequency by operator for selected model (column 5 of table 5). Data choice frequency records the number of choice situations in data where an operator (including the Vacant operator) is the incumbent, divided by the total number of choice situations (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 15: Regressions of punishment on HHI

	$\psi : \psi > 0$	$\psi : \psi > 0$	$1(\psi > 0)$	$1(\psi > 0)$
(Intercept)	1.927***	1.511***	−2.825***	−2.658***
	(0.144)	(0.374)	(0.086)	(0.227)
HHI	1.043	6.591+	4.988***	3.869
	(0.917)	(3.994)	(0.547)	(2.429)
Num.Obs.	8136	8136	71 775	71 775
R2	0.000	0.003		
R2 Adj.	0.000	0.000		
AIC	32 193.1	32 215.8	50 657.4	50 305.6
BIC	32 214.2	32 404.9	50 675.7	50 544.3
Log.Lik.	−16 093.573	−16 080.876	−25 326.686	
F	1.295	1.066		
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.

Figure 14: Winning auction bid, normalized by Peak Vehicle Requirement



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 16, 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 probabilitic 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 cuttoff values*, see Fawcett (2004) for an accessible introduction. For additional insight, Figure 15 displays an ROC curve based on our estimates, but only for a single operator.

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

Table 16: Confusion Matrix in a binary classification setting.

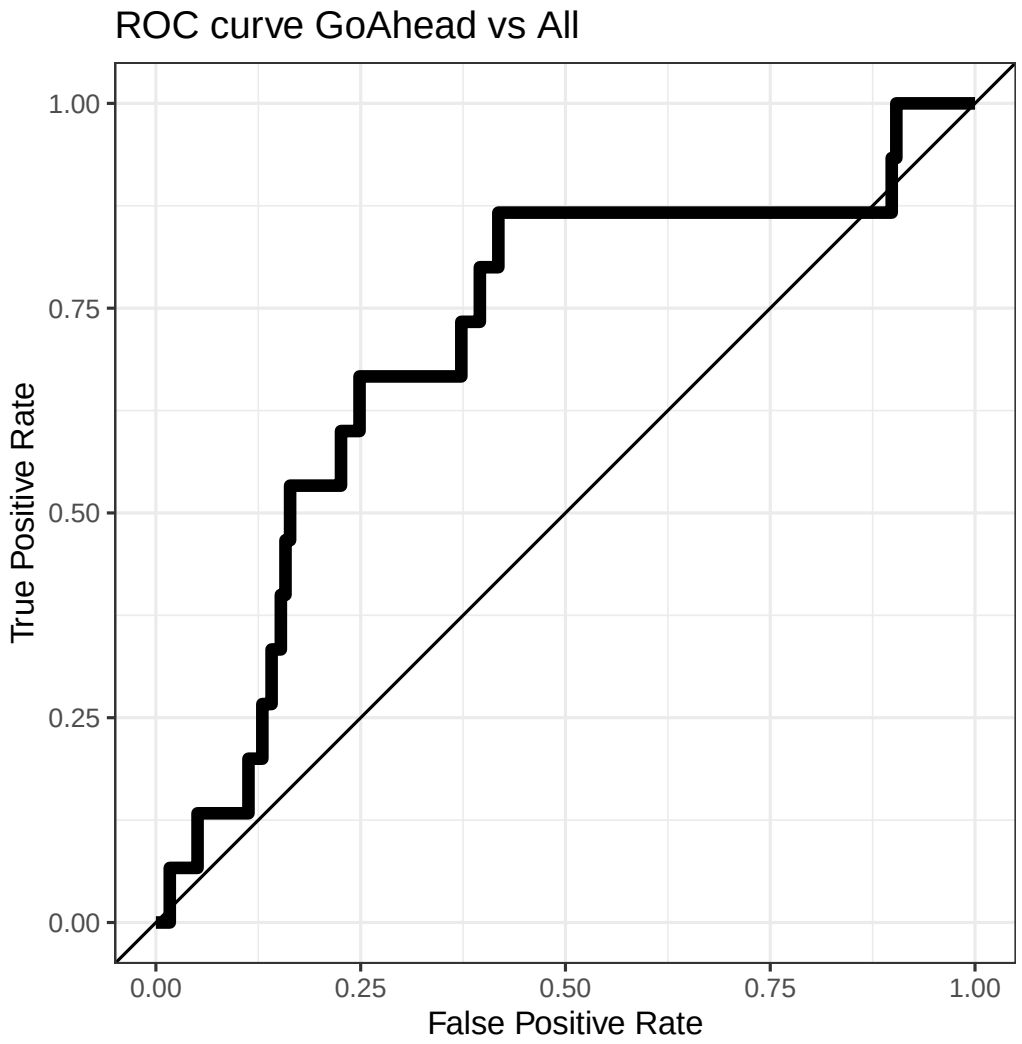


Figure 15: ROC curve choice of GoAhead vs all other alternatives. The plot decreases the cutoff level from 1 to 0 as one traverses the figure from right to left. For example, in the top right corner, and a cutoff level equal to 0, the model achieves a 100% true positive rate, because each predicition 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.