

Chapter 22: Auctioning Bus Routes: The London Experience

Estelle Cantillon and Martin Pesendorfer

August 25, 2004

Forthcoming in: Ausubel, Cramton and Steinberg, Combinatorial Auctions,
MIT Press

1 Introduction

The London bus routes market provides an early example of the use of a combinatorial auction format in public procurement. This market covers about 800 routes serving an area of 1,630 square kilometers and more than 3.5 million passengers per day. It is valued at 600 million Pounds per year (roughly US \$900 million).

Prior to deregulation, bus services in the Greater London area were provided by the publicly-owned London Buses Limited. The London Regional Transport Act of 1984 reorganized the sector. The Act designated London Regional Transport (LRT) as the authority responsible for the provision and procurement of public transport services in the Greater London area, as well as for the development and operation of bus stations and for the operational maintenance. It also advocated a franchise system by empowering LRT to invite private operators to submit bids to carry out bus services.

In order to enhance competition, LRT, which by virtue of the Transport Act acted as the holding company for the original public operator London Buses Limited, created a separate tendering division, independent from its operational division. The operational division, London Buses Limited, was split into 12 operational subsidiaries. These were privatized in 1994.

In practice, the introduction of route tendering was gradual. There was an experimentation phase as LRT built up procurement expertise and identified the routes that represented the greatest potential for cost reduction and the least risk of disruption (Kennedy, Glaister and Travers 1995). Hence, though the first auction took place in 1985, it was not until 1995 that half of the network was tendered at least once.¹ Since then, tendering has reached its steady state regime with 15-20% of the network tendered every year.

This chapter does several things. First, we describe the combinatorial auction format adopted by LRT and briefly discuss its properties (sections 2 and 3). Second, we describe the bidding patterns observed in the data that we collected for these auctions, with a special emphasis on package bidding and the effect of the auction size (section 4)

Finally, as in any practical design problem, LRT was faced with a range of options when deciding on the auction format. A critical input into any such analysis is a better understanding of bidders' preferences, i.e., in our case, their cost structure. In section 5, we

¹Non-tendered routes remained operated by the subsidiaries of London Buses Limited under a negotiated block grant. The private operators and the subsidiaries competed for the tendered services.

summarize a new method that we have developed to analyze bid data from combinatorial first price auctions to do exactly this: infer bidders' cost structure. We illustrate this method and discuss our findings.

2 Auction design

The choice of auction design in London was driven in part by circumstances and economic conditions. Early on, LRT came to the conclusion that the private sector in Britain was unlikely to have the expertise to provide full transportation services in a complex environment such as London (Kennedy, Glaister and Travers 1995). As a result, it was decided to keep the design of the bus network and all related aspects of the transport supply such as frequencies, bus types, exact routing, cross route coordination, and so on, at the network level. Only the actual provision of bus services would be outsourced.

Nevertheless, this left several important design questions open. First, what “items” would be auctioned? A contract for a route? For a set of routes? Or less than a route? Second, how should LRT auction these contracts? Should LRT seek to tender the whole network at the same time? Route by route? Or groups of routes together? If so, which criteria should they use to decide which routes to auction at the same time? Also, how much should bidders be able to express in terms of their preferences over routes and bundles of routes? Third, who should be allowed to participate? Finally, which criteria should they use to award the contracts?

In the remainder of this section, we describe LRT's actual choices with respect to these options.

2.1 Pre-auction design

Definition of the items auctioned. With two exceptions, the items auctioned are contracts for the operation of a bus route. First, “mobility routes”, which are low frequency services (one or two round trips per week) with buses especially equipped to accommodate wheelchairs, are usually auctioned as a bundle. The reason is that any such route by itself represents too small a contract. Second, the night and day portions of the same route are sometimes auctioned as separate contracts when they require different types of buses. The intention in both cases is to define an item as a self-standing homogeneous contract.

The contracts are usually for 5 years. LRT has experimented with several contractual forms over time. It has mainly used “gross cost” contracts, whereby the winner’s compensation is the amount it bid, while all revenues collected on the bus accrue to LRT.

Packaging and sequencing decisions. From the very beginning, designing the auctions to enhance competition and attract new entry to the London bus market was a clear concern. The choice of auctioned routes in the early years, peripheral and requiring a relatively small number of vehicles to operate, reflects this (Glaister and Beesley 1991). This concern may also have affected how LRT decided how many routes to auction at the same time. On the one hand, there are bundling benefits in terms of shared fixed costs, coordination efficiencies. This calls for auctioning a large amount of related routes together. On the other hand, larger auctions could discourage entry by smaller bus operators without the capacity to bid on all the routes in the auction.² The current practice may have reached a compromise between these two views. Today, LRT holds an auction every two or three weeks. An auction covers on average 3.77 routes, though the range in our data goes from one route to 21 routes in a single auction. Importantly, the routes tendered within one auction are usually in the same area of London.

Participation. Only pre-qualified operators can participate in the auctions. Pre-qualification screens for financial stability and operational capacity of potential operators. There were about 51 such operators during our sample period. All pre-qualified operators are informed of upcoming auctions.

2.2 Auction format

The auction format adopted in the London bus routes market is a variant of a combinatorial first price auction. In a standard combinatorial first price auction, bidders simultaneously submit sealed bids on any individual items and on any packages of items; the auctioneer solves the Winner Determination Problem (Lehmann, Muller, and Sandholm, Chapter 12) by determining the “best” bidder-item allocation based on these bids (since this is a procurement setting, the best allocation is the one that minimizes total cost); and the winners receive the amount they bid for the items they won.

As in the standard combinatorial first price auction, bidders in the LRT auction can

²As a benchmark, a typical garage, which is necessary to operate a route, has capacity for 50-100 buses and serves about 8 routes.

submit bids on any number of routes and route packages. There is no restriction on the number of bids placed, nor is there an obligation to bid on some routes or route packages. In particular, a bidder can submit a bid on a package without submitting a bid on the individual routes that make up that package.

The distinctive feature of the LRT auction is that each bid is a firm but non-exclusive commitment of resources. This means that two bids on different routes implicitly define a bid for the package of these routes. An important consequence of this rule is that bidders are not allowed to bid more for a package than the sum of the bids on any partition of that package. In particular, this rules out bids expressing diseconomies of scale or scope.³ The original motivation for this rule was the expectation that the market was mainly characterized by economies of scale and scope, and that by allowing bidders to express such synergies, LRT would lower its procurement costs and improve efficiency.

After verification that the bids satisfy the technical requirements of the auctioned contracts, LRT awards the contracts to the bidder-allocation that delivers the best economic value. In practice, this means that the contract is awarded to the low bidder but deviations at the margin are possible to account for operator quality, for example. The winner receives the amount he bid for the contract, indexed yearly. To allow winning operators to reorganize and order new buses if necessary, contracts start 8 to 10 months after the award date.

3 Motivations for submitting a package bid

We have built a model to study the properties of the auction format adopted in London. Our goal was twofold. First, to investigate the motivations for submitting bids on packages in the first price combinatorial auction, and their implications for welfare and efficiency.⁴ Second, to guide our analysis of the data by suggesting things to look for in the data.

An outline of the model is as follows. An agency seeks to procure m items from n risk neutral bidders. Each bidder i privately observes a cost draw, $c_s^i \in \mathbb{R}$, for each possible subset of the items, $s \subseteq S = \{1, \dots, m\}$ (that is, s represents either a single item or a package of items). Contract costs are ex-ante distributed according to some joint distribution

³In practice, the bidders in our data start expressing unease with this restriction for the auctions of six routes and more.

⁴With respect to other choices that LRT made, we have little to say from a theoretical perspective. However, the next section looks at bidder participation and prices as a function of the auction size.

$F((c_s^i)_{s \subseteq S, i=1, \dots, n} | X)$ where X denotes a vector of observable auction characteristics. Here F is common knowledge, it has bounded support, and has a well-defined strictly positive density everywhere, absolutely continuous with respect to the product of its marginals. This allows independence and correlation or affiliation in bidders' costs across bidders and contracts.

The rule for the auction replicates that of the LRT auction. Bidders may submit bids on all subsets of the set of items. Let b_s^i denote bidder i 's bid on the subset of items $s \subseteq S$, and let $b^i = (b_1^i, \dots, b_s^i, \dots, b_S^i) \in \mathbb{R}^{2^m - 1}$. Bidders receive the value of their winning bids and the auctioneer selects the winner(s) based on the allocation that minimizes his total payment. Formally, the last restriction requires that $b_{s \cup t}^i \leq b_s^i + b_t^i$ for all s and t such that $s \cap t = \emptyset$. This captures the fact that package bids must be equal to or lower than the sum of the bids on any partition of the package.

A Bayesian Nash equilibrium in this environment is a n -tuple of vector-valued bidding functions, $b^i(c_1^i, \dots, c_S^i) \in \mathbb{R}^{2^m - 1}$, that maximize bidder i 's expected payoff, for all i ,

$$\sum_{s \subseteq S} (b_s^i - c_s^i) G_s(b^i | X) \quad (1)$$

where $G_s(b^i | X)$ denotes bidder i 's probability of winning package s given his opponents' strategy and his submitted vector of bids b^i .

We have identified two distinct motivations for a bidder to submit a package bid at equilibrium. First, bidders may want to submit a bid on a package when the cost for the package differs from the sum of the costs. This is the standard synergy explanation and it motivates much of the use for combinatorial auctions.

But we also uncovered a more strategic motivation for submitting a package bid lower than the sum of the constituent bids. The reason is that the bids submitted by a bidder compete with one another in a combinatorial auction. For example, consider a two-item auction. Fix bidder i , and for each item or package of items s , define B_s^{-i} as the lowest bid submitted by bidder i 's opponents on s (from the perspective of bidder i , B_s^{-i} is a random variable; he does not know what his opponents are bidding). With two items, there are four possible winning allocations between bidder i and his opponents: bidder i wins item 1 (and one of his opponents wins item 2, if $b_1^i + B_2^{-i}$ is the cheapest allocation), he wins item 2 (and his opponents win item 1), he wins both items or he does not win anything.

Now consider item 1. Holding the distribution of the opponents' best bids $(B_1^{-i}, B_2^{-i}, B_{12}^{-i})$ fixed, decreasing b_1^i increases bidder i 's chance to win exactly that item by lowering the price

of allocation $b_1^i + B_2^{-i}$ relative to the others. But it decreases bidder i 's chance of winning the package because it could be the case that, *had bidder i not lowered b_1^i* , the cheapest bid-allocation was b_{12}^i . Another way to look at this is in terms of the following trade-off. The benefit to bidder i from lowering his bid on item 1 is that he wins item 1 more often. But the costs are twofold. First, it lowers his profit margin whenever his wins item 1. Second, it reduces his chance of winning the package of items 1 and 2.⁵ At equilibrium of course, bidder i chooses bid b_1^i such that this marginal benefit and these marginal costs exactly balance one another.⁶ A similar analysis can be carried out for bidder i 's bid on item 2 and on the package.

The result of this strategic effect is that bidder i may be tempted to submit a package bid $b_{12}^i < b_1^i + b_2^i$, even when his costs are completely additive. A simple example may illustrate the idea. Suppose that bidder i is facing two other bidders. Bidder 1 is only interested in item 1; bidder 2 is only interested in item 2. All bidder i knows is that they might submit a bid of 7 or 15 on the item they are interested in depending on whether they have a low or high cost. Moreover, whenever bidder 1 has a high cost and therefore bids 15 on item 1, bidder 2 has a low cost and bids 7 on item 2, and vice versa. Suppose finally that bidder i has a cost of 5 for each item. In this example, bidder i 's best strategy is to submit a bid just below 23 for the package, and a bid higher or equal to 15 for each of the items. This way, he wins for sure both items at a profit of $(23 - 10) = 13$.⁷ If instead, he bid only on the individual items, his best strategy would be to bid slightly less than 15 on each item for an expected profit of $(15 - 5) = 10$. Clearly, it is best in this example to submit a bid for the package. Moreover, it is not in bidder i 's interest to submit a bid on the individual

⁵Formally, the argument assumes that $b_{12}^i < b_1^i + b_2^i$ so that lowering b_1^i slightly still satisfies LRT's package bid constraint. This is also the reason why lowering b_1^i only affects the winning chances of b_{12}^i and not that of b_2^i . To see this, suppose that, originally, bidder i wins item 2, i.e. $B_1^{-i} + b_2^i \in \arg \min\{b_1^i + B_2^{-i}, b_2^i + B_1^{-i}, B_{12}^{-i}, b_{12}^i\}$. We claim that it cannot be the case that, after slightly lowering b_1^i to $b_1^i - \varepsilon$, the winning allocation now becomes to give item 1 to bidder i . Indeed, $B_1^{-i} + b_2^i < b_{12}^i$ by assumption, which itself is lower than $b_1^i + b_2^i$ by the package bid constraint. This means that $B_1^{-i} < b_1^i$ both before and after the slight decrease. But then this means that $b_1^{-i} + B_2^{-i} > B_1^{-i} + B_2^{-i} > B_{12}^{-i}$, which, in turn, is greater than $B_1^{-i} + b_2^i$ by assumption. Hence, lowering b_1^i has no effect on the winning probability of b_2^i .

⁶Interestingly, we can compare this with the trade-off a bidder faces in a single item auction. The marginal benefit is the same, but the cost is only in terms of profit margin.

⁷To simplify the argument, we are being slightly cavalier with the math. This is inessential for the argument. More formal analyses of this kind of problems are provided in Adams and Yellen (1976), McAfee, McMillan and Whinston (1989) and Armstrong (1996) for example.

items that has a chance of winning, exactly because of the strategic effect identified above. To see this, consider what happens if bidder i lower his bid below 15 on item 1 when he also submits a package bid of 23. Now, with some probability p , bidder 1 bids 15 on item 1. In that case, the cheapest allocation is to give item 1 to bidder i for a price below 15, item 2 to bidder 2 for 7, for a total amount less than the package bid. Bidder i 's profit now becomes less than $p(15 - 5) + (1 - p)(23 - 10)$ which is lower than his profit of 13 when he only submits a bid on the package. The reason is that the increased probability of winning item 1 comes at the cost of the decrease in the probability of winning the package.

This analysis generalizes. McAfee, McMillan and Whinston (1989) proved that whenever bids by bidder i 's opponents are independently distributed, it is also optimal for bidder i to submit a bid on the package that is strictly lower than the sum of the bids on the individual items. Armstrong and Rochet (1999) proved that submitting such a package bid is profitable whenever the correlation among opponents' bids is not too high.⁸

The normative implication of these two motivations - synergy and strategic - for auction design is ambiguous. In fact, there is no general theoretical result on whether a combinatorial first price auction is better than a series of independent first price auctions. It is easy to construct examples, especially when synergies are important, where a combinatorial auction is better. At the same time, Cantillon and Pesendorfer (2004) present a realistic example without cost synergy where a first price combinatorial auction leads to higher costs and lower efficiency than a simpler first price auction without package bidding.

There are several lessons from this analysis. First, we need to be careful when interpreting the data because a package bid is no guarantee that they are synergies between the items.

Second, the theory suggests a couple of observable factors likely to favor the use of package bids. Correlation or, more precisely, lack of correlation in the environment that bidders face is a driver of package bidding for strategic reasons. Underlying synergies will favor package bidding for synergy reasons.

Third, there is no clear a priori answer to the question of whether allowing package bids in the London bus routes market was a good idea or not. Any answer to this question must rely on a more careful analysis of the economic environment, in particular the drivers of costs. This is the purpose of the next two sections.

⁸These papers do not deal with the combinatorial first price auction proper but with the monopoly multiproduct pricing problem. As argued in our companion paper, the mathematical structure of the two problems is identical. Therefore the results also apply to the combinatorial first price auction.

4 Practice

This section describes our data and provides summary statistics.

We have collected data on a total of 179 auctions held between December 1995 and May 2001. For each auction and for each route in the auction the data provide the following information: (1) route characteristics including contract duration, the planned start of the contract, the start and end points of the route, the route type (day, night, school service, mobility route), the annual mileage, the bus type (single deck, midibuses, double deck or routemaster), and the peak vehicle requirement.⁹ For the routes auctioned after May 2000, we observe an internal cost estimate generated by LRT. (2) The bidders, their bids (including package bids), and the garage from which bidders plan to operate the route(s). Bids are expressed in June 1995 Pounds.

The auction format implies that bidders are committed by their route bids. Hence, route bids define implicitly a package bid with value equal to the sum of the route bids. We call a package bid “non-trivial” when it is strictly less than the sum of the component route bids. Otherwise, we call the package bid “trivial”.

Most auctions consist of a few routes only. The average number of routes per auction in our data is 3.77.

Descriptive Summary statistics are reported in Table 22.1. The table reveals that the distribution of the number of routes per auction is as follows: 50 auctions consist of a single route, 36 auctions have two routes, 32 auctions have 3 routes, 13 auctions have 4 routes, 10 have 5 routes, 12 have 6 routes, 12 have between 7 and 9 routes, and 14 auctions have more than 9 routes. The total number of routes across all auctions equals 674.

⁹The peak vehicle requirement determines how many buses the winning operator will need to commit to the contract.

Table 22.1 Descriptive Summary of Auction Data by Number of Routes^a

No of Routes:	All	1	2	3	4 - 6	7-9	10-21
No of Auctions:	179	50	36	32	35	12	14
No of Bidders per Auction	4.28 (2.04)	3.22 (1.52)	3.81 (1.82)	4.34 (1.79)	4.83 (2.01)	5.92 (2.54)	6.00 (1.84)
No of Bidders per Route	2.71 (1.40)	3.22 (1.52)	2.88 (1.68)	2.85 (1.52)	2.39 (1.20)	2.48 (1.12)	2.78 (1.33)
Log Individual Route Bid	13.31 (1.24)	12.81 (1.33)	13.22 (1.24)	13.54 (1.20)	13.34 (1.29)	13.24 (1.17)	13.27 (1.18)
No of Bids per Bidder ^b	-	1.00	1.76	2.42	3.18	5.06	-
	-	(0.00)	(0.51)	(0.70)	(1.28)	(2.97)	-
No of Package Bids per Bidder ^b	-	0	0.24	0.41	0.84	1.45	-
	-	(0.00)	(0.30)	(0.36)	(0.78)	(2.13)	-

a: Standard deviations are displayed in parenthesis. b: The No of Bids per Bidder and the No of Package Bids per Bidder exclude trivial package bids.

On average 4.28 bidders submit at least one bid on an auction. The number of bidders ranges between 1 and 13 per auction. On average 2.71 bidders submit a bid for an individual route on the auction and the number of bidders per route ranges between 1 and 7. In total, 1,818 individual route bids were submitted. The average individual route bid equals 13.31 in logarithm, which amounts to about 603,000 Pounds.

Bidder participation and auction size. The number of bidders per auction increases with the size of the auction. The average number of bidders equals 3.22 for single route auctions, 3.81 for 2 route auctions, 4.34 for 3 route auctions and increases to 6 bidders for auctions of more than 9 routes. This does not mean increased competition at the route level. Indeed, there is no correlation between the number of bidders per route and the number of routes in the auction. The number of bidders per route remains roughly constant across the range of the number of routes. Similarly, the log of individual route bids does not change significantly as the number of routes varies.

Number of bids submitted and auction size. The number of bids submitted per bidder equals 1.76 for two route auctions, it increases to 2.42 for three route auctions, 3.18 for four to six route auctions, 5.06 for seven to nine route auctions. This number underestimates the actual number of bids submitted because it does not count trivial package bids.¹⁰

¹⁰We have not coded the data on package bids for auctions with more than nine routes. Hence the entry

We also looked more specifically at non-trivial package bids. The number of package bids per bidder increases with the size of the auction. It equals 0.24 for two route auctions, it increases to 0.41 for three route auctions, 0.84 for four to six route auctions, 1.45 for seven to nine route auctions.

Next, we report the relative frequencies of route bids, trivial package bids, and non-trivial package bids for two, three and four route auctions.

Table 22.2 Frequency of Package Bids by Number of Routes

No of Routes:	2	3	4
Percent of Bidders Submitting a Package Bid	51	63	52
Percent of Non-trivial Package Bids of all Package Bids	36	59	55
Percent of Bidders Submitting a Full Set of Bids	18	3	0

The first column in Table 22.2 illustrates that for two route auctions 51% of all active bidders submitted a package bid. Non-trivial package bids account for 36% of all package bids submitted. All but one bidder who submitted a package bid also submitted a route bid. A full set of bids (individual route bids and non-trivial package bids) was submitted by 18% of active bidders.

The second column in Table 22.2 shows that for three route auctions 63% of active bidders submitted a package bid and non-trivial package bids accounted for 59% of all package bids submitted. All bidders that submitted a package bid also submitted a route bid. A full set of bids was submitted by 3% of active bidders.

For four route auctions 52% of active bidders submitted a package bid and non-trivial package bids accounted for 55% of all package bids submitted. All bidders that submitted a package bid also submitted a route bid. No bidder submits a full set of bids.

Non-trivial package bids and route attributes. The operators invoke the possibility to share spare vehicles and depot overhead costs, and more efficient organization and coordination of working schedules when they offer discounts for packages. Therefore, the frequency of package bids may be related to the attributes of the routes offered at an auction. We examined this hypothesis by looking at the sample correlation in the data between the frequency of non-trivial package bids and the similarity of the routes for sale. Specifically, we looked at correlation between the frequency of non-trivial package bids and whether the

for the column 10-21 and the overall value are left blank.

routes at the auction require the same bus type or not. We found that the correlation coefficient between submitting a package bid and whether the routes require the same bus type or not is not significantly different from zero. Furthermore, there is no significant correlation between submitting a non-trivial package bid and whether the routes require the same bus type or not when we condition on the number of routes offered at the auction.

Package bid discount. We can calculate the markdown of a non-trivial package bid relative to the sum of route bids. We considered two measures: The Total Package Bid Discount calculates the discount equal to one minus the ratio of package bid to sum of route bids. For two, three and four routes, the formula is given by:

$$1 - \frac{b_{12}}{b_1 + b_2}, 1 - \frac{b_{123}}{b_1 + b_2 + b_3}, 1 - \frac{b_{1234}}{b_1 + b_2 + b_3 + b_4}.$$

The discount equals zero if there is no package discount. It equals one half if the package bid costs half as much as the sum of route bids. If routes differ in size, then the Total Package Bid Discount will not adequately measure bid discounts. The reason is that Total Package Bid Discount measures the discount on the average route in the bid only. Hence, even if a package bid includes the smaller route for free, the Total Package Bid Discount will be almost zero, when the size difference between two routes is large.

Our second measure accounts for size heterogeneity among routes by looking at the marginal route bid. The Marginal Package Bid Discount determines the marginal discount on the smallest route bid in the package. To calculate the Marginal Package Bid Discount we sorted bids in descending order. Then, we computed one minus the ratio of the marginal route price, calculated as the difference between the package bid and the route bids on the marginal route relative to the marginal route price of the marginal route. For two, three and four route auctions the Marginal Package Bid Discount measure is given by:

$$1 - \frac{b_{12} - b_1}{b_2}, 1 - \frac{b_{123} - b_1 - b_2}{b_3}, 1 - \frac{b_{1234} - b_1 - b_2 - b_4}{b_4}.$$

where b_1 denotes the largest route bid, b_2 the second largest route bid, and so on. The Marginal Package Bid Discount equals zero if there is no package bid discount. It equals one if the marginal route is offered for free in the package bid.

Table 22.3 Package Bid Discount by Number of Items in the Package*

No of Items in the Package:	All	2	3	4
No of Observations:	421	232	131	58
Total Package Bid Discount	0.057	0.049	0.061	0.077
Marginal Package Bid Discount	0.284	0.182	0.364	0.513

* We consider non-trivial package bids only.

Table 22.3 depicts the package discount measures for two, three and four route package bids. Ignoring trivial package bids, the Total Package Bid Discount amounts to 5.7%, on average, for 2, 3 and 4 route package bids. The Marginal Package Bid Discount is substantially higher and amounts to 28.4%.

Both the Total Package Bid Discount and the Marginal Package Bid Discount increase with the number of routes in the bid. The Total Package Bid Discount equals 4.9% for two route bids, 6.1% for three route bids and 7.7% for four route bids. The Marginal Package Bid Discount equals 18.2% for two route bids, 26.4% for three route bids and 51.3% for four route bids.

5 A method to infer costs from auction data

The evidence presented so far suggests that bidders in the London bus routes market do have recourse to package bids (Table 22.2) and offer economically significant discounts (Table 22.3), even if they do not bid on all packages. Moreover, one factor suggested by the theory - the similarity between routes - did not seem to explain bidders' use of package bids much (there was no correlation between the frequency of non-trivial package bids and a measure of similarity between the routes). In this section, we go one step further and ask: Can we say anything about bidders' cost structure based on the bid data?

We have developed a method to infer the cost structure of bidders in combinatorial first price auctions that allows us to evaluate the presence of cost synergies. Our method is general and can be used to analyze any data generated from a combinatorial first price auction. It can be described heuristically as follows.

Bidders' cost structure (the cost of operating each route and each package of the routes tendered in the auction), their information (for example, about the participating bidders), together with the auction rules, describe a game of incomplete information. The information is incomplete because bidders do not know the costs of the other participating bidders in

the auction. The equilibrium of such a game, i.e. the bids submitted by each bidder given their costs, generates an (equilibrium) distribution of bids. If we make the assumption that the bidders in the London bus routes market are using equilibrium strategies, we can use the bids observed in our data to estimate this equilibrium distribution of bids.

Once we know the empirical distribution of bids, we can infer the probability of winning of any bid vector submitted by our bidders. Hence, we can replicate the optimization problem that each bidder in our data was facing when he submitted his bids. Since we assumed that our bidders followed equilibrium bid strategies, we can use the first order conditions of their optimization problem to infer their costs. A simple example may illustrate the idea. Suppose you are bidding in a single item auction and that we can estimate that your probability of winning depends on your bid according to the function $G(b)$, where b is your bid. Suppose we see you submit bid b^* . *If we assume that you are playing according to the equilibrium*, we can conclude that b^* maximizes your expected profit, that is, b^* solves $(b^* - c)G(b^*)' + G(b^*) = 0$. Now, of course, we do not know your cost, but if we reorganize this expression, we find that $c = b^* + \frac{G(b^*)}{G(b^*)'}$. Since we know everything of the right hand side of this equation, we have inferred your cost just by observing how you bid!

We now describe and illustrate each step of our approach on the subset of our data of two route auctions.¹¹

5.1 Step 1: Choice of the game that describes the data

Our first step is to choose an economic model describing bidders' behavior. We have chosen the constrained combinatorial first price auction with independent private values and risk neutral bidders as the model that we feel best represents the London bus routes auction.

The choice of the auction format deserves little comment given the description in section 2. We account for the fact that bidders are committed by their bids by placing the explicit constraint on strategies that bids on a package must be smaller or equal to the sum of the bids on any partition of the package. An additional constraint that we have included in our modeling is a reserve price on each route and package s , R_s .

We consider each auction in our data as an independent observation and assume that bidders treat each of those as independent one-shot games. Given the frequency of the auctions, this may not hold exactly. However, a mitigating factor is the fact that auctions

¹¹Additional technical details are provided in Cantillon and Pesendorfer (2004).

are held for routes in different parts of London so that bidders' overlap may be limited across any sequence of auctions. For example, we computed that an average bidder in our data participates in an auction every 5 months.

Our assumption that costs are private values is driven by our interpretation that the main source of uncertainty in this market is bidders' opportunity costs. Indeed, while it is certainly true that a large fraction of the routes' costs are common to all bidders (the labor, fuel and other material costs represent in the order of 52% of bidders' costs), asymmetries of information about these costs are unlikely because there are well-functioning markets for these inputs. By contrast, opportunity costs arising from the use of capital and rolling stock are likely to differ across bidders. Moreover, they are essentially private values.

Our assumption that costs are ex-ante independently distributed across bidders, conditional on auction characteristics, is a simplifying assumption. It is mainly driven by the small size of our dataset.

Summarizing, the primitives of the model are: (1) the 3 dimensional ex-ante distribution of costs for each bidder i given the characteristics of the auction X^{it} , $f(c_1, c_2, c_{12}|X^{it})$,¹² and (2) bidders' profit conditional on winning route or package s , that is, the difference between their winning bid and their known cost of performing that contract, $b_s - c_s$. Neither the costs nor the distribution of costs are observed in the data. Only the bids and the auction characteristics are.

5.2 Step 2: Estimation of the equilibrium distribution of bids

Our second step is to estimate the equilibrium distribution of bids. Since bidders' costs are independently distributed so must their equilibrium bids. Moreover, we have assumed that, conditional on auction characteristics, each auction was independent. This means that an observation in our analysis is a bidder-auction. Based on an analysis of bidder participation, we have considered that any bidder with a garage within 5 miles of one of the extremities of a route in an auction is a potential bidder for that auction. With this definition, we have 338 observations for the two route auctions. Not all of these bidders submitted a full vector of bids. In fact, in our data, only 131 of those 338 bidders actually participated by submitting at least one bid

¹²Auction characteristics for auction t are denoted by the vector X^t . We use the notation X^{it} to stress the fact that they are viewed from bidder i 's perspective. Hence, $f(c_1, \dots, c_S|X^{it})$ refers to $f_i(c_1, \dots, c_S|X^t)$.

Our objective is to estimate the distribution of bids based on this data. Because bids on route 1, route 2 and the package are related, we want to estimate them jointly. Denote this three-dimensional distribution by $h(b_1, \dots, b_S | X^{it})$, where the conditioning on X^{it} accounts for some auction characteristics. The marginal densities of submitted bids seemed lognormal, so we adopt the following parametric specification for our two route auctions. With probability p , latent bids are distributed lognormal:

$$\begin{bmatrix} \ln b_1 \\ \ln b_2 \\ \ln b_{12} \end{bmatrix} \sim N \left(\begin{bmatrix} \mu_1(X^{it}) \\ \mu_2(X^{it}) \\ \mu_{12}(X^{it}) \end{bmatrix}, \begin{bmatrix} \sigma_{11}(X^{it}) & \sigma_{12}(X^{it}) & \sigma_{13}(X^{it}) \\ \sigma_{12}(X^{it}) & \sigma_{11}(X^{it}) & \sigma_{13}(X^{it}) \\ \sigma_{13}(X^{it}) & \sigma_{13}(X^{it}) & \sigma_{33}(X^{it}) \end{bmatrix} \right) \quad (2)$$

and with probability $1 - p$, they are distributed with values higher than the reserve price and so they are not observed.¹³

There are several elements worth emphasizing in this specification. First, the p parameter allows the specification to follow closely the empirical distribution of observed bids, yet, at the same time, account for the fact that in the order of 40% of the bidders did not submit any bid. Second, the specification in (2) imposes some symmetry between route 1 and route 2, conditional on route characteristics. The reason is that nothing a priori distinguishes route 1 from route 2 in our sample. Both of them are single routes and therefore, conditional on route characteristics, the distribution of bids for those routes should be the same. This symmetry restriction appears in the specification of the means, $\mu_1(X^{it}) = \mu_2(X^{it})$, and in the covariance terms, $\sigma_{11}(X^{it}) = \sigma_{22}(X^{it})$, and $\sigma_{13}(X^{it}) = \sigma_{23}(X^{it})$. Third, many bidders did not submit a bid on all routes and route packages in the auctions. Our interpretation is that they did not find it worthwhile to submit a bid that would have had a positive probability of winning. Similarly, bids on packages higher than the sum of the bids for the individual items are also not observed in practice. So, referring back to (2), while we assumed that the log of the latent bid on the package was lognormal, we only observe those bids that satisfy the package bid constraint strictly.

As an illustration, Table 22.4 reports our estimates of the parameters of the three dimensional distribution of bids for the simplest specification where the covariance matrix is not a function of the covariates X^{it} , $\mu_s(X^{it}) = \mu_1 + \beta \ln(ice_{st})$ if s is a single route, and

¹³In the absence of a characterization result of the equilibrium in the combinatorial first price auction, we have no guarantee that the equilibrium takes exactly this shape. The objective here is to stay as close to the data as possible.

$\mu_s(X^{it}) = \mu_{12} + \beta \ln(ice_{st})$ if s is the package route (ice_{st} is the internal cost estimate for route s in auction t).¹⁴ The parameters were estimated using the method of moments, which yields consistent and asymptotically normal estimates.¹⁵

Table 22.4: Estimates for the three-dimensional lognormal distribution (dependent variable is $[\ln b_1, \ln b_2, \ln b_{12}]'$) :

	point estimate	std deviation*
μ_1 (constant for mean of single bids)	-0.5041	0.4412
μ_{12} (constant for mean of package bid)	-0.3127	0.4610
$\ln(\text{internal cost estimate})$	1.0299	0.0329
σ_{11}	0.1170	0.0194
σ_{12}	0.0104	0.0041
σ_{13}	0.0713	0.0131
σ_{33}	0.0805	0.0152
p	0.3785	0.0260

* Standard deviations refer to asymptotic standard deviations.

In section 3, we noted that a key driver of package bidding is the correlation in the environment that bidders are facing. Table 22.4 allows us to say more about this. The coefficient of correlation between the log of the bids on route 1 and those on route 2, $\frac{\sigma_{12}}{\sigma_{11}}$, is equal to 0.09. This measure of correlation looks at one bidder only. What matters in the objective function of a typical bidder is not the correlation among individual bids of a bidder, but the correlation in the low bids from the set of opponents on each route. In our data, the average number of potential bidders is 8.44. We computed the coefficient of correlation between the opponents' lowest bid on route 1 and the opponents' lowest bid on route 2, conditional on being lower than the reserve price, for a representative auction with 8 potential bidders and an internal cost estimate for the individual routes of 1,002,553 Pounds. The coefficient of correlation is 0.19. There is positive correlation in the environment that

¹⁴Since the log of the internal cost estimate explains 90% of the variation in the log of bids, the internal cost estimate is the most reasonable covariate to include in a sparse specification.

¹⁵Specifically, we used the first and second moments of the distribution, conditional on the bids not being censored, as well as the probability of not submitting any bid. The first moment for the individual route bids was interacted a dummy and with the $\ln(ice_{st})$ as instruments, leading to a total of 8 moment conditions i.e. an exactly identified model. For estimation purposes, the reserve price was set equal to twice the internal cost estimate.

bidders are facing. We revisit this issue in step 4.

5.3 Step 3: Computation of the probabilities of winning

Now that we have estimated the distribution of bids, we can estimate the probabilities of winning. Consider an auction with n potential bidders and characteristics X^t . Fix bidder i . His probability of winning route or package s is the probability that

$$b_s^i + B_{S \setminus s}^{-i} < \min_{\forall t \subseteq \{1,2\}} \{B_S^{-i}, b_t^i + B_{S \setminus t}^{-i}\}$$

where B_t^{-i} denotes the best bid by bidder i 's opponents on package t .¹⁶

The distribution of low bids of bidder i 's opponents and therefore the probability that bidder i wins package s given his bid vector, $G_s(b^i | X^{it})$, can be recovered from the individual distributions estimated in step 2 by integration.¹⁷

5.4 Step 4: Infer the costs

Bidder i 's optimization problem in each auction is to decide what bids to submit on each route and package in order to maximize his expected payoff. Formally, bidder i solves:

$$\max_{b_1, b_2, b_{12}} \sum_{s \subseteq \{1,2\}} (b_s - c_s) G_s(b_1, b_2, b_{12} | X^{it}) \quad (3)$$

subject to the constraint that $b_{12} \leq b_1 + b_2$, and that the bids must meet the reserve price, R_s , on each route and on the package.

This maximization problem is almost everywhere differentiable so optimal bids must solve the three first order conditions. Notice that the only ‘‘unknowns’’ in expression (3) are the three costs, c_1 , c_2 and c_{12} . The same applies for the first order conditions. Conditions under which this system of equations can be used to infer the costs can be concisely described as follows: If we observe a full set of bids, all bids have a positive probability of winning and the package bid constraint does not bind, then we can identify all the costs. The argument is similar to the one we made above concerning the single unit auction.

¹⁶ $S \setminus s$ denotes the set of both routes, S , minus the route(s) in s . By convention $B_\emptyset^{-i} = 0$.

¹⁷Practically, this involves multidimensional integrals that are difficult to evaluate numerically, so we implement this using Monte Carlo integration. Notice that our method requires that we solve the winner determination problem at each draw.

A binding constraint introduces a degree of underidentification. Specifically, if a bidder only submitted a bid on route 1 and route 2, say, but not on the package, we can only identify an upper bound on the individual route costs, a lower bound on the package cost and therefore an upper bound on the amount of synergy. Intuitively, if the cost for the package were really much lower than the sum of the costs of the individual routes, $c_1 + c_2$, the bidder would be especially eager to win both routes rather than one only. Submitting a bid on the package lower than the sum of the two individual bids, $b_{12} < b_1 + b_2$, helps him achieve this. By contrast if running both routes together is much more expensive, then bidder would prefer to win only one route. Since he cannot submit a bid on the package greater than the sum of the two, he will submit $b_{12} = b_1 + b_2$. The constraint is binding. This explains why we can bound the synergy from above (“it must be at most x %”) when the constraint is binding.

Finally, if the reserve price is binding for the package or for an individual route *but not for the package*, then we can infer that the cost for that particular route or package was higher than the reserve price. Intuitively, submitting a bid just below the reserve price for this route or for the package would not affect much the probability of winning the other routes, but because the probability of winning that route is not close to zero,¹⁸ the benefit of such a bid is positive as long as the cost is lower than the reserve price.¹⁹

Table 22.5 illustrates this cost inference method for one particular auction. Five bidders participated in this auction. Bidders 1 and 3 submitted a full set of bids. Bidder 3 only submitted a bid on the second route, and bidders 4 and 5 submitted a bid only on routes 1 and 2. Their package bid is defined by their bids on the individual routes.

The probabilities of winning were generated from the results of Table 22.4. The first order conditions for optimal bids were then used to solve for the costs.²⁰ Because bidders 1

¹⁸Indeed there is always a chance that no other bidders submitted a bid on that route or package.

¹⁹If the bidder submitted a bid on route 2 but not on route 1 and the package, we cannot say anything. Indeed, it may be that his cost for route 1 is lower than the reserve price but that his cost of operating both routes is very high. Submitting a bid on route 1 will define (through the package bid constraint) a bid on the package and run him the risk of winning both routes!

²⁰Given that the probabilities of winning were generated by Monte Carlo integration (step 3) and are therefore step functions, numerical differentiation was used. Due to the nature of the function for the first order condition, the delta method for computing the standard errors on the cost parameters could not be implemented numerically. Instead, we computed the standard errors on the basis of the cost estimates inferred from 250 draws from the bid distribution parameters. Errors due to the Monte Carlo integration

and 3 submitted a full set of bids, we can identify their costs for each route and the package. These are given in columns 6, 7 and 8. For bidder 2, we could only identify his cost for route 2 and conclude that his cost for the package was above the reserve price. We cannot say anything about his cost for route 1. Finally, because bidders 4 and 5 did not submit a non-trivial package bid, we could only identify bounds on their costs.

Table 22.5: A sample of cost estimates for one two-route auction*

	Bids (June 1995 Pounds)			Package discount (%)	Costs			Cost synergy (%)
	b_1	b_2	b_{12}		c_1	c_2	c_{12}	
1	338,175	411,925	749,763	0.045	284,908 (6,596)	347,927 (7,408)	664,733 (9,877)	-4.45 (0.80)
2	n/a	532,343	n/a	n/a	n/a	491,452 (6,504)	$>R_{12}$	n/a
3	334,457	429,771	759,174	0.661	282,993 (6,344)	368,259 (7,392)	678,515 (9,751)	-4.19 (0.75)
4	448,847	371,101	819,948	0	$<402,689$ (7,294)	$<301,861$ (7,472)	$>740,197$ (9,240)	<-5.06 (1.10)
5	381,939	439,280	821,219	0	$<341,737$ (5,329)	$<384,967$ (6,406)	$>744,290$ (8,832)	<-2.42 (0.58)

* Standard deviations in parenthesis. $R_1 = 838,626$; $R_2 = 1,867,646$. “ $<$ ” indicates that the number given is an upper bound to the true parameter. “ $>$ ” indicates that it is a lower bound.

For this particular auction, the mark-up over costs suggested by Table 22.5 is on average 16.1% for the individual routes and 11.4% for the package (due to the bounds in Table 22.5, this measure underestimates the mark-up on individual routes and overestimates the mark-up on packages). This is consistent with the announced mark-ups in bidding documents.²¹ The mark-ups are higher for the individual routes than for the package.

The last column of Table 22.5 reports the relative cost synergy estimate, $\frac{c_1+c_2-c_{12}}{c_1+c_2}$. A negative number indicates that the inferred cost for the package is higher than the sum of the individual costs. The relative synergy estimates are negative for bidders 1, 3, 4 and contribute to an additional 1% error on the coefficients.

²¹In their bidding documents, bidders were asked to provide an estimate of the costs associated with their bids. We did not have access to that information on a systematic basis as bidders report costs on a small sample only. Additionally, the measure may not be accurate as bidders may misreport costs. However, bidders’ announced margins range between 10% and 15%.

5. The negative sign is somewhat surprising given the motivation for using a combinatorial auction in the London bus routes market. They suggest that, for this particular auction at least, package bidding may have been driven by the strategic motivation.

We end with two important caveats concerning the empirical results reported in this section. First, the specification used in Table 22.4 is for illustration purposes only. Second, Table 22.5 looks at a single auction only. It permits us to illustrate the way in which we infer costs. Nevertheless, the selected auction may not be representative of our sample. A fuller analysis is carried out in our companion paper. The results also suggest that the synergies are negative.

6 Conclusions

In London, the local transportation authority has used a form of combinatorial auction for the allocation of bus routes services since the mid-1980's. The reason for this choice was expected economic synergies among routes located in the same area of London.

The London bus routes tendering is considered a success. It has led to increased quality of service and lower costs. However, there was a range of design options for LRT to choose from, from the size of the auction, the definition of the items, to the auction rule. It is unclear whether their actual choice was the most appropriate.

This chapter has documented bidding behavior in the LRT auctions and presented a new method for inferring bidders' cost structure. We did not find that auction size affected bidding behavior much: in fact, competition and prices at the item level seemed uncorrelated with the size of the auction. Moreover, when we applied our cost inference method to the data, we found evidence for slight negative synergies across routes.

These results can be used to choose the parameters in human subjects experiments (Smith, foreword) or computational experiments (Leyton-Brown and Shoham, Chapter 18) to evaluate alternative auction formats. Obvious candidates given our data analysis are a combinatorial auction that allows bidders to make exclusive bids and independent parallel or sequential single item auctions.

Acknowledgments: Funding from the Division of Research at Harvard Business School and the National Science Foundation (SES 0214222) is gratefully acknowledged.

7 References

Adams, William and Janey Yellen (1976), "Commodity Bundling and the Burden of Monopoly," *Quarterly Journal of Economics*, 90, 475-498.

Armstrong, Mark (1996), "Multi-product nonlinear pricing," *Econometrica*, 64(1), 51-75.

Armstrong, Mark and Jean Charles Rochet (1999), "Multi-dimensional Screening: A User's Guide," *European Economic Review*, 43, 959-979.

Beaumont, Helen and Michael Watson (1999), "Review of Bus Tendering in the Greater London Area," LRT Manuscript.

Cantillon, Estelle and Martin Pesendorfer (2004), "Combination Bidding in Multi-Unit Auctions," Working Paper, Harvard Business School and London School of Economics.

Kennedy, David, Stephen Glaister and Tony Travers (1995), "London Bus Tendering, Greater London Group," Report, London School of Economics.

Glaister, Stephen and Michael Beesley (1991), "Bidding for Tendered Bus routes in London," *Transportation Planning and Technology*, 15, 349-366.

McAfee, R.Preston., John McMillan and Michael D. Whinston (1989), "Multiproduct Monopoly, Commodity Bundling and Correlation of Values," *Quarterly Journal of Economics*, 102, 371-383.