Estimation and Identification of a Dynamic Multi-object Auction Model

Sam Altmann
University of Oxford

Introduction/Motivation

First-price auctions are a common auction format, but rarely occur in isolation:

- Multiple objects are often auctioned simultaneously
- Auctions are repeated regularly whenever new lots become available.

This is a problem for standard empirical methods when pay-offs are non-additive across auctions. Consider a bidder with **diminishing marginal returns**:

1. If they previously won many lots, they will bid less aggressively in future
2. If many auctions occur simultaneously, they may only bid on a small subset

It is difficult to model both repeated *and* simultaneous auctions. Previous work, including Gentry et al. (2018) and Jofre-Bonet and Pesendorfer (2003), focuses on either a static multi-object model, or a dynamic single-object model.

$\rightarrow$ It is *a-priori* unclear what effect these simplifications will have, and the magnitude of the inaccuracy created on estimates and counterfactuals.

This paper develops an empirical model of bidding in this environment, combining the models of Gentry et al. (2018) and Jofre-Bonet and Pesendorfer (2003)

Research Questions:

1. Is the model non-parametrically identified?
2. Is it computationally feasible to estimate such a model?
3. What is the inaccuracy of using a simpler single object or static model?

Model

I build an infinite horizon model of risk neutral bidders bidding in repeated simultaneous first-price auctions. The model has the following features:

- L lots are auctioned each period. At time t bidder i places bid b_{ilt} on lot l.
- There are 2^L possible outcomes for bidder i at time t, denoted $c \in \{1, ..., 2^L\}$ Combination outcomes include winning none, 1 lot, a different lot, 2 lots, *etc.*
- The state of the world at the beginning of period t is denoted $\mathbf{s}_t$. The ex-post state at the end of period t is denoted $\mathbf{s}^c_t$. This depends on the outcome c.
- If i wins lot l, they receive pay-off v_{ilt}. This value is known privately, and drawn independently from distribution F_l.
- If i ends the period in state $\mathbf{s}^c_{it}$ they receive deterministic pay-off $j(\mathbf{s}^c_{it})$
- Players believe, given the state $\mathbf{s}_t$, their bids $\mathbf{b}_{it}$, and others' strategies, that:
  1. They will win lot l with (marginal) probability $\Gamma_l(b_{ilt}|\mathbf{s}_t; \sigma_{-i})$
  2. The combination outcome will be c with probability $P_c(\mathbf{b}_{it}|\mathbf{s}_t; \sigma_{-i})$

The **Bellman equation** for this problem is given by:

$$W(\boldsymbol{v}, \mathbf{s}; \sigma_{-i}) = \max_{\mathbf{b}} \left\{ \sum_l \Gamma_l(b_l|\mathbf{s})(v_l - b_l) + \sum_c P_c(\mathbf{b}|\mathbf{s})[j(\mathbf{s}^c_i) + \beta V(\mathbf{s}^c)] \right\}$$

$\rightarrow$ Where the **Continuation Value** $V(\mathbf{s}^c) = E_{\tilde{\boldsymbol{v}}}\left[W(\tilde{\boldsymbol{v}}, \tilde{\mathrm{s}}; \sigma_{-i}) \ |\mathbf{s}^c\right]$

This optimisation problems yields L **First Order Conditions** for optimal bidding:

$$0 = \frac{\partial \Gamma_l(b_{ilt})}{\partial b_l}(v_{ilt} - b_{ilt}) - \Gamma_l(b_{ilt}) + \sum_c \frac{\partial P_c(\mathbf{b}_{it})}{\partial b_l}[j(\mathbf{s}^c_i) + \beta V(\mathbf{s}^c)] \tag{1}$$

Following Gentry et al. invert the FOCs ($\div$ by $\frac{\partial \Gamma_l}{\partial b_l}$) for the **Inverse Bid Function** ξ_l

$$\xi_l(\mathbf{b}|j + \beta V; \mathbf{s}) = b_l + \frac{\Gamma_l(b_l)}{\partial \Gamma_l(b_l)/\partial b_l} - \sum_c \frac{\partial P_c(\mathbf{b})/\partial b_l}{\partial \Gamma_l(b_l)/\partial b_l}[j(\mathbf{s}^c_i) + \beta V(\mathbf{s}^c)] \tag{2}$$

Identification

Given data $\left\{\{b_{ilt}\}_i, \mathbf{s}_t, \mathbf{s}^c_t\right\}_t$ are the distribution F and the deterministic pay-off j identified?

- As always beliefs Γ and P are identified through observed distributions of players bids.
- Gentry et al. (2018) Demonstrated that F is identified conditional on (j, β, V)
- I assume $E[v_{ilt}|\mathbf{s}] = 0$, and normalise $j(\mathbf{s}^0) = 0$ (we only identify marginal pay-offs).

  $\rightarrow$ This, plus observed variation in the state variables is sufficient for identification.

Intuitively we consider how bidding changes with observed changes to the state. If the bidder recently won, do they bid more / less aggressively than if they haven't won recently? **Algebraically**: $\xi_l(\mathbf{b}_{it}|j + \beta V; \mathbf{s}_t) = v_{ilt}$ at the true $j + \beta V$. I show that there exists a unique pay-off function j such that the moment equations $E[\xi_l(\mathbf{b}_{it}|j + \beta V; \mathbf{s})|\mathbf{s}] = 0$ for every $\mathbf{s}$.

Geometrically, consider an $L = 2$ example, with 13 possible states in the diagram below.

- Beginning in state s^0, if the bidder wins nothing their ex-post state is s^0. If they win one lot then $\mathrm{s}^c = \mathrm{s}^1$ or s^2 depending on which lot. If they win both lots then $\mathrm{s}^c = \mathrm{s}^4$. From each state, possible ex-post states are shown by the dotted arrows (ignoring self-loops).
- We want to identify the 12 marginal pay-offs of moving from s^0 to various other states. Each observation, beginning in a particular state, yields 2 pieces of information (two bids). However, there are always 3 marginal pay-offs to identify.

The **identification problem** is as follows:

By adding two (red) edges from each state, can we connect the graph, so the adjacency matrix has rank 12? (yes we can)

$\rightarrow$ Overlapping possible ex-post states are key

We **stitch** together information across states

Even though there is no edge $\mathrm{s}^0 \rightarrow \mathrm{s}^4$, this marginal pay-off is identified using our observation from s^2. Likewise $\mathrm{s}^2 \rightarrow \mathrm{s}^8$.

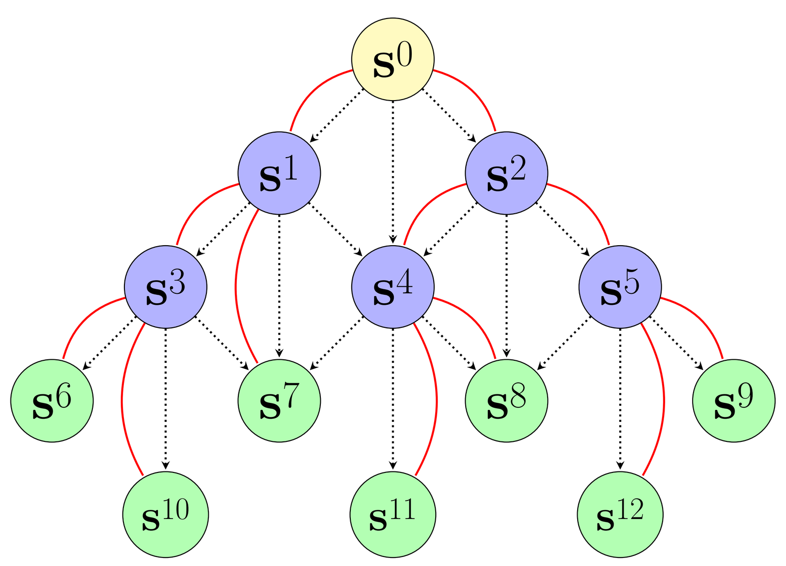

Estimation

The standard approach to estimating dynamic auctions comes from Jofre-Bonet and Pesendorfer (2003). We write $V(\mathbf{s})$ as a function of the observed distribution of bids and transition functions only, before backing out j and F.

Unfortunately in the multi-object context we cannot write $V(\mathbf{s})$ as a function of bids only. **However**, by substituting the FOCs into the bellman equation we can write the continuation value as a function of bids *and* the sum $j(\mathbf{s}^c_i) + \beta V(\mathbf{s}^c)$, which I denote as the function $k(\mathbf{s}^c)$:

$$V(\mathbf{s}) = E_{\mathbf{b}}\left[\sum_l \frac{\Gamma_l(b_l)^2}{\partial \Gamma_l(b_l)/\partial b_l} + \sum_c [P_c(\mathbf{b}) - \sum_l \Gamma_l(b_l) \frac{\partial P_c(\mathbf{b})/\partial b_l}{\partial \Gamma_l(b_l)/\partial b_l}]\, k(\mathbf{s}^c) \ |\mathbf{s}\right] \tag{3}$$

In fact, we can write both the first order conditions and the inverse bid functions as functions of bids *and* this function $k = j + \beta V$. This suggests a natural **3 step estimator**:

1. Estimate beliefs Γ and P from the distribution of bids.
2. Estimate $k(\mathbf{s})$ by GMM - setting $\hat{k}$ such that $E[\xi_l(\mathbf{b}|k, \mathbf{s})|\mathbf{s}] = 0$ for all $\mathbf{s}$

   That is, so that our identifying moment conditions hold for each observed state.

   Then estimate F from the inverse bid functions.
3. Evaluate V using equation 3, bids, and $\hat{k}$. Finally, back out $\hat{j} = \hat{k} - \beta \hat{V}$

In step 2 we almost estimate the model as if it was a static model, before giving it a dynamic interpretation at step 3. Given that step 3 is computationally very simple, the computational intensity of the estimation procedure is **little more** than that from estimating a static multi-object model. The results of a simulation study suggest this approach is effective.

Application

I investigate the extent of the inaccuracies from incorrectly estimating either a static or single-object model. I apply the model to data from **Michigan Department of Transport**'s procurement auctions for highway maintenance contracts:

- 10 years of data, with auction rounds taking place every 2-4 weeks.
- $\approx$ 45 auctions take place each round, with roughly 200 (potential) bidders.
- I investigate how a firm's backlog of contracts impacts their cost function

  $\rightarrow$ Focusing on road paving projects

This setting (and data) has been used in **numerous** previous auction studies. These previously employed either static or single-object models, and generally found evidence of dynamics and complementarities across lots.

I compare estimates of the marginal cost function across three models: $i)$ Dynamic Multi-Object, $ii)$ Dynamic Single-Object, $iii)$ Static Multi-Object

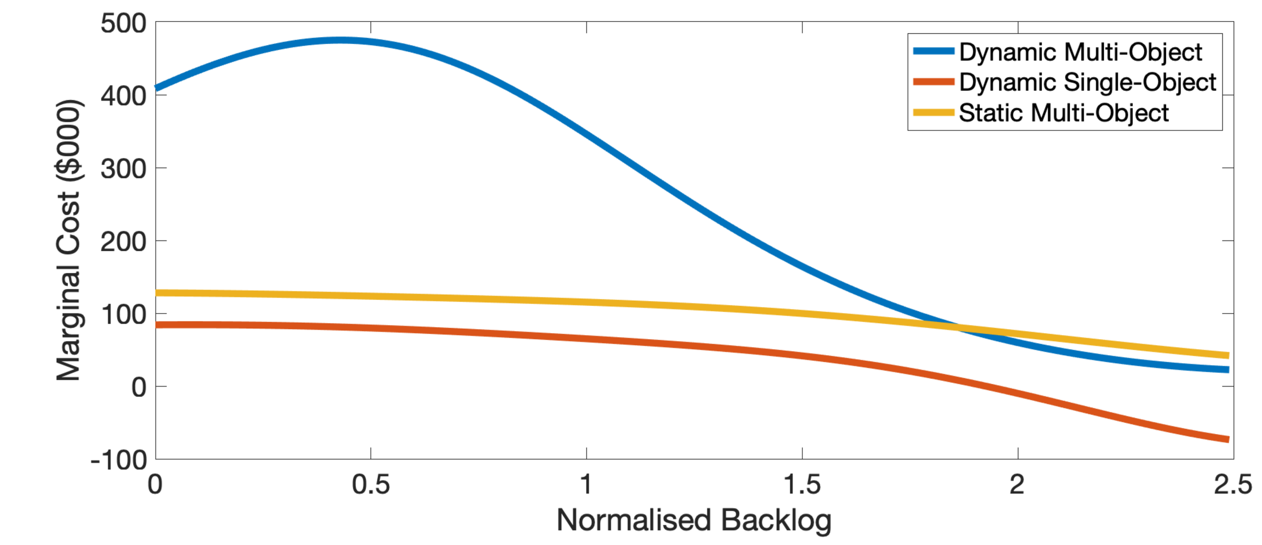

We observe that costs are decreasing in the backlog.

$\rightarrow$ This suggests firms take advantage of increasing returns to scale

When we compare the mis-specified models the costs are **under-estimated**.

The Static Multi-Object model:

Assuming myopic bidders ignores the option value of winning today vs in future:

- At low backlogs a firm that wins a contract today can afford to be more picky about contracts in future, not worrying about having no work.
- This option value off-sets the high marginal cost, so bidders are willing to bid more aggressively despite high marginal costs.

The Dynamic Single-Object model:

Assuming costs are independent across lots ignores the role of cost synergies:

- Low backlog firms bid aggressively on multiple lots to exploit returns to scale
- Aggressive bidding is mistaken for generally low costs, since we ignore that they are trying to win multiple lots.
- This effect compounds as we underestimate the option value, ignoring that firms can win multiple lots in future as well.

Conclusions:

1) Failing to account for both forward-looking behaviour *and* complementarities across auctions can lead to inaccurate cost estimates, which will lead to inaccurate counterfactuals and welfare conclusions.

2) Because the computational difficulty of estimating a dynamic multi-object model is little more than estimating a static multi-object model, it is difficult to justify not using the dynamic model.

References

References

Gentry, M. L., Komarova, T., and Schiraldi, P. (2018). Preferences and performance in simultaneous first-price auctions: A structural analysis. *Available at SSRN 2514995.*

Jofre-Bonet, M. and Pesendorfer, M. (2003). Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489.