

IDENTIFICATION AND ESTIMATION OF A DYNAMIC MULTI-OBJECT AUCTION GAME

Sam M. Altmann*
Paris School of Economics

[Click Here For The Most Recent Version](#)

Abstract

Auctions rarely take place in isolation. Often, many heterogeneous lots are auctioned simultaneously, and auctions are repeated as new lots become available. In this paper I develop an empirical model of bidding in repeated rounds of simultaneous first-price auctions. I prove non-parametric identification of primitives in this model, and introduce a computationally feasible procedure to estimate this type of game. I apply the model to data from Michigan Department of Transportation highway procurement auctions and investigate the extent of cost-synergies across auctioned contracts. Finally, I use counterfactual simulations to compare equilibrium efficiency when contracts are auctioned sequentially rather than simultaneously.

*Paris School of Economics, samuel.altmann@psemail.eu. The author also wishes to thank Ian Crawford, Emmanuel Guerre, Luke Milsom, Martin Pesendorfer, Itzhak Rasooly, and Howard Smith for useful comments and conversations, as well as seminar participants at The University of Oxford. Finally, thank you to Matt Gentry, Tatiana Komarova, and Pasquale Schiraldi for sharing their data.

1 Introduction

First-price auctions, which are regularly used to allocate government procurement contracts, rarely take place in isolation. Multiple lots (contracts) are often auctioned simultaneously, and auctions are repeated whenever new contracts become available. In real world environments bidders' values may be non-additive across different lots. For example, bidders may face capacity constraints, facing higher costs the larger their current backlog. Or, they may benefit from economies of scale, facing lower costs when working on many of the same type of contract at once. The structure of these non-additive values is highly relevant for auction design — should similar contracts be auctioned simultaneously, or spaced out over time? In this paper I develop an empirical model of forward looking bidding in repeated rounds of simultaneous first-price auctions, and study identification and estimation in this framework. I apply the model to Michigan Department of Transportation (MDOT)'s procurement auction data and investigate the empirical and policy relevance of these complementarities.

Running simultaneous auctions allows bidders to benefit from a batching effect; giving bidders information about all the contracts being auctioned that period and allowing them to focus their bidding on those in which they have a known cost advantage. However, simultaneous auctions also exhibit an exposure effect: Because firms cannot place combination bids they may risk winning too few contracts and be unable to exploit their cost synergies. Meanwhile, when capacity constraints are the dominant factor, auctioning a large number of contracts simultaneously may also create inefficiencies by depressing competition.

Previous research has either studied forward looking bidders and assumed auctions are single-object, or studied auctions of multiple objects and assumed bidders are myopic. For example, both Jofre-Bonet and Pesendorfer (2003) and Gentry et al. (2023) study synergies in bidding behaviour in repeated simultaneous first-price auctions for highway maintenance contracts.¹ Jofre-Bonet and Pesendorfer (2003) estimate a dynamic single object model, assuming that payoffs are additive in lots auctioned simultaneously, and find evidence of dynamic linkages through negative effects of capacity constraints on bids. Gentry et al. (2023) study simultaneous first-price auctions, assuming myopic bidding, and find similar capacity constraint effects across lots auction within a period. However, they also find evidence

¹Other examples of papers studying repeated and simultaneous auctions, abstracting away from either simultaneous bidding or dynamics (often justifiably), include Cantillon and Pesendorfer (2007) on London bus routes, Flambar and Perrigne (2006) on snow clearing services, Athey et al. (2011) on timber auctions, Somaini (2020) on highway procurement, Hendricks et al. (2003) on oil drilling rights, and Backus and Lewis (2016) on online marketplaces, among others.

of positive synergies among similar contracts that allow firms to exploit economies of scale. Both papers find evidence for non-additive values in one dimension, either over time or across lots, but restrict non-additivity in the other dimension. The implication is that neither paper accurately models the non-additive values. To the best of the author’s knowledge this paper is the first to unify the dynamic and multi-object approaches to empirical auctions.

I develop a structural empirical model of forward looking bidding in repeated simultaneous first-price auctions, where lots are heterogeneous and payoffs are non-additive across lots. The model is fundamentally the union of the models presented in Jofre-Bonet and Penderfer (2003) and Gentry et al. (2023), henceforth referred to as JP and GKS respectively. Bidder pay-offs are represented as the sum of privately known and potentially correlated *lot specific* values, a *combination specific* flow payoff, and a *combination specific* continuation value. Following GKS, the combination specific flow payoff is treated as a deterministic function of state variables. This is a natural framework that reflects known capacity constraints or economies of scale. The model primitives consist of the distribution of lot specific values and the combination specific flow payoff function.² The central difficulty for both identification and estimation is that there is not a one-to-one relationship between bids and values, because simultaneous first-price auctions (without combination bidding) are not a Direct Revelation Mechanism. Therefore, unlike Guerre et al. (2000), we cannot invert equilibrium bidding functions to point identify values. Likewise, unlike JP, we cannot invert equilibrium policy functions, and so write the continuation value as a function of the equilibrium distribution of bids only.

Building on this framework I make three key contributions to the empirical auction literature. First, I extend GKS’ identification framework to make use of variation in a bidder’s individual state variables, such as their backlog of contracts, to non-parametrically identify their combination specific value function *without* the need for exclusion restrictions. Intuitively, identification arises because variation in the state causes variation in bidders’ combination values, which in turn causes variation in their bidding behaviour. If lots are substitutes we expect to observe more aggressive bidding when backlogs are low. Extending the approach presented in GKS to the dynamic setting I translate the inverse bidding system, conditional on a given state, into a system of linear equations in the unknown combinatorial flow payoffs. Key to the identification argument is that we combine these systems of equa-

²Like GKS this allows me to separately identify complementarities and affiliations, the central problem studied by Kong (2021). Affiliation across lots comes through correlation in the lot specific pay-offs, while the synergies remain deterministic. Like both papers I assume the lot-specific pay-offs are independent across players.

tions across state variables, essentially stitching together observations of bidding behaviour from different states. I prove that, under mild conditions, this system has a unique solution. This result is important because it ensures the combination value is identified without the need for exclusion restrictions that prohibit identification of forward looking behaviour.³

Second, I propose a three step procedure for estimating the model and establish that it is \sqrt{T} consistent and asymptotically normal. This estimator generalises JP’s procedure: While the ex-ante value function cannot be expressed using the equilibrium bid distribution *only*, it can be written as a function of the bid distribution *and* a term that corrects for the complementarities between lots. This correction term is a function of the sum of the combinatorial flow payoff and the discounted continuation value. The novelty of the estimation procedure then concerns how we estimate this correction term. I refer to this term as the ‘pseudo-payoff’: It is the object we estimate if we incorrectly estimated a misspecified static model. This suggests a simple estimation procedure. First, one estimates bidders’ equilibrium beliefs, or the equilibrium distribution of bids. Second, we estimate the pseudo-payoff — the sum of the flow payoff and the continuation value — by essentially estimating the multi-object auction model *almost as if* it were a static model. Third, we evaluate the continuation value using the estimated correction term, before separating out the combinatorial flow payoff from the estimated pseudo-payoff. Overall, this is little more computationally costly than estimating a static multi-object model as in GKS. This estimation procedure also has broader applicability beyond the auction environment, giving a convenient way to estimate dynamic games with non-invertible policy functions or conditional choice probabilities.⁴

Finally, I apply this framework to data from Michigan Department of Transport (MDOT)’s procurement auctions. In this setting around 35 contracts for highway maintenance and construction projects are auctioned simultaneously in each round, and rounds are repeated roughly every fortnight. I focus on contracts that require use of either hot-mix asphalt, concrete, or both, and consider how firms’ backlogs of asphalt and concrete projects impact their cost functions. These backlog effects create a dynamic linkage that makes payoffs non-additive across periods, as well as driving complementarities between lots auctioned

³Without additional restrictions, excluded variables (such as the set of lots on offer or backlogs of rival bidders) enter each bidder’s continuation value, directly affecting their bidding behaviour and violating the exclusion restriction, rendering the model non-identified in GKS’s framework.

⁴In settings when policy functions, or conditional choice probabilities, are invertible this procedure is equivalent to the estimator proposed by Bajari et al. (2007). Differences only arise when invertibility fails. This occurs in many strategic settings, such as games with non-monotone equilibria, and single-agent settings, such as when agents choose among lotteries with multiple outcomes.

simultaneously. For paving firms in particular I find evidence of increasing returns to specialising in asphalt contracts: Every one standard deviation increase in their asphalt backlog increases the cost of completing a concrete contract by around 10%, and decreases the cost of an asphalt contract by roughly the same amount. I use counterfactual simulations to consider how procurement and construction costs differ when contracts are auctioned sequentially instead of simultaneously, eliminating both the batching and exposure effects. I find that the batching effect dominates, so that construction costs are 9% larger under sequential allocation. While competition appears fiercer using sequential auctions, I still find a net increase in procurement costs of around 1%.

The structure of this paper is as follows: Section 2 introduces the auction game that is the focus of this paper. Section 3 introduces the identification framework and proves that model primitives are point identified. Section 4 outlines the proposed three step estimation procedure and establishes large sample properties. Section 5 applies this procedure to data from MDOT procurement auctions. Several additional results are presented in the Appendices. Appendices A - C present technical proofs. Appendix D presents extensions to the identification and estimation framework. Appendix E presents a simulation experiment evaluating the proposed estimation procedure, and F presents additional analysis related to the empirical application.

1.1 Related Literature

My key contribution is to unify the literatures on the identification and estimation of both dynamic auction models and multi-object auction models.⁵ JP was the first to estimate a dynamic auction game, analysing sequential highway procurement auctions and find backlog effects to be determinants of future bidding behaviour. Several papers have built on this framework, including Jeziorski and Krasnokutskaya (2016) on dynamic auctions with subcontracting, Groeger (2014) on participation in repeated auctions, Balat (2013) on unobserved heterogeneity in lot quality, and Raisingh (2021) on pre-announcements. These papers study settings in which multiple auctions are held simultaneously, and assume pay-offs are additively separable across auctions within a period. This assumption is unpalatable

⁵This work is also tangentially related to the literature on empirical Multi-unit auctions, which focuses on divisible homogenous units (see e.g. Hortaçsu and McAdams (2018)). The estimation procedure presented in section 4 easily extends to dynamic multi-unit auctions. Another related literature analyses forward looking behaviour in second-price auctions, including Backus and Lewis (2016) and Bodoh-Creed et al. (2021). In Appendix D.1 I extend my identification and estimation results to the multi-object second-price setting.

given they find evidence of non-additivities across auctions held in different periods.

Cantillon and Pesendorfer (2007) were the first to estimate a model of simultaneous auctions. They use combination bids to identify complementarities in simultaneous first-price auctions, studying procurement auctions for London bus routes. Kim et al. (2014) use this framework to study the allocation of contracts for Chilean school meals. Fox and Bajari (2013) study an auction environment *without* combination bidding, using an equilibrium stability condition to allow them to identify complementarities. GKS also focus on simultaneous first-price auctions without combination bidding. They prove the model is identified using variation in ‘excluded’ variables: Variables that are excluded from the bidder’s combinatorial payoff, such as characteristics of their rivals, and only indirectly affect bidding behaviour through bidders’ equilibrium beliefs. However, exclusion restrictions fail in a dynamic environment. Bidders’ forward looking behaviour ensures every state variable directly effects their continuation value, and hence bidding behaviour. These exclusion restrictions are not necessary for identification. Arsenault Morin et al. (2022) extend GKS to allow for endogenous participation in simultaneous auctions, and study auctions for roof-maintenance contracts in Montreal.

2 The general model

2.1 Setup

Rules: Suppose that each period t , over an infinite horizon, N risk-neutral players compete in L first-price Sealed Bid auctions. Player n wins lot l in period t if $b_{ntl} \geq \max_{m \neq n} \{b_{mtl}\}$. Sealed bids are placed simultaneously, then winners are announced. Winners pay their bids, and every player observes the bids and identities of winners.

Reservation Prices and Ties: In the main text I assume reservation prices do not bind, that ties occur with probability zero, and that participation is determined exogenously.

Lots and Lot Characteristics: Each lot is characterised by a vector of characteristics, which may include the size and location of a particular contract, for example. These characteristics and other common state variables are stacked into the common state $\mathbf{s}_{0t} \in \mathbb{S}_0$. I assume $|\mathbb{S}_0|$ is finite.

Outcomes: There are N^L different ways that N players can win L lots, and so every period there are N^L possible combination outcomes. Each possible outcome, denoted by c , corresponds to bidders winning different combinations of the L lots. The set of possible

combination outcomes is given by the ‘power set’ of the set of available lots, denoted by $\{N^L\}$, so that $c \in \{N^L\}$.

2.1.1 States and Primitives

Individual States: Player n begins the period in state \mathbf{s}_{nt} . This may represent a player’s existing stock of the good, or backlog of contracts. I assume the set of possible individual states, \mathbb{S}_n , is finite.⁶ If the outcome from round of auctions in period t is c , then player i ends the period in state \mathbf{s}_{nt}^c , referred to as the ex-post state. $\mathbf{s}_{nt}^c = \mathbf{s}_{nt}$ if and only if the player does not win a single lot.

Total States: Stack the individual states $\{\mathbf{s}_{nt}\}_{n \in \{1, \dots, N\}}$, and \mathbf{s}_{0t} , into the total state variable $\mathbf{s}_t \in \mathbb{S}$, where $|\mathbb{S}| = S$ is finite. In section 3.6 I give sufficient conditions on \mathbb{S} to ensure identification. Similarly, Stack the ex-post states for $\mathbf{s}_t^c \in \mathbb{S}$.

Transition Process: At the beginning of each period, the state \mathbf{s}_t is drawn stochastically from $T_{\mathbf{s}}(\cdot | \mathbf{s}_{t-1}^c)$. Because $|\mathbb{S}|$ is finite, the transition probabilities can be described by transition matrix T , such that $P(\mathbf{s}_t = \mathbf{s}_i | \mathbf{s}_{t-1}^c = \mathbf{s}_j) = T_{ij}$.

Actions: Each player plays an L dimensional vector of bids each period, denoted \mathbf{b}_{nt} . The set of possible bids is convex and compact, so that $\mathbf{b}_{nt} \in [\underline{b}, \bar{b}]$.

Lot Specific Payoff: I focus on an independent private value framework. If n wins lot l at t they receive a privately observed lot specific payoff, v_{ntl} . Stacking these values \mathbf{v}_{nt} , a $L \times 1$ vector, is drawn from cumulative density function $F_n(\cdot | \mathbf{s}_t)$ with support $[\underline{\mathbf{v}}_n, \bar{\mathbf{v}}_n]$.

Deterministic Payoff: If a bidder ends a period in state \mathbf{s}_{nt}^c , they receive a deterministic flow payoff $\pi_n(\mathbf{s}_{nt}^c)$, where $\pi_n : \mathbb{S}_n \rightarrow \mathbb{R}$. Whereas \mathbf{v} is stochastic, I assume π_n is a deterministic function of \mathbf{s}_n and is finite. Because the state space is finite, there are $|\mathbb{S}_n|$ possible values of $\pi_n(\mathbf{s}_{nt}^c)$, which I denote by $\boldsymbol{\pi}_n$, stacking π across states. These payoffs can be interpreted as how a bidder’s payoffs depend deterministically on their state. For example, if firms face higher costs when their backlogs of contracts are larger or when they are work on many contracts simultaneously. Dynamic linkages and static complementarities are expected to be expressed through this payoff term. For example, if $\pi(\cdot)$ exhibits increasing returns to scale.

Combination Payoffs: Because there are 2^L possible combinations of lots that player n might win, there are 2^L possible ex-post deterministic payoffs $\pi_n(\mathbf{s}_{nt}^c)$, corresponding to each of the possible ex-post states. Define the 2^L vector $\Pi_n(\mathbf{s}_t)$ such that each element of this

⁶This is predominantly for mathematical convenience, but is likely to hold in practice. Highway maintenance companies likely have a maximum number of contracts they can feasibly hold at any given time, and their backlog of contracts can be arbitrarily discretised into days of work remaining.

vector corresponds to the deterministic flow payoff from winning a different combination of lots. This object varies with \mathbf{s}_t since \mathbf{s}_t determines the characteristics of the available lots, and hence the possible ex-post states. We can define a linear mapping between Π and π such that: $\Pi_n(\mathbf{s}) = B_{\mathbf{s}}\boldsymbol{\pi}_n$, where the known $2^L \times |\mathbb{S}_n|$ dimensional selection matrix $B_{\mathbf{s}}$ selects elements of $\boldsymbol{\pi}_n$ according to the possible ex-post states for player n , given the period started in state \mathbf{s} .

2.1.2 The Bidder's Problem

Strategies: A (pure) Markovian strategy σ_n consists of a mapping from a player's type (\mathbf{v}_n, Π_n) and the state of the world \mathbf{s} onto a series of bids \mathbf{b}_{nt} . Ex-ante a player's strategy admits a distribution of bids according to F_n , Π_n , and \mathbf{s} .

Marginal Win Probabilities: Denote $G_{ml}(\cdot; \sigma_m)$ and $g_{ml}(\cdot; \sigma_m)$ respectively the marginal cdf and pdf of bidder m 's bid on lot l according to their strategy σ_m . Denote $\Gamma_n(\mathbf{b}; \sigma_{-n})$ the $L \times 1$ vector where row l contains the probability that n wins lot l , given their bid and the strategies of other players. Because ties occur with zero probability we can write:

$$\Gamma_{nl}(b_{nlt}; \sigma_{-n}) = \prod_{n' \neq n} G_{n'l}(b_{nlt}; \sigma_{n'})$$

Combination Win Probabilities: Denote $P_n(\mathbf{b}; \sigma_{-n})$ the $2^L \times 1$ vector of probabilities of possible combination wins, conditional on n 's bids and σ_{-n} . Each row of this vector corresponds to the probability of n winning a different one of the 2^L possible combinations of lots. So, row c of this vector contains the probability that n 's ex-post state will be \mathbf{s}_{nt}^c .

Overall Combination Probabilities: There are N^L different ways N players can win L lots, so N^L different possible combination outcomes. Therefore, denote $Q_n(\mathbf{b}; \sigma_{-n})$ the $N^L \times 1$ vector of probabilities of possible outcomes from the round of auctions, conditional on n 's bid and σ_{-n} . Row c of this vector then contains the conditional probability that the outcome from period t is c , and so the overall ex-post state is \mathbf{s}_t^c . This object is extremely similar to the combination win probabilities P presented previously, except Q also accounts for exactly which player wins each lot.

Discounting: Players have temporally additively separable preferences, and discount future payoffs using known discount factor $\beta \in (0, 1)$.

Expected Flow Pay-off: I assume that bidders are risk neutral and payoffs are quasi-linear in payments. Consider player n with a realisation of $\mathbf{v} = \mathbf{v}_{nt}$ who places bid \mathbf{b} against

players bidding according to strategies σ_{-n} :

$$\bar{W}(\mathbf{b}|\mathbf{v}_{nt}, \mathbf{s}; \sigma_{-n}) = \Gamma_n(\mathbf{b}; \sigma_{-n})^T (\mathbf{v}_{nt} - \mathbf{b}) + P_n(\mathbf{b}; \sigma_{-n})^T \Pi_n(\mathbf{s})$$

Value Function: The Bellman equation is given by: $W_n(\mathbf{v}_{nt}, \mathbf{s}_t; \sigma_{-n}) =$

$$\max_{\mathbf{b}} \left\{ \bar{W}(\mathbf{b}|\mathbf{v}_{nt}, \mathbf{s}_t; \sigma_{-n}) + \beta \sum_{c \in \{N^L\}} Q_{nc}(\mathbf{b}; \sigma_{-n}) \int_{\bar{\mathbf{s}}} \int_{\mathbf{v}_{nt}} W_n(\mathbf{v}_n, \bar{\mathbf{s}}; \sigma_{-n}) dF(\mathbf{v}_n|\bar{\mathbf{s}}) T_{\bar{\mathbf{s}}}(\bar{\mathbf{s}}|\mathbf{s}_t^c) d\bar{\mathbf{s}} \right\} \quad (1)$$

Continuation Value: It is useful to define the continuation value: $V_{nc}(\mathbf{s}_t; \sigma_{-n}) = \int_{\bar{\mathbf{s}}} \int_{\mathbf{v}_n} W_n(\mathbf{v}_n, \bar{\mathbf{s}}; \sigma_{-n}) dF(\mathbf{v}_n|\bar{\mathbf{s}}) T_{\bar{\mathbf{s}}}(\bar{\mathbf{s}}|\mathbf{s}_t^c) d\bar{\mathbf{s}}$. The combination continuation value is given by $V_n(\mathbf{s}_t; \sigma_{-n})$, a $N^L \times 1$ vector. Each element c of this vector contains the continuation value corresponding to a different allocation, ending the period in a different state \mathbf{s}_t^c . Since bidders' continuation values depend on exactly which combination of lots they win, this is another area in which static complementarities across lots will be expressed.

2.2 Equilibrium

I now discuss equilibrium, and the assumptions required for existence of an equilibrium. A full and general proof of equilibrium existence is beyond the scope of this paper.⁷ Instead, I present a proof of existence under the conjecture that equilibrium exists in the static game. I focus on symmetric Markov Perfect Equilibria (MPE):

Definition 2.1. : *A symmetric MPE consists of a set of strategies σ^* and beliefs (Γ, P, Q) , such that for all n , and any $(\mathbf{v}, \pi, \mathbf{s})$: (Optimality) σ_n^* is a best response to σ_{-n}^* , (Consistency) Beliefs Γ_n, P_n , and Q_n are consistent with σ_{-n}^* , (Markovian) σ_n^* depends on the state \mathbf{s}_t , not on t itself, (Symmetry) $\sigma_n^* = \sigma_m^*$ for all $m \neq n$, so that σ_n^* depends on their type, (\mathbf{v}_n, Π_n) , not their identity.*

⁷To my knowledge, no complete proof of equilibrium existence exists even for the static game. This paper joins the papers studying sufficiently complex auction games in which neither existence, nor uniqueness of equilibrium can be guaranteed. For example, GKS on simultaneous first-price auctions, Fox and Bajari (2013) on simultaneous ascending auctions, or JP on dynamic single-object first-price auctions. If the bid space were discrete, then static equilibrium existence follows from Milgrom and Weber (1985).

2.2.1 Equilibrium Existence

To prove equilibrium existence in the dynamic game, I rely on the following assumption about equilibrium existence in a static game:

Assumption 1. *There exists a symmetric (non co-operative) Pure Strategy Bayesian Nash Equilibrium of the (myopic) stage game, such that for all n and l the expected pay-off is continuous in \mathbf{v}_n and Π_n .*

This conjecture takes essentially the same form as the assumption that a continuous and unique equilibrium exists in Gentry et al. (2023).

Proposition 1. *Under the assumptions of the game, and under Conjecture 1, a Symmetric Markov Perfect Equilibrium exists.*

Proof is relegated to Appendix C, as existence is not the main focus of this paper. The proof consists of showing that the equilibrium pay-off in the stage game is consistent with the continuation value, employing Kakutani’s fixed point theorem.

3 Identification

I now demonstrate that the distribution of lot specific payoffs F , and the flow payoff function π are non-parametrically point identified.⁸ The intuition is that variation in \mathbf{s} causes variation in payoffs which, in turn, cause variation in bidding behaviour. For example, if lots are complements then the more a bidder has won in the recent past, the more aggressively we expect them to bid in the present. I then use the observed bidding behaviour, as well as information about bidders’ equilibrium beliefs, to essentially ‘back out’ the distribution of values and the complementarities. My results ensure identification of F_n and π_n separately for each bidder, however I drop the n subscripts except where necessary.

I introduce the assumptions necessary for identification in subsection 3.1. In 3.2 - 3.3 I use the bidder’s optimisation problem to derive the Inverse Bid System. In 3.4 - 3.5 I combine this system across states to form a system of linear equations in Π . In 3.6 I present

⁸A model is point identified if, given the implications of equilibrium behaviour, the distribution of bidder’s pay-offs, $\{F_n, \pi_n\}_{n \in \{1, \dots, N\}}$, are uniquely determined by the distribution of observables (Athey and Haile, 2002). A model is non-parametrically identified if the identified objects are functions (Lewbel, 2019), in the sense that we do not assume a functional form, but identify $\pi_n(\mathbf{s}_n)$ for every $\mathbf{s}_n \in \mathbb{S}_n$ and $F_n(\mathbf{v}|\mathbf{s}_n)$ for every pair (\mathbf{v}, \mathbf{s}) .

sufficient conditions for this system to have a unique solution. In subsection 3.7 I consider identification under several extensions of the model.

3.1 Assumptions necessary for identification

Assumption 2. *For each t , the econometrician has a set of observations \mathbb{O}_t consisting of the initial state \mathbf{s}_t , the bids of all bidders $\{\mathbf{b}_{nt}\}_{n \in \{1, \dots, N\}}$, and the outcomes from the round of auctions.*

Assumption 3. *The data $\{\mathbb{O}_t\}_{t=1 \dots T}$ are generated by strategy profile $\boldsymbol{\sigma}^*$ which is a symmetric Markov perfect equilibrium of the dynamic auction game. Furthermore, the same MPE is played in each period.*

Define $G(\cdot|\mathbf{s})$, $\Gamma(\cdot|\mathbf{s})$, $P(\cdot|\mathbf{s})$, and $Q(\cdot|\mathbf{s})$ as the empirical counterparts to the objects presented previously. Under these assumptions G , Γ , P , Q , and T are all identified, and for the remainder of this section I treat these objects as known.

Assumption 3 also ensures the continuation value can be written as a function of the state. We can then express the continuation value in vector form as \mathbf{V} , with elements corresponding to the expectation from ending a period in any particular ex-post state. I can then define the relationship between the N^L vector $V(\mathbf{s})$ defined previously and \mathbf{V} , so that $V(\mathbf{s}) = A_{\mathbf{s}}\mathbf{V}$ using the known $N^L \times S$ selection matrix $A_{\mathbf{s}}$. This contains a 1 in entry cm if the potential outcome c yields ex-post state $\mathbf{s}^c = \mathbf{s}_m$, selecting the relevant continuation values corresponding to the possible ex-post states. This is similar to how we previously defined the relationship $\Pi(\mathbf{s}) = B_{\mathbf{s}}\boldsymbol{\pi}$.

We can define the relationship $P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}} = Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}} C$ for the $S \times S_n$ matrix C . Entry ij of C is equal to 1 if $\mathbf{s}_n^i = \mathbf{s}_n^j$, and zero otherwise.⁹

Assumption 4. *For all \mathbf{s} , n , and l $G_n(\mathbf{b}_n|\mathbf{s}; \boldsymbol{\sigma}^*)$ is absolutely continuous in b_{nl} .*

This assumption ensures that the marginal, combination, and over-all combination win probabilities are continuous and differentiable in \mathbf{b} , enabling us to take first order conditions. As shown in GKS, when this assumption does not hold we lose point-identification, though the model primitives generally remain partially identified.

⁹ $P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}}$ is a linear map $\mathbb{S}_n \rightarrow [0, 1]$ that gives the probability bidder n ends the period in a given state \mathbf{s}_n , given the period began in total state \mathbf{s} . Meanwhile, $Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}}$ is a linear map $\mathbb{S} \rightarrow [0, 1]$ that gives the probability of ending the period in any of the total states, given the period began in \mathbf{s} . C is then a transformation $\mathbb{S}_n \rightarrow \mathbb{S}$ that, for a given individual state $\mathbf{s}_n \in \mathbb{S}_n$, sums over all the total states $\mathbf{s} \in \mathbb{S}$ for which bidder n 's state is also \mathbf{s}_n .

Assumption 5. $E[\mathbf{v}|\mathbf{s}] = 0$ and $\pi(\mathbf{s}_n^1) = 0$.

\mathbf{v} is mean independent of \mathbf{s} , as we cannot separately identify $E[\mathbf{v}|\mathbf{s}]$ from $\Pi(\mathbf{s})$. This is similar to the assumption $E[\mathbf{v}|\mathbf{s}] = E[\mathbf{v}]$, except we ‘absorb’ the mean of \mathbf{v} into Π through a linear term. Finally, we must normalise $\pi(\mathbf{s}_{n1})$ because only marginal payoffs are identified. Based on these assumptions, I will prove the following proposition:

Proposition 2. *Under assumptions 2 - 5, the model primitives F and $\boldsymbol{\pi}$ are non-parametrically point identified.*

This result differs from GKS’ identification result even when bidders are myopic ($\beta = 0$), differing in the source of identifying variation. They prove identification using excluded variables which cause ‘exogenous’ variation in Γ and P . I use *included* variation in the state which directly enters $\pi(\mathbf{s}_n) + \beta V(\mathbf{s})$, and so shifts bidding behaviour directly. I show that, under a restriction on the ordering of the state space, variation in bidding behaviour that arises from variation in the state space uniquely determines $\boldsymbol{\pi}$. In this respect, the identification argument is closer in spirit to the argument presented in Pesendorfer and Schmidt-Dengler (2008).

3.2 First Order Conditions

The agent’s problem is to maximise their expected discounted pay-off, and so in each period the agent maximises the following object, with respect to \mathbf{b} :

$$\begin{aligned}\tilde{W}(\mathbf{b}|\mathbf{v};\mathbf{s}) &= \Gamma(\mathbf{b}|\mathbf{s})^T(\mathbf{v} - \mathbf{b}) + P(\mathbf{b}|\mathbf{s})^T\Pi(\mathbf{s}) + \beta Q(\mathbf{b}|\mathbf{s})^TV(\mathbf{s}) \\ &= \Gamma(\mathbf{b}|\mathbf{s})^T(\mathbf{v} - \mathbf{b}) + P(\mathbf{b}|\mathbf{s})^TB_{\mathbf{s}}\boldsymbol{\pi} + \beta Q(\mathbf{b}|\mathbf{s})^TA_{\mathbf{s}}\mathbf{V}\end{aligned}\quad (2)$$

Assumption 4 ensures that $P(\mathbf{b}|\mathbf{s})$, $Q(\mathbf{b}|\mathbf{s})$, and $\Gamma(\mathbf{b}|\mathbf{s})$ are continuously differentiable in \mathbf{b} . Necessary First Order Conditions of optimal bidding are then given as:

$$\underbrace{\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})}_{L \times L} \underbrace{(\mathbf{v} - \mathbf{b}^*)}_{L \times 1} = \underbrace{\Gamma(\mathbf{b}^*|\mathbf{s})}_{L \times 1} - \underbrace{\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})}_{L \times 2^L} \underbrace{B_{\mathbf{s}}\boldsymbol{\pi}}_{2^L \times 1} - \beta \underbrace{\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})}_{L \times n^L} \underbrace{A_{\mathbf{s}}\mathbf{V}}_{n^L \times 1} \quad (3)$$

Under the assumption of zero probability ties $\Gamma_{nl}(\mathbf{b}|\mathbf{s}) = \prod_{n' \neq n} G_{n'l}(b_{nl}|\mathbf{s})$. Therefore $\nabla\Gamma$ must be a diagonal matrix with entry ll equal to $\sum_{n' \neq n} g_{n'l}(b_{nl}|\mathbf{s}) \prod_{k \neq n', n} G_{kl}(b_{nl}|\mathbf{s})$, and so

$\nabla\Gamma$ must be invertible for most \mathbf{b} .¹⁰

3.3 The Inverse Bidding System and Identification of F

F is identified, conditional on π and βV , by inverting the first order conditions to obtain \mathbf{v} as a function of bids, π , and βV . This inversion comes from GKS and is a simple multi-object extension of Guerre et al. (2000) identification result from inverting the first order conditions. Invert the first order conditions for the inverse bid system:

$$\xi(\mathbf{b}^*|\pi, \beta V; \mathbf{s}) = \underbrace{\mathbf{b}^*}_{\text{observed}} + \underbrace{\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s}) B_{\mathbf{s}}\pi]}_{\text{Identified}} - \underbrace{\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s}) A_{\mathbf{s}}\beta \mathbf{V}}_{\text{Identified}} \quad (4)$$

This system extends the standard inverse bid function. At the optimum the lot specific value is equal to bids \mathbf{b}^* plus a lot specific markup $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\Gamma(\mathbf{b}^*|\mathbf{s})$, minus a combination markup $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\pi$, minus a dynamic markup which depends on precisely who won each combination of lots $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\beta \mathbf{V}$.

We can evaluate this inverse bid function at the observed bids, which holds for a particular candidate $(\pi, \beta V)$. If this candidate $(\pi, \beta V)$ is correct, then $\xi(\mathbf{b}^*|\pi, \beta V; \mathbf{s}) = \mathbf{v}$. From here it is simple to non-parametrically identify F , as in GKS.

3.4 Identification of V

We can write \mathbf{V} as a function of the distribution of bids and π only:

Proposition 3. *Under assumptions 2 - 5, the expected stage pay-off is given by:*

$$\begin{aligned} \tilde{W}(\mathbf{b}^*|\mathbf{v}; \mathbf{s}) = & \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \Gamma(\mathbf{b}^*|\mathbf{s}) \\ & + [P(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})] B_{\mathbf{s}}\pi \\ & + [Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})] A_{\mathbf{s}}\beta \mathbf{V} \end{aligned} \quad (5)$$

Proof of this proposition is given in Appendix A, generalising Proposition 1 in JP. The first term on the right hand side can be written as $\sum_l \frac{\prod_{n' \neq n} G_{n'l}(b_{nl})}{\sum_{n' \neq n} g_{n'l}(b_{nl})}$ — the first term in JP's proposition. Unlike in the single object case there is a correction for the non-additivity.

¹⁰In practice, $\nabla\Gamma$ may not be invertible if, for example, some lots are unavailable. In this case, the corresponding diagonal entry is zero. Instead, it is without loss to use a pseudo-inverse that consists of zeroing out the values of $\nabla\Gamma^{-1}$ that correspond to the zero values of $\nabla\Gamma$.

From Proposition 3, employing the identity $P(\mathbf{b}|\mathbf{s})^T B_s = Q(\mathbf{b}|\mathbf{s})^T A_s C$, and taking an expectation of the observed bids, we can write the ex-ante value function as:

$$V^e(\mathbf{s}) = \Phi(\mathbf{s}) + \Omega(\mathbf{s})[C\boldsymbol{\pi} + \beta\mathbf{V}]$$

$$\begin{aligned} \text{Where } \Phi(\mathbf{s}) &= E_{\mathbf{b}}[\Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \Gamma(\mathbf{b}^*|\mathbf{s}) | \mathbf{s}] \\ \Omega(\mathbf{s}) &= E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s}) | \mathbf{s}] A_s \end{aligned}$$

Stacking over \mathbf{s} write the continuation value as $\mathbf{V} = T\mathbf{V}^e = T\Phi + T\Omega[C\boldsymbol{\pi} + \beta\mathbf{V}]$, which we invert for: $\mathbf{V} = (I_S - \beta T\Omega)^{-1}[T\Phi + T\Omega C\boldsymbol{\pi}]$.¹¹ This ensures that, conditional on $\boldsymbol{\pi}$ being known, the continuation value is point identified.

3.5 Identification of $\boldsymbol{\pi}$

Impose the mean zero property of \mathbf{v} , then substitute in the expression for the Inverse Bid System given in Equation 4, for:

$$\begin{aligned} 0 &= E_{\mathbf{v}}[\mathbf{v}|\mathbf{s}] = E_{\mathbf{b}^*}[\boldsymbol{\xi}(\mathbf{b}^*; \mathbf{s}, (\boldsymbol{\pi}, \mathbf{V})) | \mathbf{s}] \\ &= E_{\mathbf{b}^*}[\mathbf{b}^* + \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \Gamma(\mathbf{b}^*|\mathbf{s}) | \mathbf{s}] - E_{\mathbf{b}^*}[\nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s}) | \mathbf{s}] A_s [C\boldsymbol{\pi} + \beta\mathbf{V}] \\ &= \Upsilon(\mathbf{s}) - \Psi(\mathbf{s})[C\boldsymbol{\pi} + \beta\mathbf{V}] \end{aligned} \tag{6}$$

Stacking over \mathbf{s} , then substituting in the expression for \mathbf{V} and simplifying, we get:

$$\begin{aligned} 0 &= \Upsilon - \Psi[C\boldsymbol{\pi} + \beta\mathbf{V}] \\ &= \Upsilon - \beta\Psi(I_S - \beta T\Omega)^{-1} T\Phi - \Psi(I_S - \beta T\Omega)^{-1} C\boldsymbol{\pi} \end{aligned} \tag{7}$$

This system of LS equations in $S_n - 1$ unknowns overcomes the standard order condition discussed in GKS. There exists a unique solution to this system ($\boldsymbol{\pi}$ is point identified) if and only if the $LS \times S_n$ matrix $\Psi(I_S - \beta T\Omega)^{-1} C$ has rank $S_n - 1$.

3.6 Rank of $\Psi(I_S - \beta T\Omega)^{-1} C$

This rank condition requires that observations of bidding behaviour, across all S states, produces sufficient information about $\boldsymbol{\pi}$ to uniquely pin down all $S_n - 1$ elements. We

¹¹Invertibility of $(I_S - \beta T\Omega)$ is standard and follows from strict diagonal dominance (Pesendorfer and Schmidt-Dengler, 2008).

gain information about $\pi(\mathbf{s}_n)$ from how bidding behaviour changes when \mathbf{s}_n is a possible outcome from the round of auctions. By stacking the moment conditions in equation 7 we stitch together the information about π across different state observations. In addition to information as \mathbf{s}_n varies, we also use information as \mathbf{s}_{-n} varies, even when this is excluded from the function π , resulting in additional identifying variation. This rank condition is essentially the same rank condition assumed for identification in most studies of dynamic games, discussed (though not proven) in Pesendorfer and Schmidt-Dengler (2008).

One additional assumption is sufficient for this rank condition to hold. Define the set valued function $\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$ as the set of possible individual ex-post states \mathbf{s}_n^c having started in state \mathbf{s}_n , given the common state \mathbf{s}_0 :

- Assumption 6.** *i) The set \mathbb{S}_n is partially ordered according to the strict partial ordering \succeq , such that if $\mathbf{s}'_n \in \mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$ then $\mathbf{s}'_n \succeq \mathbf{s}_n$.*
- ii) The maximal elements of \mathbb{S}_n do not outnumber the non-maximal elements.*
- iii) For any non-maximal $\mathbf{s}'_n, \mathbf{s}_n$ and all \mathbf{s}_0 , for any two corresponding elements of the set of possible ex-post states $\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$ and $\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$ denoted $\mathbf{s}_n^{c'}$ and \mathbf{s}_n^c respectively: If $\mathbf{s}'_n \succeq \mathbf{s}_n$ then $\mathbf{s}_n^{c'} \succeq \mathbf{s}_n^c$, and if $\mathbf{s}'_n \not\succeq \mathbf{s}_n$ then $\mathbf{s}_n^{c'} \not\succeq \mathbf{s}_n^c$.*

The partial ordering assumption only imposes the transitivity of partially ordered sets. This requires that winning an auction is monotonic: one cannot gain an object from winning one auction and give it away by winning a different auction. I limit the number of maximal elements because observations of bidding from maximal elements are not informative.¹² Part *iii)* requires that if \mathbf{s}'_n is higher in the partial ordering than \mathbf{s}_n , then each outcome in the set of possible ex-post states $\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$ is higher than the corresponding element in $\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$. For example, if $\mathbf{s}'_n \succeq \mathbf{s}_n$ then the element of $\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$ that corresponds to winning every available lot must also be higher than the element of $\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$ that corresponds to winning every available lot. This only requires that if a bidder begins a period with a larger state, winning the same set of lots means they also end the period with a larger state.

Proposition 4. *Under assumption 2 - 6 $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_n - 1$*

Proof of this proposition is given in Appendix B. The rank condition is not trivial, since Ψ is certainly rank deficient. Likewise, it is not ex-ante obvious whether stacking $\Psi(\mathbf{s})$ across

¹²An element \mathbf{s}_n is defined as maximal if there does not exist an $\mathbf{s}'_n \in \mathbb{S}_n$ such that $\mathbf{s}'_n \succ \mathbf{s}_n$. One interpretation is that these maximal elements are the largest (under \succeq) states that are observed as possible ex-post outcomes, but never as ex-ante outcomes. In this way, we want to try to identify π for these states, but do not get to use observations beginning in these states.

initial states provides information about $\pi(\mathbf{s}_n^c)$ for every possible ex-post state \mathbf{s}_n^c . The bulk of the proof establishes the rank of Ψ and finds its null space. As we stitch together observations of bidding from each state, stacking $\Psi(\mathbf{s})$ across \mathbf{s} , the rank increases by at least two each time. I then consider the image of $(I_S - \beta T\Omega)^{-1}C$, proving that the only element in the intersection of this image and Ψ 's null space is the constant vector.¹³

3.7 Extensions

Second-price auctions: In Appendix D.1 I show how this framework extends almost trivially to simultaneous second-price auctions.

Binding reservation prices: In Appendix D.2 I consider how the presence of binding reservation prices impact identification. Essentially, they cause censoring in the data so that we immediately lose point identification of both F and π . However, F remains partially identified, using a similar argument as presented in subsection 3.3. We can no longer use moment conditions to identify π , as in subsection 3.5, and instead use quantile conditions.

Endogenous Entry: In Appendix D.3 I consider an additional stage in-which the bidder chooses a subset of auctions to enter, where entering each subset has an associated cost. This creates a minor change to the representation of V as a function of π . The identification of π and F follows from previous arguments. Identification of the entry cost distribution then follows from standard results.

Stochastic Combination Value: In Appendix D.4 I allow the combination value to be a function of low dimensional ($< L$) random variables, such as unobserved states. The necessary restriction is that this function is strictly monotonic in the unobservables. Identification arises from proving that bids can be inverted to point identify the unobservables.

4 Estimation Procedure

Having established non-parametric identification, I now describe a computationally feasible procedure to estimate F and π . Because we cannot write maximised expected payoffs as a function of bids only (Proposition 3), JP's estimation method for dynamic auction models is inapplicable. I begin with a general description, outlining the key intuition. I then detail the three estimation steps and discuss asymptotics.

¹³This proof only holds for the setting when the state space is finite. However the underlying argument extends to the case with infinite states: Even though the rank of an infinitely large matrix is undefined, it is clear how the logic of combining observations across states yields identification.

4.1 The Premise

The central premise of the procedure exploits that, under the assumption that payoffs are additively separable over time, we can write the continuation value as a function of: (1) Primitives of the transition process, (2) the observed distribution of equilibrium actions, and (3) the sum of the flow pay-off function and the discounted continuation value. I refer to this sum as the pseudo-payoff; it is object we *mistake* for the payoff function if we incorrectly assume myopic bidding. This relationship is given by:

$$\begin{aligned}
 V(\mathbf{s}') &= \int_{\mathbf{s}} \int_{\mathbf{b}} \Phi(\mathbf{b}|\mathbf{s}) + \Omega(\mathbf{b}|\mathbf{s}) \bar{\kappa}(\mathbf{s}) dG(\mathbf{b}|\mathbf{s}) T_{\mathbf{s}}(\mathbf{s}|\mathbf{s}') d\mathbf{s} \\
 \text{Where: } [\bar{\kappa}(\mathbf{s})]_c &= \kappa(\mathbf{s}^c) = \pi(\mathbf{s}_n^c) + \beta V(\mathbf{s}^c) \\
 \text{and } \Phi(\mathbf{b}|\mathbf{s}) &= \Gamma(\mathbf{b}|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}|\mathbf{s})^{-1} \Gamma(\mathbf{b}|\mathbf{s}) \\
 \Omega(\mathbf{b}|\mathbf{s}) &= Q(\mathbf{b}|\mathbf{s})^T - \Gamma(\mathbf{b}|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}|\mathbf{s}) \quad (8)
 \end{aligned}$$

This equation restates Proposition 3 as a function of G and T , as well as the pseudo-payoff function κ . Both G and T can be estimated using standard methods. Therefore, if we had a consistent estimate for the function $\kappa : \mathbb{S} \rightarrow \mathbb{R}$, then we would have a consistent estimate for V , and then $\pi (= \kappa - \beta V)$. Like the distribution of equilibrium bids, κ is not a model primitive but an equilibrium object. The central estimation problem then concerns estimating this pseudo-payoff function κ .

Estimating κ is very similar to estimating a misspecified static model. If players are myopic ($\beta = 0$) then $\kappa = \pi$. However, κ generally depends on \mathbf{s}_{-n} , while π does not. The procedure involves estimating a *generalised* static model, allowing payoffs to depend on elements of the state that enter the continuation value.¹⁴

The procedure is a generalisation of the estimation procedure proposed by JP. When payoffs are additively separable $V(\mathbf{s}') = \int_{\mathbf{s}} \int_{\mathbf{b}} \Phi(\mathbf{b}|\mathbf{s}) dG(\mathbf{b}|\mathbf{s}) T_{\mathbf{s}}(\mathbf{s}|\mathbf{s}') d\mathbf{s}$ and the procedure collapses down to JP. We write the Value Function as a function of the distribution bids and this additional term $\Omega(\mathbf{b}|\mathbf{s})\bar{\kappa}(\mathbf{s})$, correcting for the non-additivity across lots. Unlike JP we require an extra estimation step to estimate this correction term. In a single-agent setting this procedure is equivalent to the Conditional Choice Probability estimator of Hotz and Miller (1993), and in other settings when the policy function is invertible it is equivalent to the procedure of Bajari et al. (2007). This equivalence arises because the pseudo-payoff function

¹⁴This permits a test of forward looking behaviour. If the model is correctly specified and $(\mathbf{s}_{-n}, \mathbf{s}_0)$ is excluded from π , then observing that κ varies with $(\mathbf{s}_{-n}, \mathbf{s}_0)$ is sufficient to reject myopia.

essentially is the inverted policy function. The procedure only differs when policy functions, or conditional choice probabilities, are not invertible. In the single-agent environment this occurs in repeated ‘choice over lotteries’ settings; when multiple (non-ordered) types may choose the same action. This also occurs in repeated games when the stage game has non-monotone equilibria, which is the case for multi-object auctions without combination bidding.¹⁵

4.2 The Procedure

The procedure can be written succinctly as:

Definition 4.1. Algorithm 1.

1. *Estimate equilibrium bid distributions G (beliefs) and the transition process T_s .*
2. *Given \hat{G} , estimate κ using the identifying conditions $E[\xi(\mathbf{b}_{it}; \mathbf{s}_t, \kappa, \hat{G}) | \mathbf{s}_t] = 0$ for each state observed in the data. Then, evaluate \hat{F} using \hat{G} and a change of variables.*
3. *Given $\hat{G}, \hat{T}, \hat{\kappa}$, evaluate \hat{V} using Equation 8. Finally, evaluate $\hat{\pi} = \hat{\kappa} - \beta \hat{V}$.*

I make the following assumption about the true underlying structure, enabling me to discuss the statistical properties of this estimator:

Assumption 7. *i) Beliefs G , the transition process T_s and the pseudo-payoff function κ are parameterised by finite parameter vectors θ^G, θ^τ , and θ^κ respectively.*

ii) $G(\mathbf{b} | \mathbf{s}; \theta^G)$, $T_s(\mathbf{s} | \mathbf{s}'; \theta^\tau)$, and $\kappa(\mathbf{s}; \theta^\kappa)$ are continuously differentiable in θ^G, θ^τ , and θ^κ respectively. Also, the spaces of parameters Θ^G, Θ^τ , and Θ^κ are compact.

This assumption ensures the consistency, asymptotic normality and \sqrt{T} convergence of the estimator.¹⁶ With a discrete state space only parameterisation of G is needed, as both $\kappa(\mathbf{s})$ and T_s can be estimated state-by-state. However, in many settings (including the application in this paper) researchers may choose to treat a particularly large state space as continuous. While this assumption rules out fully non-parametric methods, such as kernels or sieves,

¹⁵I leave discussion of the broader applicability of this estimator for future work. The identification requirements for this procedure are slightly stronger than in Bajari et al. (2007), who only require that the policy function is identified. Here, we also require that the pseudo-payoff function is identified, which is not trivial. In the simultaneous first-price setting, this is guaranteed by the rank condition on Ψ discussed in section 3.

¹⁶A similar assumption is required by Bajari et al. (2007), and is used in practice in most studies.

it permits flexible parametric and semi-nonparametric methods such as polynomials and sieve-type B-spline estimators with pre-specified knot vectors.¹⁷ Part *ii*) of the assumption ensures the standard regularity conditions hold for asymptotics of Generalised Method of Moments (GMM) estimators. The standard identification, invertibility, and finite moment assumptions are implied by the assumptions and arguments presented in Section 3. To apply this estimator in other settings requires an additional identification assumption. Proposition 5 summarises the properties of this estimator:

Proposition 5. *Under assumptions 2 - 7 \hat{F} and $\hat{\pi}$ are \sqrt{T} consistent and asymptotically normal.*

I now detail each of the three estimation steps, before demonstrating that their asymptotic properties follow from Mises (1947) and Newey and McFadden (1994).¹⁸

4.2.1 Step 1.

The First Step constitutes the standard first step in the empirical auction literature. There are several possible approaches the researcher might take. As in GKS and JP One might estimate the conditional joint distribution of bids G_i , then form $\Gamma(\mathbf{b})$, $P(\mathbf{b})$, and $Q(\mathbf{b})$ respectively. Otherwise the researcher may directly estimate these objects, essentially estimating the joint distribution of maximum rival bids as in Cantillon and Pesendorfer (2007).

Given Assumption 7 we cannot take a fully non-parametric approach as this complicates asymptotics. Instead, suppose we estimate θ^G using the estimating equation $E[\mathbf{m}_1(\mathbf{b}_t, \mathbf{s}; \theta^G)] = 0$. $\mathbf{m}_1(\mathbf{b}_t, \mathbf{s}; \theta^G)$ might be the score vector in a fully parametric specification, or $\mathbf{m}_1(\mathbf{b}_t, \mathbf{s}; \theta^G) = G(\mathbf{b}|\mathbf{s}; \theta^G) - \prod_l \mathbb{I}[b_{lt} \leq b_l]$ for all $\mathbf{b} \in \mathbb{B}$ for a moment based approach. Asymptotic properties of this GMM estimator are discussed shortly.

The parameters of the transition process θ^T must be estimated similarly. The central requirement, given Assumption 7, is that the estimator is chosen to be \sqrt{T} consistent and

¹⁷B-splines are an attractive alternative to fully non-parametric methods, allowing researchers to estimate flexible models while setting the knot vector to ensure sufficient data in each cell. The Stone Weierstrass Theorem ensures B-splines approximate any continuous function to arbitrary precision given sufficient knots. These methods are increasingly being used in the empirical auction literature (see Hickman et al. (2017) and Bodoh-Creed et al. (2021) as examples).

¹⁸Both $\hat{\pi}$ and \hat{F} are centred by their pseudo-true value, under the finite dimensional parameterisation of Assumption 7. That is, the model is likely misspecified and so there remains the possibility of asymptotic bias. This bias can always be diminished as the sample size increases by increasing the flexibility of the functional form, such as using splines with a finer grid of knots.

asymptotically normal, with analytically tractable asymptotic variance. This includes standard estimators such as maximum likelihood and GMM, really only ruling out certain non-parametric estimators. As I will discuss shortly, the choice of parameterisation depends on the parameterisation of $\kappa(\mathbf{s}; \boldsymbol{\theta}^\kappa)$, as certain flexible functional form assumptions can make estimation extremely convenient.

4.2.2 Step 2.

In the second step we estimate the pseudo-payoff function $\kappa(\mathbf{s}; \boldsymbol{\theta}^\kappa)$, estimating the (potentially large) parameter vector $\boldsymbol{\theta}^\kappa$. This broadly follows the second stage in the estimation procedure presented in GKS, estimating the model as if it were static using the identifying conditions from Section 3: $E[\mathbf{v}|\mathbf{s}] = 0$.

In practice we employ GMM using the moment condition $E[\mathbf{m}_2(\mathbf{b}, \mathbf{s}; \boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G)] = 0$, where $\mathbf{m}_2(\mathbf{b}, \mathbf{s}; \boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G) = H(\mathbf{s})\boldsymbol{\xi}(\mathbf{b}, \mathbf{s}; \boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G)$ with $H(\mathbf{s})$ as an $h \times L$ dimensional matrix of instruments that are some known function of \mathbf{s} , and so mean independent of \mathbf{v} . $\hat{\boldsymbol{\theta}}^\kappa$ minimises the standard quadratic loss function:

$$\hat{\boldsymbol{\theta}}^\kappa = \arg \min_{\boldsymbol{\theta}^\kappa} \left\{ \left(\frac{1}{T} \sum_t \mathbf{m}_2(\mathbf{b}_t, \mathbf{s}_t; \boldsymbol{\theta}^\kappa, \hat{\boldsymbol{\theta}}^G) \right)^T \hat{W} \left(\frac{1}{T} \sum_t \mathbf{m}_2(\mathbf{b}_t, \mathbf{s}_t; \boldsymbol{\theta}^\kappa, \hat{\boldsymbol{\theta}}^G) \right) \right\}$$

Where $\hat{W}^{-1} = \frac{1}{T} \sum_t H(\mathbf{s}_t)\boldsymbol{\xi}(\mathbf{b}_t, \mathbf{s}_t; \boldsymbol{\theta}^\kappa, \hat{\boldsymbol{\theta}}^G)\boldsymbol{\xi}(\mathbf{b}_t, \mathbf{s}_t; \boldsymbol{\theta}^\kappa, \hat{\boldsymbol{\theta}}^G)^T H(\mathbf{s}_t)^T$ is the multi-step asymptotically efficient weight matrix, allowing for within period correlation. Importantly, the estimate $\hat{\boldsymbol{\theta}}^\kappa$ depends on $\hat{\boldsymbol{\theta}}^G$, and so inference must take into account this multi-step estimation procedure, which I discuss in detail in Section 4.3.

In practice it is particularly convenient if the researcher fits a flexible linear in parameters parametric form to $\kappa(\mathbf{s}; \boldsymbol{\theta}^\kappa) = \mathbf{h}(\mathbf{s})^T \boldsymbol{\theta}^\kappa$, where $\mathbf{h}(\mathbf{s})$ is, for example, a vector of B-splines. We then use this vector for our instruments, so that $H_l(\mathbf{s}) = \mathbf{h}(\mathbf{s})$ for each l . This form is convenient first because of how it can simplify the third estimation step, which I discuss shortly, and second because of how it allows us to interpret this GMM step as a linear instrumental variable problem, as I now discuss.

Rewrite the Inverse Bid System as a regression equation:

$$\underbrace{b_{lt} + \frac{\Gamma_l(b_{lt}|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_{lt}|\mathbf{s}_t)}}_{y_t} = - \left[\frac{\nabla_{\mathbf{b}} Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l}\Gamma_l(b_{lt}|\mathbf{s}_t)} \right]_l \underbrace{\bar{\kappa}(\mathbf{s}_t; \boldsymbol{\theta}^\kappa)}_{\bar{H}(\mathbf{s}_t)\boldsymbol{\theta}^\kappa} + v_{lt}$$

Where row c of the known $N^L \times |\boldsymbol{\theta}^\kappa|$ matrix $\bar{H}(\mathbf{s}_t)$ is $\mathbf{h}(\mathbf{s}^c)^T$. Now, we *could* estimate $\boldsymbol{\theta}^\kappa$ using least squares; minimising the sum of squared residuals $\sum_t \mathbf{v}_t^T \mathbf{v}_t$. In general $E[v_{lt}[\frac{\nabla_{\mathbf{b}} Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l} \Gamma_l(b_{lt}|\mathbf{s}_t)}]_l] \neq 0$ because $E[v_{lt} b_{l't}] \neq 0$, an endogeneity problem. Instead, we use our instruments $\mathbf{h}(\mathbf{s})$, which are mean independent of v_l . The first stage is then:

$$-\left[\frac{\nabla_{\mathbf{b}} Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l} \Gamma_l(b_{lt}|\mathbf{s}_t)}\right]_l \bar{H}(\mathbf{s}_t) = \boldsymbol{\delta}_l \mathbf{h}(\mathbf{s}_t) + \varepsilon_{lt}$$

Existence of this first stage follows from the previous identification results. However, the instruments may be weak if $\mathbf{h}(\mathbf{s}_t)$ does not ‘cause’ sufficient variation in $\left[\frac{\nabla_{\mathbf{b}} Q(\mathbf{b}_t|\mathbf{s}_t)}{\nabla_{b_l} \Gamma_l(b_{lt}|\mathbf{s}_t)}\right]_l \bar{H}(\mathbf{s}_t)$. This occurs when the observed variation in initial states \mathbf{s}_t is less than the variation in the possible ex-post states \mathbf{s}_t^c . It is then pertinent to consider additional instruments. Fortunately, many standard packages are available for analysing the relevance and validity of our instruments in this linear instrumental variable setting.

Next, back out the distribution of lot specific values F using a change of variables:

$$\hat{F}(\mathbf{v}|\mathbf{s}) = G(\mathbf{b}^*(\mathbf{v}|\mathbf{s}; \hat{\boldsymbol{\theta}}^G, \hat{\boldsymbol{\theta}}^\kappa)|\mathbf{s}; \hat{\boldsymbol{\theta}}^G)$$

Where $\mathbf{b}^*(\cdot|\mathbf{s}; \hat{\boldsymbol{\theta}}^G, \hat{\boldsymbol{\theta}}^\kappa)$ gives the estimated bid function, which can be evaluated for a given \mathbf{v} using numerical methods. The estimated bidding function depends directly on the estimates for beliefs and the pseudo-static payoff function, as well as indirectly on \mathbf{s} . Meanwhile, to draw \mathbf{v} s from the estimated distribution, in order to perform counterfactual simulations, one convenient method is to draw $\xi_l(\mathbf{b}_t|\mathbf{s}_t; \hat{\boldsymbol{\theta}}^\kappa, \hat{\boldsymbol{\theta}}^G)$ from their empirical distribution. This has the benefit that one does not have to explicitly evaluate $\hat{F}(\mathbf{v}|\mathbf{s})$.

4.2.3 Step 3.

Given the estimated distribution of bids, transition process, and pseudo-payoffs, evaluate the ex-ante value function, then the continuation value, using equation 8. That is, given a period ends in state \mathbf{s} , estimate expected payoffs in the following period. Assumption 7 ensures the ex-ante value function can also be written as $V^e(\mathbf{s}; \boldsymbol{\theta}^V)$, where $\boldsymbol{\theta}^V$ is a finite parameter vector and the function V^e is known up to $\boldsymbol{\theta}^V$.¹⁹ We could use numerical integration to evaluate the ex-ante value function, but it is often convenient to take a finite sample approximation over observed bids and states using non-linear least squares:

¹⁹At worst this known function is given by equation 8, with $\boldsymbol{\theta}^V = (\boldsymbol{\theta}^G, \boldsymbol{\theta}^\kappa)$. Often, as in Section 5, it will be convenient to fit a flexible parametric form to $V^e(\mathbf{s}; \boldsymbol{\theta}^V)$.

$$\hat{\boldsymbol{\theta}}^V = \arg \min_{\boldsymbol{\theta}^V} \left\{ \frac{1}{T} \sum_t [V^e(\mathbf{s}_t; \boldsymbol{\theta}^V) - \bar{W}(\mathbf{b}_t, \mathbf{s}_t; \hat{\boldsymbol{\theta}}^G, \hat{\boldsymbol{\theta}}^\kappa)]^2 \right\}$$

$\bar{W}(\mathbf{b}, \mathbf{s}; \boldsymbol{\theta}^G, \boldsymbol{\theta}^\kappa)$ is the parameterised object defined in Equation 8. This is equivalent to using a third GMM step, employing the moment condition $E[\mathbf{m}_3(\mathbf{b}, \mathbf{s}; \boldsymbol{\theta}^V, \hat{\boldsymbol{\theta}}^\kappa, \hat{\boldsymbol{\theta}}^G)] = 0$, where $\mathbf{m}_3(\mathbf{b}_t, \mathbf{s}_t; \boldsymbol{\theta}^V, \hat{\boldsymbol{\theta}}^\kappa, \hat{\boldsymbol{\theta}}^G) = \nabla_{\boldsymbol{\theta}^V} V^e(\mathbf{s}_t; \boldsymbol{\theta}^V) [V^e(\mathbf{s}_t; \boldsymbol{\theta}^V) - \Pi(\mathbf{b}_t, \mathbf{s}_t; \hat{\boldsymbol{\theta}}^G, \hat{\boldsymbol{\theta}}^\kappa)]$.²⁰

Finally, we back out our estimate $\hat{\pi}(\mathbf{s}_n)$ for each \mathbf{s}_n using:²¹

$$\hat{\pi}(\mathbf{s}_n) = \kappa(\mathbf{s}; \hat{\boldsymbol{\theta}}^\kappa) - \beta \int V^e(\mathbf{s}'; \hat{\boldsymbol{\theta}}^V) T_{\mathbf{s}}(\mathbf{s}' | \mathbf{s}; \hat{\boldsymbol{\theta}}^\tau) d\mathbf{s}'$$

When states are discrete the integral can be evaluated analytically, otherwise (depending on the specification of $T_{\mathbf{s}}$) the analytic expectation may still be feasible. If not, numerical methods can be used. If the researcher fits a linear in parameters form to κ , such as $\kappa(\mathbf{s}; \boldsymbol{\theta}^\kappa) = \mathbf{h}(\mathbf{s})^T \boldsymbol{\theta}^\kappa$, then it is convenient to assume the same functional form for both V^e and $T_{\mathbf{s}}$: $V^e(\mathbf{s}; \boldsymbol{\theta}^V) = \mathbf{h}(\mathbf{s})^T \boldsymbol{\theta}^V$ and $E[\mathbf{h}(\mathbf{s}_{t+1})^T | \mathbf{s}_t] = \mathbf{h}(\mathbf{s}_t)^T \boldsymbol{\theta}^\tau$. π then inherits this linear in parameters form: $\pi(\mathbf{s}) = \mathbf{h}(\mathbf{s})^T (\boldsymbol{\theta}^\kappa - \beta \boldsymbol{\theta}^\tau \boldsymbol{\theta}^V)$, simplifying both estimation and inference.²²

4.3 Large Sample Properties

I now discuss the large sample properties of the estimator, proving proposition 5. I assume a \sqrt{T} consistent and asymptotically normal estimator is used for $\hat{\boldsymbol{\theta}}^\tau$. Next, $\hat{\boldsymbol{\theta}}^G$, $\hat{\boldsymbol{\theta}}^\kappa$, and $\hat{\boldsymbol{\theta}}^V$ result from a three step GMM procedure, and so Assumption 7 ensures we can apply Theorem 6.1 from Newey and McFadden (1994). Therefore $\hat{\boldsymbol{\theta}}^G$, $\hat{\boldsymbol{\theta}}^\kappa$, and $\hat{\boldsymbol{\theta}}^V$ are \sqrt{T} consistent and asymptotically jointly normal.

Finally, both F and π are known functions of these estimated parameters, and so their asymptotic properties follow from the delta method (Mises, 1947).²³ The asymptotic vari-

²⁰There are also large efficiency gains from weighting observations according to the estimated variance of $\hat{\boldsymbol{\theta}}^\kappa$ and $\hat{\boldsymbol{\theta}}^G$. Weight observations t using the inverse of the estimated variance of $\bar{W}(\mathbf{b}, \mathbf{s}; \boldsymbol{\theta}^G, \boldsymbol{\theta}^\kappa)$. Weighting can also be employed in the second estimation step, however in practice $\boldsymbol{\theta}^G$ will be more precisely estimated than $\boldsymbol{\theta}^\kappa$.

²¹We average $\hat{\pi}$ s over $(\mathbf{s}_{-n}, \mathbf{s}_0)$. With a correctly specified model and infinite data there will be no variation. In the spirit of Magnac and Thesmar (2002) β is identified from our exclusion restrictions on π . We *could* set β such that $\hat{\pi}$ is independent of \mathbf{s}_{-n} . This is left for future work.

²²Other approaches are also possible, such as plugging the estimated continuation value into the inverse bid system and performing a final GMM step, treating $\hat{\pi}$ as the only unknown. This is similar to the classic quasi-maximum likelihood approach to CCP estimation.

²³This requires both functions are continuously differentiable in the estimated parameters. For $\hat{\pi}$ this

ances of $\hat{\kappa}$, \hat{F} , and $\hat{\pi}$ are then as standard, and so finite sample approximations for these variances can be used for inference and hypothesis testing.

In Appendix E I present the results of a simulation study examining the performance of both semi-parametric and semi-nonparametric estimators. There are two key findings from this exercise: First, the choice of instruments in the second stage is extremely important. The initial state instruments are often weak and the resulting estimator converges slowly, particularly when the model is incorrectly specified. Surprisingly, while using additional relevant instruments is best, using no instruments performs well, exhibiting very little bias and with low variance in small samples. Second, the semi-nonparametric B-spline estimator performs very well. However computation is slow, particularly when calculating multi-step variances. Meanwhile, the misspecified semi-parametric estimators, such as simple polynomials, are subject to some misspecification bias but still provide a viable alternative; particularly in small samples when this bias is dwarfed by sampling uncertainty.

5 Application

I now apply this model and estimation procedure to data from Michigan Department of Transport’s procurement auctions for highway construction and maintenance contracts. This setting and data has been considered in several previous studies, including Groeger (2014), Somaini (2020), Raisingh (2021), and GKS. Contracts are allocated using simultaneous low-price sealed bid auctions, averaging around 35 contracts auctioned in each round, with rounds taking place every 2-4 weeks. 73 percent of regular bidders submit bids on more than one auction in a given round. I focus on road construction and paving projects. These projects either involve hot-mix asphalt, concrete construction, or both.

The MDOT data exhibits evidence of both dynamic linkages and static complementarities, which I demonstrate in Section 5.2. Similarly, both Raisingh (2021) and Groeger (2014) find evidence of dynamic linkages in MDOT auctions. They demonstrate evidence of backlog effects as firms bid less aggressively the larger their backlogs in the MDOT auctions. Meanwhile, GKS present evidence of complementarities in contracts auctioned simultaneously by MDOT. They demonstrate that firms’ costs of taking on new projects increase

is trivial, while for \hat{F} this is not. Assumptions 4 and 7 ensure G has a continuous first derivative with respect to both \mathbf{b} and $\boldsymbol{\theta}^G$. We must also assume that G has a continuous second derivative. This, in conjunction with the inverse function theorem, ensures $\mathbf{b}^*(\mathbf{v}|\mathbf{s}; \boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G)$ has a continuous first derivative with respect to $\boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G$. This is because $\boldsymbol{\xi}(\mathbf{b}|\mathbf{s}; \boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G)$ has a continuous first derivative in $\boldsymbol{\theta}^\kappa, \boldsymbol{\theta}^G$, and is (at least) locally invertible in \mathbf{b} .

in their backlogs, but the more similar their current projects the less the dis-economies of scale. This suggests the need to use a dynamic multi-object auction model to estimate costs. Understanding the degree of complementarities across lots auctioned either simultaneously or across time is important for auction design — if the cost synergies are large firms may suffer from the exposure effect of being unable to express these complementarities, and may benefit from sequential allocation of contracts.

5.1 Data

I use the same data as GKS, using their data on bids, contracts, and competing firms. This includes information on almost every auction run between 2002 and 2014. The contract data includes project descriptions, locations, the engineer’s estimate of project cost, and the list of participating firms and their bids.

The firm level data includes details on the sub-sample of firms who submit at least 50 bids. This details the number and location of plants, and a description of the type of company. Following GKS’s classification a regular bidder is one that submits more than 100 bids in the sample period, otherwise they are designated a fringe bidder. The final sample contains 44 regular bidders, and 686 fringe bidders. I further categorise regular bidders into one of three types of firm: General contractors, Paving companies, and Construction companies.

Contract level descriptives are summarised in Figure 1. Around 20 asphalt projects are auctioned simultaneously each period, predominantly highway maintenance projects. But, these tend to be smaller projects, in both duration and predicted costs, than the concrete and mixed projects. These contracts involve construction or bridge maintenance projects, and so the engineers estimates exhibits a major right skew.

Bidder level descriptive statistics are summarised in Figure 2. A firm’s backlog at t is calculated as the sum, over current contracts, of the engineer’s estimate (EE) of the cost for each project multiplied by the fraction of project duration remaining. Backlogs are calculated separately for each type of project. Unsurprisingly, regular bidders’ backlogs are much larger than fringe bidders’. Asphalt backlogs are also generally higher than concrete backlogs due to the larger number of asphalt projects. Paving firms are closer to projects than fringe bidders because they have more plants.

Figure 1: Auction level summary statistics.

	Asphalt	Concrete	Both
Number	3563	712	1974
Auctions per Round	20.13	4.02	11.15
(q25 - q75)	(5 - 30)	(1 - 6)	(3 - 17)
Project Duration (days)	134.11	216.52	200.08
	(46 - 151)	(79.75 - 261.25)	(70 - 235.25)
Engineer's Estimate (\$100,000s)	12.61	22.4	19.88
	(2.92 - 11.16)	(3.65 - 12.16)	(4.29 - 17.29)
Bidders per Auction	4.39	5.46	5.94
	(2 - 5)	(4 - 7)	(3 - 8)
Average Bid (\$100,000s)	12.75	19.93	18.28
	(3.02 - 11.46)	(3.78 - 11.85)	(4.56 - 16.96)
Winning Bid (\$100,000s)	11.98	21.19	18.69
	(2.69 - 10.46)	(3.34 - 11.46)	(3.99 - 16.27)

Note: Aside from the number of auctions, the numbers presented are means. For mixed projects the mean winning bid is higher than the mean bid. This is caused by the skewed project sizes.

5.2 Suggestive Evidence

I now present suggestive evidence of both dynamic linkages and static complementarities, evidence that we need both a dynamic and multi-object auction model. Dynamic linkages arise through backlog effects, creating non-additivities in payoffs across time. Static complementarities arise through scale effects, as there appear to be cost synergies from working on multiple similar contracts simultaneously, creating non-additivities in payoffs across contracts auctioned within a period.

To examine the dynamic linkages I consider how firm's backlogs of both asphalt and concrete projects impact their costs, and so their bidding behaviour. I regress a firm's bid on project l in period t against their backlog of both types of contracts, as well as the interaction of the two backlogs. Results are presented in Figure 3. We see that firms bid more aggressively the larger their backlog of asphalt contracts, suggesting returns to specialisation: Every one standard deviation increase in their asphalt backlog leads to a 0.025 standard deviation reduction in their normalised bids. Meanwhile, firms bid less aggressively the larger their concrete backlogs, suggesting decreasing returns to scale.

To examine the static complementarities, I consider how the set of other lots bid on in period t impacts their costs, and so their bidding behaviour. I examine how bidding

Figure 2: Bidder level summary statistics.

	General	Paving	Construction	Fringe
Plants	1.73	6.71	1.5	1.43
Bids per Round	2.07	2.8	1.8	0.24
(q25 - q75)	(0 - 3)	(0 - 4)	(0 - 3)	(0 - 0)
Backlog: Asphalt	5.57	5.61	2.97	0.24
(millions)	(0.25 - 3.88)	(0.96 - 7.6)	(0.48 - 4.39)	(0 - 0.2)
Backlog: Concrete	3.41	2.18	2.79	0.2
(millions)	(0.18 - 3.41)	(0.11 - 3.83)	(0.23 - 1.35)	(0 - 0.09)
Distance to project (mi.)	105.65	84.18	121.42	119.27
Distance given Bid	71.21	47.03	87.18	69.33
Distance given Won	65.53	45.01	82.51	58.63

Note: Project locations are coded to the centroid of the county they are based in. Distance is calculated as the minimum distance (across plant locations) between a firm and the project location. I exclude the first two years of the data to construct backlogs.

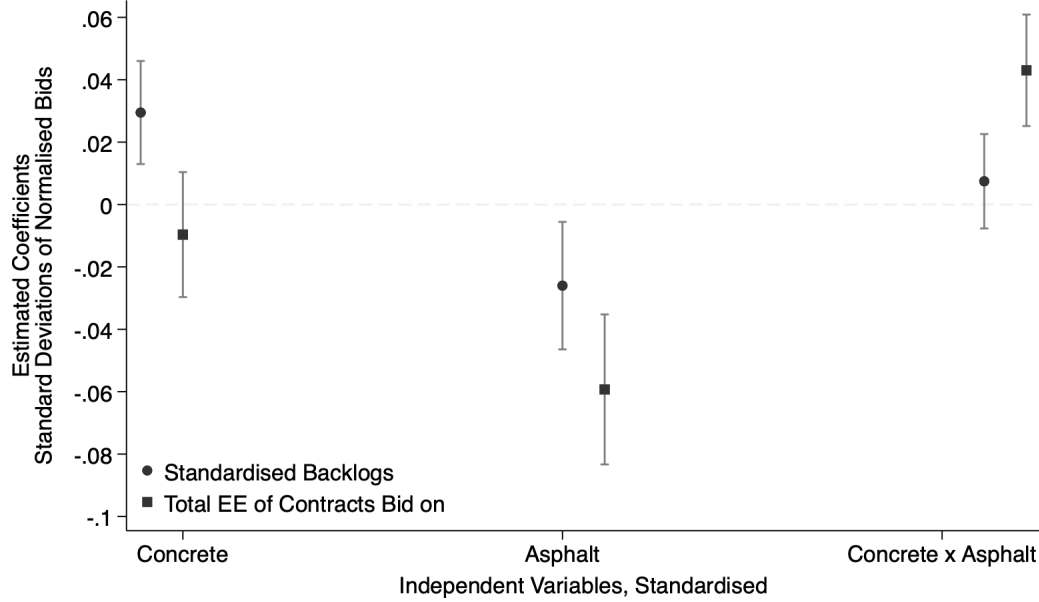
behaviour varies with the sum of the engineers estimate of other contracts they bid on in period t , calculated separately for both asphalt and concrete projects. I regress a firm's bid on project l in period t against these two sets of sums of engineers estimates, as well as their interaction. Results are presented in Figure 3. I find that firms bid more aggressively if they also bid on many different asphalt projects, but not the number of concrete projects. However, the positive coefficient on the interaction term suggests that holding constant the number of asphalt projects they bid on, the more concrete projects they bid on, the less aggressively they bid. Once more, this suggests evidence of returns to specialisation, in line with results presented in GKS.

5.3 Estimation and Results

I now apply and estimate the empirical model presented above. While the semi-nonparametric approach is possible, I follow the literature and take a parametric approach.²⁴ I apply the full dynamic multi-object model to regular bidders only, given that I need to observe sufficient observations of bidding to be able to estimate my objects of interest. I estimate separate parameters for each type of regular bidder. I assume fringe bidders are myopic, and that

²⁴The semi-nonparametric approach suffers from a curse-of-dimensionality. Unfortunately, I need to allow the pseudo-payoff κ to depend on common/rival states and auction level observables. Instead I parameterise the model to ensure parameters are interpretable and enable simple tests of additive separability and myopic bidding.

Figure 3: Evidence of Complementarities and Dynamics



Note: The plot shows the estimated coefficients from two linear regressions of bids against either bidder's backlogs (blue) or the estimated cost of other contracts bid upon (red). Bids are normalised by the EE and standardised. Independent variables are also standardised. I control for the distance between the firm and the project, and Firm \times project type Fixed Effects, with robust standard errors.

their costs are additive.

In the low-bid auction the lowest bidder receives their bid and pays their private cost, which involves a minor relabelling of the model. The individual state is the Firm's backlog of asphalt and concrete contracts. The common state consists of the set of lots on offer, including locations and other contract characteristics, such as size, duration, and type.

For simplicity I do not explicitly model firms' participation decisions. Incorrectly assuming entry is exogenous will bias estimates of the flow payoff function $\pi(\mathbf{s})$. However, as I show in Appendix D.3.3 this bias is just the expected participation costs, in this case bid preparation costs, in the following period. Raisingh (2021) found that bid preparation costs vary from \$5,000 to \$10,000, around 1% of the contract cost, suggesting this bias will be small in relative terms. This assumption may be more problematic for my counterfactuals; under counterfactual mechanisms bidders likely win different sets of contracts and may then choose to participate in different auctions accordingly. However, in Appendix F.1 I demonstrate that participation behaviour is predominantly determined by lot specific fac-

tors between firm and contract, most notably the distance between a firm and the project. None of these factors vary in my counterfactual exercises, and so we do not expect variation in participation decisions.²⁵

5.3.1 The State Space Approximation

The state \mathbf{s} should include every firms' backlogs and information on every auction held each period, which is computationally intractable. It is unlikely that firms would track such a large state space. I follow the approach taken by Raisingh (2021) and Aradillas-Lopez et al. (2022). They condense $(\mathbf{s}_{-n}, \mathbf{s}_0)$ into a one dimensional index λ_{nt} , approximating the degree of competition a firm faces on a given day. For each firm I only need to track three states — two backlogs and this competition index.

I construct λ_{nt} using a random forest to predict the minimum rival bid using $(\mathbf{s}_{-n}, \mathbf{s}_0)$. λ_{nt} is then a function of: *i*) the mean backlog of rival bidders, *ii*) the number of rival bidders, *iii*) the number of auctions held that period.²⁶ Full details of how the index is constructed, and additional results, are given in Appendix F.2.

5.3.2 First Stage

To simplify estimation I assume firms believe that, conditional on auction characteristics and firms' states, the probability they win one auction is independent of whether they win another auction. This ensures the joint probabilities P can be written as products of the marginal probabilities.²⁷ Following Athey et al. (2011) I specify the distribution of minimum rival bids as a three parameter Weibull distribution, with a support parameter as $\frac{1}{3}$ of the

²⁵Another worry concerns the independent private value assumption. Somaini (2020) finds evidence of interdependent costs in the MDOT auctions, essentially suggesting a common value component in bidder's payoffs. However their results also suggest that the private value component is the main component of payoffs. For this reason, I remain in the IPV framework.

²⁶The index assumption implies that a firm's continuation value does not depend on which combination of lots each rival bidder wins. Therefore the firm only has to consider 2^L outcomes from the round of auctions (which combination they win themselves), instead of all N^L possible outcomes. This is reasonable — it is unlikely bidders consider how their bids impact the likelihood of their rivals winning different combinations of contracts. I do not take into account sampling uncertainty in estimating the competition index.

²⁷While I can reject the null hypothesis of independence, the extent of this dependence is extremely small. I introduce dependence in the below procedure using a Gaussian Copula to allow correlation in these minimum rival bids. This correlation is allowed to depend on whether the contracts are the same type or in the same county. The maximum estimated correlation between any two winning bids is 0.0272, which I take as negligible.

engineer's estimate for that contract.²⁸ The scale is a function of auction-level characteristics and the competition index, denoted using the vector \mathbf{x}_{it} :

$$Prob(b_{nlt} \leq \min_{n' \neq n} \{b_{n'lt}\} ; \beta_1, \alpha) = 1 - e^{-\left(\frac{b_{nlt} - \frac{1}{3}}{\exp(\mathbf{x}_{it}\beta_1)}\right)^\alpha}$$

I assume that states transition according to an autoregressive order (1) process:

$$\begin{pmatrix} \lambda_{nt} \\ \mathbf{s}_{nt} \end{pmatrix} = \boldsymbol{\alpha}_n + \boldsymbol{\alpha} \begin{pmatrix} \lambda_{nt-1} \\ \mathbf{s}_{nt-1} \end{pmatrix} + \boldsymbol{\varepsilon}_{nt}$$

Where $\boldsymbol{\alpha}_n$ are firm specific intercepts, $\boldsymbol{\alpha}$ is a 3×3 dimension matrix, that is allowed to vary by firm type, and $\boldsymbol{\varepsilon}_{nt}$ is a white noise innovation.²⁹

Results from the first step are given in Figure 10. I present three specification, including varying sets of Fixed Effects. In later steps I use the County Fixed Effects specification, dropping the time fixed effects. I estimate the shape parameter well above one, ensuring that the Markup is monotonically increasing in bids. Note that mean of the distribution is increasing in the scale. For each of the scale parameters I include separate slope coefficients for each type of auction. For all three types of auction the winning bid is increasing in the competition index: When λ is large, so there is little competition, bids are less aggressive. Meanwhile the magnitude for Asphalt projects is in line with the results presented in Raisingh (2021). Magnitudes for concrete and mixed projects are similar.

We can interpret the coefficients on engineer's estimate (EE) as returns to scale, since the dependent variable (lowest rival bid) is normalised by EE. The persistent negative coefficient on asphalt suggests increasing returns, in line with GKS and Raisingh's results. In Appendix F.3 I present a kernel density plot of the minimum rival bids alongside the fitted distribution, demonstrating that the Weibull distributional assumption fits the data very well.

5.3.3 Second Stage

I assume the pseudo-static pay-off is quadratic in backlogs. I take this approach, despite the likely superior performance of a B-spline specification, due to the decreased computa-

²⁸As discussed in Raisingh (2021) this is because several projects appear to have miscalculated estimates. These are treated as outliers and removed. This occurred in around 0.1% of cases.

²⁹By construction backlogs transition deterministically. However, not all projects are completed at the same rate. Therefore I must take into account future deterministic backlogs in the state variable. I assume this transition function for simplicity, as $AR(1)$ processes are often used to model the transitions of inclusive value indices.

Figure 4: First Stage Results

		Coefficient	SE	Coefficient	SE	Coefficient	SE
Shape							
Scale	$\log(\alpha - 1)$ ($= e^{\mathbf{x}_{it}\beta_1}$)	2.029	0.001	2.083	0.001	2.093	0.001
	Concrete	-0.48	0.001	-0.484	0.002	-0.495	0.003
	Asphalt	-0.458	0.001	-0.449	0.002	-0.461	0.003
	Both	-0.44	0.001	-0.45	0.002	-0.462	0.003
	Major Road	-0.013	0.001	-0.007	0.001	-0.007	0.001
	Bridge	-0.001	0.001	0.005	0.001	0.003	0.001
	MR $\times \lambda$	0.048	0.001	0.042	0.001	0.041	0.001
	Bridge $\times \lambda$	0.027	0.001	0.024	0.001	0.021	0.001
	Concrete $\times \lambda$	0.183	0.001	0.186	0.001	0.187	0.001
	Asphalt $\times \lambda$	0.196	0.001	0.198	0.001	0.196	0.001
	Both $\times \lambda$	0.172	0.001	0.18	0.001	0.181	0.001
	Concrete $\times \log(\text{EE})$	0.008	0.001	0.001	0.001	0	0.001
	Asphalt $\times \log(\text{EE})$	-0.008	0.001	-0.011	0.001	-0.012	0.001
	Both $\times \log(\text{EE})$	-0.006	0.001	-0.009	0.001	-0.01	0.001
	Concrete $\times \lambda \times \log(\text{EE})$	-0.006	0.001	-0.004	0.001	-0.004	0.001
	Asphalt $\times \lambda \times \log(\text{EE})$	0.001	0.001	0.001	0.001	0.001	0.001
	Both $\times \lambda \times \log(\text{EE})$	-0.002	0.001	-0.001	0.001	-0.001	0.001
	Fixed Effects						
County				✓		✓	
Year						✓	
Month						✓	
	Observations	193545		193545		193545	

tional intensity as well as making it easier to interpret the parameter estimates: Testing for complementarities reduces to testing the significance of the quadratic terms. I normalise backlogs by the standard deviation of each firm's observed backlogs, so that backlog effects are estimated using within firm variation. Parameters can vary across the three firm types, so for a firm of type m the specification for the pseudo-static pay-off is:

$$k_m(\mathbf{s}_t) = \lambda_{nt}\theta_m^\lambda + \mathbf{h}(\mathbf{s}_{nt})^T\theta_m^\kappa + \lambda_{nt}\mathbf{h}(\mathbf{s}_{nt})^T\theta_m^{\kappa\lambda}$$

Where $\mathbf{h}(\mathbf{s}_{nt})^T = \left(s_{nt}^a \quad s_{nt}^c \quad (s_{nt}^a)^2 \quad (s_{nt}^c)^2 \quad s_{nt}^a \times s_{nt}^c \right)$

I make use of additional moments to facilitate estimation. If \mathbf{s}_t does not substantially shift bidding behaviour there may be a weak instrument problem. This occurs if a firm's observed backlog does not vary relatively much, but they bid on many contracts simultaneously so that the possible ex-post states \mathbf{s}_t^c vary much more than \mathbf{s}_t . In this case we are trying to

estimate κ in regions where there is little variation in our instrument. This would be a problem if firms are successfully smoothing their backlogs.³⁰

I include several additional instruments, or moment conditions, to ameliorate this problem. Write \vec{s}_l as the amount a firm’s backlog will increase if they win lot l . This is the engineer’s estimate of the project completion cost, split according to the type of contract. I make the additional assumption $E[v_{nlt}|\mathbf{s} + \vec{s}_l] = 0$, using the ex-post state from only winning lot l as an additional instrument. Many more potential instruments are available, using additional ex-post states as instruments. For illustrative purposes I also consider a specification that makes use of ex-post states from winning pairs of contracts, increasing the number of instruments ten-fold. However this risks overfitting the first stage.³¹

Figure 5 presents the results from the second estimation step and includes estimates from a least squares specification as well as three sets of instruments. Parameter interactions with the competition index are included in Appendix F.3. Estimates from the third column are used for the remainder of this application. Results are presented in thousands of dollars. So, for example, every kilometre increase in distance between a general contractor’s plant (t1) and the project increases costs by around \$170.

The coefficients on backlogs can be interpreted as their effect on the pseudo-payoffs: Every one standard deviation increase in a paving company’s (t2) backlog of asphalt projects increases their pseudo-cost (cost + expected future opportunity cost) by \approx \$870,000. Coefficients can also be interpreted as how they impact the aggressiveness of the firm’s bidding. The coefficients on linear backlogs are all positive, suggesting firms bid less aggressively on larger projects. We cannot interpret the quadratic coefficients from the second stage as evidence of returns to scale. However they give evidence of non-additivities across lots: The null hypothesis of additive values is rejected with p-value < 0.001 .

The post-estimation tests demonstrate that the choice of instruments is important. The Hansen test of over-identifying restrictions presented in column 4 rejects the null at the 1% significance level, suggesting these additional instruments are invalid. I cannot reject the validity of the additional instruments used in column 3. Likewise, the Hausman test for

³⁰This problem is alleviated if we do not use normalised backlogs, using variation across bidders to aid identification. However for this application this is undesirable.

³¹The distribution of contract sizes is very skewed, with a small number of extremely large contracts. These contracts impact backlogs much more than small contracts, and attract higher bids. These observations have a lot of leverage. To reduce the weight on these observations I weight observations by the inverse of the engineer’s estimate of lot l (EE_l). This is equivalent to using of moment conditions of the form $E[\frac{v_{nlt}}{EE_l}|\mathbf{s}_t] = 0$. Furthermore, it is standard to normalise bids and associated costs by the size of the lot, which makes a similar assumption.

endogeneity in column 1 also fails to reject. Meanwhile, the adjusted Cragg-Donald statistic in column 2 suggests that the initial state alone is a weak instrument. Therefore, even though we suspect the estimates from column 1 are inconsistent, they are almost certainly better estimates than those presented in column 2. This suggests that problems of weak instruments may be more damaging than failing to instrument at all.

Figure 5: Second Stage Results

Instruments		none (OLS)		s_{nt}		$s_{nt} + \vec{s}_{nlt}$		$s_{nt} + \vec{s}_{nlt} + \vec{s}_{nmt}$	
Type		$\hat{\theta}$	SE	$\hat{\theta}$	SE	$\hat{\theta}$	SE	$\hat{\theta}$	SE
Combinatorial									
s_t^a	t1	488	19.3	495	346	451	26.6	463	20.4
	t2	905	35	1940	2270	852	37	869	25.4
	t3	113	5.61	114	1190	113	5.96	113	4.75
s_t^c	t1	392	17.2	254	826	406	20.1	398	17.6
	t2	216	40.5	-4,380	9370	233	43.5	221	37.5
	t3	55.8	6.2	-10.1	5010	57	6.6	56.7	5.74
$(s_t^a)^2$	t1	-4.36	1.85	-21.6	67.3	-0.967	2.72	-1.76	2.27
	t2	-26	4.01	-221	481	-18	4.56	-21.5	3.35
	t3	-0.236	0.0645	-0.189	24.6	-0.26	0.0698	-0.246	0.0629
$(s_t^c)^2$	t1	-12	1.96	0.179	189	-15.9	2.4	-15.5	2.3
	t2	-18.6	6.04	741	1530	-26.1	6.91	-23.9	6.34
	t3	-0.245	0.0638	7.28	74.3	-0.344	0.119	-0.311	0.093
$s_t^a \times s_t^c$	t1	0.464	3.04	28.2	125	5.3	3.23	7.03	3.15
	t2	54.9	12.5	1.27	431	73.1	13.8	72.5	10.2
	t3	0.277	0.176	-21.6	103	0.553	0.366	0.482	0.279
Lot specific									
Distance									
	t1	0.159	0.0814	0.238	0.187	0.188	0.0779	0.164	0.0823
	t2	0.0597	0.104	-0.0511	0.599	0.0999	0.108	0.0667	0.102
	t3	0.159	0.0946	-0.018	2.87	0.166	0.0943	0.159	0.094
Tests				(stat)	(p-val)				
Hansen		36.5	(0.192)	-	(-)	19	(0.393)	637	(0)
Cragg-Donald		-		0.00257		178		119	
R^2		0.6		-10.8		0.597		0.599	

Note: Including county and firm \times contract type fixed effects. Col 1 Hansen test is a Hausman test of endogeneity, using instruments from col 3. Figures are given in 000s of dollars. The two-step consistent standard errors are clustered within bidder days. I winsorise the bottom percentile of estimated $\frac{\Gamma_l(b_{nlt})}{\nabla_b \Gamma_l(b_{nlt})}$, since the tails of the distribution are likely to be poorly estimated. Estimation uses $T = 3919$ observations.

5.3.4 Third Stage

After forming the expected maximised period pay-off $\hat{W}(\mathbf{b}_t|\hat{\kappa}; \mathbf{s}_t)$ I evaluate the ex-ante value function by approximating the conditional expectation over \mathbf{b}_t using a linear in parameters prediction of \bar{W}_t given $\mathbf{h}(\mathbf{s}_t)$.³² This ensures the ex-ante Value Function, for a firm of type m , can be written as: $E[\hat{W}_m(\mathbf{b}_t|\hat{\kappa}; \mathbf{s}_t)|\mathbf{s}_t] = \mu_n + \mathbf{h}(\mathbf{s}_t)^T \theta_m^V$. Observations are weighted according to their inverse variance, using $var(\hat{\theta}_m^\kappa)$. The quadratic form of \mathbf{h} and the AR(1) transition process implies $E[\mathbf{h}(\mathbf{s}_t)|\mathbf{s}_{t-1}] = \mathbf{h}(\mathbf{s}_{t-1})^T \theta_m^\tau$, where θ_m^τ is a $|\mathbf{h}| \times |\mathbf{h}|$ dimensional matrix function of $\boldsymbol{\alpha}_m$. This also implies I can write $\pi(\mathbf{s}_{nt}) = \mathbf{h}(\mathbf{s}_{nt})^T \theta_m^\pi$.

Figure 6 presents results from the third estimation step. Costs are increasing linear backlogs for all three types of firm. However, the magnitudes are much smaller than the linear coefficients estimated in the second stage. This suggests large anticipated opportunity costs from high backlogs. This result is sensible since projects have very long durations.

By considering the quadratic terms we see that general contractors only exhibit increasing returns to scale, or increasing returns to specialisation, in concrete contracts. Meanwhile, both paving and construction companies exhibit increasing returns for both types of contracts, but with a negative cost interaction. Taking on concrete (asphalt) projects come with additional costs for these firms already specialised in asphalt (concrete) projects. In Appendix F.4 I consider how my results compare to results from misspecified dynamic single-object, and static multi-object models. I find that the dynamic single-object model under-estimates the degree of non-additivity across lots. The static multi-object model over-estimates the effect of backlogs on costs, mistaking expected future costs for present costs.

5.4 Counterfactual

I now consider how procurement costs and efficiency change when contracts are allocated using sequential first-price auctions. This is an interesting counterfactual as it speaks to the importance of the ‘exposure problem’ as well as the value of ‘batching’. Furthermore, many empirical dynamic auction papers assume contracts are auctioned sequentially anyway, making this a useful comparison for researchers.

Theoretical results suggest sequential allocation will be less efficient than simultaneous

³²This assumption is technically incompatible with the parametric assumption made above. However we can test the extent of the misspecification error using a standard RESET test. I am unable to reject the null of no specification error (at the 10% significance level) using a RESET test of order 10. Meanwhile no explicit parametric assumptions were made on the distributions of b or v .

Figure 6: Third Stage Results

Object	Type	$\pi(\mathbf{s}_n)$		$V(\mathbf{s})$		$\kappa(\mathbf{s})$	
		$\hat{\theta}$	SE	$\hat{\theta}$	SE	$\hat{\theta}$	SE
λ	t1	0	(-)	5.39	0.63	6.79	0.763
	t2	0	(-)	15.6	1.58	6.47	1.69
	t3	0	(-)	6.81	0.713	1.81	0.228
s_t^a	t1	123	7.01	-451	26.6	451	26.6
	t2	285	11.3	-839	37.1	852	37
	t3	40	1.92	-103	6.07	113	5.96
s_t^c	t1	107	5.35	-405	20.1	406	20.1
	t2	89.1	11.9	-207	43.5	233	43.5
	t3	15.6	1.91	-57.8	6.63	57	6.6
$(s_t^a)^2$	t1	-0.337	1.29	1.74	2.73	-0.967	2.72
	t2	-9.26	2.46	18.6	4.65	-18	4.56
	t3	-1.34	0.147	-2.13	0.292	-0.26	0.0698
$(s_t^c)^2$	t1	-7.6	1.13	15.5	2.41	-15.9	2.4
	t2	-14	3.38	23.4	6.95	-26.1	6.91
	t3	-0.479	0.102	-0.262	0.202	-0.344	0.119
$s_t^a \times s_t^c$	t1	1.38	1.52	-6.7	3.25	5.3	3.23
	t2	33.4	7.12	-80.4	13.9	73.1	13.8
	t3	0.432	0.199	-0.237	0.405	0.553	0.366
Fixed Effects							
Firm				✓	✓		

allocation (batching).³³ Bidders do not know what types of contracts will be auctioned in the near future, making it more difficult to exploit cost synergies. However, batching contracts but not allowing firms to place combinatorial bids also limits their ability to exploit synergies (the exposure problem). Sequential allocation may improve efficiency by giving bidders greater control over their cost synergies, reducing the likelihood that bidders accidentally win too many or too few contracts. These effects will be more pronounced the larger the degree of complementarities (positive or negative) across lots. The effects of this alternate procurement mechanism are ex-ante unclear.

³³See Akbarpour et al. (2020) as an example. I ignore that collusion is easier to sustain in sequential auctions (Hendricks and Porter, 1989), further increasing procurement costs.

5.4.1 The Counterfactual Mechanism

I now briefly discuss how I simulate equilibrium bidding under the counterfactual mechanism. See Appendix F.5 for full details. Contracts are auctioned sequentially, in random order, within each 14 day period. Consistent with the estimated model I assume projects begin before the next auction. I use the same competition index λ_{nt} to capture changes in competition within these periods. Firms have beliefs about the probability they win any given lot, conditional on lot characteristics and λ_{nt} . Firms place bids conditional on their beliefs, backlogs, and their continuation value, defined as in the main model.³⁴ I find equilibrium beliefs and value functions using fixed point iteration: I repeatedly simulate the auction process until value functions and the distribution of winning bids converges.

5.4.2 Results

Figure 7 presents estimates of the average cost per contract for firms and MDOT, in thousands of dollars, under the simultaneous auction regime and the counterfactual sequential auction regime. The key takeaway is that sequential auctions decrease efficiency and raises procurement cost. Procurement costs are estimated to increase by an average of \$19,000 per contract (1.3%), while for firms completion costs increase by an average of \$110,000 per contract (9.4%). This suggests the batching effect dominates the exposure effect. In particular, the information benefit of batching is key: While bidding on a particular contract, firms do not know which contracts will be auctioned later in the period. In the counterfactual firms bids 1.2% less aggressively on contracts they were actually observed winning, demonstrating they are unaware that they will be able to exploit a cost synergy with a lot being auction later in the period.

However, the increase in procurement cost is much smaller than the increase in completion costs because firms face more competition for each contract. At any one time, instead of N firms competing for L contracts there are N firms competing for 1 contract, unsure of when any future contracts will be auctioned. Markups are around 40% lower under sequential auctions, but this increase in competition remains dominated by the batching effect. Furthermore, this competition finding strongly relies on the assumption of a non-

³⁴I assume firms only place bids on the set of auctions they actually bid on. Given my assumption of negligible entry costs, firms were only observed bidding on the contracts they have the largest cost advantages in. If their cost advantages were mostly additive, such as due to low v_{nlt} draws, they will have the same advantage under the sequential mechanism, and so bidding on this set of lots will remain optimal. Therefore my estimates can, to an extent, be considered lower bounds on costs.

collusive equilibrium.

Figure 7: Counterfactual Results

Mechanism	Outcome	Estimate (\$000s)	S.E.
Simultaneous Auctions	Procurement Cost	1470	-
	Completion Cost	1170	4.28
Sequential Auctions	Procurement Cost	1489	3
	Completion Cost	1280	22.6

Note: The results are based on 60 draws of parameters from their estimated asymptotic distribution. Equilibrium Beliefs and Value Functions are computed for each draw.

6 Conclusion

In this paper I did three things: First, I set-up a dynamic multi-object auction model and proved that the model primitives are identified from standard bidding data. Second, I proposed a computationally convenient estimation procedure to overcome the technical challenges of estimating model primitives in this setting. Finally, I applied the model to data from Michigan Department of Transport’s procurement data and evaluated the efficiency and revenue of holding repeated rounds of simultaneous auction relative to auctioning all contracts sequentially.

This paper was motivated by the prevalence of such repeated, multi-object auctions. Significant complementarities between auctioned objects have been found in both the dynamic single-object literature, and the static multi-object literature, most notably in JP and GKS. However, these two types of model had not, until this point, been unified in a single framework.

References

- Akbarpour, M., Li, S., and Gharan, S. O. (2020). Thickness and information in dynamic matching markets. *Journal of Political Economy*, 128(3):783–815.
- Aradillas-Lopez, A., Haile, P. A., Hendricks, K., Porter, R. H., and Raisingh, D. (2022). Testing competition in us offshore oil and gas lease auctions. *working paper*.
- Arsenault Morin, A., Arslan, H. A., and Gentry, M. L. (2022). On the timing of auctions: The effects of complementarities on bidding, participation, and welfare. *working paper*.
- Athey, S. and Haile, P. A. (2002). Identification of standard auction models. *Econometrica*, 70(6):2107–2140.

- Athey, S. and Haile, P. A. (2007). Nonparametric approaches to auctions. *Handbook of econometrics*, 6:3847–3965.
- Athey, S., Levin, J., and Seira, E. (2011). Comparing open and sealed bid auctions: Evidence from timber auctions. *The Quarterly Journal of Economics*, 126(1):207–257.
- Aumann, R. J. (1974). Subjectivity and correlation in randomized strategies. *Journal of mathematical Economics*, 1(1):67–96.
- Backus, M. and Lewis, G. (2016). Dynamic demand estimation in auction markets. Technical report, National Bureau of Economic Research.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.
- Balat, J. (2013). Highway procurement and the stimulus package: Identification and estimation of dynamic auctions with unobserved heterogeneity. *Johns Hopkins University Mimeo*.
- Bodoh-Creed, A. L., Boehnke, J., and Hickman, B. (2021). How efficient are decentralized auction platforms? *The Review of Economic Studies*, 88(1):91–125.
- Cantillon, E. and Pesendorfer, M. (2007). Combination bidding in multi-unit auctions.
- Charalambos, D. and Aliprantis, B. (2013). *Infinite Dimensional Analysis: A Hitchhiker’s Guide*. Springer-Verlag Berlin and Heidelberg GmbH & Company KG.
- Flambard, V. and Perrigne, I. (2006). Asymmetry in procurement auctions: Evidence from snow removal contracts. *The Economic Journal*, 116(514):1014–1036.
- Fox, J. T. and Bajari, P. (2013). Measuring the efficiency of an fcc spectrum auction. *American Economic Journal: Microeconomics*, 5(1):100–146.
- Fudenberg, D. and Maskin, E. (1991). On the dispensability of public randomization in discounted repeated games. *Journal of Economic Theory*, 53(2):428–438.
- Fudenberg, D. and Maskin, E. (2009). The folk theorem in repeated games with discounting or with incomplete information. In *A Long-Run Collaboration On Long-Run Games*, pages 209–230. World Scientific.
- Gentry, M., Komarova, T., and Schiraldi, P. (2023). Preferences and performance in simultaneous first-price auctions: A structural analysis. *The Review of Economic Studies*, 90(2):852–878.
- Groeger, J. R. (2014). A study of participation in dynamic auctions. *International Economic Review*, 55(4):1129–1154.
- Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal nonparametric estimation of first-price auctions. *Econometrica*, 68(3):525–574.
- Hendricks, K., Pinkse, J., and Porter, R. H. (2003). Empirical implications of equilibrium bidding in first-price, symmetric, common value auctions. *The Review of Economic Studies*, 70(1):115–145.
- Hendricks, K. and Porter, R. H. (1989). Collusion in auctions. *Annales d’Economie et de Statistique*, pages 217–230.
- Hickman, B. R., Hubbard, T. P., and Paarsch, H. J. (2017). Identification and estimation of a bidding model for electronic auctions. *Quantitative Economics*, 8(2):505–551.

- Hortaçsu, A. and McAdams, D. (2018). Empirical work on auctions of multiple objects. *Journal of Economic Literature*, 56(1):157–84.
- Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529.
- Jeziorski, P. and Krasnokutskaya, E. (2016). Dynamic auction environment with subcontracting. *The RAND Journal of Economics*, 47(4):751–791.
- Jofre-Bonet, M. and Pesendorfer, M. (2003). Estimation of a dynamic auction game. *Econometrica*, 71(5):1443–1489.
- Kim, S. W., Olivares, M., and Weintraub, G. Y. (2014). Measuring the performance of large-scale combinatorial auctions: A structural estimation approach. *Management Science*, 60(5):1180–1201.
- Kong, Y. (2021). Sequential auctions with synergy and affiliation across auctions. *Journal of Political Economy*, 129(1):148–181.
- Lewbel, A. (2019). The identification zoo: Meanings of identification in econometrics. *Journal of Economic Literature*, 57(4):835–903.
- Magnac, T. and Thesmar, D. (2002). Identifying dynamic discrete decision processes. *Econometrica*, 70(2):801–816.
- Milgrom, P. R. and Weber, R. J. (1985). Distributional strategies for games with incomplete information. *Mathematics of operations research*, 10(4):619–632.
- Mises, R. v. (1947). On the asymptotic distribution of differentiable statistical functions. *The annals of mathematical statistics*, 18(3):309–348.
- Newey, W. K. and McFadden, D. (1994). Large sample estimation and hypothesis testing. *Handbook of econometrics*, 4:2111–2245.
- Pesendorfer, M. and Schmidt-Dengler, P. (2008). Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies*, 75(3):901–928.
- Raisingh, D. (2021). The effect of pre-announcements on participation and bidding in dynamic auctions. *working paper*.
- Somaini, P. (2020). Identification in auction models with interdependent costs. *Journal of Political Economy*, 128(10):3820–3871.

Appendices

A Proof of Proposition 3

In this Appendix I essentially extend Proposition 1 from JP to the multi-object case, proving Proposition 3 from the main text.

Proof: 1. Necessary First Order Conditions are given by:

$$\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})(\mathbf{v} - \mathbf{b}^*) = \Gamma(\mathbf{b}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\boldsymbol{\pi} - \beta\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\mathbf{V}$$

2. Left multiplying by $\Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}$ yields: $\Gamma(\mathbf{b}^*|\mathbf{s})^T(\mathbf{v} - \mathbf{b}^*) =$

$$\Gamma(\mathbf{b}^*|\mathbf{s})^T\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*|\mathbf{s})B_{\mathbf{s}}\boldsymbol{\pi} - \beta\nabla_{\mathbf{b}}Q(\mathbf{b}^*|\mathbf{s})A_{\mathbf{s}}\mathbf{V}]$$

3. Substituting $\Gamma(\mathbf{b}^*|\mathbf{s})^T(\mathbf{v} - \mathbf{b}^*)$ into equation 2 gives the result.

□

Given Proposition 3 we take an expectation over the expected stage payoff, with respect to observed bids, to show that the ex-ante value function can be written as a function of the distribution of equilibrium bids and $\bar{\kappa}(\mathbf{s})$ only.

B Proof of Proposition 4

I now prove that $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_i - 1$. The proof is in three parts. First, I establish the rank of Ψ , then find its null space. I then demonstrate that the intersection of this null space and the image of $(I_S - \beta T\Omega)^{-1}C$ only contains a single element.

B.1 Rank of Ψ

First, define the partial ordering \succeq^* such that if $\mathbf{s}_n \succeq \mathbf{s}'_n$ then $\mathbf{s} \succeq^* \mathbf{s}'$. This simply extends the partial ordering of the individual state to the overall state.

Define a ‘component’ \mathbb{S}^C as a subset of \mathbb{S} that is ‘connected’ by this partial ordering. Formally, $\mathbf{s}, \mathbf{s}' \in \mathbb{S}^C$ if and only if there exists a non-directed path between the states; that is if there exists a (finite) sequence of states beginning with \mathbf{s} , ending with \mathbf{s}' where for each pair $\mathbf{s}^k, \mathbf{s}^{k+1}$ in this sequence either $\mathbf{s}^k \succeq^* \mathbf{s}^{k+1}$ or $\mathbf{s}^k \preceq^* \mathbf{s}^{k+1}$. By definition \mathbf{s}_0 does not vary within a component, and in general there is one component corresponding to each element of \mathbb{S}_0 . The \mathbb{S}^C components form a partition of \mathbb{S} .

Finally, denote $\tilde{\min}(\mathbb{S})$ as the subset of \mathbb{S} , such that $\forall \mathbf{s} \in \tilde{\min}(\mathbb{S}) : \nexists \mathbf{s}' \in \mathbb{S} : \mathbf{s} \in \mathbb{S}^c(\mathbf{s}')$. This is primarily for notational convenience, and may not coincide with the minimal elements of \mathbb{S} . Instead, this is the (potentially empty) set of states that never occur as possible ex-post states. Intuitively, pay-offs from ending in these states are be identified.

B.1.1 Additional Lemmas

Lemma B.1. *From any two distinct, non-maximal, states, \mathbf{s} and \mathbf{s}' , if $\mathbf{s}' \not\geq^* \mathbf{s}$ then there exists a state \mathbf{s}^c such that $\mathbf{s}^c \in \mathbb{S}^c(\mathbf{s})$ & $\mathbf{s}^c \notin \mathbb{S}^c(\mathbf{s}')$*

This states that if one non-maximal state is not ‘higher’ in the partial ordering than another, their set of ex-post states cannot perfectly overlap. The proof examines whether the maximal element of $\mathbb{S}^c(\mathbf{s})$ (when bidder n winning every lot, denoted \mathbf{s}^{all_n}) can be an element of $\mathbb{S}^c(\mathbf{s}')$. I exploit that $\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$ forms a lattice, with minimal state corresponding to winning no lots, and maximal state corresponding to winning every lot.

Proof:

1. Suppose $\mathbf{s}' \not\geq^* \mathbf{s}$. Therefore either $\mathbf{s} \succeq^* \mathbf{s}'$, or the states are incomparable.
2. If $\mathbf{s} \succeq^* \mathbf{s}'$ they are in the same component, so $\mathbf{s}_0 = \mathbf{s}'_0$. Assumption 6 iii) implies the maximal (win all) element of $\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)$ is ‘greater’ than the maximal element of $\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$, hence $\maximal(\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)) \notin \mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$.
3. If \mathbf{s} and \mathbf{s}' are incomparable, then \mathbf{s} and \mathbf{s}' either belong to different components, or the same component. If they belong to different components then by definition $\mathbb{S}^c(\mathbf{s})$ and $\mathbb{S}^c(\mathbf{s}')$ must be mutually exclusive.
4. If \mathbf{s} and \mathbf{s}' are incomparable but in the same component then Assumption 6 iii) ensures $\maximal(\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0))$ and $\maximal(\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0))$ are incomparable. Therefore, it cannot be that $\maximal(\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)) \in \mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$, since $\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0)$ is a lattice it requires $\maximal(\mathbb{S}_n^c(\mathbf{s}_n, \mathbf{s}_0)) \succeq \maximal(\mathbb{S}_n^c(\mathbf{s}'_n, \mathbf{s}_0))$

□

Lemma B.2. $\Psi(\mathbf{s})A_{\mathbf{s}}$ has rank at least 2 if, for all $\mathbf{s}, \mathbf{v}, l$, $\Gamma_{nl}(\mathbf{b}(\mathbf{v}, \mathbf{s})|\mathbf{s}) \in (0, 1)$

The proof proceeds by first showing that $rank(\Psi(\mathbf{s}))$ is weakly greater than two, then using the full rank property of the transformation matrix $A_{\mathbf{s}}$.³⁵

Proof:

1. Denote by $\tilde{\Psi}$ the $L \times L(N-1)^{L-1}$ sub-matrix of $\Psi(\mathbf{s})$ consisting of the columns of $\Psi(\mathbf{s})$ corresponding to outcomes in which player n wins exactly one lot.
2. Row l , column c of $\tilde{\Psi}$ is strictly positive for columns corresponding to outcomes c in which bidder n wins lot l . This is because the probability that n wins lot l , and no other lot, is strictly increasing in b_l .

³⁵In general $\Psi(\mathbf{s})A_{\mathbf{s}}$ has rank L . Essentially, each state gives us L pieces of information, rather than just two pieces of information. However, proof that the rank is always L has proven elusive.

3. Row l , column c of $\tilde{\Psi}$ is strictly negative for columns corresponding to outcomes c in which n does not win lot l . This is because the probability lot l is won, and no other, is strictly decreasing in b_m for $m \neq l$.
4. Any two rows of $\tilde{\Psi}$ are linearly independent: Each row contains one positive entry, each in a distinct column.³⁶ Therefore, $\tilde{\Psi}$, and hence $\Psi(\mathbf{s})$ have rank ≥ 2 .
5. Matrix $A_{\mathbf{s}}$ is a rank N^L transformation matrix for any non-maximal \mathbf{s} . Therefore, from step 4, $\Psi(\mathbf{s})A_{\mathbf{s}}$ for non-maximal \mathbf{s} has rank at least 2.

□

B.1.2 $Rank(\Psi) = S - S^C - |\tilde{\min}(\mathbb{S})|$

I show that as we stack these $\Psi(\mathbf{s})A_{\mathbf{s}}$ matrices for non-maximal \mathbf{s} , the rank increases by *at least* two each time. However, by definition columns corresponding to elements in $\tilde{\min}(\mathbb{S})$ are all zero, ensuring the rank is deficient by at least $|\tilde{\min}(\mathbb{S})|$. Likewise, for each submatrix of Ψ made up of rows corresponding to states that are all within the same component (denoted by Ψ^C , a $|\mathbb{S}^C| \times S$ matrix), rows all sum to zero. This ensures each Ψ^C is rank deficient by at least one, and so Ψ is rank deficient by at least S^C .

- Proof:*
1. Order elements of \mathbb{S} (likewise, columns of Ψ) according to the partial ordering \succeq^* . Incomparable states are ordered at random. So, for each \mathbf{s} , the furthest left non-zero column of $\Psi(\mathbf{s})A_{\mathbf{s}}$ is in the column corresponding to the ex-post state in which player i wins every lot \mathbf{s}^{all_n} .
 2. Focus on one component, \mathbb{S}^C . Find the ‘smallest’ state within \mathbb{S}^C , \mathbf{s}_1^C (i.e. right most column index of Ψ). This must be a minimal element of \mathbb{S}^C .
 3. Find the second smallest state \mathbf{s}_2^C , which may also be a minimal element. Vertically stack the matrices $\Psi(\mathbf{s}_1^C)A_{\mathbf{s}_1^C}$ and $\Psi(\mathbf{s}_2^C)A_{\mathbf{s}_2^C}$, for $\Psi_{\{1,2\}}^C$.
 4. $\Psi_{\{1,2\}}^C$ has rank ≥ 4 . Lemma B.2 ensures that both matrices have rank 2, while lemma B.1 ensures that each row of $\Psi(\mathbf{s}_1^C)A_{\mathbf{s}_1^C}$ is linearly independent of each row of $\Psi(\mathbf{s}_2^C)A_{\mathbf{s}_2^C}$. This last point arises because lemma B.1 ensures that since $\mathbf{s}_1^C \not\succeq^* \mathbf{s}_2^C$ there must be at least one column of non-zero entries in $\Psi(\mathbf{s}_2^C)A_{\mathbf{s}_2^C}$ that matches up to an all-zero column of $\Psi(\mathbf{s}_1^C)A_{\mathbf{s}_1^C}$.

³⁶This only holds for $L \geq 3$. For $L = 2$ we must also assume $E[\Gamma_1 + \Gamma_2] \neq 1$.

5. Continue this process for each non-maximal state in component \mathbb{S}^C . At each stage, based on the ordering of elements in \mathbb{S} at step 1, and from lemmas B.2 and B.1, $\Psi(\mathbf{s}_n^C)A_{\mathbf{s}_n^C}$ must always contain at least one non-zero column that matches up to an all-zero column of $\Psi_{\{1,2,\dots,N-1\}}^C$. Typically this is the furthest left column, corresponding to $\mathbf{s}_N^{C\text{all}_n}$. Therefore, the rank increases by at least 2 each step.
6. The final matrix $\Psi_{\{1,2,\dots\}}^C$ has non-zero entries *somewhere* in each of the $|\mathbb{S}^C|$ columns corresponding to states in this set, except for columns correspond to elements of $\tilde{\min}(\mathbb{S}^C)$. These columns are all zeros — there is always zero probability of ending in these states. As the rank of this matrix increased by \geq two at each additional non-maximal state, and because we have at least as many non-maximal states as maximal states, this matrix must have rank $\geq |\mathbb{S}^C| - |\tilde{\min}(\mathbb{S}^C)| - 1$. The rank cannot exceed this, and must be strictly less than $|\mathbb{S}^C| - |\tilde{\min}(\mathbb{S}^C)|$ because the row sum for each row of this final matrix equals zero, a property inherited from the fact that $Q^T \mathbf{1} = 1$.
7. Any two components \mathbb{S}^C and $\mathbb{S}^{C'}$ are mutually exclusive. Therefore, the two matrices for any two components $\Psi_{\{1,2,\dots\}}^C$ do not share non-zero columns. As we stack these matrices across different components, the ranks sum together at each step.
8. Therefore $\text{rank}(\Psi) = \sum_{\mathbb{S}^C \subset \mathbb{S}} |\mathbb{S}^C| - |\tilde{\min}(\mathbb{S}^C)| - 1 = S - |\tilde{\min}(\mathbb{S})| - S^C$

□

B.2 nullspace of Ψ

B.2.1 The $|\tilde{\min}(\mathbb{S})|$ elements

Ψ contains only zeros in columns corresponding to states in $\tilde{\min}(\mathbb{S})$. Any vector \mathbf{y} containing non-zero entries only in rows corresponding to elements of this set is in this null space. Denote this set of vectors \mathbb{Y}^1 , with $|\tilde{\min}(\mathbb{S})|$ distinct elements.

B.2.2 The S^c elements

Consider the vector \mathbf{y} such that $y_s = y_{s'}$ if \mathbf{s} and \mathbf{s}' belong to the same component. Denote this set of vectors \mathbb{Y}^2 , containing S^C distinct elements. As established above, columns of the submatrix $\Psi_{\{1,\dots,|\mathbb{S}^C|\}}^C$ that correspond to states in different components contain all zeros, from the definition of a component.

Therefore, for any $\mathbf{y} \in \mathbb{Y}^2$ we have $\Psi^C \mathbf{y} = 0$. Entries of \mathbf{y} are constant across rows that correspond to the non-zero entries of $\Psi_{\{1 \dots |S^C|\}}^C$. This holds for any C . Therefore, as we stack the $\Psi_{\{1 \dots |S^C|\}}^C$ s into Ψ we will have $\Psi \mathbf{y} = 0$ for any $\mathbf{y} \in \mathbb{Y}^2$.

B.3 Image of $(I_S - \beta T \Omega)^{-1} C$

I have established that the null space of Ψ is given by $\mathbb{Y}^1 \cup \mathbb{Y}^2$. I now show that the intersection of this space and the image of $(I_S - \beta T \Omega)^{-1} C$ only contains the constant vector, denoted ι_{S^n} . This result requires three additional lemmas:

Lemma B.3. *For any $\mathbf{y} \in \mathbb{Y}^1$ we have $\Omega \mathbf{y} = 0$.*

Proof: 1. Recall that $\Omega(\mathbf{s}) = E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}] A_{\mathbf{s}}$
 2. $A_{\mathbf{s}} \mathbf{y} = 0$ for $\mathbf{y} \in \mathbb{Y}^1$, since $A_{\mathbf{s}}$ selects elements of \mathbf{y} corresponding to possible ex-post states, given beginning in \mathbf{s} . But \mathbf{y} only contains non-zero entries for states that are never observed as ex-post states.

□

Lemma B.4. *For any $\mathbf{y} \in \mathbb{Y}^2$ we have $\Omega \mathbf{y} = \mathbf{y}$.*

Proof: 1. For $\mathbf{y} \in \mathbb{Y}^2$ $A_{\mathbf{s}} \mathbf{y} = y_{\mathbf{s}} \iota_{2^L}$, where ι_{2^L} is a $2^L \times 1$ vector of ones. This is because $A_{\mathbf{s}}$ selects the elements of the vector \mathbf{y} that correspond to states that are possible outcomes from an auction round beginning in state \mathbf{s} . By definition these ex-post states are all in the same component, while \mathbf{y} is constant within components.
 2. As the rows of $Q(\mathbf{b}^*|\mathbf{s})^T$ sum to one, we have $E_{\mathbf{b}}[Q(\mathbf{b}^*|\mathbf{s})^T|\mathbf{s}] \iota_{2^L} = \iota_{2^L}$.
 3. As rows of $\nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})$ sum to zero (derivative of a vector with rows summing to one) we have: $E[\Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}] \iota_{2^L} = \mathbf{0}$
 4. Therefore $\Omega(\mathbf{s}) \mathbf{y} = y_{\mathbf{s}} \iota_{2^L}$ for $\mathbf{y} \in \mathbb{Y}^2$. Stacking over \mathbf{s} yields the result.

□

Finally, for $\mathbf{y} \in \mathbb{Y}^2$ we can write $\mathbf{y} = M \bar{\mathbf{y}}$ Where $\bar{\mathbf{y}}$ is an $S^C \times 1$ vector containing the constant elements of \mathbf{y} from component. Meanwhile M is an $S \times S^C$ dimensional matrix that contains a 1 in a row corresponding to state \mathbf{s} and column corresponding to component C if $\mathbf{s} \in \mathbb{S}^C$, and zero otherwise. Each row of M contains a single 1.

Lemma B.5. *Let the matrix N be any $S^C \times S^C$ submatrix of $(I - \beta T)M$ that is formed by selecting one row from each of the S^C components. N is non-singular.*

Proof:

1. Select S^C states, one from each component, and denote the corresponding set of rows of by \mathbb{M} . The sub-matrix of interest is denoted $M_{\mathbb{M},.} - \beta T_{\mathbb{M},.}M$
2. $M_{\mathbb{M},.} = I$. This is because we chose one row of M associated with each component. Each row of M contains a single 1, therefore so must $M_{\mathbb{M},.}$. Because every row is associated with a different component, each row contains a 1 in a different column.
3. Elements of the $S^C \times S^C$ sub matrix $T_{\mathbb{M},.}M$ are just transition probabilities, so $T_{\mathbb{M},.}M \iota_{S^C} = 1$. This is because right multiplying by M causes us to sum over states within a component. For a particular row t we have element C of the row vector $T_{t,.}M$ is equal to $\sum_{\mathbf{s}: \mathbf{s}^C = \mathbf{s}^{\bar{C}}} P(\mathbf{s}|\mathbf{s}^t)$. That is, the probability, given ending a period in state \mathbf{s}^t , that they begin the next period in component C .
4. Diagonal entries of the matrix $I - \beta T_{\mathbb{M},.}M$ are strictly positive, as $\beta \times$ a probability is strictly less than 1 (for $\beta < 1$). Likewise, off diagonal entries are weakly negative, as we have $-\beta \times$ a probability. Last, rows must sum to $1 - \beta$ because rows of both I and $T_{\mathbb{M},.}M$ sum to 1. This ensures this matrix is strictly diagonally dominant. Therefore, from the Levy–Desplanques theorem, the matrix must be non-singular.

□

B.3.1 $Image((I_S - \beta T\Omega)^{-1}C) \cap null(\Psi) = \iota_{S^i}$

The proof employs the result $T\iota_S = \iota_S$ (rows of a transition matrix sum to one). The proof proceeds by first demonstrating that the image of $(I_S - \beta T\Omega)^{-1}C$ does not intersect \mathbb{Y}^1 . Next, that the intersection with \mathbb{Y}^2 only contains the constant vector.

Proof:

1. Suppose there exists an \mathbf{x} such that for some $\mathbf{y} \in \mathbb{Y}^1$ we could write $\mathbf{y} = (I_S - \beta T\Omega)^{-1}C\mathbf{x}$. Equivalently, $(I_S - \beta T\Omega)\mathbf{y} = C\mathbf{x}$.
2. From Lemma B.3 this implies $\mathbf{y} = C\mathbf{x}$. In turn, from the definition of C this requires \mathbf{x} contains zeros in every entry except the first.
3. However this cannot be the case, since we always normalise this first entry to zero. Therefore $image((I_S - \beta T\Omega)^{-1}C) \cap \mathbb{Y}^1 = \emptyset$

4. Next, suppose there exists an \mathbf{x} such that for some $\mathbf{y} \in \mathbb{Y}^2$ we could write $\mathbf{y} = (I - \beta T \Omega)^{-1} C \mathbf{x}$. Equivalently $(I - \beta T \Omega) \mathbf{y} = C \mathbf{x}$
5. From Lemma B.4 $C \mathbf{x} = (I - \beta T) \mathbf{y} = (I - \beta T) M \bar{\mathbf{y}}$. In matrix form:

$$\begin{pmatrix} M - \beta \bar{T} & -C \end{pmatrix} \begin{pmatrix} \bar{\mathbf{y}} \\ \mathbf{x} \end{pmatrix} = 0$$

Where $\bar{T} = TM$, the probability of transitioning to any component from any state. If $(M - \beta \bar{T}, -C)$, the $S \times (S^C + S_n)$ matrix has rank $S^C + S_n - 1$ then there is a unique \mathbf{y} and \mathbf{x} where this relationship holds.

6. I now show the first column of $-C$ is linearly independent of $(M - \beta \bar{T})$. $-C_{.,1}$ contains -1 in every element associated with states such that $\mathbf{s}_n = \mathbf{s}_n^1$ and zeros otherwise. No linear combination of the columns for the corresponding rows of $(M - \beta \bar{T})$ can match these zeros. Choose S^C rows of $(M - \beta \bar{T})$ such that each row is associated with a state from a different component. E.g. rows such that in each component $\mathbf{s}_n = \mathbf{s}_n^{S_n}$ — the ‘final’ individual state. Call the corresponding $S^C \times S^C$ submatrix of $M - \beta \bar{T}$ N . From Lemma B.5 N is non-singular. No $S^C \times 1$ vector \mathbf{z} exists such that $N\mathbf{z} = 0$. Therefore columns of $(M - \beta \bar{T})$ are linearly independent of $-C_{.,1}$. By concatenating this column, the rank increases by one.
7. Repeat this process for columns $k = 2 \dots S_n - 1$ of $-C$. That is, every column *except* the final column which is the only column to contain non-zeros in entries associated with $\mathbf{s}_n^{S_n}$.³⁷ Each of these columns must be linearly independent of $M - \beta \bar{T}$ — no linear combination of its columns can match the zero entries of $-C_{.,k}$, since any $S^C \times S^C$ submatrix that consists of one row from each component must be non-singular.
8. Columns of $-C$ are linearly independent. So, at each step n the rank increases by 1. Therefore $\text{rank}(M - \beta \bar{T}, -C) \geq S^C + S_n - 1$.
9. $(\bar{\mathbf{y}} = \iota_{S^C}, \mathbf{x} = (1 - \beta)\iota_{S_n})$ lies in the null space of $(M - \beta \bar{T}, -C)$. This is because $(M - \beta \bar{T})\iota_{S^C} = (1 - \beta)\iota_S$ while we also have $C(1 - \beta)\iota_{S_n} = (1 - \beta)\iota_S$. Appeal to the rank-nullity theorem for $\text{Image}((I_S - \beta T \Omega)^{-1} C) \cap \text{null}(\Psi) = \iota_{S^N}$

³⁷This assumes one individual state exists within each component (I used $\mathbf{s}_n^{S_i}$). This holds if $\mathbb{S} = \mathbb{S}^0 \times \prod_n \mathbb{S}_n$. This is not necessary — the only requirement is that at each step k I can select one state from each component such that the corresponding rows of $-C_{.,k}$ are all zero.

□

C Proof of Proposition 1

In this Appendix I prove Proposition 1, which states that under the assumptions of the game, and under Conjecture 1, a Symmetric Markov Perfect Equilibrium exists.

First I prove that, conditional on Conjecture 1, a Pure Strategy Bayesian Nash Equilibrium exists in the stage game. I then show that the equilibrium pay-off in the stage game is consistent with the continuation value, employing Kakutani's fixed point theorem. This requires showing the existence, convex-valuedness, and upper hemicontinuity of the continuation value. While I assume entry is costless in my identification framework, bidders still make a strategic decision over which auctions to enter. Therefore I consider the entry game when discussing equilibrium existence.

Proof: **Equilibrium of the entry game:** player i chooses entry decision \mathbf{d} to maximise their expected payoff, taking an expectation over rivals' entry decisions given their strategies. This is a standard game of incomplete information.

A symmetric equilibrium in distributional strategies exists (Milgrom and Weber, 1985). Because types are atomless, existence of a Pure Strategy equilibrium follows from their purification result. This equilibrium may not be unique, so the value function may not be continuous. Continuity arises by augmenting entry strategies to be a function of the realisation of a public random variable (Fudenberg and Maskin, 1991). Public randomisation enables players to coordinate equilibria. Conditional on this public random variable the set of equilibrium pay-offs is convex (Aumann, 1974).

Equilibrium existence of the dynamic game requires that the equilibrium pay-off in the stage game is consistent with the continuation value.³⁸ That is, can we write the ex-ante value function \mathbf{V}_t^E , stacked over \mathbf{s} , as a function of \mathbf{V}_{t+1}^E , so that $\mathbf{V}_t^E = \Omega(\mathbf{V}_{t+1}^E)$ (existence). Stationarity requires the correspondence Ω has a fixed point: $\mathbf{V}^E = \Omega(\mathbf{V}^E)$.

Existence of $\mathbf{V}_t^E = \Omega(\mathbf{V}_{t+1}^E)$: Taking an expectation over Equation 1 with respect to \mathbf{v}_{nt} ensures we can write the ex-ante value function recursively. Existence then

³⁸Symmetry of the dynamic equilibrium arises because equilibrium in the stage game is symmetric, with strategies depending on states not identities or time periods.

follows from the assumption that pay-offs are bounded, ensuring the set $\Omega(\mathbf{V}_{t+1}^E)$ is non-empty.

(non-)Uniqueness of $\Omega(\mathbf{V}_{t+1}^E)$: The possibility of multiple equilibria in the entry game imply the value function is non-unique. So the ex-ante value function is also non-unique. Fortunately Ω must be convex valued, as the set of equilibrium pay-offs, conditional on the public random variable, is convex.

Upper-hemi continuity of $\Omega(\cdot)$: The continuation value is continuous in \mathbf{V}_{t+1}^E , taking an expectation over the transition process. Consider the conditional value function, conditional on entry decision $\bar{\mathbf{d}}$: $\tilde{W}_n(\bar{\mathbf{d}}, \mathbf{v}_{nt}, \mathbf{s}_t; \sigma_{-i}) =$

$$\max_{\mathbf{b}} \left\{ \Gamma_n(\mathbf{b}, \bar{\mathbf{d}}; \sigma_{-n})^T (\mathbf{v}_{nt} - \mathbf{b}) + Q_n(\mathbf{b}, \bar{\mathbf{d}}; \sigma_{-n})^T [\Pi_n(\mathbf{s}_t) + \beta V_n(\mathbf{s}_t; \sigma_{-n})] \right\}$$

Continuity of $\tilde{\mathbf{W}}_t$ in \mathbf{V}_{t+1}^E is guaranteed by conjecture 1, which requires equilibrium expected pay-offs are continuous in $\Pi_n + \beta V_n$. The value function is then $W_n(\mathbf{v}_{nt}, \mathbf{s}_t; \sigma_{-n}) = \max_{\mathbf{d}} \left\{ \tilde{W}_n(\mathbf{d}, \mathbf{v}_{nt}, \mathbf{s}_t; \sigma_{-n}) \right\}$. Upper-hemi continuity of \mathbf{W}_t in $\tilde{\mathbf{W}}_t$, and hence in \mathbf{V}_{t+1}^E , arises from our public random variable (Fudenberg and Maskin, 2009).³⁹ Upper-hemi continuity of \mathbf{V}_t^E arises from the ex-ante value function taking an expectation over states.

Existence of a stationary dynamic equilibrium: In order to show existence of a stationary equilibrium we must show that there exists a fixed point of the correspondence $\mathbf{V}^E = \Omega(\mathbf{V}^E)$. As $\Omega(\cdot)$ is non-empty, convex valued, and upper-hemi continuous, we can apply Kakutani's fixed point theorem. Therefore, a Markov Perfect Equilibrium exists.

□

³⁹Public randomisation ensures that the set of equilibrium pay-offs is convex. Public randomisation means \mathbf{W}_t is the convex hull of possible equilibrium pay-offs from entry, $\tilde{\mathbf{W}}_t$. Therefore, so long as $\tilde{\mathbf{W}}_t$ is compact valued, \mathbf{W}_t is upper hemicontinuous (Charalambos and Aliprantis, 2013). Compact valuedness comes from pay-offs being drawn from a compact set.

D Extensions

D.1 Second-Price Auctions

My identification results extend, almost trivially, to second price auctions. I do not discuss estimation of the second price model. However the estimation procedure presented in Section 4 can easily be applied, making use of the inverse bid system presented below.

D.1.1 Setup

In the second price setting n wins lot l at time t if $b_{nlt} > \max_{n'} \{b_{n'lt}\}$. As in the text, let $\Gamma(\mathbf{b}|\mathbf{s})$ denote the $L \times 1$ equilibrium marginal probabilities of winning each lot. Define vectors P and Q similarly. The Value Function is given by: $W_n(\mathbf{v}_{nt}, \mathbf{s}_t; \sigma_{-n}) =$

$$\max_{\mathbf{b}} \left\{ \Gamma_n(\mathbf{b}; \sigma_{-n})^T (\mathbf{v}_n - \tilde{\mathbf{b}}(\mathbf{b}; \mathbf{s}_t)) + P_n(\mathbf{b}; \sigma_{-n})^T \Pi_n(\mathbf{s}_t) + \beta Q_n(\mathbf{b}; \sigma_{-n})^T V_n(\mathbf{s}_t; \sigma_{-n}) \right\} \quad (9)$$

Where element c of V_n is $V_{nc}(\mathbf{s}_t; \sigma_{-n}) = \int_{\mathbf{s}} \int_{\mathbf{v}} W_n(\mathbf{v}, \mathbf{s}; \sigma_{-n}) dF(\mathbf{v}|\mathbf{s}) d\Gamma(\mathbf{s}|\mathbf{s}_t^c)$. $\tilde{\mathbf{b}}(\mathbf{b}; \mathbf{s}_t)$ gives the expected second highest bid, given that b_{nlt} is the highest. Since the cdf of the highest rival bids is $\Gamma_l(x|\mathbf{s})$, we can write $\Gamma_l(b_l|\mathbf{s})\tilde{b}_l(\mathbf{b}; \mathbf{s}) = \int_{b_l}^{b_{nlt}} \bar{b}_l \nabla_{b_l} \Gamma_l(\bar{b}_l|\mathbf{s}) d\bar{b}_l$.

D.1.2 First Order Conditions and Inverse Bid System

Rewrite the maximand: $\Gamma(\mathbf{b}|\mathbf{s})^T \mathbf{v} - \sum_l \int_{b_l}^{b_{nlt}} \bar{b}_l \nabla_{b_l} \Gamma_l(\bar{b}_l|\mathbf{s}) d\bar{b}_l + P(\mathbf{b}|\mathbf{s})\Pi(\mathbf{s}) + \beta Q(\mathbf{b}|\mathbf{s})V(\mathbf{s})$

Differentiate for FOCs: $0 = \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})(\mathbf{v} - \mathbf{b}^*) + \nabla_{\mathbf{b}} P(\mathbf{b}^*|\mathbf{s})\Pi(\mathbf{s}) + \beta \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})V(\mathbf{s})$.⁴⁰

We then invert the FOCs for the inverse bid system:

$$\boldsymbol{\xi}(\mathbf{b}_{nt}|\Pi, \beta V; \mathbf{s}) = \mathbf{b}_{nt} - \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} [\nabla_{\mathbf{b}} P(\mathbf{b}^*|\mathbf{s}) B_{\mathbf{s}} \boldsymbol{\pi} + \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s}) A_{\mathbf{s}} \beta \mathbf{V}]$$

This is similar to the inverse bid system presented in text, omitting the mark-up term $\nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \Gamma(\mathbf{b}^*|\mathbf{s})$. Consequently, conditional on $\boldsymbol{\pi}$ and $\beta \mathbf{V}$, the distribution of lot specific values F is point identified from the empirical quantiles of $\boldsymbol{\xi}(\mathbf{b}_{nt}|\Pi, \beta V; \mathbf{s})$.

⁴⁰This condition can equivalently be derived by requiring that, at the optimum, b_{lt}^* equals the marginal expected pay-off from winning lot l , conditional on bids for lots $m \neq l$.

D.1.3 Extension of Proposition 3

I now extend Proposition 3 to the second price case:

$$\begin{aligned}\tilde{\Pi}(\mathbf{b}^*|\mathbf{v}; \mathbf{s}) = & \Gamma(\mathbf{b}^*|\mathbf{s})^T (\mathbf{b}^* - \tilde{\mathbf{b}}(\mathbf{b}^*; \mathbf{s})) \\ & + [P(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} P(\mathbf{b}^*|\mathbf{s})] B_{\mathbf{s}} \boldsymbol{\pi} \\ & + [Q(\mathbf{b}^*|\mathbf{s})^T - \Gamma(\mathbf{b}^*|\mathbf{s})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})] A_{\mathbf{s}} \beta \mathbf{V}\end{aligned}$$

This is similar to the expression given in Proposition 3, except that the optimal lot specific surplus term is given by $\mathbf{b}^* - \tilde{\mathbf{b}}(\mathbf{b}^*; \mathbf{s})$ instead of $\nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \Gamma(\mathbf{b}^*|\mathbf{s})$. Proof is omitted due to its simplicity — simply substitute the inverse bid function $\boldsymbol{\xi}(\mathbf{b}_{nt}|\Pi, \beta V; \mathbf{s})$ for \mathbf{v} into the maximand of the value function in equation 9.

Employing the identity $P(\mathbf{b}|\mathbf{s})^T B_{\mathbf{s}} = Q(\mathbf{b}|\mathbf{s})^T A_{\mathbf{s}} C$, and taking an expectation over the observed bids, write the ex-ante value function as:

$$V^e(\mathbf{s}) = \tilde{\Phi}(\mathbf{s}) + \Omega(\mathbf{s})[C\boldsymbol{\pi} + \beta\mathbf{V}] \quad \text{Where } \tilde{\Phi}(\mathbf{s}) = E_{\mathbf{b}}[\Gamma(\mathbf{b}^*|\mathbf{s})^T (\mathbf{b}^* - \tilde{\mathbf{b}}(\mathbf{b}^*; \mathbf{s}))|\mathbf{s}]$$

And $\Omega(\mathbf{s})$ was defined in the text. Stack this equation over \mathbf{s} for: $\mathbf{V} = T\tilde{\Phi} + T\Omega[C\boldsymbol{\pi} + \beta\mathbf{V}]$
Which we can invert for: $\mathbf{V} = (I_S - \beta T\Omega)^{-1}[T\tilde{\Phi} + T\Omega C\boldsymbol{\pi}]$.

D.1.4 Identification

As in the main text I impose the mean zero property of \mathbf{v} for:

$$\begin{aligned}0 = E_{\mathbf{b}^*}[\boldsymbol{\xi}(\mathbf{b}^*; \mathbf{s}, (\boldsymbol{\pi}, \mathbf{V}))|\mathbf{s}] &= E_{\mathbf{b}^*}[\mathbf{b}^*|\mathbf{s}] - E_{\mathbf{b}^*}[\nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbf{s})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbf{s})|\mathbf{s}] A_{\mathbf{s}} [C\boldsymbol{\pi} + \beta\mathbf{V}] \\ &= \tilde{\Upsilon}(\mathbf{s}) - \Psi(\mathbf{s})[C\boldsymbol{\pi} + \beta\mathbf{V}]\end{aligned}$$

Stack over \mathbf{s} and substitute in \mathbf{V} for: $0 = \tilde{\Upsilon} - \beta\Psi(I_S - \beta T\Omega)^{-1}T\tilde{\Phi} - \Psi(I_S - \beta T\Omega)^{-1}C\boldsymbol{\pi}$. There is a unique solution to this system ($\boldsymbol{\pi}$ is point identified) if and only if the $LS \times S_n$ matrix $\Psi(I_S - \beta T\Omega)^{-1}C$ has rank $S_n - 1$. This matrix is the same as in the main text. Proposition 4 holds in this case as well, ensuring the rank condition.

D.2 Binding Reservation Prices

I now introduce binding reservation prices. A reservation price is binding if, in equilibrium, there is non-zero probability of winning a lot at the reservation price. This also extends to

endogenous entry with zero entry costs — where reservation prices are necessary to prevent arbitrarily low bids. Binding reservation prices do not pose a substantive problem, though do introduce additional mathematical complexity.

In the presence of reservation prices a bidder with a low value may choose not to bid strictly above the reservation price. This results in corner solutions as bids clump at the reservation price. We lose point identification as the FOCs no longer point identify \mathbf{v}_n . This is a problem, even in a single object context.

The identification argument presented below diverges from the argument presented in 3. Instead, it is closer to the estimation method presented in Section 4. Identification is demonstrated in an additional step. First I show that F is (partially) identified conditional on (Π, V, β) , but in particular it is partially identified conditional on $J + \beta V$. I then show that the object $\pi(\mathbf{s}_n) + \beta V(\mathbf{s})$ is partially identified, (for some \mathbf{s}_n it is only bounded). This is shown using quantile moment conditions: Instead finding the $\pi + \beta V$ such that $E[\xi(\mathbf{b}; \mathbf{s}, \pi + \beta V)|\mathbf{s}] = 0$ I find it such that $P(\xi_l(\mathbf{b}; \mathbf{s}, \pi + \beta V) \leq 0|\mathbf{s}) = 0.5$, imposing a zero conditional median assumption. Finally, I show that conditional on the identification of F and $\pi + \beta V$, V is identified, and hence π can be backed out given an assumption about β .

D.2.1 Changes to the Model

Denote the reservation price as R , which may vary across lots, bidders, and time. Denote player n 's entry decisions by \mathbf{d}_{nt} with entry $d_{ntl} = 1$ if they enter lot l , and zero otherwise. Adjust G, Γ, P and Q to be functions of bids and entry — if a player does not enter a lot, they lose that lot with probability 1. Identification requires an additional assumption:

Assumption 8. $\frac{\partial \Gamma_{nl}(\mathbf{b}_n, \mathbf{d}_n|\mathbf{s})}{\partial b_{nm}} = 0$ for $m \neq l$

I assume the probability an individual wins any given lot, conditional on \mathbf{s} and σ_{-n} , only depends on their bid for that lot. This implies $\nabla \Gamma_n(\mathbf{b}_n, \mathbf{d}_n|\mathbf{s})$ is a diagonal matrix. This assumption was not previously necessary for identification. If ties happen with zero probability or if tie breaking is exogenous, then this assumption will hold.⁴¹ Finally, I assume

⁴¹For mathematical convenience I assume ties occur in equilibrium with zero probability. The argument below can be easily extended to allow for ties at the reservation price. All that changes is that it introduces a discontinuity in the inverse bidding system at the reservation price, so that as the bidder goes from bidding the reserve to just above it, their payoff changes discontinuously. This slightly changes how we identify F , as we must essentially introduce an additional discrete choice of whether the bidder bids the reservation compared to bidding just above it. This additional discrete choice then restores the (upper-hemi) continuity of equilibrium, payoffs.

the lot specific values have zero conditional median, replacing the previous zero conditional mean assumption. I am then able to prove the following:

Proposition 6. *Given assumption 2, 3, 4, 5, and 8, both $F_n(\cdot|\mathbf{s})$ and $\bar{\kappa}_n(\mathbf{s})$ are non-parametrically partially identified. $\kappa(\mathbf{s}^c)$ is point identified if we observe the individual bidding $b > R$ on a lot that may yield pay-off $\kappa(\mathbf{s}^c)$.*

That is, we will point identify the truncated distribution $F_n(\cdot|\mathbf{v} \geq A_1; \mathbf{s})$, as well as the objects $F_n(A_1; \mathbf{s}) - F_n(A_2; \mathbf{s})$ and $F_n(A_2; \mathbf{s})$ for known A_1, A_2 .

While I assume players play pure strategies conditional on entry, I allow for the possibility that players play mixed strategies in their entry decisions. We use bidders' entry decisions to bound the payoffs of unentered auctions, exploiting that, at the equilibrium mixing strategy, players can not *strictly* prefer to enter any other combination of auctions.

D.2.2 Identification of F , conditional on K .

Under assumptions 2 - 5, and 8, and conditional on κ being point identified, the cdf F is non-parametrically partially identified. Similar to case 6.3.1.2 described in Athey and Haile (2007), we invert observed bids such that $b_l > R$, point identifying v_l . For bids at the reservation price and for non-entered auctions we can then find bounds on v_l .

First, reformulate the problem to include entry decisions. The player's problem is to decide which auctions to enter (\mathbf{d}), then set their bids (\mathbf{b}) to maximise payoffs, subject to their bids being weakly above reservation prices. The Lagrangian and corresponding FOCs for this problem, conditional on entry \mathbf{d}^* , is given as:

$$\begin{aligned} L(\mathbf{b}, \mathbf{d}^*, \mathbf{v}, \boldsymbol{\lambda}|\mathbf{s}) &= \Gamma(\mathbf{b}, \mathbf{d}^*|\mathbf{s})^T(\mathbf{v} - \mathbf{b}) + P(\mathbf{b}, \mathbf{d}^*|\mathbf{s})^T \bar{\kappa} + \boldsymbol{\lambda}^T(\mathbf{b} - R) \\ 0 &= \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})(\mathbf{v} - \mathbf{b}^*) - \Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) + \nabla_{\mathbf{b}} P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T \bar{\kappa} + \boldsymbol{\lambda}^* \end{aligned}$$

Entry ll of $\nabla_{\mathbf{b}} \Gamma(\mathbf{b}, \mathbf{d}|\mathbf{s})$ and entry lc of $\nabla_{\mathbf{b}} P(\mathbf{b}, \mathbf{d}|\mathbf{s})$ are as they were in section 3 if $d_l = 1$, and normalised to 0 otherwise. Rearrange this equation for:

$$\begin{aligned} \boldsymbol{\xi}(\mathbf{b}^*, \mathbf{d}^*, \boldsymbol{\lambda}|\bar{\kappa}; \mathbf{s}) &= \mathbf{b}^* + \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^{-1} [\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - \nabla_{\mathbf{b}} P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) \bar{\kappa}] \\ &\quad - \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^{-1} [\boldsymbol{\lambda}^*] \end{aligned}$$

At the true κ we have $\xi_l(\mathbf{b}^*, \mathbf{d}^*, \boldsymbol{\lambda}^*|\kappa; \mathbf{s}) = v_l$. But we do not observe $\boldsymbol{\lambda}^*$. Therefore, define $\boldsymbol{\xi}(\mathbf{b}^*, \mathbf{d}^*|\kappa; \mathbf{s}) = \mathbf{b}^* + \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^{-1} [\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - \nabla_{\mathbf{b}} P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) \bar{\kappa}]$. Next, I consider what

can be inferred for each of the four possible entry/bidding possibilities: *i*) $b_l > R$, *ii*) $b_l = R$, *iii*) $d_l = 0$, and the null case $l \notin \mathbb{L}$.

i) l such that $b_l^* > R$: Any entry l such that $b_l^* > R_l$, $\lambda_l^* = 0$. By Assumption 8, entry l of $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^{-1}[\boldsymbol{\lambda}^*]$ equals zero, and $\xi_l(\mathbf{b}^*, \mathbf{d}^*|\kappa; \mathbf{s}) = v_l$ is point identified.

ii) l such that $b_l^* = R$: For entry l with $b_l^* = R_l$, $\lambda_l^* > 0$. From Assumption 8 entry l of $\nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^{-1}[\boldsymbol{\lambda}^*]$ is greater than zero, and we attain the following bound:

$v_l \leq \xi_l(\mathbf{b}^*, \mathbf{d}^*|\kappa; \mathbf{s}) = R_l + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^{-1}[\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})\bar{\kappa}]_l$ (For vector M , $[M]_l$ denotes row l). As $(\mathbf{b}^*, \mathbf{d}^*)$ maximises expected payoffs, payoffs are (weakly) higher from playing $(\mathbf{b}^*, \mathbf{d}^*)$ than not entering auction l , playing $(\mathbf{b}^{l-}, \mathbf{d}^{l-})$ (the only difference between these actions is that $d_l^{l-} = 0$). Therefore:

$\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T(\mathbf{v} - \mathbf{b}^*) + P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T\bar{\kappa} \geq \Gamma(\mathbf{b}^{l-}, \mathbf{d}^{l-}|\mathbf{s})^T(\mathbf{v} - \mathbf{b}^{l-}) + P(\mathbf{b}^{l-}, \mathbf{d}^{l-}|\mathbf{s})^T\bar{\kappa}$. This rearranges for: $v_l \geq R_l - \frac{1}{\Gamma_l(b_l^*, d_l^*|\mathbf{s})}[P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - P(\mathbf{b}^{l-}, \mathbf{d}^{l-}|\mathbf{s})]^T\bar{\kappa}$.

iii) l such that $d_l^* = 0$: Consider l such that $d_l = 0$. They must attain a greater payoff from $d_l = 0$ than from bidding the reservation price. Consider alternate action $(\mathbf{b}^{l+}, \mathbf{d}^{l+})$ where the only difference between this and $(\mathbf{b}^*, \mathbf{d}^*)$ is that $b_l^{l+} = R_l$ and $d_l^{l+} = 1$. Therefore: $\Gamma(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T(\mathbf{v} - \mathbf{b}^*) + P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s})^T\bar{\kappa} \geq \Gamma(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})^T(\mathbf{v} - \mathbf{b}^{l+}) + P(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})^T\bar{\kappa}$ Rearranging this for: $v_l < R_l - \frac{1}{\Gamma_l(b_l^{l+}, d_l^{l+}|\mathbf{s})}[P(\mathbf{b}^*, \mathbf{d}^*|\mathbf{s}) - P(\mathbf{b}^{l+}, \mathbf{d}^{l+}|\mathbf{s})]^T\bar{\kappa}$

D.2.3 Identification of κ under binding reservation prices

Under assumptions 2 - 5, and 8, the function κ is partially identified up to standard normalisations. $\kappa(\bar{\mathbf{s}})$ is point identified at $\mathbf{s} = \bar{\mathbf{s}}$ if we observe bidding strictly above R on a combination of goods that would have the outcome $\mathbf{s}^c = \bar{\mathbf{s}}$. I prove this by exploiting multiple observations for every state to establish a necessary rank condition, similar to the one presented in Section 3. Whereas the previous proof employed a condition on the mean of $\boldsymbol{\xi}(\mathbf{b}, \mathbf{d})$, this proof employs a condition on the marginal quantiles of $\boldsymbol{\xi}(\mathbf{b}, \mathbf{d})$. I set $\kappa(\mathbf{s})$ such that the median (or some other quantile) is equal to zero.

Binding reservation prices cause our FOCs to break down, so that at the true $\boldsymbol{\kappa}$ ($= C\boldsymbol{\pi} + \beta\mathbf{V}$) we can only write:

$$\mathbf{v} \leq \boldsymbol{\xi}(\mathbf{b}, \mathbf{d}|\kappa; \mathbf{s}) = \mathbf{b} + \nabla_{\mathbf{b}}\Gamma(\mathbf{b}, \mathbf{d}|\mathbf{s})^{-1}[\Gamma(\mathbf{b}, \mathbf{d}|\mathbf{s}) - \nabla_{\mathbf{b}}P(\mathbf{b}, \mathbf{d}|\mathbf{s})A_{\mathbf{s}}\boldsymbol{\kappa}]$$

Which only holds with equality for rows l with $b_l > R$. Stack these over \mathbf{s} for:

$$\underbrace{\underline{\mathbf{v}}}_{LS \times 1} \leq \underbrace{\underline{\boldsymbol{\xi}}(\underline{\mathbf{b}}, \underline{\mathbf{d}}|\kappa)}_{LS \times 1} = \underbrace{\underline{\mathbf{b}}}_{LS \times 1} + \underbrace{\nabla_{\underline{\mathbf{b}}} \Gamma(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1}}_{LS \times LS} \underbrace{[\Gamma(\underline{\mathbf{b}}, \underline{\mathbf{d}})]}_{LS \times 1} - \underbrace{\nabla_{\underline{\mathbf{b}}} P(\underline{\mathbf{b}}, \underline{\mathbf{d}})}_{LS \times S} \kappa \quad (10)$$

$$\underline{\boldsymbol{\xi}}(\underline{\mathbf{b}}, \underline{\mathbf{d}}|\kappa) = \begin{pmatrix} \boldsymbol{\xi}(\mathbf{b}_1, \mathbf{d}_1|\kappa; \mathbf{s}_1) \\ \vdots \\ \boldsymbol{\xi}(\mathbf{b}_S, \mathbf{d}_S|\kappa; \mathbf{s}_S) \end{pmatrix} \quad \underline{\mathbf{b}} = \begin{pmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_S \end{pmatrix}$$

$$\Gamma(\underline{\mathbf{b}}, \underline{\mathbf{d}}) = \begin{pmatrix} \Gamma(\mathbf{b}_1, \mathbf{d}_1|\mathbf{s}_1) \\ \vdots \\ \Gamma(\mathbf{b}_S, \mathbf{d}_S|\mathbf{s}_S) \end{pmatrix} \quad \nabla_{\underline{\mathbf{b}}} P(\underline{\mathbf{b}}, \underline{\mathbf{d}}) = \begin{pmatrix} \nabla_{\mathbf{b}} P(\mathbf{b}_1, \mathbf{d}_1|\mathbf{s}_1) A_{\mathbf{s}_1} \\ \vdots \\ \nabla_{\mathbf{b}} P(\mathbf{b}_S, \mathbf{d}_S|\mathbf{s}_S) A_{\mathbf{s}_S} \end{pmatrix}$$

I require a rank condition on $\nabla_{\underline{\mathbf{b}}} \Gamma(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1} \nabla_{\underline{\mathbf{b}}} P(\underline{\mathbf{b}}, \underline{\mathbf{d}})$. If this has full rank then each $\underline{\boldsymbol{\xi}}$ implies a unique κ , so that if I observed just one observation of $\underline{\mathbf{v}}$ I could solve for κ . Note that $E[\nabla_{\underline{\mathbf{b}}} \Gamma(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1} \nabla_{\underline{\mathbf{b}}} P(\underline{\mathbf{b}}, \underline{\mathbf{d}})] = \Psi$, the matrix presented in text. Importantly, the proof presented in B.1, that $Rank(\Psi) = S - S^C - |\min(\mathbf{S})|$ extends trivially to $\nabla_{\underline{\mathbf{b}}} \Gamma(\underline{\mathbf{b}}, \underline{\mathbf{d}})^{-1} \nabla_{\underline{\mathbf{b}}} P(\underline{\mathbf{b}}, \underline{\mathbf{d}})$. The proof never exploited the fact we had taken an expectation, and entirely used the partial ordering structure of the state space.

With binding reservation prices and entry, certain states may never be outcomes that *could have* occurred with positive probability, so the corresponding elements of κ are not point identified. These entries of κ do not appear in the above equation, having a coefficient of zero. These states will only be partially identified.

Next, fix an $LS \times 1$ vector of probabilities \mathbf{p} . By definition of the marginal CDF:

$$\begin{pmatrix} p_1 \\ \vdots \\ p_{LS} \end{pmatrix} = \begin{pmatrix} F_1(\tilde{v}_1|\mathbf{s}_1) \\ \vdots \\ F_L(\tilde{v}_{LS}|\mathbf{s}_S) \end{pmatrix} = \begin{pmatrix} E_{v_1}[\mathbb{I}[v_1 \leq \tilde{v}_1] | \mathbf{s}_1] \\ \vdots \\ E_{v_L}[\mathbb{I}[v_L \leq \tilde{v}_{LS}] | \mathbf{s}_S] \end{pmatrix}$$

Employ a change of variables, taking expectations over the observed random variables (\mathbf{b}, \mathbf{d}) instead of v_l . This change is only valid for state-lot combinations such that when $v_l = \tilde{v}_l$, $b_l > R$, when $\xi_l(\mathbf{b}, \mathbf{d}; \kappa) = v_l$ holds with equality and the mapping from \mathbf{b} to v_l is continuous, smooth, and monotonic.⁴² Drop rows where this condition fails, as we lose identifiability

⁴²This is essentially an application of the Law of the Unconscious Statistician. Monotonicity of the inverse

of corresponding elements of κ . If, even when v_l is as large as \tilde{v}_l , the elements of $\bar{\kappa}(\mathbf{s})$ corresponding to winning lot l are so small that they never bid strictly above R on lot l , these elements of $K(\mathbf{s})$ are not identified. This yields:

$$\mathbf{p} = \begin{pmatrix} E_{v_1}[\mathbb{I}[v_1 \leq \tilde{v}_1] | \mathbf{s}_1] \\ \vdots \\ E_{v_L}[\mathbb{I}[v_L \leq \tilde{v}_L] | \mathbf{s}_L] \end{pmatrix} = \begin{pmatrix} E_{\mathbf{b}, \mathbf{d}}[\mathbb{I}[\xi_1(\mathbf{b}_1, \mathbf{d}_1; \kappa) \leq \tilde{v}_1] | \mathbf{s}_1] \\ \vdots \\ E_{\mathbf{b}, \mathbf{d}}[\mathbb{I}[\xi_L(\mathbf{b}_L, \mathbf{d}_L; \kappa) \leq \tilde{v}_L] | \mathbf{s}_L] \end{pmatrix}$$

Proving point identification of κ requires we show that the \mathbf{p} th quantiles of $\underline{\xi}(\underline{\mathbf{B}}, \underline{\mathbf{D}} | \kappa)$ equals $\tilde{\mathbf{v}}$ only at the true κ . But, from our rank condition, a unique $\underline{\xi}(\underline{\mathbf{B}}, \underline{\mathbf{D}} | \kappa)$ implies a unique κ . Therefore, only a unique κ is such that the \mathbf{p} th quantiles of $\underline{\xi}(\underline{\mathbf{B}}, \underline{\mathbf{D}} | \kappa)$ equals $\tilde{\mathbf{v}}$. Therefore, there exists a unique κ such that this equation holds.⁴³

D.2.4 Identification of π and βV

I have proven the non-parametric (partial) identification of F_n and $\bar{\kappa}_n = \Pi_n + \beta V_n$. I also previously established that the ex-ante value function is known function of beliefs, F_n , and $\bar{\kappa}_n = \Pi_n + \beta V_n$, all of which are identified. Therefore so too is the ex-ante value function. The continuation value V is then a function of the ex-ante value function and the transition process, both of which I established are identified. Finally, fixing β , Π_n is a function of $\bar{\kappa}_n$, β , and V_n , ensuring that Π_n is also non-parametrically partially identified.⁴⁴

D.3 Endogenous Participation

In this Appendix I introduce endogenous entry in which entry is costly and v_{nlt} is not observed before entry, though I assume that the entry decisions of other players is observed before

bid function for bids strictly above the reservation price is discussed in A.

⁴³It should be noted that κ is unique up to $|\tilde{min}(\mathbb{S})| + S^C$ elements of κ that must be normalised due to the rank deficiency of the matrix Ψ . These elements are the entries associate with states $\mathbf{s} \in \tilde{min}(\mathbb{S})$ that are never observed as possible ex-post states, and one additional state from each component - associated with $\mathbf{s}_n = \mathbf{s}_n^1$. We will see in Appendix D.2.4 that these normalisations do not impact the identification of π .

⁴⁴If there are values of $\Pi_n(\mathbf{s}_t) + \beta V_n(\mathbf{s}_t)$ that are only bid on at the reservation price, then the value function is only partially identified. However, this non-identified region will generally be very small. Likewise, elements of κ corresponding to states which never appear as possible ex-post states are zeroed out in this equation, so it does not matter how they are normalised. Finally, the normalised elements corresponding to one (minimal, with $\mathbf{s}_n = \mathbf{s}_n^0$) element from each component $\mathbb{S}^C \subset \mathbb{S}$. These normalisations constitute location shifts of Π for all elements in that component, as we essentially made the normalisations because only marginal payoffs are identified. Finally, when we back out π , we will normalise $\pi(\mathbf{s}_n^0) = 0$, in line with these location normalisations.

bidding.⁴⁵ I focus on the case with non-binding reservation prices, though it will be clear how the results from Appendix D.2 extend to this case. I also discuss the effect of incorrectly assuming participation decisions are exogenously determined, as I do in my application in Section 5. I do not explicitly discuss estimation of the participation process as it is clear how this follows from the identification argument. In particular, the econometrician can apply the estimation procedure presented in Section 4, but estimate the entry costs using a GMM procedure after they estimate the ex-ante value function (conditional on the set of auctions entered), before using the estimated entry costs to evaluate the continuation value and proceeding as normal.

The identification argument presents a minor generalisation on the one presented in the main text. The argument proceeds as follows: F is non-parametrically point identified conditional on $\kappa = C\pi + \beta\mathbf{V}$. As in the previous Appendix, κ remains non-parametrically identified conditional on the identification of Γ and P using observed variation in \mathbf{s} , relying on our rank condition on the matrix Ψ . Given identification of κ, Γ , and P , Proposition 3 ensures that the expected payoff from each entry structure is also non-parametrically identified. Given these expected payoffs, the entry problem is then a multinomial discrete choice problem, so I rely on standard results for the identification of entry costs. Identification of expected entry payoffs and costs ensures the ex-ante value function, and hence the continuation value \mathbf{V} , is identified, thereby identifying $\pi = C^{-1}(\kappa - \beta\mathbf{V})$.

I proceed as follows: In Appendix D.3.1 I introduce changes to the main model, and demonstrate that the previous identification results for F and κ also apply. In Appendix D.3.2 I show that the distribution of entry costs is non-parametrically identified, and finally that \mathbf{V} , and hence π are also identified.

D.3.1 Changes to the Model

All objects below are functions of the state \mathbf{s} . Conditional on an entry structure \mathbb{D} and having observed the lot specific values \mathbf{v} the agent places bids to maximise the following:

$$\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D}) = \Gamma(\mathbf{b}|\mathbb{D})^T(\mathbf{v} - \mathbf{b}) + P(\mathbf{b}|\mathbb{D})^T\Pi + Q(\mathbf{b}|\mathbb{D})^T\beta\mathbf{V}$$

Given the agent's behaviour conditional on entry, the agent's problem is to choose an

⁴⁵Allowing the 'entry structure' to be unknown before bidding does not change anything substantive. We simply alter the objects Γ, P and Q to additionally take an expectation over the entry decisions of other players.

entry structure \mathbb{D}_n to maximise their expected pay-off. I assume that agent's entry costs, a $2^L \times 1$ vector \mathbf{c} , are drawn independently and privately from $C(.|\mathbf{s}_n)$ (independent of \mathbf{s}_{-n}). I assume that C is common knowledge.

The agent observes \mathbf{s} and, given knowledge of F and $\boldsymbol{\kappa}$ and their equilibrium beliefs, maximises an expected pay-off associated with any given entry structure:

$$W(\mathbb{D}_n|\mathbf{c}) = E_{\mathbb{D}_{-n}}[E_{\mathbf{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D})\} | \mathbb{D}_n] - c_{\mathbb{D}_n}]$$

The continuation value associated with ending the period in state \mathbf{s}^c is then:

$$V(\mathbf{s}^c) = E_{\mathbf{s}}[E_{\mathbf{c}}[\max_{\mathbb{D}_i} \{W(\mathbb{D}_n|\mathbf{c})\} | \mathbf{s} | \mathbf{s}^c]]$$

Identification of F conditional on the identification of $\bar{\kappa}$

The Inverse Bid System, as given in equation 4, where the state variable has simply been augmented to include the entry structure. Hence F remains non-parametrically identified conditional on the identification of Γ , Q , and $\boldsymbol{\kappa}$.

Identification of k

As in the main text, we can take a conditional expectation of the inverse bid system, setting this equal to zero: $E[\boldsymbol{\xi}|\mathbf{s}, \mathbb{D}] = 0$. We can then again stack this system of equations across states and entry structures for $0 = \Upsilon - \Psi\boldsymbol{\kappa}$. Non-parametric point identification of $\boldsymbol{\kappa}$ then requires the same rank condition on Ψ proven previously.⁴⁶

Identification of $E_{\mathbf{v}}[\tilde{\Pi}(\mathbf{b}^*|\mathbf{v}; \mathbf{s}, \mathbb{D})]$

Recognise that Proposition 3 continues to hold, and so we can write the expected maximised payoff, conditional on \mathbb{D} , as

$$\begin{aligned} \bar{\Pi}(\mathbf{s}, \mathbb{D}) &= E_{\mathbf{v}}[\tilde{\Pi}(\mathbf{b}^*|\mathbf{v}; \mathbb{D})] = \Gamma(\mathbf{b}^*|\mathbb{D})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbb{D})^{-1} \Gamma(\mathbf{b}^*|\mathbb{D}) \\ &\quad + [Q(\mathbf{b}^*|\mathbb{D})^T - \Gamma(\mathbf{b}^*|\mathbb{D})^T \nabla_{\mathbf{b}} \Gamma(\mathbf{b}^*|\mathbb{D})^{-1} \nabla_{\mathbf{b}} Q(\mathbf{b}^*|\mathbb{D})] A_{\mathbf{s}} \beta \boldsymbol{\kappa} \end{aligned}$$

⁴⁶We normalise elements of $\boldsymbol{\kappa}$ corresponding to states which are either the minimal element of their component, or never appear as possible ex-post states. By definition, there will be $S^C + |\min(\mathbb{S})|$ of these. In Appendix B.1 we found previously that Ψ has rank $S - S^C - |\min(\mathbb{S})|$.

D.3.2 Identification of C

At the entry stage, the agent sets their entry structure \mathbb{D}_i such that:

$$E_{\mathbb{D}_{-n}}[E_{\mathbf{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D})\} | \mathbb{D}_n] - c_{\mathbb{D}_n}] \geq \max_{\mathbb{D}'_n \neq \mathbb{D}_n} \left\{ E_{\mathbb{D}_{-n}}[E_{\mathbf{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D})\} | \mathbb{D}'_n] - c_{\mathbb{D}'_n}] \right\}$$

Similar to how we identify G , because we observe entry decisions, we therefore observe the equilibrium distribution of \mathbb{D}_n for all n . Therefore, following from the above, $E_{\mathbb{D}_{-n}}[E_{\mathbf{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D})\} | \mathbb{D}_n]]$ is non-parametrically point identified. Normalising that the entry cost of entering zero auctions is always zero, I now exploit the exclusion restriction that the distribution of \mathbf{c} is independent of \mathbf{s}_{-n} . Therefore, variation in \mathbf{s}_{-n} leads to known variation in $E_{\mathbb{D}_{-n}}[E_{\mathbf{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D})\} | \mathbb{D}_n]]$, thereby tracing out the distribution $C(\cdot | \mathbf{s}_n)$, ensuring we have non-parametric identification.⁴⁷

The ex-ante value function $V^e(\mathbf{s}) = E[\max_{\mathbb{D}_n} \{E_{\mathbb{D}_{-n}}[E_{\mathbf{v}}[\max_{\mathbf{b}} \{\Pi(\mathbf{b}|\mathbf{v}; \mathbb{D})\} | \mathbb{D}_n]] - c_{\mathbb{D}_n}\}]$, and hence the continuation value $V(\mathbf{s})$ are then also non-parametrically identified, which in turn yields identification of the flow payoff function π .

D.3.3 Incorrectly Assuming Exogenous Participation

I now consider what object is identified and estimated when the econometrician incorrectly assumes that firms participation decisions are exogenously determined, then applies the estimation procedure presented in section 4.

Recognise that after the second estimation step the econometrician correctly recovers both \hat{F} and $\hat{\kappa}$. Then, when they attempt to evaluate the ex-ante value function by taking an expectation over bids, they incorrectly take an expectation over the participation decisions as well. That is, they attempt to form $\hat{V}^e(\mathbf{s}) = \int_{\mathbf{b}} \bar{W}(\mathbf{b}|\mathbf{s}; \bar{\kappa}) dG(\mathbf{b}|\mathbf{s})$, but in truth evaluate:

$$\hat{V}^e(\mathbf{s}) = \sum_{\mathbb{D}_i} Prob(\mathbb{D}_i | \mathbf{s}) \int_{\mathbf{b}} \bar{W}(\mathbf{b}|\mathbf{s}, \mathbb{D}_i; \bar{\kappa}) dG(\mathbf{b}|\mathbf{s}, \mathbb{D}_i)$$

$Prob(\mathbb{D}_i | \mathbf{s})$ gives the conditional choice probability of choosing entry structure \mathbb{D}_i . This is the true ex-ante value function minus expected participation costs.

Therefore, when they evaluate the continuation value they will instead recover $\hat{V}(\mathbf{s}^c) =$

⁴⁷Technically, identification is partial: The set of states is finite, so we will only actually be point identifying $C(\cdot | \mathbf{s}_n)$ at a finite set of points across its support. We can achieve full point identification either by assuming discrete support, or introducing one continuously varying element of \mathbf{s}_{-n} .

$E_{\mathbf{s}}[\hat{V}^e(\mathbf{s})|\mathbf{s}^c] = V(\mathbf{s}^c) - E_{\mathbf{s}}[E_{\mathbf{c}}[c_{\mathbb{D}_i^*}(\mathbf{c})|\mathbf{s}|\mathbf{s}^c]]$, that is the continuation value minus the conditional expected participation costs from the following period. Back out the flow payoff from the pseudo-payoffs they recover $\hat{\pi}(\mathbf{s}^c) = \kappa(\mathbf{s}^c) - \beta\hat{V}(\mathbf{s}^c) = \pi(\mathbf{s}^c) + E_{\mathbf{s}}[E_{\mathbf{c}}[c_{\mathbb{D}_i^*}(\mathbf{c})|\mathbf{s}|\mathbf{s}^c]]$. That is, the sum of the flow payoff and the expected future participation costs, which gives the misspecification bias. Recognise that even if the distribution of participation costs does not depend on the state, because payoffs depend on the state the optimal entry decision will depend on the state, and so the distribution of realised participation costs will depend on the state. Therefore, the econometrician correctly estimates flow payoffs only if participation costs are zero.

D.4 Stochastic Combination Value

I now present two identification results for the case when the combination value is stochastic, when $\pi(\mathbf{s})$ is not a function but a probability distribution. I focus on the static setting for two reasons. First, these results are novel even in the static case. Second, as we have seen throughout this paper, identification of the primitives of a generalised static model (where primitives are allowed to depend on \mathbf{s}_0 and \mathbf{s}_{-n}), is sufficient for identification of the primitives of a dynamic model. This is because identification of the Pseudo-Static payoff function κ implies identification of π .

I focus on two cases: First, when Π is a function of low-dimensional un-observables M , such as stocks, where $M \leq L$. Second, I consider a case when $M > L$, but elements of the unobservable vector are constant over time (e.g. constant parameters).

These extensions both centre on the theme of finding some way to reduce the dimensionality of the unknowns. The key idea is this: Each observation of bidding on an auction yields L pieces of information. Therefore, in order to have any hope at point identifying unobservables, there cannot be more than L unobservables. However, as in the main text, we can combine observations of bidding across period (or bidders) to identify unobservables that remain constant across the observations.

D.4.1 Case 1: Known function of low dimensional un-observables

Suppose the combinatorial value can be written as $\boldsymbol{\pi}(\mathbf{m}_t)$ where $\mathbf{m}_t \in \mathbb{M}$ is an unobserved (potentially) stochastic random variable of dimension $M \leq L$. I require that $\boldsymbol{\pi} : \mathbb{M} \rightarrow \mathbb{K}$ is a known function (with range $\mathbb{K} \subset \mathbb{R}^{2^L}$). Importantly, some elements of \mathbf{m} may represent fixed parameters associated with the functional form π .

Normalise the first element of this vector valued function (corresponding to player i losing every lot) to zero, so that I focus on the marginal combinatorial pay-off $\boldsymbol{\pi}(\mathbf{m})_{2:2^L} - \boldsymbol{\pi}(\mathbf{m})_1$. The expected payoff is $\Pi(\mathbf{b}) = P(\mathbf{b})^T \boldsymbol{\pi}(\mathbf{m}_t) - \Gamma(\mathbf{b})^T \mathbf{b}$. Necessary first order conditions are given by: $0 = \nabla_{\mathbf{b}} P(\mathbf{b}) \boldsymbol{\pi}(\mathbf{m}_t) - \nabla_{\mathbf{b}} \Gamma(\mathbf{b}) \mathbf{b} - \Gamma(\mathbf{b})$.

The problem is then to show \mathbf{m} is point identified. I make two assumptions about this function that are sufficient for \mathbf{m}_t to be point identified:

Assumption 9. *i) $\boldsymbol{\pi}(\mathbf{m})$ is continuous and continuously differentiable for all \mathbf{m}_t .*

ii) For any \mathbf{m} and \mathbf{m}' there exists a set $\mathbb{U} \subset \{1, 2, \dots, 2^L\}$ with $|\mathbb{U}| = M$ that defines the vector value function $\mathbf{F}^{\mathbb{U}}$ where $F_n^{\mathbb{U}}(\mathbf{m}) = \Pi_{U_n}(\mathbf{m})$ such that

$$(\mathbf{m} - \mathbf{m}')^T (\mathbf{F}^{\mathbb{U}}(\mathbf{m}) - \mathbf{F}^{\mathbb{U}}(\mathbf{m}')) > 0$$

The second part of this assumption is essentially an extension of strict monotonicity to the case of 2^L dimensional functions in M dimensional variables. The assumption states that for any two distinct \mathbf{m} and \mathbf{m}' we can find a set of rows of $\boldsymbol{\pi}(\cdot)$ such that this inner product is strictly positive.⁴⁸ A key result of this property is that the function $\boldsymbol{\pi}(\cdot)$ is a bijection: Each \mathbf{m} maps onto a unique $\boldsymbol{\pi}$, and the condition ensures that for any two distinct \mathbf{m} and \mathbf{m}' it must be the case that $\boldsymbol{\pi}(\mathbf{m}) \neq \boldsymbol{\pi}(\mathbf{m}')$ (since otherwise we could not find a \mathbb{U} such that $(\mathbf{m} - \mathbf{m}')^T (\mathbf{F}^{\mathbb{U}}(\mathbf{m}) - \mathbf{F}^{\mathbb{U}}(\mathbf{m}')) > 0$). This ensures that the inverse $\boldsymbol{\pi}^{-1}(\cdot)$ exists, such that for all $\mathbf{m} \in \mathbb{M}$ $\mathbf{m} = \boldsymbol{\pi}^{-1}(\boldsymbol{\pi}(\mathbf{m}))$. Furthermore, because $\boldsymbol{\pi}(\cdot)$ is continuous and continuously differentiable everywhere, so that $\boldsymbol{\pi}^{-1}(\cdot)$ must be differentiable everywhere, $\boldsymbol{\pi}^{-1}(\cdot)$ must also be continuous.

Proposition 7. *Under assumptions 2, 3, & 9, \mathbf{m}_t is identified up to normalisation.*

For example, if the second and third elements of \mathbf{m}_t are parameters describing the mean and standard deviation of m_{1t} , then \mathbf{m}_t is identified up to location and scale.

The proof requires arguing that with L equations in only M unknowns there exists a unique solution to the system. The proof proceeds by recognising that the set of vectors $\boldsymbol{\pi}$ which satisfy the FOCs is convex. From the continuity of the inverse function $\boldsymbol{\pi}^{-1}(\cdot)$ and the (generalised) intermediate value theorem, this implies that the set of \mathbf{m} for which the FOCs hold is path connected. So, there must be a point arbitrarily close to \mathbf{m}_t for which the FOCs hold. However, that $\nabla_{\mathbf{b}} P(\mathbf{b})$ has rank L and the function $\boldsymbol{\pi}(\cdot)$ is invertible implies the system is locally unique.

⁴⁸This property is satisfied when, for example, each element of J is weakly monotone in elements of \mathbf{m} , and strictly monotonic in at least one element.

Proof: 1. Consider the set of $2^L \times 1$ dimensional vectors which satisfy the system of equations $\nabla_{\mathbf{b}}P(\mathbf{b})\boldsymbol{\pi} = \nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b})$. This set, denoted $\tilde{\mathbb{K}}$, is convex, and hence path-connected, as for two vectors $\boldsymbol{\pi}, \boldsymbol{\pi}' \in \tilde{\mathbb{K}}$:

$$\begin{aligned}\lambda \nabla_{\mathbf{b}}P(\mathbf{b})\boldsymbol{\pi} + (1 - \lambda) \nabla_{\mathbf{b}}P(\mathbf{b})\boldsymbol{\pi}' &= (\lambda + (1 - \lambda))(\nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b})) \\ \therefore \nabla_{\mathbf{b}}P(\mathbf{b})(\lambda \boldsymbol{\pi} + (1 - \lambda)\boldsymbol{\pi}') &= \nabla_{\mathbf{b}}\Gamma(\mathbf{b})\mathbf{b} + \Gamma(\mathbf{b})\end{aligned}$$

2. This implies the image of the intersection of $\tilde{\mathbb{K}}$ and \mathbb{K} defined by the continuous function $\boldsymbol{\pi}^{-1}(\cdot)$ (the set of \mathbf{m} for which the FOCs hold) is also be path connected. This follows from the generalised intermediate value theorem, which states that for a continuous function $f : \mathbb{X} \rightarrow \mathbb{Y}$, if the set \mathbb{X} is path-connected, then so is the image $f(\mathbb{X})$.
3. If the intersection of $\tilde{\mathbb{K}}$ and \mathbb{K} contains more than one element, then for any \mathbf{m} which satisfies the FOCs, there is an arbitrarily nearby \mathbf{m}' which also satisfies the FOCs.
4. However, from the inverse function theorem, the FOCs are locally unique. The Jacobian of these FOCs, with respect to \mathbf{m} are given by:

$$\nabla_{\mathbf{b}}P(\mathbf{b})\nabla_{\mathbf{m}}\boldsymbol{\pi}(\mathbf{m})$$

This has rank M because $\nabla_{\mathbf{b}}P(\mathbf{b})$ has rank L (it consists of L linearly independent rows), and $\boldsymbol{\pi}(\mathbf{m})$ is invertible (so $\nabla_{\mathbf{m}}\boldsymbol{\pi}(\mathbf{m})$ has rank M). Therefore it is locally invertible, and so the set of \mathbf{m} which satisfy the FOCs contain a single element.

□

D.4.2 Case 2: When $M > L$

When $M > L$ we can combine information across observations, instead of identifying everything from a single observation, so long as *enough* elements of M are constant across observations. This is relevant when \mathbf{m}_t can be decomposed into $(\mathbf{m}_t^1, \mathbf{m}^0)$, where \mathbf{m}^0 are fixed parameters. Suppose $M \leq 2L$, and in particular, $|\mathbf{m}_t^1| < L$. Consider a pair of FOCs from two separate periods t_1 and t_2 . Importantly, I still impose assumption 9. Combine the

two sets of first order conditions as follows:

$$\begin{pmatrix} \nabla_{\mathbf{b}} P(\mathbf{b}_{t_1}) & 0 \\ 0 & \nabla_{\mathbf{b}} P(\mathbf{b}_{t_2}) \end{pmatrix} \begin{pmatrix} \boldsymbol{\pi}(\mathbf{m}_{t_1}) \\ \boldsymbol{\pi}(\mathbf{m}_{t_2}) \end{pmatrix} = \begin{pmatrix} \nabla_{\mathbf{b}} \Gamma(\mathbf{b}_{t_1}) \mathbf{b}_{t_1} + \Gamma(\mathbf{b}_{t_1}) \\ \nabla_{\mathbf{b}} \Gamma(\mathbf{b}_{t_2}) \mathbf{b}_{t_2} + \Gamma(\mathbf{b}_{t_2}) \end{pmatrix}$$

Uniqueness of the solution to this system follows the same logic as the previous proof with the added note that $\nabla_{(\mathbf{m}_{t_1}, \mathbf{m}_{t_2})} \begin{pmatrix} \boldsymbol{\pi}(\mathbf{m}_{t_1}) \\ \boldsymbol{\pi}(\mathbf{m}_{t_2}) \end{pmatrix}$ has rank $2|\mathbf{m}_t^1| + |\mathbf{m}^0|$, so that I can appeal to the inverse function theorem for local uniqueness.

This result allows us to add a large number of additional parameters to the function $\boldsymbol{\pi}(\cdot)$ which are identified by using variation across observations. This employs a similar philosophy used to prove the identification results in the main paper.

E Monte Carlo Simulation

I now present the results of a Monte-Carlo study evaluating the estimator proposed in 4. As discussed in GKS, the difficulty in simulating these games is that solving for equilibrium bidding strategies is intractable. Meanwhile, numerically finding equilibrium bidding strategies — by iterating over equilibrium beliefs and actions until a fixed point is found — is extremely computationally intensive.

For simplicity I focus on the case where bidders are bidding against a parametric set of beliefs. That is, I essentially take the equilibrium as given. Furthermore I focus on an equilibrium in which equilibrium beliefs do not depend on each bidder's individual states $\{\mathbf{s}_{nt}\}_{n \in \mathbb{N}}$. This is similar to many applications seen in practice, including GKS, Backus and Lewis (2016), Groeger (2014), Balat (2013).

E.0.1 Set up

Every period there are two auctions ($L = 2$) and two types of object, denoted x and y . Each lot contains one type of object, and one lot of each type of good is auctioned each period. Some lots contain 10 units of the good, while other lots contain only 5. The set of available lots is denoted (z^x, z^y) : Lot 1 contains z^x units of x , lot 2 contains z^y units of y . Therefore the possible characteristics of lots $\mathbb{X}_t = \{(5, 5), (10, 5), (5, 10)\}$ give the common state. For simplicity, this transitions stochastically where each states occurs with equal probability, independent of previous states.

States consist of bidders' stocks of the two objects, which come in integer values: $s_{nt}^x \in \{0, 1, \dots, 100\}$, likewise for good y . At the end of each period bidders consume 3 units of good x with probability 0.4 and three units of good y with probability 0.3, until their stocks fall to 0. A bidder's combinatorial flow pay-off is given by:

$$\pi(s^x, s^y) = \theta_1 \log(s^x + 1) + \theta_2 \log(s^x + 1) \log(s^y + 1)$$

Where (θ_1, θ_2) are parameters set to 20 and 10 respectively. θ_1 ensures pay-offs are not additively separable over time, while $\theta_2 > 0$ ensures the lots are complements. The lot-specific pay-offs are drawn from:

$$\mathbf{v}_{nt} \sim N \begin{pmatrix} 0 & 900z_t^x & 100z_t^x z_t^y \\ 0 & 100z_t^x z_t^y & 400z_t^y \end{pmatrix}$$

I take as given the equilibrium distribution of the highest rival bids, which follows a type 2 extreme value distribution. The mean of this distribution is given by the average (across states) marginal payoff from each lot ($\approx (17.1z^x, 12.5z^y)$). The variance is tuned to the variance (across states and lot-specific payoffs) of the marginal payoffs from winning each lot. The shape is set to 0.1.

I perform value function iteration to find the continuation value under this distribution of pay-offs and these equilibrium beliefs. Having found a continuation value, I can then simulate a dataset. Given the set-up the state space consists of 30,000 unique elements. Focusing on a large number of elements is intended to simulate my real world application when the state space will be treated as continuous.

I simulate 1,000 datasets, with $T \in \{300, 1000, 10,000\}$ observations uniformly sampled from the state space. I consider 3 estimators: 1) a semi-parametric estimator using the same functional form as π , 2) a quadratic polynomial, and 3) a semi-nonparametric cubic spline. For the spline, I use uniformly spaced knots, setting the knots to ensure at least \sqrt{T} observations per knot. For each estimator I consider estimates from using no instruments, the baseline "initial state" instruments, and all the possible ex-post states as instruments. The first stage is estimated using correctly specified maximum likelihood.

E.0.2 Results

Results are presented in figure 8. Each estimator yields estimates of $\hat{\pi}(\mathbf{s}_n)$ for each $\mathbf{s}_n \in \mathbb{S}_i$. I then fit the correctly specified π across these states, extracting $\hat{\theta}_1$ and $\hat{\theta}_2$.

The semi-nonparametric estimator (3) outperforms the two semi-parametric estimators, even in relatively small samples. However, it is very computationally intensive, with estimation taking almost 100 times longer than the semi-parametric estimators. Semi-parametric estimator (1), which fits the true functional form of π to both κ and V , performs poorest. This is because we should not expect either κ or V to inherit the functional form of π . Likewise, estimator (2), the flexible polynomial, performs reasonably well despite being misspecified. The choice of instruments is found to be particularly important. Using no instruments (\emptyset) outperforms the initial state instrument. This arises for the combination of two reasons. First, except in very large samples, the initial state instrument suffers from weak instrument problems, as variation in the initial state does not induce enough variation in bidding behaviour. Second, the degree of bias in the least squares estimation is expected to be small, depending on the correlation between $\Gamma_l(b_l)$ and v_l . This correlation is relatively small because b_l varies much more with other variables, such as v_l and the state variables. Finally, using the ex-post states as instruments performs much better, but does not dominate (nor is dominated by) the no-instrument estimator.

F Estimation Details and Additional Results

F.1 Participation

Figure 9 below plots the ROC curves from considering firm participation behaviour. In particular, we see that participation is predominantly determined by the distance between a firm and the project. In fact, distance is a far better predictor of participation than it is of bidding behaviour, conditional on entry. This is potentially because firms know they can never deliver on contracts more than a certain distance away, due to the difficulty of transporting asphalt or cement such distances before they cool.

While both backlogs and the size of other contracts bid upon does help predict participation, the explanatory power is minor relative to distance and other bidder \times lot specific characteristics. I previously established that this was not the case for bidding behaviour conditional on participation. This suggests that entry behaviour is fairly insensitive to changes in firms' backlogs, and so entry decisions are unlikely to change as we consider changing the

Figure 8: Monte Carlo Study

Instrument		\emptyset			\mathbf{s}_t			$\{\mathbf{s}_t^c\}$		
θ	T	Mean	SD	rMSE	Mean	SD	rMSE	Mean	SD	rMSE
(1)	θ_1									
	300	5.12	10.9	18.5	4.08	11.9	19.9	3.48	11	19.8
	1,000	5.71	5.12	15.2	4.02	6.42	17.2	4.66	5.75	16.4
	10,000	6.03	3.09	14.3	4.71	3.48	15.7	5.14	3.27	15.2
	θ_2									
	300	5.62	1.34	4.58	6.57	1.57	3.78	6.23	1.4	4.02
	1,000	5.78	0.631	4.26	6.75	0.852	3.36	6.35	0.766	3.73
	10,000	5.85	0.348	4.16	6.82	0.439	3.21	6.4	0.411	3.63
(2)	θ_1									
	300	27.2	6.89	9.98	-75.4	126	158	24.2	14.6	15.1
	1,000	27.2	3.93	8.19	-73.7	57.5	110	24.7	7.51	8.85
	10,000	27.4	1.49	7.51	-69.9	17.7	91.6	24.6	2.64	5.29
	θ_2									
	300	12.1	0.988	2.31	39.6	20.1	35.8	12.6	2.05	3.28
	1,000	12.2	0.6	2.24	38.5	8.55	29.8	12.6	1.1	2.8
	10,000	12.2	0.221	2.23	37.5	2.7	27.7	12.7	0.361	2.7
(3)	θ_1									
	300	19.7	6.13	6.14	28	108	108	18.5	11.1	11.2
	1,000	20.1	3.26	3.26	21.9	33.1	33.2	19.4	5.81	5.84
	10,000	21.2	1.32	1.81	22.2	4.18	4.72	20.3	2	2.03
	θ_2									
	300	10.4	0.897	0.968	9.11	16	16	10.7	1.69	1.82
	1,000	10.2	0.48	0.535	10.3	5.26	5.27	10.5	0.867	0.983
	10,000	9.99	0.196	0.196	9.94	0.647	0.65	10.1	0.31	0.344

Note: The true values for θ_1 and θ_2 are 20 and 10 respectively. The three instruments are: \emptyset = no instrument (OLS), \mathbf{s}_t = initial states, $\{\mathbf{s}_t^c\}$ = all the possible ex-post states, given the period began in \mathbf{s}_t . Estimator (1) is a semi-parametric estimator, using the true functional form of π to fit κ and V . Estimator (2) fits a cubic polynomial, while Estimator (3) fits a cubic spline.


auction mechanism.

F.2 Constructing the Index Function

The index is constructed as in Aradillas-Lopez et al. (2022) and Raisingh (2021), using most of the same covariates for the random forest as in Raisingh (2021).

The aim is to predict the minimum rival bid in each auction using various elements of the state. To capture rivals' states I classify the rivals of each bidder according to their distance from the bidder using distance bins (near, 0-25km, medium, 25-50km, and far, >50km), and take the average general backlog of rivals within each bin. The features I include as predictors to form λ_{nt} are: The number and average backlog of rivals in each distance bin, the number

Figure 9: Evidence of Deterministic Participation Behaviour



Note: The plot shows the estimated ROC curves from four logit specifications, assessing how firm's entry decisions are predominantly determined by, among other things, the distance between the firm and the contract.


of asphalt / concrete projects auctioned that period, as well as interactions between the type of contract (concrete/asphalt) and the number of concrete / asphalt projects auctioned each period.

I now detail the random forest I use to estimate the competition index λ , given the covariates outlined above. For a detailed description of the algorithm, see Appendix B.2 of the full random forest algorithm Raisingh (2021). The key distinction, relative to a standard random forest, is the need to avoid over-fitting when making predictions on the training data. Broadly, the algorithm proceeds as follows:

1. Split the data into K equal sized folds.
2. Estimate K random forests, each with Q trees, on data from $K - 1$ of the folds.
3. Combine the K random forests.
4. Repeat steps 1 – 3 L times, yielding L random forests, each with $Q \times K$ trees.
5. Combine the L random forests.

Following Raisingh (2021) I set $L = 24$, $K = 2$, and $Q = 50$. So every data-point is used to train around $\frac{1}{3}$ of trees. Figure 10 gives a variable importance plot, highlighting which variables have the most predictive power for the minimum rival bid, and so what most strongly influences the competition index. As in Raisingh (2021), rival backlogs have the most predictive power, followed by the number of rivals. Further away rivals appear to more strongly influence the index, perhaps because they are likely to be larger firms.

Figure 10: Variable Importance Plot



Note: This plot shows the reduction in sum of squared residuals that occurs from splitting the data on each variable. Higher numbers demonstrate more predictive power.


Because the index is auction specific I average across auctions to form the period \times bidder specific competition index. Since the most important predictors are all period \times bidder specific the index varies much more across periods than with periods.

F.3 Additional Results

F.3.1 First Stage

Figure 11 plots the observed distribution of minimum rival bids against the estimated distribution. The three parameter Weibull distribution fits the data well.

Figure 11: First Stage Fit



F.3.2 Second Stage

Figure 12 displays additional results from the second estimation step, demonstrating how the pseudo-static cost function varies with the competition index λ_{nt} . The estimated parameters can be interpreted as follows: Holding fixed a general contractor's (t1) backlog of asphalt projects, every one standard deviation in λ , as competition *decreases*, increases the opportunity cost of winning by around \$90,000. Estimated parameters generally have the expected signs, with pseudo-costs increasing in the degree of competitiveness (coefficients are positive (negative) for positive (negative) coefficients in Figure 5).

Furthermore, the estimated interaction parameters are jointly significant ($p < 0.01$) for all but the specification with weak instruments. Under the exclusion restriction that $\pi(\mathbf{s}_n)$ is independent of λ_n , we can therefore reject the null hypothesis that $\beta = 0$, rejecting the myopic model. The association between the degree of competition and bidding behaviour is strong, even when we account for equilibrium beliefs.

F.4 Comparison to Misspecified Models

I now compare estimates of $\pi(\mathbf{s}_n)$ from the dynamic multi-object model presented above, to two misspecified models: A dynamic single object model, and a static multi-object model. Results are presented in figure 13.

Figure 12: Second Stage Results: λ interactions

Instruments	Type	none (OLS)		\mathbf{s}_{nt}		$\mathbf{s}_{nt} + \vec{\mathbf{s}}_{nlt}$		$\mathbf{s}_{nt} + \vec{\mathbf{s}}_{nlt} + \vec{\mathbf{s}}_{nmt}$	
		$\hat{\theta}$	SE	$\hat{\theta}$	SE	$\hat{\theta}$	SE	$\hat{\theta}$	SE
$s_t^a \times \lambda_t$	t1	84.6	23.8	361	510	87.2	27.4	91.9	25.7
	t2	157	35.8	1230	2670	136	37.4	146	32.5
	t3	16.4	5.45	1170	2580	16.5	6.48	15.8	5.11
$s_t^c \times \lambda_t$	t1	55	15.6	207	931	68	16.6	57.4	16.3
	t2	-56	39.6	6640	15,700	-80	45.7	-75	39
	t3	3.03	4.87	-1,260	3160	2.17	5.2	3.29	4.87
$(s_t^a)^2 \times \lambda_t$	t1	-2.5	3.03	-74.1	147	-2.19	3.5	-4.06	3.2
	t2	-8.38	4.87	-586	1370	-4.91	5.94	-7.38	4.93
	t3	0.0262	0.124	-65.4	123	0.126	0.144	0.0681	0.123
$(s_t^c)^2 \times \lambda_t$	t1	-2.19	3.38	-42.3	302	-2.65	3.48	-1.45	3.52
	t2	-2.13	2.35	-44.3	327	-3.74	2.74	-3.29	2.5
	t3	-0.24	0.139	-7.07	28.2	-0.00431	0.366	-0.19	0.2
$s_t^a \times s_t^c \times \lambda_t$	t1	-4.63	4.36	8.17	142	-8.43	4.87	-5.67	4.52
	t2	44.2	15.5	133	704	60.2	16.8	56.6	14.8
	t3	0.888	0.42	105	244	0.141	1.15	0.724	0.617
R^2		0.6		-10.8		0.597		0.599	
Observations									
	T	3919		3919		3919		3919	
	$\sum_t L_t$	14691		14691		14691		14691	

Note: Estimation includes county and firm \times contract type fixed effects. Figures are given in 000s of dollars. Holding fixed a general contractor's (t1) backlog of asphalt projects, every one standard deviation in λ , as competition *decreases*, increases the opportunity cost of winning by around \$90,000.

F.4.1 Static Model

The static model is nested within the dynamic multi-object model, imposing $\beta = 0$. Estimation involves the same first and second steps presented in section 5.

F.4.2 Single Object Model

Even though bidders place multiple bids each period, the static single-object model ignores possible cost-synergies between lots, even when it allows costs to be non-linear in backlogs. One interpretation is that separate groups within the firm bid simultaneously, without communication among one another. Therefore bidding groups do not take into account how payoffs depends not only on their own bid, but also other bids within the firm.

I estimate the model using JP's procedure. I complete the first estimation step as in

the text, then skip to the third estimation step and evaluating the continuation value as in JP, taking an expectation over observed bids instead of using estimated bid distributions. Because, in practice, multiple auctions occur each period I evaluate the expected period profit by taking the sum of the expected (additive) profit from each auction occurring that period. Finally, I back out $\pi(\mathbf{s}_n)$ from the inverse bid function.

F.4.3 Results

Estimates for the static model are off by an order of magnitude, but are extremely similar to the results for the pseudo-static pay-off presented in figure 6. This is because we essentially mistake the sum of current costs and discounted future costs (and opportunity costs) for just current costs. The results for the dynamic single-object model are more more similar to the dynamic multi-object model. However this misspecified model generally under estimates the extent of the returns to scale, generally underestimating the degree of non-additivity across lots.

Figure 13: Model comparison

Model $\pi(\mathbf{s}_n)$	Type	DMO		DSO		SMO	
		$\hat{\theta}$	SE	$\hat{\theta}$	SE	$\hat{\theta}$	SE
s_t^a	t1	123	7.01	39.5	7.77	423	23.6
	t2	285	11.3	18.7	12.7	835	36.3
	t3	40	1.92	145	7.58	108	5.77
s_t^c	t1	107	5.35	49.4	9.86	378	17.3
	t2	89.1	11.9	25.9	14.3	153	53.3
	t3	15.6	1.91	89.7	6.05	55.2	6.44
$(s_t^a)^2$	t1	-0.337	1.29	0.0669	1.59	-0.116	2.44
	t2	-9.26	2.46	-1.68	2.5	-16.2	4.4
	t3	-1.34	0.147	-2.08	1.01	-0.229	0.0872
$(s_t^c)^2$	t1	-7.6	1.13	-1.39	0.889	-14.5	2.12
	t2	-14	3.38	-2.48	1.99	-4.93	1.82
	t3	-0.479	0.102	-0.671	0.672	-0.328	0.11
$s_t^a \times s_t^c$	t1	1.38	1.52	-0.364	2	7.94	2.97
	t2	33.4	7.12	-1.6	3.44	58.8	14.4
	t3	0.432	0.199	0.801	0.876	0.534	0.34
R^2		0.597		0.595		0.581	

F.5 Counterfactual Simulations

I now detail how I simulate the sequential auction regime. Time is discrete, and each period in the simultaneous regime (14 days) is split into 100 sub-periods. Auctions are distributed randomly across sub-periods.

To map the estimated $AR(1)$ transition process from 14 day-long periods into 100 sub-periods I assume the sub-period transition process remains $AR(1)$, such that the mean and variance of the process is the same as the estimated process over the 100 sub-periods, ensuring the long run process is the same. Likewise, estimated payoffs $\pi(\mathbf{s}_n)$ are only defined on 14 day long intervals. To evaluate payoffs in the sub-periods I find a function $\tilde{\pi}(\mathbf{s}_n)$ such that the expected sum of these sub-period payoffs across 100 sub-periods equals $\pi(\mathbf{s}_n)$. Finally, I use the same estimated competition index as in the text, capturing the amount of competition for each contract.

For each parameter draw, beginning at an initial set of equilibrium beliefs, I numerically find bidders' continuation values. I iteratively loop through auctions numerically maximising bidders' payoffs. I make the simplifying assumption that bidders only enter the auctions they were actually observed entering, assuming these are the auctions they have the largest cost advantage in, regardless of the choice of mechanism. In finding the continuation value, to facilitate convergence, I fix bidders' states at their observed levels. Just as in estimation I fit a quadratic form to bidders' maximum expected payoffs, and so evaluate the next the continuation value. I continue this process until the continuation value converges. I also use Newton-Kantorovich iterations to improve convergence, employing the envelope theorem to evaluate the derivative of the maximum expected payoffs.

I then simulate the system again, allowing bidders states to vary as they win, and gradually complete, contracts. I then fit the same Weibull form to minimum rival bids as used in estimation. While the payoffs of Fringe bidders do not change in the counterfactual scenario, their beliefs do. I continue this process until achieving convergence. While there may be multiple equilibria, by beginning with the equilibrium beliefs from the simultaneous regime I try to find a equilibrium close to this regime. Therefore any equilibrium will be relatively nearby that from simultaneous auctions, ensuring estimates are conservative.