Identification and Inference in First-Price Auctions with Collusion

Karl Schurter*

August 2023

Abstract

I develop a method to identify collusive bidders and estimate their effect on the seller's revenue in first-price auctions with independent, private valuations. Though colluders can use a simple strategy to make their bids appear competitive, I exploit exogenous variation in the level of competition across auctions to construct a consistent test of the null hypothesis that a given bidder is not colluding. By controlling the probability of making one or more type I errors, the set of rejected hypotheses serves as a lower confidence bound on the set of colluders. To produce a lower confidence bound on the cost of collusion, I use consistent estimates of the bidders' valuation distributions to numerically solve for the seller's expected revenues in auctions with and without collusion. I apply this methodology to estimate the extent of collusion in British Columbia's timber auctions.

1 Introduction

Consider the possibility that some bidders in a first-price, sealed-bid auction are coordinating their bids to increase the probability that they win at a lower price. Though the seller may suspect bidders have entered into such a collusive agreement, collusion might be impossible to detect based on a statistical analysis of their bids. In fact, colluders can generally use a simple and costless strategy

*email: kschurter@psu.edu. I thank Ali Hortaçsu, Azeem Shaikh, and Brent Hickman for their advice and support throughout this research. I am also grateful to John List, Robert Porter, Michael K. Price, and seminar participants at the University of Chicago, Cornell, Minnesota, MIT, and Stanford for many helpful comments and suggestions.

to ensure their bids appear consistent with a model of competitive bidding. Therefore, the question is what additional assumptions or data are needed to identify colluders and estimate their effect on the seller's revenue.

Motivated by this problem, I propose a strategy for identifying colluders in first-price auctions under the assumption that bidders have asymmetrically distributed independent private valuations (IPV). The strategy exploits variation in the level of competition across auctions that is independent of the bidders' valuations. Such variation may be induced, for example, by an increase in a binding reserve price or the set of eligible bidders. In response to this increased competition, collusive and competitive models predict different changes in the distribution of bids. Therefore, a collusive bidder reveals itself by failing to respond appropriately to an exogenous increase in its apparent competition. Based on this identification argument, I suggest a statistic to test the null hypotheses that a bidder is not colluding.

Previous collusion detection methods have leveraged idiosyncratic auction rules, legal testimony, or more restrictive assumptions on the bidders' valuation in order to identify colluders. Thus, this paper's contribution is to suggest sources of exogenous variation that can be used to preemptively detect colluders in standard first-price auctions when bidders have asymmetrically distributed valuations. Because the seller can purposefully introduce exogenous variation in competition by, for example, varying the reserve price, this paper provides a tool for sellers to continuously screen for collusion. Moreover, if the colluders bid so as to maximize their joint expected surplus, the seller can identify all colluders with probability approaching one in sufficiently large samples. As a result, this method may also serve to deter collusion by reducing the profitability of bidding rings and increasing the probability of detection, thereby improving the efficacy of enforcement methods.

The theoretical details of the identification strategy are developed in the first half of the paper. Section 3 compares this method to previously proposed collusion detection strategies. In section 4, I present a formal model of first-price auctions with collusion and discuss assumptions on how bidders might collude. Section 5 then establishes the nonparametric identification of this model. In section 6, I derive a consistent test for collusion from a test for dependence between the competitively rationalizing valuations and an exogenous instrument. Evidence from simulations indicates this testing procedure has statistical power to detect collusion even when colluders attempt to disguise their behavior. In comparison, when colluders always bid as if they were seriously competing against the non-ring competition, tests of exchangeability and independence in bidders' strategies have power equal to size.

Detecting collusion might not be the practitioner's ultimate goal, however. Indeed, statistical evidence of collusion does not indicate per se illegal activity without additional evidence that bidders explicitly coordinated their bids. An empirical analysis of bids may help antitrust authorities direct their enforcement resources toward more suspicious activity or calculate damages after cartel members have been identified, but the hypothesis testing procedure in this paper does not substitute for the "smoking gun" in a criminal investigation.

Accordingly, I emphasize the goal of estimating the effect of collusion—explicit or tacit—on prices. In section 7, I show that sellers or antitrust authorities can obtain one such estimate by sending the family-wise error rate (FWER)—the probability of falsely rejecting one or more of the null hypotheses—to zero as the data grow. Intuitively, the testing procedure promises to become increasingly conservative, such that the probability of falsely accusing a non-collusive bidder tends toward zero as the sample of auctions grows. As long as it does not become too conservative too quickly, the probability of type II errors will also tend toward zero. Therefore, the set of rejected null hypotheses consistently estimates the members of the collusive ring. Given this estimate of the ring, standard methods from the empirical auction literature will then deliver consistent estimates of each bidder's private valuation distribution. In turn, the estimated private valuation distributions can be used to predict bids in counterfactual equilibria. Most notably, the econometrician can consistently estimate the cost of collusion by solving for the equilibrium price distribution that would prevail if none of the bidders colluded.

To provide a measure of precision of this point estimate, I note that the multiple hypothesis testing framework provides simple lower confidence bounds on the set of colluders. Indeed, controlling the FWER is equivalent to controlling the probability that the true set of colluders contains the set of bidders for whom the null hypothesis is rejected. Moreover, because the cost of collusion increases with the size of the ring, a lower confidence bound on the ring produces a corresponding lower confidence bound on the cost of collusion.

As an example of this methodology, I apply the identification strategy to British Columbia's timber auctions, where controversy surrounding the fairness of its timber prices has persisted for decades. Section 8 discusses the institutional details of the British Columbian timber industry. In section 9, I find evidence that four of the most active firms do not bid competitively. Together they form a 95%-confidence bound on the collusive bidding ring. Though the overall effect on the region's timber prices is certainly smaller, I estimate that collusion reduced the revenue at a typical auction by 6.6%.

2 An Example of the Identification Problem and Its Solution

Abstracting from estimation, I assume the econometrician directly observes the joint distribution of the bids. From these data, the goal is to infer the joint distribution of the bidders' valuations for the object to be sold. To prove colluders and their valuation distributions are not identified from bid data, it suffices to show that a joint bid distribution may be rationalized by different valuation distributions depending on who is colluding with whom.

> Example: Consider an auction with a nonbinding reserve price and three bidders whose private valuations are independently distributed. Let V_i and B_i denote bidder i's valuation and bid for $i \in \{1, 2, 3\}$. Suppose the econometrician observes that their bids have a joint distribution function $S(b_1, b_2, b_3) = \sqrt{2\,b_1} \cdot \sqrt{2\,b_2} \cdot 2\,b_3$ for $b_i \in [0, 1/2]$.
>
> Case I (No Collusion). If each bidder i privately chooses a bid to maximize its expected profits
>
> $$(v_i - b) \cdot P\{\text{Bidder } i \text{ wins if it bids } b\} \; , \tag{1}$$
>
> there exists a unique distribution of valuations that rationalizes the bid distribution. In this case, bidder 1 (or, symmetrically, bidder 2) wins with probability $(2b)^{3/2}$ when it bids b. The first-order conditions for bidders 1 and 2's profit-maximization problem then implies they will bid $b = 3\,v/5$. Their marginal valuation distributions are therefore given by
>
> $$F_1^I(v) = F_2^I(v) = P\{V_i \leq v\} = P\{B_i \leq 3v/5\} = \sqrt{6v/5}$$
>
> for v between zero and 5/6. Analogously, the probability that bidder 3 wins with a bid of b is $2\,b$, its optimal bid is $b = v/2$, and its

marginal valuation distribution must be uniform between zero and 1.

CASE II (BIDDERS 1 AND 2 COLLUDE). Assume bidders 1 and 2 jointly maximize the sum of their expected profits. Because bidder 3's profit-maximization problem is unaffected, the valuation distribution that rationalizes bidder 3's bids is the same as in Case I. On the other hand, the coalition of bidders 1 and 2 will win with probability $2\,b$ whenever the greater of their bids is b. The colluder with the higher valuation should therefore optimally bid $b = \max\{v_1, v_2\}/2$, whereas the other may submit any "phantom" bid below this amount.[1] Because these phantom bids may be unrelated to the colluders' valuations, I ignore the lower of b_1 and b_2. Under the assumption that the valuations are independent, however, the higher bid still contains enough information to infer their valuation distributions because

$$\begin{aligned} P\{B_1 \leq b,\, B_2 \leq B_1\} &= P\{V_1/2 \leq b, V_2 \leq V_1\} = b \\ P\{B_2 \leq b,\, B_1 \leq B_2\} &= P\{V_2/2 \leq b, V_1 \leq V_2\} = b\,. \end{aligned} \tag{2}$$

Under the independence assumption, this system is uniquely satisfied when V_1 and V_2 are drawn from marginal valuation distributions $F_1^{II}(v) = F_2^{II}(v) = \sqrt{v}$.[2]

This example demonstrates that bids generated from competitive and collusive auctions can be observationally equivalent. Though, if a collusive model indeed generated these data, the colluders must have been careful enough to submit phantom bids that could be rationalized without collusion. In particular, despite the fact that any phantom bid below the other's bid would have maximized the colluders' expected profits, bidders 1 and 2 ensured their marginal bid distributions were independent and solved the system of equations in (2). This fact does not determine their phantom bidding strategies as a function of their private valuations, but the fairly simple strategy of always bidding one-half of

[1] If $v_1 = v_2$, bidders 1 and 2 can arbitrarily break the tie.

[2] Interestingly, collusion affects the support of bidders 1 and 2's inferred valuation distribution. This observation also holds more generally. Therefore, colluders can be identified under the assumption that bidders are symmetric or under the much weaker assumption that their valuations share the same upper extremity of their supports. A test of the null hypothesis that bidder i is colluding could be formulated as a test of $H_0 : \bar{v}_i = \max_j \bar{v}_j$ versus $H_1 : \bar{v}_i < \max_j \bar{v}_j$, where $\bar{v}_i$ denotes the upper limit of the support of i's valuations.

their valuations could have produced these data.

Remarkably, however, a competitive bidder's marginal valuation distribution can be correctly inferred regardless of any collusion among the other bidders at the auction. Because competitive bidders do not directly care whether they are competing against a single bidding ring or the maximum bid among many independent bidders, they optimally use the same bidding strategy as long as the distribution of their highest competing bid is the same. As a result, the econometrician can infer a given bidder's valuation distribution under the hypothesis that it is not colluding. I formalize this result in Lemma 1.

On the other hand, if the bidder is colluding, a competitive model of bidding will underestimate the true valuation distribution because it underestimates how much the colluders "shade" their bids below their valuations. Furthermore, the colluders' valuation distributions will be more severely underestimated when the ring is stronger relative to its non-ring competition. Therefore, the valuations that competitively rationalize a colluder's bids will covary with the level of competition at the auction.

This argument is formalized in the proof of Theorem 2 provided in appendix A. However, an extension of the example succinctly illustrates the collusion detection strategy.

> Example (continued, Bidders 1 and 2 Collude): Suppose the collusive bidding ring faces stronger competition because bidder 4 enters the auction. Assume bidder 4's valuation is distributed uniformly between 0 and 1, while the other bidders' valuation distributions are unchanged, i.e. the arrival of bidder 4 is exogenous.
>
> In equilibrium, all serious bidders optimally bid two thirds of their valuations. If bidders 1 and 2 always bid as if they were bidding competitively against bidders 3 and 4, the joint distribution of bids will be
>
> $$S(b_1, b_2, b_3, b_4) = \sqrt{3\,b_1/2} \cdot \sqrt{3\,b_2/2} \cdot 3\,b_3/2 \cdot 3\,b_4/2$$
>
> for $b_i \in [0, 2/3]$.
>
> Under the false null hypothesis that bidder 1 is not colluding, the econometrician observes the probability it wins with a bid of b is $(3\,b/2)^{5/2}$. Its optimal bid is therefore $b = 5\,v/7$. As depicted in Figure 1, the implied distribution of bidder 1's valuation is $\sqrt{15\,v/14}$

for v between 0 and 14/15, which differs from the distribution that was inferred under the null before bidder 4 entered the auction. By symmetry, bidder 2's competitively rationalizing valuation distribution is identical to bidder 1's. Because bidder 4's entry is assumed to be exogenous, the competitive IPV model cannot rationalize the collusive bidders' responses.

Alternatively, when bidders 1 and 2 are correctly assumed to be colluding, the ring wins with probability $9\,b^2/4$. Their optimal strategy is then $b = 2\,v/3$. Their valuation distributions must therefore be equal to $\sqrt{v}$ on $[0, 1]$, which is the same as the distributions obtained before bidder 4's arrival.

Figure 1: *A collusive bidder's bid and valuation distributions before and after the exogenous entry of a fourth bidder.* The level of competition affects the valuation distribution that competitively rationalizes a colluder's bids. A colluder's competitively rationalizing valuation distribution lies above the true valuation distribution because the competitive model overestimates the colluder's competition and therefore underestimates how much the colluder bids below its valuations. Because the arrival of another competitor lessens the difference between the competitive and collusive models' predictions, the valuation distribution implied by competition shifts closer to the truth.

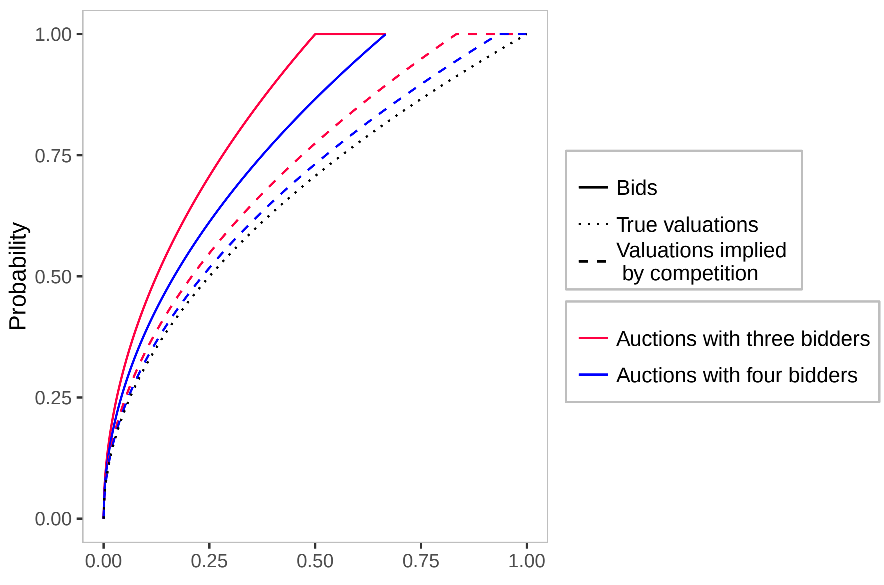

3 Related Collusion-Detection Methods

The identification strategy in this paper is based on a comparison of a given bidder's behavior across auctions, as opposed to cross-bidder comparisons within auctions. In view of this fact, Aryal and Gabrielli (2013) and Price (2008) represent the most closely related collusion-detection procedures that have been proposed in the literature. Aryal and Gabrielli suggest testing for stochastic dominance between the valuations that rationalize a firm's bids under the null and alternative hypotheses. They argue this within-bidder test asymptotically controls the probability of type I errors under their modeling assumptions, but such a test would not control size under the assumptions of section 4 because the valuations implied by collusion are always greater than the valuations implied by competition, even when the bidder is not actually colluding.

Similarly, Price (2008) looks for evidence of collusion by comparing a firm's bids across auctions. Analyzing the same data that I do in section 9, he first uses theoretically motivated criteria to identify pairs of bidders who warrant closer inspection. He then regresses bids on a vector of auction and bidder characteristics and the number of bidders at the auction. The results of these regressions indicate suspected colluders tend to bid less aggressively when another suspected colluder is nearby, which is consistent with the hypothesis that the bidders are in fact colluding.

List et al. (2007) also find suggestive evidence of collusion in the same sample of British Columbia's timber auctions. In their framework, the problem of detecting collusion is a special case of the more general problem of estimating the agents' treatment status when treatment status is not directly observed. In that sense, the bids submitted by collusive firms are the "treated" observations, and the treatment reduces the colluders' bids relative to what they would have bid if they were not colluding. Though some of their findings are inconsistent with the competitive IPV model of bidding, they conclude that the evidence of collusion is mixed. They also suggest further research is needed to quantify the impact the suspected colluders might have on expected revenues.

My proposed tests of cross-auction restrictions are also similar to the cross-mechanism analysis in Athey et al. (2011). They observe that the prices in ascending auctions were lower than predicted given the valuation distributions estimated from a sample of first-price, sealed-bid auctions. They then confirm this difference is statistically significant using a test of the null hypothesis that the average observed and predicted prices are equal. Assuming that the choice of auction mechanism is independent of the valuations, this test provides evidence against the null that all bidders are competitive.

The above identification strategies conceptually differ from the methods that attempt to detect collusion by testing the within-auction restrictions implied by the competitive model. To the extent they leverage different comparisons, these within-bidder tests complement across-bidder tests of conditional independence among bids (Porter and Zona, 1999; Bajari and Ye, 2003) and tests of the restriction that bidder i's bids depend on covariates in the same manner regardless of whether bidder i wins the auction (Porter and Zona, 1993).

Tests of independence are valid in any IPV model, but when bidders have symmetrically distributed valuations, the competitive model places further restrictions on the distribution of the data. In particular, the joint distribution of the bids must also be symmetric. By contrast, collusion among the bidders would create asymmetries if the collusive ring allocates auctions in accordance with the colluders' valuations. Thus, a test for asymmetry in bidding behavior may identify the colluders. For example, Pesendorfer (2000) observes that bidders who collude efficiently operate as though they are a single, stronger bidder whose valuation is distributed as the maximum of the individual bidders' valuations. This induced asymmetry then causes non-collusive firms to bid more aggressively, and in equilibrium, the non-collusive firms' bid distributions will stochastically dominate each of the colluders' bid distributions.

In addition, when the data include bidder-specific covariates, Bajari and Ye (2003) suggest a regression-based test of exchangeability in the bidders' strategy functions that may provide further evidence of collusion. Intuitively, this test builds on the insight that all of bidder i's competitors are exchangeable under the null hypothesis. But if bidder i is colluding with bidder j, bidder i's bid distribution will not depend on bidder j's characteristics in the same way that it depends on non-collusive bidders' characteristics.

Collusion detection methods have also been developed for other auction mechanisms, including second-price and ascending auctions (Baldwin et al., 1997; Brannman and Froeb, 2000; Marmer et al., 2017), first-price auctions with secret reserve prices and supplementary rounds of bidding (Kawai and Nakabayashi, 2014), and average-bid auctions (Conley and Decarolis, 2016). Of these methods, the nonparametric identification analysis in Marmer et al. (2017) is closest to the present paper. Under assumptions nearly identical to those in section 4, they prove the members of a collusive ring and their valuation distributions are nonparametrically identified in an ascending auction.

Interestingly, nonparametric identification in an asymmetric IPV second-price and ascending auctions is easier than in first-price auctions in the sense that detecting collusion does not require exogenous variation in competition across ascending auctions. Because collusive bidders have an incentive to manipulate their losing bids to reduce the price paid by the designated ring leader, a losing colluder's bids will be stochastically weaker than would be predicted from the distribution of its winning bids. This relative ease of detection in second-price and ascending auctions contrasts with the general perception that these auction formats are more susceptible to collusion because phantom bidders have no private incentive to defect and outbid a fellow colluder who has a higher valuation. Ironically, the same feature that facilitates collusion among the bidders—that auction prices only depend on losing bids—also facilitates the detection of collusion in second-price and ascending auctions vis-à-vis first-price auctions.

4 An Asymmetric Model of First-Price Auctions with Collusion

Let $\mathcal{N}$ denote the set of bidders bidding for an object in a first-price, sealed-bid auction. Bidder i's private valuation of the object is a random variable V_i independently distributed according to the distribution function F_i. Assume each bidder i's valuation has compact support, $\mathcal{V}_i = [\underline{v}_i, \bar{v}_i]$, and that it has a density, f_i, which is bounded away from zero on $(\underline{v}_i, \bar{v}_i]$. In addition, the seller sets a reserve price, $r < \min\{\bar{v}_i\}$.

If a bidder has not been designated to represent the collusive bidding ring or has a valuation below the reserve price, it will not submit a *serious* bid at the auction, i.e. a bid that it believes will win with positive probability. Otherwise, each bidder i chooses a serious bid b to maximize its expected payoff

$$(v_i - b) \cdot G_i(b), \tag{3}$$

where $G_i(b)$ is the probability that i wins with a bid of b. I refer to G_i as i's *competing distribution* because it is the distribution function for the highest bid among i's competitors. In equilibrium, the bidders' competing distributions depend on the strategies and marginal valuation distributions, as well as the collusive bidding ring, denoted by $\mathcal{R} \subset \mathcal{N}$.[3] I refer to the special case where $\mathcal{R}$ is the empty set or a singleton as the competitive model. When $\mathcal{R} = \mathcal{N}$, the model predicts that the all-inclusive bidding ring will obtain the object at the reserve price (McAfee and McMillan, 1992).

The third possibility is that $\mathcal{R}$ is a non-empty, non-singleton, proper subset of $\mathcal{N}$. In this case, I assume the ring colludes efficiently by nominating the member with the highest valuation to submit a serious bid that maximizes the expected profits of the ring. All other ring bidders may submit arbitrary phantom bids below the ring's serious bid. These phantom bids would never win but may be intended to create the illusion of competition.

Because each bidder knows its competing distribution, which depends on the composition and behavior of the ring, I implicitly assume bidders are sufficiently aware of any collusion to maximize their expected surplus. That said, bidders may know their competing distributions without being specifically aware of any collusion, since the number and strength of its competitors are merely proximate to the distribution of its highest competing bid. On the other hand, because comparative statics play a major role in the identification strategy, bidders must be able to respond optimally to changes in the auction environment. Unless the bidders have sufficient experience with or data from auctions spanning the full support of the covariates, this generally requires that any collusion is common knowledge among the bidders.

To summarize, the modeling assumptions are enumerated below:

Assumptions (IPV Modeling Assumptions)**.**

MA.1 The set of potential bidders, $\mathcal{N}$, is common knowledge.

MA.2 Private valuations are independently distributed over their compact supports.

MA.3 The valuation densities are bounded away from zero on $(\underline{v}_i, \bar{v}_i]$.

MA.4 Each bidder i knows G_i, the distribution of the highest bid among i's competitors.

MA.5 All serious bidders are risk neutral and bid to maximize their expected profits in (3).

[3] The model can allow for multiple rings operating in the same auction, but I focus on the case of a single ring for the sake of clarity.

In addition, the collusive ring behaves in accordance with $MA.7$

Assumptions (Modeling Assumptions Regarding Collusion).

MA.6 *A proper (possibly empty) subset of potential bidders, $\mathcal{R} \subset \mathcal{N}$ collude prior to bidding.*

MA.7 *The bidder with the highest valuation among $\mathcal{R}$ bids to maximize the expected profits in (3). All other colluders either abstain from bidding or submit a phantom bid less than the serious colluder's bid.*

Under assumptions MA.1–MA.5, Lebrun (2006) proves that, in a competitive auction, the unique Bayes-Nash equilibrium inverse bidding strategies are differentiable and strictly increasing in the bids whenever they are greater than the minimum bid. The support of the equilibrium bids share a common lower limit, though some bidders may have different upper limits. Nonetheless, the inverse bidding strategies can be continuously extended over the entire compact support of the bids, $[\underline{b}, \bar{b}]$.

Lebrun's result extends easily when some bidders collude efficiently. The competitive bidders' equilibrium strategies will be identical to the ones they would use in an auction in which the colluders were replaced by a single bidder whose valuation is distributed like the maximum of the ring's valuations. Thus, the following proposition characterizes the equilibrium bidding strategies in an auction with collusion.

Proposition 1. *Under assumptions MA.1–MA.7, there exists a unique equilibrium profile of serious bidding strategies with $\sigma_i(v_i) \leq v_i$ for each i. Moreover, each non-collusive bidder's strategy*

(i) is nondecreasing on $[\underline{v}_i, \bar{v}_i]$,

(ii) is equal to the minimum serious bid at the minimum serious bid, i.e. $\sigma_i(\underline{b}) = \underline{b}$, where $\underline{b}$ is the infimum of the set of bids that win with positive probability in equilibrium,

(iii) is differentiable and strictly increasing whenever $\sigma_i(v_i)$ is greater than the minimum serious bid, and

(iv) has an inverse that can be continuously extended over the compact support $[\underline{b}, \bar{b}]$, which is given by the bidder's first-order condition

$$\sigma_i^{-1}(b_i) = v_i = b_i + \frac{G_i(b_i)}{g_i(b_i)}, \tag{4}$$

where g_i is the density of i's competing distribution,

and each collusive bidder's equilibrium strategy

(v) satisfies properties (i)–(iv) whenever $v_i > \max_{j \neq i \in \mathcal{R}} v_j$, and

(vi) is a possibly random function of all the colluders' valuations bounded above by the serious colluder's bid whenever $v_i \leq \max_{j \in \mathcal{R}} v_j$.

In short, the serious equilibrium bids are increasing functions of the valuations, they start at a common minimum bid, and the collusive bidders adopt the same serious bidding strategies because they all face the same competing distribution. The only irregularity is that one of the non-collusive bidders' strategies or all of the colluders' strategies could be constant near the minimum bid.

5 Identification

A private value auction model is a collection of pairs $(F, \tilde{\sigma})$, where F is a joint distribution of valuations and covariates, (V, X), and $\tilde{\sigma}$ is a profile of bidding strategies $\tilde{\sigma}_i : (V, X) \mapsto B_i$ which may depend on the bidder's own valuation as well as any other bidders' valuations that it learns prior to bidding. In the competitive IPV model, F must belong to the collection of absolutely continuous valuation distributions satisfying MA.1–MA.5 that are independent conditional on X, and $\tilde{\sigma} : V \mapsto B$ is restricted to the unique strategy profile σ that satisfies the first-order condition (4) with $G_i(b|x) = \prod_{j \neq i} F_j(\sigma_j^{-1}(b; x)|x)$. When the model is enlarged to include strategy profiles that satisfy MA.1–MA.7, I refer to it as an IPV auction model with collusion.

Given the counterfactual questions of interest, the goal is to recover the F that generated the observed distribution of the data. In nonparametric models, however, data never provide information about $F(v|x)$ at valuations below the minimum bid, $\underline{b}(x)$. Thus, I say that F is identified (up to the truncation induced by the minimum bid) whenever each $F(v|x)$ is uniquely determined for v with $v_i > \sigma_i^{-1}(\underline{b}(x))$ for all i, where σ_i is the serious bidding strategy defined by (4). Letting (B, X) denote the random vector of bids and covariates, I formalize this definition as follows.

Definition 1. A model $\mathcal{M}$ is *identified* from the joint distribution of bids and covariates up to the truncation at $\underline{b}$ if, whenever $(\tilde{\sigma}(V, X), X)$ and $(\tilde{\sigma}'(V', X'), X')$ are equal in distribution for some (V, X) and (V', X') distributed according to

F and F' with $(F,\sigma),(F',\sigma') \in \mathcal{M}$, $F(v,x) = F'(v,x)$ for all continuity points of F with $v > \sigma_i^{-1}(\underline{b})$.

More generally, a function, $Y(B,X)$, of the bids and covariates may be observed. For instance, the transaction price is sometimes the only observable bid. In this case, I will say the model is identified from Y if the above definition holds when $(\tilde{\sigma}(V,X),X)$ and $(\tilde{\sigma}'(V',X'),X')$ are replaced by $Y(\tilde{\sigma}(V,X),X)$ and $Y(\tilde{\sigma}'(V',X'),X')$.

Definition 1 does not require $\tilde{\sigma}$ to be identified because the phantom bidding strategies that rationalize the data need not be unique and are not required to identify the colluders or estimate the cost of collusion. The serious bidding strategies $\sigma_i : v_i \mapsto b_i$ are identified, however, because they are uniquely determined by F, the identities of the colluders, and the bidders' first-order conditions. Note that the modeling assumptions imply $\sigma_i(v_i) = \tilde{\sigma}_i(v)$ if bidder i is competitive or is the designated serious bidder from the cartel, but the converse need not hold.

5.1 Identification from Prices and Exogenous Variation in Competition

The IPV model is not generically identified from bids alone unless the identities of colluders are known a priori. But, as suggested by the example in section 2, an instrument that induces variation in the level of competition can be used to construct a test for collusion as a test of independence between the instrument and competitively rationalizing valuation. Once all the ring members have been identified in this way, each of the bidders' true competing distributions can be computed. Their valuation distributions can then be inferred from the first-order condition of the bidders' profit maximization problems. Hence, the IPV model with collusion is identified from the distribution of winning bids and exogenous variation in the level of competition.

The proof of Theorem 2 generalizes the above argument. In the remainder of this section, I sketch the key steps toward establishing this main identification result. All proofs are in Appendix A.

As an important intermediate step, I prove Lemma 1, which establishes (i) the marginal distribution of each bidder's serious bid, $\sigma_i(V_i)$, is identified from the distribution of prices and identities of the winners, and (ii) each bidder i's valuation distribution is identified under the assumption that i is not colluding. In other words, part (i) of Lemma 1 identifies the (potentially counterfactual) distribution of the bid that bidder i would have submitted if i were trying to maximize its expected profits in (3). While trivial to show for non-collusive bidders, I construct a collusive bidder i's marginal serious bid distribution as a function of the distribution of prices, denoted by M, and the probability that bidder i wins and bids less than or equal to b, denoted by $M_i(b)$. This construction is adapted from a result in the competing risks literature that was introduced to the empirical auction literature in Athey and Haile (2002). By analogy to the competing risks literature, I refer to M_i as the cumulative incidence function for the risk that bidder i wins the auction.

Given the marginal serious bid distributions, I then construct the competing distribution for each bidder i under the null hypothesis that i is not colluding as the product of each other bidder's marginal serious bid distribution. Bidder i's valuation distribution is then identified under the assumption i is not colluding with anyone else. Note that any collusion among i's competitors does not distort the inference made about i's valuation.

Lemma 1. *Assume MA.1–MA.7.*

(i) Each bidder's marginal serious bid distribution is identified from the prices and the identities of the winners.

(ii) A bidder's competing distribution and valuation distribution are identified under the null hypothesis that it is not colluding with anyone else at the auction.

If bidder i might be colluding with an unknown subset of the other bidders, there are many competing distributions that can be constructed from the other bidders' bid distributions, each of which generally implies a different valuation distribution. Accordingly, the model is not identified from the winning bids. In fact, the example in section 2 already proved a stronger result. Namely, the model is still not identified when the full vector of bids is observed.

Theorem 1. *The IPV model with collusion is not identified from the full vector of bids.*

Interestingly, however, exogenous instruments for each bidder's competing distribution are sufficient to recover the true composition of the bidding ring even when only the winning bids are observed. Formally, the identifying assumptions are as follows.

Assumptions (Identification Assumptions)**.**

IA.1 *For each i, the reverse hazard rate of bidder i's competing distribution depends non-trivially on Z_i.*

IA.2 $Z = (Z_1, \dots, Z_n)$ *is observable, the distribution of Z_i conditional on covariates X_i is non-degenerate, and $Z_i \perp\!\!\!\perp V_i \mid X_i$ for all i.*

Assumption IA.1 guarantees that the instrument is relevant to i's optimal bidding strategy. Assumption IA.2 asserts that the econometrician observes independent variation in an instrument for each bidder i. For instance, if Z_i is equal to the exogenously varying reserve price for all i, then both identification assumptions will be satisfied without conditioning on the additional covariates in X_i. Alternatively, Z_i could be a vector of i's competitors' characteristics (e.g. distance to a job site), in which case it may be necessary to condition on i's characteristics and use the residual variation in Z_i to shift i's competing distribution.

The applied researcher must select and defend the use of an exogenous competition shifter with care. The case for exogeneity is most compelling when the seller explicitly randomizes or experiments with participation rules or reserve prices. Naturally occurring variation in the number of bidders might be useful if a model of selective entry indicates that the distribution of a bidder's valuation conditional on participating in the auction does not depend on the number of active bidders. Alternatively, if the entry model implies a negative relationship between a bidder's valuations and the set of active bidders under the null hypothesis that it is not colluding, the variation in the number of bidders might permit a conservative test for collusion.

As indicated by the example in section 2, the crux of the main identification result in Theorem 2 is that a bidder's competitively rationalizing valuations are independent of the exogenous competition shifter if and only if it is not colluding with any other bidders in the auction. After formally proving this in Lemma 2, Theorem 2 immediately follows from now standard identification results in the empirical auction literature.

Lemma 2. *Under MA.1–MA.5 and IA.2, the competitively rationalizing valuations $V_{i,\emptyset}$ are independent of the competition shifters Z_i (possibly conditional on additional covariates X_i) if i is not colluding with anyone in $\mathcal{N}$. Under the additional assumptions MA.6–MA.7 and IA.1, the competitively rationalizing valuations are independent of the competition shifters if and only if bidder i is not colluding.*

Theorem 2. *Under MA.1–MA.7 and IA.1–IA.2, the IPV model with collusion is identified from the distribution of prices, identities of the winners, and covariates.*

Remark 1. There is a unique configuration of the ring that simultaneously rationalizes all of the bidders' bid distributions and satisfies the independence assumption. Letting $V_{i,\mathcal{R}}$ denote the random valuation that rationalizes i's bids under the assumption that the bidders in $\mathcal{R}$ are colluding, this means that there is a unique non-singleton $\mathcal{R} \subset \mathcal{N}$ such that

$$V_{i,\mathcal{R}} \perp\!\!\!\perp Z_i \,|\, X_i \text{ for all } i\,. \tag{5}$$

Lemma 2 states, however, that $V_{i,\emptyset}$ and Z_i are independent if and only if $i \notin \mathcal{R}$. Thus, it is not necessary to search over all possible configurations of the ring to find the subset that satisfies the identifying restriction (5). Instead, it is sufficient to test for dependence between $V_{i,\emptyset}$ and Z_i for each i. The true set of colluders is then given by $\{i : i \in \mathcal{N},\, V_{i,\emptyset} \not\perp\!\!\!\perp Z_i \,|\, X_i\}$.

Remark 2. A slightly weaker version of Theorem 2 extends to auctions in which multiple rings are operating. In this case, the above argument can be used to show that all bidders who are colluding with anyone else will be identified. Assuming each bidder is a member of at most one bidding ring, the next step would be to argue that there is a unique partition of these collusive bidders into rival rings that rationalizes the data. This partition is unique as long as each partition implies a different competing distribution, hence different predictions about the bidders' responses to the instruments.

An exception to the identification result occurs when some of the colluders' valuation are identically distributed. For instance, suppose there are two strong bidders and two weak bidders, and each of the strong bidders is colluding with one of the weak bidders.[4] If each bidder always bids as if it is bidding competitively against the rival cartel bidder, the distribution of the reserve prices and winning bids will reject the comparative statics implied by the competitive model but will not reveal which strong bidder is colluding with which weak bidder. Even in this exceptional case, however, the configurations of the rings are identified up to permutations of the identical bidders. Moreover, the competing

[4]The same argument holds when the bidders' valuations are all drawn from the same marginal distribution. I assume that there are two types of bidders to illustrate that this problem arises even if only a subset of the bidders' valuations are exchangeable.

distributions, and hence the distribution of each bidders' private valuations, are identified because they are invariant under these permutations.

Remark 3. The test of independence between $V_{i,\emptyset}$ and Z_i is unable to detect collusion between an active bidder in $\mathcal{N}$ and a bidder who is eligible but does not submit a serious with bid positive probability. Assuming colluders maximize their joint expected surplus, such a situation could only arise if the supports of the eligible bidders' valuations do not overlap. As an example, consider a three-bidder auction with a reserve price $r < 1$ in which a weak bidder with valuations that take support on $[0, 1]$ colludes with a strong bidder whose valuations take support on $[1, 2]$, while the lone competitive bidder has valuations taking support on $[0, 2]$. In equilibrium, the minimum serious bid might be strictly less than one,[5] but the weak colluder is never observed to submit a serious bid because the strong colluder always has a greater valuation. Hence, one cannot apply Lemma 1 to recover the distribution of what the weak colluder would have bid if it were the designated cartel bidder. In this case, variation in the strong colluder's competing distribution may not be useful in detecting the collusion because there is not necessarily any difference between the competing distribution inferred under the null and alternative hypotheses. Indeed, if the weak colluder completely abstains from bidding, then it might never appear in a dataset that only records the identities of those who submit a bid, and the collusion would be undetectable based on any analysis of the bids.

6 Empirical Framework and Definition of the Estimators

6.1 Empirical Framework

Let t index auctions in which an object with characteristics X_t is available for sale to a set $\mathcal{N}_t$ of eligible bidders. Let $Z_t = (Z_{it})_{i \in \mathcal{N}_t}$ and $\tilde{B}_t = (\tilde{B}_{it})_{i \in \mathcal{N}_t}$ denote the instrument and bid for bidder i in auction t. A bidder whose payoff from winning the object, u_{it}, is less than the reserve price, $\tilde{r}_t$, will not participate in auction t, in which case I record the bid as censored at the reserve price. Therefore, for each auction, the data consist of realizations of the auction-level

[5] The minimum serious bid in the unique weakly undominated equilibrium would be the greatest maximizer of $(1 - b)F_3(b)$ on $[\max\{r, 0\}, 1]$, where F_3 is the valuation distribution of the competitive bidder. See equation (17) in Lebrun (2006).

covariates, instruments, and left-censored bids: $\left(Z_t, \tilde{R}_t, \mathcal{N}_t, X_t, \mathbb{1}\left\{\tilde{B}_t \geq \tilde{R}_t\right\}\right)$. Asymptotically, I assume the number of auctions, T, tends toward infinity while the set of bidders, $\cup_{t=1}^{\infty} \mathcal{N}_t$, is finite.

If the objects for sale are adequately described by a small number of characteristics, then nonparametrically estimating the distribution of valuations conditional on X_t may be practical. Commonly, however, some assumptions are needed to reduce the dimensionality of the objects' heterogeneity. To this end, I assume that each bidder's utility is additively separable in the object's observable characteristics and an unobservable idiosyncratic private component, V_{it}, which is independent of X_t and independently distributed across bidders and independently and identically distributed across auctions:

$$U_{it} = V_{it} + \mu(X_t)\,.$$

Additive separability in the valuations implies additive separability in the equilibrium bidding strategies and the distribution of V_{it} is identified up to location (Athey and Haile, 2002).

Homogenized bids and reserve prices can then be defined relative to an arbitrary benchmark, x_0. Let $B_{it} = \tilde{B}_{it} - \mu(X_t) + \mu(x_0)$ and $r_t = \tilde{r}_t - \mu(X_t) + \mu(x_0)$ denote those homogenized quantities.[6] In words, B_{it} is the bid that would have been observed if $\tilde{r}_t$ had been r_t and the auction-level covariates had been x_0 instead of x_t. If $\mu(x_0)$ is normalized to zero, the homogenized equilibrium bidding strategies satisfy

$$B_{it} = V_{it} - \frac{G_i(B_{it}|z_t, r_t, \mathcal{N}_t)}{g_i(B_{it}|z_t, r_t, \mathcal{N}_t)}\,.$$

In practice, the function μ is typically parameterized as $\mu(\cdot;\beta)$. The relatively fast rate of convergence for estimators of β asymptotically justifies using the homogenized bids as though they were data for the purposes of the nonparametric estimators defined below. In the next sections, I work exclusively with homogenized bids and reserve prices and refer to them simply as bids and reserve prices.

[6] In some applications, it may also be appropriate to homogenize Z_t. For instance, if Z_t is the reserve price, it should be homogenized so that variation in the instrument reflects variation in the screening level. But one would skip this step if $\mathcal{N}_t$ serves as the instrument for the level of competition.

6.2 Definition of the Estimators

The estimators that I define in this section are sample analogs to their population counterparts in the proof of Lemma 1. In each case, I condition on the set of eligible bidders and use a continuous second-order kernel function, K, to smooth over the reserve prices and Z. For instance, the sample analog to $M_i(\cdot, z, r, \mathcal{N})$—the conditional cumulative incidence function for the risk that bidder i wins—is defined by

$$\begin{aligned}\mathbb{M}_{iT}(b \mid z, r, \mathcal{N}) &= \frac{\sum_t \mathbb{1}\{b_{it} = p_t\}\mathbb{1}\{b_{it} \leq b\} \cdot \mathbb{1}\{\mathcal{N}_t = \mathcal{N}\} \cdot K_h(z - z_t, r - r_t)}{\sum_t \mathbb{1}\{\mathcal{N}_t = \mathcal{N}\} \cdot K_h(z - z_t, r - r_t)} \\ &= \frac{\sum_t \mathbb{1}\{b_{it} = p_t\}\mathbb{1}\{b_{it} \leq b\} w_t}{\sum_t w_t},\end{aligned}$$

where $K_h(u) = |h|^{-1/2} K(h^{1/2}\, u)$, h is a symmetric positive definite bandwidth matrix, and the price p_t is the maximum of the reserve price and the highest bid at auction t, and w_t is an abbreviated notation for the kernel-based weight on auction t.[7]

Let $\mathbb{M}_T(b \mid z, r, \mathcal{N}) = \sum_i \mathbb{M}_{iT}(b \mid z, r, \mathcal{N})$ be the estimator for the conditional distribution of the sale price and let $\mathbb{M}_{-iT}(b \mid z, r, \mathcal{N}) = \sum_{j \neq i} \mathbb{M}_{jT}(b \mid r, z, \mathcal{N})$ be the conditional cumulative incidence function for the risk that bidder i loses the auction. By analogy to equation (13), I then use the $\mathbb{M}_{iT}$ estimators to construct an estimator for the marginal bid distributions,

$$\begin{aligned}\mathbb{S}_{iT}(b \mid z, r, \mathcal{N}) &= \exp\left\{ - \int_b^{\bar{b}} \frac{1}{\mathbb{M}_T(\cdot \mid z, r, \mathcal{N})}\, d\mathbb{M}_{iT}(\cdot \mid z, r, \mathcal{N}) \right\} \\ &= \exp\left\{ - \frac{\sum_t \mathbb{M}_T(p_t \mid z, r, \mathcal{N})^{-1} \mathbb{1}\{b_{it} = p_t\}\mathbb{1}\{p_t > b\} w_t}{\sum_t w_t} \right\}.\end{aligned} \tag{6}$$

In accordance with part (i) of Lemma 1, this estimator does not depend on whether i is assumed to be competitive or collusive. Moreover, because it constructed from the cumulative incidence functions, it depends solely on the only bids known *a priori* to be serious—i.e. the winning bids. Consequently, $\mathbb{S}_{iT}(b|z, r, \mathcal{N})$ is consistent for the marginal distribution of bidder i's serious bid under both the null and alternative hypotheses.

Let $\frac{G_{i,\mathcal{R}}}{g_{i,\mathcal{R}}}$ denote the amount bidder i would be expected to shade its bid if $\mathcal{R}$ is the collusive bidding ring. This quantity can also be expressed in terms of

[7] If none of the bidders' valuations exceed the reserve price, I record this as a price equal to the reserve and the "winner" is the seller.

the cumulative incidence functions, as in

$$\frac{G_{i,\mathcal{R}}(b\,|\,z,r,\mathcal{N})}{g_{i,\mathcal{R}}(b\,|\,z,r,\mathcal{N})} = \begin{cases} M(b\,|\,z,r,\mathcal{N}) \Big/ \frac{\partial M_{-i}(b|z,r,\mathcal{N})}{\partial b} & i \notin \mathcal{R} \\ M(b\,|\,z,r,\mathcal{N}) \Big/ \frac{\partial M_{-\mathcal{R}}(b|z,r,\mathcal{N})}{\partial b} & i \in \mathcal{R} \end{cases}$$

When the reserve price does not bind, a consistent estimator can be obtained by substituting $\mathbb{M}_T(b\,|\,z,r,\mathcal{N})$ in the numerator and a kernel estimator for the derivative of M_{-i} or $M_{-\mathcal{R}}$ with respect to b. If the reserve price is binding and bidding strategies are all strictly increasing, this density will be unbounded near the reserve price for all i. Consequently, the typical kernel density estimator will not be consistent.

If the reserve price binds and bidders are symmetric, one can use the change of variables suggested by Guerre et al. (2000) to create the transformed data, $(\sqrt{p_t - r_t}, r_t)$, whose density can be shown to be bounded everywhere when bidders are symmetric. One can then apply a boundary correction procedure similar to Karunamuni and Zhang (2008) on these data to obtain an estimator $\mathbb{m}^*_{-iT}(b,z,r,\mathcal{N})$ or $\mathbb{m}^*_{-\mathcal{R},T}(b,z,r,\mathcal{N})$.[8] An estimator for the density of the original data is given by $\mathbb{m}_{-iT}(b,z,r,\mathcal{N}) = \mathbb{m}^*_{-iT}(\sqrt{b-r},z,r,\mathcal{N})/(2\sqrt{b-r})$.

To account for the upper boundary of the bid distribution, I first observe that the maximum bid is an increasing function of the reserve price. Therefore, the minimal non-decreasing upper envelope of the observed pairs of bids and reserve prices provides a non-parametric estimator for the upper boundary of the support of the bids conditional on the reserve price. For example, if the reserve price serves as the instrument, this envelope function is defined as $\hat{\bar{b}}(r,\mathcal{N}) = \max\{b_{it} : (b_{it}, r_t, \mathcal{N}_t),\, \mathcal{N}_t = \mathcal{N},\, r_t \le r,\}$.[9] One can then apply the same boundary correction procedure to the sample of $p_t - \hat{\bar{b}}(r_t,\mathcal{N}_t)$ to obtain a consistent estimate of the density for values of b within one bandwidth of $\hat{\bar{b}}(r)$.

Finally, the conditional density is consistently estimated by the ratio of $\mathbb{m}_{-iT,\mathcal{R}}(b,z,r,\mathcal{N})$ to a boundary corrected estimator for marginal density (r,z). Thus, the estimator for $\frac{G_{i,\emptyset}(b\,|\,z,r,\mathcal{N})}{g_{i,\emptyset}(b\,|\,z,r,\mathcal{N})}$ is given by $\frac{\mathbb{M}_T(b\,|\,z,r,\mathcal{N})}{\mathbb{m}_{-iT}(b\,|\,z,r,\mathcal{N})}$.

[8] Karunamuni and Zhang (2008) uses a combination of transformation and reflection of the data to reduce the order of the bias in kernel density estimation near the boundary. Hickman and Hubbard (2015) first demonstrate its use in application to inference in first-price auctions. Pinkse and Schurter (2019) observe that the recommended boundary bandwidth sequence in Karunamuni and Zhang (2008) leads to suboptimal rates of convergence. In this paper, I use a boundary bandwidth that is proportional to the main bandwidth.

[9] If z contains other continuously distributed variables in addition to the reserve price, one can similarly estimate an the upper envelope $\hat{\bar{b}}(r,z,N)$ as long as the upper boundary is monotonic in z, as well.

Bidder i's competitively rationalizing valuation in auction t would then be given by

$$\tilde{v}_{it,\emptyset} = b_{it} + \frac{\mathbb{M}_T(b_{it} \,|\, z_t, r_t, \mathcal{N}_t)}{\mathbb{m}_{-iT,\emptyset}(b_{it} \,|\, z_t, r_t, \mathcal{N}_t)} \,.$$

The right side of (10) may also be used to define an estimator of the inverse strategy function σ_{iT}^{-1}. Note that, even though I have previously ignored losing bids in order to avoid contamination by phantom bids, equation (10) may also be evaluated at bidder i's losing bids in order to construct a full sample of pseudo-values.

If bidders are not symmetric and the reserve price binds, however, the rate at which the bid densities diverge near the reserve price is not known *a priori*. For at most one bidder, the bid density may diverge at a rate faster than $1/\sqrt{b-r}$, while all other bidders' bid densities diverge at the same rate slower than $1/\sqrt{b-r}$ (see Proposition 2 in Appendix A). This creates a boundary issue that cannot be corrected by traditional methods because the transformed density is potentionally unbounded for one of the bidders.

Instead, I suggest using the estimates of the marginal serious bid distributions to construct estimates of the equilibrium expected payment function e_i under the null hypothesis as in Pinkse and Schurter (2022):

$$e_i(p|z, r, \mathcal{N}) = \begin{cases} p\, G_i^{-1}(p; z, r, \mathcal{N}) & , \text{ if } p \geq G_i(r|r, z, \mathcal{N}) \\ rp & , \text{ if } p < G_i(r|r, z, \mathcal{N}) \end{cases} \tag{7}$$

The expected payment function is the amount that a bidder should expect to pay the seller in equilibrium as a function of the probability with which they expect to win when they submit their bid. For $p < G_i(r|r, z, \mathcal{N})$, the expected payment is undefined because a bidder would never expect to win with that probability in equilibrium, but it will be convenient to continuously extend e_i and its first derivative by setting $e_i(p; r, z, \mathcal{N}) = rp$ for $p < G_i(r|r, z, \mathcal{N})$.

The expected payment function is useful because its slope is equal to the bidder's valuation at the bidder's optimally chosen win probability. Thus, we can estimate the inverse strategy function with an estimate of the slope of e at $G_{-iT}(b_{it}|r_t, z_t, \mathcal{N}_t)$ constructed from the estimates $\mathbb{M}_{-iT}$ and $\mathbb{M}_T$. This approach avoids estimating an unbounded derivative because the derivative of the expected payment function with respect to the probability of winning is bounded when the reserve price binds Pinkse and Schurter (2022).

A conventional kernel-smoothed estimate of the slope of e simplifies to a weighted sum of the observed prices at auctions that bidder i lost. Under the null hypothesis, an estimator for the slope of e_i is given by

$$e'_{iT}(\rho|r,z,\mathcal{N}) = \frac{1}{h^2}\int_0^1 pG^{-1}_{-iT}(p|r,z,\mathcal{N})K'\left(\frac{\rho-p}{h}\right)dp\,. \tag{8}$$

The natural estimator for G_{-i} under the null hypothesis is the product of bidder i's rivals' marginal serious bid distributions:

$$\begin{aligned} G_{-iT}(b|r,z,\mathcal{N}) &= \prod_{j\neq i}\mathbb{S}_{it}(b|r,z,\mathcal{N}) \\ &= \exp\left\{-\int_b^{\bar{b}}\frac{dM_{-iT}(x|r,z,\mathcal{N})}{M_T(x|r,z,\mathcal{N})}\right\} \\ &= \exp\left\{-\sum_t \frac{\mathbb{1}\{b\le p_t\}\mathbb{1}\{b_{it}<p_t\}w_t}{\sum_s \mathbb{1}\{p_s\le p_t\}w_s}\right\}. \end{aligned}$$

Because the estimates of the marginal serious bid distributions and cumulative incidence functions are step-functions with discontinuities at the observed auction prices, G_{-iT} is a step function with discontinuities at the observed prices in auctions in which bidder i lost. The inverse of G_{-iT} is then conventionally defined as $G^{-1}_{iT}(p|r,z,\mathcal{N}) = \inf\{b|G_{iT}(b|r,z,\mathcal{N})\ge p\}$, which is another step-function. Consequently, the integral in the definition of e'_{it} becomes a sum over prices with weights given by $\frac{1}{h^2}\int_{q_{it-1}}^{q_{it}} qK'\left(\frac{p-q}{h}\right)dq$, where $q_{it} = G_{iT}(p_{(t:T)})$ and $p_{(t:T)}$ denotes the t-th order statistic from the sample of T auction prices.[10]

Hence, an alternative estimator for the competitively rationalizing win probability and valuation is given by

$$\hat{\rho}_{it,\emptyset} = G_{-iT}(b_{it}|r_t,z_t,\mathcal{N}_t) \tag{9}$$

$$\hat{v}_{it,\emptyset} = \sum_t p_{(t:T)}\frac{1}{h^2}\int_{q_{it-1}}^{q_{it}} qK'\left(\frac{\hat{\rho}_{it,\emptyset}-q}{h}\right). \tag{10}$$

[10] As in the case of the estimator based on the approach of Guerre et al. (2000), some modifications are required near the boundaries located at zero and one because $e(p)$ is not defined outside the unit interval. Pinkse and Schurter (2022) develop a boundary correction analogous to Karunamuni and Zhang (2008), as well as a correction based on boundary kernels.

Using either estimation estimator for the pseudo-valuations, the inverse strategy estimate σ_{iT}^{-1} might be nonmonotonic in finite samples. A monotonic estimator for the strategy function can be defined as $\sigma_{iT}(v; z, r, \mathcal{N}) = \inf\{v : \sigma_{iT}^{-1}(b; z, r, \mathcal{N}) \geq v\}$ or as the slope of the greatest convex minorant of the estimated expected payment function (Pinkse and Schurter, 2022). An estimate of the competitively rationalizing valuation distribution is then given by,[11]

$$\mathbb{F}_{iT,\emptyset}\left(v|z, r, \mathcal{N}\right) = \mathbb{S}_{iT}(\sigma_{iT}(v; z, r, \mathcal{N}) \,|\, z, r, \mathcal{N})\,.$$

More generally, the valuation that rationalizes a collusive bid b_{it} when $\mathcal{R}$ is the assumed set of colluders is given by

$$\tilde{v}_{it,\mathcal{R}} = b_{it} + \frac{\mathbb{M}_T(b_{it}|z_t, r_t, \mathcal{N}_t)}{\mathbb{m}_{-\mathcal{R}T}(b_{it}|z_t, r_t, \mathcal{N}_t)}$$

or

$$\begin{aligned}
\hat{\rho}_{it,\mathcal{R}} &= \mathbb{G}_{-\mathcal{R}T}(b_{it}|r_t, z_t, \mathcal{N}_T) \\
&= \exp\left\{-\sum_s \frac{\mathbb{1}\{b_{it} \leq p_s\}\mathbb{1}\{\max_{j\in\mathcal{R}} b_{js} < p_s\}w_s}{\sum_l \mathbb{1}\{p_l \leq p_s\}w_l}\right\} \\
\hat{v}_{it,\mathcal{R}} &= \sum_t p_{(t:T)} \frac{1}{h^2} \int_{G_{-\mathcal{R}T}(p_{(t-1:T)})}^{G_{-\mathcal{R}T}(p_{(t:T)})} qK'\left(\frac{\hat{\rho}_{it,\mathcal{R}} - q}{h}\right) dq\,,
\end{aligned}$$

when the alleged colluder i outbids the other members of $\mathcal{R}$, and is otherwise bounded above by $\max_{j\in\mathcal{R}} \tilde{v}_{jt,\mathcal{R}}$ or $\max_{j\in\mathcal{R}} \hat{v}_{jt,\mathcal{R}}$, respectively.

6.3 Uniform Convergence of the Estimators

The estimators for the cumulative incidence functions, $\mathbb{M}_{iT}(\cdot|z, r, \mathcal{N})$, serve as the building blocks for the estimators defined above. Because these kernel-based estimators' asymptotic behavior is well understood, the estimates of F_i, G_i, and S_i will yield similarly well behaved asymptotics if the transformation ϕ defined by the right-hand side of equation (6) is Hadamard differentiable in the conditional cumulative incidence functions so that the functional delta method

[11] The empirical distribution of $\hat{v}_{it}$ is an alternative estimator for F_i under the null hypothesis that bidder i is not colluding. If a bidder is colluding, however, its losing bids might not have the same distribution as its winning bids, causing these two estimators to diverge. Similar to the parametric tests in Porter and Zona (1993), a test of equality between these estimators could be used to test for collusion, but it will fail to detect collusion, for example, if colluders always bid as if they were independently competing against the non-ring bidders.

applies. Fortunately, when the reserve price is binding, ϕ is differentiable as a map into the space of càdlàg functions. Otherwise, when the reserve price does not bind, some trimming near the minimum observed bid is required in order to bound the integrand in equation (6).[12] The standard uniform rates of convergence for kernel-based estimators will therefore carry through. In particular, the above estimators will uniformly converge at the optimal rates derived by Guerre et al. (2000).

Because the estimators defined above employ a second-order kernel, I make the following assumptions to ensure the cumulative incidence density has two continuous derivatives.

Assumptions (Smoothing Assumptions)**.**

SA.1 The conditional valuation densities $f_i(\cdot|\cdot,\mathcal{N})$ twice continuously differentiable in the continuous components of Z_{-i} and r and continuously differentiable in v_i for all i.

SA.2 The equilibrium bid distributions are atomless.

SA.3 $|h_M|^{1/2}(\log T/T)^{d/(4+d)} \to c < \infty$ where h_M is the bandwidth used to estimate $\mathbb{M}_{iT}$ and d is the number of continuously distributed components of (R_t, Z_t).

SA.4 $|h_m|^{1/2}(\log T/T)^{(1+d)/(5+d)} \to c < \infty$ where h_m is the bandwidth sequence used to estimate $\mathbb{m}_{iT}$.

The smoothness assumptions on f_i are sufficient for $M_i(\cdot\,|\,\cdot)$ to be twice continuously differentiable in all of its arguments for $b > r$, but the derivative with respect to the bid could be unbounded near the reserve price. In addition, one bidder could have an atom in its equilibrium bid distribution at the minimum bid (see Proposition 1). To simplify the asymptotic analysis, I assume the reserve price binds and focus on the estimator based on the expected payment function. Though I note that, if needed, the estimator based on the competing bid density could be used when the reserve price does not bind if one estimates the location and size of any atom in the equilibrium bid distributions along with the derivatives of the bid distribution conditional on a bid greater than the minimum serious bid.

[12] Marmer et al. (2017) use a trimming sequence to resolve a similar issue in the asymptotic behavior of their estimator.

Lemma 3. *Under MA.1–MA.5 and SA.1–SA.2, $\frac{\partial M_i}{\partial b}$ is twice continuously differentiable in its continuous arguments, and $\frac{\partial e_i}{\partial p}$ is bounded, continuous, and has two bounded derivatives in all of its arguments.*

Theorem 3. *Assume the above smoothness and bandwidth conditions and that $r > \underline{v}_i$ for all $i \in \mathcal{N}$. Suppose $M_i(b \mid \cdot)$ has a uniformly bounded and continuous second derivative in its continuous arguments. The estimators $\mathbb{M}_{iT}$ and $\mathbb{S}_{iT}$ converge uniformly in $(z, r, \mathcal{N})$ to tight objects in the space of càdlàg functions on $(r, \bar{b}]$, while $\mathbb{F}_i$ is uniformly consistent on compact subintervals of $(r, \bar{v}_i)$ and is pointwise asymptotically normal:*

$$\left(\frac{T}{\log T}\right)^{2/(4+d)} (\mathbb{M}_{iT}(\cdot \mid z, r, \mathcal{N}) - M_i(\cdot \mid z, r, \mathcal{N})) \rightsquigarrow W_i$$

$$\left(\frac{T}{\log T}\right)^{2/(4+d)} (\mathbb{S}_{iT}(\cdot \mid z, r, \mathcal{N}) - S_i(\cdot \mid z, r, \mathcal{N})) \rightsquigarrow \phi'_{(M_i, M)}(W_i, W)$$

$$\left(\frac{T}{\log T}\right)^{2/(5+d)} \|\mathbb{F}_{iT,\mathcal{R}}(v| z, r, \mathcal{N}) - F_{i,\mathcal{R}}(v| z, r, \mathcal{N})\|_{\infty,\mathcal{V}} = O_p(1)$$

where $\phi'_{(M_i,M)}$ denotes the Hadamard derivative of ϕ at (M_i, M), W_i and $W = \sum W_i$ are centered Gaussian processes, and $\mathcal{V} \subset (r, \bar{v}_i)$ is compact.

6.4 Test Statistics

The proof of Theorem 2 demonstrates that bidder i's competitively rationalizing valuation distribution, $F_{i,\emptyset}$, is independent of the instrument, Z_i, if and only if bidder i is not colluding. An appropriate statistic to test this prediction depends on whether Z_i is binary, discrete, or continuously distributed. If, for example, the seller randomly decides to "set aside" some auctions for a particular category of bidder, as is common practice in government procurement (see, for example, Athey et al., 2013; Krasnokutskaya and Seim, 2011), a test for collusion could be based on a Kolmogorov-Smirnov-type statistic such as

$$D_{iT} = \sup_{v>r} \left|\mathbb{F}_{iT,\emptyset}(v \mid \mathcal{N}) - \mathbb{F}_{iT,\emptyset}(v \mid \mathcal{N}')\right| ,$$

where r is minimum value for which F_i is identified for both levels of the instrument, $\mathcal{N}$ and $\mathcal{N}'$.

When the instrument takes on more than two values, generalizations of the Kolmogorov-Smirnov statistic could be used. Alternatively, Haile et al. (2003) observe that a test based on the mean valuations may converge at a faster rate.

They propose an asymptotically chi-square statistic based on the insight that the mean valuation implied by the IPV model should be strictly decreasing in the number of symmetric bidders due to the worsening winner's curse. To adapt this statistic to a test for collusion, the null hypothesis would be the same, but the alternative would be the opposite. That is, the mean valuations should be strictly increasing in the number of symmetric bidders when the bidder is colluding.

Otherwise, when the instrument is continuously distributed, I propose a test based on correlations. Although this test would not detect dependence among higher moments, covariances may be sufficient. When Z_i is equal to the reserve price, however, testing for dependence is complicated by the fact that valuations are censored below the reserve, which introduces spurious correlation between the observed $V_{i,\emptyset}$ and r. But, unlike the typical random censoring problem, auction data include the value at which the valuations would have been censored—i.e. the reserve price. Consequently, the assumption of independent censoring can be tested nonparametrically using a conditional Kendall's τ statistic.

Formally, I define the conditional Kendall's τ statistic as

$$\hat{\tau}_i^C = \hat{\tau}^C(\hat{v}_{i,\emptyset}, r) = \frac{\sum_{t \leq s} \text{sign}(\hat{v}_{it,\emptyset} - \hat{v}_{is,\emptyset}) \cdot \text{sign}(r_t - r_s) \cdot \Lambda_{its}}{\sum_{t \leq s} \Lambda_{its}},$$

where Λ_{its} is an indicator for the event that both $(\hat{v}_{it,\emptyset}, \hat{v}_{is,\emptyset})$ and (r_t, r_s) can be ordered.[13] For example, if $\hat{v}_{it} < r_t < r_s \leq \hat{v}_{is}$, then the summand in the numerator can be evaluated even though $\hat{v}_{it}$ is not observed.[14] In general, the pairs are "orderable" if and only if $\max\{r_t, r_s\} \leq \max\{\hat{v}_{it}, \hat{v}_{is}\}$. This implies that one of the pseudo-valuations must be uncensored in order for the pair to enter the summation. Because one of the observations may be censored, however, the auctions for which bidder i did not submit a bid will be represented in $\hat{\tau}_i^C$. In this sense, τ_i^C incorporates information contained in bidder i's participation

[13] A similar statistic has been proposed as a test of independent truncation in medical trials (see, for example, Tsai, 1990; Martin and Betensky, 2005), but I do not believe that this particular statistic has been studied. No doubt, this is due to the unusual nature of the censoring problem induced by binding reserve prices. In the context of a medical trial, it would correspond to the situation in which the statistician observes the times at which subjects would have dropped out of the study if they had not experienced the event under study, e.g. disease progression or death.

[14] I assume a bidder with a valuation equal to the reserve price submits a bid.

decisions.

If bidder i's valuations were directly observed, this statistic would have an expected value of zero under the null. This can be easily verified by noting the summand evaluates to 1 in six of the twelve orderable permutations and to -1 in the other six. The mean is therefore zero under the null hypothesis because each concordant permutation is equally likely to occur as one of the discordant permutations when valuations are independent of reserve prices. Moreover, because $\hat{\tau}^C(v_{i,\emptyset}, r)$ is the ratio of two U-statistics, its asymptotic normality would immediately follow.

Inference using $\hat{\tau}_i^C$ is complicated by the fact that the pseudo-observations, $\hat{v}_{it,\emptyset}$, are estimated. Nonetheless, Proposition 3 in Appendix A establishes that τ_i^C zero under the null hypothesis and the estimator is asymptotically linear and asymptotically normal if the kernel bandwidth is $o(T^{-1/4})$ so that the bias introduced by estimating the pseudo-valuations with second-order kernels is asymptotically smaller than $T^{-1/2}$. Asymptotic linearity and normality of the estimator then implies the validity of the bootstrap Gill (1989). In appendix B, I report results of simulations that demonstrate asymptotic size control using bootstrapped critical values.

Inference may also be affected by the fact that the conditional Kendall's τ estimator might converge at the same rate as the homogenized bids if a bid homogenization step is required. In the empirical application in this paper, I ignore this additional source of sampling variation when estimating critical values for the conditional Kendall's τ statistic because I use orders of magnitude more prices to homogenize the bids than there are uncensored pseudo-valuations available to compute the Kendall's τ statistic for any individual bidder i. The errors in the estimation of the τ statistic and the inverse bidding strategies given the homogenized bids is likely to be large relative to the estimation error in the homogenized bids themselves, although their variance is on the same order according to my asymptotic framework. In applications with more bids per bidder relative to the number of auctions, the sampling error in the bid homogenization step is likely to be non-negligible. Depending on the details of the bid homogenization step, one might prove consistency of the bootstrap critical values by demonstrating the bid homogenization estimator composed with the asymptotically linear conditional Kendall's τ estimator yields an asymptotically linear multistep estimator.

7 Testing and Confidence Bounds on the Cost of Collusion

To account for the fact that multiple hypotheses are tested simultaneously, I suggest a testing procedure that asymptotically controls the family-wise error rate (FWER)—i.e. the probability of making one or more false rejections. The power to detect collusion would be greater under alternative testing procedures that control less stringent error rates, such as the false discovery proportion. The FWER is convenient, however, because it lends itself to the construction of confidence bounds on the cost of collusion.

If α is the chosen tolerance for the FWER, the set of rejected hypotheses forms a lower confidence bound on the set of colluders, i.e.

$$\liminf_{T\to\infty} P\{\mathcal{R}_T \subseteq \mathcal{R}_0\} \geq 1-\alpha\,, \tag{11}$$

where $\mathcal{R}_0$ is the true collusive ring and $\mathcal{R}_T$ is the set of bidders for whom the null hypothesis is rejected given data from T auctions. Then, because the cost of collusion is monotonic in the ring (with respect to set inclusion), this translates into a lower confidence bound on the cost of collusion:

$$\liminf_{T\to\infty} P\{C(\mathcal{R}_T, F) \leq C(\mathcal{R}_0, F)\} \geq 1-\alpha\,, \tag{12}$$

where $C(\mathcal{R}, F)$ is the difference between the seller's expected revenues when the bidders in $\mathcal{R}$ are or are not colluding and bidders' valuations are jointly distributed according to F.[15] The function C does not typically have an analytic expression, but its value at $(\mathcal{R}_T, F)$, and hence a lower confidence bound on the true cost of collusion, can be found numerically. Of course, this estimator is infeasible because F is not observed. The unknown joint valuation distribution must be replaced by a consistent estimator, which will generally differ from the joint distribution that rationalizes the data under the assumption that $\mathcal{R}_T$ is the set of colluders. Because equilibrium strategies are continuous in F with respect to the weak topology (Lebrun, 2002), a consistent "plug-in" estimator for $C(\mathcal{R}_T, F)$ is obtained by numerically evaluating $C(\mathcal{R}_T, \mathbb{F}_{T,\hat{\mathcal{R}}})$ for some consistent estimator of the ring, $\hat{\mathcal{R}}$.

[15] This counterfactual assumes the set of potential bidders would not be affected by the dissolution of the collusive ring. If some firms would exit the industry in response to the stronger competition, then this counterfactual would overstate the revenue gains to eliminating collusive behavior.

8 The British Columbian Timber Market

8.1 Background

British Columbia's Ministry of Forests manages 95% of the province's timber supply. Its annual revenues averaged US$1.1 billion between 1996 and 2000, of which $210 million was raised by auctioning timber licenses under the Small Business Forest Enterprise Program (SBFEP). These licenses grant the right to harvest timber from designated areas during a specified period of time, typically lasting about one year and no more than four years. Though these auctions directly accounted for less than 20% of its revenue, the auction prices affected the index that the Ministry used to benchmark its prices for all other timber licenses. Therefore, a natural question is whether the auction prices accurately reflect the fair market value of the harvesting rights.

Two papers address this question in reduced form analyses (List et al., 2007; Price, 2008). Both find evidence consistent with the hypothesis that some of the firms in this sample are colluding. On the other hand, in technical reports prepared for the Ministry of Forests, Athey et al. (2002) and Athey and Cramton (2005) argue that their proposed auction reforms would reduce the benefit of anti-competitive bidding to the point that collusion would have to be pervasive within a local market and sustained over at least three years in order to significantly influence the market price for timber. Given the large number of small logging firms that could enter the market if firms successfully conspired to keep prices down, they conclude that collusion to this extent is implausible.

In light of this debate, I apply the identification strategy in this paper to assess the competitiveness of the SBFEP auctions between 1996 and 2000 and estimate the effect any collusion may have had on the auction prices.

8.2 Bidding Procedures and Identifying Variation

The auctions in the data were conducted under the auspices of the SBFEP. The regional SBFEP offices published a list of the timber licenses to be sold at auction. Prior to each auction, the regional office would specify which firms were eligible to bid. To be eligible to bid, a firm had to be registered with the SBFEP and hold no more than two outstanding timber licenses. In addition, the SBFEP announced whether registrants that owned or leased their own milling facilities would be eligible to participate. Firms with milling capabilities were excluded from about 80% of the auctions between 1996 and 2000.

When the regional offices solicited bids, they included several documents containing details about the timber license. These included survey maps, plans for extracting the logs, the estimated volume of merchantable timber by tree species, and projected road development costs. Along with these supporting documents, the office calculated a reserve price per cubic meter of harvested timber in accordance with one of the two appraisal methods described below. Firms were also invited to inspect the tract themselves.

At any time before the auction closed, interested firms could submit a sealed "bonus" bid equal to the amount that they would pay above the reserve price. The regional office then opened and recorded all of the bids and awarded the license to the highest bidder. All of the bids and the identities of the bidders were publicly announced at this time. Throughout logging operations, the winner paid an amount per cubic meter of harvested timber equal to the reserve price plus its bonus bid.[16]

Substantial variation in the reserve prices allows me to test whether collusion played a role in keeping the Canadian lumber prices below prices in the United States during these years immediately proceeding a lumber trade dispute between the US and Canada. For timber appraised using the old non-hedonic pricing formula, the reserve price was set at a fixed fraction of the timber's estimated price per cubic meter.[17] This estimated price was computed as the difference between the tract's appraised value and the average value of licenses sold in the region, plus a base rate that was determined by the Revenue Branch. The base rate was adjusted quarterly depending on how the prices of active licenses compared with the Ministry's target rate, which in turn was a piecewise-linear function of a weighted average of British Columbia's lumber and wood chip price indices. In addition, the regional office could increase the reserve price to reflect their silvicultural or development expenses. In practice, however, these adjustments were not made in auctions that excluded firms with milling capabilities. As a result of all these modifications to the ministry's initial appraisal, I observe auctions in which my estimate of the value of the licenses are similar while their reserve prices greatly differ.

In response to this exogenous variation in appraisals, market participants appealed to the Ministry to improve its appraisal methodology to more quickly re-

[16] As a condition of the license, the winner agrees to pay penalties for unharvested timber.

[17] The details of this appraisal process are contained in the annually updated Interior Appraisal Manual. The appraisal manual that took effect on October 1, 1999 (Canada. B.C. Ministry of Forests, 1999) describes the reserve pricing policies for timber licenses appraised using either of the two methods.

spond to changes in market conditions. Beginning in 1999, the Ministry adopted a hedonic pricing formula and generally set the reserve price at 70% of the appraised value without the option of adjusting for silvicultural and development expenses. This change in the reserve pricing policy relative to the pre-1999 method provides further exogenous variation. Moreover, the hedonic pricing formula was modified in 2001, soon after the period under investigation. Thus, to the extent that the formula needed improving, the auctions conducted in later years may still be expected to contain exogenous variation in the reserve price.

9 Analysis

9.1 The Modeling Assumptions

In the empirical framework, the valuations are assumed to be private, independent across firms, independent of the district in which the firm competes, and independent across auctions after controlling for the observed auction covariates.

In the context of British Columbia's timber auctions, the firms' valuations are likely to be private because there is no active spot market for harvested timbers and the winner only pays for the amount of timber that they actually harvest.[18] Loggers typically negotiate bilateral supply agreements with local mills before bidding in an auction. Therefore, the firms know the price per cubic meter at which they would be able to sell the logs if they were to win the auction. Furthermore, because the winners' payments to the Ministry of Forests are based on the actual merchantable volume of harvested timber, as opposed to the estimated volume of timber, they do not bear risk regarding the total merchantable volume covered by the license. Firms are also largely insured against uncertainty about the composition and quality of the timber because the Ministry fixes the stumpage rate for low-quality timber and timbers used for fence posts or other specialty products. Consequently, the bids only apply to good quality coniferous sawlogs.[19] Lumber produced from different

[18]The data come from the interior region of British Columbia. There is a spot market for timber in the coastal region.

[19]If the Ministry sets the price for low-grade and specialty timbers too high (or low) relative to a bidder's valuation, the bidder would have an incentive to decrease (increase) its bid on licenses containing a higher proportion of that timber. In an attempt to determine whether this is an issue, I test whether the proportion low-quality timber that was actually harvested is correlated with the auction price after controlling for the license characteristics as in Appendix C. Though I only observe the volume of harvested timber for a subset of the auctions

tree species within this category are typically substitutable in their commercial uses, so variation in the species composition should not greatly affect the bidders' valuations for the license.[20] Moreover, firms are not likely to have significantly different information about the species composition (Paarsch, 1997). Therefore, the winner's curse is unlikely to be an important factor in the firms' bidding strategies.

Nonetheless, the firms' willingness to pay for a license are certainly correlated through observable characteristics of the timber license. I assume, however, that the idiosyncratic components of their valuations are independent conditional on these characteristics. Given the rich set of covariates in my data, I argue that this is plausible. Indeed, I condition on the same set of covariates that the Ministry of Forests used to appraise the licenses. In Appendix C, I estimate my own appraisal of the timber licenses and find that it better predicts the winning bids than the Ministry's appraisal. I maintain, however, that the Ministry selected an appropriate set of variables to include in their hedonic pricing model even though their exact formula could be improved by accounting for the selection problem introduced by the binding reserve price.

The IPV assumptions are consistent with prior studies of these auctions, but I depart from earlier work by relaxing the symmetry assumption and allowing marginal valuation distributions to differ across firms. Such asymmetry might arise, for example, from differences in expertise or relationships with mills. These asymmetries could be significant considering the fact that firms' participation decisions vastly differ and are not explained by observable differences in their characteristics. The majority of firms participated in very few auctions, and won at most once between 1996 and 2000. In contrast, there are only nine firms that won more than 10 auctions. These firms appear to rely on the SBFEP auctions much more than the less active firms, and could have different relationships with mills than the fringe competitors do.

In order to pool data across districts in the estimation, I assume that the bidders' valuations are independent of the district in which the auction takes place. While this assumption is consistent with earlier work (Paarsch, 1997; List et al., 2007; Price, 2008), it would not be palatable if the distance from the firm's headquarters to the harvesting site were correlated with the unobservable

in the sample, the proportion of low-quality timber is not significantly correlated with the unexplained variation in the winning bids.

[20] Tree species differ in the amount of lumber that can typically be recovered from a sawlog of a given volume. The Ministry controls for this heterogeneity by using species-specific lumber recovery factors to calculate a lumber price index for each auction.

variation in the firm's harvesting costs. Because the data include the cycle time—the estimated time necessary to transport the timbers to the nearest point of appraisal—I maintain the assumption that the headquarters-to-site distance is not correlated with the cost of extracting the timber conditional on the observable covariates. See appendix D for details on the clustering algorithm used to pool districts into markets.

Finally, I ignore any inter-auction dynamics that may be induced by capacity constraints and the Ministry's cap on the number of outstanding licenses that firms can hold. A firm's bidding strategy would then depend on the current state of the market, which might be summarized by the characteristics of future auctions that have already been scheduled, expected characteristics of auctions that have not been announced, and all of the firms' backlogged workloads. I observe the time series of auction characteristics and could assume that firms have rational expectations for the licenses that are likely to come to auction in the future. But, as discussed in the previous section, I only have noisy measures of the firms' backlogs. If they were observable, auction dynamics could be accounted for as in Jofre-Bonet and Pesendorfer (2003). This approach, however, would be limited by the infrequency with which bidders participate in auctions, the high dimensionality of the state vector, and the unobservable asymmetries among the bidders.

9.2 Using Participation Decisions to Test for Collusion

The typical bidder participates in a small fraction of the auctions for which it appears to be eligible. Much of this can be explained by geography: bidders typically focus their activities in one or two neighboring districts. As a first approximation to their true participation decision process, I assume that firms never considered participating in auctions outside of the districts in which they were active. Still, some firms only bid in a nearby district a few times, so firms might not have been potential entrants in every auction in that district. I therefore define a firm as being active in a district only if it participated in at least 5% of the auctions for which my data indicates it was eligible.[21]

Yet, even in the districts where they are most active, firms typically partic-

[21] To construct a proxy for a firm's eligibility, I track all of the SBFEP licenses it has won but have not yet expired. If I observe a firm bid in an auction when my data suggest it has three outstanding licenses, I presume that it completed logging operations on one of them so that, in fact, it only has two outstanding licenses. If it does not win the present auction, then it will be eligible to bid in subsequent auctions.

ipate in less than a third of the auctions for which they appear to be eligible. To rationalize this behavior, I therefore rely on the fact that the reserve price is often binding.[22] The binding reserve price creates a censoring problem in which the valuations are left-censored at the reserve price but the value at which they would have been censored is always observed. Thus, the data for bidder i consists of $(\max\{\hat{v}_{it,\emptyset}, r_t\}, r_t)$ for auctions t in districts where bidder i was active or in which bidder i participated. Given these data, there are several ways to test for dependence between $\hat{v}_{it,\emptyset}$ and r_t, but I have found that a conditional Kendall's τ estimator performs well in simulations.

More precisely, this testing procedure performs well in simulations where each bidder is a potential entrant in all of the simulated auctions. In the present application, however, it is possible that a bidder's participation decision was affected by unobservable factors, such as the number of non-SBFEP contracts held by the bidder at the time of bidding. As a more conservative approach to the testing problem, it may be prudent to ignore any information that might or might not be contained in bidder i's participation decision.

In this case, the data for bidder i would consist only of $(\hat{v}_{it,\emptyset}, r_t) \,|\, v_{it,\emptyset} \geq r_t$, which is precisely the typical case of left-truncation discussed in Tsai (1990) and Martin and Betensky (2005). An appropriate test statistic would then be defined analogously to $\hat{\tau}_i^C$ above, except that Λ_{ts} is replaced by an indicator for the event that $\max\{r_t, r_s\} \leq \min\{\hat{v}_{it}, \hat{v}_{is}\}$. In words, the pair (t, s) is only included in the computation of $\hat{\tau}_i^{\text{Trunc}}$ if the estimated valuations are "comparable" in the sense that i's valuation was uncensored in both auctions and would have remained uncensored if the reserve prices were interchanged. Because this condition is more restrictive, $\hat{\tau}_i^{\text{Trunc}}$ uses less of the data to test the null hypothesis. As a result, it will generally be estimated with less precision but will avoid any bias that comes from erroneously assuming that valuations were censored.

I use $\hat{\tau}_i^{\text{Trunc}}$ to denote this statistic, where the superscript indicates that

[22] Alternatively, the lack of participation could be explained if participating in auctions is costly either because firms incur costs in order to evaluate the timber license and learn their valuation or because the process of submitting a bid is costly. The identification and estimation strategies disucssed above can be adapted to cases in which bidders face entry costs as long as the screening value for each bidder i, i.e. the lowest valuation for which bidder i chooses to enter the auction, can be estimated. For example, if bidder i faces a constant entry cost c_i across auctions, its screening value is identified from its conditional competing distribution $G_i(\cdot|r, c, z, \mathcal{N})$. The conditional Kendall correlation between the screening value and the competitively rationalizing valuations could then be used to test for collusion by bidder i.

In these data, there are no direct costs of participating in an auction (apart from a refundable deposit), and the cost of evaluating the timber license is likely to be significantly reduced because the Ministry of Forests' shares detailed information with all potential bidders.

$\hat{\tau}_i^{\text{Trunc}}$ is a function of the truncated data $(\hat{v}_{it,\emptyset}, r_t)|v_{it,\emptyset} \geq r_t$ as opposed to $\hat{\tau}_i^C$, which uses the full sample of censored pseudo-valuations and reserve prices. The population parameter $\tau_i^{\text{Trunc}}(V_{i,\emptyset}, r)$ is defined analogously.

Finally, another shortcoming of the data is that I do not observe auctions that did not receive any bids. Thus, the empirical distribution of prices is conditional on the event of a sale, which would lead me to overestimate the strength of the bidders' competing distributions. To avoid this bias, I estimate the probability that bidder i participates conditional on the reserve price and the event that at least one other potential bidder submits a bid. If i does not coordinate its participation decision with anyone else and valuations are independent of each other and the reserve price, this consistently estimates bidder i's unconditional probability of entry. Similarly, the probability of entry by at least one other bidder besides i can be estimated from the observed data regardless of any collusion among i's competitors. Given these probabilities, I then estimate the probability of sale conditional on r and recover the unconditional price distribution and cumulative incidence functions. The hypothesis test for bidder i built on this approach will asymptotically control size. Furthermore, one can control the FWER in a multiple hypothesis testing framework by re-estimating the conditional probability of sale under each null hypothesis that bidder i is not colluding with anyone else. In the results presented in the next section, I alternatively assume the largest bidders' entry decisions are independent of the event that at least one of the competitive fringe of relatively inactive firms chooses to participate.

9.3 Hypothesis Test Results

Following List et al. (2007) and Price (2008), I restrict my analysis to auctions in which the estimated timber volume was greater than 1,000 m^3. I further restrict the analysis to auctions that exclude firms with milling capabilities.[23]

I formulate a test of the null hypothesis that bidder i is not colluding either as a two-sided test of $\tau^C(V_{i,\emptyset}, r) = 0$ against $\tau^C(V_{i,\emptyset}, r) \neq 0$ or as a test of $\tau^e xtTrunc(V_{i,\emptyset}, r) = 0$ against $\tau^{\text{Trunc}}(V_{i,\emptyset}, r) \neq 0$.[24] To asymptotically control

[23] Upstream logging firms and vertically integrated mills might bid differently when competing against each other because the auction outcome could be relevant to price negotiations related to other timber harvest from other tracts. The static auction model presented in this paper is likely better suited to auctions among loggers.

[24] The conditional statistics, τ^C and τ^{Trunc}, need not have the same sign as the unconditional τ. Although the unconditional τ should be positive if a bidder is colluding, the sign of τ^C and τ^{Trunc} is ambiguous.

the probability of making one or more false rejections, I adopt the bootstrap-based multiple testing procedure in Romano and Wolf (2010).

As a final consideration before applying the testing procedure, I must decide which null hypotheses to test because there are not enough data to simultaneously test for collusion by all of the firms while controlling the FWER. Indeed, more than 1,400 firms bid in a SBFEP auction between 1996 and 2000, but most of these firms participated in three or fewer auctions. Moreover, if the ultimate goal is to estimate the cost of collusion, this comprehensive analysis would not be optimal because the estimated effect on expected revenue is minimally affected by whether one of the small firms is colluding. The confidence bound on the cost of collusion can be improved by judiciously allocating statistical power across the individual hypotheses.

Developing an adaptive procedure to estimate optimal weights and maximize the weighted average power of the test is outside the scope of this paper. Instead, I essentially set the weights equal to zero for small firms and simultaneously test the null hypothesis for only the 18 firms that bid in more than 30 auctions and won more than five of them. The precise thresholds are arbitrary, but as long as they are chosen in a manner that is independent of the test statistics, the procedure described above will asymptotically control the FWER.

The results of the testing procedure are reported in Table 1.[25] The marginal p-values for each firm are estimated from 40,000 bootstrap samples, while the adjusted p-values are equal to the smallest α for which the multiple testing procedure would have rejected the null hypothesis for firm i while controlling the FWER at level α.

The choice of test statistic has a large impact on the decisions to be made for each null hypothesis. I argue that $\hat{\tau}_i^T$ is more defensible because it is more robust to misspecification in the bidders' participation decisions. Thus, according to this preferred specification, the null hypothesis for firms 1–4 can be rejected at the 0.05 level of significance, while the null hypothesis for firm 5 is rejected at the 0.10 level. Firms 1–4 form a lower 95%-confidence bound on the set of firms that colluded in the SBFEP auctions.

One drawback to my detection method is that the results could be difficult to

[25]The diagonal entries in Table 2 indicate the total number of pseudo-valuations used to estimate the conditional Kendall's τ statistics. The bandwidth sequence for smoothing over reserve prices was somewhat arbitrarily chosen to be equal the scaling factor implied by Silverman's rule of thumb multiplied by $T^{-1/3.9}$ instead of the optimal rate. The bandwidth sequence used to estimate the derivative of the equilibrium expected payment function is $1.4T_{im}^{-1/3.9}$ where T_{im} is the number of auctions in market m that i lost.

Table 1: Results from simultaneous tests of the null hypotheses that bidder i is not colluding. Each test is formulated as a two-sided test for positive correlation between the reserve price and the valuations that rationalize i's bid. The p-values are estimated from 40,000 bootstrap replicates. The marginal p-values indicate the level of significance in an individual test of the hypothesis for bidder i. To generalize the concept of p-values in the multiple hypothesis testing framework, the adjusted p-values indicate the smallest level of tolerance for one or more type I errors at which the null hypothesis for bidder i can be rejected.

	$\hat{\tau}_i^T$		$\hat{\tau}_i^C$	
	Marginal	Adjusted	Marginal	Adjusted
1	$<.001$	$<.001$	$>.999$	$>.999$
2	$<.001$	.002	$>.999$	$>.999$
3	.001	.020	.998	$>.999$
4	.002	.030	.772	$>.999$
5	.006	.095	.474	$>.999$
6	.013	.208	.356	$>.999$
7	.019	.290	.986	$>.999$
8	.021	.313	.932	$>.999$
9	.022	.330	.986	$>.999$
10	.051	.600	.370	$>.999$
11	.054	.621	.978	$>.999$
12	.059	.657	.571	$>.999$
13	.072	.730	.063	.679
14	.136	.920	.199	.977
15	.193	.975	.939	$>.999$
16	.200	.978	.716	$>.999$
17	.721	$>.999$	.597	$>.999$
18	.943	$>.999$	.988	$>.999$

interpret. If, for example, only one null hypothesis is rejected, then the test does not indicate with whom they might be colluding. Similarly, if the test detects collusion by more than one firm but they only ever bid in very distant markets, then they would have very little incentive to collude. This would cast doubt on the interpretation of the test as a test for collusion. To address this concern, Table 2 tabulates the number of times that each of the 18 firms bid in a district where one of the other bidders was active. These patterns demonstrate that the firms are indeed active in the same districts as other suspected colluders. Although firms 1 and 2 did not submit bids in the same districts, they did compete in similar districts that were three of the first to be grouped into the same market by the hierarchical clustering algorithm.[26]

Table 2: Each row i tabulates the number of times firm i bid in a district where each of the other firms is active. For example, firm 1 bid 36 times in districts where firm 3 was active, whereas firm 3 bid 19 times in districts where firm 1 was active.

Firm	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
1	36	0	36	36	0	36	36	36	17	0	36	0	0	0	0	0	0	19
2	0	58	40	7	0	4	4	50	58	4	11	0	3	19	3	0	48	4
3	19	8	28	21	0	21	22	25	20	2	21	0	0	2	0	0	5	9
4	11	17	23	34	0	24	25	24	19	12	27	0	0	17	0	0	17	23
5	0	0	0	0	42	0	0	0	0	0	0	37	0	0	0	42	0	0
6	37	2	39	40	0	42	42	41	11	2	41	0	0	2	0	0	2	31
7	8	10	19	27	0	27	39	21	13	10	26	0	0	10	0	0	10	34
8	28	26	49	48	0	49	51	58	30	19	51	0	0	25	0	0	22	44
9	1	33	10	14	0	3	3	20	35	2	9	0	3	24	3	0	25	2
10	0	32	32	32	0	32	32	32	32	32	32	0	0	32	0	0	32	32
11	21	3	22	23	0	23	24	25	5	1	32	0	0	3	0	0	3	21
12	0	0	0	0	38	0	0	0	0	0	0	38	0	0	0	38	0	0
13	0	1	0	0	0	0	0	0	31	0	0	0	31	0	31	0	0	0
14	0	53	45	49	0	45	45	49	53	45	47	0	0	53	0	0	51	45
15	0	4	0	0	0	0	0	0	14	0	0	0	14	0	31	0	0	0
16	0	0	0	0	10	0	0	0	0	0	0	3	0	0	0	38	0	0
17	0	35	16	18	0	13	13	30	35	13	27	0	0	32	0	0	35	13
18	6	10	16	28	0	21	36	23	10	10	19	0	0	10	0	0	10	36

Having rejected the joint null hypothesis that all of the bidders are competitive, a logical question is whether a collusive model can better fit the data. In particular, the question is whether assuming firms 1 through 4 are members of a collusive bidding ring makes the estimated Kendall's τ statistics less significantly positive. To that end, I perform the testing procedure again, but instead use the valuations that rationalize the observed bids when $\mathcal{R} = \{1, 2, 3, 4\}$. The test

[26] Firm 1 competed exclusively in the Kamloops and Clearwater districts, while firm 2 bid in the 100 Mile House district and six other districts in three other markets.

statistic is unaffected for the non-ring bidders $i \notin \mathcal{R}$. In contrast, the collusive bidders' pseudo-valuations are greater to the extent that they win in the same market because their competing distribution is weaker than those estimated under the null hypothesis. Furthermore, the pseudo-valuations are censored at $z_{it} = \max\{r_t, \hat{v}_{jt} : j \in \mathcal{R}, j \neq i\}$ rather than r_t because any phantom bids might be unrelated to the bidders' valuations.[27] Then, under the maintained assumptions that the true valuations are mutually independent and independent of the reserve price, the test statistic should be equal to zero if $\mathcal{R}$ is the true set of colluders. The estimated p-values for the tests for firms 1–4 are reported in Table 9.3

There are two caveats in interpreting the results of this exercise. First, the p-values in Table 9.3 should be treated with caution, because they do not account for the manner in which I selected the configuration of the ring to test. Rather, this exercise provides qualitative evidence that collusion helps explain the firms' responses to variation in the reserve prices. Second, I cannot eliminate the possibility that alternative modeling assumptions might generate data that are observationally equivalent to collusive equilibria of my model.

Table 3: The same procedures are used to determine whether a model with collusion better predicts bidders' responses to variation in the reserve prices. The estimated adjusted p-values suggest that an alternative model in which firms 1–4 collude efficiently is slightly more consistent with the observed bidding behavior.

	Firm			
Model	1	2	3	4
$\mathcal{R} = \emptyset$	<.001	.002	.020	.030
$\mathcal{R} = \{1, 2\}$	.004	.002	.040	.041

To estimate a 95%-confidence bound on the cost of collusion in the typical auction where suspected colluders are active, I use the estimated valuation distributions that rationalize the bids when firms 1 and 2 are assumed to be colluding to compute the effect of collusion on the expected revenue.[28] In this

[27] Recall that $\hat{v}_{jt} > \hat{v}_{it}$ if and only if $v_{jt} > v_{it}$ when i and j are members of the ring because their estimated inverse bidding strategies are numerically identical and strictly increasing. That is, $\hat{v}_{jt} > \hat{v}_{it}$ if and only if $b_{jt} > b_{it}$. Hence, there is no error in observing the event that a colluder's valuation is censored.

[28] I cannot avoid modeling the firms' participation decisions for the purposes of estimating the private valuation distributions. Though the private valuations are identified from the winning bids, the data do not include enough wins by each firm to produce reliable estimates. Instead, I use a Kaplan-Meier estimator based on the full vector of bids and participation decisions. I assume that the firm's valuation was censored at the reserve price if it did not

typical auction in the market in which both firms were active, the reserve price is $\$15.6/m^3$ below my estimated appraisal. I therefore include firms 1–4 in the set of eligible bidders along with a group of competitors, that, for simplicity, are assumed to be symmetric and to belong to the competitive fringe consisting of firms who did not participate more than 30 times and win more than five licenses. I then solve for the equilibrium bid distributions with and without collusion among bidders 1–4 using the numerical methods described in the Appendix E. After adding enough competitive fringe bidders to match the median price in the market in which all four firms were active, I find that the median revenue increases by $\$3.10/m^3$ when firms 1–4 bid competitively, which amounts to 6.6% of the median price in that market.

10 Discussion and Extensions

In first-price auctions, the competitive and collusive models imply different comparative statics because they predict different changes in the distribution of each bidder's highest competing bid. Therefore, a test for collusion may be based upon a test of whether the competitive model rationalizes a bidder's response to exogenous variation in its competing distribution. The results of this analysis indicate four of the most active firms in British Columbia's timber auctions may have colluded. To obtain a lower 95%-confidence bound on the cost of collusion, I estimate their effect on the median revenue in a typical auction to be $\$3.10/m^3$.

Any hypothesis test is a joint test of all of the modeling and identification assumptions. Thus, a rejection of the null hypothesis could be a rejection of any of the baseline modeling assumptions or of the exogeneity of the instrument. Nonetheless, my identification strategy may be practicable when exogenous variation in competition is already available or when the seller is willing to purposely introduce randomness into its reserve price or set-aside program in order to detect collusion.

Regarding the validity of the modeling assumptions, one might be able to rule out alternative explanations based on the sign of the test statistic. For example, a colluder's competitively rationalizing valuations should be stochastically greater if the level of competition increases, but they will be negatively related if the bidder is actually competitive and risk averse. A one-sided test

bid in a district where it participates at least 5% of the time that it is eligible.

would then control the probability of type I errors even when the bidders' preferences are misspecified.

Furthermore, one can adapt the same procedures used to test the competitive model to test whether an IPV model with collusion rationalizes bidders' responses to the observed variation in the level of competition. The other modeling and identification assumptions remain the same. Thus, if the test fails to reject any of the null hypotheses when the bidders in $\mathcal{R}$ are assumed to be colluding but rejects the null hypotheses of those bidders when the ring is assumed to be empty, this finding would be consistent with the interpretation that the original hypotheses were rejected due to collusion.

Appendix F considers extensions to the model that would allow for unobserved auction-level heterogeneity and affiliation in the private valuations. Though an alternative strategy could control the probability of falsely accusing bidders of colluding when valuations are affiliated, further research is needed to determine whether a consistent test for collusion exists in this context. Moreover, the cost of collusion is generally not identified from the observables considered in this paper. On the other hand, the distribution of unobservable auction-level heterogeneity could be identified as in Krasnokutskaya (2011) or Roberts (2013) if an auxiliary measure of this heterogeneity is observed. Having isolated the distribution of auction-level heterogeneity, one could then identify the cost of collusion as above. In the appendix, I explain why I believe these methods are inappropriate in the present application. Thus, to the best of my knowledge, I employ a minimal set of assumptions to nonparametrically estimate the cost of collusion in British Columbia's timber auctions.

References

Aryal, G. and Gabrielli, M. F. (2013). Testing for Collusion in Asymmetric First-Price Auctions. *International Journal of Industrial Organization*, 31:26–35. 8

Athey, S., Coey, D., and Levin, J. (2013). Set-Asides and Subsidies in Auctions. *American Economic Journal: Microeconomics*, 5(1):1–27. 26

Athey, S. and Cramton, P. (2005). Competitive Auction Markets in British Columbia. MDI Report for British Columbia Ministry of Forests. 30

Athey, S., Cramton, P., and Ingraham, A. (2002). Auction-Based Timber Pricing and Complementary Market Reforms in British Columbia. MDI Report for British Columbia Ministry of Forests. 30

Athey, S. and Haile, P. A. (2002). Identification of Standard Auction Models. *Econometrica*, 70(6):2107–2140. 15, 19, 46, 69

Athey, S., Levin, J., and Seira, E. (2011). Comparing Open and Sealed Bid Auctions: Evidence from Timber Auctions. *Quarterly Journal of Economics*, 126(1):207–257. 8

Bajari, P. and Ye, L. (2003). Deciding Between Competition and Collusion. *The Review of Economics and Statistics*, 85(4):971–989. 9

Baldwin, L. H., Marshall, R. C., and Richard, J. F. (1997). Bidder Collusion at Forest Service Timber Sales. *Journal of Political Economy*, 105(4):657–699. 9

Berman, S. M. (1963). Note on Extreme Values, Competing Risks and Semi-Markov Processes. *The Annals of Mathematical Statistics*, 34(3):1104–1106. 46

Brannman, L. and Froeb, L. M. (2000). Mergers, Cartels, Set-Asides, and Bidding Preferences in Asymmetric Oral Auctions. *Review of Economics and Statistics*, 82(2):283–290. 9

Canada. B.C. Ministry of Forests (1999). *Interior Appraisal Manual*. 31

Conley, T. G. and Decarolis, F. (2016). Detecting Bidders Groups in Collusive Auctions. *American Economic Journal: Microeconomics*, 8(2):1–38. 9

Fibich, G. and Gavish, N. (2011). Numerical Simulations of Asymmetric First-Price Auctions. *Games and Economic Behavior*, 73(2):479–495. 65

Gill, R. D. (1989). Non- and semi-parametric maximum likelihood estimators and von mises method (part 1). *Scandinavian Journal of Statistics*, 16(2):97–128. 28

Guerre, E., Perrigne, I., and Vuong, Q. (2000). Optimal Nonparametric Estimation of First-Price Auctions. *Econometrica*, 68(3):525–574. 21, 23, 25, 50

Haile, P. A., Hong, H., and Shum, M. (2003). Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions. 26

Hickman, B. R. and Hubbard, T. P. (2015). Replacing Sample Trimming with Boundary Correction in Nonparametric Estimation of First-Price Auctions. *Journal of Applied Econometrics*, 30(5):739–762. 21

Hickman, B. R., Hubbard, T. P., and Paarsch, H. J. (2016). Identification and Estimation of a Bidding Model for Electronic Auctions. 67

Jofre-Bonet, M. and Pesendorfer, M. (2003). Estimation of a Dynamic Auction Game. *Econometrica*, 71(5):1443–1489. 34

Kaplan, T. R. and Zamir, S. (2012). Asymmetric first-price auctions with uniform distributions: Analytic solutions to the general case. *Economic Theory*, 50(2):269–302. 51, 67

Karunamuni, R. J. and Zhang, S. (2008). Some Improvements on a Boundary Corrected Kernel Density Estimator. *Statistics and Probability Letters*, 78(5):499–507. 21, 23, 52

Kawai, K. and Nakabayashi, J. (2014). Detecting Large-Scale Collusion in Procurement Auctions. 9

Krasnokutskaya, E. (2011). Identification and Estimation of Auction Models with Unobserved Heterogeneity. *The Review of Economic Studies*, 78(1):293–327. 42, 70

Krasnokutskaya, E. and Seim, K. (2011). Bid Preference Programs and Participation in Highway Procurement Auctions. *American Economic Review*, 101(6):2653–2686. 26

Lebrun, B. (2002). Continuity of the First Price Auction Nash Equilibrium Correspondence. *Economic Theory*, 20(3):435–453. 29

Lebrun, B. (2006). Uniqueness of the Equilibrium in First-Price Auctions. *Games and Economic Behavior*, 55(1):131–151. 12, 18, 68

Liang, H., Liu, X., Li, R., and Tsai, C. L. (2010). Estimation and testing for partially linear single-index models. *Annals of Statistics*, 38(6):3811–3836. 58

List, J. A., Millimet, D. L., and Price, M. K. (2007). Inferring Treatment Status When Treatment Assignment is Unknown: Detecting Collusion in Timber Auctions. 8, 30, 33, 36

Marmer, V., Shneyerov, A., and Kaplan, U. (2017). Identifying Collusion in English Auctions. 9, 25

Martin, E. C. and Betensky, R. A. (2005). Testing Quasi-Independence of Failure and Truncation Times via Conditional Kendall's Tau. *Journal of the American Statistical Association*, 100(470):484–492. 27, 35

McAfee, P. R. and McMillan, J. (1992). Bidding Rings. 82(3):579–599. 11

Paarsch, H. (1997). Deriving an Estimate of the Optimal Reserve Price: An Application to British Columbian Timber Sales. *Journal of Econometrics*, 78(2):333–357. 33, 63

Pesendorfer, M. (2000). A Study of Collusion in First-Price Auctions. *Review of Economic Studies*, 67(3):381–411. 9

Pinkse, J. and Schurter, K. (2019). Improved Bandwidth Selection for Boundary Correction using Generalized Reflection. 21

Pinkse, J. and Schurter, K. (2022). Estimation of auction models with shape restrictions. 22, 23, 24, 52

Porter, R. H. and Zona, J. D. (1993). Detection of Bid Rigging in Procurement Auctions. *Journal of Political Economy*, 101(3):518. 9, 24

Porter, R. H. and Zona, J. D. (1999). Ohio School Milk Markets: An Analysis of Bidding. *The RAND Journal of Economics*, 30(2):263. 9

Price, M. K. (2008). Using the Spatial Distribution of Bidders to Detect Collusion in the Marketplace: Evidence from Timber Auctions. *Journal of Regional Science*, 48(2):399–417. 8, 30, 33, 36

Roberts, J. W. (2013). Unobserved Heterogeneity and Reserve Prices in Auctions. *RAND Journal of Economics*, 44(4):712–732. 42, 70

Romano, J. P. and Wolf, M. (2010). Balanced control of generalized error rates. *The Annals of Statistics*, 38(1):598–633. 37

Stute, W. (1986). Conditional Empirical Processes. *Annals of Statistics*, 14(2):638–647. 51

Tsai, W.-Y. (1990). Testing the assumption of independence of truncation time and failure time. *Biometrika*, 77(1):169–177. 27, 35 van der Vaart, A. W. (1998). *Asymptotic Statistics.* Cambridge University Press, New York, NY. 52

A Propositions and Proofs

proof of Lemma 1. Let $M_i(b) = P\{B_i \leq b, \max_{j \neq i} B_j \leq B_i\}$ denote the probability of the event that bidder i wins the auction and the price is less than or equal to b.[29] If i is a member of the bidding ring, this event is equivalent to the event

$$\left\{\sigma_i(V_i) \leq b,\, \max_{j \notin \mathcal{R}} \sigma_j(V_j) \leq \sigma_i(V_i),\, \max_{k \neq i \in \mathcal{R}} v_k \leq v_i\right\},$$

where $\sigma_i(v_i)$ denotes the bid that a collusive bidder would have submitted if the other colluders' valuations had been less than v_i. Otherwise, when i is not in the ring, it is equivalent to the event

$$\left\{\sigma_i(V_i) \leq b,\, \max_{j \neq i \notin \mathcal{R}} \sigma_j(V_j) \leq \sigma_i(V_i),\, \max_{k \in \mathcal{R}} \sigma_k(V_k) \leq \sigma_i(V_i)\right\}.$$

In either case, because each member of the ring uses the same increasing bidding strategy, both events are equivalent to

$$\left\{\sigma_i(V_i) \leq b,\, \max_{j \neq i} \sigma_j(V_j) \leq \sigma_i(V_i)\right\}.$$

Under the assumption that the valuations are mutually independent, the functions M_i for $i \in \mathcal{N}$ can then be used to construct S_i, the marginal distribution of $\sigma_i(V_i)$. In words, S_i is the distribution of the bid that bidder i would submit if it were bidding competitively.

$$S_i(b) = \exp\left\{-\int_b^{\bar{b}} \frac{1}{M}\, dM_i\right\}, \tag{13}$$

where $M = \sum_i M_i$ is the distribution of auction prices. An analogous construction was proven by Berman (1963) and was introduced to the empirical auction literature by Athey and Haile (2002). This proves (i).

Under the null hypothesis that i is not colluding, the distribution of the

[29]The probability that there is a tie in the winning bid is zero, except possibly at a price equal to the minimum bid in the case where the minimum bid is an atom of the bidders' valuations distributions. In that event, ties are broken randomly.

highest bid among i's competitors is

$$G_i = \prod_{j \neq i} S_j \,, \tag{14}$$

Plugging (13) into (14), the inverse strategy can be written as

$$\sigma_i^{-1}(b) = b + \frac{M(b)}{\frac{\partial M_{-i}(b)}{\partial b}} \,,$$

where $M_{-i} = \sum_{j \neq i} M_j$. By proposition 1, σ_i^{-1} is increasing in b for $b \geq \underline{b}$. Therefore, identification of F_i on $[\sigma_i^{-1}(\underline{b}), \bar{v}_i]$ follows from

$$F_i(\sigma^{-1}(b)) = S_i(b) \,. \tag{15}$$

This proves (ii). □

proof of Lemma 2. The "if" direction follows immediately from Lemma 1 and IA.2. The proof of the "only if"direction mirrors Figure 1. First, I derive an expression for the horizontal distance between the competitively rationalizing distribution (dashed curves) and the true distribution (dotted curve) in Figure 1. In keeping with part (ii) of the lemma, this horizontal difference is zero for all values of the instrument if the bidder is not colluding. Otherwise, the collusive bidders' true valuation distribution will stochastically dominate the competitively rationalizing distribution.

To see this, I let $\hat{F}_i$, $\hat{G}_i$ and $\hat{g}_i$ denote the valuation distribution and competing distribution and density that are inferred under the null hypothesis that i is not colluding. I also suppress the argument and the conditioning on Z and let $\mathbf{1}$ denote the identity mapping $b \mapsto b$.

$$
\begin{aligned}
F_i^{-1} - \hat{F}_i^{-1} &= \left(\mathbf{1} + \frac{G_i}{g_i}\right) \circ S_i^{-1} - \left(\mathbf{1} + \frac{\hat{G}_i}{\hat{g}_i}\right) \circ S_i^{-1} \\
&= S_i^{-1} + \frac{G_i(S_i^{-1})}{g(S_i^{-1})} - \left(S_i^{-1} + \frac{\hat{G}_i(S^{-1})}{\hat{g}(S_i^{-1})}\right) \\
&= \frac{1}{\frac{g_i(S_i^{-1})}{G_i(S_i^{-1})}} - \frac{1}{\frac{g_i(S_i^{-1})}{G_i(S_i^{-1})} + \sum_{j \neq i \in \mathcal{R}} \frac{f_j(b + \frac{G_i}{g_i})}{F_j(b + \frac{G_i}{g_i})}\Big|_{S_i^{-1}} \cdot \frac{\partial}{\partial b}\left(b + \frac{G_i(b)}{g_i(b)}\right)\Big|_{S_i^{-1}}} \\
&= \frac{1}{1 \Big/ \left(F_i^{-1} - \left(\mathbf{1} + \frac{G_i}{g_i}\right)^{-1}\Big|_{F_i^{-1}}\right)} \\
&\quad - \frac{1}{1 \Big/ \left(F_i^{-1} - \left(\mathbf{1} + \frac{G_i}{g_i}\right)^{-1}\Big|_{F_i^{-1}}\right) + \sum_{j \neq i \in \mathcal{R}} \frac{f_j(F_i^{-1})}{F_j(F_i^{-1})} \Big/ \frac{\partial}{\partial v}\left(\left(\mathbf{1} + \frac{G_i}{g_i}\right)^{-1}(v)\right)\Big|_{F_i^{-1}}} .
\end{aligned}
$$

The first equality holds by the GPV equation. The S_i^{-1} cancel in the second line. In the third line, I use independence of the valuations to decompose $\hat{g}_i/\hat{G}_i$ into the reverse hazard rate of i's true competing distribution and the reverse hazard rates of the other colluders' bids. I also use the fact that the ring members would use the same inverse strategy function $(\mathbf{1} + G_i/g_i)$ if they were submitting a serious bid. The last equality again holds by the GPV equation (15).

This final expression only depends on the true valuation distribution and serious strategy function. And, because $\sigma_i = (\mathbf{1} + \frac{G_i}{g_i})^{-1} < \mathbf{1}$ and f_j/F_j is greater than zero, it implies that $F^{-1} - \hat{F}^{-1}$ is positive. Moreover, this difference is decreasing in both the slope and the level of σ_i evaluated at the quantile of i's true valuation distribution. That is, the competitive model more severely underestimates the valuation distribution when the colluders bid less aggressively in the sense that the slope and level of their strategy function is lower.

I next establish that the changes in the slope and level of a colluder's strategy function cannot exactly offset everywhere. Relative to Figure 1, this argument shows that the competitively rationalizing distribution (dashed line) must move relative to the true valuation distribution (dotted line) in response to variation in the instrument. Because IA.2 states that the dotted line is constant with respect to the instrument, the dashed lines cannot coincide for all values of Z.

More precisely, I argue that there exists a valuation, v, and realizations of the instrument, z and z', such that the horizontal difference between the distribution functions is unequal at v, i.e.

$$F_i^{-1}(F_i(v)) - \hat{F}_i^{-1}(F_i(v)\,|\,Z_i = z) \neq F_i^{-1}(F(v)) - \hat{F}_i(F_i(v)\,|\,Z_i = z')\,,$$

which implies that $\hat{F}_i(\cdot\,|\,Z_i = z) \neq \hat{F}_i(\cdot\,|\,Z_i = z')$. The existence of such a v is easy to verify near the minimum bid. To that end, I consider the following cases.

Case i. The instrument does not affect the minimum serious bid. By IA.1, bidder i's strategy function is a nontrivial function of its instrument. And by Proposition 1, $\sigma_i(\underline{b}; z) = \underline{b}$ for all z. Therefore, there is a "first" point, v_0, where strategies $\sigma_i(\cdot\,; z)$ and $\sigma_i(\cdot\,; z')$ diverge.[30] Then there is a point v_1 in the neighborhood of v_0 such that the slope and level of one of the strategies is greater than the other, e.g. the equilibrium strategy when $Z_i = z'$ is more agressive at v_1 than when $Z_i = z$. Therefore, $\hat{F}_i^{-1}(F_i(v_1)|Z_i = z') < \hat{F}_i^{-1}(F_i(v_1)|Z_i = z)$.

Case ii. The instrument affects the minimum serious bid. Let $\underline{b}$ and $\underline{b}'$ denote the minimum bids for $Z = z$ and z'. Assume without loss of generality that $\underline{b} < \underline{b}'$. Then $\hat{F}_i(\underline{b}'|Z_i = z) > F_i(\underline{b}')$ because the competitive model strictly underestimates (in the FOSD-sense) the true valuation distribution, except at the minimum bid, where $\hat{F}_i(\underline{b}'|Z_i = z') = F_i(\underline{b}')$ whenever i's strategy is strictly increasing at $\underline{b}'$. This implies that $\hat{F}_i(\underline{b}'|Z_i = z') < \hat{F}_i(\underline{b}'|Z_i = z)$.

If i's strategy is not strictly increasing at $v = \underline{b}'$, then Proposition 1 implies that i will submit a bid exactly equal to $\underline{b}'$ with positive probability. If i is not colluding, then part (i) of Proposition 1 implies that bidder i must be the only bidder with an atom in its bid distribution. In contrast, if i is colluding then all ring members will optimally use the same strategy function, and multiple bidders' bid distributions will have an atom at $\underline{b}'$. In this case, the the competitive IPV model cannot reationalize the observed bid distributions and the colluders would be identified even without exploiting variation in the level of competition at the auction. □

proof of Theorem 2. By part (i) of the lemma, the marginal distribution of bidder i's serious bidding strategy, σ_i, is identified for each i. By part (ii) of the lemma, bidder i's competitively rationalizing distribution will be equal to its true distribution for any value of Z. The competitively rationalizing valuations

[30] More precisely, let v_0 denote the infimum of the set of v on which $\sigma_i(\cdot\,; z) - \sigma_i(\cdot\,; z') = 0$.

are therefore independent of the instrument. To establish identification of the collusive model, I prove that the converse statement is also true: a colluder's competitively rationalizing valuation distribution must depend on Z. Thus, a test for independence between the competitively rationalizing valuations and the instrument can be used to detect collusion. Once all of the colluders have been identified in this way, each of the bidders' true competing distributions can be computed as $\prod_{j\neq i} S_j$ for $i \notin \mathcal{R}$ and $\prod_{j\notin \mathcal{R}} S_j$ for $i \in \mathcal{R}$. The true valuation distributions are then given by (15).

□

proof of Lemma 3. The equililbrium bid distributions solve the initial value problem in equation 16. By standard arguments for differential equations, the solution at $b > r$ is continuously differentiable in the parameters (r, z) by SA.1. Moreover, as noted by Guerre et al. (2000), the bid densities must have one more derivative than the valuation densities. In the differential system of equations, this is apparent from the fact that the bid densities appear on one side of the equation while the quantiles of the valuations are on the other. The expected payment function, $e_i(p; r, z, \mathcal{N}) = pG_i^{-1}(p|r, z, \mathcal{N})$, inherits its smoothness from the competing quantile function G_i^{-1}. Unlike the bid density which is zero below the reserve price and unbounded above the reserve price, the slope of e_i is equal to the reserve price r at $p = G_i(r|r, z, \mathcal{N})$. The expected payment function is technically undefined at $p < G_i(r|r, z, \mathcal{N})$ because there is no bid that a bidder could submit and expect to win with probability $p < G_i(r|r, z)$ in equilibrium. However, one can continously extend e_i and its first derivative by defining $e_i(p; r, z, \mathcal{N}) = pr$ for $0 \leq p < G_i(r)$. Finally, the slope of e is bounded because it is equal to the bidder's valuation, which are assumed to be bounded under MA.2. □

Proposition 2. *Under MA.1–MA.5 and SA.1–SA.2, the conditional density of the transformed serious bids* $\tilde{B} = \sqrt{B - r}$ *is bounded on* $[r, \bar{b}]$ *given* z *and* r *for all but at most one bidder.*

Proof. Fix z and r and suppress the dependence of the equilibrium bid distributions on these covariates. For simplicity, assume all n bidders in $\mathcal{N}$ are serious bidders. To prove the more general statement, one would repeat the argument below after replacing $\mathcal{N}$ with the set of serious bidders and replace the serious cartel bidders' valuation distribution with its valuation distribution conditional on being the designated serious cartel bidder.

The system of differential equations that characterizes the transformed marginal equilibrium serious bid distributions is given in equation (18). Summing (18) accross i and dividing by $n-1$

$$\sum_j \frac{\partial \tilde{S}_j(\tilde{b})}{\partial \tilde{b}} = \frac{1}{n-1} \sum_j \frac{2\tilde{b}}{F_j^{-1}(\tilde{S}_j(\tilde{b})) - \tilde{b}^2 - r}$$

and

$$\frac{\partial \log \tilde{S}_i(\tilde{b})}{\partial \tilde{b}} = \frac{1}{n-1} \sum_{j \neq i} \frac{2\tilde{b}}{F_j^{-1}(S_j(\tilde{b})) - \tilde{b}^2 - r} - \frac{n-2}{n-1} \frac{2\tilde{b}}{F_i^{-1}(S_j(\tilde{b})) - \tilde{b}^2 - r}, \tag{16}$$

for all i. Dividing the first equation by (18) by the first equation above yields

$$\frac{\frac{\partial \tilde{S}_i(\tilde{b})}{\partial \tilde{b}}}{\sum_{j \neq i} \frac{\partial \tilde{S}_j(\tilde{b})}{\partial \tilde{b}}} = \frac{1}{n-1} \sum_{j \neq i} \frac{F_i^{-1}(\tilde{S}_i(\tilde{b})) - \tilde{b}^2 - r}{F_j^{-1}(\tilde{S}_j(\tilde{b})) - \tilde{b}^2 - r} - \frac{n-2}{n-1}, \tag{17}$$

for all i.

Suppose that there exists a bidder i for whom $\frac{\partial \log \tilde{S}_i(\tilde{b})}{\partial \tilde{b}}$ is unbounded as $\tilde{b}$ tends toward zero. Then the right side of equation (18) diverges for all $j \neq i$, hence $\tilde{b}/(F_j^{-1}(\tilde{S}_j(\tilde{b})) - \tilde{b}^2 - r)$ diverges as $\tilde{b}$ approaches zero. Note that there can be at most one bidder for i for whom $\frac{\partial \log \tilde{S}_i(\tilde{b})}{\partial \tilde{b}}$ is unbounded, because the left side of equation (17) cannot be unbounded for two different bidders.

As an example of an equilibrium in which one bidder's bid density diverges at a strictly faster rate, consider the special case of two bidders with uniformly distributed valuations. Using the analytical solutions to the equilibrium derived in Kaplan and Zamir (2012), the slope of the inverse bid function is found to be proportional to $(b-r)^{\theta_i - 1}$ where $\theta_i = \frac{r - \underline{v}_i}{2r - \underline{v}_1 - \underline{v}_2}$. If $\underline{v}_1 = \underline{v}_2$, then $\theta_1 = \theta_2 = 1/2$, and both bidder's bid densities diverge at a rate of $1/\sqrt{b-r}$. If $\underline{v}_1 > \underline{v}_2$, then $\theta_1 = 1 - \theta_2 > 1/2$, and bidder 1's bid density diverges at a faster rate. □

proof of Theorem 3. Because the bids are bounded and absolutely continuous under the smoothness conditions assumed in order to prove existence of the equilibrium, the conditional process $\left(\frac{T}{\log T}\right)^{2/(4+d)} (\mathbb{M}_T(\cdot|\, z, r, \mathcal{N}) - M_T(\cdot|\, z, r, \mathcal{N}))$ and the cumulative incidence processes converge to centered Gaussian processes as a consequence of Theorem 1 in Stute (1986) and the remarks that follow on page 892. Stute (1986) also shows that the processes converge uniformly in the conditioning variables under the additional smoothness condition on the functions $M_i(b\,|\cdot)$. The cumulative distribution functions are not differentiable at bids equal to the reserve price, so the uniform convergence holds on compact intervals that do not include the reserve price.

The mapping $(M_i, M) \mapsto \exp\{-\int_{\cdot}^{\bar{b}} \frac{1}{M}\, dM_i\}$ is Hadamard differentiable by Lemma 20.10 in van der Vaart (1998) whenever M is bounded away from zero, as will be the case for all $b \geq r$ if the reserve price r is binding.

If one uses the estimator $\tilde{v}_{it,\emptyset}$ under the assumption that all bid densities diverge at the same rate at the lower extremity of their support, a boundary-corrected kernel estimator similar to Karunamuni and Zhang (2008) reduces the order of the bias to $O(h_m^2)$ over the entire support. The densities of the cumulative incidence functions, m_{iT} for each i, then converge uniformly at the rate $(T/\log)^{2/(5+d)}$ and pointwise to Gaussian limit by standard arguments for kernel density estimation. Because the density of the valuations is bounded away from zero and the inverse strategy function is assumed to be strictly increasing, the bid densities are also be bounded away from zero. The estimated inverse strategy function defined by the bidder's first-order condition therefore converges pointwise to a Gaussian limit and uniformly by the continuous mapping theorem applied (m_{-iT}, M_T) with the mapping $(m_{-i}, M) \mapsto \underline{1} + M/m_{-i}$, where $\underline{1}(b) = b$ is the identity function. Let σ_{iT}^{-1} and σ_i^{-1} denote this estimator and its limit in probability, respectively.

If one uses the estimator based on the expected payment function, Pinkse and Schurter (2022) prove as an intermediate step toward establishing the limiting behavior of their monotonized and smoothed estimator of the derviative of the expected payment function that smoothed, non-monotonized estimators such as the one in equation 10 are asymptotically equivalent. Because $e''(p; r, z, \mathcal{N})$ is bounded when the reserve price binds (Pinkse and Schurter, 2022), and G_{iT} converges to G_i at a faster rate than the estimator for the slope of the expected payment function, the pointwise and uniform rates of convergence of the slope of the expected payment function are the same as those of the inverse bidding strategy.

In either case, because the inverse strategy function converges at a slower rate than $\mathbb{S}_{iT}$, the dominant term in the asymptotic expansion of $\mathbb{F}_{iT}(v) - F_i(v)$ is $S_i(\sigma_{iT}(v)) - S_i(\sigma_i(v))$. An application of the delta method then implies $\mathbb{F}_{iT}(v) - F_i(v)$ is equal to $s_i(\sigma_v)(\sigma_{iT}(v) - \sigma_i(v))$ plus terms that converge at a faster rate, which yields the promised result since the serious bid density s_i is bounded on compact intervals that do not include the reserve price. $\square$

Proposition 3. *Under SA.1 and SA.2, the conditional Kendall's τ parameter is zero under the null hypothesis and its estimator is asymptotically normal and asymptotically linear if psuedo-valuations are estimated with second-order kernels and the kernel bandwidths are $o(T^{-1/4})$.*

Proof. The censored version of the conditional Kendall's τ is equal the fraction of orderable pairs that are concordant minus the fraction that are discordant. A pair $\{(v_{i1}, r_1), (v_{i2}, r_2)\}$ is orderable if $\max\{v_{i1}, v_{i2}\} > \min\{r_1, r_2\}$. Under the null hypothesis that v_{it} and r_t are independent and (v_{it}, r_t) is IID across auctions, the probability that $\max\{v_{i1}, v_{i2}\} = v_1$ or v_2 is the same conditional on the event that at least one of them is greater than $\min\{r_1, r_2\}$. Thus, an orderable pair is equally likely to be concordant as it is to be discordant, and the conditional Kendall's τ is zero.

Suppressing conditioning on any exogenous covariates, the population parameter can be written as

$$\tau_i^C = \frac{2\int F_i^-(v,r)\,\mathbb{1}\{r < v\}\,dF_i(v,r)}{\int F_i^-(v,v)\,\mathbb{1}\{r < v\}\,dF_i(v,r)} - 1$$

where F_i and F_i^- denote the right- and left-continuous versions of the joint distribution of $(V_{i,\emptyset} \vee r, r)$. The test statistic is numerically equivalent to the sample analog of the above equation in which the joint distribution is replaced by the joint empirical distribution of $\{(\sigma_{iT}^{-1}(b_{it}), r_t)\}_{t=1,\ldots,T}$. The numerator in the first term on the right is equal to the proportion of concordant, orderable pairs, while the denominator is equal to the proportion of orderable pairs.

The influence function of the conditional Kendall's tau estimator is

$$\psi_{F_i}(v,r) = \frac{2P_{F_i}\{(v_{it}-v)(r_t-r) > 0, \max\{r_t, r\} < \max\{v_{it}, v\}\} - (\tau_i^C + 1)P_{F_i}\{\max\{r_t, r\} < \max\{v_{it}, v\}\}}{\int F_i^-(v,v)\mathbb{1}\{r < v\}\,dF_i(v,r)}.$$

In words, the influence function evaluated at (v, r) is the twice the probability under F_i that the pair of observations $\{(v_{it}, r_t), (v, r)\}$ is concordant and orderable minus $(\tau + 1)$ times the probabilitiy the pair is orderable, divided by the probability that a random pair of independent observations drawn from F_i is orderable.[31] The numerator in the above expression for the influence function has a mean of zero and is bounded between two and negative two. The denominator is bounded away from zero because there is a strictly positive probability

[31] For comparison, the usual unconditional Kendall's τ statistic has an influence function given by $4P_{F_i}\{(v_{it} - v)(r_t - r)\} - 2\tau_i - 2$.

that two randomly selected observations are orderable, i.e. that the maximum of the two reserve prices is less than the maximum of the valuations. If valuations were observed, the central limit theorem would then imply asymptotic normality of $\sqrt{T}\int \psi_{F_i}(v,r)\,dF_{iT}(v,r)$. Hence, the estimator for τ_i^C is asymptotically linear as long as the remainder is negligible:

$$\hat{\tau}_i^C - \tau_i^C - \int \psi_{F_i}(v,r)\,dF_{iT}(v,r) = \frac{2}{\left(\int F_i^-(v,v)\mathbb{1}\{r<v\}\,dF_i(v,r)\right)^2 \int F_{iT}^-(v,v)\mathbb{1}\{r<v\}\,dF_i(v,r)} \cdot$$

$$\Bigg[\int F_i^-(v,r)\mathbb{1}\{r<v\}\,dF_i(v,r)\left(\int F_{iT}^-(v,v)\mathbb{1}\{r<v\}\,dF_{iT}(v,r) - \int F^-(v,v)\mathbb{1}\{r<v\}\,dF_i(v,r)\right)^2$$

$$-\int F_i^-(v,v)\mathbb{1}\{r<v\}\,dF_i(v,r)\left(\int F_{iT}^-(v,v)\mathbb{1}\{r<v\}\,dF_{iT}(v,r) - \int F^-(v,v)\mathbb{1}\{r<v\}\,dF_i(v,r)\right)\cdot$$

$$\left(\int F_{iT}^-(v,r)\mathbb{1}\{r<v\}\,dF_{iT}(v,r) - \int F^-(v,r)\mathbb{1}\{r<v\}\,dF_i(v,r)\right)\Bigg].$$

Indeed, if the valuations were observed, the remainder is $O_p(1/T)$ because $\int F_{iT}^-(v,r)\mathbb{1}\{r<v\}\,dF_{iT}(v,r)$ and $\int F_{iT}^-(v,v)\mathbb{1}\{r<v\}\,dF_{iT}(v,r)$ converge at the $1/\sqrt{T}$-rate.

Because the valuations are not observed, the bias in the kernel-based estimates of the pseudo-valuations must be $o(1/\sqrt{T})$ so that the bias in these functionals is also $o(1/\sqrt{T})$. In addition, $\frac{1}{\sqrt{T}}\sum_t \psi_{F_i}(\sigma_{iT}^{-1}(b_{it};r_t),r_t)$ is asymptotically normal because

$$\frac{1}{\sqrt{T}}\sum_t \psi_{F_i}(\sigma_{iT}^{-1}(b_{it};r_t),r_t) \approx \frac{1}{\sqrt{T}}\sum_t \psi_{F_i}(\sigma_i^{-1}(b_{it};r_t),r_t) +$$
$$\frac{1}{\sqrt{T}}\sum_t \frac{\partial \psi_{F_i}(\sigma_i^{-1}(b_{it};r_t),r_t)}{\partial v}\left(\sigma_{iT}^{-1}(b_{it};r_t) - \sigma_i^{-1}(b_{it};r_t)\right).$$

The first term is the dominant term in the asymptotically linear expansion when the valuations are observed. The second term is a linear combination of pseudo-valuations. Because the pseudo-valuations have an asymptotically linear representation, one can derive an influence function for the estimator for τ_i^C based on the undersmoothed pseudo-valuations. If one uses the estimator for the pseudo-valuations based on the expected payment function, the Hadamard derivative of the mapping from (M_{-i}, M) to e_i is a mapping from a pair of functions (h_{-i}, h) to a function given by the expression

$$\frac{p^2}{g_i(G_i^{-1}(p))}\left(\int_{G_i^{-1}(p)}^{\bar{b}} \frac{dh_{-i}}{M} - \int_{G_i^{-1}(p)}^{\bar{b}} \frac{h}{M^2}\, dM_{-i}\right)$$

for p between $G_i(r)$ and one. Letting $(h_{-i}, h) = (M_{-iT} - M_{-i}, M_T - M)$, the linear combination of the pseudo-valuations that rationalize all of bidder i's bids is therefore given by

$$\sum_t \frac{1}{h^2}\int \frac{\partial \psi_{F_i}(\sigma_i^{-1}(b_{it}; r_t), r_t)}{\partial v}(e_{iT}(p; r_t) - e_i(p; r_t))K'\left(\frac{G_{iT}(b_{it}|r_t) - p}{h}\right) dp + O(h^2)$$
$$\approx \frac{1}{T}\sum_t\sum_s \frac{\partial \psi_{F_i}(\sigma_i^{-1}(b_{it}; r_t), r_t)}{\partial v}\int \frac{p^2}{g_i(G_i^{-1}(p|r_t))}\frac{1}{h^2}K'\left(\frac{G_{iT}(b_{it}|r_t) - p}{h}\right)$$
$$\left(\frac{\mathbb{1}\{b_{is} < \max_{j\neq i} b_{js}\}\mathbb{1}\{G_i^{-1}(p|r_t) \leq b_{is}\}w_s}{M(b_{is}|r_t)\sum_l w_l} - \int_{G_i^{-1}(p|r_t)}^{\bar{b}} \frac{\mathbb{1}\{\max_{j\neq i} b_{js} \leq x\}\mathbb{1}\{b_{is} \leq x\}w_s}{M(x|r_t)^2 \sum_l w_l}\, dM_{-i}(x|r_t)\right) dp$$

where the $O(h^2)$ term in the first expression is the bias introduced by the kernel smoothing and w_s is the kernel-based weight used to smooth over the reserve prices.[32] The above expression is asymptotically linear because it is a double sum over a product of functions of the observations (b_{it}, r_t). The influence function involves many terms but is otherwise routine to derive under the null hypothesis because b_{it} and b_{jt} for $i \neq j$ are independent conditional on r_t.

If the estimator based on the competing density is used, a similar expression is obtained. In either case, the estimators which are asymptotically linear and converge at a $\sqrt{T}$-rate to a Gaussian limiting distribution with a mean of zero when the kernel bandwidth is $o(T^{-1/4})$.

By an analogous argument, the truncated version of the conditional Kendall's τ is asymptotically linear and normally distributed about zero under the null hypothesis. The test statistic is

$$\hat{\tau}_i^{\text{Trunc}} = \frac{2\sum_t\sum_s \mathbb{1}\{v_s < v_t\}\mathbb{1}\{r_s < r_t\}\mathbb{1}\{r_t < v_s\}}{\sum_t\sum_s \mathbb{1}\{r_s < r_t\}\mathbb{1}\{r_t < v_s\}\mathbb{1}\{r_t < v_t\}} - 1$$
$$= \frac{2\sum_t\sum_s \mathbb{1}\{v_s < v_t\}\mathbb{1}\{r_s < r_t\}\mathbb{1}\{r_t < v_t\} - \mathbb{1}\{v_s < v_t\}\mathbb{1}\{r_s < r_t\}\mathbb{1}\{r_t < v_t\}\mathbb{1}\{v_s \leq r_t\}}{\sum_t\sum_s \mathbb{1}\{r_s < r_t\}\mathbb{1}\{r_t < v_t\} - \mathbb{1}\{r_s < r_t\}\mathbb{1}\{r_t < v_t\}\mathbb{1}\{v_s \leq r_t\}} - 1$$

[32] Note that the kernel K must be a boundary kernel or the definition of e_{iT} and e_i and their derivatives need to be smoothly extended beyond the unit interval in order to reduce the bias from $O(h)$ to $O(h^2)$. If one uses a boundary kernel, the kernel is more properly written as a two-argument function $K(u, p_0)$, where p_0 is the point at which the slope of the expected payment function is being estimated. The derivative K' should then be interpreted as the total derivative of $K((p_0 - p)/h, p_0)$ with respect to p_0.

Because there are no atoms in the valuation distribution, we can ignore the distinction between $\mathbb{1}\{v_s \leq r_t\}$ and $\mathbb{1}\{v_s < r_t\}$, and the population parameter is now given by

$$\tau_i^{\text{Trunc}} = \frac{2\int (F_i^-(v,r) - F_i^-(r,r))\mathbb{1}\{v > r\}\, dF_i(v,r)}{\int (F_i^-(\infty,r) - F_i^-(r,r))\mathbb{1}\{v > r\}\, dF_i(v,r)} - 1\,.$$

The rest of the proof repeats the same steps as above. □

B Simulation Results

To assess the finite sample performance of the proposed test for collusion, I simulate auctions in which two colluders compete against a single non-ring bidder. In order to simulate samples in which the reserve price is continuously distributed, I must solve for the equilibrium bidding functions at each of many different reserve prices. Because numerically solving for these functions can be computationally expensive, I restrict attention to a special case in which the equilibrium has an analytic solution. In particular, I simulate auctions in which the non-ring bidder's valuation is uniformly distributed on $[0, 1]$, and the maximum of the ring's valuations is uniformly distributed on $[0, 1.8]$. For simplicity, I choose to make the ring members symmetric, i.e. their valuations are distributed according to $v^{1/2}$ on $[0, 1]$. The reserve price for each auction is independently drawn uniformly at random from $[0.1, 0.5]$ and all bidders use their optimal serious bidding strategy. Under this simulation design, each bidder loses with approximately equal probability, so that the number of observations used to estimate each bidder's competing distribution is roughly the same.

In the simulation results reported in Table 4, I use a bandwidth sequences, $h \propto T^{-1/3.9}$, to calculate $\hat{\tau}_i^C$.[33] I then draw T auctions with replacement and calculate the bootstrapped test statistic τ_i^{C*} many times to estimate an equal-tailed confidence interval for $\sqrt{T}\,\left(\hat{\tau}_i^{C*} - \hat{\tau}_i^C\right)$. I then reject the null hypothesis at the α level of significance if $\sqrt{T}\,\hat{\tau}_i^C$ is outside the estimated confidence interval. The fraction of tests for which the null hypothesis is rejected for the non-ring and ring bidders are reported in the size and power columns of Table 4.

Previously studied collusion detection methods are either invalid or have power equal to size under this model. As a benchmark, albeit an infeasible one,

[33] The bandwidth used to smooth over reserve prices is $2\sigma_r T^{-1/3.9}$ where σ_r is the standard deviation of the reserve prices. I use a bandwidth of $1.5T_i^{-1/3.9}$ where T_i is the number of auctions in which bidder i lost.

I also test the null using the bidders' true competitively rationalizing valuations. For the competitive bidder, the true competitively rationalizing valuations are simply equal to the bidder's valuation. For a collusive bidder, I compute the valuation that would rationalize each of its bids as a best response to the true distribution of the other bidders' bids in the collusive equilibrium. This "oracle" test statistic is a U-statistic, which means it is $\sqrt{T}$-consistent, asymptotically normal, and unbiased with a relatively simple analytical expression for the variance that is routine to estimate. As expected, the bootstrap procedure and the Guassian approximation to the oracle statistic both produce critical values that control the size of the test. The discrepancy in the power of the tests indicates the loss in power due to the fact the inverse strategy functions are nonparametrically estimated in a first stage.

Table 4: Results are based on 1,000 simulations. Critical values for the nonparametric test are estimated from 1,500 bootstrap samples. The size and power of the test are reported in the last four columns. The nominal size of each test is 10%.

	Bootstrap		Oracle	
Auctions, T	Size	Power	Size	Power
200	0.092	0.094	0.070	0.137
500	0.084	0.198	0.072	0.241
1000	0.085	0.327	0.084	0.414

C Homogenizing the Bids and Reserve Prices

I assume auction-level heterogeneity enters bidders' payoffs additively separably through a linear combination of the observable characteristics:

$$u_{it} = \beta' x_t + v_{it} \,,$$

where u_{it} is bidder i's payoff from winning auction t and v_{it} is the unobservable idiosyncratic component of i's payoff. If v_{it} is independent of x_t, then the following regression equation will hold:

$$\tilde{b}_{it} = \beta' x_t + \eta(\tilde{r}_t - \beta' x_t, \mathcal{N}_t) + \epsilon_{it} \,,$$

where

$$\epsilon_{it} = v_{it} - \eta(\tilde{r}_t - \beta' x_t, \mathcal{N}_t) - \frac{G_i}{g_i}(v_{it}|\tilde{r}_t - \beta' x_t, \mathcal{N}_t)$$

is an independent, mean-zero error term with variance $\sigma^2(\tilde{r}_t - \beta' x_t, \mathcal{N}_t)$ and η is an unknown function of the screening level—the threshold below which idiosyncratic valuations will be censored by the reserve price—and the set of participants in the same market as i. If the reserve price were never binding, then none of the bids would be censored and the above equation could be estimated via OLS with fixed effects for $\mathcal{N}_t$. On the other hand, if the reserve price binds, bidder i's competing distribution in auction t will be correlated with x_t to the extent that x_t is correlated with the screening level. The OLS estimates will therefore suffer from an omitted variable bias.

To account for the fact that the reserve price is often binding in SBFEP auctions, I use a modified version of the typical partially linear single-index model (PLSIM) (see, for example, Liang et al., 2010). The estimator, $\hat{\beta}$, minimizes the sum of squared residuals in a local linear regression of $\tilde{p}_t - \beta' x_t$ on the scalar index $\tilde{r}_t - \beta' x_t$, where $\tilde{p}_t$ is the observed price at auction t. Typically, however, the covariates that enter linearly are not constrained to be the same as those that enter nonlinearly through η. Additionally, because PLSIMs generally do not allow the function η to depend on a categorical variable, I estimate a separate local linear regression for each market, while constraining $\hat{\beta}$ to be constant across markets. Thus, my estimator is the same as in Liang et al. (2010), except I impose the obvious equality constraints across markets and on the coefficients in the linear and nonlinear components of the regression. Consequently, the restricted estimator for β inherits the nice asymptotic properties of the unrestricted estimator. In particular, $\sqrt{T}\left(\hat{\beta} - \beta\right)$ converges to a centered Gaussian distribution.

For the local linear regression estimator of η, I use a quartic kernel to ensure differentiability of the least-squares objective function. To improve the performance in sparse regions on the data, I also include ridge regression parameters, $\lambda_{\mathcal{N}} > 0$, for each market. Both the kernel bandwidths and the ridge regression parameters were selected to minimize the leave-one-out cross-validated mean integrated squared error at $\beta = \hat{\beta}$,

$$\sum_t \left(\tilde{p}_t - \beta' x_t - \hat{\eta}_{-t}(\tilde{r}_t - \beta' x_t, \mathcal{N}_t)\right)^2 \Big|_{\beta = \hat{\beta}}$$

where $\hat{\eta}_{-t}$ is the local linear ridge regression estimator computed from all auctions excluding auction t.

Descriptions and summary statistics for the variables are in Tables 5 and 6. Estimates for β are reported in Table 7. For comparison, the first column reports the OLS estimates from a regression that includes market fixed effects for ten of the eleven clusters defined in the next section. These estimates will be biased to the extent that the reserve price is binding. The PLSIM includes the covariates that the Ministry of Forests in its original hedonic pricing formula, as well as the estimated fractions of cedar and white pine volume, which were added to the model when the Ministry amended the pricing formula in 2001, immediately after the auctions in my data had been held. Originally, the Ministry used a binary variable to indicate whether 60% or more of the estimated volume came from a combination of hemlock or balsam trees, but later replaced this with a binary indicator for whether the estimated combined volume of hemlock, balsam, and cedar was greater than 50% of the total estimated volume. I opt to include these as separate continuous variables. The fitted values from the PLSIM differ substantially from the Ministry's own appraisal. Under the non-hedonic appraisal regime, the linear correlation between the Ministry's estimate of the license value and mine is only 0.6. Using the hedonic pricing formula, the correlation increases to 0.806.

Figure 2 plots the winning bonus bid against the estimated appraisal price from Model 3 minus the reserve price, $\hat{\beta}'x_t - \tilde{r}_t$.[34] As $\hat{\beta}'x_t - \tilde{r}_t$ increases, the screening level decreases and bidders are more likely to participate with a wider range of valuations. Consequently, the variance in the winning bonus bid increases. In contrast, at low values of $\hat{\beta}'x_t - \tilde{r}_t$, the bonus bids are predictably smaller and have noticeably smaller variance. Thus, the general upward trend and increasing variance in Figure 2 suggests that the estimated appraisal price accurately summarizes variation in the covariates that is relevant to bidders' valuations.[35] In addition, the wide range in the observed screening levels indicates that there is substantial variation with which it will be possible to test for collusion.[36]

[34] Any constant term in the hedonic pricing formula is not identified separately from the level of η. To make Figure 2, I normalize the level of the appraisal so that the overall mean of $\hat{\beta}'x_t - \tilde{r}_t$ is equal to the mean winning bonus bid.

[35] In contrast, before 1999, there is a negative rank correlation between the screening level implied by the Ministry's estimate of the value of the license and the winning bonus bid. After the Ministry adopted the hedonic formula the rank correlation was moderately positive, but not as strong as in Figure 2.

[36] To verify that the estimated screening level affects the bidder's participation decisions, I regress the number of participants on the estimated screening level and indicator variables for each district. On average, a one-standard-deviation increase in the estimated screening level is

Table 5: Variable descriptions.

Statistic	Description
Reserve	Reserve price ($\$/m^3$)
Price	Sale price, i.e. the winning bonus bid plus the reserve price ($\$/m^3$)
# Bids	Number of bids submitted
Volume	Estimated volume of merchantable timber (1,000 m^3)
Vol. per Week	Minimum rate of extraction to complete before license expires
Vol. per Tree	Volume per tree
Vol. per Hect.	Volume per hectare
Price Index	British Columbia's consumer price index
Lumber Price Index	Value of timbers, estimated by multiplying volume by species-specific lumber recovery factors and British Columbia's lumber price indices
Quality Index	Index computed as the ratio of the estimated volume of recoverable lumber to a fixed benchmark defined by the Ministry of Forest's
Cycle Time	Amount of time from the site to the nearest point of appraisal (hours)
Avg. Slope	Average slope of land in the tract
Horse	Percent of volume to be extracted by horse
Cable	Percent of volume to be extracted by cable yarding
Helicopter	Percent of volume to be extracted by helicopter
Blowndown	Percent of volume that has been blown down
Burned	Percent of volume that has been burned
Useless	Percent of volume that is dead or useless
Dev. Costs	Anticipated development costs to be paid by licensee ($\$/1,000m^3$)
Hemlock	Estimated percent of volume from hemlock trees
Balsam	Estimated percent of volume from balsam trees
Cedar	Estimated percent of volume from cedar trees
White Pine	Estimated percent of volume from white pine trees

Table 6: Summary statistics for Category 1 Auctions.

Statistic	Min	Pctl(25)	Median	Pctl(75)	Max	Mean
Reserve	0.250	23.175	33.750	43.090	86.540	33.192
Price	0.570	34.775	47.500	59.840	105.000	47.354
# Bids	1	2	4	6	19	4.315
Volume	1.000	3.727	6.407	10.993	61.496	8.341
Vol. per Week	0.005	0.060	0.121	0.239	29.861	0.284
Vol. per Tree	0.080	0.340	0.500	0.620	5.010	0.534
Vol. per Hect.	0.004	0.179	0.264	0.335	3.027	0.263
Lumber Price Index	38.141	102.671	115.283	127.407	173.494	115.636
Quality Index	0.538	1.164	1.240	1.310	1.544	1.241
Cycle Time	0.000	2.700	3.500	4.600	16.400	3.865
Slope	0	7	14	23	86	16.257
Horse	0	0	0	0	1	0.090
Cable	0	0	0	0	1	0.082
Helicopter	0	0	0	0	1	0.030
Blowdown	0	0	0	0	1	0.020
Burned	0	0	0	0	1	0.012
Useless	0.000	0.000	0.020	0.080	0.450	0.052
Dev. Costs	0.000	0.000	0.679	1.900	32.311	1.539
Hemlock	0	0	0	0.001	1	0.093
Balsam	0	0	0.01	0.1	1	0.121
Cedar	0	0	0	0	1	0.032
White Pine	0	0	0	0	0	0.003

$T = 1507$

Table 7: Estimates from an ordinary least-squares and the partially linear single-index model for bids in Category 1 auctions. The OLS regression includes fixed effects for each of ten markets defined using a hierarchical clustering algorithm.

	OLS	PLSIM
Volume	0.194***	0.619***
	(0.052)	(0.026)
Vol. per Tree	7.803***	13.388***
	(1.295)	(0.841)
Vol. per Hect.	21.426***	14.407***
	(2.591)	(2.095)
Lumber Price Index	0.217***	0.177***
	(0.020)	(0.011)
Quality Index	20.778***	20.990***
	(5.176)	(1.184)
Cycle Time	−2.446***	−4.350***
	(0.204)	(0.087)
Slope	−0.153***	0.169***
	(0.040)	(0.036)
Horse	−16.128***	−20.692***
	(1.276)	(0.580)
Cable	−8.008***	−13.350***
	(1.673)	(1.792)
Helicopter	−47.354***	−64.830***
	(2.029)	(0.870)
Blowdown	−9.410***	−2.260
	(2.688)	(1.380)
Burned	−21.566***	−0.114
	(2.956)	(1.947)
Useless	7.802	−4.109
	(5.649)	(4.211)
Dev. Costs	−0.721***	−0.626***
	(0.121)	(0.152)
Hemlock	−13.759***	−17.304***
	(2.211)	(1.444)
Balsam	−13.968***	−22.300***
	(1.766)	(0.973)
Cedar	−3.667	21.925***
	(3.858)	(4.073)
White Pine	36.510**	11.041
	(14.558)	(10.131)
Observations	1,507	1,507

Note: *p<0.1; **p<0.05; ***p<0.01

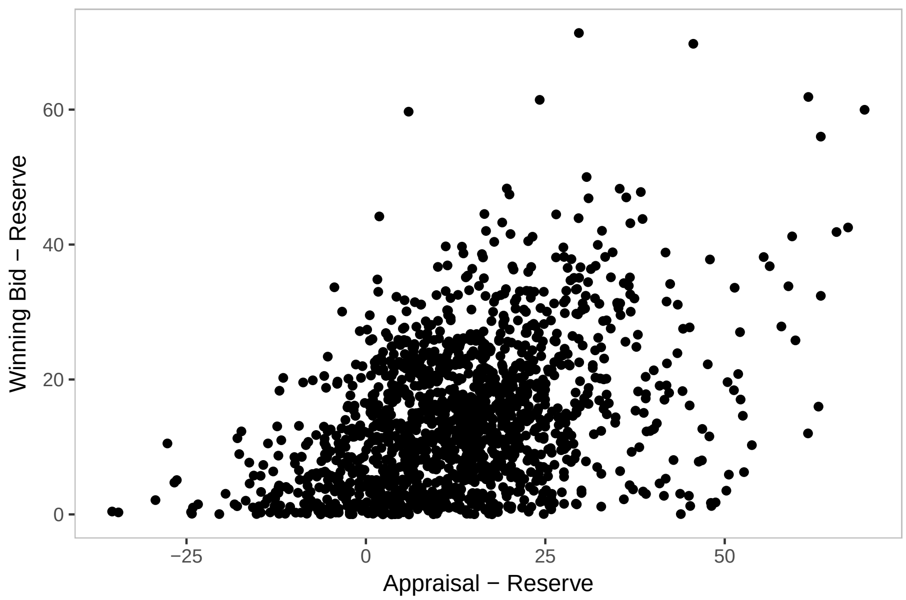

Figure 2: *Homogenized prices as a function of the homogenized reserve prices.* The winning bonus bid is positively correlated with the estimated value of the timber tract. In addition, the variance in the winning bonus bid is greater when the license is more valuable relative to the reserve price because bidders are more likely to enter with a wider range of private valuations.

D Conditioning on the Set of Potential Bidders

As a prerequisite to the estimation procedure used in this paper, I must condition on the sets of firms that each bidder i deems to be potential competitors. In line with earlier work on SBFEP auctions (Paarsch, 1997), I assume that the 31 geographic districts defined by the Ministry of Forests constitute submarkets in the sense that all bidders hold the same beliefs about the distribution of potential entrants for any auction in a given district. The bidders' equilibrium strategies can then be estimated conditional on each district.

This assumption appears to be supported by the data. There are, however, neighboring districts in which the same sets of bidders have similar participation

associated with a decrease of 0.81 in the number of bids submitted. Though measurement error might attenuate the estimated coefficient on the screening level, the estimate was significantly different from zero ($p = 0.05$). In addition, the point estimates of $\eta(\cdot, \mathcal{N})$ were monotonically increasing functions of the screening level. If the reserve prices did not bind, then η should be constant within each market.

rates. If bidders' beliefs about $\mathcal{N}_t$ are in fact the same for auctions in these districts, then they could be pooled in order to more precisely estimate the competing distributions. But this raises two key questions. First, what is the right notion of similarity between districts? And, second, how similar should districts be to justify pooling them in the estimation?

Regarding the first question, I argue that two districts should be grouped into the same market if the same bidders have similar rates of participation in both districts, in which case the distribution over competitors is approximately equal. Consequently, each bidder's competing distribution will be similar in the two districts, and it may be reasonable to pool these districts in the estimation. I therefore measure the distance between districts l and k by $\sum_i \frac{|w_{il}-w_{ik}|}{w_{il}+w_{ik}}$, where w_{il} is the frequency of participation by bidder i in district l, and the sum is taken over all bidders i for whom w_{il} and w_{ik} are not both zero. [37]

In answer to the second question of how similar is similar enough, I do not take a definitive position. Instead, I use a hierarchical clustering algorithm that produces a sequence of increasingly coarse market definitions. Initially, the algorithm assigns each of the 31 districts to its own market (or cluster). In the first step, the algorithm merges the two districts that are closest to each other and computes a measure of dissimilarity among the 30 resulting markets. It then iteratively merges the two most similar markets and recalculates the dissimilarities until only two markets remain. If the measure of dissimilarity between clusters is chosen appropriately, this algorithm produces a sequence of nested partitions of the auctions.[38] One can then perform the analysis using any one these market definitions.

The market definition consisting of eleven submarkets appears reasonable with the exception of the smallest. The lone district in this cluster is in a remote area of northeast and was the location of only one Category 1 auction during the period under study. I therefore omit this district from the empirical analysis.

Other output from the hierarchical clustering algorithm is illustrated by the

[37]This metric can be viewed as a weighted L_1-distance. I also experimented with other metrics, such as the Euclidean distance, but found that they produced less reasonable results. For example, the Euclidean distance might consider two districts to be very similar even though the identities of the firms that constitute the competitive fringe do not overlap. This underemphasis on matching based on the zeros in w_l and w_k then led the hierarchical clustering algorithm to group noncontiguous districts. The above measure appears to be better suited to the sparse participation patterns in the data.

[38]I use Ward's method of defining dissimilarities between clusters, which is designed to produce clusters that have minimal within-cluster variance.

dendrogram in Figure 3. Each "leaf" of the dendrogram represents a district as defined by the Ministry of Forests. The height at which two leaves join represents the dissimilarity between the markets that were merged to form a new market. The tree can be "cut" at different heights to produce the various market definitions. For example, the dashed horizontal line in Figure 3 shows the cut that partitions the districts into 11 markets.

At the bottom of the figure, I sum the number of Category 1 auctions in each district between 1996 and 2000. I also sum the auctions in each of the illustrated markets. This partition into 11 submarkets is the default market definition in the empirical analysis, so these sums indicate the number of independent observations available to estimate the bidders' market-specific inverse bidding strategies. The names of the districts are abbreviated as Cranbrook (Cra), Kootenay Lake (Koo), Invermere (Inv), Quesnel (Que), Robson Valley (Rob), Mackenzie (Mac), Fort Nelson (FNe), Dawson Creek (Daw), Fort St. John (FJo), Prince George (PrG), Fort St. James (FJa), Vanderhoof (Van), Burns Lake (Bur), Morice (Mor), Bulkley-Cassiar (BuC), Kispiox (Kis), Kalum (Kal), Horsefly (Hor), Williams Lake (Wil), Chilcotin (Chi), Castlegar-Arrow (CaA), Boundary (Bou), Columbia (Col), Salmon Arm (Sal), Vernon (Ver), Lillooet (Lil), Penticton (Pen), Merritt (Mer), Kamloops (Kam), Clearwater (Cle), and 100 Mile House (MiH).

E Numerical Solution to Asymmetric First-Price Auctions with Reserve Prices

The system of differential equations that characterizes the equilibrium is indeterminate or unbounded near the minimum bid, depending on whether the reserve price binds. Consequently, the "forward" methods of solving initial value problems cannot be applied because there is no way to evaluate the system at the left boundary. As an alternative, one could start with an initial guess of the maximum bid and use backward shooting algorithms to find the inverse strategies that satisfy the initial value conditions. However, Fibich and Gavish (2011) show that these methods become increasingly unstable when there are many different types of bidders. Instead, they advocate converting the initial value problem into a boundary value problem by choosing a bidder, say bidder n, as a benchmark and writing the equilibrium bids and the other bidders' strategies as a function of bidder n's valuation. Fixed-point iterations can then be used to find a solution.

While this iterative method appears to work well in several interesting cases, there is no theory to guarantee that the sequence of iterations will converge. And, in practice, I find greater success by making two modification to their solution strategy. First, rather than discretizing the solution and using finite difference approximations, I seek an approximate solution in the space of cubic splines. That is, I represent the solution as a linear combination of basis functions and solve for the coefficients that minimize the residuals in the differential equations.

Second, I solve the system (18) instead of the equivalent system of equations that characterize the inverse bidding strategies:

$$\frac{2\tilde{b}}{F_i^{-1}(\tilde{S}_i(\tilde{b})) - \tilde{b}^2 - r} = \sum_{j \neq i} \frac{\partial \log \tilde{S}_i(\tilde{b})}{\partial \tilde{b}} \quad \text{for all } i\,, \tag{18}$$

where $\tilde{b} = \sqrt{b - r}$ and $\tilde{S}_i$ is the distribution of $\tilde{B}_i$. By choosing bidder n to be the bidder whose bid density might diverge faster than $(b - r)^{-1/2}$, this change of variables ensures that the left-hand side of (18) has a finite limit as $\tilde{b}$ approaches zero. By solving for the distribution of $\tilde{B}_i$ rather than the inverse bidding strategies, evaulation of the system of equations only requires knowledge of the quantile function for bidder i. Though this modification is largely for convenience when F_i is known for all i, when applied using estimated valuation distributions, a practical advantage of this formulation is that F_i^{-1} can be nonparametrically estimated at a faster rate than the density f_i. For the purposes of computing an analytic gradient, however, it is still useful to evaluate the derivative of F_i^{-1}. To that end, I use a monotonic cubic spline to interpolate the estimated quantile function.

Thus, I approximate the solution to the boundary value problem

$$\begin{aligned}
\frac{\partial \tilde{S}_i}{\partial \tilde{S}_n} &= \frac{\tilde{S}_i(w)}{\tilde{S}_n} \frac{\sum_j \frac{1}{F_j^{-1}(\tilde{S}_j(\tilde{S}_n)) - w(\tilde{S}_n)^2 - r} - \frac{N-1}{F_i^{-1}(\tilde{S}_i(\tilde{S}_n)) - w(\tilde{S}_n)^2 - r}}{\sum_j \frac{1}{F_j^{-1}(\tilde{S}_j(\tilde{S}_n)) - w(\tilde{S}_n)^2 - r} - \frac{N-1}{F_n^{-1}(\tilde{S}_n) - w(\tilde{S}_n)^2 - r}} \\
\frac{\partial w}{\partial \tilde{S}_n} &= \frac{N-1}{2 \cdot w(\tilde{S}_n) \cdot \tilde{S}_n} \frac{1}{\sum_j \frac{1}{F_j^{-1}(\tilde{S}_j(\tilde{S}_n)) - w(\tilde{S}_n)^2 - r} - \frac{N-1}{F_n^{-1}(\tilde{S}_n) - w(\tilde{S}_n)^2 - r}} \\
& \tilde{S}_i(\underline{\tilde{S}}_n) = F_i(r) \quad , \quad \tilde{S}(1) = 1 \quad , \quad w(\underline{\tilde{S}}_n) = 0.
\end{aligned}$$

by finding the coefficients a_{il} and a_{0l} such that $\tilde{S}_i(\tilde{S}_n) = \sum_l a_{il} \beta_l(\tilde{S}_n)$ and

$w(\tilde{S}_n) = \sum_l a_{0l}\beta_l(\tilde{S}_n)$ minimize the error in the above system, where each $\beta = \{\beta_l : l = 1, \ldots, L\}$ is a basis of spline functions defined on the interval $[F_n(r), 1]$. In order to impose the boundary conditions, the knot vector used to construct the basis splines can be chosen to include ord copies of $F_n(r)$ and 1, where ord is the order of the spline functions. The boundary conditions are then satisfied whenever $a_{01} = 0$, $a_{i1} = F_i(r)$, and $a_{iL} = 1$. In addition, I impose monotonicity in the solution by restricting $a_{i\cdot}$ to be an increasing sequence for each $i = 0, \ldots, n-1$.

To be precise, let $a = (a_{il} : i = 0, \ldots, n-1, j = 1, \ldots, L)$ and $E_i(a, s; \beta)$ for $i = 0, \ldots, n-1$ denote the difference between the left- and right-hand sides of the above equations evaluated at $\tilde{S}_n = s$ in some arbitrarily fine grid, $s = s_1, \ldots, s_d$. The approximation problem is then given by

$$\begin{aligned} \min_a \sum_i \|E_i(a, \cdot\,; \beta)\| \qquad & \text{subject to} \\ F_i(r) = a_{i1} < \cdots < a_{iL} = 1 \qquad & \text{for } i = 1, \ldots, n-1 \\ 0 = a_{01} < \cdots < a_{0L}\,, \qquad & \end{aligned}$$

where $\|\cdot\|$ is some norm on $\mathbb{R}^d$. The Euclidean norm is an attractive option because it is differentiable, but the approximate solution might perform badly in a small region of the domain even when the average squared error is small. This does not appear to be an issue in simple cases, but I have found that the supnorm performs at least as well and is not vulnerable to this criticism. To preserve differentiability, I follow Hickman et al. (2016) and formulate the problem as

$$\begin{aligned} \min_{a,\epsilon} \epsilon \qquad & \text{subject to} \\ E_i(a, s; \beta) < \epsilon \qquad & \text{for } i = 0, \ldots, n-1 \text{ and } s = s_1, \ldots, s_d \\ -E_i(a, s; \beta) < \epsilon \qquad & \text{for } i = 0, \ldots, n-1 \text{ and } s = s_1, \ldots, s_d \\ F_i(r) = a_{i1} < \cdots < a_{iL} = 1 \qquad & \text{for } i = 1, \ldots, n-1 \\ 0 = a_{01} < \cdots < a_{0L} \qquad & \end{aligned}$$

Because the number of inequality constraints grows linearly with d, the grid cannot be arbitrarily fine. In general, d should be at least as large as the number of basis spline functions and could be much greater. In cases where the analytic solution has been derived by Kaplan and Zamir (2012), I find that

$d = 50$ already provides a good approximation with 12 cubic basis splines.

Figure 4 shows the approximate and the analytic equilibrium bid functions for an auction with two bidders and a reserve price of 0.3. The first bidder's valuations are uniform on $[0, 1]$ and the second benchmark bidder's valuations are uniform on $[0.2, 0.8]$. As depicted in Figure 4, the spline function provides a uniformly good approximation to the equilibrium; the maximum approximation error is 5.6×10^{-3}.

This figure also illustrates the difficulty in approximating the unbounded bid density near the reserve price. If, for example, Chebyshev polynomials in the bids were used to approximate the solution to the untransformed system, a high degree polynomial would be required in order to simultaneously approximate the steepness near the reserve price and the linearity near the maximum bid. The basis splines that I selected avoid this issue by solving the system in terms of the square-root of the bid minus the reserve price, as opposed to the bid, itself. Basis splines are also more flexible because they are defined piecewise on the partition of the domain created by the knot vector. Thus, in contrast to globally defined polynomials, they can accommodate curvature in the bid distribution functions in some regions of the domain without affecting the fit in others.

Note that the above system of equations will not be valid if the upper extremity to the support of bidder n's bids is less than the other bidders'. To avoid this situation, the benchmark bidder n can be chosen to be the bidder with the highest upper extremity to the support of its valuations. In this case, the differential equation for bidder i must be multiplied by an indicator for whether S_i is less than one. And, because the left-hand side of bidder i's differential equation might not approach zero as S_i approaches one, the knot vector for the basis spline functions should be chosen so as to allow for a discontinuous derivative at each bidder's maximum valuation.

Lastly, this solution method does not address the possibility that one of the bidders' strategy functions is constant near the reserve price. In this case the lower boundary condition for one of the n bidders must be removed or replaced with $S_i(r) = F_i\Big(r + \frac{G_i}{g_i}(r)\Big)$. This does not sacrifice uniqueness of the solution to the system of differential equations (Lebrun, 2006), but the above change of variables may no longer be appropriate. Further research is needed to determine a reliable solution method in this case.

F Extensions

F.1 Affiliated Private Valuations

One could strengthen the case for collusion by extending the basic identification strategy to more general auction models, such as models in which bidders' private valuations are affiliated. With affiliated valuations, each competitive bidder's inverse strategy function is identified under the null hypothesis it is not colluding if the two highest bids are observed (Athey and Haile, 2002).[39] Consequently, a test of independence between the competitively rationalizing valuations and an exogenous instrument would control the probability of making a type I error.

In practice, estimating these inverse strategies demands more from the data than under the independence assumption because the distribution of the highest competing bid must be estimated conditional on the bidder's own bid. A type-symmetry assumption would justify pooling bidders of the same type in order to more precisely estimate their inverse strategy. But this assumption sacrifices flexibility in the marginal distributions of bidders' valuations. Therefore, it may be necessary to choose between allowing for affiliation in valuations and allowing for unobservable heterogeneity in bidders' valuation distributions. For example, in the above analysis, I argue that valuations are independently distributed conditional on the observable covariates because the data include all the variables that the Ministry of Forests itself uses to estimate the value of the timber licenses. This assumption could be relaxed, however, if each bidder were observed to win more frequently.

On the other hand, the cost of collusion is not generally identified when bidders' private valuations are affiliated. Even if identities of the colluders were known, the possibility of phantom bidding entails that the distribution of a colluder's valuation could only be nonparametrically identified conditional on the event that its valuation is the greatest among all ring members, which would be insufficient to nonparametrically identify the equilibrium prices in a competitive auction.

[39]Under the null hypothesis that bidder i is not colluding, the highest and second-highest bids are both serious bids if one of them belongs to bidder i. Lemma 1 in Athey and Haile (2002) then implies that the inverse strategies are identified.

F.2 Unobservable Auction-Level Heterogeneity

In contrast, if at least two bids are observed in each auction, one could account for some correlation among the bids by using losing bids to estimate the distribution of auction-level heterogeneity that bidders observe but the econometrician does not. The difficulty in developing this extension is again that some of the losing bids could be phantom bids and possibly unrelated to both the bidders' private valuations and the unobserved heterogeneity. As a result, the procedure would have to simultaneously test for collusion and estimate this latent distribution. More precisely, one might consider bidder i's bid and the highest bid among bidder i's competitors as independent measures of the unobserved heterogeneity. Under the null hypothesis that i is not colluding, the deconvolution methods in Krasnokutskaya (2011) can then be used to estimate the latent distribution and the distribution of each bidder's winning bids. In a multiple testing framework, one would generally have to re-estimate the unobserved heterogeniety distribution under each null hypothesis. On the other hand, if at least one bidder in each auction is known to be competitive *a priori*, the winning bid and the other competitive bids can be used as the independent measures of the unobserved heterogeneity. Identification of the colluders would then proceed as in the baseline case.

In the application to British Columbia's timber auctions, however, the reserve price is often binding, with the result that only one bid is observed for many of the auctions. And, unlike the automobile auctions studied in Roberts (2013), the reserve price is an explicit function of observable variables and does reflect any heterogeneity the econometrician does not already observe.

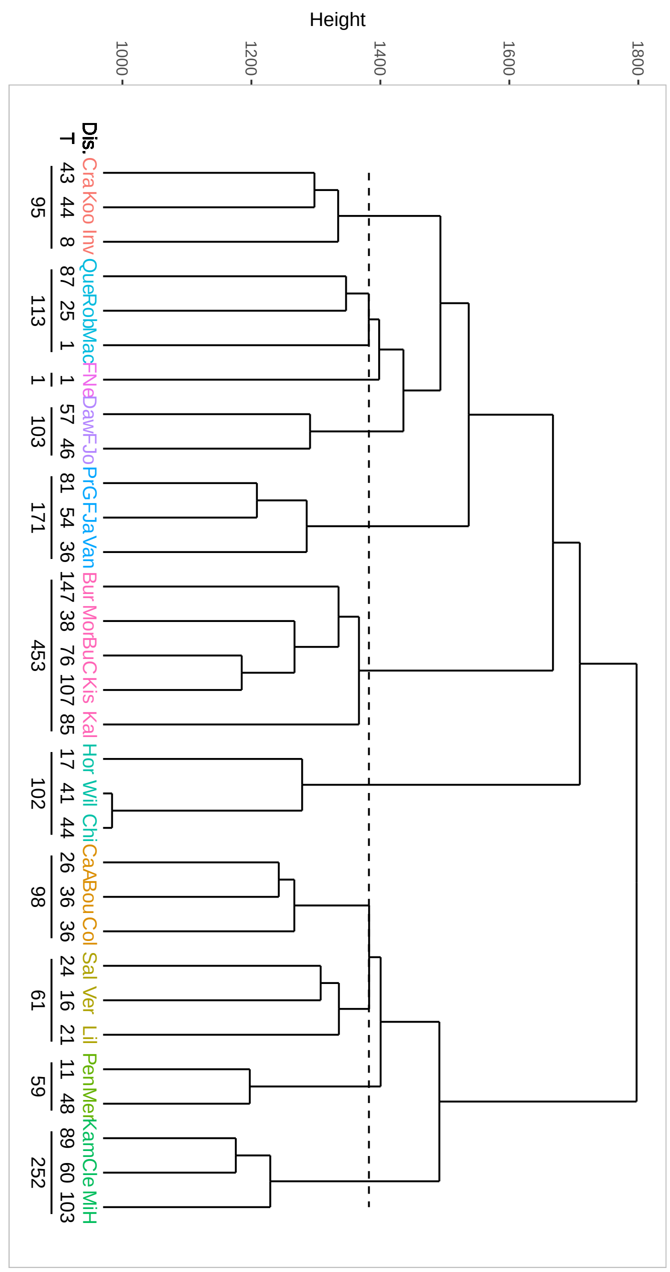

Figure 3: *A dendrogram representing the output of a hierarchical clustering algorithm.* Initially, each district is assigned to its own cluster. The two most similar districts, Williams Lake (Wil) and Chilcotin (Chi), are the first to be merged. At each juncture, the height of the dendrogram represents the dissimilarity between the two clusters that are merged.

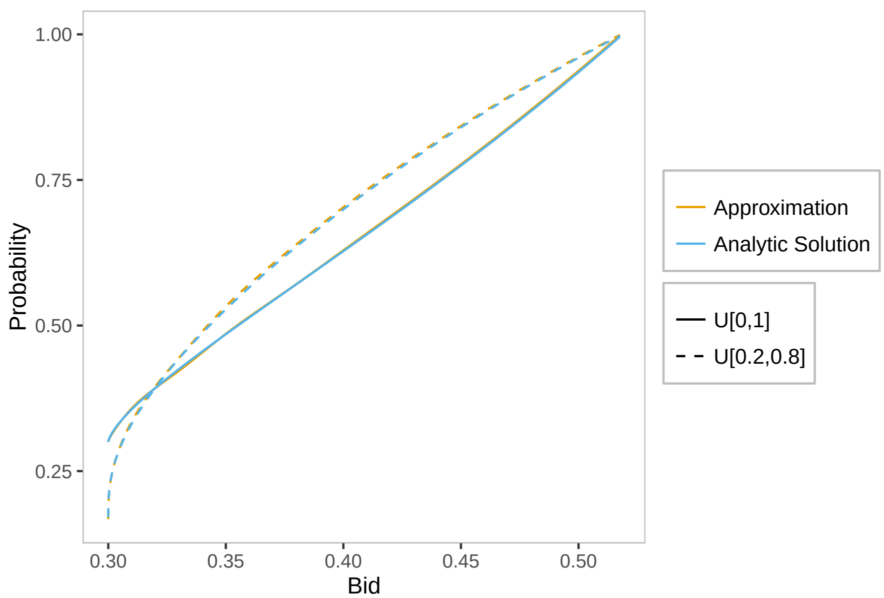

Figure 4: *A spline approximation to equilibrium bid distributions.* Twelve cubic basis spline functions were used to construct an approximate solution to the system of differential equations that defines the Bayes-Nash equilibrium for an auction with two bidders and a reserve price of 0.3. The bidders' valuations are uniformly distributed on $[0, 1]$ and $[0.2, 0.8]$. The splines are defined on a transformed domain in order to minimize the number of basis functions needed to approximate the equilibrium bid distributions near the reserve price.