

A Study of the Internal Organisation of a Bidding Cartel*

John Asker[†]

July 17, 2008

Abstract

This paper examines data on bidding in over 1700 knockout auctions used by a bidding cartel (or ring) of stamp dealers in the 1990s. The knockout was conducted using a variant of the nested knockout studied by Graham, Marshall and Richard (1990). An examination of the data reveals considerable heterogeneity among bidders, providing possible motivation for the choice of ring mechanism. Damages, induced inefficiency and the return to the ring are estimated using a structural model in the spirit of Guerre, Perrigne and Vuong (2000). A notable finding is that non-ring bidders suffered damages that were likely of the same order of magnitude as those of the sellers.

Keywords: Auction, Bidding, Ring, Cartel, Damages

JEL Codes: D44, K21,L41,L12

*I am indebted to the Antitrust Bureau of the New York Attorney General's Department for giving me the bidding data that make this study possible. Robert Hubbard and Hampton Finer were particularly generous with their time. Larry White, Ariel Pakes, Hong Luo, Matthew Gentzkow, Allan Collard-Wexler and Susan Athey provided helpful comments, as did seminar participants at the Canadian Competition Bureau, Duke, Harvard, Stanford, the UBC Summer IO conference, the US DoJ and the NYU Micro Lunch. I am indebted to Norman White for his considerable help in handling the NYU supercomputing infrastructure. Diana Mendoza provided excellent research assistance, particularly in dealing with New York's stamp collecting community. All errors are my own.

[†]Leonard N. Stern School of Business, NYU. 44 W4th St, New York, NY 10012. Email: jasker@stern.nyu.edu.

Introduction

When bidders in an auction collude - and, thus, reduce competitive bidding pressure - the resulting cartel is known as a ‘bidding ring.’ A ring can take many forms. An extensive theoretical literature exists that explores optimal ways to organize ring activity, given the form of target auction that the ring is seeking to exploit (see, for example, Graham and Marshall (1987), McAfee and McMillan (1992), Skrzyzypacz and Hopenhayn (2004) and Hendricks, Porter and Tan (2003)). Empirical work on ring activity has tended to focus on issues of ring detection (see Harrington (2005) for a survey). The detection of collusion and the determination of damages rely on drawing distinctions between market conduct with, and without, an active ring. Hence, understanding the practicalities of ring conduct is central to this endeavour. This paper contributes to this understanding by documenting and analyzing the conduct of a ring in the market for collectable stamps in North America that lasted for over 15 years.

The data used in this study comprise a record of the ring’s activities for an entire year, including sidepayments, detailed bidding behaviour in the internal ‘knockout’ auction that the ring used to coordinate its activities, and data on the associated ‘target’ auctions. These data are, as far as I am aware, unique.¹ In addition to providing documentation of the way that this long-running ring operated and the practical problems that it faced, these data allow an assessment of the damages that the ring imposed, who suffered the damages, the extent to which market efficiency was compromised, and by how much the ring benefited from its operation.

The results emphasize three aspects of ring conduct. First, the internal coordinating mechanisms employed by a ring can lead to inefficient market allocations. This runs contrary to much of the theoretical work on ring structure, which tends to assume that rings are able to implement coordinating mechanisms that lead to efficient allocation in the target auctions, since these will maximise ring returns. Second, not all ring members have equal bidding strength.² The ring studied in this paper appears to have had at least one member who was sufficiently weak, as compared to the other ring members, that his stated motivation for participating in the ring was the prospect of receiving sidepayments rather than acquiring cheap stamps. The presence of these weaker members is suggested as a justification for the particular mechanism the ring employed.³

¹The closest comprehensive data set of which I am aware are the financial records of a drug-dealing organisation analysed by Levitt and Venkatesh (2000).

²To add some precision, in an IPV auction model bidder strength would be a consequence of differences in the distribution from which bidders draw valuations. Weaker ring members would have value distributions that tend to place more mass on lower valuations.

³McAfee and McMillian (1992) introduce the notion of an ‘insincere’ bidder, which corresponds to an extreme form of heterogeneity in ring member strength.

Lastly, due to incentives created by the sidepayment structure of the ring mechanism, the ring imposed damages on both sellers and other, non-ring, bidders in the targeted auctions. Non-ring bidders were damaged by misallocation caused by the ring mechanism and also by a (somewhat counter-intuitive) tendency on the part of the ring to push prices higher than would have been the case in a competitive target auction.

Until July 1997, eleven wholesale stamp dealers were members of this bidding ring. The ring was active in auctions of collectable stamps, typically stamp collections, that occurred primarily in New York auction houses (which conduct open outcry ascending price (or English) auctions). The ring's activities started at some point during the late 1970s or early 1980s. Thus, it existed for 15 to 20 years, with very few changes in membership after the first five years. The ring collapsed when it was brought to the attention of antitrust enforcement by a stamp dealer not in the ring.

This ring operated through the use of a knockout auction that operated before the target auction. Before each target auction, ring members would submit a sealed bid on each lot that they were interested in.⁴ The highest knockout bidder on each lot received the stamps for sale if the ring won them in the target auction. The knockout bids were also used to set the bidding limit of the ring in the target auction and the sidepayments that individual ring members received. An important feature of the knockout mechanism (described in detail below) is that the greater the difference between the ring member's knockout bid and the price in the target auction (target price), the greater the sidepayment received.

An equilibrium analysis of the incentives created by the knockout mechanism makes it clear that a well-specified model is required to evaluate damages, efficiency impacts and the returns to the ring. This is because the ring members have an incentive to bid higher in the knockout than their actual valuations. This is due to the feature that sidepayments increase in the bid submitted (Graham, Marshall and Richard (1990) first noted the overbidding induced by this class of knockout auctions). Since knockout bids are not equal to underlying valuations, I propose and estimate a structural econometric model that allows the distribution of valuations to be inferred from observed bids. From this model, damages, efficiency costs and returns to the ring can be quantified. The approach taken is closest to that of Krasnokutskaya (2004), who adapts the non-parametric approach of Guerre, Perrigne and Vuong (2000) to accommodate auction-level heterogeneity observed by the bidders but not by the econometrician. Like Krasnokutskaya, I find that modelling unobserved auction-level heterogeneity has an economically significant impact on the point estimates emerging from the structural analysis. In addition, the model with unobserved heterogeneity performs much

⁴The exact timing of the knockout relative to the target seems to have varied. At one auction (the Harmer-Schau sale) the ring coordinated on bidding the night before the auction. More usually, there seems to have been at least a few days between knockout and target.

better at replicating moments observed in the raw data.

The structural econometric model provides considerable economic insight. The overbidding phenomenon introduced by the knockout mechanism results in instances where the seller actually benefits from the ring's activity since overbidding can push target prices higher than they otherwise would be. This offsets the detriment that the ring can impose on the seller, reducing damages below 'naive' estimates (computed by interpreting knockout bids as true revelations of underlying valuations). Paying careful attention to the equilibrium incentives created by the ring mechanism results in a reduction in the level of inferred damages to the seller, conditional on the ring winning a target auction, on the order of 50 percent.

Overbidding may occasionally benefit the seller, but it imposes damages on non-ring bidders and can cause misallocation and, hence, inefficiency. Non-ring bidders may suffer damages either by having their prices artificially inflated or by failing to obtain an object that they would have won in a truly competitive bidding environment. The structural model suggests that damages to other bidders may be at least as large as the damages to the seller. Interestingly, while efficiency is reduced by the ring's activity, the decrease is not economically significant, as compared to the impact of a change in participation.

Lastly, it is found that the ring, while occasionally hurting itself by overbidding, did benefit from coordinating bidding efforts.

That the realised sidepayment is an increasing function of the difference between the submitted knockout bid and the target price, places this ring mechanism in a class of knockout auction mechanisms first considered by Graham, Marshall and Richard (1990). Knockout auctions of this class have a long history, particularly in markets for collectables, such as rare books, art, rugs, stamps, coins and antiques. A notable example is the 1919 sale of the library of Ruxley Lodge, England.⁵ At this auction, 81 London booksellers combined to form a ring. This ring operated via a series of post-target knockouts in which ring members bid for the lot that had been acquired. The revenue raised in each knockout round was shared equally among the participants, with the more successful participants being invited to participate in subsequent rounds (thus increasing sidepayments as announced willingness to pay increased). This procedure resulted in at least four rounds of 81, 24, 15 and, finally, eight ring members. The subsequent knockout rounds generated additional revenues of £10,292 on top of the £3,161 paid in the target auction (see the extensive account in Freeman and Freeman (1990) and the summary by Porter (1992)). As Graham, Marshall and Richard demonstrate, the equilibrium properties of a mechanism like that used in the Ruxley Lodge ring mirror those of the mechanism used by the stamp ring considered in this paper. Other

⁵Prior to 1927, bidding rings were legal in the UK.

observed variations in ring implementation belonging to this class are recorded by Smith (1989), Cassady (1967) and Wraight (1974) in markets for farm land, collectable guns, rugs, antiques and paintings, with the earliest recorded example of a ring using a knockout belonging to this class being in 1830, in a sale for books (documented in Freeman and Freeman).

The rest of this section describes the ring mechanism in detail and relates the subsequent findings to the existing literature. Section 2 discusses the equilibrium incentives created by the ring mechanism in an IPV setting and discusses why a ring mechanism of the form used by the stamp cartel may be attractive. Section 3 then discusses the data set used in the study. Reduced-form analysis of the data is presented in section 4, a significant finding being the likely presence of weaker bidders. Section 5 then discusses the structural approach adopted in the rest of the paper. It provides an overview of the econometric methods that should allow readers less interested in econometric detail to skip section 6 and proceed straight to the results. Section 6 presents the details of the econometric approach used in estimating the structural model. The results of the structural model are presented in section 7. Section 8 concludes.

A description of the ring's internal organization

The ring used an internal auction or ‘knockout’ to coordinate bidding. Ring members would send a fax or supply a written bid to an agent (a New York taxi and limousine driver employed by the ring), indicating the lots in which they were interested, and what they were willing to bid for them in the knockout auction. The taxi driver would then collate all the bids, determine the winner of each lot, notify the ring as to the winners in the knockout and send the bids to another ring member who would coordinate the sidepayments after the target auction was concluded. Depending on who actually won the knockout, the taxi driver would, usually, either bid for the winner in the target auction, using the bid supplied in that auction as the upper limit, or organize for another auction agent to bid for the winner on the same basis. In the language of auction theory, the knockout was conducted using a sealed-bid format, with the winning bidder getting the right to own the lot should it be won by the ring in the target auction. The winning bid in the knockout set the stopping point for the ring’s bidding in the target auction.⁶ Since bidding in the target auction was handled by the ring’s agent, monitoring compliance with this policy was not a problem.

Sidepayments were used by the ring to compensate ring members for not competing for a lot, when the ring was successful in winning that lot. The determination of sidepayments is explained using the following example:

⁶Recall that that target auction is an English auction.

Example 1: Sidepayments from a successful acquisition in a target auction
Sotheby's, 24 June 1997, Lot 49

Knockout Auction	Bid (\$)	Sidepayment
Bidder A	9,000	$-\left(\frac{7,500-6,750}{2}\right) - \left(\frac{8,000-7,500}{2}\right) = -625$
Bidder G	8,000	$+\left(\frac{7,500-6,750}{2}\right) \times \frac{1}{2} + \left(\frac{8,000-7,500}{2}\right) = 437.50$
Bidder B	7,500	$+\left(\frac{7,500-6,750}{2}\right) \times \frac{1}{2} = 187.50$
Bidder J	5,100	0
Target Auction Price	6,750	

In example 1, the winner of the knockout auction was bidder A, who bid \$9,000. The ring was successful in the target auction, winning the lot for \$6,750. Since bidder J bid only \$5,100 in the knockout auction, he was not eligible for a sidepayment since his bid in the knockout was less than the target auction price. Bidders B and G bid more than the target auction price and so both are eligible for a sidepayment. The computation begins with bidder B's bid of \$7,500. The difference between \$7,500 and the target auction price is \$750. The knockout winner, bidder A, keeps half this amount. The other half is split equally between bidders B and G, resulting in each getting \$187.50. This is the only sidepayment bidder B gets. Bidder G bid higher than bidder B and so is eligible for a further payment. The winner of the knockout, bidder A, keeps half the increment between bidder G's bid and bidder B's bid and gives the balance, \$250, to bidder G.

Thus, the sidepayments involve ring members sharing each increment between bids, provided that their bids are above the target auction price. Half the increment is kept by the winner of the knockout, and the balance is shared equally between those bidders who bid equal to or more than the 'incremental' bid. The sidepayments were aggregated and settled on a quarterly basis.

Occasionally, the bids in the knockout auction tied for the winning position. In these instances, the ring's agent had discretion as to who won. He allocated the winning position after talking to the parties and 'tr[ie]d as much as possible to be fair' in who he chose, taking into account any previous tied bids.

The special position of the ring's agent (the taxi driver) in this mechanism is notable. It raises the question of what the incentives of this agent were. For instance, the ring would be vulnerable to collusion between the agent and a subset of ring members. The agent was primarily managed by one of the more active ring members, who also oversaw the payments among ring members. This member paid the agent \$30 an hour, and each ring member paid the agent a further \$50 for each target auction that they were active in. The depositions collected during the prosecution of this

ring directly addressed the issue of whether the agent ever let ring members know the knockout bid of other ring members before the knockout was completed. The depositions make it clear that this never occurred: neither the coordinating ring member nor any other ring member received this kind of assistance from the agent.

Prior to the late 1980s, the ring used a slightly different variant of this sidepayment system. The difference was that each increment was split equally among all eligible bidders. So, in the above example, the \$750 increment between the target auction price and bidder B's bid would have been shared three ways, with the winner, bidder G and bidder B, getting \$250 each. This is the same as if they treated the bids as true willingness to pay and gave each ring member their imputed shapley value (Graham and Marshall (1987) make this point in their theoretical discussion of nested knockouts).

The evidence on the enforcement of the rings rules is limited. During the early years of the ring, at least two members were ejected, one for not meeting his financial obligations and another for being felt to be only interested in collecting sidepayments. The records of the case indicate that deviant bidding behavior - behavior that did not comply with the rules of the ring - was either not a problem or not detected by the ring. For instance, there is no suggestion of members bidding and losing in the knockout and then participating in the target auction anyway. Similarly, there is no record of people getting temporarily suspended from ring membership. Instead, all accounts agree that the ring was very stable over the 15 or so years it operated.⁷

There is some evidence that one auction house was aware that some sort of ring was active in this market and communicated this to the ring. In response the ring did not operate at this auction house. It is unclear to what extent other auction houses were aware of the ring's activities. In the final few years of operation the ring become more sensitive to hiding their activities and in some target auctions would use the ring's agent, other professional bidding agents (who were not made aware of the ring), and their own members to enter bids, to make it less obvious that a ring was active. No detail exists on whether anything specific prompted this move to a slightly less transparent mode of operation.

Relationship to previous literature

A growing theoretical literature exists on the optimal design of bidding rings. Graham and Marshall (1987) and Mailath and Zemsky study collusion in second-price and English auctions, while McAfee and McMillan (1992) study first price auctions. More recent papers by Skrzyzpacz and Hopenhayn (2004), Marshall and Marx (2006), and Che and Kim (2006) provide results on dynamic collusion, rings with limited power, and the options open to a strategic seller, respectively.

⁷That said, it is also clear that ring members were not averse to bickering amongst themselves.

Hendricks, Porter and Tan (2003) extend the basic model to affiliated and common value environments and show that data from joint bidding in oil and gas drilling rights auctions are consistent with the theory.

Of particular interest is Lopomo, Marshall and Marx (2005) which shows that, when the knock-out occurs after the target auction and the mechanism has (ex-post) budget balance, inefficiency is inevitable. That is, there does not exist a scheme, however complicated and impractical, that can implement allocative efficiency. The findings of Lopomo, Marshall and Marx suggest that the qualitative empirical findings of this study are likely to be reflected in other cartels operating at English auctions.⁸

Alongside the theoretical literature, an empirical literature on bidding rings has developed. Baldwin, Marshall and Richard (1997), Porter and Zona (1993, 1999) and Bajari and Ye (2003) propose tests for the detection of cartel activity in highway construction, timber, and milk auctions, respectively. Athey, Levin and Seira (2004) also test for collusion in timber auctions.⁹

As in this study, Pesendorfer (2000) examines the activity of a known cartel in the market for milk provided to high schools. He uses extensive data on the target auctions, but says little about the internal workings of the cartel. He is limited to making inferences about the ring's internal structure from the outcomes of bidding in the target auctions. The data used in this paper mirror those used by Pesendorfer: their strength is the detail about the internal mechanism used by the cartel, and their weakness is the relative lack of data on the target auction (only the winning bid is recorded, together with whether the cartel won the target).

In the empirical auction literature, this paper is closest to Kwoka (1997). Kwoka observes the bids and sidepayments from 30 knockouts used to allocate real estate among members of a ring. The data that Kwoka uses are unclear as to the exact form of the sidepayment mechanism in all but ten auctions. However, under reasonable assumptions, Kwoka estimates that the ring distorted prices downward by up to 32 percent.

Within the wider literature on cartels, the work here is closest to that of Genesove and Mullin (2001), who use notes on cartel meetings to examine the internal operation of the US Sugar Cartel, which operated until 1936. Significantly, they find that some cheating did occur but was met with only limited punishment. Like Genesove and Mullin's sugar data, the data used here allow the inner workings of the stamp ring to be examined. Roller and Steen (2006) conduct a similar study on the Norwegian Cement Industry.

⁸Lopomo, Marshall and Marx mention several cases that fit this fact pattern.

⁹Harrington (2005) surveys the literature on cartel detection.

Properties of the ring mechanism

The ring mechanism described above creates some counterproductive incentives for the ring members. The analysis below establishes that the ring members had an incentive to bid higher than their valuations in the knockout auction, that this introduces an inefficiency in the target auction, and that other bidders and the ring itself, in addition to the seller (auctioneer), may suffer from this. It also provides the theoretical model estimated in the structural analysis. The model is an adaptation of the Graham, Marshall and Richard (1990) model.

Let the value to each bidder $i \in I$ in knockout auction k be given by $v_{ik} \in [\underline{v}_i, \bar{v}_i]$, which is drawn from a distribution $F_i(v)$. It is assumed that bidders' valuations are independently drawn from bidder-specific distributions and are private information. Ring members are assumed to know the number of bidders bidding in the knockout, but are uncertain about the identity of the other bidders. The probability of each bidder participating is assumed to be known and is denoted α_j . This structure seems to give the most reasonable approximation to the data, given that the number of bidders in a knockout is somewhat predictable from observables, but the identities of bidders, as is made clear in the accounts of the knockout in the depositions, are not known *ex ante*.

The way that the ring was organized meant that the winning knockout bid became the bidding limit of the ring in the target auction. Since the target auction was an ascending open outcry auction, the equilibrium bidding strategy of bidders outside the ring, in the target, is to raise the bid on any competing bidder so long as the price does not exceed their valuation. That is, standard arguments for ascending (English) auction formats apply. This simple equilibrium structure in the target helps make analysis of equilibrium bidding in the knockout auction tractable.¹⁰

It is assumed that there is an equilibrium of the knockout auction in which each bidder follows a strategy \tilde{b}_i^* . It is assumed that these strategies are increasing in v_i , differentiable and have inverses

¹⁰The auction houses' catalogs all report the existence of secret reserve prices. For example Christie's Oct 22 1996 sale catalog reads: "The reserve price is the minimum price the seller is willing to accept and below which a lot will not be sold. This amount is confidential and does not exceed the low pre-sale estimate. Property offered for sale subject to a reserve is indicated by the symbol * next to the lot number" Every lot is indicated as having a reserve price. However, in the data there is no evidence that this reserve price ever binds. The model does not attempt to model the reserve price, due to the lack of any data on it and the lack of evidence suggesting it has an affect on bidding (e.g. that the secret reserve price lay within the support of the distribution of valuations, or was enforced). Similarly, accounts of auction practice (e.g. Cassady (1967) and Wraight (1974)) mention the use of shill bidding by auctioneers to drive up prices. There is no evidence suggesting that this occurred in these specific auctions.

$\phi_i = \tilde{b}_i^{-1}$. The bidding problem of each ring member in the knockout can then be expressed as

$$\max_b \int_{-\infty}^b (v_{ik} - x) h_r(x) dx F_{-i}(\phi(b)) - \frac{1}{2} \int_{-\infty}^b \int_x^b (y - x) h_r(x) f_{-i}(\phi(y)) dy dx + \frac{1}{2} \int_{-\infty}^b (b - x) h_r(x) dx (1 - F_{-i}(\phi(b)))$$

where v_{ik} is the value of winning the item for sale in the target auction. r is the highest value from amongst the bidders in the target auction who are not in the ring (non-ring bidders). $H_r(\cdot)$ is the distribution function of r and $h_r(\cdot)$ is the corresponding density function. $F_{-i}(\phi(b)) = \frac{\sum_{j \neq i} \alpha_j F_j(\phi_j(b))}{\sum_{j \neq i} \alpha_j}$ and $f_{-i}(\phi(b)) = \frac{\partial F_{-i}(\phi(b))}{\partial b}$.

The first-order condition of this maximization problem is given by

$$(v_{ik} - b) h_r(b) F_{-i}(\phi(b)) + \int_{-\infty}^b (v_{ik} - x) h_r(x) dx f_{-i}(\phi(b)) - \int_{-\infty}^b (b - x) h_r(x) f_{-i}(\phi(b)) dx + \frac{1}{2} \int_{-\infty}^b h_r(x) dx (1 - F_{-i}(\phi(b))) = 0$$

With two knockout bidders, this expression simplifies to¹¹

$$v_{ik} = b - \frac{\frac{1}{2} H_r(b) (1 - G_{-i}(b))}{(h_r(b) G_{-i}(b) + H_r(b) g_{-i}(b))} \quad (2)$$

where $G_{-i}(b)$ is the distribution function of b_{-i} and $g_{-i}(b)$ is the corresponding density function.¹² Note that $F_{-i}(\phi(b)) = G_{-i}(b)$ and $f_{-i}(\phi(b)) = g_{-i}(b)$.

This mapping from bids and bid distributions to valuations is the key theoretical component of the structural empirical analysis. Crucially, equation (2) maps each observed bid to a unique valuation, in the two-bidder knockouts. The empirical analysis focuses on the two-bidder case due to the lack of identification in a well-specified empirical model of bidding when three or more bidders engage in the knockout (more detail is provided in section 5). Focusing on the two-bidder case also means that the estimated model has the same sidepayment structure as that considered by Graham, Marshall and Richard.

¹¹The three-bidder analog to (2) is:

$$v = b - \frac{1}{hG^2 + 2HGg} \left[\frac{1}{2} Gg \int_{-\infty}^b H(x) dx - \frac{1}{2} g \int_{-\infty}^b H(x) G(x) dx + \frac{1}{4} (1 - G)^2 H + (1 - G) GH \right]$$

The four-bidder analog is:

$$v = b - \frac{1}{hG^3 + 3HG^2g} \left[G^2 g \int_{-\infty}^b H(x) dx - \frac{1}{2} Gg \int_{-\infty}^b H(x) G(x) dx - \frac{1}{2} g \int_{-\infty}^b H(x) G^2(x) dx + \frac{3}{2} (1 - G) G^2 H + \frac{3}{4} (1 - G)^2 GH + \frac{1}{6} (1 - G)^3 H \right]$$

Note that g, G, h , and H are shorthand for $g_{-i}(b), G_{-i}(b), H_r(b), h_r(b)$ respectively.

¹² $G_{-i}(b) = \frac{\sum_{j \neq i} \alpha_j G_j(b)}{\sum_{j \neq i} \alpha_j}$. $g_{-i}(b)$ is similarly defined.

The general first-order condition allows persistent overbidding in the knockout auction to be established. Lemma 1 proves the first step toward this end.

Lemma 1: If the bidders objective function is denoted π_{ik} , then

$$\left. \frac{\partial \pi_{ik}}{\partial b_{ik}} \right|_{b_{ik}=v_{ik}} = \frac{1}{2} \int_{-\infty}^{v_{ik}} h_r(x) dx (1 - G_{-i}(v_{ik})) \geq 0$$

Proof: The result is immediate from evaluating the derivative of π_{ik} at the point $b_{ik} = v_{ik}$

Lemma 1 leads to the first result, that ring members' bids in the knockout auction are (weakly) greater than their valuations.

Result 1: For any number of knockout bidders $b_{ik} > v_{ik}$ if $v_{ik} \in [\underline{v}_i, \bar{v}_i)$, $b_{ik} = v_{ik}$ if $v_{ik} = \bar{v}_i$

Proof: For any bid below v_{ik} , the standard dominance argument for Vickrey and English auctions implies that $b_{ik} = v_{ik}$ is a weakly dominant bid. Provided that v_{ik} is less than the upper limit of the support of v_{-ik} , lemma 1 implies that increasing b_{ik} above v_{ik} increases expected profit (note that in lemma 1 $f_r(x)$ and $G_{-i}(b)$ are not restricted in any way). When v_{ik} is equal to the upper limit of the support of v_{-ik} , $b_{ik} = v_{ik}$ by lemma 1. QED.

That the structure of the knockout introduces inefficiency into the target auction is a consequence of overbidding in the knockout and the use of the knockout bid as an instruction to the ring's bidding agent (recall that the target auction is a standard ascending (English) style auction).¹³ This means that it is possible for the ring to win the object for sale at a price higher than the highest valuation of the ring members. Hence, when the highest non-ring valuation lies between the valuation of the highest ring member and that member's knockout bid, the target auction will generate an inefficient allocation, since the ring will win the object as opposed to the non-ring member who values it more. This conclusion is stated as a corollary.

Corollary 1: The structure of the knockout auction can lead to an inefficient allocation in the (English) target auction.

This overbidding by the ring can be a good thing for the auctioneer. It means that in some target auctions the price is bid up to a level that would not have been reached had the ring not

¹³Note that the overbidding phenomena does not rely on the presence of sidepayments. This is because there is something akin to an option value involved in having fewer competitors in the target auction.

been operating. Thus, any damages that the auctioneer suffers from the presence of the ring could be at least partially offset by overbidding.

Because the ring will bid above its highest valuation, it follows that non-ring bidders in the target auction may suffer damages. These damages may be incurred through two channels. First, when the ring wins at a price higher than its valuation, it must be the case that, were the auction competitive, a non-ring bidder would have won. The surplus that this non-ring winner fails to realize when the ring wins represents a source of damages. Second, even if a non-ring bidder wins, she may have had to face tougher bidding competition from the overbidding than that which she would have faced had the ring not been operating. This leads to a second source of damages, due to the ring artificially pushing up prices.

These factors suggest that participation decisions by all parties may be affected by the ring's design. Speculation on the implications of the ring's design for participation is put off until the empirical analysis is completed.

The following lemma, specific to the inverse bid function of the two-bidder knockout (equation 2), is useful for understanding some features of the bidding function that shapes the structural estimation of the model.

Lemma 2: The inverse bid function defined in equation (2) may have regions where $\frac{\partial v(b)}{\partial b} < 0$.

Proof: The simplest proof is by example. The mean of a lognormal distribution is given by $EX = e^{\mu + \frac{\sigma^2}{2}}$ and the variance by $VarX = e^{2(\mu + \sigma^2)} - e^{2\mu + \sigma^2}$. Let $H_r(b)$ be a lognormal distribution with $\mu = -2$ and $\sigma^2 = 0.08$ and let $G_{-i}(b)$ be lognormal with $\mu = -1$ and $\sigma^2 = 1$. It is straightforward to verify that a non-monotonic inverse bid function is the result of this parameterization. QED

Lemma 2 is a consequence of the fact that $\frac{\partial v(b)}{\partial b}$ depends on the gradients of $g_{-i}(b)$ and $h_r(b)$. The interaction of these terms can result in a non-monotonicity. Due to the fact that $v(b)$ as defined in (2) is a function, the implication of this non-monotonicity is that some critical points identified by (2) (and hence the more general first order condition preceding it) may be local minima. The most straightforward problem with this possible non-monotonicity is that it can create a situation where, in estimation, observed bids are mapped to valuations via a part of the first order condition that actually defines a local minimum.¹⁴ This violates the assumption that bidders are playing

¹⁴The discussion of equilibrium existence, monotonicity and continuity in first price auctions contained in Lizzeri and Persico (2000) is helpful in cataloging the myriad theoretical problems that this non-monotonicity would generate if it was implied by the underlying data generating process. One such problem is that, even if a pure-strategy monotonic equilibrium were assumed to exist, the inverse bid function would cease to be everywhere differentiable, which would violate an assumption needed to derive (2) and simultaneously ensure that the equilibrium bid density is properly defined.

optimally. In all that follows it is assumed that the data generating process is such that the true (population) function (2) is strictly monotonic. When estimating this model, this assumption is imposed via the selection criteria for smoothing parameters. See section 5, which outlines the structural estimation, for more details.

What motivates the use of this ring mechanism?

Central to understanding the mechanism used by the stamp ring is the relationship between bidder heterogeneity and ring design. Mailath and Zemsky (1991) show that a mechanism exists that enables efficient collusion when bidders have independent private valuations drawn from different distributions. Their mechanism has the desirable features of ex-post budget balance and coalitional stability, meaning that there is no money left on the table and that no subcoalition of bidders wishes to deviate and form a separate cartel.

Mailath and Zemsky point out that, when bidders are symmetric, a very simple mechanism is optimal: Hold a first price sealed bid auction prior to the target auction; give the winner the right to bid in the target; and equally split the winner's bid in the first price sealed bid knockout among the losers. This structure for the knockout is attractive for its simplicity and lack of dependence on the distribution function from which bidders' types are drawn.

Asymmetries between bidders in the ring makes this simple knockout design infeasible. Mailath and Zemsky provide a simple example in which it is better to exclude a weak bidder (a bidder who draws his value from a distribution that puts higher probability on lower valuations) if the ring cannot discriminate in making sidepayments. The intuition is simple: if, in expectation, the contribution of a bidder to the collusive surplus is less than the share of the total collusive surplus that that bidder would receive, then the ring should exclude that bidder.

However, if the ring can discriminate, it is better to include bidders, whether they are weak or strong. While Mailath and Zemsky characterize the optimal mechanism, it depends heavily on the distribution functions from which each bidder's valuations are drawn. This in itself is a problem, which is further compounded by the fact that a game form that implements the mechanism is, as far as I am aware, unknown.¹⁵

Graham, Marshall and Richard (1990) approach the problem from another angle. They investigate the pre-1990 design for the knockout auction described in the introduction. They motivate their theoretical investigation by arguing that this mechanism is attractive in a complete information setting since it gives each ring member his shapley value. Without asymmetric information, it provides a way to pay each ring member his marginal contribution to the ring, taking into account

¹⁵The mechanism in Mailath and Zemsky has one feature in common with that studied here, in that the share of collusive surplus is increasing in the strength of the bidder.

the differences in the magnitude of their contributions.

After introducing asymmetric information, they find that, when the ring includes all potential bidders, the mechanism introduces a persistent incentive for bids to lie above valuations in the knockout. The intuition for this result is that bidding the object's value in the knockout dominates any lower bid for the usual second price auction reasons; however, since the sidepayment is also increasing in the bid, this added effect pushes bids above valuations (Deltas (2002) shows that the overbidding incentive exists in other knockout formats).

The analyses in Mailath and Zemsky and Graham, Marshall and Richard suggest that in the face of bidder asymmetries, a ring faces a difficult set of design trade-offs. It can either reduce its potential collusive surplus by excluding weaker ring members, and use a simple knockout mechanism, or it can use a more complicated system that may be inefficient in terms of the ultimate allocation in the target auction. The stamp ring appears to have chosen the latter option, suggesting that it preferred to include weak bidders and generate some allocative inefficiency, as opposed to attempting to exclude weaker bidders from the ring altogether.

Taking into account the practical difficulties associated with the implementation of the optimal ring scheme proposed by Mailath and Zemsky, it is clear that some second-best, compromise scheme needs to be adopted. Perhaps the most curious feature of the precise scheme adopted by the stamp ring is that the winning knockout bid serves as the limit up to which the ring will bid in the target auction. To examine the effects of this feature consider an alternate design in which the winner of the knockout had discretion as to the upper limit in the target (sidepayments have to be paid on the same basis as in the actual mechanism should the cartel win in the target).

For simplicity assume the ring makes all knockout bids known before the target auction and that two ring members bid in the knockout. The ring member who won the knockout has a value of v_1^c , the second highest knockout bid is equal to k_2^c . Now consider the bidding problem of the ring member in the target auction. If $k_2^c < v_1^c$ the standard arguments for IPV English auctions imply an upper limit on bidding of v_1^c . If $k_2^c > v_1^c$ then the ring member in the target will want to stop bidding below his value to reduce the chance of winning and having to pay a sidepayment that is greater than his actual profit. His upper limit on bidding becomes $2v_1^c - k_2^c$.¹⁶

This propensity for underbidding in the absence of a link between the winning knockout bid and the ring's bidding rule in the target would likely create a series of difficult problems for the ring. First, the ring will still generate inefficient allocations (through not always acquiring lots for which they have a positive surplus at the target price). Second, the potential underbidding creates an incentive to defect from the ring either by secretly bidding against the ring in the target auction

¹⁶The objective function for the cartel bidder in the target is $\max_b \int^b (v_1^c - x) h(x) dx - \frac{1}{2} \int^b (k_2^c - x) h(x) dx$. The bidding rule arises naturally from the first order condition.

using an agent or by openly defecting from the ring. Third, it could promote entry by outside bidders due to eased competition. In contrast, the scheme actually adopted by the stamp ring has the advantage of promoting stability since no ring member has an incentive to be on the outside of a ring that can be hyperaggressive through overbidding and outside bidders may also be deterred from entry due to the same inherently aggressive bidding behaviour.

These considerations suggest that, when stability and the incentives for entry by outside bidders are taken into account, there are reasons why maintaining the link between the winning knockout bid and the ring's bidding behavior in the target auction could be advantageous to the ring.

The data

This study uses data on bidding in the ring's knockout auctions preceding 11 target sales occurring during the period 5 June 1996 to 26 June 1997.¹⁷ These data were collected by the New York Attorney General's department during its investigation of the ring. They were transcribed from faxes and written records kept by one of the ring members.^{18,19,20}

For each lot in which the ring was active, these data include the date of the target auction, the auction house where the target auction was conducted, the lot number, the transaction price in the target auction, all sidepayments made within the ring and, finally, the ring members who bid in the knockout and their corresponding bids. The data are generally very high quality. Where the data are incomplete, which is in less than three percent of observations, it is because sunlight had faded the faxes beyond legibility or a similar, plausibly random, event occurred.

These bidding data were augmented with data drawn from catalogs published by each auction house. These catalogs contain concise, but detailed, descriptions of the lots for sale, including a measure of the condition of the stamps, the countries or region of origin, an indication of the number of stamps and either the auction house's estimate of the value at auction (expressed as an interval) or the value according to one of several standard references on the retail value of stamps. The catalogs contain detailed information on their grading metrics and auction procedures.

The most significant shortcoming of this data set is the lack of information about the total

¹⁷In this industry, a sale refers to a collection of auctions of individual lots. Each auction usually takes around two to four minutes, with a sale often running over several days.

¹⁸This ring member was notorious within the ring for being paranoid about not receiving the side-payments that were owed to him.

¹⁹I am indebted to the Antitrust Bureau of the New York Attorney General's Department for making these data available to me. The case from which the data are taken is *NY et al. v. Feldman et al.*, No. 01-cv-6691 (S.D.N.Y.).

²⁰The data also include depositions from a ring member and the taxi driver who was employed to assist in coordinating the ring.

number of bidders in the target auction. This is particularly problematic for assessing damages, as it makes it hard to formulate the correct counterfactual scenarios.²¹ It is less of a problem for analyzing the knockout auction itself since at the time of bidding it is unlikely that the ring members knew exactly how many people would be bidding in any given lot in the target auction.

Tables 1 and 2 here

Tables 1 and 2 provide a summary of the bidding data. In Table 1 the mean winning bid of the target auctions is reported together with the mean bid of the knockout auctions for each of the nine auction houses at which the ring was active. It is notable that the mean bid in the knockout data is higher than the mean winning bid in the target auction data for 4 out of 9 auction houses. Pooling across auction houses (other than Harmer Schau), the mean bid in the knockout is \$2142 while the mean winning bid in the target auction is \$2042.²² The data also suggests that the ring was successful in winning lots. The ring won 58 percent of the lots in which it was active. The auctions won by the ring raised \$6,377,284 as compared to \$10,034,754 across all auctions in which the ring was active.

Table 2 reorganizes the data according to the number of ring members bidding in the knockout auction. In 623 out of 1781 lots, only one ring member was active. In 26 lots, eight ring members were active. The rest of the lots in which ring members were active had between two and seven bidders in the knockout. The mean winning bid in the target auction is weakly increasing in the number of bidders in the knockout. For lots in which there were between three and six bidders in the knockout, the mean winning price in the target seems fairly constant. There is a similar increasing trend in the average median bid in each knockout auction as the number of participants increases. The probability of winning the target auction increases as the number of active ring members increases, suggesting that either the number of non-ring bidders was not increasing at the same rate as ring participation, or that overbidding was more pronounced as more ring members participated, or some combination of these factors.

²¹Without knowing the number of bidders in the target auction it is hard to let the data drive the determination of the second-order statistic of the values of bidders not in the ring.

²²The data from the Harmer-Schau auction are recorded in this table but are not used in the structural or reduced form analysis that follows. The bidding behaviour in the Harmer-Schau auction seems to indicate that the ring was coordinating in a way that was not consistent with the rest of the auctions. The inconsistent feature is that the transaction price in the target auction is almost always the same as the second highest knockout bid. Precisely what is occurring in the Harmer-Schau auction is unclear from both the data and the depositions. All numbers reported in this paragraph also exclude Harmer-Schau.

Reduced form analysis

The stamp ring was extremely active, bidding on over 1700 lots in the course of a year. With such an active ring, it is not surprising that participation by ring members varied considerably. Since participation is likely to affect the returns to ring members and the damages suffered by sellers, I report the participation patterns found in the data together with an analysis of ring win rates. This leads to a discussion of the disproportionate share of sidepayments enjoyed by weaker ring members and, ultimately, the problems suggested by the theory in interpreting naive estimates of damages.

Participation and ring win rates

The ring had eleven members during the sample period. Table 2 gives an impression of the differential participation decisions of different ring members. The weak trend that these data suggest is that as the value of the object for sale increases, the participation of ring members also increases. The rate at which the ring won the target auction has the same pattern. The analysis here runs in two steps. First, I examine what determined the number of ring members that bid on each lot. Second, I look to see how participation rates, together with other covariates, impacted the win rate of the ring.

Table 3 here

Table 3 investigates the sensitivity of the number of ring members participating in the auction to various auction characteristics. The first column contains an OLS regression with the number of participating bidders as the dependent variable. The second column implements a (transformed) logit specification in which the dependent variable is the log of the share of ring members participating minus the log of the share not participating. Since both sets of results are similar, I focus on the logit results.

Participation is increasing in the estimated price supplied by the auction house. That is, participation is increasing in the expected price of the lot for sale (measured as the mean of the minimum and maximum estimated price, $(E_{min}+E_{max})/2$), although this effect is not statistically significant. The spread between the estimated maximum and minimum price has an much stronger (and significant) effect on participation: a unit increase in the spread has an effect an order of magnitude greater than an unit increase in the expected price. The condition of the stamps also matters (as measured by Grade Max and Grade Min).²³ Interestingly, the wider the gap between the lowest reported grade and the highest grade (holding the mean grade constant), the more ring

²³The variables in the regressions are the upper and lower limits of the estimate range published by the auction

members are likely to participate. Where no grade is supplied, participation appears to decrease (this is likely to reflect lots with a very large number of stamps).

Lastly, the auction house fixed effects suggest that ring members have a preference for some houses over others. Whether this is due to the quality to the stamps for sale, the competition in the auctions at these houses or some other characteristic of the auction house, is unclear.

The most striking thing about participation in the knockout auctions is that it is always incomplete. That is, the maximum numbers of bidders in a knockout is eight rather than the full complement of 11 ring members. Given the structure of sidepayments, this is curious, as ring members could expect to gain some small expected benefit from always entering a very low bid in the knockout. The fact that this does not occur might suggest the existence of a dynamic enforcement mechanism, whether based on reputation or explicit punishment rules.²⁴

Table 4 here

Table 4 reports the results of a logistic regression of a dummy equal to one if the ring wins, on a set of observables. The point estimates from the logistic regression are reported together with the estimated marginal effects. An increase in the number of ring members interested in the lot for sale increases the chance of winning although this effect diminishes slightly as the number of ring members increases. After controlling for other covariates neither the expected value of the lot nor the spread in estimated prices have any significant effect on win probabilities. However the ring does appear to win more often when the stamps for sale are of a higher quality (as measured by $(\text{Grade Min} + \text{Grade Max})/2$).

Even after controlling for other auction-specific observables the auction house dummies remain a significant determinant of the ring's propensity to win. Whether this was due to systematic differences in competition from outside bidders across auction houses, the effect of some unobserved lot quality or some other factor is unclear. The relative paucity of information on participation by outside bidders makes it difficult to make any informed conjecture on this issue.

Weaker ring members and outcomes in the knockout

house (Estimated Maximum and Estimated Minimum, resp.), the catalog value of the stamps (if published, stamps have 'blue books', much like cars), the upper and lower limit of the published grade of the stamps (Grade Max and Grade Min, resp.; stamps are graded on, roughly, a five-point scale), dummies for when values or grades are not published and, finally, a dummy for exclusively domestic stamp collections. The House variables are auction house dummies.

²⁴As reported earlier, there is no explicit evidence to support this speculation. A small fixed participation cost (e.g. time cost) may also explain this feature.

Table 5 here

Table 5 looks at how the bidding outcomes of individual ring members differed in the knockout auctions. Columns 1 and 3 of Table 5 report the relative frequency with which a ring member won a knockout conditional on participating in the knockout, including and excluding knockouts with just one bidder (resp.). In both columns 1 and 3, ring members D, G and H all had lower frequencies of winning than the other ring members. However, their frequency of receiving sidepayments (conditional on participating) was at least as high as the other ring members'. This suggests that D and G, in particular, were weaker bidders (H participated quite rarely and mainly on high participation knockouts). In an IPV setting, bidder strength would be a consequence of differences in the distribution from which bidders draw valuations. Weaker ring members would have value distributions that tend to place more mass on lower valuations. As a consequence any benefit that a weaker bidder received from ring participation would be more likely to be in the form of a sidepayment. Interestingly, in his deposition, D claimed, "My objective, basically was, you know, make money from these people as opposed to actually buying the stamps."

If D, G and H were weaker bidders, this should be reflected in the data on sidepayments.

Figure 1 here

Figure 1 shows the distribution of sidepayments. In Figure 1 a negative number indicates that the ring member, in aggregate, paid sidepayments *to* the ring. Aggregating over all auctions, ring member E was the largest net contributor in this sense, with an overall contribution of \$74,418. Ring members D and G received the most in net payments from the ring, getting \$22,251 and \$21,574, respectively (interestingly, G participated in only 186 knockouts, as opposed to 715 by D). This adds to the suggestion from Table 5, that D and G are the two most likely candidates for being weak in an economically significant way. Figure 1 also reports the distribution excluding the 25 lots that earned a target price of over \$10,000. This changes the distribution of payments into the ring, reducing the relative importance of E.²⁵ However, the relative efficacy of D and H in extracting sidepayments from the ring is unchanged.

If D, G and H were seeking to optimise their return from the extraction sidepayments, the degree of success they appeared to enjoyed is surprising.²⁶ Their propensity to avoid winning knockouts

²⁵E was active in 18 of these 25 high value target auctions. He made sidepayments in five of these auctions and received a sidepayment in one (of \$750). The largest sidepayment to the ring, made by E, was \$15,500 on a lot acquired for \$26,000. The next largest sidepayment to the ring by E was \$10,750 on a lot acquired for \$14,500. E's net sidepayments for the high value target auctions was \$38,000 paid to the ring.

²⁶Note that this objective is completely consistent with the objective function in the earlier theory section.

and yet maintain a probability of receiving sidepayments at least as high as those of their strong co-conspirators is especially noteworthy.²⁷

Comparing bidding behavior in the knockout and target auctions

A naive estimate of the extent of damages suffered by the seller from the ring's behavior would be given by looking at the second-highest bid in each knockout and comparing it to the transaction prices in the target auction, in those target auctions that the ring actually won. Since the target auction is conducted as an open outcry ascending auction (or English auction), the transaction price, when the ring wins, should be the valuation of the highest valuation non-ring bidder (or within a bid increment thereof). This measure would, most likely, give an incorrect measure of damages because of the strong equilibrium overbidding incentives in the knockout. This point becomes even clearer in the subsequent structural analysis. Nonetheless, Table 6 constructs this measure for various cuts of the data.

Table 6 here

Table 6 shows that the difference between the second-highest bid in the knockout and the price paid in the target auction increased as the target price increased. In the 25 auctions in which the target price exceeded \$10,000, this difference averaged \$3,820. Overall, the ring imposed \$514,900 in 'naive' damages. Across the auction houses, Matthew Bennett (a specialist in stamp auctions) stands out as suffering the greatest naive damages. This is due, in part, to the sale of some high-value lots that the ring was able to get at comparatively low prices. Lastly, as participation in the ring's knockout increased so did the naive damages.

The problem with interpreting these naive damages as real damages is that the incentive to overbid in the knockout is a confounding factor. Overbidding results in the naive damage estimates being an upper bound on the actual damages to the seller. To get a sense of the extent to which the naive damages are likely to be over-estimating these damages, it is necessary to estimate a structural model. Similarly, measures of naive damages give little indication of any damages suffered by bidders outside the cartel.

²⁷In an IPV setting this is consistent with D, G and H having value distributions that have a somewhat similar mean to that of the stronger bidders, but have a much lower variance. This configuration can mean that D, G and H are more likely to lose, but conditional on losing are more likely to have bids closer to the winning bid.

Structural Analysis

The value of imposing structure from an economic model on the data is that the structure can be used to infer the underlying distributions of valuations from the observed bids. This allows a wide range of questions to be asked that are not possible without a model, such as: How much inefficiency was introduced by the ring’s structure? How much did the ring benefit from colluding? How important were bidder asymmetries in determining the benefits from the ring structure? And, finally, what were the damages imposed by the ring?

To answer these questions, the model developed in Section 2 is estimated. The description of the structural analysis proceeds by first giving a non-technical overview of the econometric approach. This should give enough background to enable the reader to jump directly to the results and counterfactual simulations. A more-detailed treatment of the structural model and estimation procedure follows the overview.

Overview of the econometric implementation

The primary objective of estimation is to recover the distribution of valuations of cartel members in the knockout auction and the distribution of the value of the highest non-ring bidder. The valuation of bidder i in knockout auction k , u_{ik} , is modelled as

$$u_{ik} = \Gamma(x_k)(v_{ik}\varepsilon_k) \quad \text{where } \Gamma(x_k) = e^{x_k\gamma}$$

where x_k collects variables on which auctions are observed to differ; ε_k is auction-level heterogeneity observed by bidders but not the econometrician; and v_{ik} is the private value of the bidder.²⁸ In the basic model, without unobserved heterogeneity, $\varepsilon_k = 1$. The variables in x_k are: the estimated minimum and maximum values, a catalog estimate (if provided), the minimum and maximum quality grades, a dummy for lots with exclusively US stamps and a dummy for lots with no value estimates. The auction characteristics observed by bidders but not by the econometrician, ε_k , includes information not in the catalogs, but apparent from a physical inspection of the stamps as well as commonly understood conditions of the downstream stamp market. The bidders’ private information, v_{ik} , covers knowledge of the needs of specific clients, their cost of capital and private information about market conditions.

The objective is to recover the distribution of v_{ik} and, in the model accounting for unobserved

²⁸ Assuming this multiplicatively separable form for observed covariates allows both the mean and variance of the value distribution to vary with observed auction characteristics, albeit in a tightly parametrized fashion. This is attractive, as it reflects patterns observed in the data (for more details see the online appendix at <http://pages.stern.nyu.edu/~jasker/OnlineAppendix.pdf>)

heterogeneity, ε_k . With these distributions and the empirical distribution of the highest value of the non-ring bidders (r), meaningful counterfactual simulations can be run.

Due to limitations on identification in larger knockouts, the structural econometric analysis only uses data from the two-bidder knockout auctions. The ring members are classified into two groups: those who appear to be weak bidders (ring members D, G and H) and those remaining ring members who appear to be stronger. Within these groups, the distributions of bidders' valuations are assumed to be identical.

The rest of the theoretical framework for estimation follows that used to derive equation (2). It is assumed that bidders know that they are facing one other bidder but that the identity of this other bidder is uncertain. In particular, each bidder does not know whether the other bidder is strong or weak. The empirical frequencies of each bidder participating are substituted for the α_j 's used to form $G_{-i}(b)$ and $g_{-i}(b)$ in equation (2).

The econometric implementation begins by using the linear regression approach proposed in Haile, Hong and Shum (2006) to control for observed heterogeneity across auctions. This first stage is done by running an OLS regression on bids. The estimated residuals are used in the analysis that follows.

The structural analysis exploits the first-order condition (2) to estimate the mapping from observed bids to (inferred) underlying valuations. This model requires observed bids, the distribution of the bids of other ring members and the distribution of bids by non-ring members in the target auction as inputs in order to make inferences about underlying valuations. The distribution functions required as inputs can be estimated using standard nonparametric techniques. Crucial assumptions in performing this analysis are the independence of valuations across auctions and bidders, as well as the assumption that bidders know the number of other bidders in the knockout.

A problem exists in the data on the winning price in the target auction. These data only reflect the valuation of the highest-value non-ring bidder when the ring wins the target auction. Thus, these data suffer from a selection problem when being used to infer the distribution of the highest-value non-ring bidder. Since the process generating this selection problem is observed, a procedure is proposed that infers the 'un-selected' distribution of highest values from the 'selected' distribution and the distribution of winning knockout bids. This procedure exploits the independence assumption. It also assumes that the support of the 'selected' distribution is a subset of the distribution of winning knockout bids. This latter assumption seems reasonable given the overbidding phenomenon observed in equilibrium in the knockout auction.

Unobserved differences between auctions are controlled for using a deconvolution technique based on the estimation of empirical characteristic functions first proposed by Li and Vuong (1998), in the context of a classical measurement error problem, and applied to the auction setting by Li,

Perrigne and Vuong (2000) and Krasnokutskaya (2004). Here, the application of these techniques is very similar to that in Krasnokutskaya’s paper. This deconvolution technique uses variance across auctions to infer a distribution of a common component of values, ε_k , known to bidders but not to the econometrician. Variation of bids within auctions is used to infer the distribution of the private information element of valuations, v_{ik} .

Once the econometric analysis is completed, the estimated distributions of values, bids and the observed and unobserved heterogeneity can be combined to create a set of pseudo-data from which simulated damages and efficiency losses can be computed. The estimation procedure is sufficiently quick that confidence intervals can be computed for all inferred values using a standard bootstrap, resampling at the auction level.

Detail of the structural analysis

Nonparametric identification

The identification result in Guerre, Perrigne and Vuong (2000) is not immediately applicable to the problem of inferring the underlying valuations of the ring members due to the novel design of the knockout auction. However, the logic behind identification is, in spirit, the same. Since, given $F_r(\cdot)$ and $G_{-i}(\cdot)$, equation (2) provides a unique valuation, v , corresponding to each observed bid, the observed distribution of bids identifies a unique (latent) distribution of v . Guerre, Perrigne and Vuong (2000), being concerned with the identification of valuation distributions leading to a strictly increasing bidding equilibrium, require the inverse bid function to be increasing for identification. The same assumption is required here. As described earlier, assuming that observed bids are best responses, as per equation (2), implies that the potential non-monotonicity noted in Lemma 2 cannot occur in the underlying data generating process. Hence, a required maintained assumption is that the underlying inverse bid function is increasing.²⁹ As is standard in the literature, the existence of equilibrium is assumed, and it is further assumed that bids conform to a single equilibrium. In all other respects, identification is standard. For a discussion of identification of auction models with unobserved heterogeneity, see Li and Vuong (1998) and Krasnokutskaya (2004). In the model with unobserved heterogeneity, identification is a consequence of a statistical result by Kotlarski (1966).³⁰

²⁹For a small subset of the knockout bids, a problem exists in that they lie below the support of winning bids in the target auction. This means that the mapping in (2) is not defined. For this (economically unimportant) set of bids the estimation procedure sets $v = b$.

³⁰See Roa (1992).

Estimation approach: Inferring the Distributions of v_{ik} and r

The estimation approach adapts that of Guerre, Perrigne and Vuong (GPV). Estimation exploits equation (2) by estimating the various components of the first-order condition and backing out values. The steps in estimation are set out below. Steps specific to the model with unobserved auction heterogeneity are indicated with a U. The basic model, which does not account for unobserved auction-level heterogeneity, skips these steps.³¹

Step 1: Account for observed auction heterogeneity

Observed auction-level covariates can be controlled for using the first-stage regression approach in Haile, Hong and Shum (2006) and applied in Bajari, Houghton and Tadelis (2007) and elsewhere.³²

To adopt the first-stage regression approach in Haile, Hong and Shum, the following lemma proves useful.

Lemma 3: Under the maintained assumption that there exists a bidding equilibrium, there is an equilibrium in the knockout auction such that, if the optimal bid in the knockout when $u_{ik} = v_{ik}$ is b_{ik} , the optimal bid when $u_{ik} = \Gamma(x_k) v_{ik}$ is $\Gamma(x_k) b_{ik}$, provided $\Gamma(x_k) > 0$.

This Lemma is proved in the appendix. The same result is immediate for non-ring bidders in the (English) target auction. The value of lemma 3 is that, if we assume that the equilibrium described is being played in the data, it implies that bids can be normalized by estimating $\Gamma(x_k)$. The specification used here is $\Gamma(x_k) = e^{x_k \gamma}$, leading to the following first-stage regression

$$\ln b_{ik} = x_k \gamma + \eta_{ik} + \sigma_{ik} \tag{3}$$

where b_{ik} includes, for each auction k , the bids of each ring member in the knockout auction and the target price. That is, $i \in \{1, 2, \text{target}\}$. η_{ik} is a set of dummies for the target auction and individual ring members (equal to one if they are participating). σ_{ik} is an error term. The variables in x_k are: the estimated minimum and maximum values, a catalog estimate (if known), the minimum and maximum quality grades, a dummy for lots with exclusively US stamps and a dummy for lots with no value estimates. Equation (3) yields the normalized bid, $\ln \hat{b}_{ik} = \ln b_{ik} - x_k \hat{\gamma}$. The estimation of (3) is done jointly across all auction types (both knockout and target) so that the relative ordering of related bids is preserved. The reason why only the two-bidder knockouts are estimated is that lemma 3 does not hold for equilibrium bidding in the three-bidder, or higher, knockouts.³³

³¹Code is available from: <http://pages.stern.nyu.edu/~jasker>

³²For expositional clarity it is assumed that the $\varepsilon_k = 1$

³³For three-bidder knockouts, an additive version of Lemma 3 does hold. Nonetheless, this additive (rather than

Step 2(U): Separate the idiosyncratic and (unobserved) common elements of bids

This step follows the methods developed by Li and Vuong (1998) and applied in Li, Perrigne and Vuong (2000) and Krasnokutskaya (2004) to separate the unobserved (but common) heterogeneity from the idiosyncratic components of the knockout bids and target price. For expositional ease, I initially consider (hypothetical) data in which two bids are observed for each auction (that is, I initially set aside the presence of the target price). Lemma 3 can be exploited to write each (normalized) bid as a multiplicative function of a common and idiosyncratic component. This leads to

$$\ln \widehat{b}_{1k} = \ln(\varepsilon_k) + \ln \beta(v_{1k}) \quad \text{and} \quad \ln \widehat{b}_{2k} = \ln(\varepsilon_k) + \ln \beta(v_{2k}) \quad (4)$$

where the inverse of $\beta(v)$ is given by equation (2) and corresponds to that component of bids attributable to the idiosyncratic variation in bidders' valuations. The objective is to estimate the distributions of $\ln(\varepsilon)$, $\ln \beta(v_2)$ and $\ln \beta(v_1)$ from the observations of \widehat{b}_1 and \widehat{b}_2 . When there are two strong bidders in a knockout auction the distributions of $\ln \beta(v_2)$ and $\ln \beta(v_1)$ will be the same. Where there is a weak and a strong bidder these distributions are different.³⁴ The estimator exploits variation across auctions to identify the distribution of $\ln(\varepsilon)$, while the variation of bids within auctions identifies the distributions of $\ln \beta(v_1)$ and $\ln \beta(v_2)$.

The estimator proposed by Li and Vuong proceeds by estimating the characteristic functions of the joint distribution of b_1 and b_2 and then exploiting a statistical result by Kotlarski that shows that (under assumptions given below) there is a mapping from this characteristic function to the characteristic functions of the marginal distributions of interest. Densities are then recovered from these characteristic functions using an inverse Fourier transformation.

The empirical characteristic function is estimated nonparametrically using

$$\widehat{\psi}(z_1, z_2) = \frac{1}{n} \sum_{k=1}^K \exp\left(iz_1 \widehat{b}_{1k} + iz_2 \widehat{b}_{2k}\right)$$

The characteristic functions of the marginal distributions are estimated using

$$\begin{aligned} \widehat{\phi}_{\ln(\varepsilon)}(t) &= \exp \int_0^t \frac{\partial \widehat{\psi}(0, z_2) / \partial z_1}{\widehat{\psi}(0, z_2)} dz_2 \\ \widehat{\phi}_{\ln f(v_2)}(t) &= \frac{\widehat{\psi}(0, t)}{\widehat{\phi}_\varepsilon(t)} \quad \text{and} \quad \widehat{\phi}_{\ln f(v_1)}(t) = \frac{\widehat{\psi}(t, 0)}{\widehat{\phi}_\varepsilon(t)} \end{aligned}$$

This allows densities to be recovered by taking an inverse Fourier transformation

$$\widehat{g}_Y(x) = \frac{1}{2\pi} \int_{-T_n}^{T_n} d(t) \exp(-itx) \widehat{\phi}_Y(t) dt \quad \text{where } Y \in \{\ln(\varepsilon), \ln f(v_1), \ln f(v_2)\} \quad (5)$$

multiplicative) model was estimated and was found to exhibit results explicable only as reflecting strong model mis-specification.

³⁴ $\ln \beta(v_1)$ always corresponds to a strong bidder.

where $d(t)$ is a damping function (see Diggle and Hall (1993)). Assumptions that are required for this procedure are:

Assumption U: (a) \widehat{b}_{1k} and \widehat{b}_{2k} can be written as in (4) with $E(\ln \beta(v_1)) = 0$.

(b) $\ln(\varepsilon)$, $\ln \beta(v_2)$ and $\ln \beta(v_1)$ are mutually independent.

(c) The characteristic functions $\phi_{\ln(\varepsilon)}$, $\phi_{\ln \beta(v_2)}$ and $\phi_{\ln \beta(v_1)}$ are nonvanishing everywhere.

Li and Vuong provide additional regularity conditions on the smoothness and integrability of the characteristic functions.

Assumptions U(a) and U(b) have economic significance. U(b) exploits the independence of private information across bidders. U(a) follows from the identification strategy that the estimator exploits. Because the distributions of $\ln \beta(v_1)$ and $\ln \beta(v_2)$ are identified from within-auction variation, this leaves the position of the distributions ‘free.’ This is resolved by fixing the mean of $\ln \beta(v_1)$ at zero. The positions of other distributions are then determined relative to this reference point. The practical implication of this is that only distributions are identified. That is, unlike the GPV estimator, the specific valuation corresponding to an observed bid cannot be recovered. Lemma 3 is also crucial for this approach to work.

The data at hand add one complication. For each auction, an extra bid is observed: the transaction price in the target auction. When the ring wins the target auction, this reflects the highest value of the non-ring bidders (let these observations be denoted r^*). For the moment the selection problem embodied in this variable is set aside (it is dealt with in the next step). Recall that the distribution of this variable is needed to recover valuations from bids. As a consequence, the procedure outlined above is applied to both pairs $\{\ln \beta(v_2), \ln \beta(v_1)\}$ and $\{\ln \beta(r^*), \ln \beta(v_1)\}$. Being able to work with r^* has the additional advantage of providing an extra source of variation with which to identify the distribution $\ln(\varepsilon)$, as $\ln(r^*) = \ln(\varepsilon) + \ln(r)$ (the analog to (4)). Since the procedure can now be implemented on different types of pairs, this extra information is exploited by taking the average of the $\widehat{\phi}$ ’s implied by each type of pairing, weighted by the number of available observations of each combination (as suggested in Li and Vuong (1998)).

Having recovered the distributions of $\ln(\varepsilon)$, $\ln \beta(v_2)$, $\ln \beta(v_1)$ and $\ln(r)$, the distribution of the common element $\ln(\varepsilon)$ can be set aside, and step 4 (the Guerre, Perrigne and Vuong inversion of the bid function) can be done just using realizations of $\ln \beta(v_2)$ and $\ln \beta(v_1)$ and $\ln(r)$ drawn from the estimated distributions.

Step 3: Correct the selection bias in observations in r

The way the data are recorded makes the empirical distribution of r somewhat complicated to estimate. This is because the transaction price in the target auction reflects the highest value of

the non-ring bidders only when the ring is successful in the target auction. That is, the observed data on r are selected. From the data we can nonparametrically estimate $H_r(x | \max_{i \in I} b_i > x)$, $h_r(x | \max_{i \in I} b_i > x)$, $G(\max_{i \in I} b_i)$ and $g(\max_{i \in I} b_i)$ using empirical distribution functions and kernel density estimates. These, together with the assumption of independence of private information, can be used to compute $H_r(x)$ and $h_r(x)$. That is, an explicit statistical model is used to correct for the selection bias in the data on the bids from the target auction. For notational convenience, let $H_r(x | \max_{i \in I} b_i > x) = \bar{H}_r(x)$, $h_r(x | \max_{i \in I} b_i > x) = \bar{h}_r(x)$, $G(\max_{i \in I} b_i) = G_m(b)$ and $g(\max_{i \in I} b_i) = g_m(b)$. It follows that:

$$\begin{aligned} \bar{H}_r(r) &= \frac{\int_{-\infty}^r h_r(x) \int_x^{\infty} g_m(y) dy dx}{\int_{-\infty}^{\infty} h_r(x) \int_x^{\infty} g_m(y) dy dx} = \frac{1}{A} \int_{-\infty}^r h_r(x) \int_x^{\infty} g_m(y) dy dx \\ &= \frac{1}{A} \left[H_r(r) - \int_{-\infty}^r h_r(x) G_m(x) dx \right] \\ A \frac{\partial \bar{H}_r(r)}{\partial r} &= A \bar{h}_r(r) = h_r(r) [1 - G_m(r)] \end{aligned}$$

so

$$h_r(r) = A \frac{\bar{h}_r(r)}{[1 - G_m(r)]} \quad \text{and} \quad H_r(r) = A \int_{-\infty}^r \frac{\bar{h}_r(x)}{[1 - G_m(x)]} dx$$

A can be computed by evaluating $\left[\int_{-\infty}^{\infty} \frac{\bar{h}_r(x)}{[1 - G_m(x)]} dx \right]^{-1}$. This integral can be evaluated numerically. This approach employs an assumption that the support of bids in the knockout auction has an upper bound weakly greater than that of the highest valuation of non-ring bidders in the target auction. Given the incentive to ‘overbid’ created by the knockout structure, this does not seem likely to be restrictive.^{35,36}

Step 4: Recover the distribution of v_i

This is done using the procedure first suggested by Guerre, Perrigne and Vuong (2000). After de-logging, kernel density estimates of the distributions of $\beta(v_2)$ and $\beta(v_1)$ and r are obtained using a triweight kernel.³⁷ Estimates of distribution functions are obtained using the empirical dis-

³⁵In principle, the observed distribution of the target price could be used to provide a lower bound, adding more information and increasing the precision of the estimate. In the sample here, this extra information on the lower bound made no difference to the estimates.

³⁶Haile and Tamer (2003) observe that jump bidding and bid increments can break the link between the final price in an English auction and the valuation that is being exploited here. I ignore these issues since jump bidding was only very rarely observed at the stamp auctions I attended, and the data show evidence of auction houses accommodating very small bid increments relative to the levels of bids. Setting these empirical reasons aside, it is unclear, in this application, how to implement the bounds approach that Haile and Tamer develop.

³⁷The triweight kernel is defined as

$$K(u) = \frac{35}{32} (1 - u^2)^3 1(|u| < 1)$$

This kernel satisfies the conditions in Guerre, Perrigne and Vuong.

tribution functions. The α_j 's used to construct $G_{-i}(b)$ and $g_{-i}(b)$ are obtained using the empirical participation frequencies for each bidder. These estimates are then used with equation (2) to infer underlying values.

Step 5: Construct a set of simulated auctions

Having estimated the distributions of interest, it remains to construct a set of simulated auctions from which to construct estimates of damages, the return to the ring and efficiency losses. This is done by drawing from the estimated distributions of $\varepsilon, \beta(v_2), \beta(v_1)$ and r and constructing the valuations v_1 and v_2 implied by each drawn from $\beta(v_1)$ and $\beta(v_2)$ (this is the Guerre, Perrigne and Vuong procedure described in step 4). Then the model elements can be recombined to give a set of simulated auctions that correspond to the population from which the observed data are drawn.³⁸ Finally, auxiliary information from the data on the support of bids is used to exclude the 2 percent of simulation draws that have bids greater than the maximum observed bid or less than the smallest observed bid.³⁹

Practical Considerations

Executing the estimation procedure outlined above involves a series of practical issues. In Step 2, implementation of the estimator in equation (5) requires determination of the form of the damping function and determination of the smoothing parameter T_n . Step 3, in the model with unobserved heterogeneity, requires some trimming of the nonparametric density estimates. Step 4 requires selection of the bandwidth of the kernel estimator and trimming to avoid inconsistency of the estimator at the boundaries of the support of the value distribution.

The dampening function in equation (5) serves to reduce the impact of poorly estimated fluctuations in the tails of the characteristic function. Following Diggle and Hall (1993), I adopt a damping function of the form

$$d(t) = \max \left[1 - |t| / \left(\sqrt{2} T_n \right), 0 \right]$$

The smoothing parameter, T_n , in (5) operates in a similar way to the bandwidth in kernel estimation. To select the smoothing parameter, I first adopt the linear extrapolation approach proposed by Diggle and Hall. This sets a range of appropriate values of T_n . To refine the selection,

³⁸When adding in the observed auction-level heterogeneity draws are taken from the empirical joint distribution.

³⁹An implicit assumption underlying the validity of this simulated sample for the purposes of construction of counterfactuals is that the value distribution of the highest-value non-ring bidder be invariant to observed or unobserved auction-level heterogeneity. This is needed for the draws of r to be correct in expectation. Since the number of non-ring bidders is unobserved, little more can be done on this front.

I minimize a criterion function equal to

$$\text{Criterion} = \left(\frac{\widehat{ND} - ND}{ND} \right)^2 + \left(\frac{\widehat{PW} - PW}{PW} \right)^2 + P\chi_{\{non-monotonic\}}$$

where ND refers to naive damages and PW is the probability of the ring winning in the target auction.⁴⁰ $P\chi_{\{non-monotonic\}}$ is a penalty function that deters the search away from parameters that generate the non monotonicities in the bid function as highlighted in Lemma 2. The ‘hat’ notation denotes an estimated value, while the absence of a ‘hat’ denotes the quantity in the raw data. A different T_n is derived for each of $\ln(\varepsilon)$, $\ln \beta(v_s)$ and $\ln \beta(v_w)$ and $\ln(r)$, where subscript $s(w)$ refers to the strong (weak) group of bidders. In the data, $\ln(\varepsilon)$ appears to be very well identified and insensitive to the value of T_n chosen. As a result T_n is set to 50 for $\ln(\varepsilon)$. This reduces the dimensionality of the search for values of T_n that minimise the search criterion (above). This search yields values of T_n for $\ln \beta(v_s)$, $\ln \beta(v_w)$ and $\ln(r)$ of 11.0, 5.6 and 8.3 respectively.

Density estimates from the procedure in Step 2 suffer from being imprecise in the tails in finite samples. This leads to small positive densities being inferred over a very wide support. This creates particular problems for the selection procedure in Step 3, which requires that a ratio of estimated distributions and densities be taken. Hence, the imprecision in estimation in the extreme tails can lead to a very imprecise selection correction. This problem arises elsewhere in nonparametric econometrics, notably in nonparametric estimates of conditional moments (see, for instance, Robinson (1986) and Hardle and Stoker (1989)). This problem usually is dealt with by trimming in some fashion.

The approach taken here is very similar to that taken by Krasnokutskaya. The maximum within-auction difference between (logged) bids submitted by strong bidders is used to set the upper and lower bounds on the support of $\ln \beta(v_s)$, maintaining the requirement that $E \ln \beta(v_s) = 0$. The same region is used as the support for $\ln(r)$.⁴¹ To estimate the upper bound of the support of $\ln \beta(v_w)$ I use the maximal within-auction difference between a winning weak bid and losing strong bid, and add this to the lower bound of the support of $\ln \beta(v_s)$. The lower bound is derived analogously. If Θ_Y is the resulting support of variable Y , then the estimator in equation (5) is adjusted so that

$$\widehat{g}_Y(x) = \widehat{g}_Y(x) 1_{\{x \in \Theta_Y\}}$$

Kernel density estimates are estimated using a bandwidth determined by a rule-of-thumb approach. Each bandwidth is given by $h = 0.45 \times 1.06 \widehat{\sigma} n^{-\frac{1}{6}}$. Empirical distributions are also kernel-

⁴⁰The reduction of computational time demanded by this minimisation in the generation of bootstrap estimates of standard errors was greatly assisted by access to the NYU supercomputing infrastructure.

⁴¹Recall that a maintained assumption is that $\text{support}[\ln(r)] \subseteq \text{support}[\ln f(v_s)] \cup \text{support}[\ln \beta(v_w)]$. It turns out that $\text{support}[\ln \beta(v_s)] \supset \text{support}[\ln f(v_w)]$.

smoothed using a rule-of-thumb approach, with bandwidth given by $h = 1.06\hat{\sigma}n^{-\frac{1}{5}}$. Lastly, estimates of v use trimmed bid data following the method suggested in Guerre, Perrigne and Vuong and the subsequent literature.

Lastly, one extreme outlier is omitted from the sample. This auction had bids of \$42,500 and \$18,000 in the knockout (the next highest knockout bid across two-bidder auctions is \$13,500), and a recorded target price of \$18,000. The strange correspondence between the second knockout bid and the target price, and the fact that the highest knockout bid is more than three times greater than any knockout bid in any other two-bidder knockout, suggests that this auction is not drawn from the same data generating process.

Results

The structural model outlined above allows the estimation of the magnitude of damages, returns to the ring and the overall inefficiency resulting from the ring’s activity in target auctions with two ring members active in the associated knockout auction. After these results are reported, the fit of the model with unobserved auction-level heterogeneity is discussed together with various robustness considerations.

As a preliminary, the top panel of Figure 2 shows the estimated bidding function of a strong bidder (all ring members other than D, G and H) in the knockout, together with the distribution function of the highest valuation of the non-ring bidders (the bottom panel). The bids lie above the 45° line, indicating that ring members are bidding higher than their inferred valuations in the knockout auction. The estimated extent of this overbidding is the cornerstone of all the following analysis.

INSERT FIGURE 2 HERE

The lowest bids submitted by ring members lie below the lowest non-ring bid. These bids are rationalized by the estimating procedure as being equal to the bidders valuations.⁴² After this initial region, the overbidding declines steadily as the likelihood of a knockout bid winning the target auction increases. The lower panel of Figure 2 shows the probability of a bid in the knockout being enough to win in the target auction. A bidder with a valuation of \$800, who has a 42.0% chance of winning the target auction if he bids his value, bids \$852 in the knockout. This knockout bid of \$852 would have a 45.0% chance of winning the target auction. This gives some sense of

⁴²The valuations allocated to these bids are economically unimportant since these bidders never win a target auction or become eligible for sidepayments.

the economic magnitude of the overbidding effect induced by the ring mechanism. The bidding function for a weaker bidder is very similar.

Damages

Table 7 shows estimates of the damages imposed on sellers by the ring, when the ring wins. Estimates from both the model with unobserved auction-level heterogeneity and the model ignoring unobserved heterogeneity are presented. The difference between the two models is illustrated by comparing their predicted levels of naive damages ($=\max[\text{second-highest bid in the knockout minus the target price}, 0]$) and their predicted likelihoods of the ring winning. In the data, the mean naive damages, conditional on a ring win, is \$67. This compares with estimated naive damages of \$67 and \$150 for the models with and without unobserved heterogeneity (resp.). The impact of taking into account unobserved auction-level heterogeneity is easy to appreciate from a comparison of these naive damage estimates. When unobserved heterogeneity is not modelled, the level of naive damages is over 100 percent higher. This reflects, primarily, a greater within-auction variance in the simulated bids, due to the absence of an unobserved auction-specific common element. Interestingly, the estimated proportion of target auctions won by the ring in both models is close to that in the data (0.35 and 0.37 as compared to 0.36 in the data), although in both cases the confidence interval is wide.⁴³

INSERT TABLE 7 HERE

For those auctions won by the ring, the estimates of true damages are constructed by computing the price the seller receives when the ring is operating and subtracting the seller's price in the counterfactual in which all bidders (including the active ring members) bid competitively in the target auction. This counterfactual is easy to compute once the map between valuations and knockout bids has been made, since the target auction is a simple ascending-price English auction. Competitive bidding requires that bidders bid up to their valuations before dropping out.

These estimates of true damages are computed under two assumptions, labelled 'upper bound' and 'lower bound'. An exact estimate of damages requires the distribution of values of the second-highest non-ring bidder to be known. This is required to construct the price under competitive bidding since the second-highest non-ring bidder may ultimately set the price. As mentioned earlier,

⁴³The first stage OLS projection of bids on observables also plays an important role in reducing auction level heterogeneity. For each knockout auction k , I compute the mean bid, μ_k . The standard deviation of the μ_k 's prior to the OLS treatment is 1367, after treatment the standard deviation of the μ_k 's computed from the normalised bids is 370.

the relative weakness of these data is the lack of information about bidding in the target auction.⁴⁴ A consequence is that reliable estimates of the level of the second-highest non-ring valuation cannot be obtained. In the face of this, damages are estimated under the ‘upper bound’ and ‘lower bound’ assumptions (U.B. and L.B., resp.). Under U.B. the second-highest non-ring value is assumed to be equal to the highest non-ring valuation. Under L.B. the second highest non-ring value is assumed to be equal to the minimum of the highest non-ring valuation and the highest ring valuation. U.B. gives an (loose) upper bound to damages since it generates the highest model-consistent price under competitive bidding. L.B. gives a (tighter) lower bound to damages since it generates the lowest model-consistent price under competitive bidding.

In Table 7, the mean damages estimated using the model with unobserved heterogeneity are equal to \$35 and \$28 (under U.B. and L.B., resp.). When unobserved heterogeneity is ignored the estimates are \$105 and \$98. Both models illustrate the importance of appropriate equilibrium analysis of damages, in that in both instances true damages are substantially less than the naive damage estimates.⁴⁵

The rest of Table 7 decomposes the auctions won by the ring into three groups: auctions in which the competitive price and the ring price are equal; auctions in which the seller benefited from the ring ($Pr > Pc$); and auctions in which the ring hurt sellers ($Pr < Pc$). The sellers incurred harm in 23 percent of auctions that the ring won, with the ring decreasing prices by, on average, 17 percent. These estimates are invariant to the U.B or L.B. assumptions (by construction).

Under the L.B. assumption, in 18 percent of auctions won by the ring the seller actually benefited. This is because the ring pushed the price up higher than it would have been had all bidders been bidding competitively. In these auctions, the overbidding in the knockout carried over to the target auction, pushing the price up to the highest valuation across all bidders. Under the U.B. assumption, this benefit from the ring can never occur since the highest non-ring valuation is shared by two bidders by assumption. When the seller benefits, prices are inflated by five percent, under the L.B. assumption.

⁴⁴In principle, the data do contain information about the distribution of the second-highest non-ring valuation in the target prices that lie above the highest knockout bid. This raises the possibility of obtaining the desired distribution from a similar selection adjustment used to obtain the distribution of the highest non-ring valuation. In practice, however, this procedure works very poorly due to the support of the highest knockout bids not extending low enough to provide sensible estimates of the left hand tail of the distribution of the second-highest non-ring valuation. Since this is the most important part of the distribution from a damages perspective, the approach was abandoned in favour of the bounds approach.

⁴⁵The 90% confidence intervals on the ratio of true to naive damages are, for the model with unobserved heterogeneity, [0.36,0.69] (U.B. assumption) and [0.19,0.63] (L.B. assumption). For the model without unobserved heterogeneity, the corresponding intervals are [0.53,0.75] and [0.42,0.72].

It is important to note that an unmeasured source of seller benefit exists: those instances where a non-ring bidder wins, but the ring’s propensity to overbid forces the winning bidder to pay more than would be the case under competitive bidding. This means that the damage estimates overstate the damages suffered by sellers. An upper bound on the size of this effect is given in the discussion of Table 8, below.

This finding, that the ring was not always detrimental to sellers, is reflected in the case materials. In D’s deposition, his sense was:

“That probably in the end result the ring brings in as much money to the auction house as if it wasn’t there, and obviously part of that was, but there was a lot of truth in that. Because there are six or eight people that show up to an auction, that normally would not come to the auction and end up competitively bidding against the rest of the people on the floor. Did I say that right, does it make sense?”

In addition to the seller, non-ring bidders in the target auction could suffer damages. There are two sources of potential damages for these bidders: first, they could win the object but pay an inflated price from competing against the ring’s overbidding; or, second, the ring could win the object, but at a price higher than its valuation, resulting in misallocation. When misallocation occurs, the other bidders’ damages are computed as their surplus captured under competitive bidding but forgone in the presence of the ring.

INSERT TABLE 8 HERE

Table 8 reports damages to non-ring bidders. Focusing on the model with unobserved heterogeneity, the ring wins the target auction 35 percent of the time. Of the auctions that the ring wins, 18 percent result in a misallocation. Conditional on the ring winning the auction, damages to non-ring bidders average \$7, using the L.B. assumption (by construction, no damages to non-ring bidders can occur under the U.B. assumption). This amounts to 26 percent of the damages suffered by sellers under the same assumption.

Damages due to inflated prices can be assessed using the L.B. assumption (again, by construction, no damages to non-ring bidders can occur under the U.B. assumption). Using estimates from the model with unobserved heterogeneity, the L.B. assumption provides an upper bound on the extent of damages to non-ring bidders from price inflation due to ring overbidding. The average level of these damages is estimated as \$98, where the average is taken across all auctions that the ring loses. This level of damages is high as compared to the damages incurred by sellers. At least in part, this is due to the fact that, under the L.B. assumption, these damages are incurred in every auction that the ring loses. A more realistic sense of the magnitude of these damages is obtained

by noting that in the raw data, the target price is equal to the highest knockout bid in 11.5 percent of the auctions that the ring does not win. This suggests that the estimated \$98 should be deflated by a factor of five to give a number comparable to the damages from misallocation. (Damages are incurred approximately ten percent of the time the ring loses, but the ring loses twice as often as it wins.) This gives a measure of damages due to price inflation of just under \$20. When added to the damages from misallocation, this suggests that damages incurred by non-ring bidders are likely to be at least as high as damages incurred by the seller. This last point rests on the observation that damages due to price inflation represent a direct transfer from non-ring bidders to the seller, reducing the sellers' damages by the same amount as they increase the non-ring bidders' damage.

Inefficiency

The misallocation that the ring's overbidding can create leads to the possibility of inefficiency being introduced into the target auction. This is in itself interesting since, in I.P.V environments, ascending-price English auctions result in efficient allocations. Models of ring behaviour that assume that the ring is able to act efficiently do not lead to any inefficiency in English auctions, but merely change the magnitude of the transfers (see Graham and Marshall (1987) for an example). Here the ring does introduce inefficiency, suggesting an additional economic justification for antitrust enforcement.

INSERT TABLE 9 HERE

That said, Table 9 shows that the efficiency impact of the ring is small. When the ring wins, the average efficiency loss is \$7. This represents less than one percent of the value generated from the optimal allocation (reported as 'Mean proportional efficiency losses: Ring active'). To give some indication of the sense in which this is small, the proportional efficiency loss from excluding the ring bidders from the auction and the proportional efficiency loss from excluding the other non-ring bidders were computed. The model with unobserved auction-level heterogeneity estimates that these events result in a 8 percent and 27 percent decrease in efficiency relative to the optimal allocation, respectively.

Hence, while the ring does have some effect on market efficiency, the effect appears small. That said, if the damages suffered by non-ring bidders are as large as those indicated by the L.B. assumption exploited in Table 8, then this raises the possibility that the ring's activities may discourage participation. Clearly, this is something that the ring itself would be keen to accentuate. The estimates of the effect of excluding groups of bidders from the auction suggest that any impact of the ring on participation may have an economically significant effect on market efficiency. A

similar, unmeasured, welfare loss may arise if sellers are deterred from bringing lots to auction by the ring's activity. The data (and depositions), however, allow nothing more than speculation on these points.

Returns to the ring

The structural model also allows the returns to engaging in ring behaviour to be estimated. The benefit that the ring enjoys from reducing competition will be offset to some extent by the incentive to overbid. This can result in the ring winning an auction at a price higher than its value. Thus, the ring will harm itself from time to time. Focusing on the model with unobserved heterogeneity, Table 10 shows that the ring manages to harm itself in this fashion in 19 percent of the auctions it wins. Dividing the total losses by the number of auctions won, this yields an estimate of average harm of \$7, matching the mean efficiency loss reported in Table 9.⁴⁶ This capacity for self harm is offset by the benefits of easing competition among ring members. This positive effect contributes an average benefit of \$35, resulting in a net average return to the ring from winning an auction of \$28.

INSERT TABLE 10 HERE

These estimates allow the actual ring mechanism to be compared to an ideal ring that extracts all the surplus available to a cartel. The returns to an ideal cartel are represented by the \$35 estimated as the beneficial component of the ring's return. On this basis, the ring appears to be capturing 79% of the surplus available to a cartel, with the difference reflecting the costs incurred by having to deal with the practical problems imposed by bidder heterogeneity and a need to have a fairly simple mechanism. The resulting capacity for the ring to harm itself by paying too much for a lot is reflected in the case materials. D's deposition contains the following discussion of why a ring member exited the ring in the late 1980s:

Q: Did he give a reason why it would be better for him to bid himself?

A: He felt that the prices in the ring were so officially inflated that if you ended up buying it, you ended up buying [paying ?] too much.

The gains to the ring from colluding can be decomposed further into those that accrue to the strong bidders and those that accrue to the weaker bidders. Strong bidders capture \$16.2 and weaker bidders capture \$11.3 of the total \$28 that the ring as a whole earns in expectation each time that they win. Given that there are only 3 weak bidders this means that they capture a

⁴⁶These two numbers measure the same thing: the difference between the target price and the highest ring value, when the ring wins at a price above its valuation.

disproportionate amount of the ring’s gains (as is also the case in the raw sidepayments data). This raises the issue of why the strong bidders do not exclude the weaker ones. Indeed, such a move would raise the returns to stronger ring members from \$16.2 to \$16.5. The fact that there appears to be some small gain from excluding the weaker bidders, but that they remained in the ring, is consistent with the presence of a hold-up problem created by the possibility that an excluded member may inform antitrust authorities of the ring’s activities (and enjoy immunity from prosecution due to leniency provisions benefiting informants).

Model fit and robustness

The results reported above focus on the structural model that explicitly accounts for auction-level heterogeneity unobserved by the econometrician but observed by the bidders. Some aspects of the fit of this model have already been commented on, particularly in the discussion of Table 7. Figure 3 allows a more specific evaluation of the fit of the model by comparing the densities of several of the model’s simulated variables to related densities from the raw data.

INSERT FIGURE 3 HERE

Panel A reports the density of the bids across knockouts from the data and the simulation. It also shows the density of the inferred highest value of non-ring bidders. The densities of the bids have the same shape, albeit with the simulation having slightly thicker tails. A useful measure of fit in comparing these densities is the integrated absolute error, I.A.E.= $\int |f(x) - \hat{f}(x)| dx \in [0, 2]$. This is estimated for both the full model and the model ignoring unobserved heterogeneity. In the full model I.A.E. is equal to 0.201 while in the model ignoring unobserved heterogeneity it is 0.280 (hereafter quoted in parentheses). The density of the highest non-ring bidder lies to the right of the knockout bids. Given that, in the data, the ring wins only 36 percent of auctions, this is what would be expected.

In panel B, the density of the ratio of the highest non-ring value to the highest knockout bid (from the simulation) is shown along with the ratio of the observed target price to the highest knockout bid from the data. Ideally, for values of this ratio that are less than 1 these densities should match, since the observed target price in this region reflects the highest non-ring value. For the region where these ratios are greater than 1, the density from the simulation should place more mass on higher valuations than the density from the raw data since the observed target price in the raw data reflects the value of the second-highest non-ring valuation. The simulation model does fill some of these requirements, but not perfectly. As the ratios approach 1 from below, the densities start to diverge. This is the main source of unsatisfactory fit. The I.A.E. for these densities is 0.505 (0.922), although there is no suggestion that they should correspond over the entire support.

Panel C shows the same set of ratios as in panel B, but substituting the second-highest bid in the knockout (I.A.E. = 0.286 (0.899)). Panel D shows the densities of the ratio of highest strong knockout bid to second-highest strong knockout bid (I.A.E. = 0.160 (0.506)). Panel E shows the ratio of weak knockout bids to strong knockout bids (I.A.E. = 0.282 (0.401)). In all three cases the shapes of the densities in the data are reflected in the simulation, albeit with some errors.

In all these examinations of model fit the sources of disparity between the model and data are the same. First, the model is estimated from marginal distributions whereas (particularly in panels C, D and E) fit is being judged against the characteristics of joint distributions in the data. Second and related, the model puts a great deal of emphasis on the independence assumption in constructing estimates. Where the data and model diverge gives some indication of where this assumption is particularly strong.

As has been pointed out above, the econometric model estimated in this paper makes considerable use of the maintained assumption of independence across its many moving parts. In particular it is assumed that the idiosyncratic component of values are private and independent. If they were correlated in some way this would introduce mis-specification into the model. The bias in inference resulting from such mis-specification is not obvious. If some component of valuations is best modelled using a common values model, the extent to which overbidding is increased or decreased is unclear. Overbidding in this context means willingness to bid higher than would have been the case had the ring not existed. The results in Milgrom and Weber (1982) give some basis for conjecturing that the ring's propensity to overbid may be diminished via the use of a first-price style knockout (although the impact of the sidepayment scheme may add considerable complications). Examination of ring activity such as this, in a common value setting, seems a fruitful avenue for further research.

This paper focuses attention on the auctions associated with two-bidder knockouts, due to limits on identification. An assumption made in doing so is that ring bidders know that they are only faced with one other ring bidder. To the extent that some two-bidder knockouts have uncertainty associated with them as to the number of bidders, this raises the concern that the sample is selected in a way that may give rise to misleading inference due to model mis-specification.⁴⁷ To address this concern a propensity score (for the event that the auction has two bidders) for each auction based on auction observables was estimated. The model was re-estimated on two-bidder knockouts with propensity scores that places them in the top 50% of two-bidder knockouts. The full results of this exercise are reported in the online appendix.⁴⁸

⁴⁷Where such uncertainty exists ring bidders will place a bid that is a weighted average of the optimal bids in the two-, three- (and so on) knockouts that they think might be possible.

⁴⁸Available at: <http://pages.stern.nyu.edu/~jasker/OnlineAppendix.pdf>

With the full model (accounting for unobserved heterogeneity), damage estimates using the selected sample tend to be higher than those estimated using the full sample. However, each point estimate, derived from either data set, lies within the corresponding 90% confidence interval derived from the other data set. In the same vein, the results have the same qualitative features. Hence, this aspect of focusing on two-bidder knockouts appears not to have a significant effect on the economic findings.

Another general robustness issue is the extent to which the results from these two bidder knockout settings can be extrapolated to settings where more than two bidders are active in the knockout. For those auctions in which more than two ring members were active, damages suffered by the seller may increase due a possibly increased likelihood that a ring member has the second highest value, but this may be mitigated by an increase in the extent of overbidding (Graham, Marshall and Richard (1990) find that overbidding increases as the number of ring members increases) or a greater increase in participation by outside bidders. A similar trade-off needs to be weighed in forming conjectures as to the damages suffered by non-ring bidders in these auctions. This suggests that there is no clear bias in using features of these damage estimates to extrapolate to the other auctions in which the ring was active.⁴⁹

Conclusion

The ring examined in this paper faced a series of difficult challenges. Aside from the usual issues of coordination and avoiding detection, it needed to accommodate a wide range of heterogeneous bidders and it needed to be able to handle many transactions quickly. The mechanism that the ring adopted has features that seem directed at these issues, notably that sidepayments increased as the apparent importance of a bidder increased, as measured by the size of the knockout bid.

This came at a cost, however. The ring mechanism introduced inefficiency into the market that had the capacity to harm the ring. On balance, the evidence suggests that, overall, the ring benefited from coordinating bidding behaviour, but it was certainly the case that the benefits from ring activity were diminished by the overbidding incentives created by the ring design.

Another consequence of the design are the damages that appear to have been imposed on bidders who were not ring members. By introducing inefficient allocations and occasionally driving prices above the competitive level, the ring imposed damages on non-ring bidders via two channels. Although only an upper bound on these damages can be estimated, it seems likely that these

⁴⁹If the ratio of naive damages to seller damages (conditional on the ring winning) were used to extrapolate damages, this would mean that sellers lost somewhere between \$207,000 and \$269,000 in total in those auctions that the ring won.

damages are economically significant.

The implications of the ring for seller revenues are less clear. It appears possible that the damages from the ring's coordinated bidding may be offset by the capacity for the ring to push up prices from time to time.

The likely damages imposed on non-ring bidders suggest that an important issue, impossible to examine in these data, is the extent to which ring activity can discourage participation by other bidders. The results reported here suggest that adverse participation effects may have the capacity to dwarf other sources of inefficiency and damages. A deeper understanding of the extent of this possible source of market distortion requires further research.

The finding that non-ring bidders can suffer non-trivial damages also suggests that the way that damages are computed and distributed in bid-rigging cases, where the target auction is an English auction, is worthy of reconsideration. Under both the Sherman and Clayton Acts, any person "injured in his business or property by anything forbidden in the antitrust laws" is capable of suing for treble damages, subject to the damage being sufficiently proximate (a notion defined in the case law). The findings in this paper suggest that the damages imposed on other bidders by a ring can be both sufficiently large and proximate so as to justify the recovery of damages by these parties. However, to the best of my knowledge, there has yet to be a claim against a ring brought by bidders outside the ring. The findings in this paper give some support to such a claim being brought in the future.

Appendix

Proof of Lemma 3

For ease of exposition let $\Gamma(x_k) = \Gamma$. If $A(x)$ is the cdf of a random variable x and there exists a random variable y such that $y = \frac{x}{\Gamma}$ and $B(y)$ is the cdf of y , then $A(x) = B\left(\frac{x}{\Gamma}\right)$ and $a(x) = \frac{1}{\Gamma}b\left(\frac{x}{\Gamma}\right)$, where $a(x)$ and $b(y)$ are the density functions of x and y .

Having noted this, it is straightforward to derive the first order conditions for each ring member's bidding problem when the ring member's value is given by Γv_i and the highest non-ring member's value by Γr . Under these conditions each first order condition can be written as

$$\frac{1}{\frac{1}{\Gamma}(\Gamma v_i - b)} = - \frac{h\left(\frac{b}{\Gamma}\right) F\left(\frac{\phi_{\Gamma}(b)}{\Gamma}\right) + H\left(\frac{b}{\Gamma}\right) f\left(\frac{\phi_{\Gamma}(b)}{\Gamma}\right)}{\frac{1}{2}H\left(\frac{b}{\Gamma}\right) \left(1 - F\left(\frac{\phi_{\Gamma}(b)}{\Gamma}\right)\right)} \quad (\text{A1})$$

where $F_{-i}(\phi_{\Gamma}(z)) = \frac{\sum_{j \neq i} \alpha_j F_j(\phi_{j\Gamma}(z))}{\sum_{j \neq i} \alpha_j}$, $f_{-i}(\phi_{\Gamma}(b)) = \frac{\sum_{j \neq i} \alpha_j f_j(\phi_{j\Gamma}(z)) \phi'_{j\Gamma}(z)}{\sum_{j \neq i} \alpha_j}$ is the corresponding density and $\phi_{j\Gamma}(b)$ is a function mapping from the bids of ring member j to the corresponding value of Γv_j . Thus (A1) describes a system of 11 equations, one equation for each ring member.

Equation (A1) provides the first order condition defining $b_i(\Gamma v)$. What is required is to show that $\Gamma b_i(v)$ satisfies this same condition for all i . This is done by constructing a function ω_i , based on $b_i(v)$, and showing that it satisfies the above first order condition. It follows that the set of functions $\omega \ni \omega_i$ corresponds to $b(\Gamma v) \ni b_i(\Gamma v)$, which in turn allows the required relationship between $b_i(v)$ and $b_i(\Gamma v)$ to be established for all i .

For each i let $\omega_i(\Gamma v) = \Gamma b_i(v)$ so that $\omega_i(z) = \Gamma b_i\left(\frac{z}{\Gamma}\right)$. Noting that $\omega_i^{-1}(\omega_i(z)) = \Gamma b_i^{-1}\left(b_i\left(\frac{z}{\Gamma}\right)\right)$, and defining $\phi_{i1}(z) = \phi_{i\Gamma}(z)$ when $\Gamma = 1$, this lets $\frac{\omega_i^{-1}(x)}{\Gamma} = \phi_{i1}\left(\frac{x}{\Gamma}\right)$ and $(\omega_i^{-1}(x))' = \phi'_{i1}\left(\frac{x}{\Gamma}\right)$.

Substituting the set of functions $\omega^{-1}(b) \ni \omega_i^{-1}(x)$ for $\phi_{\Gamma}(b)$ in (A1) gives

$$\frac{1}{\frac{1}{\Gamma}(\Gamma v_i - b)} = - \frac{h\left(\frac{b}{\Gamma}\right) F\left(\frac{\omega^{-1}(b)}{\Gamma}\right) + H\left(\frac{b}{\Gamma}\right) f\left(\frac{\omega^{-1}(b)}{\Gamma}\right)}{\frac{1}{2}H\left(\frac{b}{\Gamma}\right) \left(1 - F\left(\frac{\omega^{-1}(b)}{\Gamma}\right)\right)}$$

which can now be rewritten as

$$\frac{1}{v_i - \frac{b}{\Gamma}} = - \frac{h\left(\frac{b}{\Gamma}\right) F\left(\phi_1\left(\frac{b}{\Gamma}\right)\right) + H\left(\frac{b}{\Gamma}\right) f\left(\phi_1\left(\frac{b}{\Gamma}\right)\right) \phi'_1\left(\frac{b}{\Gamma}\right)}{\frac{1}{2}H\left(\frac{b}{\Gamma}\right) \left(1 - F\left(\phi_1\left(\frac{b}{\Gamma}\right)\right)\right)} \quad (\text{A2})$$

If we substitute $\frac{b}{\Gamma}$ in equation (A2) with b^* , it becomes clear that (A2) is just the set of first order conditions for bids when $\Gamma = 1$. This means that the set of functions $\omega(\Gamma v)$ satisfy equation (A1). Hence, there exists an equilibrium such that:

$$b_i(\Gamma v) = \phi_{i\Gamma}^{-1}(\Gamma v) = \omega_i(\Gamma v) = \Gamma b_i\left(\frac{\Gamma v}{\Gamma}\right) = \Gamma b_i(v) \quad \forall i$$

References

- ATHEY, S., J. LEVIN and E. SEIRA (2004), Comparing Open and Sealed Bid Auctions: Theory and Evidence from Timber Auctions, Working paper.
- BAJARI, P., S. HOUGHTON and S. TADELIS (2007), Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs, working paper
- BAJARI, P. and LIXIN YE (2003), Deciding Between Competition and Collusion, *Review of Economics and Statistics*, 85, 971-989.
- BALDWIN, L., R. MARSHALL and J.-F. RICHARD (1997), Bidder Collusion at Forest Service Timber Auctions, *Journal of Political Economy*, 105(4), 657-699.
- CASSADY, R. (1967), *Auctions and Auctioneering*, University of California Press, Berkeley.
- CHE, Y.-K. and J. KIM (2006), Optimal Collusion-Proof Auctions, working paper.
- DELTAS, G. (2002), Determining damages from the operation of bidding rings: An analysis of the post-auction knockout sale, *Economic Theory*, 19, 243-269
- DIGGLE, P. and P. HALL (1993), A Fourier Approach to Nonparametric Deconvolution of a Density Estimate, *Journal of the Royal Statistical Society. Series B (Methodological)*, 55(2), 523-531.
- FREEMAN, A. and J. FREEMAN (1990), *Anatomy of an Auction: Rare Books at Ruxley Lodge 1919*, The Book Collector, London.
- GENESOVE, D. and W. MULLIN (2001), Rules, Communication and Collusion: Narrative Evidence from the Sugar Institute Case, *American Economic Review*, 91(3), 379-390.
- GRAHAM, D. and R. MARSHALL (1987), Collusive Bidder Behavior at Single-Object Second-Price and English Auctions, *Journal of Political Economy*, 95(6), 1217-1239.
- GRAHAM, D., R. MARSHALL and J.-F. RICHARD (1990), Differential Payments Within a Bidder Coalition and the Shapley Value, *American Economic Review*, 80(3), 493-510.
- GUERRE, E., I. PERRIGNE and Q. VUONG (2000), Optimal Nonparametric Estimation of First-Price Auctions, *Econometrica*, 68(3), 525-574.
- HAILE, P., H. HONG and M. SHUM (2006), Nonparametric Tests for Common Values in First-Price Sealed-Bid Auctions," NBER Working Paper 10105.
- HAILE, P. and E. TAMER (2003), "Inference with an Incomplete Model of English Auctions," *Journal of Political Economy* 111, 1-52.
- HARRINGTON, J. (2005), "Detecting Cartels," in P. BUCCIROSSI (ed.), *Handbook in Antitrust Economics*, (MIT Press), forthcoming.
- HARDLE, W. and T. STOCKER (1989), Investigating Smooth Multiple Regression by the Method of Average Derivatives, *Journal of the American Statistical Association*, 84, 986-995.

- HENDRICKS, K., R. PORTER and G. TAN (2003), Bidding Rings and the Winner's Curse: The Case of Federal Offshore Oil and Gas Lease Auctions, NBER Working Paper 9836.
- KOTLARSKI, I. (1966). "On Some Characterizations of Probability Distributions in Hilbert Spaces," *Annali di Matematica Pura et Applicata*, 74, 129-134.
- KRASNOKUTSKAYA, E. (2004), Identification and Estimation in Highway Procurement Auctions Under Unobserved Auction Heterogeneity, working paper.
- KWOKA, J. (1997), The price effects of bidding conspiracies: evidence from real estate auction knockouts, *Antitrust Bulletin*, 42(2), 503-516
- LEVITT S. and S. VENKATESH (2000), An Economic Analysis of a Drug-Selling Gang's Finances, *Quarterly Journal of Economics*, 2000, 115(3), 755-89.
- LI, T. and Q. VUONG (1998), Nonparametric Estimation of the Measurement Error Model Using Multiple Indicators, *Journal of Multivariate Analysis*, 65, 139-165.
- LI, T. I. PERRIGNE and Q. VUONG (2000), Conditionally independent private information in OCS wildcat auctions, *Journal of Econometrics*, 98, 129-161.
- LIZZERI, A. and N. PERSICO (2000), Uniqueness and Existence of Equilibrium in Auctions with a Reserve Price, *Games and Economic Behaviour*, 30, 83-114.
- LOPOMO, P., L. MARSHALL and L. MARX (2005), Inefficiency of Collusion in English Auctions, *B.E. Journals: Contributions to Theoretical Economics*, 5(1), Art. 4.
- MAILATH, G. and P. ZEMSKY (1991), Collusion in Second Price Auctions with Heterogeneous Bidders, *Games and Economic Behaviour*, 3, 467-486.
- McAFEE, R. P. and J. McMILLIAN (1992), Bidding Rings, *American Economic Review*, 82(3), 579-599.
- MARSHALL, R. and L. MARX (2007): "Bidder Collusion," *Journal of Economic Theory*, 133, 374-402.
- MILGROM, P. and R. WEBER (1982), A Theory of Auctions and Competitive Bidding, *Econometrica*, 50, 1089-1122.
- PORTER, R. (1992), Review of: FREEMAN, A. and J. FREEMAN (1990) Anatomy of an Auction: Rare Books at Ruxley Lodge 1919, *Journal of Political Economy*, 100(2), 433-436.
- PORTER, R. and D. ZONA (1993), Detection of Bid Rigging in Procurement Auctions, *Journal of Political Economy*, 101, 518-538.
- PORTER, R. and D. ZONA (1999), Ohio School Milk Markets: An Analysis of Bidding, *RAND Journal of Economics*, 30, 263-288
- PESENDORFER, M. (2000), A Study of Collusion in First Price Auctions, *Review of Economic Studies*, 67, 318-411.
- ROA, B. (1992). Identifiability in Stochastic Models: Characterization of Probability Distributions

utions, Academic Press, New York.

ROBINSON, P. (1986) On the Consistency and Finite Sample Properties of Nonparametric Kernel Time Series Regression, Autoregression and Density Estimators, *Annals of the Institute of Statistical Mathematics*, 38, 539-549.

ROLLER, L.-H. and F. STEEN (2006), On the Workings of a Cartel: Evidence from the Norwegian Cement Industry, *American Economic Review*, 96(1), 321-338.

SKRZYPACZ, A.. and H.. HOPENHAYN (2004) Tacit Collusion in Repeated Auctions, *Journal of Economic Theory*, 114, 153-169.

SMITH, C. (1989), *Auctions: The Social Construction of Value*, The Free Press (Macmillan), New York.

WRAIGHT, R. (1974), *The Art Game Again!*, Leslie Frewin Publishers, London.

Table 1: Ring activity by auction house

	Target Auction		Knockout Auction		% Of lots won by ring	% of Total Est. Value Won	Total number of lots	Number of Sales
	Mean	Standard Dev.	Mean	Standard Dev.				
Christie's	1620	1680	1223	1629	19%	51%	63	1
Daniel Kelleher	885	1203	968	2003	66%	56%	82	1
HR Harmer	1296	1767	1277	2060	63%	66%	669	2
Ivy Mader	1424	1584	1327	1757	46%	55%	153	1
Matthew Bennett	3727	4657	4474	5787	72%	75%	231	1
Robert Siegel	1829	3011	1806	4443	47%	52%	380	3
Sotheby's	4003	4270	3371	3746	42%	51%	125	1
Spink America	1729	2157	1884	2567	48%	46%	78	1
Aggregate of Above	2042	3142	2142	3942	58%	66%	1781	11
Harmer-Schau	735	1129	867	1393	87%	n/a	189	1

Notes: All auction bids are in dollars. The number of sales refers to an event in which many lots were sold. Thus, sales often run over several days. Aggregate records the values corresponding to the pooled data set, excluding the Harmer-Schau auctions. In the Harmer-Schau auctions the ring operated using an informal arrangement that differed from other recorded auctions. '% of Total Est. Value Won' is calculated by taking the average of the Estimated Min and Max values reported in catalogs (or the reported book value if Est Min and Max are not present) for lots the ring wins, taking the sum and dividing by total across all lots

Table 2: Bidding by number of bidders in the knockout

# of Bidders	Target Auction (Winning Bid)		Knockout Auction (Median Bid)		% Of lots won by ring	Total Number of lots
	Mean	Standard Dev.	Mean	Standard Dev.		
1	733	1262	616	1134	19%	623
2	1314	2016	1066	2048	36%	367
3	2014	3246	1750	3029	48%	260
4	2217	3492	2293	4082	69%	196
5	2249	3419	2092	3322	68%	144
6	2098	2628	2163	3014	74%	91
7	2979	3425	3655	4116	86%	74
8	4790	4904	6233	7726	96%	26

Notes: Does not include the Harmer-Schau auctions. All subsequent analysis also excludes these auctions.

Table 3: Participation in knockout auctions

Specification	OLS		Logit	
	Coeffs.	Std Err.	Coeffs.	Std Err.
Constant	1.596	(0.232)	-1.820	(0.117)
(Emin + Emax)/2	0.060	(0.049)	0.031	(0.024)
(Emax – Emin)	0.402	(0.204)	0.196	(0.103)
Catalog Price	-0.005	(0.007)	-0.001	(0.004)
(Grade Min + Grade Max)/2	0.112	(0.047)	0.076	(0.023)
(Grade Max – Grade Min)	0.377	(0.043)	0.198	(0.022)
No Grade	-1.181	(0.148)	-0.492	(0.074)
Exclusively US	0.875	(0.187)	0.495	(0.094)
No Value	-1.387	(0.515)	-0.756	(0.259)
House HRH	0.438	(0.213)	0.076	(0.107)
House DK	0.453	(0.274)	0.186	(0.137)
House IM	-0.365	(0.242)	-0.276	(0.121)
House MB	2.043	(0.249)	0.862	(0.125)
House RS	0.535	(0.219)	0.165	(0.110)
House S	-0.478	(0.253)	-0.322	(0.127)
House SA	-0.489	(0.273)	-0.404	(0.137)
R-squared	0.261		0.255	
SER	1.609		0.808	
Observations	1781		1781	

Notes: Both specifications were implemented via an OLS procedure. In the OLS specification the dependent variable is the number of bidders in each knockout auction. In the Logit specification the dependent variable is logged share of the ring members participating minus the logged share of those not participating. The omitted auction house is Christies. Estimated Minimum (Emin), Estimated Maximum (Emax) and Catalog Price are all divided by 1000

Table 4: Determinants of the ring winning target auctions

Specification	Point Estimates		Marginal Effects	
	Coeffs.	Std Err.	Coeffs.	Std Err.
Constant	-4.059	(0.414)	-0.988	(0.099)
# of Ring Members	0.854	(0.127)	0.208	(0.031)
(# of Ring Members)^2	-0.046	(0.017)	-0.011	(0.004)
(Emin + Emax)/2	-0.009	(0.072)	-0.002	(0.018)
(Emax - Emin)	0.291	(0.299)	0.071	(0.073)
Catalog Price	0.002	(0.010)	0.001	(0.002)
(Grade Min + Grade Max)/2	0.313	(0.069)	0.076	(0.017)
(Grade Max - Grade Min)	-0.004	(0.063)	-0.001	(0.015)
No Grade *	-0.045	(0.213)	-0.011	(0.051)
Exclusively US *	0.944	(0.268)	0.232	(0.063)
No Value *	-0.945	(1.064)	-0.200	(0.183)
House HRH *	1.094	(0.344)	0.264	(0.080)
House DK *	1.588	(0.411)	0.365	(0.075)
House IM *	1.097	(0.377)	0.267	(0.085)
House MB *	1.091	(0.394)	0.266	(0.090)
House RS *	0.653	(0.350)	0.161	(0.086)
House S *	0.929	(0.390)	0.228	(0.091)
House SA *	1.211	(0.421)	0.290	(0.090)
Pseudo R-squared	0.186			
Observations	1781			

Notes: The dependent variable is a dummy equal to 1 if the Ring won the target auction. Estimation is conducted using Logistic Regression. The omitted auction house is Christies. Estimated Minimum (Emin), Estimated Maximum (Emax) and Catalog Price are all divided by 1000. Marginal Effects for dummy variables (indicated by *) are $P(y=1|1) - P(y=1|0)$

Table 5: Knockout outcomes, by ring member

Ring Member	All auctions ($n \geq 1$)		Auctions with at least 2 ring members interested ($n \geq 2$)			
	% High KO Bid	# of Knockouts	% High KO Bid	% receive sidepayment	% pays sidepayments	# of Knockouts
A	40%	675	33%	22%	12%	607
B	57%	196	52%	21%	16%	175
C	34%	449	20%	23%	5%	368
D	14%	715	10%	20%	3%	686
E	39%	353	38%	24%	21%	348
F	31%	120	28%	28%	4%	116
G	11%	186	10%	34%	5%	184
H	14%	56	4%	34%	0%	50
I	44%	210	44%	17%	20%	209
J	45%	878	30%	22%	9%	686
K	42%	1075	28%	21%	9%	861

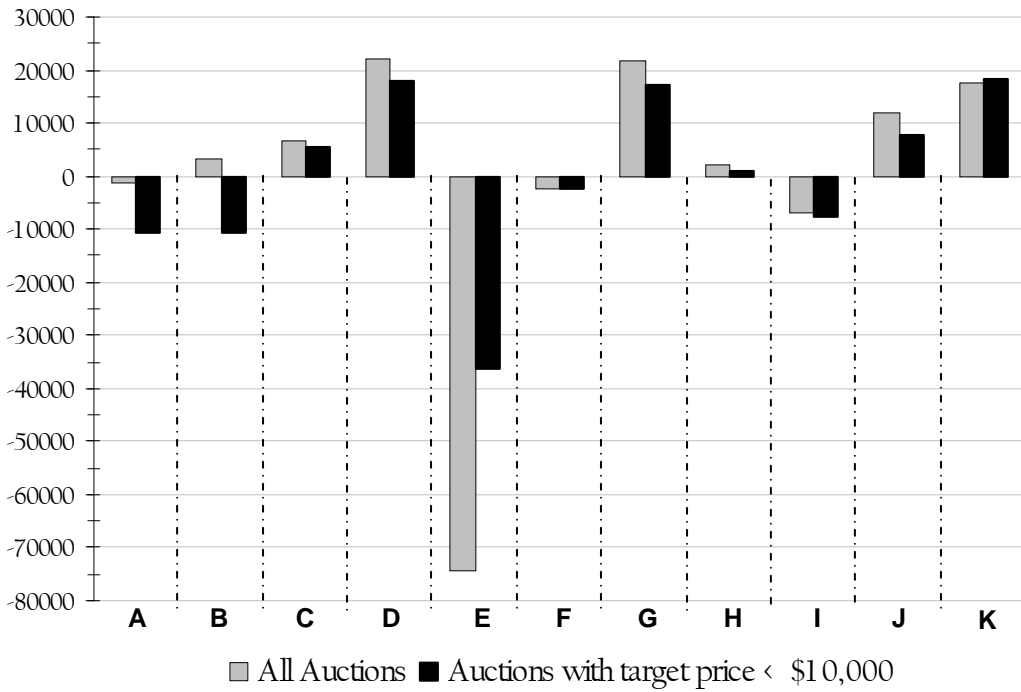


Figure 1: Net sidepayments from the ring, by member in dollars

Table 6: Naïve damages in target auctions with two or more ring members active in corresponding knockout

	By final price in target auction								Aggregate
	0-500	501-1000	1001-2000	2001-3000	3001-5000	5001-7000	7001-10000	10000+	
Mean target price (\$)	314	745	1,483	2,527	3,929	5,940	8,514	17,180	1,986
Mean highest knockout bid (\$)	471	1,066	1,996	3,187	5,918	8,041	10,428	23,840	2,718
Mean total sidepayments (\$)	42	92	154	245	622	697	526	1,910	222
Total naïve damages (\$)	28,390	53,460	68,000	51,950	113,150	65,500	38,950	95,500	514,900
Mean naïve damages (\$)	83	184	308	490	1,243	1,394	1,053	3,820	445
Number of lots won by ring	203	162	112	50	55	29	23	15	649
Total number of lots	341	290	221	106	91	47	37	25	1,158
	By number of ring members in knockout							Aggregate	
	2	3	4	5	6	7	8		
Mean target price (\$)	1,314	2,014	2,217	2,249	2,098	2,979	4,790	1,986	
Mean highest knockout bid (\$)	1,281	2,327	3,197	3,282	3,301	5,750	9,496	2,718	
Mean total sidepayments (\$)	12	96	249	211	365	895	1,898	222	
Total naïve damages (\$)	8,920	50,095	97,540	60,760	66,415	132,470	98,700	514,900	
Mean naïve damages (\$)	24	193	498	422	730	1,790	3,796	445	
Number of lots won by ring	133	126	136	98	67	64	25	649	
Total number of lots	367	260	196	144	91	74	26	1,158	
	By auction house								Aggregate
	Christies	HR Harmer	Daniel Kelleher	Ivy Mader	Matthew Bennet	Robert Siegel	Sotheby's	Spink America	
Mean target price (\$)	1,687	1,200	854	1,567	3,650	1,704	4,224	1,987	1,986
Mean highest knockout bid (\$)	1,658	1,689	1,405	1,942	5,917	2,069	4,248	2,703	2,718
Mean total sidepayments (\$)	67	110	108	46	681	169	103	225	222
Total naïve damages (\$)	6,730	88,360	13,225	7,065	283,200	88,165	14,660	13,495	514,900
Mean naïve damages (\$)	135	220	217	93	1,362	338	206	450	445
Number of lots won by ring	9	254	42	42	144	109	33	16	649
Total number of lots	50	401	61	76	208	261	71	30	1,158




Figure 2: Estimated bid function for a strong bidder in the model with unobserved auction level heterogeneity. The CDF of the highest valuation among non-ring bidders is also shown. (Constructed using an auction-specific common element of 600 and a 70% chance of facing another strong bidder in the knockout.)

Table 7: Damages to the seller

Model:	Assumption	With unobserved auction heterogeneity			No unobserved auction heterogeneity		
		Point estimate	90% Confidence interval:		Point estimate	90% Confidence interval:	
			Lower bound	Upper bound		Lower bound	Upper bound
Mean naïve damages (\$)		67.25	48.49	143.49	149.53	92.93	196.20
Mean damages (\$)	U. B.	35.16	21.00	77.58	105.02	52.07	141.70
	L. B.	27.91	11.17	69.37	98.13	42.22	135.77
Mean damage ratio	U. B.	0.96	0.92	0.98	0.88	0.84	0.93
	L. B.	0.97	0.93	0.99	0.89	0.85	0.95
Proportion of auctions with $Pr > Pc$	U. B.	0.00	0.00	0.00	0.00	0.00	0.00
	L. B.	0.18	0.05	0.23	0.097	0.040	0.17
Mean damage ratio ($Pr > Pc$)	L. B.	1.05	1.02	1.14	1.10	1.03	1.41
Proportion of auctions with $Pr < Pc$	U. B.	0.23	0.16	0.38	0.34	0.23	0.44
	L. B.	0.23	0.16	0.38	0.34	0.23	0.44
Mean damage ratio ($Pr < Pc$)	U. B.	0.83	0.75	0.89	0.64	0.57	0.74
	L. B.	0.83	0.75	0.89	0.64	0.57	0.74
Proportion of auctions with $Pr = Pc$	U. B.	0.77	0.62	0.84	0.66	0.56	0.77
	L. B.	0.59	0.46	0.73	0.57	0.46	0.68
Proportion of target auctions won		0.35	0.09	0.46	0.37	0.18	0.45
Simulated auctions		100000			100000		

Notes: Damage ratio is the ratio of the price received with the ring to the price received with competitive bidding. All means are over target auctions that the ring won (unless further conditioned as noted). L. B. = Lower Bound, U. B. = Upper Bound. Pr refers to the price sellers receive with the ring, Pc is the price with competitive bidding. Confidence intervals are bootstrapped with 5,000 iterations.

Table 8: Damages to the non-ring bidders

Model:	With unobserved auction heterogeneity			No unobserved auction heterogeneity		
	Point estimate	90% Confidence interval:		Point estimate	90% Confidence interval:	
		Lower bound	Upper bound		Lower bound	Upper bound
Damages due to misallocation:						
Proportion of target auctions ring won	0.35	0.09	0.46	0.37	0.18	0.45
Proportion of target auctions ring won with damages	0.18	0.05	0.23	0.097	0.040	0.17
Mean damages (conditional on ring winning target auction, \$)	7.24	1.59	16.03	6.89	1.65	18.46
Damages due to price inflation:						
Mean damages (conditional on ring not winning target auction, \$)	98.04	64.08	138.14	112.36	90.18	133.18
# Simulated auctions	10000			100000		

Notes: All estimates obtained using the lower bound assumption. Confidence intervals are bootstrapped with 5,000 iterations.

Table 9: Impact on market efficiency

Model:	With unobserved auction heterogeneity			No unobserved auction heterogeneity		
	Point estimate	90% Confidence interval:		Point estimate	90% Confidence interval:	
		Lower bound	Upper bound		Lower bound	Upper bound
Mean efficiency loss (\$)	7.24	1.58	16.44	6.89	1.65	18.46
Mean proportional efficiency losses:						
Ring active	0.003	0.0002	0.006	0.003	0.0005	0.007
No ring bidders	0.08	0.02	0.12	0.14	0.07	0.18
Only ring bidders	0.27	0.21	0.41	0.38	0.29	0.55
Proportion of target auctions won	0.35	0.09	0.46	0.37	0.18	0.45
# Simulated auctions	100000			100000		

Notes: Means are conditional on the ring winning. The mean proportional efficiency losses are averages over all auctions, not just those won by the ring. Confidence intervals are bootstrapped with 5,000 iterations.

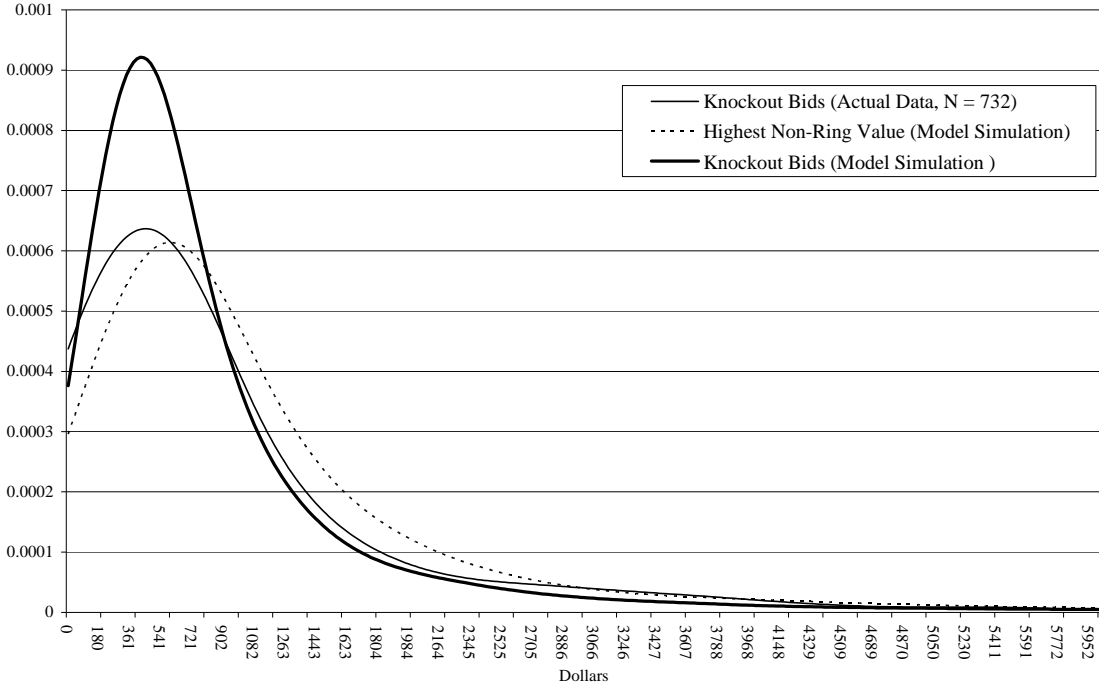
Table 10: Returns to the ring

Model:	With unobserved auction heterogeneity			No unobserved auction heterogeneity		
	Point estimate	90% Confidence interval:		Point estimate	90% Confidence interval:	
		Lower bound	Upper bound		Lower bound	Upper bound
Mean naïve return (equiv. damages, \$)	67.25	48.49	143.28	149.53	92.93	196.20
Proportion of ring wins that harmed ring	0.18	0.05	0.23	0.097	0.040	0.17
Mean return to ring (harm, \$)	-7.24	-16.03	-1.59	-6.89	-18.53	-1.65
Mean return to ring (benefit, \$)	35.16	21.00	77.56	105.02	52.07	141.70
Mean return to ring (net, \$)	27.91	11.17	70.30	98.13	42.22	135.77
Mean proportional price discount	0.96	0.92	0.98	0.88	0.84	0.93
# Simulated auctions	100000			100000		

Notes: All means are over target auctions that the ring won. Confidence intervals are bootstrapped with 5,000 iterations.

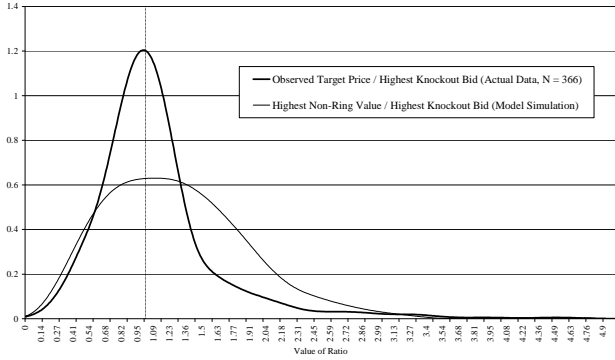
Panel A: Density of Ring Bids and Non-Ring Values

Estimates from a Triweight Kernel Density estimate with a 'rule-of-thumb' bandwidth as described in text.



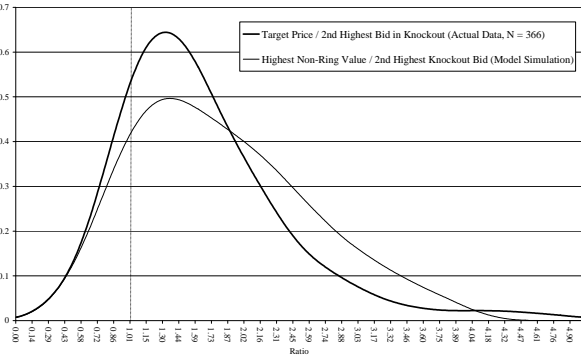
Panel B: Density of Ratio of Highest Non-Ring Value to Winning Knockout Bid

Estimates from a Triweight Kernel Density estimate with a 'rule-of-thumb' bandwidth as described in text.



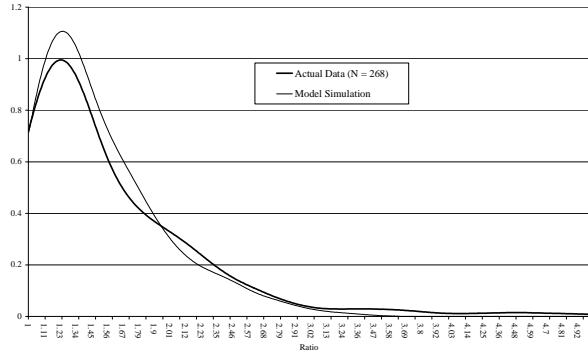
Panel C: Density of Ratio of Highest Non-Ring Value to 2nd Highest Knockout Bid

Estimates from a Triweight Kernel Density estimate with a 'rule-of-thumb' bandwidth as described in text.



Panel D: Density of Ratio of Highest Strong Knockout Bids to 2nd Highest (Strong) Bids

Estimates from a Triweight Kernel Density estimate with a 'rule-of-thumb' bandwidth as described in text.



Panel E: Density of Ratio of Weak Knockout Bids to Strong Knockout Bids

Estimates from a Triweight Kernel Density estimate with a 'rule-of-thumb' bandwidth as described in text.

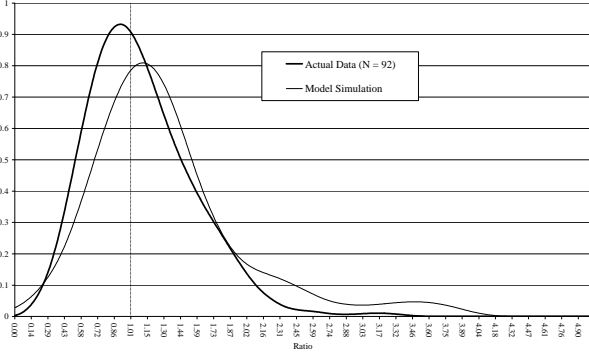


Figure 3: Indications of model fit: Model with unobserved auction heterogeneity