A Bidding Ring Protocol for First-Price Auctions

Kevin Leyton-Brown* Moshe Tennenholtz†
Navin A.R. Bhat‡ Yoav Shoham§

Abstract

We identify a self-enforcing collusion protocol (a "bidding ring") for non-repeated first-price auctions. Unlike previous work on the topic such as that by McAfee and McMillan [1992] and Marshall and Marx [2007], we allow for the existence of multiple cartels in the auction and do not assume that non-colluding agents have perfect knowledge about the number of colluding agents whose bids are suppressed by the bidding ring. We show that it is an equilibrium for agents to choose to join bidding rings when invited and to truthfully declare their valuations to a ring center, and for non-colluding agents to bid straightforwardly. Furthermore, even though our protocol is efficient, we show that the existence of bidding rings benefits ring centers and all agents, both members and non-members of bidding rings, at the auctioneer's expense.

1 Introduction

We consider the question of how agents can gain by coordinating their bidding in non-repeated single-good auctions, even when all agents still act selfishly. The case of second-price auctions is well-studied; we concentrate on the comparatively less-studied case of first-price auctions. Overall, collusion is a widespread phenomenon. Many papers in the literature offer real-world examples of cartels that have been identified and prosecuted, e.g., under the US's Sherman Act. Reduction of revenue due to bidder collusion is therefore a significant practical threat to auctioneers. Understanding the topic theoretically can help auctioneers to choose an auction type and to modify the rules of their auctions in order to make collusion more difficult. Collusion has been observed to occur in both repeated and single-auction settings; the latter is the focus of our work.

*kevinlb@cs.ubc.ca; Computer Science Department, University of British Columbia
†moshet@ie.technion.ac.il; Faculty of Industrial Engineering and Management, Technion
‡nbhat@cs.ubc.ca; Computer Science Department, University of British Columbia
§shoham@cs.stanford.edu; Computer Science Department, Stanford University

1.1 Collusion in Second-Price Auctions

Graham and Marshall [1987] wrote one of the first formal papers on collusion, considering second-price auctions. This paper described a *knockout procedure*: agents announce their bids in a knockout auction; only the highest bidder goes to the auction, but this bidder must pay a ring center the amount of his gain relative to the case where there was no collusion. The ring center pays each agent in advance; the amount of this payment is calculated so that on expectation the ring center will budget-balance *ex ante*, before knowing the agents' valuations.

Graham and Marshall's work has been extended, still considering the case of second-price auctions, to deal with variations in the knockout procedure, differential payments, and relations to the Shapley value [Graham *et al.*, 1990]. The case where only some of the agents are part of the cartel is discussed by Mailath and Zemsky [1991], who also derived a mechanism in which ring centers achieve *ex post* budget balance. Work by von Ungern-Sternberg [1988] considers collusion in second-price auctions where the designated winner of a cartel is not the agent with the highest valuation. Although we are not aware of any work that presents this result, it is also easy to extend the protocol to an environment containing both multiple cartels and independent bidders.

Less formal discussion of collusion in auctions can be found in a wider variety of papers. For example, a survey paper that discusses mechanisms that are likely to facilitate collusion in auctions, as well as methods for the detection of such schemes, can be found in [Hendricks & Porter, 1989]. A discussion and comparison of the stability of rings associated with classical auctions can be found in [Robinson, 1985], concentrating on the case where the valuations of agents in the cartel are honestly reported. Collusion is also discussed in other settings, *e.g.*, in the context of general Bertrand or Cournot competition [Cramton & Palfrey, 1990].

1.2 Collusion in First-Price Auctions

An influential paper by McAfee and McMillan [1992] presented the first theoretical results on collusion in first-price auctions. This work assumes that a fixed number of agents participate in the auction and that all agents are part of a single cartel that coordinates its behavior in the auction. The authors show optimal collusion protocols for "weak" cartels (in which transfers between agents are not permitted: all bidders bid the reserve price, using the auctioneer's tie-breaking rule to randomly select a winner) and for "strong" cartels (the cartel holds a knockout auction, the winner of which bids the reserve price in the main auction while all other bidders sit out; the winner distributes some of his gains to other cartel members through side payments). Though it was not the focus of their work, McAfee and McMillan also considered the case where in addition to a single cartel there are also additional agents. However, results are shown only for two cases: (1) where non-cartel members bid without taking the existence of a cartel into account (*i.e.*, either they are irrational or they hold the false belief that no cartels exist) and (2) where each agent i has valuation

$v_i \in \{0, 1\}$. McAfee and McMillan explain that they do not attempt to deal with general strategic behavior in the case where the cartel consists of only a subset of the agents; furthermore, they do not consider the case where multiple cartels can operate in the same auction.

The only other work of which we are aware that proposes a self-enforcing collusion protocol for single (i.e., unrepeated) first-price auctions is Marshall and Marx [2007]. This paper considers both first- and second-price auctions, addresses both the repeated and unrepeated cases, and proposes both "bid submission mechanisms" (ring protocols in which the ring center is able to submit bids on the bidders' behalf) and "bid coordination mechanisms" (ring protocols in which the ring center is only able to suggest bid amounts to agents). They assume that there is only one bidding ring involving some subset of the agents, and that the existence and formation or non-formation of the cartel is common knowledge among all agents participating in the auction.

Finally, a number of other papers, many of them written quite recently, have studied collusion in repeated first-price auctions or have presented results that bear directly on this setting [Aoyagi, 2003; Hörner & Jamison, 2007; Skrzypacz & Hopenhayn, 2004; Blume & Heidhues, 2008; Fudenberg *et al.*, 1994; Feinstein *et al.*, 1985]. Overall, these mechanisms tend to work by using folk-theorem-like constructions, incenting some bidders to rotate their participation in the auction through the threat of future punishment or the promise of future opportunities to collude.

1.3 Novelty of Our Work

Our paper[1] differs from related work on collusion in unrepeated first-price auctions by relaxing several assumptions. Comparing to [McAfee & McMillan, 1992], we allow for the possibility that some bidders will not belong to a cartel without assuming that agents are irrational or hold false beliefs in equilibrium, and we allow that more than one cartel may exist, introducing the new wrinkle that cartel members must reason about the possibility of other cartels. We model bidders' valuations as real numbers drawn from an interval according to an arbitrary distribution (as compared, *e.g.*, to the case studied in [McAfee & McMillan, 1992] where valuations take one of only two discrete values), and the decision of whether or not to join a bidding ring is part of an agent's choice of strategy.

The recent paper by Marshall and Marx [2007] is closer to our work, but still differs significantly. In their terminology our work proposes a bid submission mechanism for single first-price auctions; thus, we contrast our work with their mechanism for the same setting. Marshall and Marx [2007] extend McAfee and McMillan [1992] in some of the same ways that we do: they identify a Bayes-Nash equilibrium of a bidding ring mechanism that is not required to involve all

[1]Some of our previously published work is related to this paper. In [Leyton-Brown *et al.*, 2000] we considered bidding rings under the assumptions that only a single bidding ring exists, and that bidders who were not invited to join the ring are not aware that bidding rings could exist. [Leyton-Brown *et al.*, 2002] was a poster presentation of work that grew into this paper; we also posted a working version of the full paper on `arXiv.org` in 2002 (`arXiv:cs/0201017v1`).

bidders; they consider bidders with valuations drawn from real intervals; they include bidders' decisions about whether to join a bidding ring as part of the bidders' strategy spaces. However, in several senses their results are more limited than ours. First and most importantly, they assume that non-cartel bidders have complete information about the existence of a cartel in the auction and about the number of bids suppressed by this cartel. We find this to be a strong assumption, since the bidding ring is a clandestine arrangement that hides its existence from the seller. Second, they do not allow for the existence of multiple cartels. (It is apparently important to allow for this possibility: in an empirical analysis of forest service auctions, Marshall and Marx themselves present evidence that two cartels participated in the same first-price auction.) Third, they do not allow agents to take their valuations into account when deciding whether or not to join the cartel (i.e., they require that this choice be made *ex ante*, not *ex interim*). Fourth, they are not always able to guarantee that a decision to join the bidding ring satisfies even *ex ante* individual rationality, though they do show this for some simple valuation distributions. On the other hand, in some senses Marshall and Marx [2007] present stronger results than we do. First, their protocol still works even when the auctioneer does not announce the identity of the winner; ours needs this information in order to require the winner to make a payment to the ring center. Second and most notably, they allow for bidders with asymmetric valuation distributions, while we consider only the symmetric case. Interestingly, they write "We focus on the heterogeneous IPV model, which is important for the study of collusion because, even if bidders are homogeneous, collusion creates heterogeneity among them" [Marshall and Marx 2007, page 377]. One of our main technical results is that a collusive protocol can be constructed to cancel out this heterogeneity, allowing all agents to bid symmetrically in the auction.

In what follows, we begin by defining an auction model and establishing some notation in Section 2. In Section 3 we give our bidder model. Some prominent features are that the number of bidders is stochastic while the number of (unsuppressed) bids placed in the auction is perfectly observed; there can be multiple cartels, and all agents invited to participate in one are (privately) able to observe the number of other agents who receive invitations. We then identify a bidding ring protocol and show that responding truthfully to this protocol constitutes an equilibrium. We show, via an argument related to revenue equivalence, that the collusion benefits both ring and non-ring members as well as ring centers, at the auctioneer's expense. Finally we show (as Marshall and Marx [2007] do in their own setting) that our bidding ring protocol can be disrupted by allowing agents to place undetected shill bids.

2 Modeling First-Price Auctions

An economic environment E consists of a finite set of agents who have non-negative valuations for a good at auction, and a distinguished agent 0, the seller or center. First we define a "classical" economic environment, which we denote

E_c. Let $\mathcal{T}$ be the set of possible agent types. The type $\tau_i \in \mathcal{T}$ of agent i is the pair $(v_i, s_i) \in V \times \mathcal{S}$. v_i denotes an agent's valuation, which we assume represents a purely private valuation for the good. v_i is selected independently from the other v_j's of other agents from a known cumulative distribution, F, a continuously differentiable distribution with support on the interval $[0, 1]$ having non-cumulative distribution (density function) f. Throughout the paper we will use upper- and lower-case symbols to respectively denote such cumulative and non-cumulative distributions. By $s_i \in \mathcal{S}$ we denote agent i's private signal about the number of agents in the auction, and let $\varnothing$ denote a null signal. We will vary the set of possible signals $\mathcal{S}$ throughout the paper; in E_c let $\mathcal{S} = \{\varnothing\}$, meaning that agents receive no information about the number of other agents in the auction as part of their types. In our analysis we always assume that the economic environment is common knowledge; *e.g.*, in E_c an agent's strategy can depend on the number of agents in the auction even though all agents receive the null signal.

We assume that the utility function of agent i is linear (the agent is risk-neutral), free of externalities, and normalized (the agent achieves zero utility for not winning the good and paying nothing). When asked to pay t, let the utility of agent i (having valuation v_i) be $v_i - t$ if i is allocated the good and $-t$ otherwise. $b_i : \mathcal{T} \to \mathbb{R}^+ \cup \{\overline{P}\}$ denotes agent i's strategy, a mapping from i's type τ_i to his declaration in the auction. The declaration $\overline{P}$ indicates that i will not participate in the auction.

2.1 Classical First-Price Auctions

We argue that the choice of information structure is very important for the study of collusion in first-price auctions. The most familiar case gives rise to what we call the "classical" first-price auction, where the number of participants[2] is part of the economic environment (this is what we have called E_c). Using standard equilibrium analysis (*e.g.*, following Riley and Samuelson [1981]) the unique symmetric equilibrium can be identified.

Proposition 2.1 *If valuations are selected from a continuous distribution F having finite support, then the unique symmetric equilibrium is for each agent i to bid the amount*

$$v_i - F(v_i)^{-(n-1)} \int_0^{v_i} F(u)^{n-1} du.$$

Observe that this bidding strategy is parameterized by valuation, but also depends on information from the economic environment. It is notationally useful for us to be able to specify the amount of the equilibrium bid as a function of both v and n,

$$b^e(v_i, n) = v_i - F(v_i)^{-(n-1)} \int_0^{v_i} F(u)^{n-1} du. \tag{1}$$

[2]When we say that n agents participate in the auction we do not count the distinguished agent 0, who is always present.

We are interested in constructing a collusive agreement that has bidders with low valuations drop out of the main auction, allowing stronger bidders to reduce their bids. (For example, this is the flavor of McAfee and McMillan's protocol in which cartel non-members are assumed not to behave strategically.) However, such collusion is nonsensical in the classical first-price auction environment. When bidders' strategies already depend on the number of agents in the economic environment, the amounts these bidders bid cannot change if cartel members with low valuations fail to submit bids. This is a problem with our auction model rather than with collusion in first-price auctions *per se*—in practice bidders might *not* know the exact number of agents in the economic environment, and thus adopt a strategy that depends on the number of bidders who choose to participate in the auction.

2.2 First-Price Auctions with a Stochastic Number of Bidders

One way of modelling agents' uncertainty about the number of opponents they face is to say that the number of participants is drawn from a probability distribution; while the actual number of participants is not observed, the *distribution* is commonly known. First-price auctions of this kind were introduced by McAfee and McMillan [1987] (in an earlier paper that makes no mention of collusion).

This setting requires that we define a new economic environment. Let us denote it E_s. Let the definition of agents in E_s be the same as in E_c (in particular, again let all agents receive the null signal $\varnothing$). Let $\mathcal{D}_\ell$ be the set of all probability distributions $d : \mathbb{Z} \to \mathbb{R}$ having support on any subset of the integers greater than or equal to ℓ. Denote the distribution over the number of agents in the auction as $p \in \mathcal{D}_2$. After nature determines the number of agents by drawing from p, let the name of each agent be selected from the uniform distribution on $[0, 1]$.[3] The unique symmetric equilibrium was identified by Harstad *et al.* [1990].

Proposition 2.2 *If valuations are selected from a continuous (cumulative) distribution F having finite support, and the number of bidders is selected from the (non-cumulative) distribution p, then it is a unique symmetric equilibrium for each agent i to bid the amount*

$$b(v_i) = \sum_{j=2}^{\infty} \frac{F^{j-1}(v_i)p(j)}{\sum_{k=2}^{\infty} F^{k-1}(v_i)p(k)}\ b^e(v_i, j).$$

Observe that $b^e(v_i, j)$ is the amount of the equilibrium bid for a bidder with valuation v_i in a classical first-price auction setting with j bidders as described in Equation (1) above; p is deduced from the economic environment. We again

[3] This technical assumption prevents any agent from gaining information about the number of agents in the auction from his own identity and the fact of his selection. We make similar assumptions in Section 3.

introduce notation for the equilibrium bid, this time as a function of the agent's valuation and the probability distribution p,

$$b^e(v_i, p) = \sum_{j=2}^{\infty} \frac{F^{j-1}(v_i)p(j)}{\sum_{k=2}^{\infty} F^{k-1}(v_i)p(k)} \, b^e(v_i, j). \tag{2}$$

Although we will need these results in what follows, this auction model is still insufficient for modeling collusion in a first-price auction. If each agent knows only the distribution of agents interested in participating in the auction, he has no way of being affected by agents who drop out! Again, this is a deficiency of our model—in some settings agents may know how many agents bid in the auction, even though they may not know the number of agents who chose not to bid. For example, when an auction takes place in an auction hall, no bidder can be sure about how many potential bidders stayed home, but every bidder can count the number of people in the room before placing his or her bid.

2.3 First-Price Auctions with Participation Revelation

We model this auction hall scenario as a *first-price auction with participation revelation*, which we define as follows:

1. Agents indicate their intention to bid in the auction.
2. The auctioneer announces n, the number of agents who registered in the first phase.
3. Agents submit bids to the auctioneer. The auctioneer will only accept bids from agents who registered in the first phase.
4. The agent who submitted the highest bid is awarded the good for the amount of his bid; all other agents are made to pay 0.

When a first-price auction with participation revelation operates in E_s, the equilibrium of the corresponding *classical* first-price auction holds.

Proposition 2.3 *In E_s it is an equilibrium of the first-price auction with participation revelation for every agent i to indicate the intention to participate, and to bid according to $b^e(v_i, n)$.*

Proof. Agents always gain by participating in first-price auctions when there is no participation fee. The only way to participate in this auction is to indicate the intention to participate in the first phase. Thus the number of agents announced by the auctioneer is equal to the total number of agents in the economic environment. From Proposition 2.1 it is best for agent i to bid $b^e(v_i, n)$ when it is common knowledge that the number of agents in the economic environment is n. ■

First-price auctions with participation revelation may often be a more realistic model than classical first-price auctions, since the former allows that bidders may not know *a priori* the number of opponents they will face. When bidders are unable to collude, there is no strategic difference between these two mechanisms, justifying the common use of the simpler classical model. For the study of bidding rings, however, the difference between the mechanisms is profound—we are now able to look for a collusive equilibrium in which bidder strategies depend only on the number of other agents who "show up" for the auction. More specifically, unlike the other first-price auction models we considered, first-price auctions with participation revelation have the property that non-cartel-members' bids will change when one or more cartel members choose not to participate in the auction, even if the non-cartel bidders are rational and have true beliefs about the economic environment.

3 Bidding Rings for First-Price Auctions

We define the economic environment E_{br} as an extension of E_s, consisting of the distinguished agent 0 who offers a good for sale, a randomly-chosen set of ring centers who do not value the good, and a randomly-chosen set of agents each of whom receives an invitation from exactly one ring center. Recall that the type $\tau_i \in \mathcal{T}$ of agent i is the pair $(v_i, s_i) \in V \times \mathcal{S}$. Define v_i as in E_c. Let $\mathcal{S} \subseteq \mathbb{N} \setminus \{0\}$; $s_i \in \mathcal{S}$ represent the number of agents in i's bidding ring, which is available to i as private information. Thus, when he observes his signal s_i, an agent i learns that there are at least s_i agents in the auction, all of whom share his ring center. We model singleton bidders as bidding rings with only one invited agent; in this case we consider the ring automatically disbanded and ignore the ring center.

Again p denotes the distribution over the total number of agents.[4] Let p_{s_i} denote the distribution over the total number of agents conditional on i's signal s_i, and $p^{\sim}_{s_i}$ denote the distribution over the number of agents *not* in i's bidding ring conditional on s_i. Our results in the following sections depend on the *independent cartel* property, which states that the distributions over the numbers of agents in each cartel are independent.

Definition 3.1 (independent cartel property) *An economic environment satisfies the* independent cartel property *if* $\forall s, s' \in \mathcal{S}$, $p^{\sim}_{s} = p^{\sim}_{s'}$.

Here we provide one technical construction that achieves this property; our results also hold for other constructions. Let $\gamma_C(n_c) \in \mathcal{D}_2$ denote the probability that an auction will involve n_c ring centers. After a value is realized from $\gamma_C(n_c)$, the name of each ring center is selected from the uniform distribution on $[0, 1]$. Let $\gamma_A(n) \in \mathcal{D}_1$ denote the probability that n agents will be associated with a bidding ring. After the number of agents is determined, the name of each agent associated with a potential ring center is selected from the uniform distribution on $[0, 1]$.

[4]Observe that p is concerned with the number of agents that *exist*, not the number of agents who elect to participate.

In order to write an expression for p, we must define some notation. Let $x, y \in \mathcal{D}_0$ be independent random variables, and consider the distribution of their sum p'. Since x and y are independent, the probability of their sum being m is just the sum of the product of the individual probabilities of values of x and y that sum to m,

$$p'(m) = \sum_{j=0}^{\infty} p(m-j)\, q(j)\,. \tag{3}$$

Summing independent distributions in this way corresponds to convolution, which we denote symbolically as $p' = p * q$. Observe that convolution is associative and commutative. Denote repeated convolution of distribution d as

$$\bigotimes_n d \equiv \overbrace{d * d * d * \ldots * d}^{d \text{ repeated } n \text{ times}}\,. \tag{4}$$

We define the Kronecker delta (an indicator function) as

$$\delta_m(j) \equiv \begin{cases} 1 \text{ if } j = m; \\ 0 \text{ otherwise.} \end{cases} \tag{5}$$

Now the following identity can be inferred from Equation (3):

$$\bigotimes_j \delta_k = \delta_{(j \cdot k)}. \tag{6}$$

We can now write

$$p = \sum_{n=2}^{\infty} \gamma_C(n) \left(\bigotimes_n \gamma_A \right), \tag{7}$$

and for i's posterior beliefs about p, conditioned on his signal s_i,

$$p_{s_i} = \sum_{n=2}^{\infty} \gamma_C(n) \left(\bigotimes_{n-1} \gamma_A \right) * \delta_{s_i}. \tag{8}$$

We denote by p_{n,s_i} the distribution over the number of agents, conditioned on i's signal s_i and the additional information that there are a total of n bidding rings (and/or singleton bidders):

$$p_{n,s_i} = \left(\bigotimes_{n-1} \gamma_A \right) * \delta_{s_i}. \tag{9}$$

3.1 Symmetrizing Auctions

If we were simply to run a standard first-price auction in E_{br} (without participation revelation and without imposing a bidding ring protocol) agents would follow asymmetric strategies based on their different signals. In this section we describe a family of auction mechanisms that we dub *symmetrizing auctions*, which impose asymmetric payment rules on agents with different signals in order to give rise to symmetric equilibria.[5] Assume for this section that the auctioneer knows each agent's signal. Denote a bid from agent i as $\mu_i \in \mathbb{R}^+ \cup \{\overline{P}\}$, the tuple of bids from all agents as $\pi \in \Pi$ and an auction's transfer function for agent i (determining i's payment) as $t_i : \mathbb{R}^+ \cup \{\overline{P}\} \times \Pi \to \mathbb{R}$.

We say that an auction is *aligned* with signal s if, in an economic environment where the number of agents is drawn from p_s and all agents receive the null signal, the auction is efficient and incentive-compatible.

Definition 3.2 (auction aligned with a signal) *An auction M_s is* aligned *with signal $s \in S$ if M_s allocates the good to an agent i with $\mu_i \in \max_j \mu_j$, and M_s is a symmetric truth-revealing direct mechanism for a stochastic number of agents drawn from p_s, each of whom receives the signal $\varnothing$.*

Now we identify a class of asymmetric auctions (in which agents can receive different signals and are subject to potentially different transfer functions) that nevertheless have symmetric truthful equilibria. Intuitively, asymmetry is introduced into the transfer functions in a way that exactly balances the informational asymmetry among the agents. We call these auctions *symmetrizing*.

Definition 3.3 (symmetrizing auction) *$\bar{M}$ is a* symmetrizing auction *if it allocates the good to an agent i with $\mu_i \in \max_j \mu_j$, and if each agent i is made to transfer $t_{s_i}(\mu_i, \pi)$ to the center, with t_{s_i} taken from an auction M_{s_i} that is aligned with signal s_i.*

Using an argument related to the revelation principle, we can prove the following result.

Lemma 3.4 *Truth-revelation is an equilibrium of symmetrizing auctions.*

Proof. The payoff of agent i is uniquely determined by the allocation rule, the transfer function t_{s_i}, the distribution over the number of agents in the auction, and all agents' strategies. Assume that the other agents are truth revealing, then each other agent's behavior, the allocation rule, and agent i's payment rule are all identical in $\bar{M}$ and M_{s_i}. Conditioned on his private information s_i, agent i's posterior is that p_{s_i} is the distribution over the number of agents in the auction. Since truth-revelation is an equilibrium in M_{s_i} when the distribution of agents is p_{s_i}, truth-revelation is agent i's best response in $\bar{M}$. ■

[5] We note that symmetrizing auctions may have applications beyond the study of collusion. Furthermore, they are more general than presented here. For example, our results hold for arbitrary constructions of p and families of signals, for inefficient allocation rules, and for settings in which all agents have the same signal but agents' payment rules are taken from different truthful mechanisms.

3.2 Relating p to b^e

In classical first-price auctions, the amounts of agents' equilibrium bids increase with the number of participating agents. It is intuitive to expect that the same thing would occur in first-price auctions with a stochastic number of participants. As we show at length in Appendix A, however, simply knowing that distribution p has a smaller expected number of participants than distribution p'—or even that p stochastically dominates p'—is not enough to determine which distribution will give rise to a lower symmetric equilibrium bid for a given valuation. Nevertheless, we are able to identify a class of pairs of distributions (p, p') for which it does hold that $b^e(v_i, p)$ is always less than $b^e(v_i, p')$: those (p, p') for which p' *convolutively dominates* p.

Definition 3.5 (convolutive dominance) *Given two (non-cumulative) distributions p, p', the function p' exhibits* convolutive dominance *of p if there exists a probability density function $q \in \mathcal{D}_0$ with $q(0) \neq 1$ such that $p' = p * q$.*

Lemma 3.6 *$\forall p, p' \in \mathcal{D}_2$, if p' convolutively dominates p then $b^e(v_i, p) < b^e(v_i, p')$.*

Proof. The proof is given in Appendix A. ∎

3.3 First-Price Auction Bidding Ring Protocol

Any number of ring centers may participate in an auction. However, we assume that there is only a single collusion protocol, and that this protocol is common knowledge. What follows is the protocol of a ring center who approaches k agents and who operates in conjunction with a first-price auction with participation revelation in the economic environment E_{br}.

1. Each agent i sends a message μ_i to the ring center.
2. If all k agents accept the invitation then the ring center drops all bidders except the bidder with the highest reported valuation, whom we denote as bidder h. For this bidder the ring center indicates the intention to bid in the main auction, and places a bid of $b^e(\mu_h, p_{n,1})$.
3. Otherwise, the ring center indicates an intention to bid in the main auction on behalf of every agent who accepted the invitation to the bidding ring. For each such bidder i, the ring center submits a bid of $b^e(\mu_i, p_{n-k+1,k})$, where in this case n (the number of bidders announced by the auctioneer) will include all agents invited to the bidding ring.
4. The ring center pays each member a pre-determined payment $c_{n,k} \geq 0$ whenever all bidders participate in the ring, which is independent of the outcome of the auction and the amount each bidder bid, but which can depend on n and k.
5. If bidder h wins in the main auction, he is made to pay $b^e(\mu_h, p_{n,1})$ to the center and $b^e(\mu_h, p_{n,k}) - b^e(\mu_h, p_{n,1})$ to the ring center.

Observe that if an agent declines an invitation to participate in a bidding ring, the main auction will be asymmetric in the sense that different singleton bidders will have different information about the total number of bidders in the auction. Nevertheless, we can prove the following theorem, which is our main result.

Theorem 3.7 *It is a Bayes-Nash equilibrium for all bidding ring members to choose to participate and to truthfully declare their valuations to their respective ring centers, and for all non-bidding ring members to participate in the main auction with bids of* $b^e(v, p_{n,1})$.

Proof. We begin by partitioning the space of all possible strategies and introducing notation to describe these partitions. Since agents who are invited to join a bidding ring have a richer set of strategic choices available to them than do singleton bidders, we partition each strategy space separately. We give each set a short name which we use throughout this proof, built up of the following six symbols: $P/\overline{P}$ (participate/do not participate) and $T/\overline{T}$ (bid truthfully/do not bid truthfully), given $R/\overline{R}$ (a ring bidder/a non-ring bidder).

The space of bidding ring agent strategies is partitioned as follows:

- $(\overline{P}|R)$: the agent either chooses not to participate in the auction at all, or declines participation in the bidding ring and then bids independently in the main auction.
- $(P\overline{T}|R)$: the agent participates in the auction, accepts the invitation to join the bidding ring, and then lies to the ring center about his valuation.
- $(PT|R)$: the agent participates in the auction, accepts the invitation to join the bidding ring, and declares his true valuation to the ring center.

For the non-ring bidder, the space of strategies is partitioned as follows:

- $(\overline{P}|\overline{R})$: the agent chooses not to participate in the auction at all.
- $(P\overline{T}|\overline{R})$: the agent participates in the auction, but does not bid $b^e(v, p_{n,1})$ in the main auction.
- $(PT|\overline{R})$: the agent participates in the auction with a bid of $b^e(v, p_{n,1})$.

Given two strategy sets X and Y, and given that all agents other than agent i follow the strategy $(PT|R)$ or $(PT|\overline{R})$ (as appropriate) we denote the proposition that agent i's expected utility for following some strategy $x \in X$ is greater than his expected utility for following any strategy $y \in Y$ as $u(X) > u(Y)$.

This proof consists of two main parts, the first dealing with participation and the second dealing with bidding. In Part (1a) we show that $u(PT|\overline{R}) > u(\overline{P}|\overline{R})$, and in Part (1b) we consider the more complex case of ring bidders and show that $u(PT|R) > u(\overline{P}|R)$. In Part (2) we show simultaneously that $u(PT|\overline{R}) > u(P\overline{T}|\overline{R})$ and that $u(PT|R) > u(P\overline{T}|R)$.

Part 1a: $u(PT|\overline{R}) > u(\overline{P}|\overline{R})$**.**

Recall that we assume that all other bidders bid according to $(PT|\overline{R})$ or $(PT|R)$. If non-ring bidder i also bids according to $(PT|\overline{R})$ then all bidders follow a symmetric strategy in the main auction. Thus i has a non-zero probability of winning the good and gaining a surplus. As there is no participation fee, it is strictly better for i participate in the auction than to decline participation altogether.

Part 1b: $u(PT|R) > u(\overline{P}|R)$**.**

By the argument in Part (1a) a ring bidder i should likewise opt to participate in the auction; however, we must still consider whether i is best off accepting or rejecting his bidding ring invitation. As discussed above, if i rejects the invitation then the main auction will be asymmetric; thus, this part requires a nontrivial argument. We consider the case where $c_{n,k} = 0$, as this is the case where i has the least incentive to accept the invitation. In this discussion let n represent the *true* number of bidding rings and singleton bidders in the economic environment (*i.e.*, the value realized from the distribution γ_c).

First, consider a different setting, which we denote $(\star)$: a first-price auction with a stochastic number of participants in economic environment E_s, with the number of agents distributed according to p_{n,s_i}. In $(\star)$ all bidders have the same information as i and are subject to the same payment rule. Thus, from Proposition 2.2 it is a best response for i to bid $b^e(v_i, p_{n,s_i})$. Bidder i's expected utility is the same in $(\star)$ and when following the strategy $(PT|R)$ in the real auction, because both auctions allocate the good to the bidder who submits the highest bid, both have the same distribution over the number of agents, and both implement the same payment rule for i. Thus it suffices to show that i's expected utility after rejecting his bidding ring invitation is less than his expected utility in the equilibrium of $(\star)$.

Given that all other bidders follow the strategies $(PT|R)$ and $(PT|\overline{R})$, if the bidding ring did not alter its behavior in response to i's deviation then there would exist some distributions p and signals s_i for which i would gain by declining the ring's invitation.[6] According to the protocol, however, the bidding ring *does* change its behavior in response to deviation. If i declines the invitation the ring center will send all the other members of the ring into the main auction, causing the auctioneer to announce $n + s_i - 1$ participants. As a result there will be $s_i - 1$ bidders placing bids of $b^e(v, p_{n,s_i})$ and $n - 1$ other bidders placing bids of $b^e(v, p_{n+s_i-1,1})$. We can show that these $n - 1$ bidders will always decrease i's expected utility by bidding too high. Recall Equation (9): $p_{n,s_i} = \left(\bigotimes_{n-1} \gamma_A\right) * \delta_{s_i}$, and so $p_{n+s_i-1,1} = \left(\bigotimes_{n+s_i-2} \gamma_A\right) * \delta_1$. We can write $\gamma_A = g_A * \delta_1$, where g_A is the distribution over the number of agents in a bidding ring beyond the first agent. Then

[6] Taking into account his signal and once the auctioneer has made an announcement, i would know that the number of agents is distributed according to p_{n,s_i}; however, if he were to deviate then all agents would bid in the main auction as though the number of agents were distributed according to $p_{n+1,1}$. For certain values of p and s_i, i's expected loss from causing the auctioneer to announce one more participant is less than his expected gain from being able to bid freely and from not having to make a payment to the ring center if he wins.

$$\begin{aligned} p_{n+s_i-1,1} &= \left(\bigotimes_{n-1} \gamma_A\right) * \left(\bigotimes_{s_i-1} \gamma_A\right) * \delta_1 \\ &= \left(\bigotimes_{n-1} \gamma_A\right) * \left(\left(\bigotimes_{s_i-1} g_A\right) * \delta_{s_i-1}\right) * \delta_1 \\ &= p_{n,s_i} * \left(\bigotimes_{s_i-1} g_A\right). \end{aligned} \tag{10}$$

Since γ_A has support on a subset of the positive integers, it follows that g_A has support on a subset of the integers greater than or equal to zero. And since $\gamma_A(1) < 1$, $g_A(0) < 1$. Then Equation (10) expresses convolutive dominance of $p_{n+s_i-1,1}$ over p_{n,s_i}, so it follows from Lemma 3.6 that $b^e(v, p_{n+s_i-1,1}) > b^e(v, p_{n,s_i})$. Thus if i declines the ring's invitation, the singleton bidders and other bidding rings will bid a higher[7] function of their valuations than the equilibrium amount in ($\star$). A bidder's expected gain in a first-price auction is always reduced as other bidders' bids increase, because his probability of winning decreases while his gain in the event of winning remains constant. This is the effect of i declining the offer to join his bidding ring: the $s_i - 1$ other bidders from i's bidding ring bid according to the equilibrium of ($\star$), but the $n - 1$ singleton and bidding ring bidders submit bids that exceed this amount. Thus declining the offer to participate reduces i's expected utility.

Part 2: $u(PT|\overline{R}) > u(P\overline{T}|\overline{R})$ and $u(PT|R) > u(P\overline{T}|R)$.

Since in this part we consider only strategies in which the agent decides to participate, it is sufficient to consider the equilibrium of a simpler, one-stage mechanism in which agents are given no choice about participation. Define the one-stage mechanism M as follows:

1. The center announces n, the number of bidders in the main auction.
2. Each bidder i submits a bid μ_i to the mechanism.
3. The bidder with the highest bid is allocated the good and is made to pay $b^e(\mu_i, p_{n,s_i})$.
4. All bidders with $s_i \geq 2$ are paid c_{n,s_i}.

M has the same payment rule for bidding ring bidders as the bidding ring protocol given above, but no longer implements a first-price payment rule for singleton bidders. Observe that the original auction is efficient under the strategies

[7] Note that this occurs because the singleton bidders and other bidding rings in the main auction follow a strategy that depends on the number of bidders announced by the auctioneer; hence they bid as though all the $s_i - 1$ bidders from the disbanded bidding ring might each be independent bidding rings.

stated in the theorem because each bidder i bids $b^e(v_i, p_{n,1})$ in the main auction. Thus, in order to prove $u(PT|\overline{R}) > u(P\overline{T}|\overline{R})$ and $u(PT|R) > u(P\overline{T}|R)$, it is sufficient to show that truthful bidding is an equilibrium for all bidders in mechanism M.

Assume that all other bidders bid truthfully, and consider the strategy of bidder i. This bidder's posterior distribution over the number of other bidders he faces, given his signal s_i and the auctioneer's announcement that there are n bidders in the main auction, is p_{n,s_i}. Since agent i is made to pay $b^e(\mu_i, p_{n,s_i})$ if he wins, and since the good is always allocated to the agent who submits the highest message, M is symmetrizing. From Lemma 3.4, agent i's best response to truthful bidding in a symmetrizing auction is to bid truthfully. Observe that this analysis holds for both non-ring and ring bidders since it does not require $s_i > 1$. If i is a ring bidder then he gets the additional payment c_{n,s_i} in both scenarios, but as this payment does not depend on the amount of his bid it does not affect his strategy given his decision to participate. ■

Note that this equilibrium gives rise to an economically-efficient allocation, as was mentioned in the proof of the theorem. The highest bidder in each bidding ring always bids in the main auction, and every bidder in the main auction places a bid according to the same function, which is monotonically increasing in the bidder's valuation. Because we have included bidders' decisions about whether to participate as part of their strategy spaces, the following result follows directly from the arguments in the proof of Theorem 3.7.

Corollary 3.8 *Participating in the bidding ring protocol according to the equilibrium from Theorem 3.7 is* ex post *individually rational for all bidders.*

4 Are Bidding Rings Helpful?

So far we have defined a bidding ring protocol and shown that in equilibrium members will participate truthfully and never regret participating. However, this does not show that the protocol is beneficial. In this section we show that it benefits all participants except for the auctioneer. We begin by showing that *payment equivalence*, a property related to revenue equivalence, holds in our setting. This result is a useful tool for establishing who gains from bidding rings. Using payment equivalence, we first show that the auctioneer's revenue is lower in the presence of bidding rings, and then that a ring center would be willing to run the protocol. Next we consider the question of whether the bidding ring benefits the agents. We provide affirmative answers to several different versions of the question: agents are better off than they would be if their own rings did not exist, than if other agents' rings did not exist, and as compared (given the same information either *ex ante* or *ex interim*) to a world in which no bidding rings are possible. Finally, we show that the protocol is *un*helpful in the sense that it does not offer a unique equilibrium.

4.1 Payment equivalence

Although the revenue equivalence theorem does not quite apply to the bidding ring protocol of Section 3.3, we are able to prove a very related property, *payment equivalence*.

Lemma 4.1 (payment equivalence) *The* ex ante *expected payment by an agent in E_{br} who follows the equilibrium of Theorem 3.7 (i.e., the sum of amounts paid by this bidder to the auctioneer and the ring center) is equivalent to the expected equilibrium payment of an agent in a first-price auction in E_s, where p and F are held constant between the two environments and where $c_{n,k} = 0$ in E_{br}.*

Proof. The proof has two parts. First we introduce a transformed mechanism M' in E_s that is payment equivalent to the bidding ring mechanism. Then we argue that with $c_{n,k} = 0$, M' is payment equivalent to a first-price auction in E_s. Combination of these two parts completes the proof.

The mechanism M' in E_s is constructed to mimic the bidding ring protocol in E_{br}, but with the ring centers and signals absorbed into the mechanism, and without decisions about participation. M' is defined as follows:

1. Each agent i sends a bid μ_i to the center (auctioneer).
2. From the number of bids received the center learns j, the total number of agents in the economic environment E_s. The center artificially allocates agents to bidding rings, choosing the number of rings and the populations of the rings according to γ_c and γ_A respectively, jointly conditioned on the observed total number of agents j. Combined with nature's distribution over j, Equation (7), this implements the same distribution over the cartel numbers and populations as is implemented by E_{br}.
3. The good is allocated to the agent h with the highest bid, μ_h. This winning agent is made to pay $b^e(\mu_h, p_{n,k})$ where n is the number of rings and k is the size of the ring containing agent h.
4. Each agent i receives a payment c_{n,k_i}, where k_i is the number of agents in the ring containing agent i.

As a function of the agents' bids, M' implements the same allocation rule and same *ex ante* expected payments as the bidding ring mechanism. Since equilibria are specified in terms of these functions, and since it is an equilibrium in the bidding ring mechanism for agents to truthfully declare their valuations, it follows that it is an equilibrium for agents to bid truthfully in M'. Thus as a function of agents' valuations, in equilibrium both M' and the bidding ring mechanism induce the same *ex ante* expected payments and allocation of the good at auction.

Now we must show that with $c_{n,k} = 0$, M' is payment-equivalent to a first-price auction in E_s. Observe that the revenue equivalence theorem holds in

E_s. That is, any auction mechanism in which in equilibrium, both (1) the good is allocated to the agent with the highest valuation; and (2) any agent with valuation zero has expected utility zero, induces the same expected payment for an agent with valuation v. Standard proofs of this result (see, e.g., [Klemperer, 1999]) begin by showing that every agent makes the same expected payment under any such mechanism. The same argument (not repeated here) suffices to show that for $c_{n,k} = 0$, the *ex ante* expected payments of agents in M' and a first price auction are the same in E_s; the requirement $c_{n,k} = 0$ ensures that agents with valuation zero have expected utility zero. ■

4.2 Ring centers gain at the expense of the auctioneer

Although payment equivalence holds, revenue equivalence does not. This is because some of the revenue that would be earned by the auctioneer in an auction in E_s is instead captured by the ring centers in the bidding ring mechanism.

Theorem 4.2 (*revenue inequivalence*) *The expected revenue from a first-price auction in E_{br} where bidders follow the equilibrium of Theorem 3.7 is less than the revenue of a first-price auction in E_s, where p and F are held constant between the two environments.*

Proof. First consider $c_{n,k} = 0$. From Lemma 4.1, the expected *ex ante* payments in the two auctions are the same. But whereas the center is the sole recipient of payments in the auction with a stochastic number of participants, in the bidding ring mechanism agents' payments are partially diverted from the center to the ring centers. That a positive amount is always diverted to the ring centers is established as follows. Since the distribution $p_{n,k}$ is just $p_{n,1}$ with $k-1$ singleton agents added, $p_{n,k} = p_{n,1} * \delta_{k-1}$. Since $k \geq 2$, it follows that $p_{n,k}$ convolutively dominates $p_{n,1}$. It then follows from Lemma 3.6 that $b^e(v_i, p_{n,k}) > b^e(v_i, p_{n,1})$. This proves that a ring center always receives a positive payment when a ring member wins; since every auction has a winner, a positive amount is diverted from the center to the ring centers.

Next consider the case in which $c_{n,k}$ is positive for some (n,k). The value of $c_{n,k}$ does not affect the allocation rule, and since the size of the side payment $c_{n,k}$ is independent of agents' bids, it does not affect agents' strategies either. Since this side payment is made between ring centers and agents, it has no effect on the expected revenue of the auctioneer. Therefore the auctioneer's expected revenue for nonzero $c_{n,k}$ is the same as it is for $c_{n,k} = 0$. ■

We can also show that ring centers experience a net gain on expectation from running bidding rings as long as the unconditional payment c is small enough.

Theorem 4.3 *The ring center gains on expectation if it pays agents $c_{n,k} = \frac{1}{k}(g_{n,k} - c'_{n,k})$ with $0 < c'_{n,k} \leq g_{n,k}$ and*

$$g_{n,k} = k \int_0^{\infty} f(v_i) \sum_{j=2}^{\infty} p_{n,k}(j) F^{j-1}(v_i) \left(b^e(v_i, p_{n,k}) - b^e(v_i, p_{n,1}) \right) dv_i,$$

and is budget-balanced on expectation when $c'_{n,k} = 0$.

Proof. Recall from the proof of Theorem 4.2 that a ring center always receives a positive payment when a ring member wins. $g_{n,k}$ is the ring center's *ex ante* expected gain if all k invited agents behave according to the equilibrium in Theorem 3.7, the auctioneer announces n participants, and the ring center makes no payment to the agents. Thus the ring center will gain on expectation if each ring member's unconditional payment is less than $\frac{1}{k}g_{n,k}$, and will budget-balance on expectation when each ring member's payment is exactly $\frac{1}{k}g_{n,k}$. ■

The payment of c to all bidders follows an idea from [Graham & Marshall, 1987] for returning a ring center's profits to bidders without changing incentives. In equilibrium the ring center will have an expected profit of $c'_{n,k}$, though it will lose $kc_{n,k}$ whenever the winner of the main auction does not belong to its ring. If a ring center wants to be guaranteed never to lose money, it can set $c'_{n,k} = g_{n,k}$.

4.3 Bidders gain as compared to a world without bidding rings

There are several ways of asking whether *bidders* gain by being invited to join bidding rings. One natural question is whether bidders prefer a world with bidding rings to a world without. We consider two other settings: an auction with participation revelation in E_s and an auction with a stochastic number of bidders in E_s. First we compare the three environments *ex ante*, asking which environment an agent would prefer if he knew the distribution over types but did not know what type he would receive. Second, we compare the environments ex interim, asking which environment an agent would prefer given knowledge of his own type. (Recall that we have defined an agent's type to include his signal s_i about the number of agents in the economic environment.)

We first consider the *ex ante* case. Observe that in this case an agent does not know whether or not he will be invited to a ring, as this is part of his type.

Theorem 4.4 (*ex ante*) *For all $n \geq 2$, as long as $\exists n$, $\exists k$, $\gamma_c(n) > 0$ and $\gamma_a(k) > 0$ and $c_{n,k} > 0$, agent i obtains greater expected utility by*

Case (1) *participating in E_{br} and following the equilibrium from Theorem 3.7 than by*

Case (2) *participating in a first-price auction with participation revelation in E_s with number of bidders distributed according to p; or by*

Case (3) *participating in a first-price auction with a stochastic number of bidders in E_s with number of bidders distributed according to p.*

When $\forall n$, $\forall k$, $c_{n,k} = 0$, agent i obtains the same expected utility in all three cases.

Proof. For $c_{n,k} = 0$ this result follows immediately from Lemma 4.1 and the fact that all three mechanisms are efficient. Now consider $c_{n,k} \geq 0$: The value of $c_{n,k}$ does not affect the allocation rule, nor does it affect the relative expected utility of any agent's strategy. Thus in Case (1) an agent's utility is higher than in Cases (2) and (3) by his expectation over signals of $c_{n,k}$. It follows that agent i prefers Case (1) as long as there exists a pair (n, k) that is realized with positive probability (*i.e.*, (n, k) for which $\gamma_c(n) > 0$ and $\gamma_a(k) > 0$) and for which $c_{n,k} > 0$; otherwise, i is indifferent between the three cases. ■

We now consider the *ex interim* case.

Theorem 4.5 (*ex interim*, ring members) *For all $\tau_i \in \mathcal{T}$, for all $k \geq 2$, for all $n \geq 2$, for all $c_{n,k} > 0$, agent i obtains greater expected utility by:*

Case (1) *participating in a bidding ring of size k in E_{br} and following the equilibrium from Theorem 3.7 than by*

Case (2) *participating in a first-price auction with participation revelation in E_s with number of bidders distributed according to $p_{n,k}$; or by*

Case (3) *participating in a first-price auction with a stochastic number of participants in E_s with number of bidders distributed according to $p_{n,k}$.*

When $c_{n,k} = 0$, agent i obtains the same expected utility in all three cases.

Proof. For an efficient first-price auction, an agent i's expected utility EU_i is $\sum_{j=2}^{\infty} p(j)F^{j-1}(v_i)b$, where $p(j)$ is the probability that there are a total of j agents in the economic environment, $F^{j-1}(v_i)$ is the probability that i has the high valuation among these j agents, and b is the amount of i's bid.

First, we consider Case (1). Let $EU_{i,bc}$ denote agent i's expected utility in E_{br} as a member of a bidding ring of size k, in the equilibrium from Theorem 3.7. Recall that in this equilibrium the bidder with the globally highest valuation always wins, and if bidder i wins he is made to pay $b^e(v_i, p_{n,k})$. In any case i receives an unconditional positive payment of $c_{n,k}$. Thus,

$$EU_{i,bc} = \sum_{j=2}^{\infty} p_{n,k}(j)F^{j-1}(v_i)\left(v_i - b^e(v_i, p_{n,k})\right) + c_{n,k}. \tag{11}$$

We now consider Case (2). From Proposition 2.3 it is an equilibrium for agent i in economic environment E_s to bid $b^e(v_i, j)$ in a first-price auction with participation revelation, where j is the number of bidders announced by the auctioneer. Since the number of agents is distributed according to $p_{n,k}$, agent i's expected utility in a first-price auction with participation revelation, which we denote $EU_{i,pr}$, is

$$
\begin{aligned}
EU_{i,pr} =& \sum_{j=2}^{\infty} p_{n,k}(j) F^{j-1}(v_i)\left(v_i - b^e(v_i, j)\right) && (12)\\
=& \frac{\sum_{\ell=2}^{\infty} p_{n,k}(\ell) F^{\ell-1}(v_i)}{\sum_{\ell'=2}^{\infty} p_{n,k}(\ell') F^{\ell'-1}(v_i)} \sum_{j=2}^{\infty} p_{n,k}(j) F^{j-1}(v_i)\left(v_i - b^e(v_i, j)\right)\\
=& \sum_{\ell=2}^{\infty} p_{n,k}(\ell) F^{\ell-1}(v_i)\left(v_i - \sum_{j=2}^{\infty} \frac{p_{n,k}(j) F^{j-1}(v_i)}{\sum_{\ell'=2}^{\infty} p_{n,k}(\ell') F^{\ell'-1}(v_i)} b^e(v_i, j)\right)\\
=& \sum_{\ell=2}^{\infty} p_{n,k}(\ell) F^{\ell-1}(v_i)\left(v_i - b^e(v_i, p_{n,k})\right). && (13)
\end{aligned}
$$

Observe that we make use of the definition of $b^e(v_i, p)$ from Equation (2). Equation (13) is agent i's expected utility in Case (3), so i's expected utility is equal in Cases (2) and (3).

Combining equations (11) and (13), we obtain

$$
EU_{i,bc} - EU_{i,pr} = c_{n,k}. \tag{14}
$$

When $c_{n,k} > 0$, agent i's expected utility is strictly greater in Case (1) than in Cases (2) and (3); when $c_{n,k} = 0$ he has the same expected utility in all three cases. ■

What about agents who do not belong to bidding rings? We can show in the same way that they are not harmed by the existence of bidding rings: they are neither better nor worse off in the bidding ring economic environment than facing the same distribution of opponents in the two cases described above.

Corollary 4.6 (*ex interim*, singleton bidders) *For all $\tau_i \in \mathcal{T}$, for all $n \geq 2$, agent i obtains the same expected utility in each of the following cases:*

Case (1) *participating as a singleton bidder in E_{br} and following the equilibrium from Theorem 3.7;*

Case (2) *participating in a first-price auction with participation revelation in E_s with number of bidders distributed according to $p_{n,1}$;*

Case (3) *participating in a first-price auction with a stochastic number of participants in E_s with number of bidders distributed according to $p_{n,k}$.*

Proof. We follow the same argument as in Theorem 4.5, except that $k = 1$ and $EU_{i,bc}$ does not include $c_{n,k}$. Thus we get $EU_{i,bc} = EU_{i,pr}$. ■

4.4 Bidders gain as compared to a world where one ring doesn't exist

Another way of showing that bidding rings are helpful is to demonstrate that bidders are better off being invited to a bidding ring than being sent to the auction as singleton bidders.

Theorem 4.7 (ring members) *An agent i has higher expected utility*[8] *in a bidding ring of size k bidding as described in Theorem 3.7 than he does if the bidding ring does not exist and k additional agents (including i) participate directly in the main auction as singleton bidders, again bidding as described in Theorem 3.7, for $c_{n,k} \geq 0$.*

Proof. Consider the counterfactual case where agent i's bidding ring does not exist, and all the members of this bidding ring are replaced by singleton bidders in the main auction. We show that i is better off as a member of the bidding ring (even when $c_{n,k} = 0$) than in this case. If there were n potential ring centers in the original auction and k agents in i's bidding ring, then the auctioneer would announce $n + k - 1$ as the number of participants in the new auction. In both cases the auction is economically efficient, which means i is better off in the auction that requires him to pay a smaller amount when he wins. Under the equilibrium from Theorem 3.7, as a singleton bidder i will pay $b^e(v_i, p_{n+k-1,1})$ when he wins. If he belonged to the bidding ring and followed the same equilibrium i would pay $b^e(v_i, p_{n,k})$ when he wins. As argued in the proof of Theorem 3.7, Lemma 3.6 shows that $\forall k \geq 2, \forall n \geq 2, \forall v, b^e(v, p_{n+k-1,1}) > b^e(v, p_{n,k})$, and so our result follows. ■

Intuitively, an agent gains by not having to consider the possibility that other bidders who would otherwise have belonged to his bidding ring might themselves be bidding rings.

We can also show that singleton bidders and members of other bidding rings benefit from the existence of each bidding ring in the same sense. Following an argument similar to the one in Theorem 4.7, other bidders gain from not having to consider the possibility that additional bidders might represent bidding rings. Paradoxically, as long as $c'_{n,k} > 0$, other bidders' gain from the existence of a given bidding ring is greater than the gain of that ring's members.

Corollary 4.8 (ring non-members) *In the equilibrium described in Theorem 3.7, singleton bidders and members of other bidding rings have higher expected utility when $k \geq 2$ agents form a bidding ring than when k additional agents participate directly in the main auction as singleton bidders.*

Proof. Consider a singleton bidder i in the first case, where the ring of k agents does exist. (It is sufficient to consider a singleton bidder, since other bidding rings bid in the same way as singleton bidders.) Following the equilibrium from

[8] This is weakly higher (i.e., equal) expected utility for agents with the lowest possible valuation, and strictly higher expected utility otherwise. The same caveat also holds below.

Theorem 3.7, i would submit the bid $b^e(v_i, p_{n,1})$. In the second case, following the equilibrium from Theorem 3.7, i would bid $b^e(v_i, p_{n+k-1,1})$. In both cases, the auction is economically efficient, so i is better off in the case where he makes the smaller bid. From the argument in Theorem 4.3 we know that $\forall k \geq 2, b^e(v_i, p_{n,1}) < b^e(v_i, p_{n,k})$; from the argument in Theorem 3.7 Part (1b) we know that $\forall k \geq 2, b^e(v, p_{n+k-1,1}) > b^e(v, p_{n,k})$. Thus $\forall k \geq 2, b^e(v_i, p_{n,1}) < b^e(v_i, p_{n+k-1,1})$. ■

4.5 Another Equilibrium

So far we have considered whether agents benefit under the equilibrium from Theorem 3.7. However, we can also show that this equilibrium is not unique. There is another equilibrium under which no agents accept bidding ring invitations, and they instead bid according to the equilibrium for first-price auctions with participation revelation given in Proposition 2.3.

Proposition 4.9 *It is a Bayes-Nash equilibrium for each bidding ring invitee to decline his bidding ring invitation, and for each agent i to bid $b^e(v_i, n)$.*

Proof. If at least one agent declines the invitation to join a bidding ring, other invitees of that bidding ring are at least as well off if they decline as well. (If they decline then they can bid freely, rather than being made to submit bids of a particular form.) If no agents join bidding rings then agents' signals contain no useful information. Thus the argument from Proposition 2.3 applies, and it is a Bayes-Nash equilibrium for each bidder to submit a bid of $b^e(v_i, n)$. ■

The theorems and corollaries in Section 4 allow us to compare our first equilibrium (from Theorem 3.7) with this new equilibrium (from Proposition 4.9).

Corollary 4.10 *When $c_{n,k} > 0$, all bidders prefer the equilibrium from Theorem 3.7 to the equilibrium from Proposition 4.9* ex ante*;* ex post *bidding ring invitees prefer the equilibrium from Theorem 3.7 to the equilibrium from Proposition 4.9, while singleton bidders are indifferent between the equilibria. When $c_{n,k} = 0$, all bidders are indifferent between the equilibria both* ex ante *and* ex post*.*

Proof. A bidder's expected utility under the equilibrium from Proposition 4.9 in economic environment E_{br} is the same as his expected utility from an auction with participation revelation in economic environment E_s with the same distribution over the number of bidders, since (as given by Propositions 2.3 and 4.9) in both cases each bidder i follows the strategy $b^e(v_i, n)$. Then the result is immediate from Theorem 4.5, Corollary 4.6 and Corollary 4.4. ■

Since both bidders (and, trivially, ring centers) prefer the equilibrium from Theorem 3.7 to the equilibrium from Proposition 4.9, it follows that auctioneers have opposite preferences. It turns out that auctioneers can disrupt bidding rings by slightly changing the rules of the auction so that the strategies described in Theorem 3.7 are no longer constitute an equilibrium while the equilibrium from Proposition 4.9 is preserved. This can be achieved by making it possible for bidders to participate in their bidding rings and also place shill bids in the main auction without detection by the ring center.[9] If all agents but i followed the strategies specified in Theorem 3.7, i could declare a low valuation to the ring center but also place a competitive bid in the main auction, gaining all the benefits of the cartel without having to make any payments to the ring center and without causing the ring to change its behavior because its invitation had been declined. Note however that these defenses may not be available to all auctioneers; for example, the auctioneer might be required to verify and announce the winner's identity.

5 Conclusions

We have presented a formal model of bidding rings in first-price auctions that in many ways extends models traditionally used in the study of collusion. Most importantly, in our model all agents behave strategically and take into account the possibility that groups of other agents will collude. Other features of our setting include a stochastic number of agents and of bidding rings in each auction, revelation by the auctioneer of the number of bids received, but bidders' inability to detect the suppression of other bids by one or more bidding rings. The strategy space is expanded so that the decision of whether or not to join a bidding ring is part of an agent's choice of strategy.

We showed a bidding ring protocol for first-price auctions that leads to a (globally) efficient allocation in equilibrium. In this equilibrium all invited agents choose to participate, even when the bidding ring operates in a single auction as opposed to a sequence of auctions. This means that the protocol's stability does not rely on the threat of an agent being denied future opportunities to collude. Bidding rings make money on expectation, and can optionally be configured so they never lose money.

We asked the question of whether agents gain by participating in bidding rings in first-price auctions in three different ways:

1. Could any agent gain by deviating from the protocol?

2. Would any agent be better off if his bidding ring did not exist?

3. Would any agent would be better off (either *ex interim* or *ex ante*) in an economic environment that did not include bidding rings at all?

We have shown that agents are strictly better off in all three senses. (In the third sense, the gain is only strict when ring centers make an otherwise optional

[9] As mentioned earlier, despite studying a fairly different setting, Marshall and Marx [2007] found that their own bidding ring protocol for unrepeated first-price auctions could also be disrupted by allowing bidders to place shill bids.

side-payment to agents.) We have also shown that each bidding ring causes *non-members* to gain in the second sense, and does not hurt them in the third sense.

Our work provides many opportunities for further study. These include the following questions, all of which refer to first-price auctions in economic environment E_{br}.

- What is the optimal bidding ring protocol? (We conjecture that our protocol, with the largest possible c, gives agents the highest possible *ex ante* expected utility, as compared to all efficient protocols in which the ring center does not lose on expectation.)

- Is it possible to construct bidding ring protocols that play best responses to each other in the main auction? (Observe that under our protocol, ring centers do not behave strategically.)

- Can a bidding ring protocol be made to budget-balance *ex post*, e.g., using ideas similar to those that Mailath and Zemsky [1991] applied to second-price auctions?

- What is a full characterization of the symmetric equilibria of our protocol? (We conjecture that beyond the equilibria we described, it is an equilibrium for agents in one cartel to continue to bid truthfully when another cartel disbands.)

- Can our protocol be extended to allow agents to receive multiple invitations to join different bidding rings? If so, will agents always want to join the largest possible coalition? (We conjecture that they will.)

- Can any bidding ring protocol withstand shill bidding?

- Can a protocol be identified that works with risk-averse bidders, affiliated values, asymmetric valuation distributions, or distributions over the number of bidders that violate the independent cartel property?

Acknowledgements

We are grateful to Yossi Feinberg, Arieh Gavious, Michael Peters and Ryan Porter for their very helpful comments. Our work was supported by DARPA grant F30602-00-2-0598, a Stanford Graduate Fellowship, an NSERC discovery grant and an Ontario Graduate Scholarship.

A Relating p to b^e

This appendix provides more technical detail about the relationship between p and b^e discussed in Section 3.2.

It is intuitive to expect that in first-price auctions, the amounts of agents' equilibrium bids increase with the number of participating agents. We can easily verify that this is true in the classical case.

Lemma A.1 *$\forall v$, $\forall j \geq 2$, $b^e(v, j+1) > b^e(v, j)$.*

Proof. From Equation (1), we can write

$$b^e(v, j+1) - b^e(v, j) = \int_0^v \left(1 - \left(\frac{F(u)}{F(v)}\right)\right)\left(\frac{F(u)}{F(v)}\right)^{j-1} du. \tag{15}$$

The first factor in the integrand is clearly always positive, so the right-hand side of Equation (15) is positive. Thus $b^e(v, j)$ is strictly increasing in j. ■

This intuition does not transfer to first-price auctions with a stochastic number of bidders, in the sense that auctions with larger expected numbers of participants do not always yield higher equilibrium bids.

Example A.2 *Consider a distribution p such that only two numbers of agents have nonzero probability: j_{low} and j_{high}. Furthermore, let the probability mass be evenly divided between j_{low} and j_{high}, and denote by κ the distance between them: $j_{high} = j_{low} + \kappa$. Now consider the strategy of agent i. The classical case is recovered if $\kappa = 0$, in which case i's equilibrium bid will just be $b^e(v_i, j_{low})$. If κ is increased to 1, the equilibrium bid increases by a finite amount to some $b^e(v_i, p) \in (b^e(v_i, j_{low}), b^e(v_i, j_{high}))$, as determined by Equation (2) from Section 2.2. As κ is increased to an arbitrarily high value, $F(v_i)^{j_{high}-1}$, the probability that agent i has the highest valuation when there are j_{high} agents involved approaches zero. With arbitrarily close to unit probability, there will be j_{low} agents involved when agent i has the highest valuation, and Equation (2) indicates that i's bid will be arbitrarily close to the $\kappa = 0$ result. Thus while the $\kappa \to \infty$ distribution has a higher expected number of participants than the $\kappa = 1$ distribution, it elicits a lower equilibrium bid.*

This phenomenon also occurs among distributions of practical interest. So in a first-price auction with a stochastic number of participants, simply knowing that distribution p has a smaller expected number of participants than distribution p' is not enough to know which distribution gives rise to a lower symmetric equilibrium bid for a given valuation. The same holds for stochastic dominance. For example, in Example A.2, the distribution with very large κ stochastically dominates the distribution with $\kappa = 1$ but elicits a lower equilibrium bid. 3.6 identifies a class of pairs of distributions (p, p') for which it does hold that $b^e(v_i, p) < b^e(v_i, p')$, and this class is those for which p' convolutively dominates p. Before we can prove this lemma, we must define additional notation that was not given in Section 3.2.

Let $r_j(F, v_i, p)$ denote the probability that j agents participate conditional on agent i having the highest valuation. This is equal to the probability that j agents participate and agent i has the highest valuation among these agents, normalized by the unconditional probability that agent i has the highest valuation. Let $Z(F, v_i, p)$ be the probability that agent i has the highest valuation given that his valuation is v_i. Thus

$$Z(F, v_i, p) \equiv \sum_{k=2}^{\infty} F(v_i)^{k-1} p(k); \tag{16}$$

$$r_j(F, v_i, p) \equiv \frac{F(v_i)^{j-1} p(j)}{Z(F, v_i, p)}. \tag{17}$$

Observe[10] that Equation (2) for the equilibrium bid in a stochastic first-price auction can be written in terms of the distribution $r(F, v_i, p)$:

$$b^e(v_i, p) = \sum_{j=2}^{\infty} r_j(F, v_i, p) \, b^e(v_i, j). \tag{18}$$

The cumulative distribution $R_m(F, v_i, p)$ for the distribution r, denoting the probability that m or fewer agents participate conditional on i having the highest valuation, is simply

$$R_m(F, v_i, p) \equiv \sum_{j=2}^{m} r_j(F, v_i, p). \tag{19}$$

Lemma A.3 *$\forall p, p' \in \mathcal{D}_2$, if p' convolutively dominates p then $b^e(v_i, p) < b^e(v_i, p')$.*

Proof. The proof has two parts. First we show that for every j, the probability that no more than j bidders participate conditional on bidder i having the highest valuation is at least as high when the number of agents is drawn from p as when it is drawn from p', and that for some j this probability is higher in p than in p'. Next, we show that this relationship between conditional probabilities implies that the equilibrium bid is smaller under p than under p'.

Step 1: $\forall j$, $R_j(F, v_i, p') \leq R_j(F, v_i, p)$, and $\exists j$, $R_j(F, v_i, p') < R_j(F, v_i, p)$.

Consider the difference between the cumulative distributions:

$$\begin{aligned} \Delta R_j &\equiv R_j(F, v_i, p) - R_j(F, v_i, p') \\ &= \sum_{m=-\infty}^{j} \left(\frac{F(v_i)^{m-1} p(m)}{Z(F, v_i, p)} - \frac{F(v_i)^{m-1} p'(m)}{Z(F, v_i, p')} \right). \end{aligned} \tag{20}$$

[10] We can use 2 rather than $-\infty$ as the lower limit of the sum in Equation (16) because $p(k)$ has support which is a subset of $\{2, 3, \ldots\}$. While $Z(F, v_i, p)$ is undefined when $F(v_i) = 0$, this technicality is of no practical interest.

The denominators can be related as follows:

$$\begin{aligned}
Z(F, v_i, p') &= \sum_{k=-\infty}^{\infty} F(v_i)^{k-1} \sum_{j=0}^{\infty} p(k-j) q(j) \\
&= \sum_{j=0}^{\infty} \sum_{k=-\infty}^{\infty} \left(F(v_i) F(v_i)^{j-1} F(v_i)^{k-j-1} \right) p(k-j) q(j) \qquad (21) \\
&= F\left(v_i\right) \sum_{j=0}^{\infty} F\left(v_i\right)^{j-1} q\left(j\right) \sum_{k=-\infty}^{\infty} F\left(v_i\right)^{k-j-1} p\left(k-j\right) \\
&= F\left(v_i\right) Z\left(F, v_i, q\right) Z\left(F, v_i, p\right). \qquad (22)
\end{aligned}$$

Substituting Equation (22) into Equation (20), and making use of Equation (3),

$$\begin{aligned}
\Delta R_j &= \frac{1}{Z\left(F, v_i, p'\right)} \sum_{m=-\infty}^{j} \Bigg(Z\left(F, v_i, q\right) F\left(v_i\right)^m p\left(m\right) \\
&\quad - F\left(v_i\right)^{m-1} \sum_{k=0}^{\infty} p\left(m-k\right) q\left(k\right) \Bigg) \qquad (23) \\
&= \frac{F\left(v_i\right)}{Z\left(F, v_i, p'\right)} \Bigg(\sum_{k=0}^{\infty} q\left(k\right) F\left(v_i\right)^{k-1} \sum_{m=-\infty}^{j} F\left(v_i\right)^{m-1} p\left(m\right) \\
&\quad - \sum_{k=0}^{\infty} q\left(k\right) F\left(v_i\right)^{k-1} \sum_{m=-\infty}^{j} F\left(v_i\right)^{m-k-1} p\left(m-k\right) \Bigg) \qquad (24) \\
&= \frac{F\left(v_i\right)}{Z\left(F, v_i, p'\right)} \sum_{k=0}^{\infty} q\left(k\right) F\left(v_i\right)^{k-1} \Bigg(\sum_{m=-\infty}^{j} F\left(v_i\right)^{m-1} p\left(m\right) \\
&\quad - \sum_{m=-\infty}^{j-k} F\left(v_i\right)^{m-1} p\left(m\right) \Bigg). \qquad (25)
\end{aligned}$$

To obtain Equation (24), we have reordered the sums, made use of Equation (16) and performed factoring like that done to obtain Equation (21). To obtain Equation (25), we have factored the bracketed expression in Equation (24) and shifted the dummy indices of the second sum.

When $k = 0$, the bracketed expression in Equation (25) is zero, so that term can be dropped from the sum. The bracketed sums can then be combined, yielding

$$\Delta R_j = \frac{F\left(v_i\right)}{Z\left(F, v_i, p'\right)} \sum_{k=1}^{\infty} \left(F\left(v_i\right)^{k-1} q\left(k\right) \sum_{m=j-k+1}^{j} F\left(v_i\right)^{m-1} p\left(m\right) \right). \qquad (26)$$

Since $k \in [1, \infty)$ in Equation (26), the lower summand of the second sum is always less than or equal to the upper summand, so that sum is well-defined. Furthermore, all of the factors in Equation (26) are non-negative, so it remains only to be established whether $\Delta R_j > 0$ or $\Delta R_j = 0$. Since $p \in \mathcal{D}_2$, there exists some least element in the support of p; call this value m^*. For values of $j < m^*$ the second sum in Equation (26) gives exactly 0, and $\Delta R_j = 0$. Similarly, for all values of $j \geq m^*$, the second sum is nonzero, and since by assumption $\exists k > 0$ such that $q(k) > 0$, we have that $\Delta R_j > 0$. Thus for all j, $R_j(F, v_i, p') \leq R_j(F, v_i, p)$, and for some j, $R_j(F, v_i, p') < R_j(F, v_i, p)$.

Step 2: $(\forall j, R_j(F, v_i, p') \leq R_j(F, v_i, p)$ and $\exists j,\ R_j(F, v_i, p') < R_j(F, v_i, p))$ implies $b^e(v_i, p) < b^e(v_i, p')$.

We must show that $\Delta b > 0$, where we use Equation (18) to write

$$\begin{aligned}
\Delta b &\equiv b^e(v_i, p') - b^e(v_i, p) \\
&= \sum_{m=2}^{\infty} \left(r_m(F, v_i, p') - r_m(F, v_i, p)\right) b^e(v_i, m).
\end{aligned}$$

We rewrite this sum using summation by parts (the discrete analog of integration by parts). This yields

$$\Delta b = \sum_{m=2}^{\infty} \left(b^e(v_i, m+1) - b^e(v_i, m)\right) \sum_{j=2}^{m} \left(r_j(F, v_i, p) - r_j(F, v_i, p')\right) \quad (27)$$

$$= \sum_{m=2}^{\infty} \left(b^e(v_i, m+1) - b^e(v_i, m)\right) \left(R_m(F, v_i, p) - R_m(F, v_i, p')\right). \quad (28)$$

To obtain Equation (27), we have also used the fact that both $r(F, v_i, p)$ and $r(F, v_i, p)$ are normalized. Lemma A.1 tells us that $b^e(v_i, m)$ is strictly increasing in m; clearly it is always positive. Thus $b^e(v_i, m+1) - b^e(v_i, m) > 0$ $\forall m$. Furthermore, from Step 1, $R_m(F, v_i, p) - R_m(F, v_i, p')$ is non-negative, and for all $m \geq m^*$ it is greater than zero. The right-hand side of Equation (28) is therefore a sum of products of non-negative factors, of which at least one is a product of strictly positive factors. Thus $\Delta b > 0$, or $b^e(v_i, p) < b^e(v_i, p')$. ∎

References

Aoyagi, M. (2003). Bid rotation and collusion in repeated auctions. *Journal of Economic Theory*, *112*(1), 79–105.

Blume, A., & Heidhues, P. (2008). Modeling tacit collusion in auctions. *Journal of Institutional and Theoretical Economics*, *164*(1), 163–184.

Cramton, P., & Palfrey, T. (1990). Cartel enforcement with uncertainty about costs. *International Economic Review*, *31*(1), 17–47.

Feinstein, J., Block, M., & Nold, F. (1985). Asymmetric behavior and collusive behavior in auction markets. *American Economic Review*, *75*(3), 441–460.

Fudenberg, D., Levine, D., & Maskin, E. (1994). The Folk Theorem with Imperfect Public Information. *Econometrica*, *62*(5), 997–1039.

Graham, D., & Marshall, R. (1987). Collusive bidder behavior at single-object second-price and English auctions. *Journal of Political Economy*, *95*, 579–599.

Graham, D., Marshall, R., & Richard, J.-F. (1990). Differential payments within a bidder coalition and the Shapley value. *American Economic Review*, *80*(3), 493–510.

Harstad, R., Kagel, J., & Levin, D. (1990). Equilibrium bid functions for auctions with an uncertain number of bidders. *Economics Letters*, *33*(1), 35–40.

Hendricks, K., & Porter, R. (1989). Collusion in auctions. *Annales d'Économie et de Statistique*, *15/16*, 216–229.

Hörner, J., & Jamison, J. S. (2007). Collusion with (almost) no information. *RAND Journal of Economics*, *38*(3), 804–822.

Klemperer, P. (1999). Auction theory: A guide to the literature. *Journal of Economic Surveys*, *13*(3), 227–286.

Leyton-Brown, K., Shoham, Y., & Tennenholtz, M. (2000). Bidding clubs: institutionalized collusion in auctions. *ACM Conference on Electronic Commerce* (pp. 253–259).

Leyton-Brown, K., Shoham, Y., & Tennenholtz, M. (2002). Bidding clubs in first-price auctions. *American Association for Artificial Intelligence (AAAI)* (pp. 373–378).

Mailath, G., & Zemsky, P. (1991). Collusion in second-price auctions with heterogeneous bidders. *Games and Economic Behavior*, *3*, 467–486.

Marshall, R. C., & Marx, L. M. (2007). Bidder collusion. *Journal of Economic Theory*, *133*, 374–402.

McAfee, R., & McMillan, J. (1987). Auctions with a stochastic number of bidders. *Journal of Economic Theory*, *43*, 1–19.

McAfee, R., & McMillan, J. (1992). Bidding rings. *American Economic Review*, *82*, 579–599.

Riley, J., & Samuelson, W. (1981). Optimal auctions. *American Economic Review*, *71*, 381–392.

Robinson, M. (1985). Collusion and the choice of auction. *RAND Journal of Economics*, *16*(1), 141–145.

Skrzypacz, A., & Hopenhayn, H. (2004). Tacit collusion in repeated auctions. *Journal of Economic Theory*, *114*(1), 153–169.

von Ungern-Sternberg, T. (1988). Cartel stability in sealed bid second price auctions. *Journal of Industrial Economics*, *18*(3), 351–358.