Auction Design with Tacit Collusion*

Atila Abdulkadiroğlu
Department of Economics
Columbia University
420 West 118th Street
New York, NY 11027
U.S.A.
aa2061@columbia.edu

Kim-Sau Chung
Department of Economics
Northwestern University
2003 Sheridan Road
Evanston, IL 60208
U.S.A.
sau@northwestern.edu

June 16, 2003

Abstract

We study an auction design problem where the auctioneer anticipates that bidders collude tacitly. We model tacit collusion in the following way: whichever auction the auctioneer announces, it becomes the stage game of a corresponding repeated game, and out of all sequential equilibria of the resulting repeated game, bidders play the one that maximizes their total ex ante payoff. We show that an optimal auction in this model takes the form of a public good provision game.

JEL: D44.

Keywords: auction design, tacit collusion.

*We thank Yeon-Koo Che, Jeff Ely, and Larry Samuelson, for comments. All errors are ours.

1 Introduction

We study an auction design problem where the auctioneer anticipates that bidders will always collude tacitly when they play any announced auction mechanism. We model tacit collusion in the following way: whichever auction mechanism the auctioneer announces, it would become the stage game of a corresponding repeated game, and out of all sequential equilibria in this corresponding repeated game, bidders play the one that maximizes their total ex ante payoff. We show that an optimal auction in this model exhibits qualitative features quite opposite to those in a standard model. Instead of creating negative externalities among bidders in the sense that when one bidder bids higher every other bidder gets a lower payoff, here an optimal auction should actually create *positive* externalities among bidders in the sense that when one bidder bids higher every other bidder gets a *higher* payoff. We show that one auction that delivers such properties is one that takes the form of a public good provision game (Theorem 1).

These results echo those of a highly related paper by Che and Yoo (2001; hereafter CY). CY studied the optimal multi-agent incentive contract design problem in a similar repeated-interaction setting. They also found that an optimal contract in their model exhibits qualitative features quite opposite to those in a standard model. In particular, while optimal contracts in a standard model usually take the form of relative performance evaluation, an optimal contract in their model takes the form of joint performance evaluation. Note that the logic behind relative performance evaluation is to create negative externalities among workers (when one worker works harder every other worker gets a lower payoff), and that behind joint performance evaluation is to create *positive* externalities among workers (when one worker works harder every other worker gets a *higher* payoff). CY use this observation to explain why "teams" arise. The main difference between CY and this paper is that, while their problem is a moral harzard one, ours is an adverse selection one. But the intuition behind both papers is the same.

Our model fits better the story of an auction house who faces more or less the same set of serious bidders over time but cannot afford to change its auction mechanism frequently. As an extension to our main result, we also study an alternative model in Section 5, where the auctioneer can change her auction mechanism every period, but cannot commit to which auction mechanisms she will use in the future. In this alternatively model, the implication of our way to model tacit collusion is very strong. In particular, in the sequential equilibrium that maximizes bidders' total ex ante payoff, bidders jointly capture all the surplus. Note that in deriving this strong result, we allow the auctioneer to use secretive selling mechanisms (i.e., we allow the auctioneer to approach one of the bidders in secret and to promise this bidder not to disclose any trade details to any third party).

1.1 Related Literature

There are many studies on communicative collusion. If the communication channel is rich enough, pre-play communication can be modelled as a message game. Depending on whether the outcomes of the message game are to be enforced by side contracts or by repeated interaction, the literature can be divided into two corresponding categories. For example, Graham and Marshall (1987), McAfee and McMillan (1992), and Marshall and Marx (2002) belong to the side-contract category; whereas Athey and Bagwell (2001) and Aoyagi (2002) belong to the repeated-interaction category. These papers do not deal with mechanism design, as they restrict their attention to fixed (subsets of) mechanisms.

A relatively smaller literature touches upon tacit collusion. How the mechanism designer manages to forbid agents from talking to each other before playing the announced mechanism is usually not modelled. But the assumption that she somehow manages to do so is not unrealistic. Cramton (1995) documents probably the most entertaining example of tacit collusion in auction history, where bidders of the first FCC auction used the last few digits of their bids to communicate their preferences to each other. The fact that bidders needed to communicate their preferences *during* the auction can be viewed as an evidence that the auction designer (the FCC in this example) somehow managed to forbid bidders from communicating *before* the auction.

If the announced mechanism already possesses multiple equilibria, then tacit collusion can be modelled as an equilibrium selection rule such that agents always play the most "collusive" equilibrium. Otherwise, tacit collusion can still be studied in a model of repeated interaction. Examples of repeated-interaction model of tacit collusion include Athey, Bagwell, Sanchirico (2000), Athey and Bagwell (2001), Blume and Heidhues (2001), and Skrzypacz and Hopenhayn (2001). All these papers restrict their attention to fixed (subsets of) grand mechanisms. Once again, these papers do not deal with mechanism design, as they restrict their attention to fixed (subsets of) mechanisms.

Laffont and Martimort (1997, 2000) deal with mechanism design with communicative collusion. This paper is parallel to their papers, and deals with mechanism design with tacit collusion.

2 The Model

There is one auctioneer and a set N of $n \geq 2$ symmetric bidders. The auctioneer has one (perishable) object to sell every period. The object is indivisible, and is worth nothing to the auctioneer. Each bidder i has a private valuation $v_{it} \in [0, 1]$ over the object being sold in period t, where v_{it} is IID across i and t. For any i and t, v_{it} follows the probability distribution function F, which has a strictly positive density function f over $[0, 1]$.

There are infinitely many periods (i.e., $t = 1, 2, \ldots$). Before the beginning of period 1 (say in period 0), the auctioneer has one single chance to choose an auction game. Then, in each period t, bidders first learn their private valuations, and then play the auction game chosen in period 0.

It is important for the reader to distinguish two kinds of game here: (i) the auction game which is played period by period, and (i) the repeated game which has infinite horizon. Formally, an auction game is a selling mechanism (S, p, q, M, α), where:

- $S := \times_{i \in N} S_i$ is a vector of message spaces; each S_i is a set of possible messages bidder i can send to the auctioneer, with the restriction that $S_i = \{0_i\} \cup \hat{S}_i$, where 0_i is the "non-participation" message;

- $p : S \to \{(p_1, \ldots, p_n) \in [0,1]^n \mid \sum_{i \in N} p_i \leq 1\}$ is an allocation function that specifies the probabilities with which each bidder will get the object, with the restriction that $\forall i$, $p_i(\cdots, 0_i, \cdots) \equiv 0$ (non-participating bidders never get the object);

- $q : S \to \mathbb{R}^n$ is a payment function that specifies the amount each bidder has to pay the auctioneer, with the restriction that $\forall i$, $q_i(\cdots, 0_i, \cdots) \equiv 0$ (non-participating bidders never pay the auctioneer);

- $M := \times_{i \in N} M_i$ is a vector of message spaces; each M_i is a set of possible messages the auctioneer can send to bidder i; and

- $\alpha : S \to \Delta(M)$ is an announcement function that specifies what information the auctioneer would disclose to each bidder after the auction.

The above definition of an auction game is standard except for the announcement function. The specification of the announcement function is not important in any one-shot auction game, as it will not affect bidders' incentives. It becomes important when we study repeated auctions, as it now affects bidders' ability to collude. In the definition above, α_i (the projection of α onto $\Delta(M_i)$) specifies the (potentially random) message the auctioneer tells bidder i after the auction, and these messages can be correlated across bidders.[1]

All players are risk neutral: the auctioneer's period-t payoff u_t is simply her revenue in period t; and each bidder i's period-t payoff u_{it} is equal to $p_i v_{it} - q_i$, where p_i is his probability of obtaining the object in period t, and q_i is the amount he pays the auctioneer in period t. Notice that if a bidder does not participate in the period-t auciton game, his period-t payoff will be zero.

All players have common discount factor δ.[2] We follow the convention in the repeated-game literature and normalize a player's discounted sum of payoffs by the factor $(1-\delta)$. For example, if bidder i's period-t payoffs are u_{it}, $t = 1, 2, \ldots$, then his normalized dicounted sum of payoffs will be $(1-\delta)\sum_{t=1}^{\infty} \delta^{t-1} u_{it}$.

Once the auctioneer has chosen an auction game in period 0, the bidders are playing an induced repeated game. Typically, this repeated game will have multiple sequential equilibria. Our way of modelling tacit collusion is to focus on the sequential equilibrium that maximizes bidders' total payoff (hereafter the bidder-optimal equilibrium). In other words, we model collusive bidding by assuming that bidders can coordinate on their favorite

[1] More generally, we should allow the announcement function to depend on the realization of the winner when the winner is being picked randomly. This can easily be done with extra notations.

[2] That the auctioneer has the same discount factor bidders have is both important and unimportant. It is important because otherwise infinite surplus can be generated out of thin air by simply having players trade intertemporally. It is however *not* important for our main result that an optimal auction takes the form of a public good provision game. See also Footnote 5.

sequential equilibrium. If there are more than one sequential equilibrium that maximizes bidders' total payoff, we focus our attention to the one that maximizes the auctioneer's payoff.

Finally, the following definitions will prove useful in subsequent sections. We shall use W^* to denote each period's expected social surplus; i.e., $W^* := \mathbf{E}\max\{v_{1t}, \ldots, v_{nt}\} = \int_0^1 v\mathrm{d}F(v)^n$ (after normalization W^* is also the discounted sum of social surplus). We shall use w^* to denote each period's per-bidder expected social surplus; i.e., $w^* := W^*/n$.

When we mention the "CGK auction," we mean the following auction: all participating bidders submit sealed bids, the good is transferred to the highest bidder, and each participating bidder i pays $b_i - \frac{1}{|B|-1}\sum_{i\neq j\in B} b_j$, where B is the set of participating bidders, and the auctioneer discloses everything to every bidder after the auction. Notice that, in the CGK auction, bidder's total payment to the auctioneer is always zero. On top of this property, Cramton, Gibbons, and Klemperer (1987) prove that the CGK auction also has an efficient equilibrium. In this efficient equilibrium, even a bidder with the lowest type (i.e., a bidder with valuation 0) will receive strictly positive expected payoff. Let $\underline{w} > 0$ denote this strictly positive expected payoff a bidder with the lowest type will receive in the CGK auction's efficient equilibrium. Define $\underline{W} := n\underline{w}$.

For any probability distribution function F, we always have $W^* > \underline{W}$. For example, if $n = 2$, and F is the uniform distribution on $[0, 1]$, then $W^* = \frac{2}{3}$, and $\underline{W} = \frac{1}{3}$.

When we mention the null auction, we mean the seller simply refuses to transact; i.e., $\forall i$, $S_i = \{0_i\}$.

3 The Main Result

We arrive at the conclusion that an (almost) optimal auction takes the form of a public good provision game in two steps. The first step is to observe that there is an upper bound on the auctioneer's payoff. The second step is to verify that an auction that takes the form of a public good provision game almost achieves that upper bound.

Consider the following auction, which takes the form of a public good provision game:[3]

- Bidders simultaneously decide whether or not to contribute to a public good (say, building a temporary auction house).
- If every bidder contributes, the auctioneer runs the CGK auction.
- If at least one bidder refuses to contribute, the auctioneer runs the null auction.
- The amount of contribution is $q^* := (1-\delta)\underline{w} + \delta w^*$.

[3]Note the slight abuse of terminology below: both the CGK and the null "auctions" are actually parts of the auction we are constructing, and hence strictly speaking do not by themselves qualify as real auctions (as defined in Section 2). However, we believe there should be no confusion.

- The auctioneer discloses everything to every bidder after the auction.

Once the auctioneer has chosen this auction game in period 0, bidders will find themselves playing a repeated game. A typical approach to analyze such a repeated game would be to use the APS machinery (Abreu, Pearce, and Stacchetti (1990)) to characterize the set of equilibrium payoffs, and then identify the one that maximizes bidders' total payoff. However, since this repeated game is deliberately designed to have certain properties, we can actually derive the bidder's optimal sequential equilibrium directly without using of the APS machinery.

Let's ignore incentives for a moment and ask what the most efficient outcome (from bidders' point of view) is. We shall then show that this efficient outcome can be supported as a sequential equilibrium outcome. Finally, we shall show that, in that bidder-optimal equilibrium, the auctioneer's payoff is

$$(1-\delta)\underline{W} + \delta W^*.$$

Since W^* is an upper bound of the auctioneer's payoff, the auction described above is hence almost optimal when δ is close to 1.

Ignore incentives for a moment. Suppose a social planner can mandate some behavorial rule for each bidder, and aims at maximizing the sum of bidders' payoffs. What should such a social planner mandate each bidder to do? Since the auction game will remain the same in the next period regardless of what the planner mandates the bidders to do in the current period, the planner's problem can be collapsed into a single-period problem. If the CGK auction is ever played, the social planner should tell each bidder to use the same strictly increasing bidding function. This would guarantee that the winner of the CGK auction is the highest-valuation bidder, and hence maximize bidders' total payoff. The social planner still has to decide when the bidders should contribute, which amounts to trading-off the benefit of running the CGK auction against the cost of paying the contribution. Since each bidder's contribution decision can only be measurable with respect to his own valuation but not the others', the optimal contribution decision should take the form of a threshold rule, such that a bidder contributes if and only if his valuation is above a certain threshold.

Let $\mathbf{a} := (a_1, \ldots, a_n)$ be a vector of thresholds, and

$$W(\mathbf{a}) := \sum_{i\in N}(1-F(a_i))(-q^*) + \int_{v_1=a_1}^{1} \cdots \int_{v_n=a_n}^{1} \max\{v_1, \ldots, v_n\} \mathrm{d}F(v_n)\cdots \mathrm{d}F(v_1)$$

is the corresponding sum of bidders' payoffs. We shall first prove that the optimal thresholds must be symmetric (i.e., $a_1 = \cdots = a_n$). Suppose $\mathbf{a}$ is not symmetric. Without loss of generality assume $a_1 < a_2$. Let $b \in (a_1, a_2)$ be such that

$$2F(b) = F(a_1) + F(a_2),$$

and define $\mathbf{a}' = (b, b, a_3, \ldots, a_n)$. It suffices to prove that $W(\mathbf{a}') > W(\mathbf{a})$. Let

$$\phi(v_1, v_2) := \int_{v_3=a_3}^{1} \cdots \int_{v_n=a_n}^{1} \max\{v_1, v_2, v_3, \ldots, v_n\} \mathrm{d}F(v_n) \cdots \mathrm{d}F(v_3),$$

and notice that $\phi(\cdot,\cdot)$ is a symmetric function. Then, using the definition of b, we have

$$\begin{aligned}
W(\mathbf{a}') &= \sum_{i\in N}(1-F(a_i))(-q^*) + \int_{v_1=b}^{1}\int_{v_2=b}^{1} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&= \sum_{i\in N}(1-F(a_i))(-q^*) \\
&\quad + \int_{v_1=b}^{1}\int_{v_2=a_2}^{1} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) + \int_{v_1=a_1}^{b}\int_{v_2=a_2}^{1} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&\quad + \int_{v_1=b}^{1}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) - \int_{v_1=a_1}^{b}\int_{v_2=a_2}^{1} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&= W(\mathbf{a}) + \int_{v_1=b}^{1}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) - \int_{v_1=a_1}^{b}\int_{v_2=a_2}^{1} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&= W(\mathbf{a}) + \int_{v_1=b}^{a_2}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&\quad + \int_{v_1=a_2}^{1}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) - \int_{v_1=a_1}^{b}\int_{v_2=a_2}^{1} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&> W(\mathbf{a}) + \int_{v_1=b}^{a_2}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&\quad + \int_{v_1=a_2}^{1}\int_{v_2=b}^{a_2} \phi(v_1,b)\mathrm{d}F(v_2)\mathrm{d}F(v_1) - \int_{v_1=a_1}^{b}\int_{v_2=a_2}^{1} \phi(b,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&= W(\mathbf{a}) + \int_{v_1=b}^{a_2}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&\quad + \int_{v_1=a_2}^{1} \phi(v_1,b)(F(a_2)-F(b))\mathrm{d}F(v_1) - \int_{v_2=a_2}^{1} \phi(b,v_2)(F(b)-F(a_1))\mathrm{d}F(v_2) \\
&= W(\mathbf{a}) + \int_{v_1=b}^{a_2}\int_{v_2=b}^{a_2} \phi(v_1,v_2)\mathrm{d}F(v_2)\mathrm{d}F(v_1) \\
&> W(\mathbf{a}),
\end{aligned}$$

which verifies our claim.

From now on we can restrict our attention to symmetric thresholds, and we shall slightly abuse notation and write $W(a)$ instead of $W(a, \ldots, a)$. We want to prove that $W(a)$ is maximized at $a = 0$.

For any $a \in [0,1]$, let

$$Z_n(a) := \int_{v_1=a}^{1} \cdots \int_{v_n=a}^{1} \max\{v_1, \ldots, v_n\} \mathrm{d}F(v_n) \cdots \mathrm{d}F(v_1) \geq 0,$$

and notice that $Z_n(0) = W^*$. Define $Z_{n-1}(a)$ similarly. Also notice that $\forall a \in (0,1)$,

$$Z_n'(a) = -f(a)nZ_{n-1}(a) < 0,$$

so $Z_n(a)$ is monotone decreasing in a. Similarly $Z_{n-1}(a)$ is monotone decreasing in a as well.[4]

We can now rewrite $W(a)$ as

$$W(a) = Z_n(a) - nq^*(1 - F(a)),$$

and notice that $\forall a \in (0,1)$,

$$W'(a) = f(a)[nq^* - nZ_{n-1}(a)].$$

Since $Z_{n-1}(a)$ is monotone decreasing in a, $[nq^* - nZ_{n-1}(a)]$ is monotone increasing in a, and hence $W(a)$ is quasi-convex. $W(a)$, being a quasi-convex function, is maximized either at $a = 0$ or $a = 1$. Since

$$W(0) = W^* - nq^* = (1 - \delta)(W^* - \underline{W}) > 0 = W(1),$$

$W(a)$ is maximized at $a = 0$.

In summary, the social planner should tell the bidders to always contribute, and then play the efficient equilibrium in the CGK auction.

Can this planner's ideal outcome be supported as a sequential equilibrium outcome? The answer is affirmative, and the proof is constructive. Consider the following strategy:

1. In the first period, contribute regardless of the period-1 valuation, and if the CGK auction is run, play the efficient equilibrium.
2. In any subsequent period,
   (a) if there exists at least one bidder who once refused to contribute in the past, do not contribute regardless of the current-period valuation;
   (b) otherwise, contribute regardless of the current-period valuation, and if the CGK auction is run, play the efficient equilibrium.

We claim that it is a sequential equilibrium for every bidder to follow the above strategy. First notice that it is a sequential equilibrium for every bidder not to contribute in any

[4] This is true even when $n = 2$.

period regardless of his valuation. So the strategy described above is essentially a grim-trigger strategy. Then, by the one-stage deviation principle, it suffices to check the following two kinds of one-stage deviations:

1. *A bidder cannot benefit from not contributing when he should.* This follows from the facts that (i) if he does not contribute, his current-period and continuation payoffs will both be zero, whereas (ii) if he contributes, his current-period payoff will be at least $(1-\delta)(\underline{w}-q^*)$, his continuation payoff will be $\delta(w^*-q^*)$, and these two sum up to at least $((1-\delta)\underline{w}+\delta w^*)-q^*=0$.

2. *When a CGK auction is run, a bidder cannot benefit from not playing the efficient-equilibrium bidding strategy.* This follows from the facts that (i) his opponents' future behavior will not depend on how he bids in the current-period CGK auction, and (ii) the efficient-equilibrium bidding strategy is the best response to his opponent's current-period bidding strategies.

When bidders play the above mentioned strategies, they contribute every period. So the auctioneer's payoff is $(1-\delta)\underline{W}+\delta W^*$, which is arbitrarily close to the upper bound W^* when δ is arbitrarily close to 1 (i.e., when bidders are patient). This concludes that the auction described above is almost optimal when bidders are patient.[5]

Theorem 1 *When there are two or more bidders, and when bidders are patient, an almost optimal auction takes the form of a public good provision game (as described above).*

It is illuminating to compare the above auction with Myerson's (1981) auction. In the extreme case when $\delta=0$, bidders are completely impatient and hence cannot tacitly collude at all. In that case, the optimal auction is Myerson's auction, which boils down to the second price auction with reserve price. Note that in Myerson's auction, when a bidder bids higher, every other bidder gets a lower payoff, both because the winner may need to pay a higher second-price, and because every other bidder has a lower chance of winning the object.

However, as δ increases from 0 to 1, and bidder's ability to tacitly collude increases, Myerson's auction finally under-performs relative to the above auction. Even if Myerson's auction employs a disclosure rule that discloses *nothing* to every bidders after every round, the auctioneer's payoff in Myerson's auction is uniformly bounded away from W^* regardless of δ. This bound can be computed as follows. First notice that honest bidding is a stage-game equilibrium for the Myerson auction, and hence bidding honestly every period is a sequential

[5]This is a good place to go back to our earlier claim that our result does not depend on the assumption that the auctioneer shares the same discount factor with the bidders. (See Footnote 2.) Suppose the auctioneer has a different discount factor that is far away from 1. Suppose intertemporal trade contracts are not enforceable and hence we do not need to worry about the infinite surplus that can potentially be generated from players having different discount factors. If the auctioneer uses the public-good-provision auction, she will still collect the same contributions every period, and hence her normalized discounted sum of expected payoff–normalized and discounted by her own far-away-from-1 discount factor–will still be $(1-\delta)\underline{W}+\delta W^*$ and is arbitrarily close to the upper bound of W^* as bidders' discount factor δ goes to 1.

equilibrium in the repeated game where bidders play the Myerson auction repeatedly. Let W be the sum of bidders' payoffs in this sequential equilibrium, and by construction W does not depend on δ. Now, the bidder-optimal sequential equilibrium (which depends on δ) may give bidders an even higher total payoff: $W(\delta) \geq W$. Therefore, for any δ, the auctioneer's payoff cannot be bigger than $W^* - W(\delta)$, which in turn is no bigger than $W^* - W$. So $W^* - W$ is a uniform bound for the auctioneer's payoff regardless of δ.

4 Discussions

4.1 Renegotiation-Proofness

One apparent problem with the public-good-provision auction is that the bidder-optimal sequential equilibrium is not renegotiation-proof. Once a bidder deviates and refuses to contribute in a particular period, every bidder will receive zero payoff in the continuation equilibrium. This gives rise to the possibility of renegotiation among bidders after such a deviation. However, this problem can be fixed with some modification to the public-good-provision auction when $n \geq 3$. The idea is to allow any $(n-1)$-coalition to punish the remaining bidder by excluding him from the auction. If we can make the total expected payoff collected by this $(n-1)$-coalition the same as the total expected payoff collected by the grand coalition along the equilibrium path, then we can say that the planner's ideal outcome (i.e., to have all bidders contribute in every period regardless of their valuations) can be supported as a renegotiation-proof sequential equilibrium outcome.

Let z^* be the per-bidder expected payoff in an $(n-1)$-CGK auction (i.e., a CGK auction among a group of $(n-1)$ bidders), and $\underline{z}$ be the expected payoff a bidder with the lowest type will receive in an $(n-1)$-CGK auction. Recall that in the public-good-provision auction, bidders' total expected payoff along the equilibrium path is $n(w^* - q^*)$. Let $\hat{q}$ be implicitly defined by

$$(n-1)(z^* - \hat{q}) = n(w^* - q^*).$$

Roughly speaking, $\hat{q}$ would be the amount we should ask each member of a punishing coalition to contribute in order to maintain renegotiation-proofness. But $\hat{q}$ may be too large, and a member of a punishing coalition may not be willing to contribute this much when his valuation is low. So we need to adjust it a little bit. Let $\tilde{q}$ be the amount of per-bidder contribution if we were to apply the public-good-provision auction to a group of $(n-1)$ bidders; i.e., $\tilde{q} := (1-\delta)\underline{z} + \delta z^*$. Let $\hat{\Delta} := \min\{0, \hat{q} - \tilde{q}\}$. Now the adjusted amount of contribution, $\hat{q} - \hat{\Delta}$, would be small enough for any member of a punishing coalition to swallow. To maintain renegotiation-proofness, we need to adjust q^* as well. Let Δ^* be implicitly defined by

$$(n-1)[z^* - (\hat{q} - \hat{\Delta})] = n[w^* - (q^* - \Delta^*)].$$

Now we can describe our modified auction:

- Bidders simultaneously decide whether or not to contribute to a public good.
- If bidder i chooses to contribute, he has to report a pair $(b_i, x_i) \in \mathbb{R}_+ \times N$; i.e., he has to report a bid, and name a fellow bidder to punish.
- If at least two bidders refuse to contribute, then every contributing bidder pays a contribution of $(q^* - \triangle^*)$, and the auctioneer runs the null auction.
- If exactly one bidder (say j) refuses to contribute,
  - if $\forall i \neq j$, $x_i = j$, then every contributing bidder pays a contribution of $(\hat{q} - \hat{\triangle})$, and the auctioneer runs the $(n-1)$-CGK auction among the contributing bidders;
  - otherwise, every particpating bidders pays a contribution of $(q^* - \triangle^*)$, and the auctioneer runs the null auction.
- If every bidders contributes,
  - if there exists j such that $\forall i \neq j$, $x_i = j$, then every bidder (except j) pays a contribution of $(\hat{q} - \hat{\triangle})$, and the auctioneer runs the $(n-1)$-CGK auction among the $(n-1)$ bidders other than j (i.e., the auctioneer in effect treats j as not contributing);
  - if $\forall i$, $x_i = i$, then every bidder pays a contribution of $(q^* - \triangle^*)$, and the auctioneer runs the CGK auction;
  - otherwise, every bidder pays a contribution of $(q^* - \triangle^*)$, and the auctioneer runs the null auction.
- The auctioneer discloses everything to every bidder after the auction.

Using exactly the same argument as in Section 3, and using the definitions of $(\hat{q} - \hat{\triangle})$ and $(q^* - \triangle^*)$, we can continue to argue that one of the social planner's ideal outcomes is to have all bidders contribute in every period regardless of their valuations. A renegotiation-proof sequential equilibrium that supports this planner's ideal outcome is as follows:

- In any period that belongs to the normal state (to be specified below), every bidder i contributes regardless of his valuation, and reports $x_i = i$.
- In any period that belongs to the i-punishment state (to be specified below), i does not contribute, and every bidder $j \neq i$ contributes regardless of his valuation, and reports $x_j = i$.
- The first period belongs to the normal state.
- If in any period there exists i who deviates, then the next period belongs to the i-punishment state.

The proof that the above strategy profile is indeed a sequential equilibrium is largely the same as that in Section 3. On the equilibrium path, bidders contribute in every period regardless of their valuations. Ex post efficiency is achieved. Bidders' total expected payoff is at most $(n-1)(z^* - \tilde{q}) = (1-\delta)(n-1)(z^* - \underline{z}) \approx 0$. So the auctioneer's payoff almost achieves the upper bound of W^*, and hence the above auction is also almost optimal.

4.2 Entry and Exit

Imagine that there is a large number of potential bidders in any period, but only n of them are serious. The identities of these serious bidders may change over time, and individual bidders need not know the identities of these serious bidders in any particular period. In this situation, the auction described in Section 3 can be modified such that as long as the auctioneer manages to collect n contributions, she will run the CGK auction among those who contribute. The modified sequential equilibrium is that each bidder contributes if and only if (i) he is serious, and (ii) there was a CGK auction in the last period.

Alternatively, imagine that there is a pool of n bidders. But every once a while some bidders will die and being replaced by new comers. If bidders have constant probability of dying, then the above modification will also work.

5 When the Auctioneer is one of the Players

The model in Section 2 fits better the story of an auction house who faces more or less the same set of serious bidders over time but cannot afford to change its auction mechanism frequently. As an extension to our main result, we study an alternative model in this section, where the auctioneer can change her auction mechanism every period, but cannot commit to which auction mechanisms she will use in the future. In this alternatively model, the implication of our way to model tacit collusion is very strong. In particular, in the sequential equilibrium that maximizes bidders' total ex ante payoff, bidders jointly capture all the surplus. Note that in deriving this strong result, we allow the auctioneer to use secretive selling mechanisms (i.e., we allow the auctioneer to approach one of the bidders in secret and to promise this bidder not to disclose any trade details to any third party).

This alternative model in effect turns the whole game into a repeated game, with the auctioneer one of the players. In every stage game of this repeated game, the auctioneer moves first, announces an auction, and then bidders play this auction game. This repeated game has multiple sequential equilibria. Once again, we follow our earlier convention and focus our attention to the bidder-optimal sequential equilibrium.

The proof that, if the auctioneer is herself a player of a bigger repeated game, her profit would be driven down to zero in the bidder-optimal sequential equilibrium is contructive. Consider the following strategy profile:

- In the first period, the auctioneer announces the CGK auction.

- In any period, *unless* it is in the punishment phase (to be defined below), the auctioneer announces the CGK auction.
- In any period, *if* it is in the punishment phase, the auctioneer announces the null auction.
- In any period, every bidder refuses to participate if the auctioneer announces any auction different from the CGK auction, and plays the CGK auction's efficient equilibrium if the auctioneer announces the CGK auction.
- The transitional dynamics between the normal state and the punishment phase is as follows:
  1. The first period belongs to the normal state.
  2. In any normal-state period,
     (a) if the auctionner announces the CGK auction, the next period will belongs to the normal state;
     (b) if the auctioneer announces an auction different from the CGK auction,
        i. if no bidders participate, then the next period belongs to the normal state;
        ii. if two or more bidders participate, then the next period belongs to the normal state;
        iii. if exactly one bidder participates, then the next period becomes the first period of an m-period-long punishment phase, where m is any large but finite integer that satisfies

$$\delta w^* \geq (1-\delta) + \delta^{m+1} w^*. \tag{1}$$

  3. In any period during a punishment phase,
     (a) if the auctioneer announces the CGK auction, then the next period belongs to the normal state;
     (b) if the auctioneer announces an auction different from either the CGK auction or the null auction,
        i. if no bidders participate, then the next period belongs to the normal state;
        ii. if two or more bidders participate, then the next period belongs to the normal state;
        iii. if exactly one bidder participates, then the next period becomes the first period of an m-period-long punishment phase.
     (c) if the auctioneer announces the null auction, and if it is the m'-th consecutive period in the current punishment phase, where $m' < m$, then the next period remains in the punishment phase;
     (d) if the auctioneer announces the null auction, and if it is the m'-th consecutive period in the current punishment phase, where $m' \geq m$, then the next period belongs to the normal state.

Note that inequality (1) will be satisfied by some m if bidders are patient enough; i.e., if δ is large enough such that $\delta w^* > 1 - \delta$.

Also note that bidders' strategies do not depend on whether the current period belongs to the normal state or the punishment phase, and hence bidders do not need to know what other bidders have done in the past in order to follow these strategies. In particular, bidders can follow the prescribed strategies regardless of the disclosure policy employed by the auctioneer.

If players follow this strategy profile, then the auctioneer will announce the CGK auction every period and get zero payoff, whereas each bidder will get an expected payoff of w^*. To see that this strategy profile is indeed a sequential equilibrium, all we need is to invoke the one-stage deviation principle and check all possible one-stage deviations.

1. If the auctioneer announces the CGK auction, bidder i knows that (i) his fellow bidders will play the CGK auction's efficient equilibrium strategies, and (ii) the next period will belong to the normal state. So it does not pay to deviate from also playing the CGK auction's efficient equilibrium strategy.

2. If the auctioneer announces the null auction, bidder i has no ways to deviate anyway.

3. If the auctioneer announces an auction different from either the CGK auction or the null auction, bidder i knows that his fellow bidders will not participate. If he does not participate, the next period will belong to the normal state, and hence his continuation payoff will be δw^*. If he participates, his current-period payoff is at most 1, but he has to endure an m-period-long punishment phase. So his (normalized) payoff is at most $(1-\delta)1 + \delta^{m+1} w^*$, which by inequality (1) is no bigger than δw^*. So it does not pay to deviate from not participating.

4. In any period, the auctioneer is indifferent between announcing the CGK auction and announcing the null auction. Deviating from announcing either of these will not make any difference as bidders will not participate anyway.

This exhausts all possible one-stage deviations and verifies that the above strategy profile is indeed a sequential equilibrium.

6 Concluding Remark

In this paper we have shown that if an auctioneer anticipates tacit collusion (modelled as bidders interacting repeatedly and play their best equilibrium), then an optimal auction would take the form of a public good provision game. However, this public-good-provision auction is performing well only when bidders are patient enough. This prompts the questions of how the optimal auction would look like when δ is in the intermediate range. This optimal auction design problem is highly non-trivial because the auctioneer may want to choose an auction that does not disclose information regarding bids, identity of winners, etc., to the bidders at the end of each auction game. Such a non-disclosure policy has natural appeal to the auctioneer because it can potentially inhibit collusion (as deviations will be more difficult to detect). So solving the optimal auction design problem inevitably requires analyzing how well bidders can do when they play auctions with this kind of non- or partial-disclosure policies repeatedly, which in turn requires analyzing repeated games with private monitoring– an area of repeated games theory that we still do not know much about.[6] In this paper, we circumvented the problem of private monitoring because the public-good-provision auction always has the auctioneer discloses everything to every bidder after every auction. This did not hurt our case because we looked at cases where δ is large. But the same is unlikely to be true once we move beyond cases of large δ.

References

[1] Dilip Abreu, David Pearce, and Ennio Stacchetti (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* **58**(5): 1041-1063.

[2] Masaki Aoyagi (2002), "Bid Rotation and Collusion in Repeated Auctions," mimeo, Osaka University.

[3] Susan Athey and Kyle Bagwell (2001), "Optimal Collusion with Private Information," *Rand Journal of Economics* **32**(3): 428-465.

[4] Susan Athey, Kyle Bagwell, and Chris Sanchirico (2000), "Collusion and Price Rigidity," MIT Department of Economics Working Paper #98-23.

[5] Andreas Blume and Paul Heidhues (2001), "Tacit Collusion in Repeated Auctions," mimeo, University of Pittsburgh.

[6] Yeon-Koo Che and Seung-Weon Yoo (2001), "Optimal Incentives in Teams," *American Economic Review* **91**(3): 525-541.

[7] Peter Cramton (1995), "Money Out of Thin Air: The Nationwide Narrowband PCS Auction," *Journal of Economics and Management Strategy* **4**: 267-343.

[8] Peter Cramton, Robert Gibbons, and Paul Klemperer (1987), "Dissolving a Partnership Efficiently," *Econometrica* **55**(3): 615-632.

[9] Jeffrey C. Ely and Juuso Välimäki (2002), "A Robust Folk Theorem for the Prisoner's Dilemma," *Journal of Economic Theory* **102**:84-105.

[10] Daniel A. Graham and Robert C. Marshall (1987), "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions," *Journal of Political Economy* **95**(6): 1217-1239.

[6] See, for example, Ely and Välimäki (2002) and Matsushima (2001) for the state-of-the-art in this area.

[11] Jean-Jacques Laffont and David Martimort (1997), "Collusion Under Asymmetric Information," *Econometrica* **65**(4): 875-911.

[12] Jean-Jacques Laffont and David Martimort (2000), "Mechanism Design with Collusion and Correlation," *Econometrica* **68**(2): 309-342.

[13] Robert C. Marshall and Leslie M. Marx (2002), "Bidder Collusion," mimeo, University of Rochester.

[14] Hitoshi Matsushima (2001), "The Folk Theorem with Private Monitoring," mimeo, University of Tokyo.

[15] R. Preston McAfee and John McMillan (1992), "Bidding Rings," *American Economic Review* **82**(3): 579-599.

[16] Roger Myerson (1981), "Optimal Auction Design," *Mathematics of Operations Research* **6**(1): 58-73.

[17] Andrzej Skrzypacz and Hugo Hopenhayn (2001), "Tacit Collusion in Repeated Auctions," Stanford GSB Research Paper #1698.