

Optimal Repeated Auction with Tacit Collusion*

Atila Abdulkadirođlu Kim-Sau Chung
Columbia University Northwestern University
aa2061@columbia.edu sau@northwestern.edu

May 6, 2002

Abstract

We study the problem of optimal auction design when bidders collude tacitly. More specifically, suppose the auctioneer can only have one single chance to choose an auction game, and after then bidders will play the same auction game again and again, with their valuations drawn anew every period. Foreseeing that bidders can coordinate on their favorite equilibrium in any induced repeated game, which auction game should the auctioneer choose in the first place in order to maximize her revenue? We show that the one-buyer case differs dramatically from the multi-bidder case. In the one-buyer case, the auctioneer cannot do better than if the auction game is not repeated. In the multi-bidder case, the auctioneer can capture almost all the expected gain of trade by exploiting the fact that these multiple bidders can collude tacitly.

*We are very grateful to Yeon-Koo Che, Jeffrey Ely, and Larry Samuelson, whose many probing questions have forced us to re-think our original intuition. All errors are ours.

1 Introduction

In auction settings, tacit collusion happens when, by definition, bidders collude without speaking. How the auctioneer manages to forbid bidders from talking to each other before the auction 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 auctioneer (the FCC in this example) somehow managed to forbid bidders from communicating *before* the auction.

Of course, the lack of communication before the auction does not rule out all possibility of collusion during the auction, and this leads us to the study of tacit collusion. It is natural to conjecture that tacit collusion has the most bite when bidders interact repeatedly. Repeated-game logic suggests that repeated interaction enables bidders to sustain many otherwise non-equilibrium bidding strategies, making bidders more able to collude even though they cannot communicate explicitly.

This reasoning is of course incomplete. When we move from the one-shot auction setting to the repeated-interaction setting, a lot of details in the new setting need to be specified. Hence the above conjecture at best requires much qualification. More strikingly, this paper will demonstrate that, in some natural models of repeated interaction, tacit collusion is not a curse, but exploitable instead.

More specifically, we study the following model of repeated interaction. Suppose the auctioneer can only have one single chance to choose an auction game, and after then bidders will play the same auction game again and again, with their valuations drawn anew every period. Foreseeing that bidders can tacitly collude and coordinate on their favorite equilibrium in any induced repeated game, which auction game should the auctioneer choose in the first place in order to maximize her revenue? We show that the one-buyer case differs dramatically from the multi-bidder case. In the one-buyer case, the auctioneer cannot do better than if the auction game is not repeated. In the multi-bidder case, the auctioneer can capture almost all the expected gain of trade by exploiting the fact that these multiple bidders can collude tacitly.

We assume that bidders' valuations are drawn anew every period (i.e., IID-repetition) partly because it gives us the most dramatic result, and partly because it is exactly the model almost all previous authors have in mind when they think about repeated interaction among bidders.¹ The interaction between IID-repetition and tacit collusion can be summarized as follows: IID-repetition (by pooling individual rationality constraints) generates the opportunity for the auctioneer to extract full surplus, but the auctioneer would not be able to take advantage of this opportunity until she could exploit tacit collusion. The contrast between the one-buyer and multi-bidder cases nicely illustrates this last point.²

¹See, among others, Skrzypacz and Hopenhayn (2001) and Aoyagi (2002).

²Another way to interpret the result is that the one-buyer case is reminiscent to collusion with perfect

A paper closely related to ours is Che and Yoo (2001; hereafter CY), who design an optimal multi-agent incentive contract in a similar repeated-interaction setting. Although we differ in the kind of problem we look at (CY’s problem is a moral hazard one, whereas ours is an adverse selection one), both of our designs share the same property of creating *positive* externalities among agents. Notice that in both of our problems, when there is no repeated interaction, the optimal design is to create *negative* externalities among agents (“relative performance evaluation” in CY’s problem, “Myerson auction” in our problem). But when there is repeated interaction (and hence tacit collusion), playing agents against each other becomes ineffective. Instead, the optimal design should try to exploit tacit collusion by creating *positive* externalities among agents. CY’s trick of creating positive externalities is to use “joint performance evaluation” in place of “relative performance evaluation,” whereas our trick is to use a public good provision game in place of more traditional auction games. Although the tricks are different, the idea is the same. We believe that this same idea would have other applications as well.

The paper is organized as follows. Section 2 introduces the model. Section 3 contains the one-buyer benchmark. The main result is formally stated and proved in Section 4. Section 5 discusses several extensions.

Our model of repeated interaction is not the only way to model tacit collusion. In the appendix we outline two other possible models. In both of these alternative models, the auctioneer can choose different auctions in different periods. In Appendix A, we look at an alternative model where the auctioneer can sign long term contracts with the bidders. Not surprisingly, full surplus extraction is immediate in such a model. In Appendix B, we look at another alternative model where the auctioneer has no commitment power at all, and hence she herself is one of the players in a larger repeated game. The result goes to the other extreme: the auctioneer earns zero profit in the bidders’ favorite sequential equilibrium. After the first draft of this paper was finished, we learned that Kremer and Skrzypacz (2002) have independently worked on these two alternative models as well. So we encourage the reader to also consult their paper.

1.1 Related Literature of Repeated Auctions

[to be completed]

2 The Model

There are 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

communication which, in contrast to tacit collusion, is the kind of collusion the auctioneer cannot harness.

distribution function F , which has a strictly positive density function f over $[0, 1]$.

There are infinitely many periods (i.e., $t = 1, 2, \dots$). 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 (ii) the repeated game which has infinite horizon. Formally, an auction game is a selling mechanism (S, p, q, M, a) , 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;³
- $p : S \rightarrow \{(x_1, \dots, x_n) \in [0, 1]^n \mid \sum_{i \in N} x_i \leq 1\}$ is an allocation function that specifies the probabilities with which each bidder will get the object, with the restriction that $\forall i, p_i(\dots, 0_i, \dots) \equiv 0$ (non-participating bidders never get the object);
- $q : S \rightarrow \mathbb{R}^n$ is a payment function that specifies the amount each bidder has to pay the auctioneer, with the restriction that $\forall i, q_i(\dots, 0_i, \dots) \equiv 0$ (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
- $a : S \rightarrow \Delta(M)$ 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, a_i (the projection of a 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.⁴

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 auction game, his period- t payoff will be zero.

³It is not clear that we can appeal to the Revelation Principle and restrict our attention to direct revelation mechanisms (i.e., $S_i = [0, 1]$). As is well known in the implementation literature, if agents can defy any wish of the mechanism designer and instead choose to play their own favorite equilibrium, it is then with loss of generality for the mechanism designer to restrict her attention to direct revelation mechanisms.

⁴More 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.

All players have common discount factor δ .⁵ 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} , $t = 1, 2, \dots$, then his normalized discounted 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 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., $W^* := \mathbf{E} \max\{v_{1t}, \dots, v_{nt}\} = \int_0^1 v dF(v)^n$ (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 "Myerson auction," we mean the direct revelation mechanism that implements Myerson's (1981) optimal auction, and the auctioneer discloses everything to every bidder after the auction. Although the Myerson auction is well understood in the literature, it may still be worthwhile to remind the reader two of its major properties. First, the Myerson auction is strategyproof, meaning that it is a dominant strategy for every bidder to honestly report his true valuation. Second, although truth-telling is not the only dominant strategy, all dominant strategy equilibria are outcome-equivalent.⁶ When a bidder's true valuation is v , we shall use $D(v)$ to denote the set of all his dominant-strategy reports. Strategyproofness means that $v \in D(v)$. We say a bidder with true valuation v bids sincerely if his report b is in $D(v)$. Let π^* denote the auctioneer's expected revenue in a one-shot Myerson auction when bidders bid sincerely.

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 , even a bidder with the lowest type (i.e., a bidder with valua-

⁵That 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 the auctioneer can turn bidders' ability to collude into her advantage and capture almost all the expected gain of trade. See also Footnote 10.

⁶For example, if the probability that a bidder gets the object is $1/10$ whenever his reported valuation is between $1/3$ and $2/3$, then it is a dominant strategy for this bidder to report anything between $1/3$ and $2/3$ (and not necessarily his true valuation) as long as his true valuation is also between $1/3$ and $2/3$. However, this kind of mis-report is immaterial in the Myerson auction.

tion 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 , the three numbers (W^* , π^* , and \underline{W}) are always ranked as follows:

$$W^* > \pi^* > \underline{W}.$$

For example, if $n = 2$, and F is the uniform distribution on $[0, 1]$, then $W^* = \frac{2}{3}$, $\pi^* = \frac{5}{12}$, 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 One-Buyer Benchmark

Fix any auction game (S, p, q, M, a) . In any period t , the buyer's problem is

$$\max_{s \in S} (1 - \delta)(p(s)v_t - q(s)) + \delta(\text{CONTINUATION PAYOFF}).$$

Since neither the distribution of the buyer's future valuations nor the auction games he will face in the future depend on the past history (including the current-period message he is sending to the auctioneer), the second term is a constant. So the buyer reacts to any auction game as if it is not repeated. So the auctioneer's optimal auction design problem is the same as if the auction game is not repeated.

In the static problem, we know that the optimal selling mechanism is for the seller to post the monopoly price. Notice that, when the auctioneer posts the monopoly price, she earns the monopoly profit, which is bounded away from the expected gain of trade. In other words, although IID-repetition (by pooling individual rationality constraints) generates the opportunity for the auctioneer to extract full surplus, the auctioneer is not able to take advantage of this opportunity.

4 The Main Result

We shall now demonstrate how the auctioneer can exploit tacit collusion and capture almost all the expected gain of trade when $n \geq 2$. Consider the following auction:⁷

- Bidders simultaneously decide whether or not to participate.
- If every bidder participates, the auctioneer runs the CGK auction.

⁷Notice 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). We believe there should be no confusion.

- If at least one bidder does not participate, the auctioneer runs the null auction.
- In either case, every participating bidder pays a non-refundable contribution of $q^* := (1 - \delta)\underline{w} + \delta w^*$.
- The auctioneer discloses everything to every bidder after the auction.

Once the auctioneer has chosen this auction in period 0, bidders will find themselves playing a repeated game. We shall first 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 the upper bound of the auctioneer's payoff, the auction described above is hence almost optimal when δ is close to 1.⁸

Ignore incentives for a moment. Suppose a social planner can mandate some behavioral 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 participate, which amounts to trading-off the benefit of running the CGK auction against the cost of paying the non-refundable contribution. Since each bidder's participation decision can only be measurable with respect to his own valuation but not the others', the optimal participation decision should take the form of a threshold rule, such that a bidder participates if and only if his valuation is above a certain threshold.

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

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

is the corresponding sum of bidders' payoffs. We shall first prove that the optimal thresholds must be symmetric (i.e., $a_1 = \dots = a_n$). 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

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

⁸Actually the upper bound W^* is not achievable. As we saw in Appendix A, even when the auctioneer can sign long term contracts with the bidders, at most she can only achieve an expected payoff of $(1 - \delta)\pi^* + \delta W^*$, which is strictly smaller than W^* (but still larger than $(1 - \delta)\underline{W} + \delta W^*$).

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

$$\phi(v_1, v_2) := \int_{v_3=a_3}^1 \cdots \int_{v_n=a_n}^1 \max\{v_1, v_2, v_3, \dots, v_n\} dF(v_n) \cdots dF(v_3),$$

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

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

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 $W(a, \dots, a)$. We want to prove that $W(a)$ is maximized at $a = 0$.

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

$$Z_n(a) := \int_{v_1=a}^1 \cdots \int_{v_n=a}^1 \max\{v_1, \dots, v_n\} dF(v_n) \cdots dF(v_1) \geq 0,$$

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.⁹

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 participate, 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, participate 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 did not participate in the past, do not participate regardless of the current-period valuation;
 - (b) otherwise, participate 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 participate in any period regardless of his valuation, because participating amounts to contributing to a public good.

⁹This is true even when $n = 2$.

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 participating when he should.* This follows from the facts that (i) if he does not participate, his current-period and continuation payoffs will both be zero, whereas (ii) if he participates, 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 participate and pay the contribution every period. So the auctioneer's expected 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.¹⁰

Theorem 1 *When there 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), and the auctioneer captures almost all the expected gain of trade.*

It is illuminating to compare our auction with the Myerson auction. In the extreme case when $\delta = 0$, bidders are completely impatient and hence cannot tacitly collude at all. In that case, the Myerson auction is the optimal auction, whereas our auction reduces to the Vickrey auction. Apparently our auction is not even remotely optimal (not to mention "almost") in this case.

However, as δ increases from 0 to 1, our auction will catch up and finally surpass the Myerson auction. To see this, notice that for any $\epsilon > 0$, our auction will give the auctioneer a payoff within the ϵ -neighborhood of the upper bound W^* for δ big enough. Whereas the auctioneer's payoff in the Myerson 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

¹⁰This is a good place to go back to our earlier claim that it is the bidders' (not the auctioneer's) patience that drives this result. (See Footnote 5.) 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 our 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 total gain of trade 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 .

This argument is actually more general than it looked. Not only that it continues to apply even if we modify our definition of the Myerson auction so that the auctioneer discloses *nothing* to any bidder after the auction, it also applies to any other standard auctions that have a Bayesian Nash equilibrium (regardless of the disclosure policies). For any such an auction and any such a Bayesian Nash equilibrium, it is a sequential equilibrium for bidders to play that equilibrium every period in the corresponding repeated game. Let W be the sum of bidders' payoffs in that sequential equilibrium. Then $W^* - W$ will be a uniform bound for the auctioneer's payoff regardless of δ . Hence any fixed auction will perform worse than our auction when δ is big enough.

This argument does not apply to our auction, because our auction is not a fixed auction. It explicitly incorporates δ , which was originally a primitive parameter of the environment, into the formula of the non-refundable contribution. So its rule actually changes when the environment changes. When the environment is such that bidders are patient, it takes a form that out-performs the Myerson auction as well as any other standard auctions.

5 Discussions

5.1 Renegotiation-Proofness

One apparent problem with our auction is that the bidder-optimal sequential equilibrium is not renegotiation-proof. This problem can be fixed with some modification to our 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 participate 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 our original auction, bidders' total expected payoff along the equilibrium path is $n(w^* - q^*)$. Let \hat{q} be implicitly defined by

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

Now we can describe our modified auction:

- Bidders simultaneously decide whether or not to participate.
- If bidder i participates, 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 participate, then every participating bidder pays a contribution of q^* , and the auctioneer runs the null auction.
- If exactly one bidder (say j) refuses to participate,
 - if $\forall i \neq j, x_i = j$, then every participating bidder pays a contribution of \hat{q} , and the auctioneer runs the $(n - 1)$ -CGK auction among the participating bidders;
 - otherwise, every participating bidder pays a contribution of q^* , and the auctioneer runs the null auction.
- If every bidder participates,
 - if there exists j such that $\forall i \neq j, x_i = j$, then every bidder (except j) pays a contribution of \hat{q} , 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 participating);
 - if $\forall i, x_i = i$, then every bidder pays a contribution of q^* , and the auctioneer runs the CGK auction;
 - otherwise, every bidder pays a contribution of q^* , 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 4, and using the definition of \hat{q} , we can continue to argue that one of the social planner's ideal outcomes is to have all bidders participate 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 participates 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 participate, and every bidder $j \neq i$ participates 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 4. The only thing that remains to check is that, in any i -punishment state, any bidder $j \neq i$ cannot benefit from not participating. Let \tilde{q} be the amount of per-bidder contribution if we were to apply our original auction to a group of $(n-1)$ bidders; i.e., $\tilde{q} := (1-\delta)\underline{z} + \delta z^*$. One can prove that, in our modified auction, the amount of contribution j needs to pay in the i -punishment state is no more than this amount (i.e., $\hat{q} \leq \tilde{q}$), and hence j indeed cannot benefit from not participating.

5.2 Entry and Exit

It seems to us that any auction that explicitly exploits the cooperation (rather than competition) among bidders is bounded to be too inflexible to handle unanticipated entry and exit of bidders. Nevertheless, there exists certain situations where our original auction can be modified to accommodate 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, our original auction 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 participates if and only if (i) he is serious, and (ii) there was a CGK auction in the last period.

5.3 Other Issues

As we mentioned at the end of Section 4, our auction is performing well only when bidders are patient enough. This prompts the questions of, when δ is in the intermediate range, whether or not tacit collusion is exploitable, and if yes, how? 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. 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.¹¹ Notice how we circumvented the problem of private monitoring by always having the auctioneer disclose 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 .

¹¹See, for example, Ely and Välimäki (2002) and Matsushima (2001) for the state-of-the-art in this area.

References

- [1] MASAKI AOYAGI (2002), “Bid Rotation and Collusion in Repeated Auctions,” mimeo, Osaka University.
- [2] SUSAN ATHEY AND KYLE BAGWELL (2001), “Optimal Collusion with Private Information,” *Rand Journal of Economics* **32**(3): 428-465.
- [3] SUSAN ATHEY, KYLE BAGWELL, AND CHRIS SANCHIRICO (2000), “Collusion and Price Rigidity,” MIT Department of Economics Working Paper #98-23.
- [4] YEON-KOO CHE AND SEUNG-WEON YOO (2001), “Optimal Incentives in Teams,” *American Economic Review* **91**(3): 525-541.
- [5] PETER CRAMTON (1995), “Money Out of Thin Air: The Nationwide Narrowband PCS Auction,” *Journal of Economics and Management Strategy* **4**: 267-343.
- [6] PETER CRAMTON, ROBERT GIBBONS, AND PAUL KLEMPERER (1987), “Dissolving a Partnership Efficiently,” *Econometrica* **55**(3): 615-632.
- [7] JEFFREY C. ELY AND JUUSO VÄLIMÄKI (2002), “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory* **102**:84-105.
- [8] 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.
- [9] ILAN KREMER AND ANDRZEJ SKRZYPACZ (2002), “Optimal Selling Rules for Repeated Transactions,” mimeo, Stanford University.
- [10] HITOSHI MATSUSHIMA (2001), “The Folk Theorem with Private Monitoring,” mimeo, University of Tokyo.
- [11] R. PRESTON MCAFEE AND JOHN MCMILLAN (1992), “Bidding Rings,” *American Economic Review* **82**(3): 579-599.
- [12] ROGER MYERSON (1981), “Optimal Auction Design,” *Mathematics of Operations Research* **6**(1): 58-73.
- [13] ANDRZEJ SKRZYPACZ AND HUGO HOPENHAYN (2001), “Tacit Collusion in Repeated Auctions,” Stanford GSB Research Paper #1698.

Appendix A: Alternative Model–Full Commitment

Our model of repeated interaction is not the only way to model tacit collusion. For sake of completeness, we shall document two other relatively natural ways to set up the model in the current and the next appendices.

In this appendix, we shall look at an alternative model where the auctioneer can sign long term contracts with the bidders. Implicitly, such a model requires that the auctioneer has full commitment power over her future choices of auctions contingent on the bidding history. In other words, the auctioneer can commit to a contingent plan specifying which auction game will be announce in which period after which bidding history. Once the auctioneer has committed to any such contingent plan, the bidders will be playing a dynamic game with infinite horizon. This dynamic game may have multiple equilibria. Once again, we focus our attention on the bidder-optimal sequential equilibrium (and if there is more than one we focus our attention on the one that maximizes the auctioneer’s expected payoff). The auctioneer’s problem is: foreseeing that bidders can coordinate on their favorite sequential equilibrium in any dynamic game, which contingent plan should the auctioneer choose?

Not surprisingly, full surplus extraction is immediate in such a model. The reason why the following otherwise standard proof may look longer than one would anticipate is that it needs to handle the possibility of tacit collusion among bidders.

First observe that, if there were only one period, then the Myerson auction would have been optimal. Let π^* denote the expected revenue from a one-shot Myerson auction. Second observe that, by the first time when bidders have a chance to move, they have already learned their period-1 valuations, but not yet their valuations in subsequent periods. Hence the auctioneer has to honor the interim individual rationality constraints in the first period, but not necessarily those in subsequent periods. This suggests that $(1 - \delta)\pi^* + \delta W^*$ is an upper bound of the auctioneer’s expected payoff. Since this upper bound can be achieved by the following contingent plan, the following plan must be optimal for the auctioneer:

- In the first period, announce the following modified version of the Myerson auction:
 1. Bidders simultaneously submit two numbers (b_i, e_i) , where $b_i \in [0, 1]$ is i ’s bid in the Myerson auction, and $e_i \in \{0, 1\}$ is i ’s willingness to pay a contribution of $\frac{\delta}{1-\delta}w^*$.
 2. The object is allocated as in the Myerson auction.
 3. Bidders first pay the Myerson-auction payments. On top of that, if $[\forall i, e_i = 1]$, then every bidder also needs to pay a contribution of $\frac{\delta}{1-\delta}w^*$.
 4. The auctioneer discloses everything to every bidder after the auction.
- If $[\forall i, e_i = 1]$ in the first period, announce the CGK auction from the second period onward;
- If $[\exists i, e_i = 0]$ in the first period, announce the null auction from the second period onward.

In other words, in the first period, the auctioneer announces the Myerson auction plus the following public good provision game:¹² if every bidder agrees to contribute, the auctioneer will give away the object for free from the second period onward; and if at least one bidder refuses to contribute, the auctioneer will refuse to transact from the second period onward. Facing such a contingent plan, bidders in effect find themselves playing a dynamic game. There are multiple equilibria in this dynamic game. But we shall show that in any sequential equilibrium that does not involve dominated strategies, bidders will bid sincerely in the period-1 Myerson auction.

Let σ be any sequential equilibrium of this dynamic game. For each bidder i , the sequential equilibrium σ (together with the distribution function F) induces a joint probability distribution H_i over v_i , b_i , and e_i (i.e., i 's period-1 valuation,¹³ bid, and willingness to contribute) on the equilibrium path. Define H_{-i} and H_N similarly, and notice that these are product measures. Let $V_i \subseteq [0, 1]$ be the set of i 's period-1 valuation such that i contributes with positive probability; i.e.,

$$V_i := \{v_i \in [0, 1] \mid H_i(e_i = 1 \mid v_i) > 0\}.$$

Pick any $v_i \in V_i$, and any b_i in the support of the conditional probability $H_i(b_i \mid v_i, e_i = 1)$. When bidder i has valuation v_i , he must find that submitting the pair $(b_i, e_i = 1)$ is at least as good as submitting the pair $(b_i, e_i = 0)$. The difference between these two pairs is that, if bidder i submits the pair $(b_i, e_i = 1)$, on top of the same Myerson-allocation and Myerson-payment, with probability $H_{-i}(e_{-i} = \mathbf{1})$, he will also pay the contribution of $\frac{\delta}{1-\delta}w^*$ in exchange for the chance of playing the CGK auction from the second period onward. This is profitable only if

$$H_{-i}(e_{-i} = \mathbf{1})w^* \leq \int_{v_{-i}} \int_{b_{-i}} w_i(b_i, b_{-i})H_{-i}(db_{-i} \mid v_{-i}, e_{-i} = \mathbf{1})H_{-i}(e_{-i} = \mathbf{1} \mid v_{-i})H_{-i}(dv_{-i}), \quad (1)$$

where $w_i(b_i, b_{-i})$ is i 's continuation payoff conditional on the bidding history (b_i, b_{-i}) and the event that every bidder contributes in the first period. Multiplying both sides with $H_i(e_i = 1 \mid v_i)$, and integrating over b_i with respect to the conditional probability $H_i(b_i \mid v_i, e_i = 1)$, we can rewrite the above inequality as

$$\begin{aligned} & H_{-i}(e_{-i} = \mathbf{1}) \int_{b_i} w^* H_i(e_i = 1 \mid v_i) H_i(db_i \mid v_i, e_i = 1) \\ & \leq \int_{v_{-i}} \int_{b_N} w_i(b_N) H_N(db_N \mid v_N, e_N = \mathbf{1}) H_N(e_N = \mathbf{1} \mid v_N) H_{-i}(dv_{-i}). \end{aligned}$$

¹²More precisely, it is a game of public good provision *with* refund, and is different from the game of public good provision *without* refund that we will use in the main text of this paper.

¹³We suppress the time subscript for simplicity and write v_i instead of v_{i1} .

Integrating over v_i , and using the fact that $[v_i \notin V_i \implies H_i(e_i = 1|v_i) = 0]$, we can rewrite the above inequality as

$$H_N(e_N = \mathbf{1})w^* \leq \int_{v_N} \int_{b_N} w_i(b_N)H_N(db_N|v_N, e_N = \mathbf{1})H_N(e_N = \mathbf{1}|v_N)H_N(dv_N).$$

Summing over i , and using the fact that

$$\forall b_N, \quad \sum_{i \in N} w_i(b_N) \leq W^* = nw^*,$$

we can rewrite the above inequality as

$$H_N(e_N = \mathbf{1})W^* \leq \int_{v_N} \int_{b_N} W^* H_N(db_N|v_N, e_N = \mathbf{1})H_N(e_N = \mathbf{1}|v_N)H_N(dv_N) = H_N(e_N = \mathbf{1})W^*. \quad (2)$$

Apparently the above weak inequalities hold as equalities. Since inequality (2) comes from integrating and summing a bunch of inequalities (1), we must have for all i , for all $v_i \in V_i$, and for $H_i(b_i|v_i, e_i = 1)$ -almost all b_i , inequality (1) holds as an equality—bidder i is indifferent between submitting the pair $(b_i, e_i = 1)$ and submitting the pair $(b_i, e_i = 0)$ when his valuation is v_i . If $b_i \notin D(v_i)$, then submitted the pair $(b_i, e_i = 0)$ will be dominated, a contradiction. So we have $b_i \in D(v_i)$. The same domination argument holds for any b_i that is in the support of the conditional probability $H_i(b_i|v_i, e_i = 0)$ as well. Hence we have $b_i \in D(v_i)$ for all b_i in the support of $H_i(b_i|v_i)$.

Finally, consider any $v_i \notin V_i$. Using the same domination argument again, we have $b_i \in D(v_i)$ for all b_i in the support of $H_i(b_i|v_i)$. This completes our proof that in any sequential equilibrium that does not involve dominated strategies, bidders bid sincerely in the period-1 Myerson auction.

The above proof also shows that bidders are indifferent among all the sequential equilibria (because the contributions exactly cancel out their expected continuation payoffs on the equilibrium path of any sequential equilibrium) modulo those involve dominated strategies which we ignore for practical reason. So all these sequential equilibria are trivially bidder-optimal. According to our equilibrium selection criteria, when there are more than one bidder-optimal sequential equilibrium, we shall select the one that maximizes the auctioneer's payoff. It is easy to see which sequential equilibrium will maximize the auctioneer's payoff—the one that bidders agrees to pay the contributions regardless of their period-1 valuations:

- every bidder bids sincerely in the period-1 Myerson auction;
- every bidder submits $e_i = 1$ and pays the contribution of $\frac{\delta}{1-\delta}w^*$ regardless of his valuation in the first period; and
- in each of the CGK auctions from the second period onward, bidders play the CGK auction's efficient equilibrium.

It is straightforward to check that this is indeed a sequential equilibrium. In this sequential equilibrium, the auctioneer's expected payoff is

$$\begin{aligned} & (1 - \delta)\left(\pi^* + n\frac{\delta}{1 - \delta}w^*\right) \\ &= (1 - \delta)\pi^* + \delta W^*, \end{aligned}$$

which achieves the upper bound as we claimed earlier.

Appendix B: Alternative Model—No Commitment

In the last appendix we saw that when the auctioneer has full commitment power and can sign long term contracts with bidders, full surplus extraction is immediate. In this appendix we shall document another alternative model of repeated interaction, namely that the auctioneer has no commitment power whatsoever on her future choices of auctions.¹⁴ Such a 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 focus our attention to the bidder-optimal sequential equilibrium. We shall show that the implication of such a model is exactly the opposite of that in the long-term-contract model: when bidders are patient enough, the sum of bidders' payoffs can achieve the upper bound of W^* in the bidder-optimal sequential equilibrium. This means the auctioneer's payoff is driven down to zero.

This and the last appendices trace out the range of results we can get when we change the assumption regarding the auctioneer's commitment power from one extreme to the other. The model studied in the main text can be seen as being somewhere in between these two extremes.

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 constructive. 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 specified 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.¹⁵
- The transitional dynamics between the normal state and the punishment phase is as follows:
 1. The first period belongs to the normal state.

¹⁴This is different from saying that the auctioneer cannot commit to the auction game she announces. We maintain the assumption that, within any given period, the auctioneer's commitment power is complete. In other words, even if the announced auction is not ex post efficient, she can stick to her gun and insist not to transact anymore within that period.

¹⁵Notice 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.

2. In any normal-state period,
 - (a) if the auctioneer announces the CGK auction, the next period will belong 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

$$\delta w^* \geq (1 - \delta) + \delta^{m+1} w^*. \quad (3)$$

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.

Notice that inequality (3) will be satisfied by some m if bidders are patient enough; i.e., if δ is large enough such that $\delta w^* > 1 - \delta$.

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 (3) 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.