Endogenous Bid Rotation in Repeated Auctions

Shiran Rachmilevitch[*]

August 3, 2011

Abstract

I study collusion between two bidders in a general symmetric IPV repeated auction, without communication, side transfers, or public randomization. I construct a collusive scheme, *endogenous bid rotation*, that generates a payoff larger than the bid rotation payoff.

Keywords: Auctions; Bid rotation; Collusion; Repeated games

JEL Codes: D44; D82

1 Introduction

Repeated interaction among the same bidders in a sequence of auctions is a common occurrence. A regular group of buyers attend Sotheby's auctions, the same firms compete for government defense contracts regularly, automobile dealers meet in vehicle auctions time and again, and so on (see Nelson [12] and Philips et al. [13]). Such environments create fertile ground for collusion, because in a repeated game reputation can be built, implicit transfers can be made, and noncooperative players can be punished. Apart from being a prevalent empirical phenomenon, collusion in auctions

---

[*]Department of Economics, University of Haifa.

is an instance of a much broader class of problems, as it exhibits the defining features of any repeated adverse selection game. This class includes repeated oligopoly games, social insurance games, and more. For example, Bertrand competition with independent draws of constant marginal costs is strategically equivalent to a first-price auction with independent and private values (take the negative of the cost to be the "valuation" and the negative of the price to be the "bid").

One of the very first questions that arises naturally is what is the scope of possible collusion? That is, what equilibrium payoffs can be achieved in the repeated game above the competitive payoff, and what are the strategies delivering these payoffs? This question has been addressed under a variety of assumptions (see the literature review in subsection 1.1 below). In this paper I consider the 2-bidder independent private values (IPV) case, under the assumption that there is no public correlation device, and the bidders can neither exchange money nor communicate—all they can do, in any given round, is simply to bid in the stage auction (or refrain from bidding). This is an important environment to study, because communication is often deemed illegal by anti-trust law. Furthermore, even in the absence of legal obstacles, the cartel may prefer to collude tacitly, simply in order to minimize the probability of detection by the seller.

Apart from the infinite repetition of the static (Bayesian Nash) equilibrium, which is obviously feasible, a feasible line of behavior which does not require transfers or communication is the *bid rotation scheme*, which works as follows: in every odd numbered period bidder 1 (say) gets to be the sole participant in the auction and obtain the good for the reserve price, and in even numbered periods bidder 2 is the sole participant. For this scheme to be supported in equilibrium the bidders need to be sufficiently patient, so that the first designated loser will adhere to it. In addition, effective off-equilibrium threats are needed. Such threats are assured to exist if the bid rotation payoff exceeds the static equilibrium payoff, as simple Nash reversion can be applied. In fact, this is the case with virtually all well-behaved distributions: as

shown by McAfee and McMillan [10], the bid rotation payoff exceeds the competitive payoff if the distribution of valuations exhibits an increasing hazard rate. In Lemma 2 below I derive an alternative sufficient condition: it is sufficient that the distribution's expectation is at least as large as one half of the maximal valuation. Consequently, one would like to look for schemes that improve on bid rotation. Several such schemes have been derived in the literature, but they suffer from several drawbacks: the improvements are typically only in terms of ex ante payoff, they vanish as the discount factor tends to one, and some of the schemes involve extremely asymmetric play on the path, where one player receives a very low payoff (see subsection 1.1 for details). I demonstrate the existence of a collusive scheme, *endogenous bid rotation* (EBR, for short), that overcomes all these drawbacks.

Under this scheme, there is a particular value α, such that the period-t loser gets to obtain the good for free in period $t+1$ if his valuation in $t+1$ is above α, and the period-t winner obtains it for free otherwise. This is achieved by endowing the last loser with the privilege of determining the identity of the current winner. When the stage game is a second-price auction, the following strategy implements the scheme: the last winner is instructed to bid zero (independent of his valuation) and the last loser is instructed to bid positively if his valuation is above α and refrain from bidding otherwise. Deviations trigger a perpetual punishment phase. A similar construction is presented for a general auction format. In this way, losing the current auction is accompanied by receiving a higher expected continuation payoff (relatively to winning), because the last loser has the privilege of deciding on the current winner. This intertemporal tradeoff delivers the appropriate incentives, and sustains the scheme in equilibrium.

1.1 Literature Review

It is well-known that in the IPV case the first-best[1] is achievable if the cartel can organize side transfers. This result, which is true even for a one-shot auction, has been shown by Graham and Marshall [7] and Mailath and Zemsky [9] for second-price auctions, and by McAfee and McMillan [10] for symmetric first-price auctions. These results readily imply that with transfers the first-best can also be achieved in a repeated game—simply by repeating the static scheme in every round. By contrast, when transfers are not available, there is a world of difference between the static and the repeated auction. As shown by McAfee and McMillan [10], the best the cartel can do in a static auction is either to chose the winner randomly or to have its members bid competitively. In a repeated auction, on the other hand, the rich strategy spaces open the door for a wide range of collusive behaviors and allow for substantial collusion even without monetary side-payments, because continuation payoffs can function (partially, at least) as implicit transfers. In particular, it follows from the analysis of Fudenberg et al. [6] of repeated games with imperfect public monitoring, that in the IPV case with communication, finitely many valuations, and publicly observable bids, a folk theorem holds.[2] In particular, the first-best can be approximated as closely as one wishes when the discount factor tends to one.

Skrzypacz and Hopenhayn [15] study perfect public equilibria (due to Fudenberg et al. [6]) in repeated auctions where the identity of the winner is the only public signal. They derive schemes that improve on competitive bidding and bid rotation, and that do not require transfers or communication. Related results have been obtained by Blume and Heidhues [5]. In both [5] and [15] "one-shot improvements" are derived, where one takes the bid rotation scheme and introduces a certain change

---

[1]By *first-best* it is meant that the payment to the seller does not exceed the reserve price and the allocation is efficient.

[2]Aoyagi [2] considers a repeated auction model where bidders communicate in every round with the aid of a central mediation device. He identifies conditions under which the first-best can be approximated, even when valuations are not independent.

in the first few periods, which translates to a slight improvement in ex ante payoffs. In particular, there exists a $\bar{t}$ such that for all $t \geq \bar{t}$ the following is true: a bidder's payoff conditional on any t-history is *not* greater than the bid rotation payoff.[3] Consequently, this implies that the degree of the improvement converges to zero as the discount factor tends to one. Moreover, under some of these schemes, it is the case that conditional on some of these t-histories one of the bidders is receiving a lower-than-bid-rotation payoff.

In addition to these "one-shot improvements," Skrzypacz and Hopenhayn [15] construct two schemes that deliver nonmarginal improvements conditional on every history of play: the *exclusion scheme* and the *chips mechanism*. The first is similar to EBR in the sense that there are periods, called the *exclusive periods*, in which one bidder is the designated sole participant. In contrast to EBR, there are also periods, the *normal periods*, in which positive bidding by both bidders takes place. Transitions between the normal and exclusive periods are dictated by a public randomization device. The main shortcomings of this scheme are that it is not known how the bidding functions in the normal periods look like, it does not have a public-randomization-free analog, and it is not known whether this scheme improves on bid rotation when the distribution of valuations is other than uniform.[4]

Under the chips mechanism, each bidder starts the repeated game with T chips, and every time a bidder wins he gives one chip to his opponent. Once a bidder runs out of chips he is instructed to let his opponent win k consecutive auctions, at the end of which he receives one chip back. Skrzypacz and Hopenhayn evaluate this scheme numerically, by optimizing over T and k, for the case of a second-price auction and a uniform distribution. According to their calculations, this scheme extracts all possible collusive gains when the discount factor approaches one (earlier in their paper,

---

[3]Similar improvements can be constructed with respect to the infinite repetition of the static equilibrium.

[4]Skrzypacz and Hopenhayn compute the payoff for the special case of a uniform distribution. As we will see later, EBR does better than the exclusion scheme in the uniform case.

Skrzypacz and Hopenhayn derive an anti-folk theorem for perfect public equilibria, which establishes a bound on the cartel's payoff in any such equilibrium; the payoff of the chips mechanism converges to this bound). The scheme is too complicated for analytical analysis and it is not known what are the improvements it delivers for distributions other than uniform. It is also not known whether it can be sustained in equilibrium for auction formats other than second-price (because the second-price auction has asymmetric equilibria which are necessary in the construction). Finally, as in the case of the exclusion scheme, it is not known how the equilibrium bidding functions look like.

Under either the chips mechanism and the exclusions scheme there are periods in which there is no selection whatsoever: under the former, when one bidder is supposed to let the other win k auctions in a row, the latter obtains the good in those k rounds independent of any information; under the latter, the designated bidder wins independent of any information whenever the scheme is in an exclusive period. Under EBR, by contrast, there is always some information utilized in the selection of the winner. Additionally, under either the chips mechanism and the exclusion scheme there are periods in which nontrivial bidding takes place, and the seller receives positive payments; under EBR, by contrast, the payment to the seller is always zero.

2 Model

There are two symmetric, risk-neutral, expected utility maximizing bidders, bidder 1 and bidder 2, who participate in an infinite sequence of statistically independent IPV (sealed-bid) general *standard auctions* with a zero reserve price. A single indivisible object is sold in every period (there is one object per period). In a standard auction (see Riley and Samuelson [14]), each bidder can bid any nonnegative bid, the winner is a bidder who submitted a maximal bid, his payment to the auctioneer is nonnegative,

and the auction rules are anonymous. First-price, second-price, and all-pay auctions are examples of standard auctions.[5] Valuations are drawn every period for each bidder, and they are independent across periods and across bidders. Each such draw is from the unit interval, according to the atomless, differentiable distribution F, whose density, $f \equiv F'$, satisfies $f > 0$.[6] Whenever bidder i and bidder j are mentioned in the same sentence, it is implicitly assumed that $i \neq j$. In addition to the regular bids (i.e., nonnegative reals), a bidder can also submit a special bid, denoted 0^+, which is identical to 0 except that it wins for sure if the competing bid is 0.[7] There is no communication between the bidders, they cannot organize transfers, and they do not have a public randomization device; all they can do is simply bid in the auction (or abstain) in every round, as a function of the history leading to that round. The bidders share the discount factor $0 \leq \delta < 1$, and each bidder has a reservation utility of zero (in every period). The seller is inactive—he does not try to detect or deter collusion, nor does he behave in any other strategic way (alternatively, one can think of a sequence of short-lived sellers, one for each auction). At the end of each round, the seller publicly announces the (participation decisions and the) bids made in that round. The solution concept is *perfect Bayesian equilibrium* (PBE).

3   Endogenous bid rotation

The following collusive scheme is the main contribution of this paper:

Definition 1. *The Endogenous bid rotation scheme (EBR, for short) is defined as follows: Fix a number $\alpha \in [0,1]$, and pick the designated winner for period 0. This*

---

[5]To economize of space, I do not define a "standard auction" formally. I trust that this slight informality will not cause any confusion or unclarity.

[6]I take the interval of types to be $[0,1]$ just for notational simplicity. This is without loss of generality: all the analysis in the sequel can be carried out for an arbitrary interval of types $[a,b] \subset \mathbb{R}_+$.

[7]This special bid 0^+ is due to Blume and Heidhues [5].

bidder is instructed to bid 0, while his rival is instructed to refrain from bidding. Then, in each period $t+1$, the period-t winner is instructed to bid 0 and the period-t loser is instructed to bid 0^+ if his valuation is at least α and refrain from bidding otherwise.

In what follows, I will derive the equation that describes the value of this scheme for a given $\alpha \in [0, 1]$. Then, I will prove that there exists a unique $\alpha \in [0, 1]$ such that EBR can be supported as a repeated-game equilibrium.

For a fixed α, let $V^l(\alpha)$ be the ex ante value from starting the $(t+1)$-th period of the repeated game as the period-t loser under EBR. Similarly, let $V^w(\alpha)$ be this value for the period-t winner.[8] These numbers satisfy the following equations:

$$V^l(\alpha) = F(\alpha)\delta V^l(\alpha) + (1 - F(\alpha))[\mathbb{E}(v|v \geq \alpha) + \delta V^w(\alpha)],$$

and

$$V^w(\alpha) = F(\alpha)[v^e + \delta V^w(\alpha)] + (1 - F(\alpha))\delta V^l(\alpha).$$

Rearranging yields

$$V^l(\alpha) = \frac{(1 - F(\alpha))\mathbb{E}(v|v \geq \alpha)}{1 - \delta F(\alpha)} + \frac{(1 - F(\alpha))\delta}{1 - \delta F(\alpha)}V^w(\alpha), \tag{1}$$

and

$$V^w(\alpha) = \frac{F(\alpha)v^e}{1 - \delta F(\alpha)} + \frac{(1 - F(\alpha))\delta}{1 - \delta F(\alpha)}V^l(\alpha). \tag{2}$$

Adding up gives

$$V^l(\alpha) + V^w(\alpha) = \frac{(1 - F(\alpha))\mathbb{E}(v|v \geq \alpha) + F(\alpha)v^e}{1 - \delta}.$$

---

[8]These utility numbers depend on the discount factor. I omit it from the notation for brevity.

Therefore, the ex ante value of the scheme is:[9]

$$V(\alpha) = \frac{v^e}{2} + \frac{1}{2}\delta\left[\frac{(1 - F(\alpha))\mathbb{E}(v|v \geq \alpha) + F(\alpha)v^e}{1 - \delta}\right].$$

This is equivalent to an average per-period payoff of $(1 - \delta)V(\alpha)$. As $\delta \to 1$ this payoff converges to:

$$\frac{1}{2}(1 - F(\alpha))\mathbb{E}(v|v \geq \alpha) + \frac{1}{2}F(\alpha)v^e. \tag{3}$$

The intuition behind this formula is that conditional on being a last loser who choses to win, a bidder's expected payoff is $\mathbb{E}(v|v \geq \alpha)$, and conditional on being a last winner who is "awarded the win" by the rival it is v^e. Since on average each bidder is the last loser (and the last winner) one half of the time, the probability of the first event is $\frac{1}{2}(1 - F(\alpha))$ and the probability of the second event is $\frac{1}{2}F(\alpha)$.

Taking incentives into account, it must be that a last loser of type α is indifferent between bidding 0^+ and refraining from bidding. The corresponding indifference condition is:

$$\alpha = \delta[V^l(\alpha) - V^w(\alpha)] \tag{4}$$

Lemma 1. *For every distribution of valuations F and every discount factor $\delta \in [0, 1)$, (4) has a unique solution in $[0, 1]$.*

Equation (4) may have multiple solutions, depending on F and δ,[10] but exactly one of them belongs to the unit interval. Denote it by $\alpha(\delta)$.[11] This function α satisfies $\alpha(0) = 0$, it is strictly increasing, and takes its values in $[0, v^e)$; these properties will

---

[9]This value is calculated under the assumption that ex ante each bidder has a 50% chance of being the first winner. Strictly speaking, this is a little problematic, because I assume no public randomization. In any case, the identity of the first winner is of negligeble importance for patient bidders.

[10]This is the case when F is the uniform distribution and δ is close to one.

[11]This number depends on the discount factor *and* on the distribution, so a more complete notation would be $\alpha(\delta, F)$. Since the distribution of valuations is fixed throughout the analysis, and in order to economize on notation, I simply write $\alpha(\delta)$.

be established formally in subsection 3.1.

The fact that the right hand side of (4) is positive follows from the fact that being a last loser has a greater value, because he has the privilege of deciding on the current winner. It is also intuitive that the difference between these values is small (these values themselves are huge when δ is close to one), because all that "separates them" is one period. For future reference, it is convenient to note that by equations (1) and (2), (4) can be rewritten as:

$$\alpha = \delta \frac{\int_\alpha^1 t f(t) dt - F(\alpha) v^e}{1 + \delta - 2\delta F(\alpha)}. \tag{5}$$

From now on, let EBR denote the scheme defined in Definition 1 for $\alpha = \alpha(\delta)$.

We are now ready to turn to the first result: supporting (a slight modification of) EBR in a repeated-game equilibrium, in the special case where the auction format is second-price. In this special case, the scheme can actually be sustained independent of the discount factor.

Let π^c denote the competitive static payoff. Namely, π^c is the ex ante payoff for a bidder in the unique symmetric Bayesian Nash equilibrium of the one-shot auction.[12] It is independent of the auction format as long as we restrict attention to standard auctions; i.e., payoff equivalence holds (Myerson [11]). Given a number $x > 1$, let $E_i(x)$ be the repeated-game equilibrium of the repeated second-price auction, where bidder i bids x in every round, independent of his valuation, and bidder j bids 0 in every round.

Theorem 1. *Suppose that the auction format is second-price. Then EBR is implementable in a PBE for every discount factor $\delta \in [0, 1)$.*

Proof. Fix an $x > 1$, and consider EBR, with the following slight modification: the last loser is instructed to bid x (instead of 0^+) when his valuation is at least α.

---

[12]This payoff is computed prior to learning one's type; i.e., denoting by $\pi^c(\theta)$ the expected competitive payoff of type θ, $\pi^c = \int \pi^c(\theta) f(\theta) d\theta$.

Consider the following strategy profile: bidder i is instructed to play according to EBR as long as there is no evidence of a deviation by either bidder, in any past period. Once a deviation occurs, the bidder who deviated first, say j, is punished forever as follows: $E_i(x)$ is played forever.[13] The beliefs to accompany the strategy are obvious: a bidder believes that his opponent follows EBR as long as he does not have evidence to the contrary; off the path, any beliefs that are consistent with own behavior will do.[14]

The fact that off-path play consists of mutual best-responses is clear, so it only left to verify on-path incentives; since the latter are obvious for a last winner,[15] it only remains to verify, that the last loser has the right incentives to follow the scheme.

Consider then bidder i, who is the last loser on the path. Obviously, there is a cutoff $\beta \in [0, 1]$ such that bidder i will chose to win if his valuation is above β and he will abstain if his valuation is below β. Moreover, it is straightforward that there is such an optimal β, and that bidder i's best-response instructs him to follow this β-cutoff strategy, whenever he is the last loser, as long as there is no evidence of a deviation by bidder j. Let Z^l and Z^w be the values for bidder i under this β-strategy, when he is the last loser and last winner, respectively (on the path). These numbers satisfy

$$\beta = \delta(Z^l - Z^w),$$

[13]It may be the case that no deviation has occurred up to and including round t, but once the seller reveals the bids made in $t + 1$, it turns out that both bidders deviated (simultaneously) in $t + 1$, hence no "first deviator" exists. In such a case, any tie breaking rule for deciding on who will be punished will work.

[14]Here is an example for inconsistent beliefs: in the case where play went off the path and $E_i(x)$ is supposed to be played forever, and bidder i loses with the bid x, the belief that the opponent submitted the bid zero (as he is supposed to) is inconsistent for i with his own behavior.

[15]The last winner does not have short-term incentives to bid more than 0 because the last loser is bidding $x > 1$ when he chooses to take the good.

$$Z^w = F(\alpha)(v^e + \delta Z^w) + (1 - F(\alpha))\delta Z^l,$$

and

$$Z^l = F(\beta)\delta Z^l + (1 - F(\beta))(v^\beta + \delta Z^w).$$

where $v^\beta \equiv \mathbb{E}(v|v \geq \beta)$. Combining these equations yields $\frac{\beta}{\delta} = F(\beta)\beta - \beta + F(\alpha)\beta + (1 - F(\beta))v^\beta - F(\alpha)v^e$. This equation is satisfied for $\beta = \alpha$; moreover, this is the unique solution to this equation, because the derivative (wrt β) of the LHS is $\frac{1}{\delta} > 1$, and that of the RHS is $F(\beta) - 1 + F(\alpha) \leq \min\{F(\alpha), F(\beta)\} \leq 1$.

$\square$

Theorem 1 applies only to the second-price format. In a general auctions setting, extremely asymmetric equilibria are typically not available, and the EBR is implemented using the special bid 0^+, instead of the aforementioned modification, where the last loser takes the good by bidding a large number x. Consequently, the last winner needs to be provided with incentives not to "jump in and steal the good," a credible punishment phase needs to be introduced, and a certain degree of patience is required. The condition $\pi^c \leq \frac{v^e}{2}$ guarantees that simple Nash reversion can be invoked as the off-path threat; with this threat, the last winner will not "steal the good" if he is sufficiently patient.

Theorem 2. *Consider the general symmetric IPV repeated auction, and suppose that $\pi^c \leq \frac{v^e}{2}$. Then there exists a $\bar{\delta} \in (0,1)$ such that EBR is implementable in a PBE provided that $\delta \in (\bar{\delta}, 1)$.*

Proof. Consider the strategy that instructs each bidder to follow EBR as long as there has not been any deviation from it, and instructs perpetual competitive bidding once a deviation has occurred. Since the EBR payoff exceeds $\frac{v^e}{2}$, which, in turn, is greater than the expected competitive payoff π^c, it follows that when the bidders are sufficiently patient noone will want to make a detectable deviation. Consequently,

the only incentives that remain to be verified are those of a last loser who, when on the path, needs to decide on whether to take the good or not. The arguments from the proof of Theorem 1 establish that it is optimal for him to take it if and only if his valuation is at least $\alpha(\delta)$. □

The condition $\pi^c \leq \frac{v^e}{2}$ is satisfied for a large family of distributions. McAfee and McMillan [10] showed that it is satisfied if F has an increasing hazard rate. In addition, as the following lemma shows, it is satisfied if its expectation, v^e, is at least $\frac{1}{2}$ (more generally, one half of the maximal valuation).

Lemma 2. *Let F be a distribution on $[0, 1]$ such that $v^e \geq \frac{1}{2}$. Then $\frac{v^e}{2} \geq \pi^c$.*

Proof. It is well-known that in a general standard auction with n symmetric bidders $\pi^c = \int_0^1 [1 - F(t)]F(t)^{n-1}dt$ (see Riley and Samuelson [14]). For $n = 2$ this expression equals $\int_0^1 [1 - F(t)]F(t)dt$, and since the function $(1 - x)x$ is maximized at $x = \frac{1}{2}$ we have that $\pi^c \leq \frac{1}{4}$, which in turn is smaller than $\frac{v^e}{2}$ by assumption. □

Hence we have:

Corollary 1. *Consider the general symmetric IPV repeated auction. Suppose that either (i) $v^e \geq \frac{1}{2}$ or (ii) $\frac{f}{1-F}$ is weakly increasing. Then there exists a $\bar{\delta} \in [0, 1)$ such that for all $\delta \in [\bar{\delta}, 1)$ there exists a PBE that implements EBR. Moreover, if the auction format is second-price, then $\bar{\delta} = 0$.*

The conditions in Corollary 1 are sufficient for EBR to be sustainable in equilibrium, but they are not necessary. For example, if the distribution of valuation is $F(t) = t^\gamma$, where $\gamma \in (0, \frac{1}{2})$ and close to $\frac{1}{2}$, then $\pi^c > \frac{v^e}{2}$ and $\frac{f}{1-F}$ is not monotonic, but EBR can nevertheless be implemented by Nash reversion, because its payoff exceeds π^c.

3.1 Properties of EBR

For the sake of formulating several of the following results, the following will be useful: call a collusive scheme an *α-scheme* if it obeys Definition 1, but where α

does not necessarily equal $\alpha(\delta)$. Looking at the class of α-schemes, a question that arises naturally is what is the optimal α. Intuitively, one would expect it to equal the expectation, v^e. This intuition is correct, and is easily proved below.

Proposition 1. *The v^e-scheme has the maximal value among all α-schemes.*

Proof. Maximizing the expression in equation (3), the first-order condition is $-\alpha f(\alpha) + f(\alpha)v^e = 0$, so $\alpha = v^e$. It is easy to see that the second-order condition is also satisfied.

$\square$

The optimality of the v^e-scheme is illustrated in the following fictitious scenario. Suppose that the cartel is owned by a single agent, who instructs its subordinates how to bid in each round. Suppose further that prior to sending his instructions he only gets to observe the valuation of the previous round loser. It is clear that the optimal policy for this agent is to allocate the good to this last loser, if and only if his valuation is above v^e. With this result at hand, we may further ask whether there exist discount factors and distribution functions for which $\alpha(\delta) = v^e$. The following proposition shows that this is never the case: EBR always does worse than the first-best α-scheme.

Proposition 2. *For every distribution of valuation F and every discount factor $\delta \in [0,1)$, $\alpha(\delta) < v^e$.*

Proof. Assume by contradiction that $\alpha(\delta) \geq v^e$ for some F and δ. Then, equation (5) implies $v^e[1 + \delta - 2\delta F(\alpha(\delta))] \leq \delta[v^e - \int_0^{\alpha(\delta)} t f(t)dt - F(\alpha(\delta))v^e]$, which, upon rearranging, becomes $v^e[1 - \delta F(\alpha(\delta))] \leq -\delta \int_0^{\alpha(\delta)} t f(t)dt$, a contradiction.

$\square$

For the particular case $\alpha = \alpha(1)$, the fact stated in Proposition 2 can be seen in the following alternative way. Combining equations (3) and (5) gives that the average per-period payoff of EBR when $\delta \to 1$ is:

$$F(\alpha)v^e + (1 - F(\alpha))\alpha, \tag{6}$$

where $\alpha = \alpha(1)$. Then, $\alpha \geq v^e$ implies that the average per-period payoff *for each bidder* is weakly greater than v^e, which is of course impossible.

Another intuitive property one would expect is the monotonicity of α, because patient bidders have higher willingness to substitute current goods for continuation payoffs, relatively to impatient bidders. The following proposition shows this formally.

Proposition 3. *The function α is strictly increasing in δ.*

Proof. Multiplying both sides of equation (5) by $[1 + \delta - 2\delta F(\alpha(\delta))]$ and taking derivative with respect to δ gives:

$$\alpha'(\delta)[1 + \delta - 2\delta F(\alpha(\delta))] + \alpha(\delta)[1 - 2F(\alpha(\delta)) - 2\delta\alpha'(\delta)f(\alpha(\delta))] =$$
$$= \frac{\alpha(\delta)}{\delta}[1 + \delta - 2\delta F(\alpha(\delta))] - \delta\alpha(\delta)\alpha'(\delta)f(\alpha(\delta)) - \delta\alpha'(\delta)f(\alpha(\delta))v^e.$$

Rearranging we get:

$$\alpha'(\delta)[1 + \delta - 2\delta F(\alpha(\delta)) + \delta f(\alpha)(v^e - \alpha(\delta))] = \frac{\alpha}{\delta}.$$

By Proposition 2 the term $(v^e - \alpha(\delta))$ is positive, hence the expression on the brackets is positive. Therefore, $\alpha'(\delta) > 0$.

$\square$

Lemma 3. $\alpha(1) < \frac{v^e}{2}$.

Proof. Setting $\delta = 1$ in (5) gives:

$$\alpha(1) = \frac{v^e}{2} - \frac{\int_0^\alpha t f(t)dt}{2[1 - F(\alpha(1))]}. \tag{7}$$

$\square$

Combining Propositions 1 through 3 together with the continuity of the payoff function and Lemma 3 gives:

Corollary 2. *The EBR payoffs are bounded away from the first-best α-scheme payoffs, uniformly across δ.*

Another natural relationship to investigate is the one between the threshold α and the probability of an efficient allocation in a given auction. Letting $Q(\alpha)$ denote the probability that the good will end up in the hands of the bidder who values it the most under the α-scheme, one might think that the optimal α-scheme is the one whose α maximizes Q. That is, that "maximizing selection" is equivalent to "maximizing value" (within the class of α-schemes). This equivalence holds if the distribution of valuations is symmetric, but not in general.

Proposition 4. *The probability Q is maximized at the median; i.e., if α maximizes Q, then $F(\alpha) = \frac{1}{2}$.*

Proof. The aforementioned probability is $Q(\alpha) = \frac{1}{2} + \frac{1}{2}[\int_0^\alpha \frac{1-F(\alpha)}{1-F(t)} f(t) dt + \int_\alpha^1 \frac{F(\alpha)}{F(t)} f(t) dt]$. Maximizing the bracketed term gives the first order condition $\int_0^\alpha \frac{f(t)}{1-F(t)} dt = \int_\alpha^1 \frac{f(t)}{F(t)} dt$. Solving for the integrals gives $-\log(1 - F(\alpha)) = -\log F(\alpha) \Rightarrow F(\alpha) = \frac{1}{2}$. □

It may look surprising that the α that maximizes efficient selection does not, in general, equals the one that maximizes the scheme's value. A moment's reflection, however, reveals why this is the case. The two α's correspond to different optimization problems: the problem of "maximizing efficient selection" is defined in terms of the stage game, while the value-maximization problem is defined for the repeated game.

To see why the α that maximizes efficiency is such that $F(\alpha) = \frac{1}{2}$, consider the following fictitious scenario. Suppose that the last winner is bidding 0 independent of his valuation, and the last loser is free to do whatever he wants. Suppose further that his objective is to maximize efficient selection. That is, his payoff function is $\zeta - t$,

where t is his payment to the seller and $\zeta = 1$ if the (ex post) allocation is efficient and $\zeta = 0$ otherwise. Then, he will play a threshold-strategy for some threshold α. Furthermore, at this threshold α he is exactly indifferent between participation and abstention, the indifference condition being $F(\alpha) = 1 - F(\alpha)$, or $F(\alpha) = \frac{1}{2}$.

To see that the α that maximizes value need not satisfy $F(\alpha) = \frac{1}{2}$, consider a distribution such that $F(v^e) > \frac{1}{2}$. Suppose that the last loser draws a value $v \in (\alpha', v^e)$, where α' satisfies $F(\alpha') = \frac{1}{2}$. In this case, rejecting the good is desirable despite the fact that it increases the probability of inefficiency in the current period, because acceptance would imply acceptance in this range of values in all future periods, which, in turn, translates to a lower-than-optimal value.

3.2 Example: The Uniform Case

In order to assess the EBR payoff, it is helpful to look at some numbers. Suppose that F is uniform and that δ is close to 1. In this case, the competitive and bid rotation payoffs are 0.1666 and 0.25, respectively. Aoyagi [1] derives an interesting variant of bid rotation, *dynamic bid rotation*, where the bidders communicate their valuations in every round, and based on their reports a single designated participant for the auction is selected. The payoff from this scheme in the uniform case is approximately 0.288. The exclusion scheme of Skrzypacz and Hopenhayn gives a payoff of 0.276. The EBR payoff is higher than all of the above, 0.295.[16] The first-best payoff is 0.333, and the one of the optimal chips mechanism is 0.305.

4 Closing comments

EBR is very simple: it involves only pure strategies, no communication, no transfers, no public randomizations, and it can be applied even if the bidders have bounded

---

[16]There is no claim here that EBR is better than the above-mentioned schemes in some absolute sense. For example, Aoyagi's dynamic bid rotation, as opposed to EBR, allows for correlated types.

(and very short) memory. Furthermore, it is robust to modifications of the auctions'
monitoring technology; for example, in the case of a repeated second-price auction,
EBR can be implemented in a PBE even if there is no public information whatsoever:
that is, if all that a bidder knows is his own history of valuations, actions, and wins,
and nothing is announced by the seller.[17]

Though it is defined for a continuous distribution, EBR has an obvious analog
in a two-type model: when types are binary, the last loser is instructed to take the
good if and only if his current valuation is high. Consequently, the scheme delivers
the first-best when there are only two types.[18] Given any binary distribution, this
two-type-EBR scheme is implementable in PBE when the bidders are sufficiently
patient.

5   Appendix

Proof of Lemma 1: By equations (1) and (2), (4) can be rewritten as:

$$\alpha = \delta \frac{\int_\alpha^1 tf(t)dt - F(\alpha)v^e}{1 + \delta - 2\delta F(\alpha)}$$

For a fixed δ, let $\phi^\delta(\alpha) \equiv \delta \frac{\int_\alpha^1 tf(t)dt - F(\alpha)v^e}{1+\delta-2\delta F(\alpha)}$ and let $\psi^\delta(\alpha) \equiv \phi^\delta(\alpha) - \alpha$. The function
$\psi^\delta \colon [0,1] \to \mathbb{R}$ is continuous with $\psi^\delta(0) = \frac{\delta v^e}{1+\delta} > 0$ and $\psi^\delta(1) = -\frac{\delta v^e}{1-\delta} - 1 < 0$. By
the mean value theorem, there exists an $\alpha(\delta) \in (0,1)$ such that $\psi^\delta(\alpha(\delta)) = 0$. To
prove uniqueness, it is sufficient to show $\frac{d}{d\alpha}\phi^\delta < 1$. The derivative $\frac{d}{d\alpha}\phi^\delta$ is given by
the following expression:

---

[17]Defining and dealing with private monitoring results in much heavier notation, which, in turn,
does not translate to substantial gain in insights. An interested reader may contact the author for
details regarding the implementability of EBR under private monitoring.

[18]Related results have been obtained by Athey and Bagwell [3], [4] and Hörner and Jamison [8]
for a repeated Bertrand game.

$$\frac{\delta\{[-\alpha f(\alpha) - f(\alpha)v^e][1 + \delta - 2\delta F(\alpha)] + 2\delta f(\alpha)[(1 - F(\alpha))v^e - \int_0^\alpha tf(t)dt]\}}{[1 + \delta - 2\delta F(\alpha)]^2} \quad (8)$$

Clearly, the desired inequality holds if the numerator in (8) is negative. If it is positive, then it is sufficient to establish:

$$[-\alpha f(\alpha) - f(\alpha)v^e][1 + \delta - 2\delta F(\alpha)] + 2\delta f(\alpha)(1 - F(\alpha))v^e < [1 + \delta - 2\delta F(\alpha)]^2$$

or:

$$2\delta f(\alpha)(1 - F(\alpha))v^e < [1 + \delta - 2\delta F(\alpha)][1 + \delta - 2\delta F(\alpha) + \alpha f(\alpha) + f(\alpha)v^e]$$

For the latter, it is sufficient to establish $2\delta f(\alpha)(1 - F(\alpha))v^e < [1 + \delta - 2\delta F(\alpha)]f(\alpha)v^e$, which holds because $\delta < 1$. $\square$

6   References

[1] Aoyagi, M. (2003), "Bid Rotation and Collusion in Repeated Auctions", *Journal of Economic Theory*, **112**, 79-106.

[2] Aoyagi, M. (2007), "Efficient Collusion in Repeated Auctions with Communication", *Journal of Economic Theory*, **134**, 61-92.

[3] Athey, S. and Bagwell, K. (2001), "Optimal Collusion with Private Information", *RAND Journal of Economics*, *32*, 428-465.

[4] Athey, S. and Bagwell, K. (2008), "Collusion with Persistent Cost Shocks", *Econometrica*, *76*, 493-540.

[5] Blume, A. and Heidhues, P. (2006), "Private Monitoring in Auctions", *Journal of Economic Theory*, **131**, 179-211.

[6] Fudenberg, D., Levine, D.K., and Maskin, E. (1994), "The Folk Theorem with Imperfect Public Monitoring", *Econometrica*, **62**, 997-1039.

[7] Graham, D.A, and Marshall, R.C. (1987), "Collusive Bidder Behavior at Single Object Second-Price and English Auctions", *Journal of Political Economy*, **86**, 1217-1239.

[8] Hörner, J. and Jamison, J. (2007), "Collusion with (almost) No Information", *RAND Journal of Economics*, **38**, 804-822.

[9] Mailath, G., and Zemsky, P. (1991), "Collusion in Second Price Auction with Hetrogenous Bidders", *Games and Economic Behavior*, **3**, 467-486.

[10] McAfee, P.R., and McMillan, J. (1992), "Bidding Rings", *American Economic Review*, **82**, 579-599.

[11] Myerson, R.B., (1981), "Optimal Auction Design," *Mathematics of Operations Research*, **6**, 58-73.

[12] Nelson, J.P. (1993), "Comparative Antitrust Damages in Bid-Rigging Cases: Some Findings from a Used Car Auction", *The Antitrust Bulletin*, **38**, 369-394.

[13] Philips, O.R., Menkhaus, D.J., and Coatney, K.T. (2003), "Collusive Practice in Repeated English Auctions: Experimental Evidence on Bidding Rings", *American Economic Review*, **93**, 965-979.

[14] Riley, J.G., and Samuelson, W. (1981), "Optimal Auctions", *American Economic Review*, **71**, 381-392.

[15] Skrzypacz, A., and Hopenhayn, H., (2004), "Tacit Collusion in Repeated Auctions", *Journal of Economic Theory*, **114**, 153-169.