Games and Economic Behavior 83 (2014) 224–230

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

First-best collusion without communication

Shiran Rachmilevitch

Department of Economics, University of Haifa, Israel

ARTICLE INFO

Article history:
Received 30 December 2011
Available online 26 November 2013

JEL classification:
D21
D43
D44
D82

Keywords:
Auctions
Collusion
Repeated games

ABSTRACT

I study a 2-bidder infinitely repeated IPV first-price auction without transfers, communication, or public randomization, where each bidder's valuation can assume, in each of the (statistically independent) stage games, one of three possible values. Under certain distributional assumptions, the following holds: for every $\epsilon > 0$ there is a nondegenerate interval $\Delta(\epsilon) \subset (0, 1)$, such that if the bidders' discount factor belongs to $\Delta(\epsilon)$, then there exists a Perfect Public Equilibrium with payoffs ϵ-close to the first-best payoffs.

1. Introduction

Repeated interaction among the same set of bidders in a sequence of auctions that take place over time is a common occurrence. Car dealers who meet in vehicle auctions regularly, companies who compete for government defense contracts repeatedly, and art collectors who compete in Sothby's auctions on a regular basis, to name a few, are all instances of this phenomenon.

In a repeated game, the rich strategy sets open a door for a wide range of collusive behaviors. One of the very first questions that comes to mind, therefore, is this: *what is the scope of collusion in a repeated auction?* Namely, if the bidders were to form a cartel, what repeated-game payoffs could they achieve beyond the competitive payoff? It is known that if the cartel members can organize side transfers, then they can collude in a way that implements the *first-best outcome*—efficient allocation and reserve-price payment to the seller in every stage game; in fact, when transfers are available, the cartel can implement the first-best even in a one-shot auction.[1] If, on the other hand, transfers are not available, then there is not much that the cartel can do in the one-shot auction, but in a repeated game it can still approximate the first-best when the discount factor tends to one, provided that there are finitely many valuations and that the cartel members can communicate before every round.[2] Without communication—when all that the bidders can do is to bid (or abstain) in the stage auctions—collusion is much more difficult. In particular, it is an open question in the existing literature whether the first-best can be achieved (or at least approximated) in communication-free, transfers-free collusion, if the bidders have more than two possible valuations.[3]

[1] This was shown by McAfee and McMillan (1992) for symmetric, independent-private-values (IPV) first-price auctions and by Mailath and Zemsky (1991) for IPV second-price auctions.

[2] This follows from the analysis of Fudenberg et al. (1994) of repeated games with public monitoring.

[3] It is known that with two valuations the first-best can be achieved without communication and transfers. Hörner and Jamison (2007) approximate the first-best in a Bertrand game (which is strategically equivalent to a first-price auction) using review strategies without public randomization, and Athey and Bagwell (2001, 2008) prove that with public randomization the Bertrand game's first-best can be achieved exactly, provided that the discount factor

0899-8256/$ – see front matter
http://dx.doi.org/10.1016/j.geb.2013.11.007

In this paper I demonstrate that in a repeated 2-bidder IPV first-price auction, the first-best can be approximated as closely as one wishes when there are *three* valuations, without using communication, transfers, or public randomization. In particular, it follows that the fact that there are more than two types does not, by itself, entail an impossibility, or inefficiency result.

The idea behind the collusive strategy I propose is simple: in each round, the last period's loser has more weight in the decision on the identity of the current winner. This privilege translates to an implicit transfer—losing today is accompanied by a larger continuation value (relatively to winning) because in the next period the current loser will have greater weight in choosing the winner. More specifically, the strategy works as follows. Fix an $\epsilon > 0$, arbitrarily small. In each round, the last period's loser is instructed to abstain if his valuation is low, to bid 3ϵ if it is high, and bid ϵ if it is medium. This leads to efficient allocation, because the last period's winner is instructed to bid zero unless his valuation is high, in which case he is instructed to bid 2ϵ. Bids are observable and deviations trigger a perpetual punishment phase, Nash reversion. Under some restrictions on the distribution of valuations, this strategy is sustainable as a perfect Bayesian equilibrium (PBE) in the repeated game; moreover, this PBE is, in fact, a Perfect Public Equilibrium (PPE, due to Fudenberg et al., 1994) because in every stage the strategy conditions only on public histories. When this strategy is played the first-best is approximated, and is obtained in the limit, as $\epsilon \to 0$. Call this scheme the *ϵ-scheme*.

The sustainability of the ϵ-scheme is described in the main result, Theorem 1. This theorem, in turn, builds on two intermediate results, Proposition 1 and Proposition 2. With the type distribution denoted by $\mathbf{D}$, Proposition 1 states that if $\mathbf{D}$ is such that the medium type is between forty and fifty percent of the high type and the probability of the high type is sufficiently large, then the following holds: for every $\epsilon > 0$ there is a discount factor $\delta_{\mathbf{D}}(\epsilon) \in (\frac{2}{3}, 1)$ such that if the bidders' discount factor is $\delta_{\mathbf{D}}(\epsilon)$, then there is a PPE with payoffs ϵ-close to the first-best payoffs. Next, Proposition 2 establishes that for any $\mathbf{D}$ that satisfies the conditions from Proposition 1, the function $\delta_{\mathbf{D}}$ is strictly monotonic (its derivative is never zero). This implies that, starting from any initial level of approximation $\epsilon > 0$, one can take a discount factor, say δ', in the neighborhood of $\delta_{\mathbf{D}}(\epsilon)$, and obtain a PPE with payoffs ϵ'-close to the first-best payoffs when the discount factor is δ', where $\epsilon' = \delta_{\mathbf{D}}^{-1}(\delta') \sim \epsilon$. In particular, for every level of approximation ϵ, there is a nondegenerate interval $\Delta(\epsilon)$ such that an equilibrium with payoffs ϵ-close to the first-best payoffs exists, provided that the discount factor is in $\Delta(\epsilon)$.

The rest of the paper is organized as follows. The next section describes the model. Sections 3 and 4 contain the statements and proofs of Proposition 1, Proposition 2 and Theorem 1. Section 5 concludes.

2. Model

There are two bidders, bidder 1 and bidder 2, who are about to participate in an infinite sequence of identical first-price auctions. In each round, or stage game, there is one indivisible object up for sale; bidder i's valuation for the object in a given round is an independent draw from a distribution whose support is $\{v^H, v^M, v^L\}$, where $v^H > v^M > v^L > 0$; the probability of valuation v^k is p_k, $k \in \{H, M, L\}$. Let $\mathbf{D}$ denote this distribution. This distribution is the same for both bidders, and there is independence across bidders and across rounds: each round is an independent IPV first-price auction (that is, independent of the other rounds). Valuations (or types) are privately known by the bidders who hold them.

In every stage game (stage auction) the set of pure actions available to a bidder—the set of *generalized bids*—is $\mathbb{R}_+ \cup \{N\}$, where N denotes abstention. Thus, a bidder can chose not to show up in the auction at all, and in case that he does show up he can submit any (pure) bid in $\mathbb{R}_+$. Moreover, mixed bidding strategies are allowed, so conditional on participating in the auction the bidders can play mixed strategies.[4] The winner is the highest bidder,[5] and he pays his own bid; in case that the bids are tied, each bidder is chosen to be the winner with equal probability. The stage game utility of a type v bidder who obtains the good and pays the bid b is $v - b$, and the stage game utility from losing is zero. The bidders, who share the discount factor $\delta \in (0, 1)$, seek to maximize, in every round, their expected discounted sum of utilities from that point in the game onwards.

The reserve price is zero.[6] The seller is not strategic. Generalized bids are observable—they are announced by the seller at the end of each round. Denote this repeated game by $\mathcal{G}$.

A *public history* is a list that specifies, for each round, the list of generalized bids in all previous rounds. A *public strategy* assigns to each bidder behavior in each round, as a function of his type and of the public history leading to that round. It is a Perfect Public Equilibrium (PPE for short, due to Fudenberg et al., 1994) if it forms a Bayesian Nash equilibrium after every public history.

The term *first-best payoff* means one half of the first order statistic; namely, the ex ante stage game payoff that corresponds to the situation where the good is allocated efficiently and no payment is made to the seller.

exceeds a critical level which is strictly less than one. It follows from the analysis in Rachmilevitch (2013) that with two valuations the first-best can be achieved exactly, without public randomization, communication or transfers, in any repeated standard IPV auction (furthermore, the strategies delivering this result are simple: they are stationary with 1-recall). In their above mentioned paper, Hörner and Jamison show that with finitely many (more than two) valuations, the cartel can approximate the first-best in an ϵ-equilibrium.

4 On the path of the collusive scheme only pure actions are played. However, randomization is essential off the path, in the punishment phase.

5 If both bidders abstain then nobody wins; if only one abstains then the other bidder wins automatically, even if he bids zero.

6 This is without loss of generality in the following sense: for every reserve price $r \in (0, v^L)$, one can define adjusted valuations by $\tilde{v}^k \equiv v^k - r$, $k \in \{H, M, L\}$, and then apply the analysis to the distribution whose support is $\{\tilde{v}^H, \tilde{v}^M, \tilde{v}^L\}$.