Tacit Collusion in Auctions with Private Strategies

Galit Ashkenazi-Golan*

and

Dipjyoti Majumdar†

March 2006

Abstract

We investigate two-players private-value repeated auction, where only the winner's identity is announced. No other information is available. we show that there exist ϵ-equilibria approximating first-best collusive outcomes when the bidders are sufficiently patient and are allowed to condition their strategies on their private information. In addition, we show that when the players are restricted to using only public strategies, then their payoff is bounded away from efficiency in any ϵ-equilibria.

keywords: collusion, private strategies, imperfect private monitoring.

*Tel-Aviv University, Tel-Aviv, Israel (galit.ashkenazi@gmail.com)

†Concordia University, Montreal, Canada (dmajumda@alcor.concordia.ca)

1 Introduction

Collusive behavior among bidders is a well recognized problem in auctions. One branch of the theoretical literature tries to analyze the extent to which a group of bidders can successfully collude. Most of the early theoretical literature on collusion focuses on single period auctions (Graham and Marshall (1987), McAfee and McMillan (1992)) (here after M and M). In their important paper McAfee and McMillan show that extent to which a group of bidders can successfully collude is crucially related to the possibility of ex-post transfers. A collusion is successful if it can maximize the surplus. Surplus is maximized if the bidder with the highest valuation gets the object and pays the reserve price. In a model where a bidder's valuation is private information, the identification of the bidder with the highest valuation is an issue. McAfee and McMillan show that if the cartel members can engage in side payments, a mechanism such as running an internal auction and then bidding the reserve price can effectively sustain perfect collusion i.e., maximize the surplus (in the literature this is called efficient collusion). However in the absence of side transfers the efficiency of a cartel is severely limited.

One limitation of the analysis discussed above is that the model is one of static collusion (that is, a one-shot auction framework). However, if collusion is a result of repeated interactions, a more appropriate framework is provided by the theory of repeated games under imperfect monitoring. The analysis of collusion in auctions along these lines is relatively recent. Aoyagi (2002, 2003) considers a model of repeated auction where the agents communicate their bids at the end of each round. Aoyagi develops a *dynamic bid-rotation* scheme (as against the *static* bid-rotation scheme of M and M) and shows that even without side payment of money, it is possible to improve cartel efficiency over the static bid-rotation scheme, through intertemporal payoff transfers. Aoyagi considers explicit communication among bidders. In many real life auctions however, cartel formation among bidders is illegal and explicit communication is not allowed. In an important paper Skryzpacz and Hopenhayn (1999) (hereafter, S and H) prove the existence of a collusion scheme without communication that performs strictly better than the static scheme of M and M. They construct collusion schemes that have asymmetric continuation equilibria, and through these asymmetric equilibria they are able to generate implicit transfer schemes. S and H consider strategies of the players that depend on the history of publicly observable signals (in their case the history of wins) and constructs a Public Perfect Equilibrium (PPE) of the repeated game in which the cartel extracts a surplus that is strictly greater than the one obtainable under bid-rotation scheme. However, the equilibrium surplus is bounded away from the fully efficient surplus.

The studies so far analyzes collusion by focusing on the PPE of the resulting repeated game. In a PPE, the players are restricted to use *public strategies* i.e., strategies that depend only on the history of the publicly observable signals. In this paper, we take the analysis one step further by considering *private equilibria* of the repeated game. In a private equilibrium, the (private) strategies of a player may depend on the history of his past actions which is private information, in addition to the history of the publicly observable signals. The interest in private equilibria in repeated games whose monitoring is public is fairly recent. Mailath *et. al.* (2002) give three examples in the case of finitely repeated games where the equilibria in private strategies Pareto dominates every PPE. In the context of infinitely repeated games Kandori and Obara (2003) extends this analysis and illustrates cases where players make better use of information by using private strategies. Both the papers consider games with finite action spaces. In this paper we study collusion in auction games where the action spaces for the agents are intervals in the real line. Most of the literature on auctions deals with models where action spaces are intervals on the real line. In such a case, even with communication, under imperfect public monitoring, the best achievable pay-off configuration is bounded away from the efficient pay-off (Aoyagi, 2003). Moreover, most of the literature on repeated games deals with sequential equilibrium. However, sequential equilibrium is only defined for games with finite action spaces.

In this paper we consider a model of first price auctions where there are two bidders with Private Valuations (PV). The two bidders are engaged in repeated action. Their private valuations are randomly drawn before each round according to a common distribution over $[0, 1] \times [0, 1]$. At the end of each round the auctioneer announces the identity of the winner in that round. Although we are unable to construct an exact equilibrium that induces efficient outcome, we show that there exists ϵ-equilibria, where the outcome is close in an appropriate sense to the efficient outcome.

Specifically we show that, if players are sufficiently patient, then using private strategies it is possible to construct an ϵ-equilibrium in which the payoff to the bidders is arbitrarily close to what they would get if the object is allocated to the highest valuation bidder at the reserve price in every period. In their celebrated paper Fudenberg, Levine and Maskin (1994) show that it is possible to obtain full efficiency under imperfect public monitoring as long as the space of public signals is rich enough relative to the action spaces of the agents. Aoyagi (2003) obtains a similar result for auctions when the agents have finite action spaces. Note that, in our model the action spaces of the bidders are intervals on the real line. Moreover, the only public signals available at the end of each period are the identities of the winner. As in S and H we show that with such an uninformative public signal the cartel's surplus is strictly bounded away from the efficiency surplus.

The specific cartel mechanism we consider consists of three phases: the *cooperation* phase, the *communication* phase and the *punishment* phase. The game starts in the cooperation phase where both the bidders bid an amount that is proportional to their valuation for that period. Choosing a proportional factor close to zero makes the bids close to zero (the normalized reserve price in our model). The cooperation phase lasts for N periods (N pre-determined) at the end of which the bidders conduct a statistical test. The statistical test is to check whether the other bidder has played according to the prescribed strategy. After the statistical test the game moves to the communication phase where the bidders communicate the results of the test. In manner that will be made clear in the body of the paper, the communication takes place through bids, and hence makes use of the available public signals only. No other external means of communication is available. If both players communicate that the opponent has passed the statistical test the game returns to the cooperation phase. Otherwise the game moves to the punishment phase. The punishment phase has two stages — the punishment for player 1 (if he has failed the test) and the punishment of player 2 (if he has failed the test). Observe that the transition from cooperation phase to the punishment phase depends on statistical test, the result of which is private information. Each player knows that by mis-reporting the result of the statistical test he can induce or prevent a punishment for the other player and can favorably alter his overall payoff. The challenge is to ensure that each bidder has incentive to report the result of his statistical test truthfully. The trick is to choose a mixed action for the punished player so as to make the punishing player indifferent between punishing and not punishing. Similar ideas appear in Piccione (2002), Ely and Valimaki (2002) in the context of repeated games with private monitoring and in Kandori and Obara (2003) in the context of repeated games with imperfect public monitoring and private strategies. The details of the construction follow in the body of the paper.

The repeated game that we consider is one with discounting. Every specific discount factor δ entails a repeated game. By changing the discount factor we get a sequence of such repeated games. For each such game we construct the epsilon equilibrium and show that as the discount factor tends to 1 (i.e., players become more patient), the ϵ-equilibrium outcomes tend to the efficient outcome.

As mentioned before we are unable to give a precise equilibrium argument. The reason is that even though we consider the augmented set of private histories, we still cannot account for all possible deviations. If one has to take account of all possible deviations, it comes at the cost of efficiency. What we show that such deviations are profitable only up to ϵ. In a sense made precise later, our solution is more than an ϵ-equilibrium. It is an ϵ-consistent equilibrium (Lehrer and Sorin (1998), which is equivalent to contemporaneous perfect ϵ-equilibrium of Mailath *et. al.* (2003)).

The paper, for its major part, focuses on Independent Private Valuations model with uniform priors. However, in the appendix we also prove that our results hold for other kinds of priors.

2 Preliminaries

The set of bidders is $N = \{1, 2\}$. The bidders are symmetric and risk-neutral. A single indivisible object is sold in every period through a fixed auction format. In our model, it is a first-price auction. At the beginning of each period nature draws a valuation for each player that is independent of the past draws and each component is independently distributed of the other. The vector of private valuations is denoted by $v = (v_1, v_2)$. We assume that each player's valuation v_i is independently and uniformly distributed over the interval $[0, 1]$.

The participation in the auction is voluntary, so in any period the set of each bidder's *generalized* bids is given by the set $B = \{\emptyset\} \cup I\!R_+$, where $\emptyset$ represents no-participation.

2.1 The Stage Game

In every period the object is sold through a *first-price sealed bid* auction (FPA). We normalize the reserve price to zero. The auction mechanism is described by the measurable mappings p_i and t_i $(i = 1, 2)$ on the set B^2 of bid profiles $b = (b_1, b_2)$: $p_i(b)$ is the probability that bidder i is awarded the good and t_i is his expected payment. Each agent has a strategy $\phi : [0, 1] \to B$. The stage-game expected payoff of agent i is,

$$r_i(\phi_1, \phi_2) = \mathbb{E}(p_i(\phi_i, \phi_j)v_i - t_i(\phi_i, \phi_j)) \tag{1}$$

The functions p_i and t_i are symmetric and satisfy the following conditions:

Assumption 1:

i A bidder makes no payment when he chooses not to participate i.e., $t_i(b) = 0$ if $b_i = 0$.

ii If only one bidder participates and bids zero, then he wins the object at price zero.

A second assumption we make is that after each period only the identity of the winner is publicly announced (or, of course, that this information can costlessly be discovered).

Let r^0 be the (ex-ante) symmetric Bayseian Nash Equilibrium payoff to each bidder in the stage auction. Also let $r^\star$ be the expected payoff to each bidder under truthful information sharing and efficient allocation with bidder i winning the object and paying price 0 if and only if $v_i > v_j$. In other words,

$$r^\star = \mathbb{E}[1_{\{v_i > v_j\}}\{v_i\}]$$

2.2 The Repeated Game:

The repeated game G^δ is the repetition of the stage game G. Given a stage $t \in \{1, 2, ...\}$ let h_t denote the history of the game up to stage t. We denote by $\mathcal{H}_t$ the set of all such t-stage histories. We would like to distinguish between two kinds of histories. The *public* history of a game consists of the data up to stage t that is publicly observable. Formally let h_t^p denote the public history of a game up to stage t. In our game, the public history consists of the sequence of winners up the current period. Let $\mathcal{H}_t^p$ be the set of all such histories. For any player i, a *private* history up to stage t, on the other hand, includes the private information that player i may have up to stage t, in addition to the publicly observable data. For example, in the game we consider here, a t-stage private history for player i records for each of the first $t-1$ periods, bidder i's valuation $v_{i,\tau}$, his bids $b_{i,\tau}$ as well as the data on the publicly observable signals in period τ. Let h_{it} denote the private history of player i up to stage t and $\mathcal{H}_{i,t}$ be the set of all such private histories. A *private* behavioral strategy for player i in period t is a function $\sigma_{it} : \mathcal{H}_{i,t} \to \Delta(B)$ that maps t-stage private histories into probability distributions over set B. A *private* strategy for a player is a collection $\sigma_i = (\sigma_{it})_{t=1}^{\infty}$. Given that $\sigma = (\sigma_i, \sigma_j)$ is the strategy profile, the payoff to player i in the repeated game G^δ is,

$$\pi_i(\sigma_i, \sigma_j) = (1-\delta)\sum_{t=1}^{\infty} \delta^t r_i(\sigma_{it}, \sigma_{jt})$$

2.3 The Equilibrium Scheme

In this section we set out the collusive scheme. We are considering a scenario where the bidders' valuations are uniformly and independently distributed. In this case, the expected payoff of each bidder in the one period symmetric Nash equilibrium (NE) is $r^0 = \frac{1}{6}$. The most efficient collusive payoff is $r^\star = \frac{1}{3}$. The collusive mechanism that we consider has four phases: the cooperation phase C, the statistical test phase S, the communication phase $\bar{C}$ and lastly the punishment phases P_1 and P_2.

Cooperation Phase: Play begins in the cooperation phase which lasts for N periods

(N pre-specified). In each period $\tau \in \{1, \cdots, N\}$, the players are advised to bid $\alpha v_{i,\tau}$, $\alpha > 0$. In other words, the advised strategy $\sigma^\star$ has the following feature: for all $\tau \in \{1, \cdots, N\}$ and for all $h_{i,\tau} \in \mathcal{H}_{i,\tau}$ it is the case that,

$$\sigma^\star{}_i(h_{i,\tau}) = \alpha v_{i,\tau} \tag{2}$$

Observe that if the bidders stick to the advised strategy, the highest valuation bidder will be the winner in each of the N periods. Moreover, by fixing α close to zero, the payment to the auctioneer will be close to zero — the reserve price in our model. At the end of phase C the game moves to phase S where the players conduct a statistical test.

Statistical Test Phase S: In this phase each player conducts a statistical test to verify whether the other player has adhered to the advised strategy in phase C. The critical part here is the construction of the test function. Suppose that player i is conducting the test on j. Fix a $K \in I\!N$, $K < \infty$ and divide the interval $[0, 1]$ into K sub-intervals $\{[x_k, x_{k+1}]\}_{k=0}^{K-1}$ where $x_k = \frac{k}{K}$. Define $p_i(k, k+1, \sigma_i, \sigma_j)$ to be the number of times player i won when he bid in the interval $[\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}]$ given that the players i and j are using strategies σ_i and σ_j respectively. Similarly let, $m_i(k, k+1, \sigma_i)$ to be the number of times player i bid in the interval $[\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}]$ while using the strategy σ_i. Then the test function that player i employs to test whether player j has conformed to the advised strategy is defined as,

$$t_{ij}(\sigma_i, \sigma_j) = \max_{k \in K} \left[\left| \frac{k + 1/2}{K} - \frac{p_i(k, k+1, \sigma_i, \sigma_j)}{n_i(k, k+1, \sigma_i)} \right| \frac{n_i(k, k+1, \sigma_i)}{N} \right] \tag{3}$$

The statistical test measures a (weighted) distance between the proportion of times a player won when he bid in a certain segment and the theoretical proportion of times he should win, assuming that the opponent bids according to σ_j. A player will fail the test if this distance is too large, i.e., if t_{ij} is above some thresh-hold.

The Communication Phase: After the statistical test the game moves to the communication phase where the players communicate the results of the statistical test to their opponents. The communication phase is for two periods C_2^1 and C_1^2. In period C_j^i player j communicates to player i whether the latter has passed the statistical test or not. Player i's advised action in the period C_j^i is,

$$\sigma_i^\star() = \emptyset \tag{4}$$

For player j, the advised strategy in period C_j^i is,

$$\sigma_j^\star(.) = \begin{cases} \emptyset, & \text{if agent } i \text{ has failed } j\text{'s test;} \\ 0, & \text{otherwise.} \end{cases} \tag{5}$$

However, player i can deviate and bid any $b \in [0,1]$ and win the round. Thus there are three possible publicly observable signals:

- s_i = player i won the round
- s_j = player j won the round
- s_0 = no body won the round.

If the signal is s_0 player i's punishment phase follows after the communication phase for player j. If the signal is s_j, player i has passed player j's test. The signal s_i denotes a deviation by player i from the advised strategy and leads to the one-period NE forever.

In period C_i^j the roles of the two players are reversed.

Punishment Phases: If both agents have passed the statistical test, the play returns to the cooperation phase. If only one player, say player 1, has failed the statistical test, then punishment phase P_1 ensues that last for M periods. If both players have failed their statistical tests, then two punishment phases, P_1 and P_2, take place. The strategy of the punished player is a mixed strategy. During, for example, the punishment of player 1, in the first period of the punishment phase, player 1 (the punished player) conveys a message to the punishing player, player 2, through the sum he bids. If player 1 lets player 2 win the first bid of the punishment phase, then player 2 will be allowed to win all bids in the remaining periods of the punishment phase (and player 1 will not participate). If player 1 wins the bid, then both players will not participate in all those remaining periods. Which message to convey is decided by a randomization conducted by the player 1, with probability p to allow the opponent to win the remaining periods.

The strategies of the two players are defined below. The strategy of player 1 is,

$$\sigma_1^\star(.) = \begin{cases} 0, & \text{with probability } p \text{ in the first period of punishment;} \\ 1, & \text{with probability } (1-p) \text{ in the first period of punishment;} \\ \emptyset, & \text{for the remaining } M-1 \text{ periods.} \end{cases} \tag{6}$$

The strategy of the punishing player (in this case player 2) is,

$$\sigma_2^\star(.) = \begin{cases} \alpha, & \text{in the first period of player 1's punishment phase;} \\ 0, & \text{in the remaining } M-1 \text{ periods if player 2 won in the first period;} \\ \emptyset, & \text{otherwise.} \end{cases} \tag{7}$$

If both players fail the statistical test, the punishment phase P_1 is followed by a second punishment phase P_2 for the second player. Suppose that player 2 has also failed the statistical test. Then after player 1 is punished in phase P_1 it is the turn of player 2. The punishment phase P_2 also lasts for M periods. The strategies of the players are symmetric to those above, only the probability with which player 2 randomizes between the two options, (either that player will win the remaining bids of the punishment phase, or that non of the players will participate in those bids) is probability q, which is not equal to p. The values of p and q will be defined later.

To conclude, there are six parameters defining the scheme:

α - the (small) constant by which the players multiply their values when they bid.

N - the length of the cooperation phase

K - the number of sub-intervals to which we divide the interval $[0, 1]$

M - the length of the punishment phase

p and q - probabilities for the punished player to randomize with, at the beginning of the punishment phase.

3 ϵ-Consistent Equilibrium

When we consider games with private monitoring and a continuum of actions at each period, the "perfectness" concept which applies is that of Perfect Bayesian.[1]

Denote the game by

$$G = \langle N = \{1, 2\}, \{\mathcal{H}_i\}, \{\Theta_i\}, \{\Sigma_i\}\{\pi_i\}\rangle$$

Here $\mathcal{H}_i$ is the set of history for player i, Θ_i is the set of types, Σ_i the set of strategies and π the pay off function. The histories are private. So player i's conditional beliefs at stage t is denoted by $\mu_i(.|h_i^t)$. In our case h_i^t contains all the relevant information up to stage t including realization of player i's type in stage t. Given that in our model actions are not observable, but consequences are, the conditional probabilities are on the consequences i.e., on the public history of the game. Now given a action pair $a = (a_1, a_2)$, a stage t and a public history h_0^{t-1}, $c(h_0^{t-1}, a)$ is the unique concatenation. We denote the belief profile by $\mu = (\mu_1, \mu_2)$

Definition 1: An assessment (μ, σ) is reasonable if for all history profiles $h^t = (h_1^t, h_2^t)$:

(1) Bayes Rule is used to update beliefs whenever possible: i.e. Given any stage t and any history $h^t = (h_1^t, h_2^t)$, and any action pair a, and public history $(h_0^{t-1}$

[1]Sequential Equilibrium requires convergence of beliefs over histories, thus one would need to go into details of convergence over histories with continuum of actions.

that is *compatible* with h^t,

$$\sigma(a|h^t) = \mu(c(h_0^{t-1}, a)|h^t)$$

DEFINITION 2: A PBE of a game defined above is an assessment satisfying

1) (μ, σ) is reasonable

2) for each period t and history profile h^t, the continuation strategies $\sigma(.|h^t)$ are a Bayes Nash Equilibrium given the beliefs, $\mu(.|h^t)$.

For reasons that we will detail later, we use epsilon-equilibrium, i.e., the players' strategies are epsilon-best responses, given their beliefs. We can use, however, the more restrictive form of epsilon-equilibrium - ϵ-consistent equilibrium (Lehrer and Sorin, also known as contemporaneous epsilon equilibrium (Mailath *et. al.*)). In such an equilibrium the "allowed loss" of epsilon is weighted always from the period being played now and forward. So, for example, there is no period "far enough in the future" such that from that period on the players can be instructed to play anything since the weight (of this "future") is insignificant.

DEFINITION 3: An ϵ-Consistent Bayes Equilibrium of a game defined above is an assessment satisfying

1) (μ, σ) is reasonable

2) for each period t and history profile h^t, the continuation strategies $\sigma(.|h^t)$ are ϵ-Consistent Bayes Equilibrium given the beliefs, $\mu(.|h^t)$.

4 Main Result

In this section we derive the main result of the paper. First we show that by playing according to the advised strategy any player can pass the statistical test with a very high probability. As noted before, there are six parameters defining the scheme, $D = \langle \alpha, N, K, M, p, q \rangle$. In the following lemma we show that if player

j is following the advised strategy, the probability that he will fail player i's test goes to zero as the length of the cooperation phase increases.

The proof of Lemma 4.1 appears in the appendix. It relies on Blackwell's Approachability Theorem (Blackwell(1956))

Lemma 4.1 *For all $\eta > 0$ and $\epsilon > 0$ there exist K' and M' such that for all $K > K'$ and $N > N'$ and when player j is using the* ***advised strategy*** *${\sigma^\star}_j$, we have*

$$Prob(t_{ij}(\sigma_i, {\sigma^\star}_j)) > \eta) < \epsilon$$

We want to show that the strategy profile $\sigma^\star = ({\sigma^\star}_1, {\sigma^\star}_2)$ is a perfect ϵ-equilibrium. To that end we first start with the Punishment phase.

4.1 The Punishment Phase

We first consider the case where only one player (say player 1) has failed the test. During the punishment phase player 1 will receive nothing and in the first of his $M(\delta)$ periods of punishment he will randomize over whether player 2 will win the entire set of bids during player 1's punishment phase or not. Any profitable deviation during the punishment phase is publicly detected with probability 1. Following such a deviation the players will move to playing the one period NE for ever.

Both players announce the result of their statistical tests, and then the punishments take place. In case no body failed the test, the players restart on the equilibrium path. Assume that player 1 failed the test. We need to show that there exists a $p^\star$ such that when player 1 (the punished player) is randomizing over whether player 2 will win the bids during the punishment or not with probability $p^\star$, player 2 is indifferent between punishing and not punishing, when we disregard the period which consists the statistical test.

We prove the following two lemmata and the resulting proposition 4.1 in the appendix.

Lemma 4.2 *There exists a* $p^\star \in (0, 1)$ *such that player 2 is indifferent between punishing and not punishing (disregarding the statistical test phase)*

PROOF: (See Appendix) ■

Now we consider the case where both players fail the test. For the player who is punished first (in this case player 1) should have the same continuation payoff in case he is punishing player 2 and in case he is not. This is the same as before, (Lemma 4.2). The tricky part is when it is turn for the player who is punishing first and being punished later (player 2 in this case), to convey his message.

A player who is punishing first and being punished later can use punishing in order to delay his own punishment. Keeping the same $p^\star$ as before will make the first player favor punishing in any case when he believes that there is a positive probability that he failed his statistical test. We will choose a different probability $q^\star$, that player 1 will randomize with to decide whether player 2 will win the bids during player 2's punishment, in case both players fail the statistical test. Note that such randomization takes place during the punishment phase when the results of the statistical tests are already known. Using the different probabilities in the different situations will make player 2 indifferent to punishing whether he failed or passed the test, and hence indifferent for every belief he has regarding failing the statistical test.

Lemma 4.3 *Suppose that both players have failed the statistical tests. Suppose that player 2 is to be punished second. Also suppose that both players are following the advised strategies. Then there exists a $q^\star$ such that player 2 is indifferent between punishing and not punishing.*

Proof: (See Appendix) ■

Proposition 4.1 *At the punishment phase, the advised strategy pair $(\sigma^\star{}_1, \sigma^\star{}_2)$ are mutually ε- best responses.*

Proof: First we show that the punishing player will not gain by deviating from the advised strategy. From lemma 4.2 and 4.3 it follows that the punishing player is indifferent between punishing and not punishing. A punishing player can benefit by deviating from the advised strategy at the punishment phase, only by bidding when he is not supposed to bid. Except for the first period of punishment, the same is true also of the punished player. Such an action is immediately detected, and the game moves to the one period NE for ever. We need to show that such a deviation is unprofitable. The lowest average continuation payoff for a player is close to $\frac{1}{3}\delta^N$. The highest payoff from one period deviation is $(1-\delta)$. The one period NE is $\frac{1}{6}$. For δ close enough to 1 the following inequality holds:

$$(1-\delta) + \delta\frac{1}{6} < \frac{1}{3}\delta^N \tag{8}$$

We now come to the first period of punishment. In this period, the punished player is advised to bid 1 with positive probability. Given his valuation the player has an incentive to deviate. However, this is only one period, and as the discount factor grows, the weight of this period decreases, and for large enough delta, the gain from deviation will be less than ε.

■

We now move on to the Statistical Test phase. There are two sub-phases: the communication of the statistical test result and the test phase.

4.2 The Communication Phase

Lemma 4.4 *At the communication phase, the advised strategy pair* $({\sigma^{\star}}_1, {\sigma^{\star}}_2)$ *are mutually* ε*-best responses.*

PROOF: At the communication phase there is one player, say player 1, who should communicate the result of the statistical test by bidding zero in case the opponent failed the test, and bidding 1 otherwise. From lemma 4.2 (and lemma 4.3) player 1 (and, in turn, player 2) is indifferent between the two possible continuations. Of course, given his valuation during the communication player 1 is no longer indifferent. However, this is only one period, and as the discount factor grows, the weight of this period decreases, and for a large enough δ will be below ε.

As for player 2 (the player whose statistical-test result is communicated), he is not supposed to bid i.e. bid $\emptyset$. As mentioned before there are three possible public signals s_1, s_2 and s_0. Player 2 by deviating can only change the signal from either s_1 or s_0 to s_2. Such a deviation is publicly detected and hence, it will lead to switching to the one-period Nash equilibrium. ■

Note: At the test phase, the advised strategy pair $({\sigma^{\star}}_1, {\sigma^{\star}}_2)$ are mutually best responses. This is because that during the test phase the players simply calculates whether the opponent passed or failed the test. There are no profitable deviations here.

4.3 The Cooperation Phase

We now move to the crucial cooperation phase. We have shown so far that the advised strategy combination $(\sigma^\star{}_1, \sigma^\star{}_2)$ are mutually ϵ-best responses in the statistical test, communication and punishment phases. We want to show now that in the cooperative phase they are ϵ-best responses.

Fix a K. This implies that the interval $[0, 1]$ has K segments. Consider now the segment $[\frac{k}{K}, \frac{k+1}{K}]$, $k \in \{1, ..., K\}$. Suppose that player j is using the advised strategy $\sigma^\star{}_j$. This implies in each period t, $t = 1, ..., N$ of the cooperation phase, player j is using the strategy,

$$\sigma_{j,t} = \alpha v_{j,t}$$

If player i was using the advised strategy as well, then if we consider the interval $[\frac{k}{K}, \frac{k+1}{K}]$, player i will be bidding on average about $N\left(1 - \frac{k+1}{K}\right)$ times above the interval, about $N\left(\frac{k}{K}\right)$, below the interval and $N(\frac{1}{K})$ times within the interval.

There are two ways in which a player may deviate from the advised strategy - either bidding in the same segment as in the advised strategy (but maybe not bidding the exact sum within that segment), or bidding in a different segment. We will first show that by bidding in the same segment (perturbing the bids within the intervals) a player cannot gain much, and then we will show that a player cannot gain much by bidding in different segments altogether without being punished.

We first show that in the cooperation phase if agent i is perturbing his bids only within the interval, then he cannot gain much. For that we consider the following kind of strategies for agent i: the strategy σ_i agrees with $\sigma^\star{}_i$ in the statistical test, communication and punishment phases and in the cooperation phase is defined as follows: for all $t \in M$,

$$\sigma_{i,t} = \alpha z_{i,t} \tag{9}$$

where $z_{i,t} \in [\frac{k}{K}, \frac{k+1}{K}]$ whenever, $v_{i,t} \in [\frac{k}{K}, \frac{k+1}{K}]$. Let Σ_i^1 be the set of all such strategies, that is the set of strategies where the player is allowed to perturb hid bids within the advised segment.

Lemma 4.5 *For any strategy $\sigma_i \in \Sigma_i^1$, it is the case that*

$$\pi_i(\sigma_i, {\sigma^\star}_j) - \pi_i({\sigma^\star}_i, {\sigma^\star}_j) \leq \frac{1}{2K}$$

PROOF: Suppose that player i is using a strategy $\sigma_i \in \Sigma_i^1$. For any $t \in \{1, ..., M\}$, the stage game payoff for player i, given that player j is using the advised strategy ${\sigma^\star}_j$ is

$$r_i(\sigma_{i,t}, {\sigma^\star}_{i,t}) \leq \sum_{k=0}^{K-1} \left[\frac{1}{K} \frac{k+1}{K} - \alpha \frac{k+1}{K} \right] \tag{10}$$

Therefore,

$$r_i(\sigma_{i,t}, {\sigma^\star}_{j,t}) - r_i({\sigma^\star}_{i,t}, {\sigma^\star}_{j,t}) \leq \sum_{k=0}^{K-1} \left[\frac{1}{K} \frac{k+1}{K} - \frac{1}{K} \frac{k+\frac{1}{2}}{K} \right] \leq \frac{1}{2K}$$

The $r_i(,)$'s are the stage game payoffs. The total payoff over the cooperation phase is the discounted sum of the stage game payoffs i.e., the total payoff is the weighted average of the stage game payoffs. Hence, the result follows. ■

We now show that if a player moves "too many" of his bids above a certain segment (meaning, he will bid in one of the segments above the advised one), then he will fail the test with high probability.

Given any interval $[\frac{k}{K}, \frac{k+1}{K}]$, if player i moves at most $1.5/2K$ times above the border $\frac{k+1}{K}$, then he will be bidding at most $N\left(1 - \frac{k+0.5}{K}\right)$ times over this segment when his valuation v_i is in the segment $[\frac{k}{K}, \frac{k+1}{K}]$. He will be bidding at least $N\left(\frac{k-0.5}{K}\right)$ below the segment.

If agent i moves more than, say, $2/K$ above the border then he will be bidding at least $N\left(1-\frac{k-1}{K}\right)$ times above the segment. Let Σ_i' be the set of all such strategies for i. In other words, Σ_i' is the set of all such strategies σ_i with the property that there exists one segment $[\frac{k}{K}, \frac{k+1}{K}]$ such that player i under strategy σ_i bids at least $N\left(1-\frac{k-1}{K}\right)$ times above the segment.

Lemma 4.6 *Suppose that j is bidding according to the advised strategy. Given any interval $[\frac{k}{K}, \frac{k+1}{K}]$, $k \in \{0, ..., K-1\}$, the probability that j bids less than $\frac{N}{2K}$ times in that interval approaches 0 as $N \to \infty$*

Proof: Let $X^j_{k,k+1}(\sigma^\star{}_j) =$ the number of times agent j bids within the segment $[\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}]$. Now, $X^j_{k,k+1}(\sigma^\star{}_j) \sim Bin(N, \frac{1}{K})$. Using Chebyshev's inequality, we get,

$$P\left[\left|X^j_{k,k+1}(\sigma^\star{}_j) - \frac{N}{K}\right| \geq \frac{N}{2K}\right] \leq \frac{N\frac{1}{K}(1-\frac{1}{K})}{(\frac{N}{2K})^2} \tag{11}$$

As N gets large, the right hand side in (13) goes to zero. ■

Suppose that agent j did bid at least $\frac{N}{2K}$ times in the segment $[\frac{k}{K}, \frac{k+1}{K}]$. Suppose that player i is using a strategy $\sigma_i \in \Sigma_i'$. The next proposition shows that, if j is using the advised strategy, the probability that j will end up with a history such that i *will not* fail the test approaches zero.

Let $Y^j_{k,k+1}(\sigma_i, \sigma^\star{}_j)$ be the number of times j won when bidding in the interval $[\frac{k}{K}, \frac{k+1}{K}]$, given that the strategy profile is $(\sigma_i, \sigma^\star{}_j)$. Observe that, $Y^j_{k,k+1}(\sigma_i, \sigma^\star{}_j) \sim Bin(n, \frac{k-1}{K})$, where n is the number of times bidder j bid in the interval. From the previous lemma it follows that with a probability that approaches 1, $n \geq \frac{N}{2K}$.

In order for bidder i to pass bidder j's test, bidder j need to win at least $n\left(\frac{k-0.5}{K}\right)$ times.

Proposition 1: For all $\epsilon > 0$, there exists N_ϵ, such that for all $N > N_\epsilon$,

$$\text{Prob}\left(Y^j_{k,k+1}(\sigma_i, \sigma^\star{}_j) \geq n\left(\frac{k-0.5}{K}\right)\right) < \epsilon$$

PROOF: Observe that K is fixed. Now from Chebyshev's inequality we get that,

$$\begin{aligned}\text{Prob}\left[Y^j_{k,k+1}(\sigma_i,\sigma^\star{}_j) \geq n\left(\frac{k-0.5}{K}\right)\right] &= \text{Prob}\left[Y^j_{k,k+1}(\sigma_i,\sigma^\star{}_j) - n\left(\frac{k-1}{K}\right) \geq \frac{0.5n}{K}\right] \\ &\leq \text{Prob}\left[\left|Y^j_{k,k+1}(\sigma_i,\sigma^\star{}_j) - n\left(\frac{k-1}{K}\right)\right| \geq \frac{0.5n}{K}\right] \\ &\leq \frac{\frac{k-1}{K}\frac{K-k+1}{K}}{n(\frac{0.5}{K})^2}\end{aligned}$$

Observe that, fixing K as N goes up, so does n. Now for any $\epsilon > 0$ fix N_ϵ such that $\frac{\frac{k-1}{K}\frac{K-k+1}{K}}{n(\frac{0.5}{K})^2} < \epsilon$. This completes the argument. ■

All the previous arguments hold when the number of periods in the cooperation phase, N is large. However, as N goes up, the weight of any one period in the cooperation phase is going down. It may appear therefore, that an agent may deviate a few times in the cooperation phase, and gain in terms of payoffs, but still pass the statistical test since N is large. However, as N goes up, for any fixed K we can raise the discount factor δ such that the weight of the periods remain constant. Therefore more periods in the cooperation phase (i.e., before the test) will not mean gaining more by deviating. These ideas are formalized in the following lemma:

Lemma 4.7 *Consider any $\delta_0 \in (0,1)$ and N and N' such that $N' > N$. Consider any $t \leq N$. The weight of period t is $\frac{\delta^t}{1-\delta^N}$. Then there exists a δ_1 such that*

$$\frac{\delta_0^t}{1-\delta_0^N} = \frac{\delta_1^t}{1-\delta_1^{N'}} \tag{12}$$

PROOF: Define the function $g(.)$ as follows:

for all $\delta \in [\delta_0, 1]$,

$$g(\delta) = \frac{\delta_0^t}{1-\delta_0^N} - \frac{\delta^t}{1-\delta^{N'}}$$

The function $g(.)$ is continuous in δ. Moreover, $g(\delta_0) > 0$ and $\lim_{\delta \to 1} g(\delta) < 0$. This implies that there exists a $\delta_1 \in (\delta_0, 1)$ such that, $g(\delta_1) = 0$. ■

We are now ready to show the following theorem:

THEOREM **4.1** *For $\epsilon > 0$, there exists $\delta' \in (0,1)$ and integers N' and K' such that for all $\delta \in (\delta', 1)$, $N \geq N' \, K \geq K'$, and for all $i \in \{1,2\}$ and for all $\sigma_i \in \Sigma_i$*

$$\Pi_i(\sigma_i, \sigma^\star{}_j) - \Pi(\sigma^\star{}_i, \sigma^\star{}_j) < \epsilon$$

We first prove the following intermediate step:

STEP 1: Consider a strategy $\sigma_i \in \Sigma_i'$. Then,

$$\Pi_i(\sigma_i, \sigma^\star{}_j) < \Pi(\sigma^\star{}_i, \sigma^\star{}_j)$$

PROOF: Suppose, player i is using strategy $\sigma_i \in \Sigma_i'$. As shown in proposition 1, if player j is using the advised strategy, player i's probability of punishment is close to 1. Let $p_i(N, \sigma_i, \sigma^\star{}_j)$ be the probability that player i passes the statistical test using strategy σ_i. Note that $p_i(N, \sigma_i, \sigma^\star{}_j) \to 0$ as $N \to \infty$. With a slight abuse of notation, we drop the arguments in $p_i(.,.,.)$ and simply denote by p_i. Let M denote the length of the punishment period. Now,

$$\sup_{\sigma_i \in \Sigma_i'} \Pi_i(\sigma_i, \sigma^\star{}_j) = (1 - \delta^N) + (1 - p_i)\delta^{N+M} X + p_i \delta^N X \tag{13}$$

In equation (15), X is the pay-off to each player over the cooperation phase, if they are both following the advised strategy. Note that $X \sim \frac{1}{3}$. If both players are following the advised strategy, pay-off to player i is,

$$\Pi(\sigma^\star{}_i, \sigma^\star{}_j) = (1 - \delta^N)X + (1 - \epsilon)\delta^{N+M} X + \epsilon \delta^N X \tag{14}$$

The difference between the two payoffs is,

$$Z = (1 - \delta^N)(1 - X) + (1 - p_i - \epsilon)[\delta^{N+M} X - \delta^N X] \tag{15}$$

To show $Z < 0$, as $\delta \to 1$ is equivalent to showing,

$$\frac{M}{N} > \frac{1 - X}{X(1 - p_i - \epsilon)} \tag{16}$$

Taking $M > \frac{N}{2}$ does the job. ∎

Proof Of Theorem 4.1: Consider any $\epsilon > 0$. Fix K such that,

$$\epsilon > \frac{1.5}{2K} \tag{17}$$

From step 1 it follows that player i is not going to use a strategy $\sigma_i \in \Sigma_i'$. Therefore, he can only use a strategy $\sigma_i \in \Sigma_i \setminus \Sigma_i'$. But from lemma 3.6 it follows that, with such a strategy,

$$\Pi_i(\sigma_i, {\sigma^\star}_j) - \Pi({\sigma^\star}_i, {\sigma^\star}_j) \leq \frac{1.5}{2K} \tag{18}$$

Since, $\epsilon > \frac{1.5}{2K}$, we have the result. ∎

5 Public Strategies

In this section we will show that relying only on *public signals* will lead to a situation where the cartel's payoff is bounded away from the efficient payoff.

As mentioned in section (2), at any stage t, the public history $h^p(t)$ consists of the sequence of identities of the winners.

Skrzypacz and Hopenhayen (2004) show that, when the public history of the game consists of sequence of identities of winners in the previous rounds, the set of equilibrium payoffs for the game is bounded away from the efficiency frontier. The following result is similar in spirit to proposition 1 in Skrzypacz and Hopenhayen.

For any given $\delta > 0$, let $\Pi_i^\delta(\sigma^\star{}_1, \sigma^\star{}_2)$ be the payoff to player i from the strategy profile $(\sigma^\star{}_1, \sigma^\star{}_2)$. Let $\Pi^\delta(\sigma^\star{}_1, \sigma^\star{}_2) = \Pi_1^\delta(\sigma^\star{}_1, \sigma^\star{}_2) + \Pi_i^\delta(\sigma^\star{}_1, \sigma^\star{}_2)$. For any $\epsilon > 0$ there exists $\delta' > 0$ such that for $\delta \in (\delta', 1)$, $(\sigma^\star{}_1, \sigma^\star{}_2)$ is a ϵ-equilibrium and $\Pi^\delta(\sigma^\star{}_1, \sigma^\star{}_2) \in (\frac{2}{3} - \varepsilon, \frac{2}{3}]$.

With a slight abuse of notation, let us denote by $\mathcal{H}^p$ to be the set of public histories. Let $(\beta_1, \beta_2) : \mathcal{H}^p \times \mathcal{H}^p \to B \times B$ be a candidate public strategy profile such that for the given $\epsilon > 0$, and for a fixed δ, (β_1, β_2) is an ϵ-equilibrium. Given ϵ and δ,

let $B(\epsilon, \delta) = \{$ the set of (β_1, β_2) that are epsilon equilibria given the $\delta\}$. Let

$$\hat{V} = \sup_{\beta_1, \beta_2 \in B(\epsilon, \delta)} \Pi(\beta_1, \beta_2)$$

In the following we show that $\hat{V} < \frac{2}{3} - \varepsilon$. In order for $\hat{V}$ to be close to $\frac{2}{3}$ the following has to be true : over a weight $\hat{q}$ of the periods, $\hat{q}$ approaches 1, game, with a probability $\hat{p}$ that is close to 1, at each stage the player with the higher valuation must win. We will show that with the limited information, provided by public history alone, that is not possible.

Any public history is a sequence data on the identity of the winners. Given that, at any stage a player has 2 possible actions – $\{\emptyset, b \geq 0\}$ where as before $\emptyset$ means non-participation. If at any stage t a player decides to give up his chance to win the object he will either participate and bid 0 or he will not participate. The next lemma shows that for the scheme to be efficient, the probability of non-participation needs to go down to zero.

Lemma 5.1 *In an efficient framework, over a weight which approaches 1 of the periods the probability of Non-participation by any player approaches 0.*

Proof: Assume that the probability of Non-participation is $p > 0$, over a significant weight of the periods (a weight that does not approach zero) This implies that there exists a range of valuation, the lowest of which can be $[0, p]$ in which the player will not participate. With probability p the opponent will also have his valuation in the range $[0, p]$. So with probability $p^2/2$ the scheme is not efficient, because the bidder with the highest valuation will not win the bid. ∎

In a scheme with public strategy when both players are participating, there are only two publicly observable signals. Either player 1 won or player 2 won. Hence, given any history, with a probability that approaches 1, there are only two possible continuation payoffs for each player : the continuation payoff in case he wins the next round and the continuation payoff in case the opponent won.

Observe that, for any given public strategy pair (β_1, β_2), for $\Pi(\beta_1, \beta_2)$ to be close to $\frac{2}{3}$, it must be the case that, with a probability $q \to 1$, each player's bid b is close to zero. Otherwise the scheme is not efficient. In other words, given lemma 5.1, in order for a scheme to obtain efficient outcome, with probability that approaches one both players participate and bid at most ω, and ω should approach zero.

Let us define W to be the continuation payoff of a player who wins the next round and P to be the continuation payoff in case the opponent won. In the next lemma we will show that the probability that the difference between winning now and getting W in the future, and loosing now and getting P in the future is at most ϵ tends to zero.

Lemma 5.2 *The probability p that for any player, the absolute difference between winning now and getting W in the future, and loosing now and getting P in the future, is less than ϵ tends to zero i.e., for all $\omega > 0$ there exists $\eta > 0$ such that*

$$Prob\,(|(1-\delta)(v-b) + \delta W - \delta P| < \omega) < \eta$$

and both ω and η

PROOF: Roughly, if the difference is more than omega, then deviating by bidding more than omega would be profitable for large enough valuations ■

Observe that there may be one remaining case where one player does not participate. Since the probability of such an event is going to zero and since the continuation payoff is bounded by 1, the range of valuations for which a player is indifferent between winning and participating still goes to zero.

THEOREM **5.1** *For any discount factor $\delta < 1$, in any ϵ-PPE, the average payoffs of the players are bounded away from the efficiency frontier*

PROOF: From lemma 4.3 it follows that with a probability $\mu \to 1$, either (i) a player strictly prefers winning and paying ω over participating, or (ii) a player strictly prefers participating and paying 0.

Observe that if a player prefers to win for a certain valuation then he also prefers to win for all valuations above it. Let R_1 be such that player 1 prefers to win when his valuations are in $[R_1, 1]$ and define R_2 analogically. Observe that the range of valuations where player i strictly prefers to loose approaches $[0, R_i)$. There are four cases which will be considered separately.

CASE I: $R_1 = R_2 = R < 1$. In this case, when $v_1, v_2 < R$, both players participate and bid 0. So with a probability $R^2/2$, it will not be the player with the highest valuation who will win the object.

CASE II: $R_1 < R_2 < 1$. In this case, with a probability $\frac{1}{2}(R_2 - R_1)$ player 2 will have the higher valuation, but he strictly prefers to participate, while player 1 strictly prefers to win. So the object will not be awarded to the highest valuation player.

CASE III: In case at least one of the R_is equals 1 and the other is less than 1, say $R_1 = 1,\ R_2 < 1$ then with probability $R_2(1 - R_2)$ it is the case that player 1 has the higher valuation, but he prefers to only participate.

CASE IV: The last remaining case to consider is where both $R_1 = R_2 = 1$. In this case again, both player prefer to participate. Hence the good is given "randomly" and not, with probability 1, to the player with the highest valuation.

The four cases together completes the proof of the theorem. ■

6 References

Aoyagi, M. (2002): "Bid rotation and collusion in repeated auctions", *Journal of Economic Theory*, **112**, 79-105.

— (2003): "Bid rotation and collusion in repeated auctions", *Journal of Economic Theory*, **112**, 79-105.

Blackwell, D. (1956): "An analog of the minimax theorem for vector payoffs", *Pacific Journal of Mathematics*, **6**, 1-8.

Ely, J. and J. Valimaki (2002): "A Robust Folk Theorem for the Prisoner's Dilemma", *Journal of Economic Theory*, **102**, 84-105 .

Fudenberg, D., D. Levine and E. Maskin (1994): "The Folk Theorem with Imperfect Public Information", *Econometrica*, **58**, 1041-1063.

Gossner, O. (1995): "The Folk Theorem for Finitely Repeated Games with Mixed Strategies", *International Journal of Game Theory*, **24**, 95-107 .

Graham, D.A and R.C. Marshall (1987): "Collusive Bidder Behavior at Single-Object Second-Price and English Auctions", *Journal of Political Economy*, **95**, 1217-1239 .

Kandori, M. and I. Obara (2003): "Efficiency in Repeated Games Revisited: The Role of Private Strategies", *UCLA Economics Working Papers*, **826**.

Lehrer, E. and S. Sorin (1998): "ϵ-Consistent Equilibrium in Repeated Games", *International Journal of Game Theory*, **27**, 231-244.

Mailath J., A. Matthews, and T. Sekiguchi (2002): "Private Strategies in Finitely Repeated Games with Imperfect Public Monitoring", *Contributions to Theoretical Economics* Vol. 2: No. 1, Article 2.

Mailath J., A. Postlewaite, and L. Samuelson (2003): "Contemporaneous Perfect Epsilon-Equilibria, Second Version", *PIER Working Paper* No. 03-021.

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

Mertens, J.-F, S. Sorin and S. Zamir (1994): "Repeated games", *CORE discussion paper*, **9420** .

Piccione , M (2002): "The Repeated Prisoner's Dilemma with Imperfect Private Monitoring", *Journal of Economic Theory*, **102**, 70-83 .

7 Appendix A

7.1 proof of Lemma 4.1

The proof of this lemma uses Blackwell's Approachability Theorem (Blackwell (1956), Gossner (1995), Mertens, Sorin and Zamir (1992)).

Let k such that player i bid in period n in the segment $[\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}]$. Given any $n \in \{1, \cdots, N\}$ we consider a vector $x_n \in \Re^K$ such that for all $k \in K$

$$x_n^k = \begin{cases} \frac{k+\frac{1}{2}}{K} - 1, & \text{if } i \text{ bid in the interval } [\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}] \text{ and won;} \\ \frac{k+\frac{1}{2}}{K}, & \text{if } i \text{ bid in the interval } [\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}] \text{ and did not win;} \\ 0, & \text{otherwise.} \end{cases}$$

Let us define $\bar{x}$ to be the average of x_n over N periods, $\frac{\sum_{n=1}^{N} x_n}{N}$. Note that $\bar{x}$ is a vector in $\Re^K$. We denote the k-th element of $\bar{x}$ by $\bar{x}^k$. Observe that $\max_{k \in K} |\bar{x}^k| = t_{ji}(\sigma_i, \sigma^{\star}{}_j)$. Consider now the set,

$$S = \{x \in \Re^K | \max_{k \in K} |x^k| \leq \frac{1}{2K}\}$$

We will show that this set is approachable by player j using the advised strategy no matter what strategy player i is adopting. Consider now player j using the advised strategy. Suppose that player i has played in the segment $[\frac{\alpha k'}{K}, \frac{\alpha(k'+1)}{K}]$. The probability that player i will win is in the interval $[\frac{k'}{K}, \frac{k'+1}{K}]$. This means that, given that player i has played in the segment $[\frac{\alpha k'}{K}, \frac{\alpha(k'+1)}{K}]$, the probability that player i will win is within a distance of $\frac{1}{2K}$ from $\frac{k'+\frac{1}{2}}{K}$.

Given that at stage n, player i has played in the interval $[\frac{\alpha k}{K}, \frac{\alpha(k+1)}{K}]$, (i.e., in the k-th coordinate of x_n),

$$E(x_n^{k'}) = \begin{cases} 0 & \text{if } k' \neq k \\ \frac{k'+\frac{1}{2}}{K} - Prob(\text{player } i \text{ wins}) & \text{if } k' = k \end{cases}$$

This implies that in no coordinate k, the expected value Ex_n^k is greater than $\frac{1}{2K}$. Therefore the expected value of x_n, $E(x_n) \in S$. We can then use Blackwell's theorem to say that the set S is approachable by player j. Thus for any $\eta > 0$, $\epsilon > 0$, we can suitably select K and the corresponding N such that for all σ_i we have,

$$\text{Prob}(t_{ij}(\sigma_i, \sigma^\star{}_j)) > \eta) < \epsilon$$

■

7.2 proof of lemma 4.2

Player 2 is the punishing player. From (6) it follows that for $p = 1$, given that player 2 is following the advised strategy, player 2's payoff for the entire M periods of punishment is 0. For $p = 0$ player 2's payoff is $\frac{1}{2}$. Since the expected payoff of player 2 in the continuation in case he did not punish player 1, it follows that, there exists a $p^\star$ such that player 2 is indifferent between punishing and not punishing.

■

7.3 proof of lemma 4.3

By punishing first, player 2 postpones his punishment by M periods. This results in a gain of $\frac{1}{3}(1-\delta^M)^2$. We now prove that a probability $q^\star$ (with which player 1 randomizes over whether player 2 wins the bids or not during player 1's punishment phase in case both players fail the test) which will make player 2 indifferent to punishing when failing his own test exists. Observe that using $q = p^\star$ as before player 2 gets $\frac{1}{3}(1-\delta^M)$ during player 2's punishment. With $q = 0$ his payoff during player 2's punishment is 0. We wish his payoff to be $\frac{1}{3}(1-\delta^M) - \frac{1}{3}(1-\delta^M)^2$, and since $0 < \frac{1}{3}(1-\delta^M) - \frac{1}{3}(1-\delta^M)^2 < \frac{1}{3}(1-\delta^M)$ we have shown that there exists a $q^\star$ that does the job.

8 Appendix B - General priors

In this section we take a look at general priors. The main issue in this case is to design a suitable statistical test. Let $f(v_1, v_2)$ be the prior joint density function for the valuations of the two players over $[0, 1]$. We focus our attention to the statistical test for player 2 by player 1. The case for player 2 testing player 1 is symmetric. As before, we divide the interval $[0, 1]$ into K segments. Given that player 1's valuation v_1 is in the segment $\left[\frac{k}{K}, \frac{k+1}{K}\right]$ there is the induced conditional distribution over the valuation of player 2 given by,

$$\int_{\frac{k}{K}}^{\frac{k+1}{K}} \frac{f(v_1, v_2)}{F_1(k/K) - F_1(k+1/K)} dv_1 := \bar{f}_2 \left(v_2 | v_1 \in \left[\frac{k}{K}, \frac{k+1}{K}\right] \right) \tag{19}$$

Let $\hat{F}_2(.|.)$ be the induced distribution function. Given K, every segment $\left[\frac{k}{K}, \frac{k+1}{K}\right]$ induces a condition distribution of player 2's valuation. We divide the space of player 2's valuation [0,1] into N possible segments (N not to be confused with the length of the punishment phases). Suppose that player 1 has valuation $v_1 \in \left[\frac{k}{K}, \frac{k+1}{K}\right]$ and that he bids in the interval $b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]$. Assuming that player 2 is following advised strategy, the probability that player 1 will win the particular round is,

$$\int_{\frac{n}{N}}^{\frac{n+1}{N}} \bar{F}_2 \left(t | v_1 \in \left[\frac{k}{K}, \frac{k+1}{K}\right] \right) dt$$

$$:= \text{Prob}\left(1 \;\; wins|v_1 \in \left[\frac{k}{K}, \frac{k+1}{K}\right], b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]\right)$$
$$:= q(1 \text{ wins}|v_1, b_1)$$

Let $p_1(k, n, \sigma_1, \sigma^{\star}{}_2)$ be the number of times player 1 won when his valuation was in the interval $\left[\frac{k}{K}, \frac{(k+1)}{K}\right]$ and he bid $b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]$ where σ_1 is the strategy for player 1 that induces a bid $b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]$ when his valuation is in the interval $\left[\frac{k}{K}, \frac{(k+1)}{K}\right]$.

Let $m_1(k, n, \sigma_1)$ be the number of times player 1 had a valuation in $\left[\frac{k}{K}, \frac{(k+1)}{K}\right]$ and bid in the interval $b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]$. The test statistic is,

$$t = \max_{k \in K, n \in N} \left| \frac{[q(1\text{wins}|v_1, b_1) m_1(k, n, \sigma_1) - p_1(k, n, \sigma_1, \sigma^{\star}{}_2)]}{M} \right| \tag{20}$$

Observe that, $q(.|.,.)$ is computed on the basis of $\bar{f}_2(.|.)$. Given any $v_1 \in [0, 1]$, the actual conditional density function is $f_2(v_2|v_1)$. Suppose that individual 1 has valuation v_1 and bids $b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]$. Let us define $\dot{q}$ as follows:

$$\dot{q}(.|.,.) := \text{Prob}\left(1 \;\; wins|b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]\right)$$
$$= \int_{\frac{n}{N}}^{\frac{n+1}{N}} F_2\left(t|v_1\right) dt \tag{21}$$

It is to be noted that, if player 2 is following the advised strategy $\sigma^{\star}{}_2$, $\dot{q}(.|.,.)$ is player 1's actual probability of winning whenever player 1 is bidding $b_1 \in \left[\frac{\alpha n}{N}, \frac{\alpha(n+1)}{N}\right]$. Now we need the statistical test to become more and more accurate as we increase K and N. In other words we need $q(.|.,.)$ to be "close" to $\dot{q}(.|.,.)$ as we increase K and N. Requiring that is imposing the following condition on the distribution function:

Condition X: The distribution function $F(.|.)$ is said to be *uniformly continuous* over the interval $[0, 1]$ if for every $x \in [0, 1]$, for every $\xi > 0$, there exists, $\gamma > 0$

such that

$$|F(t|x) - F(t|x-\gamma, x+\gamma)| < \xi \tag{22}$$

In other words, as one gets finer and finer in the intervals, the conditional distributions should not be too different, and this has to hold *uniformly* over the entire interval over which a players valuation may lie i.e., over the set $[0, 1]$. Observe that if condition X is satisfied, $q(.|., .)$ converges uniformly to $\dot{q}(.|., .)$ It is to be noted that, Condition X is not required if the valuations of the players are independently distributed.

Lemma 8.1 *Suppose condition X is satisfied. Then for all $\eta > 0$ and $\epsilon > 0$ there exist $K^\star$ and $N^\star$ such that for all $K > K^\star$ and $N > N^\star$ there exists $M^\star$ such that for all $M > M^\star$,*

$$Prob(t > \eta) < \epsilon$$

PROOF: For any $l \in \{1, \cdots, M\}$ define,

$$x^l_{n,k} = \begin{cases} q(1 \text{ wins}|v_1, b_1), & \text{if player 1 lost under } v_1, b_1; \\ q(1 \text{ wins}|v_1, b_1) - 1, & \text{if player 1 won under } v_1, b_1; \\ 0, & \text{otherwise.} \end{cases} \tag{23}$$

It is to be noted that $x^l_{n,k}$ is a vector in $\Re^{NK}$. Let us define $\bar{x} = \frac{\sum_{l=1}^{l=M} x^l_{n,k}}{M}$. Note that $\bar{x}$ is also a vector in $\Re^{NK}$. We denote by $\bar{x}_{n,k}$ the (n, k)-th element of $\bar{x}$. It is then straightforward to note that

$$t = \max_{k \in K, n \in N} |\bar{x}_{n,k}|$$

Fix a $d(K, N) > 0$. Let $T = \Big\{x \in \Re^{NK} |\max_{k \in K, n \in N} |x_{n,k}| \leq d(K, N)\Big\}$. Observe that T is a closed convex set. By choosing K and N large enough we can choose $d(K,N)$ as close to zero as possible. Now from Condition X we know that as we increase K and N, $q(.|.,.)$ converges uniformly to $\dot{q}(.|.,.)$. The rest of the argument follows from the Approachability Theorem. ∎