

KIER DISCUSSION PAPER SERIES

KYOTO INSTITUTE OF ECONOMIC RESEARCH

Discussion Paper No.1055

“Full Collusion with Entry and Incomplete Information”

Ramakanta Patra
Tadashi Sekiguchi

February 2021



KYOTO UNIVERSITY
KYOTO, JAPAN

Full Collusion with Entry and Incomplete Information*

Ramakanta Patra [†]
Tadashi Sekiguchi [‡]

February 4, 2021

Abstract

This paper studies an infinitely repeated duopoly game with incomplete information and with costly entry decisions. Every period, each player learns her private type and decides whether to pay a cost in order for her to enter or not. If she enters, she plays a game belonging to a class that includes Bertrand duopoly and some auction games as special cases, either as a monopolist or as a duopolist. The players can communicate before they make their entry decisions. We study full collusion (joint profit maximization) in this environment which requires a higher-quality player to solely enter and to choose an action maximizing the stage payoff. We present a condition on the stage game which is both necessary and sufficient in order for full collusion to be an equilibrium outcome for sufficiently patient players. The condition is more likely to hold when the entry cost increases, which signifies that the entry cost is an important factor facilitating full collusion. We also show that under some parameter restrictions, asymmetric equilibria where only one player reveals her type every period sustain full collusion for a wider range of discount factors. These asymmetric equilibria reduce the total amount of communication, which makes it harder for antitrust authorities to detect collusion.

JEL classification: C73, D43, K21, L0

Keywords: Bertrand Competition; Fixed Costs; Unknown Costs; Private Information; Infinitely Repeated Game; Pre-play Communication; One-sided Communication; Full Collusion

*We thank Laurent Linnemer for stimulating conversations that prompted this research project. We also thank Masaki Aoyagi, Emmanuel Petrakis, and audience at various conferences for their helpful comments. We are grateful to Dengwei Qi for his excellent research assistance. Sekiguchi gratefully acknowledges financial supports from the Grants-in-Aid for Scientific Research (26245024, 26380238, 17K03625, 17H00761, 18K01507). Patra gratefully acknowledges financial support from Cardiff School of Management research fund. The study was supported by the Joint Research Program of KIER, Kyoto University.

[†]*Email:* rpatra@cardiffmet.ac.uk; *Web:* <https://sites.google.com/a/vt.edu/ramakanta-patra/home>; *Address:* Department of Accounting, Economics and Finance, Cardiff Metropolitan University, Cardiff, CF5 2YB, United Kingdom

[‡]*Email:* sekiguchi@kier.kyoto-u.ac.jp; *Address:* Kyoto Institute of Economic Research, Kyoto University, Yoshida-Honmachi, Sakyo-Ku, Kyoto, 606-8501, Japan

1 Introduction

Collusion is a recurring phenomenon and it takes various forms such as price-fixing in homogeneous good markets and bidding rings in procurement auctions. From theoretical standpoint, the repeated games literature has extensively studied conditions, incentives, methods and effects of collusion.¹ One important result from the literature is that, even when oligopolists have their technologies and/or other relevant variables as private information, patient firms can sustain collusive outcomes. For instance, Athey and Bagwell [5, 6] show such results for repeated Bertrand oligopoly, and Aoyagi [2, 3] and Skrzypacz and Hopenhayn [24] for repeated procurement auctions.

These papers in the literature assume that, in any period, oligopolists are always ready to operate in the market or to bid in the auction. In reality, however, ability to be active in each period may not be acquired so easily. In price competition, each firm may incur preparatory expenditures such as costs for maintenance, a recurring per-period licensing fee for operating a business, etc., for this purpose.² In procurements, the auctioneer may only admit firms paying a participation fee. This paper incorporates these aspects of a market and assumes that each firm faces a costly entry decision every period. And only when each firm has paid the entry cost, it is allowed to operate in the market or bid in the auction, after learning the other firm's entry decision. This makes our stage game an extensive-form game.

More concretely, we explicitly formulate a repeated duopoly game with costly entry decisions under incomplete information. Our model both builds on and generalizes the payoff structure of our companion paper by Patra [23], who is the first to consider this type of setting in a Bertrand oligopoly. We assume binary type space, and each player in each period is either high-quality or low-quality. Types are IID across periods and may be correlated across players. In each period, each player decides whether to make a costly entry or not, after learning her type. If she does not enter, she has no further action to choose in that period. If she enters, she learns the other player's entry decision and chooses her action under a payoff environment, which includes Bertrand duopoly and some auction games as special cases, either as a monopolist or as a duopolist.³

A primary question this paper asks is whether full collusion, or total payoff maximization, is an equilibrium outcome in this environment if the players are sufficiently patient. What kind of behavior corresponds to full collusion? If both players enter, it is inefficient because the entry costs are doubly paid. Thus, full collusion requires only one player to enter. Further, we assume that a high-quality solo entrant earns a

¹See Marshall and Marx [20] for a detailed survey.

²These costs correspond to what Sutton [25] classifies as exogenous sunk costs.

³Our formulation on the timing is thus strategically different from existing entry games where an entering firm posts its price before learning the other's decision. See Binmore [9].

greater stage payoff than her low-quality counterpart. Therefore, full collusion requires that a low-quality player solely enters only when the other player is also low-quality, and that the solo entrant chooses an action which maximizes her stage payoff.

To implement this arrangement, players will have to precisely know who possesses higher quality among themselves. Thus, we add a communication stage in each period after the players learn their private types and before they make entry decisions. This stage serves as a channel where the players may reveal their types and decide on who to enter. To summarize, we examine sustainability of full collusion in this repeated game with entry decisions, incomplete information, and the communication stage.

Main results of this paper are as follows. First, we present a necessary condition in the stage game in order for full collusion to be a repeated game equilibrium outcome under some discount factor. If this condition fails, the players cannot sustain full collusion under any level of patience. The point is that a fully collusive equilibrium must prevent any type of a player from profitably deviating by mimicking the behavior of the other type of the player. This is because this mimicry deviation is not detectable; the other player regards this as the equilibrium behavior of the mimicked type. In our model, a key deviation is a low-quality player's mimicry of her high-quality variant. We call this necessary condition "aggregate no-mimicry condition", because this is obtained by summing up the two players' incentive conditions not to mimic.

Secondly, we show that if the aggregate no-mimicry condition holds, sufficiently patient players can achieve full collusion as an equilibrium. In other words, this condition fully characterizes the stage games in which full collusion is an equilibrium outcome under patience. The condition is more likely to hold if the entry cost increases, which implies that the entry cost is a factor facilitating collusion.⁴

Thirdly, whereas our second result is based on a particular equilibrium construction which sustains an equal division of fully collusive payoffs, we also show that under some parameter restrictions some asymmetric equilibria sustain fully collusive outcomes for a wider range of discount factors. The result exemplifies importance of studying asymmetric equilibria within symmetric environments which have wide theoretical applicability in many areas of economics. Further, these asymmetric equilibria make only one player announce her type every period.⁵ This scheme reduces the total amount of communication between the colluding parties and hence, makes it harder for antitrust authorities to detect the collusive conduct.

Our model has certain resemblance with the airlines industry. Generally, airlines incur two types of costs prior to the price competition stage, which correspond to the

⁴However, the larger entry cost reduces the fully collusive payoff.

⁵This one-sided communication works because we have both two players and binary type spaces. If a player truthfully claims herself to be low-quality, then we know that the other player has (weakly) higher quality. If she truthfully claims herself to be high-quality, then we know that she has (weakly) higher quality.

entry costs in our model. One is the costs of gate-lease agreements when an airline plans to begin serving an airport.⁶ The other is the sunk costs the airlines pay to operate a new or an additional route (cost of pairing new destinations). In a recent empirical paper, Aguirregabiria and Ho [1] show that the structure of the airlines industry has a significant effect on the magnitude of those costs. Theoretically, Belford [8] studies a two-stage game where airlines first pay sunk costs and then simultaneously compete with prices. In this setting, she examines how the airlines restructure their network in response to the level of those sunk costs. We also study a duopoly game with an entry stage, but we focus on collusion which is commonplace among the airlines in recent times with many instances of illegal price fixing convictions worldwide (see [20]).

Our stage game is an extensive-form game that adds an entry stage to a duopoly game. Patra [23] first introduced this type of repeated Bertrand competition. He more explicitly studies the equilibria of the one-shot game, and presents a mechanism that supports a (partially) collusive outcome. In addition, Patra [23]’s collusive scheme leads to a welfare improving market outcome over repeated play of a one-shot equilibrium, since in his scheme the market is fully served by the least costly operating firm(s). We instead limit attention to full collusion. Elberfeld and Wolfstetter [13] consider a similar, but finitely repeated Bertrand game where firms first make a decision whether to enter or not by paying a fixed cost before competing in price subsequently. A major difference between our model and that of [13] is that while [13] assumes complete information, we consider incomplete information.⁷

Our methodology to examine sustainability of full collusion is novel in the literature on repeated oligopoly with incomplete information. Many existing papers assume no entry costs, and hence a low-quality firm is active every period and earns informational rents at the expense of high-quality opponents. This destroys incentives to fully collude. In contrast, the equilibrium in our model reveals who should enter due to the communication stage, and a low-quality player who knows that the other player is high-quality does not enter or enjoy informational rents because his continuation payoff upon entry can be designed to be too low to recover the entry cost.

The literature on repeated Bertrand oligopoly with incomplete information (without an entry stage) focus on different types of collusive schemes and related economic phenomena. Athey, Bagwell, and Sanchirico [7] study collusion where marginal cost of homogeneous good Bertrand competing firms are IID. They analyze the relationship between rigid-price and cost shocks both with and without periodic demand shocks. Hanazono and Yang [17] analyze collusive behavior and price rigidity when firms receive private signals by IID shocks affecting the demand side of the market primarily.

⁶In [14], the Federal Aviation Administration of the USA extensively documents on this.

⁷See Wen [26] for a folk-theorem type result in repeated extensive-form games with complete information.

Gerlach [16] has discussed the role of communication in collusion when firms receive IID demand shocks every period.

Communication plays an indispensable role in the construction of a fully collusive equilibrium in our environment. We have two strands of related literature which addresses the role of communication. First, in the context of entry decisions in oligopoly, cheap-talk messages facilitate coordination among potential entrants in one-shot environments under complete information (Dixit and Shapiro [12], Farrell [15]) and also under incomplete information (Park [22]). In the equilibrium we construct, communication identifies a player of higher quality and, in case their qualities are equal, serves as a tiebreaker which decides the sole entrant. Second, in the setting of repeated games, Compte [11] and Kandori and Matsushima [18] study repeated games with communication, but their idea is to elicit players' private observations, not their privately-known types as in this paper.

Another feature of our approach is that we explicitly construct fully collusive equilibria and derive the minimum discount factors sustaining the equilibria. Further, these equilibria have simple forms and may provide various empirical implications. In some repeated games with incomplete information, Olszewski and Safronov [21] also study simple equilibria which approximate efficiency without communication. However, they do not consider extensive-form stage games or full collusion.

The rest of this paper is organized as follows. Section 2 sets up our model. Section 3 presents the main result of this paper, a condition that is both necessary and sufficient for sufficiently patient players to achieve full collusion. Section 4 explores a possibility that full collusion can be sustained for a wider range of discount factors than the range identified in Section 3.

2 Model

2.1 Environment

Two players, player 1 and player 2, interact in discrete time over infinite horizon. Our environment involves incomplete information. We assume binary type space, and hence in each period, each player is of either high-quality or low-quality. We denote the type index by $\theta \in \{L, H\}$ with θ_i representing the quality of player i in a given period. We assume that the type pair of the players, denoted by $(\theta_1, \theta_2) \in \{L, H\}^2$, is IID over time. We also assume that the type pair in any given period is symmetrically distributed. Denoting the probability of a type pair (θ_1, θ_2) occurring in any given period by $\mu_{\theta_1\theta_2}$, our above assumption of symmetry implies $\mu_{HL} = \mu_{LH}$. We also assume $\mu_{\theta_1\theta_2} > 0$ for any (θ_1, θ_2) .

At the beginning of each period, each player has an opportunity to enter a “market” for that period. This entry is costly, and it costs $F_\theta > 0$ for each θ -quality player. If a player enters, she learns the other player’s entry decision and then chooses her action, either as a sole entrant or as one of the two entrants, from a common action set $A = [0, \bar{a}]$, where $\bar{a} > 0$. If she does not enter, she has no further action to choose in that period.

Note that each player must pay the entry cost every period she wants to operate in. In particular, we consider expenditures that are valid only for a single period, not a one time set-up cost incurred by a firm. Examples in oligopolistic environments include market research expenditures incurred at the beginning of a production cycle, fees paid for renewal of license or lease agreements, and sunk part deposit on input purchases that firms incur every period they participate. For simplicity, we assume complete information about the entry cost F_θ . We posit that these entry costs may be different depending on the type of a player but they do not change over time.

We first specify each player’s payoff when the other player does not enter. We assume that the players have an identical payoff function, so that the payoff of a sole entrant of type θ who chooses an action $a \in A$ is $U(a, \theta)$. This is defined as the “gross” payoff and does not count the entry cost. We call U the mono-entrant payoff function and make the following assumptions.

Assumption 1. (i) For each θ , $U(a, \theta)$ is continuous with respect to a .

(ii) $U(a, H) > U(a, L)$, for any $a < \bar{a}$.

(iii) $U(0, H) \leq 0$.

(iv) We define $\Pi_\theta \equiv \max_a U(a, \theta)$ and $\pi_\theta \equiv \Pi_\theta - F_\theta$ for each θ , and we assume $\pi_H > \pi_L > 0$.

Some comments about Assumption 1 are in order. Assumption 1(ii) states that except at the maximum action \bar{a} , a high-quality player earns more than her low-quality counterpart. Together with continuity (Assumption 1(i)), we also have $U(\bar{a}, H) \geq U(\bar{a}, L)$. Further, Assumption 1(i) ensures that the mono-entrant payoff function of each type θ has a maximum. Assumption 1(iv) states that the *mono-entrant payoff*, π_θ , defined by the maximum of the mono-entrant payoff function minus the entry cost, is positive for each type of a player, and that the mono-entrant payoff of a high-quality player is greater than that of a low-quality player. Assumption 1(iii) states that taking action ‘zero’ gives a non-positive payoff to a high-quality player and, together with Assumption 1(ii), it gives a negative payoff to a low-quality player.

Next, we specify the payoffs when both players enter. For all $i \neq j, i, j \in \{1, 2\}$, let $u_i(a_i, a_j, \theta)$ be the payoff of player i of type θ when she chooses a_i and when the other

player chooses a_j . We assume that

$$u_i(a_i, a_j, \theta) = \begin{cases} 0 & \text{if } a_i > a_j, \\ U(a_i, \theta) & \text{if } a_i < a_j, \\ \frac{1}{2}U(a_i, \theta) & \text{if } a_i = a_j. \end{cases}$$

The next subsection presents two economic examples which motivate this particular specification of the payoffs.

2.2 Examples

Bertrand Duopoly One example of our environment is a general class of the Bertrand duopoly. Here each action a corresponds to a price, and the largest action corresponds to a choke price under which the market demand is zero. The underlying market demand function $D(a)$ is continuous and satisfies both $D(\bar{a}) = 0$ and $D(a) > 0$ for any $a < \bar{a}$. The firms are identical, and a θ -type firm has a cost function $C_\theta(q) : [0, \infty) \rightarrow [0, \infty)$. Each C_θ is continuous, and we assume $C_H(q) < C_L(q)$ for any $q > 0$. The mono-entrant payoff function is represented by

$$U(a, \theta) = D(a)a - C_\theta(D(a)).$$

Notice that Assumption 1 is satisfied as long as (iv) holds.

When both firms enter, it is easy to see that the payoffs specified above fit the example if both the firms choose different prices. When they choose a common price, we assume a “winner-take-all” rationing rule. Namely, rather than both firms sharing the total demand equally, one firm is selected with a fair coin-flip and the selected firm serves the entire demand. With this rule, the specified payoffs fit our example. This may be justified when one big consumer poses the entire demands and picks a firm to buy randomly in case of a tie. Another justification for such an assumption is to imagine a situation where consumers may use a search engine for the best price of a product. If the search engine randomly picks a particular order to display the lowest prices in case of a tie, then the consumers watching the same display might buy only from the firm which happens to be displayed first, for instance.

In a special case of linear cost functions (namely, $C_\theta(q) = c_\theta q$ for any q , where $c_L > c_H > 0$), we have

$$\frac{1}{2}U(a, \theta) = \frac{1}{2}D(a)a - C_\theta\left(\frac{1}{2}D(a)\right)$$

for any a . Consequently, we may replace the winner-take-all rationing rule with a

standard, equal division of the demand.

Procurement Auctions It is well-documented that many auctioneers in procurement situations are concerned with attributes other than price.⁸ For example, time to completion is an important consideration in a road construction contract. Let us imagine a somewhat extreme scenario where the auctioneer regards time to completion as the most crucial factor. Thus, the auctioneer simply fixes a contract price and wants to give the contract to a firm that promises the quickest completion. Consequently, the set of actions, A , corresponds to the set of the possible completion times.⁹

In this set-up, we interpret entry cost in two ways. First, it may be a preparatory expenditure the bidder incurs in order to get ready for the auction. Since the bidder must keep her promise once she wins, she may need initial investments, for instance. Secondly, the entry cost may be a fee the auctioneer charges in order for a bidder to participate in the auction. As we can see, the entry cost might be type-dependent in the first interpretation whereas it would be type-independent ($F_H = F_L$) in the second interpretation.

If only one bidder entered, she chooses a completion time $a \in A$, and the auctioneer accepts it. The mono-entrant payoff function is represented by $U(a, \theta) = p - C_\theta(a)$, where p is the contract price and $C_\theta(a)$ is the cost of a completion time a for a θ -type firm. If both bidders entered, they submit their completion times simultaneously, and the bidder with a quicker completion time wins. In case of a tie, a fair coin toss selects a winning bidder.

It is reasonable to assume both that $C_H(a) < C_L(a)$ for any a and that an immediate completion, which corresponds to $a = 0$, is prohibitively costly. Assumption 1 holds if we add continuity, small entry costs, and a large contract price.

For future purposes, we point out that each C_θ may not be monotone, because any completion too soon is costly, *and* any completion too late may also be costly due to prolonged use of resources. If these are the case, the mono-entrant payoff function for each type is hump-shaped and may have an interior maximizer. Further, due to technological differences across the types, the mono-entrant payoff functions of both types may have different maximizers. As we will see, this is one of the keys to our results.¹⁰

⁸They are typical situations for scoring auctions. See Asker and Cantillon [4] for discussions.

⁹This can be seen as a special scoring auction where the score puts all weights on the time to completion.

¹⁰Note the difference from the standard first-price procurement auction, where any bidder of any type would simply prefer a higher procurement price. This makes one price, the largest price the auctioneer is willing to accept, maximize the mono-entrant payoff functions of both types.

2.3 Repeated Game

We study a class of repeated games where, given the environment specified above, players make entry decisions and choose subsequent actions in each period. We are interested in repeated game equilibria attaining full collusion, that is to say, maximization of the players' joint payoffs. These equilibria would require the players to precisely know which player is of higher quality before making their entry decisions. We, thus, allow players to send cheap-talk messages (announcements) after learning their types and before making their entry decisions. Rather than assuming a fixed message space from which each player can choose her announcement after any history, we find it convenient for the message space to depend on the *public history*, consisting of the realized announcements, the entry decisions, and the actions (if any) in all past periods. To summarize, our stage game is an extensive-form game with the following time line.

1. Players realize their type, H or L .
2. Depending on the public history, each player is given her message space. Players simultaneously choose announcements from their message spaces. We assume all message spaces to be nonempty and finite.¹¹ We also allow for a singleton message space where a player does not engage in any meaningful communication.
3. After observing each other's announcement, players simultaneously make a decision whether to enter or not.
4. Each player who enters, observes the other player's entry decision and subsequently chooses her action (simultaneously with the other player if both players entered). The action choices are publicly observable.

We study the game where the above stage game is played every period for infinite time periods. For any period, a *private history* of a player at that period consists of all pieces of information she receives in all past periods. Namely, a private history of a player includes her own realized types, announcements of all players, the entry decisions of all players, and the actions of all players in all past periods.

A player's strategy maps each possible private history of that player to a stage action consisting of (1) her announcement as a function of her current type, (2) her entry decision as a function of her current type and all players' current announcements, and (3) her action as a function of her current type, all players' current announcements

¹¹Allowing for infinite message spaces will not affect the results of this paper. This is simply because implementation of full collusion does not need infinitely many messages. Revelation of each player's type needs just choosing from two messages. As we will see, the players need some additional messages which serve as a tiebreaker, but the tiebreaker can be designed from finitely many messages.

and all players' entry decisions.¹² We allow the stage action to involve randomization over the announcements, entry decisions and action choices. Given a strategy pair, each player's repeated game payoff is the average, discounted sum of the expected stage payoffs, where $\delta \in (0, 1)$ is a common discount factor among the players.

We use sequential equilibrium by Kreps and Wilson [19] as the equilibrium concept in this paper. Note that full collusion requires players to base their play on private information (each other's current type). This makes the public perfect equilibrium concept inadequate for our model.¹³ Hereafter, we call sequential equilibrium simply *equilibrium*.

3 Necessary and Sufficient Condition for Full Collusion

The main contribution of this paper is to present a condition on the stage game which is both necessary and sufficient in order for full collusion (total payoff maximization) to be sustained by an equilibrium when the players are sufficiently patient. We start with characterization of the fully collusive behavior. Due to Assumption 1 (iv) and the fact that $F_\theta > 0$ for each θ , full collusion requires exactly one player to enter and to choose an action that attains her mono-entrant payoff corresponding to her type. Further, since $\pi_H > \pi_L$, the sole entrant may be a low-quality player only if both players are low-quality (an event whose probability is μ_{LL}). Therefore, the expected total payoff in each period under full collusion is

$$\Pi^* \equiv \mu_{LL}\pi_L + (1 - \mu_{LL})\pi_H. \quad (1)$$

We call an equilibrium *fully collusive* if its total payoff is Π^* .

Our analysis proceeds as follows. Subsection 3.1 presents a necessary condition for the existence of a fully collusive equilibrium under some discount factor. We call this condition the *aggregate no-mimicry condition*. By contraposition, failure of this condition would imply that full collusion is not sustainable under any discount factor. Hence, our next task would be to find a sufficient condition for a fully collusive

¹²Note that the stage action depends on the current message space, which is determined by the public history.

By the way, our assumption of history-dependent message spaces may cause the following problem. What happens if a message is sent, which is available at some history but not at the current history? Or what if no message is sent? We cope with this problem by implicitly assuming that each possible message space has a "representative" element, and any "wrong" or "absent" message is regarded as the representative message.

¹³However, we will later construct an equilibrium where the play within a period depends only on the *public* history at the *beginning* of that period. Thus, if we define publicity as independence of the private information in all *past* periods, our equilibria would correspond to a perfect public equilibrium.

equilibrium to exist if the players are sufficiently patient.

Subsection 3.2 provides a preliminary result which is useful in finding fully collusive equilibria. More concretely, we show existence of a *debarment equilibrium* which debars a particular player in the sense that he is never allowed to enter in any period on the path. This debarment equilibrium exists for any discount factor and gives the excluded player his minmax value. Hence, the equilibrium serves as the harshest punishment when we construct a fully collusive equilibrium. Another feature of the debarment equilibrium is that the other player is allowed to enter in all periods on the path and obtains the payoff she would have obtained had we just had a single player.

Finally, Subsection 3.3 shows that the aggregate no-mimicry condition is also a sufficient condition for the existence of a fully collusive equilibrium when the players are sufficiently patient. In other words, the condition characterizes the environments where full collusion is attainable for sufficiently patient players.

3.1 Necessary Condition for Full Collusion

Let us define

$$\rho_L \equiv \frac{\mu_{LL}}{\mu_{LL} + \mu_{LH}}, \quad (2)$$

which is the probability that the other player is low-quality given that a player is low-quality.

We also define

$$A_H^* = \{a_H \in A : \Pi_H = U(a_H, H)\}.$$

Namely, A_H^* is the set of all actions which achieve the mono-entrant payoff for a high-quality player. Then define

$$\hat{\pi}_L \equiv \min_{a_H \in A_H^*} U(a_H, L) - F_L. \quad (3)$$

We call $\hat{\pi}_L$ the mimicry payoff of a low-quality player. This idea is easiest to understand when A_H^* is singleton. In such a case, $\hat{\pi}_L$ is the payoff when a low-quality player solely enters and chooses the action which maximizes the mono-entrant payoff function of her high-quality counterpart. Thus, the low-quality player obtains the mimicry payoff when she pretends to be her high-quality counterpart. If A_H^* has two or more elements, $\hat{\pi}_L$ is the minimum payoff when a low-quality player solely enters and chooses an action that maximizes the mono-entrant payoff function of her high-quality counterpart. Thus, the low-quality player obtains the mimicry payoff when she pretends to be her high-quality counterpart in the least favorable way.

The following result shows that a fully collusive equilibrium exists only if the mimicry payoff is sufficiently small.

Proposition 3.1. *Suppose under some δ , a fully collusive equilibrium exists. Then it holds that*

$$\rho_L \pi_L \geq (1 + \rho_L) \hat{\pi}_L. \quad (4)$$

Proof. Fix δ and suppose that an equilibrium exists whose total payoff is Π^* defined by (1). Hereafter, we focus on the play in the initial period of the repeated game. Which player is prescribed to enter in the initial period depends on the message pair in that period. Since each player of each type may play randomized strategies, the message pair may also involve randomization over the message space conditional on the players' types.

For each $i \in \{1, 2\}$ and $\theta \in \{L, H\}$, let $\tau_{\theta i}$ be the probability that, conditional on both players being θ -type, the equilibrium play realizes a message pair under which player i is prescribed to enter. Under this equilibrium, each low-quality player i obtains an expected initial-period payoff of $\rho_L \tau_{Li} \pi_L$. This is because full collusion requires that she is prescribed to enter only when the other player is low-quality (an event whose conditional probability is ρ_L) and, conditional on that, a message pair is selected which prescribes player i to enter (an event whose probability is τ_{Li}). Once she enters, full collusion also requires her to choose an action leading to her mono-entrant payoff π_L .

Now suppose that a low-quality player i chooses exactly the same (possibly randomized) message, entry decision, and action as the high-quality player i in the initial period. In this case, the low-quality player i would be prescribed to enter with probability $\rho_L + (1 - \rho_L) \tau_{Hi}$. This is because full collusion requires that (1) she is prescribed to enter when the other player is low-quality (an event whose conditional probability is ρ_L), and (2) she is also prescribed to enter also when the other player is high-quality (an event whose conditional probability is $1 - \rho_L$) and, conditional on that, a message pair is selected which prescribes player i to enter (an event whose probability is τ_{Hi}). Once she enters, she chooses an action which maximizes the high-quality player's mono-entrant payoff function. Hence, her expected initial-period payoff would be at least $\{\rho_L + (1 - \rho_L) \tau_{Hi}\} \hat{\pi}_L$, since the payoff of a low-quality player from this mimicry is at least $\hat{\pi}_L$.

Hence, it follows that

$$\rho_L \tau_{Li} \pi_L \geq \{\rho_L + (1 - \rho_L) \tau_{Hi}\} \hat{\pi}_L \quad (5)$$

for any i in order for the collusion to sustain. This is because otherwise the low-quality player i could profitably deviate from the equilibrium (note that the other player regards this deviation as an act of the high-quality player i , so it does not affect the continuation payoff).

Taking the sum of (5) over i produces

$$\rho_L(\tau_{L1} + \tau_{L2})\pi_L \geq \{2\rho_L + (1 - \rho_L)(\tau_{H1} + \tau_{H2})\}\hat{\pi}_L. \quad (6)$$

Recall that full collusion requires that given any message pair on the path, exactly one player is prescribed to enter. Therefore, $\tau_{\theta 1} + \tau_{\theta 2} = 1$ for each θ . Substituting these into (6) establishes (4), which completes the proof. \square

We call (4) the *aggregate no-mimicry condition*. Note that at the initial period of a fully collusive equilibrium, a low-quality player has an undetectable deviation where she chooses exactly the same message, entry decision, and action as her high-quality self. This mimicry deviation is undetectable because the other player consistently believes it to be his high-quality opponent's conduct. The players in our model have no way to detect each other's types at any point of time during the entire play of the game. Thus, the stage payoff from the mimicry deviation may not exceed the equilibrium stage payoff for that period, which is a necessary condition in order for an equilibrium to hold. As the proof of Proposition 3.1 shows, the condition (4) follows from taking a sum of both players' individual no-mimicry conditions. This is why we call it "aggregate." Proposition 3.1 implies that if the aggregate no-mimicry condition fails, full collusion is not an equilibrium outcome under any level of discounting.

Note that the aggregate no-mimicry condition requires $\pi_L > \hat{\pi}_L$. Namely, the mimicry by a low-quality player (to the least favorable maximizer of the high-quality player's mono-entrant payoff function) causes a loss from her own mono-entrant payoff. This requires existence of an action which attains the high-quality player's mono-entrant payoff but does not attain the low-quality player's mono-entrant payoff. Without such an action, the low-quality player can profitably and undetectably deviate, which prevents full collusion.

In particular, the aggregate no-mimicry condition fails if both types' mono-entrant payoff function have the same set of maximizers. In the example of Bertrand duopoly, this is the case when both types' cost functions differ only by a constant (namely, $C_L(a) - C_H(a)$ is constant for any a), but this is not the case in many other cases such as the case of constant and different marginal costs (namely, $C_\theta(q) = c_\theta q$ for any q , where $c_L > c_H > 0$). In the example of the auction for time to completion, the aggregate no-mimicry condition fails if both $C_L(a)$ and $C_H(a)$ have the same minimizers, but it may hold otherwise.

Note also that the aggregate no-mimicry condition always holds if $\hat{\pi}_L \leq 0$. This is the case when mimicry reduces the low-quality player's payoff to the extent that it does not recover her entry cost.

The aggregate no-mimicry condition can be rewritten so that

$$\rho_L(\pi_L - \hat{\pi}_L) \geq \hat{\pi}_L. \quad (7)$$

It is easy to see from the definition of π_L and $\hat{\pi}_L$ that an increase in F_L decreases the right-hand-side of (7) but does not affect its left-hand-side. As a result, full collusion is harder to sustain under environments with relatively smaller entry costs. With a smaller entry cost the expected profit from a low-quality player's mimicry deviation would be larger and hence would provide a stronger temptation to opt for such a deviation.¹⁴

It is also clear from (7) that full collusion is easier to sustain under environments with large ρ_L . This is because, from an ex-post viewpoint, the mimicry deviation pays only when the other player is high-quality. The event is less likely when ρ_L is large, which weakens the temptation to deviate. Note that the aggregate no-mimicry condition concerns the low-quality player's incentive to mimic her high-quality version. Thus, the probabilities of the type pairs affect the condition only through ρ_L .

3.2 Debarment Equilibrium

Before we go on to examining a sufficient condition for a fully collusive equilibrium to exist if the players are sufficiently patient, this subsection constructs an equilibrium where a particular player is never allowed to enter the market in any period.¹⁵ We call this equilibrium *debarment equilibrium*.

More concretely, we consider a particular outcome path, where one player, say player 1, solely enters and chooses an action which attains her mono-entrant payoff depending on her type in all periods. The other player, player 2, does not enter in any period. Player 1's expected payoff under this outcome is

$$\Pi^m \equiv (\mu_{LL} + \mu_{LH})\pi_L + (\mu_{HL} + \mu_{HH})\pi_H, \quad (8)$$

because she is guaranteed the mono-entrant payoff according to her type in all periods. Note that if player 2 were absent in our setting, this payoff outcome would have been realized. Player 2's payoff under this outcome is zero. Note also that $\Pi^m < \Pi^*$.

The following proposition shows that under any non-zero discount factor, the outcome is sustained by an equilibrium. The result has two implications. First, the outcome corresponding to the one-player environment is always an equilibrium outcome in our two-player environment also. Second, player 2's payoff is zero, and it

¹⁴However, smaller entry costs make the collusive profits larger.

¹⁵The result holds irrespective of validity of the aggregate no-mimicry condition.

is the minmax value of every player in this repeated game. As a result, this debarment equilibrium always serves as the harshest punishment when we construct a fully collusive equilibrium.

Proposition 3.2. *For any $\delta > 0$, there exists an equilibrium where player 1's payoff is Π^m and player 2's payoff is zero.*

Proof. First, for each $\theta \in \{L, H\}$, fix an action attaining the mono-entrant payoff of that type. We denote this action by a_θ^* , so that $U(a_\theta^*, \theta) = \Pi_\theta$. Further, by Assumption 1(i)(iii), there exists the largest a such that $U(a', H) \leq 0$ for any $a' \leq a$. We denote this largest action by \underline{a} . From Assumption 1(iv), we have $\underline{a} < \bar{a}$.

Fix $\delta > 0$. For the repeated game with discount factor δ , we define a strategy pair σ as follows. This is an automaton strategy pair, represented by a quadruple consisting of a state space, an initial state, a function that maps each state to the players' stage actions, and a transition rule that maps each state and publicly observable outcomes (namely, announcements, entry decisions, and actions, if any) to a state.

More concretely, we set the state space to be $\{1, 2\}$ and the initial state as 1. Recall that each player's stage action depends on her message space. Here we let any message space to always be a singleton. That is to say, no player sends a meaningful message in any state.

The stage action given state i is prescribed as follows. Player i enters and player $j \neq i$ does not enter. The subsequent actions are specified as follows.

1. Suppose there is a single entrant (which may be player $j \neq i$). If the player is of type θ , it chooses a_θ^* .
2. Suppose both players entered. By Assumption 1(i)(ii), there exists an $\eta > 0$ such that all of the following conditions are satisfied:

$$0 < (1 - \delta)U(\underline{a} + \eta, H) \leq \delta\Pi^m, \quad (9)$$

$$U(a, L) < 0 \quad \forall a \in [\underline{a}, \underline{a} + 2\eta] \quad (10)$$

$$\max_{a \in [\underline{a}, \underline{a} + \eta]} U(a, H) = U(\underline{a} + \eta, H) < F_H. \quad (11)$$

Note that $\delta > 0$ is necessary for (9) to hold.

Each player chooses her action as follows.

- Any type of player i randomly chooses her action over the interval $(\underline{a} + \eta, \underline{a} + 2\eta)$, following a continuous distribution function $G(a)$ such that $G(\underline{a} + \eta) = 0$, $G(\underline{a} + 2\eta) = 1$, and

$$U(a, H)\{1 - G(a)\} \leq U(\underline{a} + \eta, H) \quad (12)$$

for any $a \in (\underline{a} + \eta, \underline{a} + 2\eta)$.¹⁶

- A low-quality player $j \neq i$ chooses $\underline{a} + 2\eta$.
- A high-quality player $j \neq i$ chooses $\underline{a} + \eta$.

The transition rule is defined as follows. Let the state of the current period be i .

1. If both players entered, and if player i chose an action equal to or less than $\underline{a} + \eta$, then the state in the next period is $j \neq i$.
2. If the above does not apply, then the state in the next period is i .

We combine σ with a particular consistent system of beliefs. For that purpose, we arbitrarily fix a tremble satisfying the following requirements.

- At any history, the probability of any (possibly off-the-path) stage action depends only on the state given the history.
- Suppose the state of the current period is i . An entry by a high-quality player $j \neq i$ is infinitely more likely than an entry by a low-quality player $j \neq i$.

Then, under the system of beliefs associated with the tremble, if the state of the current period is i and if both players entered, player i believes that player $j \neq i$ is high-quality.

If both players follow σ from any period whose state is i , the continuation path is such that player i always and solely enters and chooses an action attaining her mono-entrant payoff given her type. Thus player i 's continuation payoff is Π^m . Since player $j \neq i$ has no opportunity to enter, his continuation payoff is zero. Since the initial state is 1, the payoff pair of σ is $(\Pi^m, 0)$. Hence, the proof is complete if we show that σ is an equilibrium.

Our task is to verify sequential rationality of σ , combined with the system of beliefs specified above. We invoke the idea of one deviation principle. For each decision opportunity of a player in a given period (namely, an entry decision and an action after an entry), we examine the deviations where the player makes a different decision and then follows σ afterwards. If none of those deviations improves her payoff, we have verified sequential rationality. Note that we ignore incentives about messages because the players have singleton message spaces.

Suppose the state of the current period is i . We denote the two players by players i and j . We work backwards within the period.

¹⁶One example of such $G(a)$ is

$$G(a) = 1 - \frac{U(\underline{a} + \eta, H)}{\max_{a' \in [\underline{a} + \eta, \underline{a}]} U(a', H)} + \frac{(a - \underline{a} - \eta)U(\underline{a} + \eta, H)}{\eta \max_{a' \in [\underline{a} + \eta, \underline{a} + 2\eta]} U(a', H)}$$

for any $a \in [\underline{a} + \eta, \underline{a} + 2\eta]$, $G(a) = 0$ for any $a < \underline{a} + \eta$, and $G(a) = 1$ for any $a > \underline{a} + 2\eta$.

1. Suppose both players entered. Player j believes any type of player i to randomly choose her action following the distribution function $G(a)$. By the definition of \underline{a} and (11), no $a < \underline{a} + \eta$ is better than $\underline{a} + \eta$ for a high-quality player j against this randomized action. Further, due to (12), no $a > \underline{a} + \eta$ is better than $\underline{a} + \eta$ for a high-quality player j against this randomized action. This proves that $\underline{a} + \eta$ is a statically optimal action for a high-quality player j . Further, by (10), no action gives a low-quality player j a positive stage payoff, and $\underline{a} + 2\eta$ gives him a zero stage payoff against this randomized action. Hence $\underline{a} + 2\eta$ is statically optimal for a low-quality player j . Since the state in the next period is i irrespective of the action, σ prescribes an optimal action for each type of player j .

We have specified the beliefs so that any type of player i believes player j to be high-quality and therefore to choose $\underline{a} + \eta$. Due to (10), the stage payoff of any action for a low-quality player i is at most zero, which is attained by any $a \in (\underline{a} + \eta, \underline{a} + 2\eta)$. Hence it is statically optimal for a low-quality player i to randomly choose her action following the distribution function $G(a)$. Since the state in the next period is i as long as she chooses $a \in (\underline{a} + \eta, \underline{a} + 2\eta)$, σ prescribes an optimal (randomized) action for a low-quality player i .

For a high-quality player i , her continuation payoff when she chooses $a > \underline{a} + \eta$ is $\delta\Pi^m$, because the state in the next period is i in this case. If she chooses $a \leq \underline{a} + \eta$, the state in the next period is j given σ . Therefore, her continuation payoff is at most

$$(1 - \delta)U(a, H) \leq (1 - \delta)U(\underline{a} + \eta, H) \leq \delta\Pi^m,$$

where the first inequality follows from (11), and the second from (9). Therefore, it is optimal to randomly choose her action following the distribution function $G(a)$, as σ prescribes.

2. Suppose any player of type θ is a single entrant. Since the state in the next period is i irrespective of her action, it is optimal for her to choose an action which maximizes the current stage payoff. Hence, it is optimal for her to choose a_θ^* , as σ prescribes.

3. Let us consider each player's entry decision.

Player i of any type θ believes that player j will not enter and her own entry decision this period does not affect the state in the next period. Hence, her continuation payoff given σ is $(1 - \delta)\pi_\theta + \delta\Pi^m$, while her continuation payoff when she does not enter is $\delta\Pi^m < (1 - \delta)\pi_\theta + \delta\Pi^m$. Hence it is optimal for her to enter, as σ prescribes.

Next, player j of any type is prescribed not to enter, and his continuation payoff is

zero. If a high-quality player j enters, he believes he will be prescribed to choose $\underline{a} + \eta$. Due to (11), his current stage payoff is negative. Further, if a low-quality player j enters, he cannot obtain a greater stage payoff than his high-quality variant, due to Assumption 1(ii). Thus, his current stage payoff is negative. His continuation payoff from the next period on is zero irrespective of his type, because his behavior does not affect the state. Hence it is optimal for any type of player j to not enter, as σ prescribes.

These establish sequential rationality of σ , and the proof is complete. \square

While Proposition 3.2 proves existence of a debarment equilibrium where player 2 never enters, another debarment equilibrium where player 1 never enters also exists by symmetry.

Since a debarment equilibrium prescribes only one player, say player 1, to enter every period, the equilibrium construction entails three key ingredients in preventing player 2 from entering. First, if both players entered, each type of player 1 believes player 2 to be high-quality and hence, any type of player 1 chooses her action so as to hold the high-quality player 2 to a sufficiently small stage payoff. Consequently, player 1 also holds a low-quality player 2 to a sufficiently small stage payoff.

Second, player 1's action when both players entered is randomized over an open interval. The probability distribution of her action is designed so that each type of player 2 has a static best response, and hence sequential rationality is guaranteed. This idea of using randomized actions to ensure existence of a best response also appears in Blume [10].

Third, since a high-quality player 2 chooses a unique static best response against player 1's randomized actions, a high-quality player 1 is not statically best-responding given player 2's action and her belief. The equilibrium thus punishes player 1's deviation by switching to the debarment equilibrium with player 1 being excluded. Given that player 2 is held at a sufficiently small stage payoff, player 1's deviation gain is also small. Therefore, a sufficiently severe punishment can be designed for any non-zero discount factor.¹⁷

3.3 Sufficient Condition for Full Collusion

This subsection shows that if the aggregate no-mimicry condition holds and if the players are sufficiently patient, full collusion is an equilibrium outcome. In other words, the aggregate no-mimicry condition serves both as a necessary and a sufficient condition for full collusion by sufficiently patient players. The next proposition presents this result.

¹⁷Due to this feature, Proposition 3.2 does not cover the case of $\delta = 0$.

Proposition 3.3. Define $\pi^* \equiv \Pi^*/2$. If both (4) and

$$\delta \geq \underline{\delta} \equiv \frac{\pi_L}{(1 + \rho_L)\pi^* + \pi_L} \quad (13)$$

hold (ρ_L is defined by (2)), then an equilibrium exists such that each player's expected payoff is π^* . Namely, an equal division of the fully collusive total payoffs is an equilibrium outcome.

Proof. First, we define a_L^* , a_H^* , and $\underline{a} \in A$ in the exactly same way as in the proof of Proposition 3.2, except one additional requirement. Now a_H^* also satisfies $U(a_H^*, L) - F_L = \hat{\pi}_L$, where $\hat{\pi}_L$ is defined by (3). Further, fix $\delta \geq \underline{\delta}$, and we define $\eta > 0$ so that the conditions (9)–(11) and

$$(1 - \delta)U(\underline{a} + \eta, H) \leq \delta\pi^* \quad (14)$$

all hold.

For the repeated game with discount factor δ , we define an automaton strategy pair σ^* as follows. The state space is a tripleton $\{0, 1, 2\}$, and the initial state is 0. Recall that each player's message space depends on a public history. We thus let the message spaces depend on the current state. The specification of the message spaces and the stage actions are as follows.

State 0: Under this state, each player has a common message space, which has four elements and is represented by $\{L, H\} \times \{-1, 1\}$. The stage action is prescribed as follows.

1. Each player of type $\theta \in \{L, H\}$ announces $(\theta, -1)$ and $(\theta, 1)$ with equal probability.

Given a realized announcement pair, we call player 1 a *winner* if one of the following holds.

- (a) Player 1 announced either $(H, -1)$ or $(H, 1)$, and player 2 announced either $(L, -1)$ or $(L, 1)$.
- (b) Both players' announcements are the same.

Also, we call player 1 a *loser* if she is not a winner. Finally, we call player 2 a *winner* if player 1 is a loser, and we call player 2 a *loser* if player 1 is a winner. Note that we have exactly one winner and exactly one loser under any announcement pair.

2. Given an announcement pair, each player enters if she is a winner and does not enter if she is a loser. Hence, given any announcement pair, exactly one player is

prescribed to enter.

3. Given an announcement pair and entry decisions, suppose that there is a single entrant. In such a case the entrant chooses a_H^* if one of the following holds and chooses a_L^* otherwise.

- (a) The entrant is high-quality.
- (b) The entrant is low-quality, he announced either $(H, -1)$ or $(H, 1)$, and it holds that

$$(1 - \delta)\hat{\pi}_L + \delta\pi^* \geq (1 - \delta)\pi_L. \quad (15)$$

4. Given an announcement pair and entry decisions, suppose that both players entered. Each player chooses her action as follows.

- Any type of a winner randomly chooses her action over the interval $(\underline{a} + \eta, \underline{a} + 2\eta)$, following a continuous distribution function $G(a)$ such that $G(\underline{a} + \eta) = 0$, $G(\underline{a} + 2\eta) = 1$, and (12) holds.
- A low-quality loser chooses $\underline{a} + 2\eta$.
- A high-quality loser chooses $\underline{a} + \eta$.

State i with $i \geq 1$: Under this state, each player has a singleton message space. The stage action is exactly the same as the one at state i of the debarment equilibrium specified in the proof of Proposition 3.2.

Now we specify the transition rule for the states.

1. Suppose the state in the current period is 0.
 - (a) If player i is a winner but did not enter, then the state in the next period is $3 - i$.
 - (b) If player i announced either $(H, -1)$ or $(H, 1)$, solely entered, and chose an action other than a_H^* , then the state in the next period is $3 - i$.
 - (c) If no player announced either $(H, -1)$ or $(H, 1)$, if both players entered with player i being a winner, and if player i chose an action equal to or less than $\underline{a} + \eta$, then the state in the next period is $3 - i$.
 - (d) If none of the above applies, then the state in the next period is 0.
2. Suppose the state in the current period is i with $i \geq 1$.
 - (a) If both players entered, and if player i chose an action equal to or less than $\underline{a} + \eta$, then the state in the next period is $3 - i$.
 - (b) If the above does not apply, then the state in the next period is i .

We combine σ^* with a particular consistent system of beliefs. For that purpose, we arbitrarily fix a tremble satisfying the following requirements.

- At any history, the probability of any (possibly off-the-path) stage action depends only on the state given the history.
- Suppose the state in the current period is 0. Given any announcement pair, an entry by a high-quality loser is infinitely more likely than an entry by a low-quality loser.
- Suppose the state of the current period is i with $i \geq 1$. An entry by a high-quality player $j \neq i$ is infinitely more likely than an entry by a low-quality player $j \neq i$.

The system of beliefs associated with the tremble satisfies the following conditions.

1. Suppose the state in the current period is 0, and both players entered. Then a winner believes the other player to be high-quality.
2. Suppose the state of the current period is i with $i \geq 1$, and both players entered. Then player i believes player $j \neq i$ to be high-quality.

Our next objective is to compute each player's continuation payoff given a state.

State 0: First, define

$$\rho_H = \frac{\mu_{HH}}{\mu_{HL} + \mu_{HH}},$$

which is the probability that the other player is high-quality given that a player is high-quality.

Given σ^* , a high-quality player 1 becomes a winner with probability $1 - \frac{\rho_H}{2}$. This is because she becomes a winner if player 2 is low-quality (this is with probability $1 - \rho_H$) or if player 2 is high-quality (this is with probability ρ_H) and makes the same announcement as player 1 (this is with probability $\frac{1}{2}$, because the players announce one of two common messages with equal probability). By a similar argument, a high-quality player 2 becomes a winner with probability $1 - \frac{\rho_H}{2}$. Further, a low-quality player 1 becomes a winner with probability $\frac{\rho_L}{2}$. This is because she becomes a winner if player 2 is low-quality (this is with probability ρ_L) and makes the same announcement as player 1 (this is with probability $\frac{1}{2}$ by the same reason as above). Deducing similarly, a low-quality player 2 becomes a winner with probability $\frac{\rho_L}{2}$.

Given these, any player's expected stage payoff before knowing her type is

$$(\mu_{LL} + \mu_{LH})\frac{\rho_L}{2}\pi_L + (\mu_{HL} + \mu_{HH})\left(1 - \frac{\rho_H}{2}\right)\pi_H = \pi^*.$$

Since the state remains to be 0 as long as both players follow σ^* , this verifies that each player's continuation payoff when the current state is 0 is π^* . Since the initial state is 0, σ^* attains an equal division of the fully collusive total payoffs.

State i with $i \geq 1$: The continuation play given any of these states is the debarment equilibrium play with player i being the entrant. Therefore, player i 's continuation payoff is Π^m defined by (8), and the continuation payoff of player $j \neq i$ is zero.

Our last objective is to verify sequential rationality of the strategy pair σ^* , combined with the system of beliefs we have specified before.

State 0: We work backwards within the period.

1. Suppose both players entered. Irrespective of the realized announcement pair, the loser believes any type of the winner to randomly choose her action following the distribution function $G(a)$. By the definition of \underline{a} and (11), no $a < \underline{a} + \eta$ is better than $\underline{a} + \eta$ for a high-quality loser against this randomized action chosen by the winner. Further, due to (12), no $a > \underline{a} + \eta$ is better than $\underline{a} + \eta$ for a high-quality loser against this randomized action chosen by the winner. This proves that $\underline{a} + \eta$ is a statically optimal action for a high-quality loser. Further, by (10), no action gives a low-quality loser a positive stage payoff, and $\underline{a} + 2\eta$ gives him a zero stage payoff against this randomized action. Hence $\underline{a} + 2\eta$ is statically optimal for a low-quality loser. Since the state in the next period is 0 irrespective of the loser's action, σ^* prescribes an optimal action for each type of the loser.

We have specified the beliefs so that any type of the winner believes the loser to be high-quality and therefore to choose $\underline{a} + \eta$. Due to (10), the stage payoff of any action for a low-quality winner is at most zero, which is attained by any $a \in (\underline{a} + \eta, \underline{a} + 2\eta)$. Hence it is statically optimal for a low-quality winner to randomly choose her action following the distribution function $G(a)$. Since the state in the next period is 0 as long as she chooses $a \in (\underline{a} + \eta, \underline{a} + 2\eta)$, σ^* prescribes an optimal (randomized) action for a low-quality winner.

For a high-quality winner player i , her continuation payoff when she chooses $a > \underline{a} + \eta$ is $\delta\pi^*$, because the state in the next period is 0 in this case. If she chooses $a \leq \underline{a} + \eta$, the state in the next period is $3 - i$ given σ^* . Therefore, her continuation payoff is at most

$$(1 - \delta)U(a, H) \leq (1 - \delta)U(\underline{a} + \eta, H) \leq \delta\pi^*,$$

where the first inequality follows from the definition of \underline{a} and (11), and the second from (14). Therefore, it is optimal to randomly choose her action following the distribution function $G(a)$, as σ^* prescribes.

2. Suppose player i is a single entrant. Unless player i is low-quality and announced either $(H, -1)$ or $(H, 1)$, σ^* prescribes an action which maximizes the current stage payoff and ensures that the state in the next period is 0. Hence, it is

optimal to choose the action, because any other action may lead to a smaller stage payoff and may cause the state to move to $3 - i$ (and a zero continuation payoff) in the next period.

So suppose player i is low-quality and announced either $(H, -1)$ or $(H, 1)$. If she chooses a_H^* , her continuation payoff is $(1 - \delta)U(a_H^*, L) + \delta\pi^*$, since the state in the next period is 0. If she chooses $a \neq a_H^*$, the state in the next period is $3 - i$ and therefore her continuation payoff is at most $(1 - \delta)\Pi_L$, which is attained when $a = a_L^*$. Therefore, if (15) holds, it is optimal to choose a_H^* , as σ^* prescribes. Otherwise, it is optimal to choose a_L^* , as σ^* prescribes.

3. Let an announcement pair be given, and consider player i .

First, suppose player i is a loser. If she enters, she believes any type of the winner to randomly choose his action following the distribution function $G(a)$. Hence her stage payoff is $U(\underline{a} + \eta, H) - F_H$ if she is high-quality, and it is $-F_L$ if she is low-quality. Therefore, by (11), the stage payoff if she enters is always negative. Since the state in the next period is 0 irrespective of her entry decision, it is optimal for a loser not to enter, as σ^* prescribes.

Next, suppose player i is a winner. If she does not enter, her continuation payoff is zero because the state in the next period is $3 - i$. Unless player i is low-quality and announced either $(H, -1)$ or $(H, 1)$, her continuation payoff if she enters is $(1 - \delta)\pi_\theta + \delta\pi^* > 0$, because the state in the next period is 0. So suppose player i is low-quality and announced either $(H, -1)$ or $(H, 1)$. If she enters, her continuation payoff is

$$\max \{ (1 - \delta)\hat{\pi}_L + \delta\pi^*, (1 - \delta)\pi_L \} > 0, \quad (16)$$

because $\pi_L > 0$. As a result, it is optimal for a winner to enter, as σ^* prescribes.

4. At the beginning of the current period, any high-quality player's continuation payoff is $(1 - \delta)(1 - \frac{\rho_H}{2})\pi_H + \delta\pi^*$ whether she chooses $(H, -1)$ or $(H, 1)$, because the probability of being a winner is $1 - \frac{\rho_H}{2}$ in either case. If she rather chooses either $(L, -1)$ or $(L, 1)$, the probability of being a winner is $\frac{1 - \rho_H}{2} < 1 - \frac{\rho_H}{2}$. Since the state in the next period is 0 irrespective of her announcement, it is optimal for any high-quality player to choose $(H, -1)$ or $(H, 1)$. Hence, it is optimal for any high-quality player to randomize over $(H, -1)$ and $(H, 1)$, as σ^* prescribes. Any low-quality player's continuation payoff is $(1 - \delta)\frac{\rho_L}{2}\pi_L + \delta\pi^*$ whether she chooses $(L, -1)$ or $(L, 1)$, because the probability of being a winner is $\frac{\rho_L}{2}$ in either case. If, instead, she chooses either $(H, -1)$ or $(H, 1)$, the probability of being a winner is $1 - \frac{1 - \rho_L}{2}$. This is because she becomes a loser only if the other player is high-quality (this is with probability $1 - \rho_L$) and if an unfavorable

announcement pair is realized (this is with probability $\frac{1}{2}$). If she becomes a winner, her continuation payoff equals the left-hand-side of (16). If she becomes a loser, her continuation payoff is $\delta\pi^*$ because the state in the next period is 0. Therefore, the overall continuation payoff of choosing either $(H, -1)$ or $(H, 1)$ is:

$$\begin{aligned} & \left(1 - \frac{1 - \rho_L}{2}\right) \max \left\{ (1 - \delta)\hat{\pi}_L + \delta\pi^*, (1 - \delta)\pi_L \right\} + \frac{1 - \rho_L}{2} \cdot \delta\pi^* \\ &= \max \left\{ (1 - \delta)\left(1 - \frac{1 - \rho_L}{2}\right)\hat{\pi}_L + \delta\pi^*, (1 - \delta)\left(1 - \frac{1 - \rho_L}{2}\right)\pi_L + \frac{1 - \rho_L}{2} \cdot \delta\pi^* \right\}. \end{aligned}$$

From the aggregate no-mimicry condition, we obtain

$$(1 - \delta)\left(1 - \frac{1 - \rho_L}{2}\right)\hat{\pi}_L + \delta\pi^* \leq (1 - \delta)\frac{\rho_L}{2}\pi_L + \delta\pi^*.$$

Since (13) holds, we also obtain

$$(1 - \delta)\left(1 - \frac{1 - \rho_L}{2}\right)\pi_L + \frac{1 - \rho_L}{2} \cdot \delta\pi^* \leq (1 - \delta)\frac{\rho_L}{2}\pi_L + \delta\pi^*.$$

From the above conditions, we conclude that it is optimal for any low-quality player to choose $(L, -1)$ or $(L, 1)$. Hence, it is optimal for any low-quality player to randomize over $(L, -1)$ and $(L, 1)$, as σ^* prescribes.

State i with $i \geq 1$: The same line of arguments which verify sequential rationality of the debarment equilibrium (Proposition 3.2) proves sequential rationality.

Since we have examined all possibilities, the proof is complete. \square

We explain some features of the fully collusive equilibrium constructed in the proof of Proposition 3.3. The equilibrium requires a high-quality player to solely enter when the players' qualities are different *and* it decides who should enter when their qualities are equal. Thus, in each period on the equilibrium path, the players announce their types *and* verbally play a matching penny game. The idea is to use the outcome of the matching penny game as a fair tie-breaker, which is a simple way to achieve an equal division of the fully collusive total payoffs.

A key equilibrium condition is for a low-quality player not to pretend to be her high-quality counterpart. A low-quality player has two potentially profitable deviations. One is to completely mimic the behavior of a high-quality counterpart by choosing an action which would attain the high-quality counterpart's mono-entrant payoff. The aggregate no-mimicry condition ensures that this deviation is not profitable. The other is to mimic the high-quality counterpart's announcement for a greater probability of entry and, if prescribed to enter, to choose an action which attains her own mono-entrant payoff as a low-quality player. Since this deviation is observable and since the

debarment equilibrium can be used as the severest punishment, this deviation is not profitable if the players are sufficiently patient.

Similar to the debarment equilibrium, the fully collusive equilibrium also punishes a loser if he enters. Here the loser's incentives are different from the incentives of the excluded player in the debarment equilibrium, because the loser learns the winner's type from the communication stage. This additional information does not matter, however, because a winner is always prescribed to choose a common randomized action irrespective of her type. Hence, the same construction as the debarment equilibrium prevents a double entry in the fully collusive equilibrium.

We emphasize that the condition on the players' patience (13) is not so stringent. Note that $2\pi^* > \pi_L$. Substituting this into (13), we see that the condition always holds if $\delta \geq \frac{2}{3}$.

Finally, we draw some comparisons with the prior literature. At a basic level, unlike [5, 7], we characterize full collusion in this paper. A major difference is that our stage game is an extensive-form game and that the communication stage after the players learn their types and before they make entry decisions serves as a channel to decide who should solely enter. Without such a channel, full collusion is not an equilibrium outcome in [5, 7]. While this paper only focuses on full collusion, our companion paper (Patra [23]) analyzes partially collusive equilibria under similar stage game structure.

4 Full Collusion with One-sided Communication

The previous section presented a condition both necessary and sufficient in order for full collusion to be an equilibrium outcome when the players are sufficiently patient. The sufficiency result is based on a particular strategy pair, and we specified the range of discount factors under which this strategy pair formed an equilibrium. This section explores a possibility that different strategy pairs may sustain full collusion for a wider range of discount factors.

Our assumption of binary types implies that (i) if a player is high-quality, she has higher quality, and (ii) if a player is low-quality, the other player has higher quality. Hence, if just one player truthfully announces her type, a player with higher quality can be identified. Let us thus consider the following scheme of full collusion by one-sided communication. Player i announces her type, while the other player $j \neq i$ does not make any meaningful announcement. If player i announces H , she solely enters and chooses an action which attains the mono-entrant payoff as a high-quality player. If player i announces L , player j solely enters and chooses an action which attains the mono-entrant payoff corresponding to his type. Given player i 's truthful

announcements, this scheme would achieve full collusion.

While we introduce one-sided communication as a vehicle for less patient players to collude, this scheme is interesting in its own right. A clear reason is that such a scheme reduces the total amount of meaningful communication and hence, may be less likely to be detected by anti-trust regulators as compared to the equilibrium with two-sided communication constructed in the proof of Proposition 3.3.

Which player makes an announcement does not affect the efficiency of the scheme. Hence, we could think of infinitely many dynamic scenarios for full collusion. For example, the players may take turns in making announcements after every stage. Alternatively, a player's turn to make an announcement could be conditioned on the history. For example, a player may be permitted to make an announcement in two consecutive periods if and only if she announced L in the first period and hence let the other player enter in that period. Our analysis starts with a simplest scenario where one player always makes an announcement (on the path). We later consider other scenarios.

The main result of this section presents a sufficient condition for an equilibrium with one-sided communication to sustain full collusion for a wider range of discount factors than the equilibrium constructed in the previous section. As we will see in what follows, a key condition is $\hat{\pi}_L \leq 0$: the mimicry payoff is nonpositive. Note that this condition implies the aggregate no-mimicry condition.

Proposition 4.1. *Suppose $\hat{\pi}_L \leq 0$ and*

$$(\mu_{HL} + \mu_{HH})\pi_H > (1 + \rho_L)\pi^*, \quad (17)$$

where ρ_L is defined by (2). Then $\hat{\delta} \in (0, \underline{\delta})$ exists such that a fully collusive equilibrium exists for any $\delta \geq \hat{\delta}$.

Proof. We first define

$$\hat{\delta} = \frac{\pi_L}{(\mu_{HL} + \mu_{HH})\pi_H + \pi_L}.$$

It follows from (17) that $\hat{\delta} < \underline{\delta}$.

We also define a_L^* , a_H^* , and $\underline{a} \in A$ in exactly the same way as the proof of Proposition 3.3. Fix $\delta \geq \hat{\delta}$, and we choose $\eta > 0$ so that the conditions (9)–(11) and

$$(1 - \delta)U(\underline{a} + \eta, H) \leq \delta \min \{(\mu_{HL} + \mu_{HH})\pi_H, \mu_{LL}\pi_L + \mu_{LH}\pi_H\} \quad (18)$$

all hold.

Let $\hat{\sigma} = (\hat{\sigma}_1, \hat{\sigma}_2)$ be the following automaton strategy profile. The state space is a tripleton $\{0, 1, 2\}$, and the initial state is 0. The specification of the message spaces

and the stage actions are as follows.

State 0: Under this state, player 1 has a doubleton message space $\{L, H\}$, and player 2 has a singleton message space. The stage action is prescribed as follows.

1. A high-quality player 1 announces H , and a low-quality player 1 announces L . We call player 1 a *winner* if she announced H , and we call player 1 a *loser* if she announced L . Further, we call player 2 a *winner* if player 1 is a loser, and we call player 2 a *loser* if player 1 is a winner. Note that we always have one winner and one loser.
2. Given player 1's announcement, the winner enters and the loser does not enter.
3. Given player 1's announcement and both players' entry decisions, suppose that there is a single entrant. In such a case the entrant chooses a_H^* if one of the following holds and chooses a_L^* otherwise.
 - (a) The entrant is high-quality.
 - (b) The entrant is a low-quality player 1 who announced H , and it holds that

$$(1 - \delta)\hat{\pi}_L + \delta(\mu_{HL} + \mu_{HH})\pi_H \geq (1 - \delta)\pi_L. \quad (19)$$

4. Given player 1's announcement and both players' entry decisions, suppose that both players entered. Each player chooses her action as follows.
 - Any type of a winner player randomly chooses her action over the interval $(\underline{a} + \eta, \underline{a} + 2\eta)$, following a continuous distribution function $G(a)$ which satisfies $G(\underline{a} + \eta) = 0$, $G(\underline{a} + 2\eta) = 1$, and (12).
 - A low-quality loser chooses $\underline{a} + 2\eta$.
 - A high-quality loser chooses $\underline{a} + \eta$.

State i with $i \geq 1$: Under this state, each player has a singleton message space. The stage action is exactly the same as the one at state i of the debarment equilibrium specified in the proof of Proposition 3.2.

Now we specify the transition rule for the states.

1. Suppose the state in the current period is 0.
 - (a) If player 1 is a winner but did not enter, then the state in the next period is 2.
 - (b) If player 1 announced H , solely entered, and chose an action other than a_H^* , then the state in the next period is 2.

- (c) If both players entered with player i being a winner, and if player i chose an action equal to or less than $\underline{a} + \eta$, then the state in the next period is $3 - i$.
 - (d) If none of the above applies, then the state in the next period is 0.
2. Suppose the state in the current period is i with $i \geq 1$.
 - (a) If both players entered, and if player i chose an action equal to or less than $\underline{a} + \eta$, then the state in the next period is $3 - i$.
 - (b) If the above does not hold, then the state in the next period is i .

We combine $\hat{\sigma}$ with a particular consistent system of beliefs. For that purpose, we arbitrarily fix a tremble satisfying the following requirements.

- At any history, the probability of any (possibly off-the-path) stage action depends only on the state given the history.
- Suppose the state in the current period is 0. Regardless of player 1's announcement, an entry by a high-quality loser is infinitely more likely than an entry by a low-quality loser.
- Suppose the state of the current period is i with $i \geq 1$. An entry by a high-quality player $j \neq i$ is infinitely more likely than an entry by a low-quality player $j \neq i$.

The system of beliefs associated with the tremble satisfies the following conditions.

1. Suppose the state in the current period is 0, and both players entered. Then the winner believes the loser to be high-quality.
2. Suppose the state of the current period is i with $i \geq 1$, and both players entered. Then player i believes player $j \neq i$ to be high-quality.

If the players follow $\hat{\sigma}$, full collusion is achieved, because a player with higher quality always and solely enters and obtains her mono-entrant payoff. Note that the payoff of $\hat{\sigma}$ is $(\mu_{HL} + \mu_{HH})\pi_H$ for player 1 and $\mu_{LL}\pi_L + \mu_{LH}\pi_H$ for player 2. Now we verify sequential rationality of $\hat{\sigma}$ at each state, given the beliefs specified above.

State 0: We work backwards within the period.

1. Suppose both players entered. Irrespective of player 1's announcement, the loser believes any type of the winner to randomly choose her action following the distribution function $G(a)$. Then, the same line of the argument as the proof of Proposition 3.3 proves that $\hat{\sigma}$ prescribes an optimal action for each type of the loser.

We have specified the beliefs so that any type of the winner believes the loser to be high-quality and therefore to choose $\underline{a} + \eta$. Then, the same line of the argument

as the proof of Proposition 3.3 proves that $\hat{\sigma}$ prescribes an optimal (randomized) action for a low-quality winner.

Consider a high-quality winner player i . Her continuation payoff when she chooses $a > \underline{a} + \eta$ is $\delta(\mu_{HL} + \mu_{HH})\pi_H$ if $i = 1$, and $\delta(\mu_{LL}\pi_L + \mu_{LH}\pi_H)$ if $i = 2$. This is because the state in the next period is 0. If she chooses $a \leq \underline{a} + \eta$, the state in the next period is $3 - i$ given $\hat{\sigma}$. Therefore, her continuation payoff is at most

$$(1 - \delta)U(a, H) \leq (1 - \delta)U(\underline{a} + \eta, H) \leq \delta \min \{(\mu_{HL} + \mu_{HH})\pi_H, \mu_{LL}\pi_L + \mu_{LH}\pi_H\},$$

where the first inequality follows from (11), and the second from (18). Therefore, it is optimal to randomly choose her action following the distribution function $G(a)$, as $\hat{\sigma}$ prescribes.

2. Suppose player i is a single entrant. Unless the entrant is a low-quality player 1 who announced H , $\hat{\sigma}$ prescribes player i to choose an action which yields her the mono-entrant payoff and ensures that the state in the next period is 0. Hence, it is optimal for player i to choose that action, because choosing any other action may lead to a smaller stage payoff and may cause the state to move to $3 - i$ (and a zero continuation payoff) in the next period.

So, suppose the entrant is a low-quality player 1 who announced H . If she chooses a_H^* , her continuation payoff is $(1 - \delta)U(a_H^*, L) + \delta(\mu_{HL} + \mu_{HH})\pi_H$, since the state in the next period is 0. If she chooses $a \neq a_H^*$, the state in the next period is 2 and hence her continuation payoff is at most $(1 - \delta)\Pi_L$, which is attained at $a = a_L^*$. Therefore, if (19) holds, it is optimal to choose a_H^* , as $\hat{\sigma}$ prescribes. Otherwise, it is optimal to choose a_L^* , as $\hat{\sigma}$ prescribes.

3. Let player 1's announcement be given, and consider player i .

First, suppose player i is a loser. If she enters, she believes any type of the winner to randomly choose his action following the distribution function $G(a)$. Hence, the same line of the argument as the proof of Proposition 3.3 proves that it is optimal for the loser not to enter, as $\hat{\sigma}$ prescribes.

Next, suppose player i is a winner. If she does not enter, her continuation payoff is zero because the state in the next period is $3 - i$. Unless player i is a low-quality player 1 who announced H , her continuation payoff upon entry is positive. This is because the state in the next period is 0, and therefore both the current stage payoff and the continuation payoff from the next period on are positive. Thus, it is optimal to enter, as $\hat{\sigma}$ prescribes.

So suppose player i is a low-quality player 1 who announced H . If she enters, her

continuation payoff is

$$\max \{(1 - \delta)\hat{\pi}_L + \delta(\mu_{HL} + \mu_{HH})\pi_H, (1 - \delta)\pi_L\} > 0, \quad (20)$$

because $\pi_L > 0$. As a result, it is optimal to enter, as $\hat{\sigma}$ prescribes.

4. At the beginning of the current period, a high-quality player 1's continuation payoff is $(1 - \delta)\pi_H + \delta(\mu_{HL} + \mu_{HH})\pi_H$ when she chooses H as prescribed by $\hat{\sigma}$. If she rather chooses L , she will be a loser and her continuation payoff is $\delta(\mu_{HL} + \mu_{HH})\pi_H$. Since $\pi_H > 0$, it is optimal for a high-quality player 1 to choose H .

A low-quality player 1's continuation payoff is $\delta(\mu_{HL} + \mu_{HH})\pi_H$ when she chooses L as prescribed by $\hat{\sigma}$. If she rather chooses H , her continuation payoff equals the left-hand-side of (20). Since $\hat{\pi}_L \leq 0$, we have

$$(1 - \delta)\hat{\pi}_L + \delta(\mu_{HL} + \mu_{HH})\pi_H \leq \delta(\mu_{HL} + \mu_{HH})\pi_H.$$

Since (13) holds, we also obtain

$$(1 - \delta)\pi_L \leq \delta(\mu_{HL} + \mu_{HH})\pi_H.$$

From these, it is optimal for a low-quality player 1 to choose L .

State i with $i \geq 1$: The same line of arguments which verify sequential rationality of the debarment equilibrium (Proposition 3.2) proves sequential rationality.

Since we have examined all possibilities, the proof is complete. \square

The key condition for Proposition 4.1, $\hat{\pi}_L \leq 0$, is a necessary condition for any fully collusive equilibrium with one-sided communication to exist. Since only one player can reveal her type, full collusion requires the player not to enter if she is low-quality.¹⁸ Thus in equilibrium, her stage payoff is zero if she is low-quality. If $\hat{\pi}_L > 0$, the player could profitably deviate by mimicking the behavior of her high-quality variant, and therefore one-sided communication does not help at all.

Proposition 4.1 shows that under an additional condition (17), the equilibrium with one-sided communication sustains full collusion under a wider range of discount factors as compared to the equilibrium we considered in Proposition 3.3. This exemplifies importance of studying asymmetric equilibria even under symmetric environments. Note that the left-hand-side of (17) equals the equilibrium payoff of the player who makes

¹⁸If she enters with a positive probability, then an event where a low-quality player enters though the other player is high-quality would have a positive probability. This is a contradiction against total payoff maximization.

announcements every period. The condition states that this payoff is much larger than π^* which is the payoff of the equilibrium we considered in Proposition 3.3. Such larger payoff provides a stronger incentive to be truthful under one-sided communication, making full collusion an equilibrium outcome for a wider range of discount factors.

Since $\pi_H > 2\pi^*$, a sufficient condition for (17) is

$$2(\mu_{HL} + \mu_{HH}) \geq 1 + \rho_L.$$

This requires that the probability of a player being high-quality is large. This increases the equilibrium payoff of the player solely announcing her type and hence strongly incentivising her to follow the collusive equilibrium.¹⁹

Proposition 4.1 limits attention to a particular equilibrium where one player always announces her type on the path. Does any fully collusive equilibrium in which the player announcing her type depends on the history outperform the equilibrium in the sense of sustaining full collusion for a wider range of discount factors? The following result shows that the answer is ‘no’, as long as the collusion scheme where one player always announces her type outperforms the equilibrium constructed in the previous section. In other words, when we examine whether equilibria with one-sided communication help, the collusion scheme with a fixed player announcing her type every period is a representative.

Proposition 4.2. *Suppose (17) holds. Under any $\delta < \hat{\delta}$, no fully collusive equilibrium exists where at any history on the path, only one player announces her type.*

Proof. Fix $\delta < \hat{\delta}$, and suppose a fully collusive equilibrium exists where at any history on the path, only one player announces her type. Fix any period t and any history at the beginning of period t , and suppose player i announces her type at the history. Then it follows that (i) if player i is high-quality, she solely enters with probability 1, and that (ii) if player i is low-quality, the other player solely enters with probability 1. This is because, otherwise, either no entry, double-entry, or an event that a low-quality player enters when the other player is high-quality occurs with a positive probability, a contradiction against total payoff maximization. Since full collusion requires the entrant to choose an action attaining her mono-entrant payoff according to her type, it follows that (i) player i 's expected stage payoff is $(\mu_{HL} + \mu_{HH})\pi_H$, and that (ii) the expected stage payoff of player $j \neq i$ is $\mu_{LL}\pi_L + \mu_{LH}\pi_H$.

Since this argument is valid at all histories on the path, it follows that any player's continuation payoff at any history on the path is at most $(\mu_{HL} + \mu_{HH})\pi_H$. This is

¹⁹Note that this implication is purely due to the type distribution and does not depend on the payoff structure, such as the relative size of π_H and π^* .

because (17) and the definition of π^* imply

$$(\mu_{HL} + \mu_{HH})\pi_H - (\mu_{LL}\pi_L + \mu_{LH}\pi_H) > (1 + \rho_L)\pi^* - \pi^* = \rho_L\pi^* > 0. \quad (21)$$

Suppose player i announces her type at the initial history. If player i is low-quality and if she makes an announcement in the same way as her high-quality variant, then she will be prescribed to solely enter. If she then chooses an action attaining her mono-entrant payoff as a low-quality player, her continuation payoff is at least $(1 - \delta)\pi_L$, while her equilibrium continuation payoff is at most $\delta(\mu_{HL} + \mu_{HH})\pi_H$ by the above argument. Thus, an equilibrium condition is

$$(1 - \delta)\pi_L \leq \delta(\mu_{HL} + \mu_{HH})\pi_H,$$

which reduces to $\delta \geq \hat{\delta}$. This is a contradiction, and the proof is complete. \square

The intuition behind Proposition 4.2 is straightforward. As discussed before, the future continuation payoff provides the incentive for a player who solely makes an announcement to be truthful. Hence, the most effective way to provide incentives is to make the continuation payoff as large as possible. Condition (17) implies that the stage payoff of a player who announces her type is greater than that of the other player (see (21)). Hence, letting one player make announcements in all periods is most effective.

References

- [1] Aguirregabiria, V. and Ho, C.-Y. (2012): “A Dynamic Oligopoly Game of the US Airline Industry: Estimation and Policy Experiments”, *Journal of Econometrics*, 168, 156-173.
- [2] Aoyagi, M. (2003): “Bid Rotation and Collusion in Repeated Auctions”, *Journal of Economic Theory*, 112, 79-105.
- [3] Aoyagi, M. (2007): “Efficient Collusion in Repeated Auctions with Communication”, *Journal of Economic Theory*, 134, 61-92.
- [4] Asker, J. and Cantillon, E. (2008): “Properties of Scoring Auctions”, *RAND Journal of Economics*, 39, 69-85.
- [5] Athey, S. and Bagwell, K. (2001): “Optimal Collusion with Private Information”, *RAND Journal of Economics*, 32, 3, 428-465.
- [6] Athey, S. and Bagwell, K. (2008): “Collusion with Persistent Cost Shocks”, *Econometrica*, 76, 3, 493-540.
- [7] Athey, S., Bagwell, K., and Sanchirico, C. (2004): “Collusion and Price Rigidity”, *Review of Economic Studies*, 71, 2, 317-349.

- [8] Belford, C. (2007): “Entry, Airline Competition and Network Structure”, *Mimeo*, McGill University.
- [9] Binmore, K. (2007): *Playing for Real*, Oxford University Press.
- [10] Blume, A. (2003): “Bertrand without Fudge”, *Economics Letters*, 78, 2, 167-168.
- [11] Compte, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring”, *Econometrica*, 66, 597-626.
- [12] Dixit, A. and Shapiro, C. (1986): “Entry Dynamics with Mixed Strategies”. In L.G. Thomas, editor, *The Economics of Strategic Planning: Essays in Honor of Joel Dean*, Lexington, Mass.: Lexington Books.
- [13] Elberfeld, W. and Wolfstetter, E. (1999): “A Dynamic Model of Bertrand Competition with Entry”, *International Journal of Industrial Organization*, 17, 513-525.
- [14] Federal Aviation Administration, U.S. (1999): “*Airport Business Practices and Their Impact On Airline Competition*”, US Federal Aviation Administration, Task Force Study.
- [15] Farrell, J. (1987): “Cheap Talk, Coordination, and Entry”, *RAND Journal of Economics*, 18, 34-39.
- [16] Gerlach, H. (2009): “Stochastic Market Sharing, Partial Communication and Collusion”, *International Journal of Industrial Organization*, 27(6), 655-666.
- [17] Hanazono, M. and Yang, H. (2007): “Collusion, Fluctuating Demand and Price Rigidity”, *International Economic Review*, 48(2), 483-515.
- [18] Kandori, M. and Matsushima, H. (1998): “Private Observation, Communication and Collusion,” *Econometrica*, 66, 627-652.
- [19] Kreps, D. and Wilson, R. (1982): “Sequential Equilibria,” *Econometrica*, 50, 863-894.
- [20] Marshall, R. and Marx, L. (2012): *The Economics of Collusion: Cartels and Bidding Rings*, MIT Press.
- [21] Olszewski, W. and Safronov, M. (2018): “Efficient Cooperation by Exchanging Favors”, *Theoretical Economics*, 13, 1191-1231.
- [22] Park, I.-U. (2002): “Cheap-Talk Coordination of Entry by Privately Informed Firms”, *RAND Journal of Economics*, 33, 377-393.
- [23] Patra, R. (2017): “Collusion with Private Information and Fixed Costs”, *Mimeo*, University of London (Royal Holloway College).
- [24] Skrzypacz, A. and Hopenhayn, H. (2004): “Tacit Collusion in Repeated Auctions”, *Journal of Economic Theory*, 114, 1, 153-169.

- [25] Sutton, J. (1995): *Sunk Costs and Market Structure*, MIT Press.
- [26] Wen, Q. (2002): “A Folk Theorem for Repeated Sequential Games”, *Review of Economic Studies*, 69, 493-512.