

Repeated Multimarket Contact with Private Monitoring: A Belief-Free Approach*

Atsushi Iwasaki,^{*,§} Tadashi Sekiguchi,^{†,§} Shun Yamamoto,[‡] Makoto Yokoo^{‡,§}

*University of Electro-Communications, †Kyoto University, ‡Kyushu University, §RIKEN AIP
a2c.iwasaki@gmail.com, sekiguchi@kier.kyoto-u.ac.jp, syamamoto@agent.inf.kyushu-u.ac.jp, yokoo@inf.kyushu-u.ac.jp

Abstract

This paper studies repeated games where two players play multiple duopolistic games simultaneously (multimarket contact). A key assumption is that each player receives a noisy and private signal about the other's actions (*private monitoring* or *observation errors*). There has been no game-theoretic support that multimarket contact facilitates collusion or not, in the sense that more collusive equilibria in terms of *per-market* profits exist than those under a benchmark case of one market. An equilibrium candidate under the benchmark case is *belief-free* strategies. We are the first to construct a non-trivial class of strategies that exhibits the effect of multimarket contact from the perspectives of *simplicity* and *mild punishment*. Strategies must be simple because firms in a cartel must coordinate each other with no communication. Punishment must be mild to an extent that it does not hurt even the minimum required profits in the cartel. We thus focus on two-state automaton strategies such that the players are cooperative in at least one market even when he or she punishes a traitor. Furthermore, we identify an additional condition (*partial indifference*), under which the collusive equilibrium yields the optimal payoff.

Introduction

This paper investigates a simple but fundamental question: Can two players cooperate better when they confront in multiple repeated games than in a single repeated game? A typical context we concentrate on is *multimarket contact*. For example, global enterprises, such as Uber or Lyft, provide their services in multiple distinct markets. In each area, they face an oligopolistic competition, which is often modeled as a prisoners' dilemma (PD). When they repeatedly confront in a long run (*repeated games*), more cooperative or collusive behavior can be an equilibrium. They may be more likely to collude, and a regulatory agency wants to estimate the extent (Chellappaw, Sambamurthy, and Saraf 2010).

However, the answer to the above question is negative if we assume that each player can directly observe his op-

ponent's actions (*perfect* monitoring), which is an assumption often used in computer science literature. Bernheim and Whinston (1990) show that multimarket contact does not improve the most collusive per-market equilibrium profit under perfect monitoring, though some of the vast empirical studies rather suggest the opposite (Evans and Kessides 1994). To resolve this discrepancy, this paper assumes *private* monitoring where each player may observe a different signal. For example, although a firm cannot directly observe its rival's action, e.g., prices, it can observe a noisy private signal, e.g., its own sales amounts

Analytical studies on this class of games have not been very successful (Mailath and Samuelson 2006). This is because characterizing all equilibria or identifying optimal equilibria in games with private monitoring is extremely hard. Indeed, equilibrium candidates may be so complicated as to be represented only by automaton strategies with a very large state space, and the complexity, together with privacy of the other player's signals, would require very involved statistical inferences to estimate the other player's history at any period and to check optimality of the continuation strategy (Kandori 2010). Notably, a *belief-free* approach has successfully established a general characterization where an equilibrium strategy is constructed so that the statistical inferences do not matter (Ely and Välimäki 2002; Ely, Horner, and Olszewski 2005). However, it is not obvious whether the belief-free approach is helpful in examining the effects of multimarket contact. This is because we want to deal with any number of markets, which causes the number of available actions to exponentially increase and may diminish its tractability.

The main goal of this paper is to construct, under multimarket contact with private monitoring, a non-trivial class of strategies which can sustain a better per-market outcome than an equilibrium strategy for a single market. The secondary concern is twofold. First, strategies must be simple. Firms do not communicate with each other after they form a cartel so that their cartel could be escaped from monitoring by a regulatory agency. If they employ a complicated strategy, it will be hard to detect one's deviation and to restore its collusion after punishment. We thus concentrate on simple two-state automaton strategies, which are still diffi-

*A full version can be found at <http://arxiv.org/abs/1607.03583>. Supported in part by KAKENHI 17H01787, 16KK0003, and 17H00761.

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

cult to analyze. Second, we assume firms punish a traitor in only some market, instead of all markets. When a cartel forms, products therein may not be so differentiated and the profit from the market may be small. If a firm punishes a traitor in all markets, it may not be able to earn the minimum necessary profit. The more collusive equilibrium thus prescribes the players to be cooperative in at least one market even under a punishment state. Surprisingly, our analysis reveals that information from such markets is crucial to admit the multimarket contact effect.

When we employ the belief-free approach, the strategy found by the work of Ely and Välimäki is an important benchmark and it attains the optimal payoff among belief-free equilibria in PD (Ely and Välimäki 2002). This strategy, which we call EV, can form a belief-free equilibrium under private monitoring in the single market case and attains high expected payoffs with a wide range of parameter settings. Figure 1 illustrates EV, which is a variant of the well-known tit-for-tat strategy.¹ A player first cooperates and keeps cooperation as long as she observes a signal suggesting cooperation. Once she observes a signal suggesting defection, she defects with a given probability and cooperates with the remaining probability. Similarly, when she defects, she keeps defection as long as she observes a signal suggesting defection. Once she observes a signal suggesting cooperation, she returns to cooperation with another given probability and defects with the remaining probability.

Building upon those considerations, we provide a condition under which a generalization of the benchmark belief-free strategies, which we call the *generalized EV* (gEV) strategy, forms an equilibrium. To the best of our knowledge, we are the first to identify an equilibrium designed for multiple markets whose per-market equilibrium payoffs exceed one for the benchmark strategies. Furthermore, we find an additional condition, which we call *partial indifference*. It implies that each player is indifferent among all strategies which differ only in play in a given subset of the markets. Under this condition, the gEV equilibrium yields the optimal payoff.

Let us finally note a simple extension of the EV strategy to multimarket contact. We show that for any number of markets, it is an equally collusive equilibrium under the same condition as the benchmark case. The equilibrium satisfies a special case of partial indifference (*total indifference*), where the subset consists of all markets and is most collusive among such equilibria. Moreover, we reveal that, under the condition, each player's continuation play depends only on the number of signals suggesting the other player's defection in the previous period in a linear manner. This result implies that it must be designed in such a non-linear manner as the gEV strategy.

¹Here, g and b are noisy private signals suggesting that the opponent's action is C and D , respectively. ε_R or ε_P represents the transition probability between states. We omit the transitions for the remaining probabilities.

Model

Two players play M PDs simultaneously in each period. In each PD, each player chooses either C (cooperation) or D (defection). This is regarded as a model of oligopolistic competition, where C is an action increasing the total payoffs (for instance, in the case of price competition, charging a collusive high price), and D is a non-cooperative one (like a price cut). The players can choose different actions over the M PDs, so that each player's action set in each period is $\{C, D\}^M$.

Each player cannot directly observe the other player's actions, but receives an imperfect signal about them. In each PD, each player receives either a good signal g or a bad signal b . We assume that each player receives his signals individually, and cannot observe the other player's signals (private monitoring). The pair of signals they privately receive in each PD is stochastic, following a common symmetric probability distribution that depends entirely on the action pair of that PD. We denote it by $o(\omega_1, \omega_2 | a_1, a_2)$, where $(\omega_1, \omega_2) \in \{g, b\}^2$ and $(a_1, a_2) \in \{C, D\}^2$. We assume that the signals across the M PDs are independent, though the signals of a given PD may be correlated across the players. We also assume that the signal distributions are described by one parameter. There exists $p \in (1/2, 1)$ such that for any i , any ω_j ($j \neq i$) and any $a \in \{C, D\}^2$,

$$\sum_{\omega_i \in \{g, b\}} o(\omega_i, \omega_j | a) = \begin{cases} p & \text{if } (a_i, \omega_j) \in \{(C, g), (D, b)\}, \\ 1 - p & \text{otherwise.} \end{cases}$$

The marginal distribution of an individual signal in a given PD is such that the *right* signal ($\omega_j = g$ if $a_i = C$, and $\omega_j = b$ if $a_i = D$) is received with probability p . We let $s = 1 - p$, which is the probability of an *error*. The assumption is consistent with *conditionally independent* monitoring, which is a representative monitoring structure. Formally, a signal distribution is conditionally independent if $o(\omega_i, \omega_j | a) = o(\omega_i | a) o(\omega_j | a)$ for all ω_i, ω_j , and a . Also, it is consistent with nearly perfect monitoring (when p is close to 1), but inconsistent with nearly public monitoring (namely, the case where the event $\omega_1 = \omega_2$ is much more likely than $\omega_1 \neq \omega_2$).

In each PD, player i 's payoff depends only on his action and the signal of that PD. The payoff function is common to all PDs, denoted by $\pi_i(a_i, \omega_i)$. We are more interested in the expected payoff function:

$$g_i(a_1, a_2) = \sum_{(\omega_1, \omega_2)} \pi_i(a_i, \omega_i) o(\omega_1, \omega_2 | a_1, a_2).$$

We assume that their expected payoff functions are represented by the following payoff matrix:

	C	D
C	$1, 1$	$-y, 1 + x$
D	$1 + x, -y$	$0, 0$

We assume $x > 0$, $y > 0$ and $1 > x - y$, so that it indeed represents a PD.

All M PDs are played infinitely, in periods $t = 0, 1, 2, \dots$. Player i 's *private history* at the beginning of period $t \geq 1$ is an element of $H_i^t \equiv [\{C, D\}^M \times \{g, b\}^M]^t$. Let H_i^0 be

an arbitrary singleton, and let $H_i = \cup_{t \geq 0} H_i^t$ be the set of player i 's all private histories. Player i 's strategy of this repeated game is a mapping from H_i to the set of all probability distributions over $\{C, D\}^M$. That is, we allow randomized strategies. If the actual play of the repeated game is such that the action pair $(a_1^m(t), a_2^m(t))$ is played in the m -th PD in period t for each m and t , player i 's normalized average payoff is

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t \sum_{m=1}^M g_i(a_1^m(t), a_2^m(t)), \quad (1)$$

where $\delta \in (0, 1)$ is their common discount factor. The average payoff of any strategy pair is the expected value of Eq. 1, where the expectation is taken with respect to the players' randomizations and the monitoring structure.

Belief-Free Equilibrium

The solution concept for repeated games with imperfect monitoring is *sequential equilibrium* (Kreps and Wilson 1982). However, since it is still highly difficult to analyze our model, we here focus on a special class called *belief-free equilibria* (Ely, Horner, and Olszewski 2005), which is standard in the private monitoring literature (Mailath and Samuelson 2006).

Definition 1 (Belief-Free Equilibrium). *A strategy pair is a belief-free equilibrium if for any $t \geq 0$, $h_1^t \in H_1^t$ and $h_2^t \in H_2^t$, each player i 's continuation strategy given h_i^t is optimal against player j 's continuation strategy given h_j^t .*

An important property is that, while player i given her private history should, in principle, optimize her continuation payoff against her belief about player j 's history (and hence his continuation strategy), her continuation strategy is optimal even if she were to know j 's history with certainty.² In other words, the players playing a belief-free equilibrium need not compute their beliefs in the course of play. When a strategy pair is represented by finite-state automaton strategies, as will be the case in subsequent analysis, it is a belief-free equilibrium if any player's continuation strategy (behavior expanded from the automaton) starting from any state is a best response (optimal) against the other player's continuation strategy starting from any state. Note that we never restrict the other's possible strategy space, which includes strategies with an infinite number of states.

Suppose both players employ a common strategy represented by a two-state automaton with state space $\{R, P\}$. Let $V_{s_1 s_2}$, where $s_1 \in \{R, P\}$ and $s_2 \in \{R, P\}$, be player 1's continuation payoff when (i) player 2 is currently at s_2 and then follows the automaton, and (ii) player 1 always plays the action prescribed at state s_1 at any subsequent history. The strategy pair is a belief-free equilibrium if and only if there exist V_R and V_P such that

$$V_{RR} = V_{PR} = V_R, \quad V_{RP} = V_{PP} = V_P, \quad (2)$$

and that V_{s_2} ($s_2 \in \{R, P\}$) is player 1's best response payoff against player 2's continuation strategy when he is at

²We refer to player i or 1 as *her* and to player j or 2 as *him* throughout this paper.




Figure 1: EV strategy




Figure 2: sEV strategy

state s_2 . To see this, note that by Eq. 2, player 1 at any history is indifferent between her continuation strategy at state R and that at state P irrespective of her belief about player 2's state. Since the second condition implies that both continuation strategies give her best response payoff at any history, the conditions for belief-free equilibrium are all satisfied.

We shall consider a general class of two-state automaton strategies throughout this paper.

Definition 2 (Two-State Automaton Strategies). *The state space is $\{R, P\}$, and R is the initial state. At state $s \in \{R, P\}$, the player is prescribed to choose $a^s \in \{C, D\}^M$. Suppose the current state is $R(P)$. If $\omega = (\omega_k)_{k=1}^M \in \{g, b\}^M$ is observed, then the state shifts to $P(R)$ with probability $\xi_R(\omega)$ ($\xi_P(\omega)$) (and stays at $R(P)$ with the remaining probability). Note that this class is parameterized by $\xi_R(\cdot)$ and $\xi_P(\cdot)$: For any $s \in \{R, P\}$, $\xi_s : \{g, b\}^M \rightarrow [0, 1]$, which we call the transition probability functions.*

Ely-Välimäki Strategy

Let us next explain an instance of this class for a single PD case ($M = 1$). Figure 1 illustrates the EV strategy (Ely and Välimäki 2002) where $a^R = C$, $a^P = D$, $\xi_R(g) = 1$, $\xi_R(b) = \varepsilon_R$, $\xi_P(b) = 1$, and $\xi_P(g) = \varepsilon_P$. Note that

$$\varepsilon_R = \frac{(1 - \delta)x}{\delta \{2p - 1 - (1 - p)(x + y)\}}, \text{ and}$$

$$\varepsilon_P = \frac{(1 - \delta)y}{\delta \{2p - 1 - (1 - p)(x + y)\}}.$$

A solid line denotes a deterministic transition and a dashed line denotes a probabilistic transition, though, for simplicity, we omit some state transitions. EV is a representative two-state automaton strategy that forms a belief-free equilibrium under repeated games with private monitoring and attains the highest average payoff among belief-free equilibria in PD. A player first cooperates at state R , but after observing a bad signal, she punishes (defects) at the next period with probability ε_R , or keep cooperation with $1 - \varepsilon_R$. Likewise,

after she defects at P , if she observes a good signal, she returns cooperation with ε_P , or keep defection with $1 - \varepsilon_P$.

Proposition 1. *There exist $\varepsilon_R \in [0, 1]$ and $\varepsilon_P \in [0, 1]$ such that the EV strategy pair is a belief-free equilibrium if*

$$\delta[2p - 1 - (1 - p)(x + y) + \max\{x, y\}] \geq \max\{x, y\}. \quad (3)$$

The average payoff starting from state R is

$$V_R = V^{EV} \equiv 1 - \frac{(1 - p)x}{2p - 1}.$$

Simplified EV Equilibrium

What happens if there are $M (\geq 2)$ PDs, in comparison with the case of one PD? If EV forms an equilibrium, it is always an equilibrium to play it in each PD independently. Obviously, the payoff of this equilibrium is M times the EV equilibrium payoff. Under this equilibrium, a player's actions in all PDs can be quite different, depending on the histories of individual PDs. Thus, the corresponding automaton has 2^M states.

Interestingly, this equilibrium strategy can be greatly simplified so that it is an equilibrium with the same payoff and under the same condition. The simplified strategy just has two states, where a player cooperates in all PDs at one state and defects in all PDs at the other. Her actions are therefore perfectly correlated across the PDs. The class of strategies, which we call *simplified EV* (sEV), is simple, also in the sense that its transition probabilities from one state to the other depend only on the number of bad or good signals. Therefore, the player need not to know the exact configuration of the signals.

Definition 3 (sEV Strategy). *An sEV strategy for $M (\geq 2)$ PDs is a two-state automaton strategy parameterized by two numbers $\varepsilon_R \in [0, 1]$ and $\varepsilon_P \in [0, 1]$:*

- The actions at each state are prescribed by $a^R = (C, C, \dots, C)$ and $a^P = (D, D, \dots, D)$,
- The transition probabilities are defined as

$$\xi_R(\omega) = |\{k \mid \omega_k = b\}|_{\varepsilon_R}, \text{ and } \xi_P(\omega) = |\{k \mid \omega_k = g\}|_{\varepsilon_P}.$$

Figure 2 illustrates sEV for two PDs in the same manner as Figure 1. The next theorem identifies the equilibrium condition and the average payoff starting from state R .

Theorem 1. *There exist $\varepsilon_R \in [0, 1]$ and $\varepsilon_P \in [0, 1]$ such that the sEV strategy pair is a belief-free equilibrium if Eq. 3 holds. The average payoff starting from state R is MV^{EV}*

We place the proof in the full version because it is similar to Proposition 1.

Generalized EV Equilibrium and Partially Indifference

The goal of the analysis of this section is to give one answer to whether a class of strategies which can yield a better per-market payoff than the EV strategy exists or not. Let us define the class of strategies as the *generalized EV* (gEV) strategy that achieves the optimal payoff

$$\bar{Y}^d \equiv M - \frac{1 - p}{2p - 1}(d + 1)x. \quad (4)$$

We also use

$$\underline{Y}^d \equiv M - d + \frac{1 - p}{2p - 1}dy + \frac{p^{M-d}}{p^{M-d} - (1 - p)^{M-d}}(M - d)x. \quad (5)$$

Definition 4 (gEV Strategy). *A gEV strategy for $M (\geq 2)$ PDs is a two-state automaton strategy parameterized by five parameters, $d \in \{1, 2, \dots, M - 1\}$, α^R , α^P , β^R , and β^P :*

$$\begin{aligned} \alpha^R &= \frac{(1 - \delta)x}{\delta(\bar{Y}^d - \underline{Y}^d)(2p - 1)}, \\ \alpha^P &= \frac{(1 - \delta)y}{\delta(\bar{Y}^d - \underline{Y}^d)(2p - 1)}, \\ \beta^R &= \frac{(1 - \delta)x}{\delta(\bar{Y}^d - \underline{Y}^d)(2p - 1)(1 - p)^{M-d-1}}, \\ \beta^P &= \frac{(1 - \delta)(M - d)x}{\delta(\bar{Y}^d - \underline{Y}^d)\{p^{M-d} - (1 - p)^{M-d}\}}. \end{aligned}$$

Let us divide the markets into $A = \{1, 2, \dots, M - d\}$ and $B = \{M - d + 1, M - d + 2, \dots, M\}$. The actions

at each state are prescribed by $a^R = \overbrace{(C, C, \dots, C)}^M$ and $a^P = \overbrace{(C, \dots, C)}^{M-d} \overbrace{(D, \dots, D)}^d$. The transition probabilities are defined as

$$\begin{aligned} \xi_R(\omega) &= \begin{cases} \alpha^R |\{k \in B \mid \omega_k = b\}| + \beta^R \text{ if } \omega_k = b \text{ for all } k \in A, \\ \alpha^R |\{k \in B \mid \omega_k = b\}| & \text{otherwise.} \end{cases} \\ \xi_P(\omega) &= \begin{cases} \alpha^P |\{k \in B \mid \omega_k = g\}| + \beta^P \text{ if } \omega_k = g \text{ for all } k \in A, \\ \alpha^P |\{k \in B \mid \omega_k = g\}| & \text{otherwise.} \end{cases} \end{aligned}$$

Let us explain how we construct this strategy. A player cooperates in A at state P . Then, she always cooperates PDs in A regardless of which state she is in. The transition probabilities from R to P distinguish signals from A with from B . The increase of transition probabilities is constant for the number of bad signals from A . If she observes at least one good signal from A , it is zero, otherwise, β^R . The transition probabilities further increase by α^R in the number of bad signals from B . Similarly, the transition probabilities from P to R are specified. Their increase is constant for the number of good signals from A . If she observes at least one bad signal from A , it is zero, otherwise, β^P . The transition probabilities increase by α^P in the number of good signals from B .

The following theorem identifies a condition for an equilibrium outperforming the sEV equilibrium.

Theorem 2. *If $\bar{Y}^d > \underline{Y}^d$, there exists $\delta \in (0, 1)$ such that for any $\delta \geq \delta$, there is a belief-free equilibrium with $V_R = \bar{Y}^d$ and $V_P = \underline{Y}^d$.*

The payoffs equal the upper and the lower bounds derived in Theorems 3 and 4. Those bounds of equilibrium payoffs are obtained under the following condition, which we call *partial indifference*. Recall that $A = \{1, 2, \dots, M - d\}$ and $B = \{M - d + 1, M - d + 2, \dots, M\}$ by parameter d . Let $V_s(a)$ be the continuation payoff when player 2 is at state s and player 1 chooses action a .

Definition 5 (Partial Indifference). Fix $d \in [1, M - 1]$. An equilibrium is d -partially indifferent if for any $s \in \{R, P\}$, any $a', a'' \in \{C, D\}^M$, and any $k \in A$, if $a'_k = a''_k$, $V_s(a') = V_s(a'')$.

The d -partial indifference means that, given the other player's strategy, a player's repeated game payoff does not depend on the play of the B markets once he fixes a play of the A markets. For example, if the player cooperates in all A markets (as in the equilibrium), he receives a common repeated game payoff regardless of his actions in the B markets. Or if he defects in all A markets, he receives a common repeated game payoff regardless of his actions in the B markets. Similar indifference holds for any play of the A markets. In other words, this concept implies that his repeated game payoff given the other player's strategy depends solely on his play in the A markets. Further, our analysis shows that the repeated game payoff is maximized when the player cooperates in all A markets. Thus, the strategy is optimal among all possible strategies.

Let us outline the proof of Theorem 2. We calculate both the upper (lower) bound of the continuation payoffs from state R (P) among all d -partially indifferent equilibria among the class of strategies in Definition 2 (Theorems 3 and 4). Since our equilibrium uses a transition to R as a reward and a transition to P as a punishment, the upper bound must exceed the lower bound (Theorem 5). Under that condition, we construct a particular gEV equilibrium which achieves those upper and lower bounds if the discount factor is sufficiently large.

Suppose that a^R and a^P are prescribed as in Definition 4. Formally, given d , $a^R = (a_k^R)_{k=1}^M$, $a^P = (a_k^P)_{k=1}^M$ where for all k , $a_k^R = C$, for all $k \in A$, $a_k^P = C$, and for all $k \in B$, $a_k^P = D$.

Theorem 3. Fix $d \in [1, M - 1]$. Suppose a^R is prescribed as in Definition 4. Any d -partially indifferent belief-free equilibrium payoff is at most \bar{Y}^d defined in Eq. 4.

Proof. Consider an arbitrary d -partially indifferent belief-free equilibrium. Fix $k \in B$. For any $a_{-k} \in \{C, D\}^{M-1}$, d -partially indifferent belief-freeness implies

$$(1 - \delta)x = \delta(V_R - V_P) \left\{ \zeta_R(D, a_{-k}) - \zeta_R(C, a_{-k}) \right\}. \quad (6)$$

Note that $o_2(\omega_{-k} \mid a_{-k})$ is the probability that player 2 observes a signal profile ω_{-k} from the markets except k occurs when player 1 chooses an action profile a_{-k} . Besides $\zeta_s(a)$ indicates $\sum_{\omega \in \{g, b\}^M} o_2(\omega \mid a) \xi_s(\omega)$ for state s . Then we derive for any a_{-k} ,

$$(1 - \delta)x = \delta(V_R - V_P)(2p - 1) \sum_{\omega_{-k} \in \{g, b\}^{M-1}} \left\{ \xi_R(b, \omega_{-k}) - \xi_R(g, \omega_{-k}) \right\} o_2(\omega_{-k} \mid a_{-k}). \quad (7)$$

We have, for any $a_{-k} \in \{C, D\}^{M-1}$, $\zeta_R(D, a_{-k}) -$

$$\zeta_R(C, a_{-k}) = (2p - 1) \sum_{\omega_{-k} \in \{g, b\}^{M-1}} \left\{ \xi_R(b, \omega_{-k}) - \xi_R(g, \omega_{-k}) \right\} o_2(\omega_{-k} \mid a_{-k}).$$

Substituting this into Eq. 6, we obtain

$$\xi_R(b, \omega_{-k}) - \xi_R(g, \omega_{-k}) = \frac{(1 - \delta)x}{\delta(V_R - V_P)(2p - 1)}.$$

Intuitively, this implies the increase of the transition probabilities is constant when the signal observed in a market changes from the good to the bad one. Since we can arbitrarily choose $k \in B$, there exists a function $\beta^R : \{g, b\}^{M-d} \rightarrow [0, 1]$ such that

$$\xi_R(\omega) = \frac{(1 - \delta)x \{k \in B \mid \omega_k = b\}}{\delta(V_R - V_P)(2p - 1)} + \beta^R(\omega_1, \dots, \omega_{M-d})$$

for all $\omega \in \{g, b\}^M$. Note that the transition probability from state R to P is linear in the number of bad signals from B .

Let us consider the incentive condition that a player at state R defects in a single market. Recall that $a^R = (C, \dots, C)$ and let a'^R be an action that defects only in a market and cooperates in the remaining, e.g., (D, C, \dots, C) .

$$V_R \geq (1 - \delta)(M + x) + \delta V_R - \delta(V_R - V_P) \zeta_R(a'^R) \\ \Leftrightarrow (1 - \delta)x \leq \delta(V_R - V_P) \left\{ \zeta_R(a'^R) - \zeta_R(a^R) \right\}.$$

Since the left hand side of this equation is positive, we have $\zeta_R(a'^R) > \zeta_R(a^R)$. For $a \in \{a^R, a'^R\}$, it holds that

$$\zeta_R(a) = \frac{(1 - \delta)x(1 - p)d}{\delta(V_R - V_P)(2p - 1)} + \eta_R(a)$$

where $\eta_R(a) =$

$$\sum_{(\omega_1, \dots, \omega_{M-d}) \in \{g, b\}^{M-d}} \left[\prod_{k=1}^{M-d} o_2(\omega_k \mid a) \right] \beta^R(\omega_1, \dots, \omega_{M-d}).$$

Clearly, $\zeta_R(a'^R) - \zeta_R(a^R) = \eta_R(a'^R) - \eta_R(a^R) > 0$ holds. From those equations, let us transform the expected payoff starting from state R .

$$V_R = M - \frac{\delta}{1 - \delta} (V_R - V_P) \zeta_R(a^R) \\ \leq M - \frac{1 - p}{2p - 1} xd - \frac{x}{\eta_R(a'^R) - \eta_R(a^R)} \eta_R(a^R).$$

Note that $\eta_R(a)$ must be non-zero. If not, for any $(\omega_1, \dots, \omega_{M-d}) \in \{g, b\}^{M-d}$, $\beta^R(\omega_1, \dots, \omega_{M-d}) = 0$ and for any two $a, a' \in \{C, D\}^M$, $\eta(a') - \eta(a) = 0$. This contradicts $\eta_R(a'^R) - \eta_R(a^R) > 0$. Therefore, there exists a signal profile $(\omega_1, \dots, \omega_{M-d}) \in \{g, b\}^{M-d}$ such that $\beta^R(\omega_1, \dots, \omega_{M-d}) > 0$. We have

$$1 < \frac{\eta_R(a'^R)}{\eta_R(a^R)} \leq \frac{p}{1 - p}.$$

Accordingly, we finally obtain $V_R \leq \bar{Y}^d$. The proof is complete. \square

The upper bound always exceeds the benchmark, unless $d = M - 1$.

In this equilibrium, the fact that the d -partial indifference decreases the continuation payoff from R by $\frac{1-p}{2p-1}dx$ follows from the linearity of the transition probabilities in B . The remaining term of $\frac{1-p}{2p-1}x$ follows from the incentive condition that a player chooses C in A .

Let us next derive the lower bound.

Theorem 4. Fix $d \in [1, M-1]$. Suppose a^P is prescribed as in Definition 4. Among the class of strategies in Definition 2, any d -partially indifferent belief-free equilibrium payoff is at least \underline{Y}^d defined in Eq. 5.

Since the flow of this proof is the same as that of Theorem 3, we place it in the full version.

Before proving Theorem 2, we show in any d -partially indifferent belief-free equilibrium, $\bar{Y}^d \geq V_R > V_P \geq \underline{Y}^d$ must hold. Otherwise, none of the equilibrium exists. To this end, it is suffice to the following statement holds.

Theorem 5. Fix $d \in [1, M-1]$. In any d -partially indifferent belief-free equilibrium payoff, $V_R > V_P$.

Proof. Consider an arbitrary d -partially indifferent belief-free equilibrium. Recall that $\zeta_s(a)$ is $\sum_{\omega \in \{g,b\}^M} o_2(\omega | a)\xi_s(\omega)$ for state s . From Definition 2, we have

$$V_R = (1-\delta)M + \delta V_R - \delta(V_R - V_P)\zeta_R(a^R), \text{ and}$$

$$V_P = (1-\delta)(M-d) + \delta V_P + \delta(V_R - V_P)\zeta_P(a^P)$$

hold and we obtain

$$V_R - V_P = \frac{(1-\delta)d}{(1-\delta) + \delta\{\zeta_R(a^R) + \zeta_P(a^P)\}}.$$

Thus, $V_R > V_P$ holds because the right hand side is clearly positive. \square

Proof of Theorem 2. Fix $d \in [1, M-1]$ and suppose the gEV strategy. We claim that $\bar{Y}^d > \underline{Y}^d$ implies $\xi_s(\omega) \in [0, 1]$ for all $s \in \{R, P\}$ and $\omega \in \{g, b\}^M$. It is immediate that $\xi_s(\cdot)$ is always positive if $\bar{Y}^d > \underline{Y}^d$.

Next, define $\underline{\delta} \in (0, 1)$ as δ satisfying

$$\max\{\alpha^R d + \beta^R, \alpha^P d + \beta^P\} = 1.$$

Since d is the maximum value of $|\{k \in B \mid \omega_k = b\}|$ or $|\{k \in B \mid \omega_k = g\}|$, for any $\delta \geq \underline{\delta}$, it holds that $\xi_s(\omega) \leq 1$ for all $s \in \{R, P\}$ and $\omega \in \{g, b\}^M$. Therefore, the gEV strategy is verified.

By solving the following system of the equations, we obtain $V_R = \bar{Y}^d$ and $V_P = \underline{Y}^d$.

$$V_R = (1-\delta)g_R(a^R) + \delta V_R - \delta(V_R - V_P)\zeta_R(a^R)$$

$$= (1-\delta)g_R(a^P) + \delta V_R - \delta(V_R - V_P)\zeta_R(a^P).$$

$$V_P = (1-\delta)g_P(a^R) + \delta V_P + \delta(V_R - V_P)\zeta_P(a^R)$$

$$= (1-\delta)g_P(a^P) + \delta V_P + \delta(V_R - V_P)\zeta_P(a^P).$$

It suffices to verify that (i) V_R is a player's best response payoff when the other player is at state R , and (ii) V_P is a player's best response payoff when the other player is at state P . To this end, let $V_s(d_A, d_B)$ be a player's continuation payoff when he chooses D at some d_A PDs in A at some d_B PDs in B . Then it conforms to the strategy from the next period on, given that the other player is at state s and conforms to the strategy. The proof is complete if we show that $V_s(0, 0) \geq V_s(d_A, d_B)$ for any d_A , any d_B , and any s .

First, some calculation verifies

$$V_R(d_A, d_B) = (1-\delta)M + \delta V_R - \delta(V_R - V_P)d(1-p)\alpha^R$$

$$+ (1-\delta)x\left[d_A - \left(\frac{p}{1-p}\right)^{d_A} \frac{1-p}{2p-1}\right].$$

Hence, $V_R(d_A, d_B)$ does not depend on d_B . For any d_A and d_B , $V_R(d_A, d_B) - V_R(d_A + 1, d_B) =$

$$(1-\delta)x\left[-1 + \left(\frac{p}{1-p}\right)^{d_A}\right].$$

Since $p > 1-p$, $V_R(d_A, d_B) \geq V_R(d_A + 1, d_B)$ holds for any d_A and d_B . We obtain $V_R(d_A, d_B) =$

$$\sum_{k=1}^{d_A} \{V_R(k, d_B) - V_R(k-1, d_B)\} + V_R(0, d_B) \leq V_R(0, d_B)$$

for any d_A and d_B , as desired. Further, it attains the same value $-(1-\delta)x\frac{1-p}{2p-1}$ at $d_A = 0$ and $d_A = 1$. The concavity of the last term of $V_R(d_A, d_B)$ implies this consequence.

Finally, some calculation verifies

$$V_P(d_A, d_B) =$$

$$(1-\delta)(M-d-dy) + \delta V_P - \delta(V_R - V_P)dp\alpha^P$$

$$+ (1-\delta)x\left[d_A + \left(\frac{1-p}{p}\right)^{d_A} \frac{p^{M-d}(M-d)}{p^{M-d} - (1-p)^{M-d}}\right].$$

Hence, $V_P(d_A, d_B)$ does not depend on d_B . For any d_B , it attains the same value

$$(1-\delta)x\frac{p^{M-d}(M-d)}{p^{M-d} - (1-p)^{M-d}}$$

at $d_A = 0$ and $d_A = M-d$. Since $p > 1-p$, it is convex in d_A . This convexity implies that $V_P(0, d_B) \geq V_P(d_A, d_B)$ for any d_A and d_B . The proof is complete. \square

Theorem 2 reaches the main goal we raised at the beginning of this paper. Namely, if there exists an integer $d < M-1$ such that $\bar{Y}^d > \underline{Y}^d$, the gEV strategy yields the average per-market payoff greater than the EV and sEV strategies. Note that the number of markets M must be greater than three. The condition $\bar{Y}^d > \underline{Y}^d$ is equivalent to

$$d\left\{1 - \frac{1-p}{2p-1}(x+y)\right\} > \left\{\frac{1-p}{2p-1}\right.$$

$$\left. + (M-d)\frac{p^{M-d}}{p^{M-d} - (1-p)^{M-d}}\right\}x.$$

The coefficient of d in the left hand side must be positive because that of x in the right hand side is positive. This is the necessary condition that an EV equilibrium exists for some discount factor from Eq. 3. Unless EV is an equilibrium, for any d , no d -partially indifferent belief-free equilibrium exists.

sEV Equilibrium and Total Indifference

This subsection turns to the sEV equilibrium to deepen the understandings of gEV. Consider an extreme case of partial indifference. We restrict an equilibrium such that every action is indifferent in the continuation payoff regardless of the opponent's state.

Definition 6 (Total Indifference). *An equilibrium is totally indifferent if for any $s \in \{R, P\}$, any $a', a'' \in \{C, D\}^M$, $V_s(a') = V_s(a'')$.*

sEV clearly is totally indifferent. In addition, the next theorem claims the bound of the equilibrium payoff.

Theorem 6. *Suppose $a^R = (C, C, \dots, C)$. Among the class of strategies in Definition 2, any totally indifferent belief-free equilibrium payoff is at most MV^{EV} .*

We place the proof in the full version. We emphasize that the total indifference requires the signal in one market to have the same impact on the transition probabilities as the signal in any other market, as is seen from

$$\xi_R(\omega) = \frac{(1 - \delta)x|\{k \mid \omega_k = b\}|}{\delta(V_R - V_P)(2p - 1)} + \beta^R.$$

Hence, the sEV equilibrium is most collusive among all totally indifferent equilibria because it has $\beta^R = 0$ and therefore is least likely to switch to inefficient punishment. Since Theorem 6 does not specify behavior in state P , mild punishment is not sufficient to improve the per-market equilibrium payoff. Some nonlinearity of the transition probabilities is necessary. Signals from markets, where each player is prescribed to cooperate at any state, are essential to admit more collusive equilibria.

Discussions

Let us explain why gEV can form an equilibrium. A key feature involves the nonlinearity of the transition probabilities. In fact, the transition probabilities from state R to P do not depend on the outcome in A at all, as long as it contains at least one good signal. However, if all signals from A are bad, the transition probability sharply increases by β^R . Why does this nonlinearity help? Suppose the other player is at state R , and consider how a player wants to play the PDs in A . Her incentive to play C or D in one PD in A crucially depends on the probability of the event that all signals among the other PDs in A are bad. Only under that event, is her action in this PD pivotal. Naturally, the event is more likely when she defects among more PDs in A . Therefore, her temptation to defect in one PD in A is largest when she cooperates among all other PDs in A . Note that we apply a similar argument to this when we check the incentives. This observation implies that once gEV prevents a player from defecting in one PD in A , it automatically ensures that the

player has no incentive to defect in any number of PDs in A . Thus, as long as we consider gEV, we can effectively ignore all actions which defect among two or more PDs in A . This reduction in the number of incentive constraints is a key to the payoff improvement results brought about by gEV.

Next let us note why our gEV strategy outperforms the sEV strategy. First of all, if we had zero A market ($d=M$), gEV would be equal to sEV. Thus, gEV with a few A markets is a natural equilibrium candidate. Further, perpetual cooperation in the A markets leads to an improvement of the per-market equilibrium payoff. Relatedly, the presence of the A markets is a key part of our results and reflects some reality. A typical example is airline industries: cutting the fares in every route rarely occurs (except some promotion or campaign), and the fares in routes connecting major cities tend to be almost the same across airlines (Evans and Kessides 1994).

Another question is whether gEV achieves an optimal payoff among much more general equilibrium strategies than d -partially indifferent equilibria. Under perfect monitoring, it is known that a player's equilibrium payoff vector can be computed. Dynamic programming can derive the bounds of the equilibrium payoff of each player, i.e., a self-generation set (Abreu, Pearce, and Stacchetti 1990). The existence of an equilibrium strategy is guaranteed that attains a payoff vector in that set. Unfortunately, under private monitoring, this is generally impossible because the recursive structure under perfect or public monitoring does not persist. Checking the optimality is the immediate future work.

Related literature

In the literature of computer science, AI, and multi-agent systems, there are many streams associated with repeated games (Burkov and Chaib-draa 2013): The complexity of equilibrium computation (Littman and Stone 2005; Borgs et al. 2010; Andersen and Conitzer 2013), multi-agent learning (Blum and Monsour 2007; Conitzer and Sandholm 2007; Shoham and Leyton-Brown 2008), partially observable stochastic games (POSGs) (Hansen, Bernstein, and Zilberstein 2004; Doshi and Gmytrasiewicz 2006; Tennenholtz and Zohar 2009; Mescheder, Tuyls, and Kaisers 2011; Wunder et al. 2011), and so on. Among them, POSGs are most relevant because repeated games with private monitoring can be considered as a special case of POSGs. However, POSGs often impose partial observability on an opponent's strategy (behavior rule) and not on opponent's past actions (Mescheder, Tuyls, and Kaisers 2011; Wunder et al. 2011). They estimate an optimal (best reply) strategy against an unknown strategy (not always fixed) from perfectly observable actions (perfect monitoring). In contrast, we verify whether a given strategy profile is a *mutual* best reply after any history, i.e., finding an equilibrium, with partially observable actions (private monitoring). Thus, this paper also addresses understanding the gap between POSGs and repeated games with private monitoring in economics.

In fact, very few existing works have addressed verifying an equilibrium. Hansen, Bernstein, and Zilberstein (2004) develop an algorithm that iteratively eliminates dominated strategies. However, just eliminating dominated strategies

is not sufficient to find an equilibrium. Also, the algorithm is not applicable to an infinitely repeated game. Doshi and Gmytrasiewicz (2006) investigate the computational complexity of achieving equilibria in interactive POMDPs.

The economics literature has extensively studied another interesting class of repeated games where the players observe a common noisy signal (*public monitoring*). Kobayashi and Ohta (2012) studied this version of our model and showed that when multimarket contact facilitates collusion, the most collusive equilibrium payoff is attained by a variant of the trigger strategy. Hence, the players must defect in all markets in the punishment state. Our model has an opposite implication which rather favors mild punishment.

Conclusions

This paper examined equilibria in multimarket contact with a noisy signal. To the best of our knowledge, under private monitoring, we are the first to find the multimarket contact effect, i.e., the existence of more collusive equilibria than the single market case. We constructed the gEV strategy and clarified the structure of the equilibria, by finding the partial indifferent condition, which leads that strategy to the best possible payoff. In future works, we are particularly interested in an extension to asymmetric markets. We believe our equilibrium construction easily extends to the case of asymmetric markets, only at the expense of additional notations. Under asymmetry, the colluding firms may want to optimally choose the *A* markets and the *B* markets, which makes our problem more complicated.

References

- Abreu, D.; Pearce, D.; and Stacchetti, E. 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* 58:1041–1064.
- Andersen, G., and Conitzer, V. 2013. Fast equilibrium computation for infinitely repeated games. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 53–59.
- Bernheim, B. D., and Whinston, M. D. 1990. Multimarket Contact and Collusive Behavior. *RAND Journal of Economics* 21(1):1–26.
- Blum, A., and Monsour, Y. 2007. Learning, regret minimization, and equilibria. In *Algorithmic game theory*. Cambridge University Press. 79–101.
- Borgs, C.; Chayes, J.; Immorlica, N.; Kalai, A. T.; Mirrokni, V.; and Papadimitriou, C. 2010. The myth of the folk theorem. *Games and Economic Behavior* 70(1):34–43.
- Burkov, A., and Chaib-draa, B. 2013. Repeated games for multiagent systems: A survey. *The Knowledge Engineering Review* 1–30.
- Chellappaw, R. K.; Sambamurthy, V.; and Saraf, N. 2010. Competing in crowded markets: Multimarket contact and the nature of competition in the enterprise systems software industry. *Information Systems Research* 21(3):614–630.
- Conitzer, V., and Sandholm, T. 2007. AWESOME: a general multiagent learning algorithm that converges in self-play and learns a best response against stationary opponents. *Machine Learning* 67(1):23–43.
- Doshi, P., and Gmytrasiewicz, P. J. 2006. On the Difficulty of Achieving Equilibrium in Interactive POMDPs. In *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI)*, 1131–1136.
- Ely, J. C., and Välimäki, J. 2002. A robust folk theorem for the prisoner’s dilemma. *Journal of Economic Theory* 102(1):84–105.
- Ely, J. C.; Horner, J.; and Olszewski, W. 2005. Belief-free equilibria in repeated games. *Econometrica* 73(2):377–415.
- Evans, W. N., and Kessides, I. N. 1994. Living by the “Golden Rule”: Multimarket Contact in the U.S. Airline Industry. *The Quarterly Journal of Economics* 109(2):341–366.
- Hansen, E. A.; Bernstein, D. S.; and Zilberstein, S. 2004. Dynamic programming for partially observable stochastic games. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI)*, 709–715.
- Kandori, M. 2010. Repeated games. In Durlauf, S. N., and Blume, L. E., eds., *Game theory*. Palgrave Macmillan. 286–299.
- Kobayashi, H., and Ohta, K. 2012. Optimal collusion under imperfect monitoring in multimarket contact. *Games and Economic Behavior* 76(2):636–647.
- Kreps, D. M., and Wilson, R. 1982. Sequential equilibria. *Econometrica* 50(4):863–94.
- Littman, M. L., and Stone, P. 2005. A polynomial-time Nash equilibrium algorithm for repeated games. *Decision Support Systems* 39(1):55–66.
- Mailath, G., and Samuelson, L. 2006. *Repeated Games and Reputation*. Oxford University Press.
- Mescheder, D.; Tuyls, K.; and Kaisers, M. 2011. POMDP opponent models for best response behavior. In *Proceedings of the 23rd Benelux Conference on Artificial Intelligence*.
- Shoham, Y., and Leyton-Brown, K. 2008. Learning and teaching. In *Multiagent systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press. 189–222.
- Tennenholtz, M., and Zohar, A. 2009. Learning equilibria in repeated congestion games. In *Proceedings of the 8th International Joint Conference on Autonomous Agents and Multi-Agent System (AAMAS)*, 233–240.
- Wunder, M.; Kaisers, M.; Yaros, J. R.; and Littman, M. 2011. Using iterated reasoning to predict opponent strategies. In *The 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 593–600.