

Self-Evident Events and the Value of Linking*

Jimmy Chan[†] Wenzhang Zhang[‡]

August 17, 2019

Abstract

We propose a theory of linking in long-term relationships based on what information becomes self-evident in equilibrium at the end of a stage game. We obtain a tight bound on the average per-period efficiency loss that must be incurred to enforce a stage-game outcome throughout a T -period repeated game when T is large. Our results apply to all monitoring structures and strategy profiles. They encompass the inefficiency result in Abreu, Milgrom, and Pearce (1991), as well as the approximate-efficiency results in Compte (1998), Obara (2009), and Chan and Zhang (2016).

1 Introduction

In a team moral-hazard problem where it is impossible to determine which player has shirked (Holmstrom, 1982; Radner, Myerson, and Maskin, 1986), each player has an incentive to free-ride on the efforts of the other players. As a result, the Nash equilibrium outcome is typically inefficient. The inefficiency persists even when the players can

*We thank Dilip Abreu, Erik Eyster, Sambuddha Ghosh, Satoru Takahashi, and seminar participants at various seminars and conferences for helpful comments. Part of the research was done when Chan visited the Faculty of Economics of the Universidad Autónoma de Madrid. Chan thanks the Faculty for hospitality.

[†]The Chinese University of Hong Kong. Email: jimmy.hing.chan@gmail.com

[‡]Zhejiang University. Email:wenzhang.zhang@gmail.com

write a binding incentive contract among themselves so long as no budget deficit is allowed. Efficiency can be restored if the players can contract with a third party who can provide external finance to break the no-budget-deficit constraint. Thus, in his seminal paper, Holmstrom (1982) notes that “The fact that capitalistic firms feature separation of ownership and labor implies that the free-rider problem is less pronounced in such firms than in closed organizations like partnerships.”

Recent research in repeated games, however, suggests that access to external finance may not be as important as it seems when players are patient and interact repeatedly. Instead of a series of short-term contracts, the players can use a long-term contract that links incentives across periods. Such long-term contracts are common in reality. For example, in many business organizations, workers are terminated only after performing poorly in multiple periods. In a highly influential paper, Abreu, Milgrom, and Pearce (1991) show that in a repeated game with imperfect monitoring, linking incentives may reduce the cost of imperfect monitoring if the release of monitoring information can be delayed. Subsequent research shows that similar results can be obtained in repeated games of private monitoring under various information structures (Compte, 1998; Obara, 2009; Chan and Zhang, 2016).

The existing results are essentially about two polar cases: public monitoring where there is no gain from linking, and conditionally independent monitoring where, with sufficiently patient players, the efficiency loss of imperfect monitoring disappears with linking as the length of the contract goes to infinity. In many situations, one would expect that players observe both private and public signals. For example, members of a cartel may observe their own sales, which are private information, as well as a public estimate of the total industry sales. A natural question to ask is what happens then? Will it be more like public monitoring? Or will it be more like the conditionally independent monitoring?¹

In this paper we try to provide a general and intuitive explanation of how linking can improve efficiency. We consider a T -period contracting game between a principal and a group of players.² The principal’s objective is to design a contract to enforce

¹We thank a referee for the suggestion for framing our results in this way.

²Our results can be readily applied to repeated games with side-payments. Working with the T -period contracting problem allows us to focus on the mechanism of linking and abstract away from the problem

a particular stage-game outcome throughout the game, subject to the constraint that the total payments to the players be negative. Since Compte (1998), this problem has become a crucial building block in the theory of repeated game with private monitoring. We characterize the value of linking in the enforcement of any outcomes, correlated or non-correlated, in games with any monitoring structures. For outcomes that cannot be implemented efficiently in the long run, we provide a tight bound on the efficiency loss.

A central concept in our analysis is the notion of self-evident event, which is introduced by Aumann (1976) to describe beliefs in incomplete-information games. We use the concept to capture what is special about an event being “public”. In our model, a state of a period describes the private information of the players about what has happened in the period. Unlike the model in Aumann (1976), the distribution of states in our model depends on the actions of the players. Our innovation is to apply the notion of self-evident events to the players’ *equilibrium* beliefs. We show that any efficiency loss that arises from a one-shot contract can be almost entirely eliminated in the long-run if and only if it is not “self evident”. To answer our earlier questions, the results suggest, in a sense to be made precise, that public monitoring is a special case, while independent and correlated private monitoring are similar when the players are sufficiently patient.

Our results provide a unified framework to understand the value of linking. Following Abreu, Milgrom, and Pearce (1991), Compte (1998), Obara (2009), and Chan and Zhang (2016) apply the linking idea to obtain folk theorems. Taking a different approach, Fudenberg, Levine, and Maskin (1994), Kandori and Matsushima (1998), and Rahman and Obara (2010) identify conditions under which an outcome can be enforced by a budget-balanced contract. Our results naturally combine these two approaches and connect them to the inefficiency result in repeated games with public monitoring. We discuss the repeated-game literature in Section 6.1. Instead of relying on an external budget breaker, members of a partnership may hire a disinterested mediator to implement correlated strategies that virtually enforce an outcome with a budget-balanced contract (Rahman and Obara, 2010). In our setting, players may use correlated strategies to make incentives non-self-evident. We show that any strictly enforceable outcome, regardless of the monitoring structure, has a correlated outcome close to it that can be

of implementing transfers through continuation strategies.

enforced almost efficiently in the long run.

Beyond repeated games, the idea of linking also plays an important role in the literature of relational contracting and organizational economics (Fuchs, 2007; Zhao, 2008; Ke, Li, and Powell, 2018). While these models focus on a long-term principal-agent relationship given a fixed discount factor, we deal with a partnership problem when the discount factor goes to one.

The rest of the paper is organized as follows. The next section illustrates the main ideas behind our results in a repeated Prisoners' Dilemma game. Section 3 introduces the formal model. Section 4 introduces the notion of self-evident events and establishes a key lemma that is crucial for our results. Section 5 gives our main results. Section 6 characterizes the value of linking in terms of the primitive of the model and discusses the literature of repeated games with private monitoring. Section 7 shows that any strictly enforceable outcome is virtually enforceable with almost no long-run efficiency loss. Section 8 concludes.

2 Example

In this section we motivate our results by considering a T -period two-person noisy Prisoners' Dilemma game. In each period $t = 1, \dots, T$, each player $i = 1, 2$ independently chooses C or D . The expected stage-game payoff is given in Table 1. If both players choose C , then each player obtains a payoff of 1. If one player chooses C and the other chooses D , then the player who plays D receives $(1 + d)$, while the player who plays C receives $-h$, where $d, h > 0$. The unique stage-game Nash equilibrium is (D, D) .

	C	D
C	1, 1	$-h, 1 + d$
D	$1 + d, -h$	0, 0

Table 1: Payoff matrix.

At the end of each period, each player i observes a private signal $y_i \in \{H, L\}$. Table 2 describes the signal distributions conditional on the action profiles (C, C) , (C, D) , and

(D, C) .

	H	L
H	$p - \rho p(1 - p)$	$\rho p(1 - p)$
L	$\rho p(1 - p)$	$(1 - p) - \rho p(1 - p)$

Signal distribution under (C, C)

	H	L
H	$q - \rho q(1 - q)$	$\rho q(1 - q)$
L	$\rho q(1 - q)$	$(1 - q) - \rho q(1 - q)$

Signal distribution under (C, D) or (D, C)

Table 2: Signal distributions.

If both players play C , then $y_i = H$ with probability p . If one player chooses C and the other chooses D , then $y_i = H$ with probability $q < p$. The correlation between the players' signals depends on the parameter ρ . When $\rho = 0$, the signals are perfectly correlated. In this case the players are effectively observing a public signal.³ When $\rho = 1$, the signals are conditionally independent and a player cannot learn about the other player's signal from his own. When $\rho \in (0, 1)$, the signals are positively and imperfectly correlated conditional on the action profile.

A principal, without access to external funding, tries to design a contract to enforce (C, C) in every period. At the end of period T , the principal asks the players to report their signals. Write y^T for (y_1^T, y_2^T) , where $y_i^T = (y_i(1), \dots, y_i(T))$ is a T -period profile of player i 's signal. A T -period contract $w^T = (w_1^T, w_2^T)$ is a function that maps each y^T to a payment to each player, subject to the constraint that the total payment be non-positive. To simplify exposition, we assume in this section that the players' discount factor is one so that the utility of a player is equal to the total stage-game payoffs plus the contract payment.

³In our model, there is no difference between a public signal and a vector of perfectly correlated private signals.

The assumption that the total payment must be negative implies that incentives are costly. Consider the one-period case. Let $w = (w_1, w_2)$ denote a one-period contract. With a slight abuse of notation, let $w_i(H)$ and $w_i(L)$ denote player i 's payment when player j 's signal is H and L , respectively. It is straightforward to see that it is optimal for player i to choose C if and only if

$$(p - q)(w_i(H) - w_i(L)) \geq d.^4$$

Since a player's payment depends only on the report of the other player, the players have no incentive to lie about their reports. Given the constraint $w_i(H), w_i(L) \leq 0$, the most efficient way to enforce (C, C) is to set

$$\begin{aligned} w_i(H) &= 0 \\ w_i(L) &= -\frac{d}{p - q}. \end{aligned}$$

The per-player efficiency loss is thus $(1 - p)d/(p - q)$; see Figure 1. The inefficiency




Figure 1: The one-period contract and efficiency loss.

arises because a player has to destroy an amount equal to $d/(p - q)$ when the other player reports an L signal. The value cannot be transferred to the other player because doing so will interfere with the incentives of the other player. In our example, if player 1 has to pay player 2 $d/(p - q)$ when player 2's report is L , player 2 will always report L .

When the contract lasts for multiple periods, the principal can still use the one-period contract $(w_i(H), w_i(L)) = (0, -d/(p - q))$ to enforce (C, C) period by period.

⁴Because the stage game is symmetric, we can consider only symmetric contracts where $w_1 = w_2 < 0$. The single-period optimal contract does not depend on ρ .

The question is whether the principal can do better by using a non-linear contract. The existing literature has largely focused on two polar cases. When $\rho = 0$, linking has no value and the linear contract is the best that one could do, but when $\rho = 1$, linking is useful and the per-period efficiency loss disappears when T goes to infinity. Our results are about what happens in between when $\rho \in (0, 1)$. In the following, we briefly recount the two polar cases before explaining our results.

2.1 Case 1: $\rho = 0$.

We can use the two-period case to illustrate why linking has no value when the signals are perfectly correlated. To induce (C, C) in both periods, three incentive-compatibility constraints must be satisfied; namely, the first period, the second period after the players observe H , and the second period after the players observe L . See Figure 2.




Figure 2: The two-period case.

Let w^2 denote a contract that enforces (C, C) for two periods, and let $E [w_i^2 | y(1), CC]$ denote the expected payment to player i conditional on the first-period signal $y(1)$ and the second-period action profile (C, C) .⁵ Since enforcing (C, C) in the second period after $y(1) = H$ is the same as enforcing (C, C) in a single-period game,

$$E [w_i^2 | H, CC] \leq -(1-p) \frac{d}{p-q}. \quad (1)$$

The incentive-compatibility constraint in period 1 requires that

$$(p-q) (E [w_i^2 | H, CC] - E [w_i^2 | L, CC]) \geq d. \quad (2)$$

⁵Since the signals are perfectly correlated, we will mention the common signal.

Combining (1) and (2), we have

$$E [w_i^2 | L, CC] \leq -(2-p) \frac{d}{p-q}.$$

It follows that

$$pE [w_i^2 | H, CC] + (1-p)E [w_i^2 | L, CC] \leq -2(1-p) \frac{d}{p-q}.$$

Intuitively, since the continuation game after H is completely separated from the continuation game after L , the incentives in these subgames cannot be linked. Hence, enforcing (C, C) in the continuation game following H must incur a one-period efficiency loss. This, together with the fact that the value of the continuation game after L must be lower than that after H , implies that the two-period efficiency loss must be no less than twice the one-period efficiency loss.

2.2 Case 2: $\rho = 1$.

Let $-i$ denote the player who is not player i . Let $f_{-i}(L|y^T)$ denote the number of L signals that player $-i$ observes in the signal profile y^T . Consider a “linear” contract \tilde{w}^T , where for $i = 1, 2$

$$\tilde{w}_i^T(y^T) = - (f_{-i}(L|y^T) - T(1-p-v)) (d/(p-q) + \varepsilon),$$

where ε and v are small positive constants. The contract punishes player i by $(d/(p-q) + \varepsilon)$ for every L signal player $-i$ observes. It strictly enforces (C, C) but violates the constraint that the total payment be non-positive. When both players observe fewer than $T(1-p-v)$ L signals, each player will receive a strictly positive payment.

To satisfy the non-positive-payment constraint, we truncate \tilde{w}_i^T at zero to obtain a “truncated” contract \hat{w}^T , where

$$\hat{w}_i^T(y^T) = - \max(f_{-i}(L|y^T) - T(1-p-v), 0) (d/(p-q) + \varepsilon). \quad (3)$$

If both players choose C , the average number of L signals that player $-i$ should observe is $T(1-p)$. Thus, the truncated contract punishes player i for L signals in excess of a

threshold set below the mean by Tv . The truncated incentives, the difference between \widehat{w}_i^T and \widehat{w}_i^T , are equal to

$$R_i(y^T) = -\min(f_{-i}(L|y^T) - T(1-p-v), 0)(d/(p-q) + \varepsilon).$$

The distortionary effect of the truncation on player i 's incentives, however, is very small when T is large. Assuming that player i is choosing C in every period, by the law of large number, the probability that the fraction of L signals is lower than $1-p$, the ex ante mean, by v is exceedingly small when T is large. Since player i does not learn about y_{-i}^T from his own signals, the distortionary effect of the truncation remains very small throughout the contract and is compensated by the small extra punishment ε .

The expected per-player per-period efficiency loss caused by \widehat{w}_i^T is approximately

$$v(d/(p-q) + \varepsilon).$$

As T becomes large, v and ε can be chosen so that the per-player per-period efficiency loss goes to zero.

It is important to note that the truncated contract does not enforce (C, C) when $\rho = 0$. Although ex ante the probability that $f_{-i}(L|y^T)$ is less than $(1-p-v)T$ is very small, player i will come to believe that this probability is large after observing very few L signals.

2.3 Case 3: $\rho \in (0, 1)$.

In this case, as the signals are correlated, the continuation games after different signal realizations are not separated as in Case 1. Nevertheless, player i , after observing very few L signals, will come to believe that it is likely that player $-i$ has observed very few L signals as well. One may, therefore, expect that linking incentives would become less effective as ρ decreases. It turns out that $\rho = 0$ is a special case. So long as $\rho > 0$, the per-player per-period efficiency loss can be reduced to almost zero as T goes to infinity.

The idea is to distribute the truncation effect among the players in a way that distorts each player's incentive minimally. Conditional on both players choosing C in every period, each player i , on average, should observe the L signal in $T(1-p)$ periods.

Denote the “excess” L signals that player i observes in y^T by

$$\tau_i(L|y^T) = f_i(L|y^T) - T(1-p).$$

Using this notation, we can restate the payment to player i under the truncated contract \widehat{w}^T as

$$\widehat{w}_i^T(y^T) = -\max(\tau_{-i}(L|y^T) + Tv, 0)(d/(p-q) + \varepsilon).$$

The incentive is truncated when $\tau_{-i}(L|y^T) < -Tv$.

We decompose $\tau_i(L|y^T)$ into three components. For any y^T , let $f_i(y_i, y_{-i}|y^T)$ denote the number of periods in y^T in which player i observes y_i and player $-i$ observes y_{-i} . Conditional on (C, C) , player $-i$ expects player i to observe L with probability $\rho(1-p)$ when he observes H . Hence, he expects player i to observe $\rho(1-p)f_{-i}(H|y^T)$ L signals in the $f_{-i}(H|y^T)$ periods in which he observes H . The number of “excess” L signals that player i observes in the periods player $-i$ observes H is, therefore,

$$\tau_i(L|y^T, H) \equiv f_i(L, H|y^T) - \rho(1-p)f_{-i}(H|y^T).$$

Similarly, denote the number of “excess” L signals player i observes among the periods player $-i$ observes L by

$$\tau_i(L|y^T, L) \equiv f_i(L, L|y^T) - (1-\rho p)f_{-i}(L|y^T).$$

Using the fact that

$$f_{-i}(L|y^T) + f_{-i}(H|y^T) = T,$$

it is straightforward to verify that

$$\tau_1(L|y^T) = (1-\rho)\tau_2(L|y^T) + \tau_1(L|y^T, H) + \tau_1(L|y^T, L); \quad (4)$$

$$\tau_2(L|y^T) = (1-\rho)\tau_1(L|y^T) + \tau_2(L|y^T, H) + \tau_2(L|y^T, L). \quad (5)$$

Equations (4) and (5) decompose $\tau_i(y^T)$ into (1) the excess L signals observed by player $-i$, (2) the excess L signals player i observes in the periods player $-i$ observes H , and (3) the excess L signals player i observes in the periods player $-i$ observes L .

Player $-i$ can observe only the first component. By the law of iterated expectation, player $-i$ always expects $\tau_i(L|y^T, H)$ and $\tau_i(L|y^T, L)$ to be zero. That is,

$$E_{y_i^T}[\tau_i(L|y^T, H)|y_{-i}^T] = E_{y_i^T}[\tau_i(L|y^T, L)|y_{-i}^T] = 0 \quad \text{for all } y_{-i}^T.$$

Hence,

$$\begin{aligned} E_{y_i^T} [\tau_i(L|y^T) | y_{-i}^T] &= (1 - \rho) E_{y_i^T} [\tau_{-i}(L|y^T) | y_{-i}^T] \\ &= (1 - \rho) \tau_{-i}(L|y^T). \end{aligned} \quad (6)$$

That is, when $\rho > 0$, player $-i$'s expectation of $\tau_i(L|y^T)$ is less than $\tau_{-i}(L|y^T)$.

Combining (4) and (5), we have when $\rho > 0$,

$$\tau_i(L|y^T) = \frac{\tau_i(L|y^T, H) + \tau_i(L|y^T, L) + (1 - \rho) (\tau_{-i}(L|y^T, H) + \tau_{-i}(L|y^T, L))}{1 - (1 - \rho)^2}. \quad (7)$$

Equation (7) means that $\tau_i(L|y^T)$ can only be negative if one of $\tau_i(L|y^T, H)$, $\tau_i(L|y^T, L)$, $\tau_{-i}(L|y^T, H)$, and $\tau_{-i}(L|y^T, L)$ is negative. (Note that the argument critically relies on $\rho > 0$.) Recall that for every y_{-i}^T , player $-i$ expects $\tau_i(L|y^T, y_{-i})$ to be zero. Hence, any y^T with $\tau_i(L|y^T) < 0$ is “unexpected” to either player i or player $-i$.

Fix any $\nu > 0$. It is possible to pick $\zeta > 0$ such that for any y^T , $\tau_i(L|y^T) < -T\nu$ implies

$$\min(\tau_i(L|y^T, H), \tau_i(L|y^T, L), \tau_{-i}(L|y^T, H), \tau_{-i}(L|y^T, L)) < -T\zeta.$$

Start with the truncated contract \widehat{w}^T in Case 2. Add a side-bet contract $z^T = (z_1^T, z_2^T)$.

For $i = 1, 2$,

$$z_i^T(y^T) = R_i(y^T) (1 - I_i(y^T)) - R_{-i}(y^T) I_i(y^T),$$

where

$$I_i(y^T) = \begin{cases} 0 & \text{if } \min(\tau_{-i}(L|y^T, H), \tau_{-i}(L|y^T, L)) \geq -T\zeta, \\ 1 & \text{otherwise.} \end{cases}$$

Under this side-bet contract, player i receives the extra incentives R_i when

$$\min(\tau_{-i}(L|y^T, H), \tau_{-i}(L|y^T, L)) \geq -T\zeta;$$

otherwise, he pays R_{-i} , the extra incentives for player $-i$. Under the truncated contract \widehat{w}^T , the distortionary effect of the truncation of player i 's payment is entirely borne by player i . The side-bet contract reallocates the distortionary effect to a player who does not expect the distortion to occur.⁶ The total payment of this side-bet contract is always

⁶In simple words, the side-bet contract says that a player must bear the truncation effect if the distribution of the other player's signal deviates significantly from what the player expects given what he has observed.

negative. When $R_i > 0$, either I_i or I_{-i} must be equal to 1. Hence, when player i receives a strictly positive amount through the side-bet contract, player $-i$ must pay for it.

When T is large, player i believes that the probability that $I_i = 0$ is very close to one throughout the contract. As a result, player i believes, regardless of what signals he may observe during the contract, that he will almost always receive the extra incentives R_i , and almost never need to pay R_{-i} . This means that when T is large the difference between the truncated-cum-side-bet contract and the linear contract, \tilde{w}^T , is small. Since \tilde{w}^T strictly enforces (C, C) , so does the truncated-cum-side-bet contract.⁷

The efficiency loss of the side bets is very small when T is large because R_1 and R_2 are almost always equal to zero. We have already shown that the efficiency loss of the truncated contract is small. Hence, (C, C) can be enforced almost efficiently when T is sufficiently large.

Recall in the single-period case, a third-party is needed to contract efficiently. What we show is that, as T becomes large, it is possible for one player to partially compensate another player without distorting the incentive of the first player. The key observation behind this result is that when $\rho > 0$, any y^T that involves any player observing very few L signals must be “unexpected” to some player. In Section 4.1, we generalize this observation to all stage games.

We are not the first to exploit the differential beliefs between players. Fong, Gossner, Hörner, and Sannikov (2011), in a repeated Prisoners’ Dilemma similar to our example, make use of the fact each player expects the other player to observe fewer excess L signals than he does to support an approximately efficient equilibrium. As Eq. (6) makes clear, our approach can be viewed as a generalization of theirs.

⁷Although, with the side-bet contract, each player may want to mis-report his own signals, the incentive is small and the side-bet contract can be slightly modified to maintain truth-telling.

3 Model

3.1 Stage game

Consider a finite stage game endowed with a correlating device. Let $N = \{1, 2, \dots, n\}$ denote a set of players, $A = A_1 \times \dots \times A_n$ a finite set of action profiles, $\eta \in \Delta(A)$ a distribution over A , and $g = (g_1, \dots, g_n) : A \rightarrow \mathbf{R}^n$ a profile of stage-game payoff functions. In each period, the correlating device draws $\tilde{a} = (\tilde{a}_1, \dots, \tilde{a}_n) \in A$ according to η and privately recommends \tilde{a}_i to each player i . After learning \tilde{a}_i , each player $i \in N$ privately chooses $a_i \in A_i$. Player i 's expected stage-game payoff is $g_i(a)$, where $a = (a_1, \dots, a_n)$. The players do not directly observe the stage-game payoffs. Instead, each player i observes a signal y_i . The signal profile $y = (y_1, \dots, y_n)$ is drawn from a finite set $Y = Y_1 \times \dots \times Y_n$ according to a distribution $p(\cdot|a) \in \Delta(Y)$.

Since the only function of the correlating device is to allow the players to play η , modeling the correlating device as private recommendations is without loss of generality.⁸ When η is a pure or uncorrelated mixed outcome, the correlation device can be dispensed with. To avoid extra notations we shall assume that all signals are associated with distinct posterior beliefs. All results go through without this assumption, although some may have to be rephrased to allow for the possibility of redundant signals.

Assumption 1. For each $i \in N$, $a \in A$, and $y_i, y'_i \in Y_i$, $p(y_{-i}|a, y_i) \neq p(y_{-i}|a, y'_i)$ for some $y_{-i} \in Y_{-i}$.

We impose no further restriction on the correlation structure beyond Assumption 1. In general, the players' signals may be correlated and $p(\cdot|a)$ may not have full support. Hence, our model includes public monitoring as a special case.⁹

⁸As is well known, the correlating device can be replaced by pregame communication when there are more than five players (Gerardi, 2004).

⁹The game becomes one of public monitoring when $Y_1 = \dots = Y_n$ and for all $a \in A$, $p(y|a) > 0$ only if $y_1 = \dots = y_n$.

3.2 T -period contracting problem

In period 0, a principal proposes a contract. After observing the contract, the players play the stage game for T periods. At the end of period T , the players report the private signals observed, and the correlating device reports the recommendations made during the T periods. In addition to the stage-game payoffs, at the end of the T -period game, each player receives a payment as stipulated by the contract. While the correlating device always reports honestly, players may lie.

For each variable x , we use $x(t)$ to denote the period- t value of x and $x^t = (x(1), \dots, x(t))$ to denote the history of x up to period t . Hence, $\tilde{a}^T = (\tilde{a}(1), \dots, \tilde{a}(T))$ is the history of recommendations and the report of the correlating device. Let $\hat{y}_i^T = (\hat{y}_i(1), \dots, \hat{y}_i(T))$ denote the T -period signal-report of player i and $\hat{y}^T = (\hat{y}_1^T, \dots, \hat{y}_n^T)$ denote the signal-report profile. A T -period contract consists of n functions $w^T = (w_1^T, \dots, w_n^T)$, where each w_i^T maps each $(\tilde{a}^T, \hat{y}^T) \in A^T \times Y^T$ into a payment.¹⁰ The total payment must be weakly negative; i.e.,

$$\sum_{i=1}^n w_i^T(\tilde{a}^T, \hat{y}^T) \leq 0, \forall (\tilde{a}^T, \hat{y}^T) \in A^T \times Y^T.$$

Player i 's total discounted payoff is

$$\frac{1-\delta}{1-\delta^T} \left(\sum_{t=1}^T \delta^{t-1} g_i(a(t)) + w_i^T(\tilde{a}^T, \hat{y}^T) \right),$$

where $\delta \in (0, 1)$ is a common discount factor for the players.¹¹

Since N , A , and g are fixed in our analysis, we denote the T -period game by $\Gamma(\eta, T, \delta, w^T)$. A pure strategy of player i consists of two components: an action strategy α_i^T that maps each $(\tilde{a}_i^t, a_i^{t-1}, y_i^{t-1}) \in \bigcup_{s=1}^T (A_i^s \times A_i^{s-1} \times Y_i^{s-1})$ into an action $a_i \in A_i$ and a reporting strategy ρ_i^T that maps each $(\tilde{a}_i^T, a_i^T, y_i^T) \in A_i^T \times A_i^T \times Y_i^T$ into a

¹⁰The contracts in our model can be viewed as the T -period version of the ones in Rahman and Obara (2010). The only difference is that in Rahman and Obara (2010) a contract must be budget balanced, whereas in ours the total payment can be strictly negative.

¹¹The restriction to negative total transfer arises naturally in different contexts. For example, if bonus contracts are not legally enforceable, then the principal may have to commit to “burn” the difference between a lump sum and the actual bonus (MacLeod, 2003; Fuchs, 2007). In repeated games, players can enforce cooperation only by switching to inefficient continuation paths.

report $\hat{y}_i^T \in Y_i^T$.¹² A mixed strategy σ_i^T is a probability distribution over the set of pure strategies (α_i^T, ρ_i^T) . Let Σ_i^T denote the set of mixed strategies for player i .

Player i 's expected payoff conditional on $\sigma^T = (\sigma_1^T, \dots, \sigma_n^T)$ is

$$v_i^T(\sigma^T; w_i^T) \equiv \frac{1-\delta}{1-\delta^T} E \left[\sum_{t=1}^T \delta^{t-1} g_i(a(t)) + w_i^T(\tilde{a}^T, \hat{y}^T) \middle| \sigma^T \right],$$

where the expectation is taken over $(\tilde{a}^T, a^T, y^T, \hat{y}^T)$ with respect to the distribution induced by σ^T , η , and p .

The contracting problem is to choose w^T to enforce the correlated outcome η throughout the contract. By the revelation principle, we can focus on mechanisms where players play the obedient strategies that follow recommendations in every period and report signals truthfully. Let $\sigma_i^{T*} = (\alpha_i^{T*}, \rho_i^{T*})$ denote the obedient strategy of player i and $\sigma^{T*} = (\sigma_1^{T*}, \dots, \sigma_n^{T*})$.

Definition 1. A contract w^T enforces η for T periods if σ^{T*} is a Nash equilibrium in $\Gamma(\eta, T, \delta, w^T)$. That is, if for all $i \in N$ and $\sigma_i^T \in \Sigma_i^T$,

$$v_i^T(\sigma^{T*}; w_i^T) \geq v_i^T(\sigma_i^T, \sigma_{-i}^{T*}; w_i^T).$$

The enforcement is strict if the inequality is strict for σ_i^T that deviates from the recommendations with positive probability. An outcome η is (strictly) enforceable if it can be (strictly) enforced by some w^T .

Obviously, if η cannot be enforced when $T = 1$, then it cannot be enforced when $T > 1$. Conversely, if η can be enforced when $T = 1$ by w , then it can be enforced for any T by applying w period by period. Thus, it is sufficient to consider the case $T = 1$ to determine the enforceability of η .¹³

In the following we write σ for σ^1 and w for w^1 for convenience. Let μ denote the distribution over (\tilde{a}, y) induced by η and p . For all $(\tilde{a}, y) \in A \times Y$,

$$\mu(\tilde{a}, y) = p(y|\tilde{a}) \eta(\tilde{a}).$$

¹²As usual, a^0 denotes the null history \emptyset and A^0 denotes the set whose only element is a^0 . Similar notations apply for signal.

¹³Same for strict enforceability.

With a slight abuse of notation, we also use μ to denote the distribution of (\tilde{a}, \hat{y}) induced by the obedient strategy profile σ^* . Let π^{σ_i} denote the distribution of (\tilde{a}, \hat{y}) when player i deviates to σ_i , while other players choose σ_{-i}^* . For any $(\tilde{a}, \hat{y}) \in A \times Y$,

$$\pi^{\sigma_i}(\tilde{a}, \hat{y}) = \sum_{(\alpha_i, \rho_i)} \sigma_i(\alpha_i, \rho_i) \sum_{y_i: \rho_i(\tilde{a}_i, \alpha_i(\tilde{a}_i), y_i) = \hat{y}_i} p(y_i, \hat{y}_{-i} | \tilde{a}_{-i}, \alpha_i(\tilde{a}_i)) \eta(\tilde{a}).$$

Definition 2. A deviating strategy σ_i is *undetectable* if $\pi^{\sigma_i} = \mu$.

The following result due to Rahman (2012) provides a necessary and sufficient condition for enforceability.

Lemma 1 (Theorem 1, Rahman, 2012). *An action profile η is enforceable if and only if for all $i \in N$ and all undetectable σ_i ,*

$$\sum_{(\alpha_i, \rho_i)} \sigma_i(\alpha_i, \rho_i) \sum_{\tilde{a} \in A} g_i(\alpha_i(\tilde{a}_i), \tilde{a}_{-i}) \eta(\tilde{a}) \leq \sum_{\tilde{a} \in A} g_i(\tilde{a}) \eta(\tilde{a}).$$

Because the total payment must be negative, enforcing a non-stage-game Nash equilibrium may come with a cost. The per-period efficiency loss of enforcing η with w^T in $\Gamma(\eta, T, \delta, w^T)$ is

$$W(\eta, T, \delta, w^T) \equiv - \sum_{i=1}^n \frac{1-\delta}{1-\delta^T} E[w_i(\tilde{a}^T, \hat{y}^T) | \sigma^{T*}].$$

Let $\mathscr{W}(\eta, T, \delta)$ be the set of w^T that enforces η . The minimum per-period efficiency loss to enforce η is

$$W^*(\eta, T, \delta) = \min_{w^T \in \mathscr{W}(\eta, T, \delta)} W(\eta, T, \delta, w^T).$$

Our objective is to characterize $W^*(\eta, T, \delta)$ as T goes to infinity and δ goes to 1.

In the following, we often deal with the case $T = 1$. As δ does not matter, for convenience we write $\mathscr{W}(\eta)$ for $\mathscr{W}(\eta, 1, \delta)$, $W(\eta, w)$ for $W(\eta, 1, \delta, w)$ and $W^*(\eta)$ for $W^*(\eta, 1, \delta)$.

Before we proceed, a comment about the solution concept is in order. As is well known, Nash equilibrium imposes no restriction on players' responses off the equilibrium path. In our model, it is consistent with Nash equilibrium for players who observe

signals inconsistent with the equilibrium actions to report honestly. Theorem 1, which establishes a lower bound on efficiency loss, continues to hold if the stronger notion of sequential equilibrium is used instead. Following Kandori and Matsushima (1998), Theorem 2, which establishes the tightness of the bound, can be made consistent with sequential equilibrium by assuming that the support of the signal distribution is invariant with a . Extending the result without invariant support would require specifying and keeping track of the players' diverging beliefs (as well as their beliefs about other players' continuation strategies) after one or multiple players observe inconsistent signals. We do not pursue this issue in this paper.

4 Self-Evident Events

In the example in Section 2, the key property of a public signal is that no one could miss it. If a player observes that a public signal is L , he knows that every player observes L , and that every player knows that every player observes L , and so on. It is only then that the continuation game after L is entirely separated from the continuation game after H .

In a stage game with private signals, although the players do not directly observe the recommendations to and signals of the other players, they form beliefs about them on the basis of their own. Conditional on the correlated strategy profile η , the recommendation and signal pair (\tilde{a}, y) is distributed according to μ . Write $\text{supp}(\mu)$ for the support of μ . Let P_i denote player i 's information partition of $\text{supp}(\mu) \subset A \times Y$. Denote the element of P_i that contains (\tilde{a}, y) by $P_i(\tilde{a}, y)$. For any $i \in N$ and any $(\tilde{a}, y), (\tilde{a}', y') \in \text{supp}(\mu)$,

$$(\tilde{a}', y') \in P_i(\tilde{a}, y) \text{ if and only if } (\tilde{a}', y') = (\tilde{a}, y).$$

The vector (P_1, \dots, P_n) describes the players' knowledge structure when η is chosen. Conditional on observing (\tilde{a}_i, y_i) , player i believes that the realized (\tilde{a}, y) must belong to $P_i(\tilde{a}, y)$. In the terminology of interactive epistemology, a subset E of $\text{supp}(\mu)$ is an event. Player i "knows" that E occurs at (\tilde{a}, y) if

$$P_i(\tilde{a}, y) \subseteq E. \tag{8}$$

That player i knows E is itself an event that consists of all (\tilde{a}, y) where (8) is true. Thus, we can talk about player j knows that player i knows E . An event E is common belief

among the players at (\tilde{a}, y) if, when (\tilde{a}, y) occurs, every player knows E , knows that everyone knows E , and so on. An event E is self-evident if it is common belief at every $(\tilde{a}, y) \in E$.

A self-evident event is called irreducible if it contains no proper subset that is self-evident. Let P denote the meet of (P_1, \dots, P_n) (i.e., the least common coarsening). It is well known that any element of P is self-evident and irreducible (Chapter 5 of Osborne and Rubinstein, 1994).

In the noisy Prisoners' Dilemma example in Section 2, there are two irreducible self-evident events: $\{(HH)\}$ and $\{(LL)\}$ when the signals are perfectly correlated. Otherwise there is only one irreducible self-evident event: $\{(HH), (HL), (LH), (LL)\}$. In general, every realization of a public signal is self-evident conditional on *any* stage-game strategy.¹⁴ However, a self-evident event may not be related to any public signal, and an event may be self-evident conditional on one stage-game strategy but not self-evident conditional on another. In the following, when we say that an event is self-evident, it is always with respect to the equilibrium correlated strategy profile η .

4.1 A key lemma

In Section 2, we show that when the players' signals are not perfectly correlated, any y^T that deviates from the ex ante distribution must "surprise" some player (given the player's information). In this section, we generalize the result.

Write $P(\tilde{a}, \hat{y})$ for the element of P to which (\tilde{a}, \hat{y}) belongs. For any (\tilde{a}^T, y^T) , let $f(\tilde{a}, y | \tilde{a}^T, y^T)$, $f(\tilde{a}_i, y_i | \tilde{a}^T, y^T)$, and $f(P(\tilde{a}, y) | \tilde{a}^T, y^T)$ denote, respectively, the numbers of times (\tilde{a}, y) , (\tilde{a}_i, y_i) , and $P(\tilde{a}, y)$ occur in (\tilde{a}^T, y^T) .

Consider an observer who in each period observes only the element of P . Suppose (\tilde{a}^T, y^T) occurs. The observer knows that a particular outcome (\tilde{a}, y) does not occur in period t when $(\tilde{a}, y) \notin P(\tilde{a}(t), y(t))$. In each period t where $(\tilde{a}, y) \in P(\tilde{a}(t), y(t))$, the observer believes that there is a probability $\mu(\tilde{a}, y | P(\tilde{a}, y))$ that the outcome is (\tilde{a}, y) . Hence, the total number of times the observer expects (\tilde{a}, y) to occur in (\tilde{a}^T, y^T) is $\mu(\tilde{a}, y | P(\tilde{a}, y)) f(P(\tilde{a}, y) | \tilde{a}^T, y^T)$.

¹⁴To be precise, the set of all recommendation and signal profiles that are consistent with the public signal is a self-evident event.

For any $(\tilde{a}, y) \in \text{supp}(\mu)$ and any $\iota > 0$, define

$$Z^T((\tilde{a}, y), \iota) = \{(\tilde{a}^T, y^T) \mid |f(\tilde{a}, y | \tilde{a}^T, y^T) - \mu(\tilde{a}, y | P(\tilde{a}, y))f(P(\tilde{a}, y) | \tilde{a}^T, y^T)| > \iota\}$$

as the set of (\tilde{a}^T, y^T) where the number of times (\tilde{a}, y) occurs in (\tilde{a}^T, y^T) deviates by ι from the expectation of the aforementioned observer.

For each player i , define

$$Z_i^T((\tilde{a}, y), \iota) = \{(\tilde{a}^T, y^T) \mid |f(\tilde{a}, y | \tilde{a}^T, y^T) - \mu_{-i}(\tilde{a}_{-i}, y_{-i} | \tilde{a}_i, y_i)f(\tilde{a}_i, y_i | \tilde{a}^T, y^T)| > \iota\}$$

as the set of (\tilde{a}^T, y^T) where the number of times (\tilde{a}, y) occurs in (\tilde{a}^T, y^T) deviates from the mean conditional on the private information that player i receives in (\tilde{a}^T, y^T) by ι .

Write $Z^{T*}(\iota)$ for $\cup_{(\tilde{a}, y) \in A \times Y} Z^T((\tilde{a}, y), \iota)$ and $Z_i^{T*}(\iota)$ for $\cup_{(\tilde{a}, y) \in A \times Y} Z_i^T((\tilde{a}, y), \iota)$. Every (\tilde{a}^T, y^T) induces a distribution over stage-game outcomes $A \times Y$. $Z^{T*}(\iota)$ contains any (\tilde{a}^T, y^T) that induces a distribution that deviates from the expected distribution conditional on self-evident events, while $Z_i^{T*}(\iota)$ contains any (\tilde{a}^T, y^T) that induces a distribution that deviates from the expected distribution conditional on player i 's private information.

Lemma 2. *For any $\iota > 0$, there exists $c_0 > 0$ such that, for any T and any $(\tilde{a}^T, y^T) \in \text{supp}(\mu)^T$, if $(\tilde{a}^T, y^T) \in Z^{T*}(\iota)$, then $(\tilde{a}^T, y^T) \in Z_i^{T*}(c_0\iota)$ for some player i .*

Lemma 2 says that if the distribution induced by (\tilde{a}^T, y^T) deviates from the expected distribution conditional on self-evident events by ι , then it must also deviate by at least $c_0\iota$ from the expected distribution conditional on the private information of some player i .

To illustrate the lemma, recall in the Prisoners' Dilemma example in Section 2, there are four possible outcomes: (HH) , (HL) , (LH) , and (LL) . When the signals are correlated, all four outcomes belong to the same irreducible self-evident event. If in some y^T , the number of times (HH) occurs is greater than the unconditional mean but is equal to the mean conditional on the periods in which player 1 observing H , as well as the mean conditional on the periods in which player 2 observing H , then the number of times (H, L) and (L, H) occur in y^T must also be greater than their respective unconditional means. This means that the number of times (L, L) occurs in y^T must

be less than the unconditional mean, as well as the mean conditional on the periods in which player 1 observing L .

Note that ι and c_0 apply uniformly to all T and (\tilde{a}^T, y^T) in Lemma 2. Given ι and c_0 , by the law of large number both $Z^{T*}(\iota)$ and $Z_i^{T*}(c_0 \iota)$ become extremely unlikely as T becomes large.

4.2 Decomposing incentives

A single-period contract is a mapping from an outcome profile to a vector of payments. We say that the incentives provided by a single-period contract vary across two self-evident events if, given the equilibrium outcome η , the total expected payment conditional on one self-evident event is different from another. As we see in Case 1 of Section 2, efficiency loss due to incentives that vary across the self-evident events cannot be reduced by linking. In a general contract, incentives may vary both across and within self-evident events. To characterize the extent to which this efficiency loss can be reduced by linking, we need to decompose a single-period contract into a component that varies across self-evident events and a residual that is constant across self-evident events.

Write ω for a typical element of P . For any stage-game contract $w \in \mathcal{W}(\eta)$, let $E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega]$ denote player i 's expected transfer conditional on σ^* and ω , and let

$$\omega_{\max} \in \arg \max_{\omega \in P} \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega]$$

denote the element of P with the maximum expected total payment. For any stage-game contract w and any player i , we can write

$$w_i(\tilde{a}, \hat{y}) = w_{i,a}(\tilde{a}, \hat{y}) + w_{i,b}(\tilde{a}, \hat{y}), \quad (9)$$

where

$$\begin{aligned} w_{i,a}(\tilde{a}, \hat{y}) &\equiv w_i(\tilde{a}, \hat{y}) - w_{i,b}(\tilde{a}, \hat{y}) \\ w_{i,b}(\tilde{a}, \hat{y}) &\equiv E[w_i(\tilde{a}', \hat{y}') | \sigma^*, P(\tilde{a}, \hat{y})] - E[w_i(\tilde{a}', \hat{y}') | \sigma^*, \omega_{\max}]. \end{aligned}$$

We call $w_{i,b}$ the self-evident component of w_i , as it depends only on the element of P to which (\tilde{a}, \hat{y}) belongs. This component measures incentives that vary across

irreducible self-evident events. The expected value of the residual component $w_{i,a}$ is constant conditional on any $\omega \in P$. For all $i \in N$ and $\omega \in P$,

$$E[w_{i,a}(\tilde{a}, \hat{y}) | \sigma^*, \omega] = E[w_i(\tilde{a}', \hat{y}') | \sigma^*, \omega_{\max}]. \quad (10)$$

By construction, the sum of the self-evident component is always negative. For all $(\tilde{a}, \hat{y}) \in \text{supp}(\mu)$,

$$\sum_{i=1}^n w_{i,b}(\tilde{a}, \hat{y}) \leq 0,$$

with the equality holds when $P(\tilde{a}, \hat{y}) = \omega_{\max}$.

The efficiency loss of w can be similarly decomposed into two components:

$$W(\eta, w) = - \sum_{i=1}^n E[w_{i,a}(\tilde{a}, \hat{y}) | \sigma^*] - \sum_{i=1}^n E[w_{i,b}(\tilde{a}, \hat{y}) | \sigma^*]. \quad (11)$$

Denote the efficiency loss associated with the self-evident component by

$$L(\eta, w) \equiv - \sum_{i=1}^n E[w_{i,b}(\tilde{a}, \hat{y}) | \sigma^*]. \quad (12)$$

Substituting (11) and (10) into (12), we have

$$L(\eta, w) = W(\eta, w) + \sum_{i=1}^n E[w_i(\tilde{a}', \hat{y}') | \sigma^*, \omega_{\max}]. \quad (13)$$

Thus, the self-evident efficiency loss of a contract is equal to the total loss minus the loss conditional on ω_{\max} .

5 Main Results

In this section, we present our main results.

Let

$$L^*(\eta) \equiv \min_{w \in \mathcal{W}(\eta)} L(\eta, w) \quad (14)$$

denote the minimum self-evident efficiency loss of any contract that enforces η . Theorem 1 says that the per-period efficiency loss is bounded from below by $L^*(\eta)$.

Theorem 1. *For any enforceable η , $W^*(\eta, T, \delta) \geq L^*(\eta)$ for any $T \geq 1$ and $\delta \leq 1$.*

In Case 1 of the noisy Prisoners' Dilemma example in Section 2, we saw that the value of linking is limited by the need to provide separate incentives in continuation games after different realizations of a public signal. The same is true in general. Since any $\omega \in P$ is self-evident, the continuation games after different realizations of $\omega \in P$ can be treated as separate subgames. Hence, any efficiency loss associated with the self-evident component of a contract cannot be eliminated through linking.

When η is pure and the signal is public, every irreducible self-evident event is of the form $\{(\tilde{a}, y)\}$, where \tilde{a} is the unique element in the support of η . If w is an optimal one-shot contract, we must have

$$\sum_{i=1}^n E [w_i(\tilde{a}', \hat{y}') | \sigma^*, \omega_{\max}] = \max_{y \in Y} \sum_{i=1}^n w_i(\tilde{a}, y) = 0$$

and the second term in (13) is zero.¹⁵ Hence

$$L^*(\eta) = \min_{w \in \mathcal{W}(\eta)} W(\eta, w).$$

Theorem 1, therefore, implies that linking has no value when η is pure and the signal is public.

The converse of Theorem 1 holds under an additional condition. Following Blackwell (1953), we can think of a player's action as an experiment to generate information about the actions and signals of the other players. In Blackwell (1953), one experiment is more informative than another if the latter can be expressed as a garbling of the former. Let η_i denote the marginal distribution of player i 's action under η . Let $\gamma_i \in \Delta(A_i)$ denote a mixed action for player i , where $\gamma_i(a_i)$ is the probability of choosing a_i .

Definition 3. For any $\gamma_i, \gamma'_i \in \Delta(A_i)$, γ_i is more informative than γ'_i at the recommendation $\tilde{a}_i \in \text{supp}(\eta_i)$ if for any $(a_i, y_i) \in A_i \times Y_i$, there exists a distribution $\lambda_{(a_i, y_i)}(\cdot, \cdot) \in \Delta(A_i \times Y_i)$ such that for all $(\tilde{a}_{-i}, y_{-i}) \in A_{-i} \times Y_{-i}$ and all $(a'_i, y'_i) \in A_i \times Y_i$,

$$\sum_{(a_i, y_i) \in A_i \times Y_i} \lambda_{(a_i, y_i)}(a'_i, y'_i) \gamma(a_i) p(y_{-i}, y_i | \tilde{a}_{-i}, a_i) \eta(\tilde{a}) = \gamma'(a'_i) p(y_{-i}, y'_i | \tilde{a}_{-i}, a'_i) \eta(\tilde{a}). \quad (15)$$

¹⁵Otherwise, let y_{\max} maximize $\sum_{i=1}^n w_i(\tilde{a}, y)$. Then the contract $w_{i,b}(\tilde{a}, y) = w_i(\tilde{a}, y) - w_i(\tilde{a}, y_{\max})$ strictly improves upon w_i .

An action γ_i is strictly more informative than γ'_i if γ_i is more informative than γ'_i but not vice versa.

Unlike the setup in Blackwell (1953), where the distribution of states is fixed, in our model player i 's action may alter the distribution of y_{-i} . Equation (15) requires that for every \tilde{a}_{-i} with $\eta(\tilde{a}_i, \tilde{a}_{-i}) > 0$ (assuming that the other players are following the recommendations) γ_i lead to the same distribution of y_{-i} that γ'_i induces, and be more informative than γ'_i in the Blackwell sense. Since $\{\lambda_{(a_i, y_i)}(\cdot) | (a_i, y_i) \in A_i \times Y_i\}$ can be interpreted as a mixed reporting strategy, an equivalent definition is to say that γ_i is more informative than γ'_i if player i can choose γ_i and misreport y_i to mimic the distribution of y under γ'_i .

Definition 4. An action profile η satisfies the no-free-information condition if

$$\sum_{a_i \in A_i} \gamma_i(a_i) \sum_{\tilde{a}_{-i} \in A_{-i}} g(a_i, \tilde{a}_{-i}) \eta(\tilde{a}) < \sum_{\tilde{a}_{-i} \in A_{-i}} g(\tilde{a}) \eta(\tilde{a})$$

for any $i \in N$, $\tilde{a}_i \in \text{supp}(\eta_i)$, and γ_i strictly more informative than \tilde{a}_i at \tilde{a}_i .

In words, η satisfies the no-free-information condition if any deviation that generates more information for a player must strictly lower his stage-game payoff.

We can now state the converse of Theorem 1.

Theorem 2. *If η is enforceable and satisfies the no-free-information condition, then for any $\varepsilon > 0$, there exists T_0 such that, for any $T \geq T_0$ and any $\delta \geq 1 - T^{-2}$, $W^*(\eta, T, \delta) \leq L^*(\eta) + \varepsilon$.*

Theorem 2 says that the bound established in Theorem 1 is tight when η satisfies the no-free-information condition. It means that any efficiency loss due to incentives that vary within self-evident events can be eliminated in the long run. Note that when η is pure and P is a singleton, $L^*(\eta) = 0$. In the literature of repeated games with private monitoring, it is standard to assume that the signal distribution has full support, and the assumption is often treated as merely a simplifying assumption. In fact, since full support implies that P is a singleton, it, by itself, implies a pure action profile can be enforced efficiently in the long run.

Under a non-stationary contract that links incentives across periods, how an action is rewarded depends on past outcomes. Thus, players have incentives to deviate to actions that generate more information about the private information of the other players.¹⁶ The no-free-information condition ensures that no player can do so without paying a cost and without being detected. Note that the condition does not impose a lower bound on the cost to acquire more information. As T becomes large, the potential gain from having more information can be made arbitrarily small (but non-zero).

It is standard in the literature to assume that the desired outcome is strictly enforceable. Since strict enforceability rules out any profitable deviation from the recommendation, it implies the no-free-information condition. The no-free-information condition, however, is more intuitive and weaker than strict enforceability.

Definition 5 (Almost-strict enforceability). A contract w almost strictly enforces η if, for any player i and any strategy $\sigma_i \in \Sigma_i$,

$$v_i(\sigma^*; w_i) \geq v_i(\sigma_i, \sigma_{-i}^*; w_i),$$

with the inequality strict for any detectable σ_i . An action profile is almost-strictly enforceable if it can be enforced almost strictly by some w .

Unlike strict enforceability, almost-strict enforceability requires only that a player be strictly worse off when the deviation is detectable.

Lemma 3. *An enforceable action profile that satisfies the no-free-information condition is almost-strictly enforceable.*

Lemma 3 follows from the theory of alternatives. The converse of Lemma 3 is false as almost-strict enforceability does not rule out pure undetectable deviations that are more informative than the obedient strategy. Theorem 2 does not hold if the no-free-information condition is replaced with almost-strict enforceability. We show this by an example in the online Appendix.

Note that if η can be enforced by both w and w' , the latter almost strictly, then any linear combination of w and w' also enforces η almost strictly. Hence, by Lemma 3, if η satisfies the no-free-information condition, then there exists a contract w with $L(\eta, w)$ close to $L^*(\eta)$ that enforces η almost strictly.

¹⁶Under a truncated contract, a player will gain from learning whether the truncation is likely to occur.

5.1 Linking incentives

In this section, we describe the long-term contract in the proof of Theorem 2. The idea is similar to Case 3 of Section 2. Readers who are not interested in the details can skip this section.

For any $\varepsilon > 0$, let w^* denote a contract that enforces η almost strictly with

$$L(\eta, w^*) \leq L^*(\eta) + \frac{\varepsilon}{2}. \quad (16)$$

Let $w_{i,b}^*$ and $w_{i,a}^*$ denote, respectively, the self-evident and residual components of w_i^* (as defined in Section 4). Let w^{T*} , $w_{i,a}^{T*}$, and $w_{i,b}^{T*}$ denote the T -period version of w^* , $w_{i,a}^*$ and $w_{i,b}^*$, respectively. For all $i \in N$ and all (\tilde{a}^T, \hat{y}^T) ,

$$\begin{aligned} w_i^{T*}(\tilde{a}^T, \hat{y}^T) &= \sum_{t=1}^T \delta^{t-1} w_i^*(\tilde{a}(t), \hat{y}(t)); \\ w_{i,a}^{T*}(\tilde{a}^T, \hat{y}^T) &= \sum_{t=1}^T \delta^{t-1} w_{i,a}^*(\tilde{a}(t), \hat{y}(t)); \\ w_{i,b}^{T*}(\tilde{a}^T, \hat{y}^T) &= \sum_{t=1}^T \delta^{t-1} w_{i,b}^*(\tilde{a}(t), \hat{y}(t)). \end{aligned}$$

Fix some small $\kappa_T > 0$. Let

$$\begin{aligned} R_i^+(\tilde{a}^T, \hat{y}^T, \kappa_T) &= \max(0, w_{i,a}^{T*}(\tilde{a}^T, \hat{y}^T) - E[w_{i,a}^{T*}(\tilde{a}^{T'}, \hat{y}^{T'}) | \sigma^{T*}] - \kappa_T) \\ R_i^-(\tilde{a}^T, \hat{y}^T, \kappa_T) &= \min(0, w_{i,a}^{T*}(\tilde{a}^T, \hat{y}^T) - E[w_{i,a}^{T*}(\tilde{a}^{T'}, \hat{y}^{T'}) | \sigma^{T*}] - \kappa_T) \end{aligned}$$

denote, respectively, the positive and negative parts of the difference between $w_{i,a}^{T*}$ and the mean of $w_{i,a}^{T*}$ plus κ_T .

Let

$$B_i^T(\kappa_T) = \{(\tilde{a}^T, \hat{y}^T) \in A^T \times Y^T | R_i^+(\tilde{a}^T, \hat{y}^T, \kappa_T) > 0\}$$

denote set of (\tilde{a}^T, \hat{y}^T) where $R_i^+(\tilde{a}^T, \hat{y}^T, \kappa_T) > 0$. Recall that $Z^{T*}(\iota)$ denotes the set of (\tilde{a}^T, \hat{y}^T) whose distribution of outcomes deviates from the expected distribution by ι conditional on self-evident events. For $w_{i,a}^*$ (\tilde{a}^T, \hat{y}^T) to deviate from the mean, the realized outcome distribution must differ from the expected distribution as well.¹⁷ Hence,

¹⁷By construction $E[w_{i,a}(\tilde{a}, \hat{y}) | \sigma^*, \omega] = E[w_{i,a}(\tilde{a}, \hat{y}) | \sigma^*]$ for all $\omega \in P$.

there exists $\iota_0 > 0$ such that $B_i^T(\kappa_T) \subseteq Z^*(\iota_0 \kappa_T)$. Then, by Lemma 2,

$$B_i^T(\kappa_T) \subseteq \cup_{j \in N} Z_j^{T*}(c_0 \iota_0 \kappa_T). \quad (17)$$

We can therefore define a vector of indicator functions as follows. For any $i \in N$, set

$$I_i(\tilde{a}^T, \hat{y}^T, \kappa_T) = \begin{cases} 1 & \text{if } (\tilde{a}^T, \hat{y}^T) \in Z_i^*(c_0 \iota_0 \kappa_T), \\ 0 & \text{otherwise.} \end{cases}$$

Now, define a new contract w^{T**} . For all (\tilde{a}^T, \hat{y}^T) , set

$$w_i^{T**}(\tilde{a}^T, \hat{y}^T, \kappa_T) = R_i^-(\tilde{a}^T, \hat{y}^T, \kappa_T) + w_{i,b}^{T*}(\tilde{a}^T, \hat{y}^T, \kappa_T) + \left[R_i^+(\tilde{a}^T, \hat{y}^T, \kappa_T) (1 - I_i(\tilde{a}^T, \hat{y}^T, \kappa_T)) - \sum_{j \neq i} R_j^+(\tilde{a}^T, \hat{y}^T, \kappa_T) I_i(\tilde{a}^T, \hat{y}^T, \kappa_T) \right]. \quad (18)$$

Under this new contract, each player i is paid the self-evident component of w_i^{T*} , the part of the residual that is less than the mean plus κ_T , and a third component (inside the square bracket) that pays player i R_i^+ when $I_i = 0$ and $-R_j^+$ when $I_i = 1$. The total payment is negative for all (\tilde{a}^T, \hat{y}^T) . By definition, for all (\tilde{a}^T, \hat{y}^T) ,

$$\sum_{i=1}^n (R_i^-(\tilde{a}^T, \hat{y}^T, \kappa_T) + w_{i,b}^{T*}(\tilde{a}^T, \hat{y}^T, \kappa_T)) \leq 0.$$

The sum of the third component in (18) is also negative. By (17), for any player i and any $(\tilde{a}^T, \hat{y}^T) \in B_i^T(\kappa_T)$, $I_j(\tilde{a}^T, \hat{y}^T, \kappa_T) = 1$ for some player j . Intuitively, if player i is to receive some $R_i^+(\tilde{a}^T, \hat{y}^T, \kappa_T) > 0$, then some other player j must pay for it.

Rearranging the terms on the right-hand side of (18), we can write

$$w_i^{T**}(\tilde{a}^T, \hat{y}^T, \kappa_T) = w_i^{T*}(\tilde{a}^T, \hat{y}^T) - E[w_{i,a}^{T*}(\tilde{a}^T, \hat{y}^T) | \sigma^{T*}] - \kappa_T - \phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T), \quad (19)$$

where

$$\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) = \sum_{j=1}^n R_j^+(\tilde{a}^T, \hat{y}^T, \kappa_T) I_i(\tilde{a}^T, \hat{y}^T, \kappa_T)$$

measures the distortion in incentives.

The following lemma shows that we can choose κ_T so that the expected value of $\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T)$ conditional on any private information player i may learn during the game on the equilibrium path diminishes uniformly and exponentially with T .

Let H_i^{T*} denote the set of histories that player i may observe during the T -period contract under σ^{T*} .

Lemma 4. *Let $\kappa_T = T^{2/3}$. There exists $c > 0$ such that for all $i \in N$, $T \geq 1$, and $h_i \in H_i^{T*}$,*

$$E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma^{T*}, h_i] < cT \exp\left(-\frac{(c_0 t_0)^2}{2} T^{1/3}\right).$$

We prove Lemma 4 by the Hoeffding inequality (Hoeffding, 1963). Lemma 4 means that, when T is sufficiently large, a player i who has followed σ_i^{T*} up to some period $t \leq T$ (conditional on any private information that he may observe from period 1 to $(t-1)$) will believe that if he follows the equilibrium strategy σ_i^{T*} in the remaining periods, he will obtain a ϕ_i close to zero. This, together with the fact that ϕ_i is always positive, means that no deviation can reduce the expected value of ϕ_i significantly. Since w_i^{T*} enforces η almost strictly and η satisfies the no-free-information condition, a player i deviating in any period must be strictly worse off if the deviation is detectable or more informative than the recommended one. Since the effect of a single-period deviation on the total payoff is of the order $1/T$ (as δ goes to one), the players will have the incentives to play η under w^{T**} when T is sufficiently large, as the distortion in incentives due to the truncation diminishes at a rate faster than $1/T$.

The per-period efficiency loss of w^{T**} is

$$\begin{aligned} & -\frac{1-\delta}{1-\delta^T} \sum_{i=1}^n E[w_i^{T*}(\tilde{a}^T, \hat{y}^T) - E[w_{i,a}^{T*}(\tilde{a}^{T'}, \hat{y}^{T'}) | \sigma^{T*}] - \kappa_T - \phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma^{T*}] \\ & = L(\eta, w^*) + \frac{1-\delta}{1-\delta^T} \left(n\kappa_T + \sum_{i=1}^n E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma^{T*}] \right). \end{aligned} \tag{20}$$

The per-period efficiency loss converges to $L(\eta, w^*)$ as the second term in the last equation converges to zero as δ goes to one and T goes to infinity.

6 Value of Linking

In this section, we characterize $L^*(\eta)$ in terms of the primitives of the model. The following theorem provides a sufficient condition for $L^*(\eta) = 0$.

Theorem 3A. For any enforceable η , we have $L^*(\eta) = 0$ if for any player $i \in N$, any deviating strategy $\sigma_i \in \Sigma_i / \{\sigma_i^*\}$ satisfies one of the following conditions:

- (i) There exists (\tilde{a}, \hat{y}) such that $\pi^{\sigma_i}(\tilde{a}, \hat{y}) > 0$ and $\mu(\tilde{a}, \hat{y}) = 0$.
- (ii) There exists $\omega \in P$ such that $\pi^{\sigma_i}(\cdot | \omega) \neq \mu(\cdot | \omega)$.
- (iii) There exists a player $j \in N$ such that there is no σ_j with $\pi^{\sigma_j} = \pi^{\sigma_i}$.

Theorem 3A identifies three types of deviations that can be deterred almost costlessly as T becomes large: first, deviations that may result in (\tilde{a}, \hat{y}) outside of the support of μ ; second, deviations that change the distribution of (\tilde{a}, \hat{y}) conditional on some $\omega \in P$; third, deviations that lead to distributions of (\tilde{a}, \hat{y}) that cannot be caused by some other player. Rahman and Obara (2010) call the last type of deviations *attributable*. By contrast, a deviating strategy profile $(\sigma_1, \dots, \sigma_n)$ satisfying

$$\pi^{\sigma_1} = \dots = \pi^{\sigma_n}$$

is *unattributable*, as the common distribution π^{σ_1} could have been caused by any player.

The first type of deviations can be deterred costlessly by a contract that punishes all players severely when an out-of-support (\tilde{a}, \hat{y}) occurs. The second type can be deterred by a contract with a zero self-evident component. The third type can be deterred by a budget-balance contract (Rahman and Obara, 2010).

Not all outcome profiles can be enforced almost efficiently in the long run. The next theorem characterizes $L^*(\eta)$ for all η . Let

$$Q(\eta) \equiv \{\sigma \in \Sigma | \pi^{\sigma_1} = \dots = \pi^{\sigma_n} \in \text{co}(\{\mu(\cdot | \omega) | \omega \in P\}) / \{\mu\}\}$$

denote the set of unattributable deviations that are distinct from μ and undetectable with respect to any $\omega \in P$. If σ_i does not satisfy any of Conditions (i) to (iii) in Theorem 3A, then there must exist σ_{-i} such that (σ_i, σ_{-i}) belongs to $Q(\eta)$.

For any σ_i , let

$$d_i(\sigma_i) \equiv \sum_{(\alpha_i, \rho_i)} \sigma_i(\alpha_i, \rho_i) \sum_{\tilde{a} \in A} (g_i(\tilde{a}_{-i}, \alpha_i(\tilde{a}_i)) - g_i(\tilde{a})) \eta(\tilde{a})$$

denote player i 's gain from the deviation σ_i , and let

$$l(\sigma_i) \equiv \max_{\omega \in P} \frac{\pi^{\sigma_i}(\omega)}{\mu(\omega)}$$

measure the difference between π^{σ_i} and μ . Consider some $\sigma \in Q(\eta)$. Any w that enforces η must satisfy, for all player i , the incentive-compatibility constraint that

$$\sum_{(\tilde{a}, \hat{y}) \in A \times Y} (\mu(\tilde{a}, \hat{y}) - \pi^{\sigma_i}(\tilde{a}, \hat{y})) w_i(\tilde{a}, \hat{y}) \geq d(\sigma_i). \quad (21)$$

Since $\sigma \in Q(\eta)$, for all $\omega \in P$,

$$E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] = E[w_i(\tilde{a}, \hat{y}) | \pi^{\sigma_i}, \omega]. \quad (22)$$

Substituting (22) into (21), and summing over i , we have

$$\sum_{\omega \in P} (\mu(\omega) - \pi^{\sigma}(\omega)) \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] \geq \sum_{i=1}^n d(\sigma_i). \quad (23)$$

It then follows from the definition of L , (13), that

$$\begin{aligned} L(\eta, w) &= \sum_{\omega \in P} \mu(\omega) \left(- \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] + \max_{\omega' \in P} \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega'] \right) \\ &\geq \sum_{\omega \in P} \mu(\omega) \frac{\frac{\pi^{\sigma_1}(\omega)}{\mu(\omega)} - 1}{l(\sigma_1) - 1} \left(- \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] + \max_{\omega' \in P} \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega'] \right) \\ &\geq \sum_{\omega \in P} \frac{\mu(\omega) - \pi^{\sigma_1}(\omega)}{l(\sigma_1) - 1} \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] \\ &\geq \frac{\sum_{i=1}^n d(\sigma_i)}{l(\sigma_1) - 1}. \end{aligned}$$

The first inequality follows from the definition of $l(\sigma_i)$. The second inequality follows from the fact that $\sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega'] \leq 0$. The last inequality follows from (23). Intuitively, since σ is unattributable, every player must be punished, and the total punishment must be greater than $\sum_{i=1}^n d(\sigma_i)$, the total deviating gain. The resulting efficiency loss is equal to the total deviation gain multiplied by a factor that measures the difference between π^{σ_i} and μ . The smaller the difference, the harder it is to distinguish between the two distributions, and, hence, the higher the efficiency loss. Since the argument applies to every w that enforces η , for any $\sigma \in Q(\eta)$,

$$L^*(\eta) \geq \frac{\sum_{i=1}^n d(\sigma_i)}{l(\sigma_1) - 1}.$$

Theorem 3B shows that the converse is also true.

Theorem 3B. For any enforceable η ,

$$L^*(\eta) = \sup_{(\sigma_1, \dots, \sigma_n) \in Q(\eta)} \frac{\max(\sum_{i=1}^n d(\sigma_i), 0)}{l(\sigma_1) - 1}$$

if $Q(\eta)$ is nonempty. Otherwise, $L^*(\eta) = 0$.

Theorem 3B implies that $L^*(\eta) > 0$ only if there is some $\sigma \in Q(\eta)$ with $\sum_{i=1}^n d(\sigma_i) > 0$. Theorem 3B, thus, implies Theorem 3A.

In the noisy prisoners' dilemma example in Section 2, when the signals are perfectly correlated ($\rho = 0$), conditional on (C, C) there are two irreducible singleton self-evident events: HH and LL . Either player deviating to D results in the same signal distribution $(q, 1 - q)$. Hence, the only way to enforce (C, C) is to punish the players when the signal L occurs. By Theorem 3B, the per-period efficiency loss is

$$\frac{2d}{\max\left(\frac{q}{p}, \frac{1-q}{1-p}\right) - 1} = 2 \frac{(1-p)d}{p-q},$$

where d is a player's gain from deviating to D . Theorem 3B effectively shows that any long-term efficiency loss arises for the same reason as in the prisoners' dilemma example. When $Q(\eta)$ contains multiple strategy profiles, the efficiency loss is determined by the one that is the hardest to deter.

6.1 Relation to literature

Our results are closely related to the literature in repeated games with private monitoring and communication.¹⁸ The literature can be divided into two strands. One strand, following Abreu, Milgrom, and Pearce (1991), applies the linking idea to obtain approximate efficiency (Compte, 1998; Obara, 2009; Chan and Zhang, 2016; Rahman, 2014; Sugaya, 2017a,b).¹⁹ Another strand (Fudenberg, Levine, and Maskin, 1994; Kandori

¹⁸Instead of infinitely repeated games, we work with a T -period contracting problem that allows us to focus on the mechanism of linking and abstract away from the problem of implementing transfers through continuation strategies. Our results can be readily applied to repeated games with side-payments.

¹⁹The literature of repeated games with private monitoring and without communication also exploits the idea of linking. See, e.g., Matsushima (2004) and Fong, Gossner, Hörner, and Sannikov (2011).

and Matsushima, 1998; Rahman and Obara, 2010) identifies conditions that ensure that an outcome can be costlessly enforced in a stage game by a budget-balanced contract.

Theorem 3A naturally combines the two approaches. It, together with Theorem 2, implies that an outcome can be enforced in the long-run almost efficiently if every deviating strategy can be deterred *either* by a contract that links efficiently or is budget balanced. Theorems 1 and 3B further connect the efficiency results in the repeated-game literature with the inefficiency result of Abreu, Milgrom, and Pearce (1991).²⁰ It shows that the two approaches taken by the repeated-game literature are, in fact, the only approaches to obtain efficiency. If there is a deviation that cannot be deterred by either one, then long-run efficiency loss may be inevitable, and the way the efficiency loss arises is exactly the same as in Abreu, Milgrom, and Pearce (1991) (Case 1 of Section 2).

In a repeated game, the players observe private signals at the end of every period. While the players may delay revealing their signals, each may nevertheless update his beliefs about other players' signals on the basis of his own. Compte (1998) and Obara (2009) deal with this problem by imposing restrictions on the signal structure to ensure the existence of a contract that enforces the desired outcome with the property that no player can learn about his payment from his own signals.²¹ Rahman (2014) adopts a similar approach when correlated strategies are allowed. We show that these restrictions are not necessary as linking could work even when a player can learn about his payment. Our approach relies on the fact that, when T is large, it is possible to have one player to partially compensate another player for the truncation effect so long as the truncation is not self evident. Chan and Zhang (2016) adopt a similar approach to obtain efficiency in repeated games where players observe their own payoffs and the signal distribution has full support. While Chan and Zhang (2016) consider only pure outcomes, our framework applies to all monitoring structures (both public and private) and action profiles (both pure and correlated).

In two recent papers, Sugaya (2017a,b) derives upper and lower bounds in equilib-

²⁰Compte (1998), Obara (2009) and Chan and Zhang (2016) consider pure action profile and assume that the signal distribution has full support, which rules out public signals.

²¹Compte (1998) assumes independent signals. Obara (2009) considers correlated signals and identifies a condition on the signal distribution that ensures that no player can learn about his own transfer.

rium payoffs in repeated games with private monitoring and correlated strategies. In particular, Sugaya (2017a) shows that approximate efficiency can be achieved when the players can observe their own payoffs. Sugaya’s method always requires a correlating device, while ours requires one only when the outcome is correlated. While Sugaya (2017a,b) focuses on the equilibrium payoff set, we focus on the minimum efficiency loss associated with the enforcement of a given outcome.

7 Correlated Strategies

Rahman and Obara (2010) show that, instead of having a disinterested owner to break the budget constraint, a partnership may reduce the cost of noisy monitoring by using a disinterested mediator to implement a correlated strategy that identifies non-deviators. In our setting, players may also change the information structure endogenously through a correlated strategy profile.²² Since a small change in a correlated strategy can alter the support of the distribution of action-signal profiles substantially, it can have a large impact on the long-run efficiency loss as δ goes to one and T goes to infinity. The idea is first raised by Rahman (2014). He proves a folk theorem under the condition that the signal distribution satisfies conditional identifiability.

To illustrate the idea, let us return to the noisy Prisoners’ Dilemma in Case 1 ($\rho = 0$) of Section 2. As we showed in the last section, $L^*(C, C) = 2(1 - p)d/(p - q)$. Consider the correlated strategy profile $\bar{\eta}$ where

$$\bar{\eta}(C, C) = 1 - \varepsilon; \bar{\eta}(C, D) = \bar{\eta}(D, C) = 0.5\varepsilon.$$

When ε is small, $\bar{\eta}$ is close to the pure strategy profile (C, C) . Yet the support of the distribution of the action-signal profiles under $\bar{\eta}$ is very different from the support under (C, C) . Now each P_i consists of four elements. In particular,

$$\begin{aligned} P_1 & : \{CCH, CDH\}, \{CCL, CDL\}, \{DCH\}, \{DCL\} \\ P_2 & : \{CCH, DCH\}, \{CCL, DCL\}, \{CDH\}, \{CDL\}, \end{aligned}$$

²²This involves changing the correlating device correspondingly.

and the meet of P_1 and P_2 is

$$P : \{CCH, CDH, DCH\}, \{CCL, CDL, DCL\}.$$

Note that each member of the meet contains multiple elements. As a player is instructed to choose C , he is not sure whether the other player is choosing C or D . Let $p(H|DD) = r$. Conditional identifiability requires that

$$\frac{p}{q} \neq \frac{q}{r} \quad \text{and} \quad \frac{1-p}{1-q} \neq \frac{1-q}{1-r}.$$

The condition says that how player i 's action affects the relative likelihood of H and L depends on whether player j is choosing C or D . Since player i does not observe the action of player j , $\bar{\eta}$ can be “secretly” enforced so that the players do not learn about their own punishments from the public signal. Hence, by choosing a small ε , the players can obtain close to the efficient payoff $(1, 1)$, when the players are sufficiently patient and T is sufficiently large.

Using Theorems 2 and 3A, we can show that $\bar{\eta}$ can be enforced almost efficiently without assuming conditional identifiability. It is straightforward to see that $\bar{\eta}$ is enforceable. Let α_i^{xy} denote the strategy of choosing x when C is recommended and y when D is recommended. Each player has four pure action strategies: α_i^{CD} , α_i^{DD} , α_i^{CC} , and α_i^{DC} . In Table 3, each row gives the probabilities of outcomes with an H signal under a different pure strategy of player 1 (assuming that player 2 plays α_2^{CD}).

	CCH	DCH	CDH
α_1^{CD}	$(1 - \varepsilon)p$	$0.5\varepsilon q$	$0.5\varepsilon q$
α_1^{DD}	$(1 - \varepsilon)q$	$0.5\varepsilon q$	$0.5\varepsilon r$
α_1^{CC}	$(1 - \varepsilon)p$	$0.5\varepsilon p$	$0.5\varepsilon q$
α_1^{DC}	$(1 - \varepsilon)q$	$0.5\varepsilon p$	$0.5\varepsilon r$

Table 3: The probability for each outcome with an H signal.

Notice that the ratio of the relative probability of CCH over DCH is strictly higher when player 1 follows the recommendation and plays α_1^{CD} . Hence, every deviation is detectable with respect to the self-evident event $\{CCH, CDH, DCH\}$. By Theorem 3A,

$L^*(\bar{\eta}) = 0$. Intuitively, the recommendation DC serves as a “benchmark” for player 1. Given that player 2 is choosing C , player 1 choosing D minimizes the probability of H . Hence, if player 1 deviates to D when told to choose C , he must lower the relative probability of CCH over DCH . Finally, since any unilateral deviation from $\bar{\eta}$ is detectable, $\bar{\eta}$ satisfies the no-free-information condition. Hence, by Theorem 2, $\bar{\eta}$ can be enforced almost efficiently in the long run.

The following theorem generalizes the above example. It says that any strictly enforceable outcome can be virtually enforced with almost no long-run efficiency loss.²³ The proof is in the online Appendix.

Theorem 4. *For any strictly enforceable η and any $\varepsilon > 0$, there exists an enforceable correlated action profile $\bar{\eta}$ that satisfies the no-free-information condition and with $\max_{\tilde{a} \in A} |\eta(\tilde{a}) - \bar{\eta}(\tilde{a})| \leq \varepsilon$ and $L^*(\bar{\eta}) = 0$.*

8 Conclusion

Players in a long-run relationship can reduce incentive costs by linking incentives across periods, but the value of linking is limited by the information the players obtain during the course of the relationship. We show that the long-run per-period efficiency loss in enforcing an action profile is bounded from below by the incentive cost that becomes self-evident at the end of each period, and the bound is tight when players cannot obtain free information undetectably. The results extend the insights of Abreu, Milgrom, and Pearce (1991) to general stage games where players may observe both public and private signals and use a correlating device to coordinate their actions.

²³Theorem 4 does not hold if η merely satisfies the no-free-information condition but is not strictly enforceable. The strict enforceability of η , together with the fact that $\bar{\eta}$ is close to η , ensures that under $\bar{\eta}$ no player can deviate undetectably without strictly reducing his stage-game payoff when recommended to choose an action in the support of η .

A Proof of Theorem 1

By definition, $W^*(\eta) \geq L^*(\eta)$. Hence, Theorem 1 holds for $T = 1$. Suppose the theorem holds for $T - 1$. Consider the T -period case.

Let $\tilde{a}^{2,T}$ and $\hat{y}^{2,T}$ denote, respectively, the value of \tilde{a}^T and \hat{y}^T from period 2 through T . Fix $w^T \in \mathcal{W}(\eta, T, \delta)$. For each i and each $(\tilde{a}(1), \hat{y}(1)) \in A \times Y$, let

$$w_i(\tilde{a}(1), \hat{y}(1)) \equiv \sum_{(\tilde{a}^{2,T}, \hat{y}^{2,T})} w_i^T(\tilde{a}(1), \tilde{a}^{2,T}, \hat{y}(1), \hat{y}^{2,T}) \prod_{t=2}^T \mu(\tilde{a}(t), \hat{y}(t))$$

denote the expected value of w_i^T conditional on $(\tilde{a}(1), \hat{y}(1))$, assuming that all players follow the equilibrium strategy.

For each i , each $(\tilde{a}^{2,T}, \hat{y}^{2,T})$ and each $\omega \in P$, let

$$w_i^{T-1, \omega}(\tilde{a}^{2,T}, \hat{y}^{2,T}) \equiv \sum_{(\tilde{a}(1), \hat{y}(1)) \in \omega} w_i^T(\tilde{a}^T, \hat{y}^T) \mu(\tilde{a}(1), \hat{y}(1) | \omega)$$

denote the expected value of w_i^T conditional on $(\tilde{a}^{2,T}, \hat{y}^{2,T})$ and $(\tilde{a}(1), \hat{y}(1))$ being in the set ω , assuming that all players follow the equilibrium strategy.

Since w^T enforces η for T periods, $w = (w_1, \dots, w_n)$ must enforce η in the first period. Hence, $w \in \mathcal{W}(\eta)$. Furthermore, since ω is self-evident, the continuation game following ω can be treated as a $(T - 1)$ -period game with an extra randomization device $\mu(\cdot | \omega)$, and the contract $\delta^{-1} w^{T-1, \omega}$ must enforce η in this game. By the revelation principle, adding this extra randomization device will not enhance the efficiency of the contract. Hence, by the supposition that Theorem 1 holds for $T - 1$, the expected efficiency loss of this contract, which is equal to

$$-\delta^{-1} \sum_{i=1}^n w_i(\tilde{a}(1), \hat{y}(1)) \mu(\tilde{a}(1), \hat{y}(1) | \omega),$$

must be greater than $\frac{1-\delta^{T-1}}{1-\delta} L^*(\eta)$. Let

$$\omega_{\max} \in \arg \max_{\omega \in P} \sum_{i=1}^n E[w_i(\tilde{a}(1), \hat{y}(1)) | \sigma^*, \omega].$$

It follows that

$$\begin{aligned}
& - \sum_{i=1}^n E[w_i^T(\tilde{a}^T, \hat{y}^T) | \sigma^{T*}] \\
= & - \sum_{i, (\tilde{a}(1), \hat{y}(1))} w_i(\tilde{a}(1), \hat{y}(1)) \mu(\tilde{a}(1), \hat{y}(1)) \\
\geq & L^*(\eta) - \sum_{i, (\tilde{a}(1), \hat{y}(1)) \in \omega_{\max}} w_i(\tilde{a}(1), \hat{y}(1)) \mu(\tilde{a}(1), \hat{y}(1)) \omega_{\max} \\
\geq & L^*(\eta) + \delta \frac{1 - \delta^{T-1}}{1 - \delta} L^*(\eta) \\
= & \frac{1 - \delta^T}{1 - \delta} L^*(\eta).
\end{aligned}$$

The first inequality follows from the fact that $w \in \mathcal{W}(\eta)$.

B Proof of Lemma 2

Define two constants

$$c_1 \equiv \max_{(\tilde{a}, y) \in \text{supp}(\mu)} \frac{1}{\mu(\tilde{a}, y)} \quad \text{and} \quad c_0 \equiv \frac{1}{(2|\text{supp}(\mu)| + 1)c_1}.$$

We show for all $(\tilde{a}^T, y^T) \in \text{supp}(\mu)^T$, if $(\tilde{a}^T, y^T) \notin Z_i^{T*}(c_0 t)$ for all player i , then $(\tilde{a}^T, y^T) \notin Z^{T*}(t)$.

Fix any $(\tilde{a}', y') \in \text{supp}(\mu)$ and let ω be the element of P that contains (\tilde{a}', y') . Since ω is an element of the meet P , any $(\tilde{a}'', y'') \in \omega$ is reachable from (\tilde{a}', y') , i.e., there exists a sequence $(\tilde{a}', y') = (\tilde{a}^1, y^1), (\tilde{a}^2, y^2), \dots, (\tilde{a}^k, y^k) = (\tilde{a}'', y'')$ such that $(\tilde{a}^s, y^s) \in \omega$ for each $s \leq k$ and that any two consecutive profiles (\tilde{a}^s, y^s) and $(\tilde{a}^{s+1}, y^{s+1})$ have the same i_s -th component for player $i_s \in N$ (see, e.g., Aumann, 1976; Geanakoplos, 1994).

For each (\tilde{a}^s, y^s) and each $m \in N$, by supposition, we have

$$|f(\tilde{a}^s, y^s | \tilde{a}^T, y^T) - \mu_{-m}(\tilde{a}_{-m}^s, y_{-m}^s | \tilde{a}_m^s, y_m^s) f(\tilde{a}_m^s, y_m^s | \tilde{a}^T, y^T)| \leq c_0 t. \quad (24)$$

Dividing both sides of (24) by $\mu(\tilde{a}^s, y^s)$, we have

$$\left| \frac{f(\tilde{a}_m^s, y_m^s | \tilde{a}^T, y^T)}{\mu_m(\tilde{a}_m^s, y_m^s)} - \frac{f(\tilde{a}^s, y^s | \tilde{a}^T, y^T)}{\mu(\tilde{a}^s, y^s)} \right| \leq \frac{c_0 t}{\mu(\tilde{a}^s, y^s)} \leq c_1 c_0 t. \quad (25)$$

Fix any two players $i, j \in N$. It follows from (25) that

$$\begin{aligned}
& \left| \frac{f(\tilde{a}'_i, y'_i | \tilde{a}^T, y^T)}{\mu_i(\tilde{a}'_i, y'_i)} - \frac{f(\tilde{a}''_j, y''_j | \tilde{a}^T, y^T)}{\mu_j(\tilde{a}''_j, y''_j)} \right| \\
& \leq \left| \frac{f(\tilde{a}'_i, y'_i | \tilde{a}^T, y^T)}{\mu_i(\tilde{a}'_i, y'_i)} - \frac{f(\tilde{a}^1, y^1 | \tilde{a}^T, y^T)}{\mu(\tilde{a}^1, y^1)} \right| + \left| \frac{f(\tilde{a}^1, y^1 | \tilde{a}^T, y^T)}{\mu(\tilde{a}^1, y^1)} - \frac{f(\tilde{a}^1_{i_1}, y^1_{i_1} | \tilde{a}^T, y^T)}{\mu_{i_1}(\tilde{a}^1_{i_1}, y^1_{i_1})} \right| \\
& \quad + \left| \frac{f(\tilde{a}^2_{i_1}, y^2_{i_1} | \tilde{a}^T, y^T)}{\mu_{i_1}(\tilde{a}^2_{i_1}, y^2_{i_1})} - \frac{f(\tilde{a}^2, y^2 | \tilde{a}^T, y^T)}{\mu(\tilde{a}^2, y^2)} \right| + \dots + \left| \frac{f(\tilde{a}^k, y^k | \tilde{a}^T, y^T)}{\mu(\tilde{a}^k, y^k)} - \frac{f(\tilde{a}''_j, y''_j | \tilde{a}^T, y^T)}{\mu_j(\tilde{a}''_j, y''_j)} \right| \\
& \leq 2kc_1c_0\iota \leq 2|\omega|c_1c_0\iota.
\end{aligned}$$

This implies that

$$\frac{f(\tilde{a}''_j, y''_j | \tilde{a}^T, y^T)}{\mu_j(\tilde{a}''_j, y''_j)} - 2|\omega|c_1c_0\iota \leq \frac{f(\tilde{a}'_i, y'_i | \tilde{a}^T, y^T)}{\mu_i(\tilde{a}'_i, y'_i)} \leq \frac{f(\tilde{a}''_j, y''_j | \tilde{a}^T, y^T)}{\mu_j(\tilde{a}''_j, y''_j)} + 2|\omega|c_1c_0\iota. \quad (26)$$

Note that (26) holds for all j and all (\tilde{a}''_j, y''_j) . Multiplying each side of (26) by $\mu_j(\tilde{a}''_j, y''_j)$ and summing over all $(\tilde{a}''_j, y''_j) \in A_j \times Y_j$ for which there exists $(\tilde{a}''_{-j}, y''_{-j}) \in A_{-j} \times Y_{-j}$ such that $(\tilde{a}'', y'') = (\tilde{a}''_j, \tilde{a}''_{-j}, y''_j, y''_{-j}) \in \omega$, we have

$$f(\omega | \tilde{a}^T, y^T) - 2|\omega|c_1c_0\iota\mu(\omega) \leq \frac{f(\tilde{a}'_i, y'_i | \tilde{a}^T, y^T)}{\mu_i(\tilde{a}'_i, y'_i)}\mu(\omega) \leq f(\omega | \tilde{a}^T, y^T) + 2|\omega|c_1c_0\iota\mu(\omega).$$

Using (25), we have

$$\left| \frac{f(\tilde{a}', y' | \tilde{a}^T, y^T)}{\mu(\tilde{a}', y')} - \frac{f(\omega | \tilde{a}^T, y^T)}{\mu(\omega)} \right| \leq (2|\omega| + 1)c_1c_0\iota,$$

or

$$\begin{aligned}
& |f(\tilde{a}', y' | \tilde{a}^T, y^T) - \mu(\tilde{a}', y')f(\omega | \tilde{a}^T, y^T)| \\
& \leq (2|\omega| + 1)c_1c_0\iota\mu(\tilde{a}', y') \leq (2|\text{supp}(\mu)| + 1)c_1c_0\iota = \iota.
\end{aligned}$$

C Proof of Lemma 3

By Lemma 1, pure and undetectable strategies (α_i, ρ_i) are unprofitable. Hence, it suffices to show that there exists a contract w such that for any player i and any pure strategy (α_i, ρ_i) ,

$$\sum_{(\tilde{a}, \hat{y}) \in A \times Y} [\mu(\tilde{a}, \hat{y}) - \pi^{\alpha_i, \rho_i}(\tilde{a}, \hat{y})] w_i(\tilde{a}, \hat{y}) \geq \sum_{\tilde{a} \in A} (g_i(\tilde{a}_{-i}, \alpha_i(\tilde{a}_i)) - g_i(\tilde{a})) \eta(\tilde{a}) \quad (27)$$

with the inequality strict if $\pi^{\alpha_i, \rho_i} \neq \mu$. By the theory of alternatives (see, e.g., Proposition 5.6.2 of Bertsekas, 2009), (27) does not have a solution w_i if and only if there exists $\lambda_i(\alpha_i, \rho_i) \geq 0$ for each (α_i, ρ_i) such that

$$\sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) [\mu(\tilde{a}, \hat{y}) - \pi^{\alpha_i, \rho_i}(\tilde{a}, \hat{y})] = 0 \text{ for each } (\tilde{a}, \hat{y}),$$

and either one of the following two cases holds:

- (i) We have $\sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) \sum_{\tilde{a} \in A} (g_i(\tilde{a}_{-i}, \alpha_i(\tilde{a}_i)) - g_i(\tilde{a})) \eta(\tilde{a}) > 0$.
- (ii) We have $\sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) \sum_{\tilde{a} \in A} (g_i(\tilde{a}_{-i}, \alpha_i(\tilde{a}_i)) - g_i(\tilde{a})) \eta(\tilde{a}) \geq 0$ and $\lambda_i(\alpha_i, \rho_i) > 0$ for some (α_i, ρ_i) such that $\pi^{\alpha_i, \rho_i} \neq \mu$.

In either case, $\lambda_i(\alpha'_i, \rho'_i) > 0$ for some (α'_i, ρ'_i) . By dividing each $\lambda_i(\alpha_i, \rho_i)$ by $\sum_{(\alpha'_i, \rho'_i)} \lambda_i(\alpha'_i, \rho'_i)$ if necessary, we may assume that $\sum_{(\alpha'_i, \rho'_i)} \lambda_i(\alpha'_i, \rho'_i) = 1$. That is, λ_i represents a mixed strategy for player i .

Since η is enforceable, by Lemma 1, Case (i) cannot hold. Case (ii) violates the no-free-information condition. Hence (27) must have a solution.

D Proof of Lemma 4

We apply the following inequality of Hoeffding (1963) to prove the lemma. Suppose that $\xi(1), \xi(2), \dots, \xi(T)$ are independent random variables such that $|\xi(t)| \leq v$ for each $t \leq T$. Then, for any $\kappa > 0$, Hoeffding's inequality asserts that

$$\Pr \left(\sum_{t=1}^T \xi(t) \geq E \left[\sum_{t=1}^T \xi(t) \right] + \kappa \right) \leq \exp \left(-\frac{\kappa^2}{2v^2 T} \right).$$

Fix any $(\tilde{a}, y) \in \text{supp}(\mu)$ and (\tilde{a}_i^T, y_i^T) , we estimate the probability that $(\tilde{a}^T, y^T) \in Z_i^T((\tilde{a}, y), c_0 \mathbf{t}_0 \kappa_T)$ conditional on (\tilde{a}_i^T, y_i^T) . Focus on the $f(\tilde{a}_i, y_i | \tilde{a}^T, y^T)$ periods in which player i observes (\tilde{a}_i, y_i) . Note that we can rewrite $f(\tilde{a}, y | \tilde{a}^T, y^T)$ as the sum of indicators $1_{(\tilde{a}_{-i}, y_{-i})}$ for these periods that equals 1 if (\tilde{a}_{-i}, y_{-i}) occurs. Then $\mu_{-i}(\tilde{a}_{-i}, y_{-i} | \tilde{a}_i, y_i) f(\tilde{a}_i, y_i | \tilde{a}^T, y^T)$

is the mean of this sum. Hence, by Hoeffding's inequality, we have

$$\begin{aligned}
& \Pr \left(\left| f(\tilde{a}, y | \tilde{a}^T, y^T) - \mu_{-i}(\tilde{a}_{-i}, y_{-i} | \tilde{a}_i, y_i) f(\tilde{a}_i, y_i | \tilde{a}^T, y^T) \right| > c_0 \iota_0 \kappa_T \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right) \\
&= \Pr \left(\left| \sum_{t: (\tilde{a}_i(t), y_i(t)) = (\tilde{a}_i, y_i)} \mathbf{1}_{(\tilde{a}_{-i}, y_{-i})}(\tilde{a}_{-i}(t), y_{-i}(t)) \right. \right. \\
&\quad \left. \left. - \mu_{-i}(\tilde{a}_{-i}, y_{-i} | \tilde{a}_i, y_i) f(\tilde{a}_i, y_i | \tilde{a}^T, y^T) \right| > c_0 \iota_0 \kappa_T \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right) \\
&\leq 2 \exp \left(-\frac{(c_0 \iota_0 \kappa_T)^2}{2f(\tilde{a}_i, y_i | \tilde{a}^T, y^T)} \right) \leq 2 \exp \left(-\frac{(c_0 \iota_0 \kappa_T)^2}{2T} \right).
\end{aligned}$$

It follows that,

$$\begin{aligned}
& \Pr \left(Z_i^{T*}(c_0 \iota_0 \kappa_T) \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right) \\
&\leq \sum_{(\tilde{a}, y) \in \text{supp}(\mu)} \Pr \left(\left| f(\tilde{a}, y | \tilde{a}^T, y^T) - \mu_{-i}(\tilde{a}_{-i}, y_{-i} | \tilde{a}_i, y_i) f(\tilde{a}_i, y_i | \tilde{a}^T, y^T) \right| > c_0 \iota_0 \kappa_T \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right) \\
&\leq 2 |\text{supp}(\mu)| \exp \left(-\frac{(c_0 \iota_0 \kappa_T)^2}{2T} \right).
\end{aligned}$$

Let

$$c_2 = \max \{ |w_{i,a}^*(\tilde{a}, y)| \mid i \in N, \tilde{a} \in A, y \in Y \}.$$

Since $R_i^+(\tilde{a}^T, \hat{y}^T) \leq c_2 T$ and $\kappa_T = T^{2/3}$,

$$\begin{aligned}
E \left[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right] &= E \left[\sum_{j=1}^n R_j^+(\tilde{a}^T, \hat{y}^T, \kappa_T) I_i(\tilde{a}^T, \hat{y}^T, \kappa_T) \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right] \\
&\leq n c_2 T \Pr \left(Z_i^{T*}(c_0 \iota_0 \kappa_T) \left| \boldsymbol{\sigma}^{T*}, \tilde{a}_i^T, y_i^T \right. \right) \\
&< 2 c_2 |\text{supp}(\mu)| n T \exp \left(-\frac{(c_0 \iota_0)^2}{2} T^{1/3} \right).
\end{aligned}$$

Let $c = 2c_2 |\text{supp}(\mu)| n$. This completes the proof of the lemma.

E Proof of Theorem 2

We say that a pure action strategy α_i is equally informative as α_i^* if for each \tilde{a}_i that may be recommended with strictly positive probability under η , there is a one-to-one mapping $\chi_{\tilde{a}_i} : Y_i \rightarrow Y_i$ such that for any $(\tilde{a}_{-i}, y_{-i}) \in A_{-i} \times Y_{-i}$,

$$p(y_i, y_{-i} | \alpha_i(\tilde{a}_i), \tilde{a}_{-i}) = p(\chi_{\tilde{a}_i}(y_i), y_{-i} | \tilde{a}).$$

We say that a pure stage-game strategy (α_i, ρ_i) is a duplicate for (α_i^*, ρ_i^*) if α_i is as informative as α_i^* and $\rho_i(\tilde{a}_i, \alpha_i(\tilde{a}_i), \cdot) = \chi_{\tilde{a}_i}$.

Note that if some (α_i, ρ_i) is not a duplicate for (α_i^*, ρ_i^*) , then either it is detectable or α_i is strictly more informative than α_i^* . The number of pure stage-game strategies is finite. Since w_i^* is almost strict and η satisfies the no free information condition, there exists $\Delta_0 > 0$ such that for all non-duplicate (α_i, ρ_i) ,

$$v_i(\sigma^*; w_i^*) - v_i(\sigma_{-i}^*, \alpha_i, \rho_i; w_i^*) > \Delta_0. \quad (28)$$

Because η is enforceable, any duplicate action strategy must generate a lower stage-game payoff for player i than α_i^* . Player i , therefore, will receive a higher payoff if he replaces any duplicate action strategy α_i in some period t with α_i^* and then, in the reporting stage, reports the period- t signal truthfully. Hence, to prove Theorem 2, it suffices to show that any deviation to a non-duplicate strategy will make a player strictly worse off.

If σ_i^T deviates from σ_i^{T*} , there must be a first time a deviation occurs. There are two types of first-time non-duplicate deviations. First, a player may choose an action that is not equally informative as α_i^* after some history. Alternatively, the player may follow the recommendations in all T periods but lie about the signal of a particular period at the end.

We first consider the first type of deviations. Suppose σ^T first prescribes a non-equally-informative action in period t after $h_i \in H_i^*$. Let $v_i^T(\sigma^T; w_i^T, h_i)$ denote player i 's expected discounted payoff conditional σ^T and h_i . Recall that w_i^{T**} is the truncated contract with side bets in (18) and w_i^{T*} is T -period version of w_i^* .

By (19), we can write

$$v_i^T(\sigma_i^T, \sigma_{-i}^{T*}; w_i^{T**}, h_i) = \frac{1 - \delta}{1 - \delta^T} (V_i(\sigma_i^T; h_i) - E[w_{i,a}^{T*}(\tilde{a}^{T'}, \hat{y}^{T'}, \kappa_T) | \sigma^{T*}] - \kappa_T - E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma_{-i}^{T*}, \sigma_i^T, h_i])$$

where

$$V_i(\sigma_i^T; h_i) \equiv E \left[\sum_{s=1}^T \delta^{s-1} (g_i(a(s)) + w_i^*(\tilde{a}(s), \hat{y}(s))) \middle| \sigma_{-i}^{T*}, \sigma_i^T, h_i \right]$$

denote player i 's discounted payoff conditional on h_i under w_i^{T*} .

It follows that

$$\begin{aligned}
& v_i^T(\sigma^{T*}; w_i^{T**}, h_i) - v_i^T(\sigma_{-i}^{T*}, \sigma_i^T; w_i^{T**}, h_i) \\
&= \frac{1-\delta}{1-\delta^T} (V_i(\sigma_i^{T*}; h_i) - V_i(\sigma_i^T; h_i) \\
&\quad - E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma^{T*}, h_i] + E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma_{-i}^{T*}, \sigma_i^T, h_i]) \\
&\geq \frac{1-\delta}{1-\delta^T} (V_i(\sigma_i^{T*}; h_i) - V_i(\sigma_i^T; h_i) - E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma^{T*}, h_i]).
\end{aligned}$$

The last inequality follows from the fact that ϕ is always positive.

Since σ^T first prescribes a non-equally-informative action in period t , it will lower player i 's payoff (including the stage-game payment) by Δ_0 in that period. This, together with the fact that under w_i^{T*} the stage-game payoff plus payment is maximized by σ_i^{T*} in each period $s \neq t$, implies that

$$\begin{aligned}
& V_i(\sigma_i^{T*}; h_i) - V_i(\alpha_i^T, \rho_i^T; h_i) \\
&\geq \delta^{t-1} (E[g_i(a(t)) + w_i^*(\tilde{a}(t), \hat{y}(t)) | \sigma^{T*}, h_i] - E[g_i(a(t)) + w_i^*(\tilde{a}(t), \hat{y}(t)) | \sigma_{-i}^{T*}, \sigma_i^T, h_i]) \\
&\geq \delta^{t-1} \Delta_0.
\end{aligned}$$

By Lemma 4, we can choose T_0 large enough such that for all $T \geq T_0$ and $\delta \geq 1 - T^{-2}$,

$$E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa_T) | \sigma^{T*}, h_i] \leq \delta^{T-1} \Delta_0.$$

This proves that any σ_i^T first prescribes a non-equally-informative action is not optimal. The argument for following the recommendations but misreporting the signals is similar.

Finally, by (20), the per-period efficiency loss is

$$\begin{aligned}
W(\eta, T, \delta, w^{T**}) &= \frac{1-\delta}{1-\delta^T} \sum_{i=1}^n -E[w_i^{T**} | \sigma^{T*}] \\
&\leq L(\eta, w^*) + \frac{1-\delta}{1-\delta^T} \sum_{i=1}^n (E[\phi_i(\tilde{a}^T, \hat{y}^T, T^{2/3}) | \sigma^{T*}] + T^{2/3}).
\end{aligned}$$

By Lemma 4, when T is sufficiently large

$$\frac{1-\delta}{1-\delta^T} \sum_{i=1}^n (E[\phi_i(\tilde{a}^T, \hat{y}^T, T^{2/3}) | \sigma^{T*}] + T^{2/3}) \leq \frac{\varepsilon}{2}.$$

F Proof of Theorem 3B

Let

$$\bar{L} \equiv \begin{cases} \sup_{\sigma \in Q(\eta)} \frac{\sum_{i=1}^n d(\sigma_i)}{l(\sigma_i) - 1}, & \text{if } Q(\eta) \neq \emptyset; \\ 0, & \text{otherwise.} \end{cases}$$

To prove Theorem 3B, it remains to show that

$$L^*(\eta) \leq \bar{L}. \quad (29)$$

By definition, a contract w enforces η with $L(\eta, w) \leq \bar{L}$ if and only if

$$\sum_{(\tilde{a}, \tilde{y}) \in A \times Y} [\pi^{\alpha_i, \rho_i}(\tilde{a}, \tilde{y}) - \mu(\tilde{a}, \tilde{y})] w_i(\tilde{a}, \tilde{y}) \leq -d_i(\alpha_i, \rho_i) \quad \forall (\alpha_i, \rho_i, i); \quad (30)$$

$$\sum_{i=1}^n \sum_{(\tilde{a}, \tilde{y}) \in A \times Y} [-\mu(\tilde{a}, \tilde{y}) + \mu(\tilde{a}, \tilde{y} | \omega)] w_i(\tilde{a}, \tilde{y}) \leq \bar{L} \quad \forall \omega \in P. \quad (31)$$

By the theorem of alternatives (see, e.g., Proposition 5.1.2 of Bertsekas, 2009), (30) and (31) does not have a solution in w if and only if there exist $\{\lambda_i(\alpha_i, \rho_i) \geq 0 \mid (\alpha_i, \rho_i, i)\}$ and $\{v(\omega) \geq 0 \mid \omega \in P\}$ such that

$$\sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) [\pi^{\alpha_i, \rho_i}(\cdot) - \mu(\cdot)] + \sum_{\omega \in P} v(\omega) [-\mu(\cdot) + \mu(\cdot | \omega)] = 0 \quad \forall i \quad (32)$$

$$\sum_{i=1}^n \sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) d_i(\alpha_i, \rho_i) - \sum_{\omega \in P} v(\omega) \bar{L} > 0. \quad (33)$$

Suppose (32) and (33) hold. From (33), $\bar{\lambda} \equiv \max_{i \in N} \sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) > 0$. We can, therefore, define a mixed strategy σ_i for each player i such that, for all (α_i, ρ_i, i) ,

$$\sigma_i(\alpha_i, \rho_i) \equiv \begin{cases} \frac{\lambda_i(\alpha_i, \rho_i)}{\bar{\lambda}}, & \text{if } (\alpha_i, \rho_i) \neq (\alpha_i^*, \rho_i^*); \\ 1 - \sum_{(\alpha_i, \rho_i) \neq (\alpha_i^*, \rho_i^*)} \frac{\lambda_i(\alpha_i, \rho_i)}{\bar{\lambda}}, & \text{otherwise.} \end{cases}$$

Using the definition of σ_i , we can rewrite (32) and (33) as

$$\bar{\lambda} [\pi^{\sigma_i}(\cdot) - \mu(\cdot)] + \sum_{\omega \in P} v(\omega) [-\mu(\cdot) + \mu(\cdot | \omega)] = 0 \quad \text{for each } i \quad (34)$$

$$\sum_{i=1}^n \bar{\lambda} d_i(\sigma_i) - \sum_{\omega \in P} v(\omega) \bar{L} > 0. \quad (35)$$

Fix a contract w . Multiplying each (34) by $w_i(\cdot)$, then summing over all i and all $(\tilde{a}, \hat{y}) \in A \times Y$, and adding (35), we have

$$\begin{aligned} \sum_{i=1}^n \bar{\lambda} \left(\sum_{(\tilde{a}, \hat{y}) \in A \times Y} [\pi^{\sigma_i}(\tilde{a}, \hat{y}) - \mu(\tilde{a}, \hat{y})] w_i(\tilde{a}, \hat{y}) + d_i(\sigma_i) \right) \\ + \sum_{\omega \in P} v(\omega) \left(\sum_{(\tilde{a}, \hat{y}) \in A \times Y} [-\mu(\tilde{a}, \hat{y}) + \mu(\tilde{a}, \hat{y}|\omega)] w_i(\tilde{a}, \hat{y}) - \bar{L} \right) > 0. \end{aligned}$$

This means that if η cannot be enforced by any w with $L(\eta, w) \leq \bar{L}$, then there must exist σ such that, for any w with $L(\eta, w) \leq \bar{L}$,

$$v_i(\sigma_i, \sigma_{-i}^*; w_i) > v_i(\sigma^*; w_i)$$

for some player i .

We prove (29) by showing that for all $\sigma \in \Sigma$, there exists a contract w such that $v_i(\sigma_i, \sigma_{-i}^*; w_i) - v_i(\sigma^*; w_i) \leq 0$ for all i and $L(\eta, w) \leq \bar{L}$. By Theorem 4(i) of Rahman and Obara (2010), if σ is either unprofitable or attributable, then it can be deterred by a contract with total transfer summing to zero. It remains to consider σ such that $\pi^{\sigma_1} = \dots = \pi^{\sigma_n}$ and $\sum_{i=1}^n d_i(\sigma_i) > 0$. Since η is enforceable, $\pi^{\sigma_i} \neq \mu$.

Case 1. If there exists (\tilde{a}, \hat{y}) such that $\pi^{\sigma_i}(\tilde{a}, \hat{y}) > 0$ and $\mu(\tilde{a}, \hat{y}) = 0$, then σ can be deterred by a contract w that punishes every player severely whenever (\tilde{a}, \hat{y}) occurs. Clearly, $L(\eta, w) = 0$.

Case 2. Suppose that $\pi^{\sigma_i}(\cdot|\omega) \neq \mu(\cdot|\omega)$ for some $\omega \in P$. Then $\pi^{\sigma_i}(\tilde{a}, \hat{y}|\omega) > \mu(\tilde{a}, \hat{y}|\omega)$ for some $(\tilde{a}, \hat{y}) \in \omega$. We define a contract w by letting, for each i ,

$$w_i(\tilde{a}', \hat{y}') = \begin{cases} -c, & \text{if } (\tilde{a}', \hat{y}') = (\tilde{a}, \hat{y}); \\ -c \cdot \mu(\tilde{a}, \hat{y}|\omega), & \text{if } (\tilde{a}', \hat{y}') \notin \omega; \\ 0, & \text{otherwise.} \end{cases}$$

Then $E[w_i(\tilde{a}', \hat{y}')|\sigma^*, \omega'] = -c \cdot \mu(\tilde{a}, \hat{y}|\omega)$ for all $\omega' \in P$. Hence $L(\eta, w) = 0$. Moreover,

$$v_i(\sigma_i, \sigma_{-i}^*; w_i) - v_i(\sigma^*; w_i) = -c \cdot (\pi^{\sigma_i}(\tilde{a}, \hat{y}|\omega) - \mu(\tilde{a}, \hat{y}|\omega)) \pi^{\sigma_i}(\omega) + d(\sigma_i) \leq 0,$$

when c is large enough.

Case 3. Suppose that $\sigma \in Q(\eta)$. Let ω solve $\max_{\omega' \in P} \frac{\pi^{\sigma_i(\omega')}}{\mu(\omega')}$. We define a contract w by letting, for each i ,

$$w_i(\tilde{a}, \hat{y}) = \begin{cases} -\frac{d(\sigma_i)}{\pi^{\sigma_i(\omega)} - \mu(\omega)}, & \text{if } (\tilde{a}, \hat{y}) \in \omega; \\ 0, & \text{otherwise.} \end{cases}$$

Then $L(\eta, w) = \frac{\sum_{i=1}^n d(\sigma_i)}{\pi^{\sigma_1(\omega)} - \mu(\omega)} \mu(\omega) = \frac{\sum_{i=1}^n d(\sigma_i)}{l(\sigma_1) - 1}$ and

$$v_i(\sigma_i, \sigma_{-i}^*; w_i) - v_i(\sigma^*; w_i) = -\frac{d(\sigma_i)}{\pi^{\sigma_i(\omega)} - \mu(\omega)} \cdot (\pi^{\sigma_i(\omega)} - \mu(\omega)) + d(\sigma_i) = 0.$$

References

- Abreu, Dilip, Paul Milgrom, and David Pearce, 1991, Information and timing in repeated partnerships, *Econometrica* 59, 1713–1733.
- Aumann, Robert J., 1976, Agreeing to disagree, *The Annals of Statistics* 4, 1236–1239.
- Bertsekas, Dimitri P., 2009, *Convex Optimization Theory* (Athena Scientific).
- Blackwell, David, 1953, Equivalent comparisons of experiments, *The Annals of Mathematical Statistics* 24, 265–272.
- Chan, Jimmy H., and Wenzhang Zhang, 2016, Approximate efficiency in repeated games with side-payments and correlated signals, *Theoretical Economics* 11, 53–87.
- Compte, Olivier, 1998, Communication in repeated games with imperfect private monitoring, *Econometrica* 66, 597–626.
- Fong, Kyna, Olivier Gossner, Johannes Hörner, and Yuliy Sannikov, 2011, Efficiency in a repeated prisoner’s dilemma with imperfect private monitoring, Working paper, Princeton University.
- Fuchs, William, 2007, Contracting with repeated moral hazard and private evaluations, *American Economic Review* 97, 1432–1448.
- Fudenberg, Drew, David Levine, and Eric Maskin, 1994, The folk theorem with imperfect public information, *Econometrica* 62, 997–1039.

- Geanakoplos, John, 1994, Common knowledge, in R.J. Aumann, and S. Hart, ed.: *Handbook of Game Theory with Economic Applications*, vol. 2 . chap. 40, pp. 1437–1496 (Elsevier) 1 edn.
- Gerardi, Dino, 2004, Unmediated communication in games with complete and incomplete information, *Journal of Economic Theory* 114, 104 – 131.
- Hoeffding, Wassily, 1963, Probability inequalities for sums of bounded random variables, *Journal of the American Statistical Association* 58, 13–30.
- Holmstrom, Bengt, 1982, Moral hazard in teams, *The Bell Journal of Economics* 13, 324–340.
- Kandori, Michihiro, and Hitoshi Matsushima, 1998, Private observation, communication and collusion, *Econometrica* 66, 627–652.
- Ke, Rongzhu, Jin Li, and Michael Powell, 2018, Managing careers in organizations, *Journal of Labor Economics* 36, 197 – 252.
- MacLeod, W. Bentley, 2003, Optimal contracting with subjective evaluation, *The American Economic Review* 93, 216–240.
- Matsushima, Hitoshi, 2004, Repeated games with private monitoring: Two players, *Econometrica* 72, 823–852.
- Obara, Ichiro, 2009, Folk theorem with communication, *Journal of Economic Theory* 144, 120–134.
- Osborne, Martin J., and Ariel Rubinstein, 1994, *A course in game theory* (MIT press).
- Radner, Roy, Roger Myerson, and Eric Maskin, 1986, An example of a repeated partnership game with discounting and with uniformly inefficient equilibria, *The Review of Economic Studies* 53, 59–69.
- Rahman, David, 2012, But who will monitor the monitor?, *American Economic Review* 102, 2767–2797.

- , 2014, The dilemma of the cypress and the oak tree, mimeo, University of Minnesota.
- , and Ichiro Obara, 2010, Mediated partnerships, *Econometrica* 78, 285–308.
- Sugaya, Takuo, 2017a, The characterization of the limit communication equilibrium payoff set with general monitoring: Observable realized own payoff case, Working paper, Stanford University.
- , 2017b, The limit communication equilibrium payoff set with general monitoring, Working paper, Stanford University.
- Zhao, Rui R., 2008, All-or-nothing monitoring, *The American Economic Review* 98, 1619–1628.