

Self-Evident Events and the Value of Linking*

Jimmy Chan[†] Wenzhang Zhang[‡]

June 20, 2021

Abstract

We study a T -period contracting game between a group of players without access to external financing. We show that the long-term efficiency loss is bounded from below by the short-term efficiency loss that becomes self-evident at the end of a period. When T is large, the efficiency bound can be approached by a contract that involves side payments between players. 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 can 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 write a binding incentive contract among themselves so long as no budget deficit is allowed. For if one

*We thank Dilip Abreu, Erik Eyster, Sambuddha Ghosh, Satoru Takahashi, and seminar participants at various seminars and conferences for helpful comments and Pak Hang Tam for research assistance. 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

player's punishment is transferred to another player as a reward, then the latter may deliberately choose an action to get the first player punished. Hence, motivating every player may call for all players to be punished simultaneously. For example, a cartel may have to resort to costly price wars to sustain collusion. Efficiency can be restored if the players can contract with a third party who can provide external finances 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."

In this paper, we argue that the incentive problem of having side-payments between players is greatly reduced when players interact repeatedly and, as a consequence, access to external finances may not be as important as it seems in long-term partnerships. We formalize our idea in a T -period contracting game between a group of players. In each period each player chooses a private action and then observes a private signal about the chosen actions. At the end of period T , the players report their signals publicly. A contract maps the players' reports into a payment to each player. Because there is no external financing, the total payment must be non-positive and any strictly negative total payment must be destroyed.

We characterize the minimum efficiency loss in enforcing a certain stage-game action profile throughout the contract. A central concept in our analysis is the notion of *self-evident* events, which was introduced by Aumann (1976) to analyze beliefs in incomplete-information games. We apply the concept to describe information that becomes common knowledge after the players observe their private signals at the end of a period. For example, if a signal is public, then its realization is self-evident because every player observes the signal and knows that every player observes the signal and so on. Using the notion of self-evident events, we first establish an upper bound on the long-term efficiency of a partnership. We then show that when the length of the contract is sufficiently long, the efficiency bound can be approached by a simple contract whereby each player is penalized or rewarded depending on whether his average performance is below or above a performance standard.

Since there is no external financing, one player's over-performance bonus has to be paid by another player. In a short-term partnership, this may interfere with the incen-

tives of the paying player. In a long-term relationship, linking incentives across periods makes this incentive problem much less severe. This is because in equilibrium a non-deviating player is exceedingly unlikely to over-perform consistently. Unless a player's performance is self-evident (e.g., when it is public), there will always be some player who would believe that the chance of consistent over-performance by the receiving player is so small that having to pay the bonus would cause only a negligible effect on his incentives. Real-life cartels have often relied on side payments to enforce collusion (Harrington and Skrzypacz, 2011). Our results show that, under fairly weak conditions, such payments could be made incentive compatible.

Whether long-term partnerships can be efficiently run is a long-standing question in the theory of repeated games. In a seminal paper, Abreu, Milgrom, and Pearce (1991) show that in a repeated Prisoners' Dilemma, if the players observe a noisy public signal immediately at the end of each period, then the equilibrium outcome must be inefficient. However, if the signals are observed with a lag, then the players can reduce the efficiency loss by linking incentives across periods. Subsequent research has applied the insights of Abreu, Milgrom, and Pearce (1991) to repeated games with private monitoring. Following Abreu, Milgrom, and Pearce (1991), the literature has mainly focused on two polar cases: imperfect public monitoring where linking incentives has no value (Abreu, Milgrom, and Pearce, 1991; Sannikov and Skrzypacz, 2007) and conditionally independent monitoring where linking improves efficiency (Compte, 1998; Obara, 2009; Chan and Zhang, 2016). Using the notion of self-evident events, we generalize the results of Abreu, Milgrom, and Pearce (1991) to general stage games. Our results provide a unifying framework that applies not only to the two polar cases but also to the in-between cases where players observe correlated private signals.

In the setting of one-shot interaction, Rahman (2012) and Rahman and Obara (2010) characterize the action profiles that can be enforced without efficiency loss. They show that an action profile can be enforced without efficiency loss if every deviating strategy is attributable. We extend this characterization to the T -period contracting game. We show that in order for efficiency to be attainable in long-term relationships, every deviating strategy must be *either* attributable *or* detectable within a self-evident event. When efficiency loss is inevitable, we provide a similar characterization for the efficiency loss.

The rest of the paper is organized as follows. The next section uses a repeated Pris-

oners' Dilemma example to bring out the main ideas of this paper. Section 3 introduces the general model. Our main results are introduced in Sections 4 and 5. Section 6 discusses the related literature in details. Section 7 concludes.

2 Example

In this section we illustrate how side-payments between players can improve efficiency through 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) . 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 $\rho \in [0, \bar{\rho}]$ for some $\bar{\rho} > 1$.¹ When $\rho = 0$, the signals are perfectly correlated. In this case, the players are effectively observing a public signal. When $\rho = 1$, the signals are conditionally independent and a player cannot learn about the other player's signal from his own. When $\rho \neq 0, 1$, the signals are imperfectly correlated. The correlation is positive when $\rho \in (0, 1)$ and negative when $\rho \in (1, \bar{\rho})$.

The players hire a principal 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. The principal

¹Assume that $\bar{\rho}$ is not too large such that the signal distributions are well defined.

	<i>H</i>	<i>L</i>
<i>H</i>	$p - \rho p(1-p)$	$\rho p(1-p)$
<i>L</i>	$\rho p(1-p)$	$(1-p)(1-\rho p)$

Signal distribution under (C, C)

	<i>H</i>	<i>L</i>
<i>H</i>	$q - \rho q(1-q)$	$\rho q(1-q)$
<i>L</i>	$\rho q(1-q)$	$(1-q)(1-\rho q)$

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

Table 2: Signal distributions.

can neither pay the players with outside resources nor extract resources from them. He can, however, commit to destroying resources. A T -period contract $w^T = (w_1^T, w_2^T)$ is, therefore, a function that maps the players' reports to a payment to each player, subject to the constraint that the total payment be non-positive.² 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.

Since the total payment must be negative, providing incentives is costly. Consider the one-period case. Let $w = (w_1, w_2)$ denote a stage-game contract. With a slight abuse of notation, let $w_i(H)$ and $w_i(L)$ denote player i 's payment when player $-i$ reports that his signal is H and L , respectively.³ 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

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

²The 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.

³Player $-i$ is the player who is not i .

⁴In the single-period case, there is no efficiency gain by making w_i a function of player i 's own report, and the optimal contract does not depend on ρ .

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




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

captures a fundamental issue in team moral-hazard problems: since both (C,D) and (D,C) lead to the same signal distribution, it is not possible to tell which player has deviated. As a result, there is no budget-balanced contract that can enforce (C,C) . In the terminology of Rahman and Obara (2010), the deviation to D is *non-attributable*. The negative payment $d/(p-q)$ must be destroyed and cannot be transferred to the other player. If player 1 simply pays 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. 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: $\rho = 0$ and $\rho = 1$. Our contribution is to extend the analysis to $\rho \neq 0, 1$. Before proceeding to our results, we first briefly recount the two polar cases.

2.1 Case 1: $\rho = 0$.

We will use the two-period case to illustrate the result of Abreu, Milgrom, and Pearce (1991) that linking has no value when the signal is public or perfectly correlated. We derive a lower bound on the efficiency loss of a relaxed contracting problem in which the players must report their signals truthfully. Since any contract that enforces (C,C) in the original contracting problem must also enforce (C,C) in the relaxed problem, the lower bound applies to the original contracting problem as well.

Assume that the players must report truthfully. 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 .⁵ See Figure 2.




Figure 2: The two-period case.

Suppose w^2 is a two-period contract that enforces (C, C) in both periods. 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) . Since both players observe H , the values of $w_i(L(1), H(2))$ and $w_i(L(1), L(2))$ do not affect the players' incentives in the second period after H .⁶ As a result, enforcing (C, C) in the second period after H is the same as enforcing (C, C) in a single-period game. Hence,

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

The incentive-compatibility constraint in period 1 requires that

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

Combining (1) and (2), we have

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

and

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

Thus, the two-period per-player efficiency loss must be greater than twice the one-period efficiency loss. In this case, the players cannot do better than enforcing (C, C)

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

⁶Likewise, the values of $w_i(H(1), H(2))$ and $w_i(H(1), L(2))$ have no effect on the players' incentives in the second period after L .

period by period with two stage-game contracts. The argument can be extended to $T > 2$ by induction. The following observation is the key lesson of this example.

Observation 1. *Incentives following different realizations of a public signal must be separately provided. As a result, the difference between the total payments following two different realizations of a public signal in a short-term contract imposes a lower bound on the long-term efficiency loss.*

In Section 4, we apply this insight to derive a lower bound on the long-term efficiency loss for general stage games (Theorem 1). We then prove that the bound is tight under a fairly weak condition (Theorem 2). It follows that Case 1 is representative of all inefficient long-term partnerships (Theorem 3).

2.2 Case 2: $\rho = 1$.

Since the players may observe different signals, let us use H_i and L_i to denote the H and L signals observed by player i . Let $y_i^T = (y_i(1), \dots, y_i(T))$ denote the T -period profile of player i 's signal. Write y^T for (y_1^T, y_2^T) . For any y^T , denote the number of L signals that player $-i$ observes in y^T by $f(L_{-i}|y^T)$.

Consider a “linear” contract \tilde{w}^T that enforces (C, C) period by period. For $i = 1, 2$

$$\tilde{w}_i^T(\hat{y}^T) = - (f(L_{-i}|\hat{y}^T) - T(1 - p - v)) (d/(p - q) + \varepsilon), \quad (3)$$

where $\hat{y}^T = (\hat{y}_1^T, \hat{y}_2^T)$ denotes the players' reports, and ε and v are small positive constants. The contract punishes player i by $(d/(p - q) + \varepsilon)$ for every L signal that player $-i$ reports. Since player i 's payment does not depend on his own reports, it is optimal to report truthfully. Since $\varepsilon > 0$, the contract strictly enforces (C, C) . However, it 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

$$\hat{w}_i^T(\hat{y}^T) = - \max(f(L_{-i}|\hat{y}^T) - T(1 - p - v), 0) (d/(p - q) + \varepsilon). \quad (4)$$

Under \hat{w}_i^T , player i pays a penalty of $(d/(p - q) + \varepsilon)$ for every L signal that player $-i$ reports that is in excess of $T(1 - p - v)$. Note that \hat{w}_i^T is no longer linear in $f(L_{-i}|\hat{y}^T)$.

Whether player i will be punished for an L signal that player $-i$ observes in a certain period now depends on how many other L signals player $-i$ observes during the contract. Player i 's incentives are thus *linked* across periods.

Compared to the linear contract \tilde{w}_i^T , the truncated contract \hat{w}_i^T does not “reward” the players when the number of L signals is below the threshold. The truncated incentive is

$$R_i(\hat{y}^T) = -\min(f(L_{-i}|\hat{y}^T) - T(1-p-v), 0)(d/(p-q) + \varepsilon). \quad (5)$$

The distortionary effect of the truncation is very small when T is large. If player i chooses C in every period (assuming that player $-i$ also chooses C), the average number of L signals that player $-i$ should observe is $T(1-p)$. In equilibrium player i expects player $-i$ to report truthfully. By the law of large numbers, the probability that player $-i$ reports observing less than $T(1-p-v)$ L signals is exceedingly small.⁷ Since the signals are conditionally independent when $\rho = 1$, player i cannot learn about y_{-i}^T from his own signals. As a result, the distortionary effect remains very small throughout the contract and is compensated by the small extra punishment ε . Hence, it is optimal for each player to choose C in every period.

The expected per-player per-period efficiency loss caused by \hat{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.

Observation 2. *When T is large, it is “normal” for a player i to observe $(1-p)T$ L signals during the contract. Instead of punishing player $-i$ for every L_i , it suffices to punish player $-i$ when his long-term performance is worse than the mean.*

2.3 Case 3: $\rho \neq 0, 1$.

Case 2 is essentially the argument of Abreu, Milgrom, and Pearce (1991). Our approach here follows more closely the review strategies of Rubinstein (1979), Rubinstein and Yaari (1983), and Fong, Gossner, Hörner, and Sannikov (2011). Note that the truncated contract \hat{w}^T does not enforce (C, C) when $\rho \neq 1$. For example, when $\rho \in (0, 1)$, a

⁷If player i chooses D in some periods, the probability will be even lower.

player i who has observed very few L signals in the early periods will infer from his own signals that the probability that $f(L_{-i}|y^T)$ is less than $(1-p-v)T$ is non-trivial and, hence, may deviate in the remaining periods. This inference problem becomes more severe as the correlation between the two players' signals increases. One may therefore expect that linking would be less effective as ρ gets close to 0. It turns out that $\rho = 0$ is a special case. As long as $\rho \neq 0$, we can enforce (C, C) almost efficiently by supplementing \widehat{w}^T with a nexus of side payments between players.

So what is special about $\rho = 0$? When signal is public, any realization, no matter how unlikely ex ante, becomes common knowledge among the players.⁸ For example, it may be extremely unlikely that the public signal is H in every period. But if it happens, it will still be common knowledge. This property holds only when $\rho = 0$. When $\rho = 1$, each player i can learn nothing about player $-i$'s signal from his own. As a result, when T is large, each player i believes that it is extremely unlikely for the realized signals of player $-i$ to be different from the prior expectation. When $\rho \neq 0, 1$, it is no longer true that player i cannot learn about player $-i$'s signal. Yet, it remains the case that when T is large, *any* realized signal distribution that is different significantly from prior expectation must be regarded as extremely unlikely by *some* player conditional on that player's private signals.

We formally state this result as a lemma. For any signal profile y^T , let $f(y_i, y_{-i}|y^T)$ and $f(y_i|y^T)$ denote the numbers of times (y_i, y_{-i}) and y_i occur in y^T , respectively. The prior expectations of $f(y_i, y_{-i}|y^T)$ and $f(y_i|y^T)$, conditional on (C, C) being chosen in every period, are

$$E[f(y_i, y_{-i}|y^T)] = \text{prob}(y_i, y_{-i}|(C, C))T; \quad (6)$$

$$E[f(y_i|y^T)] = \text{prob}(y_i|(C, C))T. \quad (7)$$

At the end of the contract, player i observes $f(H_i|y^T)$ and $f(L_i|y^T)$. His posterior expectation about $f(y_i, y_{-i}|y^T)$ conditional on y_i^T is

$$E[f(y_i, y_{-i}|y^T) | y_i^T] = \text{prob}(y_i, y_{-i}|(C, C), y_i) f(y_i|y^T). \quad (8)$$

Let

$$Z^T(\xi) = \{y^T | \exists (y_i, y_{-i}) | f(y_i, y_{-i}|y^T) - E[f(y_i, y_{-i}|y^T)] > T\xi\}$$

⁸For our purpose, the difference between common knowledge and common belief is unimportant.

denote the set of y^T in which the occurrence of some outcome is different from the prior expectation by $T\xi$. Similarly, let

$$Z_i^T(\xi) = \{y^T | \exists (y_i, y_{-i}) | f(y_i, y_{-i} | y^T) - E[f(y_i, y_{-i} | y^T) | y_i^T] | > T\xi\}$$

denote the set of y^T in which the occurrence of some outcome is different from player i 's posterior expectation by $T\xi$. Note that the law of large numbers implies that for any $\xi > 0$, conditional on any y_i^T , it is extremely unlikely that $y^T \in Z_i^T(\xi)$ when T is large.

Lemma 1 (Posteriors Determine Prior). *In the Prisoners' Dilemma example, when $\rho \neq 0$, for any $\xi > 0$, there exists $\varsigma > 0$ such that*

$$Z^T(\xi) \subseteq Z_1^T(\varsigma) \cup Z_2^T(\varsigma). \quad (9)$$

Lemma 1 says that, when $\rho \neq 0$, if the distribution of signals in some y^T is different from the prior, it must be different from the posterior expectation of some player.

Note that if we take the realized frequencies, $f(\cdot | y^T) / T$, as a probability distribution, then the posterior distribution conditional on y_i^T is $f(y_i, y_{-i} | y^T) / f(y_i | y^T)$. By (6), (7) and (8), for any $f(y_i | y^T) > 0$, $f(y_i, y_{-i} | y^T)$ is equal to player i 's posterior expectation if and only if

$$\frac{f(y_i, y_{-i} | y^T)}{f(y_i | y^T)} = \text{prob}(y_i, y_{-i} | (C, C), y_i),$$

and $f(y_i, y_{-i} | y^T)$ is equal to the prior expectation if and only if

$$\frac{f(y_i, y_{-i} | y^T)}{T} = \text{prob}(y_i, y_{-i} | (C, C)).$$

Hence, Lemma 1 essentially says that the mapping from a full-support prior distribution to the posterior distributions it generates is continuous and one-to-one.⁹ To see that the mapping is one-to-one, note that the posterior distribution conditional on $y_1 = H_1, L_1$ pins down the relative likelihood of (y_1, H_2) and (y_1, L_2) , while the posterior distribution conditional on H_2 pins down the relative frequency of the signal pair (H_1, H_2) and (L_1, H_2) . These likelihood ratios jointly determine the prior distribution. Note that this argument applies as long as the signal distribution is connected; that is, the probability

⁹The probability distribution has full support when $\rho \neq 0$.

of either (L_1, H_2) or (H_1, L_2) is strictly positive. In Section 4, we use the concept of self-evident event to prove a general version of Lemma 1.

We now introduce a contract that enforces (C, C) when $\rho \neq 1$. Start with the contract \widehat{w}_i^T in Case 2. The set of \widehat{y}^T where player i 's incentives will be truncated under \widehat{w}_i^T is

$$B_i^T(v) = \{\widehat{y}^T | f(L_{-i} | \widehat{y}^T) < T(1 - p - v)\}.$$

If $\widehat{y}^T \in B_i^T(v)$, then either $f(H_i, L_{-i} | \widehat{y}^T)$ or $f(L_i, L_{-i} | \widehat{y}^T)$ must differ from the prior expectation. It then follows from Lemma 1 that there exists ζ such that

$$B_i^T(v) \subseteq Z_1^T(\zeta) \cup Z_2^T(\zeta).$$

Starting with the truncated contract \widehat{w}^T in Case 2, we add a side-bet contract $z^T = (z_1^T, z_2^T)$. For $i = 1, 2$,

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

where

$$I_i(\widehat{y}^T) = \begin{cases} 1 & \text{if } \widehat{y}^T \in Z_i^T(\zeta), \\ 0 & \text{otherwise.} \end{cases}$$

Under the side-bet contract, each player i acts as an ‘‘internal budget-breaker’’ that pays an over-performance bonus to the other player when $I_i = 1$. Recall that R_i represents the truncated incentives defined in (5). If player i receives R_i from the side bet, then his total incentives will be ‘‘untruncated’’. Under this side-bet contract, player i receives R_i when $\widehat{y}^T \notin Z_i^T(\zeta)$ and pays R_{-i} when $\widehat{y}^T \in Z_i^T(\zeta)$. The total payment of this side-bet contract is always negative. By (9), when $R_i > 0$, either I_i or I_{-i} must be equal to 1.¹⁰ Hence, when player i receives a strictly positive R_i (i.e., $I_i = 0$), player $-i$ must pay for it (i.e., $I_{-i} = 1$).

We first show that when T is large, it is optimal for player i to choose C in every period if he expects player $-i$ to play C in every period and report truthfully. The Hoeffding inequality (Hoeffding, 1963) implies that if player i plays C in every period and reports truthfully, conditional on *any* y_i^T the probability that $I_i = 1$ (i.e., $\widehat{y}^T \in Z_i^T(\zeta)$)

¹⁰Under the current construction, the side-bet contract is not zero-sum, as it is possible that $I_i(\widehat{y}^T) = I_{-i}(\widehat{y}^T) = 1$. This feature is not crucial. The contract can be modified (with extra notations) so that player $-i$ will pay the bonus only when player i will receive it.

converges to zero exponentially in T . As a result, in any period *during* the contract, a player i who has chosen C in all previous periods will believe that he will almost always receive R_i when $R_i > 0$ but almost never pay R_{-i} when $R_{-i} > 0$ if he continues to play C in the remaining periods. This implies that, *throughout* the contract, player i believes that if he plays C in every period, he will receive a payment that is equal to what he would receive under the linear contract, \tilde{w}^T , minus a term that converges to zero exponentially in T . By contrast, if player i deviates in any period, his payment will be *at most* equal to the payment he would receive under \tilde{w}^T (regardless of whether he mis-reports his signals or not).¹¹ Since \tilde{w}^T *strictly* enforces (C, C) , player i will still be *strictly* better off playing C in every period when T is large.

Since player i 's payment in the truncated contract \hat{w}_i^T depends only on player $-i$'s reports, player i 's report may affect his own payment only through R_{-i} , I_i , and I_{-i} . When T is large, a player i who chooses C in every period and reports truthfully will almost always receive R_i when $R_i > 0$ and almost never need to pay R_{-i} when $R_{-i} > 0$. Hence, the potential gain from lying is small and converges to zero exponentially in T .¹² To induce truthful reporting, it is sufficient to add a third component

$$e_i^T(\hat{y}^T) = \xi \sum_{t=1}^T \log(\text{prob}(\hat{y}_{-i}(t) | \hat{y}_i(t))) \quad (10)$$

to each player i 's payment.¹³ It is straightforward to verify that when (C, C) is chosen in every period and player $-i$ reports truthfully, any mis-reporting by player i will strictly reduce e_i^T .¹⁴ Because player i 's incentive to lie from the side bets is weak, the constant ξ in (10) can be made very small.¹⁵

Thus, when T is large, (C, C) can be enforced by a contract

$$w_i^{T*}(\hat{y}^T) = \hat{w}_i^T(\hat{y}^T) + z_i^T(\hat{y}^T) + e_i^T(\hat{y}^T).$$

¹¹The payment will be the same as that under \tilde{w}^T if player i receives R_i and does not need to pay R_{-i} . This is the best case scenario for player i . If the deviations reduce his chance of receiving R_i or increase his chance of paying R_{-i} , his payment will be strictly lower.

¹²The best player i can achieve through lying is to increase the probability of receiving R_i when $R_i > 0$ from almost one to one and reduce the probability of paying R_{-i} when $R_{-i} > 0$ from almost zero to zero.

¹³Note that in the proof of Theorem 2, we start with a stage-game contract that strictly induces truthful reporting. Hence, this step is not needed there.

¹⁴The component e_i^T is an example of a scoring rule that induces a player to reveal his posterior belief.

¹⁵Since ξ is small, it will still be optimal for player i to choose C in every period.

We have already shown that the efficiency loss of the truncated contract is small. Since ξ is very small, the efficiency loss due to e_i^T is very small. The efficiency loss due to the side bets is also very small because, ex ante, R_1 and R_2 are almost always equal to zero when T is large.

Observation 3. *When $\rho \neq 0$, efficiency can be enhanced by players exchanging over-performance bonuses. In a short-term contract, making one player pay another is likely to perversely affect the incentives of the paying player. Here, Lemma 1 implies that the bonus can be assigned to be paid by a player who believes in equilibrium that he will almost never have to pay. As a result, the distortion is minimal.*

Theorem 2 uses the argument in this section to show that the bound established in Theorem 1 is tight.

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 $p(\cdot|a) \in \Delta(Y)$. At the end of the stage game, each player i observes (\tilde{a}_i, y_i) . We will refer to the recommendation and signal profile (\tilde{a}, y) as the *outcome* of the stage game.

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. It is therefore possible that players observe both private and public signals. For example, each player i may observe a two-dimensional signal $y_i = (y_i^1, y_i^2)$ with $y_i^1 \in Y_i^1$ and $y_i^2 \in Y_i^2$. The first component is public if $Y_1^1 = \dots = Y_n^1$ and for all $a \in A$, $p(y|a) > 0$ only if $y_1^1 = \dots = y_n^1$. The game becomes one of public monitoring if the second component is a null signal.

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. 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. 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.

Since N , A , and g are fixed in our analysis, we denote the T -period game by $\Gamma(\eta, T, \delta, w^T)$. For player i , a pure strategy 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

report $\hat{y}_i^T \in Y_i^T$.¹⁶ 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. A strategy is called obedient if it follows recommendations in every period and reports signals truthfully. Let $\sigma_i^{T*} = (\alpha_i^{T*}, \rho_i^{T*})$ denote the obedient strategy of player i and $\sigma^{T*} = (\sigma_1^{T*}, \dots, \sigma_n^{T*})$. By the revelation principle, we can focus on contracts that enforce the obedient strategies.

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 .

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

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

¹⁶As usual, a^0 denotes the null history \emptyset and A^0 denotes the set whose only element is a^0 . Similar notations apply for signals.

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

An important special case is when $T = 1$. Write σ for σ^1 and w for w^1 . Let μ denote the distribution over (\tilde{a}, y) induced by η and p . For all $(\tilde{a}, y) \in A \times Y$,

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

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 from Rahman (2012) provides a necessary and sufficient condition for stage-game enforceability.

Lemma 2 (Theorem 1, Rahman, 2012). *An action profile η is enforceable for one period if and only if, for all $i \in N$ and all undetectable σ_i ,*

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

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, Lemma 2 is also necessary and sufficient for the enforceability of η for $T > 1$.¹⁷

Since δ does not matter when $T = 1$, we use $W(\eta, w)$ to denote the efficiency loss of enforcing η with w , and $W^*(\eta)$ to denote the minimum efficiency loss of enforcing η in a stage game. Since the principal can always choose to enforce η by means of a series of stage-game contract,

$$W^*(\eta, T, \delta) \leq W^*(\eta).$$

We say that linking is valuable if the above inequality is strict.

Before we proceed, a couple of comments are in order.

¹⁷Same for strict enforceability.

(1) The assumption that the correlating device directly recommends η is without loss of generality. Consider a correlating device that sends each player a private message and a contract that maps each message and signal-report profile to payments to the players. If this correlating device and contract induces a Nash equilibrium in which η is the equilibrium action profile, then, by the revelation principle, there exists in our set-up a contract w that enforces η with the same efficiency loss as the original contract.¹⁸

(2) 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.

3.3 Self-evident events

As we saw in Section 2, the long-term efficiency of a partnership depends critically on what the players know at the end of each period. Players' beliefs are more complicated in the general model. At the end of each period, each player i observes (\tilde{a}_i, y_i) and forms beliefs about the full outcome profile (\tilde{a}, y) . Conditional on η being played in every period, the outcome (\tilde{a}, y) is distributed identically and independently according to μ in every period. Write $\text{supp}(\mu)$ for the support of μ . Let P_i denote player i 's information partition of $\text{supp}(\mu)$. The element of P_i that contains (\tilde{a}, y) is denoted by $P_i(\tilde{a}, y)$. Conditional on (\tilde{a}_i, y_i) , player i believes that (\tilde{a}, y) belongs to $P_i(\tilde{a}, y)$; that is,

$$P_i(\tilde{a}, y) = \{(\tilde{a}', y') \in \text{supp}(\mu) : (\tilde{a}'_i, y'_i) = (\tilde{a}_i, y_i)\}.$$

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

¹⁸See Proposition 47.1 of Osborne and Rubinstein (1994).

The vector (P_1, \dots, P_n) describes the players' knowledge structure when η is chosen. A subset E of $\text{supp}(\mu)$ is called an event. What a player knows at a certain outcome depends on what he observes at the outcome. Player i "knows" event E at (\tilde{a}, y) if

$$P_i(\tilde{a}, y) \subseteq E. \quad (11)$$

That player i knows E is itself an event that consists of all (\tilde{a}, y) where (11) is true. Thus, we can talk about player j knowing that player i knows E . An event E is common belief among the players at (\tilde{a}, y) if 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 $(\tilde{a}, y) \in E$.

Self-evident events are closely related to (but more general than) public signals. Every realization of a public signal is self-evident conditional on *any* stage-game strategy.¹⁹ But, 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 conditional on another. In the following, when we say that an event is self-evident, it is always with respect to the equilibrium action profile η .

A self-evident event is irreducible if none of its proper subsets 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 Section 2, we show that Lemma 1 holds if and only if $\rho \neq 0$. The crucial difference between $\rho = 0$ and $\rho \neq 0$ is that P contains two elements: $\{(H_1, H_2)\}$ and $\{(L_1, L_2)\}$ when $\rho = 0$, and only one: $\{(H_1, H_2), (H_1, L_2), (L_1, H_2), (L_1, L_2)\}$ when $\rho \neq 0$.

4 Main Results

In Section 2 we saw a connection between long-term efficiency and short-term incentives that vary across self-evident events. In a general stage game, incentives may vary across self-evident events, as well as within self-evident events.

Recall that the efficiency loss of enforcing η in a stage game with w is

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

¹⁹To be precise, the set of outcomes consistent with a specific realization of a public signal is self-evident.

Write ω for a typical element of P . Let $\mathcal{W}(\eta)$ denote the set of stage-game contracts that enforce η . 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 payment 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. We say that the total incentives vary across ω and ω_{\max} if

$$\sum_{i=1}^n (E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega_{\max}] - E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega]) > 0.$$

The total incentives that vary across self-evident events under contract w are defined as

$$\begin{aligned} L(\eta, w) &= \sum_{\omega \in P} \left(\sum_{i=1}^n (E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega_{\max}] - E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega]) \right) \mu(\omega) \\ &= \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega_{\max}] - \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \sigma^*]. \end{aligned}$$

In Section 2, when $\rho = 0$, P has two elements: $\{(H_1, H_2)\}$ and $\{(L_1, L_2)\}$, and $L(\eta, w)$ is equal to the difference in total payments between these two elements. When $\rho \neq 0$, P is a singleton and, hence, $L(\eta, w)$ is zero for all w .

It is straightforward to see that

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

In the following, we will refer to $L(\eta, w)$ and $-\sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \sigma^*, \omega_{\max}]$ as the self-evident efficiency loss and non-self-evident efficiency loss of w , respectively.

Let

$$L^*(\eta) \equiv \min_{w \in \mathcal{W}(\eta)} L(\eta, w) \quad (13)$$

denote the minimum self-evident efficiency loss among $w \in \mathcal{W}(\eta)$.

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

Proof. See Appendix A. □

Theorem 1 is a generalization of Case 1 of Section 2. It connects the long-term efficiency of a partnership to the short-term self-evident efficiency loss.

As in Case 1, the crucial observation is that for any $\omega, \omega' \in P$, the incentives after ω and the incentives after ω' are completely separated, as players following ω assign zero probability to the first-period outcome $(\tilde{a}(1), y(1))$ not belonging to ω . As a consequence, enforcing η in the remaining $(T - 1)$ periods after ω is equivalent to enforcing η in a $(T - 1)$ -period contracting game. The only complication here is that in this new contracting game, the players observe an additional correlating device that recommends (\tilde{a}_i, y_i) to each player i according to the distribution $\mu(\cdot|\omega)$. However, as we explained in Section 3, having an extra correlating device does not improve efficiency. Hence, the efficiency loss for enforcing η in the continuation game after ω would be the same as enforcing η in a $(T - 1)$ -period contracting game. This, together with the fact that any $w^T \in \mathcal{W}(\eta, T, \delta)$ must enforce η in the first period, implies Theorem 1.

Recall that $W^*(\eta)$ is the minimum efficiency loss when $T = 1$. Since $L(\eta, w) \leq W(\eta, w)$ for every $w \in \mathcal{W}(\eta)$,

$$L^*(\eta) \leq W^*(\eta). \quad (14)$$

In general, $L^*(\eta)$ could be strictly lower than $W^*(\eta)$. A special case is when η is pure and the signal structure is public. In this case, every (\tilde{a}, y) in the support of μ is self-evident. Let w^* be a contract that minimizes $L(\eta, w)$ among all $w \in \mathcal{W}(\eta)$. Let y^* be the signal profile that maximizes the total payment $\sum_{i=1}^n w_i^*(\tilde{a}, y)$. We have

$$L^*(\eta) = \sum_{i=1}^n w_i^*(\tilde{a}, y^*) - \sum_{i=1}^n E[w_i^*(\tilde{a}, \hat{y})|\sigma^*].$$

Define a new contract w' by subtracting a constant $w_i^*(\tilde{a}, y^*)$ from the payment of every player i . It is obvious that w' also enforces η . Furthermore, for every (\tilde{a}, \hat{y}) in the support of μ ,

$$\sum_{i=1}^n w'_i(\tilde{a}, \hat{y}) = \sum_{i=1}^n (w_i^*(\tilde{a}, \hat{y}) - w_i^*(\tilde{a}, y^*)) \leq 0.$$

Hence, the contract w' also belongs to $\mathcal{W}(\eta)$. It follows that

$$L^*(\eta) = - \sum_{i=1}^n E[w'_i(\tilde{a}, \hat{y})|\sigma^*] \geq W^*(\eta).$$

This, together with (14), implies the following corollary.

Corollary 1. *When an enforceable η is pure and the signal structure is public, linking has no value, and the minimum long-term efficiency loss can be achieved by a series of short-term contracts.*

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. 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 $(a'_i, y'_i) \in A_i \times Y_i$,

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

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

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 must lead to the same distribution of y_{-i} that γ'_i induces, and must 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 mis-report 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. Under a

non-stationary contract, players have incentives to deviate to actions that generate more information about the private information of the other players. For example, under the truncated contract in Case 2 of Section 2, a player will gain if he learns whether the truncation is likely to occur. The no-free-information condition ensures that no player can do so undetectably without paying a cost.

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 $\delta \geq 1 - T^{-2}$, $W^*(\eta, T, \delta) \leq L^*(\eta) + \varepsilon$.*

Theorem 2 says that the bound established in Theorem 1 is tight when η satisfies the no-free-information condition. Note that the condition does not impose a lower bound on the cost of acquiring more information. As T becomes large, the potential gain from having more information can be made arbitrarily small (but not zero).

Theorem 2 implies that η can be enforced almost efficiently when $L^*(\eta) = 0$. The following corollary follows from the fact that $L^*(\eta) = 0$ when η is pure and P is a singleton.

Corollary 2. *An enforceable pure-action profile that satisfies the no-free-information condition can be enforced almost efficiently in the long term if P is a singleton.*

In the literature of repeated games with private monitoring, the full support of the signal distribution is often invoked as a simplifying assumption. In fact, since full support implies that P is a singleton, the assumption, by itself, implies the almost-efficient enforcement of any enforceable pure-action profile that satisfies the no-free-information condition.

We prove Theorem 2 by constructing a long-term contract that approaches the efficiency bound in the limit. The details of the proof are provided in Appendix C. Below we outline the main steps of the construction. The contract is a general version of the one in Case 3 of Section 2. In Case 3 we start with a stage-game contract that strictly enforces the desired actions (C, C) plus a scoring rule that induces truthful reporting. The no-free-information condition ensures that there exists a stage-game contract such that “almost all” deviations can be strictly deterred.

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 .

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. A formal proof is provided in an online appendix.²⁰ 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, the no-free-information condition implies that for any $\varepsilon > 0$, there exists a stage-game contract w^* that enforces η almost strictly with

$$L(\eta, w^*) < L^*(\eta) + \varepsilon. \quad (16)$$

Under w^* , any deviation (in action or reporting) that is detectable or generates more information is strictly deterred.²¹ While there may exist non-detectable deviating actions that generate the same stage-game payoff and are as informative as the obedient strategy, a player will not strictly gain from choosing such a deviating action in any period.^{22,23}

The stage-game contract w_i^* can be decomposed into two components:

$$w_i^*(\tilde{a}, \hat{y}) = w_{i,a}^*(\tilde{a}, \hat{y}) + w_{i,b}^*(\tilde{a}, \hat{y}), \quad (17)$$

²⁰The converse of Lemma 3 is false as almost-strict enforceability does not rule out pure undetectable deviations that are strictly more informative than the obedient strategy. As a result, Theorem 2 does not hold if the no-free-information condition is replaced with almost-strict enforceability. We provide an example of this in an online Appendix.

²¹Since strict enforceability rules out any profitable deviation from the recommendation, it implies the no-free-information condition. The no-free-information condition is weaker than strict enforceability.

²²Such deviation will result in the same distribution of outcomes in the current period and does not generate information that allows the player to deviate profitably in future periods.

²³Because of the possibility of such actions, the combination of enforceability and the no-free-information condition is weaker than strict enforceability.

where

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

It is straightforward to verify that

$$E[w_{i,a}^*(\tilde{a}, \hat{y}) | \sigma^*, \omega] = E[w_i^*(\tilde{a}', \hat{y}') | \sigma^*, \omega_{\max}], \quad \forall i \in N, \forall \omega \in P; \quad (18)$$

$$\sum_{i=1}^n E[w_{i,b}^*(\tilde{a}, \hat{y}) | \sigma^*] = L(\eta, w^*). \quad (19)$$

Intuitively, $w_{i,b}^*$ captures incentives that vary across self-evident events, while $w_{i,a}^*$ captures incentives that vary within self-evident events.

From Theorem 1, we know that the loss associated with the self-evident component, $w_{i,b}^*$, cannot be eliminated by linking. We prove Theorem 2 by showing that the loss associated with the self-evident component, $w_{i,a}^*$, can be reduced by linking. As in Case 3 of Section 2, we replace the T -period version of $w_{i,a}^*$ with the combination of a truncated contract and a set of side-bet contracts. To reduce efficiency loss, the truncated contract punishes player i only when the T -period value of $w_{i,a}^*$ falls below a T -period performance standard. The side-bet contracts make up for the truncation by ensuring that a player who consistently over-performs is also likely to receive a bonus from the other players.

The key to ensure that the side bets do not have a perverse effect on the incentives of the paying players is the following general version of Lemma 1 that applies to all stage games. Consider an outside observer who observes at the end of each period which element of P has occurred. With a slight abuse of terminology, we will continue to refer to the expectation of this outside observer as the “prior” and the expectations of the players conditional on their private information the “posteriors.” For any outcome history (\tilde{a}^T, y^T) and stage-game outcome (\tilde{a}, \hat{y}) , 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 occurrences of (\tilde{a}, y) , (\tilde{a}_i, y_i) , and $P(\tilde{a}, y)$ in (\tilde{a}^T, y^T) . For any (\tilde{a}, y) , the prior expectation of $f(\tilde{a}, y | \tilde{a}^T, y^T)$ is

$$E_0[f(\tilde{a}, y | \tilde{a}^T, y^T)] = \mu(\tilde{a}, y | P(\tilde{a}, y)) f(P(\tilde{a}, y) | \tilde{a}^T, y^T),$$

while player i 's posterior expectation of $f(\tilde{a}, y | \tilde{a}^T, y^T)$ is

$$E_i[f(\tilde{a}, y | \tilde{a}^T, y^T)] = \mu(\tilde{a}, y | \tilde{a}_i, y_i) f(\tilde{a}_i, y_i | \tilde{a}^T, y^T).$$

For any $\xi > 0$, let

$$Z^T(\xi) = \{(\tilde{a}^T, y^T) \mid \exists(\tilde{a}, y) \in \text{supp}(\mu) : |f(\tilde{a}, y | \tilde{a}^T, y^T) - E_0[f(\tilde{a}, y | \tilde{a}^T, y^T)]| > \xi T\}$$

denote the set of histories in which the frequency of some outcome is different from the prior expectation by ξT . Similarly, let

$$Z_i^T(\xi) = \{(\tilde{a}^T, y^T) \mid \exists(\tilde{a}, y) \in \text{supp}(\mu) : |f(\tilde{a}, y | \tilde{a}^T, y^T) - E_i[f(\tilde{a}, y | \tilde{a}^T, y^T)]| > \xi T\}$$

denote the set of histories in which the frequency of some outcome is different from player i 's posterior expectation by ξT .

Lemma 4 (Posteriors Determine Prior). *For any $\iota > 0$, there exists $\varepsilon > 0$ 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 $(\tilde{a}^T, y^T) \in Z_i^T(\varepsilon)$ for some player i .*

Proof. See Appendix B. □

In words, Lemma 4 says that when the frequencies of realized outcomes are different from the expected frequencies conditional on information that is self-evident among players, they must also be different from the expected frequencies conditional of the private information of some player.²⁴ As in Case 3 of Section 2, using Lemma 4, we can construct side-bet contracts as follows.

When a player i consistently over-performs, the frequencies of realized outcomes must be different from the “prior” expected frequencies. By Lemma 4, there exists a player j such that they are also different from j 's “posterior” expected frequencies. In the side-bet contracts, player j is asked to pay a bonus to player i . This additional bonus makes up for the truncated incentives.

By the law of large numbers, conditional on any y_i^T , it is extremely unlikely that $y^T \in Z_i^T(\varepsilon)$ when T is large. That is, player i believes that he will almost never have to

²⁴Lemma 4 is closely related to a result in Samet (1998) that shows that in an incomplete information games, if the meet of the players' information partitions is a singleton, then there is at most one common prior that can generate the posterior beliefs. Samet (1998) uses this result to show that when a common prior exists, the players' higher-order beliefs about any random variable will converge to the prior expectation of the random variable. Our innovation is to exploit the *continuity* of the mapping from the prior to posteriors to derive implications on the players' posterior expectations when they repeatedly observe the realization of a random variable.

pay the bonus and he will almost sure to receive a bonus when he needs one. The side bets, therefore, do not create a perverse effect on the incentives of the paying players.

5 Long-Term Efficiency Loss

Theorems 1 and 2 show that the long-term efficiency of a partnership depends on a simple criterion; namely, whether incentives need to vary across self-evident events to enforce the desired action in a stage game. A partnership will be inefficient if there are some deviations that can only be deterred by incentives that vary across self-evident events; for example, a deviation to D in Case 1 of Section 2.

More generally, deviations that satisfy the following two properties can be deterred only with incentives that vary across the self-evident events. First, a deviating strategy profile $(\sigma_1, \dots, \sigma_n)$ is called *unattributable* if each unilateral deviation σ_i leads to the same distribution of outcomes; that is,

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

Second, a deviation σ_i is *non-detectable within self-evident events* if it does not change the distribution of (\tilde{a}, \hat{y}) conditional on any member in P . The set of detectable deviating strategy profiles that are both unattributable and non-detectable within self-evident events is

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

For any σ_i , let

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

denote player i 's gain from the deviation σ_i . Write $l(\sigma_i) \equiv \max_{\omega \in P} \pi^{\sigma_i}(\omega) / \mu(\omega)$.

The following theorem characterizes $L^*(\eta)$ in terms of the primitives of the stage game.

Theorem 3. *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, $L^(\eta) = 0$.*

Proof. See Appendix D. □

Intuitively, since $(\sigma_1, \dots, \sigma_n) \in Q(\eta)$ 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 deviating gain multiplied by a factor that measures the difference between π^{σ_i} and μ . For example, in Case 1 of Section 2, $d(D) = d$ and $l(D) = (1 - q) / (1 - p)$. Hence,

$$\frac{\sum_{i=1}^2 d(D)}{l(D) - 1} = \frac{2(1 - p)d}{p - q}.$$

When $Q(\eta)$ contains multiple deviating strategy profiles, Theorem 3 says that $L^*(\eta)$ is entirely determined by the member in $Q(\eta)$ that is the hardest to deter.

An immediate corollary of Theorem 3 is that $L^*(\eta) = 0$ when $Q(\eta)$ is empty.

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

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

$Q(\eta)$ is empty if every deviation satisfies one of the following three conditions: first, it may result in an outcome outside of the support of μ ; second, it may change the distribution of (\tilde{a}, \hat{y}) conditional on some $\omega \in P$; third, it is attributable. The first type of deviation 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 whose total expected payment is constant across members of P . The third type can be deterred by a budget-balance contract (Rahman and Obara, 2010).

The literature on repeated games with private monitoring and communication can be divided into two strands. One strand applies the linking idea to enhance efficiency. Another strand (Fudenberg, Levine, and Maskin, 1994; Kandori and Matsushima, 1998; Rahman and Obara, 2010) identifies conditions that ensure that the desired actions can

be enforced by a budget-balanced contract. Theorem 3, together with Theorem 2, implies that an outcome can be enforced almost efficiently in the long-term if every deviating strategy can be deterred *either* by a contract that is budget-balanced or whose total expected payment is constant across self-evident events. Theorems 1 and 3 connect the efficiency results in the repeated-game literature with the inefficiency result of Abreu, Milgrom, and Pearce (1991). They imply that the two approaches in the repeated-game literature are, in fact, the only approaches to achieve long-term efficiency.

5.1 Changing information structure

Theorem 3 illustrates how a mechanism designer may improve the long-term efficiency by changing the information structure. In the following, we consider two modifications to Case 1 of Section 2.

Example 1. Suppose, in addition to the public signal $y \in \{H, L\}$, each player i also observes a private signal $z_i \in \{h_i, l_i\}$. The distribution of the private signals has full support conditional on y and (C, C) . The meet of the players' information partitions (conditional on (C, C)) is therefore equal to

$$P: \{(H, h_1, h_2), (H, h_1, l_2), (H, l_1, h_2), (H, l_1, l_2)\}, \\ \{(L, h_1, h_2), (L, h_1, l_2), (L, l_1, h_2), (L, l_1, l_2)\}.$$

Let $p(z_1, z_2 | a, y)$ denote the conditional probability of (z_1, z_2) . Theorem 3 implies that if, for some $y' \in \{H, L\}$,

$$p(\cdot | CC, y') \neq p(\cdot | DC, y') = p(\cdot | CD, y'),$$

then $L^*(C, C) = 0$. By contrast, if for all $y \in \{H, L\}$,

$$p(\cdot | CC, y) = p(\cdot | DC, y) = p(\cdot | CD, y),$$

then $L^*(C, C) = 2(1 - p)d / (p - q)$. Intuitively, adding the private signals improves efficiency only when the signals are informative about the players' actions *conditional* on the public signal.

Example 2. Instead of adding extra signals, the players may alter the information structure by changing actions. Suppose that, instead of (C, C) , the players implement

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,

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

and the meet is

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

It is straightforward to see that $\bar{\eta}$ is enforceable.²⁵ 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}).

	<i>CCH</i>	<i>DCH</i>	<i>CDH</i>
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} . 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 . If player 1 deviates to D when told to choose C , he must lower the relative probability of CCH over DCH . Hence, every deviation is detectable with respect to the self-evident event $\{CCH, CDH, DCH\}$.

²⁵Action D is enforceable because it is a best response to C .

Theorem 3 implies that $L^*(\bar{\eta}) = 0$. 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 term. The idea behind the example applies generally. We prove the following theorem in an 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, $\max_{\tilde{a} \in A} |\eta(\tilde{a}) - \bar{\eta}(\tilde{a})| \leq \varepsilon$, and $L^*(\bar{\eta}) = 0$.*

Thus, even when η cannot be enforced efficiently in the long term, there is always a correlated action profile close to it that can be enforced almost efficiently. The only restriction is that η can be strictly enforced.²⁶

The idea of using correlated strategies to enhance efficiency comes from Rahman (2014). As in our case, players in Rahman (2014) are also secretly recommended to deviate secretly. These deviations serve as “secret tests” that prevent each player from learning his own punishment. By contrast, in our method, the deviations from efficient actions serve as benchmarks that ensure that each player’s punishment is not learned by *other* players. Rahman (2014)’s method requires that an outcome be *conditionally identifiable*. Our construction, by contrast, applies to all strictly enforceable outcomes.

6 Relation to the Repeated Game Literature

In a seminal paper, Abreu, Milgrom, and Pearce (1991) demonstrate the importance of the timing of information to the efficiency of long-term partnerships. They consider two scenarios that correspond to Cases 1 and 2 in Section 2. In the first scenario, players observe a public signal at the end of each period. In this case linking has no value. In the second scenario, players observe signals in the previous T periods at the end of every T periods. When T becomes large, there exists an almost-efficient equilibrium in which incentives are linked across periods to economize on the cost of imperfect monitoring.

²⁶Theorem 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 .

Subsequent research applies the second result of Abreu, Milgrom, and Pearce (1991) to repeated games with private monitoring and communication. While players observe private signals at the end of every period in a repeated game, they may delay the arrival of information by communicating infrequently. The only problem is that players may update their beliefs about other players' signals on the basis of their own. Compte (1998) finesses this problem by assuming that signals are conditionally independent. Obara (2009) considers imperfectly correlated signals and identifies a necessary and sufficient condition on the signal distribution that ensures the existence of a performance measure that the players cannot learn from their own signals. Rahman (2014) identifies a similar condition when correlated strategies are allowed.

This paper uses the notion of a self-evident event to generalize the insights of Abreu, Milgrom, and Pearce (1991). Working with a T -period contracting problem allows us to focus on the mechanism of linking and abstract away from the problem of implementing transfers through continuation strategies. Theorem 1 extends the first result of Abreu, Milgrom, and Pearce (1991). Theorem 2 extends the second result of Abreu, Milgrom, and Pearce (1991). Taken together, the two theorems describe how the effectiveness of linking is limited by the information that becomes self-evident at the end of each period.

The proof of Theorem 2, which corresponds to Case 3 in Section 2, exploits the differential beliefs between players. Fong, Gossner, Hörner, and Sannikov (2011), in a repeated Prisoners' Dilemma similar to the example in Section 2, use the fact that each player expects the other player to observe fewer excess L signals than he does to support an approximately efficient equilibrium. Our approach can be viewed as a generalization of their result. Since our approach does not depend on players not learning their own payoffs, it is more general than Compte (1998), Obara (2009), and Rahman (2014) (which correspond to Case 2 in Section 2). In our earlier work Chan and Zhang (2016), we consider a repeated game in which each player's private signal is his own stage-game payoff and the distribution of stage-game payoffs has full support. It shows that any strictly efficient action profile can be enforced almost efficiently using bilateral side payments. Theorem 2 clarifies the logic behind Chan and Zhang (2016). It does not impose restrictions on the monitoring structure and applies to all enforceable action profiles that satisfy the no-free-information condition. In two recent papers, Sugaya (2017a,b) derives upper and lower bounds in equilibrium payoffs in repeated games

with private monitoring and correlated strategies. Sugaya's method always requires a correlating device, while ours requires one only when the outcome is correlated.

7 Conclusion

Players in a long-term 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-term 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.

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 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, $w \in \mathcal{W}(\eta)$.

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 value of w_i^T as a function of $(\tilde{a}^{2,T}, \hat{y}^{2,T})$, conditional on $(\tilde{a}(1), \hat{y}(1)) \in \omega$ (assuming that players follow the equilibrium strategy). 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, if η can be enforced by a contract and some correlating device, it can be enforced with the same efficiency loss by a correlating device that directly recommends η . Hence, having an extra randomization device does not enhance efficiency. By the supposition that Theorem 1 holds for $(T - 1)$, the expected efficiency loss of the contract $\delta^{-1}w^{T-1,\omega}$, 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 4

We will prove the contrapositive: for any $\iota > 0$, there exists $\varepsilon > 0$ such that for all $(\tilde{a}^T, y^T) \in \text{supp}(\mu)^T$, if $(\tilde{a}^T, y^T) \notin Z_i^T(\varepsilon)$ for each player i , then $(\tilde{a}^T, y^T) \notin Z^T(\iota)$.

Fix \tilde{a}^T, y^T and $\varepsilon > 0$. Suppose that for each player i and each (\tilde{a}, y) ,

$$|f(\tilde{a}, y | \tilde{a}^T, y^T) - \mu(\tilde{a}, y | \tilde{a}_i, y_i) f(\tilde{a}_i, y_i | \tilde{a}^T, y^T)| \leq \varepsilon T. \quad (20)$$

Since the game is finite, there is some constant K such that any (\tilde{a}', y') $\in \text{supp}(\mu)$ is reachable from any $(\tilde{a}'', y'') \in P(\tilde{a}', y')$ in less than K steps. That is, there exists a sequence $(\tilde{a}^1, y^1), (\tilde{a}^2, y^2), \dots, (\tilde{a}^k, y^k)$ and a sequence i_1, i_2, \dots, i_{k-1} such that (i)

$(\tilde{a}^1, y^1) = (\tilde{a}', y')$ and $(\tilde{a}^k, y^k) = (\tilde{a}'', y'')$, (ii) $k \leq K$, (iii) $(\tilde{a}^s, y^s) \in P(\tilde{a}', y')$ for each $s \leq k$ and (iv) for any $s = 1, \dots, k-1$, $(\tilde{a}^s, y^s), (\tilde{a}^{s+1}, y^{s+1}) \in P_{i_s}(\tilde{a}^s, y^s)$ (see, e.g., Aumann, 1976; Geanakoplos, 1994).²⁷

Let $c_0 \equiv \min_{(\tilde{a}, y) \in \text{supp}(\mu)} \mu(\tilde{a}, y | P(\tilde{a}, y))$. Substituting (\tilde{a}^s, y^s) for (\tilde{a}, y) and i_s for i in (20) and dividing both sides by $\mu(\tilde{a}^s, y^s | P(\tilde{a}^s, y^s))$, we have

$$\left| \frac{f(\tilde{a}^s, y^s | \tilde{a}^T, y^T)}{\mu(\tilde{a}^s, y^s | P(\tilde{a}^s, y^s))} - \frac{f(\tilde{a}_{i_s}^s, y_{i_s}^s | \tilde{a}^T, y^T)}{\mu(\tilde{a}_{i_s}^s, y_{i_s}^s | P(\tilde{a}^s, y^s))} \right| \leq \frac{\varepsilon T}{\mu(\tilde{a}^s, y^s | P(\tilde{a}^s, y^s))} \leq \frac{\varepsilon T}{c_0}. \quad (21)$$

Then, by (21) and property (iv), for any $s \leq k-1$,

$$\begin{aligned} & \left| \frac{f(\tilde{a}^s, y^s | \tilde{a}^T, y^T)}{\mu(\tilde{a}^s, y^s | P(\tilde{a}^s, y^s))} - \frac{f(\tilde{a}^{s+1}, y^{s+1} | \tilde{a}^T, y^T)}{\mu(\tilde{a}^{s+1}, y^{s+1} | P(\tilde{a}^{s+1}, y^{s+1}))} \right| \\ & \leq \left| \frac{f(\tilde{a}^s, y^s | \tilde{a}^T, y^T)}{\mu(\tilde{a}^s, y^s | P(\tilde{a}^s, y^s))} - \frac{f(\tilde{a}_{i_s}^s, y_{i_s}^s | \tilde{a}^T, y^T)}{\mu(\tilde{a}_{i_s}^s, y_{i_s}^s | P(\tilde{a}^s, y^s))} \right| \\ & \quad + \left| \frac{f(\tilde{a}^{s+1}, y^{s+1} | \tilde{a}^T, y^T)}{\mu(\tilde{a}^{s+1}, y^{s+1} | P(\tilde{a}^{s+1}, y^{s+1}))} - \frac{f(\tilde{a}_{i_s}^{s+1}, y_{i_s}^{s+1} | \tilde{a}^T, y^T)}{\mu(\tilde{a}_{i_s}^{s+1}, y_{i_s}^{s+1} | P(\tilde{a}^{s+1}, y^{s+1}))} \right| \\ & \leq \frac{2\varepsilon T}{c_0}. \end{aligned} \quad (22)$$

Applying (22) repeatedly, we have

$$\left| \frac{f(\tilde{a}', y' | \tilde{a}^T, y^T)}{\mu(\tilde{a}', y' | P(\tilde{a}', y'))} - \frac{f(\tilde{a}'', y'' | \tilde{a}^T, y^T)}{\mu(\tilde{a}'', y'' | P(\tilde{a}'', y''))} \right| \leq \frac{2K\varepsilon T}{c_0}.$$

It follows that

$$f(\tilde{a}'', y'' | \tilde{a}^T, y^T) > f(P(\tilde{a}', y') | \tilde{a}^T, y^T) \mu(\tilde{a}'', y'' | P(\tilde{a}', y')), \quad (23)$$

if

$$\frac{f(\tilde{a}', y' | \tilde{a}^T, y^T)}{\mu(\tilde{a}', y' | P(\tilde{a}', y'))} - f(P(\tilde{a}', y') | \tilde{a}^T, y^T) > \frac{2K\varepsilon T}{c_0}.$$

Note that (23) would apply to all $(\tilde{a}'', y'') \in P(\tilde{a}', y')$. But this is impossible as

$$\begin{aligned} \sum_{(\tilde{a}'', y'') \in P(\tilde{a}', y')} f(\tilde{a}'', y'' | \tilde{a}^T, y^T) &= f(P(\tilde{a}', y') | \tilde{a}^T, y^T) \\ &= \sum_{(\tilde{a}'', y'') \in P(\tilde{a}', y')} \mu(\tilde{a}'', y'' | P(\tilde{a}', y')) f(P(\tilde{a}', y') | \tilde{a}^T, y^T). \end{aligned}$$

²⁷Recall that $(\tilde{a}'_i, y'_i) = (\tilde{a}_i, y_i)$ for all $(\tilde{a}', y') \in P_i(\tilde{a}, y)$.

Thus, if (20) holds for all i and (\tilde{a}, y) , then for all (\tilde{a}', y')

$$\begin{aligned} f(\tilde{a}', y' | \tilde{a}^T, y^T) - f(P(\tilde{a}', y') | \tilde{a}^T, y^T) \mu(\tilde{a}', y' | P(\tilde{a}', y')) &\leq \frac{2K\varepsilon T \mu(\tilde{a}', y' | P(\tilde{a}', y'))}{c_0} \\ &\leq \frac{2K\varepsilon T}{c_0}. \end{aligned} \quad (24)$$

By similar logic,

$$- (f(\tilde{a}', y' | \tilde{a}^T, y^T) - f(P(\tilde{a}', y') | \tilde{a}^T, y^T) \mu(\tilde{a}', y' | P(\tilde{a}', y'))) \leq \frac{2K\varepsilon T}{c_0}. \quad (25)$$

Combining (24) and (25), we have that for any (\tilde{a}^T, y^T) , if (20) holds for all i and (\tilde{a}, y) , then for all (\tilde{a}', y')

$$|f(\tilde{a}', y' | \tilde{a}^T, y^T) - f(P(\tilde{a}', y') | \tilde{a}^T, y^T) \mu(\tilde{a}', y' | P(\tilde{a}', y'))| \leq \frac{2K\varepsilon T}{c_0}.$$

The contrapositive of the lemma can be obtained by setting $\varepsilon = \iota c_0 / 2K$.

C Proof of Theorem 2

We prove Theorem 2 in two steps. In Step 1, we construct a T -period contract w^{T**} . In Step 2, we show that w^{T**} enforces η and achieves the efficiency bound $L^*(\eta) + \varepsilon$.

C.1 Step 1. Constructing the T -period contract w^{T**} .

Recall that $(w_{i,a}^*, w_{i,b}^*)$ is the decomposition of w_i^* in (17). Let $w_{i,a}^{T*}$, $w_{i,b}^{T*}$, and w_i^{T*} denote the T -period versions of $w_{i,a}^*$, $w_{i,b}^*$, and w_i , respectively. That is, for all $i \in N$ and all (\tilde{a}^T, \hat{y}^T) ,

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

The action η can be enforced by a contract that pays $w_{i,a}^{T*} + w_{i,b}^{T*}$ to each player i . To prove Theorem 2, we show that $w_{i,a}^{T*}$ can be replaced by a truncated contract similar to the one in Case 3 of Section 2. Fix some small $\kappa > 0$. Define

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

Intuitively, $E[w_{i,a}^{T*}(\tilde{a}^{T'}, \hat{y}^{T'}) | \sigma^{T*}] + \kappa T$ can be taken as a long-term performance standard, and R_i^+ and R_i^- measure, respectively, the extent of over- and under-performance with respect to the standard.

Let

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

denote the set of (\tilde{a}^T, \hat{y}^T) where player i over-performs. By (18), the expected value of $w_{i,a}^{T*}$ is constant across P . Hence, for player i to over-perform, the distribution of outcomes must deviate from the distribution of outcomes conditional on the distribution of $\omega \in P$. It then follows from Lemma 4 that any $(\tilde{a}^T, \hat{y}^T) \in B_i^T(\kappa)$ must deviate from the posterior expectation of some player. Formally, there exists $\xi > 0$ such that

$$B_i^T(\kappa) \subseteq \cup_{j \in N} Z_j^T(\xi). \quad (26)$$

(Recall that $Z_i^T(\xi)$, defined in Section 4, is the set of (\tilde{a}^T, \hat{y}^T) that deviates from the posterior expectation of player i by ξT .)

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

$$\begin{aligned} w_i^{T**}(\tilde{a}^T, \hat{y}^T, \kappa) &= R_i^-(\tilde{a}^T, \hat{y}^T, \kappa) + w_{i,b}^{T*}(\tilde{a}^T, \hat{y}^T) \\ &+ \left[R_i^+(\tilde{a}^T, \hat{y}^T, \kappa) (1 - I_i(\tilde{a}^T, \hat{y}^T, \kappa)) - \sum_{j \neq i} R_j^+(\tilde{a}^T, \hat{y}^T, \kappa) I_i(\tilde{a}^T, \hat{y}^T, \kappa) \right], \end{aligned} \quad (27)$$

where

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

Under w^{T**} , player i receives the self-evident component $w_{i,b}^{T*}$ in full, and pays an under-performance penalty (R_i^-) when $w_{i,a}^{T*}$ falls below the performance standard. In addition, he receives an over-performance bonus when $w_{i,a}^{T*}$ is above the performance standard

and $I_i = 0$, and pays an over-performance bonus to each player j when $w_{j,a}^{T*}$ is above the performance standard and $I_i = 1$.

The total payment is negative for all (\tilde{a}^T, \hat{y}^T) . By definition, for all (\tilde{a}^T, \hat{y}^T) ,

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

The sum of the third component inside the square bracket in (27) is also negative, as, by (26), for any player i and any $(\tilde{a}^T, \hat{y}^T) \in B_i^T(\kappa)$, $I_j(\tilde{a}^T, \hat{y}^T, \kappa) = 1$ for some player j .

Rearranging the terms on the right-hand side of (27), we can write

$$w_i^{T**}(\tilde{a}^T, \hat{y}^T, \kappa) = w_i^{T*}(\tilde{a}^T, \hat{y}^T) - E[w_{i,a}^{T*}(\tilde{a}^{T'}, \hat{y}^{T'}) | \sigma^{T*}] - \kappa T - \phi_i(\tilde{a}^T, \hat{y}^T, \kappa), \quad (28)$$

where

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

measures the distortion in incentives.

C.2 Step 2. Showing that w^{T**} enforces η and achieves the efficiency bound $L^*(\eta) + \varepsilon$.

We say that a pure-action strategy α_i is equally informative to 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 $\rho_i(\tilde{a}_i, \alpha_i(\tilde{a}_i), \cdot) = \chi_{\tilde{a}_i}^{-1}$.

Note that if some (α_i, ρ_i) is not a duplicate of (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) ,

$$v_i(\sigma^*; w_i^*) - v_i(\sigma_{-i}^*, \alpha_i, \rho_i; w_i^*) > \Delta_0. \quad (29)$$

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 to 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. Let H_i^{T*} denote the set of histories that player i may observe during the T -period contract under σ^{T*} . Suppose σ_i^T first prescribes a non-equally-informative action in period t after $h_i \in H_i^{T*}$. 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 (27) and w_i^{T*} is the T -period version of w_i^* .

By (28), we can write

$$v_i^T(\sigma_i^T, \sigma_{-i}^{T*}; w_i^{T**}, h_i) = \frac{1-\delta}{1-\delta^T} (V_i(\sigma_i^T; h_i) - E[w_{i,a}^{T*}(\tilde{a}^T, \hat{y}^T, \kappa) | \sigma^{T*}] - \kappa T - E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa) | \sigma_{-i}^{T*}, \sigma_i^T, h_i]),$$

where

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

denotes player i 's discounted payoff conditional on h_i under w_i^{T*} . It follows that

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

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

Since σ_i^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

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

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

Recall that ξ is the constant defined in (26) and it depends only on κ .

Claim 1. *There exists $c_1 > 0$ such that for all $i \in N$, $T \geq 1$, and $h_i \in H_i^{T*}$,*

$$E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa) | \sigma^{T*}, h_i] < c_1 T \exp\left(-\frac{\xi^2}{2} T\right).$$

A proof of Claim 1 is provided in an online appendix. Granted Claim 1, we can choose T_0 large enough such that for all $T \geq T_0$ and $\delta \geq 1 - T^{-2}$ (which ensures that $(1 - \delta^T)/(1 - \delta)$ is on the order of T as δ tends to 1),

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

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

Finally, by (28), the per-period efficiency loss is

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

By Claim 1, when T is sufficiently large, we can choose κ small enough such that

$$\frac{1 - \delta}{1 - \delta^T} \sum_{i=1}^n (E[\phi_i(\tilde{a}^T, \hat{y}^T, \kappa) | \sigma^{T*}] + \kappa T) \leq \varepsilon_0,$$

for some $\varepsilon_0 \leq L^*(\eta) + \varepsilon - L(\eta, w^*)$. (Note that $(1 - \delta)/(1 - \delta^T)$ is of order $1/T$.)
Hence,

$$\begin{aligned} W(\eta, T, \delta, w^{T**}) &= \frac{1 - \delta}{1 - \delta^T} \sum_{i=1}^n -E[w_i^{T**} | \sigma^{T*}] \\ &\leq L^*(\eta) + \varepsilon. \end{aligned}$$

D Proof of Theorem 3

Any w that enforces η must satisfy the constraint that for each i and each $\sigma \in Q(\eta)$,

$$\sum_{(\tilde{a}, \hat{y}) \in A \times Y} (\mu(\tilde{a}, \hat{y}) - \pi^{\sigma_i}(\tilde{a}, \hat{y})) w_i(\tilde{a}, \hat{y}) \geq d(\sigma_i). \quad (30)$$

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

$$E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] = E[w_i(\tilde{a}, \hat{y}) | \pi^{\sigma_i}, \omega]. \quad (31)$$

Substituting (31) into (30), and summing over i , we have

$$\sum_{\omega \in P} (\mu(\omega) - \pi^{\sigma_i}(\omega)) \sum_{i=1}^n E[w_i(\tilde{a}, \hat{y}) | \mu, \omega] \geq \sum_{i=1}^n d(\sigma_i). \quad (32)$$

Hence,

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

where the first inequality follows from the definition of $l(\sigma_i)$, and the last inequality follows from (32). Since the argument applies to every w that enforces η ,

$$L^*(\eta) \geq \frac{\sum_{i=1}^n d(\sigma_i)}{l(\sigma_1) - 1}, \quad \forall \sigma \in Q(\eta).$$

To show the other direction of the theorem, let

$$\bar{L} \equiv \begin{cases} \sup_{\sigma \in Q(\eta)} \frac{\sum_{i=1}^n d(\sigma_i)}{l(\sigma_1) - 1}, & \text{if } Q(\eta) \neq \emptyset; \\ 0, & \text{otherwise.} \end{cases}$$

It remains to show that

$$L^*(\eta) \leq \bar{L}. \quad (33)$$

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

$$\sum_{(\tilde{a}, \tilde{y}) \in A \times Y} [\pi^{\alpha_i, \rho_i}(\tilde{a}, \tilde{y}) - \mu(\tilde{a}, \tilde{y})] w_i(\tilde{a}, \tilde{y}) \leq -d(\alpha_i, \rho_i) \quad \forall (\alpha_i, \rho_i, i); \quad (34)$$

$$\sum_{i=1}^n \sum_{(\tilde{a}, \tilde{y}) \in A \times Y} [-\mu(\tilde{a}, \tilde{y}) + \mu(\tilde{a}, \tilde{y} | \omega)] w_i(\tilde{a}, \tilde{y}) \leq \bar{L} \quad \forall \omega \in P. \quad (35)$$

By the theorem of alternatives (see, e.g., Proposition 5.1.2 of Bertsekas, 2009), (34) and (35) 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

$$\sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) [\pi^{\alpha_i, \rho_i}(\cdot) - \mu(\cdot)] + \sum_{\omega \in P} v(\omega) [-\mu(\cdot) + \mu(\cdot | \omega)] = 0 \quad \forall i \quad (36)$$

$$\sum_{i=1}^n \sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) d(\alpha_i, \rho_i) - \sum_{\omega \in P} v(\omega) \bar{L} > 0. \quad (37)$$

Suppose that (36) and (37) hold. From (37), $\bar{\lambda} \equiv \max_{i \in N} \sum_{(\alpha_i, \rho_i)} \lambda_i(\alpha_i, \rho_i) > 0$. 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 (36) and (37) as

$$\bar{\lambda} [\pi^{\sigma_i}(\cdot) - \mu(\cdot)] + \sum_{\omega \in P} v(\omega) [-\mu(\cdot) + \mu(\cdot | \omega)] = 0 \quad \text{for each } i \quad (38)$$

$$\sum_{i=1}^n \bar{\lambda} d(\sigma_i) - \sum_{\omega \in P} v(\omega) \bar{L} > 0. \quad (39)$$

Fix a contract w . Multiplying each (38) by $w_i(\cdot)$, then summing over all i and all

$(\tilde{a}, \hat{y}) \in A \times Y$, and adding (39), we have

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

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 $L(\eta, w) \leq \bar{L}$,

$$v_i(\sigma_i, \sigma_{-i}^*; w_i) > v_i(\sigma^*; w_i) \quad \text{for some player } i.$$

We prove (33) 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 $\sum_{i=1}^n d(\sigma_i) > 0$. 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.

- Harrington, Joseph E., and Andrzej Skrzypacz, 2011, Private monitoring and communication in cartels: Explaining recent collusive practices, *American Economic Review* 101, 2425–2449.
- 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.
- MacLeod, W. Bentley, 2003, Optimal contracting with subjective evaluation, *The American Economic Review* 93, 216–240.
- 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.
- Rubinstein, Ariel, 1979, An optimal conviction policy for offenses that may have been committed by accident, in Brams S.J., Schotter A., and Schwödiauer G., ed.: *Applied Game Theory* (Physica, Heidelberg).

- , and Menahem E. Yaari, 1983, Repeated insurance contracts and moral hazard, *Journal of Economic Theory* 30, 74–97.
- Samet, Dov, 1998, Iterated expectations and common priors, *Games and Economic Behavior* 24, 131 – 141.
- Sannikov, Yuliy, and Andrzej Skrzypacz, 2007, Impossibility of collusion under imperfect monitoring with flexible production, *The American Economic Review* 97, 1794–1823.
- 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.