

Robust Relational Contracts with Subjective Performance Evaluation*

V. Bhaskar[†] Wojciech Olszewski[‡] Thomas Wiseman[§]

September 2023

Abstract

We study a repeated principal-agent model with transferable utility, where the principal's evaluation of the agent's performance is subjective. Consequently, monitoring is noisy and private. The constructions of relational contracts in the literature are both unrealistic and not robust to small payoff shocks, and we argue that those two weaknesses are related. Allowing the principal and agent to make simultaneous cheap-talk announcements makes some effort sustainable in a robust equilibrium. Payoffs arbitrarily close to fully efficient ones can be achieved in a robust equilibrium if players are sufficiently patient. We identify important roles for cheap talk and self-reports in making a relational contract robust. Looking for robust constructions also leads us to other real-world features like non-trivial bonus targets and reasonable bonus sizes.

Keywords: Private monitoring, repeated games, relational contracts

*We are grateful to Marcin Pełski and to seminar audiences for very useful comments.

[†]Department of Economics, University of Texas at Austin and CEPR.
v.bhaskar@austin.utexas.edu.

[‡]Department of Economics, Northwestern University. wo@northwestern.edu.

[§]Department of Economics, University of Texas at Austin. wiseman@austin.utexas.edu.

1 Introduction

In many organizations, the tasks that employees must perform lack an objective, contractible measure of performance. Thus, performance evaluation is subjective, and the worker cannot observe the employer's evaluation of his own performance. Furthermore, the employer may be unable to commit to an incentive scheme, and need to rely on relational contracts and her subjective evaluation to incentivize the worker, as in Levin, Fuchs and Maestri.¹ Levin (2003), Fuchs (2007), and Maestri (2013) have examined the role of relational contracts in the context of subjective performance evaluation. These contracts typically specify that the worker exerts effort, and that the employer pays a bonus to the worker if and only if her subjective evaluation of the worker's performance is good. In order to provide incentives for the employer to truthfully disclose her evaluation, she is made indifferent between paying and not paying the bonus. To achieve that indifference, the equilibrium must exhibit money burning, for example by dissolving the relationship with some probability in the event that the bonus is not paid. Such an equilibrium exhibits an inefficiency, because productive relationships must be dissolved with positive probability. Fuchs (2007) shows that efficiency can be enhanced by requiring the principal to report on the agent's performance only every T periods. As in Radner (1985) and Abreu et al. (1991), the extent of inefficiency decreases with T , and when both players become arbitrarily patient, the equilibrium can attain full efficiency.²

Our paper begins with the observation that the equilibria so constructed are fragile, and they do not survive if the principal is subject to small payoff shocks to

¹Baker et al. (2002) and Malcomson (2013) examine the circumstances under which relational contracts are valuable in an organizational setting.

²Fuchs (2007) also shows that the class of optimal contracts, for any fixed discount factor, includes one where the principal provides no incentive pay and no feedback, until the agent is fired.

her flow revenues. In this case, equilibria where the principal is indifferent between paying the bonus or not do not survive, because the principal strictly prefers to pay the bonus when she learns that flow revenues in the relevant periods are high, and strictly prefers not to pay it when she learns that they will be low. In consequence, she will condition her bonus payment on the shock, and not on the agent's performance. This breaks the link between the agent's effort and bonus payments, and destroys his incentive to provide effort. More generally, we show that for a large class of equilibria, which includes the equilibria that we discussed earlier, either the agent will never exert effort or the equilibrium cannot be purified. In particular, none of the equilibria proposed in the literature survive when there are payoff shocks.

Our first goal is to construct positive effort purifiable equilibria – equilibria of the unperturbed game that are robust to small payoff shocks. In our equilibria, the agent's effort choices are random, and his continuation play depends non-trivially upon realized effort. Consequently, the principal's observations of output are informative, and affect her continuation value. We allow the principal and agent to make simultaneous cheap-talk announcements at the end of a period. The principal announces her private signal of the agent's performance, while the agent discloses his choice of effort. Strict incentives for truth telling are provided by dissolving the relationship with higher probability whenever the announcements “differ” – i.e., when the principal announces a bad signal while the agent reports that he worked, or the other way around. When the noise in monitoring is small and the discount factor is above a cutoff, we can construct an equilibrium where the agent works with arbitrarily high probability and the efficiency loss is close to zero. Moreover, both players have strict incentives for truth-telling and the agent's randomization is history-independent, so that the equilibrium is purifiable.

The reader might ask, why is it that we require cheap-talk announcements in ad-

dition to randomization by the agent? Previous work on the repeated prisoner’s dilemma with private monitoring has constructed “belief-based equilibria,” where initial randomization by both players, coupled with two-sided private monitoring, suffices to provide strict incentives in subsequent periods (Sekiguchi (1997) and Bhaskar and Obara (2002)). The difference arises due to the sequential nature of our stage game and since the principal’s actions (bonus payments) are public, rather than private. The cheap-talk announcements introduce an element of simultaneity that allows us to circumvent the induction arguments that underlie our negative results.³

For the case where the noise in monitoring is large, we explore equilibria where announcements are made only every T periods. A difficulty arises – any equilibrium that incentivizes truthful announcements with penalties for mismatched reports requires the agent to shirk with positive probability in every period. We present two approaches to resolving that difficulty. In the first, the agent randomly picks a single period of the T -period block in which to shirk. This construction allows us to “purify” a large variety of block equilibria, including the equilibria based on Abreu et al. (1991) mentioned in the first paragraph. In the second approach, the agent must be indifferent between always shirking and always working within each T -period block. The block equilibria based on Abreu et al. (1991), where penalties are imposed only if all T signals are bad, do not satisfy that requirement. Instead, we introduce a different construction, borrowing an idea from Matsushima (2004), that does yield the necessary indifference. Both approaches achieve efficiency in the limit as the players become arbitrarily patient.

³It could be argued that the principal may have a small preference for being honest, which could give her strict incentives to tell the truth regarding the agent’s performance. However, it could also be the case that the principal wants to pay the bonus if and only if she believes that the agent has worked. In this case, in any equilibrium where the worker always works, the principal’s beliefs do not depend upon the observed signals.

From a practical point of view, the types of small shocks to payoffs that we consider would be present in nearly any economic application. For example, an employer's indifference between paying a bonus or not would be broken by whether or not she has a pen and checkbook handy. That fragility motivates our focus on purifiable equilibria: theoretical predictions that are not robust to such small details are unlikely to accurately describe real-world behavior. The relational contracts that we introduce differ qualitatively from existing constructions in the literature. In the block equilibria studied previously, the agent gets a bonus at the end of the block except after extremely poor sequences of outcomes that are improbable even under no effort. Thus, the bonus payment at the end of each block must be huge compared to the effort cost. In some of our equilibria, the threshold to earn the bonus is at a level that the agent is likely to reach only with consistent effort, and bonus payments are of the size of effort cost. We view these properties as capturing more accurately how firms set, for example, sales targets for their employees. Another distinctive feature of our construction is the use of reports from the worker as well as from the employer. Many firms incorporate that sort of employee self-evaluation in their performance reviews.

On the theoretical side, the paper demonstrates how incorporating cheap-talk reports of private signals and incentivizing truth-telling through penalties for disagreement can be used to construct purifiable equilibria under private monitoring. That technique may be useful more generally in constructing equilibria where players have strict incentives to condition upon history in repeated games with private monitoring. The belief-free constructions often used in that setting rely on players' exact indifference between punishing or not in order to provide incentives. Bhaskar et al. (2013) discuss the difficulty of making those equilibria purifiable. Our technique also provides an additional role for mixed strategies in repeated games: non-trivial

randomization by both players is required so that each player’s private information is a predictor of the other’s announcement.

1.1 Related literature

There is a large literature on the repeated games with private monitoring that is relevant that we will not do full justice to. Briefly, while the early literature used both “belief-based” and “belief-free” approaches, most of the subsequent work has built on elements of the belief-free equilibria pioneered by Piccione (2002) and Ely and Välimäki (2002). Matsushima (2004) uses belief-free equilibria to show that block-strategy equilibria can ensure asymptotic efficiency in the repeated prisoners’ dilemma when the players’ private signals are independent. Ely et al. (2005) generalize belief-free equilibria for a larger class of games. Sugaya (2022), building on the constructions of Hörner and Olszewski (2006) and Matsushima (2004), proves a general version of folk theorem for repeated private monitoring games. In the work on relational contracts (Levin (2003), Fuchs (2007)), the availability of transfers makes it easier to achieve the indifferences required for belief-free constructions, without any need for randomization.

The present paper differs from the literature on repeated games with private monitoring in three aspects. First, the stage game considered here has a non-trivial extensive form, whereas the folk theorem of Sugaya (2022) obtains for simultaneous-move stage games. Second, monitoring of the agent by the principal is private but the principal’s actions are public. Third, and most important, is our insistence on equilibria that are robust to private payoff shocks and thus are purifiable. Nonetheless, we also build on these previous approaches. Since the agent shirks with positive probability in every period and his continuation strategy varies with realized effort,

our construction is “belief-based” in this respect. However, our equilibrium constructions may also be viewed as modifications of “belief-free” approaches, since they require the principal if she does not pay a bonus to the agent to burn the amount of money equal to the bonus. In addition, one of our constructions borrows an idea from Matsushima (2004).

Cheap-talk plays an important role in our analysis. In their pioneering work on repeated games with private monitoring, Compte (1998) and Kandori and Matsushima (1998) prove a folk theorem by using cheap-talk announcements to coordinate behavior. When there are two players, and signals are independent conditional on the action profile, the equilibria that they construct have a “belief-free” flavor, in that each player is made indifferent between her possible announcements. Cheap-talk plays a different role here, since it is a way for providing strict incentives for truth-telling. Furthermore, randomization by the agent plays an essential role in providing strict incentives, whereas randomization plays no such role in this earlier work.

Our negative results, that pure strategy equilibria cannot be purified, have the following antecedents. Matsushima (1991) studies pure strategy equilibria in repeated games with conditionally independent private monitoring. He imposes the restriction that a player’s strategy does not condition on the history of private signals unless there is a strict incentive to do so, and finds that players must play a Nash equilibrium of the stage game in every period. Bhaskar et al. (2013) show that in games where only one player moves at a time and where monitoring is public but subject to bounded memory, only Markov perfect equilibria can be purified. Our positive results require worker randomization and communication, and are related to Miyagawa et al. (2008), who study repeated games with costly monitoring, where each player has to be incentivized to monitor her opponents. Similarly, in Rahman (2012) both worker and monitor must randomize their effort and inspection decisions, respectively, in

order to incentivize each other.

Rahman and Obara (2010) study partnerships and assume that incentive schemes must satisfy budget balance. They show that a mediator can be used to virtually implement the efficient outcome where all partners work, by making, with a small probability, a secret recommendation to shirk to a randomly chosen partner, and by making rewards contingent both on output and on the recommendation. We note that a mediator makes it easier to construct nearly-efficient equilibria, since players do not have to be indifferent between working and shirking.

2 The basic model

Time is discrete, and the horizon is infinite. There are two players, the principal and the agent. The principal selects a base wage w , and in each period until the relationship is terminated, the principal pays the agent w , and the agent chooses between exerting effort (E) and shirking (S), with effort cost $c > 0$. The resulting output y , which is privately observed by the principal, is stochastic and takes values in the set $\{G, B\}$, where $G > B$.⁴ We assume that $\Pr(y = G|E) = p$ and $\Pr(y = G|S) = q$, satisfying $1 > p > q > 0$, so that output is a noisy signal of the agent's effort choice. After observing y , the principal may choose to pay the agent an additional bonus. The agent's outside option is \bar{w} ; we normalize the principal's outside option to 0.

Let \bar{y} and \underline{y} denote the expected values of output when E and S are chosen, respectively. We assume that $\bar{y} - c > \bar{w} > \underline{y}$. Thus, it is efficient for the agent to be employed and to exert effort (the inequalities imply that $(p - q)(G - B) > c$),

⁴Our analysis extends to any finite signal space Y . The proofs of our negative results hold in that case. For our positive results, if we order the signals y_n from lowest to highest likelihood ratio $\Pr(y_n|E) / \Pr(y_n|S)$, then we may focus on a binary partition of the signal space, $\{G, B\}$, where $G = \{y_n \in Y | n \geq \bar{n}\}$ and $B = \{y_n \in Y | n < \bar{n}\}$ for some cutoff \bar{n} .

but if the agent shirks, then it is preferable to dissolve the relationship. Both players are risk neutral, and they face no limited liability constraints. They maximize the discounted sum of payoffs, with common discount factor $\delta \in (0, 1)$.

In the interest of precision, let us consider the following stage game Γ that is played in every period, conditional on the relationship not having been terminated.

- The agent is paid the base wage w and chooses $a \in \{E, S\}$.
- The principal observes $y \in \{G, B\}$ and decides whether or not to pay a bonus, over and above the base wage w .
- The principal and agent observe the realization of a public randomization device⁵ and simultaneously decide whether or not to terminate the relationship – the relationship continues to the next period if and only if both parties want to continue.

We denote the game that is repeated infinitely, unless terminated by a player, by Γ^∞ .

The fundamental difficulty is that monitoring is imperfect and private. The principal does not observe the agent's action, and the agent does not observe the signal y . In order to incentivize effort, the agent's bonus payments (or his continuation value) must depend upon the principal's observation of output. However, because this observation is private, the principal's continuation value must be independent of the output the principal observes. MacLeod (2003) and Levin (2003) propose a solution that the principal is indifferent between paying the bonus or not paying it. This indifference can be achieved via a public randomization device that decrees that

⁵That is, the realization of a random variable that is uniformly distributed on interval $[0, 1]$.

the relationship (which is profitable for the principal) be dissolved with some probability whenever the bonus is not paid. In other words, a part of the expected surplus from the relationship must be destroyed, because the agent cannot be punished while simultaneously rewarding the principal. As usual, that surplus destruction could also take the form of the principal burning money or giving it to a third party that neither the principal nor the agent cares about.

The equilibrium is inefficient, because some surplus is destroyed. Fuchs (2007) shows that inefficiency can be mitigated if the players are patient, by dividing the interaction into blocks of T periods. The bonus is withheld only if the agent fails (that is, output is B) in every period in the block, and this way of leveraging a single bonus to incentivize effort in multiple periods reduces the loss in surplus. Fuchs (2007) also shows that the most efficient equilibrium for a fixed discount factor is an efficiency wage type equilibrium where the only feedback provided by the principal is when she terminates the relationship.

3 Purifiability

The major problem with the block equilibrium (and also the one-period construction) is that it relies on the principal's indifference between paying the bonus and not paying it, and on her breaking this indifference according to the history of private signals. Consequently, the equilibrium is fragile. In particular, if the value of the relationship or the outside option to the principal is subject to small shocks that are privately observed by the principal, then she will condition her bonus payment on the realization of these shocks, and not upon output signals. This problem also arises in the no-feedback efficiency wage construction – in any period t where the principal fires the agent with positive probability, she must be indifferent between retaining

the agent and firing, and this indifference is untenable with shocks to the principal's outside option.

To make this argument precise, we will also study the following class of ξ -perturbed games $\Gamma^\infty(\xi, \lambda)$, where ξ and λ are positive numbers, defined as follows:

- Before the principal makes her bonus decision, the principal privately observes a random shock z_t . The principal's payoff from output in period t is given by $B + \xi z_t$ and $G + \xi z_t$.⁶

We assume that z_t follow a stationary autoregressive processes of order 1, that is, $z_t = \lambda z_{t-1} + (1 - \lambda)\tilde{z}_t$, where $\lambda \in (0, 1)$ and \tilde{z}_t are independent random variables, distributed identically and continuously on a bounded interval Z with mean zero. For completeness, let $z_0 = \tilde{z}_0$. So, private shocks display some persistence, although it may be arbitrarily small. We elaborate on the reason for the persistence assumption after stating our first negative result.

The reader may wonder why variables other than output are not perturbed. For example, the agent could privately observe a random shock to his effort cost before he chooses his action, or both players could privately observe random shocks to the value of their outside options. Adding these other shocks would not affect any of our results, only making the analysis more involved in terms of notation. It should be easy to see intuitively that the problem in constructing robust equilibria comes from the principal's indifference between paying or not paying the bonus, independent of her signals. In the appendix, we show that our equilibrium construction in Section 5, where the agent randomizes across his actions, can be purified when we add a random shock to the agent's cost of effort.

⁶Output specific shocks would not make any difference.

The reader may also ask, why do we assume that the shocks to the value of output are history-independent? The reason is, we want to respect the structure of the repeated game, where histories do not directly affect stage-game payoffs. Indeed, if we assume that realized output at date t affects the distribution of z_t , then the results will depend upon the nature of this dependence. For example, if the distribution of z_t conditional on G first-order stochastically dominates that conditional on B , this will provide the principal with some incentives to report her signals truthfully in a pure strategy equilibrium – she will be more likely to report G after G than after B . In turn, if the distribution of z_t conditional on G is dominated, her reports will be negatively correlated with the truth.

Our payoff shocks have bounded support, and by taking ξ to be small, their effective range can be made arbitrarily small. Nonetheless, in their presence, it is difficult to sustain an equilibrium where the worker exerts effort on the equilibrium path. We will now present two results showing that no effort can be sustained in some classes of equilibria, which both include the equilibria studied in the existing literature, if we require in addition these equilibria to be purifiable.

To make this precise, the following preliminaries are necessary. In the perturbed game, in any period t , the agent alone observes his effort choice a_t . The principal alone observes output y_t and the shock to output, z_t . Both parties observe the public events, such as the bonus payment and the realizations of the public randomization device, which we denote by $\omega_t \in \Omega$. Termination decisions are also public, but since the game ends if one party chooses to terminate, we restrict attention to histories where both parties have chosen to continue the relationship to date. Let Ω^{t-1} denote the set of possible public histories at the beginning of period t .

Let $\sigma = (\sigma_t)_{t=1}^\infty$ denote the strategy of the principal. A strategy for the principal prescribes a bonus payment and a firing decision at all possible histories. Thus σ_t

consists of a pair (σ_t^1, σ_t^2) . $\sigma_t^1 : \Omega^{t-1} \times \{G, B\}^t \times Z \rightarrow [0, \infty)$ determines the bonus payment. This depends upon the public history to date, the history of observed outputs, and the current value of the shock to flow revenues, z_t . $\sigma_t^2 : \Omega^{t-1} \times [0, \infty) \times \{G, B\}^t \times Z \times [0, 1] \rightarrow \{F, R\}$ determines the principal's firing/retention decision, which is based additionally upon the principal's bonus payment this period and the realization of the public randomization.⁷

Let $\rho = (\rho_t)_{t=1}^\infty$ denote the strategy of the agent. A strategy for the agent prescribes an effort choice and a quitting decision at all possible histories. Thus ρ_t consists of a pair (ρ_t^1, ρ_t^2) . $\rho_t^1 : \Omega^{t-1} \times \{E, S\}^{t-1} \rightarrow \{E, S\}$ determines the agent's effort choice. This depends upon the public history to date and the agent's history of effort choices. $\rho_t^2 : \Omega^{t-1} \times [0, \infty) \times \{E, S\}^t \times [0, 1] \rightarrow \{Q, C\}$ determines his quitting/continuation decision, which is based additionally upon his current period effort choice, the principal's bonus payment this period, and the realization of the public randomization.

Note that the agent's strategies are the same mathematical objects in the original game Γ^∞ and in the perturbed game $\Gamma^\infty(\xi, \lambda)$. In turn, the principal's strategies in these two games are formally different objects. So, to compare the equilibria of the two games, we consider expected strategies of the principal in the perturbed game. More precisely, let h_t^i denote a generic history of public events and outputs observed by the principal up to period t and stage $i = 1, 2$. Let $\tilde{\sigma}_t^i(h_t^i) = \int \sigma^i(h_t^i, z_t) d\mu(z_t)$, where μ is the distribution of z_t , denote the expected behavior of the principal at h_t^i , taking expectations over the payoff shocks. Given a sequence of $(\xi^k)_{k=1}^\infty$ that converges to 0, and a sequence of perturbed game strategies $(\sigma^k, \rho^k)_{k=1}^\infty$, the associated behavior

⁷A player's private history also includes past values of the payoff shock, but since these are payoff irrelevant, the principal will not condition on these shocks in the perturbed game – see Bhaskar et al. (2013), Lemma 2.

of the latter sequence converges to a strategy profile of the unperturbed game $(\tilde{\sigma}, \tilde{\rho})$ if for all i and t and each private history h_t^i , $\tilde{\sigma}_t^k(h_t^i)$ converges in distribution to $\tilde{\sigma}_t(h_t^i)$ and $\tilde{\rho}_t^k(h_t^i)$ converges in distribution to $\tilde{\rho}_t(h_t^i)$.

An equilibrium $(\tilde{\sigma}, \tilde{\rho})$ of the original game Γ^∞ is purifiable if for any sequences of shocks $(\xi^k)_{k=1}^\infty$ and of persistence parameters $(\lambda^k)_{k=1}^\infty$ that converge to 0, there is a sequence of equilibria (σ^k, ρ^k) of the perturbed game $\Gamma^\infty(\xi^k, \lambda^k)$ such that the associated behavior converges to $(\tilde{\sigma}, \tilde{\rho})$.

We will now present negative results for two classes of equilibria. The first one concerns equilibria with deterministic effort choice, i.e. effort choice that is deterministic on path. More precisely, given the agent's strategy ρ_1 (of the original game Γ^∞), the agent's period 1 effort choice is said to be *deterministic in period 1* if ρ_1^1 prescribes a pure action.⁸ Defining recursively, the agent's period t effort choice is *deterministic in period t* if

- for every $s < t$, his effort is deterministic in period s , and
- for every public history ω_{t-1} and the unique sequence of on-path private actions along history ω_{t-1} , strategy ρ_t^1 prescribes a pure action.

The agent's effort choice is *deterministic* if it is deterministic in all periods.

Proposition 1. In any Γ^∞ purifiable equilibrium where the agent's effort choice is deterministic, the agent always shirks when hired and the principal never pays a bonus.

The idea of the proof of this result (in the appendix) is the following. If the agent's effort choice is deterministic, then the principal's privately observed outputs

⁸In other words, the agent's effort choice is deterministic from the point of view of the principal, who does not observe the agent's choice of action. From the point of view of the agent, his effort choice will always be deterministic, as required by purifiability.

convey no information regarding the agent's continuation strategy. However, the payoff shocks affect the principal's value from continuing the relationship, and she is almost never indifferent between distinct bonus payments and the agent's continuation strategies induced by the payments. So, the principal's payoff shocks, not outputs, determine her choice of bonus payments, and the agent lacks incentives for providing effort.

We could, in principle, allow the principal to report her shock as well as the output signal. However, there is no single-crossing relationship to exploit. To induce truthful reporting of the shock, those with lower shocks must be allowed to pay lower bonuses after reporting signal G , but then those with higher would mimic those with lower shocks.

This result and its proof are quite different from those in Bhaskar et al. (2013). The most important difference is that the game considered here has private monitoring, which raises new issues – the results in Bhaskar et al. (2013) assume public monitoring. The result only applies to equilibria where the worker plays a pure strategy, but does not assume finite memory.

The proposition implies that the equilibria in Fuchs (2007) – the T -period block equilibria and the no-feedback efficiency wage equilibria – are not robust to payoff shocks, no matter how small. A similar problem arises with other methods for maintaining the principal's indifference, such as various ways of burning money. For example, the principal would be willing to condition the bonus payment on her private signal of output if she is made indifferent between paying the bonus and contributing to charity. Once again, though, incentive compatibility is destroyed by the slightest deviation of the principal's preferences from exact indifference, such as a small, privately known variation in her valuation of the charity. It can be shown that Proposition 1 would apply in this version of the model as well. Thus the argument applies

also to MacLeod's (2003) analysis of formal contracting with subjective performance evaluation.

Two further observations are in order. First, it suffices to assume that there are shocks to the principal's payoff, since shocks to the worker's payoff would play no role in the proof. Second, the assumption that shocks to the principal's flow payoff are autoregressive is important. If the shocks were i.i.d, a minor modification to the T -period block strategy equilibrium, where we allow the principal to make an announcement, would suffice. The principal could announce at the end of a block whether the agent had passed or failed the test, but both the payment of the bonus and the stochastic termination of the relationship could be deferred for $K > 1$ periods. This ensures that the principal would be indifferent between paying and not paying the bonus, since she has no private information about the payoff consequences of her announcement. The assumption that shocks are ever so slightly persistent ensures that the principal is never indifferent between paying and not paying the bonus, no matter what the timing of these events.

The second negative result concerns block-strategy equilibria. An equilibrium $(\tilde{\sigma}, \tilde{\rho})$ of the repeated game Γ^∞ is a T -period *block-strategy equilibrium* if time can be partitioned into intervals of length T such that:

- If the relationship is in effect in the first period of a block, it is never terminated until the end of the block.
- Bonus payments are made by the principal only in the last period of each block.
- The strategies within a block depend only upon the events within the block.

Block-strategy equilibria are not purifiable, even if the agent's effort is not deterministic. In a block-strategy equilibrium, the agent's play in each block is independent

of events in previous blocks. Consequently, the principal can be incentivized to make distinct bonus payments as a function of her signals only by varying termination probabilities at the end of the block. However, with payoff shocks, the principal will condition her bonus payments on the shock realization and not on the output signals. The proof is in the appendix.

Proposition 2. No effort can be sustained in any purifiable T -period block-strategy equilibrium of Γ^∞ , for any finite T .

The existence of purifiable equilibria in which the agent randomizes between E and S is an open question. It is possible that such an equilibrium can be constructed as follows: Since the agent randomizes, the principal's output is an informative signal about the agent's action. The principal can be incentivized to pay a bonus contingent on output G , and to pay no bonus contingent on output B by the threat that an action inconsistent with the agent's action (i.e. not paying a bonus combined with action E , or paying the bonus combined with action S) will result in a less cooperative behavior of the agent than an action consistent with the agent's action. This, however, would require the agent to randomize in a history-dependent manner, and such a strategy is usually difficult to purify.

Our strategy in the subsequent sections will be to construct equilibria where both players – in particular, the principal – have strict incentives to condition upon their private signals. This ensures robustness to any small payoff shocks. The equilibria will require the agent to randomize effort choices, and to shirk with positive probability in every period. In addition, we will allow for simultaneous communication.

4 Simple, more efficient and purifiable equilibria with cheap talk

We modify the stage game Γ set out in Section 2 by allowing the principal and the agent to make simultaneous cheap-talk announcements. Specifically, after the agent has chosen effort and the principal has observed output, there is a cheap-talk stage within the period where the parties announce their private information. We denote this new game by $\tilde{\Gamma}$. We find that cheap talk allows effort to be sustained in equilibrium. In this section, we construct a purifiable equilibrium in $\tilde{\Gamma}^\infty$, the game $\tilde{\Gamma}$ repeated infinitely unless terminated, and we show that, as long as the discount factor is above a cutoff, then that equilibrium is approximately efficient when the noise in monitoring is small. For fixed noise, we modify the equilibrium construction in the next section so that players make announcements only every T periods; that construction will yield approximate efficiency when the players are very patient even when the noise is large.⁹

In the equilibrium that we construct, in each period the agent reports his action choice $\hat{a} \in \{E, S\}$, and the principal reports the output that she observed $\hat{y} \in \{G, B\}$. As a function of the pair of announcements (\hat{a}, \hat{y}) , the principal makes a bonus payment to the agent, and the relationship is terminated with some probability. The relationship is also terminated after any observable deviation.

We will first construct an equilibrium without message exchange, in which the principal is indifferent between paying the bonus to the agent and probabilistic termination. Then we will supplement that construction with reports in a way that will

⁹We note that our results still hold if making announcements has a small cost that may vary across announcements. The reason is that in our constructions, players have strict incentives for truthful reporting.

provide the principal with strict incentives contingent on the signal she obtains. The message exchange will also give the players strict incentives to report truthfully, but it will not affect the qualitative properties of the baseline construction, because it will have only marginal effects on actions and payoffs.

Our equilibrium requires the agent to choose both actions with positive probability. Let $\phi \in (0, 1)$ denote the probability that he chooses E . If the principal reports outcome G , then she pays a bonus τ to the agent. Otherwise the bonus is zero. To make the agent indifferent between the two actions the size of the bonus is chosen so that

$$c = (p - q)\tau; \tag{1}$$

If the principal reports B , then the relationship is terminated with a probability α . This α will be chosen so that the principal is indifferent between paying the bonus and the chance of losing continuation value V^P , i.e.

$$(1 - \delta)\tau = \delta\alpha V^P. \tag{2}$$

In other words, if the principal does not pay the bonus τ , the loss in continuation value that she suffers in consequence is equal to τ . The agent is also paid w up front as a base wage, so that his overall payoff equals \bar{w} , his outside option. That is,

$$w = \bar{w} - \frac{qc}{p - q}. \tag{3}$$

When α is chosen so that equation (2) is satisfied, the principal's value V^P is equal to:

$$V^P = (1 - \delta)[\phi(\bar{y} - \bar{w} - c) + (1 - \phi)(\underline{y} - \bar{w})] + \delta V^P - \delta[\phi(1 - p) + (1 - \phi)(1 - q)]\alpha V^P,$$

where the last term reflects the loss in continuation value when signal B is observed. Using the agent's indifference condition (1) and the principal's indifference condition (2), the principal's value can be re-written as

$$V^P = [\phi(\bar{y} - \bar{w} - c) + (1 - \phi)(\underline{y} - \bar{w})] - [\phi(1 - p) + (1 - \phi)(1 - q)]\frac{c}{p - q}.$$

Note that this expression depends only on the parameters of the model and the action of the agent. So, α can now be defined by (2) and (1) in which V^P takes the above value. Of course, α cannot exceed 1. This will be implied by two conditions. First, since we are interested in attaining efficiency, we look for equilibria in which ϕ is close to 1. By assuming that

$$\bar{V}^P := \bar{y} - \bar{w} - c - \frac{(1 - p)c}{p - q} > 0,$$

we guarantee that $V^P > 0$ for the values of ϕ that are close enough to 1. Further, if $\bar{\delta}$ is defined by

$$(1 - \bar{\delta})\frac{c}{p - q} = \bar{\delta}\bar{V}^P,$$

then it follows from (2) and (1) that $\alpha < 1$ for any $\delta > \bar{\delta}$.

We will now supplement the construction with a simultaneous exchange of messages, which will provide both the agent and the principal strict incentives for reporting signals truthfully, while keeping the agent indifferent between playing E and

playing S . In order to provide the agent with strict incentives for truth telling, we let the principal retain a portion of the base wage so that an amount x_{ij} is deducted from the agent's compensation if the agent reports $i \in \{E, S\}$ and the principal reports $j = \{G, B\}$ during the message exchange. The rest of the retained portion is returned to the agent before the realization of the public randomization is observed. The relationship is terminated with an additional probability α_{ij} . This α_{ij} will be chosen so that the flow equivalent of the loss in continuation value suffered by the principal equals $2x_{ij}$, i.e:

$$\delta\alpha_{ij}\tilde{V}^P = 2(1 - \delta)x_{ij}. \quad (4)$$

The numbers x_{ij} must satisfy the following conditions:

$$px_{EG} + (1 - p)x_{EB} < px_{SG} + (1 - p)x_{SB},$$

$$qx_{SG} + (1 - q)x_{SB} < qx_{EG} + (1 - q)x_{EB},$$

$$px_{EG} + (1 - p)x_{EB} = qx_{SG} + (1 - q)x_{SB},$$

$$\frac{\phi(1 - p)x_{EB} + (1 - \phi)(1 - q)x_{SB}}{\phi(1 - p) + (1 - \phi)(1 - q)} < \frac{\phi(1 - p)x_{EG} + (1 - \phi)(1 - q)x_{SG}}{\phi(1 - p) + (1 - \phi)(1 - q)},$$

$$\frac{\phi px_{EG} + (1 - \phi)qx_{SG}}{\phi p + (1 - \phi)q} < \frac{\phi px_{EB} + (1 - \phi)qx_{SB}}{\phi p + (1 - \phi)q}.$$

The first two conditions guarantee strict incentives of the agent for truthful reporting. The third condition maintains the agent's indifference between playing E and playing S . Finally, the last two conditions guarantee strict incentives of the principal for truthful reporting.

For ϕ strictly less than but close to 1, we can satisfy the five conditions as follows.

Set $x_{SB} = 0$ and choose $\epsilon > 0$. Then set

$$x_{SG} = \phi\epsilon, x_{EB} = r\epsilon, x_{EG} = \epsilon(\phi - 1 + r),$$

where $r = p(1 - \phi) + q\phi$. Observe that if ϕ is close to 1, then $x_{ij} \geq 0$ for all i and j , and $x_{SB} < x_{SG}$ and $x_{EG} < x_{EB}$. Roughly speaking, this construction provides incentives for truthful reporting by penalizing the report mismatch; that is, when the agent reports E while the principal reports B , or when the agent reports S and the principal reports G . Note, however, that numbers x_{ij} must be chosen carefully, not to break the agent's indifference between playing E and playing S . In addition, the base wage is slightly increased to compensate the agent for the expected wage deductions by the principal after messages have been exchanged.

We now write out the principal's value as

$$\tilde{V}^P = (1 - \delta)[\phi(\bar{y} - c) + (1 - \phi)\underline{y} - \bar{w}]$$

$$+ \delta\{1 - [\phi(1 - p) + (1 - \phi)(1 - q)]\alpha + \phi(1 - p)\alpha_{EB} + \phi p\alpha_{EG} + (1 - \phi)q\alpha_{SG} + (1 - \phi)(1 - q)\alpha_{SB}\}\tilde{V}^P.$$

Using the analogues of (1) and (2) for $\tilde{\Gamma}$ and to substitute for $\alpha\tilde{V}^P$ and using (4) to substitute for $\alpha_{ij}\tilde{V}^P$, we obtain

$$\begin{aligned} \tilde{V}^P &= \phi(\bar{y} - c) + (1 - \phi)\underline{y} - \bar{w} - [\phi(1 - p) + (1 - \phi)(1 - q)]\frac{c}{p - q} \\ &\quad - 2\{\phi[(1 - p)x_{EB} + px_{EG}] + (1 - \phi)[qx_{SG} + (1 - q)x_{SB}]\}. \end{aligned}$$

Note that this expression depends only ϕ on and ϵ , and the parameters of the model. So, α can now be defined by the analogues of (2) and (1) for $\tilde{\Gamma}$ and α_{ij} can be defined by (4) in which \tilde{V}^P takes the above value. As $\phi \rightarrow 1$ and $\epsilon \rightarrow 0$, $\tilde{V}^P \rightarrow \bar{V}^P$, so $\alpha, \alpha_{ij} < 1$ for such ϕ and ϵ , and any $\delta > \bar{\delta}$.

It remains to show that this equilibrium is purifiable. Both players have strict incentives at the cheap-talk stage, and the principal has strict incentives at the bonus payment stage. The only remaining question is whether the agent's random choice of effort can be purified. In our model, the randomness in the agent's choice imposes no problem, because we did not perturb the agent's cost. However, the agent's choice would be purified even if his cost was perturbed. In $\Gamma^\infty(\xi, \lambda)$, the agent could be simply prescribed action S if his cost exceeds some high threshold.¹⁰

For ϕ close to 1 and ϵ close to zero, the principal's value of the relationship is close to \bar{V}^P . This is less than $\bar{y} - \bar{w} - c$, the highest possible surplus. Some surplus must be burned because monitoring is imperfect, and the relationship must be terminated with positive probability since otherwise she would have an incentive to always claim that she observed low output. These payoff losses arise even in the non-purifiable equilibrium of Levin (2003), where the worker chooses effort for sure, and where there is no cheap talk. We therefore have the following proposition.

Proposition 3. Assume $\bar{V}^P > 0$ and $\delta > \bar{\delta}$. For every $\Delta > 0$, there exists a purifiable equilibrium with cheap talk such that the principal's expected payoff is at least $\bar{V}^P - \Delta$.

Proposition 3 requires the discount factor δ to be strictly greater than some $\bar{\delta} < 1$, since otherwise there are no feasible values of αV^P that satisfy the principal's

¹⁰Notice that if this threshold is fixed over time, then λ must converge to 0 as the perturbed games converge to the original game, because otherwise the agent's actions would be correlated across periods in the unperturbed game.

indifference condition 2. In other words, if the principal is too impatient, she cannot be incentivized to pay the bonus.

The additional inefficiency imposed by the purifiability requirement is two-fold – first, the worker must shirk with positive probability, and second, there is an additional probability of termination induced by the phase of message exchange. Nonetheless, these costs can be arbitrarily small, since both the shirking probability $1 - \phi$ and the additional termination probabilities α_{ij} can be arbitrarily small.

Finally, as the noise in monitoring vanishes, efficiency is achievable; more generally, we have the following corollary.

Corollary 1. Assume $\bar{V}^P > 0$ and $\delta > \bar{\delta}$. For every $\Delta > 0$, and $(1 - p)/(p - q)$ is sufficiently close to 0, then there exists a purifiable equilibrium with cheap talk such that the principal’s expected payoff is at least $\bar{y} - \bar{w} - c - \Delta$.

5 T -period block equilibria with cheap talk

With non-vanishing noise, Fuchs (2007) shows that patient players can approach efficient outcomes by dividing play into blocks of T periods. At the end of the block, the principal pays a bonus to the agent unless the output was B in all T periods. We find that adding cheap talk can deliver a similar efficiency result but in a purifiable equilibrium; in addition, the second equilibrium that we construct has qualitatively different (perhaps, more realistic) features.

As in the one-period case in the previous section, we generate incentives for truth-telling by penalizing mismatched reports. That approach requires non-trivial randomization by the agent. Let us examine the conditions that must be satisfied in a T -period construction. Suppose that the equilibrium requires the agent to choose

E with high probability in every period of the block. Then the play of E in any period t of the block must be incentivized, meaning that the principal’s reporting decision must depend on whether signal G or B is realized in that period. However, if the principal is to have strict reporting incentives, then the t -th period signal must be informative of the agent’s behavior (and hence the agent’s report). That link is possible only if the agent chooses both E and S with positive probability in the t -th period of the block. In other words, both E and S must be played with positive probability in each period of the block.

We present two approaches for constructing that period-by-period uncertainty in purifiable T -period equilibria. Our first approach builds on Fuchs (2007). The agent plays S in at most one period of the block, but it could be any period. In the second approach, the agent plays S either in all T periods or in none. The latter has the advantage that it also delivers more realistic contracts, since the bonus need not be extremely high and the base wage need not be negative. Both of those extreme approaches may be unrealistic in applications, but it is straightforward to use them to construct intermediate cases (where, for example, the agent may put in low effort for a week while unexpectedly busy at home).

5.1 Strict incentives for truth-telling in block-strategy equilibria

In this section we provide a general construction for modifying block-strategy equilibria to purifiable equilibria, with a minimal loss in terms of their efficiency. Let (σ^*, ρ^*) be a T -period block equilibrium of the baseline environment Γ^∞ with no exchange of messages, with the property that in equilibrium the agent exerts effort in every period. In particular, (σ^*, ρ^*) specifies a “test” that maps the principal’s

observed signals to a bonus and a termination probability. We then augment (σ^*, ρ^*) roughly as follows: the agent randomly chooses a period t^S in which to shirk, and he announces t^S at the end of the block. The principal simultaneously announces her signal from each of the T periods. Then the announced period t^S is “thrown out,” and the test from (σ^*, ρ^*) is run on the remaining periods. That rule, with some adjustments, preserves the effort incentives from (σ^*, ρ^*) .¹¹ Finally, in order to provide incentives for truth-telling we use small adjustments to bonus payments.

Formally, define a *T-period review strategy profile* $(\sigma^*, \rho^*) = (\tau, \chi)$ of the baseline game as follows. As a function of the number of G signals n in the principal’s private T -period history of signals, (σ^*, ρ^*) specifies whether or not the agent passes the test, $\chi(n) \in \{0, 1\}$. If the agent passes, then the principal pays him a bonus τ . If the agent fails the test, then the relationship is terminated with a probability α so that the flow value of the payoff loss to the principal equals τ .¹² The agent’s strategy is to play E in every period. A *strict T-period review strategy equilibrium* $(\sigma^*, \rho^*) = (\tau, \chi)$ is one in which the agent has a strict incentive to play E in each period.

Next, define an *augmentation* of $(\sigma^*, \rho^*) = (\tau, \chi)$ as follows.

5.1.1 Actions and Reports

The horizon is divided into blocks of length T each. The agent chooses exactly one period in which she plays S by a fair lottery over the T periods plus a fictitious $(T + 1)$ -st period. In all other periods, the agent plays E . If the fictitious period $T + 1$ is selected by the lottery, then the agent provides effort in all T periods.

¹¹These adjustments are necessary and essential, because otherwise the principal who obtained signal B in all periods of a block could not be given strict incentives for reporting her signals truthfully, and the augmented equilibrium would not be purifiable.

¹²That is, $(1 - \delta)\tau = \delta\alpha V^P$, where V^P denotes the value to the principal of the T -period game played repeatedly over the infinite horizon.

At the end of each block, the agent reports to the principal the period t^S that she chose for playing S , and the principal reveals the signals that she obtained in the T periods of the block. The “signal” of the principal in the fictitious period $T + 1$ is generated by public randomization. More precisely, if the agent plays S in one of the T periods, then the public randomization generates in the fictitious period $T + 1$ a signal in the way the signal would be generated by the agent’s action E in that period. If $t^S = T + 1$, then the public randomization generates a signal in the way the signal would be generated when the agent played action S . The reports are simultaneous, and the realization of public randomization is observed after the reports. It is understood that the agent must report action E in all but one of the $T + 1$ periods; in the remaining period she must report action S .

5.1.2 Agent’s Review

Next, the agent is subject to the test χ . The period t^S does not count for the review. All other periods count, including the fictitious one if this is not the period in which S was prescribed. The rule χ applied to the T periods that count determines whether or not the agent passes the test. If the agent passes the test, then he is paid the bonus τ . If the agent fails the test, then the principal pays no bonus, but the relationship is terminated with a probability such that the flow cost to the principal equals τ . No matter whether the agent passes or fails, he is paid an adjustment. The adjustment compensates the agent for playing S in a later rather than an earlier period, or for not playing S at all (that is, for choosing the fictitious period $T + 1$), so that the agent is indifferent regarding the period in which he plays S . More precisely, the adjustment is $(1 - \delta^{t^S - 1})c/\delta^{T-1}$ if $t^S \leq T$, and it is c/δ^{T-1} if $t^S = T + 1$.

Finally, it will be convenient to assume that the value of the agent is \bar{w} . This is

without loss of generality, because the base wage and the bonus can be adjusted to guarantee this condition.

5.1.3 Testing Principal's Reports

Finally, the principal's report is tested. More specifically, an ε is added to the agent's bonus, no matter what the outcome of the test. That is, the bonus becomes $\tau + \varepsilon$ if the agent passes the test, and it is ε if the agent fails. This ε is next subtracted if the principal's report passes the following test:

(1) A period other than t^S is drawn in a fair lottery over the remaining T periods (including the fictitious period $T + 1$ if $t^S \neq T + 1$).

(2) A fifty-fifty lottery chooses one of the two: (a) t^S ; or (b) the period chosen in (1) (in which it is assumed that the agent reported E).

(3) If the signal reported by the principal and the agent's action in the selected period coincide, that is, they are $\{G, E\}$ or $\{B, S\}$, then the principal passes. Otherwise, that is, if they are $\{G, S\}$ or $\{B, E\}$, then the principal fails.

The reader may wonder why the principal is tested in this specific manner. This specific test guarantees that at the beginning of a block the principal assigns a fifty-fifty chance to both actions in the period that will be taken for the test. Therefore the signals obtained by the principal are more likely to coincide with the action taken for the test than the opposite signals.

The reader may also wonder what is the role of the fictitious period. If we did not include the fictitious period, and the agent was prescribed to choose one of the T block periods for shirking, then the principal with all signals B (or with all signals G) would still assign a fifty-fifty chance to both actions in the period taken for the test. Therefore the principal would not have strict incentives for reporting her signals

truthfully.

With those definitions, we can present the result for this approach to constructing a robust and efficient equilibrium.

Lemma 4. Fix δ and T , and suppose that $(\sigma^*, \rho^*) = (\tau, \chi)$ is a strict T -period review strategy equilibrium of the game without cheap talk. Then for high enough δ , there is an augmentation of (σ^*, ρ^*) that is a purifiable equilibrium of the game with cheap talk.

The proof is in the appendix.

Proposition 5. Let $\Delta > 0$. Then there exists T such that for $\delta < 1$ but close enough to 1 there is a purifiable equilibrium in T -period block strategies with cheap talk where the agent exerts effort in at least $T - 1$ periods of the block, and the principal's expected payoff is at least $\bar{y} - \bar{w} - c - \Delta$.

Proof. The proposition follows from applying Lemma 4 to the strategies in Fuchs (2007), plus the observation that when players become patient and T is high enough, a single period of playing S in every block creates only a negligible efficiency loss, as well as having a negligible effect on the principal's payoff. \square

5.2 Shirking in 0 or T periods

In our second approach, the agent randomizes between exerting effort in all T periods and shirking in all T periods. That is, the agent must be indifferent between the two sequences – “always E ” and “always S ” – at the beginning of the block.

The cheapest way to make the agent indifferent between the two sequences is to punish the agent if and only if all signals in the block are B . However, if indifference between the two effort sequences is achieved in this manner, then the agent strictly

prefers to choose E in the first period, and S in subsequent periods, to either of these alternatives. Indeed, the marginal reduction in the probability of T bad signals from exerting effort in $k + 1$ rather than k periods¹³ is decreasing in k , because the probability of having gotten at least one success already is increasing in k . If the worker weakly prefers choosing E in all T periods over $T - 1$, then he must strictly prefer working in 1 period over “always S .” Conversely, if the punishment after the signal realization of “all B ” is increased, so as to deter this deviation, then the agent will find “always S ” to be inferior to “always E ,” and will not randomize in the required manner.

In other words, if the agent is to be made indifferent between these two sequences, and also deterred from deviating to other action sequences, then the agent must also be punished after intermediate sequences of signal realizations – i.e., sequences where some signals are good while the others are bad. Such punishment directly reduces the efficiency of equilibria for any fixed discount factor. Nevertheless, as δ approaches 1, there is a sequence of purifiable equilibria whose payoffs asymptote to efficient payoffs.

Proposition 4. Let $\Delta > 0$. Then there exists T such that for high enough δ there is a purifiable equilibrium of Γ^∞ with cheap talk in T -period block strategies, where the agent exerts effort in all T periods or in none, and the principal’s expected payoff is at least $\bar{y} - \bar{w} - c - \Delta$.

As before, the agent’s payoff will be held at his outside option. In addition, our equilibria will have two, presumably attractive, features: (a) the per-period payments of the principal will be close to the per-period cost of the agent’s effort; (b) getting a bonus will be highly unlikely when the agent never works, while it will be almost

¹³That reduction is $(1 - p)^k (1 - q)^{T-k} - (1 - p)^{k+1} (1 - q)^{T-1-k} = (p - q) (1 - p)^k (1 - q)^{T-1-k}$.

certain when the agent always works.

The key idea is to set the cutoff number of good outcomes for paying a bonus at an intermediate value, between fraction q and fraction p of the T periods. Then the agent's marginal benefit of choosing E in $k + 1$ rather than k periods – the increased probability of reaching the cutoff times the amount of the bonus – is first increasing and then decreasing in k . Intuitively, for low values of k , the total number of successes in the other $T - 1$ periods is likely to be well below the cutoff, so one additional period of choosing E will not make a big enough difference to matter, and for high k the cutoff is likely to be surpassed already. That inverse U -shaped pattern of marginal benefits makes it possible to set a cutoff \underline{n} and a bonus τ that satisfy two conditions: (1) the sum of all T marginal benefits just compensates the agent for the cost of working in T periods; and (2) for any $k < T$ the sum of the first k marginal benefits falls below the cost of working in k periods. The first condition means that the worker is indifferent between “always E ” and “always S ,” and the second condition means that both those choices are best responses. To satisfy the two conditions, we choose the cutoff \underline{n} so that, no matter what the size of the bonus, the agent's marginal benefit of choosing E in 1 rather than 0 falls below the agent's marginal benefit of choosing E in T rather than $T - 1$ periods. Intuitively, by moving the cutoff from a small fraction of the T periods to a large fraction of the T periods, we gradually decrease the former benefit, and we gradually increase the latter benefit, until the latter benefit exceeds the former benefit. Then we choose the bonus τ that satisfies condition (1). For this τ , the marginal benefit crosses the cost line only once, from below, and this implies that condition (2) is satisfied.

The two conditions would be satisfied for $\delta = 1$ for example if the marginal benefit of choosing E in T rather than $T - 1$ periods is equal to or is a little higher than c , as shown in Figure 1 below (where $T = 20$), and the sum of marginal benefits over

$k = 1, \dots, T$ is equal to Tc . Intuitively, we can satisfy these two conditions with two choice variables: a cutoff \underline{n} and a bonus τ .

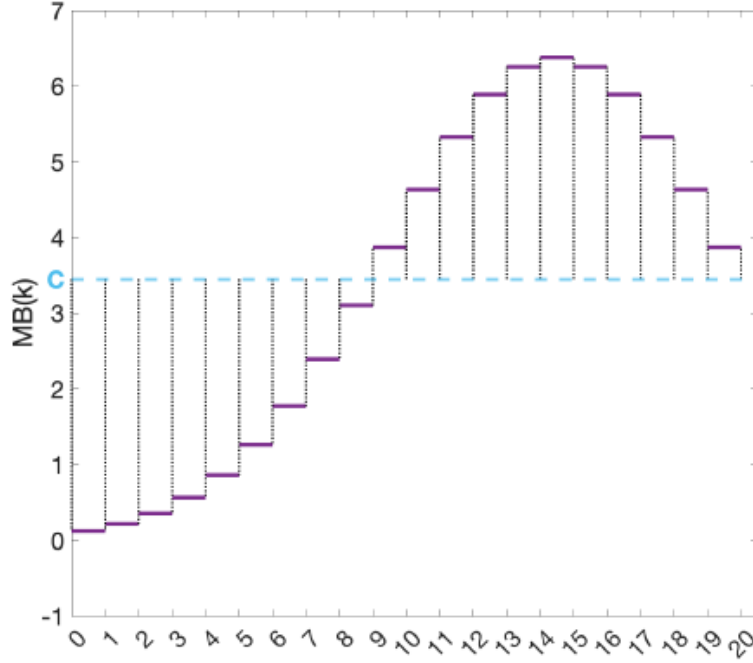


Figure 1. An example of $MB(k)$, the agent’s marginal benefit of choosing E in k rather than $k - 1$ periods, that satisfies our conditions for $\delta = 1$. The graph takes value $MB(k)$ on the interval between $k - 1$ and k . It has an inverse U shape, and the area under the graph is equal to Tc .

Before describing the equilibria more precisely and proving their properties, we state two lemmas, which display some properties of binomial distributions. They are known and easy to prove, so their proofs will be omitted. (A version of these lemmas was first noted in Matsushima (2004); see Ely et al. (2005) for perhaps simpler proofs.) Denote by $F(\underline{n}, T, k)$ the probability of the event that the principal receives more than \underline{n} signals G contingent on the agent taking action E in exactly k periods

of a block of T periods. Denote also by $f(\underline{n}, T, k)$ the probability of the event that the principal receives exactly \underline{n} signals G .

Lemma 1. For any $\varepsilon > 0$, there exists a \bar{T} such that for every $T \geq \bar{T}$, there exists an $\underline{n} = \underline{n}(T)$ such that $F(\underline{n}, T, T) > 1 - \varepsilon$, $F(\underline{n}, T, 0) < \varepsilon$, and $Tf(\underline{n}, T - 1, 0) > 1/(p - q)$.

The lemma says that there exists a cutoff number of good signals with three properties. An agent who always works will reach the cutoff with probability close to 1. An agent who always shirks will reach it with probability close to 0. And the probability of exactly hitting the cutoff in period $T - 1$ after $T - 1$ periods of shirking is sufficiently large relative to $1/T$ as T grows.

To illustrate with a numerical example, suppose that the probability of signal G contingent on action E is $p = 0.8$, and the probability of signal G contingent on action S is $q = 0.2$; let the length of the block be $T = 100$, and let the discount factor be $\delta = 0.99$. Then, the cutoff $\underline{n} = 27$ is the greatest number for which $Tf(\underline{n}, T - 1, 0) > 1/(p - q)$; $F(\underline{n}, T, 0) \approx 0.03415$ and $F(\underline{n}, T, T) \approx 1.0000$; a bonus $\tau \approx 179.5539c$ is required to make the agent indifferent between “work in all T ” and “shirk in all T ”. This ensures that the agent strictly prefers to work in every period in the block to shirking in one period and working in the remaining $T - 1$ periods.

Lemma 2. For every $T > 1$ and $\underline{n} = 0, 1, \dots, T - 1$, $f(\underline{n}, T - 1, k)$ as a function of $k = 0, \dots, T - 1$ is single-peaked, that is, if $f(\underline{n}, T - 1, k) \geq f(\underline{n}, T - 1, k + 1)$, then $f(\underline{n}, T - 1, k + 1) > f(\underline{n}, T - 1, k + 2)$ for $k = 0, \dots, T - 3$.

Let \underline{n} be any number with the properties from Lemma 1 for any given $\varepsilon > 0$. The horizon is divided into blocks of length T each. The strategy of the agent is to exert effort in every period of a block with probability ρ^* and to shirk in every period of the block with the complementary probability. At the end of the last period of

each block the agent is subject to a review. The agent passes the review when the principal obtains more than \underline{n} good signals. In this case, he obtains a bonus τ from the principal; if he fails the review, the agent obtains no bonus, and the relationship is terminated with a probability such that the flow cost to the principal equals τ . So far, we assume no message exchange, but such a message exchange will be added to the construction to guarantee purifiability.

We prescribe the bonus τ such that

$$\tau = \frac{(1 + \delta + \dots + \delta^{T-1})c}{\delta^{T-1}[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]}.$$

This guarantees that the agent is indifferent between taking action E in all periods of a block and taking action S in all periods of the block. We will next show that this implies that the agent strictly prefers taking action E in all periods of a block to taking any sequence of actions other than the two constant ones, so that the prescribed strategies are an equilibrium.

Notice first that the most profitable deviation must have the form of playing S in a number of the first periods of a block and playing E in the remaining periods of the block, because the costs of effort incurred in earlier periods have a larger effect on the agent's payoff than that of the same costs incurred in later periods, whereas the effect of playing S on the review outcome does not depend on the period in which this action is taken.

To show that it is inferior for the agent to play E in some periods and to play S in the others, notice that for any given T if δ is sufficiently large, then the difference between in the agent's payoff from playing E in exactly k periods and the agent's

payoff from playing E in all T periods is approximately

$$(T - k)c - [F(\underline{n}, T, T) - F(\underline{n}, T, k)]\tau,$$

while the difference between the agent's payoff from playing S in all T periods and the agent's payoff from playing E in all T periods is approximately

$$Tc - [F(\underline{n}, T, T) - F(\underline{n}, T, 0)]\tau.$$

We show in the next lemma that the former expression is smaller than the latter expression. This implies that if the agent is indifferent between playing S in all T periods and playing E in all T periods, then she is worse off by playing S in a positive number of periods and playing E in a positive number of periods.

Indeed,

$$\begin{aligned} (T - k)c - [F(\underline{n}, T, T) - F(\underline{n}, T, k)]\tau &< Tc \frac{[F(\underline{n}, T, T) - F(\underline{n}, T, k)]}{[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]} - \\ &- [F(\underline{n}, T, T) - F(\underline{n}, T, k)]\tau = \frac{[F(\underline{n}, T, T) - F(\underline{n}, T, k)]}{[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]} \cdot \\ &\cdot \{Tc - [F(\underline{n}, T, T) - F(\underline{n}, T, 0)]\tau\} \leq Tc - [F(\underline{n}, T, T) - F(\underline{n}, T, 0)]\tau, \end{aligned}$$

where the first inequality follows from Lemma 3, given below, and the second inequality follows from $F(\underline{n}, T, T) \geq F(\underline{n}, T, k) \geq F(\underline{n}, T, 0)$.

Thus, the prescribed strategies are equilibrium strategies.

Lemma 3.

$$\frac{T - k}{T} < \frac{[F(\underline{n}, T, T) - F(\underline{n}, T, k)]}{[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]} \quad (5)$$

for $k = 1, \dots, T - 1$.

Proof. Let

$$g(k) := \frac{[F(\underline{n}, T, k+1) - F(\underline{n}, T, k)]}{[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]} = \frac{f(\underline{n}, T-1, k)(p-q)}{[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]}$$

for $k = 0, \dots, T-1$, and let

$$h(k) = \frac{T-k}{T} - \frac{[F(\underline{n}, T, T) - F(\underline{n}, T, k)]}{[F(\underline{n}, T, T) - F(\underline{n}, T, 0)]}$$

for $k = 0, \dots, T$.

Notice that function g is singled-peaked by Lemma 2, and for large enough values of T , $g(0) > 1/T$ by the choice of \underline{n} and Lemma 1. Thus, if $h(k) \geq h(k+1)$, i.e., $g(k) \leq 1/T$, then $g(k+1) < 1/T$, i.e., $h(k+1) > h(k+2)$. This means that function h is singled-peaked. Since $h(0) = h(T) = 0$, we obtain (5) for $k = 1, \dots, T-1$. \square

Once we have the cutoff \underline{n} and transfer τ that make the principal indifferent over her reports, and make the agent indifferent between “all E ” and “all S ,” the remaining step in the construction is to slightly modify the bonuses and termination probabilities to give the players strict incentives for honest reporting. This can be done as follows.

Let $\rho(n, \rho^*)$ be the probability assigned by the principal to the event that the agent has been working in a block contingent on obtaining n good signals. Obviously, $\rho(n, \rho^*)$ is strictly increasing in n . Pick a $\underline{\rho}$ such that

$$\rho(\underline{n}, \rho^*) < \underline{\rho} < \rho(\underline{n}+1, \rho^*), \tag{6}$$

and prescribe to the principal a penalty $\eta/\underline{\rho}$ (i.e. the flow value of an associated termination probability) when the principal reports no more than \underline{n} good signals

and the agent was working; prescribe a penalty $\eta/(1 - \underline{\rho})$ when the principal reports strictly more than \underline{n} good signals and the agent was shirking, and prescribe no penalty in the remaining two cases.

Then, the principal has strict incentives to report truthfully. Indeed, suppose that she obtained exactly \underline{n} good signals. Then, her expected penalty when she reports honestly is $\rho(\underline{n}, \rho^*)\eta/\underline{\rho}$, which is smaller by (6) than $(1 - \rho(\underline{n}, \rho^*))\eta/(1 - \underline{\rho})$, her expected penalty when she reports more than \underline{n} good signals. Similarly, if the agent reports honestly $\underline{n} + 1$ good signals, her expected penalty is $(1 - \rho(\underline{n} + 1, \rho^*))\eta/(1 - \underline{\rho})$, which is smaller by (6) than $\rho(\underline{n} + 1, \rho^*)\eta/\underline{\rho}$, her expected penalty when she reports no more than \underline{n} good signals. It follows that the principal has the right incentives also when she obtains $n \neq \underline{n}, \underline{n} + 1$ good signals.

Providing incentives to the agent for truthful reporting by prescribing penalties for reports inconsistent with those of the principal is easier, because the agent who was working is almost sure that the principal obtained more than \underline{n} good signals, and the agent who was shirking is almost sure that the principal obtained no more than \underline{n} good signals. In addition, these penalties can be prescribed in the way that the agent is still indifferent between working and shirking, and the agent's payoff is held at his outside option.

The penalties guarantee strict incentives, except the agent's choice of working or shirking. This in turn guarantees that the equilibrium is purifiable, as in Section 4. Of course, the penalties can be chosen arbitrarily small to have only a negligible effect on payoffs. The agent's continuation utility in the equilibrium can be made equal to the agent's outside option by adjusting the base wage. Taking the limit as $\delta \rightarrow 1$ then completes the proof of Proposition 4.

Qualitatively, this equilibrium, where the agent randomizes between "always work" and "always shirk," differs from the T -block strategies in which the relation-

ship is not terminated as long as the outcome was G in at least one period. Under the latter strategies, when T is high, the probability of missing that cutoff is very low with or without effort, and so the size of the bonus necessary to incentivize effort is very large. These equilibria thus feature a large, negative base wage together with a bonus that is (a) much larger than the effort cost Tc , and (b) paid with probability close to 1 regardless of the agent’s effort choices. In contrast, in the second equilibria that we construct, a negative base wage is not required, the bonus is approximately equal to the effort cost Tc , and the cutoff is non-trivial. That is, the agent reaches the cutoff with high probability after choosing “always work,” but with very low probability after “always shirk.”

The equilibrium construction of our first approach, used in Proposition 4, builds on the strategies in Fuchs (2007). It utilizes self-reports by the agent, but it shares the negative base wage and large bonus of the earlier work. We note, however, that it is possible to construct a T -period block strategy equilibrium of the baseline game in which the agent shirks in at most one period, the base wage is positive and the size of the bonus τ is close to Tc . The corresponding review specifies, like the one from this section, that the agent passes the test as long as the principal observes at least \underline{n} good signals, where $\underline{n} > 1$ is non-trivial: always working ensures that the agent passes the test with high probability, but always shirking means that the agent is very likely to fail the test. We can then build on these strategies (instead of the strategies in Fuchs (2007)) to construct equilibria with properties similar to those of the equilibria constructed in this section, in which the agent shirks in at most one period, instead of always shirking or always exerting effort.

Given that result, our first approach as well as our second approach can produce purifiable, virtually efficient T -period block equilibria with positive wages, bonuses of realistic magnitudes, and nontrivial reviews.

Consider a variant of our model, where the principal can sign a binding contract whereby she commits to paying either the agent or a third-party such as a charity, but where performance evaluation remains subjective, as in McLeod (2003). Proposition 4 and 5 would then apply, but without any requirement on δ . Since the principal is committed contractually to make payments *either* to the agent *or* to charity, she cannot renege on making at least one of the payments. Our construction provides incentives for the principal to truthfully choose between these options. Furthermore, the relationship could have a finite horizon T , and the principal could obtain a per-period payoff within Δ of $\bar{y} - \bar{w} - c$, as long as T was sufficiently large.

6 Concluding comments

Requiring robustness to small payoff shocks, such as are present in many economic applications, has major consequences for relational contracting between a principal and an agent. No equilibrium where the agent always exerts effort is robust to payoff shocks. We have shown, however, that extending the basic interaction by allowing cheap talk can make effort sustainable.

We construct equilibria with realistic properties: the base wage is positive, the size of the bonus paid to the agent for good performance in order to incentivize effort is proportionate to the cost of effort, and the threshold for earning the bonus is set so that the agent can likely reach it by exerting effort but not otherwise. The optimal contracts previously studied in the literature, in contrast, feature bonuses that are much larger than the actual effort cost and thresholds that are nearly certain to be reached even without effort. Our construction also highlights the role of self review of workers' performance. The construction relies on simultaneous messages, which suggests a motivation for assigning the task of evaluating an employee's performance

to a manager rather than to the firm owner who pays the bonuses. That way, the employee and owner do not know the other's report when making their own.

From the point of view of theory, our analysis focuses on the parts that timing and communication play in purifiability. The simultaneity of the cheap-talk announcements in our constructions is crucial. If instead one player reported after seeing the other's report, then the proofs of our positive results would unravel, and effort could not be sustained. On the other hand, our results do not rely on the assumption that announcements are costless. Because players have strict incentives to make truthful reports in our constructions, those equilibria are robust to introducing small costs that may vary across announcements.

7 Appendix

7.1 Proofs of negative results

7.1.1 Proof of Proposition 1

Consider the perturbed game. Then, for every public history ω_{t-1} on equilibrium path, the action sequence $(\hat{a}^1, \dots, \hat{a}^t)$ is correctly anticipated by the principal. Consider the principal's making the decision in period t . The principal's observation of output y_t does not affect her continuation value – it only affected her flow payoff that she obtained at date t , but not her beliefs about the agent's continuation strategy. Let period $t + k$ be the first period in which the realization of y_t affects either the bonus payment or the principal's termination decision.

Suppose that the principal makes two distinct bonus payments $\tau < \tau'$ after private histories that only differ with regard to y_t . Incentive compatibility for the principal implies that the agent's continuation strategy after τ must differ from that after τ' , with τ inducing either lower effort or a greater termination probability. Thus, for the values of z_{t+k} above some cutoff the principal will prefer the agent's continuation strategy after τ' and for the values of z_{t+k} below that cutoff the principal will prefer the agent's continuation strategy after τ . Thus, the set of values of z_{t+k} such that the principal conditions her bonus payment upon y_t is negligible. If the realization of y_t does not affect the bonus payment in period $t + k$, then for almost all values of z_{t+k} , the principal must strictly prefer either the termination decision prescribed after $y_t = G$ or the termination decision prescribed after $y_t = B$.

7.1.2 Proof of Proposition 2

Consider the perturbed game, and a period t at the end of the T -period block. Suppose that the principal paid a bonus level τ (which could possibly be zero), and players observed a realization of the public randomization device θ . Conditional on the public event (τ, θ) , the agent's quitting strategy in period t is a function $\rho_t^2(\tau, \theta) : \{E, S\}^T \rightarrow \{C, Q\}$, i.e., it specifies continue or quit depending upon his effort choices in the block. The principal's firing strategy in period t is a function $\sigma_t^2(\tau, \theta) : \{G, B\}^T \times Z \rightarrow \{R, F\}$, i.e., it specifies firing or retaining the agent depending upon the outputs in the block and upon the value to the principal's payoff shock. Observe that either player's choice on whether to continue or terminate the relationship only matters when the other player chooses to continue, since the relationship terminates if one player chooses to terminate it. Suppose that the equilibrium specifies that after (τ, θ) , the relationship continues with positive probability, so that there exists some effort sequence and a non-negligible set of payoff shocks for which the agent chooses C . Since output signals have full support, for any output sequence in $\{G, B\}^T$, the principal assigns positive probability to the agent continuing. Thus, for almost any realization of the principal's payoff shock z_t , if it is optimal for the principal to continue (resp. terminate) the relationship at some sequence in $\{G, B\}^T$, it is optimal to continue (resp. terminate) at every other sequence in $\{G, B\}^T$. In other words, at any (τ, θ) where the relationship continues with positive probability, the principal's firing decision does not depend upon the output sequence in the block.

A similar argument establishes that at such (τ, θ) (where the relationship continues with positive probability), the agent's quitting decision does not depend upon the sequence of effort choices within the block. This establishes that for *any* (τ, θ) , the probability that the relationship continues does not depend either upon effort

choices within the block, nor on the output sequence.

Now consider the principal's choice between two distinct bonus values τ and τ' . If $\tau > \tau'$, then the probability of terminating the relationship must be greater after τ' . The preceding argument has established that the principal's continuation value from choosing τ (or τ') does not depend upon the observed output sequence. Thus, if it is strictly optimal to choose τ at some output sequence and some realization z_t of her payoff shock, then she will choose τ at every other output sequence for that same value of z_t . That is, the bonus payment does not depend upon the output sequence, and consequently, it is suboptimal for the agent to exert any effort.

This argument establishes that for any z_t , the set of block strategy equilibria where the agent chooses effort with positive probability is empty. Thus any block strategy equilibrium of the unperturbed game that features positive effort is not purifiable.

7.2 Proving Lemma 4

We show first that the specified strategies are an equilibrium, and then that the equilibrium is purifiable.

Proof. There are four sets of equilibrium conditions that must be verified.

The agent has strict incentives to shirk in one and only in one period: The agent is happy to shirk in one period, because the period chosen by him does not count in his review, and he saves on the cost of providing effort (in periods $t \leq T$), or he is paid an adjustment (in periods $t > 1$), or both (in periods $1 < t \leq T$). Shirking in more than one period will affect the review. By definition of a T -period review strategy, the bonus τ and test χ are such that the agent prefers to exert effort in every period except the one chosen for shirking.

The agent is indifferent across all $T + 1$ periods regarding the choice of period for shirking: This follows because all periods (except the one selected for shirking) equally matter for the agent's review, and the adjustments $(1 - \delta^{t-1})c/\delta^{T-1}$ for $t \in \{1, \dots, T\}$ and c/δ^{T-1} for $t = T + 1$ make the saving on costs equal to c (in terms of period 1) across all $T + 1$ periods. In addition, the testing procedure has the property that the chance of losing ε is independent of the choice of the period for shirking, given the honest report of the principal.

The agent has strict incentives to honestly report the period in which he has shirked: If the agent misreports the period in which he has shirked, he replaces in the review a period in which he played E with a period in which he played S . So, a misreport has the same effect on the review as taking action S instead of taking action E . In contrast to taking S instead of taking E , the agent does not save on the cost of effort by a misreport, but he reduces the probability of losing ε . However, this last benefit is assumed to be very close to zero.

The principal has strict incentives to honestly report her signals: Testing of the principal's reports is designed such that from the ex ante perspective each period (including the fictitious one) is chosen for comparing the agent's action and the principal's signal with probability $1/(T + 1)$; in addition, the probability that action E (or that action S) was taken in the chosen period is $1/2$. The principal's objective is to report G or B that is consistent with the agent's action (i.e., G if the action was E , and B if the action was S). By receiving her signal, she updates her belief in favor of E when the signal is G , and in favor of S if the signal is B . Thus, any misreport reduces the probability of attaining the principal's objective.

Of course, the principal correctly anticipates the agent's strategy of playing S in exactly one period. So, her computation of probabilities is contingent on the event

that the agent takes action S in exactly one period (possibly the fictitious one, that is, in none) and contingent on all her T signals. This affects the probabilities assigned by the principal to actions E and S in a given period, given the sequence of signals that she obtained, but it does not change which of the two probabilities is higher. For example, suppose that the principal obtained signals G in all T periods, and considers a period $t \in \{1, \dots, T\}$. If the agent had randomized 50-50 between actions E and S in period t , then the principal would have a chance of p for her report being consistent with the agent's action if she reported G , and she would have only a chance of $q < p$ if she reported B .

This is of course not the comparison that the principal is conducting. Instead, she computes the probability of: (i) t being the period in which the agent shirked, (ii) t being a period in which the agent worked and t being selected by the testing procedure, and compares these two probabilities. According to the testing procedure, she saves ε with probability $1/2$ contingent on each of these two events if her report in the event is consistent with the agent's action. The probability of the former event (i) is $q/(Tq + p)$, because if one of the actual T periods is selected for shirking, the probability of all signals being G is qp^{T-1} ; if the fictitious period is selected, the probability of all signals being G is p^T ; and each of the $T + 1$ periods is selected with probability $1/(T + 1)$. The probability of the latter event (ii) in turn is $[(T - 1)q + p]/T(Tq + p)$, because the probability of selecting one of the periods $s \in \{1, \dots, T\}$, $s \neq t$ for shirking is $q/(Tq + p)$, and then selecting t by the testing procedure is $1/T$; and the probability of selecting the fictitious period for shirking is $p/(Tq + p)$, and then selecting t by the testing procedure is $1/T$. Since $[(T - 1)q + p]/T(Tq + p) > q/(Tq + p)$, the principal saves ε with a higher probability when she reports signal G in period t .

The proof for other sequences of the principal's signals is analogous.

Purifiability: Since players have strict incentives except the agent's decision in which period to shirk, it is with no loss of generality to restrict attention to the shocks x_t that affect the agent's cost of effort. Assume that the process governing x_t is autoregressive, exactly as the process z_t from the main text. The cost of effort has the form $c + \xi x_t$. Let H denote the cdf of the cost. Consider the strategy of the agent that prescribes action S in period t , if it has not prescribed S in an earlier period, when

$$H(x_t) > 1 - \frac{1}{T + 2 - t}; \quad (7)$$

the strategy prescribes action E otherwise.

A simple induction shows the probability of shirking in each period $t \in \{1, \dots, T\}$ under the prescribed strategy is $1/(T + 1)$. To provide the agent incentives for conforming to the prescribed strategy, we modify the adjustment paid contingent on the period in which the agent shirks. For $t \in \{1, \dots, T\}$, let x_t^* be such that x_t^* satisfies condition (7) in which inequality is replaced with equality. We specify the adjustment by backward induction. First, we pick the adjustment for choosing the fictitious period $T + 1$ for shirking such that the agent with cost $c + \xi x_T^*$ in period T is indifferent between taking action E and taking action S . Next, we pick the adjustment for choosing period T for shirking such that the agent with cost $c + \xi x_{T-1}^*$ in period $T - 1$ is indifferent between taking action E and taking action S . This adjustment (multiplied by $1/\delta$) is added to the adjustment in period T in order to keep the agent with cost $c + \xi x_T^*$ in period T indifferent. Continuing in this manner, we provide the agent incentives for playing the prescribed strategy in all periods. \square

8 References

Abreu, D., P. Milgrom, and D. Pearce (1991): “Information and Timing in Repeated Partnerships,” *Econometrica*, 59(6), 1713–1733.

Baker, G., R. Gibbons, and K. J. Murphy (2002): “Relational Contracts and the Theory of the Firm,” *The Quarterly Journal of Economics*, 117(1), 39–84.

Bhaskar, V., G. J. Mailath, and S. Morris (2013): “A Foundation for Markov Equilibria in Sequential Games with Finite Social Memory,” *Review of Economic Studies*, 80(3), 925–948.

Bhaskar, V., and I. Obara (2002): “Belief-Based Equilibria in the Repeated Prisoners’ Dilemma with Private Monitoring,” *Journal of Economic Theory*, 102(1), 40–69.

Chan, J., and B. Zheng (2011): “Rewarding improvements: optimal dynamic contracts with subjective evaluation,” *The RAND Journal of Economics*, 42(4), 758–775.

Compte, O. (1998): “Communication in Repeated Games with Imperfect Private Monitoring,” *Econometrica*, 66(3), 597–626.

Ely, J. C., J. Hörner, and W. Olszewski (2005): “Belief-Free Equilibria in Repeated Games,” *Econometrica*, 73(2), 377–416.

Ely, J. C., and J. Välimäki (2002): “A Robust Folk Theorem for the Prisoner’s Dilemma,” *Journal of Economic Theory*, 102(1), 84–105.

Fuchs, W. (2007): “Contracting with repeated moral hazard and private evaluations,” *American Economic Review*, 97(4), 1432–1448.

Hörner J., and W. Olszewski (2006): “The Folk Theorem for Games with Private Almost-Perfect Monitoring,” *Econometrica*, 74(6), 1499–1544.

Kandori, M., and H. Matsushima (1998): “Private Observation, Communication

and Collusion,” *Econometrica*, 66(3), 627–652.

Levin, J. (2003): “Relational Incentive Contracts,” *American Economic Review*, 93(3), 835–857.

MacLeod, W. B. (2003): “Optimal contracting with subjective evaluation,” *American Economic Review*, 93(1), 216–240.

Maestri, L. (2012): “Bonus payments versus efficiency wages in the repeated principal-agent model with subjective evaluations,” *American Economic Journal: Microeconomics*, 4(3), 34–56.

Malcomson, J. M. (2013): “Relational incentive contracts,” *The handbook of organizational economics*, pp. 1014–1065.

Matsushima, H. (1991): “On the theory of repeated games with private information: Part I: anti-folk theorem without communication,” *Economics Letters*, 35(3), 253–256.

Matsushima, H. (2004): “Repeated Games with Private Monitoring: Two Players,” *Econometrica*, 72(3), 823–852.

Miyagawa, E., Y. Miyahara and T. Sekiguchi (2008): “The folk theorem for repeated games with observation costs,” *Journal of Economic Theory* 139(1), 192–221.

Piccione, M. (2002): “The Repeated Prisoner’s Dilemma with Imperfect Private Monitoring,” *Journal of Economic Theory*, 102(1), 70–83.

Radner, R. (1985): “Repeated Principal-Agent Games with Discounting,” *Econometrica*, 53(5), 1173–1198.

Rahman, D. (2012): “But who will monitor the monitor?,” *American Economic Review*, 102(6), 2767–2797.

Rahman, D. and I. Obara (2010): “Mediated Partnerships,” *Econometrica*, 78(1), 285–308.

Sekiguchi, T. (1997): “Efficiency in Repeated Prisoner’s Dilemma with Private

Monitoring,” *Journal of Economic Theory*, 76(2), 345–361.

Sugaya, T. (2022): “Folk Theorem in Repeated Games with Private Monitoring,”
Review of Economic Studies, 89(4), 2201-2256.