

Unraveling in a Repeated Moral Hazard Model with Multiple Agents ^{*}

Madhav Chandrasekher[†]

Current Draft: December 30, 2011

Abstract: Consider an infinite horizon repeated moral hazard problem where a single principal employs several agents. The principal cannot observe the agents' effort choices; however, agents can observe each other and can be contractually required to make observation reports to the principal. Observation reports are cheap talk, so that a primary concern is the prevention of collusion, where some agents shirk and report otherwise to the principal. The main result of the paper constructs a class of collusion-proof contracts with two properties. First, equilibrium payoffs to both the principal and the agents approach their first-best benchmarks as the discount factor tends to unity. These payoff bounds hold uniformly across all subgame perfect equilibria in the game induced by the contract. Second, while equilibria depend on the discount factor, the contract itself is independent of the discount factor.

Keywords: Repeated Games, Collusion, Communication, Statistical Testing.

^{*}Acknowledgments: I am very grateful to the co-editor, Johannes Hörner, and three anonymous referees for their feedback on this manuscript. I also thank Chris Shannon for important early guidance on this project and have received many helpful comments from Bob Anderson, Amanda Friedenberg, Natalia Kovrijnykh, Alejandro Manelli, Daniel Monte, Yuliy Sannikov, Ed Schlee, and especially Hector Chade and Sambuddha Ghosh. Any errors are my own.

[†]Mailing address: Department of Economics, W.P. Carey School of Business, P.O. Box 873806, Tempe, AZ 85287-3806; E-mail: madhav.chandrasekher@asu.edu

1 Introduction

This paper studies a contract design problem where a principal hires several agents over an infinite time horizon. The environment of the stage game has the following features. Effort is unobservable to the principal, but agents can observe each other. There is a (possibly null) communication phase added to the standard model where, between the time when effort is taken and output is realized, each agent can be required to make a publicly verifiable report of his co-workers' effort choices. Wages in any period are contingent only on the principal's information, i.e. the history of output data and observation reports. The main result of this paper constructs a class of infinite horizon collusion-proof contracts with the property that, in any equilibrium of the induced game, payoffs to the principal and agents converge to their first-best benchmarks as the discount factor tends to unity.

To place the dynamic problem in context, let us revisit the results developed for the static model. Variations of this contracting environment have been studied in a series of papers, e.g. [Holmstrom \(1982\)](#), [Mookherjee \(1984\)](#), [Ma \(1988\)](#), [Ma et al. \(1988\)](#), [Miller \(1997\)](#) among others. The model closest to ours is the one in [Ma \(1988\)](#). The following mechanism – which is a sibling of the Ma contract – solves the static optimal contract problem. The contract has two components, insurance and a stochastic bonus (reward). Each player is assigned a monitor and is, in turn, assigned to be some other player's monitor. Moreover, between the time when effort is chosen and output is realized each monitor is called upon to issue a report on the effort choice of the player he monitors. Wages are determined as follows. A player receives insurance as long as his monitor issues a positive report. Moreover, if he receives a positive report he can receive a stochastic bonus by issuing a negative report on the player he monitors. On the other hand, if he receives a negative report he is ineligible for insurance and gets his reservation wage. Since wages and bonuses are contingent only on the principal's information, under standard assumptions on the conditional output distributions (e.g. stochastic dominance) one can construct stochastic bonus payments that induce truth-telling in equilibrium. Using this construction, one can show that, in the static setting, the optimal contract returns the principal to the first-best, perfect information benchmark.

Since the observation reports are cheap talk a primary concern is preventing collusion. That is, the game induced by the contract should not possess any equilibria where some players shirk and still obtain insurance by issuing false reports to the principal. The game induced by the Ma contract has this collusion-proof property. However, a static or even a finite-horizon setting is not the best framework for addressing the possibility of collusion. To see this, imagine the time horizon is very long, but finite. The logic of the Ma contract, extended to a finite horizon, rests on a familiar unraveling argument: Collusion cannot be sustained in the last period of

a putative equilibrium path along which it allegedly occurs. Consequently, there is no last period of collusion, which implies that there could not have been any collusion at all. The existence of a deterministic end to the employment relationship is a critical feature of this argument. This assumption has the effect of choking off any incentives to collude in the final period of the contract, and arms the principal with a *de facto* incentive instrument. Thus, by imposing a deterministic end to the contract horizon we assume away a part of the problem.¹

The main goal of this paper is to provide a collusion-proof implementation of the efficient stage outcome in the infinite horizon problem. In particular, we desire an infinite horizon contract with the following three features: (i) it prevents rents from collusion, (ii) it maintains an efficient equilibrium – where agents work and report truthfully, and (iii) the principal’s payoff is nearly first-best in all equilibria of the game induced by the contract. Let us now describe the candidate contracts that have these features. The principal starts off period 0 with the null hypothesis that agents will be working hard and reporting each others’ effort choices truthfully whenever asked. In anticipation of this, he offers all agents a version of the contract in [Ma \(1988\)](#) for a fixed period of time, call this a ‘review phase’. At the end of this period he conducts an audit on output quality, and decides whether the data confirms or rejects his null hypothesis. If it confirms his null, then a new review phase is started and he continues to offer the modified [Ma \(1988\)](#) contract. On the other hand, if the data rejects the null, then this triggers a permanent punishment phase, where agents receive the subsistence wage thereafter.

The auditing mechanism in the contract takes the form of a statistical test. The principal holds a hypothesis on the level of effort that has been chosen and either accepts or rejects by comparing the empirical distribution of output to its hypothesized distribution. This idea of statistical testing in contracts can be traced to the seminal paper by [Radner \(1985\)](#). Our contract melds the static multiple-agent contract with a variant of the single agent contract in [Radner \(1985\)](#). In contrast to [Radner \(1985\)](#), it contains two incentive instruments: (i) the monitoring mechanism in the stage contract and (ii) the statistical review mechanism from which the principal infers effort choice from output data. However, setting aside the stage contract, there are some important differences between our review mechanism and the one introduced in [Radner \(1985\)](#).

The principal change is that our review mechanism, and the contract as a whole, is *independent* of the discount factor. This is a feature that is relevant to the prob-

¹This is unfair to [Ma \(1988\)](#) since the stated question, while relevant to the static model, is not the question of his paper. Ma’s paper is primarily concerned with unique implementation of the constrained efficient outcome - which had been an open question in the literature.

lem of robust design. Time preferences are subjective and the discount factor, in particular, is not outwardly observable. Hence, it is up to the principal to elicit this object through clever choice experiments. However, if the agent(s) knows that the principal is going to offer a review-style mechanism where reviews (comprised of test periods and punishment periods) are sensitive to the reported discounted factor, then he might have an incentive to misrepresent his time preference.² With our contracts, we can escape this problem. All agents receive the same contract regardless of their reported discounted factor.

Finally, let us note that the two incentive instruments in our contracts are both necessary to obtain our results. If one merely repeats the stage contract, then the infinite horizon game inherits the efficient equilibrium. However, there are now also undesirable equilibria where collusion persists in every period. These equilibria seem at least as plausible as the efficient outcome, hence we are not content with merely repeating the stage contract. In order to remove these equilibria, the principal uses the Radner-style statistical reviews. This then raises a secondary question. Since we are introducing the possibility of colluding on reports by requiring an otherwise elective communication phase, can the problem be solved by using statistical reviews or more simply – Radner’s review contracts – alone?

The [Radner \(1985\)](#) approach runs into some non-trivial difficulties. The primary obstruction regards the principal’s payoff, where we require a lower bound taken across equilibrium, as opposed to cooperative, outcomes. This is a non-issue with a single agent since team incentives and agent incentives are (vacuously) aligned. However, with multiple players this need not be the case. Moreover, what is best for the team may not even be implementable in equilibrium. There are also additional technical issues that arise once we introduce multiple agents, e.g. the fact that equilibria can be non-stationary. These issues do not offer a simple resolution via the [Radner \(1985\)](#) approach or even if we couple the [Radner \(1985\)](#) contract with the stage monitoring contract, i.e. a ‘Radner plus Ma’ approach.

A contribution of our paper is that it provides an alternative to the Radner approach that circumvents these problems and delivers payoff bounds, uniform across all (subgame perfect) equilibria, for both the principal and agent. We accomplish this by shifting away from the dynamic programming approach in [Radner \(1985\)](#) and analyzing, instead, the set of admissible probability distributions on histories, i.e. probability measures on histories that are induced by equilibrium strategies. The hypothesis of equilibrium places a tight structure on the set of admissible measures, which is then used to extract payoff bounds for both contractual parties.

²This is exactly what would happen if the principal uses Radner review contracts (which are sensitive to the discount factor) and the agent anticipates this.

An additional benefit of this approach is that it shows that the unraveling logic of the finite shot contract can be extended to the infinite horizon. Hence, in addition to obtaining payoff bounds we provide some game-theoretic intuition for *why* these bounds obtain in equilibrium: For collusion to persist in equilibrium, it cannot be so frequent that it induces detection by the statistical review. Moreover, once the statistical review registers failure the game collapses to a finite-horizon model and agreements to collude are unraveled. This unraveling phenomenon is unique to our contract and would be absent in a putative extension of the [Radner \(1985\)](#) framework. More generally, in that paper restrictions on equilibrium behavior are implicit in the results about equilibrium payoffs. In our paper, it is the other way around. We first derive explicit implications about equilibrium behavior and use these restrictions to derive payoff bounds.

1.1 Related Literature

There is a rich literature on static contracting problems with multiple agents, initiated by [Holmstrom \(1982\)](#), [Mookherjee \(1984\)](#), and subsequently developed in e.g. [Ma \(1988\)](#), [Itoh \(1991\)](#), [Ishiguro and Itoh \(2001\)](#). However, there has been comparatively little work on dynamic extensions of these contracts to the infinite horizon. Some recent exceptions are [Che and Yoo \(2001\)](#) and [Bonatti and Hörner \(2009\)](#), who also study repeated teams problems. See also [Abdulkadiroglu and Chung \(2003\)](#). Additionally, there is a related literature on experimentation in teams that we are not mentioning here, although the [Bonatti and Hörner \(2009\)](#) paper can also be considered an example of this. The principal objective of the [Che and Yoo \(2001\)](#) paper is to provide a simpler (and stationary) incentive mechanism (joint-performance evaluation contracts) that in many cases dominates more commonly used mechanisms.

[Bonatti and Hörner \(2009\)](#) study a repeated teams problem where agents must work together and exert costly and unobservable effort into a project of unknown value. Higher effort induces quicker discovery of the value of the project. Project value is realized only when discovery occurs and, moreover, when this happens the game terminates. The authors carry out a comprehensive analysis of this problem, among other things obtaining closed form solutions for equilibrium effort choice (when agents themselves value discovery) and solving for the optimal wage scheme when a principal owns the project.

2 The Stage Contract

A risk neutral principal must hire n agents to complete a task. Output assumes a finite set of values and is a function of collective effort choice. These choices are

unobservable to the principal but are observable to the agents. Between the time when effort is chosen and output is realized, each agent makes an observation report to the principal. These reports are contractible and are made publicly, so that all agents know what other agents have reported.³

The set of possible effort choices for each agent is $\{e_H, e_L\}$, i.e. each agent just chooses (if he takes a pure action) between high and low effort.⁴ Let $\mathbf{e} \in \{e_H, e_L\}^n$ denote a vector of effort choices for the labor force and let $f(x_j|\mathbf{e})$ be the probability of output x_j conditional on this choice. Lastly, let $c(e_i)$ denote each agent's (utility) cost of choosing e_i and assume the common utility index $U(w, e) = u(w) - c(e)$, where $u(\cdot)$ and $c(\cdot)$ are increasing and $u(\cdot)$ is weakly concave and C^2 . Also assume vNM preferences over wage lotteries (with money kernel $u(\cdot)$). Let \mathbf{E} be the set of all cumulative effort choice vectors. To this set-up we add the following assumptions. For $\mathbf{e} \in \{e_H, e_L\}^n$, say that $\mathbf{e} \leq \mathbf{e}'$ if at least as many agents select e_H under \mathbf{e}' than under \mathbf{e} . Let 'FOSD' abbreviate the first-order stochastic dominance relation.

Assumptions on primitives:

1. If $\mathbf{e} \leq \mathbf{e}'$, then $F(\cdot|\mathbf{e}')$ FOSD $F(\cdot|\mathbf{e})$.
2. $\mathbf{e}^* := (e_H, e_H, \dots, e_H)$ is first-best.

Definition 1 (Definition of Collusion). An outcome is said to be *collusive* if collective effort choice is less than first-best, yet some agent is earning better than his reservation wage.

Intuitively, we say that agents are colluding if they are shirking and collecting insurance and, nevertheless, lying to the principal about the effort choices of their peers. The definition above includes this possibility and more. For example, it also counts as "collusion" the situation where someone is shirking and someone else (who isn't shirking) is earning insurance. Using a broader definition of collusion only strengthens our desired conclusion since our goal is to design a contract that rules out manifestly collusive behavior without ruling out the efficient equilibrium. It suffices, then, to rule out collusive outcomes using the broader definition of collusion given above.⁵ The following proposition was previously established in [Ma \(1988\)](#). However, the mechanism given in Ma's paper is slightly different than the one we present. For this reason only, a proof is also provided in this paper (in the appendix).

³The assumptions on observability and report structure are elective. We can just as well allow private reports and/or imperfect observability of effort choice without changing the results for the stage or infinite horizon contract.

⁴The assumption on support of effort is for economy of notation. It is not necessary for any of the results.

⁵Note that the efficient outcome (i.e. all agents putting in high effort and truth-telling) is not collusive under the definition.

Proposition 1 (Ma (1988)). *Assume players cannot make inter-personal transfers (i.e. side-contracts). There exists a contract that attains the first best effort choice at the first best cost. Moreover, in the extensive form game induced by the contract there is no collusion in equilibrium.*

We first give a sketch of the argument for the illustrative case where there are two agents, two effort choices, and two values of output (Brusco (1997) refers to this as the $2 \times 2 \times 2$ model). Consider the contract represented by the following matrix of payoffs. Label the players $\{1, 2\}$. The entries in the box denote player 1's payoff.

	$r_{1,2} = +$	$r_{1,2} = -$
$r_{2,1} = +$	w^{FB}	$w^{FB} + (R_1, R_2)$
$r_{2,1} = -$	$w_0 - \epsilon$	w_0

Table 1: A sample contract

The term $r_{1,2}$ denotes player 1's report on player 2, and similarly for $r_{2,1}$. A plus (+) denotes a good report and minus (−) is a bad report. Thus, if both players issue good reports on each other, they both get the first-best wage (i.e. full insurance). If 1 gives 2 a minus and 2 gives 1 a plus, then 1 additionally obtains a stochastic reward (R_1, R_2) . The reward pays out $R_1 < 0$ if output is high and $R_2 > 0$ if the output is low. The idea behind the sign convention is that if 1 reports on 2 and is telling the truth (i.e. somebody is shirking), then low output should be more likely than high output. Thus, the reward should have positive expected value. If neither player is shirking and player 1 is untruthfully giving player 2 a thumbs down, then high output should be more likely than low so that the reward should have negative expected value. Analogously defining the payoff matrix for player 2, one can verify that there is no collusion in equilibrium under this contract. Moreover, both players putting in high effort and truth-telling is an equilibrium.

The principal difference between the mechanism we use to prove the Proposition and the one constructed in Ma (1988) is that the equilibrium outcome is *not* unique in our set-up whereas it is unique in Ma's contract. The reason for this is that the environment in our paper, while related to the one in Ma (1988), is formally distinct. In Ma (1988), the principal's information consists of a bivariate public signal – implicitly, one for each agent's action choice. In contrast, in our environment there is a single public signal for the principal with values that are correlated with the joint effort choices of the agents. The following lemma shows that (for the $2 \times 2 \times 2$ model) any mechanism that is collusion-proof and maintains an efficient SPNE must also admit another (inefficient) SPNE. Let \mathcal{E} denote the economic environment. This consists of: (i) two agents, two effort choices, and two output values, (ii) a single public signal taken to be the value of output, and (iii) with a view to the infinite horizon problem, agents are assumed to have non-negligible (but arbitrarily small)

liability. Consider contracts, $\mathcal{C}(\mathcal{E})$, that satisfy symmetry, i.e. $w_i(\cdot, (m^1, m^2)) = w_{-i}(\cdot, (m^2, m^1))$ – so that the wage is not dependent on the identity of the sender.

Lemma 1. *For any contract $\mathcal{C}(\mathcal{E})$, let \mathcal{G} denote the game form induced by the contract and let $\Sigma(\mathcal{G})$ be the set of SPNE. If the efficient outcome (i.e. high effort and full insurance) is an outcome of some equilibrium in $\Sigma(\mathcal{G})$ and no equilibrium outcomes exhibit collusion, then there must be multiple equilibria.*

Thus, if we simultaneously want collusion-proofness and (weak) implementation of the efficient outcome, then non-uniqueness of equilibrium is endemic to our single signal environment.⁶ Matters are different when the principal receives a bivariate signal (as in [Ma \(1988\)](#)) since he can use a different bonus scheme for each agent. Since agents are truth-telling in any SPNE, one can check that, in the two signal setting, the efficient equilibrium is the only equilibrium. None of this is of issue for the infinite-horizon implementation problem. Due to the presence of an (endogenous) participation constraint – where players only sign the contract if they anticipate obtaining at least a small ϵ more than reservation utility (see [proof of Theorem 1](#)) – the outcome where players shirk and report on each other in every period is “screened out” by the infinite-horizon contract.

3 The Infinite Horizon Contract

Now consider the infinite horizon setting. If the principal were to unconditionally offer the stage contract in every period, then the following trigger strategy constitutes an equilibrium: everyone shirks and covers for each other until a period when someone reports otherwise. After this happens, everyone shirks and reports truthfully in every period. It is easy to see that when players employ this strategy profile, collusion occurs in every period on the equilibrium path (for large δ).

To break this equilibrium the principal needs to offer large rewards for reporting. However, unless liability is unbounded, the only way to make the rewards large (in expectation) when agents shirk is to make the output-contingent payment large in reward states, e.g. when output is low in the $2 \times 2 \times 2$ model. But this means that the reward has large positive expected value even when nobody shirks, which destroys the efficient equilibrium. Hence, the Ma contract alone cannot yield dynamic collusion-proof implementation of the efficient outcome. We now design an infinite horizon collusion-proof contract that is insensitive to the liability bound.

⁶The proof of the lemma is not hard but requires a fair bit of notation. Moreover, the restriction to the $2 \times 2 \times 2$ model is for simplicity alone. The argument itself is not difficult to generalize. Details are available online at: <http://www.public.asu.edu/~mchandr6>.

By offering the stage contract unconditionally the principal is ignoring the information about effort choice contained in the stream of output data. With a large enough stream of data the principal should be able to infer whether agents were indeed implementing the first-best if that is what they were reporting. Hence, the additional incentive instrument takes the form of a statistical test on the hypothesis that the agents are implementing the first-best effort level. Before we explain the construction, let us make some assumptions to structure the infinite horizon problem:

- **A1:** Assume that, at time $t = 0$, the principal offers each agent a contract. This contract specifies: (i) a state space, (ii) an initial state, (iii) state-dependent payment rules, and (iv) transition rules that determine what the next period's state will be given the current state.
- **A2:** If any of the agents refuses the contract, the game ends and all agents receive their reservation utility. These rules cannot be renegotiated once the contract is signed.
- **A3:** Agents have a common (subjective) discount factor $\delta \in (0, 1)$ and payoffs in the infinite horizon game are evaluated using the δ -discounted sum of stage game payoffs.
- **A4:** Agents can accept an arbitrarily small, but positive, amount of liability. The liability bound is insensitive to the discount factor, δ .

The first three assumptions are, more or less, standard. Let us comment on A4. Note that the stage contract (Table 1 for an example) requires A4 since the reward incentives assess a small amount of punishment when high output is realized. Since we want the infinite horizon contract to inherit the efficient equilibrium, we need to be able to punish agents an arbitrarily small amount in each period, i.e. A4 requires that agents have a small, but positive, *per period* liability. This bound can be arbitrarily small, but does not change with the discount factor. Hence, we consider it a primitive of the contracting environment.

With the primitives of the infinite horizon problem described, we now move on to strategies and histories. We have not yet formally described the infinite horizon contract, which determines the game to be played. However, we just need an informal description to proceed here. The class of contracts we construct will alternate between two spot contracts, with history-dependent transitions. The first spot contract, which applies in 'work' periods, is the stage contract described in section 2. The second spot contract, which applies only in the punishment phase, pays a constant 'punishment wage' unconditionally.

Histories: Let $h^0 := \emptyset$ be the null history. Let h^t denote the history of the game up through period t . The contribution to h^t in period t itself consists of the following data: (1) the effort choices taken by the agents in period t , (2) the reports the agents have made, and (3) the realized output level. In a period in which the game is in a punishment state we denote that period's contribution to the history with the emptyset symbol, \emptyset . Let H^t denote the set of histories up through period t and put $\overline{H} := \cup_t H^t$.

Strategies: Let $R_i : \mathbf{E} \rightarrow \Delta(\{0, 1\})$ denote agent i 's observation report strategy (i.e. 0 = a shirk report, 1 = no shirk). Let Σ_i denote the set of such R_i and let $\{e_H, e_L\}$ denote the set of available effort choices. Then, agent i 's strategy space, $\mathcal{S}(i)$, in a work period can be described as $\Delta(\{e_H, e_L\}) \times \Sigma_i$. A strategy for player i (in the stochastic game) is a function $\rho_i : \overline{H} \rightarrow \mathcal{S}(i)$ that prescribes an (possibly mixed) action pair in period $t + 1$ as a function of $h^t \in \overline{H}$. If h^t is such that the spot contract in period $t + 1$ is the reservation wage contract, then set $\rho_i(h^t) = \emptyset$.⁷

Notice that each agent sees the same history at time t , consisting of the history of output, effort choices, and reports. Accordingly, the solution concept employed in this paper is subgame perfect Nash equilibrium (SPNE).

3.1 Statistical Testing

We now describe the statistical tools that are implemented in the contracts. Let \mathbf{E} be the set of collective effort choice vectors and denote the first-best choice as \mathbf{e}^* . For a given $\mathbf{e} \in \mathbf{E}$, output X is distributed as $F(\cdot|\mathbf{e})$. Consider X_1, \dots, X_T distributed independently and identically (as X). The empirical c.d.f corresponding to X_1, \dots, X_T is defined by the formula:

$$F_T(x|\mathbf{e}) := \frac{1}{T} \sum_{i=1}^T 1_{(X_i \leq x)}.$$

We now use the formula for the empirical c.d.f. to define a statistic. For $X \sim \mu_{\mathbf{e}^*}$ (i.e. X is distributed as $\mu_{\mathbf{e}^*}$) put $K = l \cdot \max_x \text{Var}(1_{(X \leq x)})$ (where l is the cardinality of the support of output). Define the following parameters of the hypothesis test

$$\gamma_n := \frac{1}{\sqrt{n}}, \quad \epsilon_n := \frac{K}{n^3}, \quad t_n := n^4.$$

These parameters serve the following roles:

⁷Reporting strategies in our stage contracts are richer than described here, since we use sequential reports. However, beyond formality, there is little conceptual content to the more elaborate definition of reporting strategies. Hence, we obscure the distinction here.

- **Margin of Error:** The quantity γ_n determines the allowable margin of error which determines the rejection region for our hypothesis test.
- **Type I error bound:** The term ϵ_n is an upper bound on the probability of a Type I error (this follows by a simple application of Chebyshev's inequality, see [proof](#) of Proposition 3).
- **Sample Size:** The quantity t_n is the sample size of the hypothesis test.

The values $\gamma_n, \epsilon_n, t_n$ are fixed for the rest of the paper. The n th *review phase*, of length $T_n := n \cdot t_n$, is partitioned into n samples of output data, $\{(i-1) \cdot t_n + 1, \dots, i \cdot t_n\}_{i=1}^n$. Define the empirical c.d.f. for each sample,

$$F_{i,T_n}(x) := \frac{1}{t_n} \cdot \sum_{j=(i-1)t_n+1}^{it_n} 1_{(X_j \leq x)}, \text{ for } i \in \{1, 2, \dots, n\}.$$

Definition 2 (Statistical Test). The *Kolmogorov-Smirnov test statistic* is given by the formula,

$$S_{T_n} := \max_i \{ \sup_x |F_{i,T_n}(x) - F(x|\mathbf{e}^*)| \}$$

The idea behind the statistical test is that the principal breaks up the n th work phase into n data gathering (sub)phases. During the entire work phase, the null hypothesis is that agents are working the first best and truth-telling in every period. For each of the n batches of output data, the principal uses a Kolmogorov-Smirnov (KS) test to match the empirical cdf of output with the hypothesized cdf. If the deviation from any of the n samples (weakly) exceeds the margin of error (γ_n defined above), then the null is rejected and a punishment phase follows.

3.2 Contract Description

State Space

The state space is divided into two classes, work states and a single (absorbing) punishment state.

1. Work states are identified by a pair (W_n, i) , where W_n is of length $T_n = n \cdot t_n$ and i denotes the current period within the given work state W_n .
2. There is a singleton and absorbing punishment state, denoted \emptyset .

Transition Rules

Transitions between work states and the punishment state are determined as follows:

1. Initial state: $(W_1, 1)$.

2. If in state (W_n, i) , where $i < T_n$, proceed to state $(W_n, i + 1)$.
3. If in state (W_n, T_n) , consider the value of the KS statistic S_{T_n} . If $S_{T_n} \geq \gamma_n$, then proceed to the unemployment state \emptyset . Otherwise, proceed to state $(W_{n+1}, 1)$.
4. If in state \emptyset , then remain in state \emptyset .

Payment Rules

Payment alternates between two types of spot contract. First, a verbal description. In any given period, either the Ma contract or the punishment contract (where everyone earns an ϵ less the reservation wage) is in place. Denote this contract by $C(\epsilon', R_1, R_2, w_\epsilon^{FB})$ with arguments: (i) the punishment quantity (ϵ'), (ii) the negative reward when high output values are realized (R_1), (iii) the positive reward when low values are realized (R_2), and (iv) the insurance payoff w_ϵ^{FB} , which equals cost of high effort plus another ϵ . Let s denote a generic state and let T denote the stopping time associated to the KS statistic, i.e. $T(h)$ is the first time (along history h) where it becomes known that the next state is the punishment state. The notation $s_t(h)$ denotes the state of the contract along history h and time t , and $C(w_0)$ denotes the spot contract that unconditionally pays every agent the reservation wage. The spot contract function is formally described:

$$C_t(h) = \begin{cases} C(\epsilon', R_1, R_2, w_\epsilon^{FB}) - t_\epsilon, & \text{if } t \leq T(h), s_t(h) \neq \emptyset \\ C(\epsilon', R_1, R_2, w_0^{FB}) - t_\epsilon, & \text{if } t > T(h), s_t(h) \neq \emptyset \\ C(w_0) - t_\epsilon, & \text{if } s_t(h) = \emptyset. \end{cases}$$

This completes the description of the infinite horizon contract.⁸ The choice parameters of the contract are:

- The stage contract parameters (i.e. $\epsilon', R_1, R_2, w_\epsilon^{FB}$),
- The KS test statistic, and
- The per-period participation fee t_ϵ .

In particular, note that *none* of these parameters depend on the discount factor. Let Φ denote the set of such contracts, with generic element \mathcal{C} .

Choose the fee t_ϵ plus the (negative) reward R_1 to sum to less than the liability bound granted by assumption A4. The fee can simply be absorbed into the wage function of the work state spot contract, $\mathcal{C}(\epsilon', R_1, R_2, w_\epsilon^{FB})$. However, we choose to write it this way since selection of an appropriate participation fee will allow us, through contract design, to elicit a bound on the induced (equilibrium) probability

⁸In the definition, w_0^{FB} denotes the insurance level where agents are insured an amount exactly equal to the cost of high effort.

that the KS test statistic never falls into the rejection region. The bound on this ‘tail’ success probability is key to getting a good lower bound on the principal’s payoff.

3.3 Equilibrium Behavior

Let $\Sigma(\delta)$ denote the set of (measurable) SPNE strategy profiles in the game induced by a contract \mathcal{C} . For a given equilibrium $\rho \in \Sigma(\delta)$, let $\bar{P}_\rho(\cdot)$ denote the conditional distribution on work phase n , where we condition on the set of histories that reach the n -th work phase.

Proposition 2. *Assume that $\limsup \bar{P}_\rho(\frac{C(T_n)}{T_n} > \alpha) > 0$ for some $\alpha > 0$. Then, $\limsup \bar{P}_\rho(\sup_x |F_{T_n}(x) - F(x|e^*)| > r) > 0$ for some $r \in (0, 1)$.*

Proposition 2 says that if agents are shirking frequently and reporting that they are working first-best, the statistical test will eventually catch on. Moreover, it tells us that when they collude often they not only fail the test, but the failure probability is bounded away from zero. Hence, whenever the conclusion of the Proposition holds it must be the case that (for large n) the null hypothesis is violated in more than one sub-phase of the work state. Taken together with the following observation, these two insights are the key to understanding equilibrium behavior and deriving bounds on equilibrium payoffs.

Observation: Collusion stops once agents are caught.

To justify this observation, let $\rho \in \Sigma(\delta)$ and assume that there is a history h such that $S_{T_n}(h)$ falls in the rejection region prior to the start of the n th sub-phase. Take the earliest of the sub-phases for which this happens, say sub-phase k where $k < n$. What can equilibrium play look like from the start of sub-phase $k + 1$ until the end of work-phase n ? We claim no collusion can occur from the start of phase $k + 1$ until the end of work phase n . This follows from an unraveling argument. If there is collusion, there is a last period in which it occurs. Since punishment states are absorbing and players receive subsistence wages after the review, the present value of reporting exceeds the value of not reporting. Thus, collusive agreements are impossible to sustain once everyone knows that a punishment phase is forthcoming. Consequently, once players have completed a sub-phase of work and know that the next state is a punishment state, then play in each period till the termination of the work state reduces to one of the stage SPNE. In so far as agents’ payoffs are concerned, we can wlog assume that play reverts to the efficient SPNE. Refer to a profile ρ with this property as a *rectified* strategy.

Proposition 3 (Equilibrium Behavior). *Let ρ be any SPNE (in the game induced by \mathcal{C}). Then, for any $\epsilon > 0$ we have $\overline{P}_\rho(\frac{C(T_n)}{T_n} \geq \epsilon) \rightarrow 0$ as $n \rightarrow \infty$.*⁹

The statement applies only to phases which are reached with positive probability on the equilibrium path. For the ‘finite punishments’ version of \mathcal{C} (see Corollary 1) all phases are reached with positive probability, so that the conditional distributions are always defined. The statement of the proposition does not invoke any restrictions on the SPNE, but we prove the proposition by reducing to rectified equilibria. By the preceding discussion, this reduction entails no loss of generality in so far as the distribution of $C(T_n)/T_n$ is concerned.

4 Main Result

Let $\mathcal{V}(\delta)$ denote the set of payoff vectors attainable through SPNE in the game induced by the contract \mathcal{C} . Our main result produces a pair of equilibrium payoff bounds. The first is an upper bound on agents’ payoffs. It says that as δ increases to one, the set $\mathcal{V}(\delta)$ converges (in the Hausdorff metric) to a point mass on the vector $(u(w^{FB}) - c(e_H), \dots, u(w^{FB}) - c(e_H))$, i.e. where every agent is earning the payoff under the first-best, perfect-information benchmark. Moreover, since sufficiently patient players can obtain close to this payoff by employing the efficient stage game equilibrium, this payoff vector is attained in the limit.

The second half of the theorem provides a lower bound on the principal’s equilibrium payoff. The lower bound is more difficult to derive than the upper bound. The key is to obtain control of (equilibrium) *tail* probabilities, i.e. the probability that agents never fail the KS test. The observation of Proposition 4 (see below) and the existence of a participation fee are critical to bounding this tail probability. Since the contract is independent of δ , the proposition delivers uniform bounds (i.e. across equilibria and across discount factors) on tail probabilities. The participation fee, which is only required for the lower bound, then ensures that these probabilities approach unity for large δ . Moreover, step 1 of the lower bound argument shows that for the principal’s payoff to approach the first-best benchmark it is both necessary and sufficient that this tail probability approaches unity.

Introduce some notation. Let Π^{FB} denote the first-best principal’s payoff and let (resp.) $W(\rho), V(\rho)$ denote the principal’s and a given agent’s discounted payoff under SPNE ρ . Note that any element $\mathcal{C} \in \Phi$ is defined independently of the discount factor. However, the SPNE set of the game induced by \mathcal{C} will typically depend on the discount factor.

⁹I am grateful to the anonymous referee who suggested this result.

Theorem 1. *Let $\Sigma(\delta)$ denote the set of all SPNE in the infinite horizon game induced by the contract \mathcal{C} . Given $\epsilon > 0$, we can find a contract $\mathcal{C}_\epsilon \in \Phi$ which yields the following bounds on equilibrium payoffs:*

1. (Upper Bound) *Put $V(\delta) = \sup_{\rho \in \Sigma(\delta)} V(\rho)$. Then, $\overline{\lim}_{\delta \uparrow 1} (1-\delta)V(\delta) = u(w^{FB}) - c(e_H)$.*
2. (Lower Bound) *Put $W(\delta) = \inf_{\rho \in \Sigma(\delta)} W(\rho)$. There exists $\delta_\epsilon > 0$ s.t. $\forall \delta \geq \delta_\epsilon$ we have $(1-\delta)W(\delta) \geq \Pi^{FB} - \epsilon$.*

The contracts in the class Φ all have an absorbing punishment state. Since agents are required to hand over a (arbitrarily small) participation fee in each period, it might seem unrealistic that they would agree upfront to this sort of liability. All the more so since a punishment phase can be triggered (albeit with very small probability) even when everyone pursues the efficient equilibrium. The following corollary shows that we recover the same result as [Theorem 1](#) even if we require punishments to be memoryless, as in [Radner \(1985\)](#), so that punishment phases are finite and of identical length. Let $\mathcal{C}(\delta)$ be identical to a contract \mathcal{C} chosen from Φ , with the only difference being that the punishment lengths are of some finite length $L(\delta)$. The state space, payoff functions, and transition rules admit obvious adjustments – hence, we omit the formal (re)definition for $\mathcal{C}(\delta)$.

Corollary 1. *Fix $\epsilon > 0$. Then, $\exists \delta_\epsilon > 0$ such that for each $\delta \geq \delta_\epsilon$ we can find a finite punishment contract $\mathcal{C}_\epsilon(\delta)$ such that the bounds in [Theorem 1](#) hold.*

The main result consists of two parts, (i) an upper bound on equilibrium payoffs for agents and (ii) a lower bound on equilibrium payoffs for the principal. To derive these bounds we require a strengthening of [Proposition 3](#). The following [Proposition](#) and its companion ([Proposition 5](#)) form the heart of the argument for [Theorem 1](#).

Proposition 4. *Given any $\alpha > 0, \beta > 0$ there is an index $N(\alpha, \beta)$ such that whenever $n \geq N(\alpha, \beta)$ we have $\overline{P}_\rho(C(T_n)/T_n \geq \alpha) < \beta, \forall \rho \in \Sigma(\delta), \forall \delta$.*

When play doesn't reach the n th review phase, the conditional distributions are not defined. Hence, we add, as in [Proposition 3](#), the qualification that the bound of the [Proposition](#) applies whenever $n \geq N(\alpha, \beta)$ and the n th phase is reached. The key is that the bound $N(\alpha, \beta)$ applies across *all* profiles $\rho \in \Sigma(\delta)$ and across *all* discount factors. This is the sense in which we are strengthening the previous [Proposition](#), which *prima facie* suggests that the bound is dependent on the given profile.

Notice that the [Proposition](#) delivers a uniform bound on the distribution of the “frequency of collusion” variable. Hence, if we can somehow pass from discounted payoffs to time-average payoffs we can leverage this result. Of course, the challenge of discounted repeated games lies precisely in the fact that dynamic incentive

provision is much more subtle with time-discounting than with time-averaging. To generate a folk theorem with discounting we require a rich enough set of ‘delayed reward sequences’. Loosely speaking, to sustain a target profile of on-path play we need to have the ability to bring sticks to the present and push carrots into the future. As the discount factor increases, to incentivize sophisticated patterns of on-path play we require access to a rich set of payoff sequences where rewards are (possibly) delayed further and further. These sequences exist so long as the feasible set of the stage game is sufficiently rich.

Let us see what happens when we try applying the folk theorem reasoning to our (stochastic) game. Put the payoff from an episode of collusion at 1 and the efficient stage outcome at 0 (any other outcome yields less than these). The Proposition says that, with probability close to 1, a payoff stream in review phases far along the time horizon is proportioned with at most ϵ ones and at least $1 - \epsilon$ zeroes. These are the only delayed reward sequences available through SPNE play. Hence, as we push the discount factor to unity it might be possible to generate a folk theorem in a vanishingly small neighborhood of the efficient stage game payoff – but that is all. This explains, in a nutshell, how we use the Proposition to obtain an ‘anti-folk’ result for the agents’ payoffs in our game.

For the lower bound argument, I require a sharper version of the Proposition. We apply the bounds to the variable $E(T_n)/T_n$, where $E(T_n)$ denotes the number of periods in phase n in which effort choice is not first-best. Notice that $C(T_n) \subseteq E(T_n)$. While the variable $C(T_n)$ is more informative about agents’ payoffs, the variable $E(T_n)$ is more informative about the principal’s payoff. Recall that the review phases of contracts in the class Φ have lengths indexed by the sequence $f(n) = n^5$.¹⁰ Let $\mathcal{C}_N \in \Phi$ denote the contract where the first review phase starts with the N th term in the sequence $f(n)$, i.e. so that the first review has length N^5 , the second $(N + 1)^5$, and so on. The Proposition stated below makes two changes from Proposition 4. First, it switches the variable of interest from $C(T_n)$ to $E(T_n)$. Second, there is now dependence of the bound $I(\alpha, \beta)$ (for convenience, we hereon switch to indexing phases with i ’s) on both the discount factor and the contract \mathcal{C}_N .

Proposition 5. *Fix any $\alpha > 0, \beta > 0$. There is an N (depending on α, β) and $\delta_{\alpha, \beta}$ such that, for the game induced by the contract $\mathcal{C}_N \in \Phi$, there is an index $I(\alpha, \beta)$ such that whenever $i \geq I(\alpha, \beta)$ we have $\overline{P}_\rho(E(T_i)/T_i \geq \alpha) < \beta, \forall \rho \in \Sigma(\delta), \forall \delta \geq \delta_{\alpha, \beta}$.*

While the Proposition is intuitively related to Proposition 4, its proof makes use of bounds on tail probabilities obtained in the lower bound argument. Hence, it is

¹⁰This may seem, at first blush, like a mysterious choice. Why such a specific choice for the sample sizes? This selection plays an important role in the proof of the lower bound for the principal. It is what enables us to equate bounds on time-average payoffs with bounds on limiting discounted payoffs. See the lower bound proof and Observation 2 in the appendix for details.

more logically correct – if somewhat non-intuitive – to present the Proposition after we go through the tail bound exercise.

4.1 Proof of Theorem 1

All discounted quantities are assumed to be pre-multiplied by $(1 - \delta)$ – we suppress this term in the notation. In the forthcoming argument, we take limits of subsequences of *measures*. Proposition 6 confirms that we can extract a convergent subsequence from any given sequence of measures, so that limits do indeed exist.¹¹

Upper Bound

Choosing an affine transformation of stage-game payoffs we assume that $u(w^{FB}) - c(e_H) = 0$. Since there is an efficient equilibrium, we clearly have $\overline{\lim}_{\delta \uparrow 1} (1 - \delta)V(\delta) \geq u(w^{FB}) - c(e_H), \forall \delta$. I check that the reverse inequality holds (in the limit, as $\delta \uparrow 1$). Take a sequence of profiles ρ_{δ_n} with $(1 - \delta_n)V(\rho_{\delta_n})$ converging to $\overline{\lim}_{\delta \uparrow 1} (1 - \delta)V(\delta)$. Let P_{ρ_n} denote the measures on H^* – the space of infinite histories – induced by these profiles. Let $u_t(h)$ (reps. $u_t^\delta(h)$) denote the function which gives payoff 1 in period t if collusion occurs in period t and 0 else. Let $u_t^\delta(h)$ be the time-discounted counterpart. For the sequence of probability measures, $P_{\rho_{\delta_n}}$, passing to a convergent subsequence if necessary, put $P(\cdot) := \lim_{\delta_n \uparrow 1} P_{\rho_{\delta_n}}(\cdot)$ and also set $u^T(h) := \frac{1}{T} \sum_t u_t(h)$. I first claim that

$$\lim_{\delta_n \uparrow 1} \int_{H^*} u^\delta(h) dP_{\rho_{\delta_n}} \leq \limsup_T E_P \hat{u}^T(h),$$

where (i) we truncate the average and put $T = \sum_{i=1}^k T_i$, the sum of the first k work phases, and (ii) $\hat{u}(h)$ is an ‘optimistic’ stream of payoffs that yields (in expectation) an upper bound on actual payoffs. In the appendix, we first find an optimistic payoff stream. Second, we verify that the discounted optimistic streams possess a long-run time average. Assuming these two steps, we verify that time-average payoffs for optimistic streams limit to the first-best value. Suppress the distinction between $\hat{u}(h)$ and $u(h)$ since we soon pass to upper bounds on both quantities.

Put ϕ_P^i equal to the probability that play reaches work phase i , i.e. that players are not absorbed into the punishment state prior to phase i . Notice that $\phi_P^i \rightarrow \phi_P$. Also let \overline{P}^i denote the conditional distribution on work phase i histories. In order for these to be well-defined we require that $\phi_P^i \neq 0$. However, if $\phi_P^i = 0$, then all play ends at the conclusion of phase $i - 1$, in which case we can trivially bound

¹¹This claim invokes Prohorov’s Theorem (e.g. Ch. 16, [Kallenberg \(2002\)](#)).

agents' payoffs as $\delta \uparrow 1$. Notice that we have

$$E_P u^T = \sum_{i=1}^k \phi_P^i E_{\bar{P}^i} u_i^T,$$

where u_i^T denote the partial sum along periods in review phase i . Also note we have:

$$P_{\rho\delta}(C(T_i)/T_i > \epsilon) \rightarrow P(C(T_i)/T_i > \epsilon)$$

and

$$\phi_{P_{\rho\delta}}^i \rightarrow \phi_P^i.$$

By Proposition 4, given any ϵ, ϵ' we can find an index $I(\epsilon, \epsilon')$ such that $\forall \rho \in \Sigma(\delta), \forall \delta$ we have

$$\bar{P}_\rho(C(T_i)/T_i > \epsilon) < \epsilon', \quad \forall i \geq I(\epsilon, \epsilon').$$

Now apply the two equalities:

- $\phi_{P_{\rho\delta}}^i \cdot \bar{P}_{\rho\delta}^i(C(T_i)/T_i > \epsilon) = P_{\rho\delta}(C(T_i)/T_i > \epsilon),$
- $\phi_P^i \cdot \bar{P}^i(C(T_i)/T_i > \epsilon) = P(C(T_i)/T_i > \epsilon).$

and obtain:

$$\bar{P}^i(C(T_i)/T_i > \epsilon) \leq \epsilon', \quad \forall i \geq I(\epsilon, \epsilon').$$

This yields an important bound on the terms $E_{\bar{P}^i} u_i^T$ for $i \geq I(\epsilon, \epsilon')$

$$E_{\bar{P}^i} u_i^T \leq \frac{(1 - \epsilon')\epsilon T_i + \epsilon' T_i}{T}. \quad (1)$$

This is the key inequality in the argument, a variant of which is used in the lower bound argument as well. Let us explain the bound on the numerator. Fix a utility stream in work phase i , think of this as a vector of 0's and 1's – where a 1 denotes the payoff from an episode of collusion and a 0 denotes the payoff from the efficient equilibrium. By Proposition 4, we know that at least a fraction $1 - \epsilon'$ of these streams is comprised of at most $\epsilon \cdot T_i$ episodes of collusion – call this the ‘almost efficient streams’. The remaining streams (of size at most ϵ') are comprised of up to T_i episodes – call these the ‘collusive streams’. The numerator is now obtained by taking the sum of the coordinates of utility streams and weighting according to whether they are efficient streams or collusive streams.

Using the trivial bound $\phi_P^i \leq 1$ we then obtain,

$$E_P u^T \leq \sum_{i=1}^{I(\epsilon, \epsilon')} \frac{T_i}{T} + \sum_{i=I(\epsilon, \epsilon')+1}^k \frac{(1 - \epsilon')\epsilon T_i + \epsilon' T_i}{T}.$$

Recall that $T = \sum_{i=1}^k T_i$, so that by choosing a small (ϵ, ϵ') and then a large k we may push the expected value of time average payoffs to zero.

Lower Bound

We obtain a uniform lower bound for the passing probabilities and verify that this lower bound can be made arbitrarily close to the passing probability in the efficient equilibrium. Let ρ_{δ_n} be any sequence with $(1 - \delta_n)W(\rho_{\delta_n})$ converging to $\liminf (1 - \delta_n)W(\delta_n)$. Let P_{ρ_n} be the sequence of measures on H^* and, passing to a subsequence if necessary, let $P(\cdot)$ denote the limit measure. Let $KS_i(h)$ denote the value of the Kolmogorov-Smirnov test in phase i along history h . Let ϕ_P denote the probability assigned to histories where agents never enter the (absorbing) punishment state.

Step 1: Bounds on passing (failure) probabilities.

We bound ϕ_P from below. Using a different normalization of (stage) payoffs, let $\pi^{FB} > 0$ denote the per-period surplus to an agent from the efficient equilibrium, i.e. equal to the ϵ of insurance that is above the cost of high effort. Let 1 denote the maximum (per-period) payoff from collusion, and let 0 be the reservation utility. Also let t denote the per-period participation fee to participate in the mechanism \mathcal{C} .¹² Then we must have,

$$\int_{H^*} u^\delta(h) dP_{\rho_{\delta_i}} \geq t, \forall \delta_i. \quad {}^{13}$$

We argue as in the upper-bound proof. Note that we have the inequality

$$\lim_{\delta} \int_{H^*} u^\delta(h) dP_\delta \leq \limsup_T E_P u^T,$$

where $P(\cdot)$ is the limit measure and $T = \sum_{i=1}^k T_i$. Putting ϕ_P equal to the probability of the set $\{h : KS_i(h) < \gamma_i, \forall i\}$ (under $P(\cdot)$), observe that we have $\phi_P^i \downarrow \phi_P$.¹⁴ Fix any $\alpha > 0$ small and find I_α such that $\phi_P^i \leq \phi_P + \alpha, \forall i \geq I_\alpha$. Put $I^* := \max\{I(\epsilon, \epsilon'), I_\alpha\}$ and use the following variation on inequality (1):

$$E_P u^T \leq \sum_{i=1}^{I^*} \frac{T_i}{T} + (\phi_P + \alpha) \cdot \sum_{i=I^*+1}^k \frac{(1 - \epsilon')(\epsilon T_i + \pi^{FB}(1 - \epsilon)T_i) + \epsilon' T_i}{T}. \quad (2)$$

¹²We will leave the choice of fee unspecified for now. Once we reach the step in the proof where this is required, we then reverse engineer it to be the quantity that delivers the required lower bound.

¹³The term on the LHS denotes discounted expected utility from wages, absent the per-period participation fee. It is in service of this step that we chose not to absorb the fee t_ϵ into the spot wage function in the definition of the stage contract $\mathcal{C}(\epsilon, R_1, R_2, w_\epsilon^{FB})$.

¹⁴Observe that participation implies $\phi_P^i \neq 0, \forall i$.

Let ϵ_1 denote the first sum and recall that participation requires

$$\limsup_k E_P u^T \geq t.$$

This yields the following lower bound on ϕ_P : (ignore the correction term $\sum_{i \geq I^*+1} T_i/T$, which limits to 1)

$$\phi_P \geq \frac{t - \epsilon_1}{(1 - \epsilon')(\epsilon + (1 - \epsilon)\pi^{FB}) + \epsilon'} - \alpha. \quad (3)$$

Since $T = \sum_{i=1}^k T_i$ we can first choose $\epsilon, \epsilon', \alpha$ small enough and then choose k large enough (which, in combination, assures that ϵ_1 is also small) to obtain:

$$\phi_P \geq \frac{t}{\pi^{FB}}.$$

Now select participation fees. Choose a small ϵ_2 and put $t := (\phi_{FB} - \epsilon_2) \cdot \pi^{FB}$, where ϕ_{FB} denotes the probability of never failing the KS-test under the efficient equilibrium. We obtain:

$$\phi_P \geq \frac{t}{\pi^{FB}} = \frac{(\phi_{FB} - \epsilon_2)\pi^{FB}}{\pi^{FB}} = \phi_{FB} - \epsilon_2.$$

It follows that, as ϕ_{FB} is close to unity, the (limiting) equilibrium probability ϕ_P is close to unity as well. In particular, the argument shows that this is achieved by controlling the contract choice parameters ϕ_{FB}, ϵ_2 .

Let us briefly summarize what has been shown. Starting with any sequence of equilibrium measures $P_{\rho_{\delta_n}}$ converging to P we have found, as a consequence of the participation constraint, a lower bound on the passing probability ϕ_P . This lower bound has two important properties. First, it can be made, by choosing appropriate participation fees, arbitrarily close to unity. Second, it is independent of the limit measure $P(\cdot)$. Finally, observe that by choosing (at the contract design stage) ϕ_{FB} large enough and ϵ_2 small enough we can ensure that, for δ large, repetition of the efficient stage SPNE satisfies the participation constraint.

The implication on the tail probability is not only sufficient, but also necessary for an efficient lower bound. To see this, consider a contract \mathcal{C} with the property that there is a sequence of equilibria ρ_δ with limit measure P such that $\phi_P \leq \kappa$. Then along this sequence we get a limiting upper bound on the principal's payoff:¹⁵

$$\int_{H^*} w(h) dP \leq \phi_P \cdot 1 + (1 - \phi_P) \cdot \epsilon' \leq \kappa + (1 - \kappa) \cdot \epsilon', \forall \epsilon'.$$

¹⁵We take limits on δ in the integral for both the integrand and the measure, which can be justified by standard convergence arguments – which we omit.

Let us justify this upper bound. For histories $h \in \text{supp}(\phi_P)$ we bound $w(h)$ from above by assuming the principal gets the normalized first-best payoff in every period (set at 1). For $h \in \text{supp}(\phi_P)^c$ there is a first time of failure. Outside a set of measure at most ϵ' , the first failure time occurs at the latest by some $N_{\epsilon'}$. All histories in this set have time-average payoff equal 0. On the exceptional set we bound the payoff at 1. Hence, the upper bound on ϕ_P controls the limiting value of the principal's payoff, so that if it is bounded away from 1, the lower bound on the principal's payoff cannot be first-best.

Step 2: Computing the principal's payoff.

Fix an (ϵ, ϵ') . Apply Proposition 5 to obtain a contract \mathcal{C}_N that attains the bounds on the distribution of the variable $E(T_i)/T_i$. In particular, for this pair (ϵ, ϵ') and the corresponding contract \mathcal{C}_N we obtain $\forall \rho \in \Sigma(\delta)$ and for all large δ

$$(\clubsuit) \quad \bar{P}_\rho\left(\frac{E(T_i)}{T_i} \geq \epsilon\right) < \epsilon'$$

This is a key bound to keep in mind, as it is important but only makes an appearance near the conclusion of the forthcoming argument. Let $w_t(h)$ denote the time t payoff to the principal along history h , and put $w^\delta(h)$ to be the discounted sum of time t payoffs. Now let ρ_{δ_n} be any sequence in $\Sigma(\delta)$ such that $(1 - \delta_n) \cdot W(\rho_{\delta_n})$ converges to $\liminf_\delta (1 - \delta)W(\delta)$ and let $P_{\rho_{\delta_n}}$ be the associated sequence of measures on H^* . Passing to a subsequence if necessary, let P denote the limit measure. We now estimate the integral

$$\lim \int_{H^*} w^\delta(h) dP_{\rho_\delta}$$

from below. We verify in the appendix that the expected value of discounted sums of payoff streams can be bounded below by (expected) discounted sums of a particular set of 'pessimistic' payoff streams. These streams possess a long-run average so that we can pass from discounted sums to time-average payoffs (by Abel's Theorem, see Radner (1985), pg. 1175). Moreover, the time-average payoff to the principal along these streams is close to first-best. With this outline, let us derive a lower bound for (a subsequence of) the limiting time T averages and show that it is arbitrarily close to the first best payoff for the principal. For simplicity, for the remainder of the argument we deal only with time averages.

Let $w^T(h)$ denote the time average of the principal's payoff from 1 to T (along history h). As before, compute this limit along the subsequence $T := \sum_{i=1}^k T_i$, i.e. look at segments of time that cover complete work phases. Let w^* denote the total wages paid out when all agents report that they are working hard, and let $x_t(h)$ denote the value of output in period t (along h). Then, we obtain:

$$(*) \quad w^T(h) \geq \frac{\sum_{t=1}^T x_t(h) - Tw^*}{T}.$$

For large T , the contribution to the average from $\sum_{t=1}^{T'} x_t$, where $T' = \sum_{i=1}^{I(\epsilon, \epsilon')} T_i$, is at most some small $\hat{\epsilon}$. From (*) we obtain:

$$(**) E_P w^T(h) \geq \frac{\sum_{t=1}^T E_P X_t - T w^*}{T} = \int_{(\mathbf{e}^1, \dots, \mathbf{e}^T)} \left[\frac{\sum_{t=T'+1}^T E X_t(\mathbf{e}) - w^* T}{T} - \hat{\epsilon} \right] dP_{(\mathbf{e}^1, \dots, \mathbf{e}^T)}.$$

The notation $X_t(\mathbf{e})$ stands for the output r.v. conditional on collective effort choice of \mathbf{e} . Note that going from the second integral to the third we have changed the domain of integration to the domain of ‘effort histories’. We now obtain a lower bound on the RHS of (**), which requires a bound on the terms $E(X_1 + \dots + X_{T_i})$. For this we will use the identity (this is an important ingredient in the proof of Proposition 2 as well)

$$E(X_1 + \dots + X_{T_i}) = \sum_{(\mathbf{e}^1, \dots, \mathbf{e}^{T_i})} \phi_P^i [E X_1(\mathbf{e}^1) + \dots + E X_{T_i}(\mathbf{e}^{T_i})] \bar{P}(\mathbf{e}^1, \dots, \mathbf{e}^{T_i}),$$

where $E X(\mathbf{e}^j)$ is the conditional expectation of output in period j conditional on effort choice \mathbf{e}^j (and on play entering phase i), $\bar{P}(\cdot)$ is the conditional distribution on phase i histories, and ϕ_P^i is the probability that play enters phase i . Let $E(T_i)$ denote the number of periods in which effort choice is not first best (in phase i). Proposition 5 and the fact that \bar{P}_{ρ_δ} limits to \bar{P} implies that (using ♣)

$$\bar{P}\left(\frac{E(T_i)}{T_i} \geq \epsilon\right) \leq \epsilon', \forall i \geq I(\epsilon, \epsilon').$$

Let $E X^*$ be the value under first-best effort, $E X_*$ is the value under the worst effort choice, and $\Pi^{FB} = E X^* - w^*$, $\underline{\Pi} = E X_* - w^*$. Ignore the correction term $\sum_{i=I(\epsilon, \epsilon')+1}^k T_i/T_i$, which limits to 1 for large k . We obtain the following bound:

$$\int_{(\mathbf{e}^1, \dots, \mathbf{e}^T)} \frac{\sum_{t=T'+1}^T E X_t(\mathbf{e}) - w^* T}{T} dP_{(\mathbf{e}^1, \dots, \mathbf{e}^T)} \geq \phi_P [(1-\epsilon') \cdot ((1-\epsilon)\Pi^{FB} + \epsilon\underline{\Pi}) + \epsilon'\underline{\Pi}] - (1-\phi_P)w^*.$$

Let us explain this bound. The $(1-\epsilon')$ term gives a lower bound on the (conditional) probability of the set

$$\{(\mathbf{e}^1, \dots, \mathbf{e}^{T_i}) : E(T_i)/T_i < \epsilon\}$$

for each $I(\epsilon, \epsilon') \leq i \leq k$, and ϕ_P is a lower bound on the probability that any of these phases are reached. The coefficient ϵ' is an upper bound on the conditional probability of the complement and the term $E X_* - w^*$ is the minimum value of the integrand on this set. Finally, as $\phi_P^i \downarrow \phi_P$ by replacing each ϕ_P^i with ϕ_P we maintain a lower bound. This yields the lower bound on (**). \square

It remains to justify some details. First, we need to verify that sequences of measures possess convergent subsequences. This is done by constructing a specific

metric on H^* whose open sets generate the domain of the measures P_ρ . Second, we need to justify passing from expected utility of normalized discounted payoffs to (limits of) expected utility of time averages. Third, in computing time averages we have just computed limits over a subsequence of times. As an aside, this is where we use the particular choice of sample sizes, viz. the fact that they have polynomial growth. These details are addressed in the appendix as (resp.) Proposition 6 and Observation 2, Observation 3.

The proof makes ample use of the statistical review mechanism, but it may not be as evident where we use the ‘short-run’ monitoring incentives. The monitoring mechanism ensures that collusion stops once the KS statistic falls into the rejection region. In the absence of this structure, for example if we unconditionally provide full insurance then players would simply (in equilibrium) take the insurance and shirk till the resumption of the next review phase. This is important for us since the proof of Proposition 4 and, consequently, Proposition 5 uses the observation that collusion stops once it has been detected by the KS test.

Thus, we require both incentive components for the proof of our result. This leaves open the issue of whether the same features are required of any contractual mechanism that attains the same result as Theorem 1. A natural question – in light of its ubiquity in repeated agency models – is whether one can obtain our main result by simply using the Radner (1985) contract. There are non-trivial obstacles to this approach. First, with multiple agents we cannot assume that equilibrium actions of a team coincide with those of a cooperative. Team equilibria in the infinite horizon game can be inefficient, and it is easy to construct examples where the set of equilibria (induced by Radner contracts) might yield an approximately efficient bound for the agents, yet yield a folk theorem for the principal. Radner’s analysis also requires stationarity, which is wlog for a cooperative problem. For example, assuming stationarity he obtains recursive expressions for the agent and principal’s payoff which allows him to solve for the (stationarity) success probability and, in turn, to verify a lower bound on the principal’s payoff. In contrast, stationarity is not without loss of generality in our model. Our approach shifts the analysis to distributions on equilibrium histories, which admits a direct analysis of equilibrium behavior.

5 Conclusion

This paper considers an infinite horizon repeated moral hazard problem with multiple agents. The environment is like the canonical principal-agent model except that between the time when effort is taken and output is realized, agents can be required to communicate their observations (of co-workers’ effort choices) to the principal.

Our main result constructs a class of contracts with the following features. First, for all (SPNE) equilibria in the game induced by the contract, payoffs for both the principal and the agents can be pushed arbitrarily close to their first best benchmarks (as the discount factor gets large). Second, while the set of SPNE is a function of the discount factor, the contract itself is independent of the discount factor.

Contrast the review contracts constructed in [Radner \(1985\)](#) with the construction in this paper. Radner review contracts are described by a sequence of hypothesis tests with two key properties. First, review tests are stationary and exhibit ‘lack of memory’, i.e. past review outcomes are by-gones and have no bearing on subsequent reviews. Second, the two main parameters of the review contract, (i) the review length and (ii) the punishment length, increase without bound as the discount factor increases. Both features are required for the Radner argument.

Our contracts maintain an abstract review structure, i.e. they are comprised of test phases and punishment phases, but there are two substantive changes. First, the reviews are becoming more refined over time. Moreover, they become refined quickly enough that we still obtain limit efficiency, e.g. Type I errors can be made arbitrarily small. Second, the same reviews apply to all discount factors. Hence, a patient agent receives the same contract as his relatively less patient counterpart. Let us elaborate on the value of this latter feature.

Most repeated games analysis assumes a common and known discount factor across players. This is innocuous from a repeated games perspective since the stage game is a primitive of the repeated games problem. However, in our case the particular repeated (more precisely, stochastic) game that is played is induced by a choice of contract between a principal and agent(s). In other words, fixing time preferences, the game itself is changing as a function of the agreed upon contract. In this setting, the assumption is more substantive. What if discount factors differ and, more practically, the principal cannot truthfully elicit the agent’s discount factor (which is subjective) before designing the contract? For example, if the agent anticipates that the principal is going to use a Radner review contract, then he has an incentive to understate his rate of time preference, i.e. make himself look much more patient than he really is. With our contracts, we can get around the preference elicitation problem as the same review mechanism applies to all agents regardless of their discount factor.

6 Appendix - Omitted Proofs

6.1 Omitted Proofs from Section 2

Proof of Proposition 1. Proceed in two steps. In the first step we show the existence of incentive compatible reward schemes that are designed to elicit truthful observation reports. In the second step, using these reward schemes we describe the contract and characterize the equilibria of the extensive form game induced by the contract.

Step 1

Put $\mathbf{e}^* := (e_H, \dots, e_H)$ and consider the set of effort choice vectors $\mathbf{E} \setminus \mathbf{e}^*$. The collection $\{F(\cdot|\mathbf{e})\}$ is partially ordered under FOSD. Put $\mathbf{e}' := (e_H, \dots, e_H, e_L)$ and note that $F(\cdot|\mathbf{e}') \succeq_{FOSD} F(\cdot|\mathbf{e}), \forall \mathbf{e} \in \mathbf{E} \setminus \mathbf{e}^*$. That is, $F(\cdot|\mathbf{e}')$ is (weakly) \succeq_{FOSD} -maximal in $\mathbf{E} \setminus \mathbf{e}^*$. Now consider the cdf's $F(\cdot|\mathbf{e}'), F(\cdot|\mathbf{e}^*)$. By FOSD these cdf's are distinct so that there is a least integer $m (\neq |X|)$ such that $F(x_m|\mathbf{e}') > F(x_m|\mathbf{e}^*)$. Thus, the system

$$\begin{aligned} (1 - F(x_m|\mathbf{e}'))R_1 + F(x_m|\mathbf{e}')R_2 &= A \\ (1 - F(x_m|\mathbf{e}^*))R_1 + F(x_m|\mathbf{e}^*)R_2 &= B. \end{aligned}$$

has a unique solution with $R_1 < 0, R_2 > 0$ for any $A > 0, B < 0$. Choosing A, B to be arbitrarily small, we can make R_1, R_2 arbitrarily small. As $u(\cdot)$ is C^2 , the quadruple $(w_\epsilon^{FB}, w_0, R_1, R_2)$, where $w_\epsilon^{FB} := u^{-1}(u_0 + c(e_H) + \epsilon), w_0 := u^{-1}(u_0)$ and R_1, R_2 are sufficiently small, solves the following system:

1. **IR:** $u(w^{FB}) - c(e_H) \geq u_0$
2. **IC₁:** $(1 - F(x_m|\mathbf{e}'))u(w^{FB} + R_1) + F(x_m|\mathbf{e}')u(w^{FB} + R_2) > u(w^{FB})$
3. **IC₂:** $(1 - F(x_m|\mathbf{e}^*))u(w^{FB} + R_1) + F(x_m|\mathbf{e}^*)u(w^{FB} + R_2) < u(w^{FB})$.

By maximality of $F(\cdot|\mathbf{e}')$, if this system holds with the given choice of R_1, R_2 then:

$$(1 - F(x_m|\mathbf{e}))u(w^{FB} + R_1) + F(x_m|\mathbf{e})u(w^{FB} + R_2) > u(w^{FB}).$$

for any $\mathbf{e} \in \mathbf{E} \setminus \mathbf{e}^*$. The two IC's are truth-telling conditions. IC_1 implies that agents report shirking to the principal. IC_2 ensures that it is not profitable to issue a false report.

Step 2

Give agents labels $1, \dots, n$ and let $r_{i,i+1} \pmod n$ denote agent i 's report on agent $i+1 \pmod n$. Require agents to make announcements sequentially. That is, i makes his

report on $i+1$. Upon hearing this, $i+1$ reports on $i+2$, etc. Let $\epsilon' > 0$ and consider the following wage contract (let $w(i) =$ agent i 's wage):

$$w(i) = \begin{cases} w^{FB}, & \text{if } r_{i-1,i} = e_H, r_{i,i+1} = e_H \\ (w_\epsilon^{FB} + R_1)1_{(x > x_m)} + (w_\epsilon^{FB} + R_2)1_{(x \leq x_m)}, & \text{if } r_{i-1,i} = e_H, r_{i,i+1} \neq e_H \\ w_0 - \epsilon', & \text{if } r_{i-1,i} \neq e_H, r_{i,i+1} = e_H \\ w_0, & \text{if } r_{i-1,i} \neq e_H, r_{i,i+1} \neq e_H. \end{cases}$$

Note that (i) choosing e_H and reporting truthfully is an equilibrium, and (ii) there is no equilibrium where collusion occurs on the equilibrium path. Let us verify (ii). Fix a history where effort choice is less than first-best, and consider agent n . By IC_1 , it is strictly dominant for him to report on agent 1. Thus, along any equilibrium history where effort choice is less than first-best, agent n is always reporting on agent 1. But then, along any such history it is a strict best-response for agent 1 to report on agent 2, and thus for 2 to report on agent 3, and so on. Hence, along any history where collective effort is less than first-best, everyone is reporting on each other and nobody is earning higher than the reservation wage. \square

6.2 Omitted Proofs from Section 3

Proof of Proposition 2. All probabilities and expectations are conditional in the forthcoming proof, although we will suppress the \bar{P} notation. Hence, for a given phase n , when we write P we mean \bar{P} and when we write $E_\rho(\cdot)$ we mean $E_{\bar{P}}(\cdot)$. Wlog assume that $\alpha > 0$ is such that $\lim P_\rho(\frac{C(T_n)}{T_n} > \alpha) = \beta > 0$. Towards a contradiction, assume that $\forall r \in (0, 1)$, $P_\rho(\sup_x |F_{T_n}(x) - F(x|\mathbf{e}^*)| > r) \rightarrow 0$. We will now compute the limit of $E_\rho(\frac{X_1 + X_2 + \dots + X_{T_n}}{T_n})$ in two different ways, obtaining two different answers as a consequence of this assumption.

Method 1

Assuming that $\forall r \in (0, 1)$, $P_\rho(\sup_x |F_{T_n}(x) - F(x|\mathbf{e}^*)| > r) \rightarrow 0$ we obtain, $E_\rho |F_{T_n}(x) - F(x|\mathbf{e}^*)| \leq 2 \cdot P_\rho(|F_{T_n}(x) - F(x|\mathbf{e}^*)| > r) + r \cdot P_\rho(|F_{T_n}(x) - F(x|\mathbf{e}^*)| \leq r) \rightarrow r$, $\forall x$. Since this holds for all $r \in (0, 1)$, it follows that we may find $\zeta_n \rightarrow 0$ such that

$$-\zeta_n \leq E_\rho(F_{T_n}(x) - F(x|\mathbf{e}^*)) \leq \zeta_n, \forall x. \quad (4)$$

Relabel the per-period output variables, in phase n , so that X_i denotes output in period i (notation for the work phase within which it occurs is suppressed). Let $X_i \sim \mu_i$ and let $\mu^*(\cdot)$ denote the measure that corresponds to $F(\cdot|\mathbf{e}^*)$. Choose a labeling of the support of output $X := \{x_1, \dots, x_l\}$ such that $x_i < x_{i+1}$. Plugging $x = x_k$ into (4) gives:

$$-\zeta_n \leq \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(\{x_j\}_{j=1}^k) - F(x_k|\mathbf{e}^*) \leq \zeta_n \quad (5)$$

equivalently,

$$F(x_k|\mathbf{e}^*) - \zeta_n \leq \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(\{x_j\}_{j=1}^k) \leq F(x_k|\mathbf{e}^*) + \zeta_n.$$

This implies

$$F(x_{k+1}|\mathbf{e}^*) - \zeta_n - \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(\{x_j\}_{j=1}^k) \leq \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(x_{k+1}) \leq F(x_{k+1}|\mathbf{e}^*) + \zeta_n - \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(\{x_j\}_{j=1}^k).$$

Using the bounds from (5) we then obtain:

$$\mu^*(x_{k+1}) - 2\zeta_n \leq \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(x_{k+1}) \leq \mu^*(x_{k+1}) + 2\zeta_n.$$

We deduce that for each j :

$$\mu^*(x_j)x_j - 2x_j\zeta_n \leq \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(x_j)x_j \leq \mu^*(x_j)x_j + 2x_j\zeta_n. \quad (6)$$

Now note that

$$E_\rho\left(\frac{X_1 + X_2 + \dots + X_{T_n}}{T_n}\right) = \frac{1}{T_n} \sum_{i=1}^{T_n} \sum_{j=1}^l \mu_i(x_j)x_j = \sum_{j=1}^l \frac{1}{T_n} \sum_{i=1}^{T_n} \mu_i(x_j)x_j. \quad (7)$$

Let $X \sim \mu^*$. Sum inequality (6) on j and apply equation (7) to obtain:

$$EX - \zeta_n \sum_{j=1}^l 2x_j \leq E_\rho\left(\frac{X_1 + X_2 + \dots + X_{T_n}}{T_n}\right) \leq EX + \zeta_n \sum_{j=1}^l 2x_j.$$

Since $\zeta_n \rightarrow 0$ we obtain that $E_\rho\left(\frac{X_1 + X_2 + \dots + X_{T_n}}{T_n}\right) \rightarrow EX$ (as $n \rightarrow \infty$).

Method 2

Let $\mathbf{E} = \{\mathbf{e}_1, \dots, \mathbf{e}_p\}$ be a labeling of the set of collective effort choice vectors. Let $X(\mathbf{e}_i)$ denote a r.v. with c.d.f. $F(\cdot|\mathbf{e}_i)$. Since \mathbf{e}^* is first-best we know that $EX(\mathbf{e}^*) > EX(\mathbf{e}_i), \forall i \neq 1$. Put $EX' := \max_{\{\mathbf{e}_i \neq \mathbf{e}^*\}} EX(\mathbf{e}_i)$. Let E_i^j be the r.v. that denotes effort choice in phase i , period j . For notational economy, suppress the i -subscript and write E_j for E_i^j . Note that the variable X_j is distributed as $X(\mathbf{e}_i)$ conditional on $E_j = \mathbf{e}_m$. This allows us to write: $E_\rho X_j = \sum_{m=1}^p P_\rho(E_j = \mathbf{e}_m) \sum_{k=1}^l P_\rho(X_j = x_k | E_j = \mathbf{e}_m) x_k$. Put $X(\mathbf{e}_m) \sim \mu_{\mathbf{e}_m}$. Since

$P_\rho(X_j = x_k | E_j = \mathbf{e}_m) = \mu_{\mathbf{e}_m}(x_k)$ we obtain: $EX_j = \sum_{m=1}^p P_\rho(E_j = \mathbf{e}_m) \cdot EX(\mathbf{e}_m)$. Let \mathbf{e}^i denote a realized value of E_i . It follows that we may write

$$E_\rho X_1 + E_\rho X_2 + \dots + E_\rho X_{T_n} = \sum_{(\mathbf{e}^1, \mathbf{e}^2, \dots, \mathbf{e}^{T_n})} P_\rho(\mathbf{e}^1, \mathbf{e}^2, \dots, \mathbf{e}^{T_n}) [EX(\mathbf{e}^1) + \dots + EX(\mathbf{e}^{T_n})].$$

By hypothesis, $\exists \alpha, \beta > 0$ such that $P_\rho(\frac{C(T_n)}{T_n} \geq \alpha) \geq \beta, \forall n$. If we define $E(T_n)$ to be the number of periods in which collective effort choice is less than first-best, then we have $E(T_n) \geq C(T_n)$. Thus, $\{\frac{E(T_n)}{T_n} \geq \alpha\} \supseteq \{\frac{C(T_n)}{T_n} \geq \alpha\}$. Note that, for each realization $(\mathbf{e}^1, \dots, \mathbf{e}^{T_n})$, either $\{(\mathbf{e}^1, \dots, \mathbf{e}^{T_n})\} \subseteq \{\frac{E(T_n)}{T_n} \geq \alpha\}$ or $\{(\mathbf{e}^1, \dots, \mathbf{e}^{T_n})\} \cap \{\frac{E(T_n)}{T_n} \geq \alpha\} = \emptyset$. Let $[\alpha T_n]$ denote the integer part of αT_n . Since $P_\rho(\frac{E(T_n)}{T_n} \geq \alpha) \geq \beta$ we have the following upper bound:

$$E_\rho X_1 + E_\rho X_2 + \dots + E_\rho X_{T_n} \leq ((1 - \beta)T_n + \beta(T_n - [\alpha T_n]))EX + \beta[\alpha T_n]EX'.$$

Thus,

$$\begin{aligned} EX - E_\rho\left(\frac{X_1 + X_2 + \dots + X_{T_n}}{T_n}\right) &\geq \left(1 - \frac{((1 - \beta)T_n + \beta(T_n - [\alpha T_n]))}{T_n}\right)EX - \frac{\beta[\alpha T_n]}{T_n}EX' \\ &= \frac{\beta[\alpha T_n]}{T_n}(EX - EX'). \end{aligned}$$

For large n the RHS is at least $\frac{\beta\alpha}{2}(EX - EX') > 0$. This contradicts the limit obtained by the first method. \square

Remark: Note that the proof only requires the hypothesis that $P_\rho(E(T_n)/T_n \geq \alpha) \geq \beta$, where $E(T_n)$ counts the number of periods in phase n where effort choice is not first-best. Since $C(T_n) \subseteq E(T_n)$, it is sufficient if the inequality is satisfied for $C(T_n)$. In the proof of Proposition 5 we will use a formulation of Proposition 2 with $E(T_n)$ in place of $C(T_n)$.

Proof of Proposition 3. Since an (equilibrium) strategy and its rectified companion (which is also an equilibrium) have the same $C(T_n)/T_n$ distribution, we will assume that all strategies are rectified. Furthermore, suppress the $\bar{P}(\cdot)$ notation - keeping in mind that all probabilities in question are conditional probabilities, where for each work phase i we condition on the set of histories that reach work phase i . Towards a contradiction, assume that $C(T_n)/T_n \not\rightarrow 0$ (in probability) and let r be the constant given by Proposition 2. Put

$$B_i(r) := \{\sup_x |F_{T_i}(x) - F(x|\mathbf{e}^*)| \geq r\}.$$

For $h \in B_n(r)$ we have

$$r \leq \sup_x |F_{T_n}(x) - F(x|\mathbf{e}^*)| \leq \sum_{i=1}^n \frac{1}{n} \cdot \sup_x |F_{i,T_n}(x) - F(x|\mathbf{e}^*)|. \quad (8)$$

Introduce the following r.v.'s,

- $M_n(h) := |\{i : \sup_x |F_{i,T_n}(x) - F(x|\mathbf{e}^*)| \geq \gamma_n\}|$,
- $C_n := \{h : M_n(h) \geq 2\}$.

Thus, $M_n(h)$ counts the number of (sub)phases within the i th work phase in which the margin of error is surpassed (along history h) and $C_n(h)$ is the set of histories for which this happens at least twice. By inequality (8), $\exists N(r) \gg 0$ such that $\forall n \geq N(r)$, $M_n(h) \geq 2$ whenever $h \in B_n(r)$. This implies $B_n(r) \subseteq C_n$, $\forall n \geq N(r)$. We claim that $P_\rho(C_n) \rightarrow 0$. Define $\{h^{N_n(h)t_n}\}$ to be the set of histories that agree with h up through the first failure time along h , $N_n(h)t_n$, and note that the sets $\{h^{N_n(h)t_n}\}$ are *disjoint* sets of histories. Put

$$C_n = \sqcup_h (\{h^{N_n(h)t_n}\} \cap C_n).$$

Observe that

$$P_\rho(C_n) = \sum P_\rho(C_n \cap \{h^{N_n(h)t_n}\} | h^{N_n(h)t_n}) \cdot P_\rho(\{h^{N_n(h)t_n}\}).$$

Note that

$$P_\rho(C_n | h^{N_n(h)t_n}) = P_\rho(\{\exists i > N_n(h) \text{ s.t. } \sup_x |F_{i,T_n}(x) - F(x|\mathbf{e}^*)| \geq \gamma_n\} | h^{N_n(h)t_n}).$$

Since ρ is rectified (wlog - as this doesn't change the \bar{P}_ρ distribution of $C(T_n)/T_n$), for histories following $h^{N_n(h)t_n}$ the output process in every period follows the law $F(\cdot|\mathbf{e}^*)$. Let $P(\cdot)$ denote the product measure on the sample space $\prod_{i=1}^{t_n} X$ - generated by t_n i.i.d draws from the distribution $F(\cdot|\mathbf{e}^*)$. By Chebyshev's inequality, we have

$$P(\sup_x |F_{t_n}(x) - F(x|\mathbf{e}^*)| \geq \gamma_n) \leq \frac{K}{\gamma_n^2 t_n} = \epsilon_n$$

where $K = l \cdot \max_x \text{Var}(1_{(X \leq x)})$ (and l is the cardinality of the range of output). Now define

$$Y_i = \begin{cases} 1 & \text{iff } \sup_x |F_{i,T_n}(x) - F(x|\mathbf{e}^*)| \geq \gamma_n \\ 0 & \text{else.} \end{cases}$$

This gives:

$$P_\rho(\{\sup_x |F_{i,T_n}(x) - F(x|\mathbf{e}^*)| \geq \gamma_n, i > N_n(h)\} | h^{N_n(h)t_n}) \leq P(Y_i = 1 \text{ for some } i, 1 \leq i \leq n).$$

Note that $\{Y_i\}_{i=1}^n$ are i.i.d and $P(Y_i = 1) \leq \epsilon_n$, by the Chebyshev bound. The fact that $n\epsilon_n \rightarrow 0$ then implies,

$$\begin{aligned} P_\rho(C_n) &= \sum P_\rho(C \cap \{h^{N_n(h)t_n}\} | h^{N_n(h)t_n}) \cdot P_\rho(h^{N_n(h)t_n}) \\ &\leq \sum P(Y_i = 1 \text{ for some } i, 1 \leq i \leq n) \cdot P_\rho(h^{N_n(h)t_n}) \\ &\leq n\epsilon_n \cdot \sum P_\rho(h^{N_n(h)t_n}) \leq n\epsilon_n \rightarrow 0. \end{aligned}$$

This contradicts that $B_n(r) \subseteq C_n, \forall n \geq N(r)$ and $\overline{\lim}_n P_\rho(B_n(r)) > 0$, by Proposition 2. \square

6.3 Omitted Proofs from Section 4

Proposition 6. *There is a metric $d : H^* \times H^* \rightarrow \mathbf{R}_+$ such that, under this metric, the space (H^*, d) is compact and the open balls under this metric are exactly the cylinder sets of histories in H^* .*

Proof. Let H^i denote the set of histories in work phase i . Choose an embedding $\kappa_i : H^i \hookrightarrow (0, 1]$. In order to make open balls coincide with cylinder sets I will need the maps κ_i to have specific properties, but for the moment this will be left unspecified. Also define $\kappa_i(\emptyset) = 0$. Now for any history $h \in H^*$ write $h = (h^1, h^2, \dots, h^i, \dots)$, where $h^i \in H^i$. This is excepting the possibility that punishment is incurred along h prior to phase i . When this happens we have $h^n = \emptyset, \forall n \geq i$. Let $\kappa(h) := (\kappa_1(h^1), \kappa_2(h^2), \dots, \kappa_i(h^i), \dots)$. Define a metric $d(\cdot, \cdot)$ on H^* via the formula:

$$d(h, h') = \sum_{i=1}^{\infty} \alpha_i \cdot |\kappa_i(h^i) - \kappa_i((h')^i)|$$

for constants $\alpha_i \in (0, 1)$ soon to be specified. For the compactness argument we just need the assumption that the sequence $\{\alpha_i\}$ is summable. Let us first check that this makes (H^*, d) a compact metric space. Let x_n denote any infinite sequence of elements of H^* . For a fixed infinite history, x_n , the term x_n^i denotes the component of the history in phase i . I will find a convergent subsequence of $\{x_n\}$ by repeated application of the ‘pigeonhole principle’. Notice that H^1 , the set of possible phase 1 histories is finite. Since $\{x_n\}$ is an infinite sequence (else, we are done), there must be some phase 1 history, h^1 such that $h^1 = x_n^1$ for infinitely many n . More precisely, there must be infinitely many distinct infinite histories, x_n , that share the same phase 1 component h^1 . Pass to this subsequence and, to avoid an infinite regress of subscripts, let us abuse notation and denote this subsequence as $\{x_n(1)\}$.

Now repeat the preceding argument. There are finitely many phase 2 histories comprising H^2 . Hence, there is some h^2 such that infinitely many elements of the sequence $\{x_n(1)\}$ share the phase 2 history h^2 . Note that since all elements of $\{x_n(1)\}$ are distinct, it cannot be the case that play terminates after h^2 , i.e. that punishment is incurred along history h^2 . Select the subsequence of $x_n(1)$ corresponding to those elements which share the same h^2 , and denote this subsequence as $\{x_n(1, 2)\}$. Inductively proceed to define, for each i , a nested collection of subsequences $\{x_n(1, 2, \dots, i)\}$. Define an element of H^* , x_* , via the formula:

$$x_*^i = x_n^i(1, 2, \dots, i)$$

That is, the phase i component of the infinite history x_* agrees with the phase i component of any element of the subsequence $\{x_n(1, 2, \dots, i)\}$, say h^i . Now choose a subsequence of the original $\{x_n\}$ converging to x_* as follows. Let $y_1 = x_{n_1}$ denote any element from the set $\{x_n(1)\}$. Let $y_2 = x_{n_2}$ denote any element from the

set $\{x_n(1, 2)\}$ with subscript larger than n_1 . Inductively proceed and let $\{y_n\}$ be the corresponding subsequence of $\{x_n\}$. Notice that $y_n \rightarrow x_*$ under $d(\cdot, \cdot)$. Hence, (H^*, d) is sequentially compact (and thus, compact).

To conclude I need to verify that the κ_i, α_i can be selected such that the topology generated by the d -open balls is exactly the one generated by cylinder sets in H^* . To this end, we inductively select constants α_i and maps κ_i . For $i = 1$, let $N_1 = |H^1|$ (the cardinality of H^1) and let $h_1^1, \dots, h_{N_1}^1$ be an enumeration of elements of H^1 . Put $\kappa_1(h_i^1) := 1/i$ and choose any summable sequence $\{\alpha_n^1\}$ with $\alpha_n^1 \in (0, 1)$ and such that

$$\frac{\alpha_1^1}{N_1(N_1 + 1)} > \sum_{n=2}^{\infty} \alpha_n^1.$$

For $i = 2$, put $N_2 = |H^2|$ and analogously define κ_2 . Choose a summable sequence α_n^2 with the property that (i) $\alpha_n^2 \leq \alpha_n^1, \forall n$ and

$$\frac{\alpha_2^2}{N_2(N_2 + 1)} > \sum_{n=3}^{\infty} \alpha_n^2.$$

At stage $i = I$, put $N_I = |H^I|$ and define κ_I as before. Choose a summable sequence $\{\alpha_n^I\}$ with the property that $\alpha_n^I \leq \alpha_n^{I-1}, \forall n$ and such that

$$\frac{\alpha_I^I}{N_I(N_I + 1)} > \sum_{n=I+1}^{\infty} \alpha_n^I.$$

Put $\alpha_n := \alpha_n^n$ and let $d(\cdot, \cdot)$ be the metric defined by the datum $(\{\alpha_i\}, \kappa_i)$. For any history h let $\{(h^1, \dots, h^n)\} \subseteq H^*$ denote the cylinder set of histories which coincide with h up through phase n . I claim that for each n there is some ϵ_n such that $B_{\epsilon_n}(h) = \{(h^1, \dots, h^n)\}$. Since

$$\frac{\alpha_n}{N_n(N_n + 1)} > \sum_{m \geq n+1} \alpha_m$$

choose

$$\epsilon_n \in \left(\sum_{m \geq n+1} \alpha_m, \frac{\alpha_n}{N_n(N_n + 1)} \right)$$

and note that $B_{\epsilon_n}(h) = \{(h^1, \dots, h^n)\}$. Furthermore, note that as we refine the cylinder sets by taking narrower cylinders, i.e. increasing n in $\{(h^1, \dots, h^n)\}$, then we can select a strictly decreasing sequence $\{\epsilon_n\}$ such that $B_{\epsilon_n}(h) = \{(h^1, \dots, h^n)\}$.

Using this fact, we can now show an equivalence of topologies. Let U_1 be an open set under the cylinder set topology and take $h \in U_1$. Since U_1 is open and cylinder

sets form a base for the topology, there is some $\{(h^1, \dots, h^n)\} \subseteq U_1$. By the preceding paragraph, $B_{\epsilon_n}(h) \subseteq U_1$. Conversely, let U_2 be an open set under the d -metric and let $h \in U_2$. By d -openness there is some $\epsilon > 0$ such that $B_\epsilon(h) \subseteq U_2$. Find ϵ_n from the preceding paragraph such that $\epsilon_n < \epsilon$ and note that $\{(h^1, \dots, h^n)\} \subseteq B_{\epsilon_n}(h) \subseteq B_\epsilon(h)$. It follows that the topologies coincide. \square

Remark: Note that measures on H^* induced by SPNE have domain the Borel σ -algebra generated by cylinder sets of histories. Since (by Proposition 6) this generating set coincides with the open sets under the constructed metric $d(\cdot, \cdot)$ and the space (H^*, d) is compact under this metric, Prohorov's Theorem (Ch.12, [Kallenberg \(2002\)](#)) can be applied – so that sequences of measures on the space $(H^*, \mathcal{B}(H^*))$ (where $\mathcal{B}(\cdot)$ denotes the Borel σ -algebra generated by cylinders) possess convergent subsequences.

Proof of Proposition 4. Proceed via contradiction. Fix an ϵ and find a sequence of SPNE $\rho_i \in \Sigma(\delta_i)$ and associated work phases n_i such that

$$\lim \bar{P}_{\rho_i} \left(\frac{C(T_{n_i})}{T_{n_i}} \geq \epsilon \right) > 0.$$

Note that we may wlog take each ρ_i to be a rectified profile and this doesn't change the \bar{P}_{ρ_i} -distribution of $C(T_{n_i})/T_{n_i}$. Now let $\mathbf{N} = \sqcup_i \{n_{i-1} + 1, \dots, n_i\}$ be the partition of \mathbf{N} induced by the sequence $\{n_i\}$ and, putting H^j equal to the space of phase j histories, define measures P_j on H^j as follows. Put

$$(*) P_j = \bar{P}_{\rho_i} \text{ iff } j \in \{n_{i-1} + 1, \dots, n_i\}.$$

Let $P_n := \otimes_{i=1}^n P_i$. Abusing notation let H^* now denote the product space $H^* = \prod_{j=1}^{\infty} H^j$. The following observation is a straightforward application of Kolmogorov's extension theorem (see Ch.6, [Kallenberg \(2002\)](#)).

Observation 1. *There is a probability space (H^*, F^*, P_*) , with σ -field $F^* \subseteq 2^{H^*}$ and p.m. P_* , that uniquely extends P_n to H^* .*

Proof of Observation 1. For each j identify each H^j with a discrete subset of $[j, j+1)$ (yielding an embedding $\kappa : H^* \hookrightarrow \mathbf{R}^{\mathbf{N}}$) with corresponding discrete measure P_{ρ_i} , where from $(*)$ we have $j \in \{n_{i-1} + 1, \dots, n_i\}$. For brevity, denote the measure as $P_{\rho(j)}^j$. Thus, we obtain a sequence of measures $P_n = \otimes_{j=1}^n P_{\rho(j)}^j$ (resp. on (\mathbf{R}^n, R^n) , where R^n is the Borel σ -algebra on \mathbf{R}^n). Moreover, note that the sequence is consistent, i.e. $P_{n+1}(A \times \mathbf{R}) = P_n(A)$, for any $A \in R^n$ of the form $A = (A_1, \dots, A_n)$, $A_i \in R$. By Kolmogorov's Extension Theorem, there is an (unique) extension of the P_n , call it μ_* , to $(\mathbf{R}^{\mathbf{N}}, R^{\mathbf{N}})$. Now we can define the infinite product measure. For the domain of the measure we take

$$F^* := \kappa^{-1}(R^{\mathbf{N}}) := \{A \subseteq H^* : \exists B \in R^{\mathbf{N}} \text{ s.t. } A = \kappa^{-1}(B)\}$$

and for $A \in F^*$ we define

$$P_*(A) := \mu_*(B)$$

where B is *any* element of $R^{\mathbf{N}}$ such that $\kappa^{-1}(B) = A$. The key is to check that this gives a well-defined function. Let $B_1, B_2 \in R^{\mathbf{N}}$ be s.t. $A = \kappa^{-1}(B_1) = \kappa^{-1}(B_2)$. Note that we have

- $\kappa(H^*)$ is $R^{\mathbf{N}}$ -measurable (as it is a countable intersection of sets which are each finite unions of G_δ 's in $R^{\mathbf{N}}$).
- $\mu_*(\kappa(H^*)) = 1$ (by the extension property and continuity of $\mu_*(\cdot)$).

Also note that $B_1 \cap \kappa(H^*) = \kappa(A) = B_2 \cap \kappa(H^*)$. It follows that

$$\mu_*(B_1) = \mu_*(B_1 \cap \kappa(H^*)) = \mu_*(\kappa(A)) = \mu_*(B_2 \cap \kappa(H^*)) = \mu_*(B_2)$$

so that P_* is well-defined. For countable additivity let $A_i \in F^*$ be disjoint, and choose any B_i such that $\kappa^{-1}(B_i) = A_i$. Put $B_i^* := B_i \setminus (\cup_{j=1}^{i-1} B_j)$ and note that since the A_i are disjoint, $\kappa^{-1}(B_i^*) = A_i$. It follows that $P_*(\cup_i A_i) = \mu_*(\cup_i B_i^*) = \sum_i \mu_*(B_i^*) = \sum_i P_*(A_i)$. \square

Now apply Proposition 2 and Proposition 3 to the composite measure P_* . Note that, by construction, we have $\overline{\lim} P_*(C(T_n)/T_n \geq \epsilon) > 0$. Hence, $\exists r > 0$ such that $\overline{\lim} P_*(\sup_x |F_{T_n}(x) - F(x|\mathbf{e}^*)| > r) > 0$. Now note that the only input of equilibrium in Proposition 3 was to be able to claim that the profiles generating the conditional measures $\overline{P}_\rho(\cdot)$ on H^{n_i} were rectified. Rectification implies the r.v's $|F_{T_n}(x) - F(x|\mathbf{e}^*)|$ cannot be bounded above some fixed r with probability bounded away from zero – as Type I error probabilities vanish. Hence, we have a contradiction, implying that there could not have been a sequence of alleged counterexample SPNE $\rho_i \in \Sigma(\delta_i)$ with $\overline{P}_{\rho_i}(C(T_{n_i})/T_{n_i} \geq \epsilon) \not\rightarrow 0$. It follows that, given $\epsilon > 0, \beta > 0$, there is some integer $N(\epsilon, \beta)$ such that $\forall n \geq N(\epsilon, \beta)$ we have

$$\overline{P}_\rho\left(\frac{C(T_n)}{T_n} \geq \epsilon\right) < \beta$$

where the bound holds for all $\rho \in \Sigma(\delta)$ and for all $\delta \in [0, 1)$. \square

Proof of Proposition 5. We use the result of Step 1 from the lower bound argument. That shows that for *any* $\mathcal{C} \in \Phi$ we have (for participation fee set equal to $t := \phi_{FB} - \epsilon_2$):

$$\lim_{\delta \uparrow 1} \phi_{P_{\rho_\delta}} \geq \phi_{FB} - \epsilon_2$$

where $\rho_\delta \in \Sigma(\delta)$. Use the following equivalent formulation of Proposition 2.

Claim 1 (Proposition 2 – Equivalent Formulation). *Fix $\alpha, \beta > 0$. Then, there are numbers $f(\alpha, \beta, \Phi), g(\alpha, \beta, \Phi)$ such that, for any sequence of measures $\{P_n\}$ on H^n whenever $\overline{\lim}_n P_n(\frac{E(T_n)}{T_n} \geq \alpha) \geq \beta$, then $\overline{\lim}_n P_n(\sup_x |F_{T_n}(x) - F(x|\mathbf{e}^*)| \geq f(\alpha, \beta, \Phi)) \geq g(\alpha, \beta, \Phi)$.*

Proof of Claim. Clearly, the claim implies Proposition 2. We check the converse via contradiction. Assume Proposition 2 (actually, we will invoke the *proof* of Proposition 2 – see *remarks* after the proof of Proposition 2) and – via contradiction – find a sequence of pairs (f_i, g_i) decreasing to $(0, 0)$, (i.e. $f_{i+1} < f_i, g_{i+1} < g_i$) such that for each (f_i, g_i) there is a measure P_{n_i} on H^{n_i} (n_i is an increasing sequence) such that (i) $P_{n_i}(\frac{E(T_{n_i})}{T_{n_i}} \geq \alpha) \geq \beta$, and (ii) $P_{n_i}(\sup_x |F_{T_{n_i}}(x) - F(x|\mathbf{e}^*)| \geq f_i) < g_i$. But now consider the sequence of measures $\{P_{n_i}\}$ viewed as conditionals on H^{n_i} , the latter as a slice of the product space $\prod_i H^{n_i}$. Apply Proposition 2 to obtain a pair $f^*, g^* > 0$ such that $\overline{\lim}_i P_{n_i}(\sup_x |F_{T_{n_i}}(x) - F(x|\mathbf{e}^*)| \geq f^*) \geq g^*$. However, for all large i we have $f_i < f^*, g_i < g^*$. Put together this yields:

$$g_i < g^* \leq P_{n_i}(\sup_x |F_{T_{n_i}}(x) - F(x|\mathbf{e}^*)| \geq f^*) \leq P_{n_i}(\sup_x |F_{T_{n_i}}(x) - F(x|\mathbf{e}^*)| \geq f_i) < g_i$$

which is a contradiction. Hence, there could have been no such counterexample sequence (f_i, g_i, P_{n_i}) , implying the result of the claim. \square

Returning to the proof of the current Proposition, note that (for fixed (ϵ, ϵ')) the constants $f(\epsilon, \epsilon', \Phi), g(\epsilon, \epsilon', \Phi)$ do not change if we change the contract to any $\mathcal{C}_N \in \Phi$. In other words, it does not matter where in the sequence $f(n) = n^5$ we start our review contract. For a fixed pair (ϵ, ϵ') the same tail constants $(f(\epsilon, \epsilon', \Phi), g(\epsilon, \epsilon', \Phi))$ apply. Hence, choose N large enough and ϵ_2 small enough so that (note: for \mathcal{C}_N we have $\phi_{FB} \geq \prod_{n \geq N} (1 - \epsilon_n)$, where ϵ_n is the Chebyshev upper bound on a Type I error)

$$(*) \quad \phi_{FB} - \epsilon_2 > 1 - g(\epsilon, \epsilon', \Phi).$$

Consider the associated contract \mathcal{C}_N (with fee $t = \phi_{FB} - \epsilon_2$). I claim that there is an index $I(\epsilon, \epsilon')$ and a threshold $\delta_{\epsilon, \epsilon'}$ such that for all $i \geq I(\epsilon, \epsilon')$ and all SPNE $\rho \in \Sigma(\delta), \forall \delta \geq \delta_{\epsilon, \epsilon'}$ we have

$$\overline{P}_\rho(\frac{E(T_i)}{T_i} \geq \epsilon) < \epsilon'.$$

Else, find a counterexample sequence of equilibria ρ_{δ_i} , where δ_i is increasing to unity, and with associated (conditional) measures $P_{n_i} := \overline{P}_{\rho_{\delta_i}}$ on H^{n_i} (n_i increasing in i) such that

$$P_{n_i}(\frac{E(T_{n_i})}{T_{n_i}} \geq \epsilon) \geq \epsilon'.$$

By the preceding claim we must have,

$$\overline{\lim}_i P_{n_i}(\sup_x |F_{T_{n_i}}(x) - F(x|\mathbf{e}^*)| \geq f(\epsilon, \epsilon', \Phi)) \geq g(\epsilon, \epsilon', \Phi).$$

Note that, for all large i , this gives a bound on the probability that play enters phase $n_i + 1$, i.e. $\phi_{P_{\rho_{\delta_i}}}^{n_i+1} \leq 1 - g(\epsilon, \epsilon', \Phi)$. By passing to a convergent subsequence of $\{P_{\rho_{\delta_i}}\}$ if necessary, we know that the limit measure P has the property that $\phi_P \geq \phi_{FB} - \epsilon_2$. Assembling these implications, we obtain:

- $\phi_{P_{\delta_i}} \rightarrow \phi_P$
- $\phi_{P_{\delta_i}}^{n_i+1} \leq 1 - g(\epsilon, \epsilon', \Phi)$
- $\phi_{P_{\delta_i}} \leq \phi_{P_{\delta_i}}^{n_i+1}$ (as, for any fixed ρ , the sequence $\phi_{P_\rho}^n \downarrow \phi_{P_\rho}$).

The second and third points above yield, $\phi_{P_{\delta_i}} \leq 1 - g(\epsilon, \epsilon', \Phi)$. Now take limits on the LHS to obtain:

$$\phi_{FB} - \epsilon_2 \leq \phi_P = \lim_{\delta_i} \phi_{P_{\delta_i}} \leq 1 - g(\epsilon, \epsilon', \Phi).$$

This contradicts inequality (*). Hence, there is no such sequence of counterexamples, concluding the proof of the Proposition. \square

Proof of Corollary 1. Let K denote the maximal payoff from the stage game. The idea is to find, for each δ , a punishment length $L(\delta)$ satisfying three inequalities. The first inequality is

$$(I) : \frac{\delta^{L(\delta)} K}{1 - \delta} < \hat{\epsilon}$$

for some appropriately small $\hat{\epsilon}$ to be specified. For the second inequality, let $-K_1$ denote the maximal (expected) stage game loss for the principal (i.e. lowest expected output plus insurance payments). Choose $L(\delta)$ such that

$$(II) : \frac{\delta^{L(\delta)} K_1}{(1 - \delta)} < \hat{\epsilon}_1$$

for some $\hat{\epsilon}_1$ to be specified. For the third inequality on $L(\delta)$ let $\hat{\epsilon}_2$ denote the expected value of the bonus payment from reporting in a given period. We need $L(\delta)$ long enough so that

$$(III) : \frac{\delta^{L(\delta)} K}{(1 - \delta)} < \hat{\epsilon}_2$$

This ensures that under $\mathcal{C}(\delta)$ collusion stops once it becomes known that the KS test registers failure. Let ϵ be fixed as in the **Corollary** and let \mathcal{C}_ϵ denote the contract produced by Theorem 1, with participation fee $t_{\mathcal{C}_\epsilon}$. Now choose $\hat{\epsilon}$ and a participation fee, $t_{\mathcal{C}(\delta)}$, such that the following two inequalities are satisfied. First:

$$(IV) : t_{\mathcal{C}_\epsilon} \leq t_{\mathcal{C}(\delta)} - \hat{\epsilon}(1 - \delta)$$

The inequality says that $t_{\mathcal{C}(\delta)}$ is to be chosen so that the difference between the participation fees $t_{\mathcal{C}_\epsilon}$ and $t_{\mathcal{C}(\delta)}$, viewed forward from a period where the KS statistic has fallen into the rejection region, is at most the maximal (normalized discounted) continuation payoff. Also select $t_{\mathcal{C}(\delta)}$ to satisfy:

$$(V) : t_{\mathcal{C}_\epsilon} \leq t_{\mathcal{C}(\delta)} - \hat{\epsilon}_1(1 - \delta)$$

Note that the principal's expected loss, viewed forward from a period where the KS statistic has fallen into the rejection region, is at most $\hat{\epsilon}_1$. Hence, if we reduce the participation fee to $t_{\mathcal{C}_\epsilon}$, his expected (normalized discounted) payoff (viewed from such a period on) is lower than in the equilibrium with the contract $\mathcal{C}(\delta)$.

Consider the class of contracts, $\mathcal{C}(\delta)$, where the KS statistic and stage parameters (e.g. R_1, R_2 , etc.) are the same as in \mathcal{C}_ϵ . Punishment lengths equal $L(\delta)$ (chosen to satisfy the above inequalities I-III) and the per-period participation fee is some $t_{\mathcal{C}(\delta)}$ that satisfies the above inequalities (IV-V). I claim that this collection $\{\mathcal{C}(\delta)\}$ satisfies the Corollary. We check this by reducing the argument to the result of Theorem 1. Let us treat the upper bound on agents' payoffs. Let ρ_{δ_n} be the sequence of SPNE in $\Sigma(\delta_n)$ and for each ρ_{δ_n} consider the associated strategy, call it $\hat{\rho}_{\delta_n}$, in the game induced by contract \mathcal{C}_ϵ . The associated strategy profile $\hat{\rho}_{\delta_n}$ just mimics ρ_{δ_n} along histories along which the null hypothesis is never rejected. Along histories in which the null hypothesis is rejected, the $\hat{\rho}$ strategy mimics ρ until the conclusion of the current review phase. From that point on, there is nothing to mimic since the punishment length is infinite under the contract \mathcal{C}_ϵ .

Note that by choice of the participation fee $t_{\mathcal{C}(\delta)}$ and the lengths $L(\delta)$ we obtain that payoffs to agents under \mathcal{C}_ϵ are weaker highly under the profile $\hat{\rho}_{\delta_n}$ than under the contract $\mathcal{C}(\delta)$ when the profile ρ_{δ_n} is played. Importantly, the participation constraint holds for the game induced by the contract \mathcal{C}_ϵ (under the profile $\hat{\rho}_{\delta_n}$). Also note that the profile $\hat{\rho}_{\delta_n}$ inherits the property that collusion stops once the null hypothesis is rejected (moreover, it is an SPNE if ρ_{δ_n} is pure – although we do not need this), hence the arguments of Theorem 1 apply to this modified profile, as the only place where we use the hypothesis of equilibrium is to claim (i) participation holds and (ii) that collusion stops once the null is failed. Finally, observe that the probability measures P_{ρ_δ} live on the product space $\prod_i H^i$. Adjust continuation probabilities as follows: include \emptyset in the space of continuation histories and whenever punishment is incurred place all mass of the continuation probability on \emptyset . Now extend to H^* – arguing as in Observation 1 – to obtain a measure $P_{\hat{\rho}_\delta}$ on the histories induced by \mathcal{C}_ϵ . We apply the result of Theorem 1 to the sequence $(\hat{\rho}_{\delta_n}, P_{\hat{\rho}_{\delta_n}})$. Payoffs to agents from this sequence approach the first-best payoff. Since payoffs under $\hat{\rho}$ are higher than under ρ , the upper bound on agents' payoffs follows.

Now for the principal's payoff. Let $\rho_{\delta_n}, \hat{\rho}_{\delta_n}$ be as above, the latter with associated measures $P_{\hat{\rho}_{\delta_n}}$. Note that by choice of the participation fees (and since the participation constraint is satisfied under $\hat{\rho}_{\delta_n}$) and punishment lengths $L(\delta)$, the principal's payoff in the contract \mathcal{C}_ϵ under the profile $\hat{\rho}_{\delta_n}$ is lower than in the game induced by contract $\mathcal{C}(\delta)$.¹⁶ Now apply the result of Theorem 1 to the pair $(\hat{\rho}_{\delta_n}, P_{\hat{\rho}_{\delta_n}})$. Since participation holds under this sequence of profiles, the principal's payoff approaches the first-best benchmark. Since payoffs under the contracts $\mathcal{C}(\delta)$ (under the equilibria ρ_{δ_n}) are higher, the lower bound follows. \square

6.4 Omitted Details from Proof of Theorem 1

We now turn to details regarding passage from discounted sums to time average sums, which was one of the steps we omitted from the proof of Theorem 1 given in the text. Define the following two payoff sequences \bar{u}_t (resp. \bar{u}_t^δ): (when $k = 1$ the sum below is, by fiat, zero)

$$\bar{u}_t = \begin{cases} 1, & \text{if } t \in (\sum_{i=1}^{k-1} T_i, \sum_{i=1}^k T_i], k < I(\epsilon, \epsilon') \\ 1, & \text{if } t \in (\sum_{i=1}^{k-1} T_i, \sum_{i=1}^{k-1} T_i + \epsilon \cdot T_k], k \geq I(\epsilon, \epsilon') \\ 0 & \text{if } t \in (\sum_{i=1}^{k-1} T_i + \epsilon \cdot T_k, \sum_{i=1}^k T_i], k \geq I(\epsilon, \epsilon'). \end{cases}$$

For \bar{u}_t^δ just multiply each case by $(1 - \delta)\delta^t$. This is the ‘‘optimistic stream’’ of expected payoffs mentioned in the text. Let $E_{P_{\rho_\delta}}^i u_t^\delta$ denote the sum of the expected values of payoffs summed over times during the i -th work phase. We have:

$$\int_{H^*} u^\delta(h) dP_{\rho_\delta} = \sum_{t=1}^{\infty} E_{P_{\rho_\delta}} u_t^\delta = \sum_i E_{P_{\rho_\delta}}^i u_t^\delta.$$

The above equality involves an interchange of integral with a summation, which can be justified as follows: truncate the normalized discounted payoff at some large T . The interchange of sum and integral is obviously justified for the truncation. Now take arbitrarily large truncations to obtain equality. Now we use the reasoning from equation (1) to bound the terms $E_{P_{\rho_\delta}}^i u_t^\delta$ for $i \geq I(\epsilon, \epsilon')$, (let $\Theta(i)$ denote the times t that cover review phase i)

$$E_{P_{\rho_\delta}}^i u_t^\delta \leq \phi^i (1 - \epsilon') \sum_{t \in \Theta(i)} \bar{u}_t^\delta + \phi^i \epsilon' \sum_{t \in \Theta(i)} 1_t^\delta$$

where ϕ^i denotes the probability of reaching this phase (we suppress the dependence on the underlying measure P_{ρ_δ} as we will be using a trivial bound on ϕ^i that applies

¹⁶Note that the inequalities defining $L(\delta), t_{\mathcal{C}(\delta)}$ imply that the reduction in the participation fee (which itself can be made arbitrarily small) outweighs the maximal potential (normalized, discounted) loss to the principal from restarting the review phases after a punishment of length $L(\delta)$.

regardless of the measure), and finally put $1_t^\delta := (1-\delta)\delta^t \cdot 1$, where 1 is the normalized stage payoff from collusion. Summing over phases i we obtain,

$$\begin{aligned} \int_{H^*} u^\delta(h) dP_{\rho_\delta} &= \sum_t E_{P_{\rho_\delta}} u_t^\delta = \sum_i E_{P_{\rho_\delta}}^i u_t^\delta \\ &\leq \sum_{i=1}^{I(\epsilon, \epsilon')} \sum_{t \in \Theta(i)} 1_t^\delta + \sum_{i \geq I(\epsilon, \epsilon') + 1} \phi^i [(1 - \epsilon') \sum_{t \in \Theta(i)} \bar{u}_t^\delta + \epsilon' \sum_t 1_t^\delta]. \end{aligned}$$

Use the trivial bound $\phi^i \leq 1$ to obtain

$$\int_{H^*} u^\delta(h) dP_{\rho_\delta} \leq \epsilon^1 + (1 - \epsilon') \cdot \sum_{i \geq I(\epsilon, \epsilon') + 1} \sum_{t \in \Theta(i)} \bar{u}_t^\delta + \epsilon'$$

where ϵ^1 is the bound on the sum over phases up until $I(\epsilon, \epsilon')$. By choosing δ large we can make ϵ^1 as small as we like (since $I(\epsilon, \epsilon')$ is *independent* of δ for all large δ). Similarly, by choosing ϵ, ϵ' to be small at the outset we can make the final term above as small as we like.¹⁷ Hence, to show that the limiting value as δ tends to unity is close to 0 it suffices to check that

$$\sum_i \sum_{t \in \Theta(i)} \bar{u}_t^\delta = \sum_t \bar{u}_t^\delta$$

can be made (by making ϵ, ϵ' small) arbitrarily close to 0 as δ gets large.

Observation 2. Let $u_t^\delta(h)$ denote the (normalized) discounted time t payoff for the agent. Then, we have (up to the ϵ^1, ϵ' terms): $\lim \int_{H^*} u_t^\delta(h) dP_\delta \leq \lim_T \frac{1}{T} \sum_{t=1}^T \bar{u}_t$.

Proof. We have already verified that (up to the ϵ', ϵ^1 terms, which we ignore here) $\lim \int_{H^*} u_t^\delta(h) dP_{\rho_\delta} \leq \sum_t \bar{u}_t^\delta$. Hence, it suffices to check that $\sum_t \bar{u}_t^\delta = \lim_T \frac{1}{T} \sum_{t=1}^T \bar{u}_t$. In the text we computed the limit of this latter term along a subsequence of times of the form $T = \sum_i T_i$ and found that it could be pushed arbitrarily close to 0. We now check that taking any other sequence of times we obtain the same limit. By Abel's Theorem (Radner (1985), pg. 1175) this proves that $\lim_{\delta \uparrow 1} \sum_t \bar{u}_t^\delta = \lim_T \frac{1}{T} \sum_{t=1}^T \bar{u}_t$. To prove the latter limit exists, take any T and put $T(k) = \sum_{i=1}^k T_i$. Find k such that $T(k) < T \leq T(k+1)$. Notice that we have

$$\frac{1}{T} \sum_{t=1}^T \bar{u}_t = \overbrace{\frac{1}{T} \sum_{t=1}^{T(k)} \bar{u}_t}^{\text{I}} + \underbrace{\frac{1}{T} \sum_{t=T(k)+1}^T \bar{u}_t}_{\text{II}}.$$

¹⁷This may require a larger $I(\epsilon, \epsilon')$, but once this is done we then take limits on δ – hence making the first and third terms as small as we like.

We check that, by the selection of sample sizes T_k , term II tends to 0 as T is large. Recall that we selected $T_k := k^5$. Notice that

$$\frac{1}{T} \sum_{t=T(k)+1}^T \bar{w}_t \leq \frac{T_{k+1}}{T(k) + 1}.$$

Now use the fact that the function $n \mapsto f(n) = n^5$ is convex and increasing. Using right-endpoint Riemann sums (with rectangles of partition length 1 and of height $f(1), f(2), \dots, f(k)$) we obtain

$$T(k) + 1 = \sum_{i=1}^k T_i + 1 \geq \int_0^k x^5 dx + 1 = \frac{1}{6}k^6 + 1.$$

Since $T_{k+1} = (k+1)^5$ we clearly obtain

$$\frac{T_{k+1}}{T(k) + 1} \rightarrow 0$$

as $k \rightarrow \infty$. Similar reasoning applies to show that

$$\frac{T(k)}{T} \rightarrow 1$$

as $k \rightarrow \infty$. Hence, term I determines the limit and, itself, limits to the average along the subsequence $T(k)$. Since we verified in the text that this term can be made as small as we like the claim follows. \square

Now for the principal's payoff. Let $w_t^\delta(h)$ denote the (normalized discounted) principal's payoff at time t along history h . We have:

$$\int_{H^*} w^\delta(h) dP_{\rho_\delta} = \sum_t E_{P_{\rho_\delta}} w_t^\delta = \sum_i \sum_{t \in \Theta(i)} E_{P_{\rho_\delta}} w_t^\delta \quad (9)$$

where $E_P w_t^\delta$ is the expected time- t discounted payoff (we are interchanging sum and integral as before). The term $\sum_{t \in \Theta(i)} E_P w_t^\delta$ denotes the discounted expected value of output summed over the i -th work phase. To find a lower bound on this sum, we find a stream of (expected) per-period payoffs that (i) yields a lower bound, (ii) possesses a time average, and (iii) has limiting time-average arbitrarily close to the first-best principal's payoff. Let w^* denote the (aggregate) insurance wage and let EX^*, EX_* (resp.) denote the expected value of output conditional on the first-best effort choice and the lowest collective effort choice. Define a stream of (expected) payoffs, call it \bar{w}_t^δ (abusing notation), as follows:

$$\bar{w}_t = \begin{cases} EX_* - w^* & \text{if } t \in (\sum_{i=1}^{k-1} T_i, \sum_{i=1}^k T_i], k < I(\epsilon, \epsilon') \\ EX_* - w^* & \text{if } (\sum_{i=1}^{k-1} T_i, \sum_{i=1}^{k-1} T_i + \epsilon T_k], k \geq I(\epsilon, \epsilon') \\ EX^* - w^* & \text{if } (\sum_{i=1}^{k-1} T_i + \epsilon T_k, \sum_{i=1}^k T_i], k \geq I(\epsilon, \epsilon'). \end{cases}$$

Ignore integer issues in the above definition (and in the previous). To obtain \bar{w}_t^δ pre-multiply by $(1 - \delta)\delta^t$ – this is the “pessimistic stream” of (expected) payoffs. Put $\underline{w}_t := EX_* - w^*$ in every period t and similarly define \underline{w}_t^δ . Let X_t denote output in period t and X_t^δ denote the normalized discounted quantity. Notice that

$$E_{\bar{P}}(X_1 + X_2 + \dots + X_T) = \sum_{(\mathbf{e}^1, \dots, \mathbf{e}^T)} [E(X(\mathbf{e}^1) + EX(\mathbf{e}^2) + \dots + EX(\mathbf{e}^T))] \cdot \bar{P}(\mathbf{e}^1, \dots, \mathbf{e}^T).$$

Replacing with normalized discounted r.v.’s we obtain:

$$(\star) \sum_{t \in \Theta(k)} E_P X_t^\delta = \sum_{(\mathbf{e}^1, \dots, \mathbf{e}^{T_k})} \phi_P^k [EX^\delta(\mathbf{e}^1) + EX^\delta(\mathbf{e}^2) + \dots + EX^\delta(\mathbf{e}^{T_k})] \cdot \bar{P}(\mathbf{e}^1, \mathbf{e}^2, \dots, \mathbf{e}^{T_k}).$$

Using Proposition 5, for $k \geq I(\epsilon, \epsilon')$ we have

$$\bar{P}(E(T_k)/T_k \geq \epsilon) < \epsilon'.$$

We apply the pessimistic streams to (\star) to bound (from below) the terms $[EX^\delta(\mathbf{e}^1) + EX^\delta(\mathbf{e}^2) + \dots + EX^\delta(\mathbf{e}^{T_k})]$ on the set $\{E(T_k)/T_k \geq \epsilon\}$. Using this bound and the fact that, $E_P w_t^\delta \geq E_P(X_t^\delta - (1 - \delta)\delta^t w^*)$, we obtain:

$$\begin{aligned} \sum_{t \in \Theta(k)} E_{P_{\rho_\delta}} w_t^\delta &\geq \sum_{t \in \Theta(k)} E_{P_{\rho_\delta}} (X_t^\delta - (1 - \delta)\delta^t w^*) \geq \phi_{P_{\rho_\delta}}^k [(1 - \epsilon') \sum_{t \in \Theta(k)} \bar{w}_t^\delta + \epsilon' \sum_{t \in \Theta(k)} \underline{w}_t^\delta] \\ &\quad - (1 - \phi_{P_{\rho_\delta}}^k) \sum_{t \in \Theta(k)} w^* 1_t^\delta. \end{aligned}$$

Now sum over all phases k to obtain, using equation (9): (ignore the contribution $k < I(\epsilon, \epsilon')$ which vanishes for large δ)

$$\int_{H^*} w^\delta(h) dP_{\rho_\delta} \geq \sum_k \phi_{P_{\rho_\delta}}^k [(1 - \epsilon') \sum_{t \in \Theta(k)} \bar{w}_t^\delta + \epsilon' \sum_{t \in \Theta(k)} \underline{w}_t^\delta] - (1 - \phi_{P_{\rho_\delta}}) w^*. \quad (10)$$

The forthcoming observation verifies that this sum has a long-run time average, and that it can be made (by choice of ϵ, ϵ') arbitrarily close to the first-best output value. The argument is nearly identical to the proof for Observation 2, hence is omitted. As before, the point is that the sequence of sample sizes that defines the reviews is (i) increasing, (ii) convex (i.e. $f(x)$ is convex), and (iii) slowly varying.¹⁸

Observation 3. Put $\bar{w}^\delta := (1 - \epsilon') \sum_t \bar{w}_t^\delta + \epsilon' \sum_t \underline{w}_t^\delta$ and let $\bar{w}^T := \frac{1}{T} \cdot \sum_{t=1}^T \tilde{w}_t$, where $\tilde{w}_t := (1 - \epsilon') \bar{w}_t + \epsilon' \underline{w}_t$ and \bar{w}_t (resp. \underline{w}_t) is the undiscounted time t payoff corresponding to \bar{w}_t^δ (resp. \underline{w}_t^δ). Then,

$$(\blacklozenge) \lim_{\delta \uparrow 1} \bar{w}^\delta = \lim_{T \rightarrow \infty} \bar{w}^T = (1 - \epsilon')[(1 - \epsilon)\Pi^{FB} + \epsilon\Pi] + \epsilon'\Pi$$

¹⁸Say that $f(x)$ is *slowly varying* if $\forall t, f(x)/f(x+t) \rightarrow 1$ as $x \rightarrow \infty$.

Clearly, the term on the RHS of (\diamond) approaches Π^{FB} as we make ϵ, ϵ' small. To conclude notice that (i) the terms ϕ_P^k are decreasing for any measure P and (ii) $\phi_{P_\delta}^k \rightarrow \phi_P^k$. Since ϕ_P is bounded below by a quantity that can be made close to unity (by selection of the participation fees), we obtain that there is some $\hat{\epsilon}$ (chosen to be as small as we like at the contract design stage) such that $\phi_P^k \geq (1 - \hat{\epsilon}), \forall P, \forall k$. Hence, we obtain from equation (10):

$$\lim_{\delta \uparrow 1} \int_{H^*} w^\delta(h) dP_{\rho_\delta} \geq (1 - \hat{\epsilon}) \cdot \lim_{\delta \uparrow 1} \bar{w}^\delta - \hat{\epsilon} w^*.$$

Using Observation 3, we obtain a lower bound on the principal's payoff that is approximately first-best.

References

- ABDULKADIROGLU, A. AND K.-S. CHUNG (2003): “Auction Design with Tacit Collusion,” *Mimeo*.
- ABREU, D., D. PEARCE, AND E. STACCHETTI (1990): “Toward a Theory of Discounted Repeated Games with Imperfect Monitoring,” *Econometrica*, 58, 1041–1063.
- BONATTI, A. AND J. HÖRNER (2009): “Collaborating,” *American Economic Review*, 101, 632–663.
- BRUSCO, S. (1997): “Implementing Action Profiles when Agents Collude,” *Journal of Economic Theory*, 73, 395–424.
- CHE, Y.-K. AND S.-Y. YOO (2001): “Optimal Incentives for Teams,” *American Economic Review*, 91, 525–541.
- CREMER, J. AND R. MCLEAN (1988): “Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions,” *Econometrica*, 56, 1247–1257.
- FUDENBERG, D., D. K. LEVINE, AND E. MASKIN (1994): “The Folk Theorem with Imperfect Public Information,” *Econometrica*, 56, 997–1039.
- HOLMSTROM, B. (1982): “Moral Hazard in Teams,” *Bell Journal of Economics*, 13, 324–340.
- HÖRNER, J. AND W. OLSZEWSKI (2006): “The Folk Theorem for Games with Private Almost-Perfect Monitoring,” *Econometrica*, 74, 1499–1544.
- ISHIGURO AND H. ITOH (2001): “Moral Hazard and Renegotiation with Multiple Agents,” *The Review of Economic Studies*, 68, 1–20.
- ITOH, H. (1991): “Incentives to Help in Multi-Agent Situations,” *Econometrica*, 59, 611–636.
- KALLENBERG, O. (2002): *Foundations of Modern Probability*, Springer-Verlag.
- MA, C. (1988): “Unique Implementation of Incentive Contracts with Many Agents,” *The Review of Economic Studies*, 55, 555–571.
- MA, C., J. MOORE, AND S. TURNBULL (1988): “Stopping Agents from Cheating,” *Journal of Economic Theory*, 46, 355–372.
- MAILATH, G. AND L. SAMUELSON (2006): *Repeated Games and Reputations: Long-Run Relationships*, Oxford University Press.

- MATSUSHIMA, H. (2004): “Repeated Games with Private Monitoring,” *Econometrica*, 72, 823–852.
- MILLER, N. H. (1997): “Efficiency in Partnerships with Joint Monitoring,” *Journal of Economic Theory*, 77, 285–299.
- MOOKHERJEE, D. (1984): “Optimal Incentive Schemes with Many Agents,” *The Review of Economic Studies*, 51, 433–446.
- RADNER, R. (1981): “Monitoring Cooperative Agreements in a Repeated Principal-Agent Relationship,” *Econometrica*, 49, 1127–1148.
- (1985): “Repeated Principal-Agent Games with Discounting,” *Econometrica*, 53, 1173–1198.
- ROGERSON, W. P. (1985): “Repeated Moral Hazard,” *Econometrica*, 53, 69–76.
- RUBINSTEIN, A. AND M. YAARI (1983): “Repeated Insurance Contracts and Moral Hazard,” *Journal of Economic Theory*, 30, 74–97.
- SANNIKOV, Y. (2007): “Games with Imperfectly Observable Actions in Continuous Time,” *Econometrica*, 75, 1285–1329.
- (2008): “A Continuous-Time Version of the Principal-Agent Problem,” *The Review of Economic Studies*, 75, 957–984.
- SPEAR, S. AND S. SRIVASTAVA (1987): “On Repeated Moral Hazard with Discounting,” *The Review of Economic Studies*, 54, 599–617.