Long Run Optimal Contracts under Adverse Selection with Limited Commitment

Subir K. Chakrabarti and Jaesoo Kim*

June 15, 2012

Department of Economics
Indiana University Purdue University Indianapolis(IUPUI)
425 University Blvd.
Indianapolis, IN 46202

Abstract: The paper studies long run optimal contracts under adverse selection with limited commitment so that the contracts are open to negotiation in every period. Thus the contracting game is repeated over multiple periods and belief about the type of the agent is updated by the principal. We study both the finite-horizon case as well as the infinite-horizon case and find that the unique perfect Bayesian equilibrium of the finite-horizon repeated contracting game is one in which the second-best optimal contract is offered in period 1, there is full revelation of the type of the agent, and from period 2 onwards the first-best contract is offered by the principal. If the agent is the least efficient type then the agent gets no informational rent but if the type of the agent is among the more efficient types, the agent receives an informational rent that has to be paid in period 1. By contrast the infinite-horizon game has multiple Perfect Bayesian Equilibrium points and the one that is optimal for the principal is the equilibrium in which the principal offers the type-separating second-best optimal contract in period 1, fully updates beliefs about the type of the agent and continues to offer the second-best contract from period 2 onwards.

Keywords: Adverse Selection, Optimal Contracts, Limited Commitment, Pooling Contracts, Separating Contracts, Perfect Bayesian Equilibrium, Relational Contract.

JEL Classification Numbers: Primary D2, D8, L1. Secondary L5

---

*We gratefully acknowledge the discussions at the economic theory workshop of the department of economics, IUPUI and thank Will Geller, Steve Russell and Jyoti Sarkar for their comments.

1 Introduction

We study long run contracts when there is limited commitment in an adverse selection model. While enforceability of the contracts is a fairly reasonable assumption in the short run and especially when the contracts are single-period contracts, it is less likely that individuals may sign contracts that are fully enforceable in the long run, and even if they do guaranteeing the enforceability of such contracts may be problematic. What might typically happen when the interaction between a principal and an agent is repeated over many periods, is that the contracts may be rewritten every period, with the contract being enforceable for that period. This of course does not preclude the possibility that the same contract may be offered over multiple periods, but the fact that the individuals are not committed to a single fixed contract over periods allows possibly a greater degree of flexibility.

In order to study the nature of long run contracts with limited commitment we study a repeated game in which a principal and an agent are able to rewrite the contract every period if they choose to do so and only single-period contracts are enforceable. We thus use the framework of repeated games and study a repeated game in which the principal offers an enforceable contract to the agent each period, but the principal is free to offer a different contract in the following period. We find this approach both interesting and useful because it opens up the possibility of renegotiation, as well as allow the principal to offer contracts that are based on the updated beliefs of the principal. Such an optimal contract problem can then be analyzed as a repeated game with incomplete information in which the payoffs of the players are private information. In the adverse selection models the cost parameter of the agent is not known to the principal and thus the principal does not know the payoff of the agent. Such repeated games have been studied in the literature, especially in the long run pricing and output strategies of firms in oligopoly markets in which the costs of the firms are not known, as for example in [4], [2] and [3]. We find that many of the observations made in these studies are useful in understanding long run optimal contracts in adverse selection models. More specifically we look for contracts that can be implemented as equilibrium points of the repeated game with incomplete information and focus on the perfect Bayesian equilibrium points of the resulting repeated game with imperfect information. Although this paper focuses on Adverse Selection it is useful to note that optimal contracts involving Moral hazard has been studied using repeated games to model repeated interaction between the Principal and the Agent. For example [13] and [14] studies the the nature of optimal contracts when there is repeated interaction between a Principal and an Agent and the agent's effort levels cannot be observed. In [16] the discounts offered by insurers who have a favorable record of past claims is explained as a mechanism that counteracts the inefficiency from Moral hazard, a mechanism that can be used when there is repeated interaction between the principal and the agent.

We find some interesting results. In a finite-horizon repeated game in which both the principal and the agent knows that the relationship between the principal and the agent will end definitely after a given number of periods, the unique perfect Bayesian equilibrium of the game is one in which the optimal second-best contract of the single-period contract problem is offered in period 1. This fully reveals the type of the agent (inferred by the principal from the output produced by the agent) so that the principal updates beliefs about the agent, and then offers the full-information first-best contract that is consistent with the agent's type to the agent. This, however, means that the principal has to compensate the agent with an informational rent in period 1, so that the agent is willing to reveal his type and forgo the informational rent in the future.

For the infinite-horizon situation we find that there are multiple perfect Bayesian equilibrium points. There is an entire set of perfect Bayesian equilibria that are pooling equilibrium points in which the principal offers a stationary contract in every period and the belief of the principal about the type of the agent is never updated. We also find that there is a perfect Bayesian equilibrium that is a separating equilibrium in which the second-best optimal contract is offered in period 1, the principal's beliefs are fully updated and the contracts from period 2 onwards is the second-best contract for the type of the agent revealed in period 1. In this separating contract, if the agent is of the more efficient type then the agent gets paid his informational rent every period and thus has no incentive to hide his true type. This particular separating equilibrium is optimal for the principal as the principal's expected payoff over the entire horizon is a the maximum over all the possible perfect Bayesian equilibrium points. This may at first seem a little counter-intuitive as this does not involve the full-information, first-best contract in any way even though beliefs are fully updated. However, it is useful to note that an information rent has to be paid to the agent in period 1 if the full-information, first-best contract is to be implemented. This information rent is higher than the informational rent of the agent in the second-best optimal contract.

Repeated interaction between a Principal and an Agent is fairly common and much of the negotiation on payments tend to be implicit and have limited commitment over multiple periods. However, although these agreements may not be fully enforceable by long-rum contracts, these agreements need to be self-enforcing in some manner. This has been recognized and has been studied in the literature. For example, [15], [12], [11] among others have studied the nature of such self-enforcing or equilibrium contracts with limited commitment in repeated games with imperfect monitoring, where the effort level of the agent is unobserved. In some cases the results also extend to the case of adverse selection where the cost parameter of the agent is private information of the agent as in [11]. Studies on adverse selection in the repeated game framework has also been extensive, [5] and [10] study adverse selection with limited commitment in the case of repeated interaction in which the parties can renegotiate a long-term contract. These study long-term contracts in which future renegotiation can be added as a constraint in the initial contract. The result obtained in the case with renegotiation differs significantly from the full-commitment case. Most notably, the renegotiation-proof contract is suboptimal compared to the full-commitment case.

In our study we do not impose any conditions on renegotiation but allow for the possibility that the principal may want to offer different contracts in the future. We examine instead entire classes of self-enforcing contracts in a repeated game framework and examine the properties of these contracts. As we have already mentioned we find that in the finite-horizon game, in which there is a definite terminal point to the relationship between the principal and the agent, there is only kind of contract that is self-enforcing. The principal in this case chooses to learn the type of the agent quickly by paying all the informational rent in period 1 and then offers only the complete information, first-best contract in the following periods.

In the infinite-horizon case we find that there are plenty of self-enforcing series of contracts that can be implemented as either pooling[1] or separating equilibrium. However, the series of contracts that the principal would be most likely to offer is a separating contract as it maximizes the expected stream of profits of the principal. This is an equilibrium in which the single-period second-best optimal contract is offered. It is of interest to note that in the case of the pooling equilibria the agent gains nothing from using his private information. In this respect the pooling equilibria are similar to the pooling equilibria in [11], which use public perfect equilibrium strategies so that the strategies of the agent and the principal are functions of only the publicly observed output levels, and the principal never updates beliefs about the type of the agent. The principal and the agent negotiate only on the available public information about the output level produced. In a separating equilibrium, the principal offers different options to the different types, learns from the output level produced by the agent, who produces

[1]We note here that in [11] the contracts in the stationary equilibria are all pooling equilibrium contracts as the contract offered by the principal does not depend on the type of the agent. This is also true of the contracts offered in our pooling equilibria.

an output based on his privately observed cost parameter, and then uses this information to make subsequent offers.

The paper is organized as follows. In section 2 we describe the details of the model. In section 3 we describe the infinite-horizon repeated game with incomplete information.In section 5 we provide a folk theorem for the pooling equilibrium points. In section 6 we discuss the separating equilibrium for both the finite-horizon and the infinite-horizon game. In section 8 we conclude.

2 The Adverse Selection Model

A principal needs to contract work out to an agent in which the work needs to be done over many periods. The principal needs to write a contract with the agent in each period although the relationship with the agent can last for many periods. This is typically the situation in many cases where a workers wage or bonuses are determined in each period during which the worker works for the principal. The total revenue of the principal from the output produced by the agent is $S(q)$ where $S(.)$ is an increasing and strictly concave function of the output q produced by the agent.

The agent can produce the output q at cost θq. The value of θ is private information to the agent and the principal only knows that θ can take finitely many values $\theta_1, \cdots, \theta_L$ with probabilities $\nu_1, \cdots, \nu_L$ with $\theta_1 < \theta_2 < \cdots < \theta_L$. We will denote by $\Theta = \{\theta_1, \cdots, \theta_L\}$ the set of possible values of θ and sometimes refer to Θ as the type set of the agent, and the probability distribution giving the belief of the type of the agent we will denote by ν. The principal's payoff is given by

$$U_P(q, T) = S(q) - T$$

where T is the amount paid by the principal to the agent in return for output q. The payoff of the agent which depends on the agent's type θ_ℓ is given by

$$U_\ell(q, T) = T - \theta_\ell q$$

where θ_ℓ is the true marginal cost of the agent. As the principal does not know the value of θ, the actual payoff of the agent is not known to the principal.

This then gives us the single -period contracting game in which the principal makes an offer T for an output q and the agent then either accepts the offer or rejects it.

3 The infinite horizon game

The infinite horizon repeated adverse selection problem is one that is generated by allowing for recontracting every period over an infinite horizon. The strategy of the principal in this sequential game is a sequence $\{\sigma_t^P\}_{t=1}^{\infty}$ such that

$$\sigma_t^P : H_{t-1} \rightarrow \mathbb{R}_+^L \times \mathbb{R}_+^L$$

where H_{t-1} is the set of histories of the game until period $t-1$ and an $h_{t-1} \in H_{t-1}$ is given by $h_{t-1} = \{(T_1, q_1), (T_2, q_2), \cdots, (T_{t-1}, q_{t-1})\}$, where q_t is the output in time period t and T_t is the payment made in period t. That is, h_{t-1} is a history that consists of a sequence of past payments and output levels until period $t-1$. The menu of choices offered by the principal in period t thus depends on the past history of payments and output levels so that given a history h_{t-1} up to time period t, the principal chooses a menu $\sigma_t^P(h_{t-1}) = \{(T_t(\theta_\ell), q_t(\theta_\ell)\}_{\ell=1}^L$ if the principal's strategy in period t is σ_t^P. A strategy of the principal will be denoted by $\sigma^P = \{\sigma_t^P\}_{t=1}^{\infty}$. The agent's strategy in any period t also depends on the past history but also on the type of the agent given by the value of θ_ℓ or the unit cost of production of the agent. Therefore, the strategy of the agent is a sequence $\{\sigma_t^{A\ell}\}_{t=1}^{\infty}$ such that

$$\sigma_t^A : H_{t-1} \times \Theta \rightarrow \mathbb{R}_+.$$

The expected payoff of the principal in period t is

$$\sum_{\ell=1}^{L} \nu_\ell (S(q_t(\theta_\ell) - T_t(\theta_\ell))$$

as the actual payoff of the principal depends on the option in the contract chosen by the agent from the menu offered by the principal. The payoff of the agent in any period t is given by

$$U_\ell(q_t, T_t) = T_t - \theta_\ell q_t.$$

The expected payoff of the principal over the entire infinite horizon is the expected discounted sum of the single-period payoffs from the sequence of offers of the principal and the offers chosen by the agent and is given by

$$\sum_{\ell=1}^{\infty} \nu_\ell [\sum_{t=1}^{\infty} \delta_P^{t-1} (S(q_t(\theta_\ell)) - T_t(\theta_\ell))],$$

where δ_P is the discount rate of the principal. Therefore, the expected payoff of the principal when the principal's strategy is σ^P and the agent's strategy is $\sigma^{A\theta}$ is

$$U_P^{\infty}(\sigma^P, \sigma^{A\ell}) = \sum_{\ell=1}^{\infty} \nu_\ell [\sum_{t=1}^{\infty} \delta_P^{t-1} U_P(\sigma_t^P(h_{t-1}), \sigma_t^{A\ell}(h_{t-1}))].$$

Similarly, the payoff of the agent over the entire infinite horizon is the discounted sum of the single-period payoffs. These single-period payoffs depend on the offers made each period by the principal and the type of the agent. Thus the payoff of agent of type θ_ℓ over the entire infinite horizon is

$$\sum_{t=1}^{\infty} \delta_A^{t-1}(T_t - \theta_\ell q_t)$$

so that the payoff of the agent over the infinite horizon, when the strategy of the principal and the agent is (σ^P, σ^A), is given by

$$U_\ell^\infty(\sigma^P, \sigma^{A_\ell}) = \sum_{t=1}^{\infty} \delta_{A_\ell}^{t-1} U_\ell(\sigma_t^P(h_{t-1}), \sigma_t^{A\ell}(h_{t-1}))$$

where $\delta_{A\ell}$ is the discount rate of the agent.

4 Long Run Optimal Contracts

In looking for optimal contracts that can be implemented in the long run, that is over the infinite horizon, we look for an optimal contract among the set of equilibrium contracts. While optimal contracts are usually derived by finding the contract that maximizes the principal's payoff subject to the participation and incentive constraints, in the case of long run contracts with limited commitment where single-period contracts are offered in each period, any long run contract should typically be an equilibrium contract in the sense that the neither the principal nor the agent has any incentive to take an action or make an offer that is different from what is proposed.

In the infinite horizon game with incomplete information the equilibrium concept that we use here is that of a Perfect Bayesian equilibrium[2]. A **Perfect Bayesian equilibrium** is a strategy combination that continues to be an optimal strategy for every player given any history and the updated beliefs of the players given that history, when the beliefs are updated using Bayes' rule.

We note that any strategy combination (σ^P, σ^A) generates histories h_t and thus generates a probability over the set of possible histories H_t up to time period t. Thus given a strategy combination, observing a history h_t the principal is able to update beliefs about the type of the agent using Bayes rule, the updated beliefs about the type then is given by the conditional probability distribution over the set Θ which we will denote as $\nu|(\sigma, h_t)$.

[2]Note that as this is a game with incomplete information it can also be viewed as a game with imperfect information, in which a chance move at the beginning of the game cannot be perfectly observed by all the players.

Definition 1 *Given the strategy combination* $\sigma^\star = (\sigma^{P\star}, \sigma^{A\ell\star})$, *the assessment* $(\sigma^\star, \nu^\star)$ *is a* **Perfect Bayesian** *equilibrium of the infinite horizon game if*
(i) $\nu^\star(.)$ *is a system of beliefs that is determined by* $\sigma^\star$ *according to the Bayesian updating rule, and*
(ii)for every time period t *and for every history* h_t *up to time period* t, *the expected payoff of the principal satisfies*

$$\sum_{\ell=1}^{\infty}(\nu_\ell|_{h_t,\sigma^\star})U_P^\infty(\sigma^{P\star}, \sigma^{A\ell\star}|_{h_t}). \geq \sum_{\ell=1}^{\infty}(\nu_\ell|_{h_t,\sigma^P,\sigma^{A\ell\star}})U_P^\infty(\sigma^P, \sigma^{A\ell\star}|_{h_t})$$

for every $\sigma^P|_{h_t}$
and the expected payoff of the agent of each type ℓ *satisfies*

$$\sum_{t=1}^{\infty}\delta_A^{t-1}U_\ell^\infty((\sigma^{P\star}, \sigma^{A\ell\star})|_{h_t}) \geq \sum_{t=1}^{\infty}\delta_A^{t-1}U_\ell^\infty((\sigma^{P\star}, \sigma^{A\ell})|_{h_t})$$

for every $\sigma^{A\ell}|_{h_t}$.

Note that the strategies are conditioned on the private information of the agent as well as the history of outputs and payments. For a detailed discussion of Perfect Bayesian equilibrium and Sequential Equilibrium one may refer to [7] and for a discussion of Sequential equilibrium see [9]. A Perfect Bayesian equilibrium will be called a **pooling equilibrium** if the strategies of the agent of different types are the same. That is, in a pooling equilibrium the agent plays the same strategy irrespective of its type. A Perfect Bayesian equilibrium will be called a **separating equilibrium** if the strategy of an agent depends on its type. It will be called a **strictly separating equilibrium** if the equilibrium strategy of the agent varies strictly with its type.

5 Pooling Equilibrium and a Folk Theorem

Here we show that there are perfect Bayesian equilibrium of the repeated game in which the equilibrium contracts are pooling contracts.

Theorem 1 A Folk Theorem for Pooling Contracts *Let* $(\hat{q}, \hat{T})$ *be any pooling contract such that* $S(\hat{q}) - \hat{T} > 0$ *and* $\hat{T} - \theta_L\hat{q} > 0$. *Then there is a perfect Bayesian equilibrium strategy* $\hat{\sigma}$ *such that* $U_P^\infty(\hat{\sigma}) = \sum_{t=1}^{\infty}\delta_P^{t-1}[S(\hat{q}) - \hat{T}] = \frac{\delta_P}{1-\delta_P}[S(\hat{q}) - \hat{T}]$ *and* $U_\ell^\infty(\hat{\sigma}) = \sum_{t=1}^{\infty}\delta_A^{t-1}[\hat{T} - \theta_\ell\hat{q}] = \frac{\delta_A}{1-\delta_A}[\hat{T} - \theta_\ell\hat{q}]$ *for all* $\ell = 1, \cdots, L$.

Proof: The claim is that the strategy combination $\{(\hat{\sigma}^P, \hat{\sigma}^{A\ell})_{\ell=1}^L\}$ described below is a pooling equilibrium.

(i)$\hat{\sigma}_t^P(h_{t-1}) = (\hat{q}, \hat{T})$ if the past history has been $(\hat{q}, \hat{T})$ in every period up to $t-1$.
(ii) If the principal offers $(q, T) \neq (\hat{q}, \hat{T})$ in any period t and both the principal and the agent had offered and produced $(\hat{q}, \hat{T})$ in all previous periods, then the agent produces $q = 0$ from time $t+1$ onwards for K periods. This is a phase I punishment strategy.
(iii) If the agent produces $q \neq \hat{q}$ in any period t and both the principal and the agent had offered and produced $(\hat{q}, \hat{T})$ in all previous periods, then the principal offers $T = \theta_L q$ if q is the output of the agent for K periods. This is a phase I punishment for the agent.
(iv) If there are no deviations during a phase I punishment by either the principal or the agent then after the length of time K the principal offers $(\hat{q}, \hat{T}+\epsilon)$, such that $\hat{T}+\epsilon < S(\hat{q})$ if the principal had been the deviator, and offers $(\hat{q}, \hat{T}-\epsilon)$ such that $(\hat{T}-\epsilon) - \theta_L q > 0$ if the agent had been the deviator.
(v) If the agent deviates during a phase I punishment for the principal, then the offer switches to $T = \theta_L q$ if q is the output for a length of time K. If the principal deviates while punishing the agent during a phase I punishment then the offer switches to $q = 0$ for the next K periods. Such a punishment is a phase II punishment.
(vi) After a phase II punishment for the principal the offer switches to $(\hat{q}, \hat{T}+\epsilon)$ and after a phase II punishment for the agent the offer becomes $(\hat{q}, \hat{T}-\epsilon)$.
(vii) If the principal deviates after a phase II punishment, then the phase I punishment for the principal is played after which the offer becomes $(\hat{q}, \hat{T}+\epsilon)$.
(viii) Finally, if the agent deviates after a phase II punishment, then the phase I punishment for the agent is played after which the offer becomes $(\hat{q}, \hat{T}-\epsilon)$.

We now proceed to show that the strategy profile $\sigma^\star$ is an equilibrium irrespective of the type of the agent.

Let M_A be the maximum "gain" the agent can make by deviating in any period irrespective of its type. If the agent deviates in any period then its maximum payoff in the subsequent periods, if it has cost θ_ℓ, is at most

$$M_A + \delta^K \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\hat{T} - \epsilon - \theta_\ell \hat{q}]$$

as for a length of time K the agent's payoff is zero or less every period. If the agent does not deviate, its payoff in the subsequent periods is

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1}[\hat{T} - \theta_\ell \hat{q}].$$

Therefore, from the construction of the strategy profile, the agent does not gain from a deviation if

$$\sum_{\nu=1}^{\infty} \delta^{\nu-1}[\hat{T} - \theta_\ell \hat{q}] \geq M_A + \delta^K \sum_{\nu=1}^{\infty} \delta^{\nu-1}[\hat{T} - \epsilon - \theta_\ell \hat{q}]. \tag{1}$$

That is,

$$\frac{1-\delta^K}{1-\delta}[\hat{T}-\theta_\ell\hat{q}] \geq M_A - \frac{\delta^K}{1-\delta}\epsilon. \tag{2}$$

Now note that the expression $\frac{1-\delta^K}{1-\delta} \to K$ as $\delta \to 1$, therefore, there is a $\delta_{A1} : 0 < \delta_1 < 1$ and K_1 sufficiently large for which equation (2) is satisfied for all $\ell = \{1, \cdots, L\}$. Choose K_1 so that

$$\frac{1-\delta^K}{1-\delta}[\hat{T}-\theta_L\hat{q}] \geq M_A - \frac{\delta^K}{1-\delta}\epsilon. \tag{3}$$

Thus, phase I punishments can deter the agent from deviating irrespective of its cost for $\delta \geq \delta_1$ and $K \geq K_1$. Similarly, the principal does not gain from a deviation if

$$\sum_{\nu=1}^{\infty}\delta^{\nu-1}[S(\hat{q})-\hat{T}] \geq M + \delta^K\sum_{\nu=1}^{\infty}\delta^{\nu-1}[S(\hat{q})-\hat{T}-\epsilon]. \tag{4}$$

That is, if

$$\frac{1-\delta^K}{1-\delta}[S(\hat{q})-\hat{T}] \geq M - \frac{\delta^K}{1-\delta}\epsilon. \tag{5}$$

Thus, if δ is chosen to be sufficiently large (say greater than δ_{P1}) and for a large enough K, the principal does not gain from a deviation.

We now consider deviations from a phase I punishment. It should be clear from the above analysis an agent cannot gain while the agent is being punished in a phase I punishment. But consider a deviation made by the agent during a phase I punishment when the principal is considered the deviator. Let L_A be the maximum loss every period that the agent sustains during a phase I punishment. Then the agent's payoff after deviating when $K-t$ $(1 \leq t < K)$ periods of the phase I punishment is left is then less than or equal to

$$M_A + \delta^{K-t}\sum_{\nu=1}^{\infty}\delta^{\nu-1}[\hat{T}-\epsilon-\theta_\ell\hat{q}],$$

and if the agent does not deviate, the payoff in the subsequent periods is:

$$\delta^{K-t}\sum_{\ell=1}^{\infty}\delta^{\nu-1}[\hat{T}+\epsilon-\theta_\ell\hat{q}] - \sum_{\nu=1}^{K-t}\delta^{\nu-1}L_A.$$

Therefore, the agent does not gain by deviating during a phase I punishment when the principal is being punished, if

$$\delta^{K-t}\sum_{\ell=1}^{\infty}\delta^{\nu-1}[\hat{T}+\epsilon-\theta_\ell\hat{q}] - \sum_{\nu=1}^{K-t}\delta^{\nu-1}L_A \geq M_A + \delta^{K-t}\sum_{\nu=1}^{\infty}\delta^{\nu-1}[\hat{T}-\epsilon-\theta_\ell\hat{q}]. \tag{6}$$

This reduces to

$$\frac{\delta^{K-t}}{1-\delta}2\epsilon \geq M_A + \frac{1-\delta^{K-t}}{1-\delta}L_A. \tag{7}$$

In equation (7) as $\delta \to 1$, the expression

$$\frac{1-\delta^{K-t}}{1-\delta}$$

goes to $K-t$ and the expression $\frac{\delta^{K-t}}{1-\delta}$ goes to ∞. Hence, there is a $\delta_{A2} : 0 < \delta_{A2} < 1$ such that equation (7) holds for all $\delta > \delta_{A2}$ and for all $\ell = 1, \cdots, L$. Again choose $K = K_2$ such that equation (7) holds[3].

Next, suppose the principal deviates while punishing the agent during a phase I punishment. Then the principal's payoff from deviating, when $K-t$ $(1 \leq t < K)$ periods of the phase I punishment is left, is less than or equal to

$$M + \delta^{K-t}\sum_{\nu=1}^{\infty}\delta^{\nu-1}[S(\hat{q}) - \hat{T} - \epsilon],$$

and if the principal does not deviate, the payoff in the subsequent periods is:

$$\delta^{K-t}\sum_{\ell=1}^{\infty}\delta^{\nu-1}[S(\hat{q}) - \hat{T} + \epsilon] - \sum_{\nu=1}^{K-t} L.$$

Therefore, the principal does not gain from deviating when the agent is being punished during a phase I punishment if

$$\delta^{K-t}\sum_{\ell=1}^{\infty}\delta^{\nu-1}[S(\hat{q}) - \hat{T} + \epsilon] - \sum_{\nu=1}^{K-t} L \geq M + \delta^{K-t}\sum_{\nu=1}^{\infty}\delta^{\nu-1}[S(\hat{q}) - \hat{T} - \epsilon]. \tag{8}$$

This reduces to

$$\frac{\delta^{K-t}}{1-\delta}2\epsilon \geq M + \frac{1-\delta^{K-t}}{1-\delta}L. \tag{9}$$

As before, as $\delta \to 1$ the left hand side of the inequality in (9) goes to ∞ and the right hand side goes to $M + (K-t)L$. Hence, there is a $\delta_{P2} : 0 < \delta_{P2} < 1$ such that for all $\delta > \delta_{P2}$ the inequality in (9) holds and the principal cannot gain by deviating during a phase I punishment.

We now consider deviations from a phase II punishment. Consider a deviation by the agent from a phase II punishment while punishing the principal. The payoff of the agent, if the agent deviates after t periods of the phase II punishment, is at most

$$M_A + \delta^{K-t}\sum_{\nu=1}^{\infty}\delta^{\nu-1}[\hat{T} - \epsilon - \theta_\ell \hat{q}],$$

[3] Note that the type of the agent enters this calculation through M_A and L_A but these are set so that (6) holds for agents of all types so if K is sufficiently large (7) will hold for agents of all types.

and if the agent does not deviate, the payoff in the subsequent periods is

$$\delta^{K-t}\sum_{\ell=1}^{\infty}\delta^{\nu-1}[\hat{T}+\epsilon-\theta_\ell\hat{q}]-\sum_{\nu=1}^{K-t}\delta^{\nu-1}L_A.$$

Therefore, the agent does not gain by deviating after t periods during a phase II punishment when the principal is being punished, if

$$\delta^{K-t}\sum_{\ell=1}^{\infty}\delta^{\nu-1}[\hat{T}+\epsilon-\theta_\ell\hat{q}]-\sum_{\nu=1}^{K-t}\delta^{\nu-1}L_A \geq M_A+\delta^{K-t}\sum_{\nu=1}^{\infty}\delta^{\nu-1}[\hat{T}-\epsilon-\theta_\ell\hat{q}]. \tag{10}$$

Note that this inequality is the same as the one in (6) and the same analysis that follows shows that for $\delta \geq \delta_{A2}$ the agent cannot gain by deviating from a phase II punishment.

Similarly, for the principal, a deviation from a phase II punishment while punishing the agent is not profitable if (8) holds and thus is not profitable for $\delta \geq \delta_{P2}$.

Finally, we consider deviations from the contracts $(\hat{q},\hat{T}-\epsilon)$ and $(\hat{q},\hat{T}+\epsilon)$ respectively. For the agent it is enough to show that the agent cannot profitably deviate from $(\hat{q},\hat{T}-\epsilon)$. If the agent deviates then the subsequent payoff of the agent of any type is at most

$$M_A+\delta^K\sum_{\nu=1}^{\infty}\delta^{\nu-1}(\hat{T}-\epsilon-\theta_\ell\hat{q})$$

and if he does not deviate then the payoff in the subsequent periods is

$$\sum_{\nu=1}^{\infty}\delta^{\nu-1}(\hat{T}-\epsilon-\theta_\ell\hat{q}).$$

Therefore, the agent does not gain from deviating if

$$\sum_{\nu=1}^{\infty}\delta^{\nu-1}(\hat{T}-\epsilon-\theta_\ell\hat{q}) \geq M_A+\delta^K\sum_{\nu=1}^{\infty}\delta^{\nu-1}(\hat{T}-\epsilon-\theta_\ell\hat{q}). \tag{11}$$

This reduces to

$$\frac{1-\delta^{K+1}}{1-\delta}(\hat{T}-\epsilon-\theta_\ell\hat{q}) \geq M_A. \tag{12}$$

As $\delta \to 1$, $\frac{1-\delta^{K+1}}{1-\delta} \to K+1$. Hence, for K such that

$$(K+1)(\hat{T}-\epsilon-\theta_\ell\hat{q}) > M_A$$

there is a δ_{A3} such that for all $\delta \geq \delta_{A3}$ the inequality in (11) holds and the agent cannot gain by deviating. A similar analysis for the principal shows that the principal cannot gain from deviating from $(\hat{q},\hat{T}+\epsilon)$ if

$$\sum_{\nu=1}^{\infty}\delta^{\nu-1}(s(\hat{q})-\hat{T}+\epsilon) \geq M+\delta^K\sum_{\nu=1}^{\infty}\delta^{\nu-1}(S(\hat{q})-\hat{T}+\epsilon) \tag{13}$$

that is if

$$\frac{1-\delta^{K+1}}{1-\delta}(S(\hat{q})-\hat{T}+\epsilon) \geq M. \tag{14}$$

Hence, there is a δ_{P3} such that for all $\delta \geq \delta_{P3}$ the inequality in (14) will hold if K satisfies

$$(K+1)(S(\hat{q})-\hat{T}+\epsilon) > M.$$

We have therefore shown that for $\delta > \max\{\delta_{A1}, \delta_{A2}, \delta_{A3}\}$, the agent cannot gain by deviating in any period t, given any history, and for $\delta > \max\{\delta_{P1}, \delta_{P2}, \delta_{P3}\}$, the principal cannot gain by deviating in any period t, given any history.

We now show that the strategy combination $\{(\hat{\sigma}^P, \hat{\sigma}^{A\ell})_{\ell=1}^L\}$ is a Perfect Bayesian equilibrium. We first note that since $(\hat{\sigma}^P|h_t, \hat{\sigma}) = (\hat{\sigma}^P|h_t, \hat{\sigma})$ and $(\hat{\sigma}^{A\ell}|h_t, \hat{\sigma}) = (\hat{\sigma}^{A\ell}|h_t, \hat{\sigma})$ for every ℓ, therefore we have $\nu_\ell|h_t, \hat{\sigma} = \nu_\ell|h_{t-1}, \hat{\sigma}$ for all $t \geq 1$. Hence, $\nu_\ell|h_t, \hat{\sigma} = \nu_\ell$ for all $\ell = 1, \cdots, L$ so that

$$\begin{aligned} \sum_{\ell=1}^{L}(\nu_\ell|h_t, \hat{\sigma})U_P^\infty(\sigma|h_t, \ell) &= U_P^\infty(\hat{\sigma}|h_t, \ell)\sum_{\ell=1}^{L}\nu|h_t, \ell \\ &= U_P^\infty(\hat{\sigma}|h_t, \ell). \end{aligned} \tag{15}$$

Since we have already shown that for any h_t and $ell = 1, \cdots, L$ and for all strategy σ^P of the principal, $U_P^\infty(\hat{\sigma}|h_t, \ell) \geq U_P^\infty((\sigma^P, \hat{\sigma}^{A\ell})|h_t, \ell)$, it now follows from (15) that

$$\begin{aligned} \sum_{\ell=1}^{L}(\nu_\ell|h_t, \hat{\sigma})U_P^\infty(\hat{\sigma}|h_t, \ell) &= U_P^\infty(\hat{\sigma}|h_t, \ell) \\ &\geq \sum_{\ell=1}^{L}(\nu_\ell|h_t, \hat{\sigma})U_P^\infty(\sigma^P, \hat{\sigma}^{A\ell})|h_t, \ell). \end{aligned} \tag{16}$$

Similarly, as we have already shown that for any h_t and for all strategy $\sigma^{A\ell}$ of the agent with cost θ_ℓ, $U_{A\ell}^\infty(\hat{\sigma}|h_t, \ell) \geq U_{A\ell}^\infty((\hat{\sigma}^P, \sigma^{A\ell})|h_t, \ell)$, for all $\ell = 1, \cdots, L$ it follows from $\nu_\ell|h_t, \hat{\sigma} = \nu_\ell$ for all $\ell = 1, \cdots, L$ that

$$\begin{aligned} \sum_{\ell=1}^{L}(\nu_\ell|h_t, \hat{\sigma})U_{A\ell}^\infty(\hat{\sigma}|h_t, \ell) &= U_{A\ell}^\infty(\hat{\sigma}|h_t, \ell) \\ &\geq \sum_{\ell=1}^{L}(\nu_\ell|h_t, \hat{\sigma})U_{A\ell}^\infty(\hat{\sigma}^P, \sigma^{A\ell})|h_t, \ell). \end{aligned} \tag{17}$$

But (16) and (17) then show that $\hat{\sigma}$ is a Perfect Bayesian equilibrium. It is by construction a pooling equilibrium. This thus concludes the proof. ∎

6 Equilibrium Contracts in the Finite-Horizon

While the result on pooling equilibrium shows that there are plenty of pooling contracts that are perfect Bayesian equilibrium of repeated games between the principal and the agent, we know that the optimal contract for the single-period is a separating contract in the sense that a menu of contracts is offered with each option in the menu meant for an agent of a particular type. Here we show that the only perfect Bayesian equilibrium of the repeated contract game with finite-horizon is the one in which the second-best optimal contract is offered in period 1, and then the first-best complete information contract is offered in the subsequent periods. It is well known that in the single-period, second-best optimal contract, the least efficient type with the highest marginal cost θ_L is offered $T_L = \theta_L q_L$ for the output q_L. That is, the single-period optimal contract is such that

$$U_L(q_L, T_L) = T_L - \theta_L q_L = 0.$$

Further, for $\ell \neq L$, the incentive compatible offers all satisfy the condition that

$$U_\ell(q_\ell, T_\ell) = T_\ell - \theta_\ell q_\ell > 0.$$

In this standard environment, the first-best outcome, in which the principal maximizes his profit, is characterized by

$$S'(q_\ell^\star) = \theta_\ell \text{ for all } \ell$$

where the efficient outcome is obtained by equating the principal's marginal benefit to the agent's marginal cost. We first examine what happens in a T-period contract. For this, we assume the monotone hazard rate property:

$$\frac{\sum_{k=1}^{i-1} \nu_k}{\nu_i} < \frac{\sum_{k=1}^{i} \nu_k}{\nu_{i+1}} \text{ for all } \ell = 1, \cdots, L-1.$$

From the literature on optimal contracts we know that this assumption implies that the second-best single-period contract will fully separate types in the sense that, in the optimal menu, the output levels assigned to the different types will be distinct; that is there will be no bunching.

Recall that when the principal offers the menu of contracts, it has to satisfy the following constraints:

$$(i)\ T_\ell - \theta_\ell q_\ell \geq T_k - \theta_k q_k \text{ for all } \ell, k \text{ and}$$

$$(ii)\ T_\ell - \theta_\ell q_\ell \geq 0 \text{ for all } \ell.$$

The first constraints (i) are the incentive compatibility constraints. The second constraints (ii) are participation constraints. This problem can be reduced greatly as follows. First, as in the two-type case, the least efficient agent's participation constraint is binding because $T_\ell - \theta_\ell q_\ell \geq T_\ell - \theta_\ell q_L \geq T_L - \theta_L q_L \geq 0$. Second, since the agent's utility function meets the single-crossing conditions, $\frac{\partial}{\partial \theta}\left[-\frac{\partial U/\partial q}{\partial U/\partial T}\right] = 1 > 0$, we can impose the monotonicity constraints such as $q_\ell \geq q_2 \geq \cdots \geq q_L)$[4]. In addition, we can restrict our attention on the local incentive constraints between two adjacent types and these incentive constraints will bind at the optimum.[5]

The reduced problem is given by

$$\text{Maximize}_{\{q_\ell, T_\ell\}} \quad \sum_{l=1}^{L} v_l(S(q_\ell) - T_\ell))$$

such that

$$\begin{aligned}
&(i)\ T_\ell - \theta_\ell q_\ell = T_{\ell+1} - \theta_\ell q_{\ell+1} \text{ for } \ell = 1, \cdots, L-1, \\
&(ii)\ T_L - \theta_L q_L = 0, \\
&(iii)\ q_1 \geq q_2 \geq \cdots \geq q_L.
\end{aligned}$$

The analysis is a quite straightforward extension of the standard two-type model. There is no distortion on the most efficient type agent's output. However, for the less efficient types, the production levels are distorted downward.

We turn to the next benchmark in which the principal makes a long term contract with the agent. In a dynamic model, the commitment is an important issue. The contractual outcome would be very different if the principal is not able to commit to renegotiating the initial contract. The reason is well-known in the literature. When the type of the agent is revealed in the first period, the principal can exploit this information fully if renegotiation is feasible. Below, we allow renegotiation in any period and investigate the renegotiation-proof contracts.

Denote the contract offered to type ℓ in period t by $(T_{\ell,t}, q_{\ell,t})$. Suppose that the principal offers a separating contract for the first time in period $\widehat{t}$. Thus from periods 1 through $\widehat{t} - 1$, the principal offers a pooling contract. After the separating contract in period $\widehat{t}$, the principal can offer the complete-information first-best contract.

[4]Consider the following the incentive constraints for $\theta_\ell \neq \theta_k$; $T_\ell - \theta_\ell q_\ell \geq T_k - \theta_\ell q_k$ and $T_k - \theta_k q_k \geq T_\ell - \theta_k q_\ell$. Summing the two constraints, we obtain $(\theta_l - \theta_k)(q_k - q_\ell) \geq 0$.

[5]Any global incentive constraint is implied by the two local incentive constraints. Consider $\theta_i < \theta_k < \theta_j$. The incentive constraint between θ_i and θ_k is $T_i - \theta_i q_i \geq T_k - \theta_i q_k$. Similarly, the incentive constraint between θ_k and θ_j is $T_k - \theta_k q_k \geq T_j - \theta_k q_j$. Adding the two constraints, we obtain $T_i - \theta_i q_i \geq T_j - \theta_i q_k + \theta_k(q_k - q_j)$. It is immediate that $T_i - \theta_i q_i \geq T_j - \theta_i q_j$.

For the pooling contract in period $t \le \widehat{t}-1$, we assume that the principal wants to induce the participation of the least efficient agent. Then, the pooling contract is simply $q_L^\star$ so that $T_{L,t}^\star - \theta_L q_{L,t}^\star = 0$. After separation in period $\widehat{t}$, the first-best contract can be offered because the type of the agent is fully revealed so that $t_\ell = \theta_\ell q_\ell^\star$ for all ℓ. This fully separating contract is robust to the possibility of renegotiation because there does not exist a Pareto-improving contract. Thus, we can apply the revelation principal to find the separating contract in period $t = \widehat{t}$.

Note however, that in order to induce the agent to produce the level of output that would reveal his true type, the principal has to compensate the agent for the revelation of his type. Thus, to induce information revelation in period $\widehat{t}$, type-ℓ agent's intertemporal incentive constraint has to satisfy

$$T_{\ell,\widehat{t}} - \theta_\ell q_{l,\widehat{t}} \quad \ge \quad T_{\ell+1,\widehat{t}} - \theta_\ell q_{\ell+1,\widehat{t}} + \sum_{t=\widehat{t}+1}^{T} \delta^{t-\widehat{t}} \left(T_{\ell+1,t}^\star - \theta_\ell q_{\ell+1,t}^\star \right).$$

Note that from period $\widehat{t}$, the agent's type is fully revealed and so each type of agent receives zero rent onward until T, i.e., $T_{l,t}^* = \theta_l q_{l,t}^*$.

In the long term contract with finite periods, a difference is that type $\ell < L$ agent will get an additional rent $\sum_{t=\widehat{t}+1}^{T} \delta^{t-\widehat{t}} (\theta_{\ell+1} - \theta_\ell)\, q_{\ell+1,t}^\star$. This tightens the IC constraints, but does not affect the other constraints. Thus, in $t = \widehat{t}$, the principal's problem is given by

$$\text{Maximize}_{\{q_{\ell,t}, T_{\ell,t}\}} \sum_{t=\widehat{t}}^{T} \delta^{t-\widehat{t}} \sum_{\ell=1}^{L} \nu_\ell (S(q_{\ell,t}) - T_{\ell,t})$$

such that

$$(i) \qquad T_{\ell,\widehat{t}} - \theta_\ell q_{\ell,\widehat{t}} = T_{\ell+1,\widehat{t}} - \theta_\ell q_{\ell+1,\widehat{t}} + \sum_{t=\widehat{t}+1}^{T} \delta^{t-\widehat{t}} \left(T_{\ell+1,t}^\star - \theta_\ell q_{\ell+1,t}^\star \right)$$

$$(ii) \qquad T_{L,t} - \theta_L q_{L,t} = 0, \text{and}$$

$$(iii) \qquad q_{1,t} \ge q_{2,t} \ge \cdots \ge q_{L,t}.$$

Lemma 1 *In period $t = \widehat{t}$, the optimal separating contract is the second-best contract given by*

$$q_{1,\widehat{t}} = q_1^\star \quad \textit{and} \quad q_{\ell,\widehat{t}} = q_\ell^{SB} \quad \textit{for} \quad \ell \ge 2.$$

The proof of this lemma is omitted because it is a straightforward extension of the single period case. The principal's expected intertemporal profit can now be written as

$$\sum_{t=1}^{T} \sum_{\ell=1}^{L} \nu_\ell U_P(\cdot; \widehat{t}) = \sum_{t=1}^{\widehat{t}-1} \delta^{t-1} \left\{ S(q_L^\star) - \theta_L q_L^\star \right\}$$

$$+\delta^{\widehat{t}-1}\left\{\nu_1(S(q_1^\star)-\theta_1 q_1^\star)+\sum_{\ell=2}^{L}\nu_\ell(S(q_\ell^{SB})-\theta_\ell q_\ell^{SB})-\left[\sum_{\ell=1}^{L}\nu_\ell(\theta_{\ell+1}-\theta_\ell)q_{\ell+1}^{SB}\right]\right\}$$
$$-\delta^{\widehat{t}-1}\left[\sum_{t=\widehat{t}+1}^{T}\delta^{t-\widehat{t}}\sum_{\ell=1}^{L}\nu_\ell(\theta_{\ell+1}-\theta_\ell)q_{\ell+1}^\star\right]+\sum_{t=\widehat{t}+1}^{T}\delta^{t-1}\sum_{\ell=1}^{L}\nu_l\left\{S(q_l^*)-\theta_l q_l^*\right\}.$$

The first term in the RHS is the principal's profit from the pooling contract from $t=1$ to $t=\widehat{t}-1$. The term in the second line of the RHS is her profit from the separating contract in period $t=\widehat{t}$. The first term in the third line is the informational rent that the principal has to pay for early revelation. Note that the principal has to pay two types of information rent to induce truth-telling. The first is the typical information rent in a single period. The second is the one for early revelation. The last term on the RHS is her profit from the first-best contract from $t=\widehat{t}+1$ to $t=T$.

What would be the optimal timing of separation? The answer is not *a priori* clear because the principal has to compensate a large amount of information rent for the agent's early revelation.

Lemma 2 *The optimal timing of separation is* $\widehat{t}=1$.

Proof: The principal's intertemporal profit can be rewritten as

$$\sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\widehat{t})=\sum_{t=1}^{\widehat{t}-1}\delta^{t-1}\left\{S(q_L^\star)-\theta_L q_L^\star\right\}$$
$$+\delta^{\widehat{t}-1}\left\{\nu_1(S(q_1^\star)-\theta_1 q_1^\star)+\sum_{\ell=2}^{L}\nu_\ell(S(q_\ell^{SB})-\theta_\ell q_\ell^{SB})-\sum_{\ell=1}^{L}\nu_\ell(\theta_{\ell+1}-\theta_\ell)q_{\ell+1}^{SB}\right\}$$
$$+\delta^{\widehat{t}-1}\sum_{t=\widehat{t}+1}^{T}\delta^{t-\widehat{t}}\left\{\sum_{\ell=1}^{L}\nu_\ell\left[(S(q_\ell^\star)-\theta_\ell q_\ell^\star)-(\theta_{\ell+1}-\theta_\ell)q_{\ell+1}^\star\right]\right\}.$$

Note that the information rent given by $\sum_{t=\widehat{t}+1}^{T}\delta^{t-\widehat{t}}\sum_{\ell=1}^{L}\nu_\ell(\theta_{\ell+1}-\theta_\ell)q_{\ell+1,t}^\star$ is now part of the third term. For expositional simplicity, let us denote the Principal's expected payoff when the expected output vector of the agent is $(q_1,q_2,\cdots,q_L)$ by

$$\widetilde{U}_P(q_1,q_2,\cdots,q_L)=\sum_{\ell=1}^{L}\nu_\ell\left[(S(q_\ell)-\theta_\ell q_\ell)-(\theta_{\ell+1}-\theta_\ell)q_{\ell+1}\right].$$

Note that the informational rent paid by the principal is included in this payoff. However, if the contract is a pooling one (q,T) then the payoff of the principal is simply $(S(q)-T)$. Thus, the principal's intertemporal profit can now be simply rewritten as

$$\begin{aligned}\sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\widehat{t}) &= \sum_{t=1}^{\widehat{t}-1}\delta^{t-1}(S(q_L^\star)-\theta_L q_L^\star)+\delta^{\widehat{t}-1}\widetilde{U}_P(q_1^*,q_2^{SB},\cdots,q_L^{SB})\\ &+ \delta^{\widehat{t}-1}\sum_{t=\widehat{t}+1}^{T}\delta^{t-\widehat{t}}\widetilde{U}_P(q_1^*,q_2^*,\cdots,q_L^*)\end{aligned} \tag{18}$$

where up to period $\widehat{t}-1$, the principal offers the optimal pooling contract, and then the second-best optimal contract, and then the complete-information separating contract. We claim that

$$\sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\widehat{t}) - \sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\widehat{t}+1) \geq 0. \tag{19}$$

To see this, observe that

$$\begin{aligned}
& \left[\sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\widehat{t}) - \sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\widehat{t}+1)\right] \\
= \; & \delta^{\widehat{t}-1}\left[\widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB}) - \widetilde{U}_P(q_L^\star, q_L^\star, \cdots, q_L^\star)\right] \\
& + \delta^{\widehat{t}}\left[\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star) - \widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB})\right].
\end{aligned} \tag{20}$$

Recall that $\widetilde{U}_P(q_1, q_2, \cdots, q_L) = \sum_{\ell=1}^{L}\nu_\ell\left[(S(q_\ell)-\theta_\ell q_\ell) - (\theta_{\ell+1}-\theta_\ell)q_{\ell+1}\right]$ is the principal's reduced maximization problem in a single period contract after inserting the binding Incentive compatibility constraints and the least efficient agent's participation constraint into the objective function. As the optimal solution of principal's constrained optimization problem is $q_1 = q_1^\star$ and $q_\ell = q_\ell^{SB}$ for $\ell > 1$ we must have

$$\widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB}) \geq \max\left\{\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star), \quad \widetilde{U}_P(q_L^\star, q_L^\star, \cdots, q_L^\star)\right\}$$

We further claim that

$$\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star) \geq \widetilde{U}_P(q_L^\star, q_L^\star, \cdots, q_L^\star).$$

Note that

$$\begin{aligned}
& \widetilde{U}_P(q_1^\star, q_2^\star, \cdots q_{L-2}^\star, q_{L-1}^\star, q_L^\star) - \widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_{L-2}^\star, q_L^\star, q_L^\star) \\
= \; & \nu_{L-1}\left[\left(S(q_{L-1}^\star) - \theta_{L-1}q_{L-1}^\star\right) - \left(S(q_L^\star) - \theta_{L-1}q_L^\star\right)\right] \geq 0.
\end{aligned}$$

This is because $q_{L-1}^\star = \operatorname{argmax}_{q_{L-1}}[S(q_{L-1}) - \theta_{L-1}q_{L-1}]$.

It now follows in a similar way that

$$\begin{aligned}
\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_{L-2}^\star, q_{L-1}^\star, q_L^\star) & \geq \widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_{L-2}^\star, q_L^\star, q_L^\star) \\
& \geq \widetilde{U}_P(q_1^\star, q_2^\star, \cdots q_{L-3}^\star, q_L^\star, q_L^\star, q_L^\star).
\end{aligned}$$

Arguing recursively in this it now follows that

$$\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star) \geq \widetilde{U}_P(q_L^\star, q_L^\star, \cdots, q_L^\star). \tag{21}$$

Using (21) in (20) we now have

$$\begin{aligned}
&\left[\sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\hat{t}) - \sum_{t=1}^{T}\sum_{\ell=1}^{L}\nu_\ell U_P(\cdot;\hat{t}+1)\right] \\
= \;& \delta^{\hat{t}-1}\left[\widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB}) - \widetilde{U}_P(q_L^\star, q_L^\star, \cdots, q_L^\star)\right] \\
& + \delta^{\hat{t}}\left[\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star) - \widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB})\right] \\
\geq \;& \delta^{\hat{t}-1}\left[\widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB}) - \widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star)\right] \\
& + \delta^{\hat{t}}\left[\widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star) - \widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB})\right] \\
\geq \;& \delta^{\hat{t}-1}(1-\delta)\left[\widetilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB}) - \widetilde{U}_P(q_1^\star, q_2^\star, \cdots, q_L^\star)\right] \\
\geq \;& 0.
\end{aligned} \tag{22}$$

This proves the claim in (19). Hence, the intertemporal contract that gives the principal the highest expected profit is the one that offers the second-best contract in period 1 together with the informational rents to the types $\ell > 1$ and then set output levels at $q_\ell^\star$ from period 2 onwards. ■

This result shows that in the contracting game with limited commitment that is repeated for T periods the only possible equilibrium in the game is one in which the principal offers the separating contract in period 1 together with the extra informational rent, learns about the type of the agent in period 1, and then sets the output at the first-best, efficient level for the type of the agent inferred from the output of the agent in period 1.

Theorem 2 *The unique Perfect Bayesian equilibrium in the T-period repeated contract game is one in which the principal offers the second-best optimal contract $\{q_\ell, T_\ell\}_{\ell=1}^{L}$ in period 1, together with the informational rent $\frac{\delta(1-\delta^{T-1})}{1-\delta}(\theta_{\ell+1}-\theta_\ell))q_\ell^\star$ for periods 2 through T, if the agent produces q_ℓ in period 1, and from period 2 onwards sets $q_t = q_\ell^\star$ and $T_t = \theta_\ell q_\ell^\star$.*

Proof: From lemma 2 it follows that the optimal strategy of the principal in the T-period game is to offer the second-best optimal separating contract in period 1 and then to offer the first-best, full information contract to the agent that is consistent with his choice of output. The principal in period 1 also offers the informational rent $\frac{\delta(1-\delta^{T-1})}{1-\delta}(\theta_{\ell+1}-\theta_\ell)q_\ell^\star$ to ensure that the agent of type ℓ for $\ell = 1, \cdots, L-1$, has no incentive to produce differently then the output consistent with his type.

The best response of the agent in period 1 to this offer is to produce the output level consistent with its type given the payment of the informational rent in period 1. The updated belief of the principal is then that $prob.(\hat{\ell}) = 1$ if $q_1 = q_{\hat{\ell}}^{SB}$ and $prob.(\ell) = 0$

otherwise. The principal then offers the contract ($q_t = q^\star_{\hat{\ell}}$, $T_t = \theta_{\hat{\ell}} q_{\hat{\ell}}$) for $t \geq 2$. The agent's best response, given this, is to produce the output level $q_t = q_{\hat{\ell}}$ in every period. ■

Thus in the standard adverse selection model with limited commitment, the principal prefers to separate in the first stage and then offer the complete-information, first-best contract after that in the finite-horizon case. A pooling contract is never part of a perfect Bayesian equilibrium in the finite-horizon case. This contrasts sharply with theorem 1.

Example 1 *An example in the Two-period Case.*

Consider the case where the principal's revenue function is given by $S(q) = \sqrt{q}$ and the marginal cost of the agent is either $\underline{\theta} = 1$ and $\bar{\theta} = 2$ with $Prob(\underline{\theta} = 1) = Prob(\bar{\theta} = 2) = \frac{1}{2}$. It can be checked that the first-best outcome is

$$\underline{q}^\star = \frac{1}{4} \text{ and } \bar{q}^\star = \frac{1}{16}$$

as $S'(\underline{q}^\star) = 1$ and $S'(\bar{q}^\star) = 2$, respectively.

Consider now the case when the principal offers a pooling contract in the first period. In this case the principal cannot learn the type of the agent. Thus, the optimal contract in the second period contract is the single-period second-best contract given by

$$\underline{q}^\star = \frac{1}{4} \text{ and } \bar{q}^{SB} = \frac{1}{36}.$$

If the principal offers a pooling contract in period 1, then that offer should be $\underline{q} = \bar{q}^\star = \frac{1}{16}$ with the payment $T_1 = 2 \times (\frac{1}{16})$. The principal's expected profit over the two periods is

$$\frac{1}{8} + \frac{\delta}{6}. \tag{23}$$

Now consider the situation in which the principal offers a separating contract in the first period. The principal now fully learns the type of the agent. As a result, the principal offers the first-best contract in the second period

$$\underline{q}^\star = \frac{1}{4} \text{ and } \bar{q}^\star = \frac{1}{16}.$$

In this case, the efficient agent would be able to get an extra information rent of $\frac{\delta}{16}$, in the second period, if he chose the contract meant for the inefficient agent. Thus, in order for the first period contract to be separating and incentive-compatible the principal has to compensate the efficient agent with this extra information rent in addition to the informational rent the efficient agent derives from the second-best optimal contract in period 1. The first period offer is thus the second-best output levels

$$\underline{q}^\star = \frac{1}{4} \text{ and } q^{SB} = \frac{1}{36},$$

with the efficient agent being offered $\underline{T}_1 = \frac{5}{18} + \frac{\delta}{16}$ in period 1. Note that this includes the informational rent $\frac{5}{18}$ of the second-best contract and the additional rent of $\frac{\delta}{16}$ so that the efficient agent does not mimic the less efficient agent in period 1. The principal's expected profit over the two periods is then

$$\frac{1}{6} + \frac{5}{32}\delta. \tag{24}$$

Comparing (23) and (24) one can verify that the principal's expected profit from separating in the first period is greater than if he pooled in the first period for all $\delta \leq 1$. This thus illustrates the result in theorem 2. In fact it is also interesting to observe that for a fairly low discount factor when $\delta = 0.5$, the increase in the expected stream of profit of the principal over the two periods from offering the separating contract is

$$\frac{\frac{1}{24} - \frac{\delta}{96}}{\frac{1}{8} + \frac{\delta}{6}} = \frac{\frac{7}{192}}{\frac{5}{24}} = \frac{7}{40}$$

or a 17 percent increase in the expected profit of the principal. ∎

The result that in the finite-horizon the principal would prefer to offer a separating contract raises the question about the nature of the optimal contract in the infinite-horizon case. In the following sections we investigate this issue, and show that the principal can do even better in the infinite-horizon, by offering a separating contract, but not of the kind that is an equilibrium contract in the finite-horizon case.

7 Separating Long Run Contracts

We show here that in the infinite-horizon game not only is a pooling contract a perfect Bayesian equilibrium but so a separating contract in which the principal offers the second-best separating contract in period 1 and then continues to offer the terms of the second-best contract for the agent-type that is revealed in period 1.

Theorem 3 *There is a perfect Bayesian equilibrium in which the optimal single-period contract is offered in period* 1, *and from period* 2 *onwards, the only contract offered in equilibrium is* (T_ℓ, q_ℓ) *if* q_ℓ *is the output produced in period* 1. *In this Perfect Bayesian equilibrium the more efficient agent continues to earn the informational rent every period.*

Proof: The claim is that the strategy combination $\{(\tilde{\sigma}^P, \tilde{\sigma}^{A\ell})_{\ell=1}^L\}$ described below is a Perfect Bayesian equilibrium.

B (i)In period 1 the principal's strategy $\tilde{\sigma}_1^P$ is to offer the menu $\{T_\ell, q_\ell\}_{\ell=1}^L$.

(ii) In period 2, the principal offers $(T_{\hat{\ell}}, q_{\hat{\ell}})$ if in period 1 the agent chose the offer $(T_{\hat{\ell}}, q_{\hat{\ell}})$

from the menu $\{T_\ell, q_\ell\}_{\ell=1}^{L}$, otherwise offer $T = \theta_L q$ for any q the agent produces in each period for the next K periods.
(iii) If the past history has been $(T_{\hat{\ell}}, q_{\hat{\ell}})$ in every period up to $t-1$ then again offer $(T_{\hat{\ell}}, q_{\hat{\ell}})$ in period t.
(iv) If the principal offers $(q, T) \neq (T_{\hat{\ell}}, q_{\hat{\ell}})$ in any period $t \geq 2$ and both the principal and the agent had offered and produced $(T_{\hat{\ell}}, q_{\hat{\ell}})$ in all previous periods, then the agent produces $q = 0$ from time $t + 1$ onwards for K periods. This is a phase I punishment strategy.
(v) If the agent produces $q \neq q_{\hat{\ell}}$ in any period t and both the principal and the agent had offered and produced $(T_{\hat{\ell}}, q_{\hat{\ell}})$ in all previous periods, then the principal offers $T = \theta_L q$ for K periods after that. This is a phase I punishment for the agent.
(vi) If there are no deviations during a phase I punishment by either the principal or the agent, then after the length of time K, the principal offers $(T_{\hat{\ell}} + \epsilon, q_{\hat{\ell}})$, such that $T_{\hat{\ell}} + \epsilon < S(q_{\hat{\ell}})$ if the principal had been the deviator, and offers $(T_{\hat{\ell}} - \epsilon, q_{\hat{\ell}})$ such that $T_{\hat{\ell}} - \epsilon - \theta_{\hat{\ell}} q_{\hat{\ell}} > 0$ and if $\hat{\ell} = 1, \cdots .L - 1$, and (T_L, q_L) if $\hat{\ell} = L$, if the agent had been the deviator.
(vii) If the agent deviates during a phase I punishment for the principal, then the offer switches to $T = \theta_L q$ if q is the output in that period, for a length of time K. If the principal deviates while punishing the agent during a phase I punishment then the offer switches to $q = 0$ for the next K periods. Such a punishment is a phase II punishment.
(viii) After a phase II punishment for the principal, the offer switches to $(q_{\hat{\ell}}, T_{\hat{\ell}} + \epsilon)$ and after a phase II punishment for the agent the offer becomes $(q_{\hat{\ell}}, T_{\hat{\ell}} - \epsilon)$.
(ix) If the principal deviates after a phase II punishment, then the phase I punishment for the principal is played after which the offer becomes $(q_{\hat{\ell}}, T_{\hat{\ell}} + \epsilon)$.
(x) Finally, if the agent deviates after a phase II punishment, then the phase I punishment for the agent is played after which the offer becomes $(q_{\hat{\ell}}, T_{\hat{\ell}} - \epsilon)$.

We note that the strategy described in (i) through (x) is similar to the strategy used for the pooling equilibria of theorem 1. Similar arguments then show that there is a δ_P and a δ_A such that for all $\delta \geq \delta_P$, the principal cannot gain by offering a different stream of contracts than the one proposed and for all $\delta \geq \delta_A$ the agent, whatever be his type, cannot gain by producing an output different from the one meant for the type in the menu.

We now show that the strategy combination $\tilde{\sigma} = \{(\tilde{\sigma}^P, \tilde{\sigma}^{A\ell})_{\ell=1}^{L}\}$ is a perfect Bayesian equilibrium of the repeated contract game. Consider first the case in which the optimal single-period contract is such that it completely separates types. That is,

$q_{\hat{\ell}} \neq q_{\hat{s}}$ if $\hat{\ell} \neq \hat{s}$. In this case consider the belief system

$$\nu_{\hat{\ell}}|_{h_t,\tilde{\sigma}} = 1, \text{ and } \nu_{\hat{s}}|_{h_t,\tilde{\sigma}} = 0 \text{ for all } \hat{s} \neq \hat{\ell} \text{ if } h_1 = (T_{\hat{\ell}}, q_{\hat{\ell}}).$$

This is an updated belief system that is consistent with $\tilde{\sigma}$ as in period 1 the the only possible outcome for this strategy is in the set $\{(q_{\hat{\ell}}, T_{\hat{\ell}})\}_{\hat{\ell}=1}^{L}$. From the construction of $\tilde{\sigma}$ it should be clear that if the principal updates beliefs such that $\nu_{\hat{\ell}} = 1$ for some $\hat{\ell} \in \{1, \cdots, L\}$ then neither the principal nor the agent, if he is type $\hat{\ell}$, can gain by deviating from $\tilde{\sigma}$ after any history h_t. Thus the updated belief system together with the strategy $\tilde{\sigma}$ is a perfect Bayesian equilibrium of the repeated contract game. ∎

Notice that the principal can infer information about the type of the agent after observing the output level of the agent, and will only offer the contract meant for the type of agent that is consistent with the output produced in period 1. It is interesting to note that the offers are different from the complete information offers, as the more efficient types can expect to receive informational rents in each period, whereas the least efficient type is never required to produce $q = q_L^\star$, where $q_L^\star$ maximizes $S(q) - \theta_L q$, but only q^{SB} from period 2 onwards. The reason for this is that if the agent suspects that the principal will renege on the implicit arrangement of *not paying the informational rent*, then in period 1 the agent will never produce anything but $\bar{q}^{SB}$. Thereafter, the agent will always react to attempts by the principal to use the information fully by reverting to a punishment phase of not producing any output for a number of periods. Even in the case of the least efficient agent, the agent may not want to produce any more than q^{SB} as the agent has nothing to gain. The principal may in some cases prefer this, as the informational rent that has to be given to the agent, is in some cases, less than the informational rent that has to be given if the complete-information first-best contract is to be implemented.

We next examine the case in which there is Bunching, that is, the same contract is offered to several distinct types. Let S_ℓ denote the set of types that are bunched together with type ℓ, that is $S_\ell = \{s | (q_{\hat{\ell}}, T_{\hat{\ell}}) = (q_s, T_s)\}$. Define the following belief system belief system.

$$\nu_s|_{h_t,\tilde{\sigma}} = \frac{\nu_s}{\sum_{\lambda \in S_\ell} \nu_\lambda} \text{ if } h_1 = (T_s, q_s) \text{ and } s \in S_\ell, \text{ and } \nu_n|_{h_t,\tilde{\sigma}} = 0 \text{ if } n \notin S_\ell.$$

Thus it is possible that the Principal may offer a sequence of second-best contracts until types have been fully separated. It could be that the perfect Bayesian equilibrium is one in which the $\tilde{\sigma}$ after any history h_t. Thus, this belief system together with the strategy $\tilde{\sigma}$ is a perfect Bayesian Equilibrium of the repeated contract game.

8 Separating and Efficient Contracts

Having examined both the nature of pooling contracts as well as separating contracts, we now investigate whether it is possible to find a perfect Bayesian equilibrium in which the complete information efficient contract is offered at some point. As before we will look for contracts that are stationary over long periods. We reconsider again the menu of the optimal single-period contract given by $\{T_\ell, q_\ell\}_{\ell=1}^{L}$. We know that this single-period optimal contract is incentive compatible and may partially separate types but there could be bunching in the sense that the same offer is made to several different types. Let $S_n = \{s : T_s = T_n \text{ and } q_s = q_n\}$ denote the types that are made the same offer as type n. In the case of repeated contracts when there is bunching and an offer (T_ℓ, q_ℓ) is taken by the agent, the updated belief of the principal is that the agent's type is in S_n. If the belief of the principal about the type of the agent is given by $\{\nu_\ell\}_{\ell=1}^{L}$, then the updated belief of the principal after observing output level q_s is that the type of the agent is among those in S_s the set of types that will take the offer (T_s, q_s). The probability distribution that then gives the updated belief of the principal is

$$\nu_s^1 = \frac{\nu_s}{\sum_{k \in S_s} \nu_k} \text{ and } \nu_\ell^1 = 0 \text{ if } \ell \notin S_s.$$

In this case the optimal single-period contract that the principal can offer the agent is the menu that solves the following problem

$$\begin{aligned} \text{maximize} \quad & \sum_{\ell=1}^{L} \nu_\ell^1 (S(q_\ell) - T_\ell) \\ \text{such that} \quad & T_\ell - \theta_\ell q_\ell \geq T_{\ell'} - \theta_{\ell'} q_\ell \text{ for all } \ell' \neq \ell, \text{ and} \\ & T_\ell - \theta_\ell q_\ell \geq 0 \text{ for all } \ell = 1, \cdots, L. \end{aligned}$$

Let (T_ℓ^1, q_ℓ^1) denote the menu that solves the principal's problem given above. Again it is possible that the offer of several types may be bunched together, in which case, in the following round, the principal after updating his belief will offer a menu of contracts that maximizes the principal's expected payoff given the updated beliefs, subject to the incentive constraints and the participation constraints. This is the optimal contract of the principal after a second round of updating beliefs and we will denote this contract as (T_ℓ^2, q_ℓ^2). In general the optimal contract of the principal will be denoted by (T_ℓ^m, q_ℓ^m) after k rounds of updating of the principal's belief. As the principal continues to update his belief it will lead to a full separation of types after at most M rounds. What we show in the next result is that there is a perfect Bayesian equilibrium in the repeated contract game in which after K periods, the only contract offered is for the agent-type that is consistent with the principal's inference about the type of the agent from observing the output levels over the K-periods. Let $\{T_\ell^\star, q_\ell^\star\}_{\ell=1}^L$ denote the complete information efficient contract for the L different types. Then we have the following.

Theorem 4 (Optimality) *The Perfect Bayesian equilibrium that gives the highest expected profit to the principal is the separating perfect Bayesian equilibrium in which the principal offers the single-period second-best optimal contract in period* 1 *and from period* 2 *onwards the agent of type* θ_ℓ *is paid* T_ℓ *and asked to produce* q_ℓ *in every period. That is, the agent is paid in each period what he would have received in the second-best optimal contract. The updated belief of the principal about the type of the agent is consistent with the type of the agent revealed in period* 1.

Proof: From the proof of lemma 2 we have that

$$\begin{aligned}\sum_{\ell=1}^{L} \nu_\ell \tilde{U}_P(q_1^\star, q_2^{SB}, \cdots, q_L^{SB}) &\geq \sum_{\ell=1}^{L} \nu_\ell \tilde{U}_P(q_1^\star, q2^\star, \cdots, q_L^\star) \\ &\geq \sum_{\ell=1}^{L} \nu_\ell \tilde{U}_P(q_L^\star, q_L^\star, \cdots, q_L^\star).\end{aligned}$$

These expected payoffs include the informational rents that have to paid in order for the agent to choose an output level consistent with its type, and thus show that the discounted payoff from the separating contract, in which the second-best optimal contract is offered in every period, is the one that is optimal for the principal. ∎

Example 2 *Pooling versus a separating contract in the infinite-horizon case.*

Consider the same set up as in example 1 in which the principal's revenue function is given by $S(q) = \sqrt{q}$ and the marginal cost of the agent is either $\underline{\theta} = 1$ and $\bar{\theta} = 2$ with $Prob(\underline{\theta} =) = Prob(\bar{\theta} = 2) = \frac{1}{2}$.

According to Theorem 1, a pooling contract $(\widehat{q}, \widehat{T})$, such that $\widehat{q} < \frac{1}{4}$ and $\sqrt{\widehat{q}} > \widehat{T} > 2\widehat{q}$, can be a perfect Bayesian equilibrium in a infinite horizon game. It is clear that the pooling contract is sub-obtimal. Note that $\widehat{q} = \bar{q}^\star$ is the payoff-maximizing pooling contract. It can be verified as in example 1 that the pooling contract that is optimal for the principal among all pooling contracts is the pooling contract given by

$$\widehat{q} = \bar{q}^\star = \frac{1}{16} \text{ and } \widehat{T} = \frac{1}{8}.$$

The expected discounted sum of profits of the principal from this pooling contract is

$$\frac{1}{8(1-\delta)}. \tag{25}$$

The second-best optimal contract is given by

$$\underline{q}^\star = \frac{1}{4},\ q^{SB} = \frac{1}{36} \text{ and } \underline{T} = \frac{5}{18}, \bar{T} = \frac{1}{18}.$$

The expected profit of the principal from the second-best optimal contract is

$$\frac{1}{2}(\frac{1}{2} - \frac{5}{18}) + \frac{1}{2}(\frac{1}{4} - \frac{1}{8}) = \frac{1}{6}.$$

By theorem 4, the expected discounted sum of profit of the principal from offering the second-best contract is

$$\frac{1}{6(1-\delta)}. \tag{26}$$

Comparing (25) and (26) shows that the increase in the expected profit of the principal when he uses the second-best contract and separate types is,

$$\frac{\frac{1}{6(1-\delta)} - \frac{1}{8(1-\delta)}}{\frac{1}{8(1-\delta)}} = \frac{1}{3},$$

or an increase in expected profit of 33 percent. ∎

9 Conclusion

The results here indicate the nature of contracts that one would expect to see in situations that involve repeated interactions between a principal and an agent. We see that learning can take place, and if it does, it happens quickly as the agent can be induced to reveal his type as long as the principal does not take too much advantage of this information. One of the more interesting feature about these long run contracts is the sharp difference that exists between the finite-horizon contracts and the infinite-horizon contracts. This is primarily because there is no credible way in which the principal can offer not to take full advantage of the information that would be revealed by the agent. Thus, all the information rent of the agent has to be paid up front in period 1 itself. In the case of the infinite-horizon, the information rent can be paid every period and the payment can be phased out over the long run. This is another important difference between the equilibrium contracts between the finite and infinite-horizon cases. While in the case of the finite horizon there is the possibility that the agent may simply walk away after being paid all the information rent in a lump sum in period 1 unless it is part of the contractual obligation. In the case of the infinite-horizon this possibility does not arise as the informational rent is part of the payment that the more efficient type of the agent is paid in every period. In fact this implicit understanding between the principal and the agent, about how the agent will be paid in a perfect Bayesian equilibrium in each period, makes it possible to have an implicit payment plan that does not require an explicit enforceable contract.

It is important to ask at this point as to why the first-best complete information contract is not part of any perfect Bayesian equilibrium in the infinite-horizon case. We have already observed that if the principal wants to implement the first-best complete information contract, the principal has to pay the more efficient type of the agent a much higher informational rent in order to induce the more efficient types to produce optimally. The principal thus has a higher expected payoff from going with the second-best contract. Once the principal decides to use the second-best optimal contract, even after the agent's type is revealed, the principal cannot use that information to implement the first-best contract as it is not part of the perfect Bayesian equilibrium of the repeated contract game, and a deviation by the principal from the proposed equilibrium contracts in the repeated contract game leads to a punishment phases. It might also be the case that once the principal and the agent knows that the principal has inferred the type of the agent, the agent and the principal may want to renegotiate and propose sharing the additional surplus that would be generated if the agent produced the *efficient* output given his type. This is certainly possible when the terms of the agreement are given by the second-best contract as in the perfect Bayesian equilibrium of the infinite-horizon repeated contract game. However, this requires both the principal and the agent to *deviate simultaneously* from the perfect Bayesian equilibrium, and cannot be done by either the principal or the agent unilaterally. These issues do not arise in the case of the pooling contracts as the type of the agent is never inferred by the principal.

One might ask as to why information about the type of the agent is not revealed more slowly over time. The answer seems to be that with discounting there is no advantage for the principal to have any delay in the revelation of the information. The agent too does not gain by waiting to disclose the information about his type as the agent is appropriately rewarded for revealing the information. It is worth noting that if there is no mechanism to deter the principal from taking full advantage of the information revealed by the , and the agent knows this, then the agent would be much less willing to reveal his type and the resulting output may be much less than optimal for the principal and the output may not even be given by the second-best optimal contract.

While we have discussed a general adverse selection model in this paper, we note that it would also be useful to study this problem to study the more specific case of adverse selection in insurance markets, using some of the results developed here, and examine the nature of optimal contracts when there is repeated interaction between the insured and the insurer. We believe that the results obtained here may tell us much about the nature of the optimal contracts in insurance markets when there is adverse selection.

References

[1] Abreu, D., D. Pearce, and E. Stachetti (1986): "Optimal cartel equilibria with imperfect monitoring," *Journal of Economic Theory*, 39, 251-269.

[2] Athey, S. and K. Bagwell (2001): "Optimal Collusion with Private Information,"*Rand Journal of Economics*, Autumn 32(3).

[3] Athey, S. and K. Bagwell (2008): "Collusion with Persistent Shocks," *Econometrica*, 76, 493 - 540.

[4] Chakrabarti, S. K. (2010): "Optimal Collusion in Cournot Oligopolies with Unknown Costs," *International Economic Review*, 51, 2010, 1209-1238.

[5] Dewatripont, M. (1989): "Renegotiation and Information Revelation over Times: The Case of Optimal Labor Contracts," *Quarterly Journal of Economics*, 104, 589-619.

[6] Fudenberg, D. and E. Maskin (1986):"The folk theorem in repeated games with discounting or with incomplete information,"*Econometrica*, 62, 997-1039.

[7] Fudenberg, D. and J. Tirole (1991): "Perfect Bayesian Equilibrium and Sequential Equilibrium," *Journal of Economic Theory*, 53, 236-260.

[8] Hart, O., and J. Tirole (1988): "Contract renegotiation and Coasian Dynamics," *Review of Economic Studies*, 55, 509-540.

[9] Kreps, D. and C. Wilson (1982): "Sequential Equilibrium," *Econometrica*, 50, 863-894.

[10] Laffont, J. J., and J. Tirole (1990): "Adverse Selection and Renegotiation in Procurement," *Review of Economic Studies*, 62, 85-125.

[11] Levin, J. (2003): "Relational Incentive Contracts," *American Economic Review*, 93, 835-857.

[12] Pearce, D.G. and E. Stachetti (1998): "The Interaction of Implicit and Explicit Contracts in Repeated Agency," *Games and Economic Behavior*, 23, 447-480.

[13] Radner, R. (1981): "Monitoring Cooperative Agreements in a Repeated Principal Relationship," *Econometrica*, 49, 1127 - 1148.

[14] Radner, R. (1985): "Repeated Principal-Agent Games with Discounting," *Econometrica*, 53, 1173 - 1198.

[15] Spear, S. E., and S. Srivastava (1987): "On repeated Moral Hazard with Discounting," *Review of Economic Studies*, 54, 599-617.

[16] Rubinstein, A. and M. Yaari (1983): "Repeated Insurance Contracts and Moral Hazard," *Journal of Economic Theory*, 30, 74-97.