

Dynamic Contracts with Moral Hazard and Adverse Selection

Alex Gershkov, Motty Perry

The International School of Economics at Tbilisi State University (ISET) is supported by BP, the Government of Georgia, the Norwegian Ministry of Foreign Affairs, Higher Education Support Program of the Open Society Institute, the Swedish International Development Agency and the World Bank.

*International School of Economics at Tbilisi State University
16 Zandukeli Street, Tbilisi 0108, Georgia
www.iset.ge e-mail: publications@iset.ge*

Dynamic Contracts with Moral Hazard and Adverse Selection*

Alex Gershkov[†]

Motty Perry[‡]

20-06-11

Abstract

We study a novel dynamic principal–agent setting with moral hazard and adverse selection (persistent as well as repeated). In the model an agent whose skills are his private information faces a finite sequence of tasks, one after the other. Upon arrival of each task the agent learns its level of difficulty and then chooses whether to accept or refuse each task in turn, and how much effort to exert. Although his decision to accept or refuse a task is publicly known, the agent’s effort level is his private information.

We characterize optimal contracts and show that the per-period utility of the agent approaches his per-period utility when his skills are publicly known, as the discount factor and the time horizon increase.

1 Introduction

We study a novel, dynamic principal–agent setting with moral hazard and adverse selection (persistent as well as repeated). In the model an agent whose quality is his private information faces a finite sequence of tasks, one after the other. Upon arrival of each task the agent discovers its level of difficulty, which is an independent random variable, and decides whether to accept or refuse the task. If he accepts, he decides how much effort to exert. Although his decision to accept or refuse a task is publicly known, his effort level is his private information—and the source of the moral hazard in the model.

There are many economic interactions for which this model might be relevant. For example, a venture capital manager receives funds from investors who

*We are grateful to Bruno Biais and two anonymous referees for very useful comments. We wish to thank Michael Borns for his invaluable editorial work. Both authors gratefully acknowledge financial support from the Israel Science Foundation. Alex Gershkov is grateful to the Falk Institute for financial support.

[†]Alex Gershkov: Department of Economics and Center for the Study of Rationality, Hebrew University of Jerusalem, alexg@huji.ac.il.

[‡]Motty Perry: Department of Economics, University of Warwick; perrymotty49@gmail.com.

desire to invest their money but lack the knowledge to do so personally. For a money manager investment opportunities arrive sequentially; some are easy to assess and manage, others, more difficult. The probability of success of a given investment is a function of the manager's quality, the complexity of the investment, and the effort exerted in first analyzing and then following and directing the investment once it is made. The investor's problem then is how to design an optimal compensation contract in light of the moral hazard and adverse selection problems that arise from the fact that the manager's quality and effort as well as the complexity of the available investment opportunities, are the manager's private information. In particular, the investor's goal is to create an incentive schemes that distinguish between low-quality money managers who cannot do well in complex investments from high-quality managers who can. Another scenario for which our set up might be of some relevance is the case of a health-care insurer or public official who employs surgeons, whose quality he does not observe, to treat a flow of patients, the severity of whose ailments is also the surgeon's private information. Our principal's problem, then, is to design a system of contracts that guarantee that surgeries are performed, and effort is exerted, if and only if the surgeon's quality matches the severity of the patient's problem, and to do so at minimal costs.

In our model, in every period, an agent encounters one task and decides whether to refuse or to accept it, and in the latter case, whether to exert a costly, unobservable effort. The probability of successful accomplishment of a task in every period is monotonic in the agent's effort at that period, but it is also a function of the agent's quality and the complexity of the task in question. While the agent's quality is determined once and for all at the start of the contract, the type of the task is drawn independently each period, and both the agent's quality and the type of the task are the agent's private information.

For simplicity, we confine our attention to a special case where the principal's preferences are lexicographic: he is concerned primarily with matching the complexity of the task and the quality of the agent, and only secondarily with payments. More precisely, depending on the agent's quality and the complexity of the task a different action of the agent is desired by the principal. If the task is simple, all types of agents (low- or high-quality) should exert effort; but if the task is complex, only high-quality agents should exert effort while low-quality agents should refuse the task.

As is always the case in solving for adverse selection, the principal offers the agent a menu of contracts, one for each quality level, that promise financial rewards as a function of the observed history of whether tasks were accepted, and if so, whether they were successfully accomplished. In this menu of contracts, the optimal contract will be the one that provides the right incentives and minimizes expected payments.

Thus, the interaction between the principal and the agent is a dynamic model

of moral hazard and repeated (and persistent) adverse selection: moral hazard arises from the fact that the principal cannot observe the agent's effort, and adverse selection, from the fact that the principal cannot observe the agent's quality or the task's complexity.

We start with a model where agents are risk-neutral and they discount the future at the same rate as the principal. With this rather stylized model we characterize the optimal contracts that take advantage of the dynamic nature of the interaction. It is shown that as the interaction time increases, the expected transfer per period decreases and in the limit approaches the optimal payment when an agent's quality is publicly known. We could obtain this result because the dynamics enables us to make the contract of the high-quality agent unattractive to the low-quality one—increasingly so as the interaction time increases. The intuition beyond this result is rather simple. Exerting effort on a task is a gamble whose probability of success is higher, the higher is the quality of the agent. Successful accomplishment of a sequence of tasks is exponentially less likely for a low-quality agent. An optimal contract for a high-quality agent takes advantage of this fact by stipulating high rewards for a long history of successes. To construct these types of sequences and at the same time preserve incentives to exert effort, the optimal contract stipulates that in every period t , a success is rewarded only if it is followed by an uninterrupted sequence of successes until period T , the end of the interaction.

The optimal contract for the high-quality agent, in which he is compensated for a success in period t only if he keeps succeeding in every period thereafter till the end of game, looks rather extreme when T becomes large. Not only is the contract very risky, but it also has the unpleasant feature that no payment is guaranteed until the end at T . Of course, these features are irrelevant when the agent and the principal are both risk-neutral and do not discount future payments, as is assumed in this basic model. Yet, we show that for T large enough, the set of optimal contracts is not a singleton and there are other contracts in the set in which these two features are relaxed dramatically.

We next turn to the case where the agent is less patient than the principal. We show that the optimal contract for every type of agent is essentially unique. Moreover, the property of using the dynamics to make the contract of the high-quality agent more risky is preserved by postponing rewards for success in t until some later period $t' > t$ and making them conditional on success in all periods in between. But unlike the case with no discounting where this procedure incurred no costs, the principal now faces a trade-off and as a result the optimal contract takes a less extreme form even for short-term contracts. That is, conditional on the agent being a high-quality one, postponing his payments for success in period t to some period $t' > t$ is costly to the principal because he must then increase the expected payment to compensate the agent for the delay. But conditional on the agent being a low-quality agent, postponing the payments in the contract for the

high-quality agent is beneficial because it makes this contract less attractive to the low-quality agent and enables the principal to lower the expected payments to this agent.

It follows that in contrast to the no-discounting case, in this case it may not be optimal to postpone payments for success to later periods (let alone period T). On the one hand, it is better for the principal to start the procedure of postponing payments only after histories that are more likely to occur in the event that the low-quality agent accepted this contract. On the other hand, applying this procedure at a later stage reduces its effectiveness. We illustrate this trade-off by means of a simple example where it is shown that sometimes payments are not postponed initially, but only after histories that are unlikely to occur when the agent is a high-quality one. We then provide a partial characterization of the unique optimal contract and a more complete characterization for a discount factor close enough to one.

The paper is organized as follows. Section 2 is a brief survey of the literature. We present the basic setup with no-discounting and risk neutrality in Section 3. In Section 4 we define the notion of an admissible contract. The optimal contracts are characterized in Section 5. In Section 6 we study the case where agents are less patient than the principal. Section 7 concludes. Most of the proofs are relegated to the Appendix.

2 Related Literature

The existing models evolved gradually from models of moral hazard only to models in which moral hazard as well as adverse selection problems are present, and from models in which only short-term contracts are offered to those in which the principal can commit to a long-term contract. The small sample of papers discussed below attest to this evolution, and no attempt is made to provide an exhaustive survey of a very productive field.

One of the first papers on dynamic agency is Rubinstein and Yaari (1983), which considered an infinitely repeated moral hazard problem and demonstrated the existence of a strategy for the principal that yields the first best in an environment in which the principal cannot commit to a strategy that governs the relation. It is worth noting, however, that the infinitely repeated aspect of their problem is crucial to the derivation of their result, which indeed falls within the realm of the theory of repeated games. In a pioneering paper on career concern and reputation, Holstrom (1982) studied the provision of incentives to exert effort when the agent's ability is unobserved in finitely repeated interactions without output-contingent multi-period contracts.

Laffont and Tirole (1988) explored a dynamic two-period model of moral hazard *and* adverse selection and identified the ratchet effect that occurs whenever the principal is constrained to offer a short-term contract. That is, the equilib-

rium is characterized by more pooling in the first period as agents internalize the cost involved in revealing their type. Baron and Besanko's (1984) model of moral hazard and adverse selection is one in which the principal can commit to a long-term strategy, but the moral hazard problem is not dynamic. In particular, they studied the case of a regulated monopoly that first invests in R&D and in subsequent periods observes privately its marginal cost, which depends stochastically on the level of investment in R&D in period zero. Thus, their model is a one-shot moral hazard problem followed by a multi-period incentive scheme under adverse selection.

An important contribution is Holmstrom and Milgrom (1987), which studied a finitely (as well as a continuous-time) repeated moral hazard problem, but, unlike the Rubinstein–Yaari model, and along the lines we are pursuing in our paper, the principal in their model can commit to a long-term strategy that governs the relations in all periods. That is, the principal pays the agent at the end of the last period based on the entire observable history. It is shown that the optimal compensation scheme is a simple linear function of observable events. Similarly, Malcomson and Spinnewyn (1988), Rey and Salanie (1990), and Fudenberg, Holmstrom, and Milgrom (1990) studied the question of when the long-term optimal contract can be replicated by a sequence of short-term (spot) contracts.

DeMarzo and Fishman (2007) analyzed a dynamic moral-hazard problem in which the principal uses multidimensional tools to provide incentives to the agents. One is through instant cash payments, and another is by affecting the continuation value of the agent. Because the agent in this model is less patient than the principal, absent moral hazard, compensation for the agent's effort at a given period should be made immediately. However, postponing the payment mitigates the agency problem presented by the moral hazard, since it increases the share of the agent in future profits. Hence, the optimal contract balances this trade-off and involves both instruments. In our model, in contrast, postponing payment and conditioning it on the revealed performance in between alleviates the screening problem and allows the principal to decrease the compensation of the low-quality agent.

Biais, Mariotti, Plantin, and Rochet (2007) extended DeMarzo and Fishman's (2007) analysis to an infinite-horizon model and showed that the agent receives cash compensation only when the accumulated performance reaches a prespecified threshold. Moreover, they illustrated how the optimal contract is implemented with standard financial instruments and showed the convergence of the discrete-time model to the continuous-time version of DeMarzo and Sannikov's (2006) model. Sannikov (2008) used a very elegant technique based on the Martingale Representation Theorem to solve a continuous-time moral hazard problem that allowed him to obtain a very clean characterization of the optimal contract. The optimal balance between immediate compensation and the contin-

uation value was shown also in Biais, Mariotti, Rochet, and Villeneuve (2010), where the principal can affect the agent's continuation value through a change in the firm's size. In a recent paper Edmans, Gabaix, Sadzik, and Sannikov (2010) analyzed a dynamic moral hazard problem with a risk-averse agent and a risk-neutral principal. In their paper, too, compensation for performance in any given period is spread over future periods. The optimality of the spread follows from the optimal risk-sharing perspective, while in our paper the spread is used to reduce the cost of screening the types of the agent.

A significant difference between our model and the multi-periods models described above is that the latter are concerned solely with a dynamic moral hazard problem where it is possible to summarize the agent's incentives using his continuation value, i.e., the agent's future expected payoff when he follows the requested sequence of actions. This method is not applicable to the environment in our model, where there is also an element of adverse selection, since the incentives of different types of agents should be taken into account.

Fong (2009) combined the problem of moral hazard and adverse selection in the dynamic environment of health care provision. As in our model, in Fong's model the principal seeks to induce the agent to follow some course of action that may depend on the type of the agent. But Fong does not allow for the use of money as an instrument in the contracts. It follows that the only available tool for providing incentives is the flow rate of tasks and Fong's first result is that there is no need to consider complicated contracts because an optimal policy takes the form of a stopping rule that specifies if and when an agent is to be permanently fired. The main result is a characterization of the optimal contract-pair that takes the form of scoring rules in which the agent's past performance is summarized by a single score and the agent is fired if his score falls below a certain threshold, and he is tenured if his score rises above some other threshold. Contracts for agents of different quality levels are different in their sensitivity to success and failure. Another study of a continuous-time model of dynamic agency with moral hazard and adverse selection is Sannikov (2007). In Sannikov's model a principal employs an agent of unknown skill, where the principal observes no information during the contract period and needs to condition his compensation only on the reports of the agent. In this environment, to prevent manipulation by the agent, the optimal contract requires very specific conditioning of the compensation for the reported information. More precisely, the agent gets a credit line and he is compensated only if the balance of the line was above the prespecified cutoff during the whole contract period.

Our model incorporates all the incentive problems mentioned above. On the one hand, it is a dynamic moral hazard model, since the agent's choice of effort in any given period is unobservable. On the other hand, there are two types of adverse selection problems to overcome: a persistent adverse selection problem due to the unobservability of the agent's competence as determined in

period zero, and a dynamic adverse selection problem due to the fact that the complexity of the task, which is different in every period, is observable only by the agent.

3 The Model

Basic setup

Consider an agent who is employed by a principal for T periods. In every period $t \in \{1, 2, \dots, T\}$, the agent receives a task and has to decide whether to accept or refuse it and in the former case whether to exert a costly effort $C \in \{0, c\}$. The probability of successfully accomplishing the task in period t is positive only if $C = c$, but it is also a function of the agent's quality, denoted by s , and the task's level of difficulty at t , denoted by p_t and referred to as the task's "type" at t .

An agent's quality is either low or high and is denoted by $s \in \{h, l\}$, respectively. Conditional on exerting effort c , an agent of type h has a higher probability of succeeding at a given task. Similarly, the task arriving in period t is either easy or difficult ($p_t \in \{e, d\}$, respectively), and conditional on the agent's quality, the probability of success is higher when the task is easy. We assume that for all $t \in \{1, 2, \dots, T\}$, the task's type $p_t \in \{e, d\}$ is independently drawn and the probability that the task is of type d is q and type e is $(1 - q)$. Finally, the quality of the agent is his private information and the task's type p_t is revealed only at t and only to the agent.

Technology

The probability $\Pi : \{0, c\} \times \{l, h\} \times \{e, d\} \rightarrow [0, 1]$ of success at a given task is

$$\Pi(C, s, p_t) = \begin{cases} 0 & \text{if } C = 0 \\ \pi_{(s, p_t)} & \text{otherwise} \end{cases}$$

where for $s \in \{h, l\}$ and $p_t \in \{e, d\}$ we have

- (i) $0 < \pi_{(s, p_t)} < 1$
- (ii) $\pi_{(h, e)} > \pi_{(h, d)}$ and $\pi_{(l, e)} > \pi_{(l, d)}$
- (iii) $\pi_{(h, e)} > \pi_{(l, e)}$ and $\pi_{(h, d)} > \pi_{(l, d)}$.

Thus, conditional on exerting effort c , the agent's probability of success is higher if he is of high quality, for any type of task; and is higher when the task is easy, for any type of agent. The analysis reveals that the nature of the optimal contract depends on whether $\pi_{(l, e)} > \pi_{(h, d)}$ or $\pi_{(h, d)} > \pi_{(l, e)}$. The bulk of the paper is devoted to the more interesting case where $\pi_{(l, e)} \geq \pi_{(h, d)}$, while the treatment of the other case, being very similar, is provided in Appendix B.

Preferences

We start by assuming that the agent and the principal do not discount future payments.¹ The case of the different time preferences is analyzed in Section 6. The agent's VNM utility is a function of efforts and payments only. In particular, the utility of an agent who exerts effort in k periods and receives a total payment of m is $m - ck$. Thus, the agent is assumed to be risk-neutral and to maximize expected payment minus costs. The outside option generates a stream of utilities, which, for simplicity, are normalized to zero per period. Consequently, due to limited liability, negative payments are ruled out.

The principal

The agent here is employed by a principal. If the agent is a low-quality agent, i.e., $s = l$, the principal would like him to refuse a difficult task and to exert effort only if the task is an easy one, $p_t = e$. If, however, the agent is a high-quality one, $s = h$, then the principal would like him to exert effort on all types of tasks, easy as well as difficult ones.

Conditional on the agent doing what is expected, the principal's objective is to minimize expected payment. Thus, the principal's preferences are lexicographic. First and foremost, he is interested in providing incentives to the agent to accept tasks and to exert effort *only* when desirable. As there are many mechanisms that lead to these incentives, the principal is interested in the one that minimizes expected payment.

The principal can fully commit at time $t = 0$ to any observable history-dependent contract governing the agent's payments. Because the effort C , the agent's quality s , and the types of task p_t , for $t \in \{1, \dots, T\}$, are not observable by the principal, the only information available to the principal at t is a specification, for every $t' \leq t$, as to whether the task was accepted by the agent, and if so whether it was accomplished successfully or not.

4 Contracts

Recall that in our setup the principal, in every period t , observes one of three possible outcomes: (i) the task was accomplished successfully, (ii) the task was not performed, and (iii) the task was not successfully accomplished, which we denote by $\{1, 0, -1\}$, respectively. A contract thus specifies for every $t \in \{1, \dots, T\}$ the payment to the agent as a function of the observable history up to (and including) t which is a sequence of t elements from $\Psi = \{1, 0, -1\}$ and is denoted by ω_t where Ω_t denotes the set of all possible histories from time zero to t .² Without loss of generality we can assume that all payments are postponed to

¹The results are essentially the same if we assume that they discount future payments at the same rate.

²Following convention, we let \emptyset denote the history in period zero.

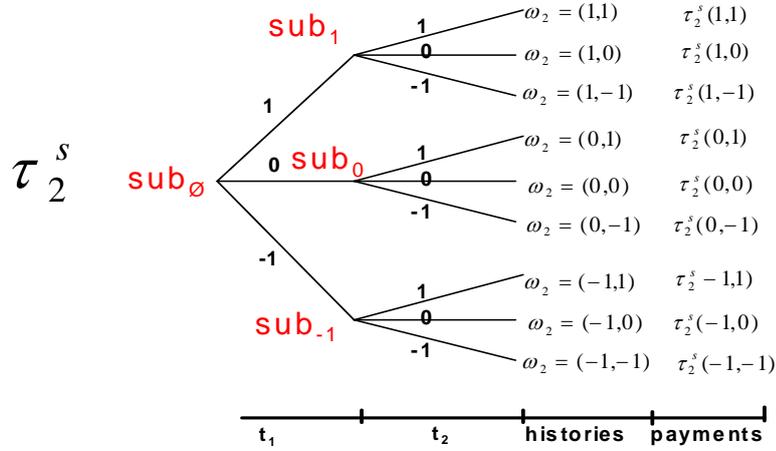


Figure 1: Two-period contract for the agent of type s , which specifies for any possible two-period history ω_2 the payment to this agent, $\tau_2^s(\omega_2)$.

the last period, T , and define a contract as follows.³

Definition 1 A T -period contract is a mapping $\tau_T : \Omega_T \rightarrow \mathbb{R}^+$ specifying the payment to the agent as a function of the observed history $\omega_T \in \Omega_T$.

As is typically the case in solving problems of adverse selection, the principal offers a menu of contracts, from which the agent chooses the contract that is best for him given his quality. Without loss of generality, we can restrict our attention to a mechanism in which only two contracts are offered by the principal: τ_T^h to the high-quality agent and τ_T^l to the low-quality one.

A two-period contract for an agent of type $s \in \{h, l\}$ is depicted above. Note that for every history $\omega_t \in \Omega_t$ we associate a subgame sub_{ω_t} that contains all possible observable histories following ω_t ⁴.

Definition 2 Admissible Contract-pair: A pair of contracts (τ_T^h, τ_T^l) is called admissible if it satisfies incentive compatibility (IC), individual rationality (IR), and efficiency (EF) where:

IC – an agent of quality h prefers the contract τ_T^h to τ_T^l , while the opposite holds for an agent of quality l .

³This definition will be modified in Section (6) when we relax the assumption that the agents do not discount future payments.

⁴The notion of subgame here, although obvious, is not exactly the one used in game theory.

IR – the contract τ_T^s yields a non-negative expected payoff to an agent of quality $s \in \{h, l\}$ starting after every history $\omega_t \in \Omega_t$ and for all $t \in \{1, 2, \dots, T\}$.

EF – for all $t \in \{1, 2, \dots, T\}$, an agent of quality h prefers to exert effort on all types of tasks, while an agent of quality l prefers to exert effort at t if and only if the task is of type e .

Remark 1 If (τ_T^h, τ_T^l) is admissible, then τ_T^h must entail taking the task in every period along the equilibrium path. It follows that if $\tau_T^h(\omega'_T) > 0$ for some ω'_T containing an outcome of zero (refuse), then there exists another contract $\tilde{\tau}_T^h$ in which $\tilde{\tau}_T^h(\omega'_T) = 0$ and $\tilde{\tau}_T^h(\omega_T) = \tau_T^h(\omega_T)$ for all $\omega_T \neq \omega'_T$ such that the new pair $(\tilde{\tau}_T^h, \tau_T^l)$ is admissible and yields, in equilibrium, the same expected payment to the principal. Thus, without loss of generality, we hereafter restrict our attention to contracts for the high-quality agent that pay zero whenever the history contains an outcome of zero. That is, if (τ_T^h, τ_T^l) is admissible, then $\tau_T^h(\omega'_T) = 0$ whenever $\{0\} \in \omega'_T$.

Remark 2 Note that if at some t and ω_t the contract provides the agent with incentives to exert effort on a given task's type, then the agent will exert effort whenever the arriving task has a higher probability of success. This implies that a contract-pair (τ_T^h, τ_T^l) satisfies EF if τ_T^h provides the high-quality agent with incentives to exert effort whenever a difficult task arrives (i.e., $p_t = d$), while τ_T^l provides the low-quality agent with incentives to exert effort only if the task is easy.

Remark 3 Since the agent can always choose to refuse, and since all payments are non-negative, all contracts satisfy IR.

Of all admissible contract-pairs, we are interested in the one that minimizes expected payment. Denote by $m^{s'}(\tau_T^s, sub_{\omega_t})$ the ex-ante (before observing the task's type in period $t + 1$) expected payment of contract τ_T^s to an agent of quality s' conditional on history ω_t being reached and conditional on playing optimally thereafter and let $u^{s'}(\tau_T^s, sub_{\omega_t})$ denote the ex-ante expected utility provided by contract τ_T^s to an agent of type s' conditional on reaching history ω_t and conditional on playing optimally thereafter. Note that $m^s(\tau_T^s, sub_{\emptyset})$ and $u^s(\tau_T^s, sub_{\emptyset})$ are monotonically related in all contracts τ_T^s satisfying EF. This is so because expected costs to an agent of quality s are the same in all contracts satisfying EF. In particular, given a T -period admissible contract-pair (τ_T^h, τ_T^l) , it is straightforward to verify that

$$u^h(\tau_T^h, sub_{\omega_t}) = m^h(\tau_T^h, sub_{\omega_t}) - c(T - t)$$

and

$$u^l(\tau_T^l, sub_{\omega_t}) = m^l(\tau_T^l, sub_{\omega_t}) - c(1 - q)(T - t)$$

where, as defined above, $(1 - q)$ is the probability that the task is easy, i.e., $p_t = e$.

We are now in a position to define an optimal contract-pair.

Definition 3 An Optimal Contract-pair. *A pair of contracts $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is called optimal if it is admissible, and if for every admissible contract-pair (τ_T^h, τ_T^l) we have*

$$m^h(\tau_T^h, sub_{\emptyset}) \geq m^h(\hat{\tau}_T^h, sub_{\emptyset}) \text{ and } m^l(\tau_T^l, sub_{\emptyset}) \geq m^l(\hat{\tau}_T^l, sub_{\emptyset}).$$

Finally, denote by p^s the ex-ante probability of a successful task by a quality s agent when effort is exerted. That is,

$$p^h = q\pi_{(h,d)} + (1 - q)\pi_{(h,e)}$$

and

$$p^l = q\pi_{(l,d)} + (1 - q)\pi_{(l,e)}.$$

5 The Optimal Contract-pair

In this section we maintain the assumption that $\pi_{(l,e)} > \pi_{(h,d)}$ and show that the optimal contract-pair is a separating pair, in the sense that agents of different quality sign different contracts. When this assumption does not hold (i.e., $\pi_{(l,e)} \leq \pi_{(h,d)}$) the unique optimal contract-pair is pooling. Since the analysis of the pooling case is very similar to that of the separating case, it is postponed to Appendix B.

We start by characterizing the set of optimal contracts when the agent is known to be a high-quality agent, and denote this set by Γ_T^h . We then show that when the agent's quality is *unobservable*, the contract offered to the high-quality agent belongs to Γ_T^h . Thus, when quality is unobservable, the contract assigned to the high-quality agent is the second-best contract as it does not generate information rents that correspond to the unobserved type of the agent and the binding constraint is the incentive constraint on the low-quality agent, whose purpose is to ensure that he prefers the contract assigned to him to the one assigned to the high-quality agent.

While the high-quality agent is indifferent between all contracts in Γ_T^h (see point 2 below), this is not the case for the low-quality agent. The main theorem of this section establishes that the optimal contract for the high-quality agent is the contract in Γ_T^h that would minimize the payoff of the low-quality agent if he pretended to be a high-quality one and accepted it. In this contract a success in period t is rewarded *only if* it is followed by a success in *every* period following t . This contract, in a way, is the riskiest contract in Γ_T^h ; however, and this is crucial, it is exponentially more risky for the low-quality agent than it is

for the high-quality one. In contrast, the optimal contract for the low-quality agent is the contract that pays a fixed amount per successful task and makes the low-quality agent indifferent between the two contracts. It is shown that as T gets larger, the per-success expected payment in the optimal contract-pair approaches the expected amount paid when quality is observable.

5.1 Agent's Quality is Known to be High

We now characterize the set of optimal contracts for an agent whose quality is known to be high. A contract $\hat{\tau}_T^h$ belongs to Γ_T^h if it satisfies *IR* and *EF* and if there is no other contract τ_T^h that also satisfies *IR* and *EF* and for which expected payment is lower, i.e., $m^h(\tau_T^h, \text{sub}_\emptyset) < m^h(\hat{\tau}_T^h, \text{sub}_\emptyset)$. Before we proceed and study the properties of Γ_T^h , a few points are worth mentioning.

1. Note that although the agent's quality is observable, there are still problems of moral hazard and adverse selection to solve because the agent's effort and the task's type are not observable by the principal. Indeed, note that if the task's type is also observable, then a first-best solution can be achieved through a simple contract that promises a payment of $c/\pi_{(h,d)}$ per success at a difficult task ($p_t = d$), and a payment of $c/\pi_{(h,e)}$ per success at an easy task ($p_t = e$). Such a contract satisfies *EF* and at the same time brings the agent to his *IR* utility. However, when the type of the task is not observable to the principal and he relies on the agent's report of the task's type, the contract is not incentive-compatible since the agent will always report that the task is difficult. As a result, when the task's type is not observable, the optimal contract does leave the agent some information rent.

Indeed, if $\hat{\tau}_T^h \in \Gamma_T^h$, then for every history $\omega_{T-1} \in \Omega_{T-1}$, $\hat{\tau}_T^h$ provides incentives for the agent to exert effort whenever the task is difficult. Consider the following feasible strategy for the agent: do not exert effort in all periods $t \in \{1, \dots, T-1\}$ and exert effort in T only if the task is easy. Note that this strategy guarantees a strictly positive expected payoff since the payment after a sequence of failures is non-negative and the payment in the last period provides incentives even if $p_T = d$. We conclude that if $\hat{\tau}_T^h \in \Gamma_T^h$, then $u^h(\tau_T^h, \text{sub}_\emptyset) > 0$.

2. The definition of Γ_T^h implies that expected payment is the same in all contracts in Γ_T^h . Since expected costs are the same in all contracts satisfying *EF* and in particular in all contracts in Γ_T^h , the agent is indifferent between all contracts in Γ_T^h .

A three-period contract for a high-quality agent, where histories containing zeroes are ignored, is described below.

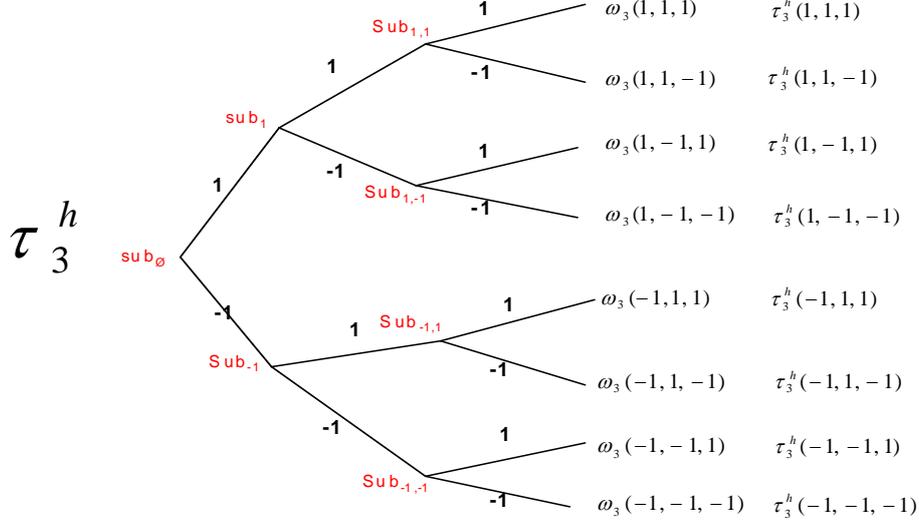


Figure 2: Three-period contract for a high-quality agent, which specifies for any possible three-period history ω_3 the payment to this agent, $\tau_3^h(\omega_3)$.

The following lemma, proved in Appendix A, lists a few properties that are satisfied by all contracts belonging to Γ_T^h . These properties are then used to characterize the set Γ_T^h of optimal contracts.

Lemma 1 *Properties of Γ_T^h*

1. If $\tau_K^h \in \Gamma_K^h$, then $\exists \tau_{K-1}^h \in \Gamma_{K-1}^h$ s.t. $\forall \omega_{K-1} \in \Omega_{K-1}$, $\tau_K^h(-1, \omega_{K-1}) = \tau_{K-1}^h(\omega_{K-1})$.
2. If $\tau_K^h \in \Gamma_K^h$, then $u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) = \frac{c}{\pi(h,d)}$.
3. If $\tau_K^h \in \Gamma_K^h$, then $m_K^h(\tau_K^h, sub_\emptyset) = Kp^h \frac{c}{\pi(h,d)}$.
4. Assume that $\tilde{\tau}_T^h$ satisfies IR and EF, but $\tilde{\tau}_T^h \notin \Gamma_T^h$. Then, there exists $\tau_T^h \in \Gamma_T^h$ such that for any history $\omega_T \in \Omega_T$, $\tilde{\tau}_T^h(\omega_T) \geq \tau_T^h(\omega_T)$ with strict inequality for at least one history $\omega_T' \in \Omega_T$.

The first property of the lemma refers to the payments restricted to sub_{-1} . In the context of Figure 2 above, it says that if a three-period contract belongs to Γ_3^h , then the induced two-period contract in sub_{-1} belongs to Γ_2^h . In other words, a failure in period one is not rewarded, and, as a result, from period two on, the agent faces a $K - 1$ -period contract.

Property 2 follows from the fact that effort is not observable and will not be exerted unless incentives are provided. In particular, if the first task to arrive turns out to be difficult, the agent will not exert effort unless the difference in

expected payoff between a success and a failure is enough to justify the risk of failure, an event that occurs with probability $(1 - \pi_{(h,d)})$ if effort is exerted. In a K -period contract, the reward for success in the first period (which occurs with probability $\pi_{(h,d)}$ if the task is difficult and effort is exerted) is given by $u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1})$. Thus, exerting effort on a difficult task is beneficial only if the expected gain is greater than the cost of exerting effort, that is, only if $\pi_{(h,d)}[u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1})] \geq c$. The content of the second property is that an optimal contract generates, in the first period, the minimal spread between the two subgames, that is needed to provide these incentives.

Property 3 follows from the one-to-one relations between expected utility and expected payment when EF is satisfied, and in particular it implies that Property 2 can be rewritten as

$$m^h(\tau_K^h, sub_1) - m^h(\tau_K^h, sub_{-1}) = \frac{c}{\pi_{(h,d)}}.$$

Of course, the exact same argument holds in every period. That is, in every period incentives to exert effort on a difficult task must be provided. Thus, for all $t \leq T$ and for every history ω_t the expected reward for success must be at least $\frac{c}{\pi_{(h,d)}}$, and it holds with equality in the first period. Finally, recall that ex-ante success occurs with probability $p^h = q\pi_{(h,d)} + (1 - q)\pi_{(h,e)}$, and you get the expected payment in a K -period contract specified in Property 3.

The first three properties are employed in the proof of the fourth property, which establishes an important characteristic property of the set Γ_K^h . That is, if a contract is not optimal, then there exists an optimal contract that pays less in *every* possible history. The proof of the following lemma, which is relegated to Appendix A, makes use of the four properties in Lemma 1 to provide a characterization of Γ_T^h and in particular to show that for all T , $\Gamma_T^h \neq \emptyset$.

Lemma 2 *Characterization of Γ_T^h .*

- i.* $\tau_1^h \in \Gamma_1^h$ if and only if $\tau_1^h(1) = c/\pi_{(h,d)}$, $\tau_1^h(-1) = 0$, and $\tau_1^h(0) = 0$.
- ii.* $\tau_{K+1}^h \in \Gamma_{K+1}^h$ if and only if τ_{K+1} can be constructed from contracts in Γ_K^h according to the following procedure:
 - ii.1* The τ_{K+1} payments restricted to sub_{-1} are a contract in Γ_K^h .
 - ii.2* The τ_{K+1} payments restricted to sub_1 are a contract in Γ_K^h inflated by an expected payment of $c/\pi_{(h,d)}$, which is allocated to the different histories of sub_1 in any way, provided that incentives to exert efforts are not distorted.

Recall that by definition the expected payment is the same in all optimal contracts. This fact together with Lemma 2 yields the following simple corollary and also establishes that the set Γ_T^h is not empty.

Corollary 1 *The set $\Gamma_T^h \neq \emptyset$ and in particular the contract $\hat{\tau}_T^h \in \Gamma_T^h$, where $\hat{\tau}_T^h(\omega_T) = \frac{c}{\pi_{(h,d)}}n(\omega_T)$, and $n(\omega_T)$ is the number of successful tasks in ω_T . Thus, a contract is optimal only if it pays in expectation $c/\pi_{(h,d)}$ for every successful task.*

5.2 Agent's Quality is Unobservable

Having characterized the set Γ_T^h we are now ready to study the case where the agent's quality is unobservable. Note that now the *IC* constraint must be taken into account since the agent will choose the contract that maximizes his expected utility, and not necessarily the one designed for him by the principal. We start by showing that if a contract-pair (τ_T^h, τ_T^l) is optimal, then $\tau_T^h \in \Gamma_T^h$.

Lemma 3 *If (τ_T^h, τ_T^l) is an optimal contract-pair, then $\tau_T^h \in \Gamma_T^h$.*

Proof. Assume by way of contradiction that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal but $\hat{\tau}_T^h \notin \Gamma_T^h$. Since $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is an optimal contract-pair, it is admissible, and, in particular, both contracts satisfy *IR* and *EF*. Hence, Property 4 in Lemma 1 implies that there exists a contract $\tilde{\tau}_T^h \in \Gamma_T^h$ such that for all history $\omega_T \in \Omega_T$, $\tilde{\tau}_T^h(\omega_T) \geq \hat{\tau}_T^h(\omega_T)$ with strict inequality for at least one history. Hence, replacing $\hat{\tau}_T^h$ with $\tilde{\tau}_T^h$ will decrease the expected utility of the low-quality agent in the event that he pretends to be a high-quality agent by adopting the high-quality agent's contract. Consider a contract $\tilde{\tau}_T^l$ that pays $r \geq c/\pi_{(l,e)}$ per success and makes the low-quality agent indifferent to the contract $\tilde{\tau}_T^h$. To see that such a contract $\tilde{\tau}_T^l$ always exists, it is enough to note that (i) a low-quality agent can always adopt the contract $\tilde{\tau}_T^h$ and then exert no effort to obtain a non-negative utility, and (ii) a contract that pays $c/\pi_{(l,e)}$ per success satisfies *EF* and yields zero expected utility to the low-quality agent.

We next argue that $(\tilde{\tau}_T^h, \tilde{\tau}_T^l)$ is admissible. That is, (i) $r \leq c/\pi_{(l,d)}$, and (ii) the high-quality agent prefers the contract $\tilde{\tau}_T^h$ to $\tilde{\tau}_T^l$. Note, however, that $r \leq c/\pi_{(h,d)}$ is sufficient for (i) and (ii). This is because (i) follows from $c/\pi_{(h,d)} < c/\pi_{(l,d)}$ and (ii) from the fact that a contract that pays $c/\pi_{(h,d)}$ per success belongs to Γ_T^h and the high-quality agent is indifferent between all contracts in Γ_T^h . Therefore, if $r \leq c/\pi_{(h,d)}$, the contract-pair $(\tilde{\tau}_T^h, \tilde{\tau}_T^l)$ is admissible and generates a lower expected payment to both agents than the pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$, which is a contradiction.

Let us therefore assume that $r > c/\pi_{(h,d)}$ and observe that a contract-pair that pays $c/\pi_{(h,d)}$ per success to both types of agents is admissible and provides both types of agents with an expected utility lower than $(\hat{\tau}_T^h, \hat{\tau}_T^l)$, which is again in contradiction to the assumed optimality of the original pair. We conclude that if a contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal, then $\hat{\tau}_T^h \in \Gamma_T^h$. ■

Note that while different contracts in Γ_T^h generate the same expected utility for the high-quality agent, they generate different expected utilities for the low-

quality one, if he chooses to adopt them. It thus follows from Lemma 3 that a contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal if the contract $\hat{\tau}_T^h$ is the one that minimizes the expected utility of the low-quality agent among all contracts in Γ_T^h . In other words, from the high-quality agent's point of view, the set Γ_T^h consists of different lotteries between which he is indifferent. However, from the point of view of the low-quality agent, these are different lotteries and $\hat{\tau}_T^h$, to be defined below, is the riskiest among them. That is, although the set Γ_T^h for $T > 1$ is not a singleton and contains many contracts, asymmetric information about the agent's type narrows down the contracts that the designer offers to the high-quality agent. The theorem also establishes that as $T \rightarrow \infty$, the optimal contract-pair converges to the second-best pair, that is, the contract-pair that is offered when the agent's quality is observable.

Prior to presenting the formal statement of the theorem, we describe its content when $T = 2$. In this case the theorem postulates that if the high-quality agent accepts the task in period one and succeeds, he is compensated for this only if he also succeeds in period two. The compensation in the event that there are two successes in a row must be high enough to cover the extra risk involved in exerting effort in period one. Specifically;

$$\hat{\tau}_2^h = \begin{cases} \frac{c}{\pi_{(h,d)}} \frac{1+p^h}{p^h} & \text{if } \omega_2 = (1, 1) \\ \frac{c}{\pi_{(h,d)}} & \text{if } \omega_2 = (-1, 1) \\ 0 & \text{if } \omega_2 = (1, -1) \\ 0 & \text{if } \omega_2 = (-1, -1) \end{cases}$$

Note that while the high-quality agent (being risk-neutral) is indifferent between this contract and the one that pays $\frac{c}{\pi_{(h,d)}}$ per success, the low-quality agent strictly prefers the latter.

The following theorem characterizes the optimal contract-pair for T periods, while making use of the following definitions:

- (i) Define $A(k)$ recursively by letting $A(0) = 0$ and $A(k) = A(k-1) + \frac{1}{(p^h)^{k-1}}$.
- (ii) Let $\tilde{k}(\omega_T)$ be the length of the longest uninterrupted sequence of successes in ω_T , starting from period T and proceeding backward.

Theorem 1 *An optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ has the following properties:*

1. *If ω_T contains an outcome of 0, then $\hat{\tau}_T^h(\omega_T) = 0$. Otherwise, if $\tilde{k}(\omega_T) = k$, then $\hat{\tau}_T^h(\omega_T) = \frac{c}{\pi_{(h,d)}} A(k)$.*
2. *There exists a constant r , such that $\hat{\tau}_T^l(\omega_T) = rn(\omega_T)$, where $n(\omega_T)$ is the number of successes in ω_T . Moreover, $\lim_{T \rightarrow \infty} r = \frac{c}{\pi_{(l,e)}}$.*

In words, the optimal contract for the high-quality agent pays zero after an history in which a task was refused at least once, and otherwise the agent is compensated only for those successes that were followed by an uninterrupted sequence of successes all the way to the end of the contract. (part (1) of the theorem). The optimal contract for the low-quality agent gives him a fixed constant payment per success (part (2) of the theorem). Theorem 1 provides the exact compensation scheme that depends on the number of uninterrupted successes starting from the end of the contract and proceeding backward. To provide incentives to the high-quality agent, the expected payment should increase by $\frac{c}{\pi(h,d)}$ per success. In order to decrease the expected utility of the low-quality agent from this contract, the payment for success in every period t should be postponed until the end of the game and provided only if it is followed by an uninterrupted sequence of successes. Paying $\hat{\tau}_T^h(\omega_T) = \frac{c}{\pi(h,d)}A(k)$ after history ω_t satisfies both requirements.

Proof. We start the proof by showing that the contract $\hat{\tau}_T^h$ described in the theorem minimizes the expected utility of the low-quality agent in all the contracts that belong to Γ_T^h . The formal argument follows from Claim 1, setting $\tilde{u} = 0$; the proof of the claim is relegated to Appendix A. First, note that if $\{0\} \in \omega_T$ and $\hat{\tau}_T^h(\omega_T) > 0$, then decreasing this payment will not affect the expected utility of the high-quality agent and will decrease (or will not affect) the expected utility of the low-quality agent from this contract. Therefore, without loss of generality we can restrict our attention to contracts in Γ_T^h , where the payments after histories containing $\{0\}$ are zero. ■

Claim 1 *Let \bar{u}_T^h denote the expected utility of the high-quality agent from any contract in Γ_T^h . For any $\tilde{u} \geq 0$ let $\Gamma_T^h(\tilde{u})$ be the set of T -period contracts such that each of them (i) provides the high-quality agent with incentives to exert effort in all T periods and (ii) generates an expected utility of $\bar{u}_T^h + \tilde{u}$ for the high-quality agent. The contract $\tau \in \Gamma_T^h(\tilde{u})$ that minimizes the expected utility of the low-quality agent in all contracts in $\Gamma_T^h(\tilde{u})$ is achieved by amending the contract $\hat{\tau}_T^h$ described in Theorem 1-1 by adding a payment of $\tilde{u}/(p^h)^T$ after a sequence of T successful tasks.*

Proof. (continued) We proceed by constructing the contract $\hat{\tau}_T^l$ described in Theorem 1. The constant r in $\hat{\tau}_T^l$ is chosen so that the low-quality agent is indifferent between choosing $\hat{\tau}_T^l$ and $\hat{\tau}_T^h$. Since the expected utility of the low-quality agent from $\hat{\tau}_T^h$ is positive (one possible strategy for him is to invest only in period T and only if the task is easy), we have $r > c/\pi(l,e)$. Moreover, since a contract that pays $c/\pi(h,d)$ per success belongs to Γ_T^h , Claim 1 implies that the utility of the low-quality agent from this contract is higher than in $\hat{\tau}_T^h$, which in turn implies that $r \leq c/\pi(h,d)$. Therefore, since $c/\pi(l,e) < r \leq c/\pi(h,d)$, the contract $\hat{\tau}_T^l$ generates the right incentives for the low-quality agent.

We complete the proof by showing the limit result. Note that to establish this result it is sufficient to show that the expected utility of the low-quality agent from the contract $\hat{\tau}_T^h$ stays bounded as $T \rightarrow \infty$. To see this, it is enough to show that as $T \rightarrow \infty$, the low-quality agent who adopts $\hat{\tau}_T^h$ exerts effort in a finite number of (last) periods. Denote by K the first period at which the agent begins exerting effort conditional on the task being easy. It is sufficient to show that as $T \rightarrow \infty$, the optimal strategy for the low-quality agent who adopts $\hat{\tau}_T^h$ is to start exerting effort only if $t \geq T - K$, where K remains bounded even if $T \rightarrow \infty$. Assume by way of contradiction that this is not the case and that, instead, $K \rightarrow \infty$ as $T \rightarrow \infty$. Observe however that whenever the agent exerts effort, it affects his utility only if it is followed by an uninterrupted sequence of successes. That is, if the agent succeeds in all remaining K periods (starting from period $(T - K)$ up to the end of the contracting period) he will, according to $\hat{\tau}_T^h$, receive a payment of

$$\frac{c}{\pi(h,d)} A(K) = \frac{c}{\pi(h,d)} \frac{\left(\frac{1}{p^h}\right)^K - 1}{\frac{1}{p^h} - 1}$$

and zero otherwise. Recall that for any strategy of the low-quality agent, the probability of success in K tasks is less than or equal to $(p^l)^K$. Since $p^l < p^h$, the expected utility of the low-quality agent from any strategy in which he starts exerting effort in period $T - K$ is bounded by

$$-c + \frac{c}{\pi(h,d)} \frac{\frac{\pi(l,e)}{p^h} \left(\frac{p^l}{p^h}\right)^{K-1} - (p^l)^{K-1}}{\frac{1}{p^h} - 1}$$

Since

$$\lim_{K \rightarrow \infty} \frac{c}{\pi(h,d)} \frac{\frac{\pi(l,e)}{p^h} \left(\frac{p^l}{p^h}\right)^{K-1} - (p^l)^{K-1}}{\frac{1}{p^h} - 1} = 0,$$

we are done. ■

Observe that when $T = 1$ (the static problem) the optimal contract-pair is actually pooling. That is,

$$\hat{\tau}_1^h(\omega) = \hat{\tau}_1^l(\omega) = \begin{cases} \frac{c}{\pi(h,d)} & \text{if } \omega = \{1\} \\ 0 & \text{otherwise} \end{cases}.$$

When $T > 1$ the contract-pair in which this pooling payment scheme is repeated satisfies *IC* and *EF*. Theorem 1, however, shows that the dynamic structure alleviates the screening problem of the principal and allows us to decrease the low-quality agent's information rents. The optimal contract uses the fact that some histories are more likely to occur when the contract is chosen by the high-quality agent, rather than the low-quality one, for any choice of effort. Increasing

the payments assigned to these histories at the expense of the payments assigned to the other histories makes this contract much less attractive to the low-quality agent.

Remark 4 *The assumption has been that the principal can commit to long-term contracts. If the principal can only commit to short-term one-period contracts, then the optimal contract described in Theorem 1 in which the high-quality agent is compensated only in the last period cannot be implemented simply because in period T the principal has no incentive to pay more than what is necessary to provide incentives to exert effort in T . The high-quality agent, knowing this, will not exert effort in all periods $t < T$. Therefore, when the principal cannot make a long-term commitment, he must pay the high-quality agent $\frac{c}{\pi(h,d)}$ per success every period.*

Non-uniqueness

The optimal contract for the high-quality agent described in Theorem 1, where he is compensated for success in period t only if he succeeds in all subsequent periods, looks rather extreme, especially when T is large. Not only does this contract become very risky, but it also has the unpleasant feature that no payment is guaranteed until T is reached. Of course, these features are irrelevant when the agent and the principal are both risk-neutral and do not discount the future. In the next section discounting is introduced, yet it is comforting to note that for T large enough, the contract described in Theorem 1 is not the unique optimal contract and there are others in which these two features are relaxed dramatically. The following corollary presents such an optimal contract when the extreme form described above is used only during the last stages of the contract. In particular, we construct a contract for the high-quality agent in which during the first $T/2$ periods he is compensated for a success in period t if his subsequent success rate is at least $\alpha \in (0.5, 1]$, while during the last $T/2$ periods the contract is described as in Theorem 1⁵. We start with a definition.

Definition 4 *A proportional contract $\tau_T^h(\alpha)$ is a contract that pays $\tau_T^h(\alpha, \omega_T)$ after a history ω_T , where $\tau_T^h(\alpha, \omega_T)$ consists of two parts. The first part defines the part of the payment that is due to successes in the first half of the contract, i.e., up to period $T/2$, while the second part defines the part of the payment that is due to successes in periods $T/2$ onward. In particular,*

$$\tau_T^h(\alpha, \omega_T) = \sum_{m=1}^{\lceil T/2 \rceil} \mathbf{1}_{\{o(m)=1, n(T)-n(m) \geq \alpha(T-m-1)\}} B(m) + \frac{c}{\pi(h,d)} A(\tilde{K}(\omega_T))$$

where the $\mathbf{1}_{\{D\}}$ in the first part is just an index function, $o(m)$ is the outcome of period m , $n(t)$ is the number of successes from the beginning of the contract

⁵Of course, the choice of $T/2$ is arbitrary and as T gets larger; the part in which the original contract is used can be reduced further.

until period t , and

$$B(m) = \frac{c}{\pi(h,d)} \frac{1}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}}.$$

where

$$\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}$$

is the ex-ante probability that a high-quality agent obtains a success rate of at least α in all future periods, conditional on exerting effort in all $T - m - 1$ remaining periods. Finally, note that the second part of $\tau_T^h(\alpha, \omega_T)$ is just the optimal contract defined in Theorem 1 restricted to periods $T/2$ onward, where $\tilde{K}(\omega_T) = \min\{K(\omega_T), T/2\}$ and $K(\omega_T)$ is the longest uninterrupted sequence of successes in ω_T starting from T and proceeding backward.

Note that this contract mimics the one defined in Theorem 1 for $\alpha = 1$, while for $\alpha = 0$ this contract pays $\frac{c}{\pi(h,d)}$ per success in every period during the first $T/2$ periods, regardless of the outcomes in other periods.

Corollary 2 For all $\alpha \in (0.5, 1)$ for which

$$(p^l)^\alpha < (1 - \alpha)^{1-\alpha} \alpha^\alpha p^h \quad (1)$$

there exists $\hat{T}(\alpha)$ and a constant $r > c/\pi(l,e)$ such that, for all $T > \hat{T}(\alpha)$, the contract-pair $(\tau_T^h(\alpha), \tau_T^l)$, where $\tau_T^h(\alpha)$ is the proportional contract described in Definition 4 and $\tau_T^l(\omega_T) = rn(\omega_T)$, is an optimal contract-pair.

The proof is provided at the end of Appendix A. Notice that since $(1 - \alpha)^{1-\alpha} \alpha^\alpha$ is a monotone increasing function with $\lim_{\alpha \rightarrow 1} (1 - \alpha)^{1-\alpha} \alpha^\alpha = 1$ for any p^l and p^h such that $p^h > p^l$ there exists $\alpha^* \in (1/2, 1)$ such that for any $\alpha \in (\alpha^*, 1)$ the inequality (1) holds.

In Corollary 2 it is shown that as long as T is large enough, there are other optimal contracts that are not as extreme as the contract presented in Theorem 1. In this class of contracts payments are not postponed all the way to the last period and the contracts are less risky than the one described in Theorem 1. The next section shows that introducing discounting pins down uniquely the equilibrium payments. Moreover, since in the case of discounting delay of payments is costly, this contract does not possess the described extreme features.

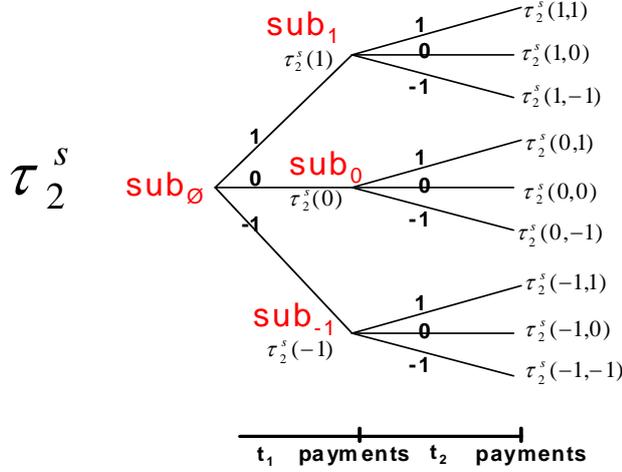


Figure 3: Two-period contract for an impatient agent of type s , which specifies for any possible histories ω_1 and ω_2 the payment to this agent, $\tau_2^s(\omega_k)$.

6 Discounting Future Payments

Up until now the assumption has been that the agents and the principal do not discount the future (or alternatively, have the same discount factor). We shall now relax this assumption to cover the case where the agents are more impatient than the principal. Unlike in the case of no discounting, where timing of payments is irrelevant and it can be assumed without loss of generality that all payments are postponed until the end of the contract, in the case here the exact timing of payments is relevant and is part of the contract, as depicted in the figure above for a two-period contract. With some abuse of notation, we denote by $\tau_T^s(\omega_k)$ the payment in the contract τ_T^s of length T after history ω_k of length k .

For ease of notation, we assume that the principal does not discount the future and hence he has the same preferences as were assumed in the previous sections, while the expected utility of the agent is

$$\sum_{t=1}^T \delta^t (m_t - C_t)$$

where $\delta \in (0, 1)$ is the agent's discount factor and m_t and C_t are his payment and effort level in period t .

Recall that the main insight from the no-discounting case was that the expected payment to the low-quality agent is reduced by employing the “*mechanism of postponing payments*” according to which the contract of the high-quality agent rewards successes only after a delay and only *conditional* on successes in *all* periods in between. While this insight is carried over to the case here, now it is costly because of the need to compensate the high-quality agent for the delay in payments. As a result the mechanism of postponing payments is now used selectively, i.e., not necessarily after every history, and not necessarily all the way to the end. In other words, when the agent is impatient the principal faces a trade-off; conditional on the agent being a low-quality one, postponing payment is beneficial, while conditional on him being high-quality one, it is costly.

When the agent is known to be a high-quality agent, the *unique* optimal contract is the one that pays for success with no delay. More precisely, the unique optimal contract pays $c/\pi_{(h,d)}$ in period t if (and only if) the task at t was successful. We refer to this contract as the *base-line contract*.

Definition 5 A T -period contract τ_T is called the *base-line contract* if

$$\tau_T(\omega_k) = \begin{cases} c/\pi_{(h,d)} & \text{if } o(k) = 1 \\ 0 & \text{otherwise} \end{cases}$$

where $o(k)$ is the outcome in period k .

Note that the expected utility of the high-quality agent from the base-line contract of length T is

$$\sum_{t=1}^T \delta^t \left(p^h \frac{c}{\pi_{(h,d)}} - c \right) = c \left(\frac{p^h}{\pi_{(h,d)}} - 1 \right) (1 - \delta^T) \frac{\delta}{1 - \delta}.$$

We denote by $C_t^s(\hat{\tau}_T^h, \omega_{t-1}, p_t)$ the action taken according to the optimal strategy of an agent of type s in period t when the contract is $\hat{\tau}_T^h$, the history is ω_{t-1} , and at t the task's type is p_t . Indeed, it is easy to verify that a contract-pair in which the base-line contract is offered to both types of agents is admissible. Note, however, that this contract-pair is not optimal (provided that the discount factor is not too small), because it pays, in expectation, too much to the low-quality agent. The following proposition is an extension of the arguments developed in the no-discounting case and its proof is provided in Appendix C. It asserts that the high-quality agent should be indifferent between the optimal contract $\hat{\tau}_T^h$ and the base-line contract. Indeed, the optimal contract for the high-quality type $\hat{\tau}_T^h$ is constructed from the base-line contract by employing the mechanism of postponing payments selectively (not after every history, and not necessarily until the end of the contract). The exact manner in which this is done is discussed in the partial characterization that follows. The low-quality agent should be indifferent between his contract $\hat{\tau}_T^l$ and the contract for the high-quality agent,

$\hat{\tau}_T^h$. However, discounting implies that all the expected utility of the low-quality agent is moved up-front and paid at the moment of signing the contract. Recall that $u^l(\hat{\tau}, \text{sub}_\emptyset)$ is the expected (discounted) utility of the low-quality agent from the contract $\hat{\tau}$.

Proposition 1 *When the agent is impatient, then the optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is such that: (i) the high-quality agent is indifferent between his contract $\hat{\tau}_T^h$ and the base-line contract, and (ii) the contract $\hat{\tau}_T^l$ pays an up-front amount $M(\hat{\tau}_T^h) = u^l(\hat{\tau}_T^h, \text{sub}_\emptyset)$ and then $c/\pi_{(l,e)}$ at t if the task at t is successful.*

When the agent does not discount future payments the set of optimal contracts may not be unique. In particular, if it is enough to eliminate incentives for the low-quality agent to exert effort in period t when the payments in period t are postponed to period $t' > t$ (and paid conditional on successes in between), then postponing payments even further to $t'' > t'$ is also optimal. Yet, this is not the case when the agent is impatient. The longer a payment is postponed, the more costly it is. It follows that when a payment is postponed, it may not be postponed all the way to period T . Indeed, it is never optimal to postpone a payment longer than is needed to eliminate the incentives of the low-quality agent to exert effort in t if he accepts the contract. Postponing payments further is of no benefit on the one hand, and is costly on the other.

Indeed, we shall now show that introducing impatience narrows down the set of equilibrium payments to a unique one. More precisely, the following proposition establishes the uniqueness of the optimal contract under the assumption that the payments to the high-quality agent are zero after all histories that include a refusal to accept a task at least once. Since any optimal contract-pair is admissible, a high-quality agent will never refuse a task in equilibrium. Hence, the proposition implies uniqueness of payments along the equilibrium path. In particular, there exist other optimal contract-pairs in which off-equilibrium payments, i.e., payments after histories containing refusals, are different.

Proposition 2 *When the agent discounts the future at some $\delta \in (0, 1)$, then the optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is generically unique among all contracts in which a refusal to accept a task by the high-quality agent in period $t \in \{1, \dots, T\}$ leads to a payment of zero in all periods $t' \geq t$.*

Proof. To prove this, we need to recall first from Proposition 1 that if the pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal, then the high-quality agent is indifferent between his contract $\hat{\tau}_T^h$ and the base-line contract. Recall, too, that the contract for the low-quality agent pays an up-front lump sum M and subsequently for all $t \in \{1, \dots, T\}$, pays $c/\pi_{(l,e)}$ at t if and only if the task at t is accomplished successfully. When $M = U_l(\hat{\tau}_T^h)$, the expected discounted utility of the low-quality agent from

the contract is $\hat{\tau}_T^h$. Also recall that $m^s(\tau_T^i, sub_\emptyset)$ denotes the expected payment to an agent of quality s in contract τ_T^i .

The principal is indifferent between the two contract-pairs $(\bar{\tau}_T^h, \bar{\tau}_T^l)$ and $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ if and only if both pairs are admissible and, moreover,

$$\mu m^l(\bar{\tau}_T^l, sub_\emptyset) + (1 - \mu) m^h(\bar{\tau}_T^h, sub_\emptyset) = \mu m^l(\hat{\tau}_T^l, sub_\emptyset) + (1 - \mu) m^h(\hat{\tau}_T^h, sub_\emptyset).$$

Recall from Proposition 1 that the expected payment to the low-quality agent in the pairs $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ and $(\bar{\tau}_T^h, \bar{\tau}_T^l)$ is

$$T(1 - q) \pi_{(l,e)} \frac{c}{\pi_{(l,e)}} + u^l(\hat{\tau}_T^h, sub_\emptyset) \quad \text{and} \quad T(1 - q) \pi_{(l,e)} \frac{c}{\pi_{(l,e)}} + u^l(\bar{\tau}_T^h, sub_\emptyset).$$

We conclude that the principal is indifferent between the two contract-pairs $(\bar{\tau}_T^h, \bar{\tau}_T^l)$ and $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ if and only if both pairs are admissible and, moreover,

$$\mu u^l(\bar{\tau}_T^h, sub_\emptyset) + (1 - \mu) m^h(\bar{\tau}_T^h, sub_\emptyset) = \mu u^l(\hat{\tau}_T^h, sub_\emptyset) + (1 - \mu) m^h(\hat{\tau}_T^h, sub_\emptyset).$$

To establish generic uniqueness of the optimal contract-pair, we now show that if two admissible contract-pairs $(\bar{\tau}_T^h, \bar{\tau}_T^l)$ and $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ are optimal, then there is only a measure zero of prior beliefs μ on which the principal is indifferent between the two contract-pairs.

From Proposition 1 it follows that the optimal contract to the high-quality agent generates an expected utility of

$$\left(-c + p^h \frac{c}{\pi_{(h,d)}}\right) \sum_{t=1}^T \delta^t.$$

Given all this, it is not difficult to show that there is a finite number of contracts that can be considered optimal. Recall that if the principal postpones some of the payments in the contract of the high-quality agent, he pays them conditional on observing an uninterrupted sequence of successes until the payment periods. Therefore the payment to the high-quality agent consists of two parts: the expected payment (ignoring delay) $T p^h \frac{c}{\pi_{(h,d)}}$, and the additional payment for compensation for delays. The additional expected payment to the high-quality agent if, following a history of l successes and m failures, a payment of ε is postponed for k periods and the amount of $\varepsilon / (p^h \delta)^k$ is paid for k successes, is given by

$$(p^h)^l (1 - p^h)^m \varepsilon \left(\frac{1}{\delta^k} - 1\right).$$

For a given strategy of the low-quality agent the decrease in his payment, if he adopts the contract for the high-quality agent, is given by

$$\Pr(l, m) (\pi_{(l,e)})^l \delta^{m+l+1} \left[(1 - q) (\pi_{(l,e)}) \varepsilon - (1 - q)^{k+1} (\pi_{(l,e)})^{k+1} \frac{\varepsilon}{(p^h)^k} \right]$$

where $\Pr(l, m)$ is the probability that the low-quality agent will reach this history. If the low-quality agent stops exerting effort even if an easy agent arrives, the decrease in his utility is

$$\Pr(l, m) (\pi_{(l,e)})^l \delta^{m+l+1} (1 - q) (\pi_{(l,e)}) \varepsilon.$$

Thus, for a given strategy of the low-quality agent, the change in the expected payment to the high-quality agent and the change in the expected utility of the low-quality agent are linear in the postponed amount. Therefore, if the principal finds it optimal to postpone some amount of money after some history, he will postpone either the maximal amount or the amount that changes the strategy of the low-quality agent. Since the low-quality agent has a finite number of strategies, and there is a finite number of potential postponements available for the principal, there is a finite number of priors μ that satisfy the indifference of the principal. ■

Having proved uniqueness, we next provide an example as well as characterization results of the optimal contract-pair that illustrate the type of trade-off faced by the principal when the agent discounts the future. The mechanism of postponing payments is more effective (and less costly) if employed after histories that are more likely to occur if the agent is of low quality. These histories are more likely to be found later in the life of the contract than earlier. However, because the agent is impatient, the mechanism is more effective the earlier it is employed. The optimal $\hat{\tau}_T^h$ balances between these two forces that are, to some extent, in conflict.

Below we develop a three-period example in which the mechanism of postponing payments is employed only after histories that are more likely to occur when the agent is of low quality. We then show that if the discount factor is smaller than the probability of failure of the high-quality agent then it pays to employ the mechanism sooner rather than later. The intuition for this result is that the lower a high-quality agent's ex-ante probability of success is, i.e., $p^h = q\pi_{(h,d)} + (1 - q)\pi_{(h,l)}$, the more difficult it becomes to use failure as a criterion for identifying which histories are more likely to occur when the agent is of low-quality.

The final result of this section is a limiting result that corresponds to the case as $\delta \rightarrow 1$. It is shown that as the cost of postponing payments is decreasing, payments for success in period t are postponed as long as the incentives of the low-quality agent to exert effort in t remain positive. For simplicity, we shall assume in the sequel that $\pi_{(l,d)}$ is small enough so that it is never optimal for a low-quality agent to exert effort when the task is difficult.⁶ Relaxing this

⁶More precisely, we assume that

$$-c + \pi_{(l,d)} \frac{c}{\pi_{(h,d)}} A(T) < 0$$

assumption would add no qualitative insight and would complicate the discussion considerably.

6.1 Three-period Example

In this section we construct an example of a three-period optimal contract in which, according to $\hat{\tau}_T^h$, if in period one the task is successful, then $c/\pi_{(h,d)}$ is paid in every subsequent period in which the respective task is successful. If, however, the first task ended in failure, then the payment for success in period two is postponed and paid only in period three and only if the third task is also successful. The contract $\hat{\tau}_3^l$, on the other hand, pays an amount M up-front and then $c/\pi_{(l,e)}$ at t if (and only if) the task at t was successful. The lump sum M is the minimal amount required to satisfy the *IC* constraint and it is the expected utility of the low-quality agent should he accept contract $\hat{\tau}_T^h$. The intuition behind this equilibrium is that postponing payments in $\hat{\tau}_T^h$ is costly to the principal because the high-quality agent must be compensated for the delay in payments; indeed, the sole purpose of postponement of payment is to make the contract less attractive to the low-quality agent. A failure in the first period is more likely to occur when the agent is of low-quality, and as the example demonstrates, there are parameters under which it pays to wait for such histories and to postpone payments, even though waiting renders this mechanism less effective due to the discounting of future utilities.

Specifically, we shall identify a set of parameters for which the contract for the high-quality agent, depicted in the figure below (where histories containing zeroes are ignored), is optimal.

To establish the optimality of $\hat{\tau}_3^h$, note first that regardless of the history a success in period three is rewarded by $c/\pi_{(h,d)}$, which is the minimal amount required to provide the needed incentives to exert effort when the task is difficult. Proceeding backward to period two, consider first the optimal policy after a failure in the first period. For the high-quality agent to exert effort when the task is difficult he must be compensated for success with the equivalent of the $c/\pi_{(h,d)}$ paid in this period. Note that postponing a payment of ε for one period requires increasing the next period payment by ε/δ . Therefore, following a failure in period one and a success in period two, the expected cost of postponing a payment of ε to period three and paying conditional on success in that period is

$$(1 - \mu) \left(1 - p^h\right) \varepsilon \frac{1 - \delta}{\delta}$$

where $(1 - \mu)$ is the prior that the agent is of high-quality, and $(1 - p^h)$ is the

which is equivalent to

$$\pi_{(l,d)} < \pi_{(h,d)} \frac{\frac{1}{p^h} - 1}{\left(\frac{1}{p^h}\right)^T - 1}.$$

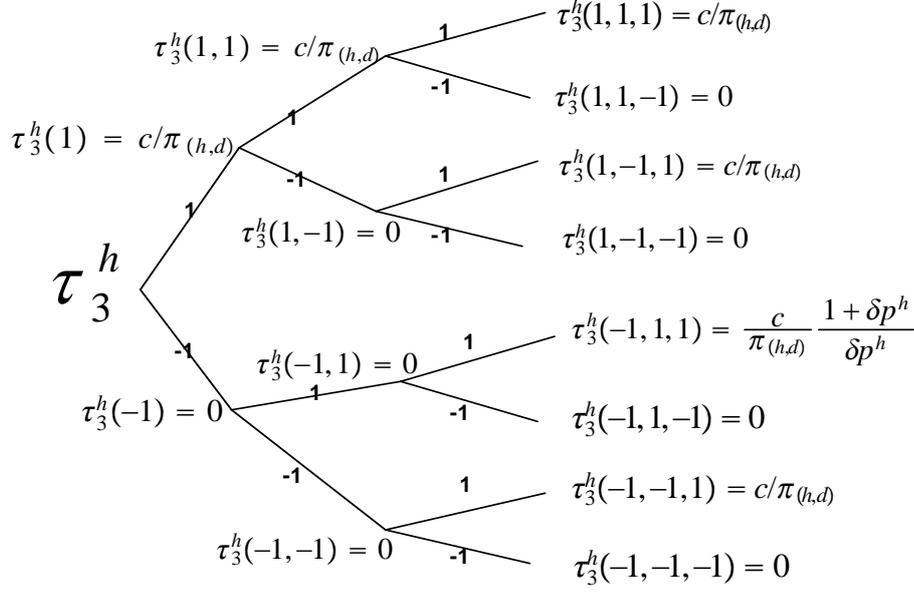


Figure 4: Three-period contract for a high-quality agent, which specifies for any possible t history ω_t the payment to this agent, $\tau_3^h(\omega_t)$.

probability of him failing in the first period (i.e., the probability of reaching this history when the agent's quality is high). The benefit of this delay is the reduction in the utility of the low-quality agent (probability μ). The utility reduction of the low-quality agent due to this change given that he failed in the first period is

$$\delta^2(1-q)\pi_{(l,e)}\varepsilon - \delta^3(1-q)^2(\pi_{(l,e)})^2\frac{\varepsilon}{\delta p^h}$$

where $\varepsilon/(\delta p^h)$ is the increased payment due to delayed payment of ε . Since the low-quality agent reaches this history with probability $q + (1-q)(1 - \pi_{(l,e)})$, the expected utility reduction of the low-quality agent due to this change is

$$\mu\delta^2(q + (1-q)(1 - \pi_{(l,e)}))(1-q)\pi_{(l,e)}\left(\varepsilon - \delta\frac{(1-q)\pi_{(l,e)}}{p^h\delta}\varepsilon\right).$$

Therefore, after a failure in period one, postponing a payment for success from period two to period three and paying conditional on success in period three is beneficial if

$$\mu\delta^2(q + (1-q)(1 - \pi_{(l,e)}))(1-q)\pi_{(l,e)}\left(1 - \frac{(1-q)\pi_{(l,e)}}{p^h}\right) > (1-\mu)(1-p^h)\frac{1-\delta}{\delta}. \quad (2)$$

Similarly, the cost of postponing a payment of ε from the first period to the second is

$$(1-\mu)\varepsilon\frac{1-\delta}{\delta}$$

and the benefit is

$$\mu\delta(1-q)\pi_{(l,e)}\left(\varepsilon - \delta\frac{(1-q)\pi_{(l,e)}}{p^h\delta}\varepsilon\right).$$

Therefore, it is profitable to postpone payments from the first period to the second if

$$\mu\delta(1-q)\pi_{(l,e)}\left(1 - \frac{(1-q)\pi_{(l,e)}}{p^h}\right) > (1-\mu)\frac{1-\delta}{\delta}. \quad (3)$$

We conclude from (2) and (3) that if

$$\mu\delta^2(q + (1-q)(1-\pi_{(l,e)}))(1-q)\pi_{(l,e)}\left(1 - \frac{(1-q)\pi_{(l,e)}}{p^h}\right) - (1-\mu)(1-p^h)\frac{1-\delta}{\delta} > 0$$

and

$$\mu\delta(1-q)\pi_{(l,e)}\left(1 - \frac{(1-q)\pi_{(l,e)}}{p^h}\right) - (1-\mu)\frac{1-\delta}{\delta} \leq 0$$

then success in the first period results in no postponement of payment for success in subsequent periods, but failure in the first period results in the delaying of payment in the second period.

In Appendix C we show that these two inequalities can coexist in the sense that there exists a range of parameters in which both inequalities hold.

Remark 5 *In general, the profitability of postponing payments is decreasing with the length of the delay because the benefits are decreasing and the costs are increasing with the length of delay. In particular, in the example above, if it is not profitable to postpone payments from the first period to the second, then it is not profitable to postpone payments from the first period to the third. Moreover, postponing payment at some period affects incentives in that period as well as in previous periods.*

Remark 6 *It is apparent from the calculation in this example that the benefits and the costs of postponing payments from period t to period $t' > t$ are linear in the amount being postponed as long as it has no effect on the incentives of the low-quality agent to exert effort in period $t'' \leq t$. In the example above, it implies that if it is profitable to transfer, say, some amount ε from period one to period two, then it is optimal to transfer $\min\{c/\pi_{(h,d)}, k\}$ where k is the transfer that makes the low-quality agent indifferent between exerting and not exerting effort on an easy task in period one, given that success is compensated by $c/\pi_{(h,d)} - k$ in this period, and conditional on success in period two an additional payment of $k/\delta p^h$ is paid.*

6.2 Intermediate Discount Factor

As we explained earlier, in postponing payments the principal faces a trade-off. Later in the life of the contract it is easier to identify histories that are more

likely to occur when the agent is of low-quality, and hence these histories are good times to postpone payments. But because agent discounts future payments, postponing payments later in the life of the contract is less effective. So whether payments are postponed early or late in the contract depends on how easy it is to identify “good” histories and on how impatient the agent is. Indeed, the following proposition, whose proof is provided in Appendix C, establishes that if the discount factor is small relative to the ex-ante probability of the failure of the high-quality agent, then payments are postponed starting from period one. The intuition for this result is that the lower a high-quality agent’s ex-ante probability of success is, i.e., $p^h = q\pi_{(h,d)} + (1 - q)\pi_{(h,l)}$, the more difficult it becomes to use failure as a criterion for identifying which histories are more likely to occur when the agent is of low-quality.

Proposition 3 *If $\delta < 1 - p^h$, then it is optimal to postpone payments in the initial periods of the contract. That is, for any history ω_t there exists k_{ω_t} such that for all $t < k_{\omega_t}$ we have $\hat{\tau}_T^h(\omega_t) = 0$ while $\hat{\tau}_T^h(\omega_t) > 0$ whenever $t \geq k_{\omega_t}$ and $o(t) = 1$.*

6.3 High Discount Factor and Limit Result

When, as was assumed in previous sections, agents do not discount future payments, postponing payments is costless and hence a contract for the high-quality agent in which all payments are postponed to the end is optimal. Yet, as was shown above, even in the case of no discounting, this may not be the unique optimal contract. In particular, when payment is postponed until after every history but only as long as the low-quality agent finds it optimal to exert effort when the task is easy, the resulting contract-pair is also optimal. We now show that when the agents are impatient but the discount factor is close enough to one, this contract is still optimal.

We start with a simple lemma that asserts that for a discount factor high enough, if the optimal contract $\hat{\tau}_T^h$, is such that the low-quality agent, if he is to adopt this contract, exerts effort in period $t \in \{1, \dots, T\}$ when the task is easy, then all payments for successes at t are postponed to future periods.

Lemma 4 *There exists $\delta^* < 1$ such that for all $\delta \in [\delta^*, 1]$ and for all histories ω_{t-1} that can be reached with positive probability when the low-quality agent signs into $\hat{\tau}_T^h$ and $C_t^l(\hat{\tau}_T^h, \omega_{t-1}, e) = c$, we have $\hat{\tau}_T^h(\omega_{t-1}, 1) = 0$.*

Proof. To establish the lemma it is enough to show that for all histories ω_{t-1} that can be reached in positive probability by the low-quality agent when he sign into $\hat{\tau}_T^h$ and $C_t^l(\hat{\tau}_T^h, \omega_{t-1}, e) = c$, if $\hat{\tau}_T^h(\omega_{t-1}, 1) > 0$ then, for δ high enough it is beneficial for the principal to postpone the payment $\hat{\tau}_T^h(\omega_{t-1}, 1)$. Observe first

that in a history that contains l successes and m failures, the cost of postponing the amount ε for one period and paying conditional on success is

$$(1 - \mu) \left(p^h\right)^l \left(1 - p^h\right)^m \varepsilon \frac{1 - \delta}{\delta} \xrightarrow{\delta \rightarrow 1} 0.$$

Next recall that since the amount ε is paid only after an additional success, the benefit of postponing payment of it is the reduction in the expected utility of the low-quality agent if he accepts the contract $\hat{\tau}_T^h$.

If the low-quality agent exerts effort after history ω_{t-1} and exerts effort also in the next period if an easy task arrives, then the benefit of the change for the principal is

$$\begin{aligned} & \mu \left(q + (1 - q) \left(1 - \pi_{(l,e)}\right)\right)^r \left(1 - q\right)^l \left(\pi_{(l,e)}\right)^l \delta^{m+l+1} \times \\ & \times \left[\left(1 - q\right) \left(\pi_{(l,e)}\right) \varepsilon - \left(1 - q\right)^2 \left(\pi_{(l,e)}\right)^2 \frac{\varepsilon}{\left(p^h\right)} \right] \xrightarrow{\delta \rightarrow 1} a > 0 \end{aligned}$$

where $m - r$ is the number of periods along the history ω_{t-1} where the low-quality agent does not exert effort even if an easy task arrives. If the low-quality agent does not exert effort in period $t + 1$ following the history ω_{t-1} and success at t even when an easy task arrives, then the benefit of the the change for the principal is

$$\mu \left(q + (1 - q) \left(1 - \pi_{(l,e)}\right)\right)^r \left(1 - q\right)^l \left(\pi_{(l,e)}\right)^l \delta^{m+l+1} \left(1 - q\right) \left(\pi_{(l,e)}\right) \varepsilon \xrightarrow{\delta \rightarrow 1} b > 0.$$

It follows that if $\hat{\tau}_T^h$ is such that if the low-quality agent adopts it, then he exerts effort at t whenever the task at t is easy, then the payment for success in t according to $\hat{\tau}_T^h$ must be zero. ■

Since payment in t is positive only when the task in period t is accomplished successfully, the previous lemma implies that if the low-quality agent exerts effort in period t (when the task is easy) and succeeds, then he must also exert effort in period $t + 1$ when the task is easy.

The following proposition establishes that for a discount factor close enough to one, the optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is such that if the low-quality agent deviates, adopts the contract $\hat{\tau}_T^h$, and then chooses a best response, then there exists some $k \leq T$ such that for all $t < k$ the agent does not exert effort, but thereafter, for $t \geq k$, he exerts effort whenever the task is easy. It follows from the lemma that the contract $\hat{\tau}_T^h$ can be split into two time intervals. In the first $k - 1$ periods, a success in t is compensated at some $t' \leq T$, while compensation for success in period $t \geq k$ is paid at the end of the contract, i.e., in period T .

Proposition 4 *There exists $\delta^* < 1$ such that for any $\delta \in [\delta^*, 1]$ there exists a period k such that, in the optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$, $C_t^l(\hat{\tau}_T^h, \omega_{t-1}, e) = 0$ if $t < k$ and $o(m) = -1$ for any $m < t$ and $C_t^l(\hat{\tau}_T^h, \omega_{t-1}, e) = c$ if $t \geq k$.*

Proof. To establish the proposition it is enough to show that if the low-quality agent exerts effort in some period t when an easy task arrives, then he also exerts effort in period $t+1$ when an easy task arrives regardless of the outcome in period t . Denote by k the first period where the low-quality agent, if he accepts $\hat{\tau}_T^h$, exerts effort when an easy task arrives. It follows that until period k the agent experiences a sequence of failures.

Assume, by way of contradiction, that there exists some period $l \geq k$ such that the agent exerts effort in l if an easy task arrives, but does not exert effort in period $l+1$ even if an easy task arrives. Observe first that if the agent succeeds in period l , the previous lemma implies that the payment for this success is postponed. Thus, compensation for exerting effort in l is paid only if the task in $l+1$ is successful, which implies that the agent must exert effort in period $l+1$ after success in l , when the task in $l+1$ is an easy task, or otherwise he shouldn't exert effort in l . In fact, it implies that after an uninterrupted sequence of successes starting from period l he should exert effort whenever the task is easy. Hence, it remains to show that it is optimal for the low-quality agent, if he adopts $\hat{\tau}_T^h$, to exert effort in period $l+1$ even after he fails in period l . Assume next that after failure in period l the low-quality agent does not exert effort in period $l+1$ when the task is easy. Note, however, that there must exist a period $m > l$ such that the low-quality agent exerts effort in m whenever the task is easy (and continues doing so as long as he succeeds), but he does not exert effort on an easy task in period $m-1$. To see this, note that it is always optimal for the low-quality agent to exert effort on an easy task in period T .

Recall that our agent exerts effort in period k and, if successful, is compensated by $\frac{c}{\delta^{T-k}\pi_{(h,d)}(p^h)^{T-k}}$ in period T with probability $(1-q)^{T-k}(\pi_{(l,e)})^{T-k}$. It follows that

$$\pi_{(l,e)} \frac{c}{\pi_{(h,d)}(p^h)^{T-k}} (1-q)^{T-k} (\pi_{(l,e)})^{T-k} - c \geq 0.$$

However, since $k < m-1$ and $p^h > (1-q)\pi_{(l,e)}$, we have that

$$\begin{aligned} -c + \frac{c}{\pi_{(h,d)}(p^h)^{T-m+1}} \pi_{(l,e)} (1-q)^{T-m+1} (\pi_{(l,e)})^{T-m+1} &> \\ -c + \frac{c}{\pi_{(h,d)}(p^h)^{T-k}} \pi_{(l,e)} (1-q)^{T-k} (\pi_{(l,e)})^{T-k} &\geq 0. \end{aligned}$$

We conclude that if the low-quality agent exerted effort in period k when an easy task arrived, he should exert effort in period $m-1$ as well, a contradiction. ■

An immediate consequence of the proposition above is that when $\delta > \delta^*$, the optimal contract $\hat{\tau}_T^h$ possesses similar features as in the case of no discounting characterized in previous sections. That is, payments are postponed as long as the incentives for the low-quality agent to exert effort are positive. Moreover, in

the history that contains $k - 1$ failures, all the payments are postponed to the last period, where k is defined as the minimal K for which

$$\pi_{(l,e)} \frac{c}{\pi_{(h,d)} (p^h)^{T-K}} (1-q)^{T-K} (\pi_{(l,e)})^{T-K} - c \geq 0. \quad (4)$$

Since $(1-q)^{T-K} (\pi_{(l,e)}) / p^h < 1$, (4) implies that there exists $P < \infty$ such that $\lim_{T \rightarrow \infty} M(\hat{\tau}_T^h) = u^l(\hat{\tau}_T^h, \text{sub}_{\emptyset}) < \mathbf{P}$.

7 Conclusion and Extensions

Risk Aversion

Throughout this paper, we have assumed that the principal and the agent are both risk-neutral. The assumption of risk neutrality simplifies the analysis considerably, but it is also important to note that the analysis of the case where the agent is more risk-averse than the principal is qualitatively similar to the one conducted above for the discounting case. In particular, the same important trade-off identified for the case where the agents are impatient is present here. Postponing payments in the contract designed for the high-quality agent is the way to reduce expected payments to the low-quality agent. However, conditional on the agent being of high-quality, postponing payment is costly because the agent, being risk averse, must be compensated for the extra risk. It follows that adding risk aversion to the model with discounting has largely the same effect as reducing the discount factor.

Pooling

Throughout this paper, we have assumed that $\pi_{(l,e)} > \pi_{(h,d)}$ and relegated to Appendix B the rather similar analysis of the case where $\pi_{(l,e)} \leq \pi_{(h,d)}$ (hereafter cases (i) and (ii) respectively). It is, however, worth describing the main result of case (ii) and providing some intuition for the sharp differences between the two cases and in particular for the fact that in case (ii) the optimal contract-pair is pooling, in the sense that regardless of the agent's type, he is paid a fixed amount $c/\pi_{(h,d)}$ per success, as is shown in Theorem 2 in Appendix B. Recall that in case (i) the contract that is offered to the high-quality agent is the one that is offered to him when his quality is observable, and it is the low-quality agent who enjoys some information rent. As we establish in Appendix B, case (ii) is different. The first difference is that in this case the high-quality agent enjoys the information rent from his privately known type (in addition he has information rents from privately observed types of tasks), and the second difference is that now the repeated nature of the relation is not helpful.

To gain some insight into the differences between the two cases, assume first that in every $t \in \{1, \dots, T\}$ the principal is constrained to propose a short-term one-period contract only. It is easy to see that in both cases the only contract-pair that satisfies *EF* and *IC* is a pooling one in which regardless of the agent's

quality he is paid a fixed amount per success: $c/\pi_{(h,d)}$ in case (i), and $c/\pi_{(l,e)}$ in case (ii). Adopting the terminology developed above for long-term contracts, and letting $n(\omega_T)$ denote the number of successes in ω_T , these contracts can be written as

$$\text{Case (i): } \bar{\tau}_T^h(\omega_T) = \bar{\tau}_T^l(\omega_T) = \frac{c \cdot n(\omega_T)}{\pi_{(h,d)}}$$

and

$$\text{Case (ii): } \tilde{\tau}_T^h(\omega_T) = \tilde{\tau}_T^l(\omega_T) = \frac{c \cdot n(\omega_T)}{\pi_{(l,e)}}.$$

Note that $\bar{\tau}_T^h(\omega_T) \in \Gamma_T^h$ and $\tilde{\tau}_T^l(\omega_T) \in \Gamma_T^l$, which implies that in case (i) the expected utility of the high-quality agent is at its lower bound (at its level when his quality is observable), while the expected utility of the low-quality agent is above its lower bound. The reverse, however, is true in case (ii), where the expected utility of the low-quality agent is at its lower bound.

As we show in the analysis of case (i) above, the important effect of long-term contracts is the availability of other contracts in Γ_T^h which, from the low-quality agent's point of view, are worse than $\bar{\tau}_T^h(\omega_T)$. The optimal contract-pair exploits this by assigning the high-quality agent the contract in Γ_T^h that is the least attractive to the low-quality agent. This enables the principal to assign to the low-quality agent a contract that yields a lower expected payment than the repeated short-term contract $\bar{\tau}_T^l(\omega_T)$. In case (ii) it is the high-quality agent who receives a level of expected utility above his lower bound. But unlike in case (i) where the short-term contract $\bar{\tau}_T^h(\omega_T)$ was, for the low-quality agent, the best in Γ_T^h , now the short-term contract $\tilde{\tau}_T^l(\omega_T)$ is the worst in Γ_T^l for the high-quality agent. It follows that in case (ii) the short-term contract is the best the principal can achieve when the low-quality agent is already at his IR , because any other contract τ_l^T that satisfies EF would yield the high-quality agent an even higher expected utility.

8 Appendix A: Proofs for the Separating Contract Case

Proof of Lemma 1:

Property 1: Assume that this property is false. $\tau_K^h \in \Gamma_K^h$, so τ_K^h provides sufficient incentives in all subgames, and in particular in sub_{-1} (the subgame following a failure in period one). Consider replacing τ_K^h with $\tilde{\tau}_K^h$, where $\tilde{\tau}_K^h$ is obtained by amending the contract τ_K^h and replacing the payments in all histories that belong to sub_{-1} by payments in one of the optimal $K-1$ -period contracts in Γ_{K-1}^h . That is, $\tilde{\tau}_K^h(-1, \omega_{K-1}) = \tilde{\tau}_{K-1}^h(\omega_{K-1})$. Clearly, the proposed change does not affect incentives in sub_1 . Also, because an optimal $K-1$ -period contract provides incentives in the $K-1$ -period problem, incentives are provided in sub_{-1} .

Since the new payment scheme in sub_{-1} is a contract in Γ_{K-1}^h , it minimizes expected payments in all schemes that provide incentives. That is,

$$m^h(\tilde{\tau}_K^h, sub_{-1}) < m^h(\tau_K^h, sub_{-1}) \quad (5)$$

because otherwise the τ_K payments restricted to sub_{-1} is a contract from Γ_{K-1}^h . Since incentives are provided by τ_K^h to exert effort in period one on a task with a major problem, it must be the case that

$$u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) \geq \frac{c}{\pi(h,d)}$$

and in particular

$$m^h(\tau_K^h, sub_1) - m^h(\tau_K^h, sub_{-1}) \geq \frac{c}{\pi(h,d)}. \quad (6)$$

This together with (5) implies that

$$u^h(\tau_K^h, sub_1) - u^h(\tilde{\tau}_K^h, sub_{-1}) > \frac{c}{\pi(h,d)}$$

which guarantees that incentives to exert effort in period one are preserved and in general incentives are provided in the revised K -period contract. Finally, note that since this revision decreases the expected utility of the agent after failure in the first period and keeps the expected utility after success in the first period, it decreases the expected payment, which is in contradiction to the claimed optimality of the original contract. ■

Property 2: Assume by way of contradiction that $\tau_K^h \in \Gamma_K^h$, but that

$$u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) \neq \frac{c}{\pi(h,d)}$$

and recall that since an optimal contract provides incentives to exert effort, it must be the case that

$$u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) > \frac{c}{\pi(h,d)}.$$

Therefore, we revise τ_K^h to $\tilde{\tau}_K^h$ so that $\tilde{\tau}_K^h(1, \omega_{K-1}) = \tau_K^h(-1, \omega_{K-1}) + \frac{c}{\pi(h,d)}$. Note that incentives to exert effort in $\tilde{\tau}_K^h$ are kept and that $u^h(\tau_K^h, sub_1)$ is now decreased to $u^h(\tau_K^h, sub_{-1}) + \frac{c}{\pi(h,d)}$ so that expected payment is decreased, which is in contradiction to τ_K^h being optimal. ■

Property 3: The proof is done by induction. Observe first that for $T = 1$ we have $\tau_1^h(1) = \frac{c}{\pi(h,d)}$ and $\tau_1^h(-1) = 0$, which implies that $m^h(\tau_1^h, sub_\emptyset) = p^h \frac{c}{\pi(h,d)}$. Next assume that if $\tau_{K-1}^h \in \Gamma_{K-1}^h$, then $m^h(\tau_{K-1}^h, sub_\emptyset) = (K-1)p^h \frac{c}{\pi(h,d)}$. From Properties 1 and 2 it follows that

$$\begin{aligned} m^h(\tau_K^h, sub_\emptyset) &= (1-p^h)(K-1)p^h \frac{c}{\pi(h,d)} + p^h[(K-1)p^h \frac{c}{\pi(h,d)} + \frac{c}{\pi(h,d)}] = \\ &= Kp^h \frac{c}{\pi(h,d)}, \end{aligned}$$

which is the desired final step of the proof. ■

Note that taken together, Properties 1, 2, and 3 imply that in any optimal contract we have

$$m^h \left(\tau_K^h, sub_1 \right) = (K - 1) p^h c / \pi_{(h,d)} + c / \pi_{(h,d)}.$$

Denote by \bar{u}_T^h the expected utility of the high-quality agent from any contract in Γ_T^h . That is,

$$\bar{u}_T^h = T p^h \frac{c}{\pi_{(h,d)}} - T c.$$

Property 4: This property is an immediate consequence of the following claim:

Claim 2 *If a T -period contract τ_T^h satisfies EF and IR and generates expected utility $u > \bar{u}_T^h$, then for any $\tilde{u} \in [\bar{u}_T^h, u)$ there exists another T -period contract $\tilde{\tau}_T^h$ that satisfies EF and IR and generates an expected utility of \tilde{u} . Moreover, for all $\omega_T \in \Omega_T$, $\tilde{\tau}_T^h(\omega_T) \leq \tau_T^h(\omega_T)$ with at least one strict inequality.*

Proof. The proof is done by inducting on the contract's length, T . Assume that $T = 1$ and observe that since τ_1^h satisfies efficiency, we have

$$\tau_1^h(1) - \tau_1^h(-1) \geq \frac{c}{\pi_{(h,d)}}.$$

Moreover,

$$p^h \tau_1^h(1) + (1 - p^h) \tau_1^h(-1) - c = u.$$

Consider two cases.

Case 1. $\tau_1^h(-1) \geq u - \tilde{u}$. In this case, we set $\tilde{\tau}_1^h(\omega_1) = \tau_1^h(\omega_1) - (u - \tilde{u})$ for $\omega_1 \in \{1, -1\}$. It can be easily verified that the new contract satisfies *EF* and *IR* and generates an expected utility of \tilde{u} , and for any $\omega_1 \in \{1, -1\}$ we have $\tilde{\tau}_1^h(\omega_1) < \tau_1^h(\omega_1)$.

Case 2. $\tau_1^h(-1) < u - \tilde{u}$. In this case set $\tilde{\tau}_1^h(-1) = 0$ and $\tilde{\tau}_1^h(1) = \frac{\tilde{u} + c}{p^h} < \frac{u + c - (1 - p^h) \tau_1^h(-1)}{p^h} = \tau_1^h(1)$, where the inequality follows from the fact that in this case $u - \tilde{u} > \tau_1^h(-1)$. Since $\tilde{u} \geq \bar{u}_1^h$, incentives are preserved. Moreover, the contract $\tilde{\tau}_1^h$ generates an expected utility of \tilde{u} and for any $\omega_1 \in \{1, -1\}$ we have $\tilde{\tau}_1^h(\omega_1) \leq \tau_1^h(\omega_1)$. This complete the proof for $T = 1$.

Having established the claim for $T = 1$, we proceed by assuming that the statement holds for $T = K - 1$ periods and we show that it holds for $T = K$ periods. Assume that there exists a K -period contract τ_K^h for which $u^h(\tau_K^h, sub_\emptyset) > \bar{u}_K^h$. As in the case of $T = 1$, we consider two cases.

Case 1. $u^h(\tau_K^h, sub_{-1}) - \bar{u}_{K-1}^h \geq u - \tilde{u}$. In this case consider two $K - 1$ -period contracts $\tau_{K-1,-1}^h$ and $\tau_{K-1,1}^h$ that satisfies EF and IR such that

1. $u^h(\tau_{K-1,-1}^h, sub_\emptyset) = u^h(\tau_K^h, sub_{-1}) - (u - \tilde{u})$ and $u^h(\tau_{K-1,1}^h, sub_\emptyset) = u^h(\tau_K^h, sub_1) - (u - \tilde{u})$
2. $\tau_K^h(-1, \omega_{K-1}) \geq \tau_{K-1,-1}^h(\omega_{K-1})$ and $\tau_K^h(1, \omega_{K-1}) \geq \tau_{K-1,1}^h(\omega_{K-1})$

Since $u^h(\tau_K^h, sub_{-1}) - (u - \tilde{u}) \geq \bar{u}_{K-1}^h$ and $u^h(\tau_K^h, sub_1) - (u - \tilde{u}) \geq \bar{u}_{K-1}^h$, the induction argument guarantees the existence of such contracts. Construct a contract $\tilde{\tau}_K^h$, such that $\tilde{\tau}_K^h(1, \omega_{K-1}) = \tau_{K-1,1}^h(\omega_{K-1})$ and $\tilde{\tau}_K^h(-1, \omega_{K-1}) = \tau_{K-1,-1}^h(\omega_{K-1})$ for any $\omega_{K-1} \in \Omega_{K-1}$. First, notice that by construction, the incentives in any subgame after the first period are guaranteed. Second, since $u^h(\tau_K^h, sub_1) - u^h(\tau_K^h, sub_{-1}) = u^h(\tilde{\tau}_K^h, sub_1) - u^h(\tilde{\tau}_K^h, sub_{-1})$, the first-period incentives are preserved. The expected utility of the high-quality agent is given by $p^h u^h(\tilde{\tau}_K^h, sub_1) + (1 - p^h) u^h(\tilde{\tau}_K^h, sub_{-1}) - c = \tilde{u}$. Finally, by construction, for all $\omega_K \in \Omega_K$ we have $\tilde{\tau}_K^h(\omega_K) \leq \tau_K^h(\omega_K)$, where the inequality is strict for at least one $\omega_K \in \Omega_K$.

Case 2. $u^h(\tau_K^h, sub_{-1}) - \bar{u}_{K-1}^h < u - \tilde{u}$. Consider two $K-1$ -period contracts that satisfy *EF* and *IR* $\tau_{K-1,-1}^h$ and $\tau_{K-1,1}^h$ such that $u^h(\tau_{K-1,-1}^h, sub_\emptyset) = \bar{u}_{K-1}^h$ and $u^h(\tau_{K-1,1}^h, sub_\emptyset) = \frac{\tilde{u} + c - (1 - p^h)\bar{u}_{K-1}^h}{p^h} < \frac{u + c - (1 - p^h)u^h(\tau_K^h, sub_{-1})}{p^h} = u^h(\tau_K^h, sub_1)$, where the inequality follows from the fact that in this case $u^h(\tau_K^h, sub_{-1}) - \bar{u}_{K-1}^h < u - \tilde{u}$. Since

$$p^h u^h(\tau_{K-1,1}^h, sub_\emptyset) + (1 - p^h) u^h(\tau_{K-1,-1}^h, sub_\emptyset) = \tilde{u} + c > \bar{u}_K^h + c > \bar{u}_{K-1}^h + c$$

and $u^h(\tau_{K-1,-1}^h, sub_\emptyset) = \bar{u}_{K-1}^h$ we get that $u^h(\tau_{K-1,1}^h, sub_\emptyset) > \bar{u}_{K-1}^h$. The induction argument guarantees the existence of the contracts with the required properties. As in the previous case we construct a contract $\tilde{\tau}_K^h$ from two $K-1$ -period contracts, $\tau_{K-1,-1}^h$ and $\tau_{K-1,1}^h$, such that $\tilde{\tau}_K^h(1, \omega_{K-1}) = \tau_{K-1,1}^h(\omega_{K-1})$ and $\tilde{\tau}_K^h(-1, \omega_{K-1}) = \tau_{K-1,-1}^h(\omega_{K-1})$ for any $\omega_{K-1} \in \Omega_{K-1}$. The rest of the proof is similar to the proof of Case 1. ■

Proof of Lemma 2:

Consider first the set Γ_1^h of one-period optimal contracts. Because effort is not verifiable, incentives must be provided to induce effort even when $p_1 = d$, where the probability of success is low. It follows that incentives to exert effort on all tasks are provided if and only if $\tau_1^h(1) - \tau_1^h(-1) \geq c/\pi(h,d)$. We conclude that Γ_1^h is a singleton and $\tau_1^h \in \Gamma_1^h$ if and only if $\tau_1^h(1) = c/\pi(h,d)$ and $\tau_1^h(-1) = 0$, which establishes (i) in the statement of the lemma. To complete the proof, note that Property 1 in Lemma 1 shows (ii.1) and to establish (ii.2) it is enough to show that for every $\tau_K^h \in \Gamma_K^h$ there exists a contract $\tilde{\tau}_{K-1}^h \in \Gamma_{K-1}^h$ such that for any history ω_{K-1} we have $\tilde{\tau}_{K-1}^h(\omega_{K-1}) \leq \tau_K^h(1, \omega_{K-1})$. This, however, follows from Property 4 in Lemma 1.

Proof of Claim 1. The proof is done by induction on T , the length of the contract. For $T = 1$, the statement holds trivially. We assume then that the

statement holds for $T = K$ and next prove it for $T = K + 1$. Denote by $\tilde{\tau}_{K+1}^h$ the contract that

1. yields a utility of $\bar{u}_{K+1}^h + \tilde{u}$ to the high-quality agent
2. induces sub contracts on sub_1 and sub_{-1} that are the contracts described in point 1 of Theorem 1 amended by some non-negative extra payments ψ_1 in sub_1 and ψ_{-1} in sub_{-1} , that are paid after a history of K uninterrupted successes.

First note that it is always possible to find ψ_1 and ψ_{-1} such that: (i) $u^h(\tilde{\tau}_{K+1}^h, sub_\emptyset) = \bar{u}_{K+1}^h + \tilde{u}$, (ii) the incentives for the high-quality agent are preserved (for example by choosing $\psi_{-1} = 0$), and (iii) by the induction argument, $\tilde{\tau}_{K+1}^h$ minimizes the expected utility of the low-quality agent in each of these subgames in all contracts that generate an expected utility of $\bar{u}_K^h + (p^h)^K \psi_1$ and $\bar{u}_K^h + (p^h)^K \psi_{-1}$, respectively.

It is left for us to show that $\psi_{-1} = 0$. Assume by way of contradiction that $\psi_{-1} > 0$ and consider decreasing ψ_{-1} (the payment after a failure following a sequence of K successes) by $\varepsilon > 0$ and increasing ψ_1 (the payment after a sequence of $K + 1$ successes) by $\varepsilon \frac{p^h}{1-p^h}$. Note that this change does not affect the expected utility of the high-quality agent, preserves his incentives and for any decrease the utility of the low-quality agent by ε . ■

Proof of Corollary 2. We first show that the contract described in Definition 4 is an optimal contract for the high-quality agent. To see this, observe that in this contract the expected compensation for success in each period is $\frac{c}{\pi(h,d)}$. It follows that this contract provides the high-quality agent with the same expected utility as did the original contract described in Theorem 1 and it generates efficient incentives. We next show that there exists a threshold $\hat{T}(\alpha)$ so that if $T > \hat{T}(\alpha)$, the contract yields the same expected utility for the low-quality agent as the one for the high-quality agent described in Theorem 1. It is enough to establish that if the low-quality agent adopts this contract then his best strategy is to exert no effort during the first $T/2$ periods. We now show now that for any period $m \in \{1, \dots, T/2\}$, if the low-quality agent does not succeed in all $t < m$ periods, his expected utility is higher if he does not exert effort in period m . Since his probability of success in any period is bounded by p^l , the change in his expected utility if he exerts effort at period $m \in \{1, \dots, T/2\}$ is bounded by

$$-c + \pi_{(l,e)} \frac{c}{\pi(h,d)} \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}}.$$

To show that this expression is negative for T big enough, it is enough to show

that

$$\lim_{T \rightarrow \infty} \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}} = 0.$$

Note that

$$\begin{aligned} & \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^h)^j (1-p^h)^{T-m-1-j}} \\ & \leq \frac{\sum_{j=\lceil \alpha(T-m-1) \rceil}^{T-m-1} \binom{T-m-1}{j} (p^l)^j (1-p^l)^{T-m-1-j}}{(p^h)^{T-m-1}} \\ & \leq \frac{\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} (p^l)^{\lceil \alpha(T-m-1) \rceil} (1-\alpha)(T-m-1)}{(p^h)^{T-m-1}} \end{aligned}$$

where the first inequality follows from the fact that $(p^h)^{T-m-1}$ is just one of the elements in summation. The second inequality follows from (i) $p^l \in (0, 1)$ and (ii). For $\alpha \geq 1/2$ the monotonicity of the binomial coefficient implies that $\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} \geq \binom{T-m-1}{j}$ for any $j \in \{\lceil \alpha(T-m-1) \rceil, \dots, T-m-1\}$.

By Stanica (2001; Corollary 2.3), it follows that for $\alpha \geq 1/2$, the binomial coefficient is bounded by

$$\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} \leq \frac{1}{\sqrt{2\pi\alpha(1-\alpha)(T-m-1)}} \frac{\left(\frac{1}{\alpha}\right)^{T-m-1}}{\left(\frac{1}{\alpha}-1\right)^{(T-m-1)(1-\alpha)}}.$$

Plugging this bound into the previous expression yields

$$\begin{aligned} & \frac{\binom{T-m-1}{\lceil \alpha(T-m-1) \rceil} (p^l)^{\lceil \alpha(T-m-1) \rceil} (1-\alpha)(T-m-1)}{(p^h)^{T-m-1}} \\ & \leq \frac{\sqrt{(1-\alpha)(T-m-1)}}{\sqrt{2\pi\alpha}} \left(\frac{\frac{1}{\alpha} (p^l)^\alpha}{\left(\frac{1}{\alpha}-1\right)^{(1-\alpha)} p^h} \right)^{T-m-1} \\ & = \frac{\sqrt{(1-\alpha)(T-m-1)}}{\sqrt{2\pi\alpha}} \left(\frac{(p^l)^\alpha \left(\frac{1}{\alpha}-1\right)^\alpha}{(1-\alpha) p^h} \right)^{T-m-1} \end{aligned}$$

Therefore, for $(p^l)^\alpha \left(\frac{1}{\alpha}-1\right)^\alpha < (1-\alpha) p^h$ we have

$$\lim_{T \rightarrow \infty} \frac{\sqrt{(1-\alpha)(T-m-1)}}{\sqrt{2\pi\alpha}} \left(\frac{(p^l)^\alpha \left(\frac{1}{\alpha}-1\right)^\alpha}{(1-\alpha) p^h} \right)^{T-m-1} = 0$$

which completes the proof. ■

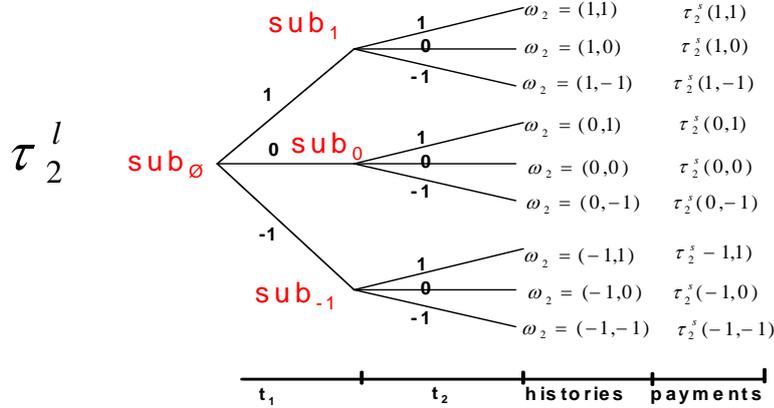


Figure 5: Two-period contract for a low-quality agent.

9 Appendix B: The Pooling Contract-pair

In this appendix we turn our attention to the second case, where $\pi_{(l,e)} \leq \pi_{(h,d)}$. We solve for the optimal contract in this case similarly to how we solved for the optimal contract in the first case, where $\pi_{(l,e)} > \pi_{(h,d)}$. That is, we start by assuming that the agent is known to be of low-quality, and define the set of optimal contracts Γ_T^l . After characterizing Γ_T^l , we drop the assumption that the agent is known to be of low-quality and show that when the agent's quality is unobservable the contract-pair (τ_T^h, τ_T^l) is optimal only if $\tau_T^l \in \Gamma_T^l$. Equipped with this result, it is rather easy to characterize the optimal contract pair $(\bar{\tau}_T^h, \bar{\tau}_T^l)$ and show that $\bar{\tau}_T^h \equiv \bar{\tau}_T^l$.

9.1 Agent's Quality is Known to be Low

First note that unlike the contract for the high-quality agent who is expected to operate on all types of tasks, the contract for a low-quality agent imposes no such requirement. It is thus necessary to consider also payments after histories along which at some t the agent's choice is not to operate. A two-period contract for a low-quality agent is depicted above.

One of the differences between the present case and the previous one is that the optimal mechanism provides no information rents to the agent. In particular, a contract that pays a constant sum of $\frac{c}{\pi_{(l,e)}}$ per success provides efficient incentives and generates an expected utility of zero to the agent, which, in particular, implies that it is optimal. The next lemma provides a characterization of the set Γ_T^l of T -period *optimal* contracts when the agent is known to be of low-quality.

Lemma 5 *Properties of Γ_T^l*

1. If a contract $\tau_T^l \in \Gamma_T^l$, then $u^l(\tau_T^l, sub_1) - u^l(\tau_T^l, sub_{-1}) = \frac{c}{\pi(l,e)}$ and $u^l(\tau_T^l, sub_0) = u^l(\tau_T^l, sub_{-1})$.
2. If a contract $\tau_T^l \in \Gamma_T^l$, then $\exists \tau_{T-1}^l \in \Gamma_{T-1}^l$ s.t. $\forall \omega_{T-1} \in \Omega_{T-1}$, $\tau_T^l(-1, \omega_{T-1}) = \tau_{T-1}^l(\omega_{T-1})$. Also $\exists \tau_{T-1}^l \in \Gamma_{T-1}^l$ s.t. $\forall \omega_{T-1} \in \Omega_{T-1}$, $\tau_T^l(0, \omega_{T-1}) = \tau_{T-1}^l(\omega_{T-1})$.
3. If a contract $\tau_T^l \in \Gamma_T^l$, then $m^l(\tau_T^l, sub_\emptyset) = T \cdot c \cdot p^l \frac{1-q}{\pi(l,e)}$.
4. Assume that a T -period contract τ_T^l satisfies EF and IR and for which $u^l(\tau_T^l, sub_\emptyset) = u > 0$. Then for any $\tilde{u} \in [0, u)$ there exists another T -period contract $\tilde{\tau}_T^l$ that also satisfies EF and IR and for which $u^l(\tilde{\tau}_T^l, sub_\emptyset) = \tilde{u}$. Moreover, for any history $\omega_T \in \Omega_T$ we have $\tau(\omega_T) \geq \tilde{\tau}(\omega_T)$ with at least one strict inequality.

Proof: Property 1. First, observe that if $u^l(\tau_T^l, sub_1) - u^l(\tau_T^l, sub_{-1}) < \frac{c}{\pi(l,e)}$, then the low-quality agent will not exert effort even if an easy task arrives in the first period. Also note that $u^l(\tau_T^l, sub_0) \geq u^l(\tau_T^l, sub_{-1})$ since otherwise the agent will accept the task and not exert effort when a difficult task arrives in the first period. Assume then that $\tau_T^l \in \Gamma_T^l$ but $u^l(\tau_T^l, sub_1) - u^l(\tau_T^l, sub_{-1}) > \frac{c}{\pi(l,e)}$. Consider then the following changes: in sub_0 adopt the same payment as in sub_{-1} and in sub_1 add a payment of $\frac{c}{\pi(l,e)}$ to every history of sub_{-1} . Note that these changes preserve incentives and decrease the expected payment, in contradiction to the fact that $\tau_T^l \in \Gamma_T^l$.

Property 2. Assume that this property is false. $\tau_T^l \in \Gamma_T^l$, so τ_T^l provides sufficient incentives in all subgames, and in particular in sub_{-1} . Consider revising the contract τ_T^l to $\tilde{\tau}_T^l$ as follows

1. replace the payments in all histories that belong to sub_{-1} by payments in one of the optimal $T-1$ -period contracts $\tilde{\tau}_{T-1}^l \in \Gamma_{T-1}^l$ (that is, $\tilde{\tau}_T^l(-1, \omega_{T-1}) = \tilde{\tau}_{T-1}^l(\omega_{T-1})$)
2. adjust the contracts in other subgames correspondingly (that is, adopt in sub_0 the same payments as in sub_{-1} , and adopt in sub_1 the same payments as in sub_{-1} and add $\frac{c}{\pi(l,e)}$ to every history).

The proposed change preserves incentives to invest in all subgames after the first period and generates efficient incentives in the first period.

Since the new payment scheme in sub_{-1} is a contract in Γ_{T-1}^l , it minimizes expected payment in all schemes that provide incentives. It follows that the proposed change strictly decreases $u^l(\tau_T^l, sub_{-1})$ because otherwise the τ_T^l payment restricted to sub_{-1} is a contract from Γ_{T-1}^l . Property 1 of the lemma implies

that this change also decreases $u^l(\tau_T^l, sub_0)$ and $u^l(\tau_T^l, sub_1)$, in contradiction to $\tau_T^l \in \Gamma_T^l$. The same argument also establishes that the payment in sub_0 is a contract in Γ_{T-1}^l .

Property 3. Consider the contract that pays $c/\pi_{(l,e)}$ per success (i.e., pays $\frac{nc}{\pi_{(l,e)}}$ after a history of n successes). Note that this is an optimal contract even when the principal observes the type of the arriving task and the effort exerted by the agent. Therefore, it is an optimal contract when the task's type and the agent's effort are not observable. Since in this contract $m^l(\tau_T^l, sub_\emptyset) = Tcp^l \frac{1-q}{\pi_{(l,e)}}$, any optimal contract should pay the same expected payment as the one described above.

Property 4. The proof is done by induction on the contract's length T . Start with $T = 1$ and observe that since τ_1^l satisfies EF we have that

$$\begin{aligned} \tau_1^l(0) &\geq \pi_{(l,d)}\tau_1^l(1) + (1 - \pi_{(l,d)})\tau_1^l(-1) - c & (7) \\ \tau_1^l(0) &\geq \tau_1^l(-1) \\ \tau_1^l(0) &\leq \pi_{(l,e)}\tau_1^l(1) + (1 - \pi_{(l,e)})\tau_1^l(-1) - c \\ \tau_1^l(-1) &\leq \pi_{(l,e)}\tau_1^l(1) + (1 - \pi_{(l,e)})\tau_1^l(-1) - c \end{aligned}$$

There are two cases to consider.

Case 1 $\tau_1^l(0) \geq \tilde{u}$. From (7) we get that

$$\pi_{(l,e)}\tau_1^l(1) + (1 - \pi_{(l,e)})\tau_1^l(-1) - c \geq \tau_1^l(0) \geq \tilde{u}. \quad (8)$$

Set $\tilde{\tau}_1^l(0) = \tilde{u}$. If $\tau_1^l(-1) \geq \tilde{u}$, then set $\tilde{\tau}_1^l(0) = \tilde{\tau}_1^l(-1) = \tilde{u}$ and $\tilde{\tau}_1^l(1) = \tilde{u} + \frac{c}{\pi_{(l,e)}} \leq \tau_1^l(1)$, where the last inequality follows from the fact that the original payment satisfied EF , which in particular implies that $\tau_1^l(1) \geq \tau_1^l(-1) + \frac{c}{\pi_{(l,e)}}$. Now note that since the expected payments are strictly lower in $\tilde{\tau}_1^l$ and both contracts satisfy EF , there exists at least one history where the payment in $\tilde{\tau}_1^l$ is strictly lower than in τ_1^l . If $\tau_1^l(-1) < \tilde{u}$, then (8) and the fact that expected utility in τ_1^l is u implies that $\tau_1^l(1) > \frac{c + \tilde{u} - (1 - \pi_{(l,e)})\tau_1^l(-1)}{\pi_{(l,e)}}$. Set $\tilde{\tau}_1^l(1) = \frac{c + \tilde{u} - (1 - \pi_{(l,e)})\tau_1^l(-1)}{\pi_{(l,e)}}$ and $\tilde{\tau}_1^l(-1) = \tau_1^l(-1)$. Recall that $\tilde{\tau}_1^l(0) = \tilde{u}$ and observe that $\tilde{\tau}_1^l$ generates an expected utility of \tilde{u} and satisfies EF .

Case 2 $\tau_1^l(0) < \tilde{u}$. We start by setting $\tilde{\tau}_1^l(0) = \tau_1^l(0)$ and proceed by decreasing the utility from exerting effort by $\frac{u - \tilde{u}}{1 - q}$, which will generate for $\tilde{\tau}_1^l$ an expected utility of \tilde{u} . If $\tau_1^l(-1) \geq \frac{u - \tilde{u}}{1 - q}$, then set $\tilde{\tau}_1^l(\omega_1) = \tau_1^l(\omega_1) - \frac{u - \tilde{u}}{1 - q}$ for $\omega_1 \in \{1, -1\}$. If, however, $\tau_1^l(-1) < \frac{u - \tilde{u}}{1 - q}$, then set $\tilde{\tau}_1^l(-1) = 0$ and $\tilde{\tau}_1^l(1) = \frac{\tilde{u} + c - q\tilde{\tau}_1^l(0)}{(1 - q)\pi_{(l,e)}} < \tau_1^l(1)$, where the last inequality follows from the fact that when $\tau_1^l(-1) < \frac{u - \tilde{u}}{1 - q}$, decreasing the utility of the agent from exerting effort by $\frac{u - \tilde{u}}{1 - q}$ implies that the payment conditional on success should be decreased by more than the amount decreased in the case where $\tau_1^l(-1) \geq \frac{u - \tilde{u}}{1 - q}$. However, since $\tilde{u} - q\tilde{\tau}_1^l(0) > 0$, the payments in $\tilde{\tau}_1^l$ satisfy EF and all payments are lower.

Having established the claim for $T = 1$, we next assume that the statement holds for $T = K - 1$ periods and show that it holds for $T = K$ periods. Assume that there exists τ_K^l for which $u^l(\tau_K^l, sub_\emptyset) = u > 0$. Similarly to the proof for $T = 1$, there are two cases to consider.

Case 1 $u^l(\tau_K^l, sub_0) \geq \tilde{u}$. We start by replacing the payment in sub_0 with a $K - 1$ -period contract $\tilde{\tau}_{(K-1)_0}^l$ that generates an expected payment of \tilde{u} and for which $\tilde{\tau}_{(K-1)_0}^l(\omega_{K-1}) \leq \tau_K^l(0, \omega_{K-1}) \forall \omega_{K-1} \in \Omega_{K-1}$. Such a contract exists by the induction argument. If $u^l(\tau_K^l, sub_{-1}) \geq \tilde{u}$, then we replace the payments in sub_{-1} by a $K - 1$ -period contract $\tilde{\tau}_{(K-1)_{-1}}^l(\omega_{K-1})$ that generates an expected payment of \tilde{u} and for which we have $\tilde{\tau}_{(K-1)_{-1}}^l(\omega_{K-1}) \leq \tau_K^l(-1, \omega_{K-1}) \forall \omega_{K-1} \in \Omega_{K-1}$ (again, such a contract exists by the induction argument). We complete this part of the argument by replacing the payments in sub_1 by a $K - 1$ -period contracts $\tilde{\tau}_{(K-1)_1}^l(\omega_{K-1})$ that generates an expected payment of $\tilde{u} + \frac{c}{\pi_{(l,e)}}$ and for which $\tilde{\tau}_{(K-1)_1}^l(\omega_{K-1}) \leq \tau_K^l(1, \omega_{K-1}) \forall \omega_{K-1} \in \Omega_{K-1}$ (again, such a contract exists by the induction argument). We still have to show that there exists a history ω_K for which the inequality is strict. However, since the new contract $\tilde{\tau}_K^l$ generated from the three contracts $\tilde{\tau}_{(K-1)_z}^l$ for $z = -1, 0, 1$ generates a strictly lower expected payment and all contracts satisfy efficiency, there must exist at least one history for which the inequality is strict.

The proof of the case where $u^l(\tau_K^l, sub_0) < \tilde{u}$ is similar. ■

Remark 7 *An immediate consequence of the lemma and in particular of Property 3 is that for all $\tau_T^l \in \Gamma_T^l$, we have $u^l(\tau_T^l, sub_\emptyset) = 0$.*

9.2 Agent's Quality is Unobservable

We are now ready to characterize the optimal contract-pair when $\pi_{(l,e)} \leq \pi_{(h,d)}$, which is shown to have a very simple structure. Namely, the two contracts are the same and they pay a fixed compensation per success. Moreover, this contract belongs to the set of optimal contracts when the agent is known to be a low-quality agent. We start by establishing the latter.

Lemma 6 *When $\pi_{(l,e)} \leq \pi_{(h,d)}$, (τ_T^h, τ_T^l) is an optimal contract-pair, only if $\tau_T^l \in \Gamma_T^l$.*

Proof. Assume by way of contradiction that (τ_T^h, τ_T^l) is an optimal contract-pair, but $\tau_T^l \notin \Gamma_T^l$. Since (τ_T^h, τ_T^l) is optimal the contract-pair is admissible and in particular satisfies *IR* and *EF*. Since $\tau_T^l \notin \Gamma_T^l$, Remark 7 implies that $u^l(\tau_T^l, sub_\emptyset) > 0$. Hence, Property 4 of Lemma 5 implies that there exists a contract $\hat{\tau}_T^l \in \Gamma_T^l$ such that for every history $\omega_T \in \Omega_T$, $\tau_T^l(\omega_T) \geq \hat{\tau}_T^l(\omega_T)$ with strict inequality for at least one $\omega_T \in \Omega_T$. Consider replacing (τ_T^h, τ_T^l) by the pair $(\tau_T^h, \hat{\tau}_T^l)$. To verify that this contract-pair satisfies *EF*, note first that since $\pi_{(h,d)} \geq \pi_{(l,e)}$ *EF* is satisfied for the high-quality agent whenever it is satisfied

for the low-quality one, and the latter holds since $\hat{\tau}_T^l \in \Gamma_T^l$. Obviously, *IC* holds as well for this new contract-pair $(\hat{\tau}_T^l, \hat{\tau}_T^l)$. By definition, the expected payments to the low-quality agent are now lower, and the same (with weak inequality) also holds for the high-quality agent. That is,

$$(i) \ m^l(\hat{\tau}_T^l, \text{sub}_\emptyset) < m^l(\tau_T^l, \text{sub}_\emptyset) \text{ and } (ii) \ m^h(\hat{\tau}_T^l, \text{sub}_\emptyset) \leq m^h(\tau_T^h, \text{sub}_\emptyset).$$

To verify *(ii)*, recall that the original contract-pair (τ_T^h, τ_T^l) was incentive-compatible, which in particular implies that the high-quality agent prefers the contract τ_T^h to τ_T^l . By Property 4 of Lemma 5 the new contract $\hat{\tau}_T^l$ generates for the high-quality agent an even lower expected utility than τ_T^l . Since this contract satisfies *EF*, the monotonicity relation between expected payment and expected utility implies *(ii)*. This establishes the contradiction to the statement that the original contract-pair (τ_T^h, τ_T^l) was optimal. ■

Theorem 2 *When $\pi_{(l,e)} \leq \pi_{(h,d)}$, the optimal contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is*

$$\hat{\tau}_T^h(\omega_T) = \hat{\tau}_T^l(\omega_T) = \frac{c \cdot n(\omega_T)}{\pi_{(l,e)}}$$

where $n(\omega_T)$ is the number of successes in ω_T .

The proof of the theorem is a simple result of the following claim and hence will be provided after the proof of the claim.

Claim 3 *Assume that the principal is asked to provide the low-quality agent with an expected utility of $\bar{u} \geq 0$ such that*

1. *incentives to exert efficient effort are preserved*
2. *the expected utility of the high-quality agent is minimized among all contracts that provide efficient incentives to the low-quality agent and generates the expected utility of $\bar{u} \geq 0$ to the low-quality agent.*

This is achieved by amending the contract $\hat{\tau}_T^l$ described in Theorem 2 and adding a payment of \bar{u} after every history. That is,

$$\hat{\tau}_T^l(\omega_T) = \frac{cn(\omega_T)}{\pi_{(l,e)}} + \bar{u} \quad \text{for all } \omega_T \in \Omega_T.$$

Proof. We prove the claim by induction on the length of the contract T . Start with one period. Recall that in this case the only optimal contract for the low-quality agent is $\tau_1^l(1) = c/\pi_{(l,e)}$, $\tau_1^l(-1) = 0$, and $\tau_1^l(0) = 0$. Denote by $u(\omega_1)$ the additional payment above $\tau_1^l(\omega_1)$ for $\omega_1 \in \{-1, 0, 1\}$. First, note that $u(1) \geq u(-1)$, because otherwise the agent would not exert effort when an

easy task arrives. In addition, observe that $u(1) \geq u(0)$, because otherwise the low-quality agent would pass the task even if an easy task arrives. Also note that $u(0) \geq u(-1)$, because otherwise the agent would accept the task (maybe without exerting effort) even when the arriving task is difficult, $p_t = d$. Recall that since $\pi_{(l,e)} \leq \pi_{(h,d)}$, if incentives are provided for the low-quality agent to exert effort on $p_t = e$, then the high-quality agent will operate on all types of tasks, if he faces the same contract. Moreover, $\pi_{(h,d)} \geq \pi_{(l,e)}$ implies that specifying $u(1) = u(0) = u(-1) = \bar{u}$ necessarily minimizes the utility of the high-quality agent from all contracts that generate efficient incentives for the low-quality agent and provides him with the additional utility of \bar{u} .

We assume that the statement holds for $T = K - 1$ periods and proceed to the proof of the statement for $T = K$ periods. Consider a contract τ_K^l that yields a utility of \bar{u} to the low-quality agent and minimizes the expected utility of the high-quality agent. We first show that the induced contracts on sub_1 , sub_0 , and sub_{-1} by τ_K^l are as described in the statement of the claim. The reason for that is as follows: assume by way of contradiction that the above statement is false and note that: (i) it is always possible to construct a contract $\tilde{\tau}_K^l$ such that the induced contracts on sub_1 , sub_0 , and sub_{-1} by $\tilde{\tau}_K^l$ are as described in the statement of the claim and for which there are \bar{u}_1 , \bar{u}_{-1} , and \bar{u}_0 such that the low-quality agent is indifferent between $\tilde{\tau}_K^l$ and τ_K^l ; (ii) the incentives for the low-quality agent are preserved in $\tilde{\tau}_K^l$; and (iii) by the induction argument, in each of the subgames, the amended contract $\tilde{\tau}_K^l$ decreases the expected utility of the high-quality agent. We still need to show that $\bar{u}_1 = \bar{u}_{-1} = \bar{u}_0 = \bar{u}$. However, this proof is identical to the proof of the one-period case. ■

Proof of Theorem 2. First observe that $\hat{\tau}_T^l \in \Gamma_T^l$ and that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ satisfies *EF* and *IR*. It follows that if we prove that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is *optimal* we are done. As in Theorem 1, we need to show that the contract $\hat{\tau}_T^l$ described in the theorem minimizes the expected utility of the high-quality agent from all contracts belonging to Γ_T^l , but the rest of the proof follows from the previous claim for $\bar{u} = 0$. ■

10 Appendix C: The Discounting Case

10.1 Proof of Proposition 1.

The line of proof is similar to the case without discounting. In particular, we first characterize the properties of the set containing all the contracts between which and the base-line contract the high-quality agent is indifferent. When the agent type is unobservable, we show that τ_T^h belongs to this set if (τ_T^h, τ_T^l) is the optimal contract-pair. Denote by $\hat{\Gamma}_T^h$ the set of contracts that satisfy *IR* and *EF* for the high-quality agent, and between which and the base-line contract, denoted by $\underline{\tau}_T^h$, the high-quality agent is indifferent. Note that the expected

utility of the high-quality agent from $\underline{\tau}_T^h$ is

$$u^h \left(\underline{\tau}_T^h, \text{sub}_\emptyset \right) = c \left(\frac{p^h}{\pi(h,d)} - 1 \right) (1 - \delta^T) \frac{\delta}{1 - \delta}$$

Lemma 7 *If contract τ_T^h satisfies EF for the high-quality agent, then $u^h (\tau_T^h, \text{sub}_\emptyset) \geq u^h (\underline{\tau}_T^h, \text{sub}_\emptyset)$.*

Proof. The proof is done by induction on T . For $T = 1$, the statement trivially holds. Assume it holds for $T = \kappa$ and we now show that it holds for $T = \kappa + 1$. Note that after a failure in the first period, the resulting contract is a κ -period contract. Since it satisfies EF, the induction argument implies that

$$u^h \left(\tau_{\kappa+1}^h, \text{sub}_{-1} \right) \geq c \left(\frac{p^h}{\pi(h,d)} - 1 \right) (1 - \delta^\kappa) \frac{\delta}{1 - \delta}.$$

Since τ_k^h satisfies EF for the high-quality agent, we have

$$\tau_k^h(1) + u^h \left(\tau_k^h, \text{sub}_{+1} \right) - u^h \left(\tau_{\kappa+1}^h, \text{sub}_{-1} \right) \geq c \frac{1}{\pi(h,d)}.$$

Since $u^h (\tau_{\kappa+1}^h, \text{sub}_\emptyset) = \delta (-c + p^h (\tau_k^h(1) + u^h (\tau_k^h, \text{sub}_{+1})) + (1 - p^h) u^h (\tau_{\kappa+1}^h, \text{sub}_{-1}))$ we have

$$\begin{aligned} u^h \left(\tau_{\kappa+1}^h, \text{sub}_\emptyset \right) &\geq c \left(\frac{p^h}{\pi(h,d)} - 1 \right) (1 - \delta^\kappa) \frac{\delta^2}{1 - \delta} + \delta \left(-c + c \frac{p^h}{\pi(h,d)} \right) \\ &= c \left(\frac{p^h}{\pi(h,d)} - 1 \right) (\delta^2 + \delta^3 + \dots + \delta^{\kappa+1} + \delta) \\ &= c \left(\frac{p^h}{\pi(h,d)} - 1 \right) (1 - \delta^{\kappa+1}) \frac{\delta}{1 - \delta} \end{aligned}$$

■

Lemma 8 *Assume that τ_T^h satisfies EF for the high-quality agent and $u^h (\tau_T^h, \text{sub}_\emptyset) > u^h (\underline{\tau}_T^h, \text{sub}_\emptyset)$; then there exists another contract $\tilde{\tau}_T^h$ that also satisfies EF and that for any $t \leq T$ and ω_t holds $\tilde{\tau}_T^h(\omega_t) \leq \tau_T^h(\omega_t)$ with at least one strict inequality.*

Proof. Proof is an immediate consequence of the following claim. ■

Claim 4 *Assume that τ_T^h satisfies EF and $u^h (\tau_T^h, \text{sub}_\emptyset) > u^h (\underline{\tau}_T^h, \text{sub}_\emptyset)$; then for any $u \in [u^h (\underline{\tau}_T^h, \text{sub}_\emptyset), u^h (\tau_T^h, \text{sub}_\emptyset)]$ there exists another contract $\tilde{\tau}_T^h$ that satisfies EF such that $u^h (\tilde{\tau}_T^h, \text{sub}_\emptyset) = u$, and for any $t \leq T$ and ω_t we have $\tilde{\tau}_T^h(\omega_t) \leq \tau_T^h(\omega_t)$ with at least one strict inequality.*

Proof. The proof is done by induction on the number of the periods. For $T = 1$, the proof is identical to the proof of Claim 2. Assume that the statement of the claim holds for $T = k - 1$ and we will show it for $T = k$. There are two cases to consider.

Case 1. $\delta (\tau_k^h(-1) + u^h(\tau_k^h, sub_{-1}) - u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)) \geq u^h(\tau_k^h, sub_\emptyset) - u$.

In this case

(i) set $\tilde{\tau}_k^h(-1) = \max \left\{ \tau_k^h(-1) - \frac{u^h(\tau_k^h, sub_\emptyset) - u}{\delta}, 0 \right\}$ and $\tilde{\tau}_k^h(1) = \max \left\{ \tau_k^h(1) - \frac{u^h(\tau_k^h, sub_\emptyset) - u}{\delta}, c/\pi(h, d) \right\}$

(ii) At sub_{-1} set a $k - 1$ period contract that is efficient and generates an expected utility of

$$u^h(\tau_k^h, sub_{-1}) - \frac{u^h(\tau_k^h, sub_\emptyset) - u - \max \left\{ \tau_k^h(-1) - \frac{u^h(\tau_k^h, sub_\emptyset) - u}{\delta}, 0 \right\}}{\delta} \geq u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)$$

where payments dominated by τ_k^h are followed by failure in the first period (such a contract exists by induction argument).

(iii) At sub_1 set a $k - 1$ period contract that is efficient and generates an expected utility of

$$u^h(\tau_k^h, sub_1) - \frac{u^h(\tau_k^h, sub_\emptyset) - u - \max \left\{ \tau_k^h(1) - \frac{u^h(\tau_k^h, sub_\emptyset) - u}{\delta}, c/\pi(h, d) \right\}}{\delta}$$

where payments dominated by contract τ_k^h are followed by success in the first period (again, such a contract exists by the induction argument).

The contract $\tilde{\tau}_k^h$ that we have constructed satisfies all the required properties.

Case 2. $\delta (\tau_k^h(-1) + u^h(\tau_k^h, sub_{-1}) - u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)) \geq u^h(\tau_k^h, sub_\emptyset) - u$.

In this case we set the new contract $\tilde{\tau}_k^h$ in the following way.

(i) Set $\tilde{\tau}_k^h(-1) = 0$ and at sub_{-1} set a $k - 1$ period contract that is efficient and generates an expected utility of $u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)$. The induction argument implies that there exists a $k - 1$ period contract that (a) satisfies EF for the high-quality agent, (b) generates an expected utility of $u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)$, and (c) payments are dominated by a $k - 1$ period contract τ_k^h followed by failure in the first period.

Note that this change satisfies EF for the high-quality agent and decreases his utility by

$$(1 - p^h) \delta (\tau_k^h(-1) + u^h(\tau_k^h, sub_{-1}) - u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)).$$

(ii) Set $\tilde{\tau}_k^h(1) = \max \left\{ \tau_k^h(1) - \frac{u^h(\tau_k^h, sub_\emptyset) - u - (1 - p^h) \delta (\tau_k^h(-1) + u^h(\tau_k^h, sub_{-1}) - u^h(\underline{\tau}_{k-1}^h, sub_\emptyset))}{\delta p^h}, c/\pi(h, d) \right\}$

and at sub_1 set a $k - 1$ period contract that satisfies EF for the high-quality agent and generates an expected utility of

$$u^h(\tilde{\tau}_k^h, sub_1) = u^h(\tau_k^h, sub_1) - \frac{1}{\delta} \frac{u^h(\tau_k^h, sub_\emptyset) - u - (1 - p^h) \delta (\tau_k^h(-1) + u^h(\tau_k^h, sub_{-1}) - u^h(\underline{\tau}_{k-1}^h, sub_\emptyset))}{p^h} - (\tau_k^h(1) - \tilde{\tau}_k^h(1)).$$

It is easy to verify that $\delta u^h(\tilde{\tau}_k^h, sub_1) + \tilde{\tau}_k^h(1) \geq c/\pi_{(h,d)} + \delta u^h(\underline{\tau}_{k-1}^h, sub_\emptyset)$. Therefore, the induction argument implies that there exists a $k-1$ period contract that (a) satisfies *EF* for the high-quality agent, (b) generates a utility of $u^h(\tilde{\tau}_k^h, sub_1)$, and (c) payments are dominated by contract τ_k^h followed by success in the first period.

The contract $\tilde{\tau}_k^h$ that we have constructed satisfies all the required properties. ■

Lemma 9 *If (τ_T^h, τ_T^l) is an optimal contract-pair, then $\tau_T^h \in \hat{\Gamma}_T^h$.*

Proof. Assume by way of contradiction that $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal but $\hat{\tau}_T^h \notin \hat{\Gamma}_T^h$. Since $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is an optimal contract-pair, it is admissible, and, in particular, both contracts satisfy *IR* and *EF*. Hence, the previous lemma implies that there exists a contract $\tilde{\tau}_T^h \in \hat{\Gamma}_T^h$ such that for any $k \leq T$ and $\omega_k \in \Omega_k$, $\hat{\tau}_T^h(\omega_k) \geq \tilde{\tau}_T^h(\omega_k)$ with strict inequality for at least one history. Hence, replacing $\hat{\tau}_T^h$ with $\tilde{\tau}_T^h$ will decrease the expected utility of the low-quality agent in the event that he pretends to be a high-quality agent and adopts the high-quality agent's contract. So, consider a contract $\tilde{\tau}_T^l$ that pays $M(\tilde{\tau}_T^h) > 0$ in the first period followed by $c/\pi_{(l,e)}$ per success, where $M(\tilde{\tau}_T^h)$ is chosen such that the low-quality agent is indifferent between this contract and the contract $\tilde{\tau}_T^h$. To see that such a payment M always exists, it is enough to note that for $M = 0$, the expected utility of the low quality agent in $\tilde{\tau}_T^l$ is zero, while $u^l(\tilde{\tau}_T^h, sub_\emptyset) \geq 0$.

We next argue that $(\tilde{\tau}_T^h, \tilde{\tau}_T^l)$ is admissible. Note that to show admissibility, it is sufficient to show that the high quality agent prefers to stick to his contract $\tilde{\tau}_T^h$ rather than switching to $\tilde{\tau}_T^l$. Recall that the expected utility of the high-quality agent from $\tilde{\tau}_T^h$ is

$$c \left(\frac{p^h}{\pi_{(h,d)}} - 1 \right) (1 - \delta^T) \frac{\delta}{1 - \delta} = c \frac{\pi_{(h,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} (1 - q) (1 - \delta^T) \frac{\delta}{1 - \delta}$$

and his expected utility from the $\tilde{\tau}_T^l$ is

$$M \left(\tilde{\tau}_T^h \right) + c \frac{\pi_{(h,e)} - \pi_{(l,e)}}{\pi_{(l,e)}} (1 - q) (1 - \delta^T) \frac{\delta}{1 - \delta}.$$

If

$$M \left(\tilde{\tau}_T^h \right) \leq c \left(\frac{\pi_{(h,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} - \frac{\pi_{(h,e)} - \pi_{(l,e)}}{\pi_{(l,e)}} \right) (1 - q) (1 - \delta^T) \frac{\delta}{1 - \delta}$$

then the contract-pair $(\tilde{\tau}_T^h, \tilde{\tau}_T^l)$ is admissible and we are done since the expected payments to both agents were reduced. Let us therefore, assume that the opposite inequality holds and consider switching $\tilde{\tau}_T^h$ to the base-line contract $\underline{\tau}_T^h$. Note that in this case

$$M \left(\underline{\tau}_T^h \right) = c \frac{\pi_{(l,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} (1 - q) (1 - \delta^T) \frac{\delta}{1 - \delta}.$$

Since for $\pi_{(h,d)} < \pi_{(l,e)}$ we have

$$\frac{\pi_{(l,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} < \frac{\pi_{(h,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} - \frac{\pi_{(h,e)} - \pi_{(l,e)}}{\pi_{(l,e)}},$$

this implies that $M(\underline{\tau}_T^h) < M(\hat{\tau}_T^h)$.

Therefore, the contract-pair where the high-quality agent gets the base-line contract, while the low-quality agent gets the payment $M(\underline{\tau}_T^h)$ in the first period followed by $c/\pi_{(l,e)}$ per success, is admissible and generates for both types of agents an expected payment lower than $(\hat{\tau}_T^h, \hat{\tau}_T^l)$. Again, this is in contradiction to the assumed optimality of the original pair. We conclude that if a contract-pair $(\hat{\tau}_T^h, \hat{\tau}_T^l)$ is optimal, then $\hat{\tau}_T^h \in \hat{\Gamma}_T^h$. ■

Lemma 10 *If (τ_T^h, τ_T^l) is an optimal contract-pair, then the contract τ_T^l pays an up-front amount $M(\tau_T^h) = u^l(\tau_T^h, \text{sub}_\emptyset)$ followed by $c/\pi_{(l,e)}$ at t if the task at t is successful.*

Proof. Assume, by way of contradiction, that (τ_T^h, τ_T^l) is an optimal contract-pair, but in the contract τ_T^l is not the one described in the statement of the Lemma. For a given contract $\tau_T^h \in \hat{\Gamma}_T^h$, denote by $\hat{\tau}_T^l$ the contract that pays an up-front amount $M(\tau_T^h) = u^l(\tau_T^h, \text{sub}_\emptyset)$ and then $c/\pi_{(l,e)}$ at t if the task at t is successful. Note that this contract is the least costly contract to the low quality agent that generates him utility of $u^l(\tau_T^h, \text{sub}_\emptyset)$ and satisfies *EF*. To see this observe that postponing any payment after the first period will require additional compensation due to the difference in the discount factors between the principal and the agent. Therefore, since $\hat{\tau}_T^l$ is not part of the optimal contract-pair, it must be the case that the utility of the high-quality agent with contract $\hat{\tau}_T^l$ is higher than with contract τ_T^h . Since the expected utility of the high quality agent in contract $\hat{\tau}_T^l$ is

$$M(\tau_T^h) + c \frac{\pi_{(h,e)} - \pi_{(l,e)}}{\pi_{(l,e)}} (1-q) (1-\delta^T) \frac{\delta}{1-\delta},$$

it must be the case that

$$M(\tau_T^h) > c \left(\frac{\pi_{(h,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} - \frac{\pi_{(h,e)} - \pi_{(l,e)}}{\pi_{(l,e)}} \right) (1-q) (1-\delta^T) \frac{\delta}{1-\delta}.$$

Consider then changing both contracts in the following way: $\tau_T^h = \underline{\tau}_T^h$ is the base-line contract, while the contract to the low-quality agent is

$$M(\underline{\tau}_T^h) = u^l(\underline{\tau}_T^h, \text{sub}_\emptyset) = c \frac{\pi_{(l,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} (1-q) (1-\delta^T) \frac{\delta}{1-\delta}$$

and $c/\pi_{(l,e)}$ at t if the task at t is successful. Since (as was shown in the proof of the previous Lemma)

$$M(\underline{\tau}_T^h) = c \frac{\pi_{(l,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} (1-q) (1-\delta^T) \frac{\delta}{1-\delta} < c \left(\frac{\pi_{(h,e)} - \pi_{(h,d)}}{\pi_{(h,d)}} - \frac{\pi_{(h,e)} - \pi_{(l,e)}}{\pi_{(l,e)}} \right) (1-q) (1-\delta^T) \frac{\delta}{1-\delta}$$

this contract-pair is admissible. Moreover, the expected payment to both types of agents is lower because the up-front payment to the low-quality agent is lower, while the per-success payment remained the same and the expected payment to the high-quality agent is the minimal one that respects the efficient incentive provision. This is a contradiction. ■

10.2 Three-period example

We now show that there exist parameters for which the two following inequalities

$$\mu\delta^2 (q + (1 - q)(1 - \pi_{(l,e)}))(1 - q)\pi_{(l,e)} \left(1 - \frac{(1 - q)\pi_{(l,e)}}{p^h}\right) - (1 - \mu) (1 - p^h) \frac{1 - \delta}{\delta} > 0$$

and

$$\mu\delta(1 - q)\pi_{(l,e)} \left(1 - \frac{(1 - q)\pi_{(l,e)}}{p^h}\right) - (1 - \mu) \frac{1 - \delta}{\delta} \leq 0$$

hold. Note first that the second inequality holds with equality when

$$\mu = \frac{\frac{1 - \delta}{\delta}}{\frac{1 - \delta}{\delta} + \delta(1 - q)\pi_{(l,e)} \left(1 - \frac{(1 - q)\pi_{(l,e)}}{p^h}\right)} = \mu^*,$$

and holds as a strict inequality for $\mu \leq \mu^*$. We next verify that there exist parameters for which, when $\mu \leq \mu^*$, the first inequality holds as well. Plugging the expression of μ^* into the first inequality and rearranging gives us

$$\frac{(1 - \delta)(1 - q)\pi_{(l,e)} \left(1 - \frac{(1 - q)\pi_{(l,e)}}{p^h}\right) [\delta(q + (1 - q)(1 - \pi_{(l,e)})) - 1 + p^h]}{\frac{1 - \delta}{\delta} + \delta(1 - q)\pi_{(l,e)} \left(1 - \frac{(1 - q)\pi_{(l,e)}}{p^h}\right)} > 0;$$

that is, when $\mu = \mu^*$, the first inequality is satisfied if $\delta(q + (1 - q)(1 - \pi_{(l,e)})) + p^h > 1$. That is,

$$q\pi_{(h,d)} + (1 - q)\pi_{(h,e)} > 1 - \delta(q + (1 - q)(1 - \pi_{(l,e)})).$$

For $\delta = 1$, this inequality requires that when the high-quality agent accepts both types of tasks, and the low-quality one accepts only easy tasks, the prior probability to see a success in a given period is always higher for the high-quality agent. Since $\pi_{(h,e)} > \pi_{(l,e)}$, this inequality holds for δ high enough.

10.3 Proof of Proposition 3

Assume, by way of contradiction, that there exists t such that after some history the payment for success in t is not postponed, while after $t' > t$ that follows the same history the payment is postponed for $d > 0$ periods. (Recall that whenever a payment is postponed for some k periods, it is paid conditional on successes in all periods in between.) Assume that the history is such that until time t there

were k failures and l successes, while between t and t' there were m failures and n successes (in addition to success in period t). The cost of postponing ε for d periods after t' is

$$(1 - \mu) (1 - p^h)^{k+m} (p^h)^{l+n+1} \left(\frac{1 - \delta^d}{\delta^d} \right) \varepsilon$$

while the benefit is

$$\begin{aligned} & \mu [q + (1 - q) (1 - \pi_{(l,e)})]^{t''} (1 - q)^{l+n+1} (\pi_{(l,e)})^{l+n+1} \delta^{k+l+m+n+2} \times \\ & \left[(1 - q) \pi_{(l,e)} \varepsilon - (1 - q)^{d+1} (\pi_{(l,e)})^{d+1} \frac{\varepsilon}{(p^h)^d} \right] \end{aligned}$$

where $t'' \leq k + m$ is the number of periods in which the low-quality agent failed while using the strategy of exerting effort if and only if an easy task arrives.

Next consider the effect of postponing the payment for success in period t for d periods. The cost of postponement of an amount ε is

$$(1 - \mu) (1 - p^h)^k (p^h)^l \left(\frac{1 - \delta^d}{\delta^d} \right) \varepsilon$$

while the benefit of the postponement is

$$\mu [q + (1 - q) (1 - \pi_{(l,e)})]^{t'''} (1 - q)^l (\pi_{(l,e)})^l \delta^{k+l+1} \left[(1 - q) \pi_{(l,e)} \varepsilon - (1 - q)^{d+1} (\pi_{(l,e)})^{d+1} \frac{\varepsilon}{(p^h)^d} \right]$$

where $t''' \leq k$ is the number of periods in which the low-quality agent failed while using the strategy of exerting effort if and only if an easy task arrives. Note that $t''' \leq t''$.

Recall that in the original mechanism a payment is postponed for d periods after period t' . Consider then decreasing the postponed amount after period t' by ε and instead postponing the amount ε' for d periods after period t where ε' is chosen such that the expected benefits are the same. That is,

$$\begin{aligned} 0 &= \mu [q + (1 - q) (1 - \pi_{(l,e)})]^{t'''} (1 - q)^l (\pi_{(l,e)})^l \delta^{k+l+1} \times \\ & \left[(1 - q) \pi_{(l,e)} \varepsilon' - (1 - q)^{d+1} (\pi_{(l,e)})^{d+1} \frac{\varepsilon'}{(p^h)^d} \right] \\ & - \mu [q + (1 - q) (1 - \pi_{(l,e)})]^{t''} (1 - q)^{l+n+1} (\pi_{(l,e)})^{l+n+1} \delta^{k+l+m+n+2} \times \\ & \left[(1 - q) \pi_{(l,e)} \varepsilon - (1 - q)^{d+1} (\pi_{(l,e)})^{d+1} \frac{\varepsilon}{(p^h)^d} \right] \\ & = \mu [q + (1 - q) (1 - \pi_{(l,e)})]^{t'''} (1 - q)^l (\pi_{(l,e)})^l \delta^{k+l+1} (1 - q) \pi_{(l,e)} \\ & \left\{ \varepsilon' - (1 - q)^d (\pi_{(l,e)})^d \frac{\varepsilon'}{(p^h)^d} - \right. \\ & \left. [q + (1 - q) (1 - \pi_{(l,e)})]^{t''-t'''} (1 - q)^{n+1} (\pi_{(l,e)})^{n+1} \delta^{m+n+1} \left(\varepsilon - (1 - q)^d (\pi_{(l,e)})^d \frac{\varepsilon}{(p^h)^d} \right) \right\} \end{aligned}$$

which can be written as

$$0 = \mu [q + (1 - q)(1 - \pi_{(l,e)})]^{t''} (1 - q)^l (\pi_{(l,e)})^l \delta^{k+l+1} (1 - q) \pi_{(l,e)} \times \\ \left\{ \varepsilon' - (1 - q)^d (\pi_{(l,e)})^d \frac{\varepsilon'}{(p^h)^d} - \right. \\ \left. [q + (1 - q)(1 - \pi_{(l,e)})]^{t''-t''} (1 - q)^{n+1} (\pi_{(l,e)})^{n+1} \delta^{m+n+1} \left(\varepsilon - (1 - q)^d (\pi_{(l,e)})^d \frac{\varepsilon}{(p^h)^d} \right) \right\}$$

or

$$0 = \varepsilon' \left[1 - \left(\frac{(\pi_{(l,e)})(1 - q)}{(p^h)^d} \right)^d \right] - \\ \varepsilon \left[1 - \left(\frac{(\pi_{(l,e)})(1 - q)}{(p^h)^d} \right)^d \right] [q + (1 - q)(1 - \pi_{(l,e)})]^{t''-t''} (1 - q)^{n+1} (\pi_{(l,e)})^{n+1} \delta^{m+n+1}.$$

That is,

$$\varepsilon' = [q + (1 - q)(1 - \pi_{(l,e)})]^{t''-t''} (1 - q)^{n+1} (\pi_{(l,e)})^{n+1} \delta^{m+n+1} \varepsilon.$$

Note that while this change does not affect the utility of the agent, the change in the expected costs are

$$(1 - \mu) (1 - p^h)^k (p^h)^l \left(\frac{1 - \delta^d}{\delta^d} \right) \varepsilon' - (1 - \mu) (1 - p^h)^{k+m} (p^h)^{l+n+1} \left(\frac{1 - \delta^d}{\delta^d} \right) \varepsilon \\ = (1 - \mu) (1 - p^h)^k (p^h)^l \left(\frac{1 - \delta^d}{\delta^d} \right) \varepsilon \times \\ \left[[q + (1 - q)(1 - \pi_{(l,e)})]^{t''-t''} (1 - q)^{n+1} (\pi_{(l,e)})^{n+1} \delta^{m+n+1} - (1 - p^h)^m (p^h)^{n+1} \right] < 0$$

where the first inequality follows from plugging the expression of ε' and the last inequality follows from the inequalities $p^h > (1 - q)\pi_{(l,e)}$, $1 - p^h \geq \delta$, and $t'' \leq t''$. Therefore, the assumed change decreases the cost of the principal while preserving the utilities of the low-quality agent. Contradiction to the assumed optimality of the mechanism.

References

- [1] Baron, D., & Besanko, D. (1984): "Regulation and Information in a Continuing Relationship," *Information Economics and Policy* 1, 267–302.
- [2] Biais, B., Mariotti, T., Plantin, G., & Rochet, J.-C. (2007): "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications," *Review of Economic Studies* 74, 345–390.

- [3] Biais, B., Mariotti, T., Rochet, J.-C., & Villeneuve, S. (2010): “Large Risks, Limited Liability, and Dynamic Moral Hazard,” *Econometrica* 78, 73–118.
- [4] DeMarzo, P., & Fishman, M. (2007): “Optimal Long-term Financial Contracting,” *Review of Financial Studies* 20, 2079–2128.
- [5] DeMarzo, P., & Sannikov, Y. (2006): “Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model,” *Journal of Finance* 61, 2681–2724.
- [6] Edmans, A., Gabaix, X., Sadzik, T., & Sannikov, Y. (2010): “Dynamic Incentive Accounts,” mimeo.
- [7] Fong, K. (2009): “Evaluating Skilled Experts: Optimal Scoring Rules for Agents,” mimeo.
- [8] Fudenberg, D., Holmstrom, B., & Milgrom, P. (1990): “Short-term Contracts and Long-term Agency Relationships,” *Journal of Economic Theory* 51, 1–31.
- [9] Holmstrom, B. (1999): “Managerial Incentive Problems: A Dynamic Perspective,” *Review of Economic Studies* 66, 169–182.
- [10] Holmstrom, B., & Milgrom, P. (1987): “Aggregation and Linearity in the Provision of Intertemporal Incentives,” *Econometrica* 55, 303–328.
- [11] Laffont, J., & Tirole J. (1988): “The Dynamic of Incentive Contracts,” *Econometrica* 56, 1153–1175.
- [12] Malcomson, J., & Spinnewyn, F. (1988): “The Multi-period Principal-agent Problem,” *Review of Economic Studies* 55, 391–407.
- [13] Rey, P., & Salanie, B. (1990): “Long-term, Short-term, and Renegotiation: On the Value of Commitment in Contracting,” *Econometrica* 58, 597–619.
- [14] Rubinstein, A., & Yaari, M. (1983): “Repeated Insurance Contracts and Moral Hazard,” *Journal of Economic Theory* 30, 74–97.
- [15] Sannikov, Y. (2007): “Agency Problems, Screening and Increasing Credit Lines,” mimeo.
- [16] Sannikov, Y. (2008): “A Continuous-time Version of the Principal-agent Problem,” *Review of Economic Studies* 75, 957–984.
- [17] Stanica, P. (2001): “Good Lower and Upper Bounds on Binomial Coefficients,” *Journal of Inequalities in Pure and Applied Mathematics* 2, 1–5.