"LEARNING-BY-SHIRKING" IN RELATIONAL CONTRACTS*

JIN LI†, ARIJIT MUKHERJEE‡, LUIS VASCONCELOS§

ABSTRACT. A worker may shirk on some of the aspects of his job in order to privately learn which ones are more critical and use this information in the future to shirk more effectively. This possibility of private learning aggravates the moral hazard problem. We study the optimal provision of relational incentives to deter the worker from such "learning-by-shirking." The firm strategically discloses information on the role of each job aspect to sharpen incentives and the optimal disclosure policy depends on the surplus in the relationship. Depending on the underlying parameters, the optimal policy may call for opacity, full disclosure, as well as partial disclosure through stochastic revelation.

1. INTRODUCTION

A common incentive problem in many employment relationships is that a worker may be tempted to cut corners at the expense of the firm. The worker, however, is often unsure about the consequences of cutting corners, and relatedly which corners to cut if the job contains many aspects. For example, the worker may not know which aspects are more critical for the success and which are not. As a result, he may start out shirking "in the dark". Importantly, once the worker shirks successfully, he learns how to shirk more effectively in the future. Such "learning-by-shirking" aggravates the moral hazard problem because shirking, when successful, provides valuable private information to the worker that he may use later on to cut corners that are harder to notice.

This type of problem is particularly prevalent in jobs where the agents are expected to perform a list of activities. Hospitals often require healthcare providers to follow a set of protocols to prevent infections, aviation companies ask pilots to follow a list of procedures to ensure that the aircrafts takeoff and land safely, civil engineers typically follow a detailed task list to avoid structural defects (see Gawande, 2010). In such a job, all associated tasks or steps may not be equally critical but omission of a critical step may lead to major failure. Moreover, due to the complexity of the job process, the agents may not readily know which steps are more essential than the others. Similar issues may also emerge in client-service

Date: August 14, 2016.

*For their helpful comments and suggestions, we thank Dan Barron, Alessandro Bonatti, Gonzalo Cisternas, Jon Eguia, William Fuchs, Robert Gibbons, Johannes Hörner, Alessandro Lizzeri, Niko Matouschek, Michael Powell, Canice Prendergast, Rani Spiegler, Jeroen Swinkels, Huanxing Yang, and the seminar audiences at Boston University, Michigan State University, Southern Methodist University, and University of Essex. Any errors that may remain are ours.

†Strategy Department, Kellogg School of Management, Northwestern University, Leverone Hall, 6th Floor, 2001 Sheridan Road, Evanston, IL 60208. Email: jin-li@kellogg.northwestern.edu. URL: http://www.kellogg.northwestern.edu/Faculty/Directory/Li_Jin.aspx.

‡Department of Economics, Michigan State University, 110 Marshall Adams Hall, East Lansing, MI 48823, USA. Email: arijit@msu.edu. URL: www.amukherjee.net.

§Department of Economics, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, UK. E-mail: lvasco@essex.ac.uk.

firms such as law and consulting. A consultant may work on a project where the client wants a detailed report within a deadline. The client is surely pleased if a thorough report is produced in time but he may care more about the quality of the report than its timely completion. But as the client's preference on these two project requirements is his private information, the agent need not know which aspect is more critical.

In this article we analyze the optimal way to provide incentives in an on-going employment relationship when the worker may shirk to learn. In complex jobs with multiple aspects verifiable performance measures well-aligned with the firm's goal could be difficult to obtain and incentives are often tied to non-verifiable measures. Following this motivation, we model the long-term employment relationship as a relational contract where the firm offers incentives through a discretionary bonus that is tied to a non-verifiable performance metric (see Malcomson, 2013). In particular, we ask the following question: when the firm could publicly obtain and divulge information on the relative importance of different job aspects, how should it reveal this information to the worker in order to sharpen incentives?

We consider a model of an infinitely repeated employment relationship where the worker performs on a job that consists of two tasks (or aspects). The first-best outcome requires the worker to exert (costly) effort in both tasks to ensure job success. The firm cannot measure the worker's performance in each task but it can observe a non-verifiable measure of the worker's overall job performance. A key aspect of the model is that one of the two tasks plays a more critical role in the production process: if the worker only performs this task and shirks on the other, with some probability, the job may still be successfully completed. At the beginning of the relationship the players do not know which task is crucial. However, the worker may privately acquire this information if he shirks and happens to pick the right task. This possibility of learning-by-shirking aggravates the moral hazard problem. The worker may be tempted to shirk not only to save on his effort cost but also to learn which task is more crucial so that he can use this information again in the future to shirk more effectively.

Incentives are provided through two channels. First, as mentioned above, the firm promises a discretionary bonus to the worker for a satisfactory job performance. As the job performance measure is not verifiable, such a promise is a part of relational contract that is sustained through a threat of future retaliation by the worker should the firm renege on its bonus payment. Therefore, in such a contract the maximum bonus that the firm promises cannot exceed the future surplus generated by the relationship (Levin, 2003). Both players discount future payoffs at a rate $\delta \in (0, 1)$. Second, the firm may also attempt to sharpen such relational incentives by accessing and voluntarily revealing information on the identity of the critical task, either at the beginning of the game or at the end of any given period if the information has not been revealed in the past. (To fix ideas, suppose that the firm may hire an expert to review its production process and identify its most essential components.)

Our analysis highlights how the amount of the surplus in the relationship—i.e., the firm's "reputational capital"—affects the firm's information disclosure policy regarding the relative priorities of the different tasks the constitute the agent's job. It turns out that if δ is too large or too small, the firm's disclosure policy does not play any major role in its incentive provision. For δ sufficiently large the firm can credibly promise a large enough bonus pay and the first-best could be attained irrespective of its disclosure policy; for δ sufficiently small, no effort could be induced as the lack of reputational capital renders any promise of bonus pay non-credible.

The more interesting case is when δ, and consequently, the firm's stock of reputational capital is in an intermediate range. Here, the optimal contract requires the firm to actively manage information to sharpen incentives.

In particular, when δ is relatively large (but still within the intermediate range), opacity is essential for attaining efficiency. The firm never reveals the critical task and implements the first-best by credibly promising a bonus payment. The intuition is as follows: When the worker knows the identity of the critical task, his temptation to shirk is stronger as he knows ex-ante which task to neglect. Therefore, the firm must offer stronger incentives to induce effort on both tasks. However, the lack of knowledge on which task is critical lowers the worker's expected payoff from shirking as he risks shirking on the critical task and getting caught by his employer. Thus, the bonus pay needed to induce the worker to work on both tasks need not have to be too large. When δ is relatively large, the maximal bonus that the firm could credibly promise may not be sufficient to induce efficiency when the worker knows which task is critical but may be large enough to induce efficiency when he does not know which task is more crucial.

In contrast, for δ relatively small (but still within the intermediate range), the optimal contract calls for full transparency. The firm must reveal the critical task at the beginning of the game and the worker performs the critical task only in all periods. In this parameter range, the maximal bonus the firm could credibly offer is also relatively small and it can induce effort in at most one of the two tasks. Hence, the firm would rather tell the worker to focus on the critical task only (and settle for a lower expected output compared to the first-best).

Finally, for moderate δ (within the intermediate range) a more nuanced disclosure policy is called for. The firm filters information by adopting a stochastic disclosure rule: At the end of every period it discloses the critical task with a constant probability (if it has not been made public yet). The worker exerts effort on both tasks until the critical task is revealed, but once it is revealed, the worker only works on the critical task only in all future periods. For such a δ the firm lacks the reputational capital needed to induce efficiency even when the critical task is not known to the worker but its reputational capital stock is still larger than what is needed to elicit effort in exactly one of the two tasks. Clearly, the firm can reveal the critical task right away and ask for effort in that task only. But it can do better. The stochastic disclosure policy dissuades the worker from "learning-by-shirking" by diluting the information advantage the worker may have over the firm when he privately learns how to cut corners. As the worker's gains from superior information may last only for a short period of time, his incentive to shirk diminishes. Consequently, such a disclosure policy bolster incentives and allows the firm to elicit effort in both tasks until the critical task gets revealed.[1]

Our analysis relates to several other human resource management policies that deter the worker from "learning-by-shirking" by diluting the value of the information that the worker may obtain in the process. For example, the firms often adopt job rotation policies where the workers are moved to a different assignment after every so few years within the firm.[2]

[1]The analysis of this case presents an interesting technical challenge as the standard recursive techniques à la Abreu et. al (1990) cannot be applied. We will elaborate on this later in our discussion on the related literature.

[2]Such a policy is routinely adopted by government organizations in many countries where the civil servants are rotated among multiple locations as an anti-corruption measure. Staying in the same location for too

We consider a variation of our main model to analyze such a job rotation policy and show that the optimal policy mirrors the optimal information disclosure policy discussed above as long as the cost of job rotation is not too large. It is worthwhile to note that a job rotation need not entail moving the worker among different geographic locations. Frequent reorganization of the job responsibilities may have a similar role in discouraging the worker from "learning-by-shirking" as any private information learned in the current job would be useless in the next.

Related Literature: This article contributes to the literature that studies how firms can use various instruments (other than discretionary bonuses) to increase the efficiency of relational contracting. Such instruments may include formal contracts (Baker, Gibbons, and Murphy, 1994), integration decision (Baker, Gibbons, and Murphy, 2002), design of ownership (Rayo, 2007), job design (Schöttner, 2008; Mukherjee and Vasconcelos, 2011; Ishihara, 2016), design of peer evaluation (Deb, Li, and Mukherjee, 2016), or delegation decisions (Li, Matouscheck and Powell, forthcoming).

The two most closely related articles in this literature are Fuchs (2007) and Fong and Li (forthcoming). As in this paper, both of these articles show that the efficiency of the relationship can be improved when less information is known. Fuchs demonstrates that by revealing the agent's performance measures every $T > 1$ periods (rather than every period), the principal can reduce the amount of surplus destroyed in the relationship. Fong and Li consider a principal-supervisor-agent relationship, and show that rather than giving truthful reports of the agent's performances to the principal, the supervisor can improve the efficiency of the relationship by sending less informative reports that help spread the reward for the agent's good performance across periods. In this paper, reduction of the information improves the efficiency by making it harder for the agent to shirk.

The articles mentioned above (including ours) share the feature that payoff-relevant information is not known publicly in every period, and therefore, the standard recursive technique à la Abreu, Pearce, and Stacchetti (1990) cannot be readily applied. The recursive structure is lost because when the agent deviates, he may be able to learn about payoff-relevant information that is not known to the principal. The loss of common knowledge implies that the multi-stage deviations must be checked to ensure that a strategy profile is an equilibrium, complicating the analysis significantly. The technical difficulty makes it in general very difficult to characterize the optimal information structure.

We are able to characterize the optimal information structure here through the introduction of an auxiliary state variable. Whereas the agent's (equilibrium) continuation payoff is typically the only state variable in standard recursive problems, our state variable also includes the agent's maximal (off-equilibrium) continuation payoff. The augmented state variable makes the problem recursive, allowing us to characterize the optimal information structure. In contrast to the existing models of relational contracting in which the information structures are typically stationary, our information structure shows that the relationship can benefit from delaying the revelation of payoff relevant information.

Our work broadly relates to several other strands of literature. There is a vast literature on long-term contracting (both relational and explicit) with persistent private information (e.g., Battaglini, 2005; Yang, 2013; Malcomson, 2016). But the focus of this literature is

long may allow the officials to become too cozy with their coworkers and some may attempt to learn how to collude and cover up a corruption racket (Bardhan, 1997).

on how to elicit and use this information in the optimal contract. In contrast, we study an environment where the agent gains such persistent private information upon deviating from the equilibrium play and the incentive problem at hand is to deter the agent from undertaking such a deviation.

A related stream of work on long-term contracting considers ex-ante symmetric information but allow the agent to manipulate the principal's posterior belief about the production environment (e.g., feasibility of a given project) through private action, see, e.g., Bergmann and Hege, 2005; Bonatti and Hörner, 2011; and Bhaskar 2014. In these models, as is the case in our setting, the posterior beliefs of the contracting parties diverge and cease to be common knowledge following a deviation by the agent. But this literature assumes the production process to be inherently uncertain; e.g., the feasibility of a given output level may depend on the underlying state of the world that is (ex-ante) unknown to all parties. And, in contrast to our setup, the key focus of these articles is to explore how the learning about the state (through the history of past production levels) shapes the dynamics of incentive provision.

Several authors have also studied strategic information disclosure in employment relationships. This literature primarily focuses on two kinds of information: the employer's private information on the agents' performance (e.g., Aoyagi, 2010; Ederer, 2010; Mukherjee, 2010; Goltsman and Mukherjee, 2011; Zabojnik, 2014; Orlov, 2016) and information on the compensation rule used by employers—i.e., what aspects of the performance are measured and how these measures affect the incentive pay (see Lazear, 2006, and Ederer, Holden and Meyer, 2014). In this literature our setting comes closest to that analyzed in Lazear (2006). Lazear considers environments where the agent's actions in all aspects of the job cannot be monitored for exogenous reasons (e.g., a test may not cover all topics taught in the course) and asks when it pays to reveal in advance what is being measured. The answer, it turns out, depends on the agent's cost of action and the principal's cost of monitoring. In our setting, however, the firm's private information is about the underlying production technology. We explore how transparency affects the feasibility relational incentives and highlight the role of the firm's reputational capital in driving its disclosure policy.

Finally, our work is also related to the literature on incentives for experimentation (see Manso, 2011; Hörner and Samuelson, 2013; Bonatti and Hörner, 2015; Halac, Kartik, and Liu, 2016; Moroni, 2016; Guo, 2016). While most of these articles do not consider relational incentives a recent exception is Chassang (2010). He analyzes experimentation in relational contracting and argues that moral hazard in experimentation, together with the principal's inability to commit, can result in a range of different actions being adopted in the long run. Note that in contrast to these settings, the incentive problem we focus on is about designing the relationship in order to discourage the agent from experimentation (i.e., selectively perform a subset of tasks to learn about the production technology). Indeed, experimentation does not occur along the equilibrium path in our model.

The rest of the paper is structured as follows. Section 2 describes our baseline model that focuses on information revelation. A benchmark case is analyzed in section 3 where the identity of the critical task is assumed to be public information. The optimal revelation of task information is studied in Section 4. In section 5 we adapt our baseline model to show how job rotation, instead of information revelation, could be used to address the moral hazard problem in our environment. A final section concludes. All proofs are provided in the Appendix.

2. Model

A principal (or "firm") P hires an agent (or "worker") A where the two parties enter in an infinitely repeated employment relationship. Time is discrete and denoted as $t \in \{1, 2, ..., \infty\}$. In each period, the firm and the agent play a stage game that is described as follows.

Stage game: We describe the stage game in terms of its three key components: *Technology, contracts,* and *payoffs.*

Technology: In any period t, the agent may be asked by the principal to perform a job that consists of two tasks: task $\mathbb{A}$ and $\mathbb{B}$. The agent must exert effort to complete the job. The effort $e_t \in \{0, 1_{\mathbb{A}}, 1_{\mathbb{B}}, 2\}$ is privately chosen by the agent and not observed by the principal. Effort is costly to the agent and we denote the cost by the function $C(e_t)$. If the agent works in both tasks $e_t = 2$ and his cost of effort is $C(2) = c_2$ but if he works on any one of the two tasks, $e_t = 1_{\mathbb{A}}$ or $1_{\mathbb{B}}$ depending on whether he works only on task $\mathbb{A}$ or on $\mathbb{B}$, and his cost of effort is $C(1_{\mathbb{A}}) = C(1_{\mathbb{B}}) = c_1$ $(< c_2)$. Also, if he shirks on both tasks then $e_t = 0$ and the cost of effort $C(0) = 0$.

The output in the job, $Y_t \in \{-z, 0, y\}$, is assumed to be observable but not verifiable. The job may be successfully completed leading to an output $y > 0$, it may remain incomplete where the output is 0, and, at the extreme, the agent may completely fail at the job leading to a negative output $-z$ (e.g., such a failure may lead to an erosion of the firm's market value).

If the agent exerts effort on both tasks, the output is always y and if the agent shirks on both tasks, the output is always $-z$. But if the agent works on only one of the two tasks, the outcome depends on which task he works on. In particular, one of the two tasks is "critical" whereas the other is not. If the agent works only on the critical task, the output is y with probability p (> 0) and 0 with probability $1 - p$. But if the agent works only on the non-critical task, the output is $-z$ with certainty. As we will see below, the negative payoff associated with the non-critical task ensures that any "experimentation" by the agent where he randomly selects and performs only one out of the two tasks is detrimental to the relationship.

The identity of the critical task is governed by the underlying production technology and remains unchanged throughout the game. At the beginning of the game, neither party knows which task is critical and both players hold a common prior belief that any of the two tasks could be critical with equal likelihood. However, if the principal so prefers, she may publicly obtain and reveal information on the identity of the critical task, either at the beginning of the game or at the end of any period. Note that the production technology described above implies that if the agent shirks by exerting effort in only one task, he may privately learn the identity of the "critical" task if he happens to pick the right task by chance. As we will see below, such possibility of private "learning-by-shirking" has significant implications for the optimal relational contract.

Contract: As the output is non-verifiable, the principal cannot offer an explicit pay-per-performance contract. In each period t, the principal decides on whether to offer a contract to the agent. We denote the principal's offer decision as $d_t^P \in \{0, 1\}$ where $d_t^P = 0$ if no offer is made and $d_t^P = 1$ otherwise. If the principal decides to make an offer, she offers a contract that specifies a commitment of wage payment w_t and a discretionary bonus $b_t = b_t(Y_t)$ that

is paid only if $Y_t = y$.[3] The principal may also reveal the identity of the critical task at the beginning of period one or may promise to reveal the task (either deterministically or stochastically) at the end of a given period. We denote the public availability of the task information by $\gamma_t \in \{0,1\}$ where $\gamma_t = 1$ if the critical task is revealed by the principal in or before period t and $\gamma_t = 0$ otherwise.[4] We assume that the identity of a task is hard information that the principal cannot mis-report should she decide to reveal such information.

The agent either accepts or rejects the contract. We denote the agent's decision as $d_t^A \in \{0,1\}$ where $d_t^A = 0$ if the offer is rejected and $d_t^A = 1$ if it is accepted. Upon accepting the offer, the agent decides on his effort level—whether to work on both tasks, shirk on both tasks, or choose one of the two tasks and work only on it.

Finally, as typical in the repeated game literature, we will assume the existence of a public randomization device to convexify the equilibrium payoff set. We assume that at the end of each period t, the principal and the agent publicly observe the realization x_t of a randomization device. This realization allows the players to publicly randomize their actions in period $t+1$. In addition, a realization x_0 is also assumed to be publicly observed at the beginning of period 1 allowing the players to randomize in period 1 as well.

PAYOFFS: Both the principal and the agent are risk neutral. If either d_t^A or d_t^P is 0, both players take their outside options in that period and the game moves on to period $t+1$. Without loss of generality, we assume that both players' outside options are 0. If $d_t^A = d_t^P = 1$, for a given effort level e_t, the payoffs are as follows. The agent and the principal earn,

$$\hat{u}_t = w_t + \mathbb{E}\left[b_t\left(Y_t\right) \mid e_t\right] - C\left(e_t\right) \quad \text{and} \quad \hat{\pi}_t = \mathbb{E}\left[Y_t - w_t - b_t\left(Y_t\right) \mid e_t\right],$$

respectively.

We assume the following restrictions on the parameters.

Assumption 1. *(i)* $y - c_2 > py - c_1 > 0$*, and (ii)* $\frac{1}{2}pc_2 > c_1$.

Under the above assumption, efficiency requires the agent to exert effort on both tasks, and exerting effort only on the critical task is more efficient than dissolving the employment relationship. Moreover, the cost of exerting effort on both tasks relative to only one is assumed to be sufficiently large.

TIME LINE: The time line of the stage game in any period t is summarized below:

- *Beginning of period* t. P decides whether to make an offer to A. If she makes an offer, she specifies the contract $(w_t, B_t\left(Y_t\right))$ and the game continues to stage $t.1$. Moreover, if $t = 1$, she decides on whether to reveal the critical task.

[3]Note that this contract specification is without loss of generality as $Y_t = -z$ is never realized on the equilibrium path. This is due to our assumption that it is never optimal for the principal to allow the agent to experiment (see Assumption 2 below).

[4]As the identity of the critical task is persistent over time, once the principal reveals the critical task, it is known to the agent in all future periods. That is, if $\gamma_\tau = 1$ for some τ, then $\gamma_t = 1$ for all $t > \tau$. So, the principal's revelation decision in any period is payoff revelant only if the task has not been disclosed in the past.

- *Period* $t.1$. A accepts or rejects the contract. If he accepts, the game continues to period $t.2$.
- *Period* $t.2$. A chooses effort e_t.
- *Period* $t.3$. P and A observe output Y_t.
- *Period* $t.4$. P pays wage w and decides on bonus payment.
- *End of period* t. Public randomization device x_t is realized. Principal decides on γ_t, i.e., whether to reveal the critical task if she has not revealed it yet, and the game moves to period $t+1$.

Repeated game: The stage game described above is repeated every period and players are assumed to have a common discount factor $\delta \in (0,1)$. At the beginning of any period t, the average payoffs of the agent and the principal in the continuation game are given by:

$$u^t = (1-\delta)\sum_{\tau=t}^{\infty}\delta^{\tau-t}\left[d_\tau \hat{u}_\tau\right] \quad \text{and} \quad \pi^t = (1-\delta)\sum_{\tau=t}^{\infty}\delta^{\tau-t}\left[d_\tau \hat{\pi}_\tau\right],$$

respectively, where $d_\tau := d_\tau^A d_\tau^P$.

Strategies and equilibrium: The extant literature defines a relational contract as a pure strategy public Perfect Equilibrium (PPE) where the players only use public strategies and the equilibrium strategies induce a Nash Equilibrium in the continuation game starting from each public history (Levin, 2003). It is important to note that in our setting, we must account for the fact that the agent may privately learn about the identity of the critical task from his past deviation and would find it profitable to use this information in future deviations. Thus, restrictions to pure strategy PPE may lead to some loss of generality and instead, we focus on the (mixed strategy) perfect Bayesian Equilibrium (PBE) of the game defined as follows:

Let $h_{t+1} = \left\{x_\tau, \gamma_\tau, d_\tau^A, d_\tau^P, Y_\tau, w_\tau, b_\tau\right\}_{\tau=1}^{t}$ denote the public history of the game at the beginning of period $t+1$ and H_{t+1} be the set of all such histories (note that, $H_1 = \{x_0\}$). The public strategy of the principal consists of a sequence of functions $\sigma_P = \left\{D_t^P, \Gamma_t, W_t, B_t\right\}_{t=1}^{\infty}$ where her participation decision is given by $D_t^P : H_t \to \Delta\{0,1\}$, the task revelation decision is given as $\Gamma_t : H_t \to \Delta\{0,1\}$, and the contract offer is given as $W_t : H_t \to \mathbb{R}$ and $B_t : H_t \times \{-z, 0, y\} \to \mathbb{R}$. Similarly, the public strategy for the agent is a sequence of functions $\sigma_A = \left\{D_t^A, E_t\right\}_{t=1}^{\infty}$ where his participation decision is given as $D_t^A : H_t \times \{0,1\} \to \Delta\{0,1\}$ that assigns to each pair (h_t, γ_t) a probability of accepting the contract offered by the principal, and his effort decision is given as $E_t : H_t \times \{0,1\}^3 \times \mathbb{R}^2 \to \{0, 1_{\mathbb{A}}, 1_{\mathbb{B}}, 2\}$ that assigns an effort level to each tuple $(h_t, \gamma_t, d_t^A, d_t^P, w_t, b_t)$. Finally, denote $\mu_t = \Pr(\textit{task } \mathbb{A} \textit{ is crucial})$ as the belief of the agent in period t. Note that $\mu_t = \frac{1}{2}$ if the agent does not have any information on which task is critical and it is either 0 or 1 if $\gamma_t = 1$ and/or if the agent has privately learned the identity of the critical task by shirking in the past.

A profile of strategies $\sigma^* = \langle \sigma_P^*, \sigma_A^* \rangle$ along with a belief $\mu^* = \{\mu_t^*\}_{t=1}^{\infty}$ constitute a PBE of this game if σ^* is sequentially rational and μ_t^* is consistent with σ^* and derived using Bayes rule whenever possible. We define an "optimal" or "efficient" relational contract as a PBE of this game where the payoffs are not Pareto-dominated by any other PBE.

In what follows, we assume that it is never optimal for the relationship to have the agent to experiment where he randomly selects one of the two tasks and performs the selected task

only. A sufficient condition for this to be the case is given by Assumption 2 below and it is trivially satisfied when z is sufficiently large.[5]

Assumption 2. $(1-\delta)\left(\frac{1}{2}(py-z)-c_1\right)+\delta(y-c_2)<0$.

The information revelation about the identity of the critical task plays a critical role in our analysis as the optimal contract must specify a revelation mechanism that maximizes the joint surplus of the principal and the agent. To clearly illustrate the role of such information revelation, we first consider a benchmark case where the identity of the critical task is assumed to be publicly known from the very beginning of the game. Next, we present a general characterization of the optimal contract when the critical task unknown to both parties but the principal can publicly obtain and reveal this information to strengthen incentives.[6]

It is worthwhile to note that while the above model highlights strategic revelation of task information as an incentive device, the model could be easily adapted to study other policies, e.g., job rotation, that may have the similar incentive implications—i.e., they sharpen incentives through diluting the agent's value of the information that he may obtain by shirking on some of the tasks. We will revisit this issue later in section 5.

3. Optimal contract when critical task is publicly known

When the identity of the critical task is publicly known, there is no room for private learning by the agent. Hence, the game boils down to the canonical form of relational contract and without loss of generality, we can restrict attention to stationary contracts (Levin, 2003). That is, we assume that the principal offers the same contract and the agent chooses the same effort level every period. If the principal fails to pay the promised bonus or if the agent is caught shirking, the players terminate the relationship and take their outside options in all future periods. As termination is the worst equilibrium payoff and deviations are never observed on equilibrium, this assumption is without loss of generality (Abreu, 1988).

There are three possible actions profiles, which we denote as a, that could be supported in an optimal stationary contract: (i) the agent exerts effort on both tasks; (ii) the agent exerts effort on the critical task only; and (iii) both players exit the relationship and take their outside option in each period. For expositional clarity, we denote these three profiles as $a=E$ ("effort"), C ("critical task only"), and O ("outside option"), respectively. Recall that by Assumption 2, it is never optimal for the relationship to have the agent exert effort only on the non-critical task.

Note that for any value of the discount factor δ, $a=O$ can always be sustained as the equilibrium action profile as both players taking their outside options in each period is trivially a PBE. Next, consider the optimal contract that sustains $a=E$ on the equilibrium path (if feasible). Let the stationary contract be (w,b) where the agent exerts effort in both tasks (i.e., $e_t=2$). As transfers between players are frictionless, without loss of generality,

[5]Assumption 2 states that it is better to dissolve the relationship than to ask the agent to perform a randomly selected task for one period only (whereupon he exerts effort on both tasks in all future periods)

[6]The benchmark case also informs our general analysis. Once the task information is made public by the principal, the features of the optimal contract in the subsequent periods mirror its benchmark-case counterpart.

we assume that in the optimal contract, the principal extracts all surplus. Thus, the agent's individual rationality constraint binds and it is given as:

$$(IR_E) \qquad w + b - c_2 = 0.$$

To sustain $a = E$, we need to ensure that the agent cannot gain by either working on the critical task only or by not working on any tasks at all. Thus, the agent's incentive compatibility constraint is:

$$(1-\delta)(-c_2 + b) \geq \max\{(1-\delta)(-c_1 + pb),\ 0\},$$

or,

$$(IC_E) \qquad b \geq \max\left\{\frac{c_2 - c_1}{1-p}, c_2\right\} = \frac{c_2 - c_1}{1-p},$$

as $(c_2 - c_1)/(1-p) > c_2$ by Assumption 1 (ii). Now, given (IR_E), on the equilibrium path, the principal earns the entire surplus. So, for the principal to not renege on the bonus, we must have the following dynamic enforceability constraint:

$$(DE_E) \qquad \delta(y - c_2) \geq (1-\delta)\, b.$$

Hence, the optimal contract sustaining $a = E$ must be a solution to the following program:

$$\mathcal{P}_E:\ \max_{w,b}\ \hat{\pi}_t = y - c_2 \quad s.t.\ (IC_E),\ (DE_E)\ \text{and}\ (IR_E)$$

Note that by combining (IC_E) and (DE_E), we get that the necessary and sufficient condition to sustain $a = E$ is:

$$(NSC_E) \qquad \frac{\delta}{1-\delta}(y - c_2) \geq \frac{c_2 - c_1}{1-p}.$$

This condition is sufficient because it allows the implementation of $a = E$ through the following feasible contract:

$$b^E = \frac{c_2 - c_1}{1-p}, \text{ and } w^E = c_2 - b^E.$$

Finally, consider sustaining $a = C$—effort only on the critical task—in a PBE. The analysis is identical to the analysis of the case of $a = E$, but with two exceptions: First, when $a = C$, the output could be either y or 0, and the only relevant deviation for the agent is to not work at all where he produces $Y_t = -z$ for sure. Hence, the optimal contract must offer a bonus whenever $Y_t = y$ or 0 and the agent's incentive compatibility constraint boils down to $b \geq c_1$. Second, the per-period surplus is now $py - c_1$ and hence, the principal's dynamic enforceability constraint becomes $\delta(py - c_1) \geq (1-\delta)\, b$. Combining the two, we can derive the necessary and sufficient condition for sustaining $a = C$:

$$(NSC_C) \qquad \frac{\delta}{1-\delta}(py - c_1) \geq c_1.$$

Notice that this condition is sufficient as it allows for the following feasible contract that implements $a = C$ on the equilibrium path: $b^C = c_1$ and $w^C = 0$.

Both (NSC_E) and (NSC_C) conditions are reminiscent of the effort implementation conditions in Levin (2003): They require that the surplus in the relationship for a given effort level needs to be at least as large as the variation in the incentive pay needed to induce that effort level at the first place. For a given δ, the optimal relational contract implements the action profile $a \in \{E, C, O\}$ that maximizes the principal's surplus subject to feasibility. The following proposition characterizes the optimal contract. (We omit the proof as it directly follows from the discussion above and the fact that the left-hand sides of both (NSC)s are increasing in δ.)

Proposition 1. *When the critical task is publicly known, the optimal relational contract is characterized as follows: There exist two thresholds $\overline{\delta}$ (δ at which (NSC_E) is binding) and $\underline{\delta}$ (δ at which (NSC_C) is binding) such that in the optimal relational contract (i) the agent exerts effort in both tasks (i.e., first-best is attained) if $\delta \geq \overline{\delta}$, (ii) the agent exerts effort only on the critical task if $\underline{\delta} \leq \delta \leq \overline{\delta}$, and (iii) no effort could be induced (and the players take their outside options) if $\delta \leq \underline{\delta}$.*

When the critical task is publicly known, the moral hazard problem is acute as the agent has a strong incentive to shirk by exerting effort only on that task and neglecting the other one. To mitigate such moral hazard problem, the principal must promise a large enough bonus payment, which is only credible when the principal has a sufficiently strong reputation concern; i.e., when δ is large enough. For smaller δ, at best, the principal can provide incentive for effort in the critical task only and for δ sufficiently small, the principal may lack the reputational capital necessary to elicit any effort at all.

We conclude this section by introducing the following notation as it will be useful in our subsequent analysis: Let $\mathcal{E}_K$ be the set of equilibrium payoffs (for a given δ) when the critical task is publicly known (hence the subscript K). From Proposition 1, it follows that:

$$\mathcal{E}_K = \left\{(u, \pi) \in \mathbb{R}^2_+ \mid u + \pi \leq s\right\}, \text{ where } s = \begin{cases} y - c_2 & \text{if } \delta > \overline{\delta} \\ py - c_1 & \text{if } \underline{\delta} \leq \delta \leq \overline{\delta} \\ 0 & \text{if } \delta \leq \underline{\delta} \end{cases}.$$

4. Optimal revelation of information on critical task

In this section, we analyze the optimal contract where the identity of the critical task is unknown to both parties but the principal can publicly avail and release this information to bolster incentives. We know from Proposition 1 that if $\delta \geq \overline{\delta}$, the first-best outcome could be attained by making the task information public at the beginning of the game. Hence, in what follows, we limit attention to the case where $\delta < \overline{\delta}$.

An important issue in such a setting is that the optimal contract need not be stationary and the game does not have a tractable recursive structure. To see this, notice that when the agent is asked to put in effort on both tasks and he shirks by working on one task only,

there is some chance that this deviation will not be detected. The agent may happen to choose the critical task and also produce a high output. But then the agent will privately learn which one is the critical task, but the principal would not know that the agent has learnt this information. Such a lack of common knowledge on the agent's belief over the task identities (in the continuation game following a deviation) makes the game lacking on any tractable recursive structure. Consequently, the standard method à la Abreu, Pearce and Stacchetti (1990) to characterize the equilibrium payoff set no longer applies. And we can no longer invoke Levin (2003) to limit attention to the class of stationary contracts without any loss of generality.

Confronting this technical challenge, we proceed along the following steps. First, we discuss the constraints that must hold in order to support the action profile $a = E$ (i.e., effort on both tasks) in a given period when the critical task is not known to the agent. Second, we derive a necessary and sufficient condition for first-best to be attainable. Finally, we characterize the optimal contract when this condition is violated.

4.1. **Constraints for sustaining effort in both tasks.** Consider a period t such that the principal has not revealed the critical task yet, and the agent has not shirked in the past. Let the set of all PBE payoffs of the repeated game starting from the period t be $\mathcal{E}$. Note that while $\mathcal{E}$ depends on δ, it is independent of t—as no information on the critical task is available to the agent, his belief over the tasks is the same as his prior.

Consider a payoff $(u, \pi) \in \mathcal{E}$ that is sustained by playing $a = E$ in the current period (i.e., period t).[7] Now, in equilibrium, at the end of any period, the strategies may call for one of the following three action profiles for the next period: the principal does not reveal the identity of the critical task and elicits effort in both tasks ($a = E$), the principal reveals which task is critical and the agent exerts effort in that task only ($a = C$), and both players take their outside options ($a = O$).[8] Also, using the public randomization device, the players could randomize over these three action profiles in the next period. Suppose that under the equilibrium strategy profile (that supports (u, π)) the action $a \in \{E, C, O\}$ is taken (in the following period) with probability α^a and the corresponding continuation payoffs for the players are given as (u^a, π^a). If any player is caught deviating without loss of generality we may assume that the players take their outside options forever.[9]

If a PBE induces the action profile $a = E$ in the current period, the associated equilibrium payoffs, contracts, and the continuation payoffs must satisfy the following constraints.

As in the case of publicly known critical task, the contract must satisfy the "incentive compatibility" constraints as the agent should not gain by deviating and shirking altogether or by performing exactly one of the two tasks. As before, if the agent shirks on both tasks, we must have $Y_t = -z$ and the relationship necessarily terminates. Hence, we require that:

[7]Note that for any δ, $\mathcal{E}_K \subseteq \mathcal{E}$ (recall that $\mathcal{E}_K$ is the set of PBE payoffs when the critical task is publicly known as defined above in Section 3) as the principal can always reveal the critical task at the beginning of the game but there may be some payoffs that can be attained only if the principal withhold this information for at least one period.

[8]As we are considering the case where $\delta < \bar{\delta}$, from Proposition 1 we know that following the revelation of the critical task, effort in both tasks ($a = E$) cannot be induced.

[9]The agent has a detectable deviation if the output does not conform to the his equilibrium effort level. Similarly, the principal's detectable deviation consists of reneging on bonus promise or failing to conform the equilibrium information revelation policy, or both.

$$(IC_0) \qquad u \geq (1-\delta)\, w.$$

But if the agent shirks by exerting effort in one of the two tasks only, the derivation of the incentive compatibility constraint is somewhat more involved. It must allow for the fact that up on deviating the agent may privately learn the identity of the critical task and he may use this information to deviate again in the future.

To address this issue, we proceed as follows. For any equilibrium payoff pair (u', π'), let $U(u', \pi')$ be the maximal continuation payoff of the agent where he privately knows which one is the critical task. That is, suppose (σ'_P, σ'_A) is the strategy profile of the players that give rise to the payoff (u', π'). Now, $U(u', \pi')$ is the agent's payoff when he deviates from σ'_A and plays his best-response to σ'_P using his knowledge on the identity of the critical task. If the payoff pair (u', π') could be implemented with different strategies leading to different maximal deviation payoff for the agent, without loss of generality, we assume that the equilibrium dictates that the players choose the strategy profile that minimizes the agent's maximal deviation payoff (after knowing the critical task). This allows (u', π') to be a sufficient statistic for the agent's maximal deviation payoff. We can now derive the agent's incentive compatibility constraints using the deviation payoff U:

$$(IC_1) \qquad u \geq (1-\delta)\left(w - c_1 + \frac{1}{2}pb\right) + \frac{1}{2}p\delta\left(\alpha^E U\left(u^E, \pi^E\right) + \alpha^C u^C + \alpha^O U\left(u^O, \pi^O\right)\right).$$

Note that if the agent shirks by working on only one of the two tasks, with probability $\frac{1}{2}p$ he would pick the critical task and produce the on-equilibrium path output y; and hence, the principal would fail to detect such a deviation.

Also as in our previous analysis, the contract must satisfy the "dynamic enforceability" constraint to ensure that neither the principal nor the agent has incentives to renege on the bonus:

$$(DE_P) \qquad -(1-\delta)\, b + \delta\left(\alpha^E \pi^E + \alpha^C \pi^C + \alpha^O \pi^O\right) \geq 0,$$

$$(DE_A) \qquad (1-\delta)\, b + \delta\left(\alpha^E u^E + \alpha^C u^C + \alpha^O u^O\right) \geq 0.$$

But in contrast to our previous analysis, there are two additional set of constraints involving the continuation payoffs. First, the consistency requirement of payoff decomposition implies that a player's payoff must be equal to the weighted sum of his current and continuation payoff. Hence, we must have the following "promise-keeping" constraints:

$$(PK_A) \qquad u = (1-\delta)(w - c_2 + b) + \delta\left(\alpha^E u^E + \alpha^C u^C + \alpha^O u^O\right),$$

$$(PK_P) \qquad \pi = (1-\delta)(y - w - b) + \delta\left(\alpha^E \pi^E + \alpha^C \pi^C + \alpha^O \pi^O\right).$$

Second, we also have the "self-enforcing contract" constraints requiring the continuation payoffs in each of the three cases (i.e., under strategies that specify $a = E,\ C$, or O to be played in the next period) to be equilibrium payoffs themselves. Notice that if the critical

task is revealed, then the analysis is identical to the case in which the critical task is publicly known and the associated payoff set is $\mathcal{E}_K$. Hence, the (SE) constraints are:

$$(SE_E) \qquad \left(u^E, \pi^E\right) \in \mathcal{E},$$

$$(SE_C) \qquad \left(u^C, \pi^C\right) \in \mathcal{E}_K,$$

and

$$(SE_O) \qquad \left(u^O, \pi^O\right) \in \mathcal{E}.$$

4.2. **A necessary and sufficient condition for the first-best.** In light of the constraints above, we can now investigate the feasibility of an efficient contract where $\alpha^E = 1$ (and hence, $a = E$) in all periods. The following lemma presents a necessary and sufficient condition for such a contract to exist.

Lemma 1. *An efficient relational contract (where $\alpha^E = 1$ in all periods) can be sustained if and only if:*

$$(NSC^*) \qquad \frac{\delta}{1-\delta}(y - c_2) \geq \frac{c_2 - c_1}{1 - p/(2 - p\delta)}.$$

That is, there exists a δ^* (δ at which (NSC^*) is binding) such that the optimal relational contract induces the agent to exert effort in both tasks in all periods if and only if $\delta \geq \delta^*$ and the principal never reveals the task identity to the agent.

Notice that the condition (NSC^*) is similar in spirit to (NSC_E) (i.e., the necessary and sufficient condition when the critical task is *publicly* known) and requires the surplus in the relationship to be larger than a threshold for the first-best to be implementable. However, the threshold is strictly smaller than the one specified in (NSC_E) and hence, we have $\delta^* < \overline{\delta}$. In other words, Proposition 1 and Lemma 1 taken together imply that efficiency is easier to attain when the critical task is not known to the agent.

The intuition is as follows. When the identity of the critical task is unknown, it lessens the moral hazard problem of the agent as he does not know which task to choose should he deviate and decide to exert effort on only one of the two tasks. In particular, the agent's gains from deviation are muted by the fact that by randomly choosing one of the two tasks, the agent would pick the critical task only half of the time. Consequently, the bonus needed to provide incentive for effort in both tasks is smaller compared to its counterpart in the public information case, and hence, the principal can induce first-best effort even when it has less reputational capital.

Lemma 1 suggests that the first-best outcome is not feasible if the (NSC^*) condition is violated (i.e., if $\delta < \delta^*$). But when $\delta < \delta^*$, what is the optimal way to reveal the task information when incentives are offered through relational contracts? We have noted earlier that if the principal reveals the critical task at the beginning of the game, then only the effort in the critical task may be sustained in the future (given that $\delta \in [\underline{\delta}, \delta^*)$), leading to a loss of surplus relative to the first-best. Thus, the analysis of the optimal contract explores the optimal disclosure of the task information to minimize such loss. In other words, if the

(NSC^*) condition is violated, can the principal improve his payoff by delaying the revelation of the critical task rather than disclosing it at the beginning of the game? We address this question next.

4.3. **Preliminary analysis.** In order to analyze the optimal contract when the first-best outcome is not feasible, we begin by presenting a set of lemmas that simplify our subsequent analysis. These lemmas state several observations about any PBE payoffs that are sustained by $a = E$ in a *given* period when the information on the critical task is unknown to the agent and they allow us to restrict attention to a specific class of contracts without any loss of generality. (The discussion below focuses on technical details; readers primarily interested in our key findings and economic intuition can omit this section and directly go to section 4.4.)

First, we show that the optimal contract need not use any bonus.

Lemma 2. *Consider a relational contract such that the identity of the critical task is not known to the agent till period t and in the game starting from period t the payoff profile $(u, \pi) \in \mathcal{E}$ is sustained by $a = E$ and $b \neq 0$ in period t. Then there exists another contract where the payoff (u, π) can be sustained by $a = E$ and $b = 0$ in period t.*

The intuition behind this observation is as follows: first, suppose that (u, π) is supported by a contract that specifies $a = E$ and a negative bonus b in a given period t. Such a contract is payoff equivalent to one where $a = E$, but the wage (w_t) in that period is reduced by b and the bonus is set to 0. (It is routine to check that this new contract is also feasible). Next, suppose that the contract specifies $a = E$ and a positive bonus b. Now, one may set $b = 0$ in period t and distribute the bonus amount among the continuation payoffs u^as (by raising the wages in period $t + 1$ that support each of the u^a payoffs) such that, in expectation, the agent continues to earn b.

Note that under Lemma 2, (DE_A) and (DE_P) are automatically satisfied, and hence, could be dropped from the set of constraints that the optimal contract must satisfy.

Next, we present three lemmas that characterize the continuation payoffs in an optimal relational contract. Lemma 3 given below claims that without loss of generality, we can restrict attention to contracts that give zero continuation value to the principal (in each period, except at the beginning of the game).

Lemma 3. *Consider a relational contract such that the identity of the critical task is not known to the agent till period t and in the game starting from period t the payoff profile $(u, \pi) \in \mathcal{E}$ is sustained by $a = E$ in period t and $\pi^a > 0$ for some $a \in \{E, C, O\}$. Then there exists another contract where the continuation payoff (u, π) can be sustained by $a = E$ in period t and $\pi^a = 0$ for all $a \in \{E, C, O\}$.*

The intuition behind this observation is similar to that of Lemma 2 discussed earlier: any contract supporting (u, π) with $\pi^a > 0$ in some period t can be replaced by one that (i) sets $\pi^a = 0$ for all $a \in \{E, C, O\}$, (ii) increases the agent's continuation payoff u^a by raising the wage w^a that supports u^a (for all $a \in \{E, C, O\}$), and (iii) reduces w by the (discounted) expected continuation payoff of the principal (π^a). It can be shown that for appropriate choices of w^as, such a contract is feasible and is payoff equivalent to the initial one.

The next lemma suggests that there is no loss in considering only those contracts that do not specify $a = O$ on the equilibrium path.

Lemma 4. *If an optimal relational contract exists where the joint surplus is strictly positive, then there exists an optimal relational contract in which $\alpha^O = 0$ in all periods.*

To see the reasoning note that a strategy profile that calls for taking the outside option in period t is payoff equivalent to an alternative strategy given as follows: in period t the strategy does not require the players to take their outside options but calls for termination of the relationship with probability $1 - \delta$. All other aspects of the new strategy are identical to the former one. Note that the above lemma implies that for our analysis, we may assume that until the information on the critical task is revealed, the equilibrium strategies always call for $a = E$.

Finally, lemma 5 below suggests that when the critical task is revealed, we may set $u^C = py - c_1$, i.e., the expected surplus when the agent exerts effort only on the critical task.

Lemma 5. *In an optimal relational contract, in any period, if $\alpha^C > 0$ then $u^C = py - c_1$.*

Lemmas 2–5 have the following important implication. In order to characterize the optimal contract, without loss of generality we may restrict attention to contracts where, in any period, $b = 0$ and

$$w = \begin{cases} y & \text{if } a = E \text{ is played} \\ py & \text{if } a = C \text{ is played} \end{cases}.$$

That is, in this class of contracts, in the continuation game following every history on the equilibrium path the agent receives all of the surplus and $\pi = 0$. Note that under such a contract, (PK_P) is trivially satisfied. Also note that for our analysis below, we can write $U(u)$ instead of $U(u, \pi)$ since the principal's continuation payoff remains 0.

Using the above observations, we are now ready to characterize the optimal relation contracts.

4.4. **Characterization of the optimal relational contract.** Using the above lemmas we can now simplify the optimal contracting problem as follows. Notice that Lemma 4 implies that when deriving the optimal contract, we can restrict attention to contracts where in each period t the principal either demands $a = E$ from the agent or with some probability, reveals the critical task and asks the agent to work on that task only. Let α_t be the probability that the principal does *not* reveal the critical task at the end of period t given that it has not been revealed yet. Note that the optimal relational contract is completely determined by the sequence $\{\alpha_t\}_{t=1}^{\infty}$. Recall that we denote by u^t the agent's (average) payoff at the beginning of period t when the critical task remains undisclosed. Using Lemma 5, we then have the following recursive relationship for the agent's payoff:

$$u^t = (1 - \delta) s_2 + \delta \left(\alpha_t u^{t+1} + (1 - \alpha_t) s_1 \right), \tag{1}$$

where $s_2 := y - c_2$ (surplus when the agent exerts effort in both tasks) and $s_1 := py - c_1$ (surplus when the agent exerts effort in the critical task only). Therefore, if there exists a contract that implements $a = E$ at least in the first period, solving for optimal contract (in the class of such contracts) tantamount to finding the optimal sequence $\{\alpha_t\}_{t=1}^{\infty}$ to maximize u^1. In other words, the optimal contract must solve the following program (denote $c := c_2 - c_1$):

$$\mathcal{P}: \begin{cases} \max_{\alpha_t \in [0,1]} \quad u^1 \quad s.t.\ \forall t, \\ u^t = (1-\delta)\, s_2 + \delta\left(\alpha_t u^{t+1} + (1-\alpha_t)\, s_1\right) & (PK_A) \\ u^t \geq (1-\delta)\, y & (IC_0) \\ u^t \geq (1-\delta)\,(s_2 + c) + \frac{1}{2} p\delta\left(\alpha_t U\left(u^{t+1}\right) + (1-\alpha_t)\, s_1\right) & (IC_1) \\ (u^t, 0) \in \mathcal{E} \quad (SE_E)\,;\ (s_1, 0) \in \mathcal{E}_K \quad (SE_C)\,;\ \textit{and}\ \left(u^O, 0\right) \in \mathcal{E} \quad (SE_O) \end{cases}$$

Note that if $\alpha_t = 1$ for all t is feasible in $\mathcal{P}$, then the optimal relational contract is efficient. As we have argued in Lemma 1, this is the case if and only if (NSC^*) is satisfied. As we are interested in the case where (NSC^*) is violated, we consider below the case where $\alpha_t = 1$ for all t is not feasible.

The optimal contracting problem $\mathcal{P}$ presents an interesting technical challenge as there is no standard method that could be used to directly compute $U(u)$, the agent's maximal deviation payoff in the continuation game (after he privately learns the task identities). The complexity stems from the fact that once the agent learns the task identities (through shirking), the profitability of his future deviations would depend on the associated disclosure policy.

For example, suppose that the disclosure policy calls for revelation of the critical task after a certain number of periods, T (say), but the agent has shirked and learned the critical task in an earlier period $t < T$. In such a scenario it may not be optimal for the agent to continue to shirk in all subsequent periods until the task information is made public. Instead, he may continue to exert effort on both tasks till some period $\bar{t}$, where $t < \bar{t} < T$, and then start to shirk again by working on the critical task only. Notice that under the above disclosure policy, the agent's continuation payoff decreases over time as we move closer to the revelation date T. Hence, it may not be worthwhile for the agent to continue to shirk immediately after his first successful deviation (as he stands to lose a larger continuation payoff) but he may be tempted to shirk again at a date closer to period T (as he risks losing a smaller continuation payoff).

We address this problem by considering a relaxed program that only allows for a specific form of deviation: if the agent deviates and (privately) learns the critical task, in all subsequent periods he always deviates by choosing the critical task only (until he is detected). Notice that for a contract to be a part of an equilibrium, it must be robust to all forms of deviation including the one specified above. Hence, the aforementioned deviation can be used to compute a lower bound on the agent's maximal deviation payoff $U(u)$ and we can characterize the optimal revelation policy that deters this type of deviation. We then show that this relaxed problem admits a stationary solution where at the end of each period, the

principal reveals the critical task with a fixed probability. Consequently, this policy is robust to all deviations and, hence, a solution to the original problem $\mathcal{P}$. Lemma 6 reports this finding.

Lemma 6. *If there exists a solution to the optimal contracting problem $\mathcal{P}$ then there also exists a stationary solution to $\mathcal{P}$ where for all t, $\alpha_t = \alpha^*$ (which may vary with δ). That is, at the end of each period the firm reveals the critical task with a constant probability α^*.*

It is instructive to elaborate on the intuition behind the above lemma. Notice that even though the relaxed problem limits attention to a specific form of deviation (the agent continues to deviate and works only on the critical task once he deviates for the first time), the exact time of deviation is still a choice variable for the agent. Hence, we have infinitely many incentive constraints: for every period t, we must have a constraint ensuring that no profitable deviation exists in that period.

It turns out that if the disclosure policy were to deter deviation in period 1 only, it takes a form that features "early revelation": there is some T such that the critical task is revealed with positive probability for $t < T$, and it is never revealed afterwards. Note that a key feature of such policy is that the agent's continuation payoff is increasing over time and it helps relaxing the agent's incentive constraint in period 1. To see why, note that compared to the agent who always works on both tasks, the agent who only works on the critical task is effectively less patient—the former discounts the future at a rate δ but the effective discount rate of the latter is $p\delta$ as the relationship is likely to terminate for him with probability $1-p$. As an increasing continuation payoff is more attractive to the more patient agent, such a revelation policy dissades the agent from shirking in the earlier periods.

However, such an early revelation policy is necessarily time-inconsistent. While the agent is deterred from deviating in period 1, he may want to deviate in the later periods when the continuation payoffs are larger and hence, the gains from shirking are larger as well. (As the principal is less likely to reveal the task in the later periods, the agent earns a larger information rent if he shirks and learns the critical task.) In other words, for every period t, the optimal policy would like to induce an increasing sequence in continuation payoff through increasing the current period's revelation probability and decrease future revelation probability. But as this needs to be done for every period, so resulting optimal policy becomes stationary and features a constant revelation probability in each period. Finally, by virtue of stationarity, this policy is necessarily robust to all possible deviations of the agent.

We are now ready to present a complete characterization of the optimal contract:

Proposition 2. *The optimal contract is characterized as follows. There exist four cutoffs $\underline{\delta} < \tilde{\delta} \leq \delta^* < \overline{\delta}$ ($\underline{\delta}$, $\overline{\delta}$ as defined in Lemma 1 and δ^* as defined in Lemma 1) such that the following holds:*

(i) For all $\delta \geq \overline{\delta}$, the optimal contract can attain the first-best outcome irrespective of the principal's decision on whether to reveal the critical task.

(ii) For all $\delta \in [\delta^*, \overline{\delta})$, the optimal contract attains the first-best outcome where the principal never reveals the identity of the critical task.

(iii) For all $\delta \in [\tilde{\delta}, \delta^*)$, the first-best cannot be attained. In the optimal contract, the principal reveals the critical task at the end of each period with a constant probability α^* (which may vary with δ). The agent works on both tasks until the critical task is revealed and works only on the critical task afterwards. Moreover, $\tilde{\delta} < \delta^*$ if and only if:

$$s_1 - \frac{2-p-p\delta^*}{2-p-p\delta^*+\frac{1}{2}p\delta^*}s_2 > 0 \text{ and } s_1 - \frac{2-p-p\delta^*}{2-p\delta^*}s_2 > \frac{1-\delta^*}{\delta^*}c_1.$$

(iv) For all $\delta \in [\underline{\delta}, \tilde{\delta})$, the first-best cannot be attained. In the optimal contract, the principal reveals the critical task at the beginning of the game and the agent works only on the critical task.

(v) Finally, for all $\delta < \underline{\delta}$, no effort could be induced and both parties take their outside options.

Proposition 2 offers a sharp characterization of the optimal contract that is arguably both realistic and intuitive. Recall that the information on the critical task makes it easier for the agent to shirk as he knows which task to perform if he decides to do only one of the two. As the agent's payoff from shirking is larger (compared to the case where the agent does not know the critical task as may choose the critical task only half of the time) his incentive compatibility constraint is harder to satisfy and the principal must provide stronger incentive to dissuade him from shirking.

As discussed earlier, for δ sufficiently large $\left(\delta \geq \overline{\delta}\right)$ the principal can credibly promise a sufficiently large bonus payment so that the agent would not shirk even when he knows the identity of the critical task. Thus, whether the principal reveals the task or not is irrelevant for the performance of the optimal contract. On the other hand, if δ is sufficiently small, i.e., if $\delta \leq \underline{\delta}$, the principal's reputational capital is too small to sustain any effort.

The more interesting case is the intermediate one where $\delta \in \left(\underline{\delta}, \overline{\delta}\right)$ as the optimal contract calls for active management of information to ensure stronger incentives. In particular, if δ is relatively large, i.e., if $\delta \in [\delta^*, \overline{\delta})$, the first-best could be attained only if the principal resorts to opacity and does not reveal the information on the critical task. That is, the maximum bonus that the principal can credibly promise is large enough to dissuade the agent from shirking when he does not know the critical task, but the bonus is too small to elicit effort on both tasks when the critical task is known to the agent. In contrast, for δ relatively small (i.e., δ close to $\underline{\delta}$) the optimal policy calls for a complete transparency on task identities. The maximum bonus that the principal can credibly promise could induce effort in at most one of the two tasks and naturally, the optimal policy is to let the agent know which task is critical.

But for a relatively large δ—i.e., if δ is not too far below δ^*—the principal may do better by not revealing the task information at the beginning of the game. A larger surplus could be attained by adopting a stochastic revelation policy where at the end of each period, the principal may reveal the critical task with a fixed probability. As the agent knows that the critical task is likely to become public information in the near future, it dilutes the value of private information that he hopes to enjoy by shirking and learning the critical task privately. Such a contract elicits effort in both tasks until the tasks are revealed, and hence, is more efficient than the one that reveals the task at the beginning of the game.

Such a stochastic revelation policy is indeed optimal for some δ if and only if the two conditions given in part (iii) are satisfied. Note that both of these conditions require the surplus with effort on critical task only (s_1) to be not too small compared to the first-best surplus (s_2). To see the intuition, observe that the revelation of the critical task has two effects. On the one hand, the benefit of revelation is that it reduces the value to the agent of experimenting and learning the identity of the critical task. On the other hand, the cost of revelation is that the total surplus in the relationship is reduced—once the critical task is revealed, the agent would perform that task only. The larger is s_1, the smaller is the cost of revelation whereas the benefit of experimentation is primarily linked to the surplus under first-best effort, since following a successful experimentation, the agent per period payoff is equal to the first-best surplus ($s_2 = y - c_2$) plus the cost of effort saved ($c_2 - c_1$). As a result, the larger is s_1 the more likely it is that a partial revelation emerges as the optimal relational contract.

5. Job rotation as an incentive device

As mentioned earlier, any organizational policy of the principal that could potentially erode the agent's information value of shirking may mitigate the moral hazard problem that we highlight here. In our main analysis, we focused on information revelation as one such policy. Another policy that one may consider is job rotation. At the end of each period, the principal may consider moving the agent to a similar yet a new job where any information that the agent may have about the tasks in his previous job is irrelevant.

For example, in a service sector firm, the agent may be asked to work on similar projects every period but for different clients. Which aspect of the project is more critical depends on the client the agent is working for. The identity of the critical task is assumed to be statistically independent across clients, i.e., the identity of the critical task in the current client's project gives no information about the critical tasks in the projects with the future clients.

The model on information revelation could be readily adapted to analyze such a job rotation policy. We keep all aspects of our initial model unchanged except for the following: assume that neither the principal nor the agent knows the identity of the critical task and, moreover, the principal cannot obtain or disclose this information either. However, at the end of each period, the principal may reassign the agent to a new project by incurring a cost ψ. Such a cost could be interpreted as an administrative cost that the principal incurs every time he assigns the agent to a new client. The projects are identical up to the identity of the critical task, which is assumed to be statistically independent across projects.

5.1. **Optimal job rotation policy.**

The qualitative features of the optimal job rotation policy bears strong resemblance with the optimal information revelation analyzed earlier in section 4. Hence, for the sake of brevity, we elaborate below only on those aspects of our analysis that differ from its information revelation counterpart.

Notice that on the equilibrium path, in any period there are three possible action profiles: (i) agent stays on the same job and exerts effort in both tasks ($a = E$), (ii) the agent is assigned to a new job (i.e., there is "job rotation") and exerts effort in both tasks ($a = R$), and finally, (iii) both parties take their outside options ($a = O$). Consider the constraints that a contract must satisfy in order to sustain $a = E$ in a given period. As before, let α^a be the probability of choosing action profile a in the subsequent period and let (u^a, π^a) be the continuation payoffs where $a \in \{E, R, O\}$. These constraints are identical to their counterpart

in section 4.1 except for the following two differences: (i) In the principal's promise-keeping constraint (PK_P) and dynamic enforceability constraint (DE_P), π^C is replaced by $\pi^R - \frac{1-\delta}{\delta}\psi$ to account for the cost of job rotation.[10] And (ii) the sequential enforceability constraint following the revelation of critical task (SE_C) is replaced by

$$(SE_R) \qquad \left(u^R, \pi^R\right) \in \mathcal{E}.$$

Observe that if the agent is assigned to a new job next period, the continuation game is identical to the one in the current period as any information on the tasks that the agent may have gained by shirking becomes irrelevant in the new job.

Now, as discussed in Lemmas 2–5 in the context of information revelation, it is routine to check that the following conditions continue to hold even in the current setting: without loss of generality, we can restrict attention to a class of contracts where (i) no bonus is used (i.e., $b = 0$), (ii) the principal's continuation payoff (net of job rotation cost) is always 0 (i.e., $\pi^E = \pi^O = \pi^R - \frac{1-\delta}{\delta}\psi = 0$), (iii) termination is never used (i.e., $\alpha^O = 0$), and finally (iv) in the continuation game following a job rotation, the principal and the agent choose the equilibrium with the highest possible joint payoff, i.e., they choose the equilibrium where $u^R + \pi^R$ is maximal amongst those in $\mathcal{E}$; denote such joint payoff as v.

Hence, as in the case of information revelation, solving for the optimal contract is equivalent to maximizing the agent's payoff and that is given by the following recursive equation:

$$(PK_A\text{-}R) \qquad u^t = (1-\delta)\, s_2 + \delta\left(\alpha_t u^{t+1} + (1-\alpha_t)\left(v - \frac{1-\delta}{\delta}\psi\right)\right),$$

where α_t is the probability that the agent will continue on the same job (recall that v is the agent's continuation payoff if a job rotation occurs at the end of period t). Similarly, (IC_1) constraint can be written as:

$$(IC_1\text{-}R) \qquad u^t \geq (1-\delta)\,(s_2 + c) + \frac{1}{2}p\delta\left(\alpha_t U\left(u^{t+1}\right) + (1-\alpha_t)\left(v - \frac{1-\delta}{\delta}\psi\right)\right).$$

We are now ready to analyze the optimal contract in this case. As v is fixed, with an abuse of notation, we may define

$$s_1 = v - \frac{1-\delta}{\delta}\psi.$$

Notice that the optimal contracting problem is now essentially the same as the principal's program $\mathcal{P}$ studied in section 4.4 in the context of information revelation. As one would expect, the optimal contracts in these two settings also share similar characteristics.

[10] Note that the cost ψ is incurred in the current periods and discounted by $1-\delta$ whereas the continuation payoff π^R is discounted by δ.

Proposition 3. *The optimal contract with job rotation is characterized as follows. There exist two cutoffs δ_R and δ^*, $\delta_R \leq \delta^*$ (δ^* as defined in Lemma 1) such that the following holds:*

(i) For all $\delta \geq \delta^*$, the optimal contract attains the first-best outcome: the agent continues to work on the same job in every period and exerts effort in both tasks.

(ii) For all $\delta \in [\delta_R, \delta^*)$, the optimal contract fails to attain first-best and is given as follows: At the end of the period, the principal assigns the agent to a new job with a constant probability α_R^* (which may vary with δ). The agent exerts effort in both tasks in the job he has been assigned to. Moreover, $\delta_R < \delta^*$ if and only if the cost of job rotation ψ is below a threshold.

(iii) Finally, for all $\delta < \delta_R$, no effort could be induced and both parties take their outside options.

Notice that unlike in the information revelation case, the optimal contract in the current setting either induces first-best effort (i.e., effort in both tasks) in all periods or no effort at all. As all jobs are assumed to be identical except for the identity of the critical task, if the threat of job rotation provides strong enough incentive for the first-best effort ($a = E$) in the current job, it provides similar incentives in the agent's next job as well (should he get re-assigned to a new job in the future). But if job rotation fails to provide strong enough incentive for effort in both tasks, it is optimal to dissolve the relationship as performing only one task chosen randomly out of the two yields an even smaller surplus (by Assumption 2).

6. Discussion and Conclusion

This article explores the optimal provision of relational incentives when the worker may shirk to learn. Workers often hold jobs that involve multiple aspects (or a set of tasks) where some aspects may be more crucial than others. An interesting moral hazard problem emerges when the worker lacks information about the relative importance of the various job aspects: he may shirk on some aspects of the job not only to save on the costly effort but also to learn more about their importance in the job and use this information in the future to shirk more effectively. We argue that the firm could strengthen incentives by strategically disclosing information about the job aspects and the optimal disclosure policy is closely tied to the amount of surplus generated by the employment relationship. A moderately large surplus calls for opacity (firm does not disclosure any information), a moderately small surplus calls for full transparency (firm disclose all information), and if the available surplus is in an intermediate range, active filtering of information through a stochastic disclosure policy is optimal.

We conclude this article with a brief discussion on the range of applicability of our analysis. Even though we consider a multitasking environment in our model, the qualitative nature of our findings could be also captured in other environments where the worker may shirk to learn and the firm may use alternative mechanisms that can discipline the worker by diminishing his gains from learning-by-shirking.

Indeed, our model closely corresponds to a canonical moral hazard problem where the worker may be unsure of the consequences of shirking. For example, we can reinterpret the effort levels $e \in \{0, 1, 2\}$ in our model as three levels of effort in a given task. If $e = 2$, the output is always good (y), if $e = 0$, the output is always bad ($-z$), but the outcome for $e = 1$ depends on the underlying state of the world that is known to the firm only. Suppose that if

the state is good, there is some chance the output would be still good and it is moderate (0) otherwise. But if the state is bad, the output would be bad for sure. Note that this model is quantitatively identical to our multitasking model where worker may shirk and choose $e = 1$ in order to learn the underlying state and the firm decides when and what information about the state to reveal to the worker.

Also, while we only highlight two possible policies aimed at dissuading learning-by-shirking—strategic disclosure of information and job rotation—one may conceive several other mechanisms that improve incentives by reducing the worker's payoff from shirking. For example, the firm can commit to introduce a new production technology after every so many periods. Adoption of new technology could be costly and such a cost plays the same role as that of the loss of surplus (from $y - c_2$ to $py - c_1$) in our model when the firm reveals the critical task. Similar incentive effects may also stem from frequent adoption of new performance measures as the existing measures "run down" and lose their ability to discriminate good from bad performances. As Meyer (2002) notes, perverse learning or learning how to "game the system" as one of the major causes of running down of performance measures. The identity of the critical tasks may depend on what job performance measures are in place and the workers may not have incentives for such perverse learning if they expect the current measures to become irrelevant in the near future.

Appendix

This appendix contains the proofs omitted in the text.

Proof of Lemma 1. The proof is given in three steps. In the first step we derive a lower bound for $U(u,\pi) - u$ in any efficient equilibrium. In the second step we use this lower bound to show that NSC^* is a necessary condition for the existence of an efficient equilibrium. In the third step we show that this condition is also sufficient for the existence of an efficient equilibrium.

Step 1 (*Lower bound for $U(u,\pi) - u$ in an efficient equilibrium*). Suppose that an efficient equilibrium exists (where $\alpha^E = 1$ in all periods). Define

$$D := \min_{(u,\pi)\in\mathcal{E}} U(u,\pi) - u, \ \ s.t.\ \pi + u = y - c_2.$$

That is, D is the minimum of the agent's superior information across all the efficient equilibria. We next derive a lower bound for D. Take an arbitrary $(u,\pi) \in \mathcal{E}$ such that $\alpha^E = 1$ in all periods, i.e., $\pi + u = y - c_2$. Then, there exist w, b, u^E, and π^E such that (PK_A) and (DE_P) are satisfied, i.e.,

$$u = (1-\delta)(w - c_2 + b) + \delta u^E, \tag{2}$$

and

$$(1-\delta)\, b \leq \delta \pi^E = \delta(y - c_2 - u^E), \tag{3}$$

where the last equality follows from the fact that in an efficient equilibrium $u^E + \pi^E = y - c_2$. Moreover, observe that

$$U(u,\pi) \geq (1-\delta)(w + pb - c_1) + p\delta U\left(u^E, \pi^E\right). \tag{4}$$

(An inequality—and not necessarily an equality—holds as the agent may choose to exert effort in both tasks even if he knows the identity of the critical task.) Using (2) and (4), it then follows that

$$\begin{aligned} U(u,\pi) - u \geq & \ (1-\delta)(c_2 - c_1) + p\left((1-\delta)b + \delta U\left(u^E, \pi^E\right)\right) - (1-\delta)b - \delta u^E \\ = & \ (1-\delta)(c_2 - c_1) - (1-p)\left((1-\delta)b + \delta u^E\right) + p\delta\left(U\left(u^E, \pi^E\right) - u^E\right) \\ = & \ (1-\delta)(c_2 - c_1) - (1-p)\delta(y - c_2) + p\delta D, \end{aligned}$$

where the last inequality follows from (3) and the definition of D. As the inequality above holds for all (u, π), we therefore have

$$D \geq (1-\delta)(c_2 - c_1) - (1-p)\delta(y - c_2) + p\delta D,$$

or,

$$D \geq \frac{1}{1-p\delta}\left((1-\delta)(c_2 - c_1) - (1-p)\delta(y - c_2)\right). \tag{5}$$

Step 2 (*Necessity of NSC**). Let $(u, \pi) \in \mathcal{E}$ such that $\alpha^E = 1$ in all periods, i.e., $\pi + u = y - c_2$. Because these payoffs correspond to an efficient equilibrium, the associated continuation payoffs u^E and π^E must also satisfy $u^E + \pi^E = y - c_2$. Now, as $\alpha^E = 1$ (and $\alpha^C = \alpha^O = 0$), combining (IC_1) and (PK_A) one obtains:

$$(1 - \frac{1}{2}p)\left((1-\delta)b + \delta u^E\right) \geq (1-\delta)(c_2 - c_1) + \frac{1}{2}p\delta\left(U\left(u^E, \pi^E\right) - u^E\right).$$

Also, from (DE_P') we find:

$$(1 - \frac{1}{2}p)\left((1-\delta)b + \delta u^E\right) \leq (1 - \frac{1}{2}p)\left(\delta\pi^E + \delta u^E\right) = (1 - \frac{1}{2}p)\delta(y - c_2).$$

Hence, we must have

$$\begin{aligned} (1 - \tfrac{1}{2}p)\delta(y - c_2) & \geq (1-\delta)(c_2 - c_1) + \tfrac{1}{2}p\delta\left(U\left(u^E, \pi^E\right) - u^E\right) \\ & \geq (1-\delta)(c_2 - c_1) + \tfrac{1}{2}p\delta D \\ & \geq (1-\delta)(c_2 - c_1) + \tfrac{p\delta}{2(1-p\delta)}\left((1-\delta)(c_2 - c_1) - (1-p)\delta(y - c_2)\right), \end{aligned}$$

(where the last inequality follows from (5)), or, equivalently,

$$\frac{\delta}{1-\delta}\left(1 - \frac{p}{2-p\delta}\right)(y - c_2) \geq c_2 - c_1. \tag{NSC^*}$$

Step 3. (*Sufficiency of NSC**) Consider the following stationary contract: in each period $w = y$ and $b = 0$, the agent is asked to exert effort in both tasks, and the relationship terminates if the agent is caught shirking. Under this arrangement, the principal's payoff is $\pi = 0$, the agent's payoff is $u = y - c_2$, and the only constraints that need to be checked in order for it to be sustained as an equilibrium are (IC_0) and (IC_1).

To check that (IC_0) is satisfied, note that under the proposed contract, $u^E = y - c_2$ and $b = 0$. Plugging these values in (IC_0), we get

$$-(1-\delta)c_2 + \delta(y - c_2) \geq 0 \Leftrightarrow \frac{\delta}{1-\delta}(y - c_2) \geq c_2,$$

which is satisfied when (NSC^*) is satisfied. To see this, observe that

$$\frac{\delta}{1-\delta}(y-c_2) \geq \frac{c_2-c_1}{1-p/(2-p\delta)} \geq \frac{c_2-c_1}{1-p/2} \geq c_2,$$

where the first inequality corresponds precisely to the (NSC^*), the second follows from the fact that $p\delta \in (0,1)$, and the third from the fact that $c_1 \leq \frac{1}{2}pc_2$ (Assumption 1 (ii)).

To check that (IC_1) is satisfied we need to analyze the agent's value from private information under the arrangement. Suppose the agent privately learns which task is critical. Given that the principal will continue to play according to the contract, the agent's problem is stationary, which implies that either the agent never shirks (by doing the critical task only) or he always shirks. Suppose first that the agent never shirks. Then, $U\left(u^E,\pi^E\right)=u^E$ and, since $u^E=y-c_2$ and $b=0$, constraint (IC_1) collapses to:

$$\frac{\delta}{1-\delta}\left(1-\frac{1}{2}p\right)(y-c_2) \geq c_2-c_1,$$

which is satisfied whenever (NSC^*) is satisfied. Suppose now the agent always shirks. Then

$$U\left(u^E,\pi^E\right)=(1-\delta)(y-c_1)+p\delta U\left(u^E,\pi^E\right),$$

or,

$$U\left(u^E,\pi^E\right)=\frac{1}{1-p\delta}(1-\delta)(y-c_1).$$

Given this and the fact that $u^E=y-c_2$ and $b=0$ under the proposed arrangement, (IC_1) is given by:

$$\left(1-\frac{1}{2}p\right)\delta(y-c_2) \geq (1-\delta)(c_2-c_1)+\frac{1}{2}p\delta\left(\frac{(1-\delta)(y-c_1)}{1-p\delta}-(y-c_2)\right),$$

or,

$$\left(1-\frac{1}{2}p(1+\delta)\right)\delta(y-c_2) \geq \left(1-\frac{1}{2}p\delta\right)(1-\delta)(c_2-c_1),$$

which is the same as the (NSC^*) above.

Finally, the existence of δ^* follows directly from the observation that the term

$$\frac{\delta}{1-\delta}\left(1-\frac{p}{2-p\delta}\right)$$

increasing in δ for $\delta \in (0,1)$ and $p \in (0,1)$. Hence, there exists a unique δ^* such that (NSC^*) is satisfied iff $\delta \geq \delta^*$. ■

Proof of Lemma 2. Consider a relational contract where, for some t, no information about the task is available to the agent till period t and in period t, the payoff profile (u,π) is sustained by $a=E$ induced by bonus $b \neq 0$. We construct another contract where, in the same period, (u,π) is sustained by $a=E$ and supported by $b=0$.

Step 1. Suppose that (u,π) is supported by a contract in which $w_t=w$ and $b_t<0$. Consider now a new contract (strategy) with wage and bonus (w',b') in period t, where $w'=w+b$ and $b'=0$. All other aspects of the contract remain the same, including past and future play. Observe that the new contract keeps (PK_P) and (PK_A) unaffected as $w'+b'=w+b$. Hence, the player's payoff remains (u,π). We claim that this contract

satisfies all other constraints as well, and hence, gives a payoff (u, π) in the game starting from period t by inducing $a = E$ in that period.

Step 1a. Notice the following about the constraints in period t: The new contract makes (IC_0), (IC_1) and (DE_A) slack and (DE_P) remains satisfied as $\pi^a \geq 0$ for all $a \in \{E, C, O\}$. Finally, this change also preserves the (IC_1) for all periods prior to t, ensuring that past play continues to be consistent with equilibrium (and hence the agent did not have any incentives to deviate in the past and learn the identity of the task). To see this, observe that under the original contract:

$$U(u,\pi) = \max \left\{ \begin{array}{l} (1-\delta)(w+b-c_2) + \delta \left(\alpha^C u^C + \sum_{a \in \{E,O\}} \alpha^a U(u^a, \pi^a) \right), \\ (1-\delta)(w+pb-c_1) + p\delta \left(\alpha^C u^C + \sum_{a \in \{E,O\}} \alpha^a U(u^a, \pi^a) \right) \end{array} \right\}, \tag{6}$$

and that the corresponding payoff under the new contract, denoted here by U', is obtained by substituting w and b in these expressions by b' and w', respectively. Clearly, with the proposed change in the contract, the first element (inside the curly brackets) remains the same and the second becomes smaller. This implies $U' \leq U(u, \pi)$. Moreover, since (6) holds for any period in which $a = E$ and

$$U(u,\pi) = \alpha^C u^C + \alpha^E U(u^E, \pi^E) + \alpha^O U(u^O, \pi^O)$$

in any period in which $a = O$, for any τ and $a \in \{E, O\}$, $U(u_\tau, \pi_\tau)$ is non-decreasing in $U(u^a_\tau, \pi^a_\tau)$ (if the identity of the critical task has not yet been revealed). Thus, $U'_\tau \leq U_\tau$ for all $\tau \leq t$. Hence, in any period prior to t, the agent's payoff on-the-equilibrium path remains the same and the payoff from deviating does not increase.

Step 2. Suppose now that (u, π) is supported by a contract in which $b > 0$. We show, again by construction, that it can also be supported by a contract in which $b = 0$.

Step 2a. Define

$$b^a = \frac{\delta}{1-\delta}\pi^a \times \frac{b}{\sum_{j \in \{E,C,O\}} \alpha^j \frac{\delta}{1-\delta}\pi^j}$$

for all $a \in \{E, C, O\}$. By construction, $\alpha^E b^E + \alpha^C b^C + \alpha^O b^O = b$. Furthermore, for all a

$$0 \leq b^a \leq \frac{\delta}{1-\delta}\pi^a, \tag{7}$$

where the second inequality follows from (DE_P).

Step 2b. Now, in the new contract, set the bonus equal to zero and adjust the continuation play as follows. First, suppose $\left(u^E, \pi^E\right)$ and $\left(u^C, \pi^C\right)$ are supported, respectively, by wages w^E and w^C. Now set the new wages

$$w^{a\prime} = w^a + \frac{b^a}{\delta}$$

for $a = E, C$. The principal's continuation payoffs become

$$\pi^{a\prime} = \pi^a - \frac{1-\delta}{\delta} b^a$$

for $a = E, C$. Observe that, by (7), $w^{a\prime} \geq w^a$ and $\pi^{a\prime} \geq 0$, which ensures that when the continuation play calls for (u^E, π^E) or $\left(u^C, \pi^C\right)$ both the principal and the agent will again accept the contract. Second, consider $\left(u^O, \pi^O\right)$. If $\pi^O = 0$, then nothing needs to be done in the new contract and we continue with the same continuation play dictated by $\left(u^O, \pi^O\right)$. If, otherwise, $\pi^O > 0$, then we know that players will engage in the relationship at some point. Let w^O be the wage the principal pays the agent the first time the relationship resumes, and assume that the parties have taken the outside option t periods before that. Note that when the relationship resumes, the principal's payoff is π^O/δ^t. Now let

$$w^{O\prime} = w^O + \frac{1}{\delta^{t+1}} b^O,$$

and this gives

$$\pi^{O\prime} = \delta^t \left[\frac{\pi^O}{\delta^t} - (1-\delta) \frac{1}{\delta^{t+1}} b^O \right] = \pi^O - \frac{1-\delta}{\delta} b^O.$$

Once again, by (7), $w^{O\prime} \geq w^O$ and $\pi^{O\prime} \geq 0$, which implies that both the principal and the agent accept the contract if continuation play calls for $\left(u^O, \pi^O\right)$. Hence, continuation play is again an equilibrium for $a \in \{E, C, O\}$.

Step 2c. Next, note that this change leaves (PK_P) and (PK_A) unchanged. It also leaves (IC_1) unchanged. To see this, notice that we can rewrite (IC_1) as

$$(1-\delta)\left(1-\frac{1}{2}p\right) b + \sum_{a\in\{E,O\}} \alpha^a \delta \left(u^a - \frac{1}{2} p U\left(u^a, \pi^a\right) \right) + \alpha^C \delta \left(1-\frac{1}{2}p\right) u^C \geq (1-\delta)(c_2 - c_1).$$

Now using the fact that $b = \sum \alpha^a b^a$ we can rewrite the left-hand side of this inequality as:

$$\left(1-\frac{1}{2}p\right) \sum_{a\in\{E,C,O\}} \alpha^a \left((1-\delta) b^a + \delta u^a\right) + \sum_{a\in\{E,O\}} \alpha^a \delta \left(u^a - U\left(u^a, \pi^a\right)\right).$$

Under the proposed contract change, the right-hand side of the inequality remains unchanged. In the left-hand side, the summation remains unchanged by construction. In addition, observe that the difference $u^E - U\left(u^E, \pi^E\right)$ remains unchanged. This is because the only change in the contract that may affect it is the change in the wage w^e which affects u^E and $U\left(u^E, \pi^E\right)$ in exactly the same way. In other words, u^E and $U\left(u^E, \pi^E\right)$ increase both by $(1-\delta)b^E/\delta$. Similarly, both u^O and $U\left(u^O, \pi^O\right)$ increase by $(1-\delta)b^O/\delta$, not affecting the difference $u^O - U\left(u^O, \pi^O\right)$. Therefore (IC_1) is satisfied under the new contract.

Step 2d. Finally, (IC_1) for all periods prior to t is also satisfied under the new contract. Under the original contract, $U(u, \pi)$ is again as stated in (6). The corresponding payoff under the new contract is obtained by substituting, in that expression, b with 0, u^C with $u^C + (1-\delta)b^C/\delta$, $U\left(u^E, \pi^E\right)$ with $U\left(u^E, \pi^E\right) + (1-\delta)b^E/\delta$, and $U\left(u^O, \pi^O\right)$ with $U\left(u^O, \pi^O\right) + (1-\delta)b^O/\delta$. It is easy to see that $U' = U(u, \pi)$. Since, as shown above, for any period τ, $U(u_\tau, \pi_\tau)$ is non-decreasing in $U(u^a_\tau, \pi^a_\tau)$ for $a = E, O$, it follows that for any period $\tau \leq t$, $U'_\tau \leq U_\tau$. Hence, in any period prior to t, the agent's payoff on-the-equilibrium path remains

the same and the payoff from deviating does not increase. This observation completes the proof. ■

Proof of Lemma 3. Consider a profile of payoffs (u, π) sustained by $a = E$ and suppose it is supported by wage w and bonus $b = 0$. By Lemma 2, there is no loss of generality in assuming that $b = 0$. Let w^a be the next period wage that supports the continuation payoffs (u^a, π^a) for all $a \in \{E, C\}$ in this equilibrium. Similarly, let w^O denote the wage paid the first time the relationship resumes (in case it resumes) that supports the continuation payoffs (u^O, π^O).

Next consider a strategy that is identical to the above equilibrium, except for the following changes in the current and next period wages. For all $a \in \{E, C\}$, let the new wage in the continuation game be

$$w^{a\prime} = w^a + \frac{\pi^a}{1-\delta}.$$

If $\pi^O > 0$, then the players will engage in the relationship at some point in the future. Suppose that the parties take the outside option t periods before engaging again in the relationship. Note that when the relationship resumes, the principal's payoff is π^O/δ^t. In this case, let

$$w^{O\prime} = w^O + \frac{\pi^O}{\delta^t (1-\delta)}.$$

Finally, let the new current wage be

$$w' = w - \frac{\delta}{1-\delta}\left(\alpha^E \pi^E + \alpha^O \pi^O + \alpha^C \pi^C\right).$$

Under these changes $\pi^{a\prime} = 0$ for all $a \in \{E, C, O\}$ and all the relevant constraints remain satisfied. Observe that the proposed changes increase the agent's continuation payoff and relax (IC_1). In particular, observe that $u^{a\prime} \geq u^a + \pi^a$ for all $a \in \{E, C, O\}$, while $U(u^a, \pi^a) - u^a = U(u^{a\prime}, \pi^{a\prime}) - u^{a\prime}$ for $a = E, O$. Moreover, they preserve (PK_P) and (PK_A). Constraints (DE_P) and (DE_A) are automatically satisfied since $b = 0$. Finally, under the proposed changes, the (IC'_A) constraint for all periods prior to t remains satisfied, ensuring that past play continues to be consistent with equilibrium. To see this, observe that under the original contract

$$U(u,\pi) = \max\left\{\begin{array}{l} (w - c_2)(1-\delta) + \delta\left(\alpha^C u^C + \sum_{a=E,O} \alpha^a U(u^a, \pi^a)\right), \\ (w - c_1)(1-\delta) + \delta p\left(\alpha^C u^C + \sum_{a=E,O} \alpha^a U(u^a, \pi^a)\right) \end{array}\right\}.$$

The corresponding payoff under the new contract, U', is obtained by replacing in this expression, w with w', $U(u^a, \pi^a)$ with $U(u^a, \pi^a) + \pi^a$ for $a = E, O$ and u^C with $u^C + \pi^C$. The first element inside the curly brackets remains the same under the new contract. The second element is the same minus

$$\delta(1-p)\left(\alpha^E \pi^E + \alpha^C \pi^C + \alpha^O \pi^O\right),$$

which implies that $U' \leq U(u, \pi)$. Since, as shown in the proof of Lemma 2, for any period τ, $U(u_\tau, \pi_\tau)$ is non-decreasing in $U(u^a_\tau, \pi^a_\tau)$ for $a = E, O$, it follows that for any period $\tau \leq t$,

$U'_\tau \leq U_\tau$. Hence, in any past period, the agent's payoff on-the-equilibrium path remains the same and the payoff from deviation does not increase. ■

Proof of Lemma 4. If the optimal contract is one in which the critical task is revealed at the beginning of the game, then the lemma trivially holds. Suppose now that the optimal contract begins with no revelation of the critical task. Then it must begin with $a = E$. Otherwise, it would not be optimal as it would begin with $a = O$, and a contract beginning with period two of this contract would have a higher associate payoff.

Now let t be the first period in which $\alpha^O > 0$. Let u be the agent's payoff at the beginning of that period. Notice $a = E$ is chosen, so (PK_A) implies that

$$u = (1-\delta)(y - c_2) + \delta\left(\alpha^E u^E + \alpha^O u^O + \alpha^C u^C\right),$$

where u^a, $a \in \{E, C, O\}$ are the appropriate continuation payoffs.

When the continuation play calls for exit, note that

$$u^O = \delta u_c,$$

where u_c is the agent's expected continuation payoff. Now consider the following alternative strategy. The new strategy is the same as that in the optimal contract we are considering here, except that in period t, if continuation play calls for exit (which happens with probability α^O), then the game continues in the following way: with probability $1-\delta$, players terminate the relationship forever; and with probability δ, the game continues with u_c (which could be sustained by randomization). Under this alternative strategy, the agent's payoff (following the contingency that exit is called for in the original equilibrium) is given by

$$u^{O\prime} = \delta u_c = u^O.$$

In addition, observe that

$$U\left(u^{O\prime}\right) = \delta U(u_c) = U\left(u^O\right).$$

As a result, (IC_1) is preserved. To see this, recall that (IC_1) can be written as (using the fact that $b = 0$, by Lemma 2):

$$(IC_A) \qquad \sum_{a\in\{E,C,O\}} \alpha^a u^a + \frac{1}{2}p \sum_{a\in\{E,O\}} \alpha^a \left(u^a - U(u^a)\right) \geq \frac{1-\delta}{\left(1-\frac{1}{2}p\right)\delta}(c_2 - c_1).$$

Since $u^{O\prime} = u^O$ and $U\left(u^{O\prime}\right) = U\left(u^O\right)$, both sides of the inequality are preserved with the proposed change of strategy.

Next, note that under the alternative strategy profile, the agent's continuation payoff at beginning of period t, denoted here by u', satisfies $u' = u$. Moreover, observe that

$$U(u) = \max\left\{\begin{array}{l}(1-\delta)(y-c_2) + \delta\left(\alpha^E U\left(u^E\right) + \alpha^O U\left(u^O\right) + \alpha^C u^C\right), \\ (1-\delta)(y-c_1) + \delta p\left(\alpha^E U\left(u^E\right) + \alpha^O U\left(u^O\right) + \alpha^C u^C\right)\end{array}\right\}.$$

Thus, from the fact that $U\left(u^{O\prime}\right) = U\left(u^O\right)$, it follows that $U(u') = U(u)$. Since, as shown in the proof of Lemma 2, for any period τ, $U(u_\tau, \pi_\tau)$ is non-decreasing in $U(u^a_\tau, \pi^a_\tau)$ for $a = E, O$, it follows that for any period $\tau \leq t$, $U'_\tau \leq U_\tau$. Hence, in any past period, the agent's payoff on-the-equilibrium path remains the same and the payoff from deviation does not increase, meaning that the (IC_1) constraint is satisfied for all periods prior to t. Therefore, the alternative strategy is also an equilibrium that gives the agent the same payoff as that originally considered. This implies that if the equilibrium asks players to take their

outside options in the next period, we can replace this with a probability of permanent exit. Proceeding this way, we see that temporary exit is never needed, and we can set $\alpha^O = 0$. ■

Proof of Lemma 5. If an optimal contract specifies $\alpha^C > 0$ in some period, then the discount factor δ is high enough that $a = C$ is sustained in every period following the revelation of the critical task. Thus, if in a given optimal contract the critical task is revealed in the first period, then $u = py - c_1$, necessarily. Suppose now the case of an optimal contract where the critical task is not revealed in the first period. Let t be the first period in which $\alpha^C > 0$, and let the agent's payoff in that period be u. By Lemma 4, we can write (PK_A) as

$$u = (1-\delta)(y - c_2) + \delta\left(\alpha^C u^C + \left(1-\alpha^C\right)u^E\right).$$

Now if $u^C < py - c_1 \equiv s_1$, we can consider an alternative strategy profile in which we replace u^C with

$$u^{C\prime} = s_1.$$

Under this new new strategy,

$$u' = (1-\delta)(y - c_2) + \delta\left(\alpha^C s_1 + \left(1-\alpha^C\right)u^E\right) = u + \delta\alpha^C\left(s_1 - u^C\right). \tag{8}$$

In addition, note that (IC_1) in period t is also satisfied. To see this note that under the original strategy (contract) (IC_1) in period t can be written as

$$\left(\alpha^C u^C + \left(1-\alpha^C\right)u^E\right) + \frac{1}{2}p\left(\left(1-\alpha^C\right)\left(u^E - U\left(u^E\right)\right)\right) \geq \frac{1-\delta}{\left(1-\frac{1}{2}p\right)\delta}(c_2 - c_1). \tag{9}$$

Following the change, (IC_1) in period t can be written as:

$$\left(\alpha^C s_1 + \left(1-\alpha^C\right)u^E\right) + \frac{1}{2}p\left(\left(1-\alpha^C\right)\left(u^E - U\left(u^E\right)\right)\right) \geq \frac{1-\delta}{\left(1-\frac{1}{2}p\right)\delta}(c_2 - c_1). \tag{10}$$

Since (9) is satisfied and $s_1 > u^C$, then (10) must also be satisfied. We next show that the proposed change also relaxes (9) for all $\tau < t$, so that the agent does not deviate in any past period under the new strategy. In what follows, let u_τ denote the agent's payoff in period τ, u'_τ the same payoff under the new strategy, and $\Delta = \delta\alpha^C\left(s_1 - u^C\right)$, i.e. Δ is the change in the agent's payoff in period t (see 8). Thus, $u'_t = u_t + \Delta$. Moreover, since period t is the first in which $\alpha^C > 0$, we can write

$$u_{t-k} = (1-\delta)(y - c_2) + \delta u_{t-k+1}$$

and

$$u'_{t-k} = (1-\delta)(y - c_2) + \delta u'_{t-k+1},$$

for all $k = 1, .., t-1$. This means that $u'_{t-k} = u_{t-k} + \delta^k\Delta$, or, equivalently,

$$u'_{t-k} - u_{t-k} = \delta^k\Delta. \tag{11}$$

Next observe that

$$U(u_t) = \max\left\{\begin{array}{l}(1-\delta)(y-c_2) + \delta\left(\alpha^C u^C + \left(1-\alpha^C\right)U(u^E)\right), \\ (1-\delta)(y-c_1) + \delta p\left(\alpha^C u^C + \left(1-\alpha^C\right)U(u^E)\right)\end{array}\right\}$$

and that $U(u'_t)$ is the same except that u^C is replaced with s_1. It follows that $U(u'_t)-U(u_t) \leq \Delta$. Moreover,

$$U(u_{t-k}) = \max \left\{ \begin{array}{l} (1-\delta)(y-c_2) + \delta U(u_{t-k+1}), \\ (1-\delta)(y-c_1) + \delta p U(u_{t-k+1}) \end{array} \right\}$$

and $U(u'_{t-k})$ can be obtained by replacing $U(u_{t-k+1})$ with $U(u'_{t-k+1})$ in this expression. Hence,

$$U(u'_{t-k}) - U(u_{t-k}) \leq \delta^k \Delta. \tag{12}$$

Next, observe that IC'_A in any period $t-k-1$ under the original strategy can be written as

$$(1-\delta)(y-c_2) + \delta u_{t-k} \geq (1-\delta)(y-c_1) + \delta p U(u_{t-k+1}) \tag{13}$$

and under the new strategy it can be written as

$$(1-\delta)(y-c_2) + \delta u'_{t-k} \geq (1-\delta)(y-c_1) + \delta p U(u'_{t-k+1}). \tag{14}$$

Since the former is satisfied and by (11) and (12), $u'_{t-k} - u_{t-k} \geq U(u'_{t-k}) - U(u_{t-k})$, the latter must also be satisfied. Finally, observe that the proposed change of strategy increases the agent's payoff at the beginning of the game. This shows that in any optimal contract $u^C = py - c_1$ in the first period in which $\alpha^C > 0$. Applying a similar procedure recursively we obtain that $u^C = py - c_1$ the second time $\alpha^C > 0$, and in any other period in which $\alpha^C > 0$. ■

Proof of Lemma 6. The proof is given by the following steps.

Step 1. (*Forming a relaxed problem by considering a specific deviation*) Let u_s^t be the agent's payoff when he privately knows which task is critical and always shirks by doing the critical task only (given that the principal continues to offer $w = y$ and $b = 0$ in all periods until the agent's deviation is detected). Note that $u_s^t \leq U(u_t)$ and satisfies the following recursive relation:

$$u_s^t = (1-\delta)(s_2 + c) + \delta p \left(\alpha_t u_s^{t+1} + (1-\alpha_t) s_1 \right). \tag{15}$$

So, if one restricts attention to only this type of deviation, (IC_1) could be simplified as:

$$u^t \geq (1-\delta)(s_2 + c) + \frac{1}{2} p\delta \left(\alpha_t u_s^{t+1} + (1-\alpha_t) s_1 \right), \tag{16}$$

or, equivalently,

$$2u^t \geq (1-\delta)(s_2 + c) + u_s^t. \tag{IC'_1}$$

Now, consider the following "relaxed" version of $\mathcal{P}$ where we replace (IC_1) with its weaker version (IC'_1) and ignore the (IC_0) and (SE_E) constraints:

$$\mathcal{P}_R : \max_{\alpha_t \in [0,1]} \quad u^1 \quad s.t. \ (1), \ (15), \text{ and } (IC'_1) \text{ hold for all } t.$$

Step 2. (*Rewriting $\mathcal{P}_R$ in terms of α_t*) By using (1) and (15), one can eliminate u^t and u^s_t in $\mathcal{P}_R$ and consider an equivalent problem in terms of α_ts. Note that (1) can be rearranged as $u^t - s_1 = (1-\delta)(s_2 - s_1) + \delta\alpha_t(u^{t+1} - s_1)$. So, one obtains:

$$u^t - s_1 = (1-\delta)(s_2 - s_1)(1 + \delta S_t),$$

where $S_t = \alpha_t + \delta\alpha_t\alpha_{t+1} + \delta^2\alpha_t\alpha_{t+1}\alpha_{t+2} + ...$ Hence,

$$u^1 = s_1 + (1-\delta)(s_2 - s_1)(1 + \delta S_1). \tag{17}$$

Next, note that

$$u^t_s - ps_1 = (1-\delta)(s_2 + c - ps_1) + \delta p\alpha_t\left(u^{t+1}_s - s_1\right),$$

and hence,

$$\begin{aligned} u^t_s - s_1 &= u^t_s - ps_1 - (1-p)s_1 \\ &= (1-\delta)(s_2 + c - ps_1) + \delta p\alpha_t(u^{t+1}_s - s_1) - (1-p)s_1 \\ &= (1-p)((1-\delta)y - \delta s_1) + \delta p\alpha_t\left(u^s_{t+1} - s_1\right). \end{aligned}$$

So,

$$u^t_s - s_1 = (1-p)((1-\delta)y - \delta s_1)(1 + \delta p D_t), \tag{18}$$

where $D_t = \alpha_t + (\delta p)\alpha_t\alpha_{t+1} + (\delta p)^2\alpha_t\alpha_{t+1}\alpha_{t+2} + ...$ Note that (IC_1') is equivalent to:

$$\begin{aligned} 2u^t - 2s_1 &\geq (1-\delta)(s_2 + c) - s_1 + u^t_s - s_1 \\ &= (1-\delta)(s_2 + c - s_1) - \delta s_1 + u^t_s - s_1, \quad \forall t, \end{aligned}$$

or,

$$k_0(1 + \delta S_t) \geq k_1 + k_2(1 + \delta p D_t) \quad \forall t.$$

where $k_0 = 2(1-\delta)(s_2 - s_1)$, $k_1 = (1-\delta)(s_2 + c - s_1) - \delta s_1$, and $k_2 = (1-\delta)(s_2 + c - s_1) - \delta(1-p)s_1$. Note that k_2 must be positive. Else, we can set $\alpha_t = 1$ for all t and the first-best would be feasible as a stationary contract that satisfies the relaxed (IC_1') also satisfies the original (IC_1). Now, the above constraint can be rewritten in the following form:

$$D_t \leq A + BS_t \quad \forall t, \tag{19}$$

where $A = (k_0 - k_1 - k_2)/k_2\delta p$ and $B = k_0/pk_2$. So, from (17) and (19), it follows that $\mathcal{P}_R$ is equivalent to the following program:

$$\mathcal{P}_R' : \quad \max_{\alpha_t \in [0,1]} \quad S_1 \quad s.t. \ (19).$$

Step 3. (*Rewriting $\mathcal{P}_R'$ in terms of α, S and D*) Note the following: (i) Any sequence of $\{\alpha_t\}_{t=1}^{\infty}$ pins down a unique sequence $\{(S_t, D_t)\}_{t=1}^{\infty}$. (ii) S_t and D_t are non-negative and $S_t \geq D_t$ with equality holding if and only if $\alpha_t\alpha_{t+1} = 0$. (iii) S_t and D_t follow the recursive relations:

$$S_t = \alpha_t(1 + \delta S_{t+1}), \text{ and } D_t = \alpha_t(1 + \delta p D_{t+1}).$$

(iv) The set of $\{\alpha\}$ that satisfy (19) gives rise to a set of (S, D) that are feasible. Call this set $\mathcal{F}$. It is not necessary for the proof to characterize $\mathcal{F}$ but by standard argument we know that it must be compact. Now, we can rewrite $\mathcal{P}'_R$ as follows:

$$\mathcal{P}''_R: \begin{cases} \max_{\alpha\in[0,1],\ S,\ D,\ S',\ D'} \quad S \\ s.t.\ \ S = \alpha\left(1+\delta S'\right);\ D = \alpha\left(1+\delta p D'\right) \quad (PK_R) \\ D \leq A + BS \quad (IC_R) \\ (S', D') \in \mathcal{F} \quad (SE_R) \end{cases}$$

(Note that the constraint (SE_R) implies (S', D') satisfies (IC_R), $D' \leq S'$, and $D \leq S$.) We will consider the case where $A > 0$. For $A \leq 0$, we will later argue that the firm's program does not have a solution (and hence, the interval $\left(\tilde{\delta}, \delta^*\right)$ does not exist).

Step 4. (*Introducing $f(S)$ function and defining S^**) Note that following about $\mathcal{P}''_R$. (i) The recursive relations suggest:

$$\frac{D}{S} = \frac{1+\delta p D'}{1+\delta S'}.$$

(ii) For any (S, D), we have

$$\frac{D}{S} \leq \frac{1+\delta p D}{1+\delta S} \text{ iff } D \leq \frac{S}{1+\delta\left(1-p\right)S} =: f\left(S\right).$$

Observe that $f(S)$ is increasing (and concave) and $f(S)/S$ is decreasing in S. Also, under the first-best solution where all $\alpha_t = 1$, $(S, D) = \left(S^{FB}, D^{FB}\right) = \left(\frac{1}{1-\delta}, \frac{1}{1-\delta p}\right)$ and it satisfies $D^{FB} = f\left(S^{FB}\right)$. (iii) Since the first best is not feasible by assumption, we must have $D^{FB} > A + BS^{FB}$. Hence, the $D = f(S)$ curve must intersect $D = A + BS$ at some point $(S^*, D^*) << \left(S^{FB}, D^{FB}\right)$ (since we have $A > 0$).

Step 5. (*S^* is the value of the program $\mathcal{P}''_R$*) We claim that S^* is the value of the program $\mathcal{P}''_R$. The proof is given by contradiction. Suppose that the value of $\mathcal{P}''_R$ is $\bar{S}_1 > S^*$. Let $\mathcal{D}(S)$ be the minimal D associated with all solutions that yield the value S. As $\mathcal{F}$ is compact, $\mathcal{D}$ is well-defined. Consider the tuple $\left(\bar{S}_1, \mathcal{D}\left(\bar{S}_1\right)\right)$. By the recursive relations, $\left(\bar{S}_1, \bar{D}_1\right) := \left(\bar{S}, \mathcal{D}\left(\bar{S}\right)\right)$ generates a sequence $\{\left(\bar{S}_2, \bar{D}_2\right), \left(\bar{S}_3, \bar{D}_3\right), ...\}$ such that each element of the sequence satisfies (i) $\bar{D}_n \leq A + B\bar{S}_n$ (if not, then (19) would be violated in some period) and (ii) the recursion relations (PK_R) for some associated sequence of $\alpha_t, \{\bar{\alpha}_t\}$ (say). We will argue in the next four sub-steps (steps 5*a* to 5*d*) that such a sequence cannot exist:

Step 5a. *We argue that* $\left(\bar{S}_1, \bar{D}_1\right) >> \left(\bar{S}_2, \bar{D}_2\right)$. First, observe that for all $S \in (S^*, S^{FB})$, $f(S) > A + BS$. As $\bar{S}_1 > S^*$, $f\left(\bar{S}_1\right) > A + B\bar{S}_1 \geq \bar{D}_1 = \mathcal{D}\left(\bar{S}_1\right)$. Next, we claim that $f\left(\bar{S}_2\right) \geq \bar{D}_2$.

The proof is given by contradiction: suppose $f\left(\bar{S}_2\right) < \bar{D}_2$. But then we have $\bar{S}_2 < S^*$. The argument is as follows: Clearly, if $\bar{S}_2 = S^*$, the highest feasible $\bar{D}_2$ that could support $\bar{S}_2$ is $f(S^*)$ and hence there is no feasible $\bar{D}_2$ such that $f\left(\bar{S}_2\right) < \bar{D}_2$. Now suppose $\bar{S}_2 > S^*$.

Since $f(S) > A + BS$ for all $S > S^*$ and $A + B\bar{S}_2 \geq \bar{D}_2$, it must be that $f\left(\bar{S}_2\right) > \bar{D}_2$. Hence, $f\left(\bar{S}_2\right) < \bar{D}_2 \Rightarrow \bar{S}_2 < S^*$.

Therefore, if $f\left(\bar{S}_2\right) < \bar{D}_2$, we obtain that:

$$\frac{\bar{D}_1}{\bar{S}_1} = \frac{1 + \delta p \bar{D}_2}{1 + \delta \bar{S}_2} > \frac{1 + \delta p f\left(\bar{S}_2\right)}{1 + \delta \bar{S}_2} = \frac{f\left(\bar{S}_2\right)}{\bar{S}_2} > \frac{f(S^*)}{S^*}, \tag{20}$$

where both equalities follow from (PK_R), the first inequality holds as $f\left(\bar{S}_2\right) < \bar{D}_2$ and the second inequality holds as $\bar{S}_2 < S^*$ (argued above) and $f(S)/S$ is decreasing in S. But as $\bar{S}_1 > S^*$ and $f\left(\bar{S}_1\right) > \bar{D}_1$ we must also have,

$$\frac{f(S^*)}{S^*} > \frac{f\left(\bar{S}_1\right)}{\bar{S}_1} > \frac{\bar{D}_1}{\bar{S}_1},$$

which contradicts (20). Hence, we must have $f\left(\bar{S}_2\right) \geq \bar{D}_2$.

As $f\left(\bar{S}_2\right) \geq \bar{D}_2$, we obtain:

$$\frac{\bar{D}_1}{\bar{S}_1} = \frac{1 + \delta p \bar{D}_2}{1 + \delta \bar{S}_2} \geq \frac{\bar{D}_2}{\bar{S}_2}.$$

As $\bar{S}_2 \leq \bar{S}_1$ (since $\bar{S}_1$ is assumed to be the highest S_1 feasible), the above inequality implies that we must have $\bar{D}_2 \leq \bar{D}_1$.

Step 5b. *We must have* $\bar{\alpha}_2 = 1$. We show this by contradiction. From (PK_R) we know that $\left(\bar{S}_2, \bar{D}_2\right) = \left(\bar{\alpha}_2\left(1 + \delta \bar{S}_3\right), \bar{\alpha}_2\left(1 + \delta p \bar{D}_3\right)\right)$. If $\bar{\alpha}_2 < 1$, increase $\bar{\alpha}_2$ to $\alpha_2' := (1 + \varepsilon)\,\bar{\alpha}_2$ for some $\varepsilon > 0$. Let $\left(S_2', D_2'\right) := (1 + \varepsilon)\left(\bar{S}_2, \bar{D}_2\right)$.

We argue that for sufficiently small ε, $\left(S_2', D_2'\right)$ is feasible. Since $\left(\bar{S}_3, \bar{D}_3\right) \in \mathcal{F}$ and (PK_R) is trivially satisfied by definition of $\left(S_2', D_2'\right)$, it is enough to show that $\left(S_2', D_2'\right)$ satisfies (IC_R). To see this, recall that $\bar{D}_1/\bar{S}_1 \geq \bar{D}_2/\bar{S}_2$ (from step 5a) and $\bar{S}_2 \leq \bar{S}_1$. So, $\left(\bar{S}_2, \bar{D}_2\right)$ must lie on or below the line joining the origin to $\left(\bar{S}_1, \bar{D}_1\right)$. Now, there are two cases: (i) If (IC_R) is slack at $\left(\bar{S}_1, \bar{D}_1\right)$, all points on this line always lie strictly below the line $D = A + BS$. So, (IC_R) is also slack at $\left(\bar{S}_2, \bar{D}_2\right)$. (ii) If (IC_R) binds at $\left(\bar{S}_1, \bar{D}_1\right)$, this is the only point on the line at which (IC_R) binds, and it is slack at all other points. But, as $f\left(\bar{S}_1\right) > \bar{D}_1$, we have:

$$\frac{1 + \delta p \bar{D}_1}{1 + \delta \bar{S}_1} > \frac{\bar{D}_1}{\bar{S}_1} = \frac{1 + \delta p \bar{D}_2}{1 + \delta \bar{S}_2}.$$

So, $\left(\bar{S}_2, \bar{D}_2\right) \neq \left(\bar{S}_1, \bar{D}_1\right)$. Therefore, (IC_R) must be slack at $\left(\bar{S}_2, \bar{D}_2\right)$. Thus, for small enough ε, $\left(S_2', D_2'\right) = (1 + \varepsilon)\left(\bar{S}_2, \bar{D}_2\right)$ always satisfies (IC_R).

Next, observe that,

$$\frac{\bar{D}_1}{\bar{S}_1} = \frac{1 + \delta p \bar{D}_2}{1 + \delta \bar{S}_2} > \frac{1 + \delta p\,(1 + \varepsilon)\,\bar{D}_2}{1 + \delta\,(1 + \varepsilon)\,\bar{S}_2} = \frac{1 + \delta p D_2'}{1 + \delta S_2'}.$$

Now, we reduce $\bar{\alpha}_1$ to some α_1' where $\alpha_1'(1+\delta S_2') = \bar{S}_1$. Let $D_1' = \alpha_1'(1+\delta p D_2')$. So, by the above inequality, we find that:

$$\frac{D_1'}{\bar{S}_1} = \frac{1+\delta p D_2'}{1+\delta S_2'} < \frac{\bar{D}_1}{\bar{S}_1}.$$

Hence, $D_1' < \bar{D}_1$. But this observation contradicts the fact that $\bar{D}_1$ is the lowest feasible D_1 that supports S_1 (as we have shown that the sequence $\{\alpha_1', \alpha_2', \bar{\alpha}_3, ...\}$ is feasible, and it yields $S_1 = \bar{S}_1$ and $D_1 = D_1' < \bar{D}_1$). Therefore, we must have $\bar{\alpha}_2 = 1$.

Step 5c. *We must have* $\bar{S}_3 < \bar{S}_2$ *and* $\bar{D}_3 < \bar{D}_2$. As $\bar{\alpha}_2 = 1$, (PK_R) implies $\bar{S}_2 = 1+\delta\bar{S}_3$ and $\bar{D}_2 = 1+\delta p\bar{D}_3$. As $\bar{S}_t < S^{FB} = 1/(1-\delta)$ and $\bar{D}_t < D^{FB} = 1/(1-\delta p)$ for any t, it is routine to check that $\bar{S}_3 < \bar{S}_2$ and $\bar{D}_3 < \bar{D}_2$.

Step 5d: Repeating steps 5b and 5c we can argue that $\bar{\alpha}_t = 1$ for all $t \geq 2$ and the sequence $\{\bar{S}_2, \bar{S}_3, ...\}$ is monotonically decreasing. So, we must have $\bar{S}_t = 1+\delta\bar{S}_{t+1}$, $t = 2, 3, ...$ But such a sequence cannot exist. First, note that this sequence cannot converge. If it converges at some $\hat{S}$, we must have $\hat{S} = 1+\delta\hat{S}$, or $\hat{S} = S^{FB} = 1/(1-\delta)$, which is not a feasible as all terms of the sequence is bounded away from $\bar{S}_1 < S^{FB}$. Therefore, some term of this sequence will be either negative or zero. But we know that $\bar{S}_t$ is non-negative. Also, suppose $\bar{S}_k = 0$. So, we must have $\bar{S}_{k-1} = \bar{D}_{k-1} = \bar{\alpha}_{k-1}$. But this is a contradiction as we know that $\bar{S}_{k-1} = \bar{D}_{k-1}$ only if $\bar{\alpha}_{k-1}\bar{\alpha}_k = 0$ but we have $\bar{\alpha}_{k-1} = \bar{\alpha}_k = 1$.

Step 6. $\mathcal{P}_R''$ does not have any solution if $A \leq 0$. Note that in this case any feasible (S, D) must be such that $D < f(S)$. But then, by argument identical to one presented in step 5a to 5d we can claim that there cannot exist a solution to $\mathcal{P}_R''$.

Step 7. **(**S^* *can be implemented by a stationary contract***)** As $D^* = f(S^*)$,

$$\frac{D^*}{S^*} = \frac{1+\delta p D^*}{1+\delta S^*}.$$

Define

$$\alpha^* := \frac{S^*}{1+\delta S^*} = \frac{D^*}{1+\delta p D^*}.$$

Notice that the stationary sequence $\alpha_t = \alpha^*$ for all t is a solution to $\mathcal{P}_R'$ as it yields $S_1 = S^*$ and the resulting sequence $\{(S_t, D_t)\} = \{(S^*, D^*)\}$ satisfies (19).

Step 8. **(**α^* *is a solution to the original problem* $\mathcal{P}$**)** We now show that optimal contract $\{\alpha^*\}$ satisfies (IC_1), (IC_0) and all (SE), and hence it is also a solution to $\mathcal{P}$. We show this in the following three sub-steps:

Step 8a: As the contract is stationary, the agent who privately learns the critical task does not have any deviation that is more profitable than always striking by doing the critical task only. That is, we must have $u_s^t = U(u_t)$. Hence, the optimal contract $\{\alpha^*\}$ satisfies (IC_1).

Step 8b: To show that (IC_0) is satisfied we need to show that under the optimal contract $u^t \geq (1-\delta)y$ for all t. If the interval $[\tilde{\delta}, \delta^*)$ exists, it must be the case that for all δ in this

interval, it is feasible to induce $e_1 = 2$ by setting $\alpha^* = 0$ (and hence $e_t = 1$ for all $t > 1$). Thus, (IC_0) must be satisfied for such a contract and we must have:

$$(1-\delta)\, s_2 + \delta s_1 \geq (1-\delta)\, y.$$

But

$$u^t = s_1 + (1-\delta)(s_2 - s_1)(1 + \delta S_t) \geq (1-\delta)s_2 + \delta s_1.$$

Thus, we obtain $u^t \geq (1-\delta)\, y$.

Step 8c: Finally, to check that (SE)s are satisfied, note that: (i) $\left(u^O, 0\right) = (\underline{u}, 0)$ which is the payoff of the trivial PBE where both players take their outside options in all periods. (b) From the definition of $\mathcal{E}_K$ in Section 3 and the fact that $u^C = py - c_1 = s_1$, we know $\left(u^C, 0\right) \in \mathcal{E}_K$. Finally, (iii) in the proposed contract, $u^t = u^*$ for all t and $(u^*, 0) \in \mathcal{E}$ by construction given in the proof above. Hence, $\{\alpha_t\} = \{\alpha^*\}$ is a solution to the original problem as well. ■

Proof of Proposition 2. Note that $\tilde{\delta} < \delta^*$ if and only if at δ^* (NSC^*) binds whereas the contract specifying $e_1 = 2$ and $\alpha^* = 0$ (hence $e_t = 1$ for all $t > 1$) remains feasible. In other words, this contract must satisfy both (IC_1) and (IC_0), i.e.,

$$(1-\delta^*)\, s_2 + \delta^* s_1 \geq (1-\delta^*)\,(y - c_1) + \frac{1}{2} p\delta^* s_1. \tag{21}$$

and

$$(1-\delta^*)\, s_2 + \delta s_1 \geq (1-\delta^*)\, y. \tag{22}$$

From (21) and (22) we have:

$$\frac{\delta^*}{1-\delta^*}\left(1 - \frac{1}{2}p\right) s_1 > \frac{\delta^*}{1-\delta^*}\left(1 - \frac{p}{2 - p\delta^*}\right) s_2,$$

and

$$\frac{\delta^*}{1-\delta^*} s_1 - c_1 > \frac{\delta^*}{1-\delta^*}\left(1 - \frac{p}{2 - p\delta^*}\right) s_2.$$

Simplifying, we obtain the two conditions:

$$s_1 - \frac{2 - p - p\delta^*}{2 - p - p\delta^* + \frac{1}{2}p\delta^*} s_2 > 0 \text{ and } s_1 - \frac{2 - p - p\delta^*}{2 - p\delta^*} s_2 > \frac{1-\delta^*}{\delta^*} c_1.$$

The rest of the proof immediately follows from Proposition 1, Lemma 1, and 6. ■

Proof of Proposition 3. By Proposition 2, we know that for any v, if there exists a solution to $\mathcal{P}$, there exists a stationary solution, $\alpha_t = \alpha_R(v)$, say. Let $u^1(v)$ be the associate value function. As the jobs are identical and in the optimal contract the agent retains all surplus we require $v = u^1(v)$ (as v is the maximum surplus in any given job). The following argument shows that such a value of v exists: Let $\mathcal{E}$ be the set of all PBE payoffs. By standard argument, $\mathcal{E}$ is compact; trivially $\min\{u \mid (u, \pi) \in \mathcal{E}\} = 0$ and let $\max\{u \mid (u, \pi) \in \mathcal{E}\} = v^*$. Note that by definition $v^* \geq u^1(v^*)$. Let $\alpha_R^* := \alpha_R(v^*)$. It could be readily shown that in the optimal contract, the agent faces a constant probability of job rotation in all periods and in all jobs that he may be placed in the course of the play. Hence, we cannot have $v^* > u^1(v^*)$ (as v^* is simply the value of the program when the agent starts at an identical job in the next period where the exact same contract is offered) and it must be that $v^* = u^1(v^*)$.

Finally, notice that the (NSC^*) condition does not depend on s_1 and hence, the δ^* cutoff for feasibility of first-best remains unaltered. From (IC_0) and (19) (representing (IC_1)) we find that $e_t = 2$ and $\alpha^t = 0$ if feasible if and only if

$$\frac{\delta}{1-\delta}(y - c_2) \geq \psi + \frac{c_2 - c_1}{1 - p/2}.$$

(It is routine to check that (IC_0) is also satisfied under this condition.) Let δ_R be the value of δ for which the above constraint is binding and notice that if ψ is sufficiently larger, we have $\delta_R > \delta^*$, i.e, the interval $[\delta_R, \delta^*)$ does not exist. ■

References

[1] Abreu, Dilip. 1988. "On the Theory of Infinitely Repeated Games with Discounting." *Econometrica*, 56 (6): pp. 383–96.

[2] Abreu, Dilip, David Pearce, and Enno Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.

[3] Aoyagi, M. 2010. "Information Feedback in a Dynamic Tournament." *Games and Economic Behavior*, 70 (2): pp. 242–260.

[4] Baker, George, Gibbons, Robert and Kevin Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *The Quarterly Journal of Economics*, 109(4); pp. 1125–56.

[5] Baker, George, Gibbons, Robert and Kevin Murphy. 2002. "Relational Contracts and the Theory of the Firm." *The Quarterly Journal of Economics* 117 (1); pp. 39–84.

[6] Bardhan, P. 1997. "Corruption and Development: A Review of Issues." *Journal of Economic Literature* 35; pp. 1320–46.

[7] Bergemann, Dirk, and Ulrich Hege. 2005. "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics* 36; pp. 719–752.

[8] Battaglini, M. 2005. "Long-Term Contracting With Markovian Consumers." *American Economic Review* 95 (3); pp. 637–658.

[9] Bhaskar, V. 2014. "The Ratchet Effect Re-examined: A Learning Perspective." Mimeo, University of Texas-Austin.

[10] Bonatti, Alessandro, and Johannes Hörner. 2011. "Collaborating." *American Economic Review* 101; pp. 632–663.

[11] Bonatti, Alessandro, and Johannes Hörner. 2015. "Learning to Disagree in a Game of Experimentation." Mimeo, Yale University.

[12] Chassang, S. 2010. "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts." *American Economic Review* 100 (1); pp. 448–465.

[13] Deb, Joyee, Jin Li, and Arijit Mukherjee. 2016. "Relational Contracts with Subjective Peer Evaluations." *The RAND Journal of Economics* 47; pp. 3–28.

[14] Ederer, Florian. 2010. "Feedback and Motivation in Dynamic Tournaments." *Journal of Economics & Management Strategy* 19 (3); pp. 733–769.

[15] Ederer, Florian, Richard Holden, and Margaret Meyer. 2014. "Gaming and Strategic Opacity in Incentive Provision." Cowles Foundation Discussion Paper No. 1935.

[16] Fong, Yuk-Fai, and Jin Li. 2016. "Information Revelation in Relational Contracts" forthcoming, *Review of Economic Studies*.

[17] Gawande, Atul. 2010. The Checklist Manifesto: How to get Things Right. New York: Metropolitan Books.

[18] Guo, Yingni. 2016. "Dynamic Delegation of Experimentation." *American Economic Review* 106; pp. 1969–2008.

[19] Goltsman, Maria, and Arijit Mukherjee. 2011. "Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure." *Journal of Labor Economics* 29; pp. 229–265.

[20] Halac, Marina, Navin Kartik, and Qingmin Liu. 2016. "Optimal Contracts for Experimentation." *Review of Economic Studies* 83, 1040–91.

[21] Hörner, Johannes, and Larry Samuelson. 2013. "Incentives for Experimenting Agents." *RAND Journal of Economics* 44; pp. 632–663.
[22] Ishihara, A. 2016. "Relational Contracting and Endogenous Formation of Teamwork." Mimeo, Kyoto University.
[23] Lazear, Edward. 2006. "Speeding, Terrorism, and Teaching test." *Quarterly Journal of Economics*; pp. 1029–61.
[24] Levin, Jonathan. 2003. "Relational Incentive Contracts." *American Economic Review* 93 (3); pp. 835–57.
[25] Li, Jin, Niko Matouschek and Mike Powell. 2016. "Power Dynamics in Organizations." forthcoming, *American Economic Journal: Microeconomics*.
[26] Manso G. 2011. "Motivating Innovation." *Journal of Finance* 66; pp. 1823–69.
[27] Mailath, George, and Larry Samuelson. 2006. Repeated Games and Reputations. Oxford: Oxford University Press.
[28] Malcomson, James. 2013. "Relational Incentive Contracts." In The Handbook of Organizational Economics, edited by Robert Gibbons and John Roberts, 1014–65. Princeton: Princeton University Press.
[29] Malcomson, James. 2016. "Relational Incentive Contracts With Persistent Private Information." *Econometrica* 84 (1); pp. 317–346.
[30] Moroni, Sofia. 2016. "Experimentation in Organizations." Working Paper, University of Pittsburgh.
[31] Mukherjee, Arijit, and Luis Vasconcelos. 2011. "Optimal Job Design in the Presence of Implicit Contracts." *RAND Journal of Economics* 42; pp. 44–69
[32] Meyer, Marshall. "Rethinking performance measurement: beyond the balanced scorecard." Cambridge; New York: Cambridge University Press, 2002.
[33] Orlov, Dmitry. 2016. "Optimal Design of Internal Disclosure." Mimeo, University of Rochester.
[34] Schöttner, Anja. 2008. "Relational Contracts, Multitasking, and Job Design." *Journal of Law, Economics, and Organization* 24; pp. 138–162.
[35] Yang, Huanxing. 2013. "Non-stationary Relational Contracts with Adverse Selection." *International Economic Review* 54 (2); pp. 525-547.
[36] Zabojnik, Jan. 2014. "Subjective evaluations with performance feedback." *The RAND Journal of Economics* 45 (2); pp. 341–369.