LEARNING-BY-SHIRKING IN RELATIONAL CONTRACTS*

JIN LI†, ARIJIT MUKHERJEE‡, LUIS VASCONCELOS§

ABSTRACT. An agent may privately learn which aspects of his responsibilities are more important by shirking on some of them and use that information in the future to shirk more effectively. We study the optimal provision of relational incentives in the presence of such learning-by-shirking. In a model of long-term employment relationship, we characterize the optimal contract and highlight how ancillary policies such as job reorganization and information disclosure on task priorities sharpen relational incentives by diluting the agent's information rents from shirking. In spite of the learning-by-shirking effect, the optimal contract is stationary and may call for stochastic reorganization/disclosure policies.

1. INTRODUCTION

A common incentive problem in many principal-agent relationships is that an agent may attempt to cut corners at the principal's expense. While the literature on incentive theory typically assumes that the agent exactly knows the consequences of shirking, in many contexts, that may not be the case. The agent may lack information on the relative importance of his assigned tasks in the overall production process and, relatedly, he may not know which corners to cut so as to minimize the risk of getting caught. As a result, he may start out shirking "in the dark." However, once he shirks successfully, he may learn how to shirk more effectively in the future. This possibility of "learning by shirking" exacerbates the incentive problem, because shirking, when successful, provides the agent with valuable private information that he may use later on to cut corners that are harder to notice.

The labor market is filled with examples where a worker has imperfect information about the consequences of shirking. A management consultant working for a client under a tight deadline may have to cut corners in some parts of his report but he may not know which parts of the report are more important to the client. Hospitals may require doctors to follow

Date: September 21, 2017.

*For their helpful comments and suggestions, we thank Ricardo Alonso, Dan Barron, Alessandro Bonatti, Gonzalo Cisternas, Jon Eguia, William Fuchs, Robert Gibbons, Johannes Hörner, Alessandro Lizzeri, Niko Matouschek, Santiago Oliveros, Michael Powell, Canice Prendergast, Andrzej Skrzypacz, Rani Spiegler, Jeroen Swinkels, Huanxing Yang, and the attendees at various seminar and conferences at Boston University, CUNEF, Indian Statistical Institute, Michigan State University, Southern Methodist University, University of British Columbia, Universidad Carlos III de Madrid, University of East Anglia and University of Essex. Any errors that may remain are ours.

†Strategy Department, Kellogg School of Management, Northwestern University, 2211 Campus Drive, Evanston, IL 60208. Email: jin-li@kellogg.northwestern.edu.

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

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

a checklist designed to reduce infection risks, but to save time, a doctor may skip some of the steps. He may not know, however, which items on the checklist are more crucial for reducing the probability of infection. A university professor considers both teaching and research as important components of her job; however, she may not know how exactly her performances in teaching and research affect her salary.[1]

Not knowing the exact consequence of shirking, and relatedly, not knowing how best to shirk, are also features common in relationships outside the labor market. In business-to-business relationships, upstream suppliers manufacturing complicated products such as engines, hard drives, and cell phone screens, may not know whether their downstream buyers care more about product compatibility, durability, or reliability. In business-to-government relationships, manufacturing plants facing government's workplace safety inspections may be required to follow a host of safety measures but may not know which measures are most critical. More broadly, any organization—universities, hospitals, banks, law firms—facing a rating agency may be imperfectly aware of how different dimensions of performance affect their final rating.

But in all these settings, if the agent cuts corners and does not face any consequence, he will privately learn which job aspects are relatively less important or less likely to be monitored. If the worker shirks on some of the job requirements and the employer fails to notice it, the worker will privately learn which aspects of his job are more important and may use this information to shirk more effectively in the future. Similarly, if a supplier neglects certain aspects of product design and receives no complains from the client, the supplier may privately learn which product features are more important for product performance or are actually being monitored more closely by the client. A similar type of learning may occur in the case of businesses-to government relationships or relationships between institutions and the agencies that rate them.

In this article we analyze the optimal incentives provision in such ongoing relationships when "learning by shirking" exacerbates the moral hazard problem. We model the long-term relationship as a relational contract between a firm and a worker where the firm offers incentives through a discretionary bonus that is tied to a non-verifiable performance metric (see Malcomson, 2013).[2] In particular, we consider an infinitely repeated employment relationship where the worker performs a job that consists of two tasks (or aspects). The first-best outcome requires the worker to exert (costly) private effort in both tasks. 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. 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, but the worker may privately acquire this information if he shirks and happens to pick the right task, and he can then use his private information to shirk more effectively in the future. To provide incentives to the

[1]Gawande (2010) provides several examples of complex production environments from a varied set of industries, e.g., healthcare, financial services, civil engineering, and aviation, where the workers are expected to follow a "check list" of tasks in order to ensure successful job completion.

[2]Relational incentives are commonplace in many industries, particularly in complex jobs with multiple aspects, where verifiable performance measures well-aligned with the firm's goal could be difficult to obtain (see Baker, Gibbons, and Murphy, 1994; Levin, 2003).

worker, the firm promises a discretionary bonus for a satisfactory job performance. As the job-performance measure is not verifiable, such a promise is a part of a relational contract that is sustained through a threat of future retaliation by the worker should the firm renege on its bonus payment.

Notice that while we model the uncertainty about the consequences of shirking in terms of the uncertainty about the underlying production process, in many real-life settings, the uncertainty may stem from the agent's lack of information about what aspects of his job are more closely monitored.[3] Indeed, we can recast our model to capture such a scenario. One may assume that both tasks are equally critical, and the principal monitors one of them more intensively than the other. The agent may privately learn about the monitoring intensities when he shirks on a given task. As we will discuss later, this setting is essentially equivalent to our model both in terms of the analysis and the findings.

We first show that in our model, the optimal relational contract has a simple characterization and it is closely tied to the future surplus in the relationship—i.e., the firm's "reputational capital"—captured by the players' common discount rate $\delta \in (0, 1)$. For δ above a cutoff, the firm can credibly offer a large bonus that induces the worker to work on both tasks, and the first-best surplus is attained. But for lower values of δ, such a strong incentive is not feasible. Since the worker might shirk by cutting corners, potentially leading to severe consequences, it is optimal to dissolve the relationship. It is important to note that while the qualitative features of the optimal contract are similar in spirit to its canonical model counterpart (see, e.g., Baker, Gibbons, Murphy, 1994), the analysis is considerably different. It must address the issue that when shirking is not detected, the beliefs about the task identities in the continuation game are no longer common knowledge. And consequently, the agent enjoys an information rent off-equilibrium when he privately learns about the tasks through shirking.

Next, we explore the following question: When the bonus incentives are not sufficient, how can the firm sharpen incentives by adopting ancillary policies that erode the worker's rents from private learning? We consider such policies in two distinct scenarios—one where the firm can access the relevant information, and one where it cannot. In the former case, the firm may directly affect the worker's rents by strategically disclosing information on the relative importance of different job aspects. And in the latter case, the firm may still affect the worker's rents by strategic job reorganization that renders the present information on the relative importance of different job aspects irrelevant in the future.

First, we consider the case of strategic job reorganization. Suppose that at the end of each period, the firm can reorganize the job environment (by incurring a cost) so as to stochastically change the identity of the critical task in the subsequent period. A common example of such reorganization is job rotation, where the agent may be assigned to different jobs over time, with the jobs being similar but for the identity of the underlying critical task. To fix ideas, consider the consulting-firm example discussed earlier, and suppose that the

[3]Since managerial attention is necessarily limited, monitoring is often imperfect and knowing which aspects the managers will focus can be valuable to the worker. Recent economic literature has emphasized the role of managers in affecting the productivity of the firm (Lazear, Shaw and Stanton, 2015; Hoffman and Tadelis, 2017), and a number of theoretical models have focused on managerial (in)attention (Dessein and Santos, 2016; Dessein, Galleotti, and Santos, 2016; Halac and Prat, 2016; also see Gibbons and Henderson, 2012, for a review).

consultant may be asked to work on similar projects every period but for potentially different clients. The projects are similar, but which aspect of a given project is more critical depends on the client for whom the agent is working. At the end of each period, the firm may assign the consultant to another project for the same client or may assign him to a new project for a different client by incurring an administrative cost. Similar policies are also common in 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 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).

We highlight how the amount of the firm's "reputational capital" affects the firm's job-reorganization policy. When δ (and consequently, the firm's stock of reputational capital) is too large or too small, the firm's job-reorganization policy does not play any role in the optimal incentive provision. For δ sufficiently large, the firm can credibly promise a large enough bonus pay that induces the agent to exert effort in both tasks (i.e., the efficient effort level) even if the job environment is never reorganized (as job reorganization is costly, it is never used in the optimal contract). For δ sufficiently small, no effort can be induced even if the job environment is reorganized every period, as the lack of reputational capital renders any promise of bonus pay non-credible.

But for an intermediate range of δ, the optimal contract sharpens relational incentives through a stochastic job-reorganization policy (provided the cost of reorganization is not too large). At the end of every period, the firm reorganizes the job environment with a constant probability, and the worker exerts effort in both tasks in every period. The possibility of reorganization dissuades the worker from "learning by shirking" by diluting his information rents from privately learning how to cut corners. As the worker's gains from superior information may only last for a short period of time, he becomes less inclined to shirk. Consequently, such a policy bolsters incentives and allows the firm to elicit effort in both tasks in every period. The optimal probability of reorganization is driven by the trade-off between the cost of such reorganization and the benefits of sharper incentives that it creates.[4]

Next, we explore strategic information disclosure as a channel of incentive provision. We consider a variation of our main model that abstracts away from the possibility of job reorganization. As the worker's job environment does not change over time, the identity of the critical task also remains unaltered. However, we assume that the firm can publicly access and reveal 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.

How does the optimal disclosure policy relate to the size of the reputational capital? As in the case of the optimal job-reorganization policy, the disclosure policy does not play any role if δ is too large or too small. When δ is too large, effort in both tasks can be induced irrespective of the firm's disclosure policy; and if δ is too small, no effort can be induced.

[4]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.

However, for a moderate δ, the firm must actively manage information to sharpen incentives. When δ is relatively large (but still within the intermediate range), opacity is essential for attaining efficiency. In contrast, for δ relatively small (but still within the intermediate range), the optimal contract calls for full transparency. The firm reveals the critical task at the beginning of the game, and the worker performs the critical task only in all periods. Finally, for an intermediate range of δ, the firm adopts a stochastic stationary disclosure policy that resembles the optimal job-reorganization policy discussed earlier: at the end of every period, the firm discloses the critical task with a constant probability (if it has not yet been made public). The worker exerts effort on both tasks until the critical task is revealed, but once it is revealed, he works only on the critical task in all future periods.

An implication of the stochastic disclosure policy is that to an outside observer, the firm may appear to be failing in the long run as, almost surely, its performance declines over time. There is vast literature on the causes of organizational failures (see Garicano and Rayo, 2016, for a review) that identifies the lack of proper incentives as a key factor. In contrast, our findings suggest that a gradual decline in organizational performance could be an unavoidable by-product of the incentive policy needed to sustain a higher surplus at the earlier stages of the relationship.

Related Literature: The key contribution of this paper is to highlight the nature of optimal incentives in the presence of "learning by shirking." This is done in the context of relational contracts when the firm may use ancillary policies—job reorganization and information disclosure—to sharpen incentives.

This paper departs from the standard incentive theory by considering that an agent may privately learn about certain aspects of his job when he shirks and use that information to shirk more effectively in the future. An important implication of such private learning is that the beliefs of the contracting parties may differ and cease to be common knowledge throughout the relationship.

Because of this feature, in terms of analytic structure, our paper is related (and contributes) to the literature on long-term and relational contracts in which the posterior beliefs of the contracting parties diverge and cease to be common knowledge following a deviation by the agent, see, e.g., Bergmann and Hege, 2005; Fuchs, 2007, Bonatti and Hörner, 2011; Bhaskar, 2014; Fong and Li, 2017. The lack of common knowledge implies that the standard recursive technique à la Abreu, Pearce, and Stacchetti (1990) cannot be readily applied.[5] While it is difficult to characterize the optimal contract in such environments, we are able to do so 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.

Our paper is also related to a few strands of literature in organization economics. The policies that we consider—job reorganization and information disclosure—have received considerable attention in recent times. In terms of job reorganization, the literature has typically focused on how policies, such as job rotation and implementation of new performance

[5]The lack of common knowledge is also na feature in models with persistent private information such as Battaglini (2005), Yang (2013), and Malcomson (2016).

measures, could be used to elicit information about a worker's productivity (Ortega, 2001; Pastorino, 2017), or to prevent the worker from gaming the performance metric (Meyer, 2002). Here, we argue that job reorganization can improve incentives by dissuading a worker from shirking to learn about certain aspects of the job. In particular, we show that it helps sustain more efficient relational contracts.[6]

The literature on strategic information disclosure in employment relationships has primarily focused on two kinds of information: the employer's private information on the agents' performance (e.g., Fuchs, 2007; Aoyagi, 2010; Ederer, 2010; Mukherjee, 2010; Goltsman and Mukherjee, 2011; Zabojnik, 2014; Orlov, 2016; Fong and Li, 2017) and information on the compensation rule used by employers—i.e., what aspects of 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 is 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. Unlike Lazear (2006), which considers monitoring and information disclosure in a static setting, we explore the role of transparency in incentive provision in a dynamic context and highlight how the firm's reputational capital drives its disclosure policy.

Finally, our work is also related to the literature on incentives for experimentation (see Bolton and Harris, 1999; Keller, Rady, and Cripps, 2005; 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. 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 strategic job reorganization. A benchmark case is analyzed in Section 3 which assumes away the possibility of job reorganization or information revelation. The optimal job reorganization is studied in Section 4. In Section 5 we adapt our baseline model to show how information revelation, instead of job reorganization, could be used to address the moral hazard problem in our environment. A final section concludes. All proofs are provided in the Appendix (and its online supplement).

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 $t \in$

[6]Other instruments used to sustain relational contract include formal contracts (Baker, Gibbons, and Murphy, 1994), integration decisions (Baker, Gibbons, and Murphy, 2002), ownership design (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, 2017).

$\{1, 2, ..., \infty\}$. 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: $\mathbb{A}$ and $\mathbb{B}$. The agent must exert effort to complete the job and privately chooses an effort level $e_t \in \{0, 1_{\mathbb{A}}, 1_{\mathbb{B}}, 2\}$. Effort is costly to the agent, and we denote the cost by the function $C(e_t)$. If the agent works on both tasks, $e_t = 2$, and his cost of effort is $C(2) = c_2$; but if he works on either one of the two tasks, $e_t = 1_{\mathbb{A}}$ or $1_{\mathbb{B}}$, depending on whether he works on task $\mathbb{A}$ or $\mathbb{B}$, 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, $e_t = 0$, and his cost of effort is $C(0) = 0$.

The job output, $Y_t \in \{-z, 0, y\}$, is assumed to be observable but not verifiable. The job may be successfully completed, leading to output $y > 0$, it may remain incomplete, with a 0 output, and, at the extreme, the agent may completely fail at the job, leading to 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 one 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 noncritical task, the output is $-z$ with certainty.

The identity of the critical task is governed by the underlying production environment. At the beginning of the game, neither player 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. Also, at the end of any period, the principal may publicly "reorganize" the job (or the production environment in general) at a cost ψ that randomly changes the identity of the critical task in the subsequent period. We assume that the identity of the critical task pre- and post-reorganization is statistically independent, but in a given job the identity of the critical task remains unchanged over time until the job is reorganized. We denote the principal's reorganization decision as $\gamma_t \in \{0, 1\}$, where $\gamma_t = 1$ if the principal reorganizes at the end of period t, and $\gamma_t = 0$ otherwise.

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 associated with his current job if he happens to pick the right task by chance. As we will see later, this possibility of private "learning by shirking" has significant implications for the optimal relational contract. Also note that in our setting, the reorganization of a job does not affect the agent's productivity and, hence, is completely wasteful but for its incentive implications, on which we will elaborate below.

CONTRACT: In each period t, the principal decides 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 $Y_t = y$. The incentives are relational as the output is assumed to be non-verifiable. Also, as discussed later, we restrict parameters such that it is never optimal for the principal to allow the agent to experiment (see Assumption 1 below). Hence, our contract specification is without loss of generality as $Y_t = -z$ is never realized on the equilibrium path.

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 that.

Finally, as is typical in the repeated game literature, we assume the existence of a public randomization device to convexify the equilibrium payoff set. In particular, 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.

The timeline of the stage game is illustrated in Figure 1 below.

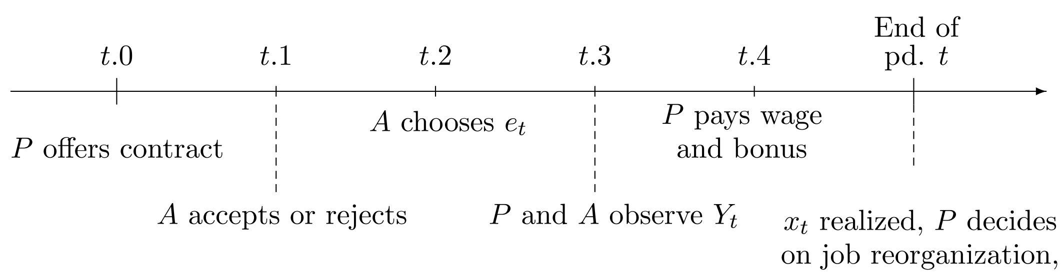

Figure 1. Timeline of the stage game.

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 $d_t^A = d_t^P = 1$, the respective expected payoffs for the agent and the principal are given as

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

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 may find it profitable to use this information in future deviations. Thus, the restriction to pure strategy PPE may lead to some loss of generality, and hence, we focus on the perfect Bayesian Equilibrium (PBE) of the game defined as follows:

Let $h_t = \left\{d_\tau^A, d_\tau^P, Y_\tau, w_\tau, b_\tau, x_\tau, \gamma_\tau\right\}_{\tau=1}^{t-1}$ denote the public history of the game at the beginning of period t and H_t be the set of all such histories (note that, $H_1 = \{x_0\}$). The strategy of the principal consists of a sequence of functions $\sigma_P = \left\{D_t^P, W_t, B_t, \Gamma_t\right\}_{t=1}^{\infty}$, where her participation decision is given by $D_t^P : H_t \to \{0,1\}$, the contract offer is given as $W_t : H_t \to \mathbb{R}$ and $B_t : H_t \cup \{Y_t\} \to \mathbb{R}$, and finally, the reorganization decision is given as $\Gamma_t : H_t \cup \left\{d_t^A, d_t^P, Y_t, w_t, b_t, x_t\right\} \to \{0,1\}$. The agent's strategy, however, may depend on his private history $\tilde{h}_t = \left\{d_\tau^A, d_\tau^P, e_\tau, Y_\tau, w_\tau, b_\tau, x_\tau, \gamma_\tau\right\}_{\tau=1}^{t-1}$, which not only records the public history but also includes information on the agent's past effort provisions. Let $\tilde{H}_t$ be the set of all such private histories. The agent's strategy is a sequence of functions $\sigma_A = \left\{D_t^A, E_t\right\}_{t=1}^{\infty}$, where his participation decision is given as $D_t^A : \tilde{H}_t \cup \left\{d_t^P, w_t, b_t\right\} \to \{0,1\}$, and his effort decision is given as $E_t : \tilde{H}_t \cup \left\{d_t^A, d_t^P, w_t, b_t\right\} \to \{0, 1_{\mathbb{A}}, 1_{\mathbb{B}}, 2\}$. Finally, denote $\mu_t = \Pr\left(\textit{task } \mathbb{A} \textit{ is crucial}\right)$ as the belief of the agent in period t about the identity of the critical task in his current job. 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 the agent has privately learned the identity of the critical task by shirking at some period in the past and the principal has not reorganized since then.

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 maintain a few restrictions on the parameters to focus on a more interesting modeling environment and to streamline the analysis.

Assumption 1. *(i)* $y - c_2 > py - c_1 > 0$*, (ii)* $\frac{1}{2}pc_2 > c_1$*, and (iii)* $(1-\delta) \times \left(\frac{1}{2}(py - z) - c_1\right) + \delta(y - c_2) < 0$.

Under Assumption 1 (i), efficiency requires the agent to exert effort on both tasks, and exerting effort only on the critical task, if known, is more efficient than dissolving the employment relationship. Assumptions 1 (ii) and (iii) simplify the analytical tractability of the optimal contracting problem. Assumption 1 (ii) stipulates that the cost of exerting effort on both tasks relative to only one is assumed to be sufficiently large. It ensures that the incentives needed to deter the agent from shirking on exactly one of the two tasks (i.e., choosing $e_t = 1_{\mathbb{A}}$ or $1_{\mathbb{B}}$ instead of $e_t = 2$) is also sufficient to deter him from shirking on both (i.e., choosing $e_t = 0$). Finally, Assumption 1 (iii) rules out the optimality of experimentation. It implies that it is always better to dissolve the relationship than to ask the agent to perform a randomly selected task in any given period (even if the efficient outcome is played in all future periods). This condition is trivially satisfied when z is sufficiently large.

We conclude this section with the following remarks about our model. First, we have been purposefully agnostic about the organizational changes that are entailed in the principal's reorganization decision. The exact nature of the reorganization is not essential to our analysis as long as the reorganization is interpreted as any organizational change that the principal may implement (at some cost) in order to change the identity of the critical task stochastically. A typical example of such a reorganization decision is job rotation, where the principal may assign the agent to different jobs over time within the organization. However, the model could be easily adapted to study other policies, e.g., disclosure of task information, that may have similar incentive implications. We will revisit this issue later in Section 5.

Second, as mentioned in the Introduction, our model could be readily adjusted to consider a setting where the agent may shirk to learn the underlying monitoring process rather than the production technology. For example, one may assume that the principal evaluates the agent's performance at each task (rather than using an aggregate output). Both tasks are equally critical for production: the consequences of neglecting a task does not depend on the task identity. However, only one the two tasks is always monitored while the other one is checked with some probability.[7] The agent may attempt to learn about the monitoring technology by shirking on a task while the principal may "reorganize" the job setting by varying the monitoring process over time. The analysis of this setup is identical to that of our model. (However, in this setting, the question of disclosure of task information is moot, as the tasks are assumed to be identical.)

Finally, our model also 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 in our model as effort in a single task where the high and low effort surely lead to high and low output, respectively. But, the output for the intermediate effort can be high or low depending on the underlying state of the world that is a priori unknown to all. Such a set up is qualitatively similar to our multitasking model when the state of the world is job specific but stationary for a given job environment. The worker may shirk in order to learn the underlying state, but the current information on the state is irrelevant in the future if the firm reorganizes the job. However, we adopt the multitasking setup, as it is more realistic in the context of strategic information revelation on relative task importance—an alternative to job reorganization that we also consider later in our analysis.

3. A benchmark case and the feasibility of first-best

We begin our analysis by considering a benchmark scenario where job reorganization is assumed to be infeasible. This benchmark case is useful for our subsequent explorations for at least three reasons: First, it characterizes the optimal relational contract when incentives

[7]In particular, suppose $Y = y$ if $e = 2$, $Y = y$ with probability μ and $-z$ with probability $1 - \mu$ if $e = 1$, *irrespective of* the task chosen, and finally, $Y = -z$ if $e = 0$. Now, if task $\mathbb{A}$ is monitored with probability one but $\mathbb{B}$ is monitored with probability $\gamma \in (0, 1)$, then task $\mathbb{A}$ essentially plays the same role as that of the "critical" task in our model since the agent can successfully shirk with probability $p := \mu (1 - \gamma)$ if he works on $\mathbb{A}$ only, but would surely be caught if he works only on $\mathbb{B}$.

can only be offered through discretionary bonuses, helping to draw out the role of job reorganization in sharpening incentives. Second, our analysis yields a necessary and sufficient condition for the feasibility of the "first-best" surplus—the surplus in the relationship when the agent exerts effort on both tasks in the job in every period and the principal never reorganizes the job environment (i.e., $e_t = 2$ and $\gamma_t = 0$ for all t. Recall that in our setting job reorganization is purely wasteful but for its impact on the agent's effort incentives). Third, it also highlights an important technical aspect of our analysis. In our model, the optimal contract need not be stationary, and the game does not have a tractable recursive structure. Consequently, the standard method to characterize the equilibrium payoff set (à la Abreu, Pearce, and Stacchetti, 1990) no longer applies. Moreover, we cannot limit attention to the class of stationary contracts (as in Levin, 2003) without any loss of generality.

We present our analysis in three steps. First, we state the set of constraints that must hold in a given period under an efficient contract. Next, given the constraints, we derive a necessary and sufficient condition for an efficient contract to be feasible. And finally, using this condition, we present a complete characterization of the optimal contract.

Let the set of PBE payoffs for a given δ be $\mathcal{E}$. Take a $(u, \pi) \in \mathcal{E}$ that is efficient, i.e., the payoffs associated in an equilibrium where the agent exerts effort on both tasks in all periods and $u + \pi = y - c_2$. Moreover, let w and b be the wage and bonus and $\left(u^N, \pi^N\right)$ be the continuation payoffs that sustain (u, π).[8] We assume that following any publicly observed deviation from the equilibrium actions, the relationship terminates.[9]

By virtue of being an equilibrium payoff with efficient actions, (u, π) must satisfy a set constraints. First, the agent and the principal's participation constraints must hold:

$$(IR) \qquad u \geq 0, \text{ and } \pi \geq 0.$$

Also, the consistency requirement of payoff decomposition implies that a player's payoff must be equal to the weighted sum of his current and continuation payoffs. Hence, we must have the following "promise-keeping" constraints:

$$(PK_A) \qquad u = (1 - \delta)(w - c_2 + b) + \delta u^N,$$

$$(PK_P) \qquad \pi = (1 - \delta)(y - w - b) + \delta \pi^N.$$

Next, 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. Recall that if the agent shirks on both tasks, $Y_t = -z$ and the relationship necessarily terminates. Hence, we must have:

---

[8]We use superscript N in the continuation payoffs to stress the fact that we are considering the case when there is *no* job reorganization. Also, note that w, b, u^N, and π^N are all functions of (u, π). We do not explicitly state these as functions in order to streamline the notation.

[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 the bonus promise or failing to conform to the equilibrium play in the continuation game, or both.

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

But if the agent shirks by exerting effort on exactly one of the two tasks, the derivation of the incentive compatibility constraint is somewhat more involved. It must allow for the fact that upon deviating, the agent may privately learn the identity of the critical task, and he may use this information to shirk again in the future. As a result, the principal and the agent (following a deviation) would have different beliefs on the task identities, and this lack of common knowledge renders the game devoid of a tractable recursive structure.

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 in his current job. That is, suppose (σ'_P, σ'_A) is the strategy profile of the players that gives 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 supported with different equilibrium strategy profiles that are associated with different maximal deviation payoff for the agent, without loss of generality, we choose the equilibrium strategy profile where the agent's maximal deviation payoff is the lowest (when the agent knows which task is critical in his current job). This allows (u', π') to be a sufficient statistic for the agent's maximal deviation payoff. We can now state the agent's incentive compatibility constraints using the deviation payoff U:

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

Two remarks are in order: First, 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; hence, the principal would fail to detect such a deviation. Second, (IC_1) highlights the information value of shirking. The key difference between this constraint and its counterpart in the standard moral hazard is that the continuation payoff following shirking is $U\left(u^N, \pi^N\right)$ instead of u^N, and their difference, $U\left(u^N, \pi^N\right) - u^N$, reflects the agent's information rents from privately learning which task is critical.

Notice that for any $(u, \pi) \in \mathcal{E}$, $U(u, \pi) - u \geq 0$, since the agent can always disregard his superior information, and such rents from learning-by-shirking aggravate the moral hazard problem. Also, if the job environment could be reorganized or if the information on the critical task can be publicly revealed, we have $U(u, \pi) - u = 0$, since the agent loses his information advantage. When this difference is strictly positive, it implies that there are further gains from shirking in the future. In other words, by acquiring knowledge about the production technology via deviation, the agent may increase his gains from future deviation.[10]

Next, the contract must satisfy the "dynamic enforceability" constraint to ensure that neither the principal nor the agent has incentives to renege on the bonus:

[10]Note that when the agent privately learns which task is critical, it may not be the case that he always shirks by just performing the critical task whenever he is asked to put in effort on both tasks. The agent may want to wait for the right time to shirk. In particular, in a period when the agent's equilibrium payoff is high, he may not want to shirk because there will be too much to lose. But the agent may be more inclined to shirk when his equilibrium payoff is low.

$$(DE_P) \qquad -(1-\delta)\, b + \delta \pi^N \geq 0,$$

$$(DE_A) \qquad (1-\delta)\, b + \delta u^N \geq 0.$$

Finally, we also have the "self-enforcing contract" constraints requiring the continuation payoffs themselves to be equilibrium payoffs in the continuation game:

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

In light of the above constraints, we can now characterize the optimal contract in this benchmark case. While an analytical expression for the deviation payoff $U(u, \pi)$ is elusive, the deviation gains $U(u, \pi) - u$ can be bounded below by using the above set of constraints. Using this bound, we obtain a necessary and sufficient condition for the existence of such an efficient contract (i.e., the agent never shirks while his job environment remains unchanged from period to period).

Lemma 1. *An efficient relational contract can be sustained if and only if*

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

The characterization of the optimal contract follows directly from the (NSC^*) condition.

Proposition 1. ***(Optimal contract in benchmark case)*** *Suppose that job reorganization is not feasible. The optimal relational contract is characterized as follows: There exists a δ^* (δ at which (NSC^*) is binding) such that if $\delta \geq \delta^*$, the agent exerts effort in both tasks in all periods. Otherwise, no effort is induced and the players take their outside options.*

Proposition 1 suggests that when δ is sufficiently large the principal can rely on the discretionary bonus alone to generate adequate effort incentives and induce the agent to exert effort on all tasks. This is intuitive, as with a high δ, there is enough surplus in the relationship such that the principal can credibly promise a large enough bonus to dissuade the agent from shirking even when the agent is sure to continue to work in the same environment. In spite of the fact that the information value of learning-by-shirking may be substantial, a sizable bonus payment mitigates the agent's temptation to shirk, as he would not risk terminating the relationship by choosing to work only on the task that happens to be the noncritical one.

An immediate implication of this finding is that even when job reorganization is feasible—as in our main model—the first-best surplus is attained if and only if $\delta \geq \delta^*$: in the optimal relational contract, the principal would never reorganize the job environment, but the agent would continue to exert effort on both tasks in all periods. But what is the optimal contract if $\delta < \delta^*$? Clearly, in our benchmark setup, termination of the relationship is optimal. But can the principal use job reorganization to sharpen incentives and attain a higher surplus? If so, what is the optimal job-reorganization policy? Also, can information disclosure play a similar role in incentive provision? These are the central questions that we address in the next two sections.

4. Optimal contract with job reorganization

Job reorganization may strengthen incentives as it depletes the information value of shirking. Any information on the critical task that the agent may learn through shirking is rendered useless if he is asked to work in a reorganized setting in the next period. Thus, if $\delta < \delta^*$ (and hence, effort on both tasks cannot be induced if the agent continues to work in the same job environment), one may expect that rather than terminating the relationship, a larger surplus may be attained through strategic use of job reorganization. If so, what is the optimal job-reorganization policy? Clearly, reorganizing every period would generate strong incentives, as the knowledge of the critical task that the agent may acquire by shirking in the current period would become irrelevant in the very next one. But such frequent reorganizations would lead to a large loss of surplus relative to the first-best as reorganization is costly.

Thus, the analysis of the optimal contract with a job-reorganization policy must minimize such loss while preserving the agent's effort incentives in every period. Below, we limit our attention to the case where $\delta < \delta^*$, and we begin with a discussion on the constraints that a contract needs to satisfy if it is to induce effort on both tasks in a given period (i.e., $e_t = 2$).

Consider a period t such that the critical task is not known to the agent, i.e., either $t = 1$ or the agent has not shirked successfully since the most recent reorganization of the production environment. Let the set of all PBE payoffs of the repeated game starting from 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 on the tasks is the same as his prior. Consider an equilibrium payoff $(u, \pi) \in \mathcal{E}$ that is sustained by eliciting effort on both tasks in the current period (i.e., $e_t = 2$).

Notice that at the end of any period, the equilibrium strategies may call for one of the following three action profiles for the next period: (i) the agent exerts effort on both tasks working in the same job environment as in the previous period; (ii) the agent exerts effort on both tasks in a new job environment (i.e., the principal has reorganized the job at the end of the previous period); and (iii) both players take their outside options in that period. For expositional clarity, we denote these three actions as $a = N$ ("no reorganization"), R ("reorganization"), and O ("outside option"), respectively. Recall that by Assumption 1 (iii), it is never optimal for the relationship to have the agent randomly select (and perform) exactly one of the two tasks. Also, using the public randomization device, the players could randomize over these three action profiles. Suppose that under the equilibrium strategy profile (that supports (u, π)), the action $a \in \{N, R, O\}$ 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.

If a PBE induces effort on both tasks in the current period, the associated equilibrium payoffs, contracts, and the continuation payoffs must satisfy a similar set of constraints to those we presented earlier in our benchmark analysis. However, the constraints need to account for the fact that the continuation game may call for any of the three possible action profiles in the following period: effort in the same job environment ($a = N$), effort following reorganization ($a = R$), or taking the outside option ($a = O$). The new set of constraints is given as follows:

The participation constraint remains the same as before:

$$(IR^*) \qquad u \geq 0, \text{ and } \pi \geq 0.$$

The "promise-keeping" constraints are modified as follows (notice that should the principal reorganize the job at the end of the period, the associated cost ψ is realized in the current period):

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

$$(PK_P^*) \qquad \pi = (1-\delta)(y - w - b) + \delta\left(\alpha^N \pi^N + \alpha^R\left(\pi^R - \frac{1-\delta}{\delta}\psi\right) + \alpha^O \pi^O\right).$$

The "incentive compatibility" requires that:

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

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

Notice that when the job environment is reorganized, the agent's information on the critical task in the current job becomes irrelevant. As he loses the informative value from shirking, we have $U\left(u^R, \pi^R\right) = u^R$.

The "dynamic enforceability" ensures neither party reneges on the bonus and the principal would not renege on his reorganization promise (notice that a relational contract not only specifies the bonus payment but also the action profiles in each period):

$$(DE_P^*) \qquad -(1-\delta)\, b + \delta\left(\alpha^N \pi^N + \alpha^R\left(\pi^R - \frac{1-\delta}{\delta}\psi\right) + \alpha^O \pi^O\right) \geq 0,$$

$$(DE_P^*\text{-}R) \qquad -(1-\delta)\,\psi + \delta \pi^R \geq 0,$$

and

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

And finally, we need to impose a "self-enforcing contract" constraint for each of the three action profiles (i.e., under strategies that specify $a = N,\ R$, or O to be played in the next period):

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

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

and

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

4.1. **Preliminary analysis.** In order to characterize the optimal contract, we begin by presenting a set of lemmas that simplify our subsequent analysis. These lemmas state several observations about any PBE payoff that is sustained by $e_t = 2$ in a *given* period when the information on the critical task (in the current job) is unknown to the agent, and they allow us to restrict attention to a specific class of contracts without any loss of generality. The proofs are given in the online appendix. (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 4.2.)

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

Lemma 2. *Consider a relational contract and take any period t and any history h_t. Suppose that the critical task for the period t job is not known to the agent, and in the game starting from period t, the payoff profile $(u, \pi) \in \mathcal{E}$ is sustained by effort on both tasks and $b \neq 0$ in period t. Then there exists another contract where the payoff (u, π) can be sustained by effort on both tasks 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 effort on both tasks and a negative bonus b in a given period t. Such a contract is payoff equivalent to one where $e_t = 2$, 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 effort in both tasks 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.

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 (net of reorganization cost, if any) to the principal in each period.

Lemma 3. *Consider a relational contract and take any period t and any history h_t. Suppose that the critical task for the period t job is not known to the agent, and in the game starting from period t, the payoff profile $(u, \pi) \in \mathcal{E}$ is sustained by effort on both tasks in period t and $\pi^a > 0$ for some $a \in \{N, R, O\}$. Then there exists another contract where the continuation payoff (u, π) can be sustained by effort on both tasks in period t and $\pi^N = \pi^R - \frac{1-\delta}{\delta}\psi = \pi^O = 0$.*

The intuition behind this observation is similar to that of Lemma 2 discussed earlier: any contract supporting (u, π) with a strictly positive continuation value (net of reorganization cost) in some period t can be replaced by one that (i) sets $\pi^N = \pi^R - \frac{1-\delta}{\delta}\psi = \pi^O = 0$, (ii) increases the agent's continuation payoff u^a by raising the wage in the continuation game, w^a, that supports u^a (for all $a \in \{N, R, O\}$), and (iii) reduces the current period wage 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. Note that by virtue of Lemma 2 and 3, (DE_A^*), (DE_P^*), and (DE_P^*-R) are automatically satisfied, and hence, could be dropped from the set of constraints that the optimal contract must satisfy.

The next lemma suggests that without loss of generality, we can consider only those contracts that never 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.*

The reason is 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.

Finally, let v be the maximal joint payoff feasible in $\mathcal{E}$. Lemma 5 below suggests that following a job reorganization, we may set $u^R = v - \frac{1-\delta}{\delta}\psi$, i.e., the maximum surplus feasible in the continuation game net of the (discounted) cost of reorganization.

Lemma 5. *In an optimal relational contract, in any period, if $\alpha^R > 0$, then $u^R = v - \frac{1-\delta}{\delta}\psi$.*

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

$$w = \begin{cases} y & \textit{if } a = N \textit{ is played} \\ y - \frac{\psi}{\delta} & \textit{if } a = R \textit{ is played} \end{cases}.$$

That is, 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. We focus on this class of contracts only for technical convenience, though other forms of implementation may be feasible.

4.2. **Characterization of the optimal relational contract.** Using the above lemmas, we can now simplify the optimal contracting problem. Notice that Lemma 4 implies that when deriving the optimal contract, we can restrict attention to contracts where in each period t the principal asks the agent to exert effort either in the same job environment as before (i.e., set $a = N$) or in a reorganized environment (i.e., set $a = R$). Let $1 - \alpha_t$ be the probability that the principal reorganizes at the end of period t. Note that the optimal relational contract is completely determined by the sequence $\{\alpha_t\}_{t=1}^{\infty}$.

To streamline exposition, we present the optimal contracting problem in two steps: first, we derive the optimal contract for an arbitrary value of the agent's continuation payoff $u^R = s_1$, say, where s_1 is taken as a parameter that satisfies two conditions: (i) $s_1 < s_2 := y - c_2$, the surplus generated when there is no reorganization and the agent exerts effort in both tasks; and (ii) $\left(u^R, \pi^R\right) = (s_1, 0)$ can be sustained as an equilibrium payoff in the continuation game. Next, we characterize the optimal contract with job reorganization by considering a specific value for s_1 as given in Lemma 5. As we will argue later, the optimal contract with strategic information disclosure can also be derived from this analysis by considering the appropriate value of s_1 in that setting.

With a slight abuse of notation, let u^t be the agent's (average) payoff at the beginning of period t when the critical task associated with the current production environment is not known to the agent. Also, below we write $U(u)$ instead of $U(u, \pi)$ since the principal's continuation payoff remains 0. Now, for a given s_1, we have the following recursive relationship for the agent's payoff:

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

Therefore, if there exists a contract that implements $a = N$ at least in the first period, solving for the optimal contract (in the class of such contracts) is tantamount to finding the optimal sequence $\{\alpha_t\}_{t=1}^{\infty}$ to maximize u^1. In other words, the optimal contract, for a given s_1, must solve the following program (denote $c := c_2 - c_1$):

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

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 in his current production setting through shirking). The complexity stems from the fact that once the agent learns the task identities, the profitability of his future deviations would depend on the associated reorganization policy.

For example, fix a production environment and suppose that the contract calls for the next reorganization after a certain number of periods, T (say), but the agent has shirked and learned the underlying 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 as long as the production environment remains unchanged. Notice that under the above reorganization policy, the agent's continuation payoff decreases over time as we move closer to date T. Hence, it may be worthwhile for the agent to continue to exert effort on both tasks until some period $\bar{t}$, where $t < \bar{t} < T$ (as he stands to lose a larger continuation payoff if he shirks immediately after his first successful deviation) and then start to shirk again by working on the critical task only (when the continuation payoff is smaller).

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 until the job environment is reorganized again, he always deviates by choosing the critical task only (if he is not detected sooner). 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 reorganizes the job at a fixed probability. We further argue that this policy, by virtue of being a stationary one, 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 reorganizes the production environment 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 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 reorganization policy were to deter deviation in period 1 only, it would take a form that features "early reorganization": there is some T such that the principal would reorganize with positive probability if $t < T$, but would never do so again afterwards. By reorganizing early, this policy backloads the rewards to the agent as much as possible. To see why it is useful to backload the rewards, note that compared to the agent who always works on both tasks, the agent who has shirked successfully and only works on the critical task is effectively less patient—the former discounts the future at rate δ, but the effective discount rate of the latter is $p\delta$, as the relationship is likely to terminate for him with probability $1-p$. Since an agent who has shirked discounts the future more (relative to an agent who has not), early reorganization most effectively discourages the agent from shirking in period 1 by backloading the rewards as much as possible.

However, such an early reorganization policy is necessarily time-inconsistent. While the agent is deterred from deviating in period 1, he may want to deviate in the later periods (while working in the same production environment), when the gains from shirking are larger. (As the principal is less likely to reorganize 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 ideally induce an increasing sequence in the continuation payoff by increasing the current period's reorganization probability and decreasing the probability of future reorganization. But as this needs to be done for every period, the resulting optimal policy becomes stationary and features a constant reorganization 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. From Lemma 6 we know that the optimal contract is stationary for any s_1, and Lemma 5 suggests that in the optimal contract with job rotation, we must have $s_1 = v - \frac{1-\delta}{\delta}\psi$, where v is the value of the optimal contracting problem under such a s_1 (notice that upon job reorganization, the principal would choose the optimal contract in the continuation game, which is identical to the game at the beginning of period one). Using these two observations, along with Lemma 1, we obtain the following proposition.

Proposition 2. ***(Optimal contract with job reorganization)*** *The optimal contract 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 principal never reorganizes the job environment, and in every period the agent exerts effort on both tasks (as he continues to work in the same production environment). The optimal contract yields the first-best surplus.

(ii) For all $\delta \in [\delta_R, \delta^*)$, the optimal contract fails to attain the first-best surplus and is given as follows: At the end of the period, the principal reorganizes the production environment with a constant probability α^* (which may vary with δ), and the agent exerts effort on both tasks in every period. Moreover, $\delta_R < \delta^*$ if and only if the cost of reorganization ψ is below a threshold.

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

The intuition for the above proposition is as follows: Recall that the optimal contract either induces effort on both tasks in all periods or no effort at all—by Assumption 1 (iii), dissolving the relationship is better than performing only one task chosen at random out of the two. Now, as noted in Lemma 1, for a large δ (i.e., if $\delta \geq \delta^*$), efficiency is feasible:

the principal never reorganizes the production environment, and in every period the agent continues to exert effort on both tasks. In contrast, for δ sufficiently small (i.e., δ below δ_R), the optimal policy dissolves the relationship. The maximum bonus that the principal can credibly promise, no matter what reorganization policy is used, fails to induce effort on both tasks.

But for a moderate δ—if $\delta \in [\delta_R, \delta^*)$—the principal may be able to induce effort on all tasks by adopting a stochastic job-reorganization policy where at the end of each period, the principal may reorganize the production environment with a fixed probability. As the agent knows that his information on the critical task for the current production setting may become irrelevant in the near future, it dilutes the value of the information that he hopes to take advantage of by shirking and learning the critical task privately. Recall that a reorganized production environment is assumed to be identical to its predecessors except for the identity of the critical task. Hence, if the threat of reorganization provides a strong enough incentive for effort on both tasks in the current production environment, it provides similar incentives in the agent's next environment as well (should the environment be reorganized in the future). Also notice that such a stochastic reorganization policy would entail a loss of surplus relative to the first-best, as reorganization is costly. Hence, such a policy could be optimal if and only if the cost of reorganization is not too large. Otherwise, it is always better to dissolve the relationship than to attempt to induce effort with the threat of job reorganization.

5. Optimal contract with information revelation

As mentioned earlier, any organizational policy 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 the reorganization of the production environment as one such policy. Another policy that one may consider is the strategic revelation of the information on the critical task if the principal can access and divulge this information at her discretion.

The model of job reorganization can be readily adapted to analyze such an information-disclosure policy. We keep all aspects of our initial model unchanged except for the following: Assume that the agent works in the same job environment every period (i.e., the identity of the critical task is invariant over time) and, at the beginning of the game, neither the principal nor the agent knows the identity of the critical task. However, at the end of each period, the principal can publicly access and disclose this information at zero cost. Using the notations of our initial model, we say that $\gamma_t = 1$ if information is accessed and disclosed at the end of period t, and $\gamma_t = 0$ otherwise.

Notice that if $\gamma_t = 1$ for some t, in contrast to our earlier model, the information on the critical task remains public in all subsequent periods. Also, if the information is revealed at the very beginning of the game, (say, $\gamma_0 = 1$), the game boils down to the canonical relational contracting model, and the optimal contract in this setting has a simple characterization. (The proof is given in the online appendix.)

Lemma 7. *If the information on the critical task is made public at the beginning of the game, the optimal relational contract is characterized as follows: There exist two cutoffs,* $\underline{\delta}$

and $\overline{\delta}$, where $\underline{\delta} \leq \delta^* < \overline{\delta}$ (δ^* as defined in Lemma 1), such that (i) if $\delta \geq \overline{\delta}$, the agent exerts effort on both tasks in every period, (ii) if $\underline{\delta} \geq \delta \geq \overline{\delta}$, the agent exerts effort only on the critical task in every period, and (iii) if $\delta < \underline{\delta}$, no effort can be induced, and the parties take their outside options in every period.

For a large δ, the relationship is sufficiently valuable and the principal can credibly promise a large bonus to induce the agent to work on both tasks even when he knows which task is critical. For a moderate δ, however, the principal cannot credibly promise such a large bonus, but she can still promise enough to induce effort on one of the two tasks. Hence, the agent is asked to work on the critical task only (as by Assumption 1 (i), the surplus is still larger than the outside option of the players). Finally, for δ sufficiently small, any credible bonus promise is too small to induce any effort, and the parties take their outside options.

An implication of the above lemma is that it is easier to induce effort on both tasks when the critical task is unknown to all than when it is public information (i.e., $\delta^* < \overline{\delta}$). When the task information is public, shirking yields a higher payoff to the agent, as he knows on which task to shirk (when the agent does not know the critical task, he expects to choose it only half of the time). Hence, the principal must offer a stronger incentive to elicit effort on both tasks. If $\delta \in [\delta^*, \overline{\delta})$, 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.

Therefore, when the principal can strategically filter the task information, the optimal contract attains the first-best surplus as long as $\delta \geq \delta^*$, but the underlying disclosure policy may vary. If $\delta \geq \overline{\delta}$, the first-best can be attained irrespective of the principal's disclosure policy. But if $\delta \in [\delta^*, \overline{\delta})$, the optimal contract calls for complete opacity.

What is the optimal disclosure policy if $\delta < \delta^*$? As Lemma 7 suggests, the principal can induce effort on the critical task by revealing the task identities at the beginning of the game. But the principal may be able to attain a larger surplus by relying on a different information-revelation policy where the disclosure of information is delayed. The qualitative features of the optimal information-disclosure policy bear strong resemblance to the optimal job-reorganization policy analyzed earlier. Hence, for the sake of brevity, we elaborate below only on those aspects of our analysis that differ from the one presented above in section 4.

If $\delta < \delta^*$, in any period, there are three possible action profiles on the equilibrium path: (i) agent exerts effort on both tasks while no information is revealed, (ii) the principal reveals which task is critical and the agent exerts effort on that task only, and finally, (iii) both parties take their outside options. As before, with a slight abuse of notation, we continue to denote these three cases as $a = N$, R, and O, respectively. 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 $a \in \{N, R, O\}$.

Consider the constraints that a contract must satisfy in order to sustain effort on both tasks in a given period when the critical task remains unknown to all. These constraints are identical to their counterpart in section 4 except for the following two differences: (i) in the principal's promise-keeping constraint (PK_P^*) and dynamic enforceability constraint (DE_P^*), the term $\pi^R - \frac{1-\delta}{\delta}\psi$ is replaced by π^R (notice that information revelation is assumed to be costless); and (ii) the sequential enforceability constraint following job reorganization (SE_R^*) is replaced by

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

where $\mathcal{E}_K$ denotes the set of equilibrium payoffs (for a given δ) when the critical task is publicly known.

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 is always 0 (i.e., $\pi^N = \pi^R = \pi^O = 0$), (iii) termination is never used (i.e., $\alpha^O = 0$), and finally, (iv) in the optimal contract, if $\alpha^R > 0$ in any period, then $u^R = py - c_1$ (i.e., the maximal surplus in the relationship when the critical task is revealed and the agent works on that task only). That is, we may restrict attention to contracts where, in any period, $b = 0$ and

$$w = \begin{cases} y & \textit{if } a = N \textit{ is played} \\ py & \textit{if } a = R \textit{ is played} \end{cases}.$$

Hence, the optimal contracting problem is now essentially the same as the principal's program $\mathcal{P}$ studied in section 4.2 in the context of job reorganization, except for the following modifications: (i) we now denote $1 - \alpha_t$ as the probability that the principal reveals the critical task at the end of period t, given that it has not yet been revealed; and (ii) the agent's continuation payoff $s_1 = py - c_1$. As one would expect, the optimal contracts in these two settings also share similar characteristics.

Proposition 3. *(**Optimal contract with information revelation**) The optimal contract is characterized as follows. There exist four cutoffs $\underline{\delta} < \tilde{\delta} \leq \delta^* < \overline{\delta}$ (δ^* as defined in Lemma 1) such that the following holds:*

(i) For all $\delta \geq \overline{\delta}$, the optimal contract attains the first-best surplus (i.e., the agent exerts effort on both tasks in all periods) 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 surplus and the principal never reveals the identity of the critical task.

(iii) For all $\delta \in [\tilde{\delta}, \delta^*)$, the first-best surplus 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

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

(iv) For all $\delta \in [\underline{\delta}, \tilde{\delta})$, the first-best surplus 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 3 closely parallels the characterization of the optimal contract with job reorganization that we discussed above and ties the optimal disclosure policy 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 discloses all information), and if the available surplus is in an intermediate range, active filtering of information through a stochastic disclosure policy is optimal.

Notice that all parts of this proposition directly follow from Lemma 7, except for part (iii): As discussed in the context of Lemma 7, for a sufficiently large or sufficiently small δ—i.e, if $\delta < \underline{\delta}$ or $\delta \geq \overline{\delta}$—the principal's information revelation policy plays little role in the optimal contract. But for an intermediate range of δ, active management of information is critical. Within this range, when δ is relatively large ($\delta \in [\delta^*, \overline{\delta})$), full opacity is optimal, whereas a relatively small δ ($\delta \in [\underline{\delta}, \tilde{\delta})$) calls for full transparency.

But for moderate values of δ ($\delta \in [\tilde{\delta}, \delta^*)$), 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 the private information that he hopes to obtain through learning-by-shirking. Such a contract elicits effort on 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.

Notice that a stochastic revelation policy can be optimal if and only if the condition given in part (iii) is satisfied—i.e., the surplus with effort on the critical task only ($s_1 = py - c_1$) is not too small compared to the first-best surplus ($s_2 = y - c_2$). Indeed, the loss of surplus due to information revelation, $s_2 - s_1$, plays the same role as the cost of reorganization, ψ, in our initial model. 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 agent's gains from 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 will 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) 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 will emerge as the optimal relational contract.

An important implication of the optimality of stochastic disclosure is that the performance of the organization decreases over time. The agent performs both tasks at the beginning of the relationship. And once the critical task is revealed, he works on the critical task only, causing the performance to fall almost surely in the long run. Our model therefore adds to the broad literature of why organizations fail (Garicano and Rayo, 2016) and, in particular, to the recent relational contracting literature in which the long-run performance of the firms may be lower than their earlier performance (Barron and Powell, 2017; Fong and Li, 2017; and Li and Matouschek, 2013). In these papers, organizational performance declines because privately observed negative shocks in the past constrain the organization's ability to make promises to its employees and therefore to motivate its workforce—the organization is burned by its past promises. In contrast, there are no privately observed shocks in this model. The decline in the performance is a by-product of information revelation, which is used to discourage the agent from learning to shirk at the beginning of the relationship.

6. Conclusion

This article explores the optimal provision of relational incentives when the worker may learn by shirking. 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 to the job and use this information in the future to shirk more effectively. We argue that in such a setting, relational incentives can be sharpened by adopting an organizational policy that dilutes the agent's information value from shirking. We highlight two such policies—job reorganization and strategic disclosure of information. We show that both could be used as a strategic tool to strengthen incentives and illustrate how the frequency of reorganization (and, in the same vein, level of transparency) is tied to the amount of surplus generated in the relationship.

While we only highlight two possible policies aimed at dissuading learning-by-shirking, our analysis could be applied to a broader range of settings as 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 certain number of periods. Adoption of new technology could be costly and such a cost plays the same role as that of the cost of reorganization (ψ) or the loss of surplus ($s_2 - s_1$) 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" is one of the major causes of running down 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 learning-by-shirking if they expect the current measures to become irrelevant in the near future.

Appendix

***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 (i.e., an equilibrium where $e_t = 2$ and $\gamma_t = 0$ in all periods) exists. Recall that in an efficient equilibrium (and only in an efficient equilibrium), $\pi + u = y - c_2$. 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 $\pi + u = y - c_2$. Then, there exist w, b, u^N, and π^N such that (PK_A) and (DE_P) are satisfied, i.e.,

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

and

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

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

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

(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

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

where the last inequality follows from (3) and the definition of D. As the inequality above holds for all (u,π), we therefore have $D \ge (1-\delta)(c_2 - c_1) - (1-p)\,\delta\,(y - c_2) + p\delta D$, or,

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

Step 2. (*Necessity of NSC^**) Let $(u,\pi) \in \mathcal{E}$ such that $\pi + u = y - c_2$. Combining (IC_1) and (PK_A) one obtains:

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

Using (DE_P) and the fact that $u^N + \pi^N = y - c_2$, we obtain:

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

Hence, we must have

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

(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, $u^N=y-c_2$ and $b=0$. Plugging these values in (IC_0), we get

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

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, $U\left(u^N,\pi^N\right)=u^N$ and, since $u^N=y-c_2$ 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 $U\left(u^N,\pi^N\right)=(1-\delta)\,(y-c_1)+p\delta U\left(u^N,\pi^N\right)$, or,

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

Given this and the fact that $u^N=y-c_2$ 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. ■

Proof of Proposition 1. From Lemma 1 and the observation that the term

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

is increasing in δ for $\delta \in (0,1)$ and $p \in (0,1)$, it follows that there exists a unique δ^* (at which NSC^* is satisfied with equality) such that for all $\delta \geq \delta^*$, in the optimal contract, the agent exerts effort in both tasks in all periods.

The fact that for $\delta < \delta^*$ no effort can be induced and it is optimal for the principal and the agent to take their outside options in every period follows directly from forthcoming results in this paper. Indeed, forthcoming Lemmas 2-4 in Section 4 of this paper are also valid when $\alpha^R = 0$ in every period (i.e., when job reorganization is not possible). Thus, when looking for the optimal contract, it is without loss of generality to restrict attention to contracts where, in each period, no bonuses are used, the principal's continuation payoff is zero and the principal and agent permanently terminate the relationship with some probability. Given this, the optimal contract when job reorganization is infeasible, must solve problem $\mathcal{P}$ that appears in Section 4.2 with $s_1 = 0$ (the continuation value of termination). However, as shown in Step 1 of the proof of Proposition 2, problem $\mathcal{P}$ has no solution when $s_1 = 0$, implying that effort cannot be sustained and it is optimal for the principal and agent to take their outside options in all periods. ■

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 or the job is reorganized. Note that $u_s^t \leq U(u_t)$ and satisfies the following recursive relation:

$$u_s^t = (1-\delta)(s_2 + c) + \delta p\left(\alpha_t u_s^{t+1} + (1-\alpha_t)s_1\right). \tag{6}$$

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

$$u^t \geq (1-\delta)(s_2 + c) + \frac{1}{2}p\delta\left(\alpha_t u_s^{t+1} + (1-\alpha_t)s_1\right), \tag{7}$$

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_N) constraints:

$$\mathcal{P}_R : \max_{\alpha_t \in [0,1]} \; u^1 \;\; s.t. \; (1), \; (6), \text{ and } (IC_1') \text{ hold for all } t.$$

Step 2. (*Rewriting $\mathcal{P}_R$ in terms of α_t*) By using (1) and (6), one can eliminate u^t and u_t^s 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 $S_t = \alpha_t + \delta\alpha_t\alpha_{t+1} + \delta^2\alpha_t\alpha_{t+1}\alpha_{t+2} + \ldots$. Hence,

$$u^1 = s_1 + (1-\delta)(s_2 - s_1)(1 + \delta S_1). \tag{8}$$

Next, note that, $u_s^t - ps_1 = (1-\delta)(s_2 + c - ps_1) + \delta p \alpha_t (u_s^{t+1} - s_1)$, and hence,

$$\begin{aligned} u_s^t - s_1 &= u_s^t - ps_1 - (1-p)s_1 \\ &= (1-\delta)(s_2 + c - ps_1) + \delta p \alpha_t (u_s^{t+1} - s_1) - (1-p)s_1 \\ &= (1-p)((1-\delta)y - \delta s_1) + \delta p \alpha_t \left(u_{t+1}^s - s_1\right). \end{aligned}$$

So,

$$u_s^t - s_1 = (1-p)((1-\delta)y - \delta s_1)(1 + \delta p D_t), \tag{9}$$

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:

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

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 $k_2 = (1-\delta)(s_2 + c - s_1) - \delta(1-p)s_1$. Since we consider the case where $\delta \leq \delta^*$, and hence, (NSC^*) is violated, it routinely follows that $k_2 > 0$. Now, the above constraint can be rewritten in the following form:

$$D_t \leq A + BS_t \quad \forall t, \tag{10}$$

where $A = (k_0 - k_1 - k_2)/k_2\delta p$ and $B = k_0/pk_2$. So, from (8) and (10), it follows that $\mathcal{P}_R$ is equivalent to the following program:

$$\mathcal{P}_R' : \max_{\alpha_t \in [0,1]} S_1 \;\; s.t. \;\; (10).$$

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 (10) 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:

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

(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 (δ_R, δ^*) does not exist).

Step 4. (*Introducing $f(S)$ function and defining S^**) Note the 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

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

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^*) where $S^* < S^{FB}$, and $D^* < D^{FB}$ (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 (10) would be violated in some period) and (ii) the recursion relations (PK_R) for some associated sequence of α_t, $\{\bar{\alpha}_t\}$ (say). We will argue in the next four sub-steps (steps $5a$ to $5d$) that such a sequence cannot exist.

Step 5a. We argue that $\bar{S}_1 > \bar{S}_2$ *and* $\bar{D}_1 > \bar{D}_2$. 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:

$$\frac{\bar{D}_1}{\bar{S}_1} = \frac{1+\delta p \bar{D}_2}{1+\delta \bar{S}_2} > \frac{1+\delta p f\left(\bar{S}_2\right)}{1+\delta \bar{S}_2} = \frac{f\left(\bar{S}_2\right)}{\bar{S}_2} > \frac{f(S^*)}{S^*}, \tag{11}$$

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 (11). 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 $(S_2', D_2') := (1+\varepsilon)\left(\bar{S}_2, \bar{D}_2\right)$.

We argue that for sufficiently small ε, (S_2', D_2') is feasible. Since $\left(\bar{S}_3, \bar{D}_3\right) \in \mathcal{F}$ and (PK_R) is trivially satisfied by definition of (S_2', D_2'), it is enough to show that (S_2', D_2') 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, $(S_2', D_2') = (1+\varepsilon)\left(\bar{S}_2, \bar{D}_2\right)$ always satisfies (IC_R).

Next, observe that,

$$\frac{\bar{D}_1}{\bar{S}_1} = \frac{1+\delta p\bar{D}_2}{1+\delta\bar{S}_2} > \frac{1+\delta p\,(1+\varepsilon)\,\bar{D}_2}{1+\delta\,(1+\varepsilon)\,\bar{S}_2} = \frac{1+\delta pD_2'}{1+\delta S_2'}.$$

Now, we reduce $\bar{\alpha}_1$ to some α_1' where $\alpha_1'\,(1+\delta S_2') = \bar{S}_1$. Let $D_1' = \alpha_1'\,(1+\delta pD_2')$. So, by the above inequality, we find that:

$$\frac{D_1'}{\bar{S}_1} = \frac{1+\delta pD_2'}{1+\delta S_2'} < \frac{\bar{D}_1}{\bar{S}_1}.$$

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. (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 $D^* = f(S^*)$,

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

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 (10).

Step 8. (*If the original problem $\mathcal{P}$ has a solution, then α^* is a solution to $\mathcal{P}$*) We now show that if $\mathcal{P}$ has a solution, the optimal contract $\{\alpha^*\}$ satisfies (IC_1^*), (IC_0^*) and all (SE^*)s, 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 shirking 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. As $\mathcal{P}'_R$ is a "relaxed" version of $\mathcal{P}$ and $\{\alpha^*\}$ is solution to $\mathcal{P}'_R$, then, for all t, the payoff u^* under the contract $\{\alpha^*\}$ must be at least as large as the payoff u^t under a contract that solves $\mathcal{P}$. Now, as any solution to $\mathcal{P}$ must satisfy (IC_0^*), i.e., it must satisfy $u^t \geq (1-\delta)y$ for all t, we must have $u^* \geq (1-\delta)y$. Hence, $\{\alpha^*\}$ also satisfies (IC_0^*).

Step 8c. Finally, to check that (SE^*)s are satisfied, note that: (i) From the definition of $\mathcal{E}$, we know that $(s_1, 0) \in \mathcal{E}$. (ii) 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 if it has a solution. ■

Proof of Proposition 2. The proof is given in two steps. In the first step, we take s_1 as given (with $s_1 < s_2$), and derive a necessary and sufficient condition for $\mathcal{P}$ to have a solution. In the second step, we use this condition to prove the proposition.

Step 1. (*A necessary and sufficient condition* on s_1 *for $\mathcal{P}$ to have a solution*) We show that for a given $\delta < \delta^*$ problem $\mathcal{P}$ has a solution if and only if

$$\frac{\delta}{1-\delta}\left(1 - \frac{1}{2}p\right) s_1 \geq c_2 - c_1. \tag{12}$$

Step 1a. We first show that (12) is necessary for $\mathcal{P}$ to have a solution. If $\mathcal{P}$ has a solution, by Lemma 6 we know it is stationary: $\alpha_t = \alpha^*$ for all t. Moreover, $\alpha^* < 1$ as we are considering the case where $\delta < \delta^*$. Now, at the solution, the following two conditions must hold: (i) (IC_1^*) in period one holds with equality; and (ii) the following inequality holds:

$$u^{t+1} - s_1 < \frac{1}{2}p\left(U\left(u^{t+1}\right) - s_1\right) \tag{13}$$

for $t = 1$. Otherwise, it would be possible to increase α_1 from α^* (keeping $\alpha_t = \alpha^*$ for $t > 1$) and increase u^1 while preserving (IC_1^*) and all other constraints in $\mathcal{P}$, contradicting the fact that α^* is solution. Now, observe that (13) implies that if under the optimal contract (IC_1^*) in period one is satisfied for $\alpha_1 = \alpha^*$ (which must be the case), then it is also satisfied for $\alpha_1 = 0$, i.e.,

$$(1-\delta)\, s_2 + \delta s_1 \geq (1-\delta)\,(s_2 + c) + \frac{1}{2}p\delta s_1,$$

which is equivalent to (12).

Step 1b. To see the sufficiency of (12), observe that if it is satisfied then clearly a contract in which $\alpha_t = 0$ for all t satisfies (IC_1^*). Moreover, such contract also satisfies (IC_0^*). To see this, observe that IC_0^* is given by

$$u \geq (1-\delta)y \Leftrightarrow (1-\delta)\, s_2 + \delta s_1 \geq (1-\delta)y \Leftrightarrow \frac{\delta}{1-\delta} s_1 \geq c_2,$$

which is implied by (12). Thus, at least the contract in which $\alpha_t = 0$ for all t is feasible, meaning that $\mathcal{P}$ has a solution.

Step 2. (*A necessary and a sufficient condition for* $\delta_R < \delta^*$) First observe that $v \in [s_2 - \psi, s_2)$ when $\delta < \delta^*$, since $v = s_2 - \psi$ when the job must be reorganized every period and $v = s_2$ when job reorganization is never used.

Step 2.1. (Necessary condition) Since $s_1 = v - \frac{1-\delta}{\delta}\psi$, the highest possible value s_1 can take is $s_1 = s_2 - \frac{1-\delta}{\delta}\psi$. From this and Step 1 it follows that $\mathcal{P}$ has a solution (i.e., $\delta_R < \delta^*$) only if

$$\frac{\delta^*}{1-\delta^*}\left(s_2 - \frac{1-\delta^*}{\delta^*}\psi\right) \geq \frac{c_2 - c_1}{1-p/2}$$

or, equivalently,

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

Step 2.2 (Sufficient condition) Problem $\mathcal{P}$ has a solution for a given $\delta < \delta^*$ if (12) is satisfied when s_1 is the lowest possible, i.e. when

$$\frac{\delta}{1-\delta}(y - c_2) \geq \frac{c_2 - c_1}{1-p/2} + \frac{\psi}{1-\delta}. \tag{14}$$

Let $u^1(v)$ be the value associated with that solution for any $v \in [s_2 - \psi, s_2)$. Since (i) $u^1(v)$ is continuous, (ii) $u^1(v) \geq v$ when $v = s_2 - \psi$ and (14) is satisfied, and (iii) $u^1(v) < v$ when $v = s_2$ (recall $\delta < \delta^*$), then there exists at least one v such that $u^1(v) = v$. The value associated with the optimal contract $v^* = \max\{v \in [s_2 - \psi, s_2) : u(v) = v\}$.

The rest of the proof follows from the fact that when problem $\mathcal{P}$ has no solution then effort in both tasks cannot be elicited (even in period one) and therefore it is optimal for the principal and agent to take their outside options in every period and from Proposition 1. ■

Proof of Proposition 3. We begin by characterizing $\tilde{\delta}$. In the case of information revelation, $s_1 = py - c_1$. From Step 1 in the proof of Proposition 2, it follows that problem $\mathcal{P}$ has a solution if and only if:

$$\frac{\delta}{1-\delta}\left(1 - \frac{p}{2}\right)(py - c_1) \geq c_2 - c_1. \tag{15}$$

Let $\tilde{\delta}$ be the value of δ for which (15) is binding. Hence, $\tilde{\delta} < \delta^*$ if and only if:

$$\frac{\delta^*}{1-\delta^*}\left(1 - \frac{p}{2}\right)(py - c_1) > \frac{\delta^*}{1-\delta^*}\left(1 - \frac{p}{2 - p\delta^*}\right)(py - c_1) = c_2 - c_1,$$

that simplifies to the condition given in part (iii) (recall that (NSC^*) is binding at δ^*.) The rest of the proof immediately follows from Proposition 1 and Lemmas 6 and 7. ■

References

[1] Abreu, Dilip, David Pearce, and Enno Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.
[2] Aoyagi, M. 2010. "Information Feedback in a Dynamic Tournament." *Games and Economic Behavior*, 70 (2): pp. 242–260.
[3] 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.
[4] 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.
[5] Bardhan, P. 1997. "Corruption and Development: A Review of Issues." *Journal of Economic Literature* 35; pp. 1320–46.
[6] Barron, Dan, and Michael Powell. 2017. "Policies in Relational Contracts." Mimeo, Northwestern University.
[7] Battaglini, M. 2005. "Long-Term Contracting With Markovian Consumers." *American Economic Review* 95 (3); pp. 637–658.
[8] Bergemann, Dirk and Ulrich Hege. 2005. "The Financing of Innovation: Learning and Stopping." *RAND Journal of Economics* 36; pp. 719-752.
[9] Bhaskar, V. 2014. "The Ratchet Effect Re-examined: A Learning Perspective." Mimeo, University of Texas-Austin.
[10] Bolton, Patrick and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67, pp. 349-374.
[11] Bonatti, Alessandro, and Johannes Hörner. 2011. "Collaborating." *American Economic Review* 101; pp. 632–663.
[12] Bonatti, Alessandro, and Johannes Hörner. 2015. "Learning to Disagree in a Game of Experimentation." Mimeo, Yale University.
[13] Chassang, S. 2010. "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts." *American Economic Review* 100 (1); pp. 448–465.
[14] Deb, Joyee, Jin Li, and Arijit Mukherjee. 2016. "Relational Contracts with Subjective Peer Evaluations." *The RAND Journal of Economics* 47; pp. 3–28.
[15] Dessein, Wouter, Andrea Galeotti, and Tano Santos. 2016. "Rational Inattention and Organizational Focus." *American Economic Review* 106; pp. 1522–36.
[16] Dessein, Wouter, and Tano Santos. 2016. "Managerial Style and Attention." Mimeo, Columbia University.
[17] Ederer, Florian. 2010. "Feedback and Motivation in Dynamic Tournaments." *Journal of Economics & Management Strategy* 19 (3); pp. 733–769.
[18] Ederer, Florian, Richard Holden, and Margaret Meyer. 2014. "Gaming and Strategic Opacity in Incentive Provision." Cowles Foundation Discussion Paper No. 1935.
[19] Fong, Yuk-Fai, and Jin Li. 2017. "Information Revelation in Relational Contracts." *Review of Economic Studies* 84; pp. 277–299.
[20] Fuchs,W. 2007. "Contracting with Repeated Moral Hazard and Private Evaluations." *American Economic Review* 97 (4); pp. 1432–48.
[21] Garicano, Luis, and Luis Rayo. 2016. "Why Organizations Fail: Models and Cases." *Journal of Economic Literature* 54 (1); pp. 137-92.
[22] Gawande, Atul. 2010. The Checklist Manifesto: How to get Things Right. New York: Metropolitan Books.
[23] Gibbons, Robert, and Rebecca Henderson. "What Do Managers Do?: Exploring Persistent Performance Differences among Seemingly Similar Enterprises." In The Handbook of Organizational Economics, edited by Robert Gibbons and John Roberts, pp. 680-731. Princeton University Press, 2012.
[24] 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.

[25] Guo, Yingni. 2016. “Dynamic Delegation of Experimentation.” *American Economic Review* 106; pp. 1969–2008.
[26] Halac, Marina, Navin Kartik, and Qingmin Liu. 2016. “Optimal Contracts for Experimentation.” *Review of Economic Studies* 83, 1040–91.
[27] Halac, Marina, and Andrea Prat. 2016. “Managerial Attention and Worker Performance.” *American Economic Review* 106 (10): 3104–32.
[28] Hoffman, Mitchell, and Steven Tadelis. 2017. “How Do Managers Matter? Evidence from Performance Metrics and Employee Surveys in a Firm.” Mimeo, University of Toronto.
[29] Hörner, Johannes, and Larry Samuelson. 2013. “Incentives for Experimenting Agents.” *RAND Journal of Economics* 44; pp. 632–663.
[30] Ishihara, A. 2016. “Relational Contracting and Endogenous Formation of Teamwork.” Mimeo, Kyoto University.
[31] Keller, Godfrey, Sven Rady and Martin Cripps. 2005. “Strategic Experimentation with Exponential Bandits.” *Econometrica* 73, pp. 39–68.
[32] Lazear, Edward. 2006. “Speeding, Terrorism, and Teaching test.” *Quarterly Journal of Economics*; pp. 1029–61.
[33] Lazear, Edward P., Kathryn L. Shaw, and Christopher Stanton. 2015. “The Value of Bosses.” *Journal of Labor Economics* 33; pp 823–861.
[34] Levin, Jonathan. 2003. “Relational Incentive Contracts.” *American Economic Review* 93; pp. 835–57.
[35] Li, Jin, Niko Matouschek and Mike Powell. 2016. “Power Dynamics in Organizations.” forthcoming, *American Economic Journal: Microeconomics*.
[36] 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.
[37] Malcomson, James. 2016. “Relational Incentive Contracts With Persistent Private Information.” *Econometrica* 84 (1); pp. 317–346.
[38] Manso G. 2011. “Motivating Innovation.” *Journal of Finance* 66; pp. 1823–69.
[39] Meyer, Marshall. “Rethinking performance measurement: beyond the balanced scorecard.” Cambridge; New York: Cambridge University Press, 2002.
[40] Moroni, Sofia. 2016. “Experimentation in Organizations.” Working Paper, University of Pittsburgh.
[41] Mukherjee, Arijit. 2010. “The Optimal Disclosure Policy when Firms Offer Implicit Contracts.” *RAND Journal of Economics* 41; pp. 549 – 573.
[42] Mukherjee, Arijit, and Luis Vasconcelos. 2011. “Optimal Job Design in the Presence of Implicit Contracts.” *RAND Journal of Economics* 42; pp. 44–69.
[43] Orlov, Dmitry. 2016. “Optimal Design of Internal Disclosure.” Mimeo, University of Rochester.
[44] Ortega, Jaime. 2001. “Job Rotation as a Learning Mechanism.” *Management Science* 147; pp. 1361–70.
[45] Pastorino, Elena. 2017. “Careers in Firms: Estimating a Model of Learning, Job Assignment, and Human Capital Acquisition.” Mimeo, University of Minnesota.
[46] Rayo, Luis. 2007. “Relational Incentives and Moral Hazard in Teams.” *Review of Economic Studies* 74 (3); pp. 937-963.
[47] Schöttner, Anja. 2008. “Relational Contracts, Multitasking, and Job Design.” *Journal of Law, Economics, and Organization* 24; pp. 138–162.
[48] Yang, Huanxing. 2013. “Non-stationary Relational Contracts with Adverse Selection.” *International Economic Review* 54 (2); pp. 525-547.
[49] Zabojnik, Jan. 2014. “Subjective evaluations with performance feedback.” *The RAND Journal of Economics* 45 (2); pp. 341–369.

Online Appendix

Below we present the proofs of several lemmas that are omitted in the main article as they are primarily technical in nature.

Proof of Lemma 2. Consider a relational contract where, for some period t and history h_t, the critical task for the period is not known to the agent and the payoff profile (u, π) is sustained by effort in both tasks ($e = 2$) and bonus $b \neq 0$ in period t. We construct another contract where, in the same period and for the same history, (u, π) is sustained by $e = 2$ and supported by $b = 0$.

Step 1. (*If* (u, π) *is supported by a contract with* $b < 0$, *then it is supported by a contract with* $b = 0$.) Suppose (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 players' 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 $e = 2$ 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 $a \in \{N, O\}$ and $\pi^R - (1 - \delta)/\delta \geq 0$. 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 since the PK_A^* is preserved, the (IC_1^*) of each one of the periods until the last job reorganization is automatically satisfied. Regarding the (IC_1^*) of the periods from the last reorganization, observe that under the original contract:

$$U(u, \pi) = \max \left\{ \begin{array}{l} (1 - \delta)(w + b - c_2) + \delta \left(\alpha^R u^R + \displaystyle\sum_{a \in \{N,O\}} \alpha^a U(u^a, \pi^a) \right), \\ (1 - \delta)(w + pb - c_1) + p\delta \left(\alpha^R u^R + \displaystyle\sum_{a \in \{N,O\}} \alpha^a U(u^a, \pi^a) \right) \end{array} \right\}, \tag{16}$$

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 (16) holds for any period in which $a = N$, and

$$U(u, \pi) = \alpha^R u^R + \alpha^N U(u^N, \pi^N) + \alpha^O U(u^O, \pi^O)$$

in any period in which $a = O$, then for any τ and $a \in \{N, O\}$, $U(u_\tau, \pi_\tau)$ is non-decreasing in $U(u_\tau^a, \pi_\tau^a)$. Thus, $U'_\tau \leq U_\tau$ for all period $\tau \leq t$ since the last reorganization. Thus, 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. (*If* (u,π) *is supported by a contract with* $b>0$, *then it is supported by a contract with* $b=0$.) 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

$$b^R = b \times \frac{\pi^R - \frac{1-\delta}{\delta}\psi}{\alpha^N\pi^N + \alpha^R\left(\pi^R - \frac{1-\delta}{\delta}\psi\right) + \alpha^O\pi^O},$$

and

$$b^a = b \times \frac{\pi^a}{\alpha^N\pi^N + \alpha^R\left(\pi^R - \frac{1-\delta}{\delta}\psi\right) + \alpha^O\pi^O},$$

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

$$0 \le b^R \le \frac{\delta}{1-\delta}(\pi^R - \frac{1-\delta}{\delta}\psi) \text{ and } 0 \le b^a \le \frac{\delta}{1-\delta}\pi^a \tag{17}$$

for all $a \in \{N,O\}$, where the second inequality in each of these two sets of inequalities 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 $\left(u^N,\pi^N\right)$ and $\left(u^R,\pi^R\right)$ are supported, respectively, by wages w^N and w^R. Now set the new wages

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

for $a = N, R$. The principal's continuation payoffs become

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

for $a = N, R$. Observe that, by (17), $w^{a\prime} \ge w^a$, $\pi^{N\prime} \ge 0$ and $\pi^{R\prime} - \frac{1-\delta}{\delta}\psi \ge 0$, which ensures that when the continuation play calls for $a = N$ or $a = R$ both the agent and the principal 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 (17), $w^{O\prime} \ge w^O$ and $\pi^{O\prime} \ge 0$, 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 $a \in \{N,R,O\}$.

Step 2c. Next, note that this change leaves (PK_P^*) and (PK_A^*) unchanged. Regarding (IC_1^*), under the new contract it is given by

$$u \geq (1-\delta)(w-c_1) + \frac{1}{2}p\delta\left(\alpha^N U\left(u^{N'},\pi^{N'}\right) + \alpha^R u^{R'} + \alpha^O U\left(u^{O'},\pi^{O'}\right)\right). \tag{18}$$

Since under the new contract, in any future periods, only the wage w^a is affected, we obtain that $U(u^{a\prime},\pi^{a\prime}) = U(u^a,\pi^a) + (1-\delta)b^a/\delta$ for all $a \in \{N,O\}$ and $u^{R'} = u^R + (1-\delta)b^R/\delta$. Using this and the fact that $b = \sum \alpha^a b^a$, it is easy to see that (18) is equivalent to the (IC_1^*) in the original 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 (16). The corresponding payoff under the new contract is obtained by substituting, in that expression, b with 0, u^R with $u^{R'}$, and $U(u^a,\pi^a)$ with $U\left(u^{a'},\pi^{a'}\right)$ for all $a \in \{N,O\}$. 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 $U(u_\tau^a,\pi_\tau^a)$ for all $a \in \{N,O\}$, 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 relational contract where, for some period t and history h_t, the critical task for the period is not known to the agent and the payoff profile (u,π) is sustained by effort in both tasks $(e=2)$, wage w and bonus $b=0$ in period t. There is no loss of generality by Lemma 2 in assuming that $b=0$. Let w^a be the next period wage that supports the continuation payoffs (u^a,π^a) for all $a \in \{N,R\}$ 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 $a \in \{N,R\}$, let the new wage in the continuation game be

$$w^{N\prime} = w^N + \frac{\pi^N}{1-\delta} \text{ and } w^{R\prime} = w^R + \frac{1}{1-\delta}\left(\pi^R - \frac{1-\delta}{\delta}\psi\right).$$

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

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

Under these changes, $\pi^{a\prime} = 0$ for all $a \in \{N,O\}$, $\pi^{R'} - (1-\delta)\psi/\delta = 0$, and all the relevant constraints remain satisfied. It is easy to see that (PK_P^*) and (PK_A^*) are preserved. Constraints (DE_P^*) and (DE_A^*) are automatically satisfied since $b=0$. Also, the proposed changes increase the agent's continuation payoff and relax (IC_1^*). More specifically, the (IC_1^*) under the original contract is given by

$$u \geq (1-\delta)(w-c_1) + \frac{1}{2}p\delta\left(\alpha^N U\left(u^N,\pi^N\right) + \alpha^R u^R + \alpha^O U\left(u^O,\pi^O\right)\right).$$

Under the new contract, the left-hand side of the constraint remains the same since PK_A^* is preserved. The right-hand side is obtained by replacing w with w', $U(u^a,\pi^a)$ with $U\left(u^{a'},\pi^{a'}\right) = U(u^a,\pi^a) + \pi^a$ for all $a \in \{N,O\}$, and u^R with $u^{R'} = u^R + (\pi^R - \frac{1-\delta}{\delta}\psi)$. Hence, it is equal to that under the original contract minus

$$\delta(1-\frac{1}{2}p)\left(\alpha^N\pi^N + \alpha^O\pi^O + \alpha^R\left(\pi^R - \frac{1-\delta}{\delta}\psi\right)\right).$$

Finally, under the proposed changes, the (IC_1^*) 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

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

The corresponding payoff under the new contract, U', is obtained by replacing in this expression, w with w', $U(u^a,\pi^a)$ with $U\left(u^{a'},\pi^{a'}\right)$ for all $a \in \{N,O\}$, and u^R with $u^{R'}$. The first element inside the curly brackets remains the same under the new contract. The second element is the same minus

$$\delta(1-p)\left(\alpha^N\pi^N + \alpha^O\pi^O + \alpha^R(\pi^R - \frac{1-\delta}{\delta}\psi)\right),$$

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 $U(u_\tau^a,\pi_\tau^a)$ for $a = N,O$, 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. Suppose there is an optimal contract that generates positive joint surplus. Such contract cannot begin with $a = O$, since a contract beginning with period two of that contract would have a higher associated payoff. Let t be the first period in which $\alpha^O > 0$ and let u be the agent's payoff at the beginning of that period. By Lemmas 2 and 3, we can restrict attention without loss of generality to contracts where, in any period, $b = 0$ and the principal's continuation payoff (net of reorganization costs) is zero. (Note that in such contracts, in any period, $w = y$ if $a = N$ is played and $w = y - \psi/\delta$ if $a = R$ is played.) Hence, if u is sustained by playing $a = N$ in period t, (PK_A^*) implies that

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

and if it is sustained by playing $a = R$, (PK_A^*) implies that

$$u = (1-\delta)(y-\psi/\delta-c_2) + \delta\left(\alpha^N u^N + \alpha^O u^O + \alpha^R u^R\right),$$

where u^a for $a \in \{N, R, O\}$ are the appropriate continuation payoffs. The analysis that follows is valid for either case.

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

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

This implies that (PK_A^*) is preserved and the agent's continuation payoff at the beginning of the period under the alternative strategy, u', satisfies $u' = u$. In addition,

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

(We omit the principal's continuation payoffs π^a for $a = N, O$ as argument of U since they are zero in the contracts considered in this proof.) Since $u' = u$ and $U\left(u^{O\prime}\right) = U\left(u^O\right)$, clearly (IC_1^*) is preserved under the alternative strategy.

Finally, if u is sustained by playing $a = R$ in period t, then for all periods prior to t, the (IC_1^*) constraint must be satisfied since $u' = u$. If instead u is sustained by playing $a = N$ in period t, then following an approach identical to that used in the proof of Lemmas 2 and 3, we obtain again that for all the periods prior to t the (IC_1^*) constraint is also satisfied. 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. Finally, in an optimal contract, permanent exit cannot be played with a positive probability since it is dominated by job reorganization. Thus, in an optimal contact, $\alpha^O = 0$ in all periods. ■

Proof of Lemma 5. Suppose there is an optimal contract that generates positive surplus. Such contract must begin with $a = N$. Let t be the first period in which $\alpha^R > 0$, and let the agent's continuation payoff at the beginning of that period be u. By Lemmas 2-4, we can restrict attention without loss of generality to contracts with no bonuses, in which the principal's continuation payoff (net of reorganization costs) are zero, and where players do not take their outside option. Hence, since u is sustained by playing $a = N$ in period t, (PK_A^*) implies that

$$u = (1-\delta)\left(y - c_2\right) + \delta\left(\alpha^R u^R + \left(1 - \alpha^R\right) u^N\right),$$

where u^R and u^N are the continuation payoffs.

Suppose $u^R < v - (1-\delta)\psi/\delta =: s_1$. Then, we can consider an alternative strategy profile in which u^R is replaced with

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

Under this new new strategy, the agent's continuation payoff at the beginning of period t is

$$u' = (1-\delta)(y-c_2) + \delta\left(\alpha^R s_1 + \left(1-\alpha^R\right)u^N\right) = u + \delta\alpha^R\left(s_1 - u^R\right) > u. \tag{19}$$

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

$$\left(\alpha^R u^R + \left(1-\alpha^R\right)u^N\right) + \frac{1}{2}p\left(\left(1-\alpha^R\right)\left(u^N - U\left(u^N\right)\right)\right) \geq \frac{1-\delta}{\left(1-\frac{1}{2}p\right)\delta}\left(c_2 - c_1\right). \tag{20}$$

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

$$\left(\alpha^R s_1 + \left(1-\alpha^R\right)u^N\right) + \frac{1}{2}p\left(\left(1-\alpha^R\right)\left(u^N - U\left(u^N\right)\right)\right) \geq \frac{1-\delta}{\left(1-\frac{1}{2}p\right)\delta}\left(c_2 - c_1\right). \tag{21}$$

Since (20) is satisfied and $s_1 > u^R$, then (21) must also be satisfied. We next show that the proposed change also relaxes (20) 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 $\Delta = \delta\alpha^R\left(s_1 - u^R\right)$, i.e. Δ is the change in the agent's payoff in period t (see 19). Thus, $u'_t = u_t + \Delta$. Moreover, since period t is the first in which $\alpha^R > 0$, 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,

$$u'_{t-k} - u_{t-k} = \delta^k\Delta. \tag{22}$$

Next observe that

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

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

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

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,

$$U(u'_{t-k}) - U(u_{t-k}) \leq \delta^k\Delta. \tag{23}$$

Next, observe that IC_1^* in any period $t-k-1$ under the original strategy can be written as

$$(1-\delta)(y-c_2) + \delta u_{t-k} \geq (1-\delta)(y-c_1) + \delta p U(u_{t-k+1}) \tag{24}$$

and under the new strategy it can be written as

$$(1-\delta)(y-c_2) + \delta u'_{t-k} \geq (1-\delta)(y-c_1) + \delta p U(u'_{t-k+1}). \tag{25}$$

Since the former is satisfied and by (22) and (23), $u'_{t-k} - u_{t-k} \geq U(u'_{t-k}) - U(u_{t-k})$, 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 $u^R = v - (1-\delta)\psi/\delta$ in the first period in which $\alpha^R > 0$. Applying a similar procedure recursively we obtain that $u^R = v - (1-\delta)\psi/\delta$ the second time $\alpha^R > 0$, and in any other period in which $\alpha^R > 0$. ■

Proof of Lemma 7. **Step 1.** When the critical task is publicly known, we can restrict attention to stationary contracts (Levin, 2003). That is, we can assume that the principal offers the same contract and the agent chooses the same effort level every period. There are three possible actions profiles 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. Recall that by Assumption 1 (iii), it is never optimal for the relationship to have the agent exert effort only on the non-critical task.

Step 2. We begin by deriving the conditions under which effort $e = 2$ in every period can be sustained in (a stationary) equilibrium. Let (w, b) be the wage and bonus in a stationary contract. As transfers between players are frictionless, without loss of generality, 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:

$$w + b - c_2 = 0. \tag{26}$$

The agent's incentive compatibility constraint is:

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

or,

$$b \geq \max\left\{\frac{c_2 - c_1}{1-p}, c_2\right\} = \frac{c_2 - c_1}{1-p}, \tag{27}$$

as $(c_2 - c_1)/(1-p) > c_2$ by Assumption 1 (ii). Now, given (26), 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:

$$\delta(y - c_2) \geq (1-\delta)\, b. \tag{28}$$

Hence, the optimal contract sustaining $e = 2$ must be a solution to the following program:

$$\mathcal{P}_E:\ \max_{w,b}\ \hat{\pi}_t = y - c_2\ \ s.t.\ (27),\ (28)\ \text{and}\ (26).$$

Note that by combining (27) and (28), we get that the necessary and sufficient condition to sustain $e = 2$ is:

$$\frac{\delta}{1-\delta}(y - c_2) \geq \frac{c_2 - c_1}{1-p}. \tag{29}$$

This condition is also sufficient because it allows the implementation of $e = 2$ through the following feasible contract:

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

Thus, $\overline{\delta}$ is the value of δ for which (29) is satisfied with equality.

Step 3. Consider now equilibria in which the agent works on the critical task only. The analysis is identical to the analysis of the case of $e = 2$, but with two exceptions. First, now 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 effort in the critical task only:

$$\frac{\delta}{1-\delta}\,(py - c_1) \geq c_1. \tag{30}$$

This condition is sufficient as it allows for the following feasible contract that implements effort in the critical task only on the equilibrium path: $b = c_1$ and $w = 0$. Thus, $\underline{\delta}$ is the value of δ for which (30) is satisfied with equality. ■