

Learning to game the system

JIN LI

University of Hong Kong

ARIJIT MUKHERJEE

Michigan State University

and

LUIS VASCONCELOS

University of Technology Sydney and Nova School of Business and Economics

First version received May 2018; Editorial decision September 2020; Accepted September 2020 (Eds.)

An agent may privately learn which aspects of his job are more important by shirking on some of them, and use that information to shirk more effectively in the future. In a model of long-term employment relationship, we characterize the optimal relational contract in the presence of such learning-by-shirking and highlight how the performance measurement system can be managed to sharpen incentives. Two related policies are studied: intermittent replacement of existing measures, and adoption of new ones. In spite of the learning-by-shirking effect, the optimal contract is stationary, and may involve stochastic replacement/adoption policies that dilute the agent's information rents from learning how to game the system.

Key words: Performance evaluation systems, Learning, Gaming, Relational contracts.

JEL Codes: C73, D83, M52

1. INTRODUCTION

A common problem in agency relationships is that the 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, and, relatedly, he may not know which corners to cut so as to minimize the risk of getting caught.

But if the agent neglects some aspects of his job and the principal fails to notice, the agent will privately learn which job aspects are relatively less important for ensuring a good performance. This possibility of “learning by shirking” exacerbates the incentive problem because shirking, when successful, informs the agent on how to game the performance evaluation system. The agent acquires valuable private information that he may use later on to cut corners in a fashion that makes shirking harder to detect. Such “perverse learning” poses a challenge for the design and management of performance measurement systems. Numerous scholars have documented how, in a wide range of organizations (including private enterprises, public institutions, and

The editor in charge of this paper was Adam Szeidl.

government agencies), performance metrics eventually lost their ability to differentiate good performances from bad ones as the agents “learned too well how to deliver what is measured rather than what is sought” (Meyer, 2002, p. xii). For example, in the 1980s, the U.S. Department of Labor implemented a set of metrics to evaluate the performance of its job training centres, but the centres eventually figured out how to strategically time their enrollees’ graduation dates so that they appear to perform better on paper relative to their actual performance. The centres’ gaming ploy was not merely an accounting scheme but it also involved diversion of resources from their training activities (Courty and Marschke, 1997, 2004, 2007). Hood (2006) reports a similar behaviour among public hospitals in the U.K. in the late 1990s. The hospitals were asked to meet certain performance targets as measured by a set of key performance indicators. But, over time, they figured out how to meet these targets through “creative compliance” where the attainment of the targets did not correspond to a positive change in the underlying output. In fact, the hospitals often gamed the system by devising new protocols that severely compromised their quality of service. Welch and Byrne (2001) document a related experience during Jack Welch’s leadership at General Electric. When the 360-degree peer evaluation system was rolled out, for the first few years it helped extricate competent managers from incompetent ones. But over time, employees learned how to game the system by manipulating their feedbacks so that everyone may get a good rating. In the same vein, Rivkin *et al.* (2010) find that in the late 2000s when the U.S. Federal Bureau of Investigation rolled out a new evaluation process (termed as “Strategic Performance Sessions”) to assess the organization’s progress, the field office personnel learned to answer questions in a way (regardless of the underlying facts) that the Headquarter would find acceptable. Such gaming through report manipulation may still require the agents to first explore which job aspects can more easily be neglected and subsequently covered up at the reporting stage.¹ One way the organizations respond to this problem is by replacing their current measures by new ones. While no measure may be perfect, different measures have different vulnerabilities, and they need not succumb to a common form of gaming. New metrics can be obtained by simply using a different timeframe for the measurements. For example, in the case of the U.S. job training centres mentioned above, as the centres started to game their performance metric by manipulating their graduation dates, the U.S. Department of Labor kept changing its rules on when an enrollee’s employment status would be recorded (Courty and Marschke, 2007). Similarly, the investors may evaluate a firm’s financial performance on the basis of their annual report instead of the quarterly reports if its managers are likely to manipulate the earnings in the interim periods (Brown and Pinello, 2007). Variations in measures may also stem from variations in granularity.² And, in some scenarios, a variation of measure can be conceived as a replacement of the agent’s supervisor, where each supervisor may be good at detecting shirking in some job aspects but not so good at others.³ However, as it may be costly to replace the metrics (*e.g.* administrative costs associated with rolling out a new evaluation system), an organization must address when and

1. The problem of gaming per se has received considerable attention not only in economics (see *e.g.* Baker, 1992; Oyer, 1998; and for more recent works, Ederer *et al.*, 2018; Jehiel and Newman, 2015; and references therein), but also in several related fields such as accounting (Demski, 1998; Brown and Pinello, 2007; Beyer *et al.*, 2014), finance (Lakonishok *et al.*, 1991; Carhart *et al.*, 2002), and public policy (*e.g.* see Goddard *et al.*, 2000; Bevan and Hood, 2004, 2006; for evidence from the healthcare sector in the U.K.; Dranove *et al.*, 2003, for examples from hospitals in the U.S.; and Jacob and Levitt, 2003, for evidence from U.S. education system). These literatures, however, typically assume that the agents always know how to manipulate the existing measures. In contrast, we focus on settings where the agents may not know how to game the metrics a priori, but attempt to learn this information over time.

2. Meyer (2002) documents how various functional measures replaced gross mortality as metrics of hospital performance, and how quality control measures in automobile manufacturing moved from counting the incidence of defects to defects weighted by severity.

3. Recent economic literature have emphasized how the managerial supervision (or lack thereof) plays a critical role in affecting the workers’ productivity (see *e.g.* Lazear *et al.*, 2015; Hoffman and Tadelis, 2018). Also, for theoretical

how often it should replace the performance measures. We analyse this question and highlight how the performance measurement system may be optimally managed in order to mitigate the learning-by-shirking problem. In particular, we argue that it may be optimal to act preemptively and replace the measures before they run down so as to dissuade the agents from learning how to game the system in the first place.

We explore this issue by modelling the long-term relationship between a firm and a worker as a relational contract where the firm offers incentives through a discretionary bonus payment.⁴ In every period, the worker performs a job that consists of two tasks (or aspects), and the first-best outcome requires the worker to exert effort in both of them. However, the firm cannot measure the worker's performance in each task. The firm observes the worker's overall job output, and it can also rely on a performance measure that is informative about the worker's effort.

We assume there is a host of such additional measures that the firm can choose from. Even though all tasks are equally important for ensuring a high output, any given performance measure is relatively more sensitive to the worker's effort in one of the two tasks. The identity of this "critical task" associated with a given performance measure is unknown to all parties. But if the worker shirks on a task and goes undetected, he privately learns which task is critical and may use this information to shirk more effectively in the future. The firm, however, can replace the existing performance measure by a new one at the end of any period (at a cost); and with such a replacement, the identity of the critical task also changes stochastically as the task identities across measures are statistically independent.

We characterize the optimal relational contract and, relatedly, the firm's optimal policy for replacing the performance measure. The analysis of the optimal contract is intricate due to the fact that when shirking goes undetected, in the continuation game, the players' beliefs about the task identity diverge and cease to be common knowledge. The worker will know which task is critical without the firm knowing that he has such information. Consequently, the worker enjoys an information rent off-equilibrium when he shirks but is not caught. The replacement of the performance measure, therefore, may help the firm precisely because it reduces such rents. When the performance measure is replaced by a new one, any information the agent might have obtained about the old measure becomes obsolete and worthless. But, in spite of the aforementioned complexity, 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 $\delta \in (0, 1)$.

More specifically, for δ sufficiently large, the firm can credibly offer strong enough incentives to induce the worker to work on both tasks, even if the same performance evaluation system is used in all periods. Hence, the first-best surplus is attained, and the performance measure is never replaced. In contrast, for sufficiently low δ , it is optimal to dissolve the relationship as no incentive could be sustained irrespective of how the performance measures are managed.

Our main result concerns with the intermediate value of δ . The optimal contract sharpens relational incentives through a stochastic replacement policy (provided the cost of replacement is not too large). At the end of every period, the firm replaces the existing performance measure by a new one with a constant probability, and the worker exerts effort in both tasks in every period. The possibility of replacement dissuades the worker from shirking by diluting the information rents he hopes to earn by privately learning how to cut corners under the current performance

models of managerial (in)attention, see [Dessein and Santos \(2020\)](#), [Dessein *et al.* \(2016\)](#), [Halac and Prat \(2016\)](#), and [Gibbons and Henderson \(2012\)](#) for a review.

4. 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 *et al.*, 1994](#); [Levin, 2003](#); [Malcomson, 2013](#), for a survey).

evaluation system. As the worker's gains from his superior information may only last for a short period of time, he becomes less inclined to shirk. The optimal replacement probability is driven by the trade-off between the cost of such replacement and the benefits of sharper incentives that it creates.

Strategically cutting corners and noting how measures respond may not be the only way in which a worker learns how to game the system. Over time, he may simply develop a better understanding of how his actions relate to his current performance measures, and subsequently exploit the measures' vulnerabilities. We study this possibility by considering a setup where, in any given period, a worker may identify the critical task for the current measure with an exogenously fixed probability, regardless of whether he works or shirks.⁵ Even in the presence of such "exogenous learning," the firm can sharpen incentives by replacing the performance measures, but the optimal replacement policy may qualitatively differ from its counterpart in our main model. In particular, if exogenous learning is the predominant source of information for the worker, then the optimal policy is deterministic: the metric is kept in place for some periods and, afterwards, it is replaced with certainty. Consequently, if the worker happens to learn which task is critical, he may game the measure until it gets replaced.

Our findings resonate with several well-documented cases of how performance measures evolve in organizations. For example, starting from 1970s General Electric Co. intermittently overhauled its performance evaluation system, not because the existing systems were being gamed, but because the company wanted to emphasize different aspects of its organizational goals—*e.g.*, profitability, worker empowerment, quality control—at different points of time to ensure that it could maintain its leadership position in the industry (Meyer, 2002). Such a policy parallels our characterization of the optimal contract where intermittent replacement is used not because the agent is suspected to exploit the current system but to ensure that he does not attempt to cut corners to the principal's detriment.

In the case of the U.S. job training centres mentioned earlier, the U.S. Department of Labor replaced the measures only after there was enough evidence that the centres had figured out a way to game them. The centres, in large parts, gamed the measures by manipulating their trainees' graduation dates (albeit they were also distorting their efforts in training activities). It is conceivable that such tactics could be learned primarily by exploring how the measures are computed and that exogenous learning played a salient role in this setting. The U.S. Department of Labor's response, therefore, aligns well with our finding that under exogenous learning, it can be optimal to keep the same measure in place until it becomes sufficiently likely that the worker has learned to game it.

Even though intermittent replacement of the existing performance measures can be an effective response to the learning-by-shirking problem, such a response may not always be feasible. In some situations, the worker's job output may be the only meaningful measure of his overall performance that the firm can avail, and some aspects of the job may indeed be more essential than others for ensuring a high output. So, learning about the performance measure effectively means learning about the production process. And since the underlying production technology does not change over time, the information on the critical task is time-persistent.

We show that in such a setting the revelation of task information has implications similar to those of a replacement of the measure (as in our main model): the firm can dissuade the worker from shirking by publicly revealing which task is critical for production, and, going forward, requiring him to perform the critical task only. In reality, the firms may do so by adopting new

5. For the sake of tractability, in this analysis we limit attention to a class of stationary contracts.

performance measures that communicate specific goals and guide the workers accordingly (see, [Gibbons and Kaplan, 2015](#), and references therein).

As before, our key finding pertains to the case of moderate δ where the optimal contract requires the firm to actively filter the information on the task identities. When δ is relatively large (but still within the moderate range), it is optimal not to reveal any information, and the worker works on both task in all periods. In contrast, for δ relatively small, the firm reveals the critical task at the beginning of the game, and the worker performs the critical task only. But, for an intermediate range of δ , at the end of each period, the firm reveals the critical task with a constant probability (if it has not been made public yet). The worker exerts effort on both tasks until the critical task is revealed, but, afterwards, he works on the critical task only. Thus, the optimal revelation policy trades off the current incentive gains with the loss of future surplus.

A key implication of such a stochastic revelation policy is that the performance of ex ante identical firms may differ over time, as the information may be revealed (and performance decline) sooner in some firms than in others. Also, to an outside observer, the firm may appear to be failing in the long run as (almost surely) its performance would decline with time. There is a large 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.

1.1. *Related literature*

Following the seminal works by [Eccles \(1991\)](#) and [Kaplan and Norton \(1992\)](#), a vast literature on the design of performance evaluation systems has developed over the last few decades. This literature primarily explores how the managers may combine information on several financial and non-financial measures, as no single performance metric may adequately reflect the organization's performance (see [Demski, 2008](#), for a review). Several authors have also studied how such collection of measures may be used in formulating incentive contracts ([Ittner et al., 1997, 2003](#)), but there is little discussion on how the performance evaluation systems should be managed over time as the agents might eventually learn how to game the system (one exception is [Meyer, 2002](#), as we have already mentioned in the introduction). The current article attempts to fill this gap.

Our article is related to a few strands of the literature in organizational economics. Several authors have studied how different auxiliary instruments can be used to sharpen relational incentives. These studies have focused on formal contracts ([Baker et al., 1994](#)), integration decisions ([Baker et al., 2002](#)), ownership design ([Rayo, 2007](#)), job design ([Schöttner, 2008](#); [Mukherjee and Vasconcelos, 2011](#); [Ishihara, 2017](#)), design of peer evaluation ([Deb et al., 2016](#)), and delegation decisions ([Li et al., 2017](#)). But, as mentioned before, the issue of design and management of performance evaluation systems has not received much attention.

There is a growing literature on strategic information disclosure in employment relationships, and it 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](#); [Fong and Li, 2017](#); [Orlov, 2018](#)) 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](#); [Ederer et al., 2018](#)). In this literature, our analysis is closest to [Lazear \(2006\)](#), who analyses when it is optimal to reveal to the agent which aspects of his performance are being measured. While [Lazear \(2006\)](#) considers monitoring and information disclosure in a static setting, we explore the role of transparency in incentive provision in a dynamic context and focus on the optimal disclosure of information over time.

Our article also relates to the literature on incentives for experimentation (see Bolton and Harris, 1999; Keller *et al.*, 2005; Manso, 2011; Hörner and Samuelson, 2013; Guo, 2016; Halac *et al.*, 2016; Moroni, 2016; Bonatti and Hörner, 2017). While most of these articles do not consider relational incentives, a recent exception is Chassang (2010). He shows that moral hazard in experimentation, combined with the lack of commitment by the principal, 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 how to design the relationship so as to dissuade the agent from experimentation (*i.e.* selectively perform only a subset of tasks to learn how to game the performance system).

From a conceptual and methodological perspective, this article belongs to the literature on dynamic relationships where the agent can learn about the environment by deviating, and, thus, obtains information rents. While this feature is common in real-life situations, the literature is small because of the technical difficulties. Once the agent deviates, his belief of the future differs from the rest of the players.⁶ Unless the information structure has certain special features (see *e.g.* Bergemann and Hege, 2005; Bonatti and Hörner, 2011), the analysis of this type of problem requires, as in our article, an estimation of the agent’s informational rent, and only a few recent papers have been able to make progress on it (Bhaskar, 2014; Prat and Jovanovic, 2014; Sannikov, 2014; DeMarzo and Sannikov, 2017; Cisternas, 2018; Bhaskar and Mailath, 2019). In contrast to these papers, the information rent in our setting can be directly controlled by replacing the performance measures, and we explore the optimal replacement policy.

2. MODEL

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

Stage game: We describe the stage game in terms of its four key components: *technology*, *performance measures*, *contracts*, and *payoffs*.

TECHNOLOGY: In any period t , the agent may perform a job that consists of two tasks: \mathbb{A} and \mathbb{B} . The agent privately exerts an effort $e_t \in \{0, 1_{\mathbb{A}}, 1_{\mathbb{B}}, 2\}$ at a cost of $C(e_t)$ in order to complete the job. 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 $Y_t \in \{-z, y\}$ is assumed to be observable but not verifiable. The job is successfully completed if the agent works on both tasks (*i.e.* $e_t = 2$), and yields an output $y > 0$. If the agent shirks on both tasks (*i.e.* $e_t = 0$) he fails at the job, leading to a negative output $-z$ (*e.g.* such a failure may lead to an erosion of the firm’s market value). But if the agent performs exactly one of the two tasks (*i.e.* $e_t = 1_{\mathbb{A}}$ or $1_{\mathbb{B}}$), the output is y with probability μ (> 0) and $-z$ with probability $1 - \mu$.

PERFORMANCE MEASURES: In addition to the output, the principal also relies on a performance measure that yields further information on the agent’s effort level. There are infinitely many

6. The lack of common knowledge can also arise in dynamic models in which the principal has private information (Fuchs, 2007; Fong and Li, 2017) and in which the agent has persistent private information (Battaglini, 2005; Malcomson, 2013, 2016; Yang, 2013).

performance measures $\{M^1, M^2, \dots\}$ that the principal can choose from. But all measures are inherently noisy, and no measure is equally sensitive to the agent's effort in all aspects of his job.

In particular, for any $i \in \{1, 2, \dots\}$, $M^i \in \{0, 1\}$, where $M^i = 1$ if the agent works on both tasks, and $M^i = 0$ if he shirks on both. But if he works on exactly one of the two tasks, the realization of M^i depends on which task M^i is more sensitive to—the “critical” task associated with the measure. If the agent only performs the critical task associated with M^i , $M^i = 1$ with probability θ (> 0) and 0 otherwise. But if the agent only performs the non-critical task, then $M^i = 0$ with certainty.⁷

The identity of the critical task is an idiosyncratic feature of a measure, and remains unchanged over time. When a measure is first put in place, neither player knows which task is critical for that measure, and both players correctly believe that any of the two tasks could be critical with equal likelihood. In other words, the identity of the critical task is i.i.d. across different performance measures. Similar to the output, we assume that all performance measures are observable but not verifiable.

Let M_t be the measure that is used in period t , chosen from the set $\{M^1, M^2, \dots\}$. At the end of any period, the principal can publicly replace the current performance measure by a new one after incurring a cost ψ (and, consequently, it randomly changes the identity of the critical task in the subsequent periods). We denote the principal's replacement decision as $\gamma_t \in \{0, 1\}$, where $\gamma_t = 1$ if the principal replaces the current measure at the end of period t , and $\gamma_t = 0$ otherwise.

Note that if the agent shirks by exerting effort in only one task, he may privately learn the identity of the critical task associated with the current performance measure. If he picks the critical task by chance, both the output and the measure turn out to be good (*i.e.* $Y_t = y$ and $M_t = 1$) with probability $p := \mu\theta$, and the principal would fail to detect the agent's shirking. 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 replacement of performance measures does not affect the agent's productivity. Hence, such a replacement is completely wasteful but for its incentive implications, on which we will elaborate below. Finally, we assume that at the beginning of the game, the principal already has a performance measure (M_1) in place.⁸

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 $b_t = b_t(Y_t, M_t)$. The incentives are relational as the output and the performance measures are assumed to be non-verifiable.

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.

7. One can also interpret a measure M^i as a performance evaluation system that combines multiple metrics. A plethora of such systems can be formulated using different combinations of various performance measures. However, all such systems tend to provide only an imprecise evaluation of the agent's overall job performance. The available metrics may be inherently noisy. Also, a system that closely tracks all aspects of the agent's performance may be hard to operationalize due to the difficulties in aggregating a large number of metrics into an overall performance assessment that can be tied to compensation.

8. This assumption streamlines the analysis and allows us to abstract away from the question of whether to adopt a performance measure at the first place, as we focus on the question of how to manage the existing performance measurement system in the face of the learning-by-shirking problem. Nevertheless, we explore the former issue of the adoption of new measures in a related environment in Section 5.1.




FIGURE 1

Timeline of the stage game.

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.

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 expected payoffs for the agent and the principal are given as

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

respectively.

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

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

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

EQUILIBRIUM CONCEPT: We use perfect Bayesian Equilibrium (PBE) in pure strategies as a solution concept. This is in contrast with the extant literature that defines a relational contract as a public Perfect Equilibrium of the game (Levin, 2003). We focus on PBE because, in our setting, a restriction to public strategies may lead to some loss of generality (as the agent can learn about the task identities privately, and use this information to shirk in the future). We define an “optimal” relational contract as a PBE of this game where the payoffs are not Pareto-dominated by any other PBE. The formal definitions of the players' strategies and the equilibrium concept are given in the [Supplementary Appendix](#).

In what follows, we maintain a few restrictions on the parameters to streamline the analysis.

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

Assumption 1 (i) states that when the agent exerts effort on both tasks, the resulting surplus is strictly larger than the contracting parties' outside option. Assumption 1 (ii) stipulates that the cost of exerting effort on both tasks is sufficiently large relative to the cost of working on only one of them. 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 $e_t = 2$) are also sufficient to deter him from shirking on both (*i.e.*, choosing $e_t = 0$). Finally, Assumption 1 (iii) ensures that it is never optimal to ask the agent to cut corners—it is always better to dissolve the relationship than to have the agent perform only one of the two tasks in any given period. This condition is trivially satisfied when z is sufficiently large. Assumptions 1 (i) and (iii) imply that production efficiency calls for the agent to exert effort on both tasks. Assumptions 1 (ii) and (iii) simplify the analytical tractability of the optimal contracting problem.

3. THE OPTIMAL CONTRACTING PROBLEM

We begin our analysis by formulating the principal's optimal contracting problem. First, we present a set of constraints that a contract must satisfy if it were to implement effort on both tasks in a given period. Next, we argue that without loss of generality, we can restrict attention to a simpler class of contracts, and frame the optimal contracting problem accordingly.

Let \mathcal{E} be the set of all PBE payoffs in the repeated game starting from any period t such that the critical task associated with the current measure M_t is not known to the agent. Consider a payoff pair $(u, \pi) \in \mathcal{E}$ that is supported by effort $e = 2$ along with wage w and bonus b in the current period.

Notice that the equilibrium strategies may call for one of the following three action profiles in the next period: (1) the agent exerts effort on both tasks while his performance is evaluated using the same measure that has been used in the previous period; (2) the agent exerts effort on both tasks but faces a new performance metric (*i.e.* the principal has replaced the metric at the end of the previous period); and (3) both players take their outside options in that period. We denote these three actions as $a = N$ ("no replacement"), R ("replacement"), and O ("outside option"), respectively.⁹

The players could also randomize over these three action profiles (using the public randomization device). Suppose that under the equilibrium strategy (that supports (u, π)), the action profile $a \in \{N, R, O\}$ is taken in the following period with probability α^a , and the corresponding continuation payoffs are (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.

Now the payoff pair (u, π) , by virtue of being equilibrium payoffs with current actions ($e = 2, w, b$) and continuation payoffs (u^a, π^a) , $a \in \{N, R, O\}$, must satisfy the following constraints.

(i) *Promise-keeping*: The players' payoffs must be equal to the weighted sum of their current and continuation payoffs:

$$u = (1 - \delta)(w + b - c_2) + \delta \sum_{a \in \{N, R, O\}} \alpha^a u^a, \quad (PK_A)$$

$$\pi = (1 - \delta) \left(y - \left(w + b + \alpha^R \psi \right) \right) + \delta \sum_{a \in \{N, R, O\}} \alpha^a \pi^a. \quad (PK_P)$$

9. By Assumption 1 (iii), it is never optimal for the relationship to have the agent perform only one of the two tasks.

Note that if the principal replaces the performance metric at the end of the current period, the associated cost ψ is realized instantaneously.

(ii) *Incentive compatibility*: The agent should not gain from deviating and shirking altogether or by performing exactly one of the two tasks. Since the agent would surely get caught if he shirks on both tasks, we must have:

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

But if the agent shirks on exactly one of the two tasks, his deviation may go undetected. The incentive constraint that deters such a deviation is more involved. The constraint must account for the fact that upon deviating, the agent may privately learn the identity of the critical task associated with the current performance measure, 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. To address this issue, we proceed as follows.

For any $(u', \pi') \in \mathcal{E}$, let $U(u', \pi')$ be the maximal payoff the agent could earn in the continuation game if he were privately informed about the critical task in the current period and chose his actions accordingly. More specifically, suppose that the payoffs (u', π') are obtained when the agent and the principal play the equilibrium strategy profile (σ'_A, σ'_P) . Let $BR_A(\sigma'_P)$ be the agent's best-response to the principal's strategy σ'_P if he were perfectly informed about the identity of the critical task associated with the current metric. We denote $U(u', \pi')$ as the agent's payoff in the continuation game under the strategy profile $(BR_A(\sigma'_P), \sigma'_P)$.¹⁰

When the agent shirks by working on only one of the two tasks, with probability $\frac{1}{2}p$ he picks the one that is critical for the current performance metric and produces the on-equilibrium path outcome of $(Y, M) = (y, 1)$. As the principal fails to detect such a deviation, the game continues. In the continuation game, as long as the same metric is in place, the agent is privately informed about the task identities and adapts his best-response to the principal's strategy so as to shirk more effectively in the future. Thus, the agent's incentive compatibility constraint can be stated as:

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

The key distinction between this constraint and its counterpart in the standard moral hazard model is that for $a \in \{N, O\}$, the continuation payoff following shirking is $U(u^a, \pi^a)$ instead of u^a . The difference between the two, $U(u^a, \pi^a) - u^a$, reflects the agent's information rents from privately learning which task is critical. For any $(u, \pi) \in \mathcal{E}$, $U(u, \pi) - u \geq 0$, since the agent can always disregard his superior information. Such rents from learning-by-shirking aggravate the moral hazard problem.¹¹ However, the agent does not get any information rents if the current metric

10. In principle, the payoff pair (u', π') could be supported by multiple equilibria, giving rise to distinct values of $U(u', \pi')$. In such case, we can select the equilibrium with the lowest value of $U(u', \pi')$ since it is the one for which the agent's incentive constraint is easiest to satisfy. As we are formulating the conditions under which a given payoff pair can be supported in an equilibrium, using this selection rule is without loss of generality.

11. 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.

is replaced at the end of the period. Since the agent's information on the current period's critical task becomes obsolete in the continuation game, his continuation payoff remains unchanged at u^R .

(iii) *Dynamic enforceability*: Neither player should renege on the bonus payment and the principal should not renege on his promise to replace the metric:

$$(1-\delta)b + \sum_{a \in \{N, R, O\}} \alpha^a u^a \geq 0, \quad (DE_A)$$

$$-(1-\delta)(b + \alpha^R \psi) + \delta \sum_{a \in \{N, R, O\}} \alpha^a \pi^a \geq 0, \quad (DE_P)$$

and

$$-(1-\delta)\psi + \delta \pi^R \geq 0. \quad (DE_{P-R})$$

(iv) *Self-enforcing contracts*: The continuation payoffs themselves must be equilibrium payoffs in the continuation game:

$$(u^a, \pi^a) \in \mathcal{E}, \text{ for } a \in \{N, R, O\}. \quad (SE_a)$$

(v) *Participation*: Both the agent's and the principal's payoff must be at least as large as their respective outside options:

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

(Also note that if we consider a period that immediately follows the replacement of the performance metric then (DE_{P-R}) implies that, on the equilibrium path, $\pi \geq \frac{1-\delta}{\delta} \psi$.)

While the optimal contract needs to abide by the aforementioned constraints, the formulation of the problem can be considerably simplified. Without loss of generality, we can restrict attention to a class of contracts where bonus is never used (*i.e.* $b=0$). Moreover, in any contract that yields a strictly positive joint surplus, the outside option is never taken on the equilibrium path (*i.e.* $\alpha^O=0$), and in any period $t > 1$, the wage

$$w = \begin{cases} y & \text{if the current metric is same as last period's} \\ y - \frac{\psi}{\delta} & \text{if the metric has been replaced in the last period} \end{cases}.$$

That is, in the continuation game following every history, the principal's payoff is zero while the agent receives all of the surplus (net of the cost of replacing the performance measure in the previous period, if any). These claims are formally stated and proved in the [Supplementary Appendix](#) (see Lemmas 1–5).¹² In this class of contracts, the promise-keeping constraint of the principal (PK_P) as well as all dynamic enforceability constraints ((DE_A) , (DE_P) , and (DE_{P-R}))

12. As these results are technical in nature and similar observations have been made in a related class of models (*e.g.* see [Fuchs, 2007](#)), we omit the rigorous treatment of these results in the main text. It is worth noting that we focus on this class of contracts only for analytical convenience, though other forms of implementation may be feasible.

are trivially satisfied. Thus, the optimal contracting problem is tantamount to solving the following program:

$$\max_{\alpha^R, u^N, u^R} u \text{ s.t. } (PK_A), (IC_0), (IC_1), (SE_N), \text{ and } (SE_R),$$

where $\alpha^N = 1 - \alpha^R$, and $\pi^N = \pi^R - \frac{1-\delta}{\delta} \psi = 0$.

4. THE OPTIMAL CONTRACT

The optimal contracting problem suggests that the agent's effort incentives can be potentially sharpened through intermittent replacement of the existing performance metrics. How and when should the principal replace a metric? Replacing the metric every period would certainly dissuade the agent from learning to game the system, but it could be too costly to do so. On the other hand, by leaving the same metric in place for too long, the principal may induce the agent to shirk as it raises the agent's information rent from learning how to game the current metric. The optimal contract is shaped by this trade-off.

The analysis of this problem presents a technical challenge: the agent's maximal payoff in the continuation game when he shirks and learns the identity of the critical task, $U(u, \pi)$, cannot be directly computed as the profitability of the agent's future shirking decisions depends on how and when the principal intends to replace a metric. Consequently, we also cannot limit attention to a class of stationary contracts a priori (as in Levin, 2003). Nevertheless, the following proposition shows that the optimal contract remains stationary and has a set of simple and intuitive characteristics (the proof is given in the Appendix).

Proposition 1 (*Optimal replacement policy*) *Under the optimal contract there exist two cutoffs, δ_R and δ^* , $\delta_R \leq \delta^*$, such that no effort can be induced if $\delta < \delta_R$, and for $\delta \geq \delta_R$ the following holds:*

(i) *If $\delta \geq \delta^*$, the principal never replaces the performance metric, and in every period the agent exerts effort on both tasks.*

(ii) *If $\delta \in [\delta_R, \delta^*)$, the principal replaces the existing performance metric at the end of each period with a constant probability α^* (that may vary with δ), and the agent exerts effort on both tasks in all periods. Moreover, $\delta_R < \delta^*$ if the cost of replacement ψ is below a threshold.*

For a large δ (i.e. if $\delta \geq \delta^*$), the first-best outcome is feasible: the principal can credibly promise a sufficiently large continuation payoff so that in every period, the agent exerts effort on both tasks even if the performance metric is never replaced. In contrast, for δ sufficiently small (i.e. δ below δ_R), the optimal policy dissolves the relationship. Regardless of the replacement policy used, the strongest relational incentives (in terms of continuation payoffs) that the principal can credibly offer fail to induce effort on both tasks.

But for a moderate δ —if $\delta \in [\delta_R, \delta^*)$ —the principal induces effort on both tasks by adopting a stochastic replacement policy where at the end of each period, the principal replaces the existing performance metric with a fixed probability. As the agent anticipates that his information on the critical task may become irrelevant in the near future, he never shirks as he becomes less inclined to learn how to game the system. But since it is costly to replace an existing performance measure, any replacement policy entails a loss of surplus. Therefore, the first-best surplus cannot be attained; moreover, such a policy is optimal if and only if the cost of replacing a metric is not too large.

In the optimal contract, the replacement probability (α^*) is invariant over time, and it is instructive to elaborate on the intuition for this finding. First, consider a relaxed problem where we assume a specific form of deviation: if the agent shirks and learns the critical task, he continues to shirk on the non-critical task in all future periods as long as the same metric is in place. Notice that the exact time of deviation is still a choice variable for the agent.

It turns out that if the replacement policy were to deter shirking in period 1 only, it would take a form that features “early replacement”: there is some T such that the principal would replace the current measure with a positive probability if $t < T$, but would never do so again afterwards. Such a policy backloads the agent’s rewards. Since in the continuation game following any history, the agent receives all surplus net of the cost of replacement of the metric, the agent is guaranteed to have a high payoff in all future periods after period T .

To see why backloading rewards is useful, note that compared to an agent who always works on both tasks, an agent who shirks successfully and then only works on the critical task is effectively less patient—the former discounts the future at rate δ , but the effective discount rate for the latter is $p\delta$ as he faces a risk that, in any period, the relationship can terminate with probability $1 - p$. Since an agent who shirks successfully discounts the future more (relative to an agent who does not shirk), an early replacement of the existing measure most effectively discourages the agent from shirking in period 1 by backloading the rewards as much as possible.

However, such an early replacement policy is necessarily time-inconsistent. While the agent is deterred from shirking in period 1, he may want to deviate in the later periods when the gains from shirking are larger. (As the principal is less likely to replace the measure 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, the optimal policy would ideally implement an increasing sequence of the agent’s continuation payoff by increasing the current period’s replacement probability and decreasing the probability of future replacements. But as this is the case in all periods, the resulting optimal policy becomes stationary and features a time-invariant replacement probability.

So far, we have argued that a stationary replacement policy is optimal in a relaxed problem where we assume a particular form of deviation: if the agent shirks and learns about the critical task, he will continue to shirk on the non-critical task as long as the same metric is in place. But, in general, there are other forms of deviation that may be more profitable for the agent; e.g., even if the agent knows which task is critical, he may still put in effort in some periods before shirking again. For the aforementioned policy to be optimal in the general contracting problem, it must also deter all such deviations where the agent shirks in different time patterns. But this is indeed the case because the policy is stationary: for stationary replacement policies, both the benefit and the cost of shirking are time-invariant. Thus, if an informed agent finds it profitable to shirk for one period, it is also profitable for him to shirk in every period in the future until the measure is replaced. And similarly, if he does not gain by continuing to shirk until the measure is replaced, he also cannot gain from any other types of deviation.

We conclude this section with a remark on a technical aspect of our analysis. In our setup, the characterization of the optimal contract is complicated by the fact that for an arbitrary replacement policy, we cannot readily apply the standard recursive approach à la [Abreu *et al.* \(1990\)](#). Such an approach is generally applicable in settings where in order to check that the proposed policy is robust to all deviations it is sufficient to check for the one-stage deviations. But in our setting the agent’s most profitable deviation plan (under a given policy) may call for shirking in multiple periods according to some specific time pattern. To deal with this issue, we first solve a relaxed problem that restricts attention to a particular form of deviation: once the agent shirks, he continues to shirk until the measure is replaced. The optimal replacement rule for the relaxed problem

is stationary, which allows us to show that it is robust to all other forms of deviation, and, consequently, we establish the optimality of the stationary contract for the general problem. A similar approach has been used in settings where multi-stage deviations are relevant (Williams, 2011; Sannikov, 2014; DeMarzo and Sannikov, 2017; Cisternas, 2018; Bhaskar and Mailath, 2019). In these papers, however, the appropriate relaxed problem is formulated by only considering some “local” deviations (*i.e.* the agent deviates only once). In contrast, we consider a multi-stage deviation in our relaxed problem, *i.e.*, once the agent shirks, he continues to shirk as long as the same measure remains in place.

5. DISCUSSION

In this section, we analyse three different environments that closely parallel our main model and highlight a broader applicability of our key insights.

5.1. *Learning about production and information revelation*

Our analysis above illustrates how strategic replacement of performance metrics can thwart the agents from learning to game the system. But in some environments, the agent’s job output may be the only reliable measure of his performance, and he may cut corners so as to privately learn which tasks are more crucial for production. In such settings, intermittent replacement of performance metrics need not be feasible; the principal must change the underlying production process altogether which could be prohibitively costly.¹³

However, the principal can still provide effort incentives, at least in the early phase of the relationship, by strategically revealing over time which tasks are more critical for production. In fact, it is a common practice for firms to adopt metrics that are narrowly focused on the critical tasks so as to reveal which tasks are more important and guide the worker towards them (Gibbons and Kaplan, 2015). The incentive implications of such a revelation policy are similar to those explored in our main model, and the optimal revelation policy bears close resemblance to the optimal replacement policy studied above.

Our model can be easily modified to reflect such an environment. As before, let $Y = y$ if $e = 2$ and $-z$ if $e = 0$. But now suppose that if the agent shirks on one of the two tasks, then the output depends on the task that is performed. One of the two tasks is “critical” for production: $Y = y$ with probability $\mu > 0$ if the agent only performs the critical task; and $Y = -z$ with certainty if he shirks on it. Efficiency requires the agent to work on both tasks, but working on the critical task only is better than dissolving the relationship and taking the outside options. The identity of the critical task is not known to either player. However, at the end of each period, the principal can publicly access this information and disclose it to the agent at zero cost. Once the information on the critical task is revealed, it remains available in all future periods (in contrast to our earlier model).

In the spirit of (Gibbons and Kaplan, 2015), we can assume that the principal discloses this information by putting in place a performance metric M that reflects effort on the critical task without any noise. Both Y and M are observable but non-verifiable. We keep all other aspects of the our main model unchanged.

The revelation of task information discourages the agent from shirking, as it dissipates the gains from privately learning which task is critical and gaming the system by shirking on the non-critical one. But once the information is revealed, it also becomes more difficult to incentivize

13. See Frankel and Kartik (2019) for a discussion on how, in a signalling environment, the principal may mitigate the agents’ gaming efforts by replacing the underlying signal-generating technology.

the agent to execute all tasks associated with his job. The optimal contract pins down if and when to reveal the task information so as to balance this trade-off.¹⁴

Proposition 2 (*Optimal contract with information revelation*) *Under the optimal contract there exist four cutoffs $\underline{\delta} < \bar{\delta} \leq \hat{\delta} < \tilde{\delta}$ such that no effort can be induced if $\delta < \underline{\delta}$, and for $\delta \geq \underline{\delta}$ the following holds:*

(i) *If $\delta \geq \bar{\delta}$, the agent exerts effort on both tasks in all periods irrespective of the principal's decision on whether to reveal information on the critical task.*

(ii) *If $\delta \in [\hat{\delta}, \bar{\delta})$, the agent exerts effort on both tasks in all periods, but the principal conceals the identity of the critical task.*

(iii) *If $\delta \in [\tilde{\delta}, \hat{\delta})$, the principal reveals the information on the critical task at the end of each period with a constant probability α^* (that may vary with δ). The agent works on both tasks until the task information is revealed, and works only on the critical task afterwards.*

(iv) *Finally, if $\delta \in [\underline{\delta}, \tilde{\delta})$, the principal reveals the identity of the critical task at the beginning of the game and the agent only works on that task.*

A salient implication of the above proposition is that for an intermediate range of δ , *i.e.*, for $\delta \in [\underline{\delta}, \bar{\delta})$, active management of information is critical. Within this range, when δ is relatively large ($\delta \in [\hat{\delta}, \bar{\delta})$), 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 ($\delta \in [\tilde{\delta}, \hat{\delta})$), the principal may do better by not revealing the task information at the beginning of the game. A larger surplus can be attained under a stochastic adoption policy where at the end of each period, the principal reveals the critical task with a fixed probability (by adopting a performance metric for that task). As the critical task is likely to become public information in the near future, the private information that the agent hopes to obtain through learning-by-shirking becomes less valuable. Such a contract elicits effort on both tasks until the critical task is revealed and, hence, is more efficient than the one that reveals this information at the beginning of the game.

In this context, two issues are worth noting. First, 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.* $\hat{\delta} < \bar{\delta}$). When the task information is public, shirking yields a higher payoff to the agent as he knows which task to shirk on. When $\delta \in [\hat{\delta}, \bar{\delta})$, the relational incentives that the principal can credibly promise are strong enough to dissuade the agent from shirking when he does not know the critical task, but too weak to elicit effort on both tasks when the critical task is known to the agent. Consequently, for such values of δ , opacity is (strictly) optimal.

Second, an important implication of the optimality of stochastic information revelation 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 on why organizations fail (Garicano and Rayo, 2016) and, in particular, to the recent relational contracting literature that explains why a firm's performance may deteriorate over time (Li and Matouschek, 2013; Fong and Li, 2017; Barron and Powell, 2019). 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. In other words, the organization is burned by its past promises. In contrast,

14. A formal analysis of the model, along with the proof of Proposition 2, is available in the Supplementary Appendix.

there are no privately observed shocks in our setting. The decline in the performance is a by-product of the information revelation that is necessary to incentivize the agent to exert effort at the beginning of the relationship. This observation is also reminiscent of [Bhaskar and Mailath \(2019\)](#) where the optimal incentive provision may also call for low performance in certain phases of the employment relationship. However, the source of such inefficiency is quite different from the ones we highlight here. Unlike our setting, Bhaskar and Mailath consider a setup where the agent can shirk to manipulate the principal’s belief about the production environment, and earn future information rents. When the employment duration is sufficiently long, the presence of such rents inflates the cost of implementing effort, and it may not be optimal to elicit high effort in all periods.

5.2. *Exogenous learning and shirking on the equilibrium path*

In our baseline model, if the agent were to learn how to game the system, he must engage in strategic shirking. As the optimal contract is designed to elicit effort on all tasks, the agent never shirks on the equilibrium path and never learns which task is critical for his performance evaluation. But as noted in the introduction, the performance measures can lose their effectiveness (and get replaced) over time as the agents eventually learn and exploit their vulnerabilities.

In order to allow for this possibility, we adapt our model and assume that even if the agent never shirks, he might still learn the identity of the critical task from an “exogenous” source. The presence of such exogenous learning can have important implications for the optimal replacement policy. In particular, when the exogenous channel is the key source of learning, it may indeed be optimal to replace the measures only after they run down over time.

Suppose that $Y_t \in \{0, y\}$ and the agent is uninformed about the critical task at the beginning of the relationship. But, at the end of each period t , he learns the identity of the critical task (associated with the measure M_t) with probability $k > 0$. That is, the agent may learn about the task from an exogenous source, reflecting the possibility that even if the agent never shirks, over time, he may gain a better understanding of how his efforts relate to the current performance measures. However, in this setting the consequence of shirking is assumed to be less severe for the principal as the output in the case of a failure is now 0 instead of $-z$.¹⁵ In order to streamline the analysis, we also assume that $C(1_{\mathbb{A}}) = C(1_{\mathbb{B}}) = 0$ and $\theta = 1$. That is, the agent incurs a cost of effort only if he works on both tasks, and a metric indicates success with certainty as long as the agent performs the corresponding critical task.

As the agent learns about the identity of the critical task over time—and an informed agent may shirk even on the equilibrium path—the future surplus of the relationship evolves over time. Consequently, a general characterization of the optimal contract is difficult to obtain, and for the sake of tractability, we restrict attention to a class of stationary contracts: In each period the agent is paid a base wage w and a bonus of b if $(Y_t, M_t) = (y, 1)$, where, without loss of generality, we assume $w + b < y$. The relationship terminates if $(Y_t, M_t) \neq (y, 1)$. All other aspects of the model remain unaltered.

In this setting, a replacement of the performance measure has two key effects. First, as in our baseline model, it dissuades the agent from learning-by-shirking. Second, if the agent has been shirking after acquiring the task information exogenously, a replacement of the metric renders his information useless and deters him from shirking in the future. The optimal replacement policy is shaped by the trade-off between the benefits of replacement stemming from these two effects and the cost of replacing the measure.

In what follows, we elaborate on a few salient characteristics of the optimal replacement policy in this environment (a more detailed analysis is available in the [Supplementary Appendix](#)). To

15. Notice that this modification of the model relaxes Assumption 1 (iii).

this effect, we limit attention to the case where p is sufficiently large, and focus on the class of equilibria with the following feature: an uninformed agent always exerts effort on both tasks, but if he exogenously learns the task identities, he shirks in all future periods as long as the same metric is in place.¹⁶

To see the implications of the exogenous learning, we first consider a relaxed problem by ignoring the incentive-compatibility constraint for effort provision by an uninformed agent. However, the agent could still learn the task identities exogenously, and use that information to shirk in the future. Thus, the optimal contract may again require an intermittent replacement of the performance measures so as to incentivize an informed agent.

Let α_t be the probability that the measure is replaced at the end of period t . The optimal replacement policy is given as:

$$\alpha_t = \begin{cases} 0 & \text{if } \pi_{t+1} > v + \frac{\delta}{1-\delta} \rho_{t+1} (1-p)v, \\ 1 & \text{otherwise} \end{cases},$$

where ρ_t is the probability that the agent is informed about the critical task at the beginning of period t , conditional on the fact that the relationship continues to period t , and the measure remains the same as in period $t-1$; π_t denotes the period t payoff of the principal if the measure is the same as in period $t-1$; and v is the average payoff of the principal in the continuation game, net of the cost of replacement, once a new performance measure is put in.

The optimal policy has two salient features. First, the principal keeps the same measure in place ($\alpha_t = 0$) as long as her expected payoff in the following period (π_{t+1}) is sufficiently large. For p sufficiently large, ρ_t is strictly increasing (and converges to 1), and π_t is strictly decreasing over time. Therefore, the measure eventually gets replaced with certainty after a certain length of time. The policy resonates with the practice observed in reality where the measures lose their effectiveness over time as the agent might learn how to game the metrics, and the principal replaces the measure only after it runs down over time.

Second, it is optimal to replace the measure when π_{t+1} is still larger than v . The wedge between the two (*i.e.* $\delta \rho_{t+1} (1-p)v / (1-\delta)$) arises due to the fact that by replacing the measure, the principal can restart the relationship and guarantee herself an average payoff of v . On the other hand, by continuing for one more period (without replacing the measure), she allows for the possibility that an informed agent may shirk and obtain $(Y_{t+1}, M_{t+1}) \neq (v, 1)$ that triggers termination.

Clearly, if the solution to the relaxed problem satisfies the incentive-compatibility constraint of the uninformed agent, it is also a solution to the original problem. However, this need not be the case in general, and the optimal contract must account for the learning-by-shirking effect. A salient feature of the agent's incentive constraint is that the timing of the replacement of the current measure affects the constraint by influencing both the equilibrium and the deviation payoffs of the agent, and, the timing of replacement, in turn, is also determined by the agent's incentive constraints. The resulting feedback effect presents a technical complexity, and the optimal contracting problem loses analytical tractability. Even though the general form of the optimal stationary contract is difficult to characterize, we show that the optimal contracting problem can be formulated as a linear programming problem and, therefore, amenable to standard numerical solution methods.

16. Recall that p is the probability that $(Y_t, M_t) = (v, 1)$ when the agent performs the critical task only; here $p = \mu$ as we have set $\theta = 1$.

5.3. Noisy performance measure and gradual learning

We have assumed in our baseline model that the agent is sure to succeed if he exerts effort on both tasks: when $e_t = 2$, the output is high ($Y_t = y$) and the performance metric reflects success ($M_t = 1$) with certainty. While this assumption is maintained for analytical tractability, it is conceivable that the agent's performance is always subject to random shocks, and he may fail regardless of his effort level. Our setup can be adapted to capture this possibility, and even in this setting, the incentive effects of a replacement of the metrics that we highlight in our model continue to hold. Below, we present a brief discussion of this case (a detailed analysis is relegated to the Supplementary Appendix).

Suppose that when $e_t = 2$, $Y_t = y$ with certainty but $M_t = 1$ with probability \bar{p} where $p < \bar{p} < 1$; all other aspects of the model are left unaltered. As both a "success" (*i.e.* $(Y_t, M_t) = (y, 1)$) and a "failure" (*i.e.* $(Y_t, M_t) = (y, 0)$) can occur when the agent exerts effort on both tasks, shirking is detected for sure only if $y = -z$.¹⁷ For tractability, we limit attention to stationary contracts where the agent is paid a wage w in every period and a discretionary bonus b is paid when $(Y_t, M_t) = (y, 1)$. Moreover, the principal replaces the existing measure following a success and a failure with probabilities α_s and α_f , respectively. The relationship is terminated if and only if $y = -z$.

Notice that as in our baseline model, when the agent deviates and performs exactly one of the two tasks, he learns the task identities with certainty if he succeeds (*i.e.* if $(Y_t, M_t) = (y, 1)$). However, in contrast to our model, the relationship continues even if the agent fails (*i.e.* $(Y_t, M_t) = (y, 0)$), and, therefore, by shirking and failing, the agent may still gradually learn about the task identities. Consequently, the analysis of the optimal policy must keep track of the agent's belief and his subsequent action choices following a deviation. The agent's incentive compatibility constraint (IC_1) is now reformulated as:

$$u \geq u_d := (1 - \delta) \left(w - c_1 + \frac{1}{2}pb \right) + \delta \left(\frac{1}{2}p((1 - \alpha_s)u_s + \alpha_s u) + \left(\mu - \frac{1}{2}p \right) ((1 - \alpha_f)u_f + \alpha_f u) \right), \quad (IC_1^*)$$

where u_s and u_f are the agent's continuation payoff when his deviation ends up in success and failure, respectively, and the performance metric is not replaced. We can compute u_s and u_f by using the fact that the optimal deviation strategy of the agent is to choose at any point in time the task that is the most likely to be the critical task.

In such an environment, the optimal replacement policy has three salient features. First, as before, the first-best surplus can be attained for δ sufficiently large. That is, there exists a cutoff δ_N^* , such that if $\delta \geq \delta_N^*$, the principal never replaces the performance metric (regardless of whether the agent succeeds or fails), and in every period, the agent exerts effort on both tasks. Second, for $\delta < \delta_N^*$, the optimal contract may call for a replacement of the metric. If the cost of replacement (ψ) is small, the principal replaces the metric with a positive probability at the end of every period and induces the agent to exert effort on both tasks. Both of these observations are reminiscent of Proposition 1, and rely on a similar argument.

Finally, one may presume that it is always better to replace the measure after a failure than after a success (as a failure is more likely to occur if the agent deviates). In the optimal contract, however, the opposite may hold. In general, by raising either α_s or α_f (or both), we can lower the agent's deviation payoff u_d and relax (IC_1^*), but the marginal impacts of α_s and α_f on u_d cannot

17. However, success occurs with probability \bar{p} when the agent works on both tasks but with probability $p := \mu\theta$ when he works on the critical task only.

be ranked *a priori*. In particular, such impacts would depend on several countervailing effects: the likelihood of success and failure following deviation, the associated continuation payoffs u_s and u_f , and how the continuation payoffs vary with the replacement probabilities. For example, while failure is more likely after deviation, the associated continuation payoff u_f is generally smaller than its counterpart following a success, u_s , and, consequently, a replacement after a success may be more effective in reducing the agent's deviation payoff u_d .

6. CONCLUSION

This article explores the optimal provision of relational incentives when the worker may attempt to learn how to game his performance measures. Workers often hold jobs that involve multiple aspects (or a set of tasks) and the performance measures in place may be more sensitive to some job aspects 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 his effort cost but also to learn more about how these job aspects affect his performance evaluations. Using a model of relational contracting, we study how the firm can sharpen incentives in such a setting by managing its performance evaluation system. We highlight two policies—frequent replacement of existing performance measures, and adoption of new measures that guides the workers towards the more critical tasks. We show that both policies could be used as a strategic tool to strengthen relational incentives, and illustrate how the frequency of replacement (and, in the same vein, the adoption of new measures) is tied to the amount of surplus generated in the relationship.

Our analysis also sheds light on how a principal, who may lack information about the model's parameters, can implement the optimal replacement policy in practice.¹⁸ When the principal interacts with several independent agents in similar production settings, the data on the agents' performance over time under different replacement rates can help in pinning down the optimal policy. This problem is somewhat straightforward in our main model as under the optimal rate, the agent never shirks and, hence, always delivers a good performance. Hence, the principal could start out by setting the rate of replacement at some arbitrary level, and can arrive at the optimal rate by trial-and-error. She may raise the rate whenever the agent performs poorly, reduce the rate if the agent delivers good performance consistently over time, and settle for the lowest replacement rate that yields good performance consistently.

But even in more complex environments, such as the ones studied in the two extensions of our main model, the empirical relationship between the agents' performance and the replacement rates can help in inferring the optimal policy. For example, consider the case of exogenous learning as discussed in Section 5.2 and suppose that the optimal policy calls for replacing the measure every T periods (as would be the case where the agent primarily learns through the exogenous source). How can the principal determine how often to replace the measures? Our analysis suggests a tight link between the duration of usage of a measure and the probability of observing a poor outcome while the measure is in place. In particular, if the duration is short, an agent is likely to be dissuaded from shirking to learn, and shirk (and may perform poorly) only if he learns about the tasks through the exogenous source. Thus, the likelihood of observing a poor performance would increase over time as long as the same measure is still in place, as it becomes more likely that the agent learns how to game the metric.¹⁹ Therefore, the principal can experiment by setting

18. It is conceivable that, in complex production environments, the principal's prior belief on some key parameters of the model is considerably diffused, and the optimal contract may be implemented only via experimentation; see *e.g.* Ortner and Chassang (2018) and Chassang and Padró i Miquel (2018) for formal treatments of this issue.

19. Such a pattern would disappear when T is sufficiently large as the agent would shirk (and may learn about the tasks) as soon as the measure is put in place.

a low T , observe how the failure rate varies as T changes moderately, and use the data to estimate some of the key parameters of the model—*e.g.*, the probability of exogenous learning (k)—that determine the optimal duration.

A similar argument can also be made in the case of noisy measures as studied in Section 5.3. When the replacement rate is sufficiently low, the agent would shirk and occasionally fail. But suppose that the rate is increased. If the agent continues to shirk, it would result in a higher rate of failure. Due to the higher replacement rate, the agent loses information on task identities more often, and as he continues to shirk, he chooses the “wrong” task more often. However, if the agent stops shirking, this pattern would disappear completely. As the agent never shirks, there is a drop in the likelihood of poor performance, and the likelihood would not change if the rate of replacement is further increased. Thus, as is the case in our main model, the principal can still ascertain the optimal rate through trial-and-error. If raising the rate leads to a higher frequency of failure, the rate should be raised further; if the frequency stays unchanged, the rate should be lowered; and the rate around which there is a discrete drop in the likelihood of failure is the optimal one.

It is worth noting that even though our model focuses on how the firm may manage the performance measures in response to the learning-by-shirking problem, the incentive effects it highlights relate to any policy that a firm may adopt to “shake up” the production environment in the future in order to dissuade the agent from learning at the present. Indeed, firms often adopt job rotation and/or reorganization policies where workers expect to be moved to different divisions or assignments within the firm after every so few months in a given job. For example, in the classic study of the leveraged buyout of RJR Nabisco by [Burrough and Helyar \(1990, p. 26\)](#), the authors note that “[The CEO, F. Ross Johnson] reorganized Standard Brands twice a year, like clockwork, changing people’s jobs, creating and dissolving divisions, reversing strategic fields. To outsiders it seemed like movement for movement’s sake. Johnson framed it as a personal crusade against specialization. ‘You don’t have a job,’ he told [...], ‘you have an assignment.’” 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 ([Bardhan, 1997](#)). Insofar as such reorganizations are costly to the firm, our model sheds light on how such a policy should be used in the optimal incentive contract.

Acknowledgments. 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 Spiegel, Jeroen Swinkels, Adam Szeidl, Huanxing Yang; the seminar and conference attendees at Boston University, CUNEF (Workshop in Relational Contracts 2016), Indian Statistical Institute (Annual Conference 2016), Michigan State University, Southern Methodist University, University of British Columbia (CETC 2017), Universidad Carlos III de Madrid, University of East Anglia, University of Essex, University of Technology Sydney, Monash University, and University of New South Wales; and the anonymous referees. We thank Peiyu Wei for excellent research assistance. Any errors that may remain are ours.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Appendix

This appendix presents the Proof of Proposition 1. As the proof is relatively elaborate, for expositional clarity, we present it in three parts. First, we present an auxiliary problem where the agent’s continuation payoff following a replacement of the metric (*i.e.* u^R) is treated as a parameter with certain specifications. Next, we present two lemmas that characterize the solution to this problem. Finally, we prove Proposition 1 by using these lemmas and by considering the specific value for u^R that would emerge in the optimal contract.

I. AN AUXILIARY PROGRAM. Recall from Section 3 that without loss of generality, we can restrict attention to the class of contracts where bonus is never used (*i.e.* $b=0$), outside option is never taken ($\alpha^O=0$ when the contract yields strictly positive joint surplus), and in every period $t > 1$, the agent's wage $w=y$ if the current metric is the same as last period's and $w=y-\psi/\delta$ otherwise. Hence, the optimal contracting problem is equivalent to finding the contract in this class that solves:

$$\mathcal{P}_O: \max_{\alpha^R, \alpha^N, u^R} u \text{ s.t. } (PK_A), (IC_0), (IC_1), (SE_N), \text{ and } (SE_R),$$

where $\alpha^N = 1 - \alpha^R$ and $\pi^N = \pi^R - \frac{1-\delta}{\delta}\psi = 0$.

Now, consider a variation of the above problem where we treat the agent's continuation payoff u^R as exogenous and set $u^R = s_1$, where s_1 is a parameter that satisfies two conditions: (1) $s_1 < s_2 := y - c_2$, the surplus generated when there is no replacement of the performance measure and the agent exerts effort in both tasks; and (2) $(u^R, \pi^R) = (s_1, \frac{1-\delta}{\delta}\psi)$ can be sustained as an equilibrium payoff in the continuation game. We can write this auxiliary problem in the following way.

With a slight abuse of notation, let u^t be the agent's (average) payoff at the beginning of period t when (1) the critical task associated with the current performance measure is not known to the agent and (2) the performance measure has not been replaced at the end of the previous period. Furthermore, let α_t be the probability that the principal replaces the performance measure at the end of period t (*i.e.* sets $a=R$).

Now, for a given s_1 , (PK_A) implies that u^t satisfies the following recursive relationship:

$$u^t = (1-\delta)s_2 + \delta(\alpha_t s_1 + (1-\alpha_t)u^{t+1}). \quad (A.1)$$

Therefore, if there exists a contract that implements effort in both tasks at least in the first period, solving for the optimal contract (for a given s_1) is tantamount to finding the sequence $\{\alpha_t\}_{t=1}^{\infty}$ that solves the following program (to simplify notation we denote $c := c_2 - c_1$, and write $U(u)$ instead of $U(u, \pi)$ as the principal's continuation payoff (net of the replacement costs) remains 0):

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

II. LEMMAS. Below we present two lemmas that characterize the solution to the auxiliary problem \mathcal{P} . These lemmas play a central role in the proof of Proposition 1.

Lemma A.1 *The program \mathcal{P} admits the solution $\alpha_t = 0$ for all t (*i.e.* the principal never replaces the performance measure) if and only if $\delta \geq \delta^*$ where δ^* is the smallest δ that satisfies the following condition:*

$$\frac{\delta}{1-\delta} \left(1 - \frac{p}{2-p\delta} \right) s_2 \geq c. \quad (FB)$$

Proof. As $s_2 := y - c_2$ is the (per period) surplus under the first-best allocation (no replacement of measure and effort in both tasks), we must have $u^1 \leq s_2$. When $\alpha_t = 0 \forall t$, from (PK_A^*) we have $u^t = s_2 \forall t$. Hence, $\alpha_t = 0 \forall t$ is a solution to \mathcal{P} , if and only if it is feasible in \mathcal{P} . We show below that $\alpha_t = 0 \forall t$ is feasible in \mathcal{P} if and only if (FB) is satisfied, *i.e.*, $\delta \geq \delta^*$.

Step 1. (Necessity of (FB)) Notice that if the agent knows the identity of the critical task in any period t , he always has the option to use his information immediately and shirk on the non-critical task. Hence, we must have

$$U(u^t) \geq (1-\delta)(y - c_1) + p\delta U(u^{t+1}).$$

But as $u^t = u^{t+1} = s_2$ when $\alpha_t = 0 \forall t$, the above condition boils down to:

$$U(s_2) \geq (1-\delta)(y - c_1) + p\delta U(s_2).$$

Rearranging and plugging $y - c_2 = s_2$, we obtain:

$$U(s_2) - s_2 \geq \frac{1}{1-p\delta} ((1-\delta)c - \delta(1-p)s_2), \quad (A.2)$$

which gives us a lower bound on the agent's expected information rents from a successful deviation.

Now, from (IC_1^*) it follows that

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

i.e.,

$$\left(1 - \frac{1}{2}p\right)\delta s_2 \geq (1-\delta)c + \frac{1}{2}p\delta(U(s_2) - s_2).$$

Using the lower bound for $U(s_2) - s_2$ as in (A.2) and simplifying, we obtain (FB); thus $\alpha_t = 0 \forall t$ is feasible only if (FB) is satisfied.

Step 2. (Sufficiency of (FB)) If $\alpha_t = 0 \forall t$, (IC_0^*) boils down to

$$s_2 \geq (1-\delta)y \Leftrightarrow -(1-\delta)c_2 + \delta s_2 \geq 0 \Leftrightarrow \frac{\delta}{1-\delta}s_2 \geq c_2,$$

which is satisfied when (FB) is satisfied. To see this, observe that

$$\frac{\delta}{1-\delta}s_2 \geq \frac{c}{1-p/(2-p\delta)} \geq \frac{c}{1-p/2} \geq c_2,$$

where the first inequality corresponds precisely to (FB), 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 also satisfied we need to analyse the agent's value from private information when $\alpha_t = 0 \forall t$. Suppose the agent privately learns which task is critical. Given that $\alpha_t = 0 \forall t$, the agent's problem is also stationary. Thus, either the agent never shirks or he always shirks (by doing the critical task only). Suppose first that the agent never shirks. Then, $U(u^{t+1}) = u^{t+1} = u^t = s_2$, and constraint (IC_1^*) collapses to:

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

which is satisfied whenever (FB) is satisfied. Suppose now the agent always shirks. Then $U(s_2) = (1-\delta)(s_2+c) + p\delta U(s_2)$, or,

$$U(s_2) = \frac{1-\delta}{1-p\delta}(s_2+c).$$

So, (IC_1^*) is given by:

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

or,

$$(2-p\delta-p)\delta s_2 \geq (2-p\delta)(1-\delta)c,$$

which is the same as the (FB) above. \square

Lemma A.2 *If there exists a solution to the 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 principal replaces the existing performance measure with a constant probability α^* .*

Proof. For $\delta \geq \delta^*$, the result holds as $\alpha_t = 0$ for all t solves \mathcal{P} (by Lemma A.1). If $\delta < \delta^*$, 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 performance measure is replaced. Note that $u_s^t \leq U(u_t)$ and satisfies the following recursive relation:

$$u_s^t = (1-\delta)(s_2+c) + \delta p \left(\alpha_t s_1 + (1-\alpha_t) u_s^{t+1} \right). \quad (\text{A.3})$$

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

$$u^t \geq (1-\delta)(s_2+c) + \frac{1}{2}p\delta \left(\alpha_t s_1 + (1-\alpha_t) u_s^{t+1} \right), \quad (\text{A.4})$$

or, equivalently,

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

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:

$$\mathcal{P}_R: \max_{\alpha_t \in [0,1]} u^1 \text{ s.t. (A.1), (A.3), and } (IC'_1) \text{ hold for all } t.$$

Step 2. (Rewriting \mathcal{P}_R in terms of α_t) By using (A.1) and (A.3), one can eliminate u^t and u_t^s in \mathcal{P}_R and consider an equivalent problem in terms of α_t s. Note that (A.1) can be rearranged as $u^t - s_1 = (1 - \delta)(s_2 - s_1) + \delta(1 - \alpha_t)(u^{t+1} - s_1)$. So, one obtains:

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

where

$$S_t = (1 - \alpha_t) + \delta(1 - \alpha_t)(1 - \alpha_{t+1}) + \delta^2(1 - \alpha_t)(1 - \alpha_{t+1})(1 - \alpha_{t+2}) + \dots$$

Hence,

$$u^1 = s_1 + (1 - \delta)(s_2 - s_1)(1 + \delta S_1). \quad (\text{A.5})$$

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

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

So,

$$u_s^t - s_1 = (1 - p)((1 - \delta)y - \delta s_1)(1 + \delta p D_t), \quad (\text{A.6})$$

where

$$D_t = (1 - \alpha_t) + (\delta p)(1 - \alpha_t)(1 - \alpha_{t+1}) + (\delta p)^2(1 - \alpha_t)(1 - \alpha_{t+1})(1 - \alpha_{t+2}) + \dots$$

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 $\delta < \delta^*$, and hence, (FB) is violated, it routinely follows that $k_2 > 0$. Hence, (IC'_1) can be rewritten as:

$$D_t \leq A + BS_t \quad \forall t, \quad (\text{A.7})$$

where $A = (k_0 - k_1 - k_2)/k_2 \delta p$ and $B = k_0/pk_2$. So, from (A.5) and (A.7), it follows that \mathcal{P}_R is equivalent to the following program:

$$\mathcal{P}'_R: \max_{\alpha_t \in [0,1]} S_1 \quad \text{s.t.} \quad (\text{A.7}).$$

Step 3. (Rewriting \mathcal{P}'_R in terms of α , S and D) Note the following: (1) Any sequence of $\{\alpha_t\}_{t=1}^\infty$ pins down a unique sequence $\{(S_t, D_t)\}_{t=1}^\infty$. (2) S_t and D_t are non-negative and $S_t \geq D_t$ with equality holding if and only if $(1 - \alpha_t)(1 - \alpha_{t+1}) = 0$. (3) S_t and D_t follow the recursive relations:

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

(4) The set of $\{\alpha_t\}_{t=1}^\infty$ sequences that satisfy (A.7) gives rise to a set of (S, D) tuples 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:

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

(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.

Step 4. (Introducing $f(S)$ function and defining S^*) Note the following about \mathcal{P}''_R : (i) By (PK_R) ,

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

(ii) For any $(S, D) \in \mathcal{F}$, we have

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

Observe that $f(S)$ is increasing (and concave) and $f(S)/S$ is decreasing in S . Also, under the first-best solution where all $\alpha_t = 0$, $(S, D) = (S^{FB}, D^{FB}) = \left(\frac{1}{1 - \delta}, \frac{1}{1 - \delta p}\right)$ and it satisfies $D^{FB} = f(S^{FB})$. (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 $(\bar{S}_1, \mathcal{D}(\bar{S}_1))$. By the recursive relations, $(\bar{S}_1, \bar{D}_1) := (\bar{S}, \mathcal{D}(\bar{S}))$ generates a sequence $\{(\bar{S}_2, \bar{D}_2), (\bar{S}_3, \bar{D}_3), \dots\}$ such that each element of the sequence satisfies (1) $\bar{D}_n \leq A + B\bar{S}_n$ (if not, then (A.7) would be violated in some period) and (2) the recursion relations (PK_R^*) for some associated sequence of $\alpha_t, \{\bar{\alpha}_t\}_{t=1}^\infty$ (say). We will argue in the next four sub-steps (Step 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 $\bar{S}_1 > S^*$, $f(\bar{S}_1) > A + B\bar{S}_1 \geq \bar{D}_1 = \mathcal{D}(\bar{S}_1)$. Next, we claim that $f(\bar{S}_2) \geq \bar{D}_2$.

The proof is given by contradiction: suppose $f(\bar{S}_2) < \bar{D}_2$. 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 $f(\bar{S}_2) < \bar{D}_2$. 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 $f(\bar{S}_2) > \bar{D}_2$. Hence, $f(\bar{S}_2) < \bar{D}_2 \Rightarrow \bar{S}_2 < S^*$.

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

$$\frac{\bar{D}_1}{\bar{S}_1} = \frac{1 + \delta p \bar{D}_2}{1 + \delta \bar{S}_2} > \frac{1 + \delta p f(\bar{S}_2)}{1 + \delta \bar{S}_2} = \frac{f(\bar{S}_2)}{\bar{S}_2} > \frac{f(S^*)}{S^*}, \quad (\text{A.8})$$

where both equalities follow from (PK_R) , the first inequality holds as $f(\bar{S}_2) < \bar{D}_2$ 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 $f(\bar{S}_1) > \bar{D}_1$ we must also have,

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

which contradicts (A.8). Hence, we must have $f(\bar{S}_2) \geq \bar{D}_2$.

As $f(\bar{S}_2) \geq \bar{D}_2$, 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 $\bar{\alpha}_2 = 0$. We show this by contradiction. From (PK_R^*) we know that $(\bar{S}_2, \bar{D}_2) = ((1 - \bar{\alpha}_2)(1 + \delta \bar{S}_3), (1 - \bar{\alpha}_2)(1 + \delta p \bar{D}_3))$. If $\bar{\alpha}_2 > 0$, decrease $\bar{\alpha}_2$ to $\alpha'_2 := \bar{\alpha}_2 - \varepsilon(1 - \bar{\alpha}_2)$ for some $\varepsilon > 0$. Note that $(1 - \alpha'_2) = (1 + \varepsilon)(1 - \bar{\alpha}_2)$. Let $(S'_2, D'_2) := (1 + \varepsilon)(\bar{S}_2, \bar{D}_2)$.

We argue that for sufficiently small ε , (S'_2, D'_2) is feasible. Since $(\bar{S}_3, \bar{D}_3) \in \mathcal{F}$ 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, (\bar{S}_2, \bar{D}_2) must lie on or below the line joining the origin to (\bar{S}_1, \bar{D}_1) .

Now, there are two cases: (1) If (IC_R) is slack at (\bar{S}_1, \bar{D}_1) , all points on this line always lie strictly below the line $D = A + BS$. So, (IC_R) is also slack at (\bar{S}_2, \bar{D}_2) . (2) If (IC_R) binds at (\bar{S}_1, \bar{D}_1) , this is the only point on the line at which (IC_R) binds, and it is slack at all other points. But, as $f(\bar{S}_1) > \bar{D}_1$, 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, $(\bar{S}_2, \bar{D}_2) \neq (\bar{S}_1, \bar{D}_1)$. Therefore, (IC_R) must be slack at (\bar{S}_2, \bar{D}_2) . Thus, for small enough ε , $(S'_2, D'_2) = (1 + \varepsilon)(\bar{S}_2, \bar{D}_2)$ always satisfies (IC_R) .

Next, observe that,

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

Now, we increase $\bar{\alpha}_1$ to some α'_1 where $(1 - \alpha'_1)(1 + \delta S'_2) = \bar{S}_1$. Let $D'_1 = (1 - \alpha'_1)(1 + \delta p D'_2)$. So, by the above inequality, we find that:

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

Hence, $D'_1 < \bar{D}_1$. But this observation contradicts the fact that \bar{D}_1 is the lowest feasible D_1 that supports S_1 (as we have shown that the sequence $\{\alpha'_1, \alpha'_2, \alpha'_3, \dots\}$ is feasible, and it yields $S_1 = \bar{S}_1$ and $D_1 = D'_1 < \bar{D}_1$). Therefore, we must have $\bar{\alpha}_2 = 0$.

Step 5c. We must have $\bar{S}_3 < \bar{S}_2$ and $\bar{D}_3 < \bar{D}_2$. As $\bar{\alpha}_2 = 0$, (PK_R) implies $\bar{S}_2 = 1 + \delta \bar{S}_3$ and $\bar{D}_2 = 1 + \delta p \bar{D}_3$. As $\bar{S}_1 < S^{FB} = 1/(1 - \delta)$ and $\bar{D}_1 < D^{FB} = 1/(1 - \delta p)$ 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 $\bar{\alpha}_t = 0$ for all $t \geq 2$ and the sequence $\{\bar{S}_2, \bar{S}_3, \dots\}$ is monotonically decreasing. So, we must have $\bar{S}_t = 1 + \delta \bar{S}_{t+1}$, $t = 2, 3, \dots$. 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 $\bar{S}_{k-1} = \bar{D}_{k-1} = 1 - \bar{\alpha}_{k-1}$. But this is a contradiction as we know that $\bar{S}_{k-1} = \bar{D}_{k-1}$ only if $(1 - \bar{\alpha}_{k-1})(1 - \bar{\alpha}_k) = 0$ but we have $\bar{\alpha}_{k-1} = \bar{\alpha}_k = 0$.

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

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

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

Define

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

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 (A.7).

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 $u'_s = U(u_t)$. 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) By definition $(s_1, \frac{1-\delta}{\delta}\psi) \in \mathcal{E}$. (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. \square

III. PROOF. Using Lemmas A.1 and A.2, we can now present a proof of Proposition 1.

Proof of Proposition 1. Step 1. Part (i) directly follows from Lemma A.1. If $\alpha_t = 0$ for all t , the constraints in \mathcal{P} are identical to those in original program. So $\alpha_t = 0$ for all t is feasible in original problem \mathcal{P}_O whenever it is feasible in \mathcal{P} . Hence, $\alpha_t = 0$ for all t is the solution to \mathcal{P}_O if and only if $\delta \geq \delta^*$.

Step 2. In order to prove part (ii) we first prove that for a given $\delta < \delta^*$ the auxiliary problem \mathcal{P} has a solution if and only if

$$\frac{\delta}{1 - \delta} \left(1 - \frac{1}{2}p\right) s_1 \geq c. \quad (\text{A.9})$$

Step 2a. (Necessity of (A.9)) If \mathcal{P} has a solution, by Lemma A.2 we know it is stationary: $\alpha_t = \alpha^*$ for all t . Moreover, $\alpha^* > 0$ as we are considering the case where $\delta < \delta^*$. Now, at the solution, the following two conditions must hold: (IC_1^*) in period one holds with equality, and

$$u^{t+1} - s_1 < \frac{1}{2}p \left(U(u^{t+1}) - s_1 \right) \quad (\text{A.10})$$

for $t = 1$. Otherwise, it would be possible to decrease 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 (A.10) implies that if under the optimal contract (IC_1^*) in period 1 is satisfied for $\alpha_1 = \alpha^*$ (which must be the case), then it is also satisfied for $\alpha_1 = 1$, 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 (A.9).

Step 2b. (Sufficiency of (A.9)) Observe that if (A.9) is satisfied then clearly a contract in which $\alpha_t = 1$ 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 (A.9). Thus, at least the contract in which $\alpha_t = 1$ for all t is feasible, meaning that \mathcal{P} has a solution.

Step 3. Using Step 2, we can now prove part (ii) by obtaining a necessary and a sufficient condition for the optimal contracting problem \mathcal{P}_O to have a solution when $\delta < \delta^*$.

Step 3a. Let $v = \max\{u + \pi \mid (u, \pi) \in \mathcal{E}\}$, i.e., the maximal joint payoff sustained in equilibrium. As the optimal contract is stationary for any s_1 (Lemma A.2), and since in the optimal contract the agent receives all surplus in the continuation game (net of cost of replacement of the metric in the current period, if applicable), we must have $s_1 = v - \frac{1 - \delta}{\delta}\psi$. (See Lemma 6 in the Supplementary Appendix for a formal proof).

Also, by definition, v must be the value of the optimal contracting problem under such a s_1 (notice that up on replacing the current measure, the principal would choose the optimal contract in the continuation game, and the continuation game is identical to the game at the beginning of period one).

Observe that when $\delta < \delta^*$, $v \in [s_2 - \psi, s_2]$; $v = s_2 - \psi$ when the performance measure must be replaced in every period and $v = s_2$ when the measure is never replaced.

Step 3b. (Necessary condition for \mathcal{P}_O to have a solution when $\delta < \delta^*$) Since $s_1 = v - \frac{1 - \delta}{\delta}\psi$ and $v \leq s_2$, the highest possible value of s_1 is $s_2 - \frac{1 - \delta}{\delta}\psi$. Hence, for \mathcal{P}_O to have a solution when $\delta < \delta^*$, it must be the case that \mathcal{P} admits a solution when s_1 is set at $s_2 - \frac{1 - \delta}{\delta}\psi$. Therefore, (A.9) implies that if $\delta < \delta^*$, \mathcal{P}_O has a solution (where $\alpha^* > 0$) only if

$$\frac{\delta^*}{1 - \delta^*} \left(s_2 - \frac{1 - \delta^*}{\delta^*} \psi \right) \geq \frac{c}{1 - \frac{1}{2}p},$$

or, equivalently,

$$\frac{\delta^*}{1 - \delta^*} s_2 \geq \frac{c}{1 - \frac{1}{2}p} + \psi. \tag{A.11}$$

Step 3c. (Sufficient condition for \mathcal{P}_O to have a solution when $\delta < \delta^*$) Since $v \geq s_2 - \psi$, the lowest possible value of s_1 is $s_2 - \frac{\psi}{\delta}$. We claim that \mathcal{P}_O has a solution (when $\delta < \delta^*$) if \mathcal{P} admits a solution when s_1 is set at $s_2 - \frac{\psi}{\delta}$. That is, if (by (A.9)):

$$\frac{\delta}{1 - \delta} s_2 \geq \frac{c}{1 - \frac{1}{2}p} + \frac{\psi}{1 - \delta}. \tag{A.12}$$

The proof of this claim is as follows. Denote $u^1(\hat{v})$ as the value associated with the solution to the auxiliary problem \mathcal{P} when $s_1 = \hat{v} - \frac{1 - \delta}{\delta}\psi$ for any $\hat{v} \in [s_2 - \psi, s_2]$. The principal's problem \mathcal{P}_O admits a solution if and only if $u^1(\hat{v})$ has a fixed point in $[s_2 - \psi, s_2]$. Such a fixed point must exist as (1) $u^1(\hat{v})$ is continuous, (2) $u^1(\hat{v}) \geq \hat{v}$ when $\hat{v} = s_2 - \psi$ and (A.12) is satisfied, and (3) $u^1(\hat{v}) < \hat{v}$ when $\hat{v} = s_2$ (as we fix $\delta < \delta^*$, $u^1 < s_2$).

Notice that both (A.11) and (A.12) are satisfied if ψ is below a threshold (given δ), and if so, the optimal contracting problem \mathcal{P}_O would admit a solution where the associated value is $v = \max\{\hat{v} \in [s_2 - \psi, s_2] : u^1(\hat{v}) = \hat{v}\}$. Let δ_R be the smallest δ (given ψ) for which such a solution exists. This observation completes the proof of part (ii).

Step 4. Finally, as \mathcal{P} has no solution when $\delta < \delta_R$, effort in both tasks cannot be elicited (even in period one). Hence, it is optimal for the principal and agent to take their outside options in every period. \square

REFERENCES

ABREU, D., PEARCE, D. and STACCHETTI, E. (1990), "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring", *Econometrica*, **58**, 1041–1063.
 AOYAGI, M. (2010), "Information Feedback in a Dynamic Tournament", *Games and Economic Behavior*, **70**, 242–260.
 BAKER, G. (1992), "Incentive Contracts and Performance Measurement", *Journal of Political Economy*, **100**, 598–614.
 BAKER, G., GIBBONS, R. and MURPHY, K. (1994), "Subjective Performance Measures in Optimal Incentive Contracts", *The Quarterly Journal of Economics*, **109**, 1125–1156.
 BAKER, G., GIBBONS, R. and MURPHY, K. (2002), "Relational Contracts and the Theory of the Firm", *The Quarterly Journal of Economics*, **117**, 39–84.
 BARDHAN, P. (1997), "Corruption and Development: A Review of Issues", *Journal of Economic Literature*, **35**, 1320–1346.
 BARRON, D. and POWELL, M. (2019), "Policies in Relational Contracts", *American Economic Journal: Microeconomics*, **11**, 228–249.
 BATTAGLINI, M. (2005), "Long-Term Contracting With Markovian Consumers", *American Economic Review*, **95**, 637–658.

- BERGEMANN, D. and HEGE, U. (2005), "The Financing of Innovation: Learning and Stopping", *RAND Journal of Economics*, **36**, 719–752.
- BEVAN, G. and HOOD, C. (2004), "Targets, Inspections, and Transparency", *British Medical Journal*, **328**, 598.
- BEVAN, G. and HOOD, C. (2006), "What's Measured is What Matters: Targets and Gaming in the English Public Health Care System", *Public Administration*, **84**, 517–538.
- BEYER, A., GUTTMAN, I. and MARINOVIC, I. (2014), "Optimal Contracts with Performance Manipulation", *Journal of Accounting Research*, **52**, 817–847.
- BHASKAR, V. (2014), "The Ratchet Effect Re-examined: A Learning Perspective" (Mimeo, University of Texas-Austin).
- BHASKAR, V. and MAILATH, G. (2019), "The Curse of Long Horizons", *Journal of Mathematical Economics*, **82**, 74–89.
- BOLTON, P. and HARRIS, C. (1999), "Strategic Experimentation", *Econometrica*, **67**, 349–374.
- BONATTI, A. and HÖRNER, J. (2011), "Collaborating", *American Economic Review*, **101**, 632–663.
- BONATTI, A. and HÖRNER, J. (2017), "Learning to Disagree in a Game of Experimentation", *Journal of Economic Theory*, **169**, 234–269.
- BROWN, L. and SPINA PINELLO, A. (2007), "To What Extent Does the Financial Reporting Process Curb Earnings Surprise Games?", *Journal of Accounting Research* **45**, 947–981.
- BURROUGH, B. and HELYAR, J. (1990), *Barbarians at the Gate: The Fall of RJR Nabisco* (New York: Harper & Row).
- CARHART, M., KANIEL, R. and MUSTO, D. K. (2002), "Leaning for the Tape: Evidence of Gaming Behavior in Equity Mutual Funds" *Journal of Finance*, **57**, 661–693.
- CHASSANG, S. (2010), "Building Routines: Learning, Cooperation, and the Dynamics of Incomplete Relational Contracts", *American Economic Review*, **100**, 448–465.
- CHASSANG, S. and PADRÓ I MIQUEL, G. (2018), "Crime, Intimidation, and Whistleblowing: A Theory of Inference from Unverifiable Reports", *The Review of Economic Studies*, **86**, 2530–2553.
- CISTERNAS, G. (2018), "Two-Sided Learning and the Ratchet Principle", *The Review of Economic Studies*, **85**, 307–351.
- COURTY, P. and MARSCHKE, G. (1997), "Measuring Government Performance: Lessons from a Federal Job-Training Program", *American Economic Review*, **87** (Papers and Proceedings), 383–388.
- COURTY, P. and MARSCHKE, G. (2004), "An Empirical Investigation of Gaming Responses to Explicit Performance Incentives", *Journal of Labor Economics*, **22**, 23–56.
- COURTY, P. and MARSCHKE, G. (2007), "Making Government Accountable: Lessons from a Federal Job Training Program", *Public Administration Review*, **67**, 904–916.
- DEB, J., LI, J. and MUKHERJEE, A. (2016), "Relational Contracts with Subjective Peer Evaluations", *RAND Journal of Economics*, **47**, 3–28.
- DEMARZO, P. and SANNIKOV, Y. (2017), "Learning, Termination, and Payout Policy in Dynamic Incentive Contracts", *The Review of Economic Studies*, **84**, 182–236.
- DEMSKI, J. (1998), "Performance Measure Manipulation", *Contemporary Accounting Research*, **15**, 261–285.
- DEMSKI, J. (2008), *Managerial Uses of Accounting Information* (New York, NY: Springer Science Business Media).
- DESSEIN, W., GALEOTTI, A. and SANTOS, T. (2016), "Rational Inattention and Organizational Focus", *American Economic Review*, **106**, 1522–1536.
- DESSEIN, W. and SANTOS, T. (2020), "Managerial Style and Attention", *American Economic Journal: Microeconomics*, forthcoming.
- DRANOVE, D., KESSLER, D., MCCLELLAN, M. et al. (2003), "Is More Information Better? The Effects of "Report Cards" on Health Care Providers", *Journal of Political Economy*, **111**, 555–588.
- ECCLES, R. (1991), "The Performance Measurement Manifesto", *Harvard Business Review*, January–February, 131–137.
- EDERER, F. (2010), "Feedback and Motivation in Dynamic Tournaments", *Journal of Economics & Management Strategy*, **19**, 733–769.
- EDERER, F., HOLDEN, R. and MEYER, M. (2018), "Gaming and Strategic Opacity in Incentive Provision", *RAND Journal of Economics*, **48**, 819–854.
- FONG, Y.-F. and LI, J. (2017), "Information Revelation in Relational Contracts", *Review of Economic Studies*, **84**, 277–299.
- FRANKEL, A. and KARTIK, N. (2019), "Muddled Information", *Journal of Political Economy*, **127**, 1739–1776.
- FUCHS, W. (2007), "Contracting with Repeated Moral Hazard and Private Evaluations", *American Economic Review*, **97**, 1432–1448.
- GARICANO, L. and RAYO, L. (2016), "Why Organizations Fail: Models and Cases", *Journal of Economic Literature*, **54**, 137–192.
- GAWANDE, A. (2010), *The Checklist Manifesto: How to get Things Right* (New York: Metropolitan Books).
- GIBBONS, R. and HENDERSON, R. (2012), "What Do Managers Do?: Exploring Persistent Performance Differences among Seemingly Similar Enterprises", in Gibbons, R., and Roberts, J. (eds) *The Handbook of Organizational Economics* (Princeton: Princeton University Press) 680–731.
- GIBBONS, R. and KAPLAN, R. (2015), "Formal Measures in Informal Management: Can a Balanced Scorecard Change a Culture?", *American Economic Review*, **105**, 447–451.
- GODDARD, M., MANNION, R. and SMITH, P. (2000), "Enhancing Performance in Health Care: A Theoretical Perspective on Agency and the Role of Information", *Health Economics*, **9**, 95–107.
- GOLTSMAN, M. and MUKHERJEE, A. (2011), "Interim Performance Feedback in Multistage Tournaments: The Optimality of Partial Disclosure", *Journal of Labor Economics*, **29**, 229–265.
- GUO, Y. (2016), "Dynamic Delegation of Experimentation", *American Economic Review*, **106**, 1969–2008.
- HALAC, M., KARTIK, N. and LIU, Q. (2016), "Optimal Contracts for Experimentation", *Review of Economic Studies*, **83**, 1040–1091.

- HALAC, M. and PRAT, A. (2016), "Managerial Attention and Worker Performance", *American Economic Review*, **106**, 3104–3132.
- HOFFMAN, M. and TADELIS, S. (2018), "People Management Skills, Employee Attrition, and Manager Rewards: An Empirical Analysis" (NBER Working Paper No. 24360).
- HOOD, C. (2006), "Gaming in Targetworld: The Targets Approach to Managing British Public Services", *Public Administration Review*, **66**, 515–521.
- HÖRNER, J. and SAMUELSON, L. (2013), "Incentives for Experimenting Agents", *RAND Journal of Economics*, **44**, 632–663.
- ISHIHARA, A. (2017), "Relational Contracting and Endogenous Formation of Teamwork", *RAND Journal of Economics*, **48**, 335–357.
- ITTNER, C., LARCKER, D. and RAJAN, M. (1997), "The Choice of Performance Measures in Annual Bonus Contracts", *The Accounting Review*, **72**, 231–255.
- ITTNER, C., LARCKER, D. and MEYER, M. (2003), "Subjectivity and the Weighting of Performance Measures: Evidence from a Balanced Scorecard", *The Accounting Review*, **78**, 725–758.
- JACOB, B. and LEVITT, S. (2003), "Rotten Apples: An Investigation of the Prevalence and Predictors of Teacher Cheating", *Quarterly Journal of Economics*, **118**, 843–877.
- JEHIEL, P. and NEWMAN, A. F. (2015), "Loopholes and the Evolution of Contract Form" (Mimeo, Boston University).
- KAPLAN, R. and NORTON, D. (1992), "The Balanced Scorecard: Measures that Drive Performance" (Harvard Business Review), January–February, 71–79.
- KAPLAN, R. and NORTON, D. (1993), "Putting the Balanced Scorecard to Work", *Harvard Business Review*, September–October, 134–147.
- KELLER, G., RADY, S. and CRIPPS, M. (2005), "Strategic Experimentation with Exponential Bandits", *Econometrica*, **73**, 39–68.
- LAKONISHOK, J., SHLEIFER, A., THALER, R. et al. (1991), "Window Dressing by Pension Fund Managers", *American Economic Review*, **81**, 227–231.
- LAZEAR, E. (2006), "Speeding, Terrorism, and Teaching to the Test", *Quarterly Journal of Economics*, **121**, 1029–1061.
- LAZEAR, E. P., SHAW, K. L. and STANTON, C. (2015), "The Value of Bosses", *Journal of Labor Economics*, **33**, 823–861.
- LEVIN, J. (2003), "Relational Incentive Contracts", *American Economic Review*, **93**, 835–857.
- LI, J. and MATOUSCHEK, N. (2013), "Managing Conflicts in Relational Contracts", *American Economic Review*, **103**, 2328–2351.
- LI, J., MATOUSCHEK, N. and POWELL, M. (2017), "Power Dynamics in Organizations", *American Economic Journal: Microeconomics*, **9**, 217–241.
- MALCOMSON, J. (2013), "Relational Incentive Contracts", in Gibbons, R., and Roberts, J. (eds) *The Handbook of Organizational Economics* (Princeton: Princeton University Press) 1014–1065.
- MALCOMSON, J. (2016), "Relational Incentive Contracts With Persistent Private Information", *Econometrica*, **84**, 317–346.
- MANSO G. (2011), "Motivating Innovation", *Journal of Finance*, **66**, 1823–1869.
- MEYER, M. (2002), *Rethinking Performance Measurement: Beyond the Balanced Scorecard* (Cambridge: New York: Cambridge University Press).
- MORONI, S. (2016), "Experimentation in Organizations" (Working Paper, University of Pittsburgh).
- MUKHERJEE, A. (2010), "The Optimal Disclosure Policy when Firms Offer Implicit Contracts", *RAND Journal of Economics*, **41**, 549–573.
- MUKHERJEE, A. and VASCONCELOS, L. (2011), "Optimal Job Design in the Presence of Implicit Contracts", *RAND Journal of Economics*, **42**, 44–69.
- ORLOV, D. (2018), "Frequent Monitoring in Dynamic Contracts" (Mimeo, University of Rochester).
- ORTEGA, J. (2001), "Job Rotation as a Learning Mechanism", *Management Science*, **147**, 1361–1370.
- ORTNER, J. and CHASSANG, S. (2018), "Making Corruption Harder: Asymmetric Information, Collusion and Crime", *Journal of Political Economy*, **126**, 2108–2133.
- OYER, P. (1998), "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality", *Quarterly Journal of Economics*, **113**, 149–185.
- PRAT, J. and JOVANOVIĆ, B. (2014), "Dynamic Contracts when Agent's Quality is Unknown", *Theoretical Economics*, **9**, 865–914.
- RAYO, L. (2007), "Relational Incentives and Moral Hazard in Teams", *Review of Economic Studies*, **74**, 937–963.
- RIVKIN, J., ROBERTO, M. and GULATI, R. (2010), "Federal Bureau of Investigation, 2009", *Harvard Business Review Case # 9-710-452*.
- SANNIKOV, Y. (2014), "Moral Hazard and Long-Run Incentives" (Mimeo, Stanford University).
- SCHÖTTNER, A. (2008), "Relational Contracts, Multitasking, and Job Design", *Journal of Law, Economics, and Organization*, **24**, 138–162.
- WELCH, J. and BYRNE, J. (2001), *Jack: Straight from the Gut* (New York: Warner Business Books).
- WILLIAMS, N. (2011), "Persistent Private Information", *Econometrica*, **79**, 1233–1275.
- YANG, H. (2013), "Non-stationary Relational Contracts with Adverse Selection", *International Economic Review*, **54**, 525–547.
- ZABOJNIK, J. (2014), "Subjective Evaluations with Performance Feedback", *The RAND Journal of Economics*, **45**, 341–369.