



HAL
open science

Value by Design: Performance Evaluation and Learning from Wages

Franz Ostrizek

► **To cite this version:**

| Franz Ostrizek. Value by Design: Performance Evaluation and Learning from Wages. 2025. ⟨hal-04938068⟩

HAL Id: hal-04938068

<https://sciencespo.hal.science/hal-04938068v1>

Preprint submitted on 10 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-NC-ND 4.0 - Attribution - Non-commercial use - No Derivative Works - International License

VALUE BY DESIGN: PERFORMANCE EVALUATION AND LEARNING FROM WAGES

Franz Ostrizek

SCIENCES PO ECONOMICS DISCUSSION PAPER

No. 2025-01

Vague by Design: Performance Evaluation and Learning from Wages*

Franz Ostrizek[†]

February 9, 2025

Abstract

We study a dynamic principal-agent setting in which both sides learn about the importance of effort. The quality of the agent's output is not observed directly. Instead, the principal jointly designs an evaluation technology and a wage schedule. More precise performance evaluation reduces current agency costs but promotes learning, which can increase future agency costs. As a result, the optimal evaluation technology is noisy. Performance information that solely pertains to effort is revealed, while the principal optimally conceals information about the agent's ability. The optimal evaluation features a censorship pattern with a base wage and tailored bonuses/penalties when the threshold is exceeded. With binary output, it is both imprecise and tough if effort and ability are strong complements: a bad performance is always sanctioned, but a good one is not always recognized.

1 Introduction

Many firms motivate their workers to exert effort with incentive pay based on objective measures of performance.¹ Such measures have become richer and easier to obtain. For example, board computers and GPS tracking allow for better monitoring of truck drivers and time tracking software in professional services simplifies billing while also logging the

*I am grateful to Roland Bénabou, Pietro Ortoleva, Wolfgang Pesendorfer and Leeat Yariv for their guidance and to Roberto Corrao, Ludmila Matysková, Sofia Moroni, Leon Musolf, Eduardo Perez-Richet, Pellumb Reshidi, Evgenii Safonov, Denis Shishkin, Nikhil Vellodi, Can Urgan and seminar audiences at Princeton University, PSE, Northwestern Kellogg, Notre Dame, University of Pittsburgh, ESMT, DICE, University of Bonn, University of Vienna, CERGE-EI, University of Pennsylvania, Sciences Po, QMUL, CREST, LEMMA, Helsinki GSE, the World Congress of the Econometric Society 2020, the World Congress of Game Theory 2020, the 2021 SEA Annual Meeting, SEAT, TTW, ESMES, and NAWMES for helpful comments and discussions. Funding by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (Project B02) is gratefully acknowledged.

[†]Department of Economics, Sciences Po, Paris. franz.ostrizek@gmail.com

¹According to data from the BLS National Compensation Survey, 39% of hours worked in US private sector firms in 2013 were in jobs with performance-related pay. 21% fall into a narrower classification of performance-related pay excluding, among other categories, referral bonuses which should arguably be excluded from a theoretical perspective, but also safety bonuses which should be retained (Gittleman and Pierce, 2013).

activities of employees; shop floor control systems monitor not only the flow of components but also workers; recent advances in language processing and AI enables data collection in applications ranging from call centers to health care.² Using such data, firms can develop better objective measures of workers' contribution to profits.

Should this additional information be used to set performance pay? What should be measured and how should the information be aggregated? The theory of incentives seems to offer a simple answer to these questions. Providing incentives to workers is costly exactly because the underlying performance measures are only partially informative about effort. Conversely, more hard information about the workers' contribution is always helpful and should be used as a basis of performance pay (Holmström, 1979; Grossman and Hart, 1983). In particular, it is not optimal to base incentives on a noisy signal instead of the contribution to output.³

In this paper, we show that this conclusion changes fundamentally when we also consider how performance evaluation also provides information to workers. Often, workers do not directly observe their own contribution to profits – which can be hard to grasp in complex organizations and occurs far from the worker on the consumer-facing side of the firm. Performance evaluations, however, are observed – certainly insofar as they determine their wage and bonus. By basing wages on a precise evaluation of output, the firm therefore reveals information about his output to the worker. Based on this information about performance, the worker can learn about his ability and fit to the job. This learning, in turn, can make it more costly to employ and motivate the agent in the future.

To capture this fundamental trade-off between incentives and information, we develop a model of twice-repeated moral hazard with learning (Section 2). The agent has a persistent match-specific ability that is unknown both to the agent and to the firm. The agent's ability affects not only average output, but also the effectiveness of effort. The quality of output is not observed directly. Instead, in every period the principal flexibly designs not only wages but also the underlying performance evaluation. The agent observes his evaluations and wages. The evaluation structure therefore determines the cost of incentives this period and the extent of learning.

We analyze the terminal period in Section 3. In the terminal period, the continuation belief is of no importance to the principal and the optimal evaluation structure is therefore fully informative (Holmström, 1979; Grossman and Hart, 1983). In the first period, however, there is a novel trade-off. A more precise evaluation structure reduces agency costs this period, but induces more agent learning. The principal's attitude towards this learning is determined by the shape of profits in the terminal period as a function of the agent's posterior belief. In general, this shape is jointly determined by the agent's utility function

²These tools allow call centers to detect the callers' mood, for example (Singer, 2013). For a survey on the use of natural language processing to extract information from health-related text, see Gonzalez-Hernandez et al. (2017).

³That it is not optimal to add noise continues to hold across a wider class of models, including multi-tasking (Holmström and Milgrom, 1991) and linear-Gaussian career concerns models (Holmström, 1999; Hörner and Lambert, 2021). We discuss some settings in which the addition of noise can be beneficial in the related literature.

and the details of the technology. We discuss the forces involved – the *filtering* effect allowing the principal to tailor the bonus scheme to the agent’s ability which increases expected profits and the *agent-learning* effect through the incentive compatibility constraint which decreases expected profits – in Section 3.2. We show that with binary output the latter effect dominates and the continuation profit is concave in the agent’s posterior belief, if the interaction between effort and ability is sufficiently strong (log-supermodular or log-submodular technology).

We return to the joint design of the evaluation and wage in the first period in Section 4. We transform the contracting problem into an information design problem with additional constraints (participation and incentive compatibility) and an additional choice variable, the wage at every posterior. We proceed in two steps. We first write the contracting problem in the agent’s posterior belief about output after observing the evaluation. We show that the optimal evaluation reveals (at least) all information contained in output that is only informative about the agent’s effort but not his ability. We can further simplify the problem when there is a one-to-one mapping between information about effort and information about ability. This is the case (essentially) if and only if the distribution over output satisfies a linear distribution function condition. In this case, we can rewrite the contracting problem directly in terms of the posterior belief about the agent’s ability.

We solve for the optimal contract and evaluation structure in Section 4.4 depending on the shape of the continuation value. If the principal prefers the agent to learn, clearly, the optimal evaluation structure is fully informative. If learning is more costly at low posteriors, the optimal evaluation features lower-censorship: output realizations below a threshold are pooled and output realizations above the threshold are exactly revealed. This corresponds to discontinuous wage scheme with a base wage and tailored bonuses above the performance threshold. The pattern is mirrored if learning is more costly at high posteriors.

In the case of binary output, we directly link the features of the contract to the technology. The optimal evaluation is binary; the additional motive of shaping learning does not add complexity relative to the fully informative evaluation. If effort and ability are strong complements the optimal evaluation structure is *tough*: The agent obtains a high evaluation and therefore a bonus only if his output is high quality. Even if output was high, however, he may receive a low evaluation and thus fail to obtain the bonus. After low quality output, he never obtains the bonus. This information structure is optimal because it avoids inducing very low posteriors. Agents with such beliefs would be very expensive to motivate in the next period and even a small increase in their posterior belief significantly reduces the required bonus. If effort and ability are strong substitutes, instead, the principal is most worried about complacent agents who are expensive to motivate because they believe themselves to be of high ability. The resulting optimal evaluation structure pools at the top, providing a bonus sometimes even if output was low in order to make the good signal less informative and avoid high posterior beliefs going into the second period.

The effects we study highlight an important consideration in the design of performance evaluation: it is not only the basis for incentives but also for workers' learning about their ability, their task, and their match to the organization. This is a rich interaction with many facets. Our model is intentionally stylized to retain tractability even though evaluations and contracts are flexibly designed and focuses on learning about match-specific ability. Nevertheless, our results can speak to several notable patterns in incentive systems. First, we often see performance bonuses that are conditional on reaching a performance target. Importantly, these bonus schemes are discontinuous as the bonus upon hitting the target is substantial. This pattern directly follows from the informational aspect of performance evaluation analyzed in our model but is difficult to rationalize otherwise. Second, when output is a coarse measure of performance, we would expect tough evaluations in professions with a strong complementarity between skills and effort such as law or investment banking. Indeed, these professions are generally associated with an uncompromising mentality. This is especially prevalent in evaluating the work of fresh associates, which is in line with our model as well: As the trade-off is intertemporal, evaluations become more informative and less tough over time (Section 5.4). Finally, performance evaluations often focus on the conduct of workers on the job and less so on available measures of output. In our setting, this is optimal since such evaluations allow the firm to motivate without revealing information about ability (Proposition 4) while it is hard to explain in the standard model where direct and indirect information about effort is used equally.

In Section 5 we consider several extensions of our model. We show that the optimal evaluation remains the principal-preferred equilibrium if the principal can acquire private information about the agent's performance, and we discuss when this can be supported as an equilibrium of the endogenously-informed principal game when the principal can deviate privately. The structure of the optimal evaluation is also preserved when effort is unobserved in addition to being noncontractible (in the binary case), and when the principal can commit to a continuation value. We also study the long-run evolution of beliefs. Section 6 concludes. The proofs not given in the text are collected in the Appendix.

Related Literature

This paper contributes to the large literature on information in moral-hazard models. We offer a counterpoint to the classic results establishing that more precise evaluation reduces agency costs (Holmström, 1979; Grossman and Hart, 1983; Kim, 1995) by providing a setting in which the principal prefers to base wages on a noisy information structure.

We show that noisy evaluation is optimal even though verifiable information about the agent's true performance would be available. Several strands of the literature show that coarse or noisy evaluation is optimal when such information is not available, for instance with multitasking (Holmström and Milgrom, 1991) or when the agent has private information that would allow him to game a deterministic incentive scheme (Ederer et al., 2018). Similarly, coarse rewards emerge when the evaluation is subjective, i.e. based on unverifiable private information of the principal (MacLeod, 2003; Fuchs, 2007).

That more information about the technology can reduce profits in a moral hazard setting has been noted in the literature in several settings. In general, it is well understood that ex-post incentive compatibility is more demanding than ex-ante incentive compatibility. [Lizzeri et al. \(2002\)](#) show that interim performance evaluation is not optimal when there is no learning.⁴ [Nafziger \(2009\)](#) demonstrates that it can be optimal to conceal information until after the agent’s effort choice, even though this precludes the principal from adjusting the implemented action. Indeed, such situations occur generically if the problem is sufficiently rich ([Jehiel, 2015](#)). In all these papers, the wage is still allowed to depend on the true realization of the signal, even if it is not revealed ex-ante. We show that less information about the technology increases profits even if this implies that the wage cannot depend on the state even ex-post.⁵

To our knowledge, this is the first paper to combine the three key features of explicit incentives, learning about a persistent type, and information design. There are several literatures combining each two of these features.

A growing literature investigates the design of information structures in one-shot moral hazard problems with commitment to a wage scheme. The older literature ([Dye, 1986](#); [Feltham and Xie, 1994](#); [Datar et al., 2001](#)) considers the optimal acquisition and aggregation of information within a parametric class.⁶ [Demougin and Fluet \(2001\)](#) study the case with limited liability which results in a binary evaluation. In [Georgiadis and Szentes \(2020\)](#) and [Li and Yang \(2020\)](#) the costs of information acquisition are assumed as part of the monitoring technology. [Dai et al. \(2019\)](#) study the optimal contract when the principal can allocate attention between finding good and bad news. [Hoffmann et al. \(2021\)](#) analyze a setting where the agent takes a single action, but information about his performance arrives over time. Information acquisition requires delayed payments, which creates endogenous costs because of impatience and imperfect risk sharing. [Ely et al. \(2024\)](#) study the joint design of feedback and reward when effort is exerted over time to obtain a breakthrough. Perhaps the closest reference regarding the analyzed trade-off is [Orlov et al. \(2020\)](#), which studies the optimal contract and intensity of monitoring in a dynamic setting with limited liability. The central trade-off is between monitoring to avoid wasteful investment and thereby revealing information about the continuation value which is costly. We analyze not only the intensity but also the shape of the optimal evaluation when monitoring is required for incentives and its costs stem from the agent’s learning about his type.

Learning about a persistent state and information design are combined in a growing literature. Most closely related to our moral-hazard setting are [Smolin \(2021\)](#) and [Ely and Szydlowski \(2019\)](#) in which the principal uses information, which is valuable for the agent, as an incentive. We analyze the role of information design when – since the principal sets incentives and ability is match-specific – information itself is not valuable. Information and

⁴[Ederer \(2010\)](#) shows that the optimality of interim performance evaluation in a tournament setting with exogenous rewards depends on the shape of the effort cost function.

⁵Under this assumption, [Fang and Moscarini \(2005\)](#) show that information is detrimental if it erodes profitable overconfidence.

⁶When restricting attention to linear contracts, it can be optimal to leave information unused ([Feltham and Xie, 1994](#); [Datar et al., 2001](#)) This is a consequence of the restricted space of contracts, however.

incentive design constrain each other, as the principal reveals at least as much information as is contained in wages. Information and implicit incentives are also linked in models of career concerns (Holmström, 1999), a connection we return to in the conclusion.

The literature on learning in moral hazard models (Adrian and Westerfield, 2009; Giat et al., 2010; Prat and Jovanovic, 2014; Demarzo and Sannikov, 2017) studies learning based on output while we study learning based on an information structure that is designed endogenously by the principal. Another important distinction is that we consider learning about the importance of effort as opposed to learning about a state that affects only the level of output, which is often considerably more tractable (see Bhaskar and Mailath (2019) and Bhaskar and Roketskiy (2023) for notable exceptions).

The formal approach to the analysis in our paper relates to the literature on information design, in particular recent contributions to information design problems with constraints (Boleslavsky and Kim, 2017; Le Treust and Tomala, 2019; Doval and Skreta, 2023) and additional choice variables (Georgiadis and Szentes, 2020) – in our case wages. We consider a setting where the information designer chooses a signal structure about one variable – output – in order to affect beliefs about another – the ability of the agent. We also contribute to the literature studying solutions to linear persuasion problem by generalizing the optimality of censorship information structures (Kolotilin et al., 2022) to a setting without regularity restrictions on the dominating distribution.

2 The Model

A principal (she) employs an agent (he) for two periods. The principal is risk neutral, the agent is risk averse with a strictly increasing utility index $u : [0, \infty) \rightarrow [0, \infty)$ which we assume to be unbounded with $u'(x) \rightarrow 0$ as $x \rightarrow \infty$. Both share a common discount factor $\delta \in (0, 1]$.

Technology

Each period, the agent exerts nonverifiable effort $e_t \in \{0, 1\}$ at cost $c \cdot e_t$, with $c > 0$. The worker has a time-invariant ability $\theta \in \Theta$ with finite Θ . The principal and the agent share a common prior belief $\mu_0 \in \Delta\Theta$ about the agent's ability. Output $y \in Y$ realizes each period in a compact set $Y \subset \mathbb{R}$ according to a distribution $F(y|e, \theta)$. We assume that the distributions $\{F(y|e, \theta)\}_{e, \theta \in \{0, 1\} \times \Theta}$ are mutually absolutely continuous and we write $F(y|e, \mu)$ for the expected distribution over output for a belief $\mu \in \Delta\Theta$. We assume that the principal wants to implement high effort in both periods and after all histories.⁷

⁷It is easy to see that implementing high effort after all histories is optimal for the principal for a sufficiently high gain from high effort. Focusing on this case sharpens the trade-off between incentives and learning we aim to investigate, as the principal derives no instrumental value of information. Furthermore, implementing a given effort level is a standard focus in the contracting literature.

Information, Contracts and Commitment

In the main sections, we assume that the principal has full commitment within each period, but no commitment across periods. Within every period, the timing is as follows: The principal proposes a contract (S, p, w) , comprising a measurable signal space S , a distribution over signals $p : Y \rightarrow \Delta S$, where we denote by $p(\cdot|y)$ the distribution over signals conditional on output y , and a mapping $w : S \rightarrow \mathbb{R}$ from signals to wages.⁸ Having observed the contract, the agent decides whether to quit and obtain outside utility U or to work, choosing effort level e_t . The outside utility is independent of the agent's match-specific type. At the end of the period, output, signals and wages realize.

Output is informative about the agent's type, but not directly observed by the agent or the principal.⁹ The principal and the agent observe (noncontractible) effort, signals and wages and update their beliefs about the agent's type according to Bayes rule.¹⁰ Therefore, the evaluation designed by the principal has the dual role of serving as the basis for incentive pay and determining the learning environment.

The principal's problem in period $t \in \{1, 2\}$ is given by

$$\Pi_t(\mu_t) = \max_{S, p, w} \iint_{S \times Y} (y - w(s) + \delta \Pi_{t+1}(\mu_{t+1}(s))) dp(s|y) dF(y|1, \mu_t) \quad (1)$$

$$\text{s.t.} \quad \iint_{S \times Y} u(w(s)) dp(s|y) dF(y|1, \mu_t) - c \geq U \quad (\text{P})$$

$$\iint_{S \times Y} u(w(s)) dp(s|y) dF(y|1, \mu_t) - c \geq \iint_{S \times Y} u(w(s)) dp(s|y) dF(y|0, \mu_t) \quad (\text{IC})$$

where $\mu_{t+1}(s)$ denotes the posterior belief after signal s and the terminal continuation value satisfies $\Pi_3 \equiv 0$.

This is a standard moral hazard problem with two added features. First, the principal *chooses an evaluation structure* and the wage cannot be more informative about the agent's output than the evaluation structure it is based on. In particular, the principal can choose to condition the wage on partially informative signals of output instead of output directly. Second, there is a *belief-dependent continuation profit* $\Pi_{t+1}(\mu(s))$.

⁸Concretely, it is without loss of generality to set $S = \Delta Y \times [0, 1]$. The restriction to deterministic wages conditional on the signal is without loss, as the principal can simply extend the signal space (as we just did) to generate any desired randomness in the wage. Furthermore, such randomness is never optimal in our setting.

⁹In a large organization, this assumption is consistent with the firm observing statistics such as output, profit, and sales *in aggregate*, as long as it is difficult to link aggregate shortfalls to the individual worker. Formally, consider the model with a continuum of agents. Through its regular accounting activities, the firm observes aggregate outcomes such as profits, revenues or the average quality of output. These outcomes are not informative about the performance about an individual, infinitesimal agent.

¹⁰The agent observes not only the wage but also the evaluation. This allows the principal to reveal more information about output to the agent than is used in wage-setting. This will not be essential, but is allowed for additional generality. It also helps to ensure that the information revealed does not depend on the wage schedule in a "discontinuous" way as the principal could always reveal additional information through the wage schedule directly by infinitesimal changes.

Discussion

An important feature of our model is that the firm cannot separate providing incentives from revealing the information about the agent's performance that underlies them. If the principal could record the output of the agent without revealing this information and credibly commit to contingent payments at the end of the relationship and if such a delay was costless (as with risk-neutrality and common discounting), a fully informative and fully delayed evaluation would be optimal.¹¹ Our mechanism comes into effect when such complete backloading is costly or infeasible. In the main sections, this is achieved in a tractable manner by restricting the principal to period-by-period contracts. This lack of intertemporal commitment ensures that incentives for effort have to be provided in the concurrent period. Wages in the first period are thus informative about output. More generally, risk aversion itself makes backloading information costly: Period-by-period contracting makes the model tractable, but is not the source of the trade-off between incentive provision and learning. We return to this commitment assumption and how it can be relaxed in Section 5.3.

The present model features learning based on an endogenous signal distribution with hidden actions and therefore has the potential to create subtle issues of endogenous private information, both for the principal and the agent. For tractability and to focus on the main trade-off between incentives and learning, our assumptions in the baseline model ensure that no endogenous private information arises. Regarding the principal, she does not acquire private information about the type of the agent, since she does not privately observe the quality of output itself but only the result of the public evaluation. We relax this assumption in Section 5.1 and show how our results generalize. Regarding the agent, the benchmark model ensures that his posterior belief remains common knowledge after a deviation to lower effort, since effort is noncontractible but observed. If effort was unobserved, the agent would acquire private information about his belief after a deviation and double deviations to low effort in both periods may be profitable. In Section 5.2, we derive the resulting dynamic incentive compatibility constraint and show how our results generalize to this case.

Taken together, our assumptions ensure that the common belief of the agent and principal is the only state variable of the problem and that the constraints are static as well. The participation constraint is static since the agent will be at his outside utility on path. The incentive compatibility constraint is static and ratchet effects are absent since the agent does not have and cannot acquire private information about his ability and is hence held at his outside utility even after a deviation.

¹¹An instructive comparison is with Orlov (2022): Due to common discounting and risk-neutrality, the cost of incentives is minimized in his setting when all payments and also all information about performance is delayed until the end of the relationship. Monitoring instead has an efficiency benefit (weeding out bad projects), which drives the main trade-off. In our setting of risk aversion, by contrast, the cost of incentives is minimized when all information about performance is revealed as early as possible and payments are distributed over time, while fully delayed information revelation would be best to avoid the costs associated with learning about match-specific ability. Period-by-period contracting makes this trade-off more tractable, but is not its source.

3 The Contracting Problem in the Terminal Period

We proceed our analysis in the second period to determine the properties of the continuation value which influence the design of the performance evaluation in the first period.

3.1 Performance Evaluation in the Terminal Period

In the second period, the principal has no continuation value from the relationship. Absent any reason to manipulate the agent's learning, the only objective in designing the signal structure is to provide incentives cheaply and there is no reason to leave information about output unused. It is thus optimal to use the most informative signal structure (Grossman and Hart, 1983).

Proposition 1. *The optimal contract in the second period uses the fully informative evaluation structure.*

The principal's continuation value is therefore the solution of the the standard moral hazard problem

$$\begin{aligned} \Pi_2(\mu) = \max_w \int_Y y - w(y) dF(y|1, \mu) \\ \text{s.t. } \int_Y u(w(y)) dF(y|1, \mu) \geq U \\ \int_Y u(w(y)) dF(y|1, \mu) - c \geq \int_Y u(w(y)) dF(y|0, \mu) \end{aligned}$$

3.2 The Shape of the Continuation Value

From the perspective of the first period, the profit in the terminal period induces a continuation value

$$\int_{\Delta\Theta} \Pi_2(\mu) dm(\mu), \tag{2}$$

where $m \in \Delta\Delta\Theta$ is the distribution over posteriors induced by learning from the evaluation in the first-period. The shape of this continuation value is crucial in determining the optimal evaluation structure for the first period. If Π_2 is concave, the agent's learning about his ability imposes a cost on the principal. This cost is larger the more concave the continuation value. In particular, it is more costly to allow agent learning at the top (bottom) if Π_2'' is increasing (decreasing) in μ .

How does information about the agent's ability affects the continuation value? On the one hand, it allows the principal to adapt the contract to the agent's ability. The contract filters out this nuisance parameter more effectively and provides incentives for effort more precisely. As a consequence of this *filtering channel*, the wage can be less risky and it is cheaper to provide incentives. On the other hand, the agent also has more information in the second period when he decides whether to shirk or exert effort. Consequently, the wage has to be more risky on average in order to satisfy the IC constraint posterior-by-posterior and it is more expensive to provide incentives. In other words, it is easier to satisfy

the incentive compatibility constraint in expectation (“ex-ante”) rather than for a more informed agent (“interim”) and this *agent-learning channel* depresses the principal’s profit. The shape of the continuation value and the resulting preferences of the principal over the information structure depend on the balance of these forces.

To gain intuition into the determinants of the shape of this continuation value, it is instructive to consider two special cases.

Square-Root Utility

Suppose that the agent has square-root utility, $u(w) = \sqrt{2w}$. Then, we can solve for the continuation value in closed form. Let $\mathbf{x}(y; \mu) = 1 - \frac{dF(y|0, \mu)}{dF(y|1, \mu)}$ denote the score at y , i.e. a measure of how indicative an output realization is about high as opposed to low effort. The continuation value is

$$\Pi_2(\mu) = \int_Y y dF(y|1, \mu) - \frac{1}{2}U^2 - \frac{1}{2} \frac{c^2}{\int_Y \mathbf{x}(y; \mu)^2 dF(y|1, \mu)} \quad (3)$$

The denominator $\int_Y \mathbf{x}(y; \mu)^2 dF(y|1, \mu)$ is the χ^2 -divergence measuring the distance between the output distributions with low and high effort. Analogously to Fisher-information and the Cramer-Rao bound, $\frac{1}{\int_Y \mathbf{x}(y; \mu)^2 dF(y|1, \mu)}$ is a lower bound on the variance of an unbiased estimator of effort based on output. The cost of providing incentives, therefore, is directly proportional to the variance of this ideal estimator.

Generating more information about θ is desirable if on average it decreases the difficulty of the incentive problem in the second period. In other words, that is if it decreases the average variance of the optimal estimator. There are two forces determining the expected cost. First, with more information the expectation of the χ^2 -divergence increases by the joint convexity of divergences. This *filtering channel* makes the principal prefer more information about θ : She can filter out the noise caused by uncertainty about θ and obtain more precise information about effort and provide cheaper bonuses in expectation. Second, because the agent also learns about his type, the incentive compatibility constraint has to be satisfied at every induced belief. Therefore, the distribution of the χ^2 -divergence matters, not just its expectation. This agent-learning channel makes the principal prefer less information about θ : The cost of incentives is a convex function of the χ^2 -divergence.

If the second channel dominates, the principal faces a trade-off. The costs are higher for providing precise good or bad news depending on whether the continuation value is more concave at the top or at the bottom. That is, where variance and therefore the average informativeness is more responsive to information.

To see how this is determined by the technology, consider as a second example the binary case.

The Binary Case

Suppose that the worker has a binary ability $\theta \in \{\theta_L, \theta_H\}$ and that the resulting output is either high or low, $y \in \{y_L, y_H\}$. The probability of high output depends on the agent's effort and type, $\mathbb{P}(y = y_h | e, \mu) = a + \mu\Delta a + (b + \mu\Delta b)e$. Effort and ability are both productive for both types and levels of effort, $\min\{b, b + \Delta b\} \geq 0$ and $\min\{\Delta a, \Delta a + \Delta b\} \geq 0$. We say that we are in the binary case when these restrictions are satisfied. Then

$$(a\Delta b - b\Delta a)(b\Delta a + \Delta b(1 - a)) > 0 \implies \Pi_2'' < 0 \quad (4)$$

if we are at an interior solution of the problem (which is the case if $U > \frac{a + \mu\Delta a}{b + \mu\Delta b}c$).

In the binary case, the two forces discussed above directly correspond to the parameters Δa and Δb . The direct effect of learning on the continuation value through the filtering channel depends the effect of ability on the probability of high output (Δa). The effect of the agent-learning channel is stronger the larger the effect of ability on the impact of effort (Δb). The agent-learning channel dominates and learning reduces profits whenever the interaction between effort and ability is sufficiently strong. With complementarities ($\Delta b > 0$), we require that are sufficiently strong, i.e. when the technology is log-supermodular in effort and ability ($a\Delta b - b\Delta a > 0$). With substitutes ($\Delta b < 0$), we require that the substitutability is sufficiently strong, i.e. that the technology is log-submodular in effort and ability ($b\Delta a + \Delta b(1 - a) < 0$).¹² Between those two cases of strong interaction, for example at an additive technology, the agent-learning channel is dominated by the filtering channel and the expected cost of providing incentives is reduced by learning. Indeed, the above condition is tight if $u(x) = \sqrt{x}$.

The technology also determines which kind of information is more costly. To see how, consider the second period IC,

$$u(w_H) - u(w_L) \geq \frac{c}{b + \Delta b\mu}, \quad (5)$$

where w_H, w_L denote the optimal wage after y_L and y_H . The impact of effort, $b + \Delta b\mu$, and the bonus (in utility space) required to satisfy incentive compatibility are inversely proportional. As we have seen above, this effect dominates and the continuation value is concave if the interaction is sufficiently strong. Furthermore, this effect of learning is stronger when the agent is pessimistic about the impact of his effort, $b + \Delta b\mu$. Then even a small change in his belief has a large relative effect and causes large changes to the bonus. If effort and ability are log-complements, it is workers with low posterior beliefs who are disproportionately difficult to motivate. They believe that their effort only has a small impact on the likelihood of high output, the principal wants to avoid such beliefs and therefore avoid "information at the bottom". If effort and ability are log-substitutes, it is instead the workers with high posteriors that are disproportionately difficult to motivate.

¹²We say that the technology is log-supermodular, if $\frac{\partial^2}{\partial e \partial \mu} \log(1 - F(y|e, \mu)) \geq 0$ and log-submodular if $\frac{\partial^2}{\partial e \partial \mu} \log(F(y|e, \mu)) \geq 0$, for all $y \in Y$, respectively.

Such workers are willing to rest on their laurels thinking that their effort is unlikely to matter since their high ability is likely to guarantee success in any case. The principal wants to avoid such beliefs and therefore avoid “information at the top”. This effect determines the shape of the continuation value if the curvature of the utility function doesn’t interfere by changing too rapidly. Formally let $v := u^{-1}$ denote the cost of providing utility. The curvature of the continuation value is determined by the technology parameters if the condition

$$\frac{v'''}{v''} \in \left(-\frac{3(b+\mu\Delta b)}{c(b\Delta a+(1-a)\Delta b)} \left(\Delta b + \frac{a\Delta b-b\Delta a}{a+b+\mu(\Delta a+\Delta b)} \right), \frac{3(b+\mu\Delta b)}{c(a\Delta b-b\Delta a)} \left(\Delta b + \frac{b\Delta a+\Delta b(1-a)}{(1-a-b-\mu(\Delta a+\Delta b))} \right) \right) \quad (6)$$

is satisfied at the wages w_H, w_L determined by the constraints. Intuitively, the condition bounds the relative change in the curvature of the cost of providing utility, ensuring that rapid changes in the curvature of utility do not overpower the impact of information.¹³

Let us summarize the discussion.

Proposition 2. *Suppose that we are in the binary case and that the contract is interior ($U > \frac{a+\mu\Delta a}{b+\mu\Delta b}c$). Π_2 is strictly concave if the technology is either log-supermodular or log-submodular.*

Furthermore, if condition (6) is satisfied, $\Pi_2''' > 0$ if the technology is log-supermodular and $\Pi_2''' < 0$ if the technology is log-submodular.

4 The Optimal Evaluation Structure

In this section, we transform the contracting problem of finding the optimal (S, p, w) into the space of posterior beliefs and derive some general results about optimal evaluation structures. We then characterize the optimal evaluation structure.

Throughout, we suppress time subscripts when no confusion can arise.

4.1 Transformation to Posterior Space

For further analysis, it is convenient to rewrite the principal’s problem as the choice of an information structure about output. As is customary in the literature of information design, we can represent such an information structure as a distribution over posterior beliefs over output. Formally, let $\phi \in \Delta Y$ denote a belief about output and $\Phi \in \Delta \Delta Y$ a distribution over such beliefs. Every signal structure (S, p) induces such a distribution satisfying the martingale condition $\mathbb{E}_\Phi \phi = F(\cdot | 1, \mu_0)$ and for every such Φ there exists a signal structure that induces it. We can therefore identify a signal realization s with the associated posterior belief ϕ .

While rewriting the choice of a signal structure as a choice of a distribution over posteriors is standard in the literature on Bayesian persuasion, applying this transformation

¹³For CRRA utility, $u(x) = \frac{x^{1-\gamma}}{1-\gamma}$, we have $\frac{v'''}{v''} = \frac{1}{x} \frac{2\gamma-1}{1-\gamma}$ and the condition is satisfied as long as the outside utility is sufficiently high. This condition is not related to the [Jewitt \(1988\)](#) condition, which instead requires that $v''' \geq 0$ without an upper bound.

to our contracting problem requires two adaptations. First, note that the principal designs an information structure about *output*, while we also need to track beliefs about the agent's *ability*. Second, after a deviation to low effort, the distribution over output and hence signals changes. Incentive compatibility requires the principal to pay the agent a higher wage after signal realizations that are more likely under high effort than under low effort. To rewrite the incentive compatibility constraint we need to track of this shift, namely the score statistic between high and low effort associated to every signal realization.

Both these complications can be dealt with in the same manner: Note that we can easily compute the posterior about the agent's type as well as the score statistic as a function of the true realized output (if it were observed). Formally, fixing a prior belief μ_0 , let $\boldsymbol{\mu} : Y \rightarrow \Delta\Theta$ denote the posterior as a function of output, given by $\boldsymbol{\mu}(y)(\theta) = \mu_0(\theta) \frac{dF(y|1,\theta)}{dF(y|1,\mu_0)}$, and $\boldsymbol{x} : Y \rightarrow \mathbb{R}$ the score statistic of effort, given by $\boldsymbol{x}(y) = 1 - \frac{dF(y|0,\mu_0)}{dF(y|1,\mu_0)}$.¹⁴ The posterior (score) associated with a signal realization s is the expected posterior (score) conditional on this signal, that is, a linear function of the belief over output. Formally, $\boldsymbol{\mu}$ and \boldsymbol{x} extend by linearity to ΔY ; the posterior and score corresponding to a belief over output ϕ are given by $\boldsymbol{\mu}(\phi) = \int \boldsymbol{\mu}(y) d\phi(y)$ and $\boldsymbol{x}(\phi) = \int \boldsymbol{x}(y) d\phi(y)$.

We can then write the principal's problem—using a stochastic wage $w : \Delta Y \rightarrow \Delta\mathbb{R}$ if necessary—as

$$\begin{aligned} \max_{w, \Phi \in \Delta\Delta Y} \int_{\Delta Y} y - w(\phi) + \delta \Pi_2(\boldsymbol{\mu}(\phi)) d\Phi(\phi) \\ \text{s.t. } \int_{\Delta Y} u(w(\phi)) - c d\Phi(\phi) \geq U \\ \int_{\Delta Y} \boldsymbol{x}(\phi) u(w(\phi)) d\Phi(\phi) \geq c \\ \int_{\Delta Y} \phi d\Phi(\phi) = F(\cdot|1, \mu_0) \end{aligned} \quad (7)$$

The rewriting of the objective and of the participation constraint simply amount to a change of variables from s to ϕ . To understand the form of the incentive compatibility constraint, note that the original constraint in (1) can be written as requiring that utility is on average sufficiently higher after signal realizations that are more likely after high effort; equivalently, that there is a sufficient covariance between the score at a certain signal and the associated utility. From this form of the constraint, the above is again just a change of variables, requiring now a sufficient covariance between the score at a certain posterior and the associated utility. Finally, we require that Φ agrees with the true distribution over output on average.

Proposition 3. *An evaluation contract (S, p, w) solves the principal's problem (1) if and only if it induces a (w, Φ) that solves the belief-space problem (7). Furthermore, it is without loss of generality to take the optimal wage to be deterministic in both problems.*

¹⁴Note that the updating rule relies on the presumption that high effort was exerted and is therefore only valid if there is no deviation from the effort proposed in the contract.

4.2 General Properties of Optimal Performance Evaluation and Incentive Contracts

Before proceeding to further simplify and analyze the joint contracting and evaluation design problem, we derive some general properties of the optimal evaluation structure.

In our problem, the benefit from a more informative performance evaluation is that it provides more information about effort; any cost, by contrast, stems from learning about the agent’s ability. It is then intuitively clear that if the evaluation can be made more precise in a way that reveals more information about effort without revealing additional information about ability, the principal would desire to do so. We call a signal realization decomposable if starting at this signal the principal can learn more about effort without revealing additional information about ability.

Definition 1. A posterior ϕ is **decomposable** if there exist $\phi', \phi'' \in \Delta Y$ such that $\phi = \alpha\phi' + (1 - \alpha)\phi''$, $\alpha \in (0, 1)$, $\boldsymbol{\mu}(\phi) = \boldsymbol{\mu}(\phi') = \boldsymbol{\mu}(\phi'')$, and $\boldsymbol{x}(\phi') \neq \boldsymbol{x}(\phi'')$.

We say an evaluation structure Φ is **indecomposable** if it does not put positive mass on decomposable signals.

Indecomposability is a quite stringent geometric condition. The level sets of $\boldsymbol{\mu}$ define hyperplanes in ΔY . When a signal ϕ is in the relative interior of such a hyperplane, it will typically be decomposable. More generally, signals that are supported on more than $|\Theta|$ output realizations will typically be decomposable. Intuitively, as Y becomes rich relative to Θ , there will be many ways to learn about effort without learning about ability. An indecomposable evaluation structure, therefore, must be supported on the boundary of ΔY , and only on small-support edges if $|\Theta|$ is small. In this case, the signal will reveal a lot of information about Y , but it may hide all information about Θ .

Proposition 4. *Any optimal signal structure is indecomposable.*

This result formalizes the above intuition. If there is a “dimension” of output¹⁵ that is informative about effort but not indicative of ability, such as for example the conduct, punctuality and presence of the worker, it will be revealed in the optimal evaluation and become the basis of performance pay.

The proof of the proposition follows directly from the decomposition. If a signal is decomposable, it can be split in a way such that there is more information about the agent’s effort this period, without any additional learning about the agent’s type. Therefore, the principal can incentivize the agent more cheaply this period without changing the continuation value. An optimal structure, therefore, must be indecomposable. The formal proof (p.31) ensures this can be achieved in a measurable way.

4.3 Monotonic Environment

As we have seen, the optimal evaluation when learning about effort and ability can be separated will do so and extract all information about effort conditional on the level of

¹⁵While we don’t formally allow for multidimensional output, this is just for notational simplicity until this point.

learning about the agent’s ability. The evaluation design problem is most interesting, therefore, when effort and ability are inseparably intertwined. Furthermore, it turns out that the problem also becomes more tractable in this case.¹⁶ Instead of identifying an evaluation with a distribution over output and deriving the associated posterior belief and score statistic, we can directly identify an evaluation with the posterior belief it induces. The following proposition characterizes when this identification is without loss of generality with binary types.¹⁷

Proposition 5. *Let $\Theta = \{\theta_L, \theta_H\}$. The following are equivalent*

1. $\mathbf{x} = \alpha_0 + \alpha_1 \boldsymbol{\mu}(\cdot)(\theta_H)$ F a.e. for some $\alpha_0, \alpha_1 \in \mathbb{R}^2$
2. *There is a bijection between scores and posteriors.*

If these conditions are satisfied for all $\mu_0 \in \Delta\Theta$, then $F(\cdot|e, \mu) = F(\cdot|0, 0) + g(e, \mu)\Delta F(\cdot)$ for a bilinear $g : E \times [0, 1] \rightarrow [0, 1]$

The linear structure of both $\boldsymbol{\mu}$ and \mathbf{x} implies that any bijection between the two has to be an affine map. Imposing this restriction for all priors imposes a strong restriction on the distribution over output. It needs to satisfy a linear distribution condition in which effort and ability enter bilinearly. In analogy with the binary case, we parametrize $g(e, \mu) = \mu\Delta a + b \cdot e + \mu\Delta b \cdot e$ without loss of generality. We will focus on the case when the condition of Proposition (5) is satisfied.

Definition. We say that we are in a **monotonic environment** if $|\Theta| = 2$, the distribution over output satisfies $F(\cdot|e, \mu) = F(\cdot|0, 0) + (\mu\Delta a + b \cdot e + \mu\Delta b \cdot e) \Delta F(\cdot)$ with $\int d\Delta F(y) = 0$, $\int y d\Delta F(y) > 0$, $\min\{\Delta a, \Delta a + \Delta b\} > 0$ and $\min\{b, b + \Delta b\} > 0$ and F satisfies **MLRP** (monotone likelihood ration property) i.d. $\frac{d\Delta F(y)}{dF(y|0,0)}$ is increasing in y .

In addition to restricting the distribution over output to the case covered by Proposition (5), the definition restricts the technology to be monotonic in the sense that a higher ability and a higher effort increases expected output: High levels of the aggregator $(\mu\Delta a + b \cdot e + \mu\Delta b \cdot e)$ are associated with higher output since $\int y d\Delta F(y) > 0$. It is ensured that higher ability is beneficial no matter the level of effort ($\min\{\Delta a, \Delta a + \Delta b\} > 0$) and that higher effort is always beneficial no matter the level of ability ($\min\{b, b + \Delta b\} > 0$). The MLRP condition ensures that μ and x are increasing in output, which inessential but makes the results more interpretable.

¹⁶The general problem can be analyzed as an information design problem dependent on the two moments μ, x . Indeed, assuming square-root utility and regularity conditions, the problem can be recast as moment persuasion with a quadratic objective (Dworczak and Kolotilin, 2024, Section 5) in the induced moments $x, \frac{1}{\sqrt{\int x(y;\mu)^2 dF(y|1,\mu)}}$. We instead focus on the monotonic environment, which allows a more concrete and interpretable solution.

¹⁷The restriction to binary ability is mostly for simplicity. In general, there exists a mapping from the posterior to the score if and only if there exists a linear aggregator functional on $\Delta\Theta$ that determines the score. The proposition goes through replacing μ by this aggregator and the distribution satisfies a linear distribution function condition in the aggregator. The arguments of the following section also generalize to this case, noting that we can formulate the information design problem on this summary statistic.

In such an environment, we can further simplify the problem. Let $\bar{m} \in \Delta\Delta\Theta$ denote the distribution over posterior beliefs induced by observing output. Any evaluation structure $\Phi \in \Delta\Delta Y$ induces a distribution over posterior beliefs $m \in \Delta\Delta\Theta$ that is a mean-preserving contraction of \bar{m} , denoted as $m \leq_{MPS} \bar{m}$. Conversely, for every such distribution, there exists an evaluation structure that induces it. By Proposition 5, we can write the score x as a function of the posterior belief, concretely

$$x(\mu) = \frac{1}{\mu_0(1 - \mu_0)} \frac{b + \Delta b \mu_0}{\Delta a + \Delta b} (\mu - \mu_0). \quad (8)$$

A signal is indicative of high effort if and only if and to the degree that it is indicative of high ability. Using this representation, we can rewrite the contracting problem 7 as the choice of a distribution over posteriors about ability. This rewriting simply involves a change of variables, noting that in the original problem, wages are deterministic as a function of the score. We thus get

$$\begin{aligned} & \max_{w, m \in \Delta[0,1]} \mathbb{E}_m [y - w(\mu) + \delta \Pi_2(\mu)] \\ & \text{s.t. } \mathbb{E}_m [u(w(\mu))] - c \geq U \\ & \mathbb{E}_m \left[\frac{1}{\mu_0(1 - \mu_0)} \frac{b + \Delta b \mu_0}{\Delta a + \Delta b} (\mu - \mu_0) u(w(\mu)) \right] \geq c \\ & m \leq_{MPS} \bar{m} \end{aligned} \quad (9)$$

Proposition 6. *Suppose that we are in a monotonic environment. An evaluation contract (S, p, w) solves the principal's problem (1) if and only if it induces a (w, m) that solves the μ -space problem (9). Furthermore, it is without loss of generality to take the optimal wage to be deterministic.*

4.4 The Optimal Evaluation Structure

In the first period, our main trade-off is in effect. Providing incentives for the agent is cheaper in this period if the evaluation structure is more informative, while the resulting learning can be costly by increasing the expected cost of incentivizing the agent in the next period. How is this trade-off resolved in the optimal contract? We employ the tools of information design to characterize the optimal evaluation structure without imposing any exogenous restrictions.

Definition 2. We call an evaluation structure **generalized upper-censorship** if there exists a cutoff y^* such that it reveals output strictly below y^* , it pools output $(y^*, 1]$, and there is a probability $\sigma \in [0, 1]$ such that y^* is revealed with probability $1 - \sigma$ and pooled with the interval $(y^*, 1]$ with probability σ . We call an evaluation structure **generalized lower-censorship** if there exists a cutoff y^* such that it reveals output above y^* , it pools

output $[0, y^*)$, and there is a probability $\sigma \in [0, 1]$ such that y^* is revealed with probability $1 - \sigma$ and pooled with the interval $[0, y^*)$, with probability σ .

This notion generalizes the censorship policies commonly encountered in linear persuasion problems (see [Kolotilin et al., 2022](#), for a discussion) to account for the possibility of atoms in the underlying distribution. In terms of our application, censorship evaluation structures correspond to commonly encountered bonus schemes. Lower-censorship corresponds to a base wage plus fine-grained performance bonuses that are paid only upon reaching a certain performance threshold. Upper-censorship corresponds to a base wage with fine-grained deductions when the performance requirement is not met.

Theorem 1. *Suppose that we are in a monotonic environment. The optimal evaluation structure is*

1. *fully informative if $\Pi_2'' > 0$,*
2. *generalized lower-censorship if $\Pi_2''' > 0$ and $v''' \leq 0$, and*
3. *generalized upper-censorship if $\Pi_2''' < 0$, $v''' \geq 0$ and wages are interior.*

The theorem characterizes the structure of the optimal evaluation for the cases in which the continuation value has an unambiguous shape. If there is no cost from learning ($\Pi_2'' > 0$), a fully informative evaluation is optimal both from the perspective of incentives and from the perspective of learning. If learning is costly ($\Pi_2'' < 0$), the principal desires to introduce noise into the evaluation. The optimal performance evaluation takes the form of partial pooling as determined by the shape of the continuation value. If pessimistic agents are disproportionately expensive to motivate ($\Pi_2''' > 0$), the principal avoids strong pessimism by pooling at the bottom at a base wage and providing precise evaluation and tailored performance bonuses above a threshold. If by contrast optimistic agents are disproportionately expensive to motivate ($\Pi_2''' < 0$), the principal avoids strong optimism by pooling at the top at a capped high-performance wage and providing precise evaluation and tailored low-performance deductions below a threshold. This upper-censorship pattern holds as long as the implicit lower bound on wages that is part of our utility specification (recall that $u : \mathbb{R}_+ \rightarrow \mathbb{R}$) is satisfied. If this constraint is binding for a positive measure of output realizations, this would create another pooling region at the bottom, the per-posterior profit becomes concave-to-convex-to-concave and our proof approach fails.

In general, the optimal evaluation depends both on the shape of the continuation value Π_2 and of the utility function u . The conditions restricting the sign of v''' are sufficient conditions that ensure that the information design problem is governed by the shape of the continuation value and not by changes in the curvature of the utility function. To illustrate the intuition behind the conditions, consider the second case of the theorem: Since the continuation value is less concave for high posteriors, the principal prefers revealing information about high output while shrouding information about low output. This means that incentives will be provided and payments are risky for high output, which goes together

with high wage payments. If the agent's risk aversion and hence the cost of providing risky payments were much larger in this range than for lower utility levels, it might not be optimal to provide incentives at the top. The condition ensures that this is not the case.

In the optimal contract, wages are discontinuous when viewed as a function of the true performance of the agent. Under lower-censorship, for example, the bonus the agent receives when his performance just clears the threshold of the pooling region can be substantial. This discontinuity of the bonus follows naturally from the informational considerations driving our result, as the pooled evaluation is strictly less indicative of high effort than any other evaluation outcome.

The optimal evaluation structure and wage contract are unique up to deviations of measure zero and informationally equivalent modifications of the evaluation structure. For example, in the case of lower censorship with optimal cutoff y^* , the principal could split the pooling interval in an uninformative way, revealing whether output is in $[\underline{y}, y_1^*) \cup [y_2^*, y^*)$ or in $[y_1^*, y_2^*)$ where y_1^*, y_2^* are chosen such that the posterior after observing the two evaluations is the same as after observing the full pooling set $[\underline{y}, y^*)$.

In the binary case, we can characterize the shape of the optimal evaluation structure based on the primitives of the technology using Proposition 2.

Corollary 1. *Suppose that we are in the binary case and that u satisfies condition (6). Then,*

1. *if the technology is log-supermodular ($a\Delta b - b\Delta a > 0$) and $v''' \leq 0$, the optimal evaluation is binary and tough, with $S = \{G, B\}$ and*

$$p(G|y_H) = 1 - \sigma, \quad p(B|y_H) = \sigma, \quad p(G|y_L) = 0 \quad p(B|y_L) = 1, \quad (10)$$

for $\sigma \in [0, 1)$;

2. *if the technology is log-submodular ($b\Delta a + (1 - a)\Delta b < 0$) and $v''' \geq 0$, the optimal evaluation is binary and lenient, with $S = \{G, B\}$ and*

$$p(G|y_H) = 1, \quad p(B|y_H) = 0, \quad p(G|y_L) = \sigma \quad p(B|y_L) = 1 - \sigma, \quad (11)$$

for $\sigma \in [0, 1)$.

First, the motive to control learning does not increase the complexity of the evaluation structure. While the most informative evaluation is binary, a noisy evaluation can take many forms. The theorem implies that the optimal evaluation remains binary. Note that this result is not an immediate consequence of the assumption binary output or binary types. Since the contracting problem has two constraints – participation and incentive compatibility – results from constrained information design suggest that the optimal evaluation structure may involve up to four signals (Le Treust and Tomala, 2019; Doval and Skreta, 2023) even in the binary case. The joint design of wages and information is crucial, this result may not hold when the wage function is fixed exogenously.

Second, the binary structure allows a different interpretation of upper- and lower-censorship evaluation structures. With strong complementarities, the evaluation appears tough: Even high output does not ensure a good evaluation while low output always results in a bad one. The principal does not engage in “grade inflation”, but instead measures performance against an “unreasonably” high standard. With strong substitutes instead, the evaluation appears lenient: Even low output is sometimes rewarded with a good evaluation. To avoid costly complacency, the principal only rarely notices the shortcomings of the agents output, which in turn implies that not getting called out is not very informative about high ability.

Sketch of the Proof of Theorem 1

The proof of Theorem 1 poses the challenge of jointly designing an information structure and a wage scheme. Given a wage scheme, the information design problem can in principle be solved by the usual concavification or linear programming approaches (Aumann and Maschler, 1995; Kamenica and Gentzkow, 2011; Kolotilin, 2018; Dworzak and Martini, 2019) taking into account the participation and incentive compatibility constraints (Boleslavsky and Kim, 2017; Le Treust and Tomala, 2019). The constraints make the problem multidimensional so that, although in some cases conceptually tractable, a direct approach is analytically difficult (Doval and Skreta, 2023). Conversely, given an information structure, the problem of finding wages is a standard moral hazard problem. This tractable problem provides the starting point for a duality-based approach to such a joint information and incentive design problem (Georgiadis and Szentes, 2020). In the main text, we sketch the main steps of the argument, while we relegate the explicit duality arguments that are required to justify the existence of a solution to the appendix.

Consider the Lagrangian \mathcal{L} associated to the contracting problem (9), where we retain the condition that $m \leq_{MPS} \bar{m}$ as a constraint and $\lambda = (\lambda_P, \lambda_{IC})$ denote the Lagrange multipliers associated to the participation and incentive constraint, respectively,

$$\begin{aligned} \mathcal{L}(m, w; (\lambda_P, \lambda_{IC})) = & \int \left\{ \delta \Pi_2(\mu) - w(\mu) \right. \\ & + \lambda_P (u(w(\mu)) - c - U) \\ & \left. + \lambda_{IC} \left(\frac{1}{\mu_0(1 - \mu_0)} \frac{b + \Delta b \mu_0}{\Delta a + \Delta b} (\mu - \mu_0) u(w(\mu)) - c \right) \right\} dm(\mu). \end{aligned} \quad (12)$$

We have omitted expected output, since it is a constant in the principal’s problem.

Wage Setting Fix a feasible distribution $m \in \Delta \Delta \Theta$ and consider the problem of finding optimal wages subject to the participation and incentive constraint. This is a standard moral hazard problem and the optimal wage schedule follows from pointwise optimization of the Lagrangian.

$$w^*(\lambda, \mu) = \max\{0, u'^{-1}\left(\left(\lambda_P + \lambda_{IC} \frac{1}{\mu_0(1-\mu_0)} \frac{b + \Delta b \mu_0}{\Delta a + \Delta b} (\mu - \mu_0)\right)^{-1}\right)\} \quad (13)$$

Plugging this function back into the Lagrangian of the problem, it is written as an expectation of a function of the posterior

$$\sup_w \mathcal{L}(m, w; \lambda) = \int \ell^*(\mu; \lambda) dm(\mu) \quad (14)$$

where ℓ^* is the integrand of (12) evaluated at (13).

Information Design Therefore, the information design problem *for a given* λ is of standard form. The principal simply maximizes the expectation of a function of posteriors,

$$\sup_{m \leq_{MPS} \bar{m}} \int \ell^*(\mu; \lambda) dm(\mu). \quad (15)$$

subject to the constraint that the chosen distribution m is a mean-preserving contraction of the distribution \bar{m} which would be induced by the fully revealing performance evaluation. Such an information design problem can be solved via the usual linear programming approaches (Kolotilin, 2018; Dworzak and Martini, 2019; Kolotilin et al., 2022). In order to determine the solution, we need to determine the shape of ℓ^* as a function of μ . Using an envelope argument, it is straightforward to show¹⁸ that

$$\begin{aligned} \frac{\partial^2}{\partial \mu^2} \ell^*(\mu; \lambda) = & \lambda_{IC}^2 \left[\frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \right]^2 \rho' \left(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) \right) \\ & + \delta \Pi_2''(\mu) \end{aligned} \quad (16)$$

where $\rho(x) := u(u'^{-1}(\frac{1}{x}))$ denotes the function that translates multipliers and scores to utilities, a function commonly encountered in moral hazard problems. The first term of (16) corresponds to the cost of providing incentives in the first period. It is positive, indicating convexity: the principal prefers the most informative evaluation structure in order to reduce agency costs. The second term corresponds to the impact of beliefs on the continuation value. If the continuation value is also convex, then so is the per-posterior payoff overall, and a fully informative evaluation structure is optimal. If it is concave, there is a trade-off: the principal wants to keep the agent uninformed in order to reduce agency costs in the next period.

To obtain the optimal evaluation, we then need to further consider the shape of ℓ^* . We have that

$$\begin{aligned} \frac{\partial^3}{\partial \mu^3} \ell^*(\mu; \lambda) = & \lambda_{IC}^3 \left[\frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \right]^3 \rho'' \left(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) \right) \\ & + \delta \Pi_2'''(\mu). \end{aligned} \quad (17)$$

¹⁸In the main text, we suppress boundary conditions related to the non-negativity constraint on wages.

The third derivative of ℓ^* has two components. The second term is determined by the shape of the continuation value. The first term is the impact of the shape of the utility function. For given Lagrange multipliers, it is cheaper to provide incentives where ρ' is small, or, equivalently, where the curvature of v is low. Hence, the principal prefers dispersion at the top (bottom) if $v''' \leq 0$ ($v''' \geq 0$) or, equivalently, $\rho'' \geq 0$ ($\rho'' \leq 0$).¹⁹ This structure also reveals why Theorem 1 relies on bounds on v''' as a sufficient condition: The proof establishes that the optimal evaluation structure has the required shape for any λ , while the result only relies on the shape at the λ that solves the Lagrangian problem.

To solve the information design problem (15), we generalize the result from Kolotilin et al. (2022) to accommodate a distribution that may not admit a density and may not have full support. This required for our application, since the full-information distribution \bar{m} may have atoms and gaps in its support, e.g. in the binary case.

Theorem 2. *Suppose $V''' > (<)0$. Then, generalized lower- (upper-)censorship is the essentially unique solution to*

$$\begin{aligned} \max \int_0^1 V(s) dH(s) \\ \text{s.t. } H \leq_{MPS} F \end{aligned}$$

The proof of the theorem closely follows Kolotilin et al. (2022), generalizing it by working in quantile space (to deal with atoms) and dealing with the nondifferentiability that arises if there is a gap in the support.




Figure 1: The optimal evaluation structure.

In the relevant case with $\Pi_2'' < 0$ there are three possibilities. If λ_{IC} is sufficiently small, the objective ℓ^* is strictly concave and the optimal information structure is uninformative. Clearly, this cannot be the case in the solution of (9), since the incentive constraint cannot be satisfied without any information. As λ_{IC} increases, we reach a region where, depending

¹⁹To see this, note that $v''(u(x)) = -\frac{u''(x)}{[u'(x)]^3}$ and that $\rho'(x) = -\frac{[u'(f(x))]^3}{u''(f(x))}$ for the strictly increasing $f(x) = u^{-1}(\frac{1}{x})$. Hence, the former is decreasing if and only if the latter is increasing ($\rho'' \geq 0$ as required).

on the sign of v''' and Π_2''' , ℓ^* is concave for low posteriors and convex for high posteriors or vice versa. The optimal information structure is fully informative where ℓ^* is convex and uses partial pooling at the other side, leading to a point mass in belief space (Fig. 1a). Finally, as λ_{IC} increases further, ℓ^* becomes globally convex as the costs of incentives overwhelm the gains from concealing information. The resulting evaluation structure becomes fully informative (Fig. 1b).

The arguments in the appendix establish through a series of lemmas that a solution to the problem exists and is characterized by the two-step procedure above. QED.

4.5 Analysis of the Solution and Comparative Statics

The properties of an interior solution are pinned down by tangency condition (assuming that \bar{m} admits a non-zero density). The censorship cutoff in belief space, μ^* , and the posterior after the pooling evaluation, $\bar{\mu}$, are determined by

$$\ell^*(\mu^*, \lambda(\mu^*)) = \ell(\bar{\mu}, \lambda(\mu^*)) + \frac{\partial \ell^*}{\partial \mu} \Big|_{(\mu, \lambda) = (\bar{\mu}, \lambda(\mu^*))} (\mu^* - \bar{\mu}) \quad (18)$$

(Fig. 1). Note, however, that this condition does not correspond to information design on a given function. Instead, there is an additional dependence on $\lambda(\mu^*)$. This term is present because we design payoffs and information jointly, subject to a participation and incentive compatibility constraint. For the graphical representation of our analysis this implies that, as we vary the cutoff point in Figure 1 to find the optimal μ^* , not only the tangent line but the whole function ℓ^* shifts.

Under the assumption that $u(x) = \sqrt{2x}$ we can eliminate the multipliers and transform (18) into a more concrete form (assuming upper-censorship, $\Pi_2''' < 0$):

$$\frac{1}{2} \frac{c^2}{\left(\frac{b + \Delta b \mu}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \right)^2} \frac{(\mu^* - \bar{\mu})^2}{\left(\int_0^{\mu^*} (\mu - \mu_0)^2 d\bar{m}(\mu) + (1 - m(\mu^*)) (\bar{\mu} - \mu_0)^2 \right)^2} = \quad (19)$$

$$= \delta [\Pi_2(\mu^*) - \Pi_2(\bar{\mu}) + \Pi_2'(\bar{\mu})(\bar{\mu} - \mu^*)] \quad (20)$$

The LHS is the benefit from a more informative evaluation structure in period one. A more precise signal about output decreases agency costs today. This effect is larger if agency costs $\left(\frac{c}{b + \Delta b \mu} \right)$ are already high and if a large dispersion of posteriors is required for a given level of information about output (since output is very informative: $(\Delta a + \Delta b) \mu_0 (1 - \mu_0)$ large). The RHS is the cost of a more informative information structure through learning. A more precise signal today allows learning and thereby may increase average agency costs in the next period. Indeed, the RHS is a measure of the concavity of the continuation value.

Proposition 7. *Suppose that learning is costly ($\Pi_2'' < 0$) and that \bar{m} admits a strictly positive density. The size of the pooling region is*

1. *increasing in the discount factor δ and*

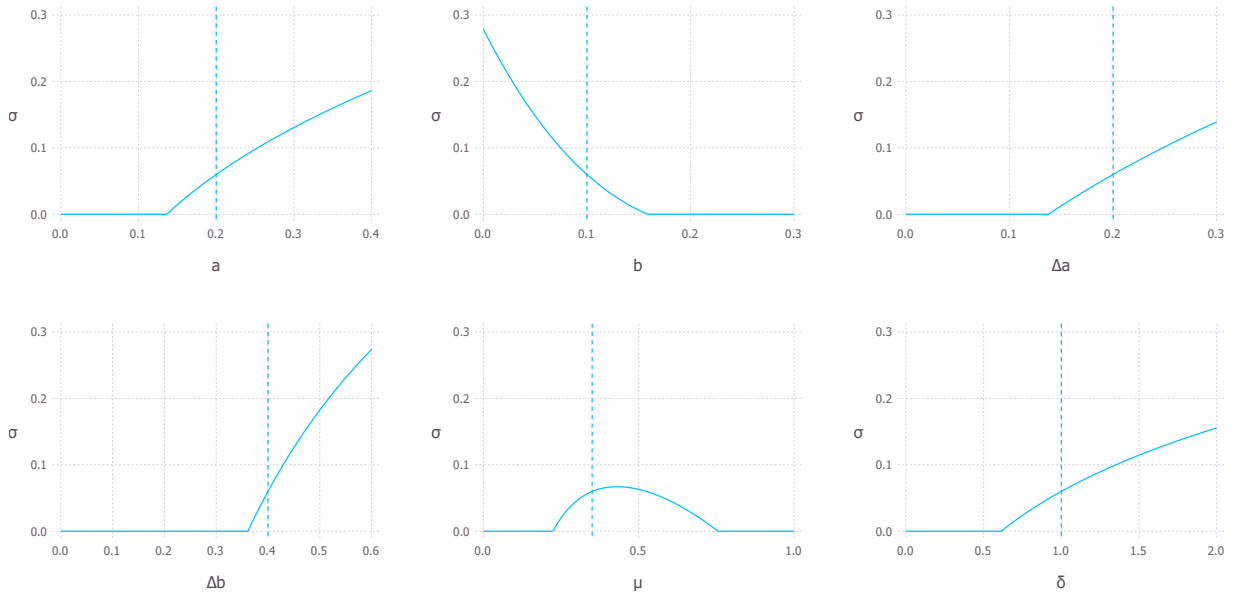


Figure 2: Comparative statics in the binary case (Corollary 1) with $u(x) = \sqrt{2x}$ around $(a, b, \Delta a, \Delta b) = (0.2, 0.1, 0.2, 0.4)$, $\mu = 0.35$, $\delta = 1$.

2. for $u(x) = \sqrt{2x}$, decreasing in the costs of effort in the first period, increasing in the costs of effort in the second period and independent of a common proportional increase in the cost of effort.

Both comparative statics illustrate the trade-off between the cost of incentives in the first and second period. As the second period becomes more important, the evaluation structure becomes less informative. Higher costs of effort in the first period make economizing on agency costs in that period more important, thus the evaluation structure becomes more informative. Numerical computations in the binary case with complementarities (Figure 2) show that the noise in the evaluation is highest when there is a lot of uncertainty about the agent's type (intermediate priors), when the impact of ability on output is high (both directly, Δa , and through the impact of effort, Δb), and when high ability is relatively more essential for the impact of effort (low Δb) and less essential for the baseline level of output (high a).

5 Discussion and Extensions

In the previous sections, we made several assumptions to ensure that the agent's posterior belief is the only state variable of the problem and that no party can acquire endogenous private information. In this section, we discuss how these assumptions can be relaxed. Finally, we comment on the problem with more than two periods.

5.1 Private Information Acquisition

In some settings, it may be possible for the firm to privately observe additional information about the worker’s output without disclosing it or using it as a basis of wages in the first period. In this section, we argue that the principal would like to commit not to use additional private information. We then sketch two settings in which she can achieve this commitment indirectly, namely when the private information structure is chosen and observed; and when the space of contracts is not too rich. In those cases there exist natural equilibria that replicate the optimal contract.

When the principal can acquire private information, the contracting problem now induces a dynamic game with incomplete information. In the first period, the principal now also designs a *private* evaluation structure. Its realizations are not observed by the agent. Writing the problem in belief space, the principal designs a feasible joint distribution of agent and principal posteriors $m_P(\mu_P, \mu_A)$.²⁰ The marginal on the agent’s posterior, $m(\mu_A) = \int m_P(\mu_P, \mu_A) d\mu_P$, is observed by the agent.

A perfect Bayesian equilibrium consists of (1) a feasible evaluation structure m_P , (2) a wage function $w : \mu_A \rightarrow w(\mu_A) \in \mathbb{R}_+$, (3) a first-period strategy of the worker mapping the evaluation and wage scheme to participation and effort choices, $m, w \rightarrow \{0, 1\}^2$ ($m_P, w \rightarrow \{0, 1\}^2$ if the information structure is observable), (4) a second period contract offer, $(\mu_P, \mu_A) \rightarrow w_{\mu_P, \mu_A} \in \mathbb{R}_+^Y$, and (5) a belief system for the agent over his type and (in the unobserved information case) the information structure chosen by the principal, as a function of the posterior and the contract offer, $(\mu_A, w_{\mu_P, \mu_A}) \rightarrow \Delta([0, 1] \times \Delta[0, 1]^2)$, satisfying sequential rationality and consistency.²¹ We say that a PBE satisfies *no-holdup* if the agent’s participation constraint is binding in all on-path second period contract offers.²² A PBE is said to have *passive beliefs* if the second period belief of the agent is independent of the contract offer and equal to the posterior μ induced by the first period signal.²³

Remark 1. The principal’s profit in any no-holdup PBE is weakly smaller than in the optimal contract solving the contracting problem (1).

In other words, the principal would like to commit not to use additional private information in designing the contract in the second period. The intuition behind this result is simple. Any information used and thereby revealed in equilibrium through the contract

²⁰To the best of our knowledge, a fundamental characterization of feasible joint beliefs is not known. In a symmetric setting without a mean-preserving contraction condition, see [Arieli et al. \(2021\)](#).

²¹For simplicity of exposition, we assume that the principal uses a fully informative evaluation in the second period, as shown to be optimal in [Proposition 1](#). This restriction is without loss on path, as a fully informative evaluation structure remains optimal for the principal. Off-path, the restriction reduces the degrees of freedom for deviations, but the equilibrium we study can be extended naturally.

²²Without such a refinement, the equilibrium could grant intertemporal commitment. This would allow the principal to smooth out bonus payments across periods, yielding higher profits through a channel orthogonal to the acquisition of private information.

²³Common refinements for signaling games, such as the intuitive criterion ([Cho and Kreps, 1987](#)) or D1 ([Cho and Sobel, 1990](#)), do not apply as they require the set of types of the principal to be fixed. See also [Ekmecki and Kos \(2023\)](#) for an application of NWBR when the sender chooses whether to acquire full information about his type or not.

offer in the second stage could have been revealed and used to provide incentives already at the first stage, which would have increased profits then.

When can the principal achieve this bound? First, suppose that the choice of the private information structure is observed by the agent. Then, the principal can achieve the profit in the optimal contract by choosing not to acquire additional information. Indeed, she can deviate to this outcome under a natural restriction on out-of-equilibrium beliefs.

Remark 2. When the information structure is observed, any equilibrium²⁴ is outcome equivalent to the optimal contract characterized in Theorem 1.

Second, suppose that the principal's choice of her private information structure is private. Then she can deviate by privately acquiring perfect information about output and using this information to offer more profitable contracts in the second stage. We now show that under a condition on the contracting environment, a natural restriction on out of equilibrium beliefs can provide sufficient commitment to the principal to avoid this deviation. We call the second-period contract *agent-determined* if the optimal contract for an agent of posterior μ subject to the participation and incentive compatibility constraint is independent of the principal's belief μ_P for all μ_P in the support of \bar{m} . This is a stringent condition, essentially requiring that the number of active (participation and incentive compatibility) constraints is greater or equal to the number of output realizations. It is satisfied in the binary case as long as the amount of private information of the principal is not too large (see Dumav et al., 2021, Observation 4.1).

Remark 3. The (essentially unique) equilibrium with passive beliefs is outcome-equivalent to the optimal contract characterized in Theorem 1 if the second-period contract is agent-determined. This equilibrium is principal preferred among all no-holdup PBE of the game.

To see why this is the case, note that since the optimal contract is agent-determined, the private information of the principal is of no use facing an agent with passive belief μ , and the continuation value induced on the first period is the same as in Propositions 1 and 2. Consequently, the principal's choice of m_P is equivalent to the first-period problem. We provide further details of the arguments in Appendix A.1.

5.2 Unobservable Effort

In the main sections, we assume that effort is observed but not contractible. This ensures that even after a deviation, the principal and the agent share a common belief over the agent's type. Assume instead that effort is not observed by the principal. This does not affect beliefs on equilibrium path, since the conjectured effort is correct. After a deviation to $e_1 = 0$, however, the agent updates his beliefs according to

$$\tilde{\mu}(\phi)(\theta) = \int \mu_0(\theta) \frac{dF(y|0, \theta)}{dF(y|0, \mu_0)} d\phi(y), \quad (21)$$

²⁴Among no-holdup PBE that satisfy the following natural restriction, a form of no-signaling-what-you-don't-know: After observing the information structure m_P and signal μ_A , the agent's belief is supported on the convex hull of the support of $m_P(\cdot, \mu_A)$.

while the principal continues to use the on-path updating rule μ . Hence, depending on the signal realization, the agent will be less (resp. more) optimistic about his type in the second period and the contract offered by the principal will violate (resp. over-satisfy) the participation and incentive compatibility constraint.²⁵ A deviation in the first period is more profitable for the agent because of this belief-manipulation effect.²⁶ We now sketch this model, assuming that we are in the monotonic case.

Note that the problem in the second period is unchanged. In the first period, we need to modify the incentive-compatibility constraint because of the belief-manipulation effect. The incentives to work are unchanged in the current period, but the constraint becomes dynamic. Let $w_\mu(y)$ denote the optimal wage in the second period problem with common belief μ . The first period IC then reads

$$\begin{aligned} & \int x(\phi)u(w(\phi)) d\Phi(\phi) - c + \delta U \\ & \geq \delta \int \max\left\{ \underbrace{U}_{(i)}, \underbrace{\int u(w_{\mu(\phi)}(y)) dF(1, \tilde{\mu}(\phi)) - c}_{(ii)}, \underbrace{\int u(w_{\mu(\phi)}(y)) dF(0, \tilde{\mu}(\phi))}_{(iii)} \right\} d\Phi(\phi). \end{aligned} \quad (22)$$

If the agent does not deviate, he will obtain his reservation utility δU in the final period. If effort were observable, this would also be the case after a deviation, so this term would cancel and we obtain our familiar constraint. Since effort is not observable, the agent has private information about his type after a deviation and has a nontrivial choice in the second period between: (i) not participating, (ii) exerting effort, and (iii) shirking. The former is optimal if he is more pessimistic after the deviation ($\tilde{\mu}(\phi) < \mu(\phi)$): The wage $w_\mu(y)$ is monotonic in output and just calibrated to ensure that an agent with belief $\mu(\phi)$ is willing to participate. If the agent is more optimistic after a deviation ($\tilde{\mu}(\phi) > \mu(\phi)$), he participates. He exerts effort in the second period and IC is over-satisfied if effort and ability are complements ($\Delta b > 0$) and shirks if they are substitutes ($\Delta b < 0$).

We can translate this dynamic IC into a generalized belief space and write it as

$$\int \frac{(b+\mu_0\Delta b)}{\mu_0(1-\mu_0)\Delta b} (\mu - \mu_0) u(w(\mu)) - \delta \max\left\{0, c \frac{\Delta b}{b+\mu\Delta b} (\tilde{\mu} - \mu), c \frac{\Delta a + \Delta b}{b+\mu\Delta b} (\tilde{\mu} - \mu)\right\} dm(\mu, \tilde{\mu}) \geq c \quad (23)$$

Even in the linear case, however, we cannot write the deviation posterior $\tilde{\mu}$ as a function of the on-path posterior μ . The problem therefore becomes multidimensional, with a kink in the Lagrangian at $\mu = \tilde{\mu}$, the analysis of which is left for future research.

Our analysis can be extended when we restrict attention to the log-supermodular binary case and assume that $\Delta a = 0$. This condition keeps the problem one-dimensional

²⁵This assumes that the principal does not elicit the agent's belief at the beginning of the second period. In such a mechanism, however, truth-telling would need to be preferable to imitating the type that realizes on path. Hence, a screening mechanism in the second period cannot reduce the agent's post-deviation payoff and therefore does not affect the optimal contract.

²⁶This effect is central in the analysis of many models of moral hazard with learning, e.g. [Prat and Jovanovic \(2014\)](#); [Demarzo and Sannikov \(2017\)](#). [Bhaskar and Mailath \(2019\)](#) show that this motive implies that the costs of providing incentives using spot contracts grows unboundedly with the length of the time horizon in a model similar to ours, but with learning from output. It is unknown whether the design of the information structure can reverse this conclusion.

since the agent does not learn after a deviation, which therefore pins down $\tilde{\mu} \equiv \mu_0$. The optimal contract then has the following structure: With some probability, the evaluation is uninformative. Conditional on receiving an informative evaluation, the tough binary structure we found in the main section is preserved. The details of this analysis can be found in Appendix A.2.

5.3 Long-Run Commitment

In the main sections, we assumed that the principal does not have commitment across periods. This is not crucial for our results, but was assumed because of tractability. What is crucial, however, is that providing incentives goes hand in hand with providing information about the agents performance.

The principal could sever this link in a specific setting with full commitment. Suppose that the principal can commit to wages that depend on output in both periods and are revealed and paid only at the end of the employment relationship. Furthermore, suppose that doing so does not make risk-sharing between the principal and the agent less efficient, e.g. because the agent only consumes at the end or because the agent is risk neutral. Then, the principal can costlessly postpone the bonus payment, learning is no longer a concern and hence using a fully informative evaluation is optimal. Any feature of the model, however, that makes it costly or impossible to postpone the bonus reinstates the trade-off between learning and incentives analyzed in this paper. Suppose for instance that the agent is less patient than the principal. In the extreme case of a myopic agent, only the current concurrent payments of the principal matter and the problem is equivalent to period-by-period contracting. Our theorem 1 extends verbatim to this case. Risk aversion itself already implies that it is costly to delay informative incentive payments, as it would be optimal to smooth them across both periods.

When adding intertemporal commitment to the principal in our setting, we have to distinguish two cases. Our results continue to hold if the principal can postpone payments by committing to a continuation value. For a proof of this result, see Appendix A.3.

Remark 4. Suppose that $u(w) = \sqrt{2w}$ and that in the first period, the principal can commit to signal-contingent wages and continuation values. The optimal evaluation is as characterized in Theorem 1.

The case with full commitment instead raises considerable difficulties. This is because of the interaction between full commitment and the belief-manipulation problem discussed in the previous subsection. The full commitment contract cannot condition on the true effort exerted in the first period, as this would resolve the moral hazard problem. Therefore, the full commitment contract has to deal with the dynamic constraint (22). As a result, the principal may find it optimal to commit to excessive bonus payments in the final period to relax this constraint by inducing a learning motive in the agent. To analyze this problem,

the contracts in both periods need to be designed jointly with the information structure, which is intractable.²⁷

5.4 Many Periods: Simulations

In the main sections, we analyze the twice-repeated problem. Consider now the same period-by-period problem, but repeated for $T > 2$ periods. Characterizing the solution of this problem is not tractable. Numerical analysis of the binary case, however, indicates that the results from the two-period problem generalize: The optimal evaluation under log-supermodularity is binary with partial pooling at the bottom in all periods. The longer the remaining time horizon, the more important the dynamic learning channel. Therefore, the optimal evaluation is less precise in early periods and becomes more precise over time (Figure 3). The worker is left in the dark through a noisy and tough (under log-supermodularity) evaluation early in the relationship when additional information affects many future incentive compatibility constraints. There is significant uncertainty under the optimal evaluation even after ten periods, with fewer very low and more moderately high posterior beliefs compared to a fully informative one (Figure 4b).




Figure 3: The probability of a false negative evaluation decreases over time for any prior.

6 Concluding Remarks

Our model demonstrates why it can be in a principal’s interest to base incentives on a noisy evaluation of the agent’s performance, even when the principal could measure true output and commit to contingent wages. The underlying insight is that output contains information both about effort, which she wants to ascertain and incentivize, and the agent’s match-specific ability, which she would like to keep shrouded.

²⁷One possible work-around is to consider the problem with full commitment when the contract terms in the second period can condition on true effort in the first period, the expected utility of the agent, however, is restricted to be independent of this information conditional on the evaluations. This problem is akin to commitment to a continuation value and our results continue to hold as in Remark 4.

(a) The evolution of beliefs under the dynamically optimal evaluation. (b) The CDF at $T = 10$ (posteriors on the y -axis).




Figure 4: Belief dynamics in the Problem with $T = 10$ (parameters as in Fig. 2)

The optimal resolution of the trade-off between incentives and learning pools evaluations to avoid creating agents who are expensive to motivate. This may involve a fixed wage with tailored deductions for low performance to avoid revealing top performance, which would cause complacency in the future, or a base wage with bonuses to avoid revealing detailed information about low performance which is demotivating. When the underlying measure of output is coarse, the optimal evaluation randomizes. When effort and ability are log-complements, it is tough: Good performance is not always recognized, but bad performance is always punished. Again, the reason is learning, as a tough evaluation ensures that even after a bad evaluation, the agent is not too pessimistic about his type.

In order to implement the optimal contract, the principal requires commitment. After effort is exerted, she otherwise prefers to pay low wages while trying to convince the agent that his effort is essential for his success to reduce agency costs in the next period, all independently of the agent’s performance. That the principal requires commitment to the wage and bonus scheme is standard in moral hazard problems. The design of the evaluation structure shares the commitment assumption with information design. In our setting, the principal often can commit to performance indicators in the employment contract. More broadly, she designs the organizational structure, the incentives of managers and their evaluation guidelines. Persistent practices and internal rules determine how data that is collected incidentally in the course of business operations, be it at the workplace or from the consumer facing side of the firm, is aggregated and linked back to an individual worker. The principal can also leverage reputation across time (Best and Quigley, 2024) and across the workers in the organization (Lin and Liu, 2024). A further way a firm can commit to the style of evaluation is through the selection and training of evaluators. Interpreted from this angle, our results show for example that in environments with strong complementarities between effort and ability unreasonably strict supervisors and drill-sergeant mentality may part of the optimal organization design.

Our results inform not only the optimal evaluation of employee performance, but are also suggestive of the selection of information sources. Among measures that combine

information about effort and ability, the principal prefers measure that are less sensitive to ability. This may motivate secrecy about salary differences between employees in the same job as opposed to differences in bonuses, as base-pay reflects the principal's estimate of ability while bonuses more directly reward effort. Monitoring effort itself and the conduct of employees more generally instead of output is highly desirable. Both patterns are in sharp contrast with models of implicit incentives through career concerns (Holmström, 1999). In such models, the fact that a signal combines information about effort and ability is not a friction but the source of incentives, as the agent exerts effort to avoid being perceived as low-ability. To achieve the right mixture of information about effort and ability the principal might add noise (Dewatripont et al., 1999) or combine information from different sources or vintages (Hörner and Lambert, 2021).²⁸ In our setting with explicit incentives, by contrast, such entangled information is the source of the friction. The analysis of evaluation design when both explicit and implicit incentives are present is an interesting avenue for future research.

²⁸Rodina (2018) provides conditions for additional information about output and less prior information about ability to increase effort.

A Proofs

Proof of Proposition 1: Full information is Blackwell more informative than any other information structure. Then, the result follows from Proposition 13 in [Grossman and Hart \(1983\)](#). Since both the Blackwell comparison as well as the concavity of the utility function are strict, uniqueness follows from their proof (up to pooling of informationally equivalent output realizations, if they exist, and modifications of measure zero). \square

Proof of Proposition 2: For the binary case, we denote the probability of good output as $P_\mu := a + b + (\Delta a + \Delta b)\mu$. By standard arguments, both the participation and the incentive constraint are binding. Hence, solving the problem in utility space, let $v := u^{-1}$, then

$$\Pi_2(\mu) = P_\mu Y - P_\mu v(u_H) - (1 - P_\mu)v(u_L),$$

where $u_L = U - \frac{a+\mu\Delta a}{b+\mu\Delta b}c$ and $u_H = U + \frac{1-a-\mu\Delta a}{b+\mu\Delta b}c$. Note that we require $U - \frac{a+\mu\Delta a}{b+\mu\Delta b}c > 0$ to satisfy the implicit nonnegativity constraint in the agent's utility function. It is easy to verify that at an interior solution

$$\begin{aligned} \Pi_2''(\mu) &\propto 2(b\Delta a - a\Delta b)(b\Delta a + \Delta b(1 - a))(b + \Delta b\mu) (v'(u_H) - v'(u_L)) \\ &\quad - cP_\mu (b\Delta a + \Delta b(1 - a))^2 v''(u_H) \\ &\quad - c(1 - P_\mu) (b\Delta a - a\Delta b)^2 v''(u_L) \end{aligned}$$

The two latter terms are negative, and so is the first, since $b\Delta a - a\Delta b < 0$ or $b\Delta a + \Delta b(1 - a) < 0$.

For the third derivative of the continuation value, it is elementary but tedious to show that

$$\begin{aligned} \Pi_2'''(\mu) &= \frac{c}{(b + \mu\Delta b)^6} \left[6\Delta b (b\Delta a + (1 - a)\Delta b) (a\Delta b - b\Delta a) (b + \mu\Delta b)^2 (v'(u_H) - v'(u_L)) \right. \\ &\quad + 3c (a\Delta b - b\Delta a)^2 (b + \mu\Delta b) (b\Delta a + \Delta b(2 - 2a - b - \mu(\Delta a + \Delta b))) v''(u_L) \\ &\quad + 3c (b\Delta a + (1 - a)\Delta b)^2 (b + \mu\Delta b) (a\Delta b - b\Delta a + \Delta b(a + b + \mu(\Delta a + \Delta b))) v''(u_H) \\ &\quad - c^2 (a\Delta b - b\Delta a)^3 (1 - a - b - \mu(\Delta a + \Delta b)) v'''(u_L) \\ &\quad \left. + c^2 (b\Delta a + (1 - a)\Delta b)^3 (a + b + \mu(\Delta a + \Delta b)) v'''(u_H) \right]. \end{aligned}$$

The condition on $\frac{v'''}{v''}$ implies that the higher order terms are dominated and that the sign is determined by Δb , with the desired property. \square

Proof of Proposition 3: The identification between signals and posteriors is standard and the equivalence of the constraints is an immediate rewriting following the discussion in the text. It remains to show that restricting to deterministic wages is without loss in both problems. In either problem, suppose that an optimal solution involves a stochastic wage. Replacing this wage function with the one offering the certainty equivalent conditional on each s , respectively ϕ , keeps the constraints unchanged. Since the agent is risk averse, this deterministic wage is pointwise weakly below the conditional expectation of the original wage conditional on s (resp. ϕ), weakly increasing profits. \square

Proof of Proposition 4: Suppose towards a contradiction that Φ is optimal but that it is decomposable. It follows from Theorem 2.1 and Theorem 3.1 [Winkler \(1988\)](#) that for every $\phi \in \Delta Y$ there exists a decomposition $\Upsilon_\phi \in \Delta \Delta Y$ with $\phi = \int \psi d\Upsilon_\phi$ satisfying $\mu(\psi) = \mu(\phi)$ and $|\text{supp}(\psi)| \leq |\Theta|$ for all $\psi \in \text{supp}(\Upsilon_\phi)$. All the $\psi \in \text{supp}(\Upsilon_\phi)$ are indecomposable as they are extreme points of

ΔY under the restriction that $\mu(\psi) = \mu(\phi)$. Following the proof of Theorem 2 in [Kolotilin et al. \(2023\)](#)—using the above slight generalization of their Lemma 4 and the condition that $\mu(\psi) = \mu(\phi)$ as the moment condition throughout their argument—there exists a distribution Φ' that induces the same distribution over $\Delta\Theta$ as Φ and is supported on indecomposable signals.

It remains to show that Φ' yields a higher profit if Φ puts positive mass on decomposable signals. Since it induces the same distribution over $\Delta\Theta$, we have $\int \Pi(\mu(\phi)) d\Phi = \int \Pi(\mu(\phi)) d\Phi'$. If Φ put a positive mass on decomposable signals, the distribution over x induced by Φ' is a (strict) mean-preserving spread of that induced by Φ . Note that (IC) is binding for any information structure since a deterministic wage cannot implement high effort. It then follows from [Kim \(1995\)](#) that the expected wage required to induce effort is strictly lower under Φ' than under the original Φ ; a contradiction. \square

Proof of Proposition 5: With slight abuse of notation, we write μ for the implied probability of θ_H , $\mu(\cdot)(\theta_H)$ and we write $F_{e,\theta}$ for $F(\cdot|e, \theta)$.

1 \implies 2: Suppose that there exist $\alpha_0, \alpha_1 \in \mathbb{R}^2$ such that $\mathbf{x} = \alpha_0 \mathbf{1} + \alpha_1 \mu$. Note that $\alpha_0, \alpha_1 \neq 0$ since $\int \mathbf{x}(y) dF(y|e=1, \mu_0) = \alpha_0 + \alpha_1 \mu_0 = 0$, so that $\mathbf{x} = -\alpha_1 \mu_0 \mathbf{1} + \alpha_1 \mu$. Given a distribution $m \in \Delta\Delta\Theta$ the induced distribution over posteriors is given by $n(x) = m(\frac{x-\alpha_0}{\alpha_1})$. Conversely, given a distribution $n \in \Delta\mathbb{R}$, the corresponding distribution over posteriors is $m(\mu) = n(\alpha_0 + \alpha_1 x)$.

2 \implies 1: Suppose \mathbf{x} is not in the span, i.e. there exist y_1, y_2, y_3 such that $\mathbf{x} \neq \alpha_0 \mathbf{1} + \alpha_1 \mu$ on y_1, y_2, y_3 . Let $\alpha \in [0, 1]$ such that $\alpha \mu(y_1) + (1-\alpha) \mu(y_2) = \mu(y_3)$. By the failure of linearity, we have that $\alpha \mathbf{x}(y_1) + (1-\alpha) \mathbf{x}(y_2) \neq \mathbf{x}(y_3)$. Let $\delta_y \in \Delta Y$ denote the point measure on y . Then, δ_{y_3} and $\alpha \delta_{y_1} + (1-\alpha) \delta_{y_2}$ map to the same posterior, but different scores, contradicting the existence of such a map.

To show the LCDF condition, assume that $\mathbf{x} = -\alpha_1 \mu_0 \mathbf{1} + \alpha_1 \mu$ for all μ_0 (with α_1 depending on μ_0 , as we will derive). Then, we have

$$\begin{aligned} \mathbf{x} &= 1 - \frac{dF_{e_L, \mu_0}}{dF_{e_H, \mu_0}} = -\alpha_1 \mu_0 \mathbf{1} + \alpha_1 \mu_0 \frac{dF_{e_H, \theta_H}}{dF_{e_H, \mu_0}} \\ \implies F_{e_H, \mu_0} - F_{e_L, \mu_0} &= \alpha_1 \mu_0 (F_{e_H, \theta_H} - F_{e_H, \mu_0}) = \alpha_1 \mu_0 (1 - \mu_0) (F_{e_H, \theta_H} - F_{e_H, \theta_L}) \end{aligned}$$

To find α_1 , consider

$$\mathbb{E}[y|e_H, \mu_0] - \mathbb{E}[y|e_L, \mu_0] = \int \mathbf{x}(y)y dF_{e_H, \mu_0} = \int y \alpha_1 (\mu(y) - \mu_0) dF_{e_H, \mu_0} = \alpha_1 \mu_0 (\mathbb{E}[y|e_H, 1] - \mathbb{E}[y|e_H, \mu_0])$$

and hence, by linearity in μ_0

$$\alpha_1 = \frac{\mathbb{E}[y|\Delta_e, 0] + \mu_0 [\mathbb{E}[y|\Delta_e, 1] - \mathbb{E}[y|\Delta_e, 0]]}{\mu_0 (1 - \mu_0) \mathbb{E}[y|e_H, \Delta_\theta]},$$

where we let $\mathbb{E}[y|\Delta_e, \theta] := \mathbb{E}[y|e_H, \theta] - \mathbb{E}[y|e_L, \theta]$ and $\mathbb{E}[y|e_H, \Delta_\theta] := \mathbb{E}[y|e_H, \theta_H] - \mathbb{E}[y|e_H, \theta_L]$. Consequently,

$$F_{e_H, \mu_0} - F_{e_L, \mu_0} = \frac{\mathbb{E}[y|\Delta_e, 0] + \mu_0 [\mathbb{E}[y|\Delta_e, 1] - \mathbb{E}[y|\Delta_e, 0]]}{\mathbb{E}[y|e_H, \Delta_\theta]} (F_{e_H, \theta_H} - F_{e_H, \theta_L})$$

We have established the linear distribution function condition if we can express both $F(\cdot|e_H, 0)$ and $F(\cdot|e_L, 1)$ as convex combinations of $F(\cdot|e_L, 0)$ and $F(\cdot|e_H, 1)$ (since linearity in μ is immediate).

To see this, let $\mu_0 = 0$ and note that

$$F_{e_H, \theta_L} = F_{e_L, \theta_L} + \frac{\mathbb{E}[y|\Delta_e, 0]}{\mathbb{E}[y|e_H, \Delta_\theta]} (F_{e_H, \theta_H} - F_{e_H, \theta_L})$$

establishes this for $F(\cdot|e_H, 0)$. Then, let $\mu_0 = 1$ and note that

$$F_{e_H, \theta_H} - F_{e_L, \theta_H} = \frac{\mathbb{E}[y|\Delta_e, 1]}{\mathbb{E}[y|e_H, \Delta_\theta]} (F_{e_H, \theta_H} - F_{e_H, \theta_L})$$

expresses the relation for $F(\cdot|e_L, 1)$ by using the representation of $F(\cdot|e_H, 0)$. We hence have the linear distribution function condition.

Finally, suppose that $F(\cdot|e, \mu) = F(\cdot|0, 0) + g(e, \mu)\Delta F(\cdot)$. Then

$$\boldsymbol{\mu}(y)(\theta_H) = \mu_0(\theta_H) \frac{dF_{e_H, \mu_0} + (g(1, 1) - g(1, \mu_0)) d\Delta F(y)}{dF_{e_H, \mu_0}} = \mu_0(\theta_H) \left(1 + (g(1, 1) - g(1, \mu_0)) \frac{d\Delta F(y)}{dF_{e_H, \mu_0}} \right)$$

and the score is given by

$$\boldsymbol{x}(y) = 1 - \frac{dF_{e_L, \mu_0}}{dF_{e_H, \mu_0}} = (g(1, \mu_0) - g(0, \mu_0)) \frac{d\Delta F(y)}{dF_{e_H, \mu_0}}$$

Then, note that

$$\boldsymbol{x} = \frac{1}{\mu_0(\theta_H)} \frac{g(1, \mu_0) - g(0, \mu_0)}{g(1, 1) - g(1, \mu_0)} (\boldsymbol{\mu} - \mu_0) \quad (24)$$

thus concluding the proof. \square

Proof of Proposition 6: Plugging in the parametrization for g from Definition 4.3 to (24) yields (8). From (8), we see that we can write incentive compatibility in (7) as a function of the posterior μ . The objective and all constraints therefore depend only on μ and we can change variables from ϕ to $\mu = \boldsymbol{\mu}(\phi)(\theta_H)$, where the fact that wages are deterministic follows as in the proof of Proposition 3. We can therefore formulate the problem in this space. A distribution m is feasible if and only if it is a mean-preserving contraction of the fully informative \bar{m} (Gentzkow and Kamenica, 2016; Kolotilin, 2018). \square

Proof of Theorem 2: A function V is strictly S-shaped on $[0, 1]$ if there exists a $x \in [0, 1]$ such that V is strictly convex on $[0, x]$ and strictly concave on $(x, 1]$. We focus on this case in the proof, i.e. $V''' < 0$ and upper censorship; the other case follows from applying the result to $V^-(\mu) := V(1 - \mu)$.

We follow the general approach of the proof of Theorem 1 in Kolotilin et al. (2022) (in short KMZ), extending the argument to a distribution \bar{m} which may not admit a density. For easy comparability, we adopt their notation conventions. The basic outline is as follows. We define the profit W from a censorship policy with a given cutoff and show that it is (strictly) quasi-concave in the cutoff. This part of the proof needs to be adapted to take into account potential splitting of atoms, and we do so below, writing the objective in quantile space. From there, KMZ study the properties of the derivative of W and construct an auxiliary value function which dominates the original one and is attained by the generalized cutoff property. This portion of the proof is largely independent of the distribution and goes through verbatim, so we are very brief in discussing it, only modifying certain steps to take account of the fact that W may not be everywhere differentiable in our case.

Let $\omega(q) = \sup\{\mu : F(\mu) \leq q\}$ denote the quantile function associated to F . Where F is continuous, it is given by $F^{-1}(q)$, the above generalized inverse ensures that ω is increasing with

a limit from the right and continuous from the left. We can then define the value of an upper censorship policy with cutoff q^* as

$$W(q^*) = \int_0^{q^*} V(\omega(q)) dq + (1 - q^*)V(m^*(q^*))$$

where $m^*(q^*) := \frac{\int_{q^*}^1 \omega(q) dq}{1 - q^*}$ denotes the mean state when the highest $1 - q^*$ realizations are pooled. Note that m^* and W are continuous; they are differentiable except at points where ω has a discontinuity, which happens when F is constant (i.e. the distribution has a gap in its support).

Then, adapting KMZ Lemma 4, we have

$$\begin{aligned} W'(q^*) &= V(\omega(q^*)) - V(m^*(q^*)) + (1 - q^*)V'(m^*(q^*)) \frac{-(1 - q^*)\omega(q^*) + \int_{q^*}^1 \omega(q) dq}{(1 - q^*)^2} \\ &= V(\omega(q^*)) - V(m^*(q^*)) + V'(m^*(q^*)) (m^*(q^*) - \omega(q^*)) \\ &= \int_{\omega(q^*)}^{m^*(q^*)} V''(z)(z - \omega(q^*)) dz \end{aligned} \quad (25)$$

when the derivative exists; when it does not exist the left and right derivatives are given by

$$W'^-(q^*) = \int_{\omega^-(q^*)}^{m^*(q^*)} V''(z)(z - \omega^-(q^*)) dz; \quad W'^+(q^*) = \int_{\omega(q^*)}^{m^*(q^*)} V''(z)(z - \omega(q^*)) dz,$$

where ω^- is the left limit $\lim_{q \nearrow q^*} \omega(q)$, which exists by construction. The properties of these integral expressions established in KMZ Lemma 5 extend to our setting, W' cannot cross 0 from below, in particular, $W'^-(q^*) \leq 0$ implies that $W'^+(q^*) < 0$. Indeed,

$$\begin{aligned} W'^-(q^*) &= \int_{\omega^-(q^*)}^{\omega(q^*)} V''(z)(z - \omega^-(q^*)) dz + \int_{\omega(q^*)}^{m^*(q^*)} V''(z)(z - \omega^-(q^*)) dz \\ &< \int_{\omega(q^*)}^{m^*(q^*)} V''(z)(z - \omega(q^*)) dz = W'^+(q^*) \end{aligned}$$

since the second component of W'^- has to be strictly below zero if $W'^-(q^*) \leq 0$ since $V''' < 0$.

Therefore, the Proof of Lemma 1 in KMZ applies verbatim and W is strictly quasi-concave. It is easy to see that KMZ Lemma 2 also generalizes: Let q^* be the maximizer of W , then there exists an $x \in [\omega(q^*), m^*(q^*)]$ such that V is convex on $[0, x]$ and concave on $(x, 1]$.

To generalize KMZ Lemma 6, consider $\omega(q^*)$. If ω is continuous at $\omega(q^*)$, W has a derivative and $W'(q^*) = \int_{\omega(q^*)}^{m^*(q^*)} V''(z)(z - \omega(q^*)) dz = 0$. In this case, define $\mu^* = \omega(q^*)$. If ω has a jump at q^* , then $0 \in [W'^-(q^*), W'^+(q^*)]$ and there exists a $\mu^* \in [\omega(q^*), \omega^+(q^*)]$ such that $\int_{\mu^*}^{m^*(q^*)} V''(z)(z - \mu^*) dz = 0$. Then, define the auxiliary value \bar{V} as

$$\bar{V}(s) = \begin{cases} V(s) & s < \mu^* \\ V(m^*(q^*)) + V'(m^*(q^*)) (s - m^*(q^*)) & s \geq \mu^* \end{cases}$$

Then, as above, we have

$$V(s) - V(m^*(q^*)) + V'(m^*(q^*)) (s - m^*(q^*)) = - \int_s^{m^*(q^*)} V''(z)(z - s) dz$$

Therefore, \bar{V} is convex on $[0, 1]$ and $\bar{V}(s) \geq V(s)$, with equality only on $[0, \omega(q^*)] \cup \{m^*(q^*)\}$. Adapting the proof of KMZ Lemma 7, suppose $H \leq_{MPS} F$. Then

$$\int_0^1 V(s) dH(s) \leq \int_0^1 \bar{V}(s) dH(s) \quad (26)$$

$$\leq \int_0^1 \bar{V}(s) dF(s) \quad (27)$$

$$= \int_0^{\mu^*} V(s) dF(s) + \int_{\mu^*}^1 \bar{V}(s) dF(s) \quad (28)$$

$$= \int_0^{\mu^*} V(s) dF(s) + \int_{\mu^*}^1 V(m^*(q^*)) dF(s) = W(q^*) \quad (29)$$

The inequalities follow since $\bar{V} \geq V$ pointwise and since \bar{V} is convex. The splitting of the integral requires a careful interpretation if F has an atom at $\omega(q^*)$: We allocate a weight of $q^* - F(\omega(q^*))$ to the first integral and the remaining weight to the second. This is irrelevant in the first equality, since $V(\mu^*) = \bar{V}(\mu^*)$ but is essential in the second equality. Generalized censorship is hence optimal, as it attains the value (29) which is an upper bound on the profit.

To see that it is (essentially) unique, note that any other optimal policy needs to satisfy the two inequalities above with equality. Since \bar{V} is strictly convex on $[0, \omega(q^*))$, the second holds with equality only if H reveals all states in this interval almost surely. Then, note that $\bar{V}(s) > V(s)$ on $(\omega(q^*), m^*(q^*)) \cup (\omega(q^*), m^*(q^*))$ such that the first inequality holds with equality only if the remaining states are either revealed to be $\omega(q^*)$ or pooled to induce an expected state $m^*(q^*)$, which up to equal-mean splittings of the pooling region, is only satisfied by upper censorship with cutoff $\omega(q^*)$, with suitable splitting in the case of an atom at the cutoff. \square

Proof of Theorem 1: The contracting problem (9) is equivalent to

$$\sup_{w, m \leq_{MPS} \bar{m}} \inf_{\lambda \geq 0} \mathcal{L}(m, w; (\lambda_P, \lambda_{IC})). \quad (30)$$

To see this, note that since

$$\inf_{\lambda \geq 0} \mathcal{L}(m, w; \lambda) = \begin{cases} \int (\delta \Pi_2(\mu) - w(\mu)) dm(\mu) & \text{if (P)\&(IC) are satisfied} \\ -\infty & \text{else} \end{cases}, \quad (31)$$

the infimum simply wraps the constraints into the objective function. It is always the case that $\inf \sup \mathcal{L} \geq \sup \inf \mathcal{L}$, where the supremum is taken over the choice variables and the infimum over the multipliers. If this condition holds with equality, i.e. if we can exchange sup and inf, we say that the optimization problem satisfies *strong duality*.

Fix a distribution $m \leq_{MPS} \bar{m}$ and consider the problem of finding optimal wages subject to the participation and incentive constraint. It can be solved by point-wise optimization, arriving at (13).

Lemma 1. *The wage setting problem satisfies strong duality, i.e.*

$$\sup_w \inf_{\lambda \geq 0} \mathcal{L}(m, w; \lambda) = \inf_{\lambda \geq 0} \sup_w \mathcal{L}(m, w; \lambda).$$

Proof of Lemma: If m is degenerate (a point mass on μ), the problem is infeasible and hence both the primal and dual value are $-\infty$. If m is nondegenerate, it is easy to see that the problem can be written in utility space where the objective is a concave functional and the constraints are

linear. Furthermore, a strictly feasible utility promise exists (e.g., pay U after every posterior with a suitable large bonus if and only if $\mu > \mu_0$). Therefore, by [Luenberger \(1969, p. 224\)](#), the problem satisfies strong duality. \triangle

Consider now the Lagrangian that results from plugging in for w from (13).

$$\sup_w \mathcal{L}(m, w; \lambda) = \int \ell^*(\mu; \lambda) dm(\mu) \quad (32)$$

The information design problem given λ reads

$$\sup_{m \leq_{MPS} \bar{m}} \int \ell^*(\mu; \lambda) dm(\mu), \quad (33)$$

and can therefore be solved via [Theorem 2](#).

Lemma 2. *The optimal evaluation structure in (33) is unique and*

1. *fully informative for every λ if $\Pi_2'' > 0$,*
2. *lower-censorship for every λ if $\Pi_2''' > 0$ and $v''' \leq 0$, and*
3. *upper-censorship for every λ such that the non-negativity constraint is slack if $\Pi_2''' < 0$ and $v''' \geq 0$.*

Proof of Lemma: From (16), it is easy to see that $\ell^{*'} is continuous at the nonnegativity constraint as the wage is locally zero, that$

$$\frac{\partial^2}{\partial \mu^2} \ell^*(\mu; \lambda) = \delta \Pi_2''(\mu) + \lambda_{IC}^2 \left[\frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \right]^2 \rho'(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0))$$

and that $\ell^{*''}$ jumps upwards as $\rho' \geq 0$. Furthermore,

$$\frac{\partial^3}{\partial \mu^3} \ell^*(\mu; \lambda) = \delta \Pi_2'''(\mu) + \lambda_{IC}^3 \left[\frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \right]^3 \rho''(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0))$$

at an interior wage and $\frac{\partial^3}{\partial \hat{\mu}^3} \ell^*(\hat{\mu}; \lambda) = \delta \Pi_2'''(\hat{\mu})$ when the wage is zero.

If $\Pi_2'' \leq 0$, ℓ^* is convex and $m = \bar{m}$ is uniquely optimal. If $\Pi_2''' > 0$ and $\rho'' \geq 0$, ℓ^* is concave to convex for all λ , and lower-censorship is optimal. If $\Pi_2''' < 0$ and $\rho'' \leq 0$, ℓ^* is convex to concave, except when $\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) < 0$. In this region, ℓ^* is concave and $\ell^{*''}$ jumps up at the boundary. For λ such that $\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) \geq 0$, upper-censorship is uniquely optimal. \triangle

Applying this result to the original problem requires another step of duality.

Lemma 3. *The information design problem satisfies strong duality, i.e.*

$$\sup_{m \leq_{MPS} \bar{m}} \inf_{\lambda \geq 0} \int \ell^*(\mu; \lambda) dm(\mu) = \inf_{\lambda \geq 0} \sup_{m \leq_{MPS} \bar{m}} \int \ell^*(\mu; \lambda) dm(\mu). \quad (34)$$

Proof of Lemma: The space of posterior distributions satisfying $m \leq_{MPS} \bar{m}$ is compact in the weak* topology, and, as ℓ^* is continuous and bounded for any λ , the objective is continuous and linear in m . The objective is continuous in λ . To see quasi-convexity in λ , note that by an envelope

argument

$$\frac{\partial \int \ell^*(\mu; \lambda) dm(\mu)}{\partial \lambda_P} = \int \rho(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0)) - U - c dm(\hat{\mu})$$

and similarly for λ_{IC} and hence the Hessian of $\int \ell^*(\mu; \lambda) dm(\mu)$ is given by

$$\begin{pmatrix} \int f(\mu) dm(\mu) & \int f(\mu) \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) dm(\mu) \\ \int f(\mu) \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) dm(\hat{\mu}) & \int f(\mu) \left[\frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0) \right]^2 dm(\mu) \end{pmatrix} \quad (35)$$

where $f(\mu) := \rho'(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0))$ is a positive kernel and the range of integration is over μ such that $u'^{-1}(\lambda_P + \lambda_{IC} \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0)) \geq 0$. Hence, the integral $\int f(\mu) g_1(\mu) g_2(\mu) dm(\mu)$ defines an inner product $\langle \cdot, \cdot \rangle_f$ (between functions that share support with m), and $\lambda \rightarrow \int \ell^*(\mu; \lambda) dm(\mu)$ is weakly convex by Cauchy-Schwartz, as the determinant of the Hessian reads

$$\langle g_1, g_1 \rangle_f \langle g_2, g_2 \rangle_f - \langle g_1, g_2 \rangle_f^2 \geq 0$$

for $g_1 = 1$ and $g_2 = \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} (\mu - \mu_0)$.

Therefore, the problem satisfies the conditions of Sion's Minimax Theorem and we have

$$\inf_{\lambda \geq 0} \sup_{w, m \leq_{MPS} \bar{m}} \mathcal{L}(m, w; (\lambda_P, \lambda_{IC})) = \sup_{m \leq_{MPS} \bar{m}} \inf_{\lambda \geq 0} \sup_w \mathcal{L}(m, w; (\lambda_P, \lambda_{IC})). \quad \Delta$$

Using Lemma 1 and 3, we have

$$\Pi_1 = \sup_{w, m \leq_{MPS} \bar{m}} \inf_{\lambda \geq 0} \mathcal{L}(m, w; \lambda) = \sup_{m \leq_{MPS} \bar{m}} \inf_{\lambda \geq 0} \sup_w \mathcal{L}(m, w; \lambda) = \inf_{\lambda \geq 0} \sup_{w, m \leq_{MPS} \bar{m}} \mathcal{L}(m, w; \lambda).$$

We can therefore simplify the general problem (9) using the properties of optimal evaluation structures from Lemma 2, i.e. we can restrict attention to censorship information structures. We restrict attention to the case of lower censorship, the case of upper-censorship is closely analogous. This simplified problem is

$$\max_{q \in [0, 1], \bar{w}, w} q [\delta \Pi_2(\bar{\mu}(q)) - \bar{w}] + (1 - q) \int_{\bar{m}^{-1}(q)}^1 \delta \Pi_2(\mu) - w(\mu) d\bar{m}(\mu) \quad (36)$$

$$\text{s.t. } qu(\bar{w}) + (1 - q) \int_{\bar{m}^{-1}(q)}^1 u(w(\mu)) d\bar{m}(\mu) - c \geq U \quad (\text{P}^S)$$

$$\frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \left[q(\bar{\mu}(q) - \mu_0)u(\bar{w}) + (1 - q) \int_{\bar{m}^{-1}(q)}^1 (\mu - \mu_0)u(w(\mu)) d\bar{m}(\mu) \right] \geq c \quad (\text{IC}^S)$$

where q denotes the mass of output realizations pooled at the low evaluation, \bar{w} is the associated wage, $\bar{\mu}(q) = \int_0^{\bar{m}^{-1}(q)} \mu d\bar{m}(\mu)$ the associated posterior and the integrals $\int_0^{\bar{m}^{-1}(q)} \cdot d\bar{m}(\mu)$ and $\int_{\bar{m}^{-1}(q)}^1 \cdot d\bar{m}(\mu)$ are understood to have measure q and $1 - q$.

It remains to show that this simplified problem has a solution, i.e. that the principal cannot increase his profit by choosing a sequence of information structures converging to no information.

Lemma 4. *The simplified contracting problem (36) has a unique solution. The optimal information structure is non-degenerate ($q < 1$).*

Proof of Lemma: Suppose not, we will show that the costs of providing incentives diverge as $q \rightarrow 1$. To see this, first note that, by the martingale property of beliefs, the belief after the pooled signal has to be close to the prior

$$\begin{aligned} q\bar{\mu}(q) &\leq \mu_0 \leq 1 - q + q\bar{\mu}(q) \\ \implies \bar{\mu}(q) - \mu_0 &\in \left[-\frac{1-q}{q}(1-\mu_0), \frac{1-q}{q}\mu_0\right]. \end{aligned}$$

In the IC constraint, we have

$$\begin{aligned} c &\leq \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \left[q(\bar{\mu}(\epsilon) - \mu_0) u(\bar{w}(q)) + (1 - q) \int_{\bar{m}^{-1}(q)}^1 (\mu - \mu_0) u(w(\mu)) d\bar{m}(\mu) \right] \\ &\leq \frac{b + \Delta b \mu_0}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)} \left[(1 - q) \mu_0 u(\bar{w}(\epsilon)) + (1 - q) (1 - \mu_0) \int_{\bar{m}^{-1}(q)}^1 u(w(\mu)) d\bar{m}(\mu) \right] \end{aligned}$$

Hence, as $q \rightarrow 1$, we require $u(w) \geq c_0(1-q)^{-1}$, for a suitable constant $c_0 > 0$, for either the pooling utility or the expected utility after revelation. But then, the objective is $\leq c_1 - \epsilon \cdot c_2 w ((1-q)^{-1}) \rightarrow -\infty$ for suitable constants $c_1, c_2 > 0$, which is clearly not optimal. Hence, the optimal distribution of posteriors is nondegenerate and we can restrict the set of censoring policies to those pooling output with at most probability $\bar{q} < 0$.

Then, we can write

$$\Pi_1 = \max_{q^* \in [0, \bar{q}]} \inf_{\lambda \geq 0} \sup_w \mathcal{L}(m, w; \lambda)$$

where writing max is justified since $\inf_{\lambda \geq 0} \sup_w \mathcal{L}(m, w; \lambda)$ is the well-defined and continuous value of the contracting problem given the information structure. Hence, a unique solution exists. \triangle

Since a wage function and a distribution over posteriors solve the original problem if and only if they induce a solution in the simplified problem, this concludes the proof. The contract is unique in the sense that the signal structure is unique up to splittings that don't reveal any additional information about effort or ability and the wage function is determined up to modifications on realizations of measure zero. \square

Proof of Proposition 7: If \bar{m} admits a density, the value of a censorship information structure W is differentiable and the optimality condition (18) is immediate from condition (25). Recall that the profit of the principal in the case of square-root utility is given by (3). Using (8) and the censorship structure (assuming upper-censorship for concreteness) we can write

$$\begin{aligned} \Pi_1 = \max_{\mu^*} & -\frac{1}{2} \frac{c^2}{\left(\frac{b + \Delta b \mu}{(\Delta a + \Delta b) \mu_0 (1 - \mu_0)}\right)^2} \frac{1}{\int_0^{\mu^*} (\mu - \mu_0)^2 d\bar{m}(\mu) + (1 - m(\mu^*)) (\bar{\mu}(\mu^*) - \mu_0)^2} \\ & + \delta \left[m(\mu^*) \Pi_2(\bar{\mu}(\mu^*)) + \int_{\mu^*}^1 \Pi_2(\mu) d\bar{m}(\mu) \right], \end{aligned}$$

where $\bar{\mu}(\mu^*) = \frac{\int_{\mu^*}^1 \mu d\bar{m}(\mu)}{1-m(\mu^*)}$ is the mean posterior in the bunching region. Assuming that \bar{m} admits a density, the first-order condition is

$$-\frac{1}{2} \frac{c^2}{\left(\frac{b+\Delta b\mu}{(\Delta a+\Delta b)\mu_0(1-\mu_0)}\right)^2} \frac{(\mu^* - \bar{\mu}(\mu^*))^2}{\left(m(\mu^*) (\bar{\mu}(\mu^*) - \mu_0)^2 + \int_{\mu^*}^1 (\mu - \mu_0)^2 d\bar{m}(\mu)\right)^2} + \delta [\Pi_2(\bar{\mu}(\mu^*)) - \Pi_2(\mu^*) + \Pi_2'(\bar{\mu}(\mu^*))(\mu^* - \bar{\mu})] = 0 \quad (37)$$

which establishes (19).

The comparative statics follow from implicitly differentiating the first-order condition. The denominator of the implicit derivatives is positive since the objective is concave local to the optimal solution. We have

$$\frac{d\mu^*}{d\delta} = \frac{\Pi_2(\bar{\mu}(\mu^*)) - \Pi_2(\mu^*) + \Pi_2'(\bar{\mu}(\mu^*))(\mu^* - \bar{\mu})}{-\frac{d}{d\mu^*} [(37)]} > 0$$

since the denominator is positive (by the concavity of Π_2). The optimal evaluation structure pools a larger interval of states if the discount rate is higher. The proof for the case of upper-censorship is analogous. It is easy to see that the argument generalizes to an arbitrary u , noting that we can again write the expected wage as a now implicit function of μ^* , that the problem has a unique solution (whence the denominator of the implicit derivative is positive) and that the numerator is unchanged.

To see the statement about costs, let c_t denote the cost of effort in period t . Then

$$\frac{d\mu^*}{dc_1} = \frac{-2 \frac{c_1}{\left(\frac{b+\Delta b\mu}{(\Delta a+\Delta b)\mu_0(1-\mu_0)}\right)^2} \frac{(\mu^* - \bar{\mu}(\mu^*))^2}{\left(\int_0^{\mu^*} (\mu - \mu_0)^2 d\bar{m}(\mu) + (1-m(\mu^*))(\bar{\mu}(\mu^*) - \mu_0)^2\right)^2}}{[\cdot]} < 0.$$

To see that

$$\frac{d\mu^*}{dc_2} = \frac{\delta \frac{d}{dc_2} [\Pi_2(\mu^*) - \Pi_2(\bar{\mu}) + (\bar{\mu} - \mu^*) \Pi_2'(\mu^*)]}{[\cdot]} > 0$$

note that by (3) the part of the continuation value nonlinearly depending on beliefs is quadratic in c_2 . This also establishes that the problem is independent of cost shifts that leave c_1/c_2 fixed. \square

A.1 Private Information of the Principal

Consider the game as described in the text.

First, consider any weak PBE satisfying no-holdup. We will show that the principal profit is smaller than Π^* . On path, it induces a distribution over agent posteriors $m(\mu)$ and conditional on the posterior μ a distribution over information structures and wage schedules. As the participation and incentive compatibility constraints are satisfied conditional on the agent's second stage information set, they are satisfied conditional on μ . Hence, this distribution over information structures and wage schedules satisfies the constraints of the second period problem for μ . Furthermore, in equilibrium, the principal's expected belief conditional on the agent's posterior has to be μ by Bayes-plausibility. Therefore, the continuation profit Π_2^{EQ} satisfies $\iint \Pi_2^{EQ}(\mu_p, \mu) dm(\mu_p|\mu) dm(\mu) \leq \int \Pi_2^*(\mu) dm(\mu)$. The principal's continuation value is dominated by that under the optimal contract. Similarly, in the first period, the equilibrium induces a distribution over wages and agent posteriors on-path that

satisfies the conditions of the first period problem (9).²⁹ This implies that the first-period profit under the equilibrium is dominated by that under the optimal contract, as we set out to argue.

We say that the second period contract is agent-determined (for a lower-censorship contract with pooling posterior μ^* and cutoff $\bar{\mu}$) if for all μ_P the optimal contract in the second period with an agent with belief μ^* is independent of the principal's belief μ_P .

Second, consider a weak PBE with passive beliefs. Formally, we require that the agent does not update his beliefs about θ based on the contract offer in either period. We will show that if the second period contract is agent-determined any such PBE induces a joint distribution over agent beliefs and wages that is identical to the one induced by the optimal contract up to a set of measure zero. Consider the contract offer stage in the second period. If the agent has posterior belief μ^* and the optimal contract is independent of the principal's belief, the principal offers this contract in any such PBE. Consequently, by the martingale property, the principal's value of inducing posterior beliefs μ_p, μ^* is $\Pi_2(\mu^*)$. In the first period, the agent has belief μ_0 on path (by Bayes' rule) and off-path (by passive beliefs). Therefore, the principal's best response is the solution to the contracting problem. Therefore, the equilibrium is outcome equivalent to the optimal contract. We also have a partial converse: If the optimal contract depends on the principal's belief conditional on the agent's belief after the pooling outcome, consider a passive belief equilibrium replicating the optimal contract. Then, the principal strictly prefers learning the true output and offering this optimal contract after the pooling signal realized in period 1. This deviation is feasible by passive beliefs and yields a higher payoff, hence there is no passive beliefs equilibrium replicating the optimal contract.

Third, suppose that the agent observes the principal's information structure and the agent's posterior satisfies the restriction of FN 24. We will show that any such PBE induces a joint distribution over agent beliefs and wages that is identical to the one induced by the optimal contract up to a set of measure zero. Note that the principal can achieve Π^* by offering the optimal first-period contract and choosing not to acquire private information. Furthermore, she can achieve no higher profit by the first remark. By the uniqueness of the optimal contract (up to measure zero events), the equilibrium has to induce this outcome, otherwise the principal would obtain a strictly lower profit.

A.2 Private Effort Choice

First, we provide details for the rewriting of the dynamic IC constraint to (23). Note that

$$\begin{aligned} \int u(w(y, \mu)) dF(y|1, \bar{\mu}) - c &= \int u(w(y, \mu)) dF(y|e_H, \mu) + \int u(w(y, \mu)) d\Delta F(y)(\Delta a + \Delta b)(\bar{\mu} - \mu) - c \\ &= U + c + \frac{c}{b + \Delta b\mu}(\Delta a + \Delta b)(\bar{\mu} - \mu) - c \\ &= U + \frac{c}{b + \Delta b\mu}(\Delta a + \Delta b)(\bar{\mu} - \mu) \end{aligned}$$

where we used the binding second-period P and IC constraint. Similarly, $\int u(w(y, \mu)) dF(y|e_L, \bar{\mu}) = U + \frac{c}{b + \Delta b\mu} \Delta a (\bar{\mu} - \mu)$. The transformation that follows from plugging in these representations and rearranging.

²⁹In a weak no-holdup PBE the agent may be misguided about the contract offered in the second period after a hypothetical deviation on the first period. Such beliefs can only strengthen the first-period IC constraint: On-path, the agent is held to the participation constraint (no-holdup); after a deviation, he might obtain a positive continuation surplus.

In the binary case with $\Delta a = 0$, we know that $\tilde{\mu} = \mu_0$, and we can follow the analysis using the dynamic IC. The partially maxed out Lagrangian then reads

$$\int \tilde{\ell}(\mu; \lambda) dm(\mu).$$

with

$$\tilde{\ell}(\mu; \lambda) = \begin{cases} \ell^*(\mu; \lambda) - \lambda_{IC} \left(1 - \frac{(b+\mu_0\Delta b)}{\mu_0(1-\mu_0)\Delta b} (\mu - \mu_0)\right) (\mu_0 - \mu) \Delta b \frac{c}{b+\Delta b\mu} & \mu \leq \mu_0 \\ \ell^*(\mu; \lambda) & \mu > \mu_0 \end{cases}$$

We have

$$\frac{\partial^3}{\partial \mu^3} \left(-\lambda_{IC} \left(1 - \frac{(b+\mu_0\Delta b)}{\mu_0(1-\mu_0)\Delta b} (\mu - \mu_0)\right) (\mu_0 - \mu) \Delta b \frac{c}{b+\Delta b\mu} \right) = \lambda_{IC} \frac{6c\Delta b(b+\mu_0\Delta b)(b^2+\mu_0(2b+\Delta b))}{\mu_0(1-\mu_0)(b+\Delta b\mu)^4} > 0$$

and therefore it remains the case that $\tilde{\ell}''' > 0$ wherever it is continuously differentiable. The kink is concave, as

$$\frac{\partial}{\partial \mu} \left(-\lambda_{IC} \left(1 - \frac{(b+\mu_0\Delta b)}{\mu_0(1-\mu_0)\Delta b} (\mu - \mu_0)\right) (\mu_0 - \mu) \Delta b \frac{c}{b+\Delta b\mu} \right) \Big|_{\mu=\mu_0} = \frac{c\Delta b\lambda_{IC}}{b+\Delta b\mu_0} > 0$$

and the curvature of the function increases, as

$$\frac{\partial^2}{\partial \mu^2} \left(-\lambda_{IC} \left(1 - \frac{(b+\mu_0\Delta b)}{\mu_0(1-\mu_0)\Delta b} (\mu - \mu_0)\right) (\mu_0 - \mu) \Delta b \frac{c}{b+\Delta b\mu} \right) \Big|_{\mu=\mu_0} = -2c\lambda_{IC} \left(\frac{1}{1-\mu_0} + \frac{1}{\mu_0} + \frac{\Delta b^2}{(b+\Delta b\mu)^2} \right) < 0$$

(Note that the signs are flipped relative to their intuitive interpretation, since the component is part of the Lagrangian for $\mu \leq \mu_0$.) This establishes the result, as the pasted Lagrangian is concave to convex, with a concave kink at $\mu = \mu_0$. If $\tilde{\ell}$ is concave at $\mu = \mu_0$ with a single support point of concavification, the optimal information structure is uninformative and therefore the IC cannot be satisfied. Therefore, from the shape of $\tilde{\ell}$, the concavification is supported at max support \bar{m} and at a point $\mu^* < \mu_0$, possibly in addition to μ_0 .

A.3 Commitment to a Continuation Value

Suppose that $u(w) = \sqrt{2w}$ and that in the first period, the principal can commit to a continuation value, i.e. $U(s)$. Note that this does not change the transformation to belief space and we can hence write $w(\mu)$, $U(\mu)$. The problem reads

$$\Pi_1(\mu_0) = \max_{m \leq_{MPS} \bar{m}, w} \int y dF(y|e_H, \mu_0) + \int \delta \Pi_2(\mu, U(\mu)) - w(\mu) dm(\mu) \quad (38)$$

$$\text{s.t.} \int [u(w(\mu)) + U(\mu) - U] dm(\mu) - c \geq U \quad (P)$$

$$\int (b + \Delta b\mu_0) \frac{\mu - \mu_0}{(\Delta a + \Delta b)\mu_0(1 - \mu_0)} [u(w(\mu)) + U(\mu)] dm(\mu) \geq c \quad (IC)$$

Straightforward computation establishes that

$$\Pi_2(\mu, U(\mu)) = \int y dF(y|e_H, \mu) - \frac{U(\mu)^2}{2} - \frac{1}{2} \frac{c^2}{\int x(y; \mu)^2 dF(y|1, \mu)}.$$

That is, the continuation value is additively separable in the posterior independent cost of providing the continuation value and the cost of providing incentives. This is a feature of the utility function and greatly simplifies the analysis.

Rewriting the contracting problem in utility space, we see that the first period objective reads (dropping expected output terms)

$$\int \delta \left(-\frac{U(\mu)^2}{2} - \frac{1}{2} \frac{c^2}{\int x(y; \mu)^2 dF(y|1, \mu)} \right) - \frac{u(\mu)^2}{2} dm(\mu)$$

Equating marginal costs of providing utility to the agent, the optimal contract satisfies $\delta U(\mu) = u(\mu)$. Hence, the problem is equivalent to the period by period contracting problem with a cost of utility of $w(u) = (\delta + \delta^2) \frac{u^2}{2}$, or, equivalently, a utility function $u(w) = \sqrt{\frac{2}{\delta + \delta^2} w}$. Therefore, Theorem 1 applies and we have the desired result.

Remark. With a general utility function, the costs of providing the continuation utility and the posterior belief are not separable in the principal's continuation profit. This introduces cross-terms in the derivatives of the Lagrangian which are hard to control without making restrictions on the Lagrange multipliers. Hence, the proof strategy of Theorem 1 does not easily generalize to this case.

References

- Adrian, Tobias and Mark M. Westerfield**, “Disagreement and Learning in a Dynamic Contracting Model,” *Review of Financial Studies*, October 2009, *22* (10), 3873–3906. [6](#)
- Arieli, Itai, Yakov Babichenko, Fedor Sandomirskiy, and Omer Tamuz**, “Feasible Joint Posterior Beliefs,” *Journal of Political Economy*, September 2021, *129* (9), 2546–2594. [24](#)
- Aumann, Robert J. and Michael Maschler**, *Repeated Games with Incomplete Information*, Cambridge, Mass: MIT Press, 1995. [19](#)
- Best, James and Daniel Quigley**, “Persuasion for the Long Run,” *Journal of Political Economy*, May 2024, *132* (5), 1740–1791. [29](#)
- Bhaskar, V. and George J. Mailath**, “The Curse of Long Horizons,” *Journal of Mathematical Economics*, May 2019, *82*, 74–89. [6](#), [26](#)
- Bhaskar, V and Nikita Roketskiy**, “The Ratchet Effect: A Learning Perspective,” *Journal of Economic Theory*, December 2023, *214*, 105762. [6](#)
- Boleslavsky, Raphael and Kyungmin Kim**, “Bayesian Persuasion and Moral Hazard,” *SSRN Electronic Journal*, 2017. [6](#), [19](#)
- Cho, In-Koo and David M. Kreps**, “Signaling Games and Stable Equilibria,” *The Quarterly Journal of Economics*, May 1987, *102* (2), 179–221. [24](#)
- **and Joel Sobel**, “Strategic Stability and Uniqueness in Signaling Games,” *Journal of Economic Theory*, April 1990, *50* (2), 381–413. [24](#)
- Dai, Liang, Yenan Wang, and Ming Yang**, “Dynamic Contracting with Flexible Monitoring,” 2019. [5](#)

- Datar, Srikant, Susan Cohen Kulp, and Richard A. Lambert**, “Balancing Performance Measures,” *Journal of Accounting Research*, June 2001, 39 (1), 75–92. 5
- Demarzo, Peter M. and Yuliy Sannikov**, “Learning, Termination, and Payout Policy in Dynamic Incentive Contracts,” *The Review of Economic Studies*, January 2017, 84 (1), 182–236. 6, 26
- Demougin, Dominique and Claude Fluet**, “Monitoring versus Incentives,” *European Economic Review*, October 2001, 45 (9), 1741–1764. 5
- Dewatripont, Mathias, Ian Jewitt, and Jean Tirole**, “The Economics of Career Concerns, Part I: Comparing Information Structures,” *The Review of Economic Studies*, January 1999, 66 (1), 183–198. 30
- Doval, Laura and Vasiliki Skreta**, “Constrained Information Design,” *Mathematics of Operations Research*, February 2023, 49 (1), 78–106. 6, 18, 19
- Dumav, Martin, Urmeem Khan, and Luca Rigotti**, “Moral Hazard with Heterogeneous Beliefs,” October 2021. 25
- Dworczak, Piotr and Anton Kolotilin**, “The Persuasion Duality,” *Theoretical Economics*, 2024, 19 (4), 1701–1755. 15
- **and Giorgio Martini**, “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy*, October 2019, 127 (5), 1993–2048. 19, 20
- Dye, Ronald A.**, “Optimal Monitoring Policies in Agencies,” *The RAND Journal of Economics*, 1986, 17 (3), 339–350. 5
- Ederer, Florian**, “Feedback and Motivation in Dynamic Tournaments,” *Journal of Economics & Management Strategy*, September 2010, 19 (3), 733–769. 5
- **, Richard Holden, and Margaret Meyer**, “Gaming and Strategic Opacity in Incentive Provision,” *The RAND Journal of Economics*, December 2018, 49 (4), 819–854. 4
- Ekmekci, Mehmet and Nenad Kos**, “Signaling Covertly Acquired Information,” *Journal of Economic Theory*, December 2023, 214, 105746. 24
- Ely, Jeffrey and Martin Szydlowski**, “Moving the Goalposts,” *Journal of Political Economy*, May 2019, p. 704387. 5
- Ely, Jeffrey C, George Georgiadis, and Luis Rayo**, “Feedback Design in Dynamic Moral Hazard,” 2024. 5
- Fang, Hanming and Giuseppe Moscarini**, “Morale Hazard,” *Journal of Monetary Economics*, May 2005, 52 (4), 749–777. 5

- Feltham, Gerald A. and Jim Xie**, “Performance Measure Congruity and Diversity in Multi-Task Principal/Agent Relations,” *The Accounting Review*, 1994, 69 (3), 429–453. 5
- Fuchs, William**, “Contracting with Repeated Moral Hazard and Private Evaluations,” *American Economic Review*, September 2007, 97 (4), 1432–1448. 4
- Gentzkow, Matthew and Emir Kamenica**, “A Rothschild-Stiglitz Approach to Bayesian Persuasion,” *American Economic Review*, May 2016, 106 (5), 597–601. 33
- Georgiadis, George and Balazs Szentes**, “Optimal Monitoring Design,” *Econometrica*, 2020, p. 55. 5, 6, 19
- Giat, Yahel, Steve T. Hackman, and Ajay Subramanian**, “Investment under Uncertainty, Heterogeneous Beliefs, and Agency Conflicts,” *Review of Financial Studies*, April 2010, 23 (4), 1360–1404. 6
- Gittleman, Maury and Brooks Pierce**, “How Prevalent Is Performance-Related Pay in the United States? Current Incidence and Recent Trends,” *National Institute Economic Review*, November 2013, 226 (1), R4–R16. 1
- Gonzalez-Hernandez, G., A. Sarker, K. O’Connor, and G. Savova**, “Capturing the Patient’s Perspective: A Review of Advances in Natural Language Processing of Health-Related Text,” *Yearbook of Medical Informatics*, August 2017, 26 (1), 214–227. 2
- Grossman, Sanford J. and Oliver D. Hart**, “An Analysis of the Principal-Agent Problem,” *Econometrica*, January 1983, 51 (1), 7. 2, 4, 9, 31
- Hoffmann, Florian, Roman Inderst, and Marcus Opp**, “Only Time Will Tell: A Theory of Deferred Compensation,” *The Review of Economic Studies*, May 2021, 88 (3), 1253–1278. 5
- Holmström, Bengt**, “Moral Hazard and Observability,” *The Bell Journal of Economics*, April 1979, 10 (1), 74–91. 2, 4
- , “Managerial Incentive Problems: A Dynamic Perspective,” *The Review of Economic Studies*, January 1999, 66 (1), 169–182. 2, 6, 30
- **and Paul Milgrom**, “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,” *Journal of Law, Economics, & Organization*, January 1991, 7, 24–52. 2, 4
- Hörner, Johannes and Nicolas S. Lambert**, “Motivational Ratings,” *Review of Economic Studies*, July 2021. 2, 30
- Jehiel, Philippe**, “On Transparency in Organizations,” *The Review of Economic Studies*, April 2015, 82 (2), 736–761. 5

- Jewitt, Ian**, “Justifying the First-Order Approach to Principal-Agent Problems,” *Econometrica*, September 1988, 56 (5), 1177–1190. 12
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, October 2011, 101 (6), 2590–2615. 19
- Kim, Son Ku**, “Efficiency of an Information System in an Agency Model,” *Econometrica*, 1995, 63 (1), 89–102. 4, 32
- Kolotilin, Anton**, “Optimal Information Disclosure: A Linear Programming Approach,” *Theoretical Economics*, May 2018, 13 (2), 607–635. 19, 20, 33
- , **Roberto Corrao, and Alexander Wolitzky**, “Persuasion and Matching: Optimal Productive Transport,” 2023. 32
- , **Timofiy Mylovanov, and Andriy Zapechelnyuk**, “Censorship as Optimal Persuasion,” *Theoretical Economics*, 2022, 17 (2), 561–585. 6, 17, 20, 21, 33
- Li, Anqi and Ming Yang**, “Optimal Incentive Contract with Endogenous Monitoring Technology,” *Theoretical Economics*, 2020, 15, 1135–1173. 5
- Lin, Xiao and Ce Liu**, “Credible Persuasion,” *Journal of Political Economy*, July 2024, 132 (7), 2228–2273. 29
- Lizzeri, Alessandro, Margaret A. Meyer, and Nicola Persico**, “The Incentive Effects of Interim Performance Evaluations,” Technical Report, Penn Economics Department September 2002. 5
- Luenberger, David G.**, *Optimization by Vector Space Methods* Series in Decision and Control, New York: J. Wiley, 1969. 36
- MacLeod, W. Bentley**, “Optimal Contracting with Subjective Evaluation,” *The American Economic Review*, 2003, 93 (1), 216–240. 4
- Nafziger, Julia**, “Timing of Information in Agency Problems with Hidden Actions,” *Journal of Mathematical Economics*, December 2009, 45 (11), 751–766. 5
- Orlov, Dmitry**, “Frequent Monitoring in Dynamic Contracts,” *Journal of Economic Theory*, December 2022, 206, 105550. 8
- , **Andrzej Skrzypacz, and Pavel Zryumov**, “Persuading the Principal To Wait,” *Journal of Political Economy*, July 2020, 128 (7). 5
- Prat, Julien and Boyan Jovanovic**, “Dynamic Contracts When the Agent’s Quality Is Unknown: Dynamic Contracts,” *Theoretical Economics*, September 2014, 9 (3), 865–914. 6, 26
- Rodina, David**, “Information Design and Career Concerns,” Technical Report 220 November 2018. 30

Singer, Natasha, “In a Mood? Call Center Agents Can Tell,” *The New York Times*, October 2013. [2](#)

Smolin, Alex, “Dynamic Evaluation Design,” *American Economic Journal: Microeconomics*, 2021, *13* (4), 300–331. [5](#)

Treust, Maël Le and Tristan Tomala, “Persuasion with Limited Communication Capacity,” *Journal of Economic Theory*, November 2019, *184*, 104940. [6](#), [18](#), [19](#)

Winkler, Gerhard, “Extreme Points of Moment Sets,” *Mathematics of Operations Research*, November 1988, *13* (4), 581–587. [31](#)