

Monitoring with Rich Data*

Mira Frick Ryota Iijima Yuhta Ishii

First posted version: April 8, 2023

This version: December 19, 2023

Abstract

We consider moral hazard problems where a principal has access to rich monitoring data about an agent's action. Rather than focusing on optimal contracts (which are known to in general be complicated), we characterize the optimal rate at which the principal's payoffs can converge to the first-best payoff as the amount of data grows large. Our main result suggests a novel rationale for the widely observed binary wage schemes, by showing that such simple contracts achieve the optimal convergence rate. Notably, in order to attain the optimal convergence rate, the principal must set a lenient cutoff for when the agent receives a high vs. low wage. In contrast, we find that other common contracts where wages vary more finely with observed data (e.g., linear contracts) approximate the first-best at a highly suboptimal rate. Finally, we show that the optimal convergence rate depends only on a simple summary statistic of the monitoring technology. This yields a detail-free ranking over monitoring technologies that quantifies their value for incentive provision in data-rich settings and applies regardless of the agent's specific utility or cost functions.

*Frick: Yale University (mira.frick@yale.edu); Iijima: Yale University (ryota.iijima@yale.edu); Ishii: Pennsylvania State University (yxi5014@psu.edu). For helpful comments, we thank Nageeb Ali, Dirk Bergemann, Alex Bloedel, Hector Chade, Drew Fudenberg, George Georgiadis, Nima Haghpanah, Marina Halac, Johannes Hörner, Michihiro Kandori, Vijay Krishna, Giacomo Lanzani, Yingkai Li, Akihiko Matsui, Larry Samuelson, Ran Shorrer, Philipp Strack, Takuo Sugaya, Balázs Szentés, Alex Wolitzky, and numerous seminar audiences. We thank Haoning Chen for excellent research assistance. Frick gratefully acknowledges the financial support from a Sloan Research Fellowship.

1 Introduction

1.1 Motivation and Overview

A longstanding problem in contract theory is why wage schemes observed in practice are often “simple,” even though textbook models predict more complicated optimal contracts.¹ In this paper, we suggest a novel perspective on this problem. We consider a standard static moral hazard setting, where a principal (e.g., employer) designs a wage scheme to incentivize an agent (e.g., worker) whose costly action choice is not directly observable to the principal. As a key feature, we assume that the principal has access to *rich data* about the agent’s action, as may increasingly be the case in many workplaces (e.g., due to technological advances such as automated quality control systems or productivity tracking software). In the limit where the principal is able to perfectly monitor the agent’s action, there is no moral hazard and the principal can achieve her first-best payoff. Away from this limit, we are interested in which contracts allow the principal to efficiently exploit rich but imperfect monitoring data, and whether, and if so which, simple contracts may be enough for this purpose.

Importantly, to capture the efficient exploitation of rich data, we do not focus on which contracts are optimal (i.e., maximize the principal’s payoff). Rather, we analyze the *rate of convergence* of the principal’s payoff to the first-best as the amount of data grows rich, and we focus on the less demanding criterion of which contracts achieve the *optimal convergence rate*. The convergence rate to the first-best is a natural measure of efficiency in data-rich settings: Contracts with a higher convergence rate yield higher payoffs for the principal than contracts with a lower convergence rate whenever data is rich enough. Moreover, while optimal contracts necessarily converge to the first-best at the optimal rate, there may be simpler and more realistic classes of contracts that achieve the same optimal convergence rate.

Indeed, our main result is that the optimal convergence rate is achieved by a particular class of simple contracts that are widely used in practice: *binary* payment schemes, i.e., contracts with only two possible wage levels. Notably, in order to attain the optimal convergence rate, the principal must set a *maximally lenient* cutoff for when the agent receives the high vs. low wage. In contrast, we find that other common contracts where wages vary more finely with observed data (e.g., linear

¹See, e.g., the recent survey by [Georgiadis \(2022\)](#).

contracts) approximate the first-best at a highly suboptimal rate. Finally, we show that the optimal convergence rate depends only on a simple summary statistic of the monitoring technology. This yields a detail-free ranking over monitoring technologies that quantifies their value for incentive provision in data-rich settings and applies regardless of the agent’s specific utility or cost functions.

In our baseline model (Section 2), the principal has access to a monitoring technology that generates n i.i.d. signals about the agent’s (one-shot) action, where n parametrizes the richness of the principal’s data. The principal seeks to implement a target action by designing a contract that specifies a wage payment contingent on each realized signal sequence, subject to standard individual rationality (IR) and incentive compatibility (IC) constraints. The principal is risk-neutral while the agent is risk-averse. We study the optimal rate at which the principal’s implementation cost converges to the first-best (i.e., the observable action case) as the amount of data n grows large.

Our main result (Theorem 1) shows that, regardless of whether the principal optimizes over general contracts or binary contracts, her payoffs converge to the first-best at the same exponential rate. This optimal convergence rate is given by the Kullback-Leibler (KL) divergence between the signal distribution under the target action and the closest signal distribution under any less costly deviation. Thus, this rate depends only on the detectability of the hardest-to-detect deviation from the target action.

The binary contracts that achieve the optimal convergence rate take the form of statistical tests: The principal uses a score to partition signal sequences into a “pass” (high wage) and “fail” (low wage) region. In choosing how demanding it is to pass the test, the principal trades off the risk of *false negatives* (i.e., failing the test under the target action) and *false positives* (i.e., passing the test under a deviation). We show that false negatives become the dominant source of inefficiency as the amount of data grows large. Thus, in order to achieve the optimal convergence rate, the principal must make the pass–fail cutoff as lenient as possible subject to IC.

Binary wage schemes are frequently used in practice (e.g., [Murphy, 1999](#); [Prendergast, 1999](#)): A prominent and widely studied example are single-bonus contracts (a base wage, plus a fixed bonus if performance is sufficiently good), which are common in many professions.² Complementing existing explanations in the literature (see

²Another example highlighted by the literature (e.g., [Prendergast, 1999](#)) are jobs that feature

Section 1.2), Theorem 1 suggests that access to rich data about workers’ actions may provide a novel rationale for such contracts: Any benefit from using more complex, non-binary wage schemes has a negligible effect on the convergence to the first-best as the amount of data grows large. Our finding of leniency may also conform with evidence about binary contracts used in practice: For example, in the context of single-bonus contracts, Joseph and Kalwani (1998) observe that organizations tend to use bonuses to reward “acceptable” rather than “exceptional” performance.³

To illustrate that contracts with more fine-grained wage variation may approximate the first-best at a highly suboptimal rate, Proposition 1 considers another widely observed class of contracts, linear wage schemes. Under these contracts, we show that the convergence to the first-best is much slower than optimal, viz. subexponential. Thus, compared with binary contracts, linear contracts perform quite poorly at exploiting rich data. Proposition 2 further adds to the rationale for binary contracts in data-rich settings, by showing that, as n grows large, any sequence of optimal contracts approximates a binary contract of the maximally lenient form we identify above.

Our characterization of the optimal convergence rate in Theorem 1 also yields a ranking over monitoring technologies that quantifies their value for incentive provision (Corollary 1): In data-rich settings, monitoring technologies with higher KL-divergence between the signal distributions under target vs. non-target actions are more valuable to the principal than those with lower KL-divergence. Notably, this ranking is independent of the agent’s utility over money and specific cost function, providing detail-free guidance to a principal choosing between different monitoring technologies.

We conclude by discussing some variants of our baseline model. Section 4.1 introduces a severe form of limited liability under which IR cannot bind. We show that binary contracts continue to achieve the optimal convergence rate to the first-best, but the rate of convergence is slower, reflecting a tradeoff between risk aversion and rent extraction that is absent in the main model. Finally, beyond our baseline model with n i.i.d. signal draws about a one-shot action, Section 4.2 extends our main result to two formulations of rich monitoring data that allow for (i) general sequences of increasingly precise monitoring technologies (e.g., vanishing observation noise, serially

essentially flat wages but entail the threat of being fired in case of sufficiently poor performance.

³See also Johnston and Marshall (2016) (p. 152).

correlated data) or (ii) repeated but infrequent action adjustments.

1.2 Related Literature

In the standard moral hazard setting à la [Holmström \(1979\)](#) with a risk-averse agent, binary contracts are optimal only under restrictive, binary monitoring technologies (i.e., whose distribution of signal likelihood ratios has binary support).⁴ Several papers derive the optimality of binary contracts by enriching this setting. For example, [Georgiadis and Szentes \(2020\)](#) consider a principal who can flexibly design a monitoring technology at a cost.⁵ They show that the principal optimally chooses a binary monitoring technology, by establishing a connection with an information design problem. [Herweg, Müller, and Weinschenk \(2010\)](#) and [Lopomo, Rigotti, and Shannon \(2011\)](#) consider agents with non-expected utility preferences. Complementary to these papers, we revisit a standard moral hazard setting with an exogenous monitoring technology and risk-averse expected-utility agent, but we relax the criterion of exact optimality. Instead, we provide a rationale for binary contracts based on the idea that in data-rich settings, they allow the principal to approximate the first-best at the optimal rate. We also highlight the importance of a lenient high–low wage cutoff for achieving the optimal convergence rate, which is not a general feature of the optimal binary contracts in the aforementioned settings.⁶ Several papers identify natural forces that favor linear contracts (e.g., [Holmström and Milgrom, 1987](#); [Carroll, 2015](#); [Barron, Georgiadis, and Swinkels, 2020](#)). In contrast, we show that in our data-rich setting, linear contracts perform significantly worse than binary contracts.

Convergence rates as a measure of the performance of simple but suboptimal mechanisms have also been analyzed in other settings. For example, in the context of large markets, a classic literature studies the rate at which simple trading mechanisms approximate efficiency as the number of market participants grows large.⁷ While much

⁴With a risk-neutral agent, binary contracts can be optimal under various other assumptions (e.g., [Oyer, 2000](#); [Palomino and Prat, 2003](#); [Levin, 2003](#)), but these findings rely on exact risk neutrality.

⁵Specifically, the principal chooses when to stop observing the output of a diffusion process whose drift is the agent’s action, at a cost proportional to the principal’s stopping time.

⁶[Herweg, Müller, and Weinschenk \(2010\)](#) highlight that the set of signals that result in a high wage sometimes (but not always) contains some “bad” signals that are more indicative of lower effort.

⁷See, e.g., [Rustichini, Satterthwaite, and Williams \(1994\)](#); [Satterthwaite and Williams \(2002\)](#); [Hong and Shum \(2004\)](#); [Pakzad-Hurson \(2023\)](#). In other settings, some papers show that simple mechanisms can approximate the first-best, without analyzing convergence rates (e.g., [Radner, 1985](#), in the context of dynamic moral hazard problems where the discount factor grows large).

of this literature does not focus on identifying mechanisms that achieve the optimal convergence rate, [Satterthwaite and Williams \(2002\)](#) show that double auctions achieve the optimal worst-case convergence rate (i.e., when evaluated at the least favorable trading environment). In contrast, we find that binary contracts can be used to achieve the optimal convergence rate in all moral hazard environments we consider, not just the worst-case environment. The optimal convergence rate to the first-best is not directly comparable to a different notion of approximate optimality that is often analyzed in the computer science literature: worst-case guarantees, i.e., lower bounds on the performance ratio of simple vs. optimal mechanisms that are uniform with respect to some features of the environment (for a survey, see [Roughgarden and Talgam-Cohen, 2019](#)). In moral hazard environments, [Dütting, Roughgarden, and Talgam-Cohen \(2019\)](#) obtain such a guarantee for linear contracts.

Our ranking over monitoring technologies ([Corollary 1](#)) relates to the literature on the value of monitoring in moral hazard problems (e.g., [Holmström, 1979](#); [Gjesdal, 1982](#); [Singh, 1985](#); [Kim, 1995](#); [Jewitt, 2007](#)). Several papers derive partial orders over monitoring technologies in settings where the principal observes a single signal draw about the agent’s action. As we discuss in [Section 3.3](#), by quantifying the convergence rate to the first-best as the amount of data grows large, we obtain a ranking over monitoring technologies that is a completion of the orders in [Kim \(1995\)](#) and [Jewitt \(2007\)](#). As such, [Corollary 1](#) is an analog for incentive problems of [Moscarini and Smith \(2002\)](#): They quantify the value of rich data in learning settings, based on a different index that characterizes the rate at which a decision-maker who observes repeated i.i.d. signal draws from an information structure learns the state; this yields a completion of [Blackwell’s \(1951\)](#) order over information structures.⁸ Comparisons of monitoring structures have also been studied in the context of repeated games (e.g, [Kandori, 1992](#)). Some recent papers analyze how the monitoring structure affects the convergence rate of equilibrium payoffs to the efficient frontier as players become patient ([Hörner and Takahashi, 2016](#); [Sugaya and Wolitzky, 2023a,b](#)).

⁸Convergence rates are also analyzed to measure efficiency in other learning settings, including social learning (e.g., [Vives, 1993](#); [Harel, Mossel, Strack, and Tamuz, 2021](#)), higher-order beliefs (e.g., [Frick, Iijima, and Ishii, 2023a](#)), and misspecified learning (e.g., [Frick, Iijima, and Ishii, 2023b](#)).

2 Model

Environment. Consider the following static moral hazard setting. There is a principal (“she”) and an agent (“he”). The agent chooses a one-shot action from a finite action set A .⁹ The principal does not observe the agent’s action, but has access to a monitoring technology μ : This specifies a signal space X , assumed to be a subset of a Euclidean space, and, for each chosen action a , a distribution $\mu_a \in \Delta(X)$ over signals, where $\Delta(X)$ denotes the set of Borel probability measures over X .

For all distinct actions $a \neq a'$, we impose the identification assumption that the signal distributions μ_a and $\mu_{a'}$ are different. We also assume that $\mu_{a'}$ is absolutely continuous with respect to μ_a (with corresponding Radon-Nikodym derivative $\frac{d\mu_{a'}}{d\mu_a}$), which implies that no signals perfectly reveal the chosen action. Finally, we impose the regularity condition that the moment-generating functions of signal log-likelihood ratios are well-defined, i.e., $\int \left(\frac{d\mu_{a'}}{d\mu_a}(x)\right)^\lambda d\mu_a(x) < \infty$ for all $\lambda > 0$.

To capture that the principal has access to rich data about the agent’s action, we assume that after the agent chooses his action a , the principal observes n i.i.d. draws of signals $x^n = (x_1, \dots, x_n)$ from μ_a . Here, n parametrizes the richness or precision of the principal’s data, and we will be interested in settings where n is large. Let \mathbb{P}_a denote the distribution over signal sequences x^n conditional on action a ; let \mathbb{E}_a and Var_a denote the corresponding expectation and variance operators.

Our assumption that the agent chooses a one-shot action which then generates multiple i.i.d. signals can be seen as an approximation of some real-world settings with rich monitoring data, such as automated quality control or teaching evaluations.¹⁰ However, as Section 4.2 discusses, the analysis generalizes readily to less stylized settings: First, moving beyond the i.i.d. signal formulation, we can consider general sequences μ^n of increasingly precise monitoring technologies (e.g., one-shot observation of the action perturbed by some vanishing noise, or repeated but serially correlated signals). Second, moving beyond one-shot action choice, we can allow the agent to repeatedly adjust his action, as long as action adjustments are infrequent relative to the number n of signals.

⁹We do not impose any order assumptions on A ; for example, actions may be multi-dimensional.

¹⁰In the former case, a represents a factory worker’s assembly of a widget, and upon completion n signals are generated by a machine that repeatedly scans the widget for various possible errors. In the latter, a represents an instructor’s effort, and signals take the form of reviews submitted by the n students at the end of the course.

Payoffs. The principal is risk-neutral and seeks to implement some target action $a^* \in A$ by designing a contract, i.e., a wage scheme $w : X^n \rightarrow [\underline{w}, \infty)$ that specifies a payment $w(x^n)$ contingent on each realized signal sequence x^n .¹¹ To rule out “shoot the agent” arguments à la Mirrlees (1999) (see Remark 1), we impose a lower bound $\underline{w} \in \mathbb{R}$ on wages, which can be arbitrarily low.

The agent’s payoffs are additively separable, consisting of a consumption utility $u : [\underline{w}, \infty) \rightarrow \mathbb{R}$ over money minus a cost $c : A \rightarrow \mathbb{R}$ for each action. The agent has an outside option, whose payoff we normalize to 0. We assume the following:

Assumption 1.

1. u is twice continuously differentiable with $u'(w) > 0$ and $u''(w) < 0$ for all $w \in [\underline{w}, \infty)$ and $\lim_{w \rightarrow \infty} u'(w) = 0$;
2. $c(A) \subseteq \text{int}(u([\underline{w}, \infty))$;
3. $c(a^*) > c(a)$ for some $a \in A$ and $c(a^*) \neq c(a)$ for all $a \in A \setminus \{a^*\}$.

The first condition requires the agent to be strictly risk-averse, which will play an essential role in our analysis.¹² Throughout, we let $h := u^{-1}$. The second condition requires the consumption utility range to be rich enough that (under perfect monitoring) suitable wage payments can make the the agent indifferent between each action and the outside option; Section 4.1 considers the case when this condition is violated. To avoid trivialities, the third condition assumes that the target action a^* is not the least costly action.

Second-best problem. In the optimal contract, the principal chooses a wage scheme w to minimize the implementation cost of a^* ,¹³

$$C_n^{\text{SB}}(\mu, u, c, a^*) = \inf_{w: X^n \rightarrow [\underline{w}, \infty)} \mathbb{E}_{a^*}[w(x^n)], \quad (1)$$

¹¹Our results extend straightforwardly to the case of a risk-averse principal. As is common in the literature, we focus on the principal’s incentive design problem and do not explicitly model her optimization over the target action a^* .

¹²Assumption 1.1 is satisfied under standard families of utility functions, e.g., CARA or CRRA. If instead u is linear, standard arguments imply that the principal can achieve the first-best provided there is an (IC) contract that satisfies (IR) with equality. Under Assumption 1.2, this is the case provided either the range of u is large enough or n is large enough.

¹³At all large enough n , the inf in (1) is attained by some wage scheme; see Appendix A.1.2.

subject to standard incentive compatibility (IC) and individual rationality (IR) constraints:

$$\mathbb{E}_{a^*}[u(w(x^n))] - c(a^*) \geq \mathbb{E}_a[u(w(x^n))] - c(a), \quad \forall a \in A, \quad (\text{IC})$$

$$\mathbb{E}_{a^*}[u(w(x^n))] - c(a^*) \geq 0. \quad (\text{IR})$$

We refer to the induced minimal cost $C_n^{\text{SB}}(\mu, u, c, a^*)$ as the *second-best cost*. As is standard, it is convenient to rewrite the principal's problem in utility terms. That is, instead of choosing wage schemes w , the principal equivalently chooses maps $v : X^n \rightarrow u([\underline{w}, \infty))$ from signal sequences to consumption utilities,

$$C_n^{\text{SB}}(\mu, u, c, a^*) = \inf_{v: X^n \rightarrow u([\underline{w}, \infty))} \mathbb{E}_{a^*}[h(v(x^n))]$$

subject to the IC and IR constraints

$$\mathbb{E}_{a^*}[v(x^n)] - c(a^*) \geq \mathbb{E}_a[v(x^n)] - c(a), \quad \forall a \in A, \quad (\text{IC})$$

$$\mathbb{E}_{a^*}[v(x^n)] - c(a^*) \geq 0. \quad (\text{IR})$$

Convergence rate analysis. As is well-known, the second-best problem gives rise to wage schemes that depend finely on realized signals, in a potentially complicated manner that does not in general resemble contracts observed in practice (e.g., binary or linear contracts). Thus, instead of focusing on the second-best cost for a given n , we analyze the convergence rate of the second-best cost to the first-best as n grows large. Formally, let $C^{\text{FB}}(u, c, a^*)$ denote the *first-best cost*, i.e., the minimal cost of implementing a^* under perfect monitoring. Under Assumption 1.2, this is the same as the minimal implementation cost of a^* when only (IR) is imposed and is given by

$$C^{\text{FB}}(u, c, a^*) = h(c(a^*)).$$

Clearly, with infinitely many signals, the second-best cost coincides with the first-best, $C_\infty^{\text{SB}}(\mu, u, c, a^*) = C^{\text{FB}}(u, c, a^*)$, as $n = \infty$ corresponds to perfect monitoring given the identification assumption on μ . Instead, we are interested in environments where the amount n of data is large but finite. To understand such settings, we will characterize the rate at which the second-best cost $C_n^{\text{SB}}(\mu, u, c, a^*)$ converges to the first-best $C^{\text{FB}}(u, c, a^*)$ as $n \rightarrow \infty$, i.e., the optimal rate at which the principal can

approximate the first-best.

This will allow us to address two questions: First, are there simple classes of contracts that attain this same optimal convergence rate? As noted, achieving the optimal convergence rate to the first-best is a natural criterion for efficient exploitation of rich data: Whenever n is large enough, contracts with a higher convergence rate yield a lower implementation cost than contracts with a lower convergence rate. Second, which monitoring technologies are more valuable for incentive provision in data-rich settings, i.e., how does the optimal rate of convergence vary across μ ?

3 Analysis

3.1 Optimal Rate of Convergence

Our main result characterizes the rate at which the second-best cost converges to the first-best as the amount of data grows large. Moreover, we show that this optimal rate of convergence can be achieved by a particular class of simple contracts: **binary contracts**, i.e., wage schemes w with $|w(X^n)| = 2$. Let $C_n^{\text{bin}}(\mu, u, c, a^*)$ denote the second-best cost when, instead of optimizing over all IC and IR wage schemes, the principal is restricted to optimizing over IC and IR binary contracts. Denote by $A^-(c, a^*) := \{a \in A : c(a) < c(a^*)\}$ the set of actions that are less costly to the agent than the target action a^* , which is nonempty by Assumption 1.3. For all $\nu, \nu' \in \Delta(X)$, let $\text{KL}(\nu, \nu')$ denote the Kullback-Leibler (KL) divergence of ν relative to ν' , i.e.,

$$\text{KL}(\nu, \nu') := \begin{cases} \int \log\left(\frac{d\nu}{d\nu'}(x)\right) d\nu(x) & \text{if } \nu \text{ is absolutely continuous w.r.t. } \nu' \\ \infty & \text{otherwise.} \end{cases}$$

Theorem 1. *Under both general and binary contracts, the second-best cost converges to the first-best exponentially at rate $\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})$: We have*

$$C_n^{\text{SB}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = \exp\left[-\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})n + o(n)\right]; \quad (2)$$

$$C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = \exp\left[-\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})n + o(n)\right]. \quad (3)$$

By (2), when the principal optimizes over all IC and IR wage schemes, she ap-

proximates the first-best at an exponential rate given by $\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})$.¹⁴ To interpret, note that $\text{KL}(\mu_a, \mu_{a^*})$ is a statistical measure that quantifies how dissimilar the signal distribution under a deviation to action $a \neq a^*$ is from the signal distribution under the target action a^* . Thus, (2) shows that, for any given monitoring technology μ , the optimal rate of convergence to the first-best depends only on a simple statistic: the detectability $\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})$ of the hardest-to-detect deviation to a less costly action.

Crucially, (3) implies that the principal can achieve this same optimal rate of convergence using binary contracts. Thus, Theorem 1 offers a novel rationale for binary contracts, which (as discussed in the Introduction) are widely observed in practice: While for any finite n , binary contracts are in general suboptimal, Theorem 1 shows that binary contracts are an effective way to exploit rich data. As the amount of data n grows large, any benefit from using more general, non-binary contracts has at most a second-order effect on the convergence to the first-best: The only potential difference between (2) and (3) is in the terms $o(n)$, but as these are sublinear (i.e., $\lim_{n \rightarrow \infty} \frac{o(n)}{n} = 0$) they become negligible as $n \rightarrow \infty$; in terms of the principal’s payoffs, this can be quantified, for instance, by saying that at large n , the benefit of optimizing over general vs. binary contracts is smaller than the benefit of having access to an arbitrarily small fraction ε of additional signals.¹⁵ As such, in data-rich settings, the benefit of optimizing over general contracts may plausibly be outweighed by other (unmodeled) benefits of using simple, binary contracts.

Section 3.2 below illustrates the idea behind Theorem 1 and sheds light on the structure of the binary contracts that achieve the optimal convergence rate. These contracts take the form of statistical tests, where the principal uses a score to partition signal sequences x^n into a “pass” (high wage) and “fail” (low wage) region. Importantly, as Section 3.2 will formalize, to achieve the optimal convergence rate, the pass–fail cutoff must be chosen in the *maximally lenient* way subject to (IC). As noted in the Introduction, such leniency may also be in line with binary contracts observed in practice. Figure 1 illustrates the importance of leniency in an example with binary actions and signals. We compare binary contracts (with binding (IC) and (IR)) that pay a high wage if and only if the fraction of high signals exceeds a cutoff. Here, a lenient binary contract, whose cutoff 0.31 is only slightly above the expected

¹⁴By the assumptions on μ , $\text{KL}(\mu_a, \mu_{a^*}) \in (0, \infty)$ for all $a \neq a^*$, so this rate is positive and finite.

¹⁵Formally, for any $\varepsilon > 0$, there is N such that for all $n \geq N$, $C_n^{\text{SB}}(\mu, u, c, a^*) > C_{(1+\varepsilon)n}^{\text{bin}}(\mu, u, c, a^*)$.




Figure 1: Implementation cost of $a^* = 1$ as a function of n when $A = \{0, 1\} = X$, $\mu_0(0) = \mu_1(1) = 0.7$, $c(0) = 0$, $c(1) = 2$, $u(w) = \log w$. Dashed: first-best cost. Blue: cost under the “lenient” binary contract that pays a high wage iff the fraction of high signals exceeds 0.31. Red: cost under the “strict” binary contract that pays a high wage iff the fraction of high signals exceeds 0.6. Yellow: cost under the optimal utility-linear contract. Wages in the three contracts are pinned down by requiring (IC) and (IR) to bind.

fraction 0.3 of high signals under the *low* (non-target) action, already comes close to the first-best at small values of n . In contrast, under a strict binary contract, whose cutoff 0.6 is closer to the expected fraction 0.7 of high signals under the *high* (target) action, convergence is substantially slower.

Section 3.2 also illustrates that, unlike binary contracts, contracts where the wage varies finely with the signal realizations may approximate the first-best at a highly suboptimal rate. Below, we formalize this point for another class of frequently observed wage schemes, **linear contracts**: Here, there exists $b_n : X \rightarrow [\underline{w}, \infty)$ such that $w(x^n) = \sum_{i=1}^n b_n(x_i)$ for all $x^n \in X^n$.¹⁶ Let $C_n^{\text{lin}}(\mu, u, c, a^*)$ denote the second-best cost when, instead of optimizing over all IC and IR contracts, the principal is restricted to optimizing over IC and IR linear contracts. The following result shows that the convergence rate to the first-best under linear contracts is subexponential.

Proposition 1. *Suppose that $\inf_v h''(v) > 0$.¹⁷ Under linear contracts, the second-best cost converges to the first-best subexponentially: There exists a constant $K > 0$*

¹⁶That is, we define linear contracts as linear in empirical frequencies, where the empirical frequency $\nu_{x^n} \in \Delta(X)$ associated with signal sequence x^n is given by $\nu_{x^n}(x) := \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{\{x_i=x\}}$ for all $x \in X$. If $X \subseteq \mathbb{R}$, then one example of linear contracts are wage schemes w that are linear in average signals, i.e., there exist $\alpha_n > 0$, $\beta_n \in \mathbb{R}$ such that $w(x^n) = \alpha_n \frac{1}{n} \sum_{i=1}^n x_i + \beta_n$ for all $x^n \in X^n$.

¹⁷This is a slight strengthening of the assumption that $\lim_{w \rightarrow \infty} u'(w) = 0$. We conjecture that it can be relaxed.

such that

$$C_n^{\text{lin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) \geq \frac{1}{n}K + o\left(\frac{1}{n}\right). \quad (4)$$

Moreover, (4) holds with equality if $\mu_{a^*} \notin \text{co}\{\mu_a : a \neq a^*\}$ and \underline{w} is sufficiently low.¹⁸

Note that, in the perfect monitoring limit where $n = \infty$, both binary and linear contracts can achieve the first-best cost (the latter claim requires the regularity assumption on μ in the “moreover” part of Proposition 1).¹⁹ However, Proposition 1 shows that, away from this limit, linear contracts are less effective at exploiting rich data than binary contracts, because they approximate the first-best at a much slower rate. Thus, regardless of the agent’s utility and cost function and the principal’s monitoring technology and target action, in the current setting the principal is always better off using binary contracts when the amount of data is rich enough: That is, there exists N such that for all $n \geq N$,

$$C_n^{\text{bin}}(\mu, u, c, a^*) < C_n^{\text{lin}}(\mu, u, c, a^*).$$

Proposition 1 remains valid under utility-linear contracts, which instead require the utility contract $v = u \circ w$ to be linear. Reflecting this, in Figure 1, the lenient binary contract approximates the first-best significantly faster than does the optimal utility-linear contract.²⁰

3.2 Illustration of Theorem 1

For expositional simplicity, we focus on binary action sets, $A = \{0, 1\}$ with $a^* = 1$.

Step 1: Variance minimization. Without loss of optimality, consider any uniformly bounded sequence of contracts v_n with binding (IR) (i.e., $\mathbb{E}_1[v_n(x^n)] =$

¹⁸Here co denotes the convex hull operator.

¹⁹The first-best is achieved by any w such that $w(x^\infty) = h(c(a^*))$ for x^∞ whose empirical frequency is μ_{a^*} and $w(x^\infty) \leq h(c(a))$ for x^∞ whose empirical frequency is μ_a for some $a \neq a^*$. Clearly, such w can always be made binary, and can be made linear under the regularity assumption in Proposition 1.

²⁰Figure 1 focuses on utility-linear contracts for ease of computation. While here the lenient binary contract outperforms the linear contract at all n , under different parameters linear contracts can outperform the optimal binary contract at small n . Moreover, while the linear contract outperforms the strict binary contract in the depicted range of n , the strict binary contract eventually does better as it also converges to the first-best exponentially (albeit at a lower rate than the lenient contract).

$c(1)$). Then

$$\underbrace{\mathbb{E}_1[h(v_n(x^n))] - C^{\text{FB}}(u, c, a^*)}_{\text{inefficiency}} = \underbrace{\mathbb{E}_1[h(v_n(x^n))] - h(\mathbb{E}_1[v_n(x^n)])}_{\text{Jensen-inequality gap}} \approx \text{Var}_1[v_n(x^n)], \quad (5)$$

where “ \approx ” denotes convergence to 0 at the same exponential rate. That is, the efficiency loss relative to the first-best is equal to the Jensen-inequality gap of the convex function h with respect to the random variable $v_n(x^n)$; by Taylor expansion arguments (e.g., [Liao and Berg, 2019](#)), the latter has the same exponential decay rate as the variance of $v_n(x^n)$. Thus, to obtain the optimal convergence rate to the first-best, the principal must choose v to maximize the decay rate of the agent’s utility variance $\text{Var}_1[v_n(x^n)]$ subject to (IC). This reflects the familiar tradeoff in moral hazard problems (e.g., [Laffont and Martimort, 2009](#)) of seeking to make contracts as “safe” for the risk-averse agent as possible without violating incentive compatibility.

Step 2: Binary test contracts. Next, we construct simple binary “test” contracts under which $\text{Var}_1[v_n(x^n)]$ decays at rate $\text{KL}(\mu_0, \mu_1)$. Denote by

$$L_n := \frac{1}{n} \sum_{i=1}^n \log \frac{d\mu_1}{d\mu_0}(x_i)$$

the *log-likelihood score* of the realized signal sequence x^n , defined as the average log-likelihood ratio of action 1 vs. 0. Note that the expected scores under both actions, $\mathbb{E}_1[L_n] > 0$ and $\mathbb{E}_0[L_n] < 0$, are independent of n . By standard arguments (e.g., [Laffont and Martimort, 2009](#)), it is without loss to restrict to contracts v_n that are nondecreasing functions of L_n . Consider a sequence of binary contracts of the form

$$v_n(x^n) = \begin{cases} v_n^+ & \text{if } L_n \geq \gamma \\ v_n^- & \text{if } L_n < \gamma. \end{cases} \quad (6)$$

Here γ is some threshold with $\gamma \in (\mathbb{E}_0[L_n], \mathbb{E}_1[L_n])$, and the utility payments $v_n^+ > v_n^-$ are pinned down by requiring (IR) and (IC) to bind, which can be ensured at all large enough n .²¹ That is, the agent is rewarded with the higher utility v_n^+ whenever his score is above the threshold (i.e., he “passes the test”), and he is punished with the lower utility v_n^- if his score is below the threshold (i.e., he “fails the test”).

²¹That is, $v_n^+ := \frac{c(1)\mathbb{P}_0[L_n < \gamma] - c(0)\mathbb{P}_1[L_n < \gamma]}{\mathbb{P}_1[L_n \geq \gamma] - \mathbb{P}_0[L_n \geq \gamma]}$ and $v_n^- := \frac{c(0)\mathbb{P}_1[L_n \geq \gamma] - c(1)\mathbb{P}_0[L_n \geq \gamma]}{\mathbb{P}_1[L_n \geq \gamma] - \mathbb{P}_0[L_n \geq \gamma]}$, which converge to $c(1)$ and $c(0)$, respectively, as $n \rightarrow \infty$. By Assumption 1.2, $v_n^+, v_n^- \in u(\underline{w}, \infty)$ for large enough n .




Figure 2: Contour curves of $\text{Var}_1[v_n(x^n)]$ under contract (6). Given γ , n determines the probabilities of false positives and false negatives, which jointly determine the utility variance.

Clearly, any such sequence approximates the first-best as $n \rightarrow \infty$, because the test becomes arbitrarily precise, i.e., $\lim_{n \rightarrow \infty} \mathbb{P}_1[L_n \geq \gamma] = 1$ and $\lim_{n \rightarrow \infty} \mathbb{P}_0[L_n \geq \gamma] = 0$, by the law of large numbers and the fact that $\gamma \in (\mathbb{E}_0[L_n], \mathbb{E}_1[L_n])$. However, we need to analyze which choice of the threshold γ achieves the fastest convergence rate.

To this end, observe that for any fixed n , the optimal choice of γ must trade-off reducing two mistakes: On the one hand, the probability of **false positives**, $\mathbb{P}_0[L_n \geq \gamma]$, is lower the *higher* γ ; reducing this is relevant for (IC), as it makes action 0 less attractive to the agent. On the other hand, the probability of **false negatives**, $\mathbb{P}_1[L_n < \gamma]$, is lower the *lower* γ ; reducing this is relevant for both (IC) and (IR), as it makes action 1 more attractive to the agent. Crucially, we show that as n grows large, false negatives are the dominant force affecting the agent's utility variance, because

$$\text{Var}_1[v_n(x^n)] \approx \mathbb{P}_1[L_n < \gamma].$$

Thus, as $n \rightarrow \infty$, to minimize utility variance, the principal should optimally choose a **maximally lenient** threshold $\gamma \searrow \mathbb{E}_0[L_n]$.²² To illustrate, Figure 2 shows the contour curves of $\text{Var}_1[v(x^n)]$ as a function of the probabilities of false positives and false negatives. Observe that at large values of $\text{Var}_1[v_n(x^n)]$ (i.e., at small n), false posi-

²²For simplicity, (6) considered sequences of contracts with a fixed threshold γ . In this case, the optimal decay rate of $\text{Var}_1[v_n(x^n)]$ is approximated by choosing $\gamma \in (\mathbb{E}_0[L_n], \mathbb{E}_1[L_n])$ arbitrarily close to $\mathbb{E}_0[L_n]$, but $\gamma = \mathbb{E}_0[L_n]$ is ruled out by (IC). Exactly attaining the optimal decay rate instead requires using an appropriate sequence of thresholds $\gamma_n \in (\mathbb{E}_0[L_n], \mathbb{E}_1[L_n])$ with $\gamma_n \rightarrow \mathbb{E}_0[L_n]$.




Figure 3: Contract (6) with binary signals. The space of empirical signal frequencies $\Delta(X)$ is divided into two regions, “pass” ($L_n \geq \gamma$) and “fail” ($L_n < \gamma$). As γ approaches the maximally lenient threshold $\mathbb{E}_0[L_n]$, the cutoff between the two regions approaches μ_0 .

tives and false negatives have a fairly symmetric effect on utility variance; however, as $\text{Var}_1[v_n(x^n)]$ approaches 0 (i.e., at large n), the contour curves become arbitrarily steep, illustrating the dominant effect of false negatives in this region. Intuitively, when monitoring is near-perfect, (IC) becomes less important than (IR) (as captured by a vanishing shadow value in the principal’s optimization problem). As a result, reducing false positives becomes relatively unimportant, as this only affects (IC).

Finally, we observe that as γ approaches the maximally lenient threshold $\mathbb{E}_0[L_n]$, the probability of false negatives, and hence the utility variance, decays exponentially at rate $\text{KL}(\mu_0, \mu_1)$. This observation essentially corresponds to Stein’s lemma, a classical result in hypothesis testing (e.g., Cover and Thomas, 1999) and can be proved using Sanov’s theorem from large deviation theory. Sanov’s theorem states that for any set $D \subseteq \Delta(X)$ that is equal to the closure of its interior, the probability $\mathbb{P}_1[\nu_n \in D]$ of observing an empirical frequency ν_n in D conditional on action 1 decays exponentially at rate $\inf_{\nu \in D} \text{KL}(\nu, \mu_1)$ (i.e., it decays faster the farther D is from the theoretical signal distribution μ_1 under action 1). Figure 3 illustrates this in the case of binary signals: Applying Sanov’s theorem to the “fail” region D , the probability $\mathbb{P}_1[\nu_n \in D]$ of false negatives under the maximally lenient threshold decays at rate $\inf_{\nu \in D} \text{KL}(\nu, \mu_1) = \text{KL}(\mu_0, \mu_1)$.

Step 3: General contracts cannot do better. It remains to show that general, non-binary contracts cannot converge to the first-best faster than at the exponential rate $\text{KL}(\mu_0, \mu_1)$. To illustrate the idea, consider sequences of contracts of the form

$$v_n(x^n) = f(L_n),$$

for some function f that is continuously differentiable at $\mathbb{E}_1[L_n]$. For example, this includes linear contracts where the wage is a fixed affine function of the score L_n

at all n .²³ If $f'(\mathbb{E}_1[L_n]) \neq 0$ (i.e., there is some wage variation around $\mathbb{E}_1[L_n]$, as is the case for linear contracts), then the delta method implies that $\text{Var}_1[v_n(x^n)]$ is of order $\frac{1}{n}f'(\mathbb{E}_1[L_n])$, as L_n is the sample average of n i.i.d. draws of $\log \frac{d\mu_1(x_i)}{d\mu_0(x_i)}$. Thus, the utility variance, and hence, by (5), the efficiency loss relative to the first-best, vanishes subexponentially, i.e., more slowly than in Step 2.

Therefore, to achieve the optimal convergence rate, f must be flat around $\mathbb{E}_1[L_n]$, as is the case for the above binary test contracts. Moreover, as we formalize in Appendix A.2.2, the convergence rate is higher the greater a flat region f displays around $\mathbb{E}_1[L_n]$. At the same time, to satisfy (IC), the flat region cannot extend below $\mathbb{E}_0[L_n]$. Since the maximally lenient binary contract maximizes the flat region subject to this constraint, no other contract can outperform its convergence rate $\text{KL}(\mu_0, \mu_1)$.

Non-binary actions. The above logic extends readily to non-binary action sets. Specifically, we use binary contracts of the following form: For each $a \in A^-(c, a^*)$, consider the log-likelihood score $L_n(a) := \frac{1}{n} \sum_{i=1}^n \log \frac{d\mu_{a^*}}{d\mu_a}(x_i)$ and a threshold $\gamma(a)$. Award a high utility payment v_n^+ if $L_n(a) \geq \gamma(a)$ for all $a \in A^-(c, a^*)$ and a low payment v_n^- otherwise, where v_n^+ and v_n^- are again pinned down by requiring (IR) and (IC) to bind. Here, the optimal rate of convergence is again achieved by setting each threshold $\gamma(a) \searrow \mathbb{E}_a[L_n(a)]$ to be maximally lenient.

Remark 1 (Contrast with “shoot the agent” contracts). In the setting where $n = 1$ and there is no lower bound on wages, Mirrlees (1999) shows that, even when the monitoring technology is highly imprecise, the first-best can (essentially) be achieved as long as signal likelihood ratios are unbounded. For this purpose, he uses binary contracts that severely punish the agent at signals with very low log-likelihood scores (i.e., he considers the limit where both the low wage and the score cutoff become arbitrarily negative). This argument does not apply in our setting at any n , as we imposed a lower bound \underline{w} on wages. As a result, the binary contracts we constructed above are qualitatively quite different from Mirrlees’ “shoot the agent” contracts. In particular, Mirrlees’ contracts punish deviations extremely severely but with probability close to 0; in contrast, our binary contracts punish deviations with a moderately low wage but with probability close to 1 as $n \rightarrow \infty$. Indeed, note that our binary contracts use a low wage $v_n^- \rightarrow c(0) > \underline{w}$ that remains bounded away from \underline{w} at large

²³The logic extends to general sequences of linear contracts $w_n(x^n) = \sum_{i=1}^n b_n(x_i)$, where the functions $b_n : X \rightarrow [\underline{w}, \infty)$ can vary with n .

n (see footnote 21).²⁴ Thus, even for arbitrarily low values of \underline{w} , our binary contracts are not an approximation of Mirrlees’ construction. \blacktriangle

Remark 2 (Limit of optimal contracts). While maximally lenient binary contracts achieve the optimal convergence rate to the first-best, the illustration above shows that the same is true for some other contracts: To converge at the optimal rate, we saw that wage schemes must become flat at likely signal realizations (i.e., empirical signal frequencies ν within distance $\text{KL}(\mu_0, \mu_1)$ of μ_1). But beyond maximally lenient binary contracts, this allows for contracts that display additional wage variation at very rare empirical signal frequencies (farther than distance $\text{KL}(\mu_0, \mu_1)$ from μ_1).

However, even under the general second-best problem, the benefit to introducing such additional wage variation vanishes as $n \rightarrow \infty$: Any convergent sequence of optimal contracts has as its limit a maximally lenient binary contract.

To state this formally, again assume for simplicity that $A = \{0, 1\}$ with $a^* = 1$ (Appendix A.4 extends the result to non-binary action sets). As noted, we can identify optimal contracts with functions v_n^* of log-likelihood scores L_n .²⁵ Endow the space of such functions with the topology of weak convergence (under the L^2 -norm).

Proposition 2. *For any weakly convergent sequence of optimal contracts (v_n^*) ,*

$$\lim_{n \rightarrow \infty} v_n^*(L) = \begin{cases} c(1) & \text{if } L > \gamma^*, \\ u(\underline{w}) & \text{if } L < \gamma^*, \end{cases}$$

where $\gamma^* := \mathbb{E}_0[L_n]$ is the maximally lenient threshold.

Proposition 2 adds to the rationale for maximally lenient binary contracts in data-rich settings. We note that Proposition 2 is neither implied by nor implies Theorem 1.²⁶ \blacktriangle

3.3 Ranking over Monitoring Technologies

Theorem 1 immediately implies the following ranking over monitoring technologies:

²⁴In contrast, the lower bound \underline{w} does bind at the general second-best contract at large n .

²⁵More precisely, we treat v_n^* as a function on the domain $\text{co}\{\log \frac{d\mu_1}{d\mu_0}(x) : x \in X\}$.

²⁶To see the latter point, note that Proposition 2 is silent about the rate of convergence of v_n^* to its binary limit contract. In particular, it does not rule out that this convergence is slower than the rate at which v_n^* approximates the first-best.

Corollary 1. *Take any nonempty $A^- \subsetneq A$ and any two monitoring technologies μ and μ' with $\min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*}) > \min_{a \in A^-} \text{KL}(\mu'_a, \mu'_{a^*})$. Then, for any u and c with $A^-(c, a^*) = A^-$, there exists N such that for all $n \geq N$,*

$$C_n^{\text{SB}}(\mu, u, c, a^*) \leq C_n^{\text{bin}}(\mu, u, c, a^*) < C_n^{\text{SB}}(\mu', u, c, a^*).$$

Holding fixed a target action a^* , Corollary 1 yields a (generically complete) ranking over monitoring technologies that quantifies their value to the principal in data-rich settings: Whenever the amount of data n is rich enough, monitoring technologies μ are more valuable to the principal the higher is the index $\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})$. This is because, by Theorem 1, this index captures the optimal convergence rate to the first-best under μ . Moreover, since the optimal convergence rate is achieved using binary contracts, the principal is better off using a superior monitoring technology μ along with a binary contract than optimizing over general contracts under any (even slightly) lower-ranked monitoring technology μ' .

Notably, Corollary 1 provides the principal with *detail-free* guidance for selecting between monitoring technologies: The ranking is independent of the agent's utility function u and depends on his cost function c only through the set $A^-(c, a^*)$ of actions that are less costly than the target action.

The following example illustrates a qualitative implication of the ranking:

Example 1 (Precise bad vs. good news). *Assume binary actions and signals, $A = X = \{0, 1\}$ with $a^* = 1$, $c(1) > c(0)$. Consider monitoring technologies μ and μ' with*

$$\mu_0(0) = 0.8, \mu_1(1) = 0.99; \quad \mu'_0(0) = 0.99, \mu'_1(1) = 0.8.$$

Under both μ and μ' , signal 1 is “good news” (more indicative of the target action $a^ = 1$) and signal 0 is “bad news” (more indicative of the deviation $a = 0$). The only difference is that the error probabilities under the two actions are flipped across μ and μ' : μ provides precise bad news and relatively less precise good news (i.e., $\frac{\mu_0(0)}{\mu_1(0)} > \frac{\mu_1(1)}{\mu_0(1)}$), while μ' provides precise good news and relatively less precise bad news.*

Observe that $\text{KL}(\mu_0, \mu_1) \approx 3.2 > \text{KL}(\mu'_0, \mu'_1) \approx 1.5$. Thus, by Corollary 1, whenever data is sufficiently rich, then regardless of the agent's utility and cost functions, μ is strictly more valuable to the principal than μ' . Intuitively, as we saw in Section 3.2, the principal seeks to design tests under which false negatives decay as fast as possible, and for this monitoring technologies that provide precise bad news about the agent's action are more valuable than ones that provide precise good news. \blacktriangle

The ranking in Corollary 1 relies on the principal observing rich data, which reduces the comparison of monitoring technologies to a comparison of the corresponding convergence rates to the first-best. If the principal only observes a single signal draw ($n = 1$), then Kim (1995) and Jewitt (2007) show that requiring μ to be more valuable to the principal than μ' regardless of the agent's preferences yields an order over monitoring technologies that extends Blackwell's (1951) order. However, their order is more conservative than our ranking:²⁷ For instance, the monitoring technologies in Example 1 are incomparable based on the $n = 1$ order.

As noted, Corollary 1 can be viewed as an analog for incentive problems of Moscarini and Smith's (2002) rich-data ranking over information structures in learning settings (see Section 1.2). However, reflecting the difference between the value of information for incentive provision vs. learning, our ranking is different from their ranking: For instance, the two monitoring structures in Example 1 are equally valuable in the learning setting of Moscarini and Smith (2002).

4 Discussion

4.1 Severe Limited Liability

Our main model (Assumption 1.2) assumed that $c(A) \subseteq \text{int}(u([\underline{w}, \infty)))$, which ensured that the (IR) constraint could be made to bind at large enough n . In this section, we suppose this assumption is violated, i.e., limited liability is so severe that (IR) does not bind. We show that binary contracts continue to achieve the optimal convergence rate to the first-best. However, the analysis highlights a tradeoff between rent extraction and risk aversion that was absent in the main setting.

For expositional simplicity, we focus on binary actions, $A = \{0, 1\}$ with $a^* = 1$ and $c(0) < c(1)$; Appendix A.5 considers non-binary A . The assumption $c(A) \subseteq \text{int}(u([\underline{w}, \infty)))$ implies $u(\underline{w}) < c(0)$. Departing from this, we impose the severe limited liability restriction that $u(\underline{w}) \geq c(0)$; thus, even if taking action 0 leads the agent to be punished with the lowest wage \underline{w} , he prefers this to his outside option of 0. This,

²⁷Based on the $n = 1$ order, μ dominates μ' if, for all $a \in A^-(c, a^*)$, the distribution of $\frac{d\mu_{a^*}}{d\mu_a}(x)$ under μ_a is a strict mean-preserving spread of the distribution of $\frac{d\mu'_{a^*}}{d\mu'_a}(x)$ under μ'_a . (This is an adaptation to our setting of the condition in Kim (1995) and Jewitt (2007), who studied a continuous-action setting assuming the first-order approach). In this case, by Jensen's inequality, $\text{KL}(\mu_a, \mu_{a^*}) > \text{KL}(\mu'_a, \mu'_{a^*})$ holds for all $a \in A^-(c, a^*)$, so μ dominates μ' in the sense of Corollary 1.

together with (IC), implies that (IR) does not bind at any n . Moreover, the first-best cost, i.e., the implementation cost of action 1 under perfect monitoring, is now

$$C^{\text{FB}}(u, c, a^*) = h(u(\underline{w}) + c(1) - c(0)).$$

We assume $u(\underline{w}) + c(1) - c(0) < \lim_{w \rightarrow \infty} u(w)$, so that $C^{\text{FB}}(u, c, a^*)$ is well-defined.

The following result shows that the optimal rate of convergence to the first-best can again be achieved using binary contracts. However, the rate of convergence is slower than in Theorem 1: Rather than being given by the KL-divergence $\text{KL}(\mu_0, \mu_1)$, the convergence rate is now given by the Chernoff distance

$$\text{Ch}(\mu_0, \mu_1) := \min_{\nu \in \Delta(X)} \max\{\text{KL}(\nu, \mu_0), \text{KL}(\nu, \mu_1)\}. \quad (7)$$

Note that $\text{Ch}(\mu_0, \mu_1)$ captures the KL-divergence from distributions μ_0 and μ_1 to their KL-midpoint, as any minimizer ν in (7) must satisfy $\text{KL}(\nu, \mu_0) = \text{KL}(\nu, \mu_1)$. Thus, $\text{Ch}(\mu_0, \mu_1)$ is smaller than $\text{KL}(\mu_0, \mu_1)$ (and, unlike KL-divergence, is symmetric).

Theorem 2. *Assume $A = \{0, 1\}$ with $a^* = 1$. Under both general and binary contracts, the second-best cost converges to the first-best exponentially at rate $\text{Ch}(\mu_0, \mu_1)$:*

$$C_n^{\text{SB}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = \exp[-\text{Ch}(\mu_0, \mu_1)n + o(n)];$$

$$C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = \exp[-\text{Ch}(\mu_0, \mu_1)n + o(n)].$$

The difference between Theorems 1 and 2 reflects that different binary contracts yield the optimal rate of convergence to the first-best in each setting. While the optimal convergence rate in Theorem 1 was achieved by a binary test contract with maximally lenient threshold $\gamma \searrow \mathbb{E}_0[L_n]$, in the current setting it is achieved by a binary test contract with *symmetric* threshold $\gamma = 0$. Indeed, as discussed in Section 3.2, using a maximally lenient threshold maximizes the rate at which false negatives decay, which was the dominant consideration in our main model at large n . In contrast, in the current setting, it turns out that both false positives and false negatives remain important sources of inefficiency even at large n . This necessitates using a symmetric binary test, as such a test equalizes the decay rates of false negatives and false positives, both of which equal $\text{Ch}(\mu_0, \mu_1)$ (e.g., Cover and Thomas, 1999).

The equal concern for false positives and false negatives in the current setting re-

reflects a tradeoff between risk aversion and rent extraction that was absent in the main model. To illustrate, we can decompose the difference between the implementation cost under any contract v and the first-best cost as follows:

$$\begin{aligned} & \mathbb{E}_1[h(v(x^n))] - C^{\text{FB}}(u, c, a^*) \\ = & \underbrace{\mathbb{E}_1[h(v(x^n))] - h(\mathbb{E}_1[v(x^n)])}_{\text{Jensen-inequality gap}} + \underbrace{h(\mathbb{E}_1[v(x^n)]) - h(u(C^{\text{FB}}(u, c, a^*)))}_{\text{information rent}}. \end{aligned}$$

The first term is the Jensen-inequality gap under v , which reflects the agent’s risk aversion. As in Section 3.2, this gap decays at the same rate as $\text{Var}_1[v(x^n)]$ as $n \rightarrow \infty$, which under binary contracts is governed by the decay rate of false negatives. The second term reflects that the agent’s consumption utility $\mathbb{E}_1[v(x^n)] = \mathbb{E}_0[v(x^n)] + c(1) - c(0)$ (by (IC)) exceeds his utility $u(C^{\text{FB}}(u, c, a^*)) = u(\underline{w}) + c(1) - c(0)$ under the first-best; that is, even at large n , the agent receives an information rent. Reducing the probability of false positives reduces the size of this rent. The second term was absent in our main model, as there (IR) could be made to bind at large n , giving the agent zero rent.

Notably, with non-binary actions, Appendix A.5 shows that the above tradeoff applies only to the least costly deviation $\hat{a} \in \text{argmin}_A c$, while the logic for all other deviations is analogous to our main model. As a result, the optimal convergence rate is achieved by a binary contract that employs a mix of maximally lenient and symmetric tests.²⁸ Depending on which deviation is the dominant source of inefficiency at large n , the convergence rate is either $\min_{a \in A^-(c, a^*) \setminus \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})$ as in Theorem 1 or $\text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$ as in Theorem 2.

Remark 3 (Risk-neutral case). Under severe limited liability (unlike in our baseline model), there is a gap between the first-best and second-best at each n even when u is linear. As is well-known, the optimal contract in this setting (if it exists) is binary if $|A| = 2$, but otherwise is in general complicated (e.g., Dütting, Roughgarden, and Talgam-Cohen, 2019). Moreover, even if $|A| = 2$, the optimal binary contract takes an extreme, “minimally lenient” form very different from the contracts used in Theorems 1–2: It pays a high wage only if x^n *maximizes* the log-likelihood score L_n , an event whose probability vanishes as $n \rightarrow \infty$ even under the target action. \blacktriangle

²⁸Specifically, as in Section 3.2, this contract sets thresholds $\gamma(a)$ for each $a \in A^-(c, a^*)$, and pays a high wage iff the score $L_n(a)$ exceeds $\gamma(a)$ for all a . Here $\gamma(a) \searrow \mathbb{E}_a[L_n(a)]$ is maximally lenient for all $a \in A^-(c, a^*)$, except for the least costly action \hat{a} , for which $\gamma(\hat{a}) = 0$ is the symmetric threshold.

4.2 More General Formulations of Rich Monitoring Data

4.2.1 Non-i.i.d. Signals

While our main model assumed that the principal observes n i.i.d. draws of signals from a fixed monitoring technology μ , Theorem 1 extends to settings where n indexes more general sequences of increasingly precise monitoring technologies.

Appendix A.6.1 formalizes this by considering sequences (μ^n) of monitoring technologies on a signal space Z with the following key property: Under each chosen action a , the distributions of the log-likelihood scores $L_n(a') = \frac{1}{n} \log \frac{d\mu_{a'}^n}{d\mu_{a^*}^n}(z)$ ($a' \neq a^*$) concentrate on deterministic limits as $n \rightarrow \infty$, where these limits are distinct across different chosen actions a and the convergence is described by a well-behaved rate function I_a (in the sense of the large deviation principle). Beyond our main model, this setting nests other natural examples of rich or precise monitoring data, such as:

Vanishing observation noise: When the agent chooses action $a \in A \subseteq \mathbb{R}$, the principal observes a single signal $x = a + \kappa_n \varepsilon$, where ε is a standard normal noise term and the scaling factor $\kappa_n \rightarrow 0$ as $n \rightarrow \infty$.

Serially correlated signals: When the agent chooses action a , the principal observes n signals x_1, \dots, x_n . However, instead of being i.i.d., these signals follow a (well-behaved) Markov chain whose transition kernel depends on a .

In this more general setting, Theorem A.2 in Appendix A.6.1 shows that maximally lenient binary contracts again achieve the optimal convergence rate to the first-best. The only novelty is that the KL-based convergence rate in Theorem 1 is replaced by a more general expression based on the rate function I_{a^*} .

4.2.2 Adjustable Actions

Our main model assumed that the principal observes n i.i.d. signals about a fixed, one-shot action a . However, Theorem 1 generalizes to settings where the agent can repeatedly adjust his action, as long as adjustments are infrequent relative to the frequency of signals.²⁹

Appendix A.6.2 formalizes this by assuming that the agent can adjust his action T times, i.e., sequentially chooses actions $a_1, \dots, a_T \in A$ at a total cost of $\frac{1}{T} \sum_{t=1}^T c(a_t)$. Each action a_t generates $\frac{n}{T}$ i.i.d. draws of signals from μ_{a_t} that are observed by both

²⁹This may capture, for instance, the use of digital productivity monitoring software, which tracks employees' on-screen activity at extremely small time intervals.

the principal and agent. The agent can condition his choice a_t on all past signal realizations. To keep the departure from the main model minimal, we continue to assume that the principal chooses contracts that offer a (one-shot) payment $w(x^n)$ as a function of the entire sequence x^n of n signals, and that the principal’s first-best payoff is attained by a constant action profile with $a_t = a^*$ for all t . To capture that action adjustments are infrequent relative to the number of signals, we consider the rate of convergence of the principal’s second-best payoff to her first-best payoff in the data-rich limit as $n \rightarrow \infty$, while holding fixed T .

Theorem A.3 in Appendix A.6.2 shows that this convergence rate is again the same under general and binary contracts, and is given by $\frac{1}{T} \min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})$. Thus, inefficiency is greater the more frequently the agent can adjust his action. Similar to our main model, the optimal convergence rate is achieved by a binary contract that employs a sequence of maximally lenient tests (one for each possible deviation at each t) and requires the agent to pass all tests to receive the high wage. Notably, the optimal convergence rate is the same as in the setting where the agent chooses an action sequence (a_1, \dots, a_T) but does not observe the signals generated by past actions; the latter can be viewed as a special case of our main model with T -dimensional actions and signals. Thus, whether or not the agent can condition his action choices on past data has a negligible effect on efficiency as the amount of data grows rich.

4.3 Concluding Remarks

In this paper, we studied moral hazard problems in settings where the principal has access to rich monitoring data. Rather than focusing on optimal contracts, we analyzed the optimal convergence rate of the principal’s payoffs to the first-best as data grows rich. We found that this optimal rate is achieved by a particular class of widely observed contracts—binary wage schemes—that are simple in the sense of featuring the coarsest possible wage variation. In contrast, linear contracts, another common compensation scheme that requires wages to vary finely with observed data (albeit in an arguably simple manner), approximate the first-best at a highly suboptimal rate.

In settings such as ours where optimal contracts are known to be complicated, convergence rates are a useful measure to quantify and compare the performance of simple but suboptimal contracts. Moreover, as a one-dimensional statistic, conver-

gence rates do not depend on all the specifics of the environment (e.g., the agent’s preferences or the fine details of the monitoring technology), thus allowing for sharp predictions. At the same time, a disadvantage of the convergence rate approach is that it is silent about outcomes at quite small n . This can be partly remedied by conducting numerical analysis: For instance, Figure 1 and related examples exhibit settings where lenient binary contracts start to be highly efficient at small n . Alternatively, it may be valuable to derive analytical bounds on the second-order terms $o(n)$. Such bounds would depend on the details of the environment but would allow one to bound the number n beyond which the conclusions of the convergence rate analysis hold (e.g., the finding that binary contracts outperform linear contracts).

Finally, an important assumption for our analysis was that the agent’s action space is discrete (in the sense of $\text{KL}(\mu_a, \mu_{a^*})$ being bounded away from 0). With this, we intend to capture settings where monitoring data is rich enough that it can effectively detect even relatively small deviations from the target action a^* : In particular, as n grows large, the principal can design statistical tests that detect deviations with exponentially small errors. Our analysis applies for any arbitrarily small value of $\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*}) > 0$, but more data n is needed to design such accurate tests the smaller is $\min_{a \in A^-(c, a^*)} \text{KL}(\mu_a, \mu_{a^*})$, leading to a slower convergence rate. In the extreme case where A is continuous (and μ_a is appropriately continuous in a), the convergence to the first-best would be subexponential under both optimal and binary contracts. An analogous remark applies to the adjustable action setting in Section 4.2.2, which relied on signal arrivals being frequent relative to action adjustments (i.e., letting $n \rightarrow \infty$ while fixing T). If instead action adjustments and signals are equally frequent (i.e., $n = T$), then the impact of a single deviation becomes too small as $n \rightarrow \infty$ to be detected effectively by a statistical test. The analysis of the continuous action or frequent action adjustment settings would require a different approach from the current paper, as the large-deviation tools we relied on no longer apply. Thus, we leave this as a direction for future work.

A Appendix

A.1 Preliminaries

Throughout the Appendix, we denote by $L_n(a) := \frac{1}{n} \sum_{i=1}^n \log \frac{d\mu_{a^*}}{d\mu_a}(x_i)$ the log-likelihood score for each $a \in A \setminus \{a^*\}$ and by $L_n := (L_n(a))_{a \neq a^*}$ the vector of log-likelihood scores.

A.1.1 Cramér's Theorem and Other Statistical Preliminaries

For each $a \in A$, $a' \in A \setminus \{a^*\}$ and $\ell \in \mathbb{R}$, define

$$I_{a,a'}(\ell) := \sup_{\lambda \in \mathbb{R}} \left(\lambda \ell - \log \int \left(\frac{d\mu_{a^*}}{d\mu_{a'}}(x) \right)^\lambda d\mu_a(x) \right).$$

By Cramér's theorem (Dembo and Zeitouni, 2010, Theorem 2.2.3), function $I_{a,a'}$ describes the rate at which, under action a , the distribution of the log-likelihood score $L_n(a')$ concentrates on its expectation $\hat{L}_a(a') := \mathbb{E}_a[L_n(a')]$ as $n \rightarrow \infty$: Specifically, $I_{a,a'}$ is continuous (on its effective domain), convex, and uniquely minimized at $\hat{L}_a(a')$ with $I_{a,a'}(\hat{L}_a(a')) = 0$, and for any measurable set $B \subseteq \mathbb{R}$,

$$\begin{aligned} - \inf_{\ell \in \text{cl}B} I_{a,a'}(\ell) &\geq \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_a[L_n(a') \in B] \\ &\geq \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_a[L_n(a') \in B] \geq - \inf_{\ell \in \text{int}B} I_{a,a'}(\ell). \end{aligned} \quad (8)$$

For each $\ell \in \mathbb{R}$, Lemma 6.2.3(f) in Dupuis and Ellis (2011) implies that

$$I_{a^*,a}(\ell) = \inf_{\nu \in \Delta(X)} \text{KL}(\nu, \mu_{a^*}) \text{ s.t. } \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) = \ell. \quad (9)$$

The following lemma characterizes the two statistical distances used in the paper:

Lemma A.1. *For all $a \in A \setminus \{a^*\}$, we have the following two equalities:*

$$\text{KL}(\mu_a, \mu_{a^*}) = \min_{\nu \in \Delta(X)} \text{KL}(\nu, \mu_{a^*}) \text{ s.t. } \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) \leq \mathbb{E}_a[L_n(a)]; \quad (10)$$

$$\begin{aligned} \text{Ch}(\mu_a, \mu_{a^*}) &= \min_{\nu \in \Delta(X)} \text{KL}(\nu, \mu_{a^*}) \text{ s.t. } \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) \leq 0 \\ &= \min_{\nu \in \Delta(X)} \text{KL}(\nu, \mu_a) \text{ s.t. } \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) \geq 0. \end{aligned} \quad (11)$$

Proof. To show (10), take any $\nu \in \Delta(X)$. Since $\int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) = \text{KL}(\nu, \mu_a) - \text{KL}(\nu, \mu_{a^*})$ and $\mathbb{E}_a[L_n(a)] = \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\mu_a(x) = -\text{KL}(\mu_a, \mu_{a^*})$, the constraint can be written as $\text{KL}(\nu, \mu_a) + \text{KL}(\mu_a, \mu_{a^*}) \leq \text{KL}(\nu, \mu_{a^*})$. By Gibbs' inequality, this implies that the minimum is achieved at $\nu = \mu_a$. Equation (11) follows from standard information-theoretic arguments (e.g., Cover and Thomas, 1999, Section 11.9). \square

A.1.2 Optimal Contracts

We briefly note some features of optimal contracts. First, for all large enough n , an optimal contract, i.e., a solution to the second-best problem (1), exists. To see this, note that for each large enough n , there is a contract under which (IC) and (IR) hold with strict inequality; for example, such a contract can be obtained by adding a positive constant to the binary test contract v_n defined in the proof of Theorem 1 below. Given this, the existence of an optimal contract v_n^* at all large enough n follows from Section 4 in Ke and Xu (2023). By standard arguments, v_n^* can be written as a weakly increasing function of log-likelihood vectors $L_n \in \mathbb{R}^{A \setminus \{a^*\}}$.

Second, given the existence of contracts where (IC) and (IR) hold strictly, standard arguments (e.g., Luenberger, 1997, Section 8.3) imply that the optimal contracts v_n^* satisfy the Kuhn-Tucker conditions

$$\begin{aligned} h'(v_n^*(L_n)) &= \lambda_n + \sum_{a \in A \setminus \{a^*\}} \kappa_n(a) (1 - \exp[-nL_n(a)]) & \text{if } v_n^*(L_n) > u(\underline{w}) \\ 0 &\geq \lambda_n + \sum_{a \in A \setminus \{a^*\}} \kappa_n(a) (1 - \exp[-nL_n(a)]) & \text{if } v_n^*(L_n) = u(\underline{w}), \end{aligned} \quad (12)$$

where $\lambda_n \geq 0$ is the multiplier for (IR) and $\kappa_n(a) \geq 0$ is the multiplier for the (IC) constraint with respect to deviation $a \neq a^*$. This implies

$$v_n^*(L_n) = \max \left\{ (h')^{-1} \left(\lambda_n + \sum_{a \in A \setminus \{a^*\}} \kappa_n(a) (1 - \exp[-nL_n(a)]) \right), u(\underline{w}) \right\}. \quad (13)$$

Finally, we note that the sequence (v_n^*) is uniformly bounded. To see this, it is sufficient by (13) to verify that $\limsup_{n \rightarrow \infty} \lambda_n + \sum_{a \neq a^*} \kappa_n(a) < \infty$. Suppose for a contradiction that $\lim_{k \rightarrow \infty} \lambda_{n_k} + \sum_{a \neq a^*} \kappa_{n_k}(a) = \infty$ for some subsequence (n_k) . Then for every strictly positive $L \in \mathbb{R}^{A \setminus \{a^*\}}$ (i.e., with strictly positive entries), (12) implies $\lim_{k \rightarrow \infty} h'(v_{n_k}^*(L)) = \infty$ and thus $\lim_{k \rightarrow \infty} h(v_{n_k}^*(L)) = \infty$. By the weak law of large numbers, L_n converges in probability to $\mathbb{E}_{a^*}[L_n]$, which is strictly positive. Thus, $\limsup_{n \rightarrow \infty} \mathbb{E}_{a^*}[h(v_n^*(L_n))] = \infty$, which contradicts optimality.

A.2 Proof of Theorem 1

Henceforth, let $A^-(c, a^*) := \{a \in A : c(a) < c(a^*)\}$, as in the main text, and $A^+(c, a^*) := \{a \in A : c(a) > c(a^*)\}$. We omit the dependency of A^- and A^+ on c and a^* when there is no risk of confusion. By Assumption 1.3, $A = A^- \cup A^+ \cup \{a^*\}$.

A.2.1 Upper Bound on Inefficiency under Binary Contracts

We first show that, as $n \rightarrow \infty$, $C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes at least as fast as at the exponential rate $\min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

For each $a \in A^-$, fix some $\gamma(a) \in (\mathbb{E}_a[L_n(a)], \mathbb{E}_{a^*}[L_n(a)])$. For each $a' \in \{a^*\} \cup A^-$, define $p_{a',n} := \mathbb{P}_{a'}[L_n(a) \geq \gamma(a), \forall a \in A^-]$. Observe that, by the weak law of large numbers and the choice of $\gamma(a)$, we have

$$\lim_{n \rightarrow \infty} p_{a^*,n} = 1 \quad \text{and} \quad \lim_{n \rightarrow \infty} p_{a,n} = 0 \quad \text{for all } a \in A^-.$$

Consider the sequence of binary contracts (v_n) given by

$$v_n(x^n) = \begin{cases} v_n^+ & \text{if } L_n(a) \geq \gamma(a) \text{ for all } a \in A^- \\ v_n^- & \text{otherwise,} \end{cases}$$

where v_n^+ and v_n^- are defined by

$$v_n^+ := \frac{(1 - \max_{a \in A^-} p_{a,n})c(a^*) - (1 - p_{a^*,n}) \min_{a \in A^-} c(a)}{p_{a^*,n} - \max_{a \in A^-} p_{a,n}},$$

$$v_n^- := \frac{p_{a^*,n} \min_{a \in A^-} c(a) - \max_{a \in A^-} p_{a,n} c(a^*)}{p_{a^*,n} - \max_{a \in A^-} p_{a,n}}.$$

Note that, up to restricting to large enough n , these contracts are well-defined (i.e., $v_n^+, v_n^- \in u([\underline{w}, \infty))$), because $\lim_{n \rightarrow \infty} v_n^+ = c(a^*)$, $\lim_{n \rightarrow \infty} v_n^- = \min_{a \in A^-} c(a)$ and $c(A) \subseteq \text{int}u([\underline{w}, \infty))$ by Assumption 1.2. Moreover, the choice of v_n^+ and v_n^- implies

$$p_{a^*,n} v_n^+ + (1 - p_{a^*,n}) v_n^- - c(a^*) = 0 \quad \text{and}$$

$$p_{a^*,n} v_n^+ + (1 - p_{a^*,n}) v_n^- - c(a^*) = \max_{a \in A^-} p_{a,n} v_n^+ + (1 - \max_{a \in A^-} p_{a,n}) v_n^- - \min_{a \in A^-} c(a).$$

The first line ensures that (IR) holds with equality, as the LHS is the agent's expected payoff under v_n to choosing action a^* . The second line ensures that choosing any

$a \in A^-$ is (weakly) suboptimal for the agent. Moreover, choosing any $a \in A^+$ yields expected payoff strictly less than 0 for all large enough n , as $\lim_{n \rightarrow \infty} v_n^+ = c(a^*) < c(a)$. Thus, v_n also satisfies (IC) for all large enough n .

Observe that

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log(1 - p_{a^*,n}) &= \lim_{n \rightarrow \infty} \frac{1}{n} \log(\mathbb{P}_{a^*}[\exists a \in A^- \text{ s.t. } L_n(a) < \gamma(a)]) \\ &\leq - \min_{a \in A^-} \inf_{\ell \leq \gamma(a)} I_{a^*,a}(\ell), \end{aligned} \quad (14)$$

where the inequality holds by Cramér's theorem (8). Thus, for all large enough n ,

$$\begin{aligned} C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) &\leq \mathbb{E}_{a^*}[h(v_n(x^n))] - h(c(a^*)) \\ &= p_{a^*,n} (h(v_n^+) - h(c(a^*))) + (1 - p_{a^*,n}) (h(v_n^-) - h(c(a^*))) \\ &\leq p_{a^*,n} (h(v_n^+) - h(c(a^*))) \leq \exp[-n \min_{a \in A^-} \inf_{\ell \leq \gamma(a)} I_{a^*,a}(\ell) + o(n)]. \end{aligned}$$

Here, the penultimate inequality uses $\lim_{n \rightarrow \infty} v_n^- = \min_{a \in A^-} c(a) < c(a^*)$, and the last inequality uses (14) and the fact that $v_n^+ - c(a^*) = \frac{(1-p_{a^*,n})(\min_{a \in A^-} c(a) - c(a^*))}{p_{a^*,n} - \max_{a \in A^-} p_{a,n}}$ vanishes at the same exponential rate as $1 - p_{a^*,n}$.

Finally, for each $a \in A^-$, we have

$$\sup_{\gamma > \mathbb{E}_a[L_n(a)]} \inf_{\ell \leq \gamma} I_{a^*,a}(\ell) = \inf_{\ell \leq \mathbb{E}_a[L_n(a)]} I_{a^*,a}(\ell) = \text{KL}(\mu_a, \mu_{a^*}),$$

where the second equality holds by (9)–(10) and the continuity of $I_{a^*,a}$. Thus, by choosing $\gamma(a)$ to be arbitrarily close to the maximally lenient threshold $\mathbb{E}_a[L_n(a)]$ for each $a \in A^-$, it follows that $C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes at least as fast as at the exponential rate $\min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

A.2.2 Lower Bound on Inefficiency under General Contracts

Second, to complete the proof of Theorem 1, we show that, as $n \rightarrow \infty$, $C_n^{\text{SB}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes no faster than at the exponential rate $\min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

Pick any action $\hat{a} \in \text{argmin}_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$. Consider the relaxed problem where the agent's action set is restricted to $\{a^*, \hat{a}\}$. Below, we show that the difference between the principal's value under this relaxed problem and $C^{\text{FB}}(u, c, a^*)$ cannot vanish faster than at the exponential rate $\text{KL}(\mu_{\hat{a}}, \mu_{a^*}) = \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$. This

yields the desired conclusion, because under the original problem, $C_n^{\text{SB}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ must vanish at a weakly slower rate than under the relaxed problem.

Consider any sequence of contracts (v_n) under the relaxed problem, where (without loss of optimality) we write each v_n as a weakly increasing function of $L_n := L_n(\hat{a})$. Let $\hat{L}_a := \mathbb{E}_a[L_n]$ for each $a \in \{a^*, \hat{a}\}$.

We first show that in order for the agent's utility variance to decay (at least) exponentially, the contracts v_n must become asymptotically flat on some interval around \hat{L}_{a^*} , where this interval must be larger the greater the rate of decay:

Lemma A.2. *Consider any sequence of contracts (v_n) such that $\text{Var}_{a^*}[v_n(L_n)] \leq \exp[-\lambda n + o(n)]$ for some $\lambda > 0$. Then for any interval $[\underline{\gamma}, \bar{\gamma}] \ni \hat{L}_{a^*}$ with $\inf_{\gamma < \underline{\gamma}} I_{a^*, \hat{a}}(\gamma) < \lambda$ and $\inf_{\gamma > \bar{\gamma}} I_{a^*, \hat{a}}(\gamma) < \lambda$,*

$$\lim_n (v_n(\bar{\gamma}) - v_n(\underline{\gamma})) = 0.$$

Proof. Suppose for a contradiction that the result fails. Then, since each v_n is weakly increasing, there is a subsequence (n_k) such that $\lim_k (v_{n_k}(\bar{\gamma}) - v_{n_k}(\underline{\gamma})) > 0$. Thus, writing $m_{n_k} := \mathbb{E}_{a^*}[v_{n_k}(L_{n_k})]$, we have

$$\begin{aligned} & \liminf_{k \rightarrow \infty} \frac{1}{n_k} \log \text{Var}_{a^*}[v_{n_k}(L_{n_k})] = \liminf_{k \rightarrow \infty} \frac{1}{n_k} \log \mathbb{E}_{a^*}[(v_{n_k}(L_{n_k}) - m_{n_k})^2] \\ & \geq \liminf_{k \rightarrow \infty} \frac{1}{n_k} \log \left(\min \{ \mathbb{P}_{a^*}[L_{n_k} \leq \underline{\gamma}], \mathbb{P}_{a^*}[L_{n_k} \geq \bar{\gamma}] \} \left(\frac{v_{n_k}(\bar{\gamma}) - v_{n_k}(\underline{\gamma})}{2} \right)^2 \right) \\ & \geq -\max \left\{ \inf_{\gamma < \underline{\gamma}} I_{a^*, \hat{a}}(\gamma), \inf_{\gamma > \bar{\gamma}} I_{a^*, \hat{a}}(\gamma) \right\} > -\lambda, \end{aligned}$$

where the second inequality uses Cramér's theorem (8). This contradicts the assumption that $\text{Var}_{a^*}[v_{n_k}(L_{n_k})] \leq \exp[-\lambda n_k + o(n_k)]$. \square

Next, we show that under (IC) the agent's utility variance cannot vanish faster than at the exponential rate $\text{KL}(\mu_{\hat{a}}, \mu_{a^*})$:

Lemma A.3. *Consider any uniformly bounded sequence (v_n) of contracts satisfying (IC) under the relaxed problem. Then $\text{Var}_{a^*}[v_n(L_n)] \geq \exp[-\text{KL}(\mu_{\hat{a}}, \mu_{a^*})n + o(n)]$.*

Proof. Suppose otherwise. Then there exists $\lambda > \text{KL}(\mu_{\hat{a}}, \mu_{a^*})$ such that $\text{Var}_{a^*}[v_n(L_n)] \leq \exp[-\lambda n + o(n)]$. By (9)–(10), $\text{KL}(\mu_{\hat{a}}, \mu_{a^*}) = I_{a^*, \hat{a}}(\hat{L}_{\hat{a}})$, so $\lambda > I_{a^*, \hat{a}}(\hat{L}_{\hat{a}})$. Thus, by continuity of $I_{a^*, \hat{a}}$, there is $\underline{\gamma} < \hat{L}_{\hat{a}}$ with $\inf_{\gamma < \underline{\gamma}} I_{a^*, \hat{a}}(\gamma) < \lambda$. Moreover, by continuity

of $I_{a^*, \hat{a}}$ and the fact that $I_{a^*, \hat{a}}(\hat{L}_{a^*}) = 0$, there is $\bar{\gamma} > \hat{L}_{a^*}$ with $\inf_{\gamma > \bar{\gamma}} I_{a^*, \hat{a}}(\gamma) < \lambda$. By Lemma A.2, $\lim_{n \rightarrow \infty} (v_n(\bar{\gamma}) - v_n(\underline{\gamma})) = 0$.

At the same time, for each $a \in \{a^*, \hat{a}\}$, the fact that \hat{L}_a is in the interior of $[\underline{\gamma}, \bar{\gamma}]$ and the weak law of large numbers implies that $\lim_{n \rightarrow \infty} \mathbb{P}_a [L_n \in [\underline{\gamma}, \bar{\gamma}]] = 1$. Since the sequence of contracts is bounded, this implies $\lim_{n \rightarrow \infty} (\mathbb{E}_{a^*}[v_n(L_n)] - \mathbb{E}_{\hat{a}}[v_n(L_n)]) = 0$. But then v_n violates (IC) for all large enough n , a contradiction. \square

Finally, consider a sequence (v_n^*) of optimal contracts under the relaxed problem. By Appendix A.1.2, such a sequence exists (up to restricting to large n) and is uniformly bounded. The difference between the principal's cost and the first-best is

$$\begin{aligned} \mathbb{E}_{a^*}[h(v_n^*(L_n))] - h(c(a^*)) &= \mathbb{E}_{a^*}[h(v_n^*(L_n))] - h(\mathbb{E}_{a^*}[v_n^*(L_n)]) + h(\mathbb{E}_{a^*}[v_n^*(L_n)]) - h(c(a^*)) \\ &\geq \mathbb{E}_{a^*}[h(v_n^*(L_n))] - h(\mathbb{E}_{a^*}[v_n^*(L_n)]) \\ &\geq \text{Var}_{a^*}[v_n^*(L_n)] \frac{\inf_{L_n} h''(v_n(L_n))}{2}, \end{aligned}$$

where the first inequality holds because $\mathbb{E}_{a^*}[v_n^*(L_n)] \geq c(a^*)$ by (IR) and the second inequality uses the variance-based estimate of the Jensen-inequality gap (Liao and Berg, 2019). Here $\inf_{L_n} h''(v_n(L_n)) > 0$, where the infimum is taken over all possible realizations of L_n , is bounded away from zero uniformly in n . This follows by observing that (i) the contract wages are uniformly bounded, by the Kuhn-Tucker conditions (12) and the assumption that $\lim_{w \rightarrow \infty} u'(w) = 0$, and (ii) u is C^2 . Thus, by Lemma A.3, the difference relative to the first-best cannot vanish faster than at the exponential rate $\text{KL}(\mu_{\hat{a}}, \mu_{a^*})$.

A.3 Proof of Proposition 1

Proof of (4): Suppose for a contradiction that (4) fails, i.e.,

$$C_n^{\text{lin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = o(1/n). \quad (15)$$

Thus, we can pick a sequence of IR and IC linear contracts (w_n) with $w_n(x_n) =$

$\frac{1}{n} \sum_{i=1}^n b_n(x_i)$ such that $\mathbb{E}_{a^*}[w_n(x^n)] - C^{\text{FB}}(u, c, a^*) = o(1/n)$. Note that

$$\begin{aligned} \mathbb{E}_{a^*}[w_n(x^n)] - C^{\text{FB}}(u, c, a^*) &= \mathbb{E}_{a^*}[w_n(x^n)] - h(\mathbb{E}_{a^*}[u(w_n(x^n))]) + h(\mathbb{E}_{a^*}[u(w_n(x^n))]) - h(c(a^*)) \\ &\geq \frac{\inf_v h''(v)}{2} \text{Var}_{a^*}[u(w_n(x^n))] + \underbrace{h(\mathbb{E}_{a^*}[u(w_n(x^n))]) - h(c(a^*))}_{\geq 0 \text{ by IR}} \end{aligned}$$

by the variance-based estimate of the Jensen-inequality gap (Liao and Berg, 2019). Thus, $\mathbb{E}_{a^*}[w_n(x^n)] - C^{\text{FB}}(u, c, a^*) = o(1/n)$ requires that $\mathbb{E}_{a^*}[u(w_n(x^n))] \rightarrow h(c(a^*))$ and $\text{Var}_{a^*}[u(w_n(x^n))] = o(1/n)$ since $\inf_v h''(v) > 0$. By a Taylor expansion argument, the latter implies $\text{Var}_{a^*}[w_n(x^n)] = o(1/n)$. As $\text{Var}_{a^*}[w_n(x^n)] = \frac{1}{n} \text{Var}_{a^*}[b_n(x_1)]$ for each n , this implies $\lim_{n \rightarrow \infty} \text{Var}_{a^*}[b_n(x_1)] = 0$. Thus, $\mathbb{E}_{a^*}[u(w_n(x^n))] \rightarrow h(c(a^*))$ implies that for any $\varepsilon > 0$, $\lim_{n \rightarrow \infty} \mu_{a^*}(\{x : |b_n(x) - h(c(a^*))| > \varepsilon\}) = 0$. This in turn implies that for all $a \neq a^*$,

$$\lim_{n \rightarrow \infty} \mu_a(\{x : |b_n(x) - h(c(a^*))| > \varepsilon\}) = 0, \quad (16)$$

because, by the data-processing inequality, $\mu_a(X') \log \frac{\mu_a(X')}{\mu_{a^*}(X')} + \mu_a(X \setminus X') \log \frac{\mu_a(X \setminus X')}{\mu_{a^*}(X \setminus X')} \leq \text{KL}(\mu_a, \mu_{a^*}) < \infty$ for all measurable $X' \subseteq X$.

The agent's expected payoff to choosing a is then

$$\mathbb{E}_a \left[u \left(\frac{1}{n} \sum_{i=1}^n b_n(x_i) \right) \right] - c(a) \geq \mathbb{E}_a [u(b_n(x_1))] - c(a),$$

by concavity of u . By (16) and the fact that u is bounded below, the lim sup of the RHS as $n \rightarrow \infty$ is no less than $h(c(a^*)) - c(a)$. Thus, since $\mathbb{E}_{a^*}[u(w_n(x^n))] - h(c(a^*)) \rightarrow 0$, (IC) fails for any a with $c(a) < h(c(a^*))$ at all large enough n .

“Moreover” part: We show that (4) holds with equality if $\mu_{a^*} \notin \text{co}\{\mu_a : a \neq a^*\}$ and \underline{w} is sufficiently low. We first exhibit a linear contract that achieves the first-best cost at $n = \infty$. This can be done by finding $b : X \rightarrow [\underline{w}, \infty)$ such that (i) $\int b(x) d\mu_{a^*}(x) = h(c(a^*))$ and (ii) $\int b(x) d\mu_a(x) < h(c(a))$ for each $a \neq a^*$. Indeed, (i) ensures that (IR) binds, i.e., $u(\int b(x) d\mu_{a^*}(x)) - h(c(a^*)) = 0$. Given this, (ii) ensures that (IC) holds strictly, i.e., $u(\int b(x) d\mu_a(x)) - c(a) < 0$ for each $a \neq a^*$. To find such a b , observe that, by a separation theorem (Dunford and Schwartz, 1988) and the assumption that $\mu_{a^*} \notin \text{co}\{\mu_a : a \neq a^*\}$, there is a bounded function $b : X \rightarrow \mathbb{R}$ with $\int b(x) d\mu_{a^*}(x) > \int b(x) d\mu_a(x)$ for all $a \neq a^*$. By scaling b if necessary, we can ensure $\int b(x) d\mu_{a^*}(x) - h(c(a^*)) > \int b(x) d\mu_a(x) - h(c(a))$ for all $a \neq a^*$. By adding a constant

to b if necessary, we can ensure $0 = \int b(x)d\mu_{a^*}(x) - h(c(a^*)) > \int b(x)\mu_a(x) - h(c(a))$ for all $a \neq a^*$. Since b is bounded, $b(X) \subseteq [\underline{w}, \infty)$ if \underline{w} is sufficiently low.

To establish that (4) holds with equality, we modify the linear contract obtained above for each n by setting $w_n(x^n) = \frac{1}{n} \sum_{i=1}^n (b(x_i) + \varepsilon_n) = \sum_x \nu_{x^n}(x)b(x) + \varepsilon_n$, where $\varepsilon_n \in \mathbb{R}_+$ is a constant and $\nu_{x^n} \in \Delta(X)$ denotes the realized empirical signal frequency. The constant ε_n is pinned down for all large enough n by requiring (IR) to bind, i.e., $\mathbb{E}_{a^*}[u(\sum_x \nu_{x^n}(x)b(x) + \varepsilon_n)] - c(a^*) = 0$.

Note that $\lim_{n \rightarrow \infty} \varepsilon_n = 0$ by the law of large numbers and construction of b . We moreover show that $\varepsilon_n = \frac{K}{n} + o(1/n)$ for some $K > 0$. To see this, observe that

$$\begin{aligned} & \inf_{w \in \text{co}\{b(x) + \varepsilon_n : x \in X\}} -\frac{u''(w)}{2} \text{Var}_{a^*} \left[\sum_x \nu_{x^n}(x)b(x) \right] \\ & \geq u \left(\mathbb{E}_{a^*} \left[\sum_x \nu_{x^n}(x)b(x) + \varepsilon_n \right] \right) - \mathbb{E}_{a^*} \left[u \left(\sum_x \nu_{x^n}(x)b(x) + \varepsilon_n \right) \right] \\ & \geq u \left(\mathbb{E}_{a^*} \left[\sum_x \nu_{x^n}(x)b(x) \right] \right) + \varepsilon_n u' \left(\mathbb{E}_{a^*} \left[\sum_x \nu_{x^n}(x)b(x) + \varepsilon_n \right] \right) - c(a^*) \\ & = \varepsilon_n u' \left(\mathbb{E}_{a^*} \left[\sum_x \nu_{x^n}(x)b(x) + \varepsilon_n \right] \right), \end{aligned}$$

where the first inequality uses the variance-based estimate of the Jensen-inequality gap (note $\inf -u'' > 0$ on the relevant domain of w , as b is bounded), the second inequality uses the concavity of u and construction of ε_n , and the last equality uses the construction of b . Since the variance term in the first line is $o(\frac{1}{n})$, we have $\varepsilon_n = \frac{K}{n} + o(\frac{1}{n})$ for some $K > 0$.

Because $\varepsilon_n \rightarrow 0$ and the original linear contract satisfies (IC) strictly at $n = \infty$, the modified contracts satisfy (IC) for all large enough n . Under these contracts, the principal's cost difference relative to the first-best cost is $\varepsilon_n = \frac{K}{n} + o(\frac{1}{n})$, as required.

A.4 Proof of Proposition 2

We prove the following generalization of Proposition 2 to an arbitrary finite action space. As usual, we write contracts v_n as functions of log-likelihood scores.

Proposition A.1. *For any weakly convergent sequence of optimal contracts (v_n^*) ,*

there exists $\gamma : A^- \rightarrow \mathbb{R}$ such that

$$\lim_{n \rightarrow \infty} v_n^*(L) = \begin{cases} c(a^*) & \text{if } L(a) > \gamma(a) \text{ for each } a \in A \setminus \{a^*\} \\ u(\underline{w}) & \text{if } L(a) < \gamma(a) \text{ for some } a \in A \setminus \{a^*\}, \end{cases}$$

where $\gamma(\hat{a}) = \mathbb{E}_{\hat{a}}[L_n(\hat{a})]$ is maximally lenient for all $\hat{a} \in \operatorname{argmin}_{a \in A^-} \operatorname{KL}(\mu_a, \mu_{a^*})$.

A.4.1 The Limit Contract is Binary

By Appendix A.1.2, for all large enough n , v_n^* is given by (13) for multipliers λ_n and $\kappa_n(a)$ for each $a \neq a^*$. Since (v_n^*) approximates the first-best, $\lim_{n \rightarrow \infty} \mathbb{E}_{a^*}[v_n^*(L_n)] = c(a^*)$. Since L_n converges in probability to $\mathbb{E}_{a^*}[L_n] > 0$, (13) implies

$$\lim_{n \rightarrow \infty} \left(\lambda_n + \sum_{a \neq a^*} \kappa_n(a) \right) = h'(c(a^*)). \quad (17)$$

Up to restricting to an appropriate subsequence, we can assume that $\lim_{n \rightarrow \infty} \frac{1}{n} \log \kappa_n(a) =: \gamma(a)$ exists (allowing for $-\infty$) for each $a \neq a^*$. Since $\kappa_n(a)$ is bounded by (17),

$$\lim_{n \rightarrow \infty} \kappa_n(a) (1 - \exp[-nL(a)]) = \begin{cases} \lim_{n \rightarrow \infty} \kappa_n(a) & \text{if } L(a) > \gamma(a) \\ -\infty & \text{if } L(a) < \gamma(a). \end{cases}$$

Thus, by (13),

$$\lim_{n \rightarrow \infty} v_n^*(L) = \begin{cases} c(a^*) & \text{if } L(a) > \gamma(a) \text{ for each } a \in A \setminus \{a^*\} \\ u(\underline{w}) & \text{if } L(a) < \gamma(a) \text{ for some } a \in A \setminus \{a^*\}. \end{cases} \quad (18)$$

A.4.2 Maximal Leniency of Limit Thresholds

It remains to show that in (18), $\gamma(\hat{a}) = \mathbb{E}_{\hat{a}}[L_n(\hat{a})]$ for all $\hat{a} \in \operatorname{argmin}_{a \in A^-} \operatorname{KL}(\mu_a, \mu_{a^*})$.

By the weak convergence of (v_n^*) and (18), for any $\varepsilon > 0$ there is n_ε such that

$v_n^*(L) \leq u(\underline{w}) + \varepsilon$ if $L(a) \leq \gamma(a) - \varepsilon$ for some $a \neq a^*$ and $n \geq n_\varepsilon$. Thus

$$\begin{aligned} & \liminf_n \frac{1}{n} \log \text{Var}_{a^*}[v_n^*(L_n)] = \liminf_n \frac{1}{n} \log \mathbb{E}_{a^*}[(\mathbb{E}_{a^*}[v_n^*(L_n)] - v_n^*(L_n))^2] \\ & \geq \liminf_n \frac{1}{n} \log \mathbb{P}_{a^*}[L_n(a) < \gamma(a) - \varepsilon \text{ for some } a \neq a^*](c(a^*) - u(\underline{w}) - 2\varepsilon)^2 \\ & = \liminf_n \frac{1}{n} \log \mathbb{P}_{a^*}[L_n(a) < \gamma(a) - \varepsilon \text{ for some } a \neq a^*] \geq -\min_{a \neq a^*} \inf_{\ell < \gamma(a) - \varepsilon} I_{a^*,a}(\ell) \end{aligned}$$

where the first inequality uses that $\mathbb{E}_{a^*}[v_n^*(L_n)] \geq c(a^*) - \varepsilon$ for all large enough n and the second inequality uses Cramér's theorem (8). Since $\varepsilon > 0$ is arbitrary, we have $\liminf_n \frac{1}{n} \log \text{Var}_{a^*}[v_n^*(L_n)] \geq -\min_{a \neq a^*} \inf_{\ell < \gamma(a)} I_{a^*,a}(\ell)$.

Since $\text{Var}_{a^*}[v_n^*(L_n)]$ vanishes at exponential rate $\min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$ (by the proof of Theorem 1), it follows that (using (9)) that

$$\min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*}) \leq \min_{a \neq a^*} \inf_{\ell < \gamma(a)} I_{a^*,a}(\ell) = \inf_{\substack{a \neq a^*, \nu \in \Delta(X) \text{ s.t.} \\ \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) < \gamma(a)}} \text{KL}(\nu, \mu_{a^*}). \quad (19)$$

Suppose for a contradiction that $\gamma(\hat{a}) \neq \mathbb{E}_{\hat{a}}[L_n(\hat{a})]$ for some $\hat{a} \in \text{argmin}_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$. If $\gamma(\hat{a}) > \mathbb{E}_{\hat{a}}[L_n(\hat{a})]$, then since $\mathbb{E}_{\hat{a}}[L_n(\hat{a})] = -\text{KL}(\mu_{\hat{a}}, \mu_{a^*})$ by definition of the log-likelihood score, Lemma A.1 and the strict convexity of KL-divergence implies that

$$\text{KL}(\mu_{\hat{a}}, \mu_{a^*}) > \inf_{\nu \in \Delta(X) \text{ s.t. } \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x) < \gamma(\hat{a})} \text{KL}(\nu, \mu_{a^*}),$$

which contradicts (19). If instead $\gamma(\hat{a}) < \mathbb{E}_{\hat{a}}[L_n(\hat{a})]$, there exists $a \neq a^*$ such that

$$\gamma(a) \geq \mathbb{E}_{\hat{a}}[L_n(a)] = \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\mu_{\hat{a}}(x),$$

as otherwise $\lim_{n \rightarrow \infty} \mathbb{P}_{\hat{a}}[L_n(a) > \gamma(a) \text{ for all } a \neq a^*] = 1$ by the weak law of large numbers, which leads to a violation of (IC) at all large enough n . Let $B := \{\nu \in \Delta(X) : \text{KL}(\nu, \mu_{a^*}) \leq \text{KL}(\mu_{\hat{a}}, \mu_{a^*})\}$. Then (19) implies that, for every $\nu \in B$, $\gamma(a) \leq \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x)$.³⁰ Thus, $\gamma(a) = \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\mu_{\hat{a}}(x)$, as $\mu_{\hat{a}} \in B$. By the strict convexity of KL-divergence, this shows that $\mu_{\hat{a}}$ is the unique point at which B is tangent to the hyperplane $\{\nu \in \Delta(X) : \gamma(a) = \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x)\}$. Likewise, by

³⁰If $\gamma(a) > \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu(x)$, one can find ν' close to ν with $\text{KL}(\nu', \mu_{a^*}) < \text{KL}(\mu_{\hat{a}}, \mu_{a^*})$ and $\gamma(a) > \int \log \frac{d\mu_{a^*}}{d\mu_a}(x) d\nu'(x)$ (based on the convexity of KL-divergence), which violates (19).

construction of B and Lemma A.1, $\mu_{\hat{a}}$ is the unique element in B that is tangent to the hyperplane $\{\nu \in \Delta(X) : \text{KL}(\mu_{\hat{a}}, \mu_{a^*}) = \int \log \frac{\mu_{a^*}(x)}{\mu_{\hat{a}}(x)} d\nu(x)\}$. By $\mu_{a^*} \neq \mu_a$, these two hyperplanes are not parallel to each other. As every boundary point of B has a unique tangent space (by strict convexity of KL-divergence), this is a contradiction.

A.5 Proof of Theorem 2

We prove the following generalization of Theorem 2 that allows for non-binary actions. For simplicity, we assume there is a unique action \hat{a} that minimizes c . We consider the case where Assumption 1.2 is violated, in the sense that $c(\hat{a}) \leq u(\underline{w})$. Then the first-best cost is $C^{\text{FB}}(u, c, a^*) = h(u(\underline{w}) + c(a^*) - c(\hat{a}))$. As in the main text, we assume that $u(\underline{w}) + c(a^*) - c(\hat{a}) < \lim_{w \rightarrow \infty} u(w)$, so $C^{\text{FB}}(u, c, a^*)$ is well-defined.

Theorem A.1. *Under both general and binary contracts, the second-best cost converges to the first-best exponentially at rate $\rho(\mu) := \min\{\text{Ch}(\mu_{\hat{a}}, \mu_{a^*}), \min_{a \in A^- \setminus \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})\}$:*

$$C_n^{\text{SB}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = \exp[-\rho(\mu)n + o(n)];$$

$$C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*) = \exp[-\rho(\mu)n + o(n)].$$

To understand the convergence rate $\rho(\mu)$, it is helpful to distinguish two cases. First, suppose $\text{Ch}(\mu_{\hat{a}}, \mu_{a^*}) < \min_{a \in A^- \setminus \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})$; intuitively, a deviation to the least costly action \hat{a} is relatively harder to detect than other deviations. In this case, we show based on a similar intuition as in the binary-action case (Theorem 2) that the convergence rate is $\text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$. Second, suppose $\text{Ch}(\mu_{\hat{a}}, \mu_{a^*}) > \min_{a \in A^- \setminus \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})$; intuitively, a deviation to \hat{a} is easier to detect than some other deviations. In this case, we show based on a similar logic as in Theorem 1 that the convergence rate is given by $\min_{a \in A^- \setminus \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})$.

A.5.1 Upper Bound on Inefficiency under Binary Contracts

We first show that, as $n \rightarrow \infty$, $C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes at least as fast as at the exponential rate $\rho(\mu)$. For each $a \in A^-$, fix some $\gamma(a) \in (\mathbb{E}_a[L_n(a)], \mathbb{E}_{a^*}[L_n(a)])$. For each $a' \in \{a^*\} \cup A^-$, define $p_{a',n} := \mathbb{P}_{a'}[L_n(a) \geq \gamma(a) \forall a \in A^-]$. Observe that, by the weak law of large numbers and the choice of $\gamma(a)$, we have

$$\lim_{n \rightarrow \infty} p_{a^*,n} = 1 \quad \text{and} \quad \lim_{n \rightarrow \infty} p_{a,n} = 0 \quad \text{for all } a \in A^-.$$

Consider the sequence of binary contracts (v_n) given by

$$v_n(x^n) = \begin{cases} v_n^+ & \text{if } L_n(a) \geq \gamma(a) \text{ for all } a \in A^- \\ u(\underline{w}) & \text{otherwise,} \end{cases}$$

where $v_n^+ := u(\underline{w}) + \frac{c(a^*) - c(\hat{a})}{p_{a^*,n} - p_{\hat{a},n}}$.

Note that, up to restricting to large enough n , these contracts are well-defined, because $\lim_{n \rightarrow \infty} v_n^+ = u(\underline{w}) + c(a^*) - c(\hat{a})$. Moreover, by choice of v_n^+ , we have $\mathbb{E}_{a^*}[v_n(x^n)] - c(a^*) = \mathbb{E}_{\hat{a}}[v_n(x^n)] - c(\hat{a})$, i.e., (IC) holds with equality at deviation \hat{a} . Furthermore, $\lim_{n \rightarrow \infty} \mathbb{E}_{a^*}[v_n(x^n)] - c(a^*) = u(\underline{w}) - c(\hat{a}) > \lim_{n \rightarrow \infty} \mathbb{E}_a[v_n(x^n)] - c(a) = u(\underline{w}) - c(a)$ for all $a \in A^- \setminus \{\hat{a}\}$, so deviations to any other action in A^- are also suboptimal at large enough n . Finally, the agent's expected payoff to choosing any $a \in A^+$ is bounded above by $u(\underline{w}) + c(a^*) - c(\hat{a}) - c(a) < u(\underline{w}) - c(\hat{a}) = \lim_{n \rightarrow \infty} \mathbb{E}_{a^*}[v_n(x^n)] - c(a^*)$. Thus, (IC) holds for all large enough n .

For all large enough n , the difference between the principal's cost under v_n and the first-best can be decomposed as

$$\begin{aligned} & \mathbb{E}_{a^*}[h(v_n(x^n))] - h(u(\underline{w}) + c(a^*) - c(\hat{a})) \\ &= \mathbb{E}_{a^*}[h(v_n(x^n))] - h(\mathbb{E}_{a^*}[v_n(x^n)]) + h(\mathbb{E}_{a^*}[v_n(x^n)]) - h(u(\underline{w}) + c(a^*) - c(a)) \\ &= \mathbb{E}_{a^*}[h(v_n(x^n))] - h(\mathbb{E}_{a^*}[v_n(x^n)]) + h(\mathbb{E}_{\hat{a}}[v_n(x^n)] + c(a^*) - c(\hat{a})) - h(u(\underline{w}) + c(a^*) - c(\hat{a})), \end{aligned}$$

where the second equality uses the binding (IC) against \hat{a} .

As in the proof of Theorem 1, $\mathbb{E}_{a^*}[h(v_n(x^n))] - h(\mathbb{E}_{a^*}[v_n(x^n)])$ vanishes at least as fast as at the exponential rate $\min_{a \in A^-} \inf_{\ell \leq \gamma(a)} I_{a^*,a}(\ell)$. Moreover, for each $a \in A^-$, Lemma A.1, (9), and the continuity of $I_{a^*,a}$ implies that

$$\sup_{\gamma > \mathbb{E}_a[L_n(a)]} \inf_{\ell \leq \gamma} I_{a^*,a}(\ell) = \inf_{\ell \leq \mathbb{E}_a[L_n(a)]} I_{a^*,a}(\ell) = \text{KL}(\mu_a, \mu_{a^*}) \quad \text{and} \quad \inf_{\ell \leq 0} I_{a^*,a}(\ell) = \text{Ch}(\mu_a, \mu_{a^*}).$$

Thus, by choosing $\gamma(a)$ to be arbitrarily close to the maximally lenient threshold $\mathbb{E}_a[L_n(a)]$ for each $a \in A^- \setminus \{\hat{a}\}$ and setting $\gamma(\hat{a}) = 0$, $\mathbb{E}_{a^*}[h(v_n(x^n))] - h(\mathbb{E}_{a^*}[v_n(x^n)])$ can be made to vanish at a rate arbitrarily close to the exponential rate $\rho(\mu)$.

Next, note that the rate at which $h(\mathbb{E}_{\hat{a}}[v_n(x^n)] + c(a^*) - c(\hat{a})) - h(u(\underline{w}) + c(a^*) - c(\hat{a}))$ vanishes is the same as the rate at which $\mathbb{P}_{\hat{a}}[L_n(a) \geq \gamma(a), \forall a \in A^-]$ vanishes. Moreover, $\mathbb{P}_{\hat{a}}[L_n(a) \geq \gamma(a), \forall a \in A^-] \leq \mathbb{P}_{\hat{a}}[L_n(\hat{a}) \geq \gamma(\hat{a})]$, and the latter vanishes at

rate $\text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$ since we set $\gamma(\hat{a}) = 0$.

The previous two paragraphs imply the desired conclusion.

A.5.2 Lower Bound on Inefficiency under General Contracts

Next, we show that $C_n(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes at most as fast as the exponential rate $\rho(\mu)$. We consider two cases:

Case (i): $\rho(\mu) = \min_{a \in A - \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})$. Take any $a' \in \arg\min_{a \in A - \{\hat{a}\}} \text{KL}(\mu_a, \mu_{a^*})$. Consider the relaxed problem where the agent's action space is $A = \{a^*, a'\}$ and his outside option yields payoff $u(\underline{w}) - c(\hat{a})$. This is indeed a relaxation of the original problem, as the agent's expected utility of choosing \hat{a} is no less than $u(\underline{w}) - c(\hat{a})$. Take a sequence of optimal contracts (v_n) under the relaxed problem, where we write each v_n as a weakly increasing function of $L_n := L_n(a')$. Suppose for a contradiction that the principal's cost in this problem converges to the first-best $h(u(\underline{w}) + c(a^*) - c(\hat{a}))$ faster than at some exponential rate $\lambda > \rho(\mu) = \text{KL}(\mu_{a'}, \mu_{a^*})$. Observe that

$$\begin{aligned} & \mathbb{E}_{a^*}[h(v_n(L_n))] - h(u(\underline{w}) + c(a^*) - c(\hat{a})) \\ &= \mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)]) + h(\mathbb{E}_{a^*}[v_n(L_n)]) - h(u(\underline{w}) + c(a^*) - c(\hat{a})) \\ &\geq \mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)]), \end{aligned}$$

where the inequality uses the (IR) constraint of the relaxed problem. Thus, $\mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)])$ vanishes faster than at rate $\lambda > \text{KL}(\mu_{a'}, \mu_{a^*})$. As in Appendix A.2.2, this implies that (IC) fails for all large enough n , a contradiction.

Case (ii): $\rho(\mu) = \text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$. Consider the relaxed problem with action space $A = \{a^*, \hat{a}\}$ and outside option 0. Take a sequence of optimal contracts (v_n) under the relaxed problem, where we write each v_n as a weakly increasing function of $L_n := L_n(\hat{a})$. Let $\hat{L}_a := \mathbb{E}_a[L_n]$ for $a \in \{a^*, \hat{a}\}$. Suppose for a contradiction that the principal's cost in this problem converges to the first-best $h(u(\underline{w}) + c(a^*) - c(\hat{a}))$ faster than at some exponential rate $\lambda > \rho(\mu) = \text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$. Observe that

$$\begin{aligned} & \mathbb{E}_{a^*}[h(v_n(L_n))] - h(u(\underline{w}) + c(a^*) - c(\hat{a})) \\ &= \mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)]) + h(\mathbb{E}_{a^*}[v_n(L_n)]) - h(u(\underline{w}) + c(a^*) - c(\hat{a})) \\ &\geq \mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)]) + h(\mathbb{E}_{\hat{a}}[v_n(L_n)] + c(a^*) - c(\hat{a})) - h(u(\underline{w}) + c(a^*) - c(\hat{a})), \end{aligned}$$

where the inequality uses (IC). Thus, both $\mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)])$ and $h(\mathbb{E}_{\hat{a}}[v_n(L_n)] +$

$c(a^*) - c(\hat{a}) - h(u(\underline{w}) + c(a^*) - c(\hat{a}))$ vanish faster than at rate $\lambda > \text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$. The latter also implies that $\mathbb{E}_{\hat{a}}[v_n(L_n)] - u(\underline{w})$ vanishes faster than at rate λ .

By (9) and (11), we have $I_{a^*, \hat{a}}(0) = I_{\hat{a}, \hat{a}}(0) = \text{Ch}(\mu_{\hat{a}}, \mu_{a^*})$. Thus, by the continuity of the rate functions, there is an interval $\Gamma \ni 0$ such that $\lambda > I_{a^*, \hat{a}}(\ell), I_{\hat{a}, \hat{a}}(\ell)$ for all $\ell \in \Gamma$. By the same argument as in Appendix A.2.2, in order for $\mathbb{E}_{a^*}[h(v_n(L_n))] - h(\mathbb{E}_{a^*}[v_n(L_n)])$ to vanish faster than at rate λ , we must have $\lim_{n \rightarrow \infty} v_n(\hat{L}_{a^*}) - v_n(0) = 0$. Thus, $u(\underline{w}) + c(a^*) - c(\hat{a}) = \lim_{n \rightarrow \infty} v_n(\hat{L}_{a^*}) = \lim_{n \rightarrow \infty} v_n(0)$. By an analogous argument, in order for $\mathbb{E}_{\hat{a}}[v_n(L_n)] - u(\underline{w})$ to vanish faster than at rate λ , we must have $\lim_{n \rightarrow \infty} v_n(0) - u(\underline{w}) = 0$. This contradicts $c(a^*) > c(\hat{a})$.

A.6 Details for Section 4.2

A.6.1 Non-i.i.d. Signals

Consider a sequence $\mu = (\mu^n)$ of monitoring technologies indexed by n . Fix a measurable signal space Z that is common across n (without loss of generality) and endowed with some σ -finite measure ν . For each n and $a \in A$, the monitoring technology μ^n specifies a (not necessarily full-support) signal distribution $\mu_a^n \in \Delta(Z)$ that is absolutely continuous with respect to ν , so that μ_a^n admits a density function g_a^n . Assume that for each $a, a' \in A \setminus \{a^*\}$, the log-likelihood score $L_n(a') := \frac{1}{n} \log \frac{g_{a^*}^n(z)}{g_{a'}^n(z)}$ is well-defined $\mu_{a'}^n$ -almost surely.

The remaining setup and assumptions on payoffs are the same as in Section 2. In particular, the principal's second-best cost solves

$$C_n^{\text{SB}}(\mu, a^*, u, c) = \inf_{w: Z \rightarrow [\underline{w}, \infty)} \int w(z) d\mu_{a^*}^n(z)$$

subject to the IC and IR constraints

$$\int u(w(z)) d\mu_{a^*}^n(z) - c(a^*) \geq \int u(w(z)) d\mu_a^n(z) - c(a), \forall a \in A, \quad (\text{IC})$$

$$\int u(w(z)) d\mu_{a^*}^n(z) - c(a^*) \geq 0. \quad (\text{IR})$$

The following key requirement captures that monitoring technologies become very precise as $n \rightarrow \infty$: We assume that, under any chosen action $a \in A$, the distributions of the log-likelihood score $L_n(a')$ for each $a' \neq a^*$ obey the *large-deviation principle*

with respect to some **rate function** $I_{a,a'} : \mathbb{R} \rightarrow \overline{\mathbb{R}}_+$ as $n \rightarrow \infty$. That is, for each measurable set $B \subseteq \mathbb{R}$,

$$\begin{aligned} - \inf_{\ell \in \text{int}B} I_{a,a'}(\ell) &\leq \liminf_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_a[L_n(a') \in B] \\ &\leq \limsup_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}_a[L_n(a') \in B] \leq - \inf_{\ell \in \text{cl}B} I_{a,a'}(\ell). \end{aligned} \quad (20)$$

Intuitively, $I_{a,a'}(\ell)$ measures how atypical it is, under the chosen action a , for the score $L_n(a')$ to take value ℓ when n is large. We impose an identification condition: For each $a \in A$ and $a' \neq a^*$, $I_{a,a'}$ is uniquely minimized by some value $\hat{L}_a(a')$, where $\hat{L}_a(a') \neq \hat{L}_{\hat{a}}(a')$ for $a \neq \hat{a}$. Thus, the distribution of $L_n(a')$ under a concentrates on the deterministic limit $\hat{L}_a(a')$ as n becomes large, where the limit is distinct across different actions a , thus allowing the principal to identify the agent's chosen action as $n \rightarrow \infty$. We also impose the following regularity conditions: For each $a \in A$ and $a' \neq a^*$, $I_{a,a'}$ has compact level sets,³¹ and $I_{a^*,a'}$ is continuous at $\hat{L}_{a'}(a')$ and $\hat{L}_{a^*}(a')$.

This setting encompasses several natural examples of rich/precise monitoring data:

Example A.1 (Many signals). *Suppose $Z = \bigcup_n X^n$, i.e., n represents the number of signals observed by the principal, as in our main model. However, beyond i.i.d. signals, we can allow for correlated signals: The signal sequence (x_1, \dots, x_n) is generated by a Markov transition matrix $M_a(\cdot|x) \in \Delta(X)$ that depends on the chosen action a . Assume for simplicity that X is finite, each $M_a(\cdot|x)$ has full support, and the initial signal x_1 is drawn from some full-support distribution that can depend on a and n .*

This setting satisfies our assumptions, with a rate function that is the solution to the following optimization problem (e.g., Dembo and Zeitouni, 2010):

$$I_{a,a'}(\ell) = \inf_{\nu \in \Delta(X \times X)} \sum_x \nu(x) \text{KL}(\nu(\cdot|x), M_a(\cdot|x)) \quad \text{subject to}$$

$$\sum_{x,x'} \nu(x,x') \log \frac{M_{a^*}(x'|x)}{M_{a'}(x'|x)} = \ell \quad \text{and} \quad \sum_{x' \in X} \nu(x,x') = \sum_{x' \in X} \nu(x',x) \quad \text{for each } x \in X,$$

where $\nu(x) := \sum_{x'' \in X} \nu(x,x'')$ for each x and $\nu(x'|x) := \frac{\nu(x,x')}{\nu(x)}$ for each $x, x' \in X$ with $\nu(x) > 0$. In our main model where each $M_a(\cdot|x)$ is independent of x , this expression reduces to the rate function (9) from Cramér's theorem.³² ▲

³¹That is, $\{L \in \mathbb{R} : I_{a,a'}(L) \leq k\}$ is compact for each $k \in \mathbb{R}_+$.

³²This is because in this case it is without loss of generality to minimize over product measures

We can also capture a principal who, instead of observing many signals, observes a single signal, which becomes increasingly precise as $n \rightarrow \infty$:

Example A.2 (Vanishing observation noise). *Suppose $A \subseteq \mathbb{R}$. If the agent chooses action a , the principal observes signal $x = a + \kappa_n \varepsilon$, where ε is a random shock that follows the standard normal distribution.³³ Under the scaling factor $\kappa_n = \frac{1}{\sqrt{n}}$, we have*

$$I_{a,a'}(\ell) = \inf_{x \in \mathbb{R}} \frac{(x - a)^2}{2} \quad \text{subject to} \quad \frac{(x - a')^2 - (x - a^*)^2}{2} = \ell. \quad \blacktriangle$$

The following result extends Theorem 1 to the current setting:

Theorem A.2. *Under both general and binary contracts, the second-best cost converges to the first-best exponentially at rate $\rho(\mu) := \min_{a \in A^-} I_{a^*,a}(\hat{L}_a(a))$:*

$$C_n^{\text{SB}}(\mu, a^*, u, c) - C_n^{\text{FB}}(a^*, u, c) = \exp[-\rho(\mu)n + o(n)];$$

$$C_n^{\text{bin}}(\mu, a^*, u, c) - C_n^{\text{FB}}(a^*, u, c) = \exp[-\rho(\mu)n + o(n)].$$

To understand the convergence rate, recall that, for each $a \neq a^*$, $I_{a^*,a}(\hat{L}_a(a))$ measures how atypical it is, if the chosen action is a^* , to observe the score value $L_n(a) = \hat{L}_a(a)$ (i.e., the limit score under a deviation to a). Thus, $\rho(\mu) := \min_{a \in A^-} I_{a^*,a}(\hat{L}_a(a))$ is again a measure of the detectability of the hardest-to-detect deviation. Under i.i.d. signals, $I_{a^*,a}(\hat{L}_a(a)) = \text{KL}(\mu_a, \mu_{a^*})$, so Theorem A.2 reduces to Theorem 1. The proof of Theorem A.2 follows similar steps as that of Theorem 1 and is presented in Online Appendix B.1. To achieve the optimal convergence rate, we again construct binary contracts based on the maximally lenient thresholds $\gamma(a) \searrow \hat{L}_a(a)$ for each $a \in A^-$.

A.6.2 Adjustable Actions

Fix $T \in \mathbb{N} \setminus \{0\}$ and assume that the agent sequentially chooses T actions, $a_1, \dots, a_T \in A$. Each action a_t generates $\frac{n}{T}$ i.i.d. draws of signals, $x_t^n = (x_{t,1}, \dots, x_{t,\frac{n}{T}}) \in X^{\frac{n}{T}}$, from $\mu_{a_t} \in \Delta(X)$, where we restrict attention to n that are divisible by T and impose the same assumptions on the monitoring technology μ as in Section 2. Signals are

$\nu \in \Delta(X \times X)$, due to the convexity of KL-divergence.

³³We conjecture that we can also allow for shocks ε that follow other distributions (subject to some assumptions).

observed by both the principal and agent. Thus, a strategy of the agent is described by a collection of mappings $\alpha = (\alpha_t)_{t=1}^T$, where $\alpha_t : X^{\frac{(t-1)n}{T}} \rightarrow A$ specifies the agent's action a_t after each history of past signal realizations $(x_1^n, \dots, x_{t-1}^n)$ (and x_0^n denotes the empty history).

We continue to assume that the agent has additively separable payoffs, with a consumption utility $u : [\underline{w}, \infty) \rightarrow \mathbb{R}$ over money and a total cost $\frac{1}{T} \sum_{t=1}^T c(a_t)$ over action sequences a_1, \dots, a_T , where u and c satisfy Assumption 1.1–2. The principal designs a contract $v : X^n \rightarrow [\underline{w}, \infty)$ that specifies a one-shot utility payment $v(x^n)$ as a function of the entire sequence $x^n = (x_1^n, \dots, x_T^n)$ of n signals. In the current setting, the principal may want the agent to choose different actions a_t depending on past signal realizations $(x_1^n, \dots, x_{t-1}^n)$. To accommodate this, instead of exogenously fixing a target action profile, we assume that the risk-neutral principal receives a monetary payoff $\frac{1}{T} \sum_{t=1}^T g(a_t)$ from the agent's actions, for some function $g : A \rightarrow \mathbb{R}$.

In the second-best problem, the principal jointly chooses a contract v and agent-strategy α to maximize her payoff,

$$G_n^{\text{SB}} := \sup_{v, \alpha} \mathbb{E}_\alpha \left[\frac{1}{T} \sum_{t=1}^T g(a_t) - h(v(x^n)) \right],$$

subject to the IC and IR constraints

$$\mathbb{E}_\alpha \left[v(x^n) - \frac{1}{T} \sum_{t=1}^T c(a_t) \right] \geq \mathbb{E}_{\alpha'} \left[v(x^n) - \frac{1}{T} \sum_{t=1}^T c(a_t) \right], \quad \forall \alpha', \quad (\text{IC})$$

$$\mathbb{E}_\alpha \left[v(x^n) - \frac{1}{T} \sum_{t=1}^T c(a_t) \right] \geq 0. \quad (\text{IR})$$

Let G_n^{bin} denote the principal's value under the restriction to binary contracts.

As before, we also consider the first-best problem, which corresponds to only imposing (IR). We assume there is $a^* \in A$ such that it is strictly optimal to set $\alpha_t(\cdot) = a^*$ for all t under the first-best problem at each n .³⁴ Thus, the principal's first-best payoff is $G^{\text{FB}} = g(a^*) - h(c(a^*))$.

Given a^* , we impose Assumption 1.3 and let $A^- := \{a \in A : c(a) < c(a^*)\}$ and

³⁴For some specifications of g , u , and c , the first-best solution may require the agent to choose different actions a_t depending on past signal realizations. We rule out these cases to keep the departure from the main model parsimonious.

$A^+ := \{a \in A : c(a) > c(a^*)\}$. For technical convenience, we also impose the following richness condition on μ : For each $p \in (0, 1)$ and $a \in A$, there is a set of signals $Y \subseteq X$ such that $\mu_a(Y) = p$.³⁵

The following result fixes T and characterizes the optimal convergence rate to the first-best in the limit as $n \rightarrow \infty$:

Theorem A.3. *Under both general and binary contracts, the second-best payoff converges to the first-best exponentially at rate $\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$ as $n \rightarrow \infty$:*

$$G^{\text{FB}} - G_n^{\text{SB}} = \exp\left[-\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})n + o(n)\right];$$

$$G^{\text{FB}} - G_n^{\text{bin}} = \exp\left[-\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})n + o(n)\right].$$

Theorem A.3 shows that that the optimal convergence rate is again achieved by binary contracts, extending Theorem 1, which can be viewed as the special case $T = 1$. In the proof (Online Appendix B.2), we show that achieving the optimal convergence rate again relies on binary contracts that employ a sequence of maximally lenient tests. As noted in the main text, the convergence rate $\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$ is decreasing in the frequency T of the agent’s action adjustments, but is the same as in the setting where the agent chooses an action sequence (a_1, \dots, a_T) without observing the signals generated by past actions.

References

- BARRON, D., G. GEORGIADIS, AND J. SWINKELS (2020): “Optimal contracts with a risk-taking agent,” *Theoretical Economics*, 15(2), 715–761.
- BLACKWELL, D. (1951): “Comparison of experiments,” in *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102. University of California Press.
- CARROLL, G. (2015): “Robustness and linear contracts,” *American Economic Review*, 105(2), 536–63.
- COVER, T. M., AND J. A. THOMAS (1999): *Elements of Information Theory*. John Wiley & Sons.
- DEMBO, A., AND O. ZEITOUNI (2010): *Large Deviations Techniques and Applications*. Springer.
- DUNFORD, N., AND J. T. SCHWARTZ (1988): *Linear Operators, Part 1: General Theory*, vol. 10. John Wiley & Sons.
- DUPUIS, P., AND R. S. ELLIS (2011): *A weak convergence approach to the theory of large deviations*, vol. 902. John Wiley & Sons.

³⁵This is satisfied, for example, if we augment the monitoring technology μ by an uninformative signal that is drawn uniformly (independently of the agent’s action).

- DÜTTING, P., T. ROUGHGARDEN, AND I. TALGAM-COHEN (2019): “Simple versus optimal contracts,” in *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 369–387.
- FRICK, M., R. IJIMA, AND Y. ISHII (2023a): “Learning Efficiency of Multi-Agent Information Structures,” *Journal of Political Economy*.
- (2023b): “Welfare Comparisons for Biased Learning,” working paper.
- GEORGIADIS, G. (2022): “Contracting with Moral Hazard: A Review of Theory & Empirics,” working paper.
- GEORGIADIS, G., AND B. SZENTES (2020): “Optimal monitoring design,” *Econometrica*, 88(5), 2075–2107.
- GJESDAL, F. (1982): “Information and incentives: The agency information problem,” *Review of Economic Studies*, 49(3), 373–390.
- HAREL, M., E. MOSSEL, P. STRACK, AND O. TAMUZ (2021): “Rational groupthink,” *The Quarterly Journal of Economics*, 136(1), 621–668.
- HERWEG, F., D. MÜLLER, AND P. WEINSCHENK (2010): “Binary payment schemes: Moral hazard and loss aversion,” *American Economic Review*, 100(5), 2451–77.
- HOLMSTRÖM, B. (1979): “Moral hazard and observability,” *The Bell Journal of Economics*, 10(1), 74–91.
- HOLMSTRÖM, B., AND P. MILGROM (1987): “Aggregation and linearity in the provision of intertemporal incentives,” *Econometrica*, 55(2), 303–328.
- HONG, H., AND M. SHUM (2004): “Rates of information aggregation in common value auctions,” *Journal of Economic Theory*, 116(1), 1–40.
- HÖRNER, J., AND S. TAKAHASHI (2016): “How fast do equilibrium payoff sets converge in repeated games?,” *Journal of Economic Theory*, 165, 332–359.
- JEWITT, I. (2007): “Information order in decision and agency problems,” working paper.
- JOHNSTON, M. W., AND G. W. MARSHALL (2016): *Sales force management: Leadership, innovation, technology*. Routledge.
- JOSEPH, K., AND M. U. KALWANI (1998): “The role of bonus pay in salesforce compensation plans,” *Industrial Marketing Management*, 27(2), 147–159.
- KANDORI, M. (1992): “The use of information in repeated games with imperfect monitoring,” *The Review of Economic Studies*, 59(3), 581–593.
- KE, R., AND X. XU (2023): “The existence of an optimal deterministic contract in moral hazard problems,” *Economic Theory*, 76(2), 375–416.
- KIM, S. K. (1995): “Efficiency of an information system in an agency model,” *Econometrica*, 63(1), 89–102.
- LAFFONT, J.-J., AND D. MARTIMORT (2009): *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press.
- LEVIN, J. (2003): “Relational incentive contracts,” *American Economic Review*, 93(3), 835–857.
- LIAO, J., AND A. BERG (2019): “Sharpening Jensen’s inequality,” *The American Statistician*, 73(3), 278–281.

- LOPOMO, G., L. RIGOTTI, AND C. SHANNON (2011): “Knightian uncertainty and moral hazard,” *Journal of Economic Theory*, 146(3), 1148–1172.
- LUENBERGER, D. G. (1997): *Optimization by Vector Space Methods*. John Wiley & Sons.
- MIRRLEES, J. A. (1999): “The theory of moral hazard and unobservable behaviour: Part I,” *The Review of Economic Studies*, 66(1), 3–21.
- MOSCARINI, G., AND L. SMITH (2002): “The law of large demand for information,” *Econometrica*, 70(6), 2351–2366.
- MURPHY, K. J. (1999): “Executive compensation,” *Handbook of Labor Economics*, 3, 2485–2563.
- OYER, P. (2000): “A theory of sales quotas with limited liability and rent sharing,” *Journal of Labor Economics*, 18(3), 405–426.
- PAKZAD-HURSON, B. (2023): “Crowdsourcing and optimal market design,” working paper.
- PALOMINO, F., AND A. PRAT (2003): “Risk taking and optimal contracts for money managers,” *RAND Journal of Economics*, pp. 113–137.
- PRENDERGAST, C. (1999): “The provision of incentives in firms,” *Journal of Economic Literature*, 37(1), 7–63.
- RADNER, R. (1985): “Repeated principal-agent games with discounting,” *Econometrica*, 53(5), 1173–1198.
- ROUGHGARDEN, T., AND I. TALGAM-COHEN (2019): “Approximately optimal mechanism design,” *Annual Review of Economics*, 11, 355–381.
- RUSTICHINI, A., M. A. SATTERTHWAITE, AND S. R. WILLIAMS (1994): “Convergence to efficiency in a simple market with incomplete information,” *Econometrica*, 62(5), 1041–1063.
- SATTERTHWAITE, M. A., AND S. R. WILLIAMS (2002): “The optimality of a simple market mechanism,” *Econometrica*, 70(5), 1841–1863.
- SINGH, N. (1985): “Monitoring and hierarchies: The marginal value of information in a principal-agent model,” *Journal of Political Economy*, 93(3), 599–609.
- SUGAYA, T., AND A. WOLITZKY (2023a): “Monitoring versus Discounting in Repeated Games,” *Econometrica*, 91(5), 1727–1761.
- (2023b): “Performance Feedback in Long-Run Relationships: A Rate of Convergence Approach,” working paper.
- VIVES, X. (1993): “How fast do rational agents learn?,” *The Review of Economic Studies*, 60(2), 329–347.

Online Appendix to “Monitoring with Rich Data”

Mira Frick, Ryota Iijima, and Yuhta Ishii

B Proofs for Appendix A.6

B.1 Proof of Theorem A.2

B.1.1 Properties of Rate Functions

Let $I_{a,a'}$ denote the rate functions associated with sequence (μ^n) .

Lemma B.1. *For each $a \neq a^*$ and $\ell \in \mathbb{R}$, we have:*

1. $I_{a^*,a}(\ell) - I_{a,a}(\ell) = -\ell$.
2. $I_{a^*,a}(\ell) \geq I_{a^*,a}(\hat{L}_a(a))$ if $\ell \leq \hat{L}_a(a)$.

Proof. **First part:** For any $\varepsilon > 0$, let $\Gamma_\varepsilon := [\ell - \varepsilon, \ell + \varepsilon]$. By the median inequality,

$$\ell - \varepsilon \leq \frac{1}{n} \log \frac{\mathbb{P}_{a^*}[L_n(a) \in \Gamma_\varepsilon]}{\mathbb{P}_a[L_n(a) \in \Gamma_\varepsilon]} \leq \ell + \varepsilon.$$

Fix any $\varepsilon' > 0$. Since $I_{a^*,a}$ is lower-semicontinuous, by choosing ε sufficiently small, we have $I_{a^*,a}(\ell') \geq I_{a^*,a}(\ell) - \varepsilon'$ for all $\ell' \in \Gamma_\varepsilon$. Then

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log \frac{\mathbb{P}_{a^*}(L_n(a) \in \Gamma_\varepsilon)}{\mathbb{P}_a(L_n(a) \in \Gamma_\varepsilon)} &= - \inf_{\ell' \in \Gamma_\varepsilon} I_{a^*,a}(\ell') + \inf_{\ell' \in \Gamma_\varepsilon} I_{a,a}(\ell') \\ &\leq -I_{a^*,a}(\ell) + \varepsilon' + I_{a,a}(\ell), \end{aligned}$$

where the equality uses (20). Thus, $\ell - \varepsilon \leq -I_{a^*,a}(\ell) + \varepsilon' + I_{a,a}(\ell)$. Since $\varepsilon, \varepsilon'$ can be chosen arbitrarily small, we have $\ell \leq -I_{a^*,a}(\ell) + I_{a,a}(\ell)$. A symmetric argument yields $\ell \geq -I_{a^*,a}(\ell) + I_{a,a}(\ell)$.

Second part: Observe that for any $\ell \leq \hat{L}_a(a)$ with $I_{a^*,a}(\ell) < \infty$, we have

$$I_{a^*,a}(\ell) = I_{a,a}(\ell) - \ell \geq I_{a,a}(\hat{L}_a(a)) - \hat{L}_a(a) = I_{a^*,a}(\hat{L}_a(a)),$$

where the two equalities use the first part and the inequality uses $I_{a,a}(\hat{L}_a(a)) = 0 \leq I_{a,a}(\ell)$. \square

B.1.2 Upper Bound on Inefficiency under Binary Contracts

We first show that, as $n \rightarrow \infty$, $C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes at least as fast as at the exponential rate $\rho(\mu)$.

By the first part of Lemma B.1 and the identification condition, we have $-\hat{L}_a(a) = I_{a^*,a}(\hat{L}_a(a)) - I_{a,a}(\hat{L}_a(a)) = I_{a^*,a}(\hat{L}_a(a)) > 0$ and $-\hat{L}_{a^*}(a) = I_{a^*,a}(\hat{L}_{a^*}(a)) - I_{a,a}(\hat{L}_{a^*}(a)) = -I_{a,a}(\hat{L}_{a^*}(a)) < 0$ for each $a \neq a^*$. By the choice of $\gamma(\cdot)$, for each $a \in A^-$,

$$\inf_{\ell < \gamma(a)} I_{a^*,a}(\ell) > 0 \text{ and } \inf_{\ell \geq \gamma(a)} I_{a,a}(\ell) > 0.$$

Thus, by (20),

$$\lim_{n \rightarrow \infty} p_{a^*,n} = 1 \text{ and } \lim_{n \rightarrow \infty} p_{a,n} = 0 \text{ for each } a \in A^-. \quad (21)$$

As in the proof of Theorem 1, we can use (21) to construct binary test contracts (v_n) with thresholds $\gamma(a)$ that satisfy (IR) with equality and (IC) for large enough n . Note that

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \log(1 - p_{a^*,n}) &= \lim_{n \rightarrow \infty} \frac{1}{n} \log(\mathbb{P}_{a^*}[L_n(a) < \gamma(a) \text{ for some } a \in A^-]) \\ &= \max_{a \in A^-} \lim_{n \rightarrow \infty} \frac{1}{n} \log(\mathbb{P}_{a^*}[L_n(a) < \gamma(a)]) \\ &\leq - \min_{a \in A^-} \inf_{\ell \leq \gamma(a)} I_{a^*,a}(\ell). \end{aligned}$$

Thus, as in the proof of Theorem 1,

$$\begin{aligned} C_n^{\text{bin}}(\mu, u, a^*, c) - C^{\text{FB}}(u, a^*, c) &\leq \mathbb{E}_{a^*}[h(v_n(x^n))] - h(c(a^*)) \\ &\leq \exp[-n \min_{a \in A^-} \inf_{\ell \leq \gamma(a)} I_{a^*,a}(\ell) + o(n)]. \end{aligned}$$

Finally, note that for each $a \in A^-$,

$$\sup_{\gamma(a) > \hat{L}_a(a)} \inf_{\ell \leq \gamma(a)} I_{a^*,a}(\ell) = \inf_{\ell \leq \hat{L}_a(a)} I_{a^*,a}(\ell) = I_{a^*,a}(\hat{L}_a(a)),$$

where the second equality uses the second part of Lemma B.1 and the local continuity of $I_{a^*,a}$. Thus, by choosing $\gamma(a)$ arbitrarily close to $\hat{L}_a(a)$ for each $a \in A^-$, it follows that $C_n^{\text{bin}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes at least as fast as at the exponential rate $\min_{a \in A^-} I_{a^*,a}(\hat{L}_a(a)) = \rho(\mu)$.

B.1.3 Lower Bound on Inefficiency under General Contracts

Next, we show that, as $n \rightarrow \infty$, $C_n^{\text{SB}}(\mu, u, c, a^*) - C^{\text{FB}}(u, c, a^*)$ vanishes no faster than at the exponential rate $\rho(\mu)$.

Let $\hat{a} \in \operatorname{argmin}_{a \in A^-} I_{a^*,a}(\hat{L}_a(a))$, i.e., $\rho(\mu) = I_{a^*,\hat{a}}(\hat{L}_{\hat{a}}(\hat{a}))$. We consider a relaxed problem with action set $\{a^*, \hat{a}\}$. As in the proof of Theorem 1, it is sufficient to show that the difference between the second-best cost under the relaxed problem and $C^{\text{FB}}(u, c, a^*)$ cannot vanish faster than at the exponential rate $I_{a^*,\hat{a}}(\hat{L}_{\hat{a}}(\hat{a}))$.

As in the proof of Theorem 1, we consider a sequence of optimal contracts (v_n) for the relaxed problem, where each v_n is written as a weakly increasing function of L_n . The sequence is uniformly bounded by the same argument as in Appendix A.1.2.

Lemma A.2 from the proof of Theorem 1 and its proof generalizes to the current setting by using the general large-deviation principle (20) instead of Cramér's theorem. Lemma B.2 below generalizes Lemma A.3 to the current setting. Given this, the remaining proof follows the same arguments as the proof of Theorem 1.

Lemma B.2. *Consider any uniformly bounded sequence (v_n) of contracts satisfying (IC) under the relaxed problem. Then $\text{Var}_{a^*}[v_n(L_n)] \geq \exp[-I_{a^*,a}(\hat{L}_a(\hat{a}))n + o(n)]$.*

Proof. Suppose to the contrary that $\text{Var}_{a^*}[v_n(L_n)] \leq \exp[-\lambda n + o(n)]$ for some $\lambda > I_{a^*,\hat{a}}(\hat{L}_a(\hat{a}))$. By local continuity of $I_{a^*,\hat{a}}$, there is $\underline{\gamma} < \hat{L}_a(\hat{a})$ with $I_{a^*,\hat{a}}(\underline{\gamma}) < \lambda$. Moreover, by local continuity of $I_{a^*,\hat{a}}$ and the fact that $I_{a^*,\hat{a}}(\hat{L}_a(\hat{a})) = 0$, there is $\bar{\gamma} > \hat{L}_a(\hat{a})$ with $I_{a^*,\hat{a}}(\bar{\gamma}) < \lambda$. By (the generalized) Lemma A.2, $\lim_{n \rightarrow \infty} (v_n(\bar{\gamma}) - v_n(\underline{\gamma})) = 0$.

At the same time, for each $a \in \{a^*, \hat{a}\}$, the fact that $I_{a,\hat{a}}$ achieves 0 only at $\hat{L}_a(\hat{a}) \in (\underline{\gamma}, \bar{\gamma})$ and that the lower level sets are compact implies $\inf_{\gamma \in (-\infty, \underline{\gamma}] \cup [\bar{\gamma}, \infty)} I_{a,\hat{a}}(\gamma) > 0$. Thus, $\mathbb{P}_a[L_n(\hat{a}) \in [\underline{\gamma}, \bar{\gamma}]] \rightarrow 1$ by (20). As the sequence of contracts is uniformly bounded, this implies $\lim_{n \rightarrow \infty} (\mathbb{E}_{a^*}[v_n(L_n)] - \mathbb{E}_{\hat{a}}[v_n(L_n)]) = 0$. Then (IC) is violated for all large enough n , a contradiction. \square

B.2 Proof of Theorem A.3

B.2.1 Upper Bound on Inefficiency under Binary Contracts

We first show that, as $n \rightarrow \infty$, $G^{\text{FB}} - G_n^{\text{bin}}$ vanishes at least as fast as at the exponential rate $\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

For each $a \neq a^*$ and $t = 1, \dots, T$, define

$$L_{t,n}(a) := \frac{T}{n} \sum_{i=1}^{\frac{n}{T}} \log \frac{d\mu_{a^*}}{d\mu_a}(x_{t,i}).$$

For each $a \in A^-$, fix some $\gamma(a) \in (\mathbb{E}_a[L_{t,n}(a)], \mathbb{E}_{a^*}[L_{t,n}(a)])$; note that the expectations $\mathbb{E}_{a'}[L_{t,n}(a)]$ do not depend on n and t . For each $a' \in A$, define

$$p_{a',n} := \mathbb{P}_{a'}[L_{t,n}(a) \geq \gamma(a), \forall a \in A^-, t = 1, \dots, T].$$

By the weak law of large numbers and the choice of $\gamma(a)$, we have

$$\lim_{n \rightarrow \infty} p_{a^*,n} = 1 \text{ and } \lim_{n \rightarrow \infty} p_{a,n} = 0 \text{ for all } a \in A^-. \quad (22)$$

Consider the recommended strategy α^* given by $\alpha_1^* \equiv a$ and, for all $t = 2, \dots, T$,

$$\alpha_t^*(x_1^n, \dots, x_{t-1}^n) = \begin{cases} a^* & \text{if } L_{\tau,n}(a) \geq \gamma(a) \text{ for all } a \in A^-, \tau = 1, \dots, t-1 \\ \underline{a} & \text{otherwise} \end{cases}$$

for some $\underline{a} \in \arg \min_{a \in A} c(A)$. Consider the sequence of binary contracts (v_n) of the form

$$v_n(x^n) = \begin{cases} v_n^+ & \text{if } L_{k,n}(a) \geq \gamma(a) \text{ for all } a \in A^-, k = 1, \dots, T \\ v_n^- & \text{otherwise.} \end{cases}$$

Here

$$v_n^+ := \frac{(1 - \bar{p}_n p_{a^*,n}^{T-1})\beta_n^* - (1 - p_{a^*,n}^T)\beta_n}{p_{a^*,n}^{T-1}(p_{a^*,n} - \bar{p}_n)}, \quad v_n^- := \frac{p_{a^*,n}^T \beta_n - \bar{p}_n p_{a^*,n}^{T-1} \beta_n^*}{p_{a^*,n}^{T-1}(p_{a^*,n} - \bar{p}_n)},$$

where

$$\bar{p}_n := \max_{a \in A^-} p_{a,n}, \quad \beta_n^* := \mathbb{E}_{\alpha^*} \left[\frac{\sum_{t=1}^T c(a_t)}{T} \right], \quad \beta_n := \min_{a \in A^-} \mathbb{E}_{\alpha_a^*} \left[\frac{\sum_{t=1}^T c(a_t)}{T} \right],$$

and α_a^* denotes the strategy defined by a one-shot deviation from α^* to $a \neq a^*$ at $t = 1$. Note that, by (22),

$$\lim_{n \rightarrow \infty} v_n^+ = \lim_{n \rightarrow \infty} \beta_n^* = c(a^*) \text{ and } \lim_{n \rightarrow \infty} v_n^- = \lim_{n \rightarrow \infty} \beta_n = c(\underline{a}), \quad (23)$$

so $v_n^+, v_n^- \in u([\underline{w}, \infty))$ for all sufficiently large n . Moreover, contracts (v_n) satisfy

$$p_{a^*,n}^T v_n^+ + (1 - p_{a^*,n}^T) v_n^- - \beta_n^* = 0,$$

$$p_{a^*,n}^T v_n^+ + (1 - p_{a^*,n}^T) v_n^- - \beta_n^* = \bar{p}_n p_{a^*,n}^{T-1} v_n^+ + (1 - \bar{p}_n p_{a^*,n}^{T-1}) v_n^- - \beta_n.$$

The first line implies that (IR) binds under strategy α^* , as the LHS is the agent's expected payoff under α^* . The second line ensures that one-shot deviations from α^* to any action in A^- at $t = 1$ are not strictly profitable for the agent.

Conditional on any history at which α^* prescribes \underline{a} , the wage becomes constant, so the agent has no incentive to deviate. Let $U_{n,t}(a)$ denote the agent's expected total payoff from a one-shot deviation from α^* to $a \neq a^*$ at some t and history $(x_1^n, \dots, x_{t-1}^n)$ at which α^* prescribes a^* . (Note that this value is independent of the particular history). Observe that, absent any deviation, the agent's expected total payoff conditional on such histories tends to 0 as $n \rightarrow \infty$ by (22)-(23) and binding (IR), where the convergence is uniform across histories.

For any $a \in A^-$ and $t \geq 2$, (22)-(23) imply

$$\lim_{n \rightarrow \infty} U_{n,t}(a) = \frac{tc(\underline{a}) - c(a) - (t-1)c(a^*)}{T} < 0,$$

where the inequality follows from $c(\underline{a}) \leq c(a) < c(a^*)$. Thus, there is no incentive to deviate to actions in A^- at large enough n . For any $a \in A^+$ and $t \geq 1$, (22)-(23) imply

$$\limsup_{n \rightarrow \infty} U_{n,t}(a) = \limsup_{n \rightarrow \infty} p_{a,n} \frac{c(a^*) - c(a)}{T} + (1 - p_{a,n}) \frac{tc(\underline{a}) - c(a) - (t-1)c(a^*)}{T} < 0,$$

where the inequality follows from $c(\underline{a}) < c(a^*) < c(a)$. Thus, there is no incentive to deviate to actions in A^+ at large enough n .

As in the proof of Theorem 1, we can show that $1 - p_{a^*,n} = \exp[-\frac{n}{T} \min_{a \in A^-} \inf_{\ell < \gamma(a)} I_{a^*,a}(\ell) + o(\frac{n}{T})]$. Moreover, observe that $1 - p_{a^*,n}^T$ and $\beta_n^* - c(a^*)$ also vanish at the same rate as $1 - p_{a^*,n}$. Thus,

$$\begin{aligned} & G^{\text{FB}} - G_n^{\text{bin}} \\ & \leq \sum_{t=1}^{T-1} p_{a^*,n}^{t-1} (1 - p_{a^*,n}) \frac{t(g(a^*) - g(\underline{a}))}{T} + p_{a^*,n}^T (h(v_n^+) - h(c(a^*))) \\ & \quad + (1 - p_{a^*,n}^T) (h(v_n^-) - h(c(a^*))) \\ & \leq \sum_{t=1}^{T-1} p_{a^*,n}^{t-1} (1 - p_{a^*,n}) \frac{t(g(a^*) - g(\underline{a}))}{T} + p_{a^*,n}^T (h(v_n^+) - h(c(a^*))) \\ & = \exp[-\frac{n}{T} \min_{a \in A^-} \inf_{\ell < \gamma(a)} I_{a^*,a}(\ell) + o(\frac{n}{T})], \end{aligned}$$

where the last line follows from observing that

$$v_n^+ - c(a^*) = \frac{(1 - p_{a^*,n}^T)(\beta_n^* - \beta_n)}{p_{a^*,n}^{T-1}(p_{a^*,n} - \bar{p}_n)} + \beta_n^* - c(a^*)$$

vanishes at the same exponential rate as $1 - p_{a^*,n}$.

The remaining step of letting $\gamma(a) \searrow \mathbb{E}_{a^*}[L_{t,n}(a)]$ for each $a \in A^-$ is analogous to the proof of Theorem 1.

B.2.2 Lower Bound on Inefficiency under General Contracts

Next, we show that, as $n \rightarrow \infty$, $G^{\text{FB}} - G_n^{\text{SB}}$ vanishes no faster than at the exponential rate $\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

By standard arguments, G_n^{SB} can be written in a recursive manner by taking the agent's continuation value as a state variable: Let $G_{n,T+1}^{\text{SB}}(z) := -h(z)$ for each

$z \in u([\underline{w}, \infty))$. For each $t = T, T-1, \dots, 1$ and $z \in u([\underline{w}, \infty)) - \frac{T-t+1}{T}c(A)$, we inductively define $G_{n,t}^{\text{SB}}(z) := \max_{a \in A} G_{n,t}^{\text{SB}}(z, a)$, where for each $a \in A$,

$$G_{n,t}^{\text{SB}}(z, a) := \sup_{\psi: X^{\frac{n}{T}} \rightarrow u([\underline{w}, \infty)) - \frac{T-t}{T}c(A)} \frac{b(a)}{T} + \mathbb{E}_a[G_{n,t+1}(\psi(x_t^n))]$$

subject to IC

$$-\frac{c(a)}{T} + \mathbb{E}_a[\psi(x_t^n)] \geq -\frac{c(a')}{T} + \mathbb{E}_{a'}[\psi(x_t^n)] \quad \forall a' \in A \quad (24)$$

and promise keeping

$$-\frac{c(a)}{T} + \mathbb{E}_a[\psi(x_t^n)] = z, \quad (25)$$

where $\psi(\cdot)$ specifies the agent's continuation value as a function of the realized x_t^n .³⁶ The optimal value of the principal is then given by $G_n^{\text{SB}} = G_{n,1}^{\text{SB}}(0)$.

The first-best value can be written analogously: Let $G_{n,T+1}^{\text{FB}}(z) := -h(z)$ for each $z \in u([\underline{w}, \infty))$. For each $t = T, T-1, \dots, 1$ and $z \in u([\underline{w}, \infty)) - \frac{T-t+1}{T}c(A)$, we inductively define $G_{n,t}^{\text{FB}}(z) := \max_{a \in A} G_{n,t}^{\text{FB}}(z, a)$, where for each $a \in A$,

$$G_{n,t}^{\text{FB}}(z, a) := \sup_{\psi: X^{\frac{n}{T}} \rightarrow u([\underline{w}, \infty)) - \frac{T-t}{T}c(A)} \frac{b(a)}{T} + \mathbb{E}_a[G_{n,t+1}^{\text{FB}}(\psi(x_t^n))]$$

subject to promise keeping (25). Then $G^{\text{FB}} = G_{n,1}^{\text{FB}}(0)$. By the richness condition on μ_a , in the above problem the principal can induce any randomization over continuation values subject to promise keeping by appropriately choosing ψ . Thus, by standard arguments for information design problems,

$$G_{n,t}^{\text{FB}}(z, a) = \frac{b(a)}{T} + \text{cav} G_{n,t+1}^{\text{FB}} \left(z + \frac{c(a)}{T} \right),$$

where cav denotes the concavification operator. Consequently, $G_{n,t}^{\text{FB}}$ is independent of n , so we drop the subscript n and simply write G_t^{FB} .

As we show in Appendix B.2.3 below, there is a C^2 function $f : u([\underline{w}, \infty)) - \frac{(T-1)c(A)}{T} \rightarrow \mathbb{R}$ that upper-bounds G_2^{FB} and such that $f' < 0$, $f'' < 0$, $\inf_w f'(w) = -\infty$, and $f(\frac{c(a^*)}{T}) = \frac{(T-1)b(a^*)}{T} + h(c(a^*))$.

Since $f(\cdot) \geq G_{n,2}^{\text{SB}}(\cdot)$, we have that $G_{n,1}^{\text{SB}}(0, a^*)$ is weakly lower than the value of the

³⁶The existence of an optimal policy in this problem is guaranteed by analogous arguments as in Appendix A.1.2 when n is sufficiently large.

problem

$$\max_{\psi: X^{\frac{n}{T}} \rightarrow u(\underline{w}, \infty) - \frac{T-1}{T}c(A)} \frac{b(a^*)}{T} + \mathbb{E}_{a^*}[f(\psi(x_1^n))] = \frac{b(a^*)}{T} - \min_{\psi: X^{\frac{n}{T}} \rightarrow u(\underline{w}, \infty) - \frac{T-1}{T}c(A)} \mathbb{E}_{a^*}[-f(\psi(x_1^n))] \quad (26)$$

subject to (24) and (25) with $a = a^*$ and $z = 0$.

Observe that the minimization problem in (26) can be viewed as the second-best problem of our baseline model in Section 2, by treating $\frac{c(\cdot)}{T}$ and $-f(\cdot)$ as the agent's cost function and inverse utility function, respectively. Thus, by Theorem 1, the difference between $\frac{b(a^*)}{T} + f(\frac{c(a^*)}{T}) = b(a^*) + h(c(a^*))$ and the value in (26) is $\exp[-\frac{n}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*}) + o(\frac{n}{T})]$. Therefore, $G_{n,1}^{\text{SB}}(0, a^*)$ cannot converge to $b(a^*) + h(c(a^*)) = G^{\text{FB}}$ faster than at rate $\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

Finally, for any $a \neq a^*$, we have $G_1^{\text{FB}}(0, a) < G_1^{\text{FB}}(0, a^*) = G^{\text{FB}}$ by the assumption that it is strictly optimal to implement $a_1 = a^*$ under the first-best problem. Thus, $G_{n,1}^{\text{SB}}(0, a)$, which is weakly lower than $G_1^{\text{FB}}(0, a)$, is bounded away from the first-best G^{FB} . Combined with the above observations, this shows that $G_n^{\text{SB}} = \max_{a \in A} G_{n,1}^{\text{SB}}(a^*)$ cannot converge to $b(a^*) + h(c(a^*)) = G^{\text{FB}}$ faster than at rate $\frac{1}{T} \min_{a \in A^-} \text{KL}(\mu_a, \mu_{a^*})$.

B.2.3 Details of Bounding G_2^{FB}

Since by construction $\text{cav}G_2^{\text{FB}}$ is a concave and decreasing function, it is dominated by some linear and decreasing function. To ensure that the bound f can be chosen to be appropriately concave, we make some preliminary observations.

Because it is strictly optimal under the first-best problem to implement a^* at all t and $\frac{(t-1)c(a^*)}{T}$ is the agent's (on-path) continuation value, we have

$$\text{cav}G_t^{\text{FB}} \left(\frac{(t-1)c(a^*)}{T} + \varepsilon \right) = G_t^{\text{FB}} \left(\frac{(t-1)c(a^*)}{T} + \varepsilon \right) = \frac{(T-t+1)g(a^*)}{T} - h(c(a^*) + \varepsilon)$$

whenever $|\varepsilon|$ is sufficiently small. This in particular implies that $\text{cav}G_2^{\text{FB}} \left(\frac{c(a^*)}{T} \right) = \frac{(T-1)g(a^*)}{T} - h(c(a^*))$, and hence

$$\text{cav}(G_2^{\text{FB}})' \left(\frac{c(a^*)}{T} \right) = -h'(c(a^*)) < 0 \quad \text{and} \quad \text{cav}(G_2^{\text{FB}})'' \left(\frac{c(a^*)}{T} \right) = -h''(c(a^*)) < 0. \quad (27)$$

Let \hat{a} be an action that maximizes $b(\cdot)$ among $\arg \min_{a \in A} c(a)$. Since $h'(z) \rightarrow \infty$ as $z \rightarrow \sup_w u(w)$, we have $\text{cav}G_T^{\text{FB}}(z) = G_T^{\text{FB}}(z) = \frac{g(\hat{a})}{T} - h(z + \frac{c(\hat{a})}{T})$ for all large enough z . By induction, for each $t = T, T-1, \dots, 1$ and sufficiently large z ,

$$\text{cav}G_t^{\text{FB}}(z) = G_t^{\text{FB}}(z) = \frac{(T-t+1)g(\hat{a})}{T} - h(z).$$

Thus,

$$\text{cav}(G_2^{\text{FB}})'(z) \rightarrow -\infty \quad (28)$$

as z tends to the upper-bound of the range of $\text{cav}(G_2^{\text{FB}})$. Then based on (27)-(28), there is a C^2 function f that upper-bounds G_2^{FB} and such that $f' < 0$, $f'' < 0$, $\inf_w f'(w) = -\infty$, and $f(\frac{c(a^*)}{T}) = G_2^{\text{FB}}(\frac{c(a^*)}{T}) = \frac{(T-1)b(a^*)}{T} + h(c(a^*))$.