What to Measure

Optimal Monitoring under Adverse Selection and Moral Hazard

Xinrong Liu*

September 15, 2026

Abstract

A principal often chooses not only how much information to collect, but also what to measure. I study this choice in a contracting problem with both hidden type and hidden action, where an existing performance measure is informative about both dimensions and the principal can design an additional, scarce monitoring signal. I show that the optimal measurement need not target the dimension that is statistically least informative. Instead, it depends jointly on the incentive requirements of the contract and the information already contained in existing performance. These forces are summarized by an information-adjusted incentive requirement: the local incentive requirements premultiplied by the inverse Fisher information matrix of the existing measure. The optimal direction of scarce monitoring is the leading eigenvector of the second-moment matrix of these adjusted requirements. Existing information can therefore make it optimal for a new signal to respond to type and action with opposite signs, even when both underlying incentive requirements are positive. Under strict local and global implementation margins, the same direction governs the fully incentive-compatible problem. The results distinguish statistical informativeness from contracting value and characterize what a principal should measure first when monitoring capacity is scarce.

1 Introduction

Principals increasingly choose not only how much information to collect, but also what to measure. An insurer can monitor speeding, hard braking, mileage, or driving conditions. An employer can measure output, work processes, or worker characteristics. These measures need not reveal the same hidden dimension. Some are more informative about an agent's action, others about his type, and many respond to both.

Which information should the principal acquire?

A natural answer is to improve what is currently measured least precisely. But this answer treats monitoring as a statistical problem. A principal uses information to solve an incentive problem. What matters is therefore not only which hidden dimension is difficult to infer, but also which distinctions are valuable for implementing the desired allocation. A direction can be statistically poorly measured and yet have little contracting value. Conversely, information about a relatively well-measured direction can remain valuable when incentives place strong demands on it.

*Department of Economics, The University of Texas at Austin.

This distinction becomes particularly important when the same performance measure contains information about several hidden dimensions. Consider an insurer contracting with a driver whose risk characteristics are privately known and whose care is unobserved. Driving records may respond to both. Hard braking, for example, may reflect behavior, underlying risk, or both. Improving the measurement of the statistically least precise dimension need not be the best use of monitoring capacity because the insurer uses the data both to provide incentives and to screen different types.

This paper studies this problem in a contracting environment with adverse selection and moral hazard. An agent privately observes his type and then chooses an unobserved action. The principal already observes a performance measure whose distribution depends on both type and action. She can design an additional signal that responds to a linear combination of the two hidden dimensions, subject to a scarce monitoring capacity. The question is which direction in the joint action–type space should receive the first units of additional measurement.

The answer depends jointly on what incentives require and what the existing performance measure already reveals. Local action and reporting incentives generate a vector of incentive requirements, denoted by $\Delta(\theta)$. If $I^y(\theta)$ is the Fisher information matrix of the existing performance measure, the relevant object for monitoring is

$$w(\theta) = I^y(\theta)^{-1}\Delta(\theta).$$

This information-adjusted incentive requirement captures the interaction between the economic demand for information and its existing statistical supply. The optimal direction of scarce monitoring is determined by the leading direction of the second moment of $w(\theta)$ across types.

This characterization has several implications. First, the statistically weakest direction need not be the economically most valuable direction to measure. Statistical precision matters, but only after being weighted by the incentive problem the information is used to solve. Even when the existing performance measure is statistically symmetric across hidden dimensions, incentive requirements alone can make one direction more valuable than another.

Second, the relevant adjustment is joint rather than coordinate by coordinate. When the existing performance measure responds similarly to type and action, information about the two dimensions is confounded. The best new measure may then respond to them in opposite directions. In particular, a signal can optimally be positively sensitive to action and negatively sensitive to type even when the underlying incentive requirements are positive in both dimensions. The negative sensitivity does not mean that type information is undesirable. It allows the new measure to separate variation that the existing performance measure mixes together.

Third, the value of monitoring is locally quadratic in measurement capacity. The first units of a common-index signal change Fisher information only at second order, and the direction of that information determines how much it reduces contracting costs. This gives both a direction and a cardinal value for scarce monitoring, rather than only a ranking of exogenously given information systems.

The analysis first obtains this characterization from the local incentive conditions associated with obedience and truthful reporting. With hidden type and hidden action, however, an agent can deviate jointly in his report and action, so these conditions need not by themselves guarantee implementation. I show that when the benchmark allocation is strictly locally implementable and deviations farther from the truthful allocation are uniformly unprofitable, sufficiently small changes in monitoring preserve these properties. The optimal direction obtained from the local problem then also governs the fully incentive-compatible problem. When an implementation

margin binds, by contrast, the cost of preserving implementation can affect the optimal direction at the same order as the informational gain.

A numerical insurance illustration shows that these forces can generate substantially different prescriptions. The existing driving record in the illustration is informative about both type and care. A statistical criterion recommends improving measurement almost entirely in the type direction. The contracting criterion instead favors a signal that is primarily sensitive to care and negatively sensitive to type. The example also shows that the local approximation remains accurate over a nontrivial range of monitoring capacity while the relevant implementation constraints remain slack.

The paper contributes to the literature on information and monitoring in agency problems by shifting the question from how informative a performance measure is to what it should be informative about. Classical informativeness results compare exogenously given information systems, while monitoring-design problems typically study information about a single hidden action. Here the principal chooses a direction of information when both hidden type and hidden action matter for contracting. This makes the value of a measurement depend on the interaction between its statistical content and multiple incentive requirements. The resulting information-adjusted incentive requirement summarizes how existing information affects the cost of satisfying these incentive requirements.

The paper also relates to the literature on multidimensional performance measurement, where the value of a measure depends on both its sensitivity and its precision, and to models combining adverse selection and moral hazard. The distinction here is that the sensitivities of the new measure are themselves design variables. Existing information can therefore affect not only how much a new measure is worth, but also which combination of hidden dimensions it should respond to. Section 2 discusses these connections in detail.

The remainder of the paper is organized as follows. Section 2 discusses the related literature. Section 3 introduces the model and monitoring-design problem. Section 4 derives the local incentive requirements and the cost of providing them with the existing information structure. Section 5 characterizes the optimal direction of scarce monitoring. Section 6 develops its economic implications. Section 7 studies local and full implementation. Section 8 provides the insurance illustration, and Section 9 discusses extensions beyond the baseline environment.

2 Related literature

The value of information in agency. The principle that a contract should condition on any signal carrying incremental information about a hidden action is due to Holmström (1979) and Shavell (1979). A subsequent literature ranks *exogenous* information systems by their value for contracting. Gjesdal (1982) shows that informativeness for incentive purposes is not the same as Blackwell informativeness; Kim (1995) establishes that a mean-preserving spread of the distribution of the likelihood ratio is sufficient to rank systems for a broad class of preferences; and Jewitt (1988) provides conditions on the score that validate the first-order approach. More recently, Chaigneau et al. (2019) show that the informativeness condition can fail without the first-order approach and give a generalized condition based on non-local constraints.

We differ from this literature in both the question and the instrument. Rather than ranking given information systems, we take the existing information as given and ask what additional information the principal should acquire. In particular, we ask which direction in the two-dimensional space of hidden action and hidden type a marginal unit of measurement should

target. Its value depends jointly on the incentive requirements and on what the existing performance measure already reveals about the two dimensions.

Monitoring design. The closest antecedent is Georgiadis and Szentes (2020), who study a principal who acquires signals about hidden effort at constant marginal cost. Their problem is organized around a one-dimensional score—the derivative of the log-density with respect to effort—and can be recast as an information-design problem in which the principal chooses a distribution over scores at a cost equal to its variance. The optimal acquisition policy has two thresholds and implements a single-bonus contract. Earlier work on the principal's monitoring decision includes Dye (1986) and Lambert (1985). In recent work, Sun and Zhang (2026) bound the number of outcomes of an optimal experiment by an "effective policy dimension" that counts independent post-signal control margins.

Our problem differs in that monitoring can provide information about both a hidden action and a hidden type. These two dimensions generate distinct incentive requirements, and the existing performance measure may already provide correlated information about them. We therefore fix the amount of additional monitoring capacity and ask what the new signal should measure. When capacity is small, the answer is determined by the interaction between the incentive requirements and the Fisher information of the existing measure, with the optimal direction characterized by a leading eigenvector. Relative to Sun and Zhang (2026), who characterize how coarse an optimal experiment can be, we characterize what a marginal experiment should measure.

Sensitivity and precision. The idea that the value of a performance measure depends on what it is sensitive *to*, and not only on how precisely it measures it, originates in the performance measurement literature. Banker and Datar (1989) show that signals are optimally aggregated linearly with weights proportional to the product of sensitivity and precision. Feltham and Xie (1994) define congruity as the alignment between a measure's vector of marginal sensitivities and the principal's objective, and show that a precise but incongruent measure has limited value; Holmström and Milgrom (1991) and Baker (2002) develop the associated distortion. Our information-adjusted incentive requirement

$$w_\sigma(\theta) = \left(I^y\right)^{-1}\Delta_\sigma(\theta)$$

extends this sensitivity–precision logic to a setting with two hidden dimensions. The vector Δ_σ describes the incentive requirements, while $(I^y)^{-1}$ accounts for what the existing performance measure already reveals about them. Because this adjustment is determined by a matrix rather than a scalar, the two dimensions cannot in general be evaluated separately. In particular, when the action and type scores of the existing measure are positively correlated, the optimal new measure can respond to the two dimensions in opposite directions even though both incentive requirements are positive. Such a reversal cannot arise from a scalar sensitivity–precision comparison. Our design problem therefore asks not only how informative a performance measure should be, but what it should measure.

Adverse selection with moral hazard. Models combining hidden type and hidden action go back to Laffont and Tirole (1986, 1993), Guesnerie and Laffont (1984), Picard (1987), and McAfee and McMillan (1987). Gottlieb and Moreira (2022) show that with limited liability

a single contract is often optimal because menus create gaming opportunities, while Castro-Pires et al. (2024) develop a decoupling method that solves a moral-hazard cost problem type by type and then incorporates the resulting cost into a screening problem. Our formulation of incentive compatibility—truthful reporting and obedience jointly against all deviations in report and action—follows the coordination-mechanism structure of Myerson (1982). Myerson also observes that with hidden action, unlike in pure screening, the output message may need to conceal the agent's report. Doval and Skreta (2022) make what the designer learns an explicit object of design. Under limited commitment, the designer may prefer to learn less, whereas our principal can commit and chooses where to allocate scarce monitoring capacity.

We do not solve the joint allocation problem. Instead, we extract its two local incentive requirements and characterize their cost using the same information matrix. The interaction between adverse selection and moral hazard then appears through the off-diagonal terms of this matrix rather than through a sequencing of two separate problems. Our result therefore concerns the demand for information generated by the joint incentive problem, rather than the shape of the optimal contract.

Fisher information and the cost of incentives. The methodological antecedent is Chade and Swinkels (2020), who show that in a pure moral-hazard problem the shadow value of the incentive constraint is proportional to the reciprocal of the Fisher information of the signal about effort. The same object appears implicitly in Holmström and Milgrom (1987), where the optimal piece rate varies inversely with the noise of the aggregate and, in the Gaussian case, the signal-to-noise ratio coincides with Fisher information. Our cheapest-contract theorem extends this relation to two incentive requirements. Because the score has mean zero, the cost of a contract separates into a mean component and an incentive component, with the latter equal to $\Delta_\sigma^\top I^{-1}\Delta_\sigma$ and the optimal contract affine in the score. To our knowledge, this Cramér–Rao form has not been used to characterize the joint cost of action and reporting incentives. Conditional independence makes Fisher information additive across the existing and new signals, while the common-index monitoring technology contributes a rank-one information increment. Together, these properties yield a simple characterization of the optimal monitoring direction when capacity is small.

Signals responsive to both type and action. Finally, our monitoring technology is related to the signal structures studied in the muddled-information literature. In Frankel and Kartik (2019), an observed signal responds jointly to an agent's natural action and gaming ability, so that equilibrium behavior confounds the two dimensions, and Frankel and Kartik (2022) show that a decision-maker should commit to underutilizing such data. Ball (2025) characterizes the accuracy-maximizing score for a strategic agent, which underweights some features to deter distortion and overweights others to maintain accuracy on average. Closest in objective is Hörner and Lambert (2021), who solve for the effort-maximizing rating in the standard career-concerns model. They show that the optimal rating is linear in past observations and that higher effort can reduce some future ratings. This sign reversal is related to our confounding result, although it arises from the dynamics of reputation rather than from correlation between the action and type scores of an existing measure.

These papers either take the sensitivities of the signal as given, choose them to improve predictive accuracy, or design a signal when incentives are provided by the market rather than by a contract. We instead choose how the new signal responds to type and action in order to reduce

the cost of providing incentives. The key tradeoff is between measuring action and separating action variation from the type variation already contained in the existing performance measure.

3 Model and design problem

3.1 Environment

Type and action. A principal contracts with an agent with private information. The agent's type is $\theta \in \Theta = [\underline{\theta}, \bar{\theta}]$, drawn from F with continuous density $f > 0$ and privately observed. After learning his type, the agent chooses an action $a \in A = [\underline{a}, \bar{a}]$. His cost is $c(\theta, a)$, twice continuously differentiable with

$$c_a(\theta, a) > 0, \qquad c_{aa}(\theta, a) > 0.$$

Existing performance measure. The principal observes a performance measure $y \in \mathcal{Y}$ with distribution

$$q(y \mid \theta, a) > 0.$$

The probabilities are twice continuously differentiable in (θ, a). The performance measure is exogenous: the principal can condition transfers on y, but cannot change how informative y is about type or action. We take $\mathcal{Y}$ to be finite for simplicity.

Preferences. The agent's utility from a transfer $t \geq 0$ is

$$\sqrt{t} - c(\theta, a),$$

and his reservation utility is zero. The principal is risk neutral and obtains gross benefit $b(y)$ when a type-θ agent takes action a, where $b : \mathcal{Y} \to \mathbb{R}$. Her payoff is

$$b(y) - t.$$

The gross benefit and the existing performance measure play distinct roles. The former determines the productive value of the agent's type and action, while the latter provides information that can be used for contracting. The standard case in which the observed performance measure is itself and $\mathcal{Y} \subset \mathcal{R}$, output can be nested by setting

$$b(y) = \mathbb{E}_{\theta,a}[y].$$

It is convenient to write transfers in utility units,

$$v = \sqrt{t} \geq 0, \qquad t = v^2.$$

The agent's payoff is then linear in v, while the principal's cost of delivering v is v^2.

3.2 Monitoring design

In addition to y, the principal can design a signal $s \in \mathcal{S}$ that is informative about both the agent's action and his type. Fix reference points $a_0 \in A$ and $\theta_0 \in \Theta$, and define

$$x(\theta, a) = \begin{pmatrix} a - a_0 \\ \theta - \theta_0 \end{pmatrix}.$$

A monitoring design is a vector

$$p = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}, \qquad \|p\| \leq r, \tag{1}$$

where α determines the sensitivity of the new signal to action and β its sensitivity to type. The signal distribution depends on (θ, a) through the index

$$z(\theta, a; p) = p^\top x(\theta, a) = \alpha(a - a_0) + \beta(\theta - \theta_0), \tag{2}$$

with

$$\Pr(s = j \mid \theta, a; p) = g_j\big(z(\theta, a; p)\big), \qquad j \in \mathcal{S}. \tag{3}$$

Each g_j is positive and twice continuously differentiable, with $\sum_j g_j(z) = 1$.

Conditional on (θ, a), the existing performance measure and the designed signal are independent. Their joint distribution is therefore

$$\phi(y, s \mid \theta, a; p) = q(y \mid \theta, a) g_s\big(z(\theta, a; p)\big). \tag{4}$$

The constraint $\|p\| \leq r$ limits the sensitivity of the additional measurement. For $p \neq 0$, write

$$p = r_p \hat{u}, \qquad r_p = \|p\|, \qquad \|\hat{u}\| = 1. \tag{5}$$

The magnitude r_p describes how much measurement is used, while $\hat{u}$ describes what the additional signal measures. Our interest is in environments in which the available capacity r is small and in the direction selected as $r \downarrow 0$.

3.3 Contracts and incentives

By the generalized revelation principle (Myerson 1982), we restrict attention to direct coordination mechanisms in which the agent reports his type and receives an action recommendation. A mechanism is a pair

$$M = (e, v),$$

where

$$e : \Theta \to A$$

specifies the recommended action for each reported type, and

$$v : \Theta \times \mathcal{Y} \times \mathcal{S} \to \mathbb{R}_+$$

specifies utility from the transfer as a function of the report and the realized observables. We write

$$v_{\hat{\theta}}(y, s) \equiv v(\hat{\theta}, y, s),$$

so that the corresponding monetary transfer is $v_{\hat{\theta}}(y,s)^2$.

After observing his type θ, the agent reports $\hat{\theta} \in \Theta$, receives the recommendation $e(\hat{\theta})$, and then privately chooses an action $\hat{a} \in A$. Given a monitoring design p, his expected utility from $(\hat{\theta}, \hat{a})$ is

$$\Pi_p(\theta, \hat{\theta}, \hat{a}; M) = \sum_{y \in \mathcal{Y}} \sum_{s \in \mathcal{S}} \phi(y, s \mid \theta, \hat{a}; p) v_{\hat{\theta}}(y, s) - c(\theta, \hat{a}). \tag{6}$$

His utility under truthful reporting and obedience is

$$U_p(\theta; M) = \Pi_p\big(\theta, \theta, e(\theta); M\big). \tag{7}$$

A mechanism is incentive compatible under design p if truthful reporting and obedience jointly dominate every deviation in report and action:

$$U_p(\theta; M) \geq \Pi_p(\theta, \hat{\theta}, \hat{a}; M) \qquad \text{for all } \theta \in \Theta \text{ and } (\hat{\theta}, \hat{a}) \in \Theta \times A. \tag{IC}$$

It is individually rational if

$$U_p(\theta; M) \geq 0 \qquad \text{for all } \theta \in \Theta. \tag{IR}$$

3.4 The principal's design problem

For a fixed monitoring design p, the principal chooses a direct mechanism $M = (e, v)$ to maximize

$$V(p) = \sup_M \int_\Theta \left\{ b\big(\theta, e(\theta)\big) - \mathbb{E}_{\theta, e(\theta)} \left[v_\theta(y,s)^2 \right] \right\} dF(\theta) \tag{8}$$

subject to (IC) and (IR). The monitoring problem is

$$W(r) = \sup_{\|p\| \leq r} V(p), \tag{9}$$

and let $p^{F,*}(r)$ denote an optimal design.

Our object of interest is the direction of optimal measurement as capacity becomes scarce:

$$\lim_{r \downarrow 0} \frac{p^{F,*}(r)}{r}. \tag{10}$$

This asks which combination of action and type should receive the first units of additional measurement capacity.

4 Incentive requirements and existing information

4.1 Local incentive requirements

The joint incentive constraint in (IC) restricts all report and action deviations. We begin by extracting the local restrictions implied by this constraint. These restrictions define the first-order relaxation that will be used to characterize the value of monitoring.

For a fixed monitoring design p, define

$$\ell(y, s \mid \theta, a; p) = \begin{pmatrix} \ell_a(y, s \mid \theta, a; p) \\ \ell_\theta(y, s \mid \theta, a; p) \end{pmatrix} = \begin{pmatrix} \partial_a \log \phi(y, s \mid \theta, a; p) \\ \partial_\theta \log \phi(y, s \mid \theta, a; p) \end{pmatrix}. \tag{11}$$

The two components measure how the likelihood of an observation (y, s) changes locally with the agent's action and type. In statistical terminology, ℓ is the score vector of the experiment.

Write $\mathbb{E}_{\theta,a}$ for expectation under $\phi(\cdot, \cdot \mid \theta, a; p)$. For any function $h : \mathcal{Y} \times \mathcal{S} \to \mathbb{R}$ that does not directly depend on (θ, a),

$$\partial_a \mathbb{E}_{\theta,a}[h] = \mathbb{E}_{\theta,a}[h\ell_a], \qquad \partial_\theta \mathbb{E}_{\theta,a}[h] = \mathbb{E}_{\theta,a}[h\ell_\theta]. \tag{12}$$

Taking $h \equiv 1$ also gives

$$\mathbb{E}_{\theta,a}[\ell_a] = \mathbb{E}_{\theta,a}[\ell_\theta] = 0.$$

These identities give a convenient representation of the local restrictions implied by incentive compatibility. Obedience requires the recommended action $e(\theta)$ to be optimal for a truthful type θ. Since (12) gives

$$\partial_a \mathbb{E}_{\theta,a}[v_\theta] = \mathbb{E}_{\theta,a}[v_\theta \ell_a],$$

the action first-order condition is

$$\mathbb{E}_{\theta,e(\theta)}\left[v_\theta \ell_a\right] = c_a\big(\theta, e(\theta)\big). \tag{OB}$$

Truthful reporting similarly gives the envelope condition

$$U'(\theta) = \mathbb{E}_{\theta,e(\theta)}\left[v_\theta \ell_\theta\right] - c_\theta\big(\theta, e(\theta)\big). \tag{EN}$$

Proposition 1 (Local incentive requirements)**.** *Suppose that $e(\theta) \in \operatorname{int} A$ and that the mechanism satisfies* (IC)*. Then* (OB) *holds for every θ. Moreover, U is Lipschitz and* (EN) *holds at every point at which U is differentiable.*

These two restrictions can be written as a single vector requirement. An allocation is a pair $\sigma = (e, U)$, where $e : \Theta \to \operatorname{int} A$ is measurable and $U : \Theta \to \mathbb{R}_+$ is Lipschitz. Let Σ denote the set of such allocations. For $\sigma \in \Sigma$, define

$$\Delta_\sigma(\theta) = \begin{pmatrix} c_a\big(\theta, e(\theta)\big) \\ U'(\theta) + c_\theta\big(\theta, e(\theta)\big) \end{pmatrix}. \tag{13}$$

Then any incentive-compatible mechanism inducing σ must satisfy

$$\mathbb{E}_{\theta,e(\theta)}\left[v_\theta \ell\right] = \Delta_\sigma(\theta) \tag{14}$$

for almost every θ.

The two components of Δ_σ correspond to the two incentive problems. The first captures the payment sensitivity needed to induce the recommended action and is determined by its marginal cost. The second captures the payment sensitivity needed for truthful reporting and depends on both the direct effect of type on cost and the slope of equilibrium utility $U'(\theta)$.

4.2 The cost of providing incentives

The cost of satisfying the first-order incentive requirements depends on the information contained in the observed outcomes. Define

$$I(\theta, a; p) = \mathbb{E}_{\theta,a}\left[\ell\ell^\top\right]. \tag{15}$$

This is the Fisher information matrix of the joint experiment about (a, θ).

For the existing performance measure y, let

$$\ell^y(y \mid \theta, a) = \begin{pmatrix} \partial_a \log q(y \mid \theta, a) \\ \partial_\theta \log q(y \mid \theta, a) \end{pmatrix},$$

with Fisher information matrix

$$I^y(\theta, a) = \sum_{y \in \mathcal{Y}} q(y \mid \theta, a) \ell^y(y \mid \theta, a) \ell^y(y \mid \theta, a)^\top. \tag{16}$$

Assumption 1 (Nondegenerate baseline information)**.** There exists $\varepsilon > 0$ such that

$$I^y(\theta, a) \succeq \varepsilon I_2$$

for all $(\theta, a) \in \Theta \times A$.

Thus the existing performance measure provides locally nondegenerate information about both hidden dimensions.

For an allocation $\sigma = (e, U)$, equilibrium utility requires

$$\mathbb{E}_{\theta, e(\theta)}[v_\theta] = U(\theta) + c\big(\theta, e(\theta)\big).$$

Define

$$\bar{v}_\sigma(\theta) = U(\theta) + c\big(\theta, e(\theta)\big). \tag{17}$$

Together with (14), implementing σ in the first-order relaxation requires

$$\mathbb{E}_{\theta, e(\theta)}[v_\theta] = \bar{v}_\sigma(\theta), \qquad \mathbb{E}_{\theta, e(\theta)}[v_\theta \ell] = \Delta_\sigma(\theta). \tag{18}$$

Theorem 1 (Cheapest contract)**.** *Fix an allocation $\sigma = (e, U)$, a type θ, and a monitoring design p. Let*

$$I = I\big(\theta, e(\theta); p\big)$$

and suppose that I is nonsingular. Ignoring the nonnegativity constraint on v, the unique solution to

$$\min_v \mathbb{E}_{\theta, e(\theta)}[v^2] \quad \textit{subject to} \quad \mathbb{E}_{\theta, e(\theta)}[v] = \bar{v}_\sigma(\theta), \qquad \mathbb{E}_{\theta, e(\theta)}[v\ell] = \Delta_\sigma(\theta) \tag{19}$$

is

$$v^*_\sigma = \bar{v}_\sigma(\theta) + \Delta_\sigma(\theta)^\top I^{-1} \ell. \tag{20}$$

The minimum cost is

$$\bar{v}_\sigma(\theta)^2 + \Delta_\sigma(\theta)^\top I^{-1} \Delta_\sigma(\theta). \tag{21}$$

Proof. Since $\mathbb{E}[\ell] = 0$, any feasible contract can be written as

$$v = \bar{v}_\sigma(\theta) + \lambda^\top \ell + \varepsilon,$$

where $\mathbb{E}[\varepsilon] = 0$ and $\mathbb{E}[\varepsilon \ell] = 0$. The incentive requirement implies

$$I\lambda = \Delta_\sigma(\theta),$$

so

$$\lambda = I^{-1}\Delta_\sigma(\theta).$$

Orthogonality then gives

$$\mathbb{E}[v^2] = \bar{v}_\sigma(\theta)^2 + \Delta_\sigma(\theta)^\top I^{-1}\Delta_\sigma(\theta) + \mathbb{E}[\varepsilon^2],$$

which is uniquely minimized at $\varepsilon = 0$. □

The cheapest contract separates into a constant component that determines its mean and a variable component that provides incentives. Since the score has mean zero, these two components are orthogonal. The resulting contracting cost therefore separates into the cost of the mean payment and the additional cost of providing incentives. The cost of the latter is

$$\Delta_\sigma(\theta)^\top I^{-1}\Delta_\sigma(\theta).$$

An increase in Fisher information lowers the cost of satisfying a fixed incentive requirement.

Using the contracting cost in (21), the value of allocation σ in the first-order relaxation is

$$V_\sigma^{\mathrm{FOC}}(p) = \int_\Theta \left\{ b\big(\theta, e(\theta)\big) - \bar{v}_\sigma(\theta)^2 - \Delta_\sigma(\theta)^\top I\big(\theta, e(\theta); p\big)^{-1}\Delta_\sigma(\theta) \right\} dF(\theta). \tag{22}$$

Optimizing over allocations gives the value of the first-order relaxation:

$$V^{\mathrm{FOC}}(p) = \sup_{\sigma\in\Sigma} V_\sigma^{\mathrm{FOC}}(p). \tag{23}$$

The solution in Theorem 1 ignores limited liability. We assume that this constraint remains uniformly slack near the benchmark optimum.

Assumption 2 (Local slackness of limited liability)**.** There exist $\varepsilon_0 > 0$, $\bar{r} > 0$, and $\eta > 0$ such that, for every allocation σ satisfying

$$V_\sigma^{\mathrm{FOC}}(0) \geq V^{\mathrm{FOC}}(0) - \varepsilon_0,$$

every monitoring design p with $\|p\| \leq \bar{r}$, and almost every θ, the unconstrained optimal contract in (20) satisfies

$$v_\sigma^*(y, s) \geq \eta$$

for every $(y, s) \in \mathcal{Y} \times \mathcal{S}$.

The first-order relaxation imposes the local restrictions implied by incentive compatibility but not the full joint-deviation constraints. Hence its value bounds the value of the original contracting problem.

Proposition 2 (Relaxation bound)**.** *For every feasible monitoring design p,*

$$V(p) \leq V^{\mathrm{FOC}}(p). \tag{24}$$

Proof. Any mechanism satisfying (IC) induces an allocation satisfying (18) by Proposition 1. Its expected transfer is therefore bounded below by (21) type by type. Integrating over types and optimizing over allocations gives $V(p) \leq V^{\mathrm{FOC}}(p)$. □

4.3 The information-adjusted incentive requirement

At the benchmark $p = 0$, the additional signal is independent of (θ, a) and therefore contributes no Fisher information. Hence

$$I(\theta, a; 0) = I^y(\theta, a). \tag{25}$$

For a fixed allocation σ, the incentive cost for type θ is therefore

$$\Delta_\sigma(\theta)^\top I^y\big(\theta, e(\theta)\big)^{-1}\Delta_\sigma(\theta). \tag{26}$$

The value of additional information depends not only on the incentive requirement $\Delta_\sigma(\theta)$, but also on how well the existing performance measure already provides information in the relevant directions. This motivates the information-adjusted incentive requirement

$$w_\sigma(\theta) = I^y\big(\theta, e(\theta)\big)^{-1}\Delta_\sigma(\theta). \tag{27}$$

The vector $w_\sigma(\theta)$ adjusts the incentive requirement for the information already available from the existing performance measure.

For example, if

$$I^y = \begin{pmatrix} \iota_a & 0 \\ 0 & \iota_\theta \end{pmatrix},$$

then

$$w_\sigma = \begin{pmatrix} \Delta_{\sigma,a}/\iota_a \\ \Delta_{\sigma,\theta}/\iota_\theta \end{pmatrix}. \tag{28}$$

Thus a dimension receives greater weight when incentives require a stronger response along that dimension or when the existing performance measure is less informative about it. When I^y is not diagonal, the adjustment is joint: the relevant object is $(I^y)^{-1}\Delta_\sigma$ rather than separate coordinate-by-coordinate comparisons.

Indeed, for a fixed incentive requirement, a perturbation dI of the information matrix changes incentive cost by

$$d\left(\Delta_\sigma^\top I^{-1}\Delta_\sigma\right) = -\Delta_\sigma^\top I^{-1}(dI)I^{-1}\Delta_\sigma = -w_\sigma^\top(dI)w_\sigma$$

at the benchmark.

5 The optimal direction of scarce monitoring

5.1 Information supplied by the monitoring technology

We first characterize how a small monitoring design changes the information available about the agent's action and type. Since the signal distribution depends on (θ, a) only through

$$z = p^\top x(\theta, a),$$

its score with respect to (a, θ) is

$$\nabla_{(a,\theta)} \log g_s(z) = \frac{g_s'(z)}{g_s(z)} p.$$

Thus every realization of the signal provides information along the same direction p.

Define

$$I_g(z) = \sum_{j \in \mathcal{S}} \frac{g_j'(z)^2}{g_j(z)}, \tag{29}$$

and assume $I_g(0) > 0$. The scalar $I_g(z)$ measures how informative the signal is about its index z, while p determines which combination of action and type the index represents.

Because y and s are conditionally independent given (θ, a), their Fisher information matrices add.

Proposition 3 (Information supplied by the monitoring technology)**.** *Let*

$$I^s(\theta, a; p) = I(\theta, a; p) - I^y(\theta, a).$$

Then

$$I^s(\theta, a; p) = I_g(z)\, pp^\top. \tag{30}$$

In particular, I^s has rank at most one and

$$\operatorname{tr} I^s(\theta, a; p) = I_g(z)\|p\|^2. \tag{31}$$

The rank-one structure follows from the common-index specification: regardless of the number of signal realizations, the monitoring technology provides information about only one linear combination of the two hidden dimensions.

Write

$$p = r_p \hat{u}, \qquad r_p = \|p\|, \qquad \|\hat{u}\| = 1.$$

Since $z = r_p \hat{u}^\top x(\theta, a)$, uniformly over (θ, a) and $\hat{u}$,

$$\operatorname{tr} I^s(\theta, a; p) = I_g(0) r_p^2 + O(r_p^3). \tag{32}$$

Hence, near the uninformative benchmark, the magnitude r_p determines how much information the technology supplies, while $\hat{u}$ determines what that information is about. At finite r_p, however, $I_g(z)$ may depend on both the state and the direction, so the Euclidean capacity constraint should not be interpreted as an exact trace constraint.

5.2 The value of a rank-one information update

The effect of the new information on incentive cost follows from a rank-one inverse update.

Lemma 1 (Exact gain from a rank-one update)**.** *Let B be symmetric positive definite, let $\hat{u}$ be a unit vector, let $\kappa > 0$, and let $\Delta \in \mathbb{R}^2$. Then*

$$\Delta^\top B^{-1}\Delta - \Delta^\top \big(B + \kappa \hat{u}\hat{u}^\top\big)^{-1}\Delta = \frac{\kappa\big(\hat{u}^\top B^{-1}\Delta\big)^2}{1 + \kappa \hat{u}^\top B^{-1}\hat{u}}. \tag{33}$$

Proof. By the Sherman–Morrison formula,

$$\big(B + \kappa \hat{u}\hat{u}^\top\big)^{-1} = B^{-1} - \frac{\kappa B^{-1}\hat{u}\hat{u}^\top B^{-1}}{1 + \kappa \hat{u}^\top B^{-1}\hat{u}}.$$

Pre- and post-multiplying by $\Delta^\top$ and Δ gives (33). □

Apply Lemma 1 at the benchmark with

$$B = I^y\big(\theta, e(\theta)\big), \qquad w_\sigma(\theta) = B^{-1}\Delta_\sigma(\theta).$$

For a design $p = r_p\hat{u}$, the new signal contributes

$$I^s = I_g(0)r_p^2\hat{u}\hat{u}^\top + O(r_p^3).$$

It follows that, for a fixed allocation, the reduction in incentive cost for type θ is

$$I_g(0)\big(p^\top w_\sigma(\theta)\big)^2 + O(\|p\|^3). \tag{34}$$

Equation (34) gives the economic meaning of the information-adjusted incentive requirement. For a single type, the leading value of additional information is determined by the projection of $w_\sigma(\theta)$ onto the direction measured by the technology. Among directions with the same capacity, the principal therefore prefers to measure along $w_\sigma(\theta)$.

With heterogeneous types, however, the vectors $w_\sigma(\theta)$ generally point in different directions. A single monitoring technology must choose one direction for all types. For a fixed allocation σ, define

$$M_\sigma = \int_\Theta w_\sigma(\theta)w_\sigma(\theta)^\top \, dF(\theta). \tag{35}$$

Integrating (34) across types gives the leading gain

$$I_g(0)p^\top M_\sigma p. \tag{36}$$

Thus, conditional on an allocation, the preferred monitoring direction solves

$$\max_{\|\hat{u}\|=1} \hat{u}^\top M_\sigma \hat{u}. \tag{37}$$

It is therefore a leading eigenvector of M_σ.

The matrix M_σ aggregates the information-adjusted incentive requirements across types. A direction receives greater weight when the requirements are large in that direction and when they are aligned across types.

5.3 Optimal monitoring with an endogenous allocation

The preceding argument holds the allocation fixed. The principal, however, re-optimizes both the allocation and the contract when monitoring changes. We now show that this re-optimization does not change the leading-order value of scarce monitoring.

Let σ_0 be a benchmark-optimal allocation and define

$$M = M_{\sigma_0} = \int_\Theta w_{\sigma_0}(\theta)w_{\sigma_0}(\theta)^\top \, dF(\theta). \tag{38}$$

Assumption 3 (Well-posed benchmark)**.** The supremum defining $V^{\mathrm{FOC}}(0)$ is finite and attained. Moreover, for every sequence $(\sigma_n) \subset \Sigma$ satisfying

$$V^{\mathrm{FOC}}_{\sigma_n}(0) \longrightarrow V^{\mathrm{FOC}}(0),$$

we have

$$M_{\sigma_n} \longrightarrow M,$$

where $M = M_{\sigma_0}$ for any benchmark-optimal allocation σ_0.

Assumption 3 does not require the benchmark allocation to be unique. It requires only that benchmark near-optimizers agree asymptotically on the matrix that determines the leading value of monitoring. If two benchmark-optimal allocations generated different matrices, an arbitrarily small amount of monitoring could select between them, and the local value would instead involve the maximum of several quadratic forms.

The following result makes the fixed-allocation expansion uniform enough to allow the allocation to be re-optimized.

Lemma 2 (Uniform bound on the gain from monitoring)**.** *There exists a finite constant C such that, for every allocation σ and all sufficiently small r,*

$$0 \leq \sup_{\|p\| \leq r} \left\{ V_\sigma^{\mathrm{FOC}}(p) - V_\sigma^{\mathrm{FOC}}(0) \right\} \leq C r^2 \left\{ \bar{b} - V_\sigma^{\mathrm{FOC}}(0) \right\}, \tag{39}$$

where

$$\bar{b} = \max_{(\theta,a) \in \Theta \times A} b(\theta, a).$$

The bound implies that an allocation that is substantially suboptimal at the benchmark cannot become optimal after an arbitrarily small improvement in monitoring. Hence allocations selected by sufficiently small designs must be benchmark near-optimizers.

We also use the following standard implication of a uniform quadratic expansion.

Lemma 3 (A quadratic expansion selects the leading eigenspace)**.** *Let $c > 0$ and let M be symmetric positive semidefinite with $\nu_1(M) > 0$. Suppose*

$$W(p) = W_0 + c\, p^\top M p + \epsilon_r(p),$$

where

$$\sup_{\|p\| \leq r} |\epsilon_r(p)| = o(r^2).$$

Then

$$\sup_{\|p\| \leq r} W(p) = W_0 + c r^2 \nu_1(M) + o(r^2).$$

Every maximizer $p^*(r)$ satisfies

$$\frac{\|p^*(r)\|}{r} \longrightarrow 1.$$

If $\nu_1(M)$ is simple, with unit eigenvectors $\{\hat{u}_1, -\hat{u}_1\}$, then

$$\operatorname{dist}\left(\frac{p^*(r)}{r}, \{\hat{u}_1, -\hat{u}_1\} \right) \longrightarrow 0.$$

Proof. The value statement follows from

$$\max_{\|p\| \leq r} p^\top M p = r^2 \nu_1(M).$$

Since $\nu_1(M) > 0$, any design using a fixed fraction less than one of the available capacity loses a fixed fraction of the leading r^2 gain. Uniformity of the remainder therefore implies

$$\frac{\|p^*(r)\|}{r} \longrightarrow 1.$$

Optimality then implies that the Rayleigh quotient of $p^*(r)/\|p^*(r)\|$ converges to $\nu_1(M)$. When the leading eigenvalue is simple, the maximizers of the Rayleigh quotient are exactly $\{\hat{u}_1, -\hat{u}_1\}$. □

We can now characterize the optimal direction of scarce monitoring.

Theorem 2 (Optimal direction of scarce monitoring)**.** *Under Assumptions 1–3, uniformly over* $\|p\| \leq r$*,*

$$V^{\mathrm{FOC}}(p) = V^{\mathrm{FOC}}(0) + I_g(0) p^\top M p + o(r^2). \tag{40}$$

If $\nu_1(M) > 0$*, then*

$$\sup_{\|p\| \leq r} V^{\mathrm{FOC}}(p) = V^{\mathrm{FOC}}(0) + I_g(0) r^2 \nu_1(M) + o(r^2), \tag{41}$$

and every maximizer $p^{\mathrm{FOC},*}(r)$ *satisfies*

$$\frac{\|p^{\mathrm{FOC},*}(r)\|}{r} \longrightarrow 1. \tag{42}$$

If, in addition, $\nu_1(M)$ *is simple, then*

$$\operatorname{dist}\left(\frac{p^{\mathrm{FOC},*}(r)}{r}, \{\hat{u}_1, -\hat{u}_1\}\right) \longrightarrow 0, \tag{43}$$

where $\{\hat{u}_1, -\hat{u}_1\}$ *are the unit eigenvectors associated with* $\nu_1(M)$*.*

The leading term separates the technological and economic determinants of the value of monitoring. The scalar $I_g(0)$ measures how much information a unit of monitoring capacity generates and therefore scales the value of monitoring. It does not affect the optimal direction. The matrix M determines where monitoring should be directed. Its leading eigenvector identifies the combination of action and type for which additional information is most valuable for contracting.

Proof sketch. For a fixed allocation σ, Proposition 3 and Lemma 1 imply, uniformly over $\|p\| \leq r$,

$$V_\sigma^{\mathrm{FOC}}(p) = V_\sigma^{\mathrm{FOC}}(0) + I_g(0) p^\top M_\sigma p + O(r^3).$$

By Lemma 2, an allocation that is optimal under a design p with $\|p\| \leq r$ must be a benchmark near-optimizer as $r \downarrow 0$. Assumption 3 therefore implies

$$M_\sigma \longrightarrow M$$

along the relevant allocations. This gives the uniform expansion (40). Lemma 3 then gives (41)–(43). □

Remark 1 (Sign). The leading objective

$$\hat{u}^\top M \hat{u}$$

is unchanged when $\hat{u}$ is replaced by $-\hat{u}$. The second-order analysis therefore identifies an axis rather than an orientation. What is determined is the relative sign of the action and type components of the optimal direction.

This does not imply exact symmetry at finite capacity. For a general common-index technology, $g_j(z)$ need not be equivalent to $g_j(-z)$ after relabeling signal realizations, so higher-order terms may distinguish $\hat{u}$ from $-\hat{u}$.

6 Understanding the result

The optimal monitoring direction is determined jointly by the incentive requirements and the information already contained in the existing performance measure. This section develops three implications of that result. First, statistical difficulty and economic value are distinct. Second, when the existing performance measure is jointly informative about action and type, the two dimensions cannot be evaluated separately. Third, adverse selection and moral hazard generate their local incentive requirements through different economic channels, even though both requirements are evaluated using the same information matrix.

Throughout this section, M refers to the matrix associated with the first-order relaxation. It is constructed from the local obedience and envelope conditions (OB) and (EN). These conditions are necessary for full incentive compatibility, but they do not incorporate the second-order or global restrictions associated with joint deviations. Accordingly, the interpretation below concerns the local value of information in the first-order relaxation. Conditions under which the same direction governs the full problem are given in Section 7.

6.1 Statistical difficulty is not economic value

A natural statistical benchmark is to measure where the existing performance measure is least informative. For a unit direction d, the quantity

$$d^\top I^y(\theta, a)^{-1} d$$

measures local statistical difficulty along d. Under the quadratic criterion

$$\max_{\|d\|=1} d^\top I^y(\theta, a)^{-1} d, \tag{44}$$

the statistically weakest direction is the leading eigenvector of $I^y(\theta, a)^{-1}$, equivalently the smallest-eigenvalue direction of $I^y(\theta, a)$.[1]

The contracting problem generally selects a different direction. For the benchmark allocation σ_0,

$$M = \int_\Theta I^y\big(\theta, e_0(\theta)\big)^{-1} \Delta_{\sigma_0}(\theta) \Delta_{\sigma_0}(\theta)^\top I^y\big(\theta, e_0(\theta)\big)^{-1} \, dF(\theta). \tag{45}$$

Thus the value of measurement depends not only on what the existing performance measure reveals poorly, but also on the incentive requirements that additional information helps satisfy. A statistically weak direction can have little contracting value when incentives place little demand on it. Conversely, a direction that is already measured relatively well can remain valuable when the corresponding incentive requirement is sufficiently strong.

The distinction is especially transparent when the existing information matrix is proportional to the identity:

$$I^y\big(\theta, e_0(\theta)\big) = \iota I_2. \tag{46}$$

There is then no statistical asymmetry between action and type. Nevertheless,

$$w_{\sigma_0}(\theta) = \frac{1}{\iota} \Delta_{\sigma_0}(\theta),$$

[1] This comparison is relative to the chosen normalization of the action and type coordinates.

and hence

$$M = \frac{1}{\iota^2} \int_{\Theta} \Delta_{\sigma_0}(\theta) \Delta_{\sigma_0}(\theta)^\top \, dF(\theta). \tag{47}$$

The optimal monitoring direction is therefore determined entirely by the magnitude and alignment of the incentive requirements across types.

This aggregation is important. The optimal direction is not generally obtained by comparing the action and type components of $\Delta_{\sigma_0}(\theta)$ pointwise. Different types may generate requirements pointing in different directions, and the relevant object is their second-moment matrix. A coordinate-by-coordinate comparison is valid only in special cases, such as a single type, a common direction of the requirements across types, or a diagonal aggregate matrix.

The optimal monitor therefore does not target statistical uncertainty by itself. It targets directions in which additional information is valuable for satisfying the incentive requirements of the contracting problem.

6.2 Confounding can make the optimal signal respond oppositely to type and action

The information-adjusted requirement

$$w(\theta) = \left(I^y\right)^{-1} \Delta(\theta)$$

need not point in the same direction as the underlying incentive requirement. In particular, correlation between the existing action and type scores can make the optimal new signal respond in opposite directions to action and type even when both incentive requirements are positive. This effect can be characterized in closed form when incentive requirements differ across types only in magnitude.

Suppose that the existing information matrix is constant across types and write

$$I^y = \begin{pmatrix} \iota_a & \rho\sqrt{\iota_a \iota_\theta} \\ \rho\sqrt{\iota_a \iota_\theta} & \iota_\theta \end{pmatrix}, \qquad \iota_a, \iota_\theta > 0, \qquad |\rho| < 1. \tag{48}$$

Here ρ is the correlation between the action and type scores of the existing performance measure. A positive value means that the existing performance measure tends to respond in the same direction to changes in action and type.

Consider the special case in which the incentive requirements have a common direction across types:

$$\Delta(\theta) = h(\theta)\delta, \qquad \delta = \begin{pmatrix} \delta_a \\ \delta_\theta \end{pmatrix}, \qquad h(\theta) > 0, \tag{49}$$

where

$$\delta_a > 0, \qquad \delta_\theta > 0.$$

Different types may therefore generate incentive requirements of different magnitudes, but the relative action and type requirements are the same for all types. In the absence of correlation in the existing scores, both components of the information-adjusted requirement are then positive.

Proposition 4 (Confounding and opposite signal sensitivities)**.** *Suppose* (48) *and* (49) *hold. Then M has rank one, and the optimal measurement direction is uniquely determined up to sign by*

$$\hat{u}_1 \propto \left(I^y\right)^{-1} \delta \propto \begin{pmatrix} \iota_\theta \delta_a - \rho\sqrt{\iota_a \iota_\theta}\delta_\theta \\ \iota_a \delta_\theta - \rho\sqrt{\iota_a \iota_\theta}\delta_a \end{pmatrix}. \tag{50}$$

In particular, define

$$\rho_\theta = \sqrt{\frac{\iota_a}{\iota_\theta}}\frac{\delta_\theta}{\delta_a}, \qquad \rho_a = \sqrt{\frac{\iota_\theta}{\iota_a}}\frac{\delta_a}{\delta_\theta}. \tag{51}$$

If

$$\rho_\theta < \rho < \rho_a, \tag{52}$$

then the optimal direction can be oriented so that

$$\hat{u}_{1,a} > 0 \qquad \textit{and} \qquad \hat{u}_{1,\theta} < 0. \tag{53}$$

Hence the optimal new signal is positively sensitive to action and negatively sensitive to type even though both components of the underlying incentive requirement are positive.

Moreover,

$$\rho_a \rho_\theta = 1. \tag{54}$$

Consequently, whenever $\rho_\theta < 1$, there is a nonempty range

$$\rho \in (\rho_\theta, 1)$$

over which (53) *holds.*

Proof. Under (49),

$$w(\theta) = \big(I^y\big)^{-1}\Delta(\theta) = h(\theta)\big(I^y\big)^{-1}\delta.$$

All information-adjusted requirements therefore point in the same direction. It follows that

$$M = \int_\Theta w(\theta)w(\theta)^\top \, dF(\theta) = \left(\int_\Theta h(\theta)^2 \, dF(\theta)\right)\big(I^y\big)^{-1}\delta\delta^\top\big(I^y\big)^{-1}.$$

Since $\delta \neq 0$, this matrix has rank one, and its unique nonzero eigenspace is spanned by

$$\big(I^y\big)^{-1}\delta.$$

From (48),

$$\big(I^y\big)^{-1} = \frac{1}{\det I^y}\begin{pmatrix} \iota_\theta & -\rho\sqrt{\iota_a\iota_\theta} \\ -\rho\sqrt{\iota_a\iota_\theta} & \iota_a \end{pmatrix}.$$

Since $\det I^y > 0$, the action component of $\big(I^y\big)^{-1}\delta$ is positive if and only if

$$\iota_\theta\delta_a > \rho\sqrt{\iota_a\iota_\theta}\delta_\theta,$$

or equivalently

$$\rho < \rho_a.$$

Its type component is negative if and only if

$$\iota_a\delta_\theta < \rho\sqrt{\iota_a\iota_\theta}\delta_a,$$

or equivalently

$$\rho > \rho_\theta.$$

Both inequalities therefore hold exactly on the interval (52). Finally,

$$\rho_a\rho_\theta = \left(\sqrt{\frac{\iota_\theta}{\iota_a}}\frac{\delta_a}{\delta_\theta}\right)\left(\sqrt{\frac{\iota_a}{\iota_\theta}}\frac{\delta_\theta}{\delta_a}\right) = 1,$$

which gives (54). If $\rho_\theta < 1$, then $\rho_a = 1/\rho_\theta > 1$. Since admissible correlations satisfy $\rho < 1$, every $\rho \in (\rho_\theta, 1)$ therefore satisfies (52). □

The proposition separates two forces that are otherwise combined in M. The vector δ describes what the contracting problem requires, while I^y describes how the existing performance measure already responds to the two hidden dimensions. With uncorrelated scores, $\rho = 0$, equation (50) becomes

$$\hat{u}_1 \propto \begin{pmatrix} \delta_a/\iota_a \\ \delta_\theta/\iota_\theta \end{pmatrix},$$

so positive action and type requirements translate into positive sensitivities of the new signal. Correlation changes this conclusion because the inverse information matrix adjusts for information that is already shared across the two dimensions.

To see the economic force, suppose $\rho > 0$. The existing performance measure then responds similarly to higher action and higher type, making the two changes difficult to separate. If the action requirement is sufficiently important relative to the type requirement, adding a signal that also increases with both dimensions continues this confounding. Once

$$\rho > \sqrt{\frac{\iota_a}{\iota_\theta}}\frac{\delta_\theta}{\delta_a}, \tag{55}$$

the principal instead benefits from a signal that increases with action but decreases with type. The opposite type sensitivity helps distinguish the action variation relevant for incentives from the type variation already embedded in the existing performance measure.

This does not mean that type information has negative value. Indeed, $\delta_\theta > 0$ throughout the proposition. Nor does it mean that the principal wants to conceal type. The result concerns the direction in which an *additional* signal should respond once the existing information structure is taken as given. A negative type sensitivity can be optimal precisely because the existing performance measure already mixes type and action in the same direction.

The threshold in (55) also shows how the contracting problem and the information structure interact. A larger relative type requirement δ_θ/δ_a raises the amount of confounding needed before opposite sensitivities become optimal. Greater existing precision about action relative to type, measured by ι_a/ι_θ, does the same. Stronger positive correlation ρ works in the opposite direction. Thus whether the optimal signal should respond to type and action in the same or in opposite directions cannot be determined from the incentive requirements or from the existing information matrix separately.

The common-direction assumption makes this interaction transparent because it removes disagreement across types: M is rank one and aggregation introduces no additional force. With heterogeneous directions of $\Delta(\theta)$, the optimal direction instead depends on the second-moment matrix

$$M = \int_\Theta w(\theta)w(\theta)^\top \, dF(\theta),$$

so there need not be a single closed-form threshold. The same mechanism remains, however, and will also appear in the numerical illustration after the implementation analysis.

6.3 Why action and type generate different requirements

Although action and type enter the same information matrix, their incentive requirements arise from different economic problems. Recall that

$$\Delta_\sigma(\theta) = \begin{pmatrix} c_a\big(\theta, e(\theta)\big) \\ U'(\theta) + c_\theta\big(\theta, e(\theta)\big) \end{pmatrix}. \tag{56}$$

The first component comes from obedience. The agent chooses his action after reporting, and the contract must make the recommended action optimal. The corresponding local requirement is

$$\mathbb{E}_{\theta,e(\theta)}[v_\theta \ell_a] = c_a\big(\theta, e(\theta)\big).$$

Action-sensitive information is valuable because it allows payments to respond more effectively to deviations from the recommended action.

The second component comes from truthful reporting. Type itself is not chosen; the relevant choice is the report. The envelope condition gives

$$\mathbb{E}_{\theta,e(\theta)}[v_\theta \ell_\theta] = U'(\theta) + c_\theta\big(\theta, e(\theta)\big).$$

Type-sensitive information is therefore valuable through the screening problem: it allows observed outcomes to support payoff differences required for truthful reporting.

Once these two requirements have been derived, however, the cheapest-contract problem treats them symmetrically as two moment restrictions:

$$\mathbb{E}_{\theta,e(\theta)}[v_\theta \ell] = \Delta_\sigma(\theta).$$

Both are priced by the same Fisher information matrix through

$$\Delta_\sigma(\theta)^\top I\big(\theta, e(\theta); p\big)^{-1} \Delta_\sigma(\theta).$$

Adverse selection and moral hazard therefore generate different incentive requirements, but both must be satisfied using the same observed outcomes. Once expressed as moment restrictions, their costs are determined by the same Fisher information matrix. This is why the optimal direction admits a symmetric eigenvector characterization even though the two incentive problems are economically different.

6.4 The rent channel in type incentives

The type requirement contains an additional endogenous component:

$$\Delta_{\sigma,\theta}(\theta) = \underbrace{c_\theta\big(\theta, e(\theta)\big)}_{\text{direct effect of type on cost}} + \underbrace{U'(\theta)}_{\text{rent slope}}. \tag{57}$$

The action requirement

$$\Delta_{\sigma,a}(\theta) = c_a\big(\theta, e(\theta)\big)$$

also depends endogenously on the allocation through $e(\theta)$. The difference is that the type requirement additionally depends on the slope of equilibrium utility. Screening therefore creates an additional channel through which the allocation determines the demand for information.

This does not imply that the type requirement is necessarily larger. The two terms in (57) may reinforce or offset each other. In particular,

$$U'(\theta) = -c_\theta\big(\theta, e(\theta)\big)$$

implies

$$\Delta_{\sigma,\theta}(\theta) = 0.$$

In that case there is no direct local type incentive requirement.

A zero type component of Δ_σ, however, does not imply that type-sensitive measurement has no value. When the existing performance measure is jointly informative about action and type, the information adjustment mixes the two requirements. Write

$$I^y = \begin{pmatrix} \iota_{aa} & \iota_{a\theta} \\ \iota_{a\theta} & \iota_{\theta\theta} \end{pmatrix}. \tag{58}$$

Then

$$(I^y)^{-1} = \frac{1}{\det I^y} \begin{pmatrix} \iota_{\theta\theta} & -\iota_{a\theta} \\ -\iota_{a\theta} & \iota_{aa} \end{pmatrix}, \tag{59}$$

so

$$w_\sigma = \frac{1}{\det I^y} \begin{pmatrix} \iota_{\theta\theta}\Delta_{\sigma,a} - \iota_{a\theta}\Delta_{\sigma,\theta} \\ \iota_{aa}\Delta_{\sigma,\theta} - \iota_{a\theta}\Delta_{\sigma,a} \end{pmatrix}. \tag{60}$$

If $\Delta_{\sigma,\theta} = 0$, for example,

$$w_\sigma = \frac{\Delta_{\sigma,a}}{\det I^y} \begin{pmatrix} \iota_{\theta\theta} \\ -\iota_{a\theta} \end{pmatrix}. \tag{61}$$

The optimal direction can therefore contain a type component even when there is no direct type incentive requirement. The reason is not an independent demand for type information. Rather, the existing performance measure is jointly informative about the two hidden dimensions, so additional measurement can use sensitivity to type to improve the information available for satisfying the action requirement.

More generally, the type component of the information-adjusted requirement is negative whenever

$$\iota_{a\theta}\Delta_{\sigma,a} > \iota_{aa}\Delta_{\sigma,\theta}. \tag{62}$$

Thus even when both components of Δ_σ are positive, the preferred measurement direction need not have positive components in both dimensions. The sign reflects the interaction between the incentive requirements and the joint information contained in the existing performance measure, rather than an intrinsic ranking of action and type information.

Optimal monitoring therefore depends neither on statistical information I^y nor on incentive requirements Δ_σ alone. It is their interaction,

$$w_\sigma = (I^y)^{-1}\Delta_\sigma,$$

aggregated across types through

$$M = \int_\Theta w_{\sigma_0}(\theta) w_{\sigma_0}(\theta)^\top \, dF(\theta),$$

that determines where scarce monitoring capacity is most valuable.

7 From local incentives to full implementation

The optimal direction in Theorem 2 is derived from the first-order relaxation. We now ask when the same direction remains optimal under full incentive compatibility.

Let $V^L(p)$ denote the value when local joint deviations are imposed, and let

$$V^F(p) = V(p)$$

denote the value under full incentive compatibility. The three problems are nested:

$$V^{\mathrm{FOC}}(p) \geq V^L(p) \geq V^F(p). \tag{63}$$

Since the value of scarce monitoring is of order r^2, the direction derived from the first-order problem remains valid whenever both gaps in (63) are $o(r^2)$ uniformly over monitoring directions.

For a type θ, define the payoff from reporting $\hat{\theta}$ and then choosing action $\hat{a}$ by

$$R(\theta, \hat{\theta}, \hat{a}; p) = \mathbb{E}_{\theta,\hat{a}}\left[v_{\hat{\theta}}(y, s)\right] - c(\theta, \hat{a}). \tag{64}$$

Truthful obedience corresponds to $(\hat{\theta}, \hat{a}) = (\theta, e(\theta))$, where

$$R(\theta, \theta, e(\theta); p) = U(\theta).$$

7.1 Local joint deviations

The first-order relaxation imposes local obedience and truthful reporting. These conditions are necessary for full incentive compatibility, but they do not by themselves ensure that the truthful report and recommended action are jointly optimal. In particular, an agent who changes his report may also adjust his action.

Let $\nabla R(\theta)$ and $H_R(\theta)$ denote the gradient and Hessian of $(\hat{\theta}, \hat{a}) \mapsto R(\theta, \hat{\theta}, \hat{a}; p)$ evaluated at the truthful-obedient pair. Thus

$$H_R(\theta) = \begin{pmatrix} R_{\hat{\theta}\hat{\theta}}(\theta) & R_{\hat{\theta}a}(\theta) \\ R_{a\hat{\theta}}(\theta) & R_{aa}(\theta) \end{pmatrix}.$$

Definition 1 (Local joint implementability)**.** A mechanism is locally jointly implementable if, for almost every θ,

$$\nabla R(\theta) = 0, \tag{L1}$$

and

$$H_R(\theta) \preceq 0. \tag{L2}$$

Let $V^L(p)$ denote the supremum of the principal's payoff over mechanisms satisfying (L1)–(L2) and IR.

Condition (L1) reproduces the local first-order requirements used in the relaxation. Condition (L2) adds the curvature restriction needed for the truthful report and recommended action to be a local joint optimum.

The interaction between report and action deviations can be seen by allowing the agent to re-optimize his action after a nearby report. Suppose $R_{aa}(\theta) < 0$, so that the action first-order condition defines a local best-response branch

$$\hat{a} = A(\theta, \hat{\theta})$$

through $A(\theta, \theta) = e(\theta)$. Define the corresponding loss from a report deviation by

$$G^L(\theta, \hat{\theta}) = U(\theta) - R\big(\theta, \hat{\theta}, A(\theta, \hat{\theta}); p\big). \tag{65}$$

Proposition 5 (Local joint-incentive curvature)**.** *Suppose that the action problem is strictly locally concave at the truthful-obedient pair, so that*

$$R_{aa}(\theta) < 0.$$

Then

$$\partial^2_{\hat{\theta}} G^L(\theta, \theta) = -R_{\hat{\theta}\hat{\theta}}(\theta) - \frac{R_{\hat{\theta}a}(\theta)^2}{|R_{aa}(\theta)|}. \tag{66}$$

The second term in (66) is the effect of joint deviations. If the agent changes his report and is then allowed to re-optimize his action, the local report-incentive margin falls by

$$\frac{R_{\hat{\theta}a}(\theta)^2}{|R_{aa}(\theta)|}.$$

The cross derivative $R_{\hat{\theta}a}$ therefore measures the local interaction between screening and obedience. When it is zero, action re-optimization does not change report curvature to second order.

The relevant question is whether imposing this additional curvature condition changes value at order r^2. Small monitoring designs change the optimal signal-contingent component of the contract at order r. At the uninformative benchmark, however, its first-order contribution to expected payoffs vanishes because

$$\sum_{j \in \mathcal{S}} g_j'(0) = 0.$$

Consequently, the local joint-incentive margins change only at order r^2. If they are strictly satisfied at the benchmark, they therefore remain satisfied under sufficiently small monitoring designs.

Assumption 4 (Strict local implementability)**.** There exists $\eta > 0$ such that, for almost every θ, the benchmark mechanism satisfies

$$R_{aa,0}(\theta) \leq -\eta, \qquad \partial^2_{\hat{\theta}} G^L_0(\theta, \theta) \geq \eta. \tag{67}$$

Along the relaxed near-optimal mechanisms associated with designs $p = r\hat{u}$, these quantities converge to their benchmark values as $r \downarrow 0$, uniformly over unit directions $\hat{u}$.

Theorem 3 (Local joint deviations do not change the leading value)**.** *Under Assumptions 1–3, 6, and 4, the relaxed near-optimal mechanisms can be chosen to be locally jointly implementable for all sufficiently small r, uniformly over unit directions $\hat{u}$. Moreover,*

$$V^L(r\hat{u}) = V^{\mathrm{FOC}}(r\hat{u}) + o(r^2) \tag{68}$$

uniformly in $\hat{u}$.

The theorem does not require joint deviations to be irrelevant. They affect the curvature condition through (66). What matters for the leading monitoring result is that imposing this additional condition does not reduce value at order r^2.

If a local margin binds at the benchmark, this conclusion need not follow. Measurement still perturbs the margin only at order r^2, but restoring local feasibility may itself have an order-r^2 cost that depends on the monitoring direction. In that case the direction selected by the first-order relaxation need not remain optimal.

7.2 From local to global incentive compatibility

Local joint implementability rules out profitable deviations in a neighborhood of the truthful-obedient pair. Full incentive compatibility also requires that deviations farther away remain unprofitable. We impose a uniform benchmark margin over these global deviations.

For a deviation $(\hat{\theta}, \hat{a})$, write

$$d(\theta, \hat{\theta}, \hat{a}) = \begin{pmatrix} \hat{\theta} - \theta \\ \hat{a} - e(\theta) \end{pmatrix}.$$

Assumption 5 (Robust global implementability)**.** There exist $\delta > 0$ and $\varepsilon_\delta > 0$ such that the benchmark mechanism satisfies

$$U_0(\theta) - R_0(\theta, \hat{\theta}, \hat{a}) \geq \varepsilon_\delta \tag{69}$$

whenever

$$\|d(\theta, \hat{\theta}, \hat{a})\| \geq \delta.$$

This slack is preserved along the relaxed near-optimal mechanisms as $r \downarrow 0$, uniformly over monitoring directions. In addition, the strict local conditions in Assumption 4 hold uniformly throughout the neighborhood

$$\|d(\theta, \hat{\theta}, \hat{a})\| \leq \delta.$$

The assumption separates the deviation space into two regions. Near the truthful-obedient pair, strict local curvature prevents profitable deviations. Away from it, the positive benchmark payoff gap prevents profitable global deviations. Stability under small monitoring designs preserves both margins.

Theorem 4 (Scarce monitoring preserves full implementation)**.** *Under the hypotheses of Theorem 3 and Assumption 5, the relaxed near-optimal mechanisms can be chosen to satisfy full incentive compatibility for all sufficiently small r, uniformly over unit directions $\hat{u}$. Moreover,*

$$V^F(r\hat{u}) = V^L(r\hat{u}) + o(r^2) \tag{70}$$

uniformly in $\hat{u}$.

Near the truthful-obedient pair, the vanishing gradient and strict negative curvature imply

$$R(\theta, \hat{\theta}, \hat{a}; p) \leq U(\theta)$$

for sufficiently small deviations. Outside that neighborhood, (69) and its stability preserve a strictly positive payoff gap. Together the two regions cover the full joint deviation space.

7.3 The optimal direction under full implementation

Theorems 3 and 4 imply that tightening the first-order relaxation first to local joint incentive compatibility and then to full incentive compatibility changes value only by $o(r^2)$. The quadratic term that determines the optimal monitoring direction is therefore unchanged.

Corollary 1 (Optimal direction under full implementation)**.** *Under the hypotheses of Theorem 4, uniformly over* $\|p\| \le r$*,*

$$V^F(p) = V^F(0) + I_g(0)p^\top Mp + o(r^2). \tag{71}$$

If $\nu_1(M) > 0$*, then every maximizing design* $p^{F,*}(r)$ *satisfies*

$$\frac{\|p^{F,*}(r)\|}{r} \longrightarrow 1. \tag{72}$$

If, in addition, $\nu_1(M)$ *is simple, with associated unit eigenvectors* $\{\hat{u}_1, -\hat{u}_1\}$*, then*

$$\operatorname{dist}\left(\frac{p^{F,*}(r)}{r}, \{\hat{u}_1, -\hat{u}_1\}\right) \longrightarrow 0. \tag{73}$$

Proof. Theorems 3 and 4 imply

$$V^F(p) = V^{\mathrm{FOC}}(p) + o(r^2)$$

uniformly over $\|p\| \le r$. Combining this with Theorem 2 gives (71). The remaining conclusions follow from Lemma 3. □

Thus the passage from the relaxed problem to the fully incentive-compatible problem can be summarized as

$$V^{\mathrm{FOC}} \longrightarrow V^L \longrightarrow V^F, \qquad V^{\mathrm{FOC}} - V^L = o(r^2), \qquad V^L - V^F = o(r^2).$$

Under strict local and robust global implementability, the additional incentive constraints do not alter the leading quadratic value of monitoring. The optimal direction under full implementation is therefore the same leading eigenvector of M identified by the first-order relaxation.

8 An insurance illustration

Consider an insurer contracting with a driver whose type $\theta \in [\underline{\theta}, \bar{\theta}]$ is uniformly distributed. The driver chooses care $a \in [\underline{a}, \bar{a}]$ at cost

$$c(\theta, a) = \frac{a^2}{2[1 + \gamma(\theta - \underline{\theta})]}. \tag{74}$$

Higher types therefore have a lower marginal cost of care. The insurer's gross benefit is

$$b(\theta, a) = [b_a + b_{a\theta}(\theta - \theta_0)]\, a. \tag{75}$$

The productive value of care may therefore vary with the driver's type but is distinct from the performance record used for contracting.

Table 1: Calibration of the insurance illustration

Primitive	Parameter	Value
Type support	$(\underline{\theta}, \bar{\theta})$	$(1, 2)$
Action support	$(\underline{a}, \bar{a})$	$(0.05, 2.75)$
Cost heterogeneity	γ	0.20
Marginal benefit	b_a	4.25
Benefit–type interaction	$b_{a\theta}$	0.40
Performance index 1	$(\lambda_{1a}, \lambda_{1\theta})$	$(1.00, 1.02)$
Performance index 2	$(\lambda_{2a}, \lambda_{2\theta})$	$(-0.55, 1.18)$
Monitoring center	(a_0, θ_0)	$(1.20, 1.50)$

Existing performance information. The insurer initially observes a categorical performance measure $y \in \{0, 1, 2\}$. Its distribution is multinomial logit:

$$q_j(\theta, a) = \frac{\exp\{\eta_j(\theta, a)\}}{1 + \exp\{\eta_1(\theta, a)\} + \exp\{\eta_2(\theta, a)\}}, \qquad j = 1, 2, \tag{76}$$

with $q_0(\theta, a) = 1 - q_1(\theta, a) - q_2(\theta, a)$ and

$$\eta_j(\theta, a) = \lambda_{ja} a + \lambda_{j\theta} \theta, \qquad j = 1, 2. \tag{77}$$

The two performance outcomes respond differently to action and type. Consequently, the existing record contains information about both hidden dimensions but does not measure them separately.

The insurer can add a binary monitoring signal $s \in \{0, 1\}$ satisfying

$$\Pr(s = 1 \mid \theta, a; p) = \frac{\exp\{p_a(a - a_0) + p_\theta(\theta - \theta_0)\}}{1 + \exp\{p_a(a - a_0) + p_\theta(\theta - \theta_0)\}}, \tag{78}$$

where

$$p = \begin{pmatrix} p_a \\ p_\theta \end{pmatrix}, \qquad \|p\| \le r.$$

At $p = 0$, the new signal is an independent fair coin and hence $I_g(0) = 1/4$.

Table 1 reports the calibration used below.

What should the insurer measure? We solve the relaxed contracting problem at the uninformative benchmark $p = 0$. Aggregating the information-adjusted incentive requirements gives

$$M = \int_1^2 w(\theta) w(\theta)^\top \, dF(\theta), \qquad w(\theta) = I^y(\theta, e(\theta))^{-1} \Delta(\theta). \tag{79}$$

The leading eigenvalue of M is

$$\nu_1(M) = 1.8255,$$

and the corresponding unit direction, oriented to have a positive action component, is

$$\hat{u}_1 = \begin{pmatrix} 0.834 \\ -0.552 \end{pmatrix}. \tag{80}$$

Writing a unit monitoring direction as

$$u(\phi) = \begin{pmatrix} \cos\phi \\ \sin\phi \end{pmatrix},$$

the optimal angle is

$$\phi^* = -33.5^\circ.$$

The optimal monitor therefore places most of its sensitivity on action and responds to type with the opposite sign.

A purely statistical criterion gives a very different answer. To aggregate statistical weakness across types, define

$$J = \int_1^2 I^y(\theta, e(\theta))^{-1}\, dF(\theta), \tag{81}$$

and let $\hat{u}^{\text{stat}}$ denote its leading eigenvector. Numerically,

$$\hat{u}^{\text{stat}} = \begin{pmatrix} 0.037 \\ 0.999 \end{pmatrix}, \qquad \phi^{\text{stat}} = 87.9^\circ. \tag{82}$$

Thus the statistical criterion recommends almost pure type information, whereas the contracting criterion recommends a signal that is primarily action-sensitive and negatively type-sensitive.

Figure 1 displays the two criteria geometrically. The statistical-information ellipse is generated by J, while the contracting-value ellipse is generated by M. Each ellipse is normalized to have unit major semi-axis, so its orientation and eccentricity represent the preferred direction and relative directional strength of the corresponding criterion rather than their absolute magnitudes. The major axes point in sharply different directions: statistical weakness is concentrated near the type dimension, while contracting value is maximized by a mixture of positive action sensitivity and negative type sensitivity.

The source of this difference can be seen in the information-adjusted incentive requirements. Figure 2 plots the two components of $w(\theta)$. The action component is positive throughout the type distribution, while the type component is negative and approaches zero for high types. Existing performance information therefore transforms the incentive requirements in a way that makes a signal positively sensitive to action and negatively sensitive to type particularly valuable.

Figure 3 shows the economic consequence of choosing different monitoring directions. For each $u(\phi)$, the leading-order contracting value is proportional to

$$u(\phi)^\top M u(\phi).$$

The figure reports this quantity relative to its maximum. The maximum occurs at $\phi^* = -33.5^\circ$, whereas the statistical direction lies close to pure type measurement and delivers substantially less contracting value. The distinction between statistical and economic informativeness therefore affects not only which direction is selected but also the value created by scarce monitoring.

The value of scarce monitoring. The local result predicts that, along the optimal direction,

$$V^{FOC}(r\hat{u}_1) - V^{FOC}(0) = I_g(0)\nu_1(M)r^2 + o(r^2) = 0.4564\, r^2 + o(r^2). \tag{83}$$

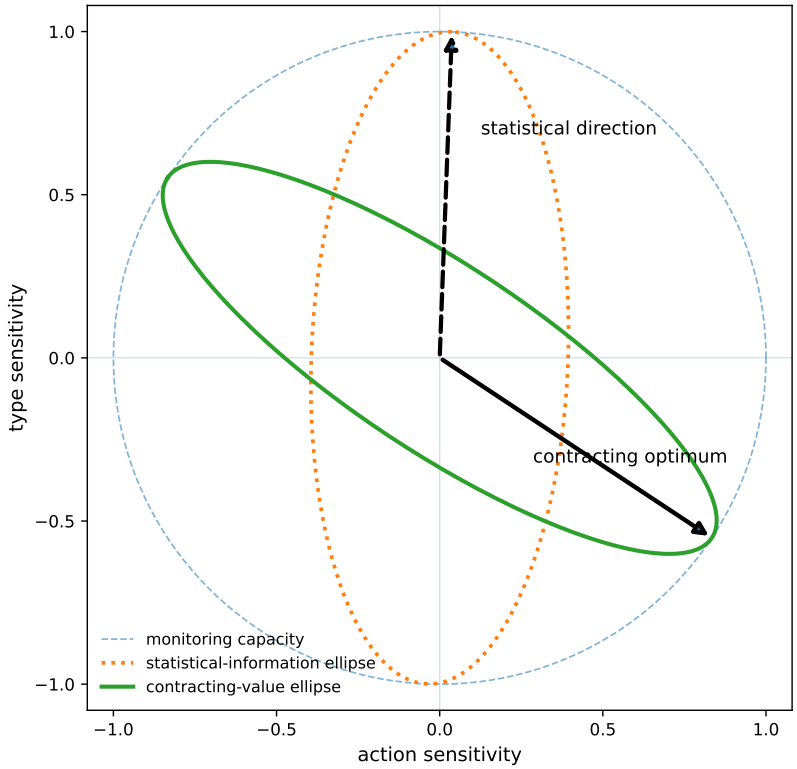


Figure 1: Statistical informativeness and contracting value. The dotted ellipse represents the statistical criterion J, and the solid ellipse represents the contracting-value matrix M. Each ellipse is normalized to have unit major semi-axis. The dashed circle represents the monitoring capacity set. The statistically preferred direction is almost entirely type-sensitive, whereas the contracting optimum combines positive action sensitivity with negative type sensitivity.

Figure 4 compares this approximation with the value obtained after reoptimizing the allocation for each monitoring capacity. The two curves are nearly indistinguishable over the displayed range. At $r = 0.05$, for example, the reoptimized gain is 0.0011404, compared with the quadratic prediction 0.0011409.

Implementation. The relaxed solution also satisfies the implementation conditions discussed in Section 7. At the uninformative benchmark, the largest eigenvalue of the local deviation Hessian is strictly negative, far deviations generate a strictly positive payoff loss, and the limited-liability constraint is slack.

Figure 5 follows these margins as monitoring capacity increases along the optimal direction. To make their movements comparable, each margin is normalized by its value at $r = 0$. The local curvature margin and the far-deviation incentive margin remain essentially unchanged over the displayed range. The limited-liability margin declines with monitoring capacity but remains strictly positive. Thus the value gains shown in Figure 4 are attained in a region in which the relaxed allocation remains locally implementable, the tested nonlocal deviations remain unprofitable, and limited liability remains slack.

The illustration makes the distinction between statistical informativeness and contracting

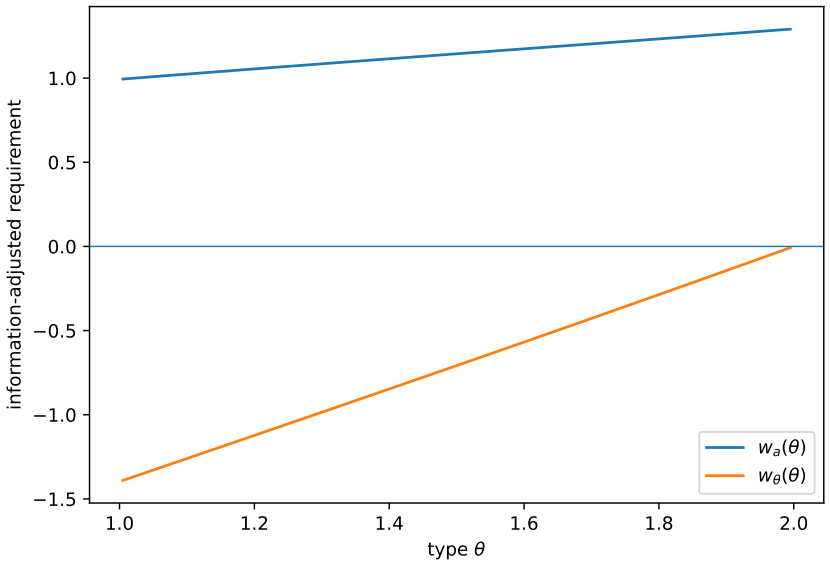


Figure 2: Information-adjusted incentive requirements. The figure plots the two components of $w(\theta) = I^y(\theta, e(\theta))^{-1}\Delta(\theta)$ at the uninformative benchmark.

value concrete. The existing performance measure is statistically weakest in a direction close to type, but additional information in that direction is not the most valuable for contracting. Once the existing information structure is combined with the incentive requirements, the optimal monitor instead emphasizes action and responds to type with the opposite sign. The resulting direction also accurately predicts the value of scarce monitoring, while the relevant implementation margins remain slack over the local region considered here.

9 Discussion

The analysis focuses on the design of a single, scarce monitoring technology around an uninformative benchmark. This section discusses several cases that fall outside this environment and how they may change the design problem.

Monitoring away from the uninformative benchmark. The main results characterize which direction should be measured first when monitoring capacity is scarce. This local restriction is important. At the uninformative benchmark, the value of monitoring admits a quadratic approximation, so the design problem is summarized by the matrix M. With a non-negligible amount of monitoring, both the allocation and the information-adjusted incentive requirements generally change with the monitoring technology. The optimal direction may therefore rotate as monitoring capacity increases.

This creates a distinction between the optimal initial direction and the optimal monitoring system at a larger scale. The leading eigenvector of M answers the former question. Characterizing the latter requires following how the allocation, the incentive requirements, and the information matrix change jointly with the monitoring technology. In particular, information

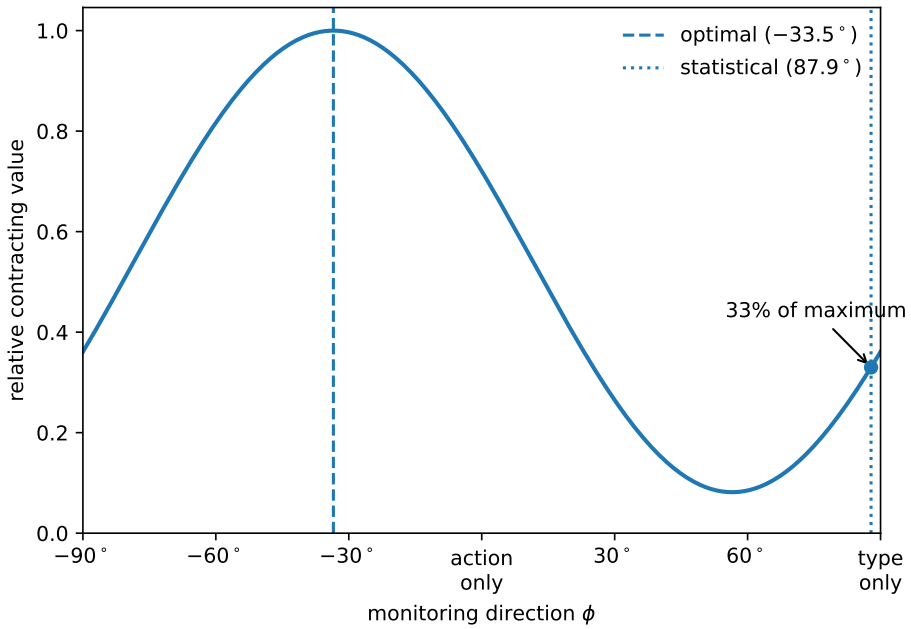


Figure 3: Contracting value by monitoring direction. For $u(\phi) = (\cos\phi, \sin\phi)^\top$, the curve reports $u(\phi)^\top M u(\phi)/\nu_1(M)$. The contracting optimum maximizes this value, while the statistical direction is determined by (81).

that is most valuable initially need not remain the most valuable after that direction has already been measured precisely.

Multidimensional monitoring. The paper considers a single new signal whose local information contribution is rank one. More generally, a principal may be able to introduce several signals or a multidimensional monitoring technology. If the principal can allocate a small information budget across several independent directions, the relevant design object becomes a positive semidefinite information increment rather than a single vector p.

The matrix M continues to describe the local value of different directions, but the solution then depends on the constraint placed on the information increment. Under a constraint that permits concentration of all capacity in one direction, the leading eigenvector remains optimal. Under constraints that require diversification or limit the amount of information that can be placed in any one direction, capacity may instead be allocated across several eigendirections. The spectrum of M, rather than only its leading eigenvector, then becomes economically relevant.

Dependence between existing and new information. The baseline model assumes that the existing performance measure y and the designed signal s are conditionally independent given (θ, a). This makes their Fisher information contributions additive and isolates the choice of what the new signal should measure. In many applications, however, new monitoring technologies are constructed from data that overlap with existing performance measures. A new driving score, for example, may use some of the same underlying events that enter an insurer's

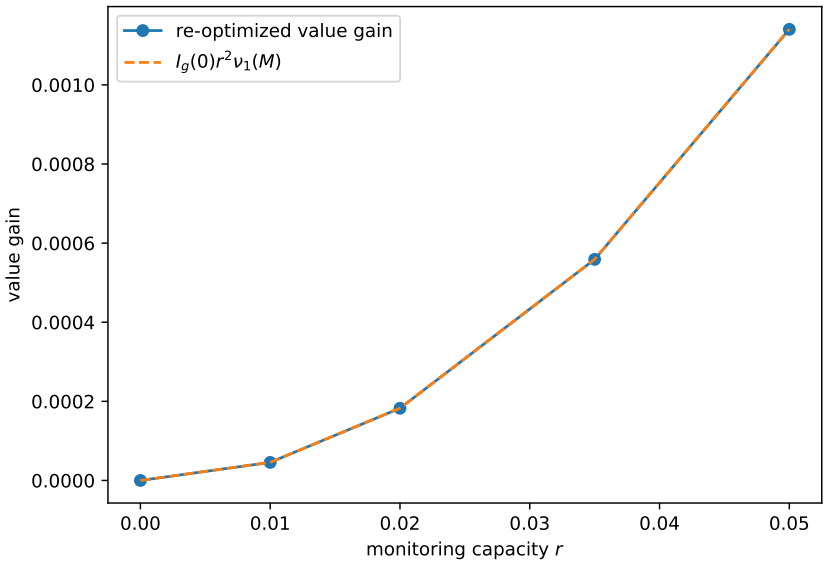


Figure 4: The value of scarce monitoring. The solid curve is the value gain after reoptimizing the relaxed contracting problem along the optimal direction $\hat{u}_1$. The dashed curve is the second-order approximation $I_g(0)r^2\nu_1(M)$.

existing record.

With conditional dependence, the marginal value of the new signal depends not only on its own score but also on its covariance with the score of the existing performance measure. A signal that is highly informative in isolation may add little information if its innovations are largely redundant with y. The relevant design problem would therefore concern the incremental score after accounting for existing information. The basic economic force remains the interaction between information and incentive requirements, but the rank-one update used in the paper would have to be replaced by a characterization of incremental information under correlated signals.

Binding implementation constraints. The implementation results focus on benchmarks at which the relevant local and global incentive conditions have sufficient slack. This allows the first-order relaxation to determine the value of scarce monitoring locally. If an implementation constraint binds at the benchmark, the monitoring design problem can change at the same order as the information gain.

In that case, a monitoring direction must be evaluated not only by how much it reduces the cost of satisfying the first-order incentive requirements, but also by how it affects the binding implementation constraint. The optimal direction may therefore differ from the leading eigenvector of M. A natural extension would attach shadow values to the binding local or global incentive constraints and incorporate their responses to monitoring into the local design criterion. Such an extension would characterize how the optimal direction changes when implementation, rather than information alone, is locally scarce.

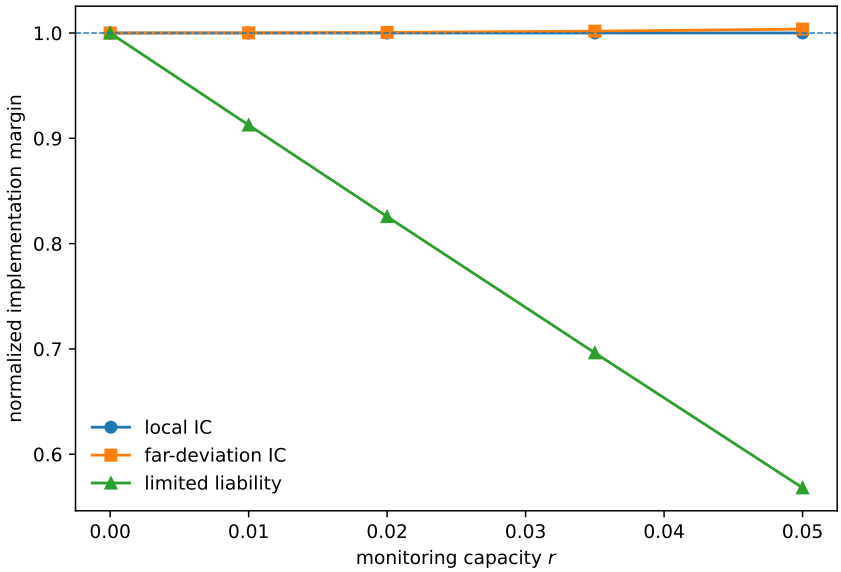


Figure 5: Implementation under scarce monitoring. The figure reports the local IC, far-deviation IC, and limited-liability margins along the optimal monitoring direction, normalized by their values at $r = 0$. Local IC is measured by $-H_{\max}$, far-deviation IC by the minimum payoff loss over the nonlocal deviations considered in the numerical check, and limited liability by the minimum transfer.

References

Baker, G. (2002). Distortion and risk in optimal incentive contracts. *Journal of Human Resources*, 37(4):728–751.

Ball, I. (2025). Scoring strategic agents. *American Economic Journal: Microeconomics*, 17(1):97–129.

Banker, R. D. and Datar, S. M. (1989). Sensitivity, precision, and linear aggregation of signals for performance evaluation. *Journal of Accounting Research*, 27(1):21–39.

Castro-Pires, H., Chade, H., and Swinkels, J. (2024). Disentangling moral hazard and adverse selection. *American Economic Review*, 114(1):1–37.

Chade, H. and Swinkels, J. (2020). The moral hazard problem with high stakes. *Journal of Economic Theory*, 187:105032.

Chaigneau, P., Edmans, A., and Gottlieb, D. (2019). The informativeness principle without the first-order approach. *Games and Economic Behavior*, 113:743–755.

Doval, L. and Skreta, V. (2022). Mechanism design with limited commitment. *Econometrica*, 90(4):1463–1500.

Dye, R. A. (1986). Optimal monitoring policies in agencies. *RAND Journal of Economics*, 17(3):339–350.

Feltham, G. A. and Xie, J. (1994). Performance measure congruity and diversity in multi-task principal/agent relations. *The Accounting Review*, 69(3):429–453.

Frankel, A. and Kartik, N. (2019). Muddled information. *Journal of Political Economy*, 127(4):1739–1776.

Frankel, A. and Kartik, N. (2022). Improving information from manipulable data. *Journal of the European Economic Association*, 20(1):79–115.

Georgiadis, G. and Szentes, B. (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.

Gottlieb, D. and Moreira, H. (2022). Simple contracts with adverse selection and moral hazard. *Theoretical Economics*, 17(3):1357–1401.

Guesnerie, R. and Laffont, J.-J. (1984). A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm. *Journal of Public Economics*, 25(3):329–369.

Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10(1):74–91.

Holmström, B. and Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303–328.

Holmström, B. and Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization*, 7:24–52. Special Issue.

Hörner, J. and Lambert, N. S. (2021). Motivational ratings. *Review of Economic Studies*, 88(4):1892–1935.

Jewitt, I. (1988). Justifying the first-order approach to principal-agent problems. *Econometrica*, 56(5):1177–1190.

Kim, S. K. (1995). Efficiency of an information system in an agency model. *Econometrica*, 63(1):89–102.

Laffont, J.-J. and Tirole, J. (1986). Using cost observation to regulate firms. *Journal of Political Economy*, 94(3):614–641.

Laffont, J.-J. and Tirole, J. (1993). *A Theory of Incentives in Procurement and Regulation*. MIT Press, Cambridge, MA.

Lambert, R. A. (1985). Variance investigation in agency settings. *Journal of Accounting Research*, 23(2):633–647.

McAfee, R. P. and McMillan, J. (1987). Competition for agency contracts. *RAND Journal of Economics*, 18(2):296–307.

Myerson, R. B. (1982). Optimal coordination mechanisms in generalized principal-agent problems. *Journal of Mathematical Economics*, 10(1):67–81.

Picard, P. (1987). On the design of incentive schemes under moral hazard and adverse selection. *Journal of Public Economics*, 33(3):305–331.

Shavell, S. (1979). Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics*, 10(1):55–73.

Sun, R. and Zhang, Y. (2026). Coarse screening. Working paper, arXiv:2604.04405.

A Proofs for Sections 4 and 5

Proof of Proposition 1

Proof. We first establish the obedience condition. Fix a type θ and set $\hat{\theta} = \theta$ in (IC). Incentive compatibility implies that the recommended action $e(\theta)$ maximizes

$$a \mapsto \mathbb{E}_{\theta,a}[v_\theta] - c(\theta, a)$$

over A. Since $e(\theta) \in \operatorname{int} A$, the first-order condition is

$$\partial_a \mathbb{E}_{\theta,a}[v_\theta]\big|_{a=e(\theta)} = c_a\big(\theta, e(\theta)\big).$$

Because $v_\theta(y, s)$ does not directly depend on (θ, a), differentiating the finite sum defining the expectation gives

$$\partial_a \mathbb{E}_{\theta,a}[v_\theta] = \sum_{y,s} v_\theta(y, s) \partial_a \phi(y, s \mid \theta, a; p).$$

Using

$$\partial_a \phi = \phi \, \partial_a \log \phi = \phi \, \ell_a,$$

we obtain

$$\partial_a \mathbb{E}_{\theta,a}[v_\theta] = \mathbb{E}_{\theta,a}[v_\theta \ell_a].$$

Evaluating at $a = e(\theta)$ gives

$$\mathbb{E}_{\theta,e(\theta)}[v_\theta \ell_a] = c_a\big(\theta, e(\theta)\big),$$

which is (OB).

We next establish the envelope condition. By (IC),

$$U_p(\theta; M) = \max_{(\hat{\theta},\hat{a}) \in \Theta \times A} \Pi_p(\theta, \hat{\theta}, \hat{a}; M).$$

The maximization is over both the report and the action, so joint deviations are included in the envelope problem.

For fixed $(\hat{\theta}, \hat{a})$,

$$\Pi_p(\theta, \hat{\theta}, \hat{a}; M) = \mathbb{E}_{\theta,\hat{a}}[v_{\hat{\theta}}] - c(\theta, \hat{a}).$$

Differentiating with respect to the true type gives

$$\partial_\theta \Pi_p(\theta, \hat{\theta}, \hat{a}; M) = \mathbb{E}_{\theta,\hat{a}}[v_{\hat{\theta}} \ell_\theta] - c_\theta(\theta, \hat{a}).$$

Under the maintained regularity conditions, these derivatives are uniformly bounded over the compact domain. Hence $U_p(\theta; M)$, as the pointwise maximum of a family of functions with uniformly bounded derivatives, is Lipschitz and therefore differentiable almost everywhere.

At every point of differentiability, the envelope theorem gives

$$U_p'(\theta; M) = \partial_\theta \Pi_p(\theta, \hat{\theta}, \hat{a}; M)$$

evaluated at any optimal pair. By incentive compatibility, $(\hat{\theta}, \hat{a}) = (\theta, e(\theta))$ is such a pair. Therefore

$$U_p'(\theta; M) = \mathbb{E}_{\theta,e(\theta)}[v_\theta \ell_\theta] - c_\theta\big(\theta, e(\theta)\big),$$

which is (EN). □

Proof of Theorem 1

Proof. Fix $\sigma = (e, U)$ and θ, and suppress the dependence on $(\theta, e(\theta); p)$ in the notation. Let

$$\bar{v} = \bar{v}_\sigma(\theta), \qquad \Delta = \Delta_\sigma(\theta), \qquad I = I\big(\theta, e(\theta); p\big).$$

Equip the finite-dimensional space of contracts with the inner product

$$\langle g, h \rangle = \mathbb{E}_{\theta, e(\theta)}[gh].$$

The score has mean zero. Indeed, for either component $k \in \{a, \theta\}$,

$$\mathbb{E}[\ell_k] = \sum_{y,s} \phi(y, s \mid \theta, e(\theta); p) \partial_k \log \phi(y, s \mid \theta, e(\theta); p) = \partial_k \sum_{y,s} \phi(y, s \mid \theta, e(\theta); p) = 0.$$

Hence the constant function is orthogonal to both components of ℓ.

Write any feasible contract as

$$v = \bar{v} + w, \qquad \mathbb{E}[w] = 0.$$

Then

$$\mathbb{E}[v^2] = \bar{v}^2 + \mathbb{E}[w^2],$$

and, since $\mathbb{E}[\ell] = 0$, the incentive requirement in (18) becomes

$$\mathbb{E}[w\ell] = \Delta.$$

Let

$$L = \operatorname{span}\{\ell_a, \ell_\theta\}.$$

Decompose

$$w = \Pi_L w + w^\perp,$$

where $w^\perp$ is orthogonal to L. Since each component of ℓ belongs to L,

$$\mathbb{E}[w\ell] = \mathbb{E}[(\Pi_L w)\ell].$$

Thus $w^\perp$ does not contribute to the incentive requirement. By orthogonality,

$$\mathbb{E}[w^2] = \mathbb{E}[(\Pi_L w)^2] + \mathbb{E}[(w^\perp)^2],$$

so any cost-minimizing contract must satisfy $w^\perp = 0$.

It follows that the minimizing variable component can be written as

$$w = \lambda^\top \ell$$

for some $\lambda \in \mathbb{R}^2$. The incentive requirement then becomes

$$\Delta = \mathbb{E}[w\ell] = \mathbb{E}[\ell\ell^\top]\lambda = I\lambda.$$

Since I is nonsingular,

$$\lambda = I^{-1}\Delta.$$

Therefore

$$v_\sigma^* = \bar{v}_\sigma(\theta) + \Delta_\sigma(\theta)^\top I^{-1}\ell,$$

which gives (20).

Finally,

$$\begin{aligned}\mathbb{E}[(v_\sigma^*)^2] &= \bar{v}_\sigma(\theta)^2 + \mathbb{E}\left[\left(\Delta_\sigma(\theta)^\top I^{-1}\ell\right)^2\right] \\ &= \bar{v}_\sigma(\theta)^2 + \Delta_\sigma(\theta)^\top I^{-1}\mathbb{E}[\ell\ell^\top]I^{-1}\Delta_\sigma(\theta) \\ &= \bar{v}_\sigma(\theta)^2 + \Delta_\sigma(\theta)^\top I^{-1}\Delta_\sigma(\theta),\end{aligned}$$

which is (21). Uniqueness follows from the nonsingularity of I and the fact that any nonzero orthogonal component $w^\perp$ strictly increases the objective. □

Proof of Proposition 3

Proof. By conditional independence in (4),

$$\log \phi(y, s \mid \theta, a; p) = \log q(y \mid \theta, a) + \log g_s\big(z(\theta, a; p)\big).$$

Hence the joint score decomposes as

$$\ell = \ell^y + \ell^s,$$

where

$$\ell^y = \nabla_{(a,\theta)} \log q(y \mid \theta, a)$$

and, using (2),

$$\ell^s = \frac{g_s'(z)}{g_s(z)} p.$$

Both score components have conditional mean zero. For the designed signal,

$$\begin{aligned}\mathbb{E}_{\theta,a}[\ell^s] &= \sum_{j\in\mathcal{S}} g_j(z)\frac{g_j'(z)}{g_j(z)}p \\ &= \left(\sum_{j\in\mathcal{S}} g_j'(z)\right)p = 0,\end{aligned}$$

because

$$\sum_{j\in\mathcal{S}} g_j(z) = 1.$$

The same score identity gives

$$\mathbb{E}_{\theta,a}[\ell^y] = 0.$$

Since y and s are conditionally independent given (θ, a), the cross terms vanish:

$$\mathbb{E}_{\theta,a}[\ell^y(\ell^s)^\top] = \mathbb{E}_{\theta,a}[\ell^y]\mathbb{E}_{\theta,a}[\ell^s]^\top = 0,$$

and similarly for its transpose. Therefore

$$I(\theta, a; p) = I^y(\theta, a) + I^s(\theta, a; p),$$

where

$$\begin{aligned} I^s(\theta, a; p) &= \mathbb{E}_{\theta,a}[\ell^s(\ell^s)^\top] \\ &= \sum_{j\in\mathcal{S}} g_j(z) \left(\frac{g_j'(z)}{g_j(z)}\right)^2 pp^\top \\ &= \left(\sum_{j\in\mathcal{S}} \frac{g_j'(z)^2}{g_j(z)}\right) pp^\top \\ &= I_g(z) pp^\top. \end{aligned}$$

This proves (30). Since $pp^\top$ is an outer product, I^s has rank at most one. Finally,

$$\operatorname{tr} I^s(\theta, a; p) = I_g(z) \operatorname{tr}(pp^\top) = I_g(z)\|p\|^2,$$

which gives (31). □

B Proof of Theorem 2 and Lemma 2

Proof of Lemma 2

Proof. Fix an allocation $\sigma = (e, U)$ and a design

$$p = \rho\hat{u}, \qquad 0 \le \rho \le r, \qquad \|\hat{u}\| = 1.$$

At the benchmark $p = 0$,

$$V_\sigma^{\mathrm{FOC}}(0) = \int_\Theta \left\{ b\big(\theta, e(\theta)\big) - \bar{v}_\sigma(\theta)^2 - \Delta_\sigma(\theta)^\top I^y\big(\theta, e(\theta)\big)^{-1} \Delta_\sigma(\theta) \right\} dF(\theta).$$

Since

$$b\big(\theta, e(\theta)\big) \le \bar{b}$$

and $\bar{v}_\sigma(\theta)^2 \ge 0$, it follows that

$$\int_\Theta \Delta_\sigma(\theta)^\top I^y\big(\theta, e(\theta)\big)^{-1} \Delta_\sigma(\theta)\, dF(\theta) \le \bar{b} - V_\sigma^{\mathrm{FOC}}(0). \tag{84}$$

For each θ, Proposition 3 gives

$$I^s\big(\theta, e(\theta); p\big) = I_g(z)\rho^2 \hat{u}\hat{u}^\top.$$

Let

$$\kappa_\sigma(\theta; p) = I_g(z)\rho^2.$$

By Lemma 1, the reduction in incentive cost for type θ is

$$\frac{\kappa_\sigma(\theta; p)\big(\hat{u}^\top w_\sigma(\theta)\big)^2}{1 + \kappa_\sigma(\theta; p)\hat{u}^\top I^y\big(\theta, e(\theta)\big)^{-1}\hat{u}},$$

where

$$w_\sigma(\theta) = I^y\big(\theta, e(\theta)\big)^{-1}\Delta_\sigma(\theta).$$

The denominator is at least one, so this gain is bounded above by

$$\kappa_\sigma(\theta; p)\big(\hat{u}^\top w_\sigma(\theta)\big)^2.$$

By Assumption 1,

$$I^y(\theta, a) \succeq \varepsilon I_2,$$

and hence

$$I^y(\theta, a)^{-1} \preceq \varepsilon^{-1} I_2.$$

Cauchy–Schwarz in the inner product generated by $I^y(\theta, e(\theta))^{-1}$ gives

$$\begin{aligned}\big(\hat{u}^\top w_\sigma(\theta)\big)^2 &= \Big(\hat{u}^\top I^y\big(\theta, e(\theta)\big)^{-1}\Delta_\sigma(\theta)\Big)^2 \\ &\le \Big(\hat{u}^\top I^y\big(\theta, e(\theta)\big)^{-1}\hat{u}\Big)\Big(\Delta_\sigma(\theta)^\top I^y\big(\theta, e(\theta)\big)^{-1}\Delta_\sigma(\theta)\Big) \\ &\le \frac{1}{\varepsilon}\Delta_\sigma(\theta)^\top I^y\big(\theta, e(\theta)\big)^{-1}\Delta_\sigma(\theta).\end{aligned}$$

For all sufficiently small r, the index z remains in a compact neighborhood of zero uniformly over σ, θ, and $\|p\| \le r$. Continuity of I_g therefore implies the existence of a finite constant $\overline{I}_g$ such that

$$I_g(z) \le \overline{I}_g.$$

Since $\rho \le r$,

$$\kappa_\sigma(\theta; p) \le \overline{I}_g r^2.$$

Integrating the preceding bounds and using (84) yields

$$0 \le V_\sigma^{\mathrm{FOC}}(p) - V_\sigma^{\mathrm{FOC}}(0) \le \frac{\overline{I}_g}{\varepsilon} r^2 \left\{\bar{b} - V_\sigma^{\mathrm{FOC}}(0)\right\}.$$

Taking the supremum over $\|p\| \le r$ proves (39), with

$$C = \frac{\overline{I}_g}{\varepsilon}.$$

□

Proof of Theorem 2

Proof. For an allocation $\sigma = (e, U)$, define the gain from monitoring by

$$G_\sigma(p) = V_\sigma^{\mathrm{FOC}}(p) - V_\sigma^{\mathrm{FOC}}(0). \tag{85}$$

By Lemma 1,

$$G_\sigma(p) \ge 0.$$

Step 1: Expansion for a fixed allocation. Write

$$p = \rho\hat{u}, \qquad 0 \le \rho \le r, \qquad \|\hat{u}\| = 1.$$

For type θ, Proposition 3 gives

$$I^s\big(\theta, e(\theta); p\big) = \kappa_\sigma(\theta; p)\hat{u}\hat{u}^\top,$$

where

$$\kappa_\sigma(\theta; p) = I_g\big(z(\theta, e(\theta); p)\big)\rho^2.$$

Since $\Theta \times A$ is compact, there exists $\bar{x} < \infty$ such that

$$\|x(\theta, a)\| \le \bar{x}$$

for all (θ, a). Hence

$$|z(\theta, e(\theta); p)| \le \rho\bar{x} \le r\bar{x}.$$

Because I_g is continuously differentiable near zero,

$$I_g\big(z(\theta, e(\theta); p)\big) = I_g(0) + O(r)$$

uniformly over θ, σ, and $\|p\| \le r$. Therefore

$$\kappa_\sigma(\theta; p) = I_g(0)\rho^2 + O(r\rho^2) \tag{86}$$

uniformly.

Applying Lemma 1 with

$$B = I^y\big(\theta, e(\theta)\big)$$

gives

$$G_\sigma(p) = \int_\Theta \frac{\kappa_\sigma(\theta; p)\big(\hat{u}^\top w_\sigma(\theta)\big)^2}{1 + \kappa_\sigma(\theta; p)\hat{u}^\top I^y\big(\theta, e(\theta)\big)^{-1}\hat{u}}\, dF(\theta). \tag{87}$$

Assumption 1 implies

$$0 < \hat{u}^\top I^y\big(\theta, e(\theta)\big)^{-1}\hat{u} \le \frac{1}{\varepsilon}.$$

Thus the denominator in (87) equals

$$1 + O(r^2)$$

uniformly over θ, σ, and $\|p\| \le r$.

It follows from (86) that

$$G_\sigma(p) = I_g(0)p^\top M_\sigma p + O(r^3)\|M_\sigma\|, \tag{88}$$

uniformly over $\|p\| \le r$. In particular, the remainder is $O(r^3)$ uniformly on every collection of allocations for which $\|M_\sigma\|$ is uniformly bounded.

We next obtain such a bound for benchmark near-optimizers. Since M_σ is positive semidefinite,

$$\begin{aligned}
\|M_\sigma\| &\le \operatorname{tr} M_\sigma \\
&= \int_\Theta \|w_\sigma(\theta)\|^2\, dF(\theta) \\
&= \int_\Theta \Delta_\sigma(\theta)^\top I^y\big(\theta, e(\theta)\big)^{-2}\Delta_\sigma(\theta)\, dF(\theta).
\end{aligned}$$

Assumption 1 implies

$$I^y(\theta,a)^{-2} \preceq \frac{1}{\varepsilon} I^y(\theta,a)^{-1},$$

so by (84),

$$\|M_\sigma\| \leq \frac{1}{\varepsilon}\left\{\bar{b} - V_\sigma^{\mathrm{FOC}}(0)\right\}. \tag{89}$$

Hence (88) is uniform over every fixed set of benchmark near-optimal allocations.

Step 2: Allocations selected by small monitoring designs are benchmark near-optimizers. For each $\|p\| \leq r$, choose an allocation σ_p satisfying

$$V_{\sigma_p}^{\mathrm{FOC}}(p) \geq V^{\mathrm{FOC}}(p) - r^3\eta(r), \tag{90}$$

where $\eta(r) \downarrow 0$.

Let σ_0 be a benchmark-optimal allocation and define

$$D_p = V^{\mathrm{FOC}}(0) - V_{\sigma_p}^{\mathrm{FOC}}(0) \geq 0$$

and

$$B = \bar{b} - V^{\mathrm{FOC}}(0).$$

Because monitoring weakly reduces the incentive cost of any fixed allocation,

$$G_{\sigma_0}(p) \geq 0.$$

Therefore

$$\begin{aligned} V^{\mathrm{FOC}}(0) &= V_{\sigma_0}^{\mathrm{FOC}}(0) \\ &\leq V_{\sigma_0}^{\mathrm{FOC}}(p) \\ &\leq V^{\mathrm{FOC}}(p) \\ &\leq V_{\sigma_p}^{\mathrm{FOC}}(p) + r^3\eta(r) \\ &= V_{\sigma_p}^{\mathrm{FOC}}(0) + G_{\sigma_p}(p) + r^3\eta(r). \end{aligned}$$

Hence

$$D_p \leq G_{\sigma_p}(p) + r^3\eta(r). \tag{91}$$

Lemma 2 implies

$$G_{\sigma_p}(p) \leq Cr^2\left\{\bar{b} - V_{\sigma_p}^{\mathrm{FOC}}(0)\right\} = Cr^2(B + D_p).$$

Combining this with (91),

$$D_p \leq Cr^2(B + D_p) + r^3\eta(r).$$

For sufficiently small r, $Cr^2 \leq 1/2$, so

$$0 \leq D_p \leq 2Cr^2B + 2r^3\eta(r) = O(r^2) \tag{92}$$

uniformly over $\|p\| \leq r$.

Thus every allocation selected by a sufficiently small monitoring design is a benchmark near-optimizer.

Step 3: The associated matrices converge uniformly to M. We claim that

$$\sup_{\|p\|\le r} \|M_{\sigma_p} - M\| \longrightarrow 0. \tag{93}$$

Suppose not. Then there exist $\delta > 0$, a sequence $r_n \downarrow 0$, and designs p_n with $\|p_n\| \le r_n$ such that

$$\|M_{\sigma_{p_n}} - M\| \ge \delta$$

for every n. But (92) implies

$$V^{\mathrm{FOC}}_{\sigma_{p_n}}(0) \longrightarrow V^{\mathrm{FOC}}(0).$$

The sequence (σ_{p_n}) is therefore a benchmark maximizing sequence. Assumption 3 then gives

$$M_{\sigma_{p_n}} \longrightarrow M,$$

a contradiction. This proves (93).

Step 4: Uniform expansion of the optimized value. For the lower bound, evaluate the problem at the benchmark-optimal allocation σ_0. By (88),

$$\begin{aligned} V^{\mathrm{FOC}}(p) &\ge V^{\mathrm{FOC}}_{\sigma_0}(p) \\ &= V^{\mathrm{FOC}}(0) + I_g(0) p^\top M p + o(r^2), \end{aligned}$$

uniformly over $\|p\| \le r$.

For the upper bound, use

$$V^{\mathrm{FOC}}_{\sigma_p}(0) \le V^{\mathrm{FOC}}(0)$$

together with (90):

$$\begin{aligned} V^{\mathrm{FOC}}(p) &\le V^{\mathrm{FOC}}_{\sigma_p}(p) + r^3 \eta(r) \\ &= V^{\mathrm{FOC}}_{\sigma_p}(0) + G_{\sigma_p}(p) + r^3 \eta(r) \\ &\le V^{\mathrm{FOC}}(0) + I_g(0) p^\top M_{\sigma_p} p + o(r^2). \end{aligned}$$

By (93),

$$p^\top M_{\sigma_p} p = p^\top M p + o(r^2)$$

uniformly over $\|p\| \le r$. Hence

$$V^{\mathrm{FOC}}(p) \le V^{\mathrm{FOC}}(0) + I_g(0) p^\top M p + o(r^2)$$

uniformly over the ball.

Combining the lower and upper bounds gives

$$V^{\mathrm{FOC}}(p) = V^{\mathrm{FOC}}(0) + I_g(0) p^\top M p + o(r^2)$$

uniformly over $\|p\| \le r$, which is (40).

Finally, apply Lemma 3 with

$$c = I_g(0) > 0.$$

If $\nu_1(M) > 0$, it gives

$$\sup_{\|p\| \le r} V^{\mathrm{FOC}}(p) = V^{\mathrm{FOC}}(0) + I_g(0) r^2 \nu_1(M) + o(r^2),$$

as well as

$$\frac{\|p^{\mathrm{FOC},*}(r)\|}{r} \longrightarrow 1.$$

If the leading eigenvalue is simple, Lemma 3 further gives

$$\operatorname{dist}\left(\frac{p^{\mathrm{FOC},*}(r)}{r}, \{\hat{u}_1, -\hat{u}_1\} \right) \longrightarrow 0.$$

This proves the remaining conclusions of Theorem 2. □

C Sufficiency of the information-symmetry condition

Condition (**??**) is sufficient but not necessary in any given environment. Expanding $\mathcal{I}_g(r\,\hat{u}^\top x) = \mathcal{I}_g(0) + r\,\mathcal{I}_g'(0)\,\hat{u}^\top x + O(r^2)$, the r^3 contribution to the fixed-allocation gain is proportional to

$$\mathcal{I}_g'(0) \int_\Theta \big(\hat{u}^\top x\big)\big(\hat{u}^\top w_\sigma\big)^2 \, dF,$$

which can vanish for a particular allocation, type distribution or direction even when $\mathcal{I}_g'(0) \neq 0$. What (**??**) buys is that it vanishes *uniformly*, which is what the proof of Theorem 2 uses. The condition concerns the local Fisher-information function $\mathcal{I}_g$, not the symmetry of the experiment as a whole, which is why we call it information symmetry rather than symmetry.

D Uniformity of the allocation limit

This appendix supplies Step 2′ in the proof of Theorem 2.

Step 2′: the convergence is uniform in the direction. Pointwise convergence for each $\hat{u}$, together with compactness of the circle, does not by itself give uniformity, which Step 4 requires. We claim

$$\sup_{\|\hat{u}\|=1} \big\| M_{\sigma_r(\hat{u})} - M \big\| \longrightarrow 0. \tag{94}$$

Suppose not. Then there are $\varepsilon > 0$, $r_n \downarrow 0$ and unit vectors $\hat{u}_n$ such that $\sigma_n \equiv \sigma_{r_n}(\hat{u}_n)$ satisfies $\|M_{\sigma_n} - M\| \ge \varepsilon$ for all n. Because (**??**) is uniform in $\hat{u}$, $\bar{V}(0) - \bar{V}_{\sigma_n}(0) = O(r_n^2) \to 0$, so (σ_n) is a maximizing sequence for $\bar{V}.(0)$. The last clause of Assumption 3 then gives $M_{\sigma_n} \to M$, a contradiction.

E Smoothness used in Section 7

The results below differentiate objects that Sections 3.1–5.3 never differentiate, and the assumptions of Section 3.1 are not strong enough for them. The relevant orders are as follows. The deviation payoff $R(\theta, \hat{\theta}, a) = \mathbb{E}_{\theta,a}[v_{\hat{\theta}}] - c(\theta, a)$ must be C^2 jointly in $(\hat{\theta}, a)$ for Proposition 5: C^2 gives $R_a \in C^1$, hence a C^1 maximizer $\hat{a}^*$ by the implicit function theorem, hence a C^2 indirect

payoff $\widetilde{R}$. That requires q and each g_j to be C^2 in (θ, a) — already assumed — *and* $\hat{\theta} \mapsto v_{\hat{\theta}}(y, s)$ to be C^2, which is a restriction on the mechanism, not on primitives. the sufficient condition in Appendix G additionally requires $R_{\hat{\theta}\hat{\theta}}$, $R_{\hat{\theta}a}$ and R_{aa} to be well defined and continuous along the family, which needs ℓ to be C^1 in (θ, a) and hence $\phi \in C^3$, that is q and the g_j three times continuously differentiable. We impose the following for this section only.

Assumption 6 (Additional smoothness, Section 7 only)**.** $q(y \mid \cdot, \cdot)$ and each g_j are three times continuously differentiable, and the mechanisms considered have $\hat{\theta} \mapsto v_{\hat{\theta}}(y, s)$ twice continuously differentiable for each (y, s).

Remark 2. Assumption 6 is used nowhere in Sections 4–5.3, whose results require only the C^2 primitives of Assumption 1. Keeping the two separate matters, because the relaxed value theorem is the part of the paper that is unconditional and it should not inherit smoothness it does not need.

F Displaced material from Section 7

The ordering $V^{\mathrm{FOC}} \geq V^{L} \geq V^{F}$

Lemma 4 ((L1) is exactly the first-order relaxation)**.** *Let $e(\cdot)$, $U(\cdot)$ and $\hat{\theta} \mapsto v_{\hat{\theta}}$ be differentiable. Then $\nabla R(\theta) = 0$ for a.e. θ if and only if* (OB) *and* (EN) *hold.*

Proof. The $\hat{a}$-component of ∇R is $\partial_{\hat{a}} \mathbb{E}_{\theta,\hat{a}}[v_\theta] - c_a$ at $\hat{a} = e(\theta)$, which by Lemma **??** is $\mathbb{E}[v_\theta \ell_a] - c_a$; its vanishing is (OB). The $\hat{\theta}$-component is $D_{\hat{\theta}} \mathbb{E}_{\theta,e(\theta)}[v_{\hat{\theta}}]\big|_{\hat{\theta}=\theta}$. The chain-rule identity displayed in the proof of Lemma 5 expresses $U'(\theta)$ as the sum of a true-type term, this report-index term, and an action term that vanishes under (OB); so, given (OB), the report-index term vanishes if and only if (EN) holds. □

Proposition 6 (The ordering)**.** *For every feasible design p, $V^{\mathrm{FOC}}(p) \geq V^{L}(p) \geq V^{F}(p)$.*

Proof. A mechanism satisfying (L1) satisfies (OB) and (EN) by Lemma 4, hence has value at most $V^{\mathrm{FOC}}(p)$ by Proposition 2; this is the first inequality. For the second, let $\mathcal{M}$ satisfy (IC). Then $(\theta, e(\theta))$ maximizes $R(\theta, \cdot, \cdot)$ over $\Theta \times A$ and is interior by Assumption 1, so the first- and second-order necessary conditions for an interior maximum give $\nabla R(\theta) = 0$ and $H_R(\theta) \preceq 0$. Every fully incentive compatible mechanism is therefore locally jointly implementable, and the supremum over the smaller feasible set is no larger. □

Lemma 5 (Truthful reporting is stationary)**.** *Let a mechanism satisfy* (OB) *and* (EN)*, suppose $e(\cdot)$, $U(\cdot)$ and $\hat{\theta} \mapsto v_{\hat{\theta}}$ are differentiable, and suppose the hypothesis of global action optimality holds, so that $\hat{a} = e(\theta)$ is the unique maximizer in* (64) *at $\hat{\theta} = \theta$. Then*

$$\frac{\partial}{\partial \hat{\theta}} G(\theta, \hat{\theta})\bigg|_{\hat{\theta}=\theta} = 0.$$

Proof. By the envelope theorem applied to the maximization over $\hat{a}$ in (64), $\partial_{\hat{\theta}} \widetilde{R}(\theta, \hat{\theta}) = \partial_{\hat{\theta}} \mathbb{E}_{\theta,\hat{a}^*}[v_{\hat{\theta}}]$ with $\hat{a}^*$ held fixed; uniqueness at $\hat{\theta} = \theta$ gives $\hat{a}^* = e(\theta)$, so

$$\partial_{\hat{\theta}} \widetilde{R}(\theta, \theta) = D_{\hat{\theta}}\, \mathbb{E}_{\theta,e(\theta)}\big[v_{\hat{\theta}}\big]\Big|_{\hat{\theta}=\theta}, \tag{95}$$

the derivative with respect to the report index alone. Now differentiate $U(\theta) = \mathbb{E}_{\theta,e(\theta)}[v_\theta] - c(\theta, e(\theta))$ totally in θ. The report index, the true type inside the distribution, and the action all move, so

$$U'(\theta) = \underbrace{\partial_\theta \mathbb{E}_{\theta,e(\theta)}[v_\theta]}_{\text{true type}} + \underbrace{D_{\hat\theta}\mathbb{E}_{\theta,e(\theta)}[v_{\hat\theta}]\big|_{\hat\theta=\theta}}_{\text{report index}} + \Big(\partial_a \mathbb{E}_{\theta,e(\theta)}[v_\theta] - c_a\Big)e'(\theta) - c_\theta.$$

By (OB) the bracket vanishes; by (EN) and Lemma **??**, $U'(\theta) = \partial_\theta \mathbb{E}_{\theta,e(\theta)}[v_\theta] - c_\theta$. Subtracting, $D_{\hat\theta}\mathbb{E}[v_{\hat\theta}]|_{\hat\theta=\theta} = 0$, which by (95) is $\partial_{\hat\theta}\widetilde{R}(\theta,\theta) = 0$; and $\partial_{\hat\theta}G = -\partial_{\hat\theta}\widetilde{R}$. □

Lemma 5 is the reason (L1) does genuine work: the first-order relaxation already delivers stationarity of the report deviation, so local joint implementability reduces to the curvature condition (L2).

The local action branch

Definition 2 (Local and global action branches)**.** Suppose $R_{aa} \le -\kappa < 0$ near $(\theta, \theta, e(\theta))$. The implicit function theorem applied to $R_a(\theta, \hat\theta, a) = 0$ yields a unique C^1 branch $a^L(\theta, \hat\theta)$ with $a^L(\theta,\theta) = e(\theta)$, defined for $\hat\theta$ near θ. Set

$$\widetilde{R}^L(\theta,\hat\theta) = R\big(\theta,\hat\theta, a^L(\theta,\hat\theta)\big), \qquad \widetilde{R}^F(\theta,\hat\theta) = \max_{\hat a \in A} R(\theta,\hat\theta,\hat a),$$

and correspondingly $G^L = U - \widetilde{R}^L$, $G^F = U - \widetilde{R}^F$.

The local branch and the globally optimal action solve different problems. Since a^L is one feasible action, $\widetilde{R}^L \le \widetilde{R}^F$ and $G^L \ge G^F$.

Corollary 2 ((L2) is the curvature condition)**.** *Suppose $R_{aa} < 0$ at the truthful pair. Then*

$$H_R(\theta) \preceq 0 \quad\Longleftrightarrow\quad R_{\hat\theta\hat\theta} - \frac{R_{\hat\theta a}^2}{R_{aa}} \le 0 \quad\Longleftrightarrow\quad \partial_{\hat\theta}^2 G^L(\theta,\theta) \ge 0.$$

Assumption 5 in components

Assumption 5 combines the following. With d and δ as in the text:

Assumption 7 (Joint benchmark slack outside a neighbourhood)**.** There are $\delta > 0$ and $\varepsilon_\delta > 0$ such that the benchmark mechanism satisfies

$$U_0(\theta) - R_0(\theta,\hat\theta,\hat a) \;\ge\; \varepsilon_\delta \qquad \text{whenever } \|d(\theta,\hat\theta,\hat a)\| \ge \delta.$$

Assumption 8 (Neighbourhood strictness and stability)**.** With δ as above there is $\eta > 0$ such that, for the benchmark mechanism, $H_{R,0}(\theta,\hat\theta,\hat a) \preceq -\eta I$ for all $\|d\| \le \delta$; and, along the mechanisms $\mathcal{M}_{r,\hat u}$ and uniformly over unit $\hat u$,

$$\sup_{\|d\|\le\delta} \big\| H_{R,r,\hat u} - H_{R,0} \big\| \longrightarrow 0, \qquad \sup_{\|d\|\ge\delta} \big|\big(U_r - R_r\big) - \big(U_0 - R_0\big)\big| \longrightarrow 0.$$

The near-optimal mechanisms

Throughout Section 7 $\mathcal{M}_{r,\hat{u}}$ denotes a mechanism implementing an allocation that is r^5-near-optimal for the relaxed problem at design $r\hat{u}$,

$$\bar{V}_{\sigma_r(\hat{u})}(r\hat{u}) \;\geq\; V^{\mathrm{FOC}}(r\hat{u}) - r^5; \tag{96}$$

such mechanisms exist by definition of a supremum, and since $r^5 = o(r^2)$ nothing below requires the relaxed supremum to be attained.

Local stability in components

Assumption 4 folds together a benchmark margin and a stability requirement. Written separately:

Assumption 9 (Local stability along the selected mechanisms)**.** Along the near-optimal mechanisms $\mathcal{M}_{r,\hat{u}}$ of above, and uniformly over unit $\hat{u}$,

$$\sup_{\theta}\big|R_{aa,r,\hat{u}}(\theta) - R_{aa,0}(\theta)\big| \longrightarrow 0, \qquad \sup_{\theta}\big|\partial^2_{\hat{\theta}} G^L_{r,\hat{u}}(\theta,\theta) - \partial^2_{\hat{\theta}} G^L_0(\theta,\theta)\big| \longrightarrow 0$$

as $r \downarrow 0$.

The second concerns only two scalars per type, evaluated on the diagonal, and by Lemma 6 it holds automatically at rate $O(r^2)$ when the allocation is held fixed; what it assumes is that the movement of the allocation does not undo that.

The direction under local incentive compatibility

Corollary 3 (The optimal direction is unchanged)**.** *Under the hypotheses of Theorem 3, with $\nu_1(M) > 0$ simple, any maximizing* design $p^*_L(r)$ *of V^L over the capacity ball satisfies*

$$\operatorname{dist}\left(\frac{p^*_L(r)}{r}, \{\pm\hat{u}_1\}\right) \longrightarrow 0.$$

Here p denotes the measurement design and $\hat{u} = p/\|p\|$ its direction; the normalization is by the capacity r, not by the design's own norm.

Proof. By (68) and Proposition **??** the expansion $V^L(p) = V_0 + \mathcal{I}_g(0)\, p^\top M p + o(r^2)$ holds uniformly on $\|p\| \leq r$.

Capacity is asymptotically exhausted. Since $\nu_1(M) > 0$, a design using only a fixed fraction $c < 1$ of capacity delivers at most $\mathcal{I}_g(0)c^2 r^2 \nu_1(M) + o(r^2)$, whereas $r\hat{u}_1$ delivers $\mathcal{I}_g(0) r^2 \nu_1(M) + o(r^2)$. Leaving a nonvanishing fraction of capacity unused therefore sacrifices the fraction $1 - c^2 > 0$ of the leading $O(r^2)$ gain, which optimality forbids for small r; hence $\|p^*_L(r)\|/r \to 1$.

The direction. Writing $p^*_L(r) = \rho_r \hat{u}_r$ and dividing the optimality inequality by $\mathcal{I}_g(0) r^2$ gives $(\rho_r/r)^2\, \hat{u}_r^\top M \hat{u}_r \to \nu_1(M)$; with $\rho_r/r \to 1$ this forces $\hat{u}_r^\top M \hat{u}_r \to \nu_1(M)$, and when ν_1 is simple the maximizers of the Rayleigh quotient are $\{\pm\hat{u}_1\}$. □

G Proofs for Section 7

The second-order effect on deviation incentives

Lemma 6 (Second-order effect on deviation incentives)**.** *Fix an allocation and let the contract be the cheapest one of Theorem 1 at design $p = r\hat{u}$. Then, uniformly in $\hat{u}$,*

$$v_r - v_0 = O(r), \qquad R_r - R_0 = O(r^2), \qquad \partial^2_{\hat{\theta}} G^L_r = \partial^2_{\hat{\theta}} G^L_0 + O(r^2),$$

and likewise $R_{aa,r} = R_{aa,0} + O(r^2)$.

The full statement proved here is the following. Fix $\sigma \in \Sigma$ and let the contract be $v_{r,\hat{u},\hat{\theta}}(y,j) = \bar{v}_\sigma(\hat{\theta}) + \Delta_\sigma(\hat{\theta})^\top \mathcal{I}_r(\hat{\theta})^{-1} \ell_r(y,j;\hat{\theta})$. Write $h_j = g_j'(0)/g_j(0)$ and $\beta_{\hat{\theta}} = \hat{u}^\top w_\sigma(\hat{\theta})$. Then, uniformly in (y,j) and $\hat{u}$:

(i)
$$v_{r,\hat{u},\hat{\theta}}(y,j) = v_{0,\hat{\theta}}(y) + r\,\beta_{\hat{\theta}} h_j + \varrho_r(y,j,\hat{\theta}) \tag{97}$$
with $\|\varrho_r\|_{C^2(\hat{\theta})} = O(r^2)$;

(ii) $\mathbb{E}_{\theta,\hat{a},r\hat{u}}[v_{r,\hat{u},\hat{\theta}}] - \mathbb{E}_{\theta,\hat{a}}[v_{0,\hat{\theta}}] = O(r^2)$, hence $R_{r,\hat{u}} - R_0 = O(r^2)$ and, for the globally optimized gap, $G^F_{r,\hat{u}} - G^F_0 = O(r^2)$;

(iii) under Assumption 6 and C^2 dependence of $(\bar{v}_\sigma, \Delta_\sigma, e)$ on $\hat{\theta}$, $R_{aa,r} = R_{aa,0} + O(r^2)$, $R_{\hat{\theta}a,r} = R_{\hat{\theta}a,0} + O(r^2)$ and $R_{\hat{\theta}\hat{\theta},r} = R_{\hat{\theta}\hat{\theta},0} + O(r^2)$;

(iv) consequently, since R_{aa} stays bounded away from zero and $(x,y,z) \mapsto -x + y^2/z$ is Lipschitz on the relevant range, $\partial^2_{\hat{\theta}} G^L_r = \partial^2_{\hat{\theta}} G^L_0 + O(r^2)$.

Proof. (i). Write $z_{\hat{\theta}} = z\big(\hat{\theta}, e(\hat{\theta}); r\hat{u}\big)$ for the on-path index. Since $e \in C^2$ and z is affine in its arguments,
$$\big|\partial^k_{\hat{\theta}} z_{\hat{\theta}}\big| = O(r), \qquad k = 0,1,2, \tag{98}$$
uniformly: *every* $\hat{\theta}$-derivative of z, not merely z itself, carries a factor r. This is what makes the remainder C^2-small rather than merely small.

By Proposition 3, $\mathcal{I}_r = \mathcal{I}^y + \mathcal{I}_g(z_{\hat{\theta}}) r^2 \hat{u}\hat{u}^\top$, and Sherman–Morrison gives the *exact* identity

$$\mathcal{I}_r^{-1} = (\mathcal{I}^y)^{-1} - r^2 \frac{\mathcal{I}_g(z_{\hat{\theta}})\,(\mathcal{I}^y)^{-1}\hat{u}\hat{u}^\top(\mathcal{I}^y)^{-1}}{1 + \mathcal{I}_g(z_{\hat{\theta}}) r^2\, \hat{u}^\top (\mathcal{I}^y)^{-1}\hat{u}} \;\equiv\; (\mathcal{I}^y)^{-1} + r^2 B_r(\hat{\theta}). \tag{99}$$

Because $\mathcal{I}^y \succeq \varepsilon I$ and $\mathcal{I}_g$ is C^2 and bounded on $[-\zeta,\zeta]$, and because $\mathcal{I}^y(\hat{\theta})$ and $z_{\hat{\theta}}$ are C^2 in $\hat{\theta}$, the matrix B_r is C^2 in $\hat{\theta}$ with $\|B_r\|_{C^2}$ bounded uniformly in r and $\hat{u}$; the denominator is bounded below by 1.

By (**??**), $\ell_r = \ell^y + \big(g_j'(z_{\hat{\theta}})/g_j(z_{\hat{\theta}})\big) r\hat{u}$. Each g_j is C^3 and bounded below on $[-\zeta,\zeta]$ by Assumption 6 and (**??**), so g_j'/g_j is C^2 there and

$$\frac{g_j'(z_{\hat{\theta}})}{g_j(z_{\hat{\theta}})} = h_j + z_{\hat{\theta}}\,\psi_j(z_{\hat{\theta}}), \qquad \psi_j \in C^1, \text{ bounded.} \tag{100}$$

Substituting (99) and (100) into the formula for v_r and using $\Delta^\top(\mathcal{I}^y)^{-1} = w_\sigma^\top$,

$$\varrho_r = \underbrace{r^2\,\Delta^\top B_r\,\ell^y}_{\text{(a)}} \;+\; \underbrace{r\,z_{\hat\theta}\psi_j(z_{\hat\theta})\,\Delta^\top\mathcal{I}_r^{-1}\hat u}_{\text{(b)}} \;+\; \underbrace{r^3 h_j\,\Delta^\top B_r\hat u}_{\text{(c)}}.$$

Each term is a product of functions that are C^2 in $\hat\theta$ with uniformly bounded C^2 norms — Δ_σ and $\bar v_\sigma$ by hypothesis, ℓ^y by Assumption 6, B_r by the above — multiplied by an explicit power of r. Term (a) carries r^2 outright and (c) carries r^3; term (b) carries r explicitly and a further r through $z_{\hat\theta}$, and by (98) that second factor survives differentiation, since $\partial_{\hat\theta}$ and $\partial^2_{\hat\theta}$ applied to $z_{\hat\theta}\psi_j(z_{\hat\theta})$ again produce terms each containing at least one derivative of z, hence a factor r. Applying the product rule to each term therefore gives $\|\varrho_r\|_{C^2(\hat\theta)} = O(r^2)$, which is (97).

(ii). Write $\mathbb{E}_{\theta,\hat a,r\hat u}[v] = \sum_y q(y\mid\theta,\hat a)\sum_j g_j(z_{\theta\hat a})\,v(y,j)$ with $z_{\theta\hat a} = r\,\hat u^\top x(\theta,\hat a)$. Insert (97). The $O(r^2)$ term contributes $O(r^2)$ because $\sum_j g_j \equiv 1$. The order-r term contributes

$$r\big(\hat u^\top w_\sigma\big)\sum_j g_j(z_{\theta\hat a})\frac{g_j'(0)}{g_j(0)}.$$

At $z=0$ the sum is $\sum_j g_j'(0) = 0$, because $\sum_j g_j \equiv 1$; and $g_j(z_{\theta\hat a}) = g_j(0) + O(r)$ uniformly. Hence the sum is $O(r)$ and the whole term is $O(r^2)$. The statement for R follows. For G^F, the bound just obtained is uniform in $\hat a$, so $\sup_{\hat a}|R_r - R_0| = O(r^2)$ and hence $|\max_{\hat a} R_r - \max_{\hat a} R_0| = O(r^2)$, a maximum being 1-Lipschitz in the sup norm; this is $G^F_r - G^F_0 = O(r^2)$.

(iii). Here a second mechanism operates and should be made explicit, because differentiating in a also differentiates z and so could in principle restore a term of order r. Write $\beta_{\hat\theta} = \hat u^\top w_\sigma(\hat\theta)$ and

$$S(z) = \sum_j g_j(z)\,\frac{g_j'(0)}{g_j(0)}, \qquad S(0) = \sum_j g_j'(0) = 0, \qquad S'(0) = \sum_j \frac{g_j'(0)^2}{g_j(0)} = \mathcal{I}_g(0),$$

so that by the computation in (ii)

$$R_r(\theta,\hat\theta,a) = R_0(\theta,\hat\theta,a) + r\,\beta_{\hat\theta}\,S\big(z(\theta,a)\big) + \rho_r(\theta,\hat\theta,a), \tag{101}$$

where $z(\theta,a) = r\,\hat u^\top x(\theta,a)$ and ρ_r collects the contribution of the $O(r^2)$ term of (97). Note that $S'(0) \neq 0$: the payoff-level cancellation of (ii) rests on $S(0)=0$ and does *not* survive differentiation. What makes the derivatives behave is instead that z carries its own factor of r, namely

$$\partial_a z = r\,\hat u_a, \qquad \partial_{\hat\theta} z = 0.$$

Differentiating the middle term of (101):

$$\partial_a^2\big[r\beta S(z)\big] = r\beta\,S''(z)\,(r\hat u_a)^2 = O(r^3), \quad \partial_{\hat\theta a}\big[r\beta S(z)\big] = r\beta'\,S'(z)\,r\hat u_a = O(r^2), \quad \partial^2_{\hat\theta}\big[r\beta S(z)\big] = r\beta''\,S(z) = O(r^2$$

the last because $S(z) = \mathcal{I}_g(0)z + O(z^2) = O(r)$ by $S(0)=0$. So each of the three derivatives picks up $O(r^2)$: the a-derivatives because every differentiation of z brings down a factor r, the $\hat\theta$-derivatives because z does not depend on $\hat\theta$ at all and $S(z)$ is itself $O(r)$. For ρ_r — the contribution of ϱ_r to the payoff — part (i) gives $\|\varrho_r\|_{C^2(\hat\theta)} = O(r^2)$, which is what licenses differentiating it twice in $\hat\theta$; and in (θ,a) it is constant, so those derivatives fall on q and on $g_j(z)$, both bounded with bounded derivatives by Assumption 6 and (**??**), leaving the $O(r^2)$ prefactor intact. Adding the three contributions gives the stated bounds. □

Proposition 5

The branch a^L and the indirect payoff $\widetilde{R}^L$ satisfy

$$\partial_{\hat{\theta}} a^L = -\frac{R_{\hat{\theta}a}}{R_{aa}}, \qquad \partial^2_{\hat{\theta}} \widetilde{R}^L = R_{\hat{\theta}\hat{\theta}} - \frac{R^2_{\hat{\theta}a}}{R_{aa}}, \tag{102}$$

and (66) follows since $G^L = U - \widetilde{R}^L$ and $R_{aa} < 0$. Strict local concavity is imposed separately from Assumption 6, which supplies differentiability but not concavity.

Proof. Differentiating $R_a(\theta, \hat{\theta}, a^L) = 0$ in $\hat{\theta}$ gives $R_{a\hat{\theta}} + R_{aa}\partial_{\hat{\theta}} a^L = 0$, which is the first identity. By the envelope theorem along the branch, $\partial_{\hat{\theta}} \widetilde{R}^L = R_{\hat{\theta}}$ evaluated at a^L; differentiating once more and accounting for the movement of a^L, $\partial^2_{\hat{\theta}} \widetilde{R}^L = R_{\hat{\theta}\hat{\theta}} + R_{\hat{\theta}a}\partial_{\hat{\theta}} a^L = R_{\hat{\theta}\hat{\theta}} - R^2_{\hat{\theta}a}/R_{aa}$. Since $U(\theta)$ does not depend on $\hat{\theta}$, $\partial^2_{\hat{\theta}} G^L = -\partial^2_{\hat{\theta}} \widetilde{R}^L$; the last expression uses $R_{aa} < 0$. □

Corollary 2

Proof. For a symmetric 2×2 matrix with negative $(2,2)$ entry, $H_R \preceq 0$ holds if and only if $\det H_R \geq 0$, that is $R_{\hat{\theta}\hat{\theta}} R_{aa} - R^2_{\hat{\theta}a} \geq 0$; dividing by $R_{aa} < 0$ reverses the inequality and gives the Schur complement condition, which by (66) is $\partial^2_{\hat{\theta}} G^L \geq 0$. (The remaining requirement $R_{\hat{\theta}\hat{\theta}} \leq 0$ is implied, since $R_{\hat{\theta}\hat{\theta}} \leq R^2_{\hat{\theta}a}/R_{aa} \leq 0$.) □

Theorem 3

Proof. Feasibility. $\mathcal{M}_{r,\hat{u}}$ satisfies (OB) and (EN) by construction, hence (L1) by Lemma 4. By Assumption 9 there is $\bar{r} > 0$ with $|R_{aa,r} - R_{aa,0}| \leq \eta/2$ and $|\partial^2_{\hat{\theta}} G^L_r - \partial^2_{\hat{\theta}} G^L_0| \leq \eta/2$ for all $r < \bar{r}$, uniformly in θ and $\hat{u}$. With Assumption 4 this gives $R_{aa,r} \leq -\eta/2 < 0$ and $\partial^2_{\hat{\theta}} G^L_r \geq \eta/2 \geq 0$, which by Corollary 2 is $H_R \preceq 0$, that is (L2).

Value. Being feasible for the V^L problem, $\mathcal{M}_{r,\hat{u}}$ gives $V^L(r\hat{u}) \geq \Pi(\mathcal{M}_{r,\hat{u}}, r\hat{u}) \geq V^{\mathrm{FOC}}(r\hat{u}) - r^5$ by (96). Proposition 6 gives $V^L(r\hat{u}) \leq V^{\mathrm{FOC}}(r\hat{u})$. Since $r^5 = o(r^2)$, the two values agree to second order; Theorem 2 and Assumption **??** then supply the expansion. □

Theorem 4

Proof. Near the truthful pair. By Theorem 3, $\mathcal{M}_{r,\hat{u}}$ satisfies (L1), so $\nabla R_r(\theta) = 0$. By Assumption 5 there is $\bar{r}_1 > 0$ with $H_{R,r} \preceq -\frac{\eta}{2} I$ on $\{\|d\| \leq \delta\}$ for $r < \bar{r}_1$. Taylor's theorem with the integral form of the remainder gives, for such d,

$$R_r(\theta, \hat{\theta}, \hat{a}) - U_r(\theta) = \int_0^1 (1-t)\, d^\top H_{R,r}(\theta; td)\, d\, dt \;\leq\; -\tfrac{\eta}{4}\|d\|^2 \;\leq\; 0.$$

Away from it. By Assumption 5 there is $\bar{r}_2 > 0$ with $U_r - R_r \geq \varepsilon_\delta/2 > 0$ on $\{\|d\| \geq \delta\}$ for $r < \bar{r}_2$.

The two regions cover $\Theta \times A$, and the thresholds do not depend on $\hat{u}$ because the suprema in Assumption 5 are taken over $\hat{u}$ as well. Hence $U_r(\theta) \geq R_r(\theta, \hat{\theta}, \hat{a})$ for all $(\hat{\theta}, \hat{a})$, which is (IC).

Value. Being feasible for the full problem, $\mathcal{M}_{r,\hat{u}}$ gives $V^F(r\hat{u}) \geq \Pi(\mathcal{M}_{r,\hat{u}}, r\hat{u}) \geq V^{\mathrm{FOC}}(r\hat{u}) - r^5$ by (96), while Proposition 6 gives $V^F \leq V^L \leq V^{\mathrm{FOC}}$. Since $r^5 = o(r^2)$ all three agree to second order. □