

Buying Opinions

Mark Whitmeyer* Kun Zhang[†]

September 28, 2022

Abstract

A principal hires an agent to acquire soft information about an unknown state. Even though neither *how* the agent learns (the experiment chosen by the agent) nor *what* the agent discovers (the realization of the experiment) are contractible, the principal is unconstrained as to what information the agent can be induced to acquire and report honestly. When the agent is risk neutral, and a) is not asked to learn too much, b) can acquire information sufficiently cheaply, or c) can face sufficiently large penalties, the principal can attain the first-best outcome. Risk aversion (by the agent) introduces inefficiencies: the first-best is unattainable, though whether the agent obtains rents depends on whether he may exit to take his outside option after learning.

Keywords: Moral hazard, Information acquisition, Rational inattention, Bayesian persuasion, Information design

JEL Classifications: D81; D82; D83; D86

*Arizona State University, Email: mark.whitmeyer@gmail.com

[†]Arizona State University, Email: kunzhang@asu.edu

We thank Brian Albrecht, Hector Chade, Vasudha Jain, Andreas Kleiner, Alejandro Manelli, Teemu Pekkari-nen, Ludvig Sinander, Eddie Schlee, Ina Taneva, Can Urgan, Joseph Whitmeyer, and Thomas Wiseman for their advice. We also appreciate the useful feedback from conference audiences at AMES, ESEM, MWET, NASMES, and Stony Brook 2022; and seminar audiences at ASU, Arizona, Fordham, and Princeton.

1 Introduction

People buy advice: investors pay for stock picks, politicians and executives in firms employ advisors, and bettors at the race track ask for winners. In some situations this advice can be backed up with hard, verifiable, evidence; whereas in others advice is merely cheap talk and honesty is supported only by the advisor's incentives to be truthful. This paper studies the latter situation: we analyze a contracting problem in which a principal hires an agent to acquire unverifiable evidence, which cannot be credibly disclosed or contracted upon.¹

Information acquisition is costly for the agent—after observing the contract, he chooses what information to obtain before reporting his findings to the principal. The agent has significant freedom in his learning: he may choose any distribution over posterior beliefs whose mean is the prior. We allow the agent to exit the relationship at any point: the agent has an outside option, which he can take both before learning and after. This is realistic: an advisor who learns privately before giving advice usually has the option of declining to report to his employer and seeking employment elsewhere.²

One profession that fits our setting is talent scouting. In sports, for instance, teams have hard evidence about prospects (goals scored, batting average, shooting percentage etc...) but nevertheless send scouts to obtain soft (unquantifiable) information: a report from an FC Barcelona scout describes a player's running style, balance, and control, among other things (Vidal (2019)). The scout writes about the player's positioning, "Excellent. It is undoubtedly his best quality. He is always where he should be..."

Remuneration for talent scouts is also based on the realized state. Baseball scouts receive bonuses if a prospect they recommended makes it to a team. Headhunters are rewarded if they identify a candidate that is hired—it is common for firms to pay recruiters a fraction of a hired worker's salary (crucially, *if they are hired*, i.e., make it past the firm's final screening). In addition, these bonuses may not be part of an explicit agreement—

¹This is the key difference between this paper and Rappoport and Somma (2017), who explore a similar problem but specify that evidence is observable and contractible.

²We discuss throughout how results change when the agent may only exit the relationship at the outset.

one baseball scout states “I had a contract that was sort of word-of-mouth that if they drafted or signed anybody they didn’t already have info on, they’d give me \$5,000” (Silvy (2014))—which supports our assumption that the expert may leave at any time.

As in the classical setup (e.g., Holmström (1979)), our principal’s problem can be decomposed into two parts: first, given a desired distribution over posteriors, how can the principal implement such a distribution as cheaply as possible? Second, given the answer to the first question, what distribution over posteriors would the principal like to implement?

We begin by observing that any contract induces a decision problem for an agent. This allows us to show, in Lemma 4.1, that the principal can implement any feasible learning: she can write a contract such that the agent is willing to learn precisely as desired *and* report honestly. That is, the agency problem does not impede the principal’s ability to acquire information. This result stands in stark contrast to the classical setting, wherein not all effort levels or distributions thereof can be implemented in the second-best world.

Next, we show that the required optimality of the agent’s learning produces a number of conditions whose structure allows us to simplify the principal’s problem. Theorem 4.3 states that for any state k , each message contingent transfer in that state can be written as the difference between the transfer paid in that state for a “benchmark message” and a constant that depends only on exogenous values and the posteriors themselves. Not only do the relative incentives completely pin down the agent’s optimal learning, but the converse is also true: the agent’s optimal learning specifies the relative incentives.

As in the classical moral hazard environment, there is a natural benchmark in our model: the first-best problem in which learning (our analog of effort) is observable and contractible. Our Proposition 5.1 establishes that when the agent is risk neutral and negative transfers are allowed, any distribution over posteriors can be implemented efficiently, even in our main setting with hidden learning and unverifiable evidence. Moreover, this holds even though the agent may exit the relationship after acquiring information, which renders the “selling the project to the agent” contract generically ineffective. This highlights another essential difference between the canonical setting and ours. In the classical setting, the possibility of an interim exit is more-or-less equivalent to a limited liability

constraint, which allows the agent to accrue rents.

If negative transfers are forbidden and the outside option is sufficiently low, the principal cannot efficiently acquire information through the agent. Nevertheless, we show that optimal incentives take simple forms in a number of cases. In Proposition 6.2, we provide a full characterization of the optimal contract when the agent's outside option is sufficiently small. There, it is only the limited liability constraint that binds, which allows us to pin down the optimal contract for any desired distribution over posteriors. We also fully characterize optimal implementation with an arbitrary outside option and limited liability in the binary-state case when the agent is risk neutral. In particular, implementation is efficient if and only if the agent is not asked to learn too much (in relation to her cost of acquiring information and outside option).

We also show that the agent's risk aversion introduces inefficiencies: providing incentives for the agent to learn requires that he be exposed to risk, which is surplus destroying when the agent is risk averse. Similar to the classical setting, we establish that with only an *ex ante* participation constraint, the agent gets zero rents. On the other hand, the possibility of an interim exit grants the agent surplus (generically).

We finish this section by discussing related literature. Section 2 lays out the model before Section 3 states the principal's problem and discusses the first-best benchmark. Section 4 presents some preliminary results, and Sections 5 and 6 contain the main results in the absence and presence of limited liability constraints, respectively. We wrap things up in Section 7.

1.1 Related Literature

Our study belongs to the literature on delegated expertise, pioneered by Lambert (1986), Demski and Sappington (1987) and Osband (1989), in which a principal hires an agent to collect payoff relevant information. The central theme of this literature is incentive design for effective information acquisition and communication.

There are three recent papers that are close to ours. Rappoport and Somma (2017) also study contracting for flexible information acquisition; crucially, they assume that the posterior generated by the agent's choice of distribution is verifiable and contractible.

They show that the first best can be achieved whenever the agent is risk averse and is not subject to limited liability, or risk neutral and subject to limited liability.³ [Zermeño \(2011\)](#) and [Clark and Reggiani \(2021\)](#) explore contracting environments in which both information acquisition and decision making are delegated to an agent. [Zermeño’s](#) focus is the interaction between the variables on which the transfer schemes can depend and whether contracts specify transfer scheme menus. [Clark and Reggiani \(2021\)](#) show that any Pareto-optimal contract can be decomposed into a fraction of output, a state-dependent transfer, and an optimal distortion.

[Carroll \(2019\)](#) studies a robust contracting problem in which the principal has limited knowledge about how the agent can learn and evaluates each possible contract by its worst-case guarantee. In [Häfner and Taylor \(2022\)](#) the agent acquires information to help the principal decide how much she should invest in a project. The distribution over posteriors and its cost are primitives of the model, and the agent’s report of the realized posterior is unverifiable. Their focus is on finding the optimal contract—which can depend on the report and the outcome of the project—that motivates the agent to conduct the experiment and report truthfully.⁴

[Gromb and Martimort \(2007\)](#) consider a problem of delegated expertise with two agents, where the agents may collude among themselves or with the principal. In their model, the state space is binary, and the agent is restricted to a fixed message space containing two signals whose meaning is common knowledge. Their one-agent/one-signal case is similar to our model: the agent is risk neutral and protected by limited liability, the compensation can be conditioned on both the report and the realized state, and incentives must be provided for the agent to gather information and report truthfully. [Chade and Kovrijnykh \(2016\)](#) study a dynamic model of contracting for information acquisition

³[Bizzotto et al. \(2020\)](#) consider a similar problem. However, they only allow the agent to deviate to a “default” distribution, instead of any Bayes-plausible distribution. Also related is [Yoder \(2022\)](#), who studies a model of adverse selection without risk aversion and limited liability in which the agent’s marginal cost of acquiring information (κ) is the agent’s private information.

⁴[Terovitis \(2018\)](#) tackles a similar problem. In his framework, the outcome is deterministically pinned down by the action and state, and the decision is delegated to the agent.

in a two state-two (fixed) signals environment. The more effort the agent exerts, the more informative the signal he acquires. They assume that the realized signals are contractible, but the true state is not.

Since in our model every contract induces a decision problem with a posterior separable cost of the agent, our work is naturally related to the rational inattention literature pioneered by Sims (1998, 2003). To analyze the agent’s problem, we use insights from Caplin et al. (2022). Maćkowiak et al. (Forthcoming) provides an excellent review of this literature that covers both theory and applications.

Our principal also needs to elicit information from the agent, which connects our paper to the belief elicitation literature. Indeed, our transfer scheme is a scoring rule. The most important distinction is that the beliefs in our work are endogenously determined through the agent’s learning. While the papers in that literature study what scoring rules induce truthful reporting (Gneiting and Raftery (2007) and Schlag et al. (2015) are good surveys) and what properties of a state distribution can be elicited (see Lambert (2019) and references therein), our focus is on deriving incentive contracts that induce the agent to learn *and* report truthfully.

Finally, because we study the motivation of an agent to acquire costly and unverifiable information, our work also connects to the moral hazard literature. In the canonical moral hazard problem (see, for example, Mirrlees (1999), Holmström (1979), Grossman and Hart (1983), and Holmström and Milgrom (1987)), the agent is impelled to exert costly effort that yields some (distribution over) output; whereas in ours, he must be coerced into choosing a much more complicated object (a particular probability distribution) then reporting honestly.

2 The Model

The principal (she) is faced with a decision problem in which she chooses an action $a \in A$, where A is compact. The payoff to each action depends on an unknown state of the world $\theta \in \Theta$, where Θ is a finite set; $|\Theta| = n < \infty$. The principal’s utility from taking action a in state θ is given by $u(a, \theta)$, where u is continuous in a . $\mu \in \Delta(\Theta)$ denotes the common

(full-support) prior belief about the state.

The principal cannot acquire information herself but instead must rely on the assistance of an agent (he), who acquires information flexibly at a cost. Specifically, the agent may choose any Bayes-plausible (Kamenica and Gentzkow (2011)) distribution over posteriors, $F \in \Delta\Delta(\Theta)$, subject to a posterior separable cost C à la Caplin et al. (2022). That is, the cost of acquiring F is

$$C(F) = \kappa \int_{\Delta(\Theta)} c(\mathbf{x}) dF(\mathbf{x}) ,$$

where $\kappa > 0$ is a scaling parameter, $c: \Delta(\Theta) \rightarrow \mathbb{R}_+$ is a strictly convex and twice continuously differentiable function bounded on the interior of $\Delta(\Theta)$, and $c(\mu) = 0$. This class of information costs includes the entropy-based cost function (see e.g. Sims (1998, 2003), and Matějka and McKay (2015)); the log-likelihood cost of Pomatto et al. (2020); and the quadratic (posterior variance) cost function.

After acquiring information, the agent sends a message to the principal, who then takes an action. The true state is eventually observable to both parties after the action is taken and can be contracted upon. A contract specifies the set of messages available to the agent, and a transfer paid to the agent which can be contingent on both the realized state and the message sent. Formally, the principal proposes a pair (M, t) consisting of a compact set of messages M available to the agent, and a transfer $t: M \times \Theta \rightarrow \mathbb{R}$ ($t: M \times \Theta \rightarrow \mathbb{R}_+$ when the agent is protected by limited liability). We assume the principal's payoff is quasi-linear in the transfer. The agent's payoff is additive separable in his utility from the transfer and the cost of acquiring information, and he values the transfer according to a continuously differentiable, weakly concave, and strictly increasing function v , with $v(0) = 0$. To ease presentation, transfer t is expressed in utils.

We further assume the agent has access to an outside option of value $v_0 \geq 0$, and that there are two chances for him to leave with his outside option: he can choose not to accept the contract, or walk away after acquiring information by reporting nothing.⁵ More generally, the agent's outside option could derive from some salvage value for information

⁵We term this the null message, \emptyset . Thus, the set of messages available to the agent is $M \cup \{\emptyset\}$, where sending message \emptyset yields the agent a state-independent payoff of v_0 .

that is an arbitrary upper semicontinuous function of the posterior $p(\mathbf{x})$. In the supplementary appendix, we explain that this does not alter our analysis in a meaningful way.

Unless otherwise noted, we assume throughout that the principal suffers a penalty that is strictly greater than v_0 if the agent takes his outside option. This ensures that it is not optimal for the principal to offer the agent a contract in which he ever takes his outside option. In the discussion following Theorem 4.3, we argue that our framework can accommodate “shoot the messenger” contracts—in which the agent is asked to exit the relationship with positive probability on path—with ease.

The timing of the game is as follows:

- (i) The principal proposes a contract (M, t) ;
- (ii) If the agent does not accept, the game ends; otherwise the agent chooses a Bayes-plausible distribution F , from which a posterior $\mathbf{x} \in \Delta(\Theta)$ is drawn and privately observed by the agent;
- (iii) The agent chooses whether to report. If he reports, he sends a message $m \in M$; and if he does not report, he takes his outside option v_0 (and the principal observes the null message);
- (iv) The principal takes an action $a \in A$, the true state $\theta \in \Theta$ realizes, and payoffs accrue: the principal gets $u(a, \theta) - v^{-1}(t(m, \theta))$ and the agent $t(m, \theta) - c(F)$.

3 The Principal’s Problem

3.1 The First Best Benchmark

As is standard, our specification allows us to write the principal’s (expected) gross payoff as a function of the posterior \mathbf{x} ; denote it by $V(\mathbf{x})$. V is convex, and if A is finite V is piecewise affine and convex (being the maximum of affine functions). Denote the set of Bayes-plausible distributions over posteriors by $\mathcal{F}(\mu)$. It is a convex and compact subset of $\Delta\Delta(\Theta)$. If the principal controlled the information acquisition, she would solve

$$\max_{F \in \mathcal{F}(\mu)} \int (V - \kappa c) dF ,$$

which is a linear functional of F , guaranteeing the existence of a maximizer.

In our context, “first best” refers to the situation where the principal can observe the distribution over posteriors chosen by the agent, so the principal can specify transfer $t: \Delta(\Theta) \rightarrow \mathbb{R}_+$. When the distribution is observable, the following contract implements any distribution F and is optimal: the principal pays the agent precisely the amount that makes him indifferent between learning and walking away with his outside option if and only if the agent acquires F . Otherwise, the principal pays the agent nothing. Evidently, the transfer is never strictly negative, and the agent is willing to acquire F . Therefore, at the first best, the principal’s cost of acquiring information is $v^{-1}(C(F) + v_0)$.

3.2 The Contracting Problem

A contract must guarantee that the agent chooses the right distribution and report honestly. Following [Caplin et al. \(2022\)](#), we define a *decision problem* (μ, D, w) as the choice over a compact set of actions D given the prior μ over states in Θ , and $w: D \times \Theta \rightarrow \mathbb{R}$ is the decision maker’s utility function. Given a decision problem (μ, D) , the decision maker chooses a Bayes-plausible distribution over posteriors G and an action strategy $\sigma: \text{supp}(G) \rightarrow \Delta(D)$.

A contract (M, t) induces a decision problem (μ, M, t) of the agent; and we say that a distribution F is *implementable* if there exists a contract (M, t) such that $M = \text{supp}(F)$, and the agent’s optimal strategy is $(F, \{\delta_{\mathbf{x}}\}_{\mathbf{x} \in \text{supp}(F)})$.⁶ Equivalently, the contract (M, t) *implements* F . In particular, we say that F can be *implemented efficiently* if it can be implemented at the first-best cost.

For any $d \in M$, we define the agent’s *net utility* $N(\mathbf{x}|d)$ as the expected utility of the message d net of the cost of \mathbf{x} :

$$N(\mathbf{x}|d) = \mathbb{E}_{\mathbf{x}}[t(d, \theta)] - \kappa c(\mathbf{x}) .$$

The agent chooses a distribution over posteriors G to maximize his value function $W(\mathbf{x}) = \max_{d \in M} N(\mathbf{x}|d)$. The agent’s optimal choice is determined by concavifying the value function: \mathcal{H} denotes the hyperplane tangent to the hypograph of W at the support points of

⁶ $\delta_{\mathbf{x}}$ denotes the degenerate distribution at \mathbf{x} .

F . We can identify this supporting hyperplane \mathcal{H} by an affine function $f_{\mathcal{H}}(\mathbf{x}): \Delta(\Theta) \rightarrow \mathbb{R}$. The set of optimal posteriors is the set of points at which $f_{\mathcal{H}}$ and W intersect, which we denote by $P_{(M,t)}$. By construction, at every optimal posterior \mathbf{x}_j , $W(\mathbf{x}_j) = N(\mathbf{x}_j | \mathbf{x}_j)$; that is, it is optimal for the agent to report the realized posterior honestly. Therefore, a necessary condition for a distribution F to be implemented by a contract (M, t) is that $\text{supp}(F) = P_{(M,t)}$.

The above condition need not be sufficient for implementation: the contract must also prevent the agent from walking away at any point in the interaction. In particular, no matter what the realized posterior is, the agent cannot deviate profitably by taking his outside option without making a report; this requires (note that this constraint holding also prevents potential double deviations in which the agent learns differently before taking his outside option)

$$f_{\mathcal{H}}(\mathbf{x}) \geq v_0 - \kappa c(\mathbf{x}) \quad \text{for all } \mathbf{x} \in \Delta(\Theta). \quad (IR - v_0)$$

This constraint is stronger than the *ex ante* participation constraint $f_{\mathcal{H}}(\mu) \geq v_0$. Thus,

Lemma 3.1. *A contract (M, t) implements distribution F if and only if*

- (i) **Incentive Compatibility:** $\text{supp}(F) = P_{(M,t)}$; and
- (ii) **Individual Rationality:** Constraint $IR - v_0$ holds; and
- (iii) **Limited Liability:** if there is limited liability, $t(m, \theta) \geq 0$ for all $\theta \in \Theta$ and $m \in M$.

To solve the principal's contracting problem, we adopt a two-step approach: first, for every implementable distribution F , we solve the principal's cost minimization problem:

$$\min_{(M,t)} \mathbb{E}_{F,\mathbf{x}} \left[v^{-1}(t(\mathbf{x}, \theta)) \right],$$

subject to (i), (ii), and (iii) in Lemma 3.1; denote its value by $\Gamma(F)$. Second, the principal chooses an implementable distribution F to maximize her payoff under agency, $\int V(\mathbf{x}) dF(\mathbf{x}) - \Gamma(F)$. Like most papers studying moral hazard, we focus on the first step.

4 Preliminary Results

We begin by arguing that any distribution over posteriors with support on n or fewer points can be implemented by some contract.

Lemma 4.1. *If F is a distribution over posteriors with $|\text{supp}(F)| \leq n$ and $\text{supp}(F) \subseteq \text{int } \Delta(\Theta)$, there exists a contract (M, t) that implements F , and the expected cost to the principal is finite.*

The proof of Lemma 4.1, and all other proofs omitted from the main text, are collected in Appendix A. For each F supported on n or fewer interior points of $\Delta(\Theta)$, because the cost function c is bounded and differentiable on $\Delta(\Theta)$, Lemma 2 of Caplin et al. (2022) guarantees that there is a decision problem such that F is optimal. Therefore, we can construct a contract—which induces a decision problem for the agent—with bounded transfers such that the agent finds it optimal to first acquire F then report the realized posterior truthfully. Moreover, by adding a finite constant to the transfer, we can make Constraint $IR - v_0$ hold. Therefore, every such distribution can be implemented at finite cost.

Because the support of any extreme point of $\mathcal{F}(\mu)$ is on n or fewer points, every $F \in \mathcal{F}(\mu)$ can be expressed as a convex combination of distributions with support on n or fewer points. Therefore, any $F \in \mathcal{F}(\mu)$ can be obtained by randomizing over a set of contracts each of which implements a distribution with support on at most n points—consequently, any distribution whose support is on the interior of $\Delta(\Theta)$ can be induced at a finite expected cost. As it is weakly less costly for the principal to randomize first rather than implement F directly, it is without loss of generality for the principal to implement a distribution over posteriors with support on at most n points.

Corollary 4.2. (i) *Every $F \in \mathcal{F}(\mu)$ with $\text{supp}(F) \subseteq \text{int } \Delta(\Theta)$ can be implemented at a finite cost.*

(ii) *Without loss of generality, the principal only implements distributions with support on at most n points.*

By Corollary 4.2 (ii), we can restrict our attention to distributions over posteriors supported on $\{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_s\}$, where n is the number of states, and $s \leq n$. In our next result, we discover that optimality allows us to reduce transfers to a single variable for each state.

For each state $k = 1, \dots, n$, define $\Omega^k(i, j) := t_i^k - t_j^k$ ($i, j = 1, \dots, s$). Each $\Omega^k(i, j)$ specifies the difference between the payoff to the agent from sending any (on path) message i versus message j in state k . Importantly, because (on path) each message corresponds to a different posterior, the collection of differences $(\Omega^k(i, j))_{k=1}^n$ captures the relative benefit

to the agent from obtaining posterior j rather than posterior i .

Theorem 4.3 (Identification/Non-identification). *Given a distribution over posteriors F chosen by an agent and an information acquisition cost function c , only the relative incentives $(\Omega^k(i, j))_{i, j=1, \dots, s; k=1, \dots, n}$ are identified.*

The principal's problem of optimally inducing a distribution over posteriors reduces to an n -variable optimization problem, where n is the number of states. For each state k , the principal fixes a benchmark message $j(k)$, then chooses $(t_{j(k)}^k)_{k=1}^n$; the payoff to the agent from sending message $j(k)$ in state k .

Theorem 4.3 is reminiscent of the standard result that truth-telling only identifies relative payments in adverse selection settings. Here; however, the relative incentives are pinned down jointly by the optimality of the desired distribution *ex ante* and truthful reporting *ex interim*. As part (i) of Lemma 3.1 states, incentive compatibility for the agent requires that the value function of the agent, W , intersects the concavifying hyperplane at the support points of the distribution over posteriors the contract aims to implement. Such a hyperplane pins down the transfers in each state for each posterior (in the support of the agent's learning). Consequently, the principal's problem is equivalent to one of choosing a hyperplane, which is an n -variable optimization problem.⁷

That was a technical explanation, here is an economic one. Fix a desired distribution; if the agent wants to deviate by slightly increasing the probability of a message realization in a certain state, basic probability implies that there must be a commensurate decrease in the probability of another message realization in that state. At the optimum, no such local deviation in the agent's information acquisition strategy can be profitable. Hence for any two posteriors in the support of the desired distribution, any deviation of the sort described above must generate a marginal value to the agent equal to the marginal cost. Because the marginal value of varying the probability of a message realization in a state is determined by the transfer for sending that message in that state, this "zero net marginal gain" observation generates an equality that connects the transfers for sending

⁷Framed in this manner, this theorem is closely related to Lemma 2 in Caplin et al. (2022), which states that when constructing a decision problem the tangent hyperplane is arbitrary.

two distinct posteriors in the support of the desired distribution in the same state.

Recall that we specified early on that the principal suffers a disutility greater than v_0 should the agent take his outside option. This ensures that the principal does not want to replace one of the messages with the null message, i.e., have the agent exit the relationship, sending the null message with positive probability. In principle, if the principal is not hurt (severely) by the agent's exit, it could be optimal for the principal to write a contract in which the null message is sent with positive probability (inducing the desired posterior) thereby allowing the principal to save on paying the agent. By Theorem 4.3 the belief to which the null message corresponds pins down the other transfers. Thus, if the principal's penalty from an agent's exit is less than v_0 , one must check at most s additional contracts (other than those in which the null message is never sent), in which the null message is sent after each belief, in turn.⁸

5 Main Results I. No Limited Liability

5.1 Risk-Neutral Agent

Recalling that $f_{\mathcal{H}}$ is the function that identifies the concavifying hyperplane \mathcal{H} , the agent's value from acquiring information for the principal (from an *ex ante* perspective) is $f_{\mathcal{H}}(\mu)$. Thus, efficient implementation requires $f_{\mathcal{H}}(\mu) = v_0$, which implies that Constraint $IR - v_0$ must bind at $\mathbf{x} = \mu$. In this case, Constraint $IR - v_0$ reduces to

$$t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j) = -\kappa c_k(\mu) \quad \text{for all } k = 1, \dots, n-1, \quad (IR - R)$$

where the index for posterior j is arbitrary.⁹ By Theorem 4.3, $(t_j^k)_{k=1}^n$ identifies a contract.

When there is no limited liability and the agent is risk neutral, whether a distribution can be implemented efficiently boils down to whether the system of equations defined by the $n-1$ equations in Constraint $IR - R$ and $f_{\mathcal{H}}(\mu) = v_0$ has a solution. In fact,

⁸In the supplementary appendix we discuss an example in which it is optimal for the agent to exit the relationship with positive probability.

⁹This is because optimality requires $t_i^k - t_i^n - \kappa c_k(\mathbf{x}_i) = t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j)$ for all $i, j = 1, \dots, s$ and each $k = 1, \dots, n-1$.

Proposition 5.1. *If the agent is risk neutral and not protected by limited liability, every (feasible) distribution F with $\text{supp}(F) \subseteq \text{int } \Delta(\Theta)$ can be implemented efficiently.¹⁰*

When there is no limited liability, the amount of incentive constraints is “just right” such that there exists a transfer scheme that delivers the right incentives and keeps the agent’s surplus at his outside option. Figure 1 illustrates this construction. The gross payoffs (as a function of posterior x) to the agent from sending each one of the contract’s two messages are the blue and purple lines. The maximum of these functions, net of the agent’s learning cost, is the agent’s induced value function, W , depicted in black. The concavifying line $f_{\mathcal{H}}$ —which pins down optimal learning—is in orange. Finally, the agent’s net payoff from taking the outside option v_0 is the red curve.

It is instructive to compare Proposition 5.1 to Proposition 2 in [Rappoport and Somma \(2017\)](#), which states that when the realized posteriors are contractible (but the true state is not), efficient implementation is possible when the agent is risk neutral, even if he is protected by limited liability. This is made possible in their problem by assigning a transfer for each posterior in the support of the distribution as a divergence from the prior—which is by construction nonnegative and hence satisfies limited liability—and the expected transfer equals the cost of generating the distribution. In fact, in their setting, when the agent is risk-neutral and there is no limited liability, selling the project to the agent is optimal. Not so here, as we discuss shortly.

If we require limited liability, for some distributions and outside option values, efficient implementation cannot be achieved (irrespective of the interim IR constraint’s presence). Because posteriors are unverifiable and hence noncontractible, [Rappoport and Somma](#)’s construction does not work. In our model, to induce the agent to gather information, the transfers must be “rewarding” when the agent “gets the state right” and “punishing” when he is wrong. This gap between the two outcomes must be large enough to justify the cost of learning. Therefore, when v_0 is small enough, to achieve an expected transfer of $\Gamma(F) = C(F) + v_0$, some “punishing” transfer(s) must be negative.

¹⁰The supplementary appendix reveals that this holds even when there are uncountably many states.

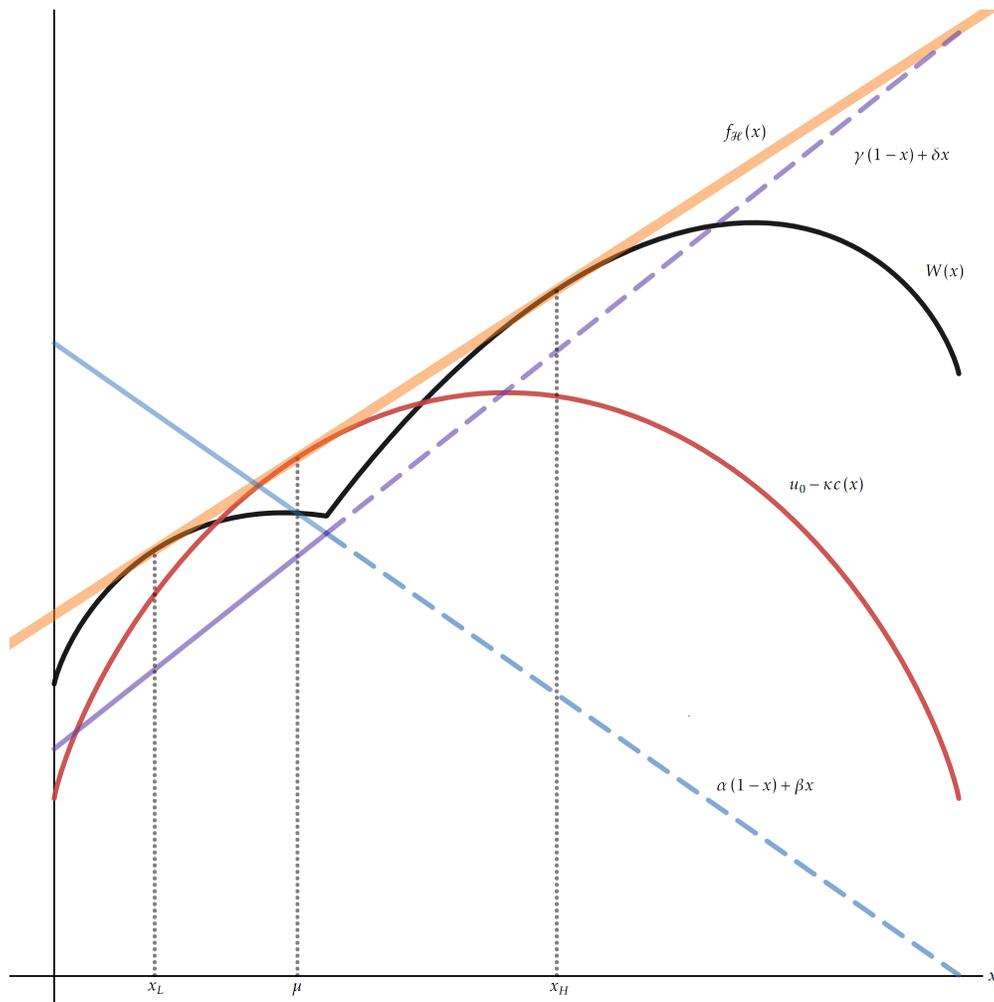


Figure 1: Efficient implementation of $x_L = 1/9$, $x_H = 5/9$ when $\mu = 1/(1 + e)$, $\kappa = 1$, and $v_0 = \log\{9/(1 + e)\}$, and with entropy cost. This contract satisfies the limited liability constraints—as stated in Proposition 6.3, the specified ratio v_0/κ is the minimum such ratio such that efficient implementation is feasible under limited liability.

5.1.1 Comparison to Classical Moral Hazard with an Interim Participation Constraint

A natural question is whether an analog of Proposition 5.1 holds if we add in an interim participation constraint to the canonical moral hazard problem. That is, is the ability of the principal to accommodate the interim participation constraint a special feature of our information acquisition problem or does it also hold in the classical environment?

As we show in the supplementary appendix, in the classical environment, the ability of the agent to exit the relationship after (privately) observing his output realization is tantamount to limited liability: clearly a contract cannot promise the agent a payoff less than his outside option for any (divulged) output. Moreover, because a contract must be incentive compatible, unless the principal implements the lowest possible effort, i.e., pays a constant wage, the agent must get strictly positive rents.

5.1.2 Selling the Project to the Agent?

One might also wonder whether Proposition 5.1 is really needed. In the standard moral hazard problem, when the agent is risk neutral and there are no limited liability constraints, the principal can attain the first best by “selling the project to the agent” (henceforth the **STP** contract). In our setting,¹¹ that corresponds to the principal writing the contract such that the agent’s net utility as a function of his posterior \mathbf{x} is $V(\mathbf{x}) - \kappa c(\mathbf{x}) - \tau$, where $\tau = f_{\mathcal{H}}(\mu) - v_0$. That is, the principal writes a contract so that the agent’s decision problem, gross of the cost, is precisely that faced by the principal, then lowers the agent’s uniformly to leave his *ex ante* expected payoff equal to his outside option.

It is easy to see that absent the interim IR constraint ($IR - v_0$), the first best can be attained in this way. However, when the agent can take his outside option *ex interim*, the principal cannot implement a distribution over posteriors at the first-best cost generically by selling the project. Indeed, consider the concavifying hyperplane corresponding to the agent’s optimal learning under the STP contract when there is no interim participation constraint. For the agent’s optimal learning to remain unchanged after adding in the in-

¹¹We are also assuming here that the set of actions in the principal’s decision problem is finite. The construction when the principal has infinitely many actions is analogous but more ungainly, so we omit it.

terim IR constraint, this hyperplane must be tangent to the function $v_0 - \kappa c(\mathbf{x})$ at μ . That is, the optimal contract must be robust to double deviations in which the agent learns differently then takes her outside option with positive probability. Unless the curves are tangent, such a deviation exists. This is, however; a non-generic property of the principal's decision problem, which we establish formally in the supplementary appendix.

5.2 Risk-Averse Agent

When the agent is risk averse, but unprotected by limited liability, characterizing the optimal contract is more involved. Fix an arbitrary benchmark message for all states, say j ; the principal's payoff is strictly decreasing in each of the n control variables $(t_j^k)_{k=1}^n$ and so the principal wants to set each one as low as possible. Unencumbered by limited liability, the lone constraint is $IR - v_0$, which necessarily binds (since otherwise, the principal could reduce the control variables). Thus,

Remark 5.2. When the agent is risk averse, there exists an $\mathbf{x}^* \in \Delta(\Theta)$ such that $f_{\mathcal{H}}(\mathbf{x})$ is tangent to $v_0 - \kappa c(\mathbf{x})$ at \mathbf{x}^* .

Given this, solving for the optimal implementation of a distribution over posteriors F can be turned into an $n - 1$ variable optimization problem by using the tangency conditions to substitute in for each t_j^k .¹² This yields the principal an objective that is a function of \mathbf{x}^* . Unless $\mathbf{x}^* = \mu$, which does not hold in general, the agent obtains positive rents. This finding is a consequence of the interim participation constraint, which requires that $f_{\mathcal{H}}$ lie above $v_0 - c$ everywhere. Otherwise—with only *ex ante* IR—the agent would not obtain rents. Indeed, without the interim IR constraint, the lone constraint is the *ex ante* participation constraint, which obviously binds; hence, $f_{\mathcal{H}}(\mu) = v_0$.

Summarizing things,

Proposition 5.3. *Suppose the agent is risk averse and not protected by limited liability.*

- (i) *If the agent can exit ex interim, he gets strictly positive rents unless $\mathbf{x}^* = \mu$. If the agent cannot exit ex interim, he gets zero rents.*

¹²More precisely, we have $t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j) = -\kappa c_k(\mathbf{x}^*)$ for all $k = 1, \dots, n - 1$, and $f_{\mathcal{H}}(\mathbf{x}^*) = v_0 - \kappa c(\mathbf{x}^*)$.

(ii) Only the degenerate distribution of posteriors can be implemented efficiently.

In choosing \mathbf{x}^* , the principal optimally trades off between risk sharing and conceding rents: when a contract that makes the agent break even entails too much risk, moving \mathbf{x}^* away from μ mitigates this risk. Then, although the agent receives strictly positive rents, implementing the new contract can be cheaper to the principal. This is reminiscent of the trade-off studied in Proposition 5 in [Rappoport and Somma \(2017\)](#) though the exact mechanisms are different: in their work, the most cost-efficient way for compelling the agent to choose the right distribution is to have the hyperplane determined by the wage contract (which, in their setting, is a function of the *verifiable* posterior) to be tangent to the agent's value function. In our problem, averting double deviations to the outside option is what begets the tangency condition mentioned in Remark 5.2.

When either the agent cannot exit at the interim stage or can but $\mathbf{x}^* = \mu$, the principal faces a trade-off between insurance and incentives like in the classical moral hazard model. Because the agent is risk averse, the cheapest way to generate a value of v_0 for the principal is a flat wage; but a flat wage fails to provide the correct (relative) incentives for the agent to choose the desired (nondegenerate) distribution and report truthfully. Because the agent, therefore, must be exposed to risk, the first best is unattainable. This result is an analog of the celebrated result in the canonical moral hazard model with a risk-averse agent that only the lowest effort level can be implemented efficiently.

In their Proposition 1, [Rappoport and Somma \(2017\)](#) show that when the agent is risk averse and not protected by limited liability, efficient implementation is always possible. This result contrasts sharply with ours, even if the agent is not allowed to exit *ex interim*. Without loss of generality, assume that the desired distribution F has linearly independent support. Then F is an extreme point of $\mathcal{F}(\mu)$, and hence for any other distribution, say F' , there exists a posterior $\mathbf{x}' \in \text{supp}(F') \setminus \text{supp}(F)$. Then because the posteriors are observable in their model, a contract that offers a constant transfer for all posteriors in the support of F , and specifies a severe enough punishment for all other posteriors implements F at the first-best cost. In our model; however, the agent can secretly acquire a different distribution but still report a posterior in the support of the desired distribution, which breaks down the “moving support property” they leverage.

5.2.1 Entropy Cost & Logarithm Utility Example

It is straightforward to solve for the optimal contract when there are just two states, the agent's utility $v(t) = \log(t + 1)$ and the agent's cost $c(x) = x \log x + (1 - x) \log(1 - x) - \mu \log \mu - (1 - \mu) \log(1 - \mu)$. The point of tangency, x^* , that pins down the optimal contract (Remark 5.2) is the solution to the first-order condition (FOC) in x :

$$(\mu - x_L) \left[\left(\frac{1 - x_H}{1 - x} \right)^{\kappa+1} - \left(\frac{x_H}{x} \right)^{\kappa+1} \right] + (x_H - \mu) \left[\left(\frac{1 - x_L}{1 - x} \right)^{\kappa+1} - \left(\frac{x_L}{x} \right)^{\kappa+1} \right] = 0.$$

When $\kappa = 1$ we can go even further. For $\kappa = 1$ the FOC yields

$$x^* = \begin{cases} \frac{-\sqrt{((x_L - \mu)x_H + (2 - x_L)\mu - 1)((x_L - \mu)x_H - x_L\mu) - (x_L - \mu)x_H + x_L\mu}}{2\mu - 1}, & \text{if } \mu \neq \frac{1}{2} \\ \mu = \frac{1}{2}, & \text{if } \mu = \frac{1}{2} \end{cases}.$$

It can be checked that $x^* \neq \mu$ unless $\mu = 1/2$, and therefore the agent generically gets strictly positive rents. When $\mu \neq 1/2$, x^* is increasing in x_H and decreasing in x_L : as either x_H increases or x_L decreases, at the initial x^* , the marginal benefit of better risk sharing dominates the marginal loss in conceding rents.

Substituting the expression for x^* into the Principal's objective we obtain that the cost to the principal of implementing a posterior distribution with support $\{x_L, x_H\}$ is increasing in x_H and decreasing in x_L : more informative distributions are more expensive to implement. Less obvious is the fact that the surplus accrued by the agent is also increasing in the informativeness of the distribution he is asked to acquire. This can be seen by computing the agent's payoff from an *ex ante* perspective:

$$\mu \log\left(\frac{\mu}{x^*}\right) + (1 - \mu) \log\left(\frac{1 - \mu}{1 - x^*}\right).$$

Intuitively, if $x^* \neq \mu$, as x_H increases, x^* moves in the same direction, and hence $f_{\mathcal{H}}(\mu) - v_0$ is increasing since $v_0 - c$ is strictly concave. A similar argument applies when x_L decreases.

6 Main Results II. Limited Liability

Throughout this section, we assume that the agent is protected by limited liability. In Subsection 6.1, we solve for the optimal incentives when the agent's value for his outside

option is sufficiently small. In Subsection 6.2, we allow for an arbitrary outside option but impose that the agent is risk neutral.

6.1 Low Outside Option

For simplicity, we set $v_0 = 0$; it is not hard to see that all the results in this subsection go through for all sufficiently small $v_0 > 0$. By Theorem 4.3, for any desired distribution, the relative incentives are identified. Consequently, for each state we can pinpoint a benchmark message that determines the lowest payment.

Lemma 6.1. *For every state $k = 1, \dots, n$, there exists $j^*(k)$ such that $t(\mathbf{x}_{j^*(k)}, \theta_k) \leq t(\mathbf{x}_i, \theta_k)$ for all $i = 1, \dots, s$.*

Lemma 6.1 allows us to completely identify the optimal transfers when the agent's outside option is sufficiently low.

Proposition 6.2. *Suppose $v_0 = 0$, and the agent is protected by limited liability. Then for each state $k = 1, \dots, n$, there exists $j^*(k)$ such that $t(\mathbf{x}_{j^*(k)}, \theta_k) = 0$, and all other transfers are nonnegative and determined by optimal learning.*

Proposition 6.2 is intuitive: for a sufficiently small outside option, Constraint $IR - v_0$ always holds, so the transfer scheme is pinned down by optimal learning and limited liability. Optimal learning leaves, for each state, one degree of freedom to the principal; and to satisfy limited liability, the best that the principal can do is to find the smallest transfer in each state and set it to zero.

6.2 Risk-Neutral Agent

Now, we dispense with the assumption that the outside option is small— v_0 can take any value. For expository ease, we start with the two state case and then argue that our results generalize when there are more than two states.

6.2.1 Two States

When there are just two states, $\Theta = \{\theta_1, \theta_2\}$. By Corollary 4.2 (ii), we can identify a distribution by its support $\{x_L, x_H\}$. Our first result characterizes the distributions over posteriors that a principal can implement efficiently; *viz.*, at the first-best cost. Defining

$$\eta(x_L, x_H) := \max\{-\mu c'(\mu) - c(x_H) + c'(x_H)x_H, (1 - \mu)c'(\mu) - c(x_L) - (1 - x_L)c'(x_L)\},$$

we have

Proposition 6.3. *The principal can implement $\{x_L, x_H\}$ efficiently if and only if $v_0/\kappa \geq \eta(x_L, x_H)$.*

Proposition 6.3 states that a given distribution can be implemented efficiently if and only if either the agent has a sufficiently high outside option, or he can acquire information sufficiently cheaply. Recall from the discussion in Section 5.1 that, when the agent is risk neutral and his outside option is zero, we can always find a transfer scheme that implements any distribution efficiently—it is just that the transfer is negative for some message in some states. When the agent’s outside option increases to some $v_0 > 0$, to make sure that the agent’s expected payoff equals v_0 while maintaining the incentive for acquiring the same distribution, the new transfer scheme must be “lifted up” by v_0 for every pair of state and posterior realization.¹³ Consequently, for v_0 sufficiently high, all transfers become nonnegative and hence limited liability is satisfied.

The left-hand side of Proposition 6.3’s necessary and sufficient condition is strictly increasing in the outside option v_0 and strictly decreasing in the cost of information κ . Moreover, it is easy to calculate that η is decreasing in x_L and increasing in x_H . This suggests the following corollary:

Corollary 6.4. *(i) For any pair of posteriors $\{x_L, x_H\}$ with $0 < x_L \leq \mu \leq x_H < 1$, if v_0/κ is sufficiently large, $\{x_L, x_H\}$ can be implemented efficiently.¹⁴*

¹³More precisely, optimal learning only leaves one degree of freedom on the transfers for each state, and the tangency conditions in Constraint $IR - R$ connect different states, so the entire transfer scheme is determined by these up to a constant.

¹⁴If $c'(0)$ and $c'(1)$ are finite, this is true for any $0 \leq x_L \leq \mu \leq x_H \leq 1$.

- (ii) *Efficient implementation is monotone with respect to the Blackwell order: if $\{x_L, x_H\}$ can be implemented efficiently, then any distribution that corresponds to a less informative experiment can be implemented efficiently.*
- (iii) *If $v_0 > 0$ then any distribution that corresponds to a sufficiently uninformative experiment can be implemented efficiently.*

In the canonical moral hazard problem with a risk-averse agent, no matter what outside option the agent has, only the lowest action can be implemented efficiently. Corollary 6.4 has a flavor of that classical result: efficient implementation is possible whenever the agent is not asked to learn too much. For distributions more spread than some set of threshold distributions; however, positive rents must be provided to the agent. To implement such distributions efficiently it must be that the relative incentives are high enough for the agent to acquire that much information and so when v_0 is small limited liability is always violated.

When the first-best implementation of $\{x_L, x_H\}$ is infeasible, there are three other possibilities, listed in our next proposition. Denoting $\gamma := t_2^1$ the transfer from sending message x_L in state θ_2 , and $\beta := t_1^2$ the transfer from sending message x_H in state θ_1 , we have

Proposition 6.5. *One of the following must occur at the optimum. Either*

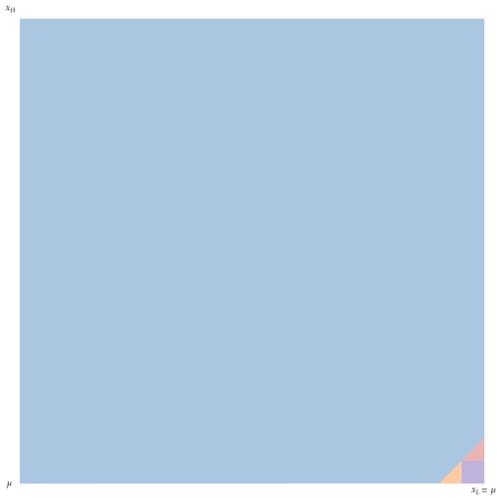
- (i) *$\{x_L, x_H\}$ can be implemented efficiently (and Constraint $IR - v_0$ binds); or*
- (ii) *$\{x_L, x_H\}$ cannot be implemented efficiently; and either*
 - (a) *Constraint $IR - v_0$ binds and $\beta = 0$; or*
 - (b) *Constraint $IR - v_0$ binds and $\gamma = 0$; or*
 - (c) *Constraint $IR - v_0$ does not bind and $\gamma = \beta = 0$.*

When the cost function is the entropy cost, it is straightforward to characterize the four regions of $\{x_L, x_H\}$ pairs. They are depicted in Figure 2.

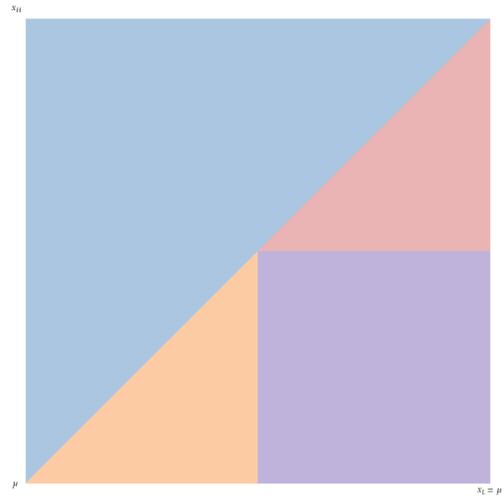
6.2.2 Two States and No Interim IR

Things are even simpler when there is no interim IR constraint. Now, individual rationality reduces to $f(\mu) \geq v_0$. Defining

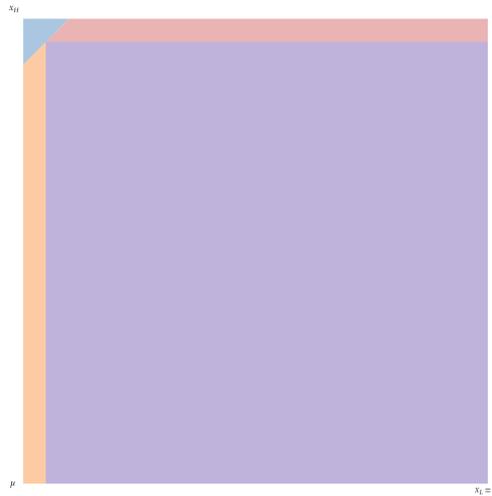
$$\zeta(x_L, x_H) := (1 - \mu)(x_H c'(x_H) - c(x_H)) - \mu((1 - x_L)c'(x_L) + c(x_L)) ,$$



(a) $v_0/\kappa = .05$.



(b) $v_0/\kappa = \log(1/(1-\mu))$.



(c) $v_0/\kappa = 3$.

Figure 2: **Implementation Regions for $\mu = 1/2$:** Pairs (x_L, x_H) in the purple region can be implemented efficiently, (x_L, x_H) in the blue region are optimally implemented by $\gamma = \beta = 0$, (x_L, x_H) in the orange region are optimally implemented by $\beta = 0$ and some $\gamma \geq 0$; and (x_L, x_H) in the red region are optimally implemented by $\gamma = 0$ and some $\beta \geq 0$.

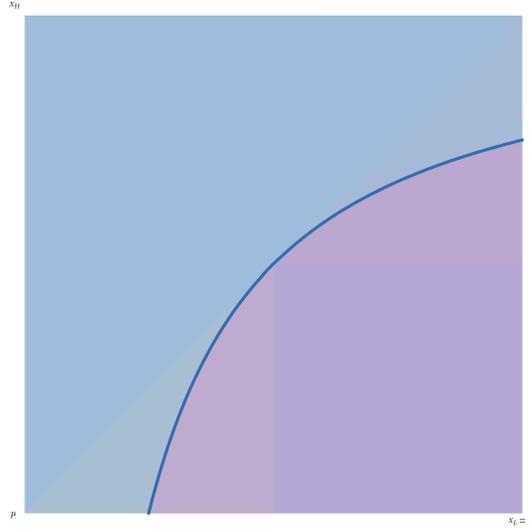


Figure 3: **Implementation Regions for $\mu = 1/2$ and $v_0/\kappa = \log(1/(1-\mu))$ (No Interim IR):** Pairs (x_L, x_H) in the purple region can be implemented efficiently and (x_L, x_H) in the blue region are optimally implemented by $\gamma = \beta = 0$.

we have

Proposition 6.6. *The principal can implement $\{x_L, x_H\}$ efficiently if and only if $v_0/\kappa \geq \zeta(x_L, x_H)$. Otherwise, $\gamma = \beta = 0$.*

It is obvious that an exact analog of Corollary 6.4 holds when there is no interim IR constraint. *Viz.*, any pair of posteriors can be implemented efficiently if the outside option is sufficiently large and the cost of acquiring information κ is sufficiently small. Moreover, the more information an agent is asked to acquire, the more difficult it is to implement the distribution efficiently.

When the cost function is the entropy cost, it is straightforward to characterize the two regions of $\{x_L, x_H\}$ pairs. They are depicted in Figure 3, superimposed over the four regions present when there is an interim participation constraint.

6.2.3 More Than Two States

By Lemma 6.1, for each state k we can find a message $j^*(k)$ that delivers the lowest payment; and by Theorem 4.3, to pin down the transfer scheme, it suffices to determine $t_{j^*(k)}^k$

for each state k . Thus, there are n unknowns. When we impose the interim IR constraint, there are n equations: efficient implementation is equivalent to $f_{\mathcal{H}}(\mu) = v_0$, and the other $n - 1$ equations are given by Constraint *IR - R*:

$$t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j) = -\kappa c_k(\mu) ,$$

where $k = 1, \dots, n - 1$ indicates the state, and j , which indexes the posterior, is arbitrary. Then the distribution can be efficiently implemented if and only if $t_{j^*(k)}^k \geq 0$ for each k ; consequently, Proposition 6.3, Corollary 6.4 (i), Proposition 6.5, and Proposition 6.6 naturally extend to more than two states.

7 Discussion

We conclude by discussing a few of our assumptions.

The true state is contractible. A key stipulation in our model is that the true state is observable and contractible *ex post*. While this assumption is standard in the literature, and fits some applications well, it can be extended. If we assume that the state is not observable, but there is an outcome variable that is determined probabilistically by the state, the principal could merely treat this outcome (random) variable as the state and contract upon that. Even in applications in which we cannot find such an outcome variable, our results may still apply: if distinct action-state pairs lead to different payoffs for the principal, the state can be inferred from the realized payoff.¹⁵ If only the unverifiable message that the agent sends can be contracted upon; however, the agent cannot be coerced into gathering any nontrivial information: he never learns anything and sends the message that yields the highest reward.

Agent has no intrinsic preferences for learning. We make this assumption to zero in the incentive provision problem when the principal has to delegate information acquisi-

¹⁵This assumption is not as restrictive as it first seems: for instance, different actions taken by firm executives usually result in nonidentical performance under dissimilar market conditions, and distinct portfolio choices typically yield unequal revenue under different stock market trends.

tion to an agent who cannot make verifiable reports. To allow the agent to have intrinsic motivation, we assume that the agent's intrinsic value from posterior \mathbf{x} is $\phi(\mathbf{x})$, which is known to both parties. Then it is as if that the agent's cost of arriving at posterior \mathbf{x} is $\kappa c(\mathbf{x}) - \phi(\mathbf{x})$, and all of our results survive intact. It is also reasonable in some applications to assume that the agent has intrinsic preferences over the action to be taken by the principal. This is also easy to accommodate in a similar manner.

Other objectives of the principal. For the sake of exposition, we assume that the principal would like to acquire information to improve her decision making. Our model allows for a more general use of information produced by the agent: for all of our results to hold, we only need to assume that the principal's indirect utility is a function of the posterior. For example, the principal might seek to influence the choice of a third party decision maker as in the literature on Bayesian persuasion and information design.

Restricted learning. Our agent is unconstrained in how she learns: she may choose any Bayes-plausible distribution. How might our results change if the agent instead could only choose from some subset thereof? In general, such restrictions must make implementation (of those available distributions) weakly cheaper, as the agent has fewer possible deviations. A corollary of this observation (proved in the supplementary appendix), is that the principal can implement any (feasible) distribution efficiently if agent is risk neutral and there are no limited liability constraints.

References

Jacopo Bizzotto, Eduardo Perez-Richet, and Adrien Vigier. Information design with agency. *Mimeo*, 2020.

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130(6):1676–1715, 2022.

- Gabriel Carroll. Robust incentives for information acquisition. *Journal of Economic Theory*, 181:382–420, 2019.
- Hector Chade and Natalia Kovrijnykh. Delegated Information Acquisition with Moral Hazard. *Journal of Economic Theory*, 162:55–92, 2016.
- Aubrey Clark and Giovanni Reggiani. Contracts for acquiring information. *arXiv: 2103.03911*, 2021.
- Joel S Demski and David E M Sappington. Delegated expertise. *Journal of Accounting Research*, 25:68–89, 1987.
- Tilmann Gneiting and Adrian E Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Denis Gromb and David Martimort. Collusion and the organization of delegated expertise. *Journal of Economic Theory*, 137(1):271–299, 2007.
- Sanford J Grossman and Oliver D Hart. An Analysis of the Principal-Agent Problem. *Econometrica*, 51(1):7, 1983.
- Samuel Häfner and Curtis R Taylor. On young turks and yes men: Optimal contracting for advice. *The RAND Journal of Economics*, 53(1):63–94, 2022.
- Bengt Holmström. Moral hazard and observability. *The Bell Journal of Economics*, 10(1):74–91, 1979.
- Bengt Holmström and Paul Milgrom. Aggregation and Linearity in the Provision of Intertemporal Incentives. *Econometrica*, 55(2):303–328, 1987.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, 2011.
- Nicolas S Lambert. Elicitation and Evaluation of Statistical Forecasts. *Mimeo*, 2019.
- Richard A Lambert. Executive Effort and Selection of Risky Projects. *The RAND Journal of Economics*, 17(1):77–88, 1986.

- Bartosz Maćkowiak, Filip Matějka, and Mirko Wiederholt. Rational inattention: A review. *Journal of Economic Literature*, Forthcoming.
- Filip Matějka and Alisdair McKay. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review*, 105(1):272–98, 2015.
- J A Mirrlees. The theory of moral hazard and unobservable behaviour: Part i. *The Review of Economic Studies*, 66(1):3–21, 1999.
- Kent Osband. Optimal Forecasting Incentives. *Journal of Political Economy*, 97(5):1091–1112, 1989.
- Luciano Pomatto, Philipp Strack, and Omer Tamuz. The cost of information. *Mimeo*, 2020.
- Daniel Rappoport and Valentin Somma. Incentivizing Information Design. Available at SSRN 3001416, 2017.
- Karl H Schlag, James Tremewan, and Joël J van der Weele. A Penny for Your Thoughts: a Survey of Methods for Eliciting Beliefs. *Experimental Economics*, 18:457–490, 2015.
- Tyler Silvy. What’s it like to be an mlb scout? *The Coloradoan*, 2014. URL <https://www.coloradoan.com/story/sports/2014/06/06/like-mlb-scout/10076017/>.
- Christopher A Sims. Stickiness. In *Carnegie-rochester conference series on public policy*, volume 49, pages 317–356, 1998.
- Christopher A Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- Spyros Terovitis. Motivating Information Acquisition Under Delegation. Available at SSRN 3264177, 2018.
- Cristian Martín Vidal. *La Masia. Developing People Beyond Sport*. Editorial Base, 2019.

Nathan Yoder. Designing incentives for heterogeneous researchers. *Journal of Political Economy*, 130(8):2018—2054, 2022.

Luis Zermeno. A principal-expert model and the value of menus. *Mimeo*, 2011.

A Omitted Proofs

A.1 Lemma 4.1 Proof

Proof. Let $\text{supp}(F) = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_s\}$, where $s = |\text{supp}(F)| \leq n$. Consider a contract (M, t) where $M = \text{supp}(F)$, and for each $j = 1, \dots, s$,

$$t(\mathbf{x}_j, \theta_k | \tau) = \kappa c(\mathbf{x}_j) - \sum_{i=1}^{n-1} x_j^i \kappa c_i(\mathbf{x}_j) + \kappa c_k(\mathbf{x}_j) + \tau \text{ for all } k = 1, \dots, n-1,$$

$$t(\mathbf{x}_j, \theta_n | \tau) = \kappa c(\mathbf{x}_j) - \sum_{i=1}^{n-1} x_j^i \kappa c_i(\mathbf{x}_j) + \tau,$$

where x_j^i is the i -th entry of \mathbf{x}_j , c_i is the partial derivative of c with respect to its i -th entry, and τ is a constant that scales the transfers. Now the agent is facing a decision problem (μ, M, t) . Let G be a distribution over posteriors, and let $\sigma : \Delta(\Theta) \rightarrow \Delta(M)$ denote a reporting strategy. Then, the agent's value of (G, σ) is given by

$$Y(G, \sigma) = \sum_{\mathbf{x} \in \text{supp}(G)} \sum_{d \in M} G(\mathbf{x}) \sigma(d | \mathbf{x}) N(\mathbf{x} | d).$$

We claim that (F, σ^*) , where $\sigma^*(\cdot | \mathbf{x}_j) = \delta_{\mathbf{x}_j}$ is an optimal strategy for the agent. By Lemma 1 in [Caplin et al. \(2022\)](#), it suffices to show that, for every \mathbf{x}_j , $j = 1, \dots, s$, there exists a $n-1$ dimensional vector $\lambda = (\lambda_1, \dots, \lambda_s)$ such that

$$N(\mathbf{x} | d) - \sum_{i=1}^{n-1} \lambda_i x^i \leq N(\mathbf{x}_j | \mathbf{x}_j) - \sum_{i=1}^{n-1} \lambda_i x_j^i,$$

for all $\mathbf{x} \in \Delta(\Theta)$ and $d \in M$. We set λ to be the zero vector, so the above inequality reduces to $N(\mathbf{x} | d) \leq N(\mathbf{x}_j | \mathbf{x}_j)$. We first show that for any fixed $d \in M$, $N(\mathbf{x} | d) \leq N(\mathbf{x}_j | d)$, and then

we show that $N(\mathbf{x}_j|d) \leq N(\mathbf{x}_j|\mathbf{x}_j)$. To establish the first inequality, since $c(\mathbf{x})$ is strictly convex, the first-order conditions (FOCs) are sufficient; the FOCs are

$$t(d, \theta_i|\tau) - t(d, \theta_n|\tau) - \kappa c_i(\mathbf{x}) = \kappa (c_i(\mathbf{x}_j) - c_i(\mathbf{x})) = 0 \text{ for all } i = 1, \dots, n-1,$$

clearly setting $\mathbf{x} = \mathbf{x}_j$ makes all of them hold. For the second inequality,

$$N(\mathbf{x}_j|\mathbf{x}_j) - N(\mathbf{x}_j|d) = \kappa \left(c(\mathbf{x}_j) - c(d) - \sum_{i=1}^{n-1} (x_j^i - d_i) c_i(d) \right) \geq 0,$$

where the inequality follows from the convexity of c . Therefore, (F, σ^*) is indeed optimal, and it is direct that the agent's payoff is $\Upsilon(F, \sigma^*) = \tau$. Moreover, there exists $\tau^* < \infty$ large enough, since c is bounded and differentiable on $\text{int} \Delta(\Theta)$, such that Constraint $IR - v_0$ holds. Thus, contract (M, t) implements F . The principal's expected cost is finite since $t(\mathbf{x}_j, \theta_k|\tau^*)$ is finite for all j, k . ■

A.2 Corollary 4.2 Proof

Proof. Let $\text{ext} \mathcal{F}(\mu)$ denote the set of extreme points of $\mathcal{F}(\mu)$. Because $\mathcal{F}(\mu)$ is convex and compact, by Choquet's theorem, for any $G \in \mathcal{F}(\mu)$ there exists a probability measure Λ_G that puts probability 1 on $\text{ext} \mathcal{F}(\mu)$, and

$$G = \int_{\text{ext} \mathcal{F}(\mu)} H d\Lambda_G(H). \quad (R)$$

Therefore, any distribution G with support on $\text{int} \Delta(\Theta)$ can be obtained by randomizing over distributions supported on at most n points. Then by Lemma 4.1, G can be implemented at a finite cost by randomizing over contracts we constructed therein. This establishes part (i).

For part (ii), suppose there exists a contract (M, t) under which the agent chooses G , where $|\text{supp}(G)| > n$, and $(G, \hat{\sigma})$ is the induced optimal strategy of the agent. Without loss of generality, $M = \text{supp}(G)$ and $\hat{\sigma}(\cdot|\mathbf{x}) = \delta_{\mathbf{x}}$ for all $\mathbf{x} \in \text{supp}(G)$. Then for every posterior $\mathbf{x} \in \text{supp}(G)$ and every $d \in M$ with $\hat{\sigma}(d|\mathbf{x}) > 0$,

$$N(\mathbf{x}|d) + \sum_{i=1}^{n-1} (t(d, \theta_i) - t(d, \theta_n) - \kappa c_i(\mathbf{x})) (\tilde{x}_i - x_i) \geq N(\tilde{\mathbf{x}}|d') \quad (H)$$

for all $\tilde{\mathbf{x}} \in \Delta(\Theta)$ and $d' \in M$. By Equation *R*, for every $F \in \text{supp}(\Lambda_G)$, and every posterior \mathbf{x} , $\mathbf{x} \in \text{supp}(G)$. Hence, the strategy $(F, \hat{\sigma}|_{\text{supp}(F)})$ is also optimal for the agent since Inequality *H* holds for every $\mathbf{x} \in \text{supp}(F)$ and every $d \in M$ with $\hat{\sigma}|_{\text{supp}(F)}(d|\mathbf{x}) > 0$. Now it is direct that each $F \in \text{supp}(\Lambda_G)$ can be implemented by the contract (M_F, t_F) where $M_F = \text{supp}(F)$, and t_F is the restriction of t to M_F ; thus, G can be implemented at the same cost by randomizing over $\text{supp}(\Lambda_G)$.

Note that; however, for all $F \in \text{supp}(\Lambda_G)$, (M_F, t_F) need not be the least costly contract under which the agent chooses F : randomizing over $\text{supp}(\Lambda_G)$ and finding the least costly contract for each F is at least weakly cheaper than (M, t) . Therefore, without loss of generality, the principal only implements distributions with support on at most n points. This concludes the proof of part (ii). \blacksquare

A.3 Theorem 4.3 Proof

Proof. The principal wants to implement a distribution F using some contract (M, t) . By part (i) of Lemma 3.1, a necessary condition for implementation is that $\text{supp}(F) = P_{(M,t)}$; this condition holds if and only if the contract is such that the following s expressions

$$\begin{aligned} & \sum_{k=1}^{n-1} x_1^k t_1^k + \left(1 - \sum_{k=1}^{n-1} x_1^k\right) t_1^n - \kappa c(\mathbf{x}_1) + \sum_{k=1}^{n-1} (t_1^k - t_1^n - \kappa c_k(\mathbf{x}_1))(x_k - x_1^k) \\ & \sum_{k=1}^{n-1} x_2^k t_2^k + \left(1 - \sum_{k=1}^{n-1} x_2^k\right) t_2^n - \kappa c(\mathbf{x}_2) + \sum_{k=1}^{n-1} (t_2^k - t_2^n - \kappa c_k(\mathbf{x}_2))(x_k - x_2^k), \\ & \vdots \\ & \sum_{k=1}^{n-1} x_s^k t_s^k + \left(1 - \sum_{k=1}^{n-1} x_s^k\right) t_s^n - \kappa c(\mathbf{x}_s) + \sum_{k=1}^{n-1} (t_s^k - t_s^n - \kappa c_k(\mathbf{x}_s))(x_k - x_s^k) \end{aligned}$$

define the same hyperplane, where $t_j^k := t(\mathbf{x}_j, \theta_k)$, x_j^i is the i -th entry of \mathbf{x}_j , and c_i is the partial derivative of c with respect to its i -th entry. Accordingly, for all $k = 1, \dots, n-1$ and $i, j = 1, \dots, s$

$$t_i^k - t_i^n - \kappa c_k(\mathbf{x}_i) = t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j) \quad \text{and} \quad t_i^n = t_j^n + \Xi_{ij},$$

where Ξ_{ij} is some function of the primitives (but not directly of the t s). Combining these two equations, we obtain

$$\Omega^k(i, j) = \kappa c_k(\mathbf{x}_i) - \kappa c_k(\mathbf{x}_j) + \Xi_{ij}, \text{ for } k = 1, \dots, n-1 \text{ and } \Omega^n(i, j) = \Xi_{ij}.$$

Accordingly, for each state $k = 1, \dots, n$, once the principal chooses the transfer for one of the messages in state k , the transfers for all other messages are automatically pinned down. In other words, the principal has one degree of freedom for each of the states. In every state $k = 1, \dots, n$, and for every $i, j = 1, \dots, s$, we can write $t_i^k = t_j^k + \Omega^k(i, j)$. ■

A.4 Proposition 5.1 Proof

Proof. Let F be such that $\text{supp}(F) = \{\mathbf{x}_1, \dots, \mathbf{x}_s\} \subseteq \text{int} \Delta(\Theta)$, where $s \leq n$. As noted in the main text, there are $n-1$ equations given by Constraint $IR-R$: $t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j) = -\kappa c_k(\mu)$ for all $k = 1, \dots, n-1$, and efficient implementation requires $f_{\mathcal{H}^e}(\mu) = v_0$, which can be written as

$$\sum_{k=1}^{n-1} (t_j^k - t_j^n - \kappa c_k(\mathbf{x}_j)) \mu_k + t_j^n = Q,$$

where μ_k is the k -th entry of μ , and Q does not depend on t 's. To show that F can be efficiently implemented, it suffices to find a solution of this system of n equations. Using $IR-R$, the equality above can be reduced to $t_j^n = Q + \sum_{k=1}^{n-1} \kappa \mu_k c_k(\mu)$; plugging this into the other $n-1$ equations, we get $t_j^k = Q + \sum_{i=1}^{n-1} \kappa \mu_i c_i(\mu) + \kappa (c_k(\mathbf{x}_j) - c_k(\mu))$ for each $k = 1, \dots, n$. We have thus found a solution. Because F is an arbitrary distribution over posteriors supported on at most n points, the principal can implement any distribution G with $\text{supp}(G) \subseteq \text{int} \Delta(\Theta)$ efficiently by randomizing *ex ante*. ■

A.5 Proposition 5.3 Proof

Proof. Suppose first that the agent can exit *ex interim*. Because c is strictly convex, $v_0 - c(\mathbf{x})$ is strictly concave. By Remark 5.2, $f_{\mathcal{H}^e}(\mathbf{x})$ is tangent to $v_0 - c(\mathbf{x})$ at \mathbf{x}^* . Thus, $\mathbf{x}^* \neq \mu$ implies that $f_{\mathcal{H}^e}(\mu) > v_0$, and hence the agent gets strictly positive rents. When the agent cannot exit *ex interim*, the fact that he gets zero rents is almost immediate: if $f_{\mathcal{H}^e}(\mu) > v_0$, because there is no limited liability, the transfer can be lowered by some small $\varepsilon > 0$.

Let F be a nondegenerate distribution with $\text{supp}(F) = \{\mathbf{x}_1, \dots, \mathbf{x}_s\}$, where $\mathbf{x}_i \neq \mathbf{x}_j$ for all $i, j = 1, \dots, s$ with $i \neq j$. Suppose to the contrary that F can be efficiently implemented. To simplify notation, let $x_j^n := 1 - \sum_{k=1}^{n-1} x_j^k$ for each $j = 1, \dots, s$. The principal wishes to minimize $\sum_{j=1}^s \sum_{k=1}^n p_j x_j^k v^{-1}(t_j^k)$. Because F can be efficiently implemented,

$$\sum_{j=1}^s \sum_{k=1}^n p_j x_j^k v^{-1}(t_j^k) = v^{-1}(C(F) + v_0) = v^{-1}\left(\sum_{j=1}^s \sum_{k=1}^n p_j x_j^k t_j^k\right).$$

Because v is strictly concave, v^{-1} is strictly convex; Jensen's inequality then implies $t_j^k = \tilde{t}$ for all $k = 1, \dots, n$ and $j = 1, \dots, s$. But then since c is strictly convex, learning according to the degenerate distribution is uniquely optimal to the agent, and hence the contract with constant transfer cannot implement F . A contradiction. ■

A.6 Lemma 6.1 Proof

Proof. Fix any state k and an arbitrary message, say s . Define $N(k) = \{i : \Omega^k(i, s) < 0\}$. If $N(k) = \emptyset$, let $j^*(k) = s$; then since $t_i^k = t_{j^*(k)}^k + \Omega^k(i, s)$, we have $t_{j^*(k)}^k \leq t_i^k$ for all $i = 1, \dots, s$. Otherwise, let $j^*(k)$ be an arbitrary selection of $\arg \min_{j \in N(k)} \Omega^k(j, s)$. Optimal learning requires, for any i , $t_i^k = t_s^k + \Omega^k(i, s)$ and $t_{j^*(k)}^k = t_s^k + \Omega^k(j^*(k), s)$, which implies $t_{j^*(k)}^k - t_i^k = \Omega^k(j^*(k), s) - \Omega^k(i, s) \leq 0$. Again, $t_{j^*(k)}^k \leq t_i^k$ for all $i = 1, \dots, s$. ■

A.7 Proposition 6.2 Proof

Proof. By Lemma 6.1, for every state $k = 1, \dots, n$, there exists $j^*(k)$ such that $t_{j^*(k)}^k \leq t_i^k$ for all $i = 1, \dots, s$. Then by setting $t_{j^*(k)}^k = 0$, the agent's honesty is not affected, and the limited liability constraints are satisfied. For every $i \neq j^*(k)$, we have

$$t_i^k = \Omega^k(i, j^*(k)) = \kappa c_k(\mathbf{x}_i) - \kappa c_k(\mathbf{x}_{j^*(k)}) + \Xi_{ij^*(k)}$$

for each $k = 1, \dots, n-1$; and $t_i^n = \Xi_{ij^*(n)}$ for $i \neq j^*(n)$. ■

A.8 Proposition 6.3 Proof

Proof. Without loss of generality $\alpha := t_1^1 \geq t_2^1 =: \gamma$; and $\delta := t_2^2 \geq t_1^2 =: \beta$. In this case, it is convenient to write down the agent's value function:

$$W(x) = \begin{cases} \alpha(1-x) + \beta x - \kappa c(x), & \text{if } 0 \leq x \leq \frac{\alpha-\gamma}{\alpha-\gamma+\delta-\beta} \\ \gamma(1-x) + \delta x - \kappa c(x), & \text{if } \frac{\alpha-\gamma}{\alpha-\gamma+\delta-\beta} \leq x \leq 1 \end{cases}.$$

Consequently, the equations that pin down the agent's optimal learning simplify to

$$\kappa(c'(x_H) - c'(x_L)) = A + B \quad \text{and} \quad A + \kappa(c(x_H) - c(x_L)) = \kappa(c'(x_H)x_H - c'(x_L)x_L),$$

where $A := \alpha - \gamma \geq 0$, $B := \delta - \beta \geq 0$. Because c is strictly convex, both A and B are strictly positive if $x_L < \mu < x_H$, and zero if $x_L = x_H = \mu$. Furthermore, the concavifying line is

$$f(x) = (\beta - \gamma - A - \kappa c'(x_L))x + \gamma + A - \kappa(c(x_L) - x_L c'(x_L)). \quad (\star)$$

The principal chooses γ and β in order to maximize

$$-\gamma(1-\mu) - \beta\mu - px_H B - (1-p)(1-x_L)A,$$

where $p = (\mu - x_L)/(x_H - x_L)$ is the (unconditional) probability that posterior x_H realizes, subject to limited liability: $\beta, \gamma \geq 0$, and

$$f(x) \geq v_0 - \kappa c(x) \quad \text{for all } x \in [0, 1], \quad (IR-v_0)$$

where f is given in Equation \star . By construction the agent cannot deviate profitably by learning differently *and* reporting to the principal. Constraint $IR-v_0$ ensures that the agent cannot deviate profitably by learning differently and taking his outside option.

Using the concavifying line (\star) , $\{x_L, x_H\}$ can be implemented efficiently iff

- (i) $(\beta - \gamma - A - \kappa c'(x_L))\mu + \gamma + A - \kappa(c(x_L) - x_L c'(x_L)) = v_0$; and
- (ii) $\beta - \gamma - A - \kappa c'(x_L) = -\kappa c'(\mu)$; and
- (iii) $\beta, \gamma \geq 0$.

From (i) and (ii),

$$\gamma = v_0 + \kappa c'(\mu)\mu - A - \kappa(c'(x_L)x_L - c(x_L)) = v_0 + \kappa c'(\mu)\mu - \kappa(c'(x_H)x_H - c(x_H)),$$

and

$$\beta = v_0 - \kappa(1-\mu)c'(\mu) + \kappa(1-x_L)c'(x_L) + \kappa c(x_L).$$

(iii) requires $v_0/\kappa \geq \eta(x_L, x_H)$, as stated in the result. ■

A.9 Proposition 6.5 Proof

Proof. (i) is a consequence of Proposition 6.3. Suppose that $v_0/\kappa < \eta(x_L, x_H)$ so that efficient implementation is infeasible. Recall that P wants to maximize $-\gamma(1-\mu) - \beta\mu$. Thus, if $\gamma = \beta = 0$ is implementable, they are obviously optimal. Substituting them into the concavifying line (★) we get

$$h(x) = -(A + \kappa c'(x_L))x + A - \kappa(c(x_L) - x_L c'(x_L)) .$$

We need to check for which values of x_L and x_H h lies above $v_0 - \kappa c(x)$. To that end, we define function $g(x) := h(x) - v_0 + \kappa c(x)$. Then,

$$g'(x) = -(A + \kappa c'(x_L)) + \kappa c'(x) ,$$

and observe that g is strictly convex in x . Evidently, $g'(0) < 0$, so f is either minimized at $x^\circ = x^\circ(x_L, x_H)$, implicitly defined as $g'(x^\circ) = 0$ (if such an $x \leq 1$ exists) or $x = 1$. Define $x^\dagger := \min\{x^\circ, 1\}$. Thus, $\gamma = \beta = 0$ is optimal if and only if $g(x^\dagger) \geq 0$. Note that there is a knife-edge case where $v_0/\kappa = \eta(x_L, x_H)$, $x^\dagger = \mu$, and $\beta = \gamma = 0$ (and the first-best is attained). This is the only way for all three constraints to bind.

Can we have one of the non-negativity constraints bind, $\gamma = 0$, say; and the other constraints all be slack, i.e., $\beta > 0$ and $f(x) > v_0 - \kappa c(x)$ for all $x \in [0, 1]$? No: otherwise the principal could decrease β by a sufficiently small $\varepsilon > 0$, strictly increasing her payoff and still leaving Constraint $IR-v_0$ satisfied. This yields (ii)a and (ii)b of the result. ■

A.10 Proposition 6.6 Proof

Proof. Simply rearrange the inequality $f(\mu) \geq v_0$ to get

$$\gamma \geq \frac{v_0 - \kappa((1-\mu)(x_H c'(x_H) - c(x_H)) - \mu((1-x_L)c'(x_H) + c(x_L)))}{1-\mu} - \frac{\mu}{1-\mu}\beta ,$$

then set $\beta = 0$ and solve for when the right-hand side of this inequality is positive. ■