# Buying Opinions

Mark Whitmeyer[*] Kun Zhang[†]

February 11, 2022

**Abstract**

A principal hires an agent to acquire a distribution over *unverifiable* posteriors before reporting to the principal, who can contract on the realized state. An agent's optimal learning and truthful disclosure completely specify the marginal incentives the principal must provide, which radically simplifies the principal's problem. When the agent i. is risk neutral, and iia. has a sufficiently high outside option, or iib. can face sufficiently large penalties, the principal can attain the first-best outcome. We also explore in detail the general problem of cheaply implementing distributions over posteriors with limited liability constraints and a risk-averse agent.

***Keywords:*** Moral hazard, Principal–agent model, Information acquisition, Rational Inattention

***JEL Classifications:*** D81; D82; D83; D86

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

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

We thank Vasudha Jain, Andreas Kleiner, Joseph Whitmeyer and Thomas Wiseman for their comments.

# 1 Introduction

There are many instances in which decision makers buy advice. Investors pay for stock picks, politicians employ advisors, bettors at the race track ask for winners, and executives in firms appoint subordinates to suggest policies and actions. In some situations this advice can be backed up with hard, verifiable, evidence; whereas in others advice is merely cheap talk and is supported only by the advisor's incentives.

We analyze a contracting problem in which a principal hires an agent to acquire information. We focus on the situation in which the evidence an agent acquires is unverifiable and cannot be credibly disclosed (or contracted upon). That is, the agent merely provides advice to the principal. The principal may; however, condition the agent's remuneration on both the realized state and the agent's report. Equivalently, the principal offers the agent a menu of state dependent lotteries from which the agent selects once she has acquired information.

Information acquisition is costly for the agent–after observing the contract, she chooses what information to obtain, before reporting her findings to the principal. The agent has significant freedom in her learning: she may choose any distribution over posterior beliefs whose mean is the prior. In doing so she incurs a posterior separable cost (Caplin et al. (Forthcoming)); i.e., one that can be expressed as the expectation of a convex function of the beliefs she obtains. We also assume that the agent may exit the relationship, should she so choose, at any point. The agent has an outside option of value $v_0$, which she 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 her employer and seeking employment elsewhere. Naturally, our attainability of the first best result still goes through even if the agent could only leave for her outside option before learning–there are fewer constraints for the principal, making the first best even easier to attain.

There are clear parallels between this setting and the canonical model of moral hazard with hidden effort. Like the basic moral hazard setup, the 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 induce? In our first result, Lemma 3.2, we note that the principal is *unconstrained* in the distribution over posteriors the agent acquires: for any distribution over posteriors the principal would optimally choose in her decision problem should she control the information acquisition directly, she can always write a contract such that the agent acquires that distribution (and reports her findings honestly). Keep in mind; however, that due to the agency problem the *cost* to the principal of obtaining such a distribution may be greater, though we note in the lemma that any such distribution acquired via an agent still costs only a finite amount.

We observe that an analog of the standard first-order approach can be used to characterize the cheapest way to implement a distribution over posteriors. Given the contract written by the principal, the agent faces a decision problem. The characterization of the agent's optimal learning is well known: it corresponds to the concavification of the agent's value function. This requirement (optimality of the agent's learning) produces a number of conditions that the contract must satisfy, which are the analog of the standard incentive compatibility conditions. Furthermore, we discover that these conditions have a particular structure that allows us to radically simplify the principal's problem of implementing a posterior distribution. 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 (which are not control variables for this problem). Proposition 3.4 summarizes this finding: the principal's optimization problem reduces $m$-fold–from one with $n \times m$ variables, where $m$ is the number of posteriors and $n$ is the number of states, to one with just $n$ variables. This dramatic simplification sets the stage for the remainder of our results.

This result carries with it significant economic content. Not only do the marginal incentives completely pin down the agent's optimal learning, but the *converse* is also true: the agent's optimal learning specifies the marginal incentives. Accordingly, all the principal needs to determine is the state-dependent payoffs for a benchmark message.

Our direct approach described above ensures that a collection of incentive compatibility constraints are satisfied by construction–the agent cannot benefit by learning differently before sending a message, and/or by misrepresenting her findings. However, even absent limited liability, an additional constraint remains, engendered by the agent's outside option. The contract must be such that the agent does not want to deviate by taking her outside option with positive probability both before and after learning. In particular, we need to rule out double deviations, in which the agent both learns differently *and* takes her outside option at some of the resulting beliefs. All in all, the principal faces an $n$ variable optimization problem subject to this constraint (plus limited liability, if so required).

As in the standard moral hazard problem, there is a natural benchmark in our model: the first-best problem in which effort (in our case, learning) is observable and contractible.[1] Perhaps surprisingly, 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. That is, Proposition 4.1 argues that efficient learning is feasible. We derive this result by construction: a principal can always write a contract so that the resulting concavifying hyperplane of the agent's payoff function is tangent to the payoff from taking the outside option at the prior. Importantly, this is not a "shoot the agent contract" (Mirrlees (1999))–the principal can *attain* (not approximate) the first best even if the agent's utility function is bounded. In fact, as long as the agent's outside option is sufficiently high, efficient implementation is feasible even when negative transfers are forbidden.

On the other hand, 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 with limited liability, optimal incentives take simple forms in a number of special cases. In Proposition 5.1, we provide a full characterization of the optimal contract when the agent's outside option is sufficiently small. There, it is only the non-negativity constraint that binds, which allows us to pin down the optimal contract for any desired distribution over posteriors. We also fully characterize optimal imple-

[1] If the agent is risk neutral, the first-best problem is (strategically) equivalent to one in which the principal can control the information acquisition (and incur the costs) herself.

mentation with an arbitrary outside option in the binary-state case when the agent is risk neutral. Of special interest is our finding that less informative distributions (in the Blackwell sense) are easier to implement (Corollary 5.3).

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

## 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 this one. Rappoport and Somma (2017) also study contracting for flexible information acquisition; while the true state cannot be contracted upon, they assume that the posterior generated by the agent's choice of distribution is verifiable and contractible. Among other things, 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.[2] Zermeño (2011) and Clark and Reggiani (2021) study contracting environments in which both information acquisition and decision making are delegated to the agent. In Zermeño (2011), the action is assumed to be observable and contractible; and his main focus is the interaction between the variables on which the transfer schemes can depend (e.g., the true state and the principal's payoff from the action) and whether contracts specify transfer scheme menus. Clark and Reggiani (2021) assume that the agent acquires information flexibly subject to an upper bound on the cost and that both action and the state are contractible. Their main result is that any Pareto

[2] Bizzotto et al. (2020) consider a similar problem. However, they do not require the message space to be the set of posteriors, and they only allow the agent to deviate to a "default" distribution, instead of any Bayes-plausible distribution.

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 on the set of distributions available to the agent as well as their costs. Similar to our work, the agent chooses a distribution and sends a message to the principal, and both the message and the true state are contractible. The principal evaluates each possible contract by its worst-case guarantee. In Häfner and Taylor (Forthcoming) 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 realized posterior is unverifiable. Their focus is on finding the optimal contract, which can depend on the report and the outcome of the project, in order to motivate the agent to conduct the experiment and report truthfully.[3] Chade and Kovrijnykh (2016) study a dynamic model of contracting for information acquisition in a two state-two (fixed) signals environment. The more effort the agent exerts, the more informative the signal she acquires. They assume that the realized signals are contractible, but the true state is not. Azrieli (2021) builds on the static version of Chade and Kovrijnykh (2016)'s model–the key difference is that the signals are not observable and that multiple agents acquire information for the principal. He shows that the least costly contract utilizes cross-checking: the agents are paid only when they all report the same signal.

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. (Forthcoming). Maćkowiak et al. (Forthcoming) provides an excellent review of this literature that covers both theory and applications.

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), and Grossman and Hart (1983)), the agent is impelled to exert costly effort that yields some output;

[3] 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.

whereas in ours, she must be coerced into choosing a much more complicated object (a particular probability distribution) then reporting honestly. That being said, there are some interesting analogies between some of our results and classical insights in the moral hazard problem, which we discuss as we encounter them.

## 2 The Model

The principal 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 $\Theta$ is a finite set; $|\Theta| = n < \infty$. The principal's utility from taking action $a$ in state $\theta$ is given by $u(a, \theta)$, where $u$ is continuous in $a$ (and therefore bounded). Let $\mu \in \Delta(\Theta)$ denote the principal's prior belief about the state; without loss of generality, assume $\mu$ has full support.

The principal cannot acquire information herself but instead must rely on the assistance of an agent, who shares the same prior and 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. (Forthcoming). That is, the cost of acquiring $F$ is

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

where $\kappa > 0$ is a scaling parameter, $c \colon \Delta(\Theta) \to \mathbb{R}_+$ 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, which is a special case of the Tsallis-entropy-based cost function (see Tsallis (1988), who introduces this form of entropy, and Lipnowski et al. (2021) for an economic application).[4]

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

[4]Further discussion of this cost function in economic contexts may be found in Bloedel and Segal (2020) and Caplin et al. (Forthcoming).

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 $t\colon M\times\Theta\to\mathbb{R}$ ($t\colon M\times\Theta\to\mathbb{R}_+$ 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 her utility from the transfer and the cost of acquiring information, and she values the transfer according to a continuously differentiable and strictly increasing function (which is therefore invertible) $v$, 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 there are two chances that she can leave with her outside option: she can choose not to accept the contract, or walk away after acquiring information by reporting nothing.[5]

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 $F$;

(iii) A posterior $\mathbf{x}\in\Delta(\Theta)$ is drawn from $F$, which is privately observed by the agent;

(iv) The agent chooses whether to report, and if she reports, she sends a message $m\in M$;

(v) The principal takes an action $a\in A$;

(vi) The true state $\theta\in\Theta$ realizes;

(vii) Payoffs accrue: the principal gets $u(a,\theta)-v^{-1}(t(m,\theta))$, and the agent gets $t(m,\theta)-c(F)$.

# 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 $V(\mathbf{x})$. $V(\cdot)$ is piecewise affine and convex

[5] Alternatively, the agent may report "I did not get anything."

(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 $\Delta(\Theta)$. If the principal controlled the information acquisition, she would solve

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

Let $\mathcal{X}$ denote the set of optimal distributions over posteriors that have at most $n$ elements in support. That is, an element of $\mathcal{X}$, is a collection of points, $X = \left(\mathbf{x}_1^*, \ldots, \mathbf{x}_k^*\right)$, $k \le n$, in the simplex $\Delta(\Theta)$ such that $\mu$ lies in the convex hull of these points. The Fenchel-Bunt extension[6] of Carathéodory's theorem guarantees that $\mathcal{X}$ is nonempty.

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 $t\colon \Delta\Delta(\Theta) \to \mathbb{R}_+$. 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 her indifferent between learning and walking away with her 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

To get the same distribution as if she were able to generate it by herself, the principal needs to guarantee that the agent chooses the right distribution, and every message she sends must represent a unique posterior realization. Without loss of generality, for every distribution that the principal would like to implement, she sets the message space to be the support of the distribution.

Following Caplin et al. (Forthcoming), we define a *decision problem* $(\mu, D, w)$ as the choice over a compact set of actions $D$ given the prior $\mu$ over states in $\Theta$, and $w : D\times\Theta \to \mathbb{R}$ is the decision maker's utility function. Given a decision problem $(\mu, D)$, the decision maker chooses a Bayes-plausible distribution over posteriors $G$ and an action strategy $\sigma : \mathrm{supp}(G) \to \Delta(D)$. A contract $(M, t)$ induces a decision problem $(\mu, M, t)$ of the agent;

[6] Reference, e.g., Theorem 1.3.7 of Hiriart-Urruty and Lemaréchal (2001).

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 $\left(F, \{\delta_{\mathbf{x}}\}_{\mathbf{x}\in\text{supp}(F)}\right)$.[7] 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* $N(\mathbf{x} \mid d)$ as the expected utility of the message $d$ net of the cost of $\mathbf{x}$:

$$N(\mathbf{x} \mid d) = \sum_{i=1}^{n-1} x_i t(d,\theta_i) + \left(1 - \sum_{i=1}^{n-1} x_i\right) t(d,\theta_n) - \kappa c(\mathbf{x}) \ ,$$

where $x_i$ is the $i$-th entry of $\mathbf{x}$. The agent chooses a distribution over posteriors $G$ to maximize her value function $W(\mathbf{x}) = \max_{d\in M} N(\mathbf{x} \mid d)$. It is well-known that the agent's optimal choice is determined by concavifying the value function: let $\mathcal{H}$ denote the hyperplane tangent to the hypograph of $W$ at the support points of $F$. We can identify this supporting hyperplane $\mathcal{H}$ by an affine function $f_{\mathcal{H}}(\mathbf{x}) : \Delta(\Theta) \to \mathbb{R}$. 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 $\mathbf{x}_j$, $W\left(\mathbf{x}_j\right) = N\left(\mathbf{x}_j \mid \mathbf{x}_j\right)$; 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 her 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 her outside option)

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

We also need to ensure that the agent's value from accepting the contract exceeds her outside option, that is, $f_{\mathcal{H}}(\mu) \geq v_0$; but since $c(\mu) = 0$, this is already included in Constraint $IC$. Thus,

**Lemma 3.1.** *A distribution $F$ can be implemented by a contract $(M,t)$ if and only if*

[7] $\delta_{\mathbf{x}}$ denotes the degenerate distribution at $\mathbf{x}$.

*(i)* $\mathrm{supp}(F) = P_{(M,t)}$*; and*

*(ii) Constraint* $IC$ *holds; and*

*(iii) 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: for every implementable distribution $F$, we solve the principal's cost minimization problem:

$$\min_{(M,t)} \sum_{j=1}^{m} p_j \left[ \sum_{j=1}^{n-1} x_i^k v^{-1}\left(t\left(\mathbf{x}_j, \theta_k\right)\right) + \left(\left(1 - \sum_{k=1}^{n-1} x_j^k\right) v^{-1}\left(t\left(\mathbf{x}_j, \theta_n\right)\right)\right)\right],$$

where $p_j$ is the (unconditional) probability that posterior $\mathbf{x}_j$ realizes, subject to (i), (ii), and (iii) in Lemma 3.1; denote its value by $\Gamma(F)$. Then 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 problems, we focus on the first step.

## 3.3 Preliminary Results

We finish this section by establishing a few preliminary results. We begin by arguing that any distribution over posteriors with support on $n$ or fewer points can be induced by some contract.

**Lemma 3.2.** *Let F be a distribution over posteriors with* $|\mathrm{supp}(F)| \leq n$ *and* $\mathrm{supp}(F) \subseteq \mathrm{int}\,\Delta(\Theta)$*, then there exists a contract* $(M,t)$ *that implements* $F$*, and the expected cost to the principal is finite.*

The proof of Lemma 3.2, 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 $\Delta(\Theta)$, we can (i) construct a contract with bounded transfer such that the agent finds it optimal to first acquire $F$, and then report the realized posterior truthfully; (ii) find a finite constant such that by adding it to the transfer, Constraint $IC$ holds. Therefore, every such distribution can be implemented, and the expected transfer is finite.

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. In fact, it is weakly less costly for the principal to randomize first rather than implement $F$ directly. Thus, it is without loss of generality for the principal to offer a contract that results in a distribution over posteriors chosen by the agent with support on at most $n$ points.

**Corollary 3.3.** *(i) Every $F \in \mathcal{F}(\mu)$ with $\operatorname{supp}(F) \subseteq \operatorname{int}\Delta(\Theta)$ can be induced at a finite cost.*

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

By Corollary 3.3 (ii), we can restrict our attention to distributions over posteriors supported on $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_m\}$, where $n$ is the number of states, and $m \le n$. Suppose the principal would like to implement a distribution $F$ using some contract $(M, t)$. By part (i) of Lemma 3.1, a necessary condition for implementation is that $\operatorname{supp}(F) = P_{(M,t)}$; this condition holds if and only if the contract is such that the following $m$ expressions

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

define the same hyperplane, where $t_j^k := t\left(\mathbf{x}_j, \theta_k\right)$, $x_j^i$ 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 $k = 1, \ldots, n-1$ and $i, j = 1, \ldots m$

$$t_i^k - t_i^n - \kappa c_k(\mathbf{x}_i) = t_j^k - t_j^n - \kappa c_k\left(\mathbf{x}_j\right) ,$$

and

$$t_i^n = t_j^n + \Xi_{ij} ,$$

where $\Xi_{ij}$ is some function of the primitives (but not directly of the $t$s).

For each state $k = 1,\ldots,n$, define $X^k(i,j) := t_i^k - t_j^k$ $(i,j = 1,\ldots,m)$. Then for any one of the first equations, we have

$$t_i^k - t_j^k - \kappa c_k(\mathbf{x}_i) = t_i^n - t_j^n - \kappa c_k\left(\mathbf{x}_j\right) ,$$

for all $k = 1,\ldots,n-1$ and $i,j = 1,\ldots m$, so

$$X^k(i,j) - \kappa c_k(\mathbf{x}_i) = X^n(i,j) - \kappa c_k\left(\mathbf{x}_j\right) .$$

The second set of equations becomes $X^n(i,j) = \Xi_{ij}$; plugging this equation into the one above, we get

$$X^k(i,j) = \kappa c_k(\mathbf{x}_i) - \kappa c_k\left(\mathbf{x}_j\right) + \Xi_{ij} .$$

Therefore, in every state $k = 1,\ldots,n$, and every $j = 1,\ldots,m$, we can write $t_j^k = t_m^k + X^k(j,m)$. We have established

**Proposition 3.4.** *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. The principal fixes a benchmark message $m$, then chooses $\left(t_m^k\right)_m^k$, $k = 1,\ldots,n$; the payoff to the agent from sending message $m$ in each state $k$.*

That is, optimality allows us to reduce transfers to a single variable for each state.

# 4 Main Results I. No Limited Liability

## 4.1 Risk-Neutral Agent

Recall that $f_{\mathcal{H}}$ is the function that identifies the concavifying hyperplane $\mathcal{H}$, the agent's value from acquiring information for the principal can be written as $f_{\mathcal{H}}(\mu)$. Then efficient implementation requires $f_{\mathcal{H}}(\mu) = v_0$, which implies that Constraint $IC$ must bind at $\mathbf{x} = \mu$. It is not difficult to see that, in this case, Constraint $IC$ can be reduced to

$$t_m^k - t_m^n - \kappa c_k(\mathbf{x}_m) = -\kappa c_k(\mu) \quad \text{for all} \quad k = 1,\ldots,n-1 . \qquad (IC-R)$$

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

**Proposition 4.1.** *If the agent is risk neutral and not protected by limited liability, every distribution $F$ with $\operatorname{supp}(F) \subseteq \operatorname{int}\Delta(\Theta)$ can be implemented efficiently.*

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 her outside option. Figure 1 illustrates this construction. There, the agent's induced value function, $W$, is depicted in black and its concavification in orange. The gross payoffs (as a function of posterior $x$) to agent from the two messages are the blue and purple lines. Finally, the agent's net payoff from taking the outside option $v_0$ is the red curve. The state dependent payoffs for each message are chosen in such a way that the agent is willing to choose the desired distribution (and report truthfully) and cannot gain by taking her outside option at any point.

It is instructive to compare Proposition 4.1 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 she 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.

If we require limited liability, for some distributions and outside option values, efficient implementation cannot be achieved. In particular, because posteriors are unverifiable and hence noncontractible, Rappoport and Somma's construction does not work. A simple intuition is that, for unverifiable posteriors, we need to guarantee that for every possible posterior realization, on or off path, that the agent obtains, she prefers reporting truthfully to walking away, which constraint is absent in the verifiable posteriors world. This is correct yet incomplete: in fact, for $v_0$ close enough to 0, even if we do not impose

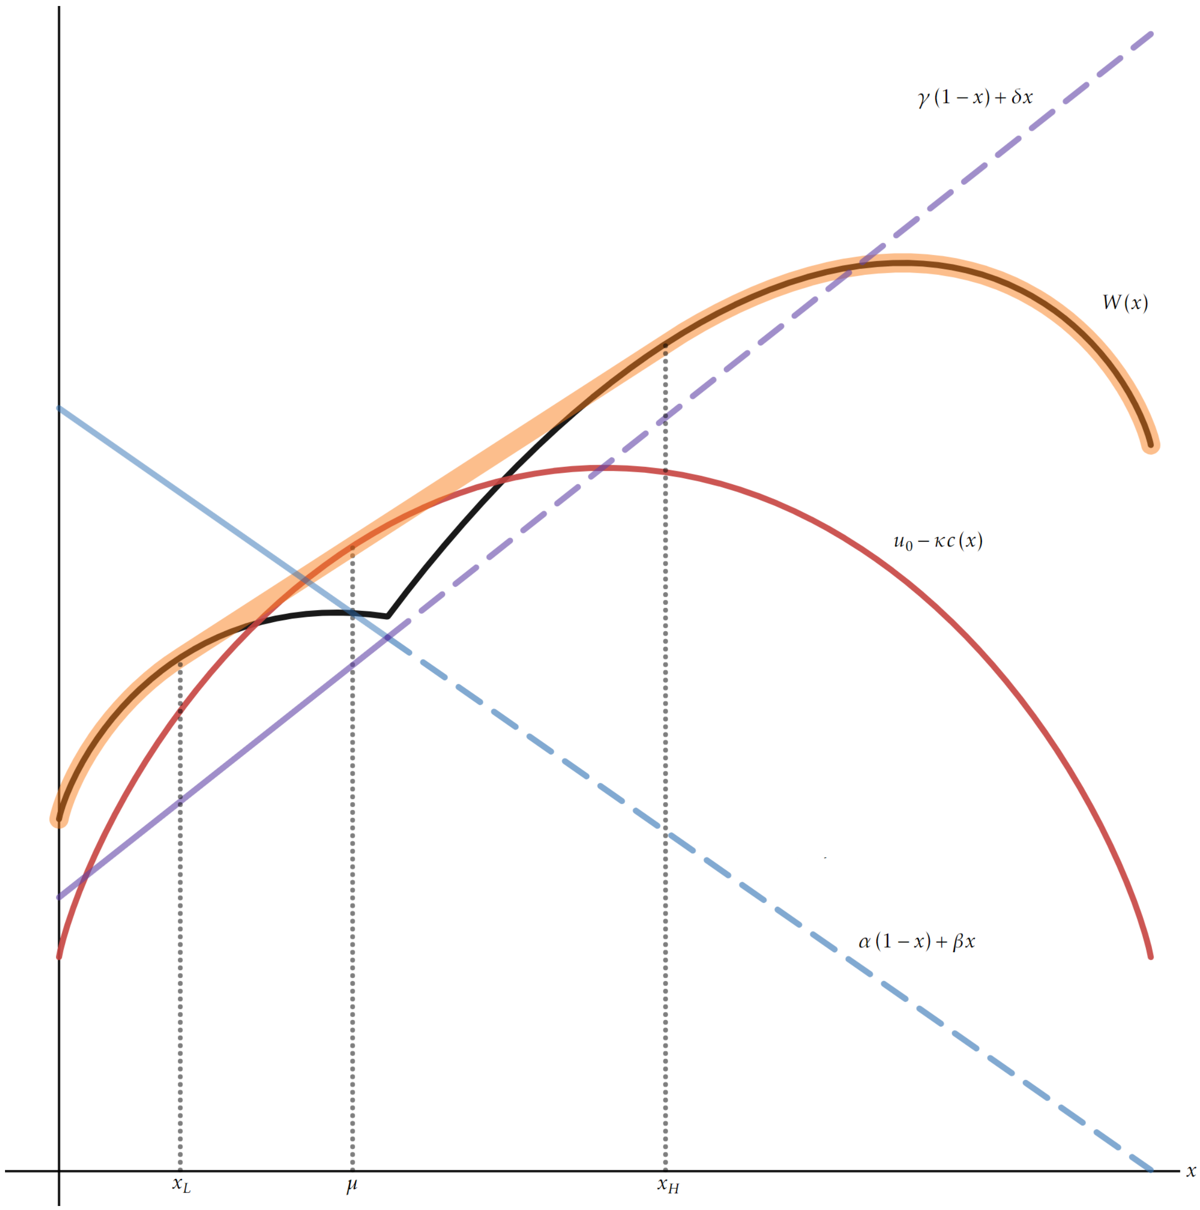


Figure 1: Efficient implementation of $x_L = 1/9$, $x_H = 5/9$ when $\mu = 1/(1+e)$, $\kappa = 1$, and $v_0 = \log\{9/(1+e)\}$, and with entropy cost. This contract satisfies the limited liability constraints–as stated in Proposition 5.2, the specified ratio $v_0/\kappa$ is the minimum such ratio such that efficient implementation is feasible under limited liability.

Constraint $IC-R$, efficient implementation is still not possible for any non-degenerate distribution. This is because, to induce the agent to gather information, the transfers must be "rewarding" when the agent "gets the state right" and "punishing" when she is wrong. The gap between the two scenarios 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 must be negative.

## 4.2 Risk-Averse Agent

When the agent is risk averse, but unprotected by limited liability, characterizing the optimal contract is more involved. The following remark is obvious: the principal's payoff is strictly decreasing in each of the $n$ control variables $\left(t_m^k\right)_{k=1}^n$ and so the principal wants to set each one as low as possible. Unencumbered by limited liability, the lone constraint is $IC$, which necessarily binds (since otherwise, the principal could reduce the control variables). Thus,

**Remark 4.2.** 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 condition to substitute in for each $t_m^k$.[8] This yields the principal an objective that is a function of $\mathbf{x}^*$, with the lone constraint that $\mathbf{x}^*$ lies in $\Delta(\Theta)$. Unless $\mathbf{x}^* = \mu$, which does not hold in general, the principal does not attain the first best, and the agent obtains positive rents.

This finding contrasts interestingly with the standard finding in the classic moral hazard problem that the agent is left with zero rents in both the first- and second-best worlds. Not so here, since the agent's risk aversion alters the optimal way for the principal to prevent deviations to the outside option. Here the principal optimally trades off between risk sharing and efficiency: when a contract that makes the agent break even entails too much risk, by moving $\mathbf{x}^*$ away from $\mu$, the risk in the contract is mitigated; then although the

[8] More precisely, we have $t_m^k - t_m^n - \kappa c_k(\mathbf{x}_m) = -\kappa c_k(\mathbf{x}^*)$ for all $k = 1, \ldots, n-1$.

agent receives strictly positive rent, implementing the new contract can be much cheaper to the principal.

This is similar to the trade-off studied in Proposition 5 in Rappoport and Somma (2017). Both this problem and theirs require some tangency conditions to hold at optimum. Beyond this superficial resemblance, the exact conditions differ: 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 this tangency condition.

### 4.2.1 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 $c(x) = x\log x + (1-x)\log(1-x) - \mu\log\mu - (1-\mu)\log(1-\mu)$. After some algebra, the principal's problem reduces to

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

This objective is strictly convex and equals $+\infty$ as $x \downarrow 0$ and $x \uparrow 1$. Consequently, there is a unique interior minimizer pinned down by the first-order condition

$$(\mu - x_L)\left(\frac{(1-x_H)^{\kappa+1}}{(1-x)^2} - \frac{x_H^{\kappa+1}}{x^2}\right) + (x_H - \mu)\left(\frac{(1-x_L)^{\kappa+1}}{(x-1)^2} - \frac{x_L^{\kappa+1}}{x^2}\right) = 0 .$$

It can be checked that when $\kappa = 1$, $x^* = \mu$ (and so the agent obtains zero surplus) if and only if $\mu = 1/2$. When $\kappa = 2$, $x^* = \mu = 1/2$ if and only if $x_L = 1 - x_H$.

# 5 Main Results II. Limited Liability

In this section, we assume that the agent is protected by limited liability.

## 5.1 Low Outside Option

In this subsection we solve for the optimal incentives when the agent's value for her outside option is sufficiently small; we set $v_0 = 0$ for simplicity.[9] For expository ease, we start with the two state case, and then we show that most of our results generalize when there are more than two states.

### 5.1.1 Two States

When there are just two states, $\Theta = \{\theta_1, \theta_2\}$; by Corollary 3.3 (ii), we can identify a distribution by its support $\{x_L, x_H\}$. 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:

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

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

$$\kappa\left(c'(x_H) - c'(x_L)\right) = A + B ,$$

and

$$A + \kappa\left(c(x_H) - c(x_L)\right) = \kappa\left(c'(x_H)x_H - c'(x_L)x_L\right) ,$$

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

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

The principal's objective function is

$$p\Big(v^{-1}(\gamma)(1-x_H) + v^{-1}(B+\beta)x_H\Big) + (1-p)\Big(v^{-1}(A+\gamma)(1-x_L) + v^{-1}(\beta)x_L\Big) ,$$

where $p = (\mu - x_L)/(x_H - x_L)$ is the (unconditional) probability that posterior $x_H$ realizes.

It is easy to see that the principal's objective function is strictly decreasing in both $\beta$ and $\gamma$. Moreover, the limited liability constraint guarantees that the agent never wants to

[9] In the next subsection (Section 5.2) we assume the agent is risk neutral and allow for an arbitrary outside option.

deviate to take her outside option: the agent's (on path) payoff gross of information costs is strictly above 0 and hence the concavification of her objective lies strictly above $-\kappa c(x)$ for all $x$. Therefore, it is optimal for the principal to set $\gamma = \beta = 0$; and so $\alpha = A$ and $\delta = B$.

Given this, we can easily back out the principal's cost of implementing any binary posterior $\{x_L, x_H\}$. For instance, if the agent is risk neutral and her cost of information acquisition is the quadratic cost,[10]

$$\alpha = \kappa\left(x_H^2 - x_L^2\right) \quad \text{and} \quad \delta = \kappa\left(x_L\left(x_L - 2\right) - x_H\left(x_H - 2\right)\right) .$$

After some algebra, the principal's cost of implementing a binary posterior $\{x_L, x_H\}$ is

$$\kappa\left(x_H - x_L\right)\left(x_H\left(1-\mu\right) + \mu\left(1 - x_L\right)\right) .$$

When the agent's cost of information is entropy,[11]

$$\alpha = \kappa \log \frac{1 - x_L}{1 - x_H} \quad \text{and} \quad \delta = \kappa \log \frac{x_H}{x_L} .$$

### 5.1.2 More Than Two States

When there are $n > 2$ states, consider a distribution $F$ whose support is $\{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_m\}$, where $m \leq n$. Recall from the discussion precedes Proposition 3.4 that for each state $k = 1, \ldots, n$ and posterior $\mathbf{x}_j$, optimal learning requires

$$t_j^k = t_m^k + X^k\left(j, m\right) .$$

Importantly, note that unlike the two state case it is possible that $X\left(j, m\right) < 0$. Define $N(k) = \left\{j : X^k(j,m) < 0\right\}$. If $N(k) = \varnothing$, let $j^*(k) = m$; otherwise let $j^*(k)$ be an arbitrary selection of $\arg\min_{j \in N(k)} X^k(j,m)$. By setting $t^k_{j^*(k)} = 0$, the agent's honesty is not affected, and it is not hard to check that the limited liability constraints are satisfied. For every $i \neq j^*(k)$, we have

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

for each $k = 1, \ldots, n-1$; and $t_i^n = \Xi_{i j^*(n)}$ for $i \neq j^*(n)$. Since the argument above works for an arbitrary state $k$, we have thus completely identified the optimal transfers:

[10]That is, $c(x) = (x-\mu)^2$.

[11]That is, $c(x) = x \log x + (1-x)\log(1-x) - \mu \log \mu - (1-\mu)\log(1-\mu)$.

**Proposition 5.1.** *Suppose $v_0 = 0$, and the agent is protected by limited liability. Then for each state $k = 1, \ldots, n$, there exists $j^*(k)$ such that $t\left(\mathbf{x}_{j^*(k)}, \theta_k\right) = 0$, and all other transfers are determined by optimal learning.*

Proposition 5.1 is intuitive: for sufficiently small outside option, Constraint $IC$ 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.

## 5.2 Risk Neutral Agent

Now, we dispense with the assumption that the outside option is small–$v_0$ can take any value.

### 5.2.1 Two States

Again, we begin with just two states. When the agent is risk neutral ($v(\cdot) = \cdot$), the principal's objective function can be simplified: she chooses $\gamma$ and $\beta$ in order to maximize

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

subject to limited liability: $\beta, \gamma \geq 0$, and

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

By construction the agent cannot deviate profitably by learning differently *and* reporting to the principal. Constraint $IC$-$v_0$ ensures that the agent cannot deviate profitably by learning differently and taking her outside option.

Our first result characterizes the distributions over posteriors that a principal can implement efficiently; *viz.*, at the first-best cost. Defining

$$\eta_1(x_H) := -\mu c'(\mu) - c(x_H) + c'(x_H) x_H, \quad \text{and} \quad \eta_2(x_L) := (1-\mu) c'(\mu) - c(x_L) - (1-x_L) c'(x_L) \ ,$$

and $\eta(x_L, x_H) := \max\{\eta_1(x_H), \eta_2(x_L)\}$, we have

**Proposition 5.2.** *The principal can implement $\{x_L, x_H\}$ efficiently if and only if $v_0/\kappa \geq \eta(x_L, x_H)$.*

Recall from the discussion in Section 4.1 that, when the agent is risk neutral and her outside option is zero, we can always find a transfer scheme which implements any distribution efficiently, just that the transfer is negative for some posterior realization 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.[12] Consequently, for $v_0$ sufficiently high, all transfers become nonnegative, and hence limited liability is satisfied.

The left-hand side of Proposition 5.2's necessary and sufficient condition is strictly increasing in the outside option $v_0$ and strictly decreasing in the cost of information $\kappa$. Moreover, it is easy to calculate that $\eta$ is decreasing in $x_L$ and increasing in $x_H$, strictly so if $\eta = \eta_2$ or $\eta = \eta_1$, respectively. This suggests the following corollary:

**Corollary 5.3.** *(i) For any pair of posteriors $\{x_L, x_H\}$ with $0 < x_L \leq \mu \leq x_H < 1$, if $v_0/\kappa$ is sufficiently large, $\{x_L, x_H\}$ can be implemented efficiently.*[13]

*(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 5.3 has a flavor of that classical result: for any outside option level, any distribution that is Blackwell less informative than some threshold distribution can be implemented efficiently. In particular, when $v_0 = 0$, the only distribution that can be implemented effi-

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

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

ciently corresponds to the uninformative experiment. For $v_0 > 0$, however, a continuum of nontrivial distributions can be implemented efficiently.

When the first-best implementation of $\{x_L, x_H\}$ is infeasible, there are three other possibilities, listed in our next proposition.

**Proposition 5.4.** *One of the following must occur at the optimum. Either*

*(i)* $\{x_L, x_H\}$ *can be implemented efficiently (and Constraint* $IC$-$v_0$ *binds); or*

*(ii)* $\{x_L, x_H\}$ *cannot be implemented efficiently; and either*

  *(a) Constraint* $IC$-$v_0$ *binds and* $\beta = 0$*; or*

  *(b) Constraint* $IC$-$v_0$ *binds and* $\gamma = 0$*; or*

  *(c) Constraint* $IC$-$v_0$ *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. Curiously, note that the boundaries of the four regions of posterior pairs are line segments. This is not true in general but is special to the entropy cost.

One last comment about Proposition 5.4: if $\{x_L, x_H\}$ is in the region corresponding to (ii)c then $\gamma = \beta = 0$ is optimal even when the agent is risk averse. That is, this portion of the result does not depend at all on the risk preferences of the agent but is completely determined by incentive compatibility (in which transfers are expressed as utils). Consequently the necessary and sufficient conditions described in the proposition's proof hold even absent risk neutrality. If $\beta = \gamma = 0$ is not feasible, on the other hand, then risk aversion (in the presence of limited liability) complicates things. In particular, it may now be optimal to give the agent rents even if it is possible to write a contract eliminating them.

### 5.2.2 More Than Two States

The discussion preceding Proposition 5.1 argues that for each state $k$ we can find a lottery $j^*(k)$ that delivers the lowest payment; and by Proposition 3.4, to pin down the transfer scheme, it suffices to determine $t^k_{j^*(k)}$ for each state $k$. Thus, there are $n$ unknowns. Moreover, we have $n$ equations: efficient implementation is equivalent to $f_{\mathcal{H}}(\mu) = v_0$, and the

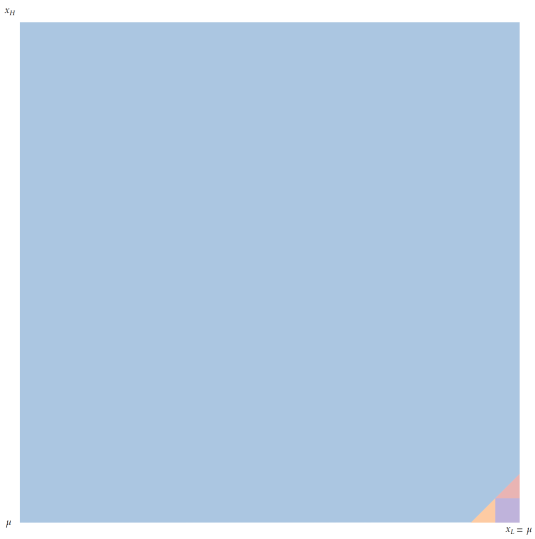


(a) $v_0/\kappa = .05$.

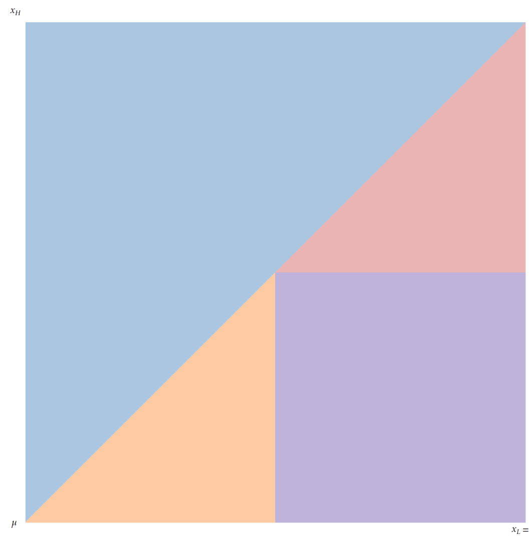


(b) $v_0/\kappa = \log 1/(1-\mu)$.

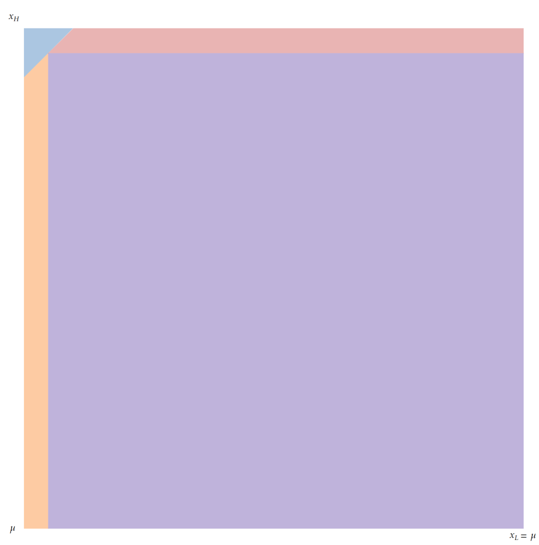


(c) $v_0/\kappa = 3$.

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 $\beta \geq 0$.

other $n-1$ equations are imposed by Constraint $IC-R$:

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

where $k = 1, \ldots, n-1$ indicates the state, and $j$, which indexes the posterior, is arbitrary. This is because optimality requires

$$t_i^k - t_i^n - \kappa c_k(\mathbf{x}_i) = t_j^k - t_j^n - \kappa c_k\left(\mathbf{x}_j\right)$$

for all $i, j$ and each $k = 1, \ldots, n-1$. Then the distribution can be efficiently implemented if and only if $t_{j^*(k)}^k \geq 0$ for each $k$; consequently, Proposition 5.2, Corollary 5.3 (i), and Proposition 5.4 extend naturally to more than two states.

# 6 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, we 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.[14] If only the unverifiable message that the agent sends can be contracted upon; however, the agent cannot be coerced into gathering any nontrivial information: she never learns anything and sends the message that yields the highest reward.

[14]This 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.

**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 acquisition to an agent who cannot make verifiable reports. To allow for the agent to have intrinsic motivation, we can 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 to assume that the agent has intrinsic preferences over the action to be taken by the principal; we leave this extension to future research.

**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 posterior. For example, the principal might seek to influence the choice of a third party decision maker à la the literature on Bayesian persuasion and information design.

## References


Yaron Azrieli. Monitoring experts. *Theoretical Economics*, 16(4):1313–1350, 2021.

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

Alexander W Bloedel and Ilya Segal. Persuading a rationally inattentive agent. *Mimeo*, 2020.

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, Forthcoming.

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.

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*, Forthcoming.

Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Fundamentals of Convex Analysis*. Springer, 2001.

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

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *The American Economic Review*, 101(6):2590–2615, 2011.

Richard A Lambert. Executive Effort and Selection of Risky Projects. *The RAND Journal of Economics*, 17(1):77–88, 1986.

Elliot Lipnowski, Laurent Mathevet, and Dong Wei. Optimal attention management: A tractable framework. *Mimeo*, 2021.

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.

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.

Constantino Tsallis. Possible generalization of boltzmann-gibbs statistics. *Journal of Statistical Physics*, 1988.

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

# A Omitted Proofs

## A.1 Lemma 3.2 Proof

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

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

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

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 $\tau$ is a constant that scales the transfers. Now the agent is facing a decision problem $(\mu, M, t)$. Let $G$ be a distribution over posteriors, and let $\sigma : \Delta(\Theta) \to \Delta(M)$ denote a reporting strategy. Then, the agent's value of $(G, \sigma)$ is given by

$$\Upsilon(G, \sigma) = \sum_{\mathbf{x} \in \operatorname{supp}(G)} \sum_{d \in M} G(\mathbf{x}) \sigma(d \mid \mathbf{x}) N(\mathbf{x} \mid d) \ .$$

We claim that $(F, \sigma^*)$, where $\sigma^*(\cdot \mid \mathbf{x}_j) = \delta_{\mathbf{x}_j}$ is an optimal strategy for the agent. By Lemma 1 in Caplin et al. (Forthcoming), it suffices to show that, for every $\mathbf{x}_j$, $j = 1, \ldots, m$, there exists a $n-1$ dimensional vector $\lambda = (\lambda_1, \ldots, \lambda_m)$ such that

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

for all $\mathbf{x} \in \Delta(\Theta)$ and $d \in M$. We set $\lambda$ to be the zero vector, so the above inequality reduces to $N(\mathbf{x} \mid d) \leq N\left(\mathbf{x}_j \mid \mathbf{x}_j\right)$. We first show that for any fixed $d \in M$, $N(\mathbf{x} \mid d) \leq N\left(\mathbf{x}_j \mid d\right)$, and then we show that $N\left(\mathbf{x}_j \mid d\right) \leq N\left(\mathbf{x}_j \mid \mathbf{x}_j\right)$. To establish the first inequality, since $c(\mathbf{x})$ is strictly convex, the first-order conditions (FOC) are sufficient; the FOCs are

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

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

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

where the inequality follows from the convexity of $c$. Therefore, $(F, \sigma^*)$ 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 $\operatorname{int}\Delta(\Theta)$, such that Constraint $IC$ holds. Thus, contract $(M, t)$ induces $F$. The principal's expected cost is finite since $t\left(\mathbf{x}_j, \theta_k \mid \tau^*\right)$ is finite for all $j, k$. ■

## A.2 Corollary 3.3 Proof

*Proof.* Let $\operatorname{ext}\mathcal{F}(\mu)$ 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 $\Lambda_G$ that puts probability 1 on $\operatorname{ext}\mathcal{F}(\mu)$, and

$$G = \int_{\operatorname{ext}\mathcal{F}(\mu)} H \,\mathrm{d}\Lambda_G(H)\,. \tag{R}$$

Therefore, any distribution $G$ with support on $\operatorname{int}\Delta(\Theta)$ can be obtained by randomizing over distributions supported on at most $n$ points. Then by Lemma 3.2, $G$ can be induced 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 $|\operatorname{supp}(G)| > n$, and $(G, \hat{\sigma})$ is the induced optimal strategy of the agent. Without loss of generality, $M = \operatorname{supp}(G)$ and $\hat{\sigma}(\cdot \mid \mathbf{x}) = \delta_{\mathbf{x}}$ for all $\mathbf{x} \in \operatorname{supp}(G)$. Then for every posterior $\mathbf{x} \in \operatorname{supp}(G)$ and every $d \in M$ with $\hat{\sigma}(d \mid \mathbf{x}) > 0$,

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

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

Note that; however, for all $F \in \operatorname{supp}(\Lambda_G)$, $(M_F, t_F)$ need not be the least costly contract under which the agent chooses $F$: randomizing over $\operatorname{supp}(\Lambda_G)$ 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 induces distributions with support on at most $n$ points. This concludes the proof of part (ii). ■

## A.3 Proposition 4.1 Proof

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

$$\sum_{k=1}^{n-1}\left(t_m^k - t_m^n - \kappa c_k(\mathbf{x}_m)\right)\mu_k + t_m^n = Q\ ,$$

where $\mu_k$ is the $k$-th entry of $\mu$, 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 $IC-R$, the equality above can be reduced to $t_m^n = Q + \sum_{k=1}^{n-1} \kappa \mu_k c_k(\mu)$; plugging this into the other $n-1$ equations, we get $t_m^k = Q + \sum_{i=1}^{n-1} \kappa \mu_i c_i(\mu) + \kappa(c_k(\mathbf{x}_m) - c_k(\mu))$ for each $k = 1, \ldots, n$. We have thus found a solution. Because $F$ is an arbitrary distribution over posteriors supported on at most $n$ points, then by randomizing *ex ante*, any distribution $G$ with $\mathrm{supp}(G) \subseteq \mathrm{int}\,\Delta(\Theta)$ can be implemented efficiently. ∎

## A.4 Proposition 5.2 Proof

*Proof.* Using the concavifying line ($\star$), $\{x_L, x_H\}$ can be implemented efficiently if and only if

(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 $v_0/\kappa \geq \max\{\eta_1(x_H), \eta_2(x_L)\}$, as stated in the result. ∎

## A.5 Proposition 5.4 Proof

*Proof.* (i) is a consequence of Proposition 5.2. 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 ($\star$) 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 $g(x) \coloneqq h(x) - v_0 + \kappa c(x)$. 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 $x^\dagger \coloneqq \min\{x^\circ, 1\}$. Thus, $\gamma = \beta = 0$ is optimal if and only if $g\left(x^\dagger\right) \geq 0$. 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 $\beta$ by a sufficiently small $\varepsilon > 0$, strictly increasing her payoff and still leaving Constraint $IC\text{-}v_0$ satisfied. This yields (ii)a and (ii)b of the result. ■