

Robust Incentives for Information Acquisition

Gabriel Carroll, Stanford University

`gdc@stanford.edu`

July 7, 2015

Abstract

A principal needs to make a decision, and hires an expert, who can exert costly effort to supply useful information. The principal can incentivize effort by paying a reward based on the expert's reported information and on the true state of nature, which is revealed ex post. We assume risk-neutrality and limited liability. The principal is uncertain about the expert's information acquisition technology: she knows some actions (experiments) that he can take to obtain information, but there may also be other experiments available. The principal seeks robustness to this uncertainty, and so evaluates any incentive contract using a worst-case criterion. Under quite general conditions, we show that the optimal contract is a *restricted investment contract*, in which the expert chooses from a subset of the decisions available to the principal, and is then rewarded proportionally to the value of his designated decision in the realized state.

Keywords: information acquisition; principal-expert problem; restricted investment contract; robustness; scoring rules; worst case

Thanks to (in random order) David Parkes, Sylvain Chassang, Bengt Holmström, Iván Werning, Alex Wolitzky, Drew Fudenberg, Ben Golub, Yusuke Narita, Abhijit Banerjee, Daron Acemoglu, and Luis Zermeno for helpful discussions and advice, as well as participants at many seminars and conferences.

1 Introduction

How should one pay for information that requires effort to produce? An extensive theoretical literature on proper scoring rules (e.g. [18, 12]) studies how one can incentivize agents to truthfully report pre-existing knowledge, such as beliefs about the probability of some event. By contrast, the question of how best to give incentives to discover this information has been much less systematically studied. Yet there are plenty of situations where people are employed precisely to produce information: market researchers trying to forecast demand for new products; stock analysts, meteorologists, policy advisors, and so forth. We examine here the question of what shape and size of incentives are best for generating information, trading off the benefits of inducing effort against the fact that providing strong incentives can be expensive.

We adopt an agency model. A principal needs to make a decision that depends on an unknown state of nature, and so hires an expert, who can privately obtain information about the state by exerting effort. The expert has no intrinsic preferences over the decision being made, but the principal can incentivize him to exert effort by making his payment depend on how well the information he reports corresponds to the true state, which is publicly revealed *ex post*. We assume both parties are risk-neutral, and we impose limited liability — payments to the expert can never be less than zero.

We build on the recent work of Zermeno [20, 21], who gave an extremely general formulation of the principal-expert problem. However, we depart from that work, and from most of the existing agency literature, by assuming that the expert’s technology for acquiring information is not common knowledge. Instead, as in this author’s previous work [4], we take a robust-contracting approach: The principal knows some actions (here termed *experiments*) that the expert can perform to acquire information, but there may be other, unknown experiments available. The principal does not have a probabilistic belief about which experiments are and are not available. Rather, she wishes for a contract that can assure her a high expected net payoff without requiring any further knowledge about the available experiments. Accordingly, she evaluates incentive contracts based on their worst-case performance over all possible experiments the expert may have access to.

The above-mentioned previous work [4] took this worst-case approach to a standard moral hazard problem, in which an agent’s action directly produces output for the principal, and the agent can be paid based on the observed output. In that simpler setting, the concern for robustness results in linear contracts — in which the agent is paid a constant share of output — being optimal. Briefly, the intuition is as follows: the principal eval-

uates contracts based on the worst-case guarantee they provide on her expected payoff, and her limited knowledge provides only a worst-case guarantee on the agent’s expected payoff (via the actions she knows he has available); linearity provides a tight connection between these two.

In the present problem, a linear contract could be defined as follows: the expert recommends a decision, and is rewarded with some fixed fraction of the resulting payoff. The same forces leading to linearity in the simpler model apply here too. But the problem here is more complex, and it turns out the optimal contract is a variant which we call a *restricted investment contract*. Roughly, instead of paying the expert a constant share of the principal’s own realized payoff, she pays him proportionally to the payoff that *would* have accrued (in the realized state) if she had been restricted to a certain subset of her possible decisions. By excluding extreme decisions, the principal can make the limited liability constraint slack and thus pay the expert less while maintaining incentives.

This paper aims to make two main points. One is literal: we give a recipe that can be used to write contracts in a particular agency setting, and the resulting contracts are optimally robust in a precise sense. From this point of view, the principal-expert application is one where our assumption of extreme ignorance by the principal seems particularly natural. It is inherent in the nature of expertise that the principal who hires an expert does not fully understand how he does his job.

The other, broader point is methodological: we show how using a maxmin objective leads to a tidy and tractable model. By contrast, a traditional Bayesian approach, where the principal knows the expert’s information acquisition technology (or has a probabilistic prior about it), is unlikely to be tractable without much more specific functional form assumptions. (For example, Zermeño [20] gives an explicit solution in the special case when there are two states, effort is one-dimensional, and the first-order approach is valid.) Here, we can give results on the shape of the optimal contract that are quite general, with essentially *no* structural assumptions needed on the set of known experiments.

Beyond the worst-case criterion, our other key assumption is that the state is fully revealed *ex post*. This means that the principal finds out not only her actual payoff, but also what the payoff from any other decision would have been. This assumption is not always realistic, but it is reasonable in some applications. For example, the principal may be a financial investor, seeking advice as to how to divide her wealth among several assets; *ex post*, she can observe the realized returns on all assets (including the ones that she didn’t invest in). For another application, the principal may be a firm deciding how to price a new product, and the uncertainty may be about some additive demand shifter; *ex*

post, the value of the shifter can be inferred from the quantity sold at the chosen price. Or the principal may be a firm deciding which of several investments to undertake, and the profitability of each investment depends on various news events that may or may not occur and that the expert is supposed to help predict. If these events are reported in public media, then this application fits with our assumption that the state is revealed.

The full revelation assumption allows us to separate the problem of choosing the decision from that of providing incentives to the expert. This distinguishes our work from some prior literature on incentives for experts, including [20] as well as the computer science literature on decision markets [16, 7], in which optimal provision of incentives involves distorting decisions so as to reveal more information about the state. In the model here, the principal’s decision plays no role in the incentives faced by the expert. Nonetheless, we need to model the decision problem explicitly in order to describe how information is valued.

In the next section, we present the formal model. We then give a more detailed discussion of the intuition behind restricted investment contracts, before proceeding to the formal analysis showing that such contracts are optimal. The main proof is essentially an application of the same linear separation techniques used in [4]. In Section 3, we discuss a number of extensions and variations. In particular, we discuss when an *unrestricted* investment contract — which, up to a normalization, is just a linear contract — is or is not optimal; this generally can happen when the principal’s optimal decision is not too sensitive to the posterior, and in particular always holds if the decision problem is binary.

There is some previous work on agency problems that considered incentives to acquire information. Besides the work of Zermeno [20, 21] cited above, other examples include Demski and Sappington [9] (who introduced the term “expert”) and Malcolmson [14]. However these models assumed that the decision is delegated to the expert, and that only the realized payoff is observed, not the entire state. Osband [15] considered a model more similar to the present paper in these respects, but focused on optimal screening with a prior belief over the expert’s technology, with specific functional form assumptions. The scoring rules literature has also recognized at least informally that higher stakes give more incentives to acquire information (e.g. [13]), but without actually formulating the optimal contracting problem.

In addition to these strands of literature, the present paper also contributes to the growing literature on mechanism design using a maxmin objective in uncertain environments, e.g. [1, 3, 8, 10, 11]. The earlier paper [4] contains more discussion of this literature, and of the interpretation of the worst-case objective.

2 Model and results

2.1 Notation

For $X \subseteq \mathbb{R}^k$, $\Delta(X)$ is the space of Borel probability distributions on X , with the weak topology. For $x \in X$, we write δ_x for the degenerate distribution putting probability 1 on x . \mathbb{R}^+ is the set of nonnegative real numbers.

2.2 The setup

As already indicated, the model is based on that of [20], with the main difference being the principal's non-quantifiable uncertainty about the expert's information acquisition technology.

The principal needs to choose a decision from some compact space D . The payoff to each decision depends on a state of nature, which will be realized in the future. We assume the set Ω of possible states is finite. Payoffs are represented by a continuous function $u : D \times \Omega \rightarrow \mathbb{R}$. We may refer to the triple (D, Ω, u) as the principal's *decision problem*. The principal also has a prior belief about the state, $p_0 \in \Delta(\Omega)$, with $p_0(\omega) > 0$ for each $\omega \in \Omega$.

Before the principal makes her decision, she can hire the expert to obtain information about the state. The expert initially shares the prior p_0 , but can obtain more information by performing an *experiment*. One can think of an experiment as producing a signal, observed by the expert; signals and states can follow any joint distribution such that the marginal distribution over states is p_0 . However, the signal will matter only through the expert's resulting posterior belief about the state, so we take the notational shortcut of representing experiments directly in terms of posteriors. Thus, we define an *experiment* to be a pair $(F, c) \in \Delta(\Delta(\Omega)) \times \mathbb{R}^+$, such that F has mean p_0 . The interpretation is that the expert can, at a cost of c , perform the experiment, and obtain a posterior (an element of $\Delta(\Omega)$) drawn from distribution F . The requirement that F should have mean p_0 is simply the law of iterated expectations — the posterior should, in expectation, be equal to the prior. We will typically use the variable p for a posterior.

We give $\Delta(\Delta(\Omega)) \times \mathbb{R}^+$ the natural product topology, and define an *information acquisition technology* (IAT) to be a nonempty, compact subset of $\Delta(\Delta(\Omega)) \times \mathbb{R}^+$, with every element (F, c) satisfying $E_F[p] = p_0$. An IAT then describes the set of experiments available to the expert. There is some exogenously given IAT, \mathcal{I}_0 , consisting of the experiments which are mutually known to be available. From the principal's point of view,

the true set of experiments (known only to the expert) may potentially be any IAT \mathcal{I} such that $\mathcal{I}_0 \subseteq \mathcal{I}$. We will sometimes be specifically interested in \mathcal{I}_0 that satisfy the following *full-support condition*: for every $(F, c) \in \mathcal{I}_0$, the support of F is all of $\Delta(\Omega)$, or else $F = \delta_{p_0}$.

After the contract is chosen, the expert can acquire information, and make a report to the principal, who then chooses a decision $d \in D$. After the decision is made, the true state is revealed, and payments can be made contingent on all relevant observable information (the report and the realized state). Limited liability means that payments can never be negative. Thus, a *contract* is a pair (M, w) , consisting of a *message space* M , some nonempty compact space; and a *payment function* $w : M \times \Omega \rightarrow \mathbb{R}^+$, which must be continuous. (The topological assumptions ensure that the expert’s behavior is well-defined.)

We could potentially apply the revelation principle and assume that the expert’s message is simply his posterior belief (this approach is used in [20, 21], and is also in keeping with the proper scoring rules literature). However, we will not do so, because we will be interested in contracts that are naturally described as indirect mechanisms.

On the other hand, we can safely assume that the principal learns the expert’s posterior when he makes his report. This assumption requires a bit of justification. It is not necessarily true given the way we have modeled contracts thus far; if the message space M is small, then the expert may choose the same message m at many posteriors p and, more germanely, the message may not be informative enough for the principal to infer what her optimal decision given p is. But any such “pooling” contract can be (weakly) improved upon as follows: expand the message space by having the expert also report his posterior p , in addition to the original message m ; but for purposes of calculating payment, the posterior report is ignored. Then the expert’s incentives for information acquisition, and the for m component of his report, are exactly the same as in the original contract; and he is indifferent about the p component of his report, so he is willing to report the true posterior, which clearly allows the principal to choose the best possible decision d . Because we are concerned specifically with optimal contracts for the principal, we may as well assume she uses this modified contract.

This argument is more fully formalized in Appendix A. For the rest of the main paper, we just assume without further comment that the expert reports both the explicitly-modeled $m \in M$ and his posterior p , so that the principal can make the optimal decision given p .

We can now summarize the timing of the game:

1. the principal offers a contract (M, w) ;
2. the expert, knowing \mathcal{I} , chooses experiment $(F, c) \in \mathcal{I}$;
3. a posterior $p \sim F$ is realized, and privately observed by the expert;
4. the expert chooses a message $m \in M$ to send (along with his posterior p);
5. the principal chooses decision d ;
6. the state ω is revealed;
7. payoffs are received: $u(d, \omega) - w(m, \omega)$ for the principal; $w(m, \omega) - c$ for the expert.

It remains only to describe behavior. We give a brief summary here, and will introduce formal notation shortly. The expert knows \mathcal{I} , and he chooses (F, c) and then m to maximize his expected payoff. If the expert is indifferent between several choices, then we assume he acts so as to maximize the principal's expected payoff. The above behavior by the expert gives rise to an expected payoff for the principal for each possible IAT \mathcal{I} ; we will notate this expected utility by $V_P(M, w|\mathcal{I})$. The principal evaluates each contract (M, w) in terms of the guarantee it can provide on V_P , across all possible IAT's \mathcal{I} , knowing only that $\mathcal{I}_0 \subseteq \mathcal{I}$. Finally, the question is how to design a contract to maximize this worst-case value.

Before developing the formal notation, we introduce some other useful objects. Suppose that the principal learns the expert's posterior is $p \in \Delta(\Omega)$. Then, clearly, she will choose d to maximize $E_p[u(d, \omega)]$. We denote this decision and the expected payoff by

$$d(p) = \arg \max_{d \in D} E_p[u(d, \omega)]; \quad U(p) = \max_{d \in D} E_p[u(d, \omega)].$$

(If there are several maximizers then $d(p)$ can be chosen arbitrarily among them.) Note that U is a convex function on the simplex $\Delta(\Omega)$, since it is the maximum of the affine functions $p \mapsto E_p[u(d, \omega)]$.

Similarly, when the expert has posterior p , he will choose message m so as to maximize $E_p[w(m, \omega)]$. We denote

$$W(p) = \max_{m \in M} E_p[w(m, \omega)],$$

and call this function $W : \Delta(\Omega) \rightarrow \mathbb{R}^+$ the *reduced form* of the given contract (M, w) . W is likewise convex in p ; it is the upper envelope of the nonnegative-valued affine functions

given by $p \mapsto E_p[w(m, \omega)]$, for each $m \in M$. As we shall see, the expert's incentives to acquire information depend only on the reduced form of the contract.

Now we can describe behavior formally. The expert, given posterior p , chooses message m as above, leading to expected payment $W(p)$. At the earlier, experiment-choosing stage, he knows the IAT \mathcal{I} and so chooses $(F, c) \in \mathcal{I}$ to maximize expected payoff $E_F[W(p)] - c$. We will write the value of the contract (M, w) to the expert as

$$V_E(M, w|\mathcal{I}) = \max_{(F, c) \in \mathcal{I}} (E_F[W(p)] - c)$$

and the expert's choice set as

$$I^*(M, w|\mathcal{I}) = \arg \max_{(F, c) \in \mathcal{I}} (E_F[W(p)] - c).$$

When the principal learns the posterior p , she will make decision $d(p)$, gaining expected gross payoff $U(p)$. Thus, her expected net payoff from the contract under IAT \mathcal{I} is

$$V_P(M, w|\mathcal{I}) = \max_{(F, c) \in I^*(M, w|\mathcal{I})} (E_F[U(p)] - W(p)).$$

(The maximization reflects the fact that the expert may be indifferent between several optimal experiments; we assume he then chooses the one that is best for the principal. Alternative tie-breaking assumptions would give substantively similar results, but would add some inconvenient technical details.) The principal evaluates each possible contract ex ante by its worst-case guarantee on her expected payoff, over all IAT's \mathcal{I} :

$$V_P(M, w) = \inf_{\mathcal{I} \supseteq \mathcal{I}_0} V_P(M, w|\mathcal{I}).$$

The principal's problem, which we analyze, is then how to choose the contract (M, w) to maximize V_P .

From here on, we will maintain the *non-triviality* assumption that there exists some contract (M, w) with $V_P(M, w) > U(p_0)$. That is, the principal benefits from hiring the expert. We shall shortly see conditions on primitives that ensure that this assumption is satisfied. (This assumption is not formally needed for our results. But if it is not satisfied, the problem is uninteresting — the optimal contract is clearly always to pay zero.)

2.3 Intuitions

How can the principal write a contract to guarantee herself a reasonably high expected payoff? A natural first try would be to use a linear contract: The expert recommends a decision, and the principal follows the recommendation and then pays the expert some fixed fraction $\alpha \in (0, 1]$ of her gross payoff. We could compute a payoff guarantee from such a contract exactly as in [4] (see also [6]): The principal pays fraction α of her gross payoff to the expert, and keeps the remaining $1 - \alpha$ for herself, so her net payoff is at least $(1 - \alpha)/\alpha$ times the expert's. (This relation need not be an equality, since the expert also suffers the disutility of effort.) A lower bound on the expert's expected payoff can be obtained, by taking the best among the experiments that are known to be available; the expert would only use an unknown experiment if it was even better for him. Then, multiplying by $(1 - \alpha)/\alpha$ gives a lower bound on the principal's expected payoff.

Actually, this is not quite the way to define linear contracts in our setting: Since the principal's gross payoffs $u(m, \omega)$ may be negative (in which case the contract above violates limited liability) or may all be very high (in which case limited liability is slack), we should first add some constant adjustment $\beta \in \mathbb{R}$. Thus, the proper definition of linear contracts in our formalism is given by

$$M = D, \quad w(m, \omega) = \alpha(u(m, \omega) + \beta). \quad (2.1)$$

When the expert has posterior p , she sends a message describing her recommended decision, $m = d(p)$. The resulting reduced form of the contract is

$$W(p) = \alpha(U(p) + \beta).$$

(Explicitly, the optimal choice of β in such a contract would be such that limited liability just binds: $\beta = -\min_{d, \omega} u(d, \omega)$.)

It is easy to incorporate this β adjustment into our calculation of the guarantee from such a contract. Let us write out the algebra explicitly.

If the expert chooses experiment (F, c) , his expected payoff will be $E_F[W(p)] - c = \alpha(E_F[U(p)] + \beta) - c$. Hence we have

$$\alpha(E_F[U(p)] + \beta) \geq \alpha(E_F[U(p)] + \beta) - c = V_E(M, w|\mathcal{I}) \geq V_E(M, w|\mathcal{I}_0).$$

Here the inequality at the end comes from the fact that $\mathcal{I} \supseteq \mathcal{I}_0$, so the expert certainly

does at least as well with \mathcal{I} as \mathcal{I}_0 .

Whenever the expert produces a gross payoff of y , he receives $\alpha(y + \beta)$, and the principal receives $(1 - \alpha)y - \alpha\beta$. Hence, the principal's expected payoff will be

$$\begin{aligned}
(1 - \alpha)E_F[U(p)] - \alpha\beta &\geq \frac{1 - \alpha}{\alpha}V_E(M, w|\mathcal{I}_0) - \beta \\
&= \frac{1 - \alpha}{\alpha} \left[\max_{(F, c) \in \mathcal{I}_0} \alpha(E_F[U(p)] + \beta) - c \right] - \beta \\
&= \max_{(F, c) \in \mathcal{I}_0} \left((1 - \alpha)E_F[U(p)] - \frac{1 - \alpha}{\alpha}c \right) - \alpha\beta. \quad (2.2)
\end{aligned}$$

Thus the guarantee from the linear contract, $V_P(M, w)$, is at least the right-hand side of (2.2).

At this point, we can fulfill a promise made earlier: to give a sufficient condition on primitives to ensure the non-triviality assumption is satisfied. Let F be any distribution over posteriors such that $E_F[U(p)] > U(p_0)$ — that is, any distribution that potentially provides useful information for the decision. One can check that, as long as

$$c < \left(\sqrt{E_F[U(p)] + \beta} - \sqrt{U(p_0) + \beta} \right)^2, \quad (2.3)$$

there exists α such that the right side of (2.2) is greater than $U(p_0)$. Hence, if \mathcal{I}_0 contains some such (F, c) , non-triviality holds.

Now, the logic so far is identical to that of the pure moral hazard setting in [4]. In that simpler setting, linear contracts were optimal: the only way to turn a guarantee on the agent's expected payoff (provided by the known technology) into a guarantee on the principal's payoff was to have a linear relationship between the two. Here, however, there is scope for improving on linear contracts, in two ways.

To illustrate the ideas, we refer to Figure 1, which shows an example. Panel (a) shows the principal's decision problem. This represents a situation with two states, ω_0 and ω_1 , and four decisions d_1, d_2, d_3, d_4 with the following payoffs:

	ω_0	ω_1
d_1	10	2
d_2	8	8
d_3	6	10
d_4	0	12

In this case, a posterior p is simply a probability of state ω_1 , so the space of posteriors is identified with the unit interval. Each decision is represented as a (thin) line — plotting the principal’s expected payoff as a function of the posterior. The upper envelope of these lines is then $U(p)$, as shown.

The linear contract that simply pays the expert $1/2$ of the principal’s gross payoff is shown in panel (b). The four thin lines here correspond to messages the expert can send in this contract — each line showing the expert’s expected payment from sending the given message, as a function of his posterior — and their upper envelope is the reduced form $W(p)$.

One simple way to improve on such a contract is to replace the single constant adjustment β by an adjustment that depends on the state, $\beta(\omega)$. Explicitly, instead of (2.1), we define the contract by

$$M = D, \quad w(m, \omega) = \alpha(u(m, \omega) + \beta(\omega)).$$

Note that this change can have no effect on the expert’s incentives as to which experiment to perform, because he has no control over which state arises and so which adjustment he receives. This generalization gives the principal more degrees of freedom, and she optimally chooses the state-by-state adjustments so that limited liability binds in each state separately: $\beta(\omega) = -\min_{d \in D} u(d, \omega)$. This reduces the amount paid to the expert without changing his incentives. In our example, this change in the contract is illustrated in Figure 1(c).

The second reason why linear contracts may not be optimal is more substantive: By cutting out risky decisions, the principal can relax the limited liability constraint, and thus further lower β . This is illustrated in Figure 1(d), which differs from (c) in that the two extreme messages in the contract — corresponding to decisions d_1 and d_4 — have been removed. Notice that there is now lots of slack in the limited liability constraint; the principal can further adjust payments downward by 3 in each state before limited liability binds, and thereby further save money.

Now, cutting out the extreme decisions makes W less convex, and so may weaken the expert’s incentives for information acquisition. However, if the known experiments rarely yield posteriors in the ranges where d_1 or d_4 are optimal, then this effect on incentives is small, and the effect of lowering β dominates. Note that under the new contract, the principal’s *decisions* would not be restricted — she could still choose d_1 or d_4 if the expert happened to report an extreme posterior, but the expert’s payment would be based on d_2

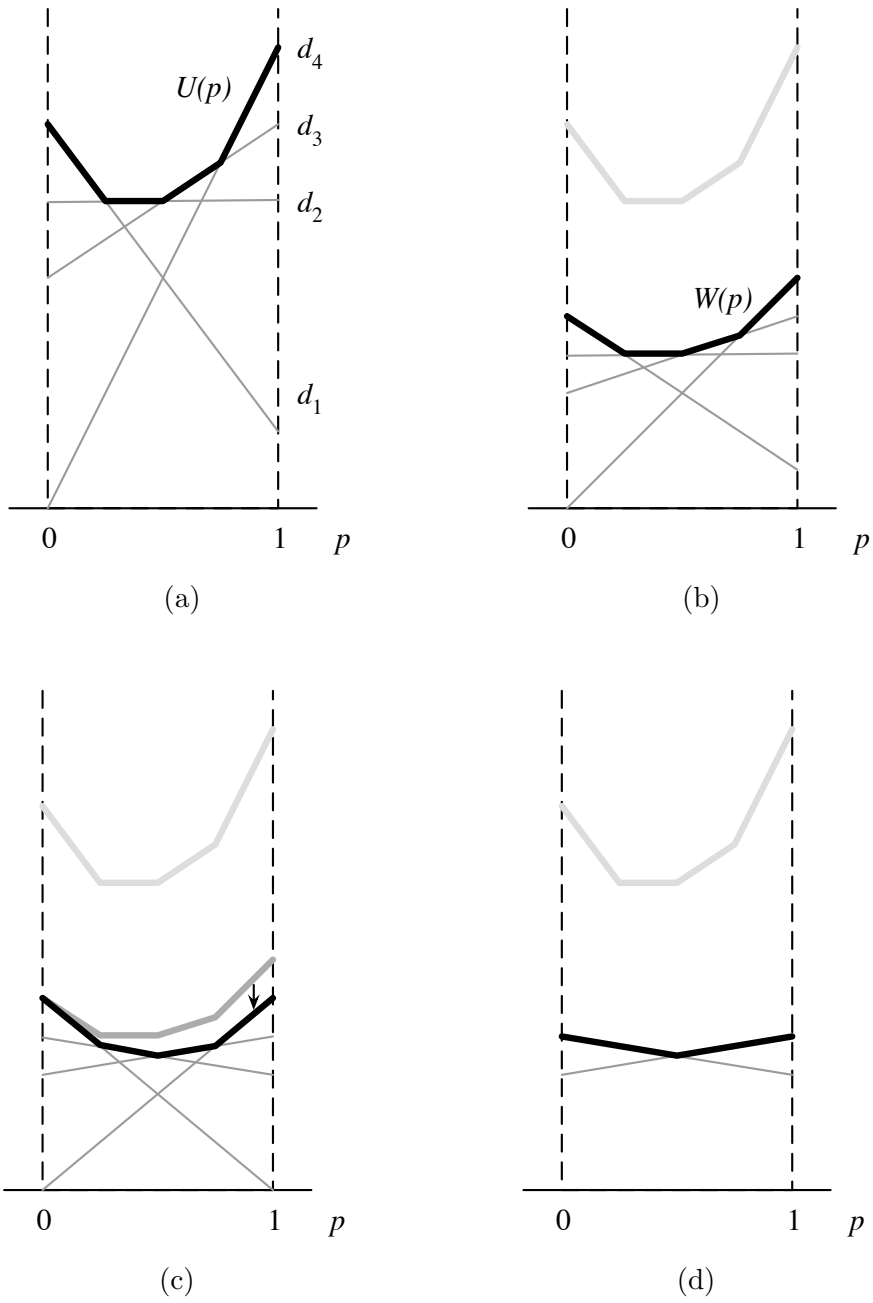


Figure 1: Intuition for improving on linear contracts. (a) The decision problem. (b) A linear contract. (c) State-by-state adjustment. (d) Trimming extreme decisions.

or d_3 instead.

This idea leads us to the notion of a *restricted investment contract*, defined formally as follows: The message space is some compact subset of decisions, $M \subseteq D$, and the payment function is given by

$$w(m, \omega) = \alpha(u(m, \omega) + \beta(\omega)) \quad (2.4)$$

for some $\alpha \in [0, 1]$ and some $\beta : \Omega \rightarrow \mathbb{R}$, such that the resulting payments are always nonnegative. The name emphasizes that the expert is allowed to “invest” in a restricted subset of decisions, and once the state is revealed, he is compensated according to the payoff that his chosen decision would have produced in that state.

What makes our setting different from the pure moral hazard setting, in which linear contracts are optimal? In that setting, limited liability applies directly to the incentive instrument. Here, there is a separation between the two: Limited liability applies to the state-by-state payments $w(m, \omega)$, but the expert’s incentives are determined by the reduced form $W(p)$. In our setting, a departure from pure linearity can relax limited liability appreciably while having only a small effect on the expert’s incentives.

2.4 Full analysis

So far, the discussion has emphasized restricted investment contracts. However, for the formal analysis, it will be useful to introduce a closely related class, which we call *transform-bounded contracts*. Such a contract is parameterized by a value of $\alpha \in [0, 1]$ and $\beta : \Omega \rightarrow \mathbb{R}$, and is defined as follows: First, define a transformed version of the principal’s indirect utility function, $\widehat{U}_{\alpha, \beta} : \Delta(\Omega) \rightarrow \mathbb{R}$, by

$$\widehat{U}_{\alpha, \beta}(p) = \alpha(U(p) + E_p[\beta(\omega)]).$$

Then, the expert is allowed to choose any nonnegative-valued affine function on the simplex that is bounded above by $\widehat{U}_{\alpha, \beta}$, and to be rewarded according to the function he chooses. Formally, the message space M of the transform-bounded contract is the set of functions $m : \Omega \rightarrow \mathbb{R}^+$ satisfying

$$E_p[m(\omega)] \leq \widehat{U}_{\alpha, \beta}(p) \quad \text{for all } p \in \Delta(\Omega),$$

and the payment function is simply

$$w(m, \omega) = m(\omega).$$

As long as the values of β are large enough so that $U(p) + E_p[\beta(\omega)] \geq 0$ for all p , the resulting M is nonempty, and so we do obtain a contract. The concept is illustrated in Figure 2. Panel (a) shows a possible function U (the medium-gray thick curve), a corresponding $\widehat{U}_{\alpha, \beta}$ (the lighter curve), and some examples of affine functions bounded by $\widehat{U}_{\alpha, \beta}$ (the thin lines). Panel (b) further shows, in reduced form, the corresponding transform-bounded contract, which is the upper envelope of these thin lines.




Figure 2: (a), (b) Construction of a transform-bounded contract. (c) The transformed utility function itself may not be the reduced form of a contract, due to limited liability.

We will say that two contracts are *equivalent* if they have the same reduced form. Notice that the values of $V_E(M, w|\mathcal{I})$, $I^*(M, w|\mathcal{I})$, $V_P(M, w|\mathcal{I})$, and therefore also $V_P(M, w)$ all depend on the contract (M, w) only through its reduced form; hence, any two equivalent contracts give the same payoff both to the expert and to the principal. Our analysis will show that transform-bounded contracts are optimal, and then that, under appropriate circumstances, transform-bounded contracts are equivalent to restricted investment contracts.

Thus, we state the first main step of the analysis:

Theorem 2.1. *There is an optimal contract that is transform-bounded. Moreover, if the known IAT \mathcal{I}_0 satisfies the full-support assumption, then every optimal contract is equivalent to a transform-bounded contract.*

For the second portion of the analysis, we will need one more definition. We will say that the decision problem (D, Ω, u) is *convex* if, for all $d, d' \in D$ and all $\lambda \in [0, 1]$, there exists $\hat{d} \in D$ such that

$$u(\hat{d}, \omega) \geq \lambda u(d, \omega) + (1 - \lambda)u(d', \omega)$$

for all ω .

This is satisfied, for example, whenever D is a convex subset of some vector space and $u(d, \omega)$ is concave in d for each ω . So it holds in the asset allocation application mentioned in the introduction, if the total return on each asset (conditional on the state) is weakly concave in the amount invested. It holds in the price selection application if profit is a concave function of price, for example, in a model with a linear demand curve. Or, more generally, if we consider any arbitrary decision problem and extend D by explicitly including randomized decisions, then convexity holds trivially.

Convexity is a richness condition that serves to ensure that the messages used in a transform-bounded contract correspond to actual decisions available to the principal. Thus the following, purely technical result, which constitutes the second part of our analysis:

Proposition 2.2. *Suppose the decision problem is convex. Then every transform-bounded contract is equivalent to a restricted investment contract. In fact, the transform-bounded contract with parameters α, β is equivalent to the restricted investment contract with the same parameters, and message space M equal to*

$$D_R(\beta) = \{d \in D \mid u(d, \omega) + \beta(\omega) \geq 0 \text{ for all } \omega\}, \quad (2.5)$$

unless $\alpha = 0$ in which case we can just take $M = D$.

Combining these two results immediately gives the main result of the paper:

Corollary 2.3. *Suppose the decision problem is convex. Then there is some restricted investment contract that is optimal. If \mathcal{I}_0 satisfies the full-support assumption, then every optimal contract is equivalent to a restricted investment contract.*

Now that the results are stated, we begin the proofs. The main steps will be described here, but a few details will be left to Appendix B.

The proof of Theorem 2.1 closely parallels the analysis in [4]. We take any given contract, and show that it can be improved to a transform-bounded contract whose worst-case guarantee for the principal is at least as good.

To be more specific, we use a linear separation argument to show that, for any given contract, there is some linear relation between the principal's and expert's reduced-form payoffs that drives its worst-case guarantee: there exist parameters α and β such that $W(p) \leq \widehat{U}_{\alpha,\beta}(p)$ for all p , with equality on the support of the worst-case F . If we could replace this contract with another contract satisfying the relation with equality, $W'(p) = \widehat{U}_{\alpha,\beta}(p)$, then this new contract would be better for the expert. Then, using the linear relation between the two parties' payoffs, and a calculation similar to that used for (2.2) above, it would follow that the new contract is better for the principal as well. Unfortunately, $\widehat{U}_{\alpha,\beta}$ may not be the reduced form of any contract: As shown in Figure 2(c), if we try to construct messages whose upper envelope equals $\widehat{U}_{\alpha,\beta}$, we may end up violating the limited liability constraint. Fortunately, the same calculation succeeds if we use the corresponding transform-bounded contract instead.

We begin by characterizing the worst-case payoff for any given contract. We actually need to deal separately with *zero contracts*: those whose reduced form W satisfies $W(p) = 0$ for all p . We denote the guarantee of a zero contract by $V_P(0)$. If there exists some experiment $(F, c) \in \mathcal{I}_0$ with $c = 0$, then for any \mathcal{I} , the expert will choose whichever such experiment maximizes the principal's expected payoff $E_F[U(p)]$; hence, $V_P(0)$ is simply $\max_{(F,0) \in \mathcal{I}_0} E_F[U(p)]$. If there is no such experiment in \mathcal{I}_0 , then when the IAT is $\mathcal{I} = \mathcal{I}_0 \cup \{(\delta_{p_0}, 0)\}$, the expert will choose the latter experiment, and we see that $V_P(0) = U(p_0)$. (The principal will not do worse than this under any other IAT, by convexity of U .)

For the rest of the analysis, it is useful to define *eligible* contracts as those satisfying both $V_P(M, w) \geq V_P(0)$ and $V_P(M, w) > U(p_0)$. We know that some eligible contract exists — the contract guaranteed by the non-triviality assumption is eligible, unless it is worse than the zero contract, in which case the zero contract is eligible. Therefore, when searching for the optimal contract, we can focus our attention on eligible contracts.

Now we can describe the worst-case payoff in general:

Lemma 2.4. *Let (M, w) be any eligible contract, different from the zero contract. Let W*

be its reduced form. Then,

$$V_P(M, w) = \min E_F[U(p) - W(p)] \quad \text{over } F \in \Delta(\Delta(\Omega)) \text{ such that} \quad (2.6)$$

$$E_F[p] = p_0 \text{ and } E_F[W(p)] \geq V_E(M, w|\mathcal{I}_0).$$

Moreover, for any F attaining the minimum, the constraint holds with equality: $E_F[W(p)] = V_E(M, w|\mathcal{I}_0)$.

Intuitively, the principal faces great uncertainty about the experiments available to the expert, and she knows only two things about the distribution F on posteriors that he will choose: she knows that $E_F[p] = p_0$, and she knows a lower bound on $E_F[W(p)]$ due to the experiments that are known to be available. Lemma 2.4 translates these directly into the principal's worst-case guarantee. The actual proof requires a little more work than this brief explanation, due to subtleties introduced by the assumption of tie-breaking in favor of the principal. We leave the details to the full proof in Appendix B.

Next, we can give the key step in the proof of Theorem 2.1: a linear inequality relating the principal's and expert's payoffs under the given contract, with equality holding for the worst case. This follows from Lemma 2.4 by applying a linear separation theorem.

Lemma 2.5. *Let (M, w) be any eligible contract, different from the zero contract, and let W be its reduced form. Then there exist $\alpha \in (0, 1)$ and $\beta : \Omega \rightarrow \mathbb{R}$ such that*

$$W(p) \leq \alpha(U(p) + E_p[\beta(\omega)]) \quad \text{for all } p \in \Delta(\Omega); \quad (2.7)$$

$$V_E(M, w|\mathcal{I}_0) = \alpha(V_P(M, w) + V_E(M, w|\mathcal{I}_0) + E_{p_0}[\beta(\omega)]). \quad (2.8)$$

Proof: Consider the following two subsets of $\Delta(\Omega) \times \mathbb{R} \times \mathbb{R}$:

- S is the convex hull of points $(p, W(p), U(p) - W(p))$;
- T is the set of all points (p_0, y, z) such that $y \geq V_E(M, w|\mathcal{I}_0)$ and $z < V_P(M, w)$.

These two sets are disjoint — otherwise there would be some distribution F such that $E_F[p] = p_0$, $E_F[W(p)] \geq V_E(M, w|\mathcal{I}_0)$, and $E_F[U(p) - W(p)] < V_P(M, w)$, contradicting Lemma 2.4. Applying a proper separating hyperplane theorem¹ to these sets gives us

¹E.g. [17, Theorem 11.3]: given two disjoint convex sets in Euclidean space, there exists a hyperplane that weakly separates them *and* such that the two sets are not both contained in the hyperplane.

$\lambda \in \mathbb{R}^\Omega$, $\mu, \nu, \xi \in \mathbb{R}$ such that

$$\sum_{\omega} \lambda_{\omega} p(\omega) + \mu W(p) + \nu(U(p) - W(p)) \leq \xi \quad \text{for all } p \in \Delta(\Omega); \quad (2.9)$$

$$\sum_{\omega} \lambda_{\omega} p_0(\omega) + \mu V_E(M, w|\mathcal{I}_0) + \nu V_P(M, w) \geq \xi; \quad (2.10)$$

$\mu \geq 0$, $\nu \leq 0$; and we do not simultaneously have $\mu = \nu = 0$ and all the λ_{ω} equal to each other.

Moreover, we can let F^* be the distribution attaining the minimum in (2.6), and take the expectation over $p \sim F^*$ in (2.9). Since $E_{F^*}[W(p)] = V_E(M, w|\mathcal{I}_0)$ and $E_{F^*}[U(p) - W(p)] = V_P(M, w)$, we must have equality in (2.9) for all p in the support of F^* , and also have equality in (2.10).

At least one of the inequalities $\mu \geq 0$, $\nu \leq 0$ must hold strictly: otherwise $\sum \lambda_{\omega} p(\omega) \leq \xi \leq \sum \lambda_{\omega} p_0(\omega)$ for all $p \in \Delta(\Omega)$, which would only be possible if all λ_{ω} were equal (because p_0 has full support), but proper separation means this cannot happen. Let's show that in fact both inequalities are strict: $\mu > 0$ and $\nu < 0$.

- Suppose $\nu = 0$. Then $\mu > 0$, and (2.9) implies $W(p)$ is bounded above by the affine function $Z(p) = (\xi - \sum_{\omega} \lambda_{\omega} p(\omega))/\mu$, with equality throughout the support of F^* . The last statement of Lemma 2.4 then assures us that $V_E(M, w|\mathcal{I}_0) = E_{F^*}[W(p)] = E_{F^*}[Z(p)] = Z(p_0)$. On the other hand, if (F, c) is the experiment the expert chooses under IAT \mathcal{I}_0 , then

$$V_E(M, w|\mathcal{I}_0) = E_F[W(p)] - c \leq E_F[Z(p)] - c = Z(p_0) - c,$$

so it must be that $c = 0$. Thus, there is an experiment available for free in \mathcal{I}_0 that gives the principal $V_P(M, w|\mathcal{I}_0) \geq V_P(M, w)$. But then the principal could have gotten a strictly higher payoff by using a zero contract, since she would get at least $E_F[U(p)] > E_F[U(p) - W(p)] \geq V_P(M, w)$. (The strict inequality holds because W is nonzero, so the expert can assure himself positive expected payment.) This contradicts eligibility. Therefore, $\nu < 0$ strictly.

- Suppose $\mu = 0$. Then (2.9) for p_0 and the equality in (2.10) imply

$$U(p_0) - W(p_0) \geq \frac{\sum_{\omega} \lambda_{\omega} p_0(\omega) - \xi}{-\nu} = V_P(M, w),$$

so $V_P(M, w) \leq U(p_0)$, again contradicting eligibility. Thus $\mu > 0$ strictly.

Now we can finish the proof of the lemma. We can rearrange (2.9) to obtain

$$W(p) \leq \frac{-\sum_{\omega} \lambda_{\omega} p(\omega) - \nu U(p) + \xi}{\mu - \nu} \quad (2.11)$$

for all p . Put

$$\alpha = \frac{-\nu}{\mu - \nu}, \quad \beta(\omega) = \frac{-\lambda_{\omega} + \xi}{-\nu} \text{ for each } \omega.$$

Then (2.7) is just (2.11), and (2.8) follows from equality in (2.10). Note also that indeed $\alpha \in (0, 1)$. So the lemma is proven. \square

To complete the proof of Theorem 2.1, showing that a transform-bounded contract is optimal, we first need to know:

Lemma 2.6. *An optimal contract exists.*

The proof is topological — we show that we can restrict attention to a compact set of contracts (under an appropriate topology), and that V_P is upper semi-continuous on this set. (It is not continuous in general.) The details are in Appendix B.

Now we can finish the proof of Theorem 2.1. We can start with any given contract, and use the values of α and β given by Lemma 2.5 to construct a transform-bounded contract that improves on it. In particular, if we start from an optimal contract, then the transform-bounded contract we construct must again be optimal.

Proof of Theorem 2.1: By Lemma 2.6, there exists an optimal contract, call it (M, w) ; and we know it must be eligible. Write W for the reduced form. We can assume it is nonzero, since otherwise it is already equivalent to a transform-bounded contract.

Let α, β be the values given by Lemma 2.5. Let (M', w') be the transform-bounded contract with these parameters, and let W' be its reduced form.

Let m be any message in the original contract. Then for all p , we have

$$E_p[w(m, \omega)] \leq W(p) \leq \alpha(U(p) + E_p[\beta(\omega)])$$

where the first inequality is by definition of $W(p)$ and the second is from (2.7). Thus, the map $\omega \mapsto w(m, \omega)$ is in M' . In particular, for all p , $W'(p) \geq E_p[w(m, \omega)]$. Taking the maximum over m gives $W'(p) \geq W(p)$: the new contract dominates the original one (for the expert).

On the other hand, we still have

$$W'(p) \leq \alpha(U(p) + E_p[\beta(\omega)]) \quad (2.12)$$

by definition of the transform-bounded contract. This can be rearranged to give

$$U(p) - W'(p) \geq \frac{1 - \alpha}{\alpha} W'(p) - E_p[\beta(\omega)]. \quad (2.13)$$

The relation $W'(p) \geq W(p)$ for all p implies immediately that $V_E(M', w' | \mathcal{I}_0) \geq V_E(M, w | \mathcal{I}_0)$, since any experiment the effort chooses under (M, w) and \mathcal{I}_0 will pay at least as much under (M', w') . Moreover, under the new contract, for any IAT \mathcal{I} , the expert will choose an experiment (F, c) such that $E_F[W'(p)] \geq E_F[W'(p)] - c = V_E(M', w' | \mathcal{I}) \geq V_E(M', w' | \mathcal{I}_0)$. And so

$$\begin{aligned} E_F[U(p) - W'(p)] &\geq E_F \left[\frac{1 - \alpha}{\alpha} W'(p) - E_p[\beta(\omega)] \right] && \text{(by (2.13))} \\ &\geq \frac{1 - \alpha}{\alpha} V_E(M', w' | \mathcal{I}_0) - E_{p_0}[\beta(\omega)] \\ &= V_P(M, w) + \frac{1 - \alpha}{\alpha} (V_E(M', w' | \mathcal{I}_0) - V_E(M, w | \mathcal{I}_0)) && \text{(by (2.8)).} \end{aligned}$$

The left-hand side equals $V_P(M', w' | \mathcal{I})$. Taking the infimum over all \mathcal{I} , we have

$$V_P(M', w') \geq V_P(M, w) + \frac{1 - \alpha}{\alpha} (V_E(M', w' | \mathcal{I}_0) - V_E(M, w | \mathcal{I}_0)) \geq V_P(M, w). \quad (2.14)$$

Since (M, w) was assumed to be an optimal contract, this must be an equality, and the transform-bounded contract (M', w') is again optimal.

It remains to prove the final statement: we will show that under the full-support assumption, the initial optimal contract is equivalent to our transform-bounded contract, that is, W' is identically equal to W . Suppose not. Then, we have $V_E(M', w' | \mathcal{I}_0) > V_E(M, w | \mathcal{I}_0)$, since the experiment chosen by the expert under (M, w) and \mathcal{I}_0 has full support and so gives the expert strictly higher expected payoff under (M', w') . (The expert's chosen experiment cannot be distribution δ_{p_0} , because then $V_P(M, w) \leq U(p_0)$ contradicting eligibility.) Together with $\alpha \in (0, 1)$, this implies that (2.14) is a strict inequality. But this contradicts the optimality of (M, w) . Thus, W' must be equal to W after all, which completes the proof. \square

Now, to complete our main analysis (aside from the details that we have left to Appendix B), it remains only to make the leap from transform-bounded contracts to restricted investment contracts, Proposition 2.2. We need to show that, under convexity, for any given transform-bounded contract, the message optimally chosen by the expert at any given posterior p actually represents some decision available in the corresponding

restricted decision space.

Proof of Proposition 2.2: Suppose that the decision problem is convex. Let (M, w) be the transform-bounded contract with parameters α, β , and (M', w') the corresponding restricted investment contract identified in the proposition statement. Write W, W' for the corresponding reduced forms. We wish to show that W and W' are identical. The $\alpha = 0$ case is trivial (both W and W' are zero), so assume $\alpha > 0$.

First the easy direction: Every message available in the restricted investment contract is also available in the transform-bounded contract. More explicitly, take any decision $d \in M'$, and consider the function $m_d : \Omega \rightarrow \mathbb{R}^+$ defined by $m_d(\omega) = w'(d, \omega) = \alpha(u(d, \omega) + \beta(\omega))$. This is nonnegative, and its expectation over any posterior p is bounded above by $\alpha(U(p) + E_p[\beta(\omega)])$ (by definition of U), so $m_d \in M$. Therefore, for every posterior p , $W(p) \geq E_p[w'(d, \omega)]$ for any $d \in M'$. And so $W(p) \geq \max_{d \in M'} E_p[w'(d, \omega)] = W'(p)$.

Now for the reverse inequality. Consider the following subset S of \mathbb{R}^Ω : S is the set of all functions $m : \Omega \rightarrow \mathbb{R}$ such that there exists some $d \in D$ with $\alpha(u(d, \omega) + \beta(\omega)) \geq m(\omega)$ for all ω . S is closed, and the convexity assumption on the decision problem implies S is also convex. Now fix any posterior p . By definition of the reduced form W , there is some $m^* \in M$ such that $E_p[m^*(\omega)] = W(p)$.

We will show that $m^* \in S$. Suppose not. Then we can apply a strict separating hyperplane theorem to conclude the existence of nonzero $\lambda \in \mathbb{R}^\Omega$ and ξ such that

$$\sum_{\omega} \lambda_{\omega} m(\omega) < \xi \quad \text{for all } m \in S, \quad (2.15)$$

$$\sum_{\omega} \lambda_{\omega} m^*(\omega) > \xi. \quad (2.16)$$

(2.15) implies that $\lambda_{\omega} \geq 0$ for all ω . Then, we can normalize to assume that $\sum_{\omega} \lambda_{\omega} = 1$, so that λ equals some probability distribution $q \in \Delta(\Omega)$. Consider the decision $d = d(q)$, and $m \in S$ given by $m(\omega) = \alpha(u(d, \omega) + \beta(\omega))$; then (2.15) and (2.16) give us

$$E_q[\alpha(u(d, \omega) + \beta(\omega))] < \xi < E_q[m^*(\omega)].$$

But the definition of the transform-bounded contract requires that

$$E_q[m^*(\omega)] \leq \alpha(U(q) + E_q[\beta(\omega)]) = E_q[\alpha(u(d, \omega) + \beta(\omega))],$$

and so we have a contradiction.

Thus, $m^* \in S$, which means that there is some $d \in D$ satisfying

$$\alpha(u(d, \omega) + \beta(\omega)) \geq m^*(\omega) \geq 0$$

for all ω . Dividing through by $\alpha > 0$, we see that this decision d lies in the restricted decision space M' as well. Then the value of the restricted investment contract at posterior p satisfies

$$W'(p) \geq E_p[\alpha(u(d, \omega) + \beta(\omega))] \geq E_p[m^*(\omega)] = W(p).$$

Now we have $W(p) \geq W'(p)$ and $W'(p) \geq W(p)$ for every p , so we are finished. \square

Before wrapping up this section, we should comment a bit further on how to identify the parameters α, β of the optimal (transform-bounded) contract. Although it is not possible to express all the parameters in closed form, we can give an implicit characterization that will be useful for further investigation in Subsection 3.6.

Given $\beta : \Omega \rightarrow \mathbb{R}$, put

$$M(\beta) = \{m : \Omega \rightarrow \mathbb{R}^+ \mid E_p[m(\omega)] \leq U(p) + E_p[\beta(\omega)] \text{ for all } p \in \Delta(\Omega)\}$$

and then put

$$U_R(p; \beta) = \max_{m \in M(\beta)} (E_p[m(\omega)] - E_p[\beta(\omega)]).$$

(If $M(\beta) = \emptyset$ we put $U_R(p; \beta) = -\infty$.) Clearly $U_R(p; \beta) \leq U(p)$. In the convex case, we have an especially straightforward interpretation for U_R : Proposition 2.2 implies that $U_R(p; \beta) = \max_{d \in D_R(\beta)} E_p[u(d, \omega)]$, where $D_R(\beta)$ is the restricted decision space defined in (2.5).

Proposition 2.7. *The payoff guarantee from the optimal contract is equal to*

$$\max_{\substack{\alpha \in [0,1] \\ \beta : \Omega \rightarrow \mathbb{R}}} \left(\max_{(F,c) \in \mathcal{I}_0} \left((1 - \alpha) E_F[U_R(p; \beta)] - \frac{1 - \alpha}{\alpha} c \right) - \alpha \cdot E_{p_0}[\beta(\omega)] \right), \quad (2.17)$$

where, if $\alpha = 0$, we interpret the $-((1 - \alpha)/\alpha)c$ term as 0 for $c = 0$ and $-\infty$ otherwise. (In particular, the maximum in (2.17) exists.)

Moreover, when α, β attain the maximum, the corresponding transform-bounded contract is optimal.

The derivation follows our earlier calculation (2.2) that gave a payoff guarantee for linear contracts. A very similar calculation applies for any transform-bounded contract,

and shows that for any α, β , the corresponding transform-bounded contract guarantees at least the value in (2.17). In general, this is only a lower bound for the guarantee, but by retracing the steps of the proof of Theorem 2.1, we can show that it holds with equality at the maximum. The full proof is in Appendix B.

One extra note on transform-bounded contracts. The discussion so far has focused primarily on restricted investment contracts, treating transform-bounded contracts mainly as a technical stepping stone along the way. However, the notion of transform-bounded contracts is also computationally convenient. Indeed, while arbitrary contracts are very high-dimensional objects, and restricted investment contracts as we have defined them are still quite complicated (because the restricted decision space $M \subseteq D$ can be anything), transform-bounded contracts are parameterized by just $|\Omega| + 1$ numbers. Hence, Theorem 2.1 makes the search for the optimal contract potentially computationally tractable. Moreover, Proposition 2.7 makes a further step in this direction. In particular, if \mathcal{I}_0 consists of only a small number of experiments, one can try out a grid of values of α and β , and calculate the value of (2.17) in each case by considering each $(F, c) \in \mathcal{I}_0$ in turn, and then pick the best one.

3 Extensions and variations

We now consider extending the model in several directions. These extensions are independent of each other; we do not attempt to incorporate them all into a single general model. Several of these extensions are analogous to extensions of the basic robust moral hazard model from [4], and the arguments are identical to arguments in that paper. In these cases, we only describe the ideas briefly; the reader is referred to the earlier paper for details.

3.1 Participation constraint

The original model assumes that the expert will accept whatever contract is offered. Suppose, instead, that the principal needs to guarantee the expert some minimum expected payoff $\bar{U}_E > 0$ in order to hire him. Thus, the principal is restricted to contracts that satisfy $E_F[W(p)] - c \geq \bar{U}_E$ for some $(F, c) \in \mathcal{I}_0$. We maintain non-triviality: assume there exists some such contract (M, w) with $V_P(M, w) > U(p_0)$.

As in [4], adding a participation constraint does not change our main results. The compactness argument that ensures existence of an optimal contract still holds when

we add the participation constraint; and since the main step of the proof of Theorem 2.1 replaced the given contract (M, w) by a transform-bounded contract that is weakly better for both the expert and the principal, we see that if the original contract satisfies the participation constraint, so does the new one. Then the main results (Theorem 2.1, Corollary 2.3) still hold.

3.2 Smaller sets of experiments

We have written the model so that the principal evaluates contracts by their worst case over all possible IAT's $\mathcal{I} \supseteq \mathcal{I}_0$. This is perhaps an unrealistically large class of IAT's. In fact, we do not need such drastic uncertainty: As in [4], the results hold as long as every IAT of the form $\mathcal{I}_0 \cup \{(F, c)\}$, for some (F, c) , is considered possible. Moreover, we can restrict to distributions F whose support consists of at most $2|\Omega| + 2$ posteriors. This holds because the minimum in (2.6) is attained by some distribution whose support has size at most $|\Omega| + 2$, by Carathéodory's Theorem (e.g. [17, Theorem 17.1]). The extra $|\Omega|$ support points are needed to construct the auxiliary distributions used in the proof of Lemma 2.4.

3.3 Screening on technology

In general, the principal's minmax payoff typically is strictly greater than her maxmin payoff: that is, there is some payoff \bar{V}_P strictly greater than the guarantee of the maxmin-optimal contract, such that, if she knew the IAT \mathcal{I} with certainty when choosing the contract, she could achieve an expected payoff at least \bar{V}_P , no matter what $\mathcal{I} \supseteq \mathcal{I}_0$ she faced. This suggests the possibility of screening experts according to their IATs, by offering a *menu* of contracts (M, w) , which the expert chooses from before performing any experiment, so that experts with different IATs may choose different contracts. It is natural to imagine that this additional tool would be better for the principal than in the basic model, where she can only offer one contract.

It turns out, however, that the principal cannot actually achieve a better worst-case guarantee with screening than she can without screening. The argument is exactly the same as that given in [4]. If the principal can guarantee herself some payoff V_P^* using a menu of contracts, we show she can also do so using just the contract (M_0, w_0) that the expert with IAT \mathcal{I}_0 would choose from this menu. If this were not the case, then there would be some experiment (F', c') such that, under $\mathcal{I}_1 = \mathcal{I}_0 \cup \{(F', c')\}$ and contract (M_0, w_0) , the expert strictly prefers to perform experiment (F', c') , and the principal's

resulting expected payoff is less than V_p^* . Then, given the choice from the menu, the expert with IAT \mathcal{I}_1 might choose a different contract, but he would again perform experiment (F', c') , since we know his payoff under \mathcal{I}_1 is strictly higher than the best he could do using experiments in \mathcal{I}_0 . This means the principal's resulting payoff can only be even worse than under (F', c') and (M_0, w_0) (since the principal gets the same distribution over posteriors but now has to pay the expert more). This means that the menu of contracts does not guarantee V_p^* , a contradiction.

3.4 Hard information

Our model has assumed that information acquisition is private: when the expert performs an experiment, only he observes the outcome. What if instead we assume that the information acquired becomes publicly verifiable, so that the principal updates to the same posterior p as the expert, and p can be contracted on? What contract provides the optimal worst-case guarantee in this setting?

In this version of the model, there is no strategic choice of message, so no reason to have contracts depend on the realized state. Instead, we can have them depend directly on the posterior: a contract is now any continuous function $W : \Delta(\Omega) \rightarrow \mathbb{R}^+$. The model now looks like the simple robust moral hazard model of [4], since we can treat the expected gross payoff $U(p)$ as being the observed output variable, and the expert's experiment (F, c) directly determines the distribution of output via $p \sim F$. This suggests that the optimal contract should simply be affine in $U(p)$ — that is, of the form $W(p) = \alpha(U(p) + \beta)$.

This is almost correct, except that the actions potentially available to the expert do not produce all possible distributions over posteriors, only those distributions with mean p_0 . Hence, our observable-posteriors model does not collapse to the main model of [4], but rather to the extension described in Section 3.1 of that paper, with a lower bound on costs. Here all of p is observable, and the lower bound on the cost of an experiment is 0 if it has mean p_0 and ∞ otherwise. The results from that model then show that the optimal contract is of the form $W(p) = \alpha(U(p) + E_p[\beta(\omega)])$ for some function $\beta : \Omega \rightarrow \mathbb{R}$.

Since the choice of β can have no effect on incentives (as discussed in Subsection 2.3), its only role can be to relax the limited liability constraint. Therefore the optimal β is chosen by minimizing $E_{p_0}[\beta(\omega)]$ subject to $U(p) + E_p[\beta(\omega)] \geq 0$ for all p . Clearly, if we can make equality hold at p_0 , then this β is optimal; this is achieved by taking $\beta(\omega) = -u(d(p_0), \omega)$. The resulting contract has a simple interpretation: The expert is paid a share α of the principal's expected gain, conditional on the newly discovered

information, from choosing the new optimal decision rather than the decision she would have chosen in his absence.

3.5 Influencing states

What if we adopt the main model, except that we allow the expert's actions to influence the distribution of states, in addition to providing information? That is, we drop the requirement $E_F[p] = p_0$ from the definition of an experiment, so an experiment is now any element of $\Delta(\Delta(\Omega)) \times \mathbb{R}^+$; where we originally imposed a full-support common prior p_0 , we now simply require that for each $(F, c) \in \mathcal{I}_0$, $E_F[p]$ should have full support. The rest of the model is left unchanged. (Zermeño [21] also considers allowing the expert to influence the distribution. See also [19].)

We can repeat the analysis, and find that there is an optimal contract that is transform-bounded with the adjustment term $\beta(\omega)$ constant across all states ω . In the convex case, this means likewise that a restricted investment contract with β constant across all states is optimal. (For an overview of what changes in the analysis: Lemma 2.4 is the same as before, except that the constraint $E_F[p] = p_0$ is dropped from (2.6). Lemma 2.5 is the same but with β constant across states. The separation argument used to prove Lemma 2.5 is now simply done in $\mathbb{R} \times \mathbb{R}$, so that the λ_ω terms in inequalities (2.9) and (2.10) are absent.)

3.6 Restricted versus unrestricted investment

In this subsection we briefly consider the question of when the optimal contract is an *unrestricted* investment contract, with $M = D$. Such contracts have a natural interpretation, as delegating the decision directly to the expert, who is then paid a fraction α of the (gross) payoff plus the state-by-state adjustment β . If it also happens that the worst possible payoff in each state is zero ($\min_d u(d, \omega) = 0$ for each ω), then the interpretation is even simpler: the optimum is a linear contract. Identifying situations where these contracts are optimal helps to bridge the gap between our model and some of the prior principal-expert literature [9, 14] which assumed that the decision must be delegated to the expert and compensation could depend only on the realized payoff. These contracts also seem more in line with reality, where linear incentive contracts based on realized payoffs are quite widespread (see e.g. the references in [2, pp. 763–4]) whereas our notion of restricted investment contracts seems less familiar.

We cannot give a complete characterization of when restriction is or is not optimal, but we can give some partial results. Roughly, if the principal's optimal decision is extremely sensitive to her posterior belief, then unrestricted investment contracts are generally not optimal: There are some extreme decisions that are very unlikely to come into play, and trimming them out of the message space is useful in relaxing limited liability. Conversely, if the optimal decision is not so sensitive, then unrestricted investment contracts may be optimal. In the special case where the principal's decision problem is binary, unrestricted investment contracts are always optimal.

For formal statements, we maintain several assumptions. We assume throughout this subsection that the zero contract is not optimal. We also assume that in every state, not all decisions give the same payoff. Furthermore, for Proposition 3.1 below, we assume that the decision problem is convex; thus Proposition 2.2 applies, and we can describe the optimal contract either as a transform-bounded contract or as a restricted investment contract with the same parameters α, β . Finally, we assume that dominated decisions have been eliminated a priori: For any distinct $d, d' \in D$, there exists ω such that $u(d, \omega) > u(d', \omega)$. This eliminates uninteresting cases where there is some decision giving extremely low payoffs in every state, which would make an unrestricted investment contract extremely costly (via limited liability forcing $\beta(\omega)$ to be high) even though the expert would never invest in that decision.

Also, we need a little more terminology. Say that a decision d is ω -*extremal* if it minimizes $u(d, \omega)$ among all decisions; and it is *extremal* if it is ω -extremal for some ω .

Now we can state our main result relating the structure of the principal's decision problem to the optimality of unrestricted investment:

Proposition 3.1. *Fix the decision problem (D, Ω, u) , satisfying convexity.*

- (a) *Suppose that, for every decision d , there is only one posterior p for which d is optimal for the principal. Then, for any p_0 and \mathcal{I}_0 , there is a properly restricted investment contract (i.e. $M \subset D$ strictly) that is optimal.*
- (b) *Suppose that, for every state ω , there is an ω -extremal decision $d \in D$, and a nonempty open set of posteriors $P \subseteq \Delta(\Omega)$, such that d is optimal for any posterior $p \in P$. Then, for any prior p_0 , it is possible to choose \mathcal{I}_0 so that an unrestricted investment contract is optimal. We can do this while satisfying the non-triviality and full-support assumptions.*
- (c) *Suppose that there exists a nonempty open set of posteriors $P \subseteq \Delta(\Omega)$ such that there*

is no one decision that is optimal for all $p \in P$, and such that extremal decisions are not optimal for any $p \in P$. Then, it is possible to choose $p_0 \in P$ and \mathcal{I}_0 , satisfying non-triviality and full support, such that a properly restricted investment contract is optimal.

To aid in understanding this proposition, consider Figure 3, which shows a couple of possible decision problems. Each panel shows a decision problem with state space $\Omega = \{\omega_1, \omega_2\}$; each decision is represented by the pair of payoffs $(u(d, \omega_1), u(d, \omega_2))$. The thick curve depicts the set of available decisions D . Within each state, payoffs have been normalized so that the payoff of the worst decision is 0. A posterior p can be represented as a vector $(p(\omega_1), p(\omega_2))$; the best decision for this posterior is then the one where the tangent to the decision set is orthogonal to the vector.




Figure 3: Example decision problems.

Panel (a) of the figure depicts part (a) of the proposition. In this case, for every possible posterior there is a different optimal decision, and the proposition tells us that properly restricted investment contracts are always optimal.

In panel (b), by contrast, the leftmost decision (where the tangent line is illustrated) is an ω_1 -extremal decision that is optimal for all posteriors that put sufficiently low weight on ω_1 . Hence, part (b) of the proposition tells us that unrestricted investment contracts may be optimal in this case.

For a more general example in the same spirit, suppose that the decision space is the efficient frontier of the convex hull of a finite set of decisions. That is, suppose we begin

with some finite set of decisions D_0 ; we identify each decision d with the corresponding vector of payoffs in \mathbb{R}^Ω , and form the convex hull D_1 ; and then D consists of all $d \in D_1$ that are not weakly dominated by some other vector in D_1 . In this case, one can show, using a separating hyperplane argument, that the condition in (b) is satisfied. So at least for some choices of \mathcal{I}_0 , an unrestricted investment contract is optimal. This, in turn, is the same as an unrestricted investment contract without allowing mixed decisions — simply letting the expert choose from among all decisions in D_0 that are not weakly dominated by some mixture. We omit the details here.

Finally, panel (b) of the figure also illustrates part (c) of the proposition: The smooth bump on the right side represents a set of non-extremal decisions, which correspond to an open interval of posteriors, with the optimal decision varying as the posterior moves. Thus, by part (c), properly restricted investment contracts may also be optimal for this decision problem.

Are there simple conditions on the decision problem sufficient to ensure that unrestricted investment contracts are optimal? The logic of part (c) is quite general, and shows that this cannot happen unless “most” decisions are extremal. But there is one important special case where the conclusion does hold: when the principal’s decision problem is binary.

Proposition 3.2. *Suppose that D consists of just two decisions. Then, there is an unrestricted investment contract that is optimal.*

The proofs of the two propositions from this subsection are in Appendix C. Proposition 3.1 can be proven by analyzing formula (2.17) for the optimal β . Proposition 3.2 holds because, in the binary-decision case, restricted and unrestricted investment contracts are actually the same under a suitable change of parameters.

4 Conclusion

We wrap up by briefly recapitulating our results and putting them in context. We considered a principal-expert model, with risk-neutrality and limited liability, and ex-post revelation of the state of nature. We assumed unquantifiable uncertainty about the expert’s information acquisition technology, and sought out contracts that are robust to this uncertainty, as expressed by a maxmin objective. This led quite generally to a novel form of contracts being optimal: transform-bounded contracts, or equivalently (under a convexity assumption) restricted investment contracts, in which the expert chooses to invest

in one of a restricted subset of less-risky decisions, and is paid proportionally to the payoff of his invested-in decision in the realized state, even though the decision that the principal actually makes may be different. The result reflects the intuition that linear payment rules are robust because they tightly align the principal’s and expert’s payoffs, turning an expected-payoff guarantee for the expert into a guarantee for the principal. But there is the added wrinkle that by prohibiting investment in risky decisions, the principal can relax the limited liability constraints and so pay the expert less.

The direct interpretation of our results is that they show how one can go about optimally providing incentives to experts in uncertain environments. They offer qualitative insights into the shape of an optimal contract, and they also show that it is characterized by a small number of parameters, which can be helpful in actually computing the optimum.

More broadly, this work illustrates the value of the worst-case methodology in contract design. The standard Bayesian version of the principal-expert problem, even under risk-neutrality and limited liability as we have assumed here, seems to be intractable except in special cases [20]. The corresponding maxmin version offers traction quite generally — without any functional-form assumptions on the known information acquisition technology — by extending the linear separation methods of the earlier paper [4]. As in that earlier paper, we also saw that the model can be extended in various directions without changing the underlying technical machinery. The worst-case environments are not extreme. And the analysis has led us to a new qualitative insight — that a contract tying the expert’s payoff closely to the principal’s can be improved on by cutting out extreme messages, thereby relaxing limited liability — that is not knife-edge sensitive to the worst-case formulation of the objective.

It seems likely that the methods used here generalize to other kinds of robust moral hazard contracting problems. A natural task for future work is to find as large a class of models as possible that are amenable to the same proof approach — writing a linear program to identify the worst case for a given contract, using duality to infer a linear bound relating the principal’s and agent’s payoffs, and from there producing a new contract that improves on the original. At the same time, the worst-case modeling framework may also bring traction to difficult problems by other routes. For example, ongoing work by the author [5] takes a worst-case approach to formulate a class of multidimensional screening problems with a simple solution. There, the proof also works via duality, but rather than starting with any given mechanism and finding an improvement, it involves a direct (and more involved) construction of an adversarial environment for the principal.

A On direct reporting of posteriors

Here we elaborate formally the revelation-principle-style argument outlined in Subsection 2.2: when looking for optimal contracts, we may assume that the expert directly reports his posterior as one component of his message.

First we should clarify why this is not trivial. In fact, with arbitrary contracts, where multiple posteriors may pool on the same message, it is not even clear how to write down the model of subsequent behavior. The principal should choose her decision based on her updated belief about the posterior, given the expert's strategy. This means that, to describe the principal's behavior, we quickly run into the thorny problem of Bayesian updating by a non-Bayesian player. Moreover, this difficulty leads to trouble in describing the expert's strategy as well, if we want to maintain our assumption that the expert breaks ties by choosing the message that is best for the principal.

In this appendix, we will avoid this issue by showing that any contract, together with *any* decision strategy for the principal, is weakly outperformed by a contract in which the expert reports his posterior truthfully as part of his message, and the principal takes the corresponding optimal decision. And to avoid the ambiguity in the expert's tie-breaking, we instead take a more conventional approach to the definition of contracts: requiring a contract to specify a recommended strategy for the expert, which must be incentive-compatible, in addition to specifying the compensation rule. We then think of the model formulated in the main paper as expressing the same concept, but in a more compact notation; we elaborate at the end of this appendix.

Let \mathcal{IAT} denote the set of all possible IAT's \mathcal{I} with $\mathcal{I}_0 \subseteq \mathcal{I}$. We properly define a *contract* to consist of four parts:

- a compact message space M ;
- a continuous payment function $w : M \times \Omega \rightarrow \mathbb{R}^+$;
- an experiment strategy $\sigma : \mathcal{IAT} \rightarrow \Delta(\Delta(\Omega)) \times \mathbb{R}^+$, such that $\sigma(\mathcal{I}) \in \mathcal{I}$ for each \mathcal{I} ;
- a reporting strategy $\rho : \mathcal{IAT} \times \Delta(\Omega) \rightarrow M$, which should be measurable in its second argument.

Thus (σ, ρ) describes the strategy recommended for the expert. (We could also allow for mixed strategies; this would change nothing below.)

We want to impose that the specified strategy for the expert be incentive-compatible. This requires defining payoffs. So given (M, w, σ, ρ) , define

$$W(p) = \max_{m \in M} E_p[w(m, \omega)], \quad V_E(M, w | \mathcal{I}) = \max_{(F, c) \in \mathcal{I}} (E_F[W(p)] - c)$$

just as in the main paper. We now say that the contract (M, w, σ, ρ) is *incentive-compatible* if

- for all \mathcal{I} , the experiment $(F, c) = \sigma(\mathcal{I})$ satisfies $E_F[W(p)] - c = V_E(M, w | \mathcal{I})$; and moreover
- for all \mathcal{I} and all p ,

$$E_p[w(\rho(\mathcal{I}, p), \omega)] = W(p).$$

Now define a *decision rule* for the principal to be any measurable $r : M \rightarrow D$. Given the contract and the decision rule, we define the principal's payoff from a given \mathcal{I} to be

$$V_P(M, w, \sigma, \rho; r | \mathcal{I}) = E_F[E_p[u(r(\rho(\mathcal{I}, p)), \omega)] - W(p)]$$

where F is the distribution chosen by $\sigma(\mathcal{I})$; and then define

$$V_P(M, w, \sigma, \rho; r) = \inf_{\mathcal{I} \supseteq \mathcal{I}_0} V_P(M, w, \sigma, \rho; r | \mathcal{I})$$

as before.

Now we can give the formal justification for the modeling of contracts in the main text, including the assumption that the expert directly communicates his posterior.

Proposition A.1. *Let (M, w, σ, ρ) be any incentive-compatible contract, and $r : M \rightarrow D$ a decision rule. Then define a new contract (M', w', σ', ρ') and decision rule r' by taking:*

- $M' = M \times \Delta(\Omega)$;
- $w'((m, p), \omega) = w(m, \omega)$;
- $\rho'(\mathcal{I}, p) = (\rho(\mathcal{I}, p), p)$;
- $r'(m, p) = d(p)$ (recall this was defined as $\arg \max_{d \in D} E_p[u(d, \omega)]$); and
- $\sigma'(\mathcal{I})$ is taken to be any $(F, c) \in \mathcal{I}$ that lexicographically maximizes $E_F[W(p)] - c$ and then $E_F[U(p) - W(p)]$, where W is the reduced form of the original contract.

Then the new contract is also incentive-compatible, and (with r') has at least as high a payoff guarantee V_P as the original contract (with r).

Proof: At any posterior p , (m, p) is an optimal report under the new contract iff m was optimal under the old contract; and the reduced form of the new contract is the same as the old, $W'(p) = W(p)$. This in turn implies that the new contract has the same V_E as the old one. Incentive-compatibility for the new contract immediately follows.

As for the comparison of payoff guarantees, consider any \mathcal{I} . Let $(F, c) = \sigma(\mathcal{I})$ and $(F', c') = \sigma'(\mathcal{I})$. Since (F', c') is an optimal experiment, as is (F, c) (by incentive-compatibility of the original contract), lexicographic maximization in the definition of σ' implies that $E_{F'}[U(p) - W(p)] \geq E_F[U(p) - W(p)]$. Thus,

$$\begin{aligned}
V_P(M', w', \sigma', \rho'; r'|\mathcal{I}) &= E_{F'} [E_p[u(r'(\rho'(\mathcal{I}, p)), \omega)] - W'(p)] \\
&= E_{F'} [E_p[u(d(p), \omega)] - W(p)] \\
&= E_{F'} [U(p) - W(p)] \\
&\geq E_F [U(p) - W(p)] \\
&\geq E_F [E_p[u(r(\rho(\mathcal{I}, p)), \omega)] - W(p)] \\
&= V_P(M, w, \sigma, \rho; r|\mathcal{I}).
\end{aligned}$$

Taking the infimum over \mathcal{I} gives the result. \square

Now we can think of the modeling approach taken in the main text simply as using (M, w) as a notational shorthand for the full contract (M', w', σ', ρ') (and decision rule r') above. Indeed, the choice of experiment and the principal's payoff as defined in the main text agree with the definitions given here for that full contract. Since we are concerned specifically with identifying optimal contracts for the principal, it is sufficient to look at these contracts that are constructed in Proposition A.1 (with the appropriate decision rule).

B Omitted proofs for main optimal-contract results

Proof of Lemma 2.4:

First, for any $\mathcal{I} \supseteq \mathcal{I}_0$, and any experiment (F, c) chosen by the expert, $E_F[W(p)] \geq E_F[W(p)] - c \geq V_E(M, w|\mathcal{I}_0)$, so F satisfies the constraints in (2.6), and hence $V_P(M, w|\mathcal{I})$ is at least the indicated minimum. This holds for any \mathcal{I} , so $V_P(M, w)$ is at least the minimum in (2.6). (Note that this minimum is well-defined.)

To prove the reverse inequality, we begin by defining the affine function $Z : \Delta(\Omega) \rightarrow \mathbb{R}^+$ by $Z(p) = \sum_{\omega} p(\omega)W(\delta_{\omega})$. By convexity, $W(p) \leq Z(p)$ for every p . Now let F be a distribution attaining the minimum in (2.6). Suppose F places positive probability on posteriors p such that $W(p) < Z(p)$. Then we have

$$E_F[W(p)] < E_F[Z(p)] = Z(p_0)$$

where the equality holds because Z is affine. For small $\epsilon > 0$, define a distribution F' as follows: with probability $1 - \epsilon$, F' chooses a posterior according to F ; with the remaining probability ϵ , F' picks a state $\omega \sim p_0$ and gives posterior δ_{ω} . Evidently, $E_{F'}[p] = p_0$, and

$$E_{F'}[W(p)] = (1 - \epsilon)E_F[W(p)] + \epsilon Z(p_0) > E_F[W(p)] \geq V_E(M, w|\mathcal{I}_0).$$

So if the IAT is $\mathcal{I} = \mathcal{I}_0 \cup \{(F', 0)\}$, the expert's unique optimal choice of experiment is $(F', 0)$. The principal's expected payoff $V_P(M, w|\mathcal{I})$ is then

$$E_{F'}[U(p) - W(p)] = (1 - \epsilon)E_F[U(p) - W(p)] + \epsilon E_{p_0}[U(\delta_{\omega}) - W(\delta_{\omega})].$$

By taking $\epsilon \rightarrow 0$, we see that the principal cannot be guaranteed a payoff higher than $E_F[U(p) - W(p)]$, which is exactly the amount in (2.6).

Also, if $W(p) = Z(p)$ throughout the support of F , but $E_F[W(p)] > V_E(M, w|\mathcal{I}_0)$ strictly, then we can give a similar argument by simply taking $\mathcal{I} = \mathcal{I}_0 \cup \{(F, 0)\}$.

This leaves us with the case where

$$V_E(M, w|\mathcal{I}_0) = E_F[W(p)] = E_F[Z(p)] = Z(p_0).$$

For this to happen, it must be that whatever experiment (F_0, c_0) is chosen under \mathcal{I}_0 satisfies $W(p) = Z(p)$ throughout the support of F_0 , and $c_0 = 0$. However, in this case, the expert would be willing to perform the same experiment under the zero contract, so

$$V_P(0) \geq E_{F_0}[U(p)] > E_{F_0}[U(p) - W(p)] = V_P(M, w|\mathcal{I}_0) \geq V_P(M, w).$$

(The strict inequality follows from the assumptions that W is nonzero, and F_0 has mean p_0 which has full support.) This contradicts eligibility, so this case cannot happen.

Finally, let F be a distribution attaining the minimum in (2.6). Then certainly

$$E_F[U(p) - W(p)] = V_P(M, w) > U(p_0) \geq U(p_0) - W(p_0).$$

If $E_F[W(p)] > V_E(M, w|\mathcal{I}_0)$ strictly, then define another distribution F' by drawing a posterior from F with probability $1 - \epsilon$, and placing the remaining probability mass ϵ on p_0 . For small ϵ , F' still satisfies the constraints of (2.6), and

$$E_{F'}[U(p) - W(p)] = (1 - \epsilon)E_F[U(p) - W(p)] + \epsilon(U(p_0) - W(p_0)) < E_F[U(p) - W(p)],$$

contradicting minimality for F . □

Proof of Lemma 2.6:

Let $\bar{U} = \max_p U(p)$, and $\bar{W} = (\bar{U} - U(p_0))/(\min_\omega p_0(\omega))$. We may restrict attention to contracts whose reduced form satisfies $W(\delta_\omega) \leq \bar{W}$ for all ω . To see this, note that if there were some message guaranteeing the expert payoff higher than $(\bar{U} - U(p_0))/p_0(\omega)$ in some state ω , then no matter what experiment the expert performs, he can always force the principal to pay more than $\bar{U} - U(p_0)$ in expectation (by just always sending this message), so the principal's expected payoff must be less than $U(p_0)$, and the principal is worse off than by not hiring the expert.

Since W is convex, this restriction implies $W(p) \leq \bar{W}$ for all p . Now say that a function $W : \Delta(\Omega) \rightarrow \mathbb{R}^+$ is a *reduced-form contract* if it is the reduced form of some contract. We make two claims:

- Claim 1: The set of reduced-form contracts $W : \Delta(\Omega) \rightarrow [0, \bar{W}]$ is compact in the sup-norm topology.
- Claim 2: V_P is upper semi-continuous on the set of reduced-form contracts (with respect to the sup-norm topology).

Together, these claims imply that V_P attains a maximum over the reduced-form contracts whose values never exceed \bar{W} , which is then a global maximum, as needed.

To prove Claim 1, let W_1, W_2, \dots be a sequence of reduced-form contracts taking values in $[0, \bar{W}]$. Note that each W_k must be a Lipschitz function with constant \bar{W} (relative to the L_1 metric on $\Delta(\Omega)$). By passing to a subsequence, we may assume that W_k converges pointwise at each *rational* point $p \in \Delta(\Omega)$. Let $W_\infty(p) = \lim_k W_k(p)$ for each such p . Define M as the set of all functions $m : \Omega \rightarrow [0, \bar{W}]$ such that $E_p[m(\omega)] \leq W_\infty(p)$ for each rational p . Notice that M is a nonempty, compact subset of \mathbb{R}^Ω . Define $w : M \times \Omega \rightarrow [0, \bar{W}]$ by $w(m, \omega) = m(\omega)$. Then (M, w) is a contract, with reduced form $W(p) = \max_{m \in M} E_p[m(\omega)]$.

For each rational p , we will show $W(p) = W_\infty(p)$. The direction $W(p) \leq W_\infty(p)$ is immediate from the definition of W . For the reverse inequality, fix p . For each k ,

since W_k is a reduced-form contract, there is some affine $m_k : \Delta(\Omega) \rightarrow [0, \overline{W}]$ such that $m_k(p) = W_k(p)$ and $m_k(p') \leq W_k(p')$ for all other rational p' . There is some subsequence along which the m_k converge to some m_∞ . Then, $m_\infty(p) = W_\infty(p)$, and $m_\infty(p') \leq W_\infty(p')$ for each other rational p' , so that $m_\infty \in M$. Therefore $W(p) \geq W_\infty(p)$, and equality follows.

Now we claim that $W_k \rightarrow W$ in sup norm. If not, there exists some $\epsilon > 0$ and a subsequence of k 's and points p_k along which $|W_k(p_k) - W(p_k)| > \epsilon$. Again by taking a subsequence, we may assume the p_k converge to some point p . Now, pick a rational point q such that the L_1 distance between p and q is less than $\epsilon/4\overline{W}$. Then, for k high enough, $|W_k(q) - W(q)| \leq \epsilon/4$ (because $W(q) = W_\infty(q)$ by the previous paragraph), and for any reduced-form contract with values at most \overline{W} , its values at p_k and q differ by at most $\epsilon/4$. Therefore,

$$|W_k(p_k) - W(p_k)| \leq |W_k(p_k) - W_k(q)| + |W_k(q) - W(q)| + |W(q) - W(p_k)| \leq \frac{3\epsilon}{4},$$

a contradiction.

This shows that the set of reduced-form contracts taking values in $[0, \overline{W}]$ is sequentially compact, proving Claim 1.

For Claim 2, suppose we have a convergent sequence of reduced-form contracts $W_k \rightarrow W$. We need to show that $V_P(W) \geq \limsup_k V_P(W_k)$. By passing to a subsequence, we may assume $V_P(W_k)$ converges (to the lim sup of the original sequence). Fix any IAT \mathcal{I} , and let (F_k, c_k) be the expert's chosen experiment under contract W_k . Then, again by taking a subsequence, we may assume that (F_k, c_k) converges to some limit (F, c) . It follows that $(F, c) \in \mathcal{I}$, and since

$$E_F[W(p)] - E_{F_k}[W_k(p)] = (E_F[W(p)] - E_{F_k}[W(p)]) + (E_{F_k}[W(p)] - E_{F_k}[W_k(p)]) \rightarrow 0$$

(the first parenthesized expression goes to 0 by weak convergence of F_k , and the second by sup-norm convergence of W_k), we can conclude that (F, c) is an optimal experiment for the expert under W and \mathcal{I} : If there were some strictly better experiment (F', c') , then this experiment would also be preferred to (F_k, c_k) under W_k for high enough k , a contradiction.

This same double-convergence argument also implies that

$$E_F[U(p) - W(p)] - E_{F_k}[U(p) - W_k(p)] \rightarrow 0$$

from which

$$V_P(W|\mathcal{I}) \geq E_F[U(p) - W(p)] = \lim_k E_{F_k}[U(p) - W_k(p)] = \lim_k V_P(W_k|\mathcal{I}) \geq \limsup_k V_P(W_k).$$

Since this holds for all \mathcal{I} , we have

$$V_P(W) \geq \limsup_k V_P(W_k)$$

which proves Claim 2. □

Proof of Proposition 2.7: Let (M, w) be the transform-bounded contract with parameters α, β . We will show that its payoff guarantee satisfies

$$V_P(M, w) \geq \max_{(F, c) \in \mathcal{I}_0} \left((1 - \alpha) E_F[U_R(p; \beta)] - \frac{1 - \alpha}{\alpha} c \right) - \alpha \cdot E_{p_0}[\beta(\omega)]. \quad (\text{B.1})$$

We will also show that there are some α, β for which the corresponding transform-bounded contract is optimal and (B.1) is an equality. This will imply the result. (Note that once we show this, if there exist additional maximizers α', β' for the right side of (B.1), then (B.1) must be an equality for those parameters as well: otherwise the contract with parameters α', β' is strictly better than the one with parameters α, β , contradicting optimality.)

First suppose $\alpha = 0$. If there does not exist any $(F, 0) \in \mathcal{I}_0$, the right side of (B.1) is $-\infty$; and if there is such an experiment, the right side is $\max_{(F, 0) \in \mathcal{I}_0} E_F[U_R(p; \beta)] \leq \max_{(F, 0) \in \mathcal{I}_0} E_F[U(p)]$ which is exactly the payoff guarantee from that contract. Also, if a zero contract is optimal, we can obtain equality in (B.1) by taking $\alpha = 0$ and β large enough so that $U_R(p; \beta) = U(p)$.

Now suppose $\alpha > 0$. The reduced form of the transform-bounded contract is precisely

$$W(p) = \alpha(U_R(p; \beta) + E_p[\beta(\omega)]). \quad (\text{B.2})$$

Once the expert attains posterior p , his expected payment is $W(p) \leq \alpha(U(p) + E_p[\beta(\omega)])$, while the principal's is

$$\begin{aligned} U(p) - W(p) &\geq (1 - \alpha)U_R(p; \beta) - \alpha E_p[\beta(\omega)] \\ &= \frac{1 - \alpha}{\alpha} W(p) - E_p[\beta(\omega)]. \end{aligned}$$

Therefore, whatever experiment (F, c) the expert performs, the principal's expected payoff

$V_P(M, w|\mathcal{I})$ is

$$\begin{aligned} E_F[U(p) - W(p)] &\geq E_F\left[\frac{1-\alpha}{\alpha}W(p) - E_p[\beta(\omega)]\right] \\ &= \frac{1-\alpha}{\alpha}E_F[W(p)] - E_{p_0}[\beta(\omega)]. \end{aligned}$$

Since

$$E_F[W(p)] \geq E_F[W(p)] - c \geq V_E(M, w|\mathcal{I}_0) = \max_{(F,c) \in \mathcal{I}_0} (E_F[W(p)] - c),$$

we can plug in to obtain

$$V_P(M, w|\mathcal{I}) \geq \frac{1-\alpha}{\alpha} \cdot \max_{(F,c) \in \mathcal{I}_0} (E_F[W(p)] - c) - E_{p_0}[\beta(\omega)].$$

Now plugging in from (B.2) and rearranging gives exactly the right side of (B.1). Since this applies for all \mathcal{I} , (B.1) follows.

Now consider an optimal contract (M, w) . Assume it is nonzero, since we already dealt with the case where a zero contract is optimal. Lemma 2.4 identifies its payoff guarantee to the principal; let F^* be the worst-case distribution given by that lemma. In the proof of Theorem 2.1, we obtained parameters α, β , such that the corresponding transform-bounded contract (M', w') , with reduced form W' , satisfies:

- $W'(p) \geq W(p)$ for all p , hence $V_E(M', w'|\mathcal{I}_0) \geq V_E(M, w|\mathcal{I}_0)$;
- $V_P(M', w') \geq V_P(M, w)$ (from which (M', w') is again an optimal contract); and
- $\alpha(U(p) + E_p[\beta(\omega)]) \geq W'(p) \geq W(p)$, with equality on the support of F^* (this holds because in the proof of Lemma 2.5 we showed equality in (2.7) throughout the support of the worst-case distribution).

The proof of Theorem 2.1 also showed that we cannot have $V_E(M', w'|\mathcal{I}_0) > V_E(M, w|\mathcal{I}_0)$ strictly, because that would imply $V_P(M', w') > V_P(M, w)$, contradicting optimality of the original contract. Hence $V_E(M', w'|\mathcal{I}_0) = V_E(M, w|\mathcal{I}_0)$, and so F^* satisfies the constraints of (2.6) for the new contract (M', w') . Since $E_{F^*}[U(p) - W'(p)] = E_{F^*}[U(p) - W(p)]$, F^* must actually attain the minimum in (2.6) (for the new contract) — otherwise the minimum would be strictly lower, and we would have $V_P(M', w') < V_P(M, w)$, a contradiction.

Now we show that (B.1) is an equality for the new contract (M', w') . From (2.6) we

have

$$\begin{aligned}
V_P(M', w') &= E_{F^*}[U(p) - W'(p)] \\
&= E_{F^*}[(1 - \alpha)U(p) - \alpha E_p[\beta(\omega)]] \\
&= \frac{1 - \alpha}{\alpha} E_{F^*}[W'(p)] - E_{p_0}[\beta(\omega)] \\
&= \frac{1 - \alpha}{\alpha} V_E(M', w' | \mathcal{I}_0) - E_{p_0}[\beta(\omega)] \\
&= \frac{1 - \alpha}{\alpha} \max_{(F, c) \in \mathcal{I}_0} (E_F[W'(p)] - c) - E_{p_0}[\beta(\omega)].
\end{aligned}$$

Now plugging in $W'(p) = \alpha(U_R(p; \beta) + E_p[\beta(\omega)])$ gives us equality in (B.1) for the new contract. \square

C Proofs for restricted versus unrestricted investment

Here we prove Proposition 3.1. We first develop some preliminary tools. Recall Proposition 2.7, which gives the formula (2.17) for the parameters of the optimal contract, and solve out for α : given β and (F, c) , the α that attains the maximum is $\alpha = \sqrt{c / (E_F[U_R(p; \beta)] + E_{p_0}[\beta(\omega)])}$. Plugging in this value of α , (2.17) turns into

$$\max_{\substack{\beta: \Omega \rightarrow \mathbb{R} \\ (F, c) \in \mathcal{I}_0}} \left(E_F[U_R(p; \beta)] + c - 2\sqrt{c(E_F[U_R(p; \beta)] + E_{p_0}[\beta(\omega)])} \right). \quad (\text{C.1})$$

Moreover, the optimal contract uses the corresponding value of β .

Let β_0 be given by $\beta_0(\omega) = -\min_{d \in D} u(d, \omega)$, for each ω . The optimal β certainly satisfies $\beta(\omega) \leq \beta_0(\omega)$ for each ω , since otherwise β can be decreased without changing the restricted decision space M in the corresponding transform-bounded contract, thus saving money without affecting the expert's incentives for information acquisition. Thus, in the maximization (C.1), we can restrict the domain to $\beta \leq \beta_0$ componentwise.

The resulting contract is unrestricted investment iff $\beta = \beta_0$. Thus, if the maximum in (C.1) is not attained at $\beta = \beta_0$, then there is a properly restricted investment contract that is optimal; if it is attained at $\beta = \beta_0$, then an unrestricted investment contract must be optimal.

Lemma C.1. $U_R(p; \beta)$ is concave and increasing in β .

Proof: For values β and β' , let $m \in M(\beta)$ and $m' \in M(\beta')$ attain the respective maxima in the definitions of $U_R(p; \beta), U_R(p; \beta')$. Suppose $\beta'' = \lambda\beta + (1 - \lambda)\beta'$ with $\lambda \in [0, 1]$. It is immediate that $m'' = \lambda m + (1 - \lambda)m' \in M(\beta'')$, from which

$$U_R(p; \beta'') \geq E_p[m''(\omega)] - E_p[\beta''(\omega)] = \lambda U_R(p; \beta) + (1 - \lambda)U_R(p; \beta').$$

And to see that U_R is increasing in β : If $\beta \leq \beta'$ componentwise, and m attains the maximum for β , then $m'(\omega) = m(\omega) + \beta'(\omega) - \beta(\omega)$ is in $M(\beta')$ and gives the same value for the maximand in defining $U_R(p; \beta')$; thus $U_R(p; \beta') \geq U_R(p; \beta)$. \square

It follows that, for each posterior p and each state ω , the one-sided partial derivative

$$\psi(p; \omega) = \left. \frac{\partial^- U_R}{\partial \beta(\omega)} \right|_{(p; \beta_0)}$$

is well-defined and nonnegative. We will use this throughout the proof below.

Proof of Proposition 3.1:

Part (a). Fix any state ω^* . We first show that, under the assumptions of this part, $\psi(p, \omega^*) = 0$ for all p . Suppose not; pick p^* with $\psi(p^*, \omega^*) > 0$. Write $\kappa = \psi(p^*, \omega^*)$. Then, for each $\epsilon \geq 0$, if $\beta_\epsilon : \Omega \rightarrow \mathbb{R}$ is defined by $\beta_\epsilon(\omega) = \beta_0(\omega)$ for $\omega \neq \omega^*$ and $\beta_\epsilon(\omega^*) = \beta_0(\omega^*) - \epsilon$, we get $U_R(p^*; \beta_\epsilon) \leq U_R(p^*; \beta_0) - \kappa\epsilon$. (This follows by concavity of U_R .) That is, for any decision d satisfying $u(d, \omega^*) \geq -\beta_0(\omega^*) + \epsilon$, we have $E_{p^*}[u(d, \omega)] \leq E_{p^*}[u(d(p^*), \omega)] - \kappa\epsilon$. So for any $d \in D$, we can take $\epsilon = u(d, \omega^*) + \beta_0(\omega^*)$ and obtain

$$E_{p^*}[u(d, \omega)] \leq E_{p^*}[u(d(p^*), \omega)] - \kappa(u(d, \omega^*) + \beta_0(\omega^*)). \quad (\text{C.2})$$

Plugging in $d = d(p^*)$, transposing, and dividing by $\kappa > 0$ gives $u(d(p^*), \omega^*) + \beta_0(\omega^*) \leq 0$. But since $u(d, \omega^*) + \beta_0(\omega^*) \geq 0$ for all d by definition of β_0 , we must have equality: $\beta_0(\omega^*) = -u(d(p^*), \omega^*)$. Then, (C.2) rewrites as

$$E_{p^*}[u(d, \omega)] + \kappa u(d, \omega^*) \leq E_{p^*}[u(d(p^*), \omega)] + \kappa u(d(p^*), \omega^*) \quad (\text{C.3})$$

for all $d \in D$.

Hence, if we define $q(\omega) = p^*(\omega)$ for each $\omega \neq \omega^*$ and $q(\omega^*) = p^*(\omega^*) + \kappa$, (C.3) tells us that $d(p^*)$ maximizes $\sum_\omega q(\omega)u(d, \omega)$ over all $d \in D$. Hence, scaling down q by a factor of $1 + \kappa$ gives another posterior belief, distinct from p^* , for which $d(p^*)$ is optimal. This contradicts the assumption that each decision can only be optimal for one belief. (This argument would fail if p^* puts weight 1 on ω^* , since then $q/(1 + \kappa)$ is actually equal to

p^* . But this cannot be the case: $d(p^*)$ is the best decision for p^* , but we saw above that it is the worst decision in state ω^* .)

This contradiction completes the proof that $\psi(p, \omega) = 0$ for all p and ω .

Now, the concavity of U_R implies that $(U_R(p; \beta_0) - U_R(p; \beta_\epsilon))/\epsilon$ decreases as $\epsilon \rightarrow 0$, so we can invoke the monotone convergence theorem to conclude that

$$\frac{d^+}{d\epsilon} E_F[U_R(p; \beta_\epsilon)] \Big|_{\epsilon=0} = E_F \left[\frac{d^+}{d\epsilon} U_R(p; \beta_\epsilon) \Big|_{\epsilon=0} \right] = 0 \quad \text{for any } F \in \Delta(\Delta(\Omega)).$$

Now consider any \mathcal{I}_0 , and consider the maximand in (C.1), corresponding to the optimal contract. The maximizing choice of (F, c) must have $c > 0$, since otherwise the zero contract is optimal, contrary to assumption. Suppose that the maximizing choice of β is β_0 . Consider replacing β by β_ϵ . Then the derivative of the maximand in (C.1) with respect to ϵ (at 0) is equal to

$$p_0(\omega) \sqrt{\frac{c}{E_F[U_R(p; \beta_0)] + E_{p_0}[\beta_0(\omega)]}} > 0.$$

This contradicts maximality at β_0 . So the maximum is not attained at β_0 after all, and so our optimal contract is a properly restricted investment contract.

Part (b). In this case, we first choose, for each state ω , a posterior p_ω lying in the corresponding open set P . We claim that $\psi(p_\omega, \omega) > 0$ for each state ω .

This is essentially the mirror image of the argument at the beginning of (a). Fix any state ω^* , and put $p^* = p_{\omega^*}$. By assumption, there is an ω^* -extremal decision $d^* = d(p^*)$ that is optimal throughout the corresponding P . In particular, if we take $\kappa > 0$ sufficiently small, and define a belief q by $q(\omega) = p^*(\omega)/(1 + \kappa)$ for all $\omega \neq \omega^*$ and $q(\omega^*) = (p^*(\omega^*) + \kappa)/(1 + \kappa)$, then d^* will be optimal for each such belief q . Multiplying through by $1 + \kappa$ tells us that, for every $d \in D$, we have

$$\sum_{\omega} p^*(\omega) u(d, \omega) + \kappa u(d, \omega^*) \leq \sum_{\omega} p^*(\omega) u(d^*, \omega) + \kappa u(d^*, \omega^*)$$

or

$$\begin{aligned} E_{p^*}[u(d, \omega)] &\leq E_{p^*}[u(d^*, \omega)] - \kappa(u(d, \omega^*) - u(d^*, \omega^*)) \\ &= E_{p^*}[u(d^*, \omega)] - \kappa(u(d, \omega^*) + \beta_0(\omega^*)) \end{aligned}$$

where the equality holds by assumption that d^* is ω^* -extremal. In particular, if d is a

decision satisfying $u(d, \omega^*) \geq -\beta_0(\omega^*) + \epsilon$, then we have $E_{p^*}[u(d, \omega)] \leq E_{p^*}[u(d^*, \omega)] - \kappa\epsilon$. It follows that $\psi(p^*, \omega^*) \geq \kappa$, proving the claim.

So now let $\eta = \min_{\omega} \psi(p_{\omega}, \omega) > 0$. Let F be any distribution with full support, and mass at least ϵ on each p_{ω} , whose mean equals the given p_0 . We will show that an unrestricted investment contract is optimal under $\mathcal{I}_0 = \{(\delta_{p_0}, 0), (F, c)\}$, for sufficiently small $c > 0$. Note that full support and non-triviality are satisfied.

We claim that for any β that is $\leq \beta_0$ componentwise,

$$E_F[U_R(p; \beta)] \leq E_F[U_R(p; \beta_0)] - \frac{\epsilon\eta}{|\Omega|} \sum_{\omega} (\beta_0(\omega) - \beta(\omega)). \quad (\text{C.4})$$

Indeed, choose the state ω^* for which $\beta_0(\omega) - \beta(\omega)$ is largest. Put $\beta_1(\omega^*) = \beta(\omega^*)$ and $\beta_1(\omega) = \beta_0(\omega)$ for other states ω . Then increasingness and concavity imply

$$U_R(p; \beta) \leq U_R(p; \beta_1) \leq U_R(p; \beta_0) - (\beta_0(\omega^*) - \beta(\omega^*))\psi(p, \omega^*).$$

Applying expectations under F , noting that $p = p_{\omega^*}$ arises with probability at least ϵ , and $\beta_0(\omega^*) - \beta(\omega^*) \geq \sum_{\omega} (\beta_0(\omega) - \beta(\omega))/|\Omega|$ by choice of ω^* , leads to (C.4).

Thus, on the domain $\beta \leq \beta_0$, the function $E_F[U_R(p; \beta)]$ is bounded above by an affine function of β that is uniquely maximized at β_0 , and this upper bound (C.4) holds with equality at $\beta = \beta_0$. Since the quantity $\sqrt{E_F[U_R(p; \beta)] + E_{p_0}[\beta(\omega)]}$ is locally Lipschitz in β near β_0 , adding a small multiple of it will not change this fact. Hence, for c sufficiently small, the maximand in (C.1) is still uniquely maximized over β by taking $\beta = \beta_0$. (The optimum with respect to choice of experiment must indeed be given by (F, c) , not $(\delta_{p_0}, 0)$, by non-triviality.) It follows that an unrestricted investment contract is optimal.

Part (c). We may shrink P if necessary to assume it is an open ball. Let F_0 be any continuous distribution over posteriors, whose support is the closure of P . The assumption that no decision is optimal throughout P ensures that $E_{F_0}[U(p)] > U(E_{F_0}[p])$. By continuity, this remains true when we replace P by a sufficiently large closed sub-ball \bar{P} , and replace F_0 with the distribution conditioned on $p \in \bar{P}$; call this distribution \bar{F} . Let the prior be $p_0 = E_{\bar{F}}[p]$. Then $p_0 \in P$ by convexity.

As noted at the beginning of Subsection 2.3, if $c > 0$ satisfies the bound (2.3), then $\mathcal{I}_0 = \{(\delta_{p_0}, 0), (\bar{F}, c)\}$ satisfies the non-triviality assumption. Let F' be any distribution with full support on $\Delta(\Omega)$ and with mean p_0 . By continuity of the formula (2.17), we can take $F = (1 - \epsilon)\bar{F} + \epsilon F'$ for sufficiently small $\epsilon > 0$, and then $\mathcal{I}_0 = \{(\delta_{p_0}, 0), (F, c)\}$ still satisfies non-triviality. It clearly satisfies full support as well.

We will show that for ϵ sufficiently small, under this \mathcal{I}_0 , a properly restricted investment contract is optimal. Compactness implies that for all $p \in \overline{P}$, optimal decisions are uniformly bounded away from extremal decisions; that is, one can choose $\beta_1 : \Omega \rightarrow \mathbb{R}$ with $\beta_1(\omega) < \beta_0(\omega)$, such that for all $p \in \overline{P}$, any optimal decision d satisfies $u(d, \omega) + \beta_1(\omega) \geq 0$ for each ω .

In particular, $U_R(p; \beta_1) = U(p) = U_R(p; \beta_0)$ for each such p .

Now, as long as ϵ is small enough, we know that the optimum in (C.1) is attained with (F, c) rather than $(\delta_{p_0}, 0)$ (by non-triviality). On the other hand, it is immediate that when $\epsilon = 0$, the value of (C.1) is strictly higher with $\beta = \beta_1$ than β_0 , since the $U_R(p; \beta)$ terms are unchanged and the $E_{p_0}[\beta(\omega)]$ term is smaller. By continuity, this is still true for sufficiently small ϵ . So in this case, the value of the maximum in (C.1) is attained at some $\beta \neq \beta_0$, which means that the corresponding optimal contract is a properly restricted investment contract. □

Proof of Proposition 3.2: Let $D = \{d_0, d_1\}$, and let D' be the space of all randomized decisions, which we notate as $\{(1 - \gamma)d_0 + \gamma d_1 \mid \gamma \in [0, 1]\}$, thereby making the problem convex. So we know that there are some α and β such that the corresponding restricted investment contract, described in Proposition 2.2, is optimal.

Note that the set of γ 's corresponding to decisions satisfying (2.5) is closed and convex, so it is an interval, say $[\underline{\gamma}, \overline{\gamma}]$. Also, if $\overline{\gamma} = \underline{\gamma}$ then the restricted investment contract allows only a single message, hence there is zero incentive for information acquisition and the optimal contract is the zero contract, contrary to assumption. Hence, $\overline{\gamma} > \underline{\gamma}$ (and $\alpha > 0$).

Now, notice that

$$\alpha \left(u((1 - \gamma)d_0 + \gamma d_1, \omega) + \beta(\omega) \right) \quad \text{with } \gamma \in [\underline{\gamma}, \overline{\gamma}]$$

is equal to

$$\alpha' \left(u((1 - \delta)d_0 + \delta d_1, \omega) + \beta'(\omega) \right)$$

where

$$\alpha' = \alpha(\overline{\gamma} - \underline{\gamma}), \quad \beta'(\omega) = \frac{1}{\overline{\gamma} - \underline{\gamma}} \left((1 - \overline{\gamma})u(d_0, \omega) + \underline{\gamma}u(d_1, \omega) + \beta(\omega) \right)$$

and δ is linearly related to γ by

$$\delta = \frac{\gamma - \underline{\gamma}}{\bar{\gamma} - \underline{\gamma}} \in [0, 1].$$

Therefore, the restricted investment contract with parameters α, β , and restricted decision space given by $\{(1 - \gamma)d_0 + \gamma d_1 \mid \gamma \in [\underline{\gamma}, \bar{\gamma}]\}$, is equivalent to the restricted investment contract with parameters α', β' , and the full decision space D' — that is, the unrestricted investment contract with these parameters α', β' .

Finally, since the optimal decision in D' is always either d_0 or d_1 , this is equivalent to the unrestricted investment contract with parameters α', β' over the original pure decision space D . \square

References

- [1] Dirk Bergemann and Stephen Morris (2005), “Robust Mechanism Design,” *Econometrica* 73 (6), 1771-1813.
- [2] Sugato Bhattacharyya and Francine Lafontaine (1995), “Double-Sided Moral Hazard and the Nature of Share Contracts,” *RAND Journal of Economics* 26 (4), 761–781.
- [3] Benjamin Brooks (2013), “Surveying and Selling: Belief and Surplus Extraction in Auctions,” unpublished paper, University of Chicago.
- [4] Gabriel Carroll (2015), “Robustness and Linear Contracts,” *American Economic Review* 105 (2), 536–563.
- [5] Gabriel Carroll (2015), “Robustness and Separation in Multidimensional Screening,” unpublished paper, Stanford University.
- [6] Sylvain Chassang (2013), “Calibrated Incentive Contracts,” *Econometrica* 81 (5), 1935–1971.
- [7] Yiling Chen, Ian Kash, Mike Ruberry, and Victor Shnayder (2011), “Decision Markets with Good Incentives,” in *Internet and Network Economics* (WINE-11), 72–83.
- [8] Kim-Sau Chung and J. C. Ely (2007), “Foundations of Dominant-Strategy Mechanisms,” *Review of Economic Studies* 74 (2), 447-476.

- [9] Joel S. Demski and David E. M. Sappington (1987), “Delegated Expertise,” *Journal of Accounting Research* 25 (1), 68-89.
- [10] Alexander Frankel (2014), “Aligned Delegation,” *American Economic Review* 104 (1), 66–83.
- [11] Daniel Garrett (2014), “Robustness of Simple Menus of Contracts in Cost-Based Procurement,” *Games and Economic Behavior* 87, 631–641.
- [12] Tilmann Gneiting and Adrain E. Raftery (2007), “Strictly Proper Scoring Rules, Prediction, and Estimation,” *Journal of the American Statistical Association* 102 (477), 359–378.
- [13] Edi Karni (2009), “A Mechanism for Eliciting Probabilities,” *Econometrica* 77 (2), 603–606.
- [14] James M. Malcolmson (2009), “Principal and Expert Agent,” *B.E. Journal of Theoretical Economics*, 9 (1) (Contributions), Article 17.
- [15] Kent Osband (1989), “Optimal Forecasting Incentives,” *Journal of Political Economy* 97 (5), 1091-1112.
- [16] Abraham Othman and Tuomas Sandholm (2010), “Decision Rules and Decision Markets,” in *Autonomous Agents and Multi-Agent Systems (AAMAS '10)*, 625–632.
- [17] R. Tyrrell Rockafellar (1970), *Convex Analysis*. Princeton: Princeton University Press.
- [18] Leonard J. Savage (1971), “Elicitation of Personal Probabilities and Expectations,” *Journal of the American Statistical Association* 66 (336), 783–801.
- [19] Peng Shi, Vincent Conitzer, and Mingyu Guo, “Prediction Mechanisms That Do Not Incentivize Undesirable Actions,” in *Internet and Network Economics (WINE-09)*, 89-100.
- [20] Luis Zermeno (2011), “A Principal-Expert Model and the Value of Menus,” unpublished paper, Massachusetts Institute of Technology.
- [21] Luis Zermeno (2012), “The Role of Authority in a General Principal-Expert Model,” unpublished paper, Massachusetts Institute of Technology.