



Minimum-Regret Contracts for Principal-Expert Problems

Caspar Oesterheld^(✉)  and Vincent Conitzer^{}

Department of Computer Science, Duke University, Durham, NC, USA
{ocaspar, conitzer}@cs.duke.edu

Abstract. We consider a principal-expert problem in which a principal contracts one or more experts to acquire and report decision-relevant information. The principal never finds out what information is available to which expert, at what costs that information is available, or what costs the experts actually end up paying. This makes it challenging for the principal to compensate the experts in a way that incentivizes acquisition of relevant information without overpaying. We determine the payment scheme that minimizes the principal's worst-case regret relative to the first-best solution. In particular, we show that under two different assumptions about the experts' available information, the optimal payment scheme is a set of linear contracts.

1 Introduction

A company has to choose one of a number of different projects, where a project might be to develop a particular product. While the company's personnel is suited to successfully execute any of these projects, the company lacks expertise in market research to decide which of the projects will yield the highest expected profit. To make an informed choice, the company (henceforth, the *principal*) would like to contract faculty members from a nearby business school to give advice on which project to pursue and to make a prediction about the outcome of that project.

While the business school's faculty members (henceforth, the *experts*) have relevant expertise, they need to invest some effort into conducting one relevant research project or another before they can give useful advice. The so-called *first-best solution* is to acquire the information that maximizes expected profit net of the costs of that information. The principal would have to reimburse the experts for those costs, but could keep the rest of the project's profits. However, in general, the principal is unaware of what information can be acquired at what costs and cannot verify the experts' effort or report. The principal can use a payment scheme or *contract* that compensates the experts based on both their final collective report and the outcome of pursuing the recommended project (but not on what would have happened if another project had been chosen). What contract should the principal use?

One way to arrive at a solution would be for the principal to assign some prior probability distribution over configurations of available evidence and select

the contract that maximizes expected profit net of payment to the experts [2, 8, 19, 25]. However, determining such a prior is often impractical. For many priors, it may also be computationally infeasible to identify the optimal contract. We therefore ask what contract ensures the minimum worst-case regret relative to the first-best solution.

Outline. After describing our setup and goals in more detail (Sects. 2 and 3), we show (in Sects. 4 and 5) how linear contracts – which simply pay each expert some fixed fraction of the company’s profits – ensure regret bounds. In Sect. 6, we go on to show that the optimal regret bound is achieved only by a particular linear contract: the one that pays each of the n experts $1/(n+1)$ of the profit obtained. This ensures a regret bound of $v(\mathbf{E}^*)n/(n+1)$, where $v(\mathbf{E}^*)$ is the expected profit (prior to subtracting costs) of the first-best solution. Under stronger assumptions, the approach of this paper can be used to derive different linear contracts to achieve better optimal bounds. In Sect. 7, we give an example of this. Section 8 puts our work in the context of the literature.

2 Setup

Principal and Experts. We consider a *principal* (“she”) who has to choose one of a finite set of projects or *actions* A , each of which probabilistically gives rise to outcomes from some finite set Ω . The principal would like to maximize the expected value of some utility function $u : \Omega \rightarrow \mathbb{R}$. To figure out which action is best, she may interact (in ways specified below) with n *experts*. An important special case is $n = 1$. This case has received the most attention in the literature. Many of the assumptions that we will make later (e.g., about how the experts coordinate) are very weak or even vacuous in the case of $n = 1$, while for $n \geq 2$ they are realistic in some but not all applications.

Each expert $i = 1, \dots, n$ can choose to observe the value of a random variable in some set of random variables $\mathbb{H}_i$. We will refer to these variables as *evidence variables*. We also require that these sets of values are finite. To observe $E_i \in \mathbb{H}_i$, expert i must pay a *cost* (or effort) of $c_i(E_i)$, where $c_i : \mathbb{H}_i \rightarrow \mathbb{R}_{\geq 0}$ is some cost function. We assume that each $\mathbb{H}_i$ contains the constant (trivial) random variable E^0 and that $c_i(E^0) = 0$ for all i . That is, each expert has the option to acquire no information and expend no cost. The experts, on the other hand, all know what evidence variables the other experts have access to and at what costs. They also have a common prior P which, for any vector of random variables $\mathbf{E} \in \mathbb{H} := \times_{i=1}^n \mathbb{H}_i$ and any vector $\mathbf{e}$ of values of $\mathbf{E}$, assigns a probability $P(\mathbf{e}) := P(\mathbf{E} = \mathbf{e})$, as well as for any outcome $\omega \in \Omega$ and action $a \in A$, the probability $P(\omega \mid a, \mathbf{e})$ of obtaining outcome ω after taking action a if $\mathbf{E} = \mathbf{e}$ was observed. For simplicity, we also assume that every observation of $\mathbf{E} = \mathbf{e}$ is consistent, i.e., that for all $\mathbf{E} \in \mathbb{H}$ and $\mathbf{e}$ in the Cartesian product of the sets of values of $E_1, \dots, E_n$, we have $P(\mathbf{e}) > 0$. Some common-knowledge assumptions such as these are necessary to determine the experts’ strategies within standard game-theoretic paradigms. Of course, as is usually the case in such

models, the common-knowledge assumptions – in particular, exact knowledge of one another’s cost of acquisition – are only approximately realistic in practice. Alternatively, one might imagine that they have probabilistic beliefs about each other’s costs or perhaps that they can communicate about each other’s cost. However, this adds an additional layer of complications in expert coordination, which is beyond the scope of the present paper.

The principal knows little about the experts. In particular, she does not know what the $\mathbb{H}_i$ or c_i are, nor does she know the probability distribution P which specifies the probabilities $P(\mathbf{e})$ and $P(\omega \mid a, \mathbf{e})$.

We require that u is normalized s.t. $\max_{a \in A} \mathbb{E}[u(O) \mid a] = 0$, where O is the random variable distributed according to the (prior) probability distribution $P(\cdot \mid a)$ that arises from conditioning only on null evidence E^0 .

Contracts for Information Elicitation. The principal wants the experts to acquire and honestly report useful information. Since acquiring information is costly, the principal has to set some kind of incentive. If she could observe expended costs, then this problem would be easy: simply reimburse costs and pay some small bonus that is positive affine in the utility obtained by the principal net of the overall reimbursements for the experts’ acquisition costs. However, we assume that effort is unobservable to the principal. We furthermore assume that the information obtained is unverifiable.

We will consider a simple class of mechanisms in which the experts only submit (potentially dishonest) reports $\hat{\mathbf{e}}$ on what information they obtained. The principal then takes the best action given $\hat{\mathbf{e}}$, i.e., takes $a_{\hat{\mathbf{e}}} := \arg \max_{a \in A} \mathbb{E}[u(O) \mid \hat{\mathbf{e}}, a]$, where ties are broken arbitrarily and O is the random variable distributed according to $P(\cdot \mid \hat{\mathbf{e}}, a)$. Of course, to determine $a_{\hat{\mathbf{e}}}$ based on $\hat{\mathbf{e}}$, one has to know (at least partially) $P(\cdot \mid \cdot, \hat{\mathbf{e}})$, which so far we have assumed the principal not to know. For example, we could imagine that the experts convene to summarize their evidence into a report that the principal can interpret.

Some authors have allowed the principal to randomize between projects – giving the most probability to the best ones – to have some chance of testing the predictions made for suboptimal actions [7, 27, 28]. Of course, randomization comes at the cost of sometimes taking suboptimal actions. Indeed, our negative results (see Sect. 6 and Theorem 5) can be extended to show that to minimize worst-case regret, the principal must always select the best action given the report.

Finally, each expert i is rewarded only based on the probability distribution resulting from the overall report and the observed outcome, i.e., based on $s_i(P(\cdot \mid \hat{\mathbf{e}}, a_{\hat{\mathbf{e}}}), \omega)$, where s_i is some *scoring rule* or *contract*. Again, we have to imagine that the principal somehow learns about $P(\cdot \mid \hat{\mathbf{e}}, a_{\hat{\mathbf{e}}})$, e.g., by having the experts provide that distribution. Note that the payoff depends only on the prediction about the recommended action $a_{\hat{\mathbf{e}}}$. Other predictions are not tested and it is therefore futile to ask for predictions about them, as pointed out by Othman and Sandholm [22, Theorems 1 and 4] and Chen et al. [7, Theorem 4.1]. It is

easy to show that the results of this paper generalize to a setting in which the principal's scoring rule can depend on all of $P(\cdot \mid \hat{\mathbf{e}}, \cdot)$.

More importantly, we assume that the principal scores only according to the aggregated expert report. That is, we assume that the principal does not know the experts' information structure and therefore cannot determine the *relative* value of individual experts' contributions. Similarly, we assume that the principal does not ask the experts for the cost of their information. In principle, in the case of multiple experts (i.e., $n \geq 2$), different kinds of mechanisms could also be considered. In particular, the principal could ask the experts to report on the value and cost of each other's information. However, this will often be unrealistic. For instance, consider the members of a team in a firm. The members of the team may have a good understanding of each other's abilities and contributions as well as of how costly these contributions are to the different members, but the firm will generally not ask the team members to report on these things and instead determine salaries based on relatively little information. (Note that none of these considerations are relevant to the single-expert case.)

The Principal's and Experts' Goals. We assume that the principal accounts for her payments to the experts quasilinearly, so that her overall utility after payments is given by $u(\omega) - \sum_{i=1}^n s_i(P(\cdot \mid \hat{\mathbf{e}}, \mathbf{a}_{\hat{\mathbf{e}}}), \omega)$.

As for the experts, a configuration of available evidence $\mathbb{H}$ with prior P and costs $(c_i)_{i=1, \dots, n}$, and a (multi-expert) scoring rule $\mathbf{s}$ induce an n -player game played by the experts. Each player's strategy σ_i consists of two parts, one determining which evidence he obtains and one determining how observed evidence is mapped onto reports. Throughout this paper, we use $\mathbf{E} \in \mathbb{H}$ to denote the strategy profile in which each player i obtains and honestly reports E_i . A strategy profile σ gives rise to an expected payoff $\text{EU}_{\mathbf{s}}^i(\sigma)$ for expert (or player) i and an expected utility net of payments $\text{EU}_{\mathbf{s}}(\sigma)$ for the principal.

Since the experts play a strategic game, we use Nash equilibrium to describe their behavior. We say that σ is a *Nash equilibrium* iff for each i and each alternative strategy σ'_i for i , we have $\text{EU}_{\mathbf{s}}^i(\sigma) \geq \text{EU}_{\mathbf{s}}^i(\sigma_{-i}, \sigma'_i)$. In general, the game resulting from a configuration and scoring rule will have many equilibria. For $n \geq 2$, it is futile to ask for regret bounds that hold for *all* Nash equilibria. For example, imagine that the value for the principal of $\mathbf{E}$ being obtained is high if $\mathbf{E} = \mathbf{E}^*$ and low otherwise. Imagine further that $c(E_i^*)$ is small but positive for all i . Then in the first-best solution, $\mathbf{E}^*$ is acquired. But, if there are multiple experts, everyone obtaining E^0 (no information) is also a Nash equilibrium with (arbitrarily close to) maximum regret. Throughout the rest of this paper, we therefore ask: what is the regret in the Nash equilibrium that is *best* for the principal? (Cf. the notion of *price of stability* [1, 24, Section 1.3], as opposed to the price of anarchy [18, 23].) Our negative results, of course, are made stronger by the fact that they say that *no* Nash equilibrium can exceed a certain bound. Our positive results, on the other hand, are mostly about a particular kind of Nash equilibria (see Lemma 1) which arise from maximizing the experts' profit.

In what follows, we do not require our scoring rules to be proper, i.e., we do not require that they incentivize the experts to report honestly. However, our results will show that the optimal contract is indeed proper. We do require that our scoring rules satisfy an individual rationality constraint. In particular, we require that each expert i receives an expected payoff of at least 0 in the strategy profile $\mathbf{E}^0 = (E^0, \dots, E^0)$ where everyone honestly reports the null information, i.e., we require that for all i , $\text{EU}_s^i(\mathbf{E}^0) \geq 0$. Note that this is a fairly weak notion of individual rationality. For instance, it does *not* say that the expected payoff for the expert is nonnegative if others truthfully report non-null information. This makes our negative results stronger. The linear contracts of our positive results will in fact satisfy stronger versions of individual rationality. For instance, they do ensure nonnegative ex-ante expected scores whenever all experts submit information honestly.

3 Competitive Analysis

In this paper, we analyze scoring rules in the style of competitive analysis, a technique for analyzing algorithms that combines two ideas. The first is worst-case analysis. To avoid dependence on some prior probability distribution over, in our case, configurations of costs and available evidence, we consider how a scoring rule performs in the worst case. The second idea of competitive analysis is to consider worst-case expected utility *relative to some benchmark* for the problem. Similar approaches have been used in the literature on principal-expert and -agent problems before [4–6, 16].

As is common in principal-agent problems, we use the *first-best solution* as a benchmark, i.e., the utility (net of information acquisition costs) that the principal could obtain if she had full control over the experts and knew everything about the information structure that the experts know. Formally, let $v(\mathbf{E}) := \mathbb{E}_{\mathbf{E}}[\mathbb{E}_O[u(O) \mid a_{\mathbf{E}}, \mathbf{E}]]$ be the expected utility obtained from acquiring $\mathbf{E}$ and then taking the best action according to it. Also, let $c(\mathbf{E}) := \sum_{i=1}^n c_i(E_i)$ be the overall cost of acquiring $\mathbf{E}$. Then the expected utility net of costs of the first-best solution is $\text{EU}_{\text{OPT}} := \max_{\mathbf{E} \in \mathbb{H}} v(\mathbf{E}) - c(\mathbf{E})$. We will use $\mathbf{E}^*$ to denote a first-best solution itself, i.e., a maximizer of $v(\mathbf{E}) - c(\mathbf{E})$.

There are two ways in which the performance of an algorithm is commonly compared against the benchmark: competitive ratios and regret. Unfortunately, we cannot derive any nontrivial competitive ratio. Consider the case where there is just one expert and only one available piece of evidence E with $v(E) = 1$. Then to be competitive (i.e., to get positive utility at all), if the cost of E is $c(E) = 1 - \epsilon$, the principal has to reward the expert with almost 1. To be reasonably competitive at $c(E) = \epsilon$, on the other hand, she cannot give away anything close to 1. Because the rewards cannot depend on the cost function (which the principal does not know), obtaining a non-trivial competitive ratio is generally impossible, even in the single-expert case. That said, we will give two competitive-ratio-like results (Proposition 2 and Theorem 3) in which EU_{OPT} is replaced with a weaker benchmark.

Our primary focus will be on regret, which is the *difference* between the first-best solution’s utility (net of costs) and the utility (net of payments to the experts) achieved by using the scoring rule. So, for any strategy profile σ we define the regret for that strategy profile as $\text{REGRET}_s(\sigma) := \text{EU}_{\text{OPT}} - \text{EU}_s(\sigma)$. We will also use REGRET_s to denote the lowest regret achieved in any Nash equilibrium σ for s , i.e., $\text{REGRET}_s = \min_{\sigma \in \text{NE}(s)} \text{REGRET}_s(\sigma)$. Roughly, the regret is what is sometimes called the agency cost in the literature on principal-agent and -expert problems, or the price of stability in mechanism design [1, 24, Sect. 1.3].

4 Linear Contracts

In this paper, we study and justify the use of a particular type of scoring rule: *linear contracts*. For any $\alpha \in (0, 1]^n$ with $\sum_{i=1}^n \alpha_i \leq 1$, define the linear scoring rule $\mathbf{q}^\alpha$ as scoring according to $q_j^\alpha(\hat{P}, \omega) = \alpha_j u(\omega)$ for all outcomes ω and reported probability distributions over outcomes $\hat{P}$. That is, each expert receives a fixed fraction of the total payoff generated. Requiring $\alpha_i > 0$ for all i is done for simplicity. All the positive results about linear scoring rules can easily be generalized to linear contracts in which $\alpha_i = 0$ for some i .

Before proceeding with our detailed analysis of linear contracts, it is worth pointing out some immediately obvious and appealing properties. Most importantly, by rewarding according to a positive affine transformation of the principal’s utility, they align the experts’ interests with the principal’s. In contrast, if one were to, say, reward one expert in proportion to $\exp(u(\omega))$, then that expert would sometimes want the principal to take a risky (high variance) rather than a safe action, even if the risky action has lower expected utility. When using linear scoring rules, the only misalignment between experts and principal is that the experts only receive a fraction of the utility obtained and therefore do not value information as highly as the principal would in the first-best solution. Many other desirable properties have been pointed out in the literature; see the discussion of related work in Sect. 8.

From the definition of linear contracts, it is immediately clear that, while they reward the choice of a good action, beyond that they do not reward accurate probabilistic forecasts about the outcome. Because the principal may additionally like to know what to expect for the chosen action, this is an undesirable aspect of linear scoring rules. Note, however, that Oosterheld and Conitzer [21, Sect. 2.5.1] show that linear scoring rules are the only ones which incentivize honest reporting of the best action without incentivizing the expert to sometimes prefer acquiring decision-irrelevant over decision-*relevant* evidence variables.

A more substantial issue with linear contracts is that (in some configurations of available evidence) they violate *ex-interim* individual rationality constraints. After acquiring some piece of evidence E_i , an expert i may come to believe that the expected utility of the principal is negative. Expert i may then wish to withdraw from the mechanism. Also, because utilities can end up being negative, linear contracts cannot be used if the experts are protected by limited liability.

However, these concerns do not apply in cases where the principal always has an option to walk away with utility 0, regardless of the evidence.

5 General Regret and Ratio Bounds for Linear Scoring Rules

In this section, we give positive results about what regret (and competitive ratio-like) bounds linear scoring rules achieve. Because linear contracts do not score experts on their reported beliefs, all of these results carry over to generic principal-agent problems. We start with a lemma on which the subsequent results of this section are based.

Lemma 1. *Let $\mathbf{q}^\alpha$ be a linear contract. Then for all configurations of available evidence, any*

$$\hat{\mathbf{E}} \in \arg \max_{\mathbf{E} \in \mathbb{H}} v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c_i(E_i) \quad (1)$$

is a Nash equilibrium of the game induced by $\mathbf{q}^\alpha$.

Based on Lemma 1 we now give a bound on the regret of using any linear scoring rule. Let $\alpha_{\min} := \min_i \alpha_i$.

Theorem 1. *For all configurations of available evidence, the Nash equilibria $\hat{\mathbf{E}}$ of Lemma 1 satisfy*

$$\text{REGRET}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) \leq \max \left(\sum_{i=1}^n \alpha_i, 1 - \alpha_{\min} \right) v(\mathbf{E}^*). \quad (2)$$

In particular, setting $\alpha_j = 1/(n+1)$ for all j achieves a regret bound of $\text{REGRET}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) \leq nv(\mathbf{E}^)/(n+1)$.*

The regret bound $\text{REGRET}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) \leq nv(\mathbf{E}^*)/(n+1)$ is the best bound that a linear contract can achieve without any assumptions about the configuration of available evidence. One might have hoped for a better bound, at least for larger n . Also, it requires the principal to give each expert a share of the proceeds equal to her own, which means that unless a large fraction of the experts pay an amount close to $v(\mathbf{E}^*)/(n+1)$, regret is generally high. However, we will see (in Sect. 6) that the regret bound is tight not only for linear scoring rules but that no scoring rule can achieve a better bound. We will also consider two ways of making assumptions about the configuration of available evidence to achieve better bounds. One is based on a competitive-ratio-type bound from the literature and is discussed in the rest of this section. The other targets regret and will be the subject of Sect. 7.

Theorem 1 gives a regret bound for a specific equilibrium. It is natural to ask whether this equilibrium is a plausible one. If it was a bad equilibrium for the experts, we might not expect that equilibrium to be played. The first thing to

note is that in the case of $n = 1$, there is only one Nash equilibrium, anyway, and in this Nash equilibrium the single expert maximizes his expected profit. For the multi-expert case, notice first that the equilibrium of Lemma 1 explicitly maximizes a term that is closely tied to the experts' expected utility. A more formal point is the following.

Proposition 1. *Let $\mathbf{q}^\alpha$ be a linear contract and $\tilde{\mathbf{E}}$ be a Nash equilibrium of the game induced by $\mathbf{q}^\alpha$. If $\tilde{\mathbf{E}}$ is not strongly Pareto-dominated (for the experts) by $\mathbf{E}^*$, then $\tilde{\mathbf{E}}$ satisfies the regret bound of Ineq. 2.*

Intuitively, this means that if some equilibrium does not satisfy Ineq. 2, then the expert dislikes this equilibrium in the sense of it being strictly Pareto dominated by $\mathbf{E}^*$. Unfortunately, $\mathbf{E}^*$ itself may not be a Nash equilibrium. In fact, it may be that all Nash equilibria of the game induced by $\mathbf{q}^\alpha$ are strictly Pareto-dominated by $\mathbf{E}^*$.

Lemma 1 also gives us the following result, which is a generalization to the multi-expert case of a result shown by Chassang [6, Theorem 1.i] and Carroll [4, Sect. 2.3].

Proposition 2. *For all configurations of available evidence, the Nash equilibria $\hat{\mathbf{E}}$ of Lemma 1 for the linear scoring rule $\mathbf{q}^\alpha$ satisfy*

$$\text{EU}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) \geq \left(1 - \sum_{i=1}^n \alpha_i\right) \max_{\mathbf{E}} \left(v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c_i(E_i) \right). \quad (3)$$

Proposition 2 is essentially a competitive-ratio-type result, except that the benchmark is lower than the first-best solution. Chassang [6, Theorem 1.ii] shows how in a single-expert version of this result, the principal can optimize α if she knows a bound on the cost-to-value ratio of information. If information is known to be cheap, then α can be small. Chassang's proof only operates on the $n = 1$ special case of Ineq. 3. A similar line of reasoning applies to our multi-expert setting. Such a result is useful for practical purposes. It also shows how the existing results can be used to give better bounds and recommendations that are to some extent tailored to specific settings. Unfortunately, it seems that if the cost-to-value bounds vary between experts, no succinct expression for the optimal contracts can be given.

6 Unique Optimality of Linear Scoring Rules

Having proven bounds on the regret of linear contracts, the natural next question is: can we do any better by using a different scoring rule? In particular, can we do better by eliciting predictions of what outcome will materialize, in addition to recommendations of what action to take? It is easy to come up with examples of particular prior probability distributions over configurations of available evidence under which the answer is yes. But it turns out that in the worst case and without further assumptions, we cannot get any better regret bounds; moreover, linear

contracts are in fact the *only* ones that achieve the optimal regret bound in general. This is true even if the principal knows the pre-cost expected utility $v(\mathbf{E}^*)$ of the information acquired in the first-best solution.

Theorem 2. *Let $0 < H < \max_{\omega \in \Omega} u(\omega)$ and let $\mathbf{s}$ be a scoring rule. Then if for all configurations with $v(\mathbf{E}^*) = H$, $\text{REGRET}_{\mathbf{s}} \leq nH/(n+1)$, then it must be that for all $j = 1, \dots, m$, $s_j(\hat{P}, \omega) = u(\omega)/(n+1)$, whenever $\omega \in \text{supp}(\hat{P})$. There is no scoring rule $\mathbf{s}$ s.t. for all configurations, $\text{REGRET}_{\mathbf{s}} < nH/(n+1)$.*

We briefly give a sketch of the proof, which consists of two parts. In the first part, we identify “critical cases” for any $\mathbf{s}$, i.e., a small set of classes of configurations on which the bound is tight and which together determine s_j to be the hypothesized linear scoring rule. One critical case is that in which $v(\mathbf{E}^*) = H$ and $\mathbf{E}^*$ is in fact free to acquire. To keep regret low in this case, the principal has to make sure that she does not give away too much. Overall, she can only give away $nH/(n+1)$ in expectation. The other critical case is that in which $v(\mathbf{E}^*) = H$ and in $\mathbf{E}^*$ exactly one expert j acquires information at a price of $H/(n+1) - \epsilon$. To achieve low regret in these cases, the principal must make sure that whenever an expectation of H is achieved, any expert j receives an expected payoff of at least $H/(n+1)$ (or, gets at least $H/(n+1)$ more than it gets for reporting the prior). The critical cases together imply that if information $\mathbf{E}^*$ with value $v(\mathbf{E}^*) = H$ is acquired, each expert receives an expected payoff of $H/(n+1)$ (and that if the prior is reported, each expert receives an expected payoff of 0). The second part of the proof shows that this (across all possible $\mathbf{E}^*$ with $v(\mathbf{E}^*) = H$) implies that s_j is as claimed in the theorem. Roughly, in this part we show that the scoring rule must be linear, using the fact that the expected payoff is constant across different distributions with the same mean.

The different aspects of this result depend on the details of our setup to different extents. The result that worst-case regret is $nH/(n+1)$ generalizes far beyond our setting. In particular, even if the principal knows the experts’ information structure, there will still be cases with regret $nH/(n+1)$ if the principal cannot obtain reliable information about the different experts’ costs of acquisition. The uniqueness of linear scoring rules in minimizing worst-case regret, on the other hand, does hinge on our assumption that the principal does not know the information structure. With knowledge of the specific information structure, the principal can use very different contracts. As a straightforward example, if it is known that one expert cannot obtain sufficiently useful information, the scoring rule need not pay that expert at all.

A result analogous to Theorem 2 holds true for the competitive ratio-based bound and can be proven with very similar ideas.

Theorem 3. *Let $\alpha \in (0, 1)^n$ with $\sum_{j=1}^n \alpha_j < 1$ and $\mathbf{s}$ be a scoring rule. Then if for all configurations there is Nash equilibrium $\hat{\mathbf{E}}$*

$$\text{EU}_{\mathbf{s}}(\hat{\mathbf{E}}) \geq \left(1 - \sum_{i=1}^n \alpha_i\right) \max_{\mathbf{E}} \left(v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c_i(E_i)\right), \quad (4)$$

then it must be the case that for all $j = 1, \dots, m$, $s_j(\hat{P}, \omega) = \alpha_i u(\omega)$, whenever $\omega \in \text{supp}(\hat{P})$. There is no scoring rule s s.t. Ineq. 4 is always strict.

7 Restrictions on the Configurations of Available Evidence

In this section we consider a setting in which the principal is assumed to have a particular type of knowledge about the configuration of available evidence (similar to Chassang’s [6, Theorem 1.ii] result, mentioned at the end of Sect. 5). With this we would like to show that (as one would expect) under stronger assumptions, substantially better bounds can be derived. Perhaps more importantly, it shows that the strategy in the proof of Theorem 2 of using critical cases to derive linear contracts and their optimality generalizes to settings with additional assumptions.

Arguably, much of the reason why our general bound is not better than it is that we do not know who has access to decision-relevant information. While we use the term “experts”, we allow for configurations in which almost all of the “experts” cannot acquire decision-relevant information at a reasonable cost. Indeed, these cases drive the proof of Theorem 2. In many real-world settings, the principal is able to select a set of experts who all can acquire relevant information. We will model this by introducing the assumption that all experts have access to the same set of evidence variables – though note that of this set each expert can still only obtain one element.

Assumption 1. $\mathbb{H}_1 = \mathbb{H}_2 = \dots = \mathbb{H}_n$.

Furthermore, we assume that there is some known bound on how much acquisition costs differ.

Assumption 2. There is some known $\Lambda \in (0, 1]$ such that for any two experts i, j and non-trivial evidence variables E_i, E_j we have $c_j(E_j) > 0$ and $\Lambda \leq c_i(E_i)/c_j(E_j)$.

If $\Lambda = 1$, then all experts pay the exact same price for all pieces of information. If Λ is small, then some experts may be able to acquire information much cheaper than others. Note that Assumption 2 not only restricts how costs differ between experts but also between different evidence variables (both across experts and for a single expert).

We add another assumption:

Assumption 3. For all vectors of information $\mathbf{E} \in \mathbb{H}$ and any expert i , we have $v(\mathbf{E}_{-i}) \in \{0, v(\mathbf{E})\}$.

Roughly, this means that any set of evidence variables is either fully complementary (in which case $v(\mathbf{E}_{-i}) = 0$ for all i that acquire non-trivial information) or has some redundant piece of information (in which case $v(\mathbf{E}_{-i}) = v(\mathbf{E})$ for some i). There are some settings in which such an assumption is (at least approximately) natural. For instance, we may imagine that the principal and experts

are morally or legally obliged to pay due diligence and cannot pursue projects unless they are fully researched. In the context of this paper, another reason we consider this assumption is that it allows for an equilibrium analysis that is more powerful than that of Lemma 1.

As before (Sects. 5 and 6), we first provide the positive result. That is, we show that a particular linear scoring rule achieves a particular regret bound. We then show (Theorem 5) that this scoring rule is optimal and the only one that achieves the given regret bound. It turns out that in this case the optimal scoring rule is much harder to guess. We hope that the proof sketch of Theorem 5 makes clear where its parameters come from.

Theorem 4. *Let $n \in \mathbb{N}$ be the number of experts. Given Assumptions 1 to 3, define*

$$B_{\Lambda,n} := 1 - \frac{1}{1 + \sum_{i=1}^n \frac{1}{1+(i-1)\Lambda}}. \quad (5)$$

and for $j = 1, \dots, n$

$$\alpha_j = \frac{1}{(1 + (j-1)\Lambda) \left(1 + \sum_{i=1}^n \frac{1}{1+(i-1)\Lambda}\right)}. \quad (6)$$

Then, $\text{REGRET}_{\mathbf{q}^\alpha} \leq B_{\Lambda,n} v(\mathbf{E}^*)$.

We now prove that the scoring rule of Theorem 4 is the only one that achieves its regret bound. Our strategy is the same as the strategy behind the proof of Theorem 2 and the omitted proof of Theorem 3. Very roughly, the idea is as follows. For any given linear contract q^α , we guess the cases where the regret is highest. The first such case is – as in the proof of Theorem 2 – the one in which information is free to the experts and regret is entirely a result of the principal having to give away some fraction of her profits that she can keep in the first-best solution. Second, there is a critical case for each $k = 1, \dots, n$, in which k pieces of information are needed and the expert i with the k -th highest α_i cannot quite afford a relevant piece of information. One can then find the given bound and parameters of the linear contract by minimizing worst-case regret across these cases. Using these cases, one can prove as in the proof of Theorem 2 that to obtain the bound, one has to use this linear rule.

Theorem 5. *Let $0 < H < \max_{\omega \in \Omega} u(\omega)$, and $\mathbf{s}$ be a scoring rule. Then, if for all configurations with $v(\mathbf{E}^*) = H$ that satisfy Assumptions 1 to 3, we have $\text{REGRET}_{\mathbf{s}} \leq B_{\Lambda,n} H$, then – up to permutation of the experts – for all $j = 1, \dots, n$: $s_j(\hat{P}, \omega) = \alpha_j u(\omega)$ whenever $\omega \in \text{supp}(\hat{P})$, where the α_j are as defined in Eq. 6. There is no scoring rule $\mathbf{s}$ s.t. $\text{REGRET}_{\mathbf{s}} < B_{\Lambda,n} H$ for all configurations.*

If $\Lambda = 0$, then $B_{\Lambda,n} = n/(n+1)$ and $\alpha_j = 1/(n+1)$ for $j = 1, \dots, n$. That is, as the restriction on the cost ratios becomes vacuous, the optimal bound and scoring rule approach the optimal *general* bound and scoring rule of Theorems 1 and 2. If $\Lambda = 1$ (i.e., all costs are the same), then $B_{\Lambda,n} = h_n/(h_n + 1)$ and $\alpha_j = 1/((h_n + 1)j)$, where $h_n = \sum_{i=1}^n 1/i$ is the n -th harmonic number.

Note that even though – for all the principal knows – the experts are all identical, the minimum-regret contract varies the numbers of shares in the project given to different experts. Theorem 5 therefore provides another (and quite different) demonstration of a point made by Winter [26], who shows that the optimal reward structure for a principal-(multi-)agent problem sometimes has to treat identical agents differently. To understand why in our setting optimal rewards are asymmetric despite symmetry between agents, consider only the cases where $\Lambda = 1$, i.e., where all experts pay exactly the same price for all pieces of information. Consider the question of how many experts we should give enough shares to overcome some given acquisition cost of c . If that number is k , then our worst-case regret at cost c from no information being acquired is $H - (k + 1)c$ and occurs in the case where $k + 1$ pieces of information (all at cost c) are needed. Since this number decreases with k , giving k experts sufficiently many shares to outweigh a cost of a sufficiently large c at some point becomes non-critical for minimizing regret. Given regret considerations in other cases (in particular the one where all information is essentially free), the minimum-regret value of k will therefore be smaller than n but bigger than 0 for many values of c .

8 Related Work

The most closely related strand of literature is that on principal-expert (and more generally principal-agent) problems. Our results merely concern one of many possible variants of and approaches to such problems. For example, much of the literature on principal-expert problems differs from the present work in that they do not let the expert submit (or reveal by selection of a contract from a contract menu) any information apart from a recommendation. We are not the first to approach the problem from a worst-case perspective [4–6, 16]; but many others have derived very different kinds of results without the worst-case assumption, for instance by considering specific (types of) distributions or other restrictions [2, 8, 12, 14, 19, 25, 27]. Also, many papers have richer problem representations and specialized foci on issues that do not arise in the present framework. For instance, most authors take into account that the expert is protected by limited liability. With a few exceptions [2, 13], existing work only considers settings with a single expert. While, as we have noted, some of our results can be seen as generalizations of corresponding single-expert results (one of which – Proposition 2 – was already given in the literature for the single-expert case), Sect. 7 discusses issues that are very specific to the multi-expert case. To our knowledge, our main optimality arguments (the proofs of Theorems 2, 3 and 5) and most of our results are also unique. At the same time, our results support other work which has aimed to discuss and explain the use of linear contracts [4–6, 10, 11, 16, 21, 25, Sect. 2.5.1].

In mechanism design, a few authors have worked to characterize scoring rules that incentivize experts to honestly report *existing* (or free) decision-relevant information [7, 21, 22]. The setups of these papers do not give any objective that allows one to identify particular scoring rules as optimal; they allow for rewards of tiny scale (say, giving the experts a trillionth of the principal’s profit).

The introduction of information acquisition costs into the model forces the use of nontrivial rewards, and allows us to ask meaningful questions about what scoring rule is optimal. Overcoming acquisition costs is one way to introduce a target for optimization among scoring rules that gives a reason to give larger-scale scores. The same can be achieved by introducing conflicts of interest that arise if the expert has (contrary to the setup of this paper) an intrinsic interest in the principal's decision. The expert may have an incentive to misreport (or not report anything if information is verifiable) to make the principal take the *expert's* (rather than the principal's) favorite option [3, 9, 15, 20]; cf. the literature on Bayesian persuasion [17].

9 Conclusion

We have shown how competitive analysis can be used to derive the optimality of particular linear contracts in principal-expert problems. We demonstrated that when adding specific assumptions about the structure and cost of available information, the analysis can also provide optimal scoring rules for specific settings. The optimal scoring rules in all of these settings give away a substantial fraction of the principal's profit. The present work therefore motivates the use of more complicated mechanisms when dealing with multiple experts. For instance, the principal may look to save money by asking the experts to reveal each other's costs of acquisition. Can similar arguments as in this paper then still be used to justify the use of linear scoring rules? Further, it is worth asking what the cost of the worst-case simplification is: how much better can we do if the principal formulates a prior over configurations of available evidence and optimizes the expected utility over the set of contracts [2, 8, 19, 25]?

Acknowledgments. We are thankful for comments from our anonymous reviewers and support from NSF under award IIS-1814056.

References

1. Anshelevich, E., Dasgupta, A., Kleinberg, J., Éva Tardos, Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. *SIAM J. Comput.* **38**(4), 1602–1623 (2008)
2. Barron, J.M., Waddell, G.R.: Executive rank, pay and project selection. *J. Financ. Econ.* **67**, 305–349 (2003)
3. Boutilier, C.: Eliciting forecasts from self-interested experts: scoring rules for decision makers. In: Winikoff, C., van der Hoek, P. (eds.) *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, Richland, South Carolina, pp. 737–744 (2012)
4. Carroll, G.: Robustness and linear contracts. *Am. Econ. Rev.* **105**(2), 536–563 (2015)
5. Carroll, G.: Robust incentives for information acquisition. *J. Econ. Theory* **181**, 382–420 (2019)
6. Chassang, S.: Calibrated incentive contracts. *Econometrica* **81**(5), 1935–1971 (2013)

7. Chen, Y., Kash, I.A., Ruberry, M., Shnayder, V.: Eliciting predictions and recommendations for decision making. *ACM Trans. Econ. Comput.* **2**(2), 6:1–6:27 (2014)
8. Core, J.E., Qian, J.: Project selection, production, uncertainty, and incentives (2002). <https://doi.org/10.2139/ssrn.297461>
9. Crawford, V.P., Sobel, J.: Strategic information transmission. *Econometrica* **50**(6), 1431–1451 (1982)
10. Diamond, P.: Managerial incentives: on the near linearity of optimal compensation. *J. Polit. Econ.* **106**(5), 931–957 (1998)
11. Dütting, P., Roughgarden, T., Talgam-Cohen, I.: Simple versus optimal contracts. In: *Proceedings of the 2019 ACM Conference on Economics and Computation*, New York, pp. 369–387 (2019)
12. Feess, E., Walzl, M.: Delegated expertise - when are good projects bad news? *Econ. Lett.* **82**(1), 77–82 (2004)
13. Gromb, D., Martimort, D.: Collusion and the organization of delegated expertise. *J. Econ. Theory* **137**(1), 271–299 (2007)
14. Häfner, S., Taylor, C.R.: On young turks and yes men: Optimal contracting for advice (2019). <https://doi.org/10.2139/ssrn.3229927>
15. Holmström, B.: On the theory of delegation. Discussion Papers 438, Northwestern University, Center for Mathematical Studies in Economics and Management Science (1980). <https://ideas.repec.org/p/nwu/cmsems/438.html>
16. Hurwicz, L., Shapiro, L.: Incentive structures maximizing residual gain under incomplete information. Technical report, pp. 77–83 (1977). <https://conservancy.umn.edu/bitstream/handle/11299/54918/1977-83.pdf>
17. Kamenica, E., Gentzkow, M.: Bayesian persuasion. *Am. Econ. Rev.* **101**, 2590–2615 (2011)
18. Koutsoupas, E., Papadimitriou, C.: Worst-case equilibria. In: Meinel, C., Tison, S. (eds.) *STACS 1999. LNCS*, vol. 1563, pp. 404–413. Springer, Heidelberg (1999). https://doi.org/10.1007/3-540-49116-3_38
19. Lambert, R.A.: Executive effort and selection of risky projects. *RAND J. Econ.* **17**(1), 77–88 (1986)
20. Milgrom, P., Roberts, J.: Relying on the information of interested parties. *RAND J. Econ.* **17**(1), 18–32 (1986)
21. Oosterheld, C., Conitzer, V.: Eliciting information for decision making from individual and multiple experts (2019). <https://users.cs.duke.edu/~ocaspar/DecisionScoringRules.pdf>. Short version appeared in WINE 2020
22. Othman, A., Sandholm, T.: Decision rules and decision markets. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pp. 625–632 (2010)
23. Papadimitriou, C.H.: Algorithms, games, and the internet. In: *Talk at STOC 2001*, Hersonissos, Crete, Greece, 6–8 July 2001 (2001)
24. Roughgarden, T., Tardos, É.: Introduction to the inefficiency of equilibria. In: Nisan, N., Roughgarden, T., Tardos, É., Vazirani, V.V. (eds.) *Algorithmic Game Theory*, Chap. 17, pp. 443–459. Cambridge University Press, Cambridge (2007)
25. Stoughton, N.M.: Moral hazard and the portfolio management problem. *J. Finance* **48**(5), 2009–2028 (1993)
26. Winter, E.: Incentives and discrimination. *Am. Econ. Rev.* **94**(3), 764–773 (2004)
27. Zermeno, L.: A principal-expert model and the value of menus (2011). <http://economics.mit.edu/files/7299>
28. Zermeno, L.: The role of authority in a general principal-expert model (2012). <http://economics.mit.edu/files/7627>