

Minimum-regret contracts for principal-expert problems

Caspar Oesterheld and Vincent Conitzer

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 over-paying. 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

Optimally acquiring information is a well-recognized problem in AI and beyond. An agent that sets out to always acquire all relevant information before taking any decision will quickly deplete its resources, or find that the opportunity has passed by the time it has completed its analysis. A good agent trades off the value of the information against the cost of acquiring it. However, in many circumstances, it is not possible for the agent to acquire the information on its own; rather, it needs to rely on another agent to do so. For example, the former agent may have the relevant actuators, but the latter the relevant sensors. If the agents are self-interested, the former agent will need to provide an *incentive* to the latter to acquire the information.

Consider the following illustrative example. 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 for different choices). 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 [cf. Barron and Waddell, 2003; Core and Qian, 2002; Lambert, 1986; Stoughton, 1993]. 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 (Sections 2 and 3), we show (in Sections 4 and 5) how linear contracts – which use the reports only for making the decision and then simply pay each expert some fixed fraction of the company’s profits regardless of their reports – ensure regret bounds. In Section 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 Section 7, we give an example of this. Section 8 puts our work in the context of the literature. Briefly: the most closely related strand is work on principal-expert and -agent problems. In particular, the single-expert version of our Proposition 3 has been given in earlier work by Chassang [2013, Theorem 1.i] and Carroll [2015, Section 2.3]. One may also view

our paper as contributing to the literature on decision scoring rules [Othman and Sandholm, 2010; Chen *et al.*, 2014; Oosterheld and Conitzer, 2019]. Relative to that line of work, the distinguishing feature of our setup is that it explicitly asks which decision scoring rules strike the optimal balance between incentivizing experts to acquire costly information and minimizing the overall payment to the experts. Finally, Section 9 concludes by pointing out implications of the present work for a common topic of debate in the literature: simple (linear) versus optimal (potentially complex) contracts.

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 principal does not know what the \mathbb{H}_i or c_i are. The experts, on the other hand, all know what evidence variables the other experts have access to and at what costs. Furthermore, the experts 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 each other’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.

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 . This normalization requires the principal to know what utility she can achieve without any additional information. For instance, in the scenario of the introduction, the baseline may be to develop no new project or to sell the entire firm for some known price. For the positive results of the paper it is only necessary that $\max_{a \in A} \mathbb{E}[u(O) \mid a] \geq 0$.

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 after taking into account 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 $\hat{\mathbf{E}} = \hat{\mathbf{e}}$, i.e., takes

$$a_{\hat{\mathbf{e}}} := \arg \max_{a \in A} \mathbb{E}[u(O) \mid \hat{\mathbf{E}} = \hat{\mathbf{e}}, a], \quad (1)$$

where ties are broken in some arbitrary way and O is the random variable distributed according to $P(\cdot \mid \hat{\mathbf{e}}, a)$. 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 [Zermeño, 2011; Zermeño, 2012; Chen *et al.*, 2014]. Of course, randomization comes at the cost of sometimes taking suboptimal actions. Indeed, our negative results (see Section 6 and Theorem 7) 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*. 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 [2010, Theorem 1 and 4] and Chen *et al.* [2014, 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 generally not be realistic. 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

$$u(\omega) - \sum_{i=1}^n s_i(P(\cdot \mid \hat{e}, a_{\hat{e}}), \omega). \quad (2)$$

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 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 $\text{EU}_s^i(\sigma)$ for expert (or player) i and an expected utility net of payments $\text{EU}_s(\sigma)$ 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:

$$\text{EU}_s^i(\sigma) \geq \text{EU}_s^i(\sigma_{-i}, \sigma'_i). \quad (3)$$

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 the following reason. Imagine that the available evidence is very “complementary” (in the terminology of Chen and Waggoner [2017]). That is, imagine that $v(\mathbf{E})$ is high iff $\mathbf{E} = \mathbf{E}^*$ 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* [Anshelevich *et al.*, 2008; Roughgarden and Tardos, 2007, Section 1.3], rather than price of anarchy [Papadimitriou, 2001; Koutsoupiasa and Papadimitriou, 2009].) 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 (Lemma 1) which arise from maximizing the experts’ profit.¹

¹In the mechanism design literature, it is common to consider the best equilibrium for the principal (or, more generally, for the studied objective). The reason is that, by the revelation principle, any equilibrium can be turned into the truthful equilibrium of a (truthful) mechanism, and presumably the truthful equilibrium is *focal*, i.e., natural to coordinate on. In our setting, experts do not only re-

port but also take information acquisition actions. To apply similar reasoning in such a setting, one may imagine that before each of their actions, the experts submit their beliefs to a mediator who recommends actions to each of the experts [cf. Myerson, 1986]. It then seems plausible that if reporting honestly to the mediator and following her recommendations is an equilibrium, the experts are most likely to play that focal equilibrium. In our setting, the mediator may be an additional manager who coordinates the experts.

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. This raises a problem: there is not much we can do to maximize worst-case expected utility. After all, it may be that no decision-relevant information is available to the experts. The second idea of competitive analysis is therefore to not consider worst-case expected utility period, but 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 [Hurwicz and Shapiro, 1977; Chassang, 2013; Carroll, 2015; Carroll, 2019].

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

$$v(\mathbf{E}) := \mathbb{E}_{\mathbf{E}} [\mathbb{E}_O [u(O) \mid a_{\mathbf{E}}, \mathbf{E}]] \quad (4)$$

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: com-

port but also take information acquisition actions. To apply similar reasoning in such a setting, one may imagine that before each of their actions, the experts submit their beliefs to a mediator who recommends actions to each of the experts [cf. Myerson, 1986]. It then seems plausible that if reporting honestly to the mediator and following her recommendations is an equilibrium, the experts are most likely to play that focal equilibrium. In our setting, the mediator may be an additional manager who coordinates the experts.

petitive 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 3 and Theorem 5) 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

$$\text{REGRET}_{\mathbf{s}}(\sigma) := \text{EU}_{\text{OPT}} - \text{EU}_{\mathbf{s}}(\sigma). \quad (5)$$

We will also use $\text{REGRET}_{\mathbf{s}}$ to denote the lowest regret achieved in any Nash equilibrium σ for \mathbf{s} , i.e.,

$$\text{REGRET}_{\mathbf{s}} = \min_{\sigma \in \text{NE}(\mathbf{s})} \text{REGRET}_{\mathbf{s}}(\sigma). \quad (6)$$

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 [Anshelevich *et al.*, 2008; Roughgarden and Tardos, 2007, Section 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}^α as follows:

$$q_j^\alpha(\hat{P}, \omega) = \alpha_j u(\omega). \quad (7)$$

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 Section 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.

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 [2019, Section 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 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}^α be a linear contract. Then for all configurations of available evidence, any*

$$\hat{\mathbf{E}} \in \arg \max_{\mathbf{E} \in \mathbb{H}} v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(E_i) \quad (8)$$

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

Proof. As noted earlier, \mathbf{q}^α incentivizes honest reporting if everyone else reports honestly. So we only need to show that no expert j can profit by deviating from acquiring \hat{E}_j to acquiring E'_j . From the definition of $\hat{\mathbf{E}}$ it follows that

$$v(\hat{\mathbf{E}}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(\hat{E}_i) \geq v(\hat{\mathbf{E}}_{-j}, E'_j) - \frac{1}{\alpha_j} c(E'_j) - \sum_{i \neq j} \frac{1}{\alpha_i} c(\hat{E}_i). \quad (9)$$

Adding $\sum_{i \neq j} c(\hat{E}_i)/\alpha_i$ and then multiplying by α_j yields

$$\alpha_j v(\hat{\mathbf{E}}) - c(\hat{E}_j) \geq \alpha_j v(\hat{\mathbf{E}}_{-j}, E'_j) - c(E'_j), \quad (10)$$

which means deviating is not profitable for j . \square

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

$$\text{REGRET}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) \leq \max \left(\sum_{i=1}^n \alpha_i, 1 - \alpha_{\min} \right) v(\mathbf{E}^*). \quad (11)$$

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

Proof. First, if $v(\hat{\mathbf{E}}) > v(\mathbf{E}^*)$, then the regret bound is trivially satisfied.² From now on, we assume $v(\hat{\mathbf{E}}) \leq v(\mathbf{E}^*)$. By definition of $\hat{\mathbf{E}}$ we have:

$$v(\hat{\mathbf{E}}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(\hat{E}_i) \geq v(\mathbf{E}^*) - \sum_{i=1}^n \frac{1}{\alpha_i} c(E_i^*). \quad (12)$$

Hence,

$$\begin{aligned} \alpha_{\min} v(\hat{\mathbf{E}}) &\geq \alpha_{\min} v(\hat{\mathbf{E}}) - \sum_{i=1}^n \frac{\alpha_{\min}}{\alpha_i} c(\hat{E}_i) \\ &\stackrel{\text{Ineq. 12}}{\geq} \alpha_{\min} v(\mathbf{E}^*) - \sum_{i=1}^n \underbrace{\frac{\alpha_{\min}}{\alpha_i}}_{\leq 1} c(E_i^*) \\ &\geq \alpha_{\min} v(\mathbf{E}^*) - c(\mathbf{E}^*), \end{aligned}$$

or

$$\alpha_{\min} v(\hat{\mathbf{E}}) + c(\mathbf{E}^*) \geq \alpha_{\min} v(\mathbf{E}^*). \quad (13)$$

With this, we can prove the desired bound:

$$\begin{aligned} \text{REGRET}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) &= v(\mathbf{E}^*) - c(\mathbf{E}^*) - \left(1 - \sum_{i=1}^n \alpha_i\right) v(\hat{\mathbf{E}}) \\ &= v(\mathbf{E}^*) - \left(\alpha_{\min} v(\hat{\mathbf{E}}) + c(\mathbf{E}^*)\right) \\ &\quad - \left(1 - \alpha_{\min} - \sum_{i=1}^n \alpha_i\right) v(\hat{\mathbf{E}}) \\ &\stackrel{\text{Ineq. 13}}{\leq} (1 - \alpha_{\min}) v(\mathbf{E}^*) - \left(1 - \alpha_{\min} - \sum_{i=1}^n \alpha_i\right) v(\hat{\mathbf{E}}). \end{aligned}$$

If $1 - \alpha_{\min} - \sum_{i=1}^n \alpha_i \geq 0$ and therefore $\max(1 - \alpha_{\min}, \sum_{i=1}^n \alpha_i) = 1 - \alpha_{\min}$, then the subtrahend is positive. Hence, we can drop it to obtain the desired upper bound. If $1 - \alpha_{\min} - \sum_{i=1}^n \alpha_i \leq 0$ and therefore $\max(1 - \alpha_{\min}, \sum_{i=1}^n \alpha_i) = \sum_{i=1}^n \alpha_i$, we continue by using our assumption that $v(\hat{\mathbf{E}}) \leq v(\mathbf{E}^*)$:

$$\begin{aligned} &\leq (1 - \alpha_{\min}) v(\mathbf{E}^*) - \left(1 - \alpha_{\min} - \sum_{i=1}^n \alpha_i\right) v(\mathbf{E}^*) \\ &= \left(\sum_{i=1}^n \alpha_i\right) v(\mathbf{E}^*), \end{aligned}$$

which allows us to conclude the bound in the theorem. \square

It is worth noting that any equilibrium of the game induced by \mathbf{q}^α that is not strongly Pareto-dominated by \mathbf{E}^* will also

²Configurations with $v(\hat{\mathbf{E}}) > v(\mathbf{E}^*)$ are a possibility if (and only if) the α_i are not all the same. For example, imagine $\alpha_1 = 0.01, \alpha_2 = 0.5$ and that expert 1 has exclusive access to E_1 with $v(E_1) = 0.5$ and $c_1(E_1) = 0.1$ and expert 2 has exclusive access to E_2 with $v(E_2) = 0.51$ and $c_2(E_2) = 0.15$. Assume $v(E_1, E_2) = v(E_2) = 0.51$. Then $\mathbf{E}^* = (E_1, E^0)$ and $\hat{\mathbf{E}} = (E^0, E_2)$ and hence $v(\mathbf{E}^*) = 0.5 < 0.51 = v(\hat{\mathbf{E}})$.

satisfy the regret bound of ineq. 11. In particular, if $\hat{\mathbf{E}}$ is not strongly Pareto-dominated by \mathbf{E}^* , this means there must be some expert j such that

$$\alpha_j v(\hat{\mathbf{E}}) - c(\hat{E}_j) \geq \alpha_j v(\mathbf{E}^*) - c(E_j^*). \quad (14)$$

From this we can derive ineq. 13 from the above proof:

$$\begin{aligned} &\alpha_{\min} v(\hat{\mathbf{E}}) + c(\mathbf{E}^*) \\ &= \alpha_j v(\hat{\mathbf{E}}) - c(\hat{E}_j) + c(\hat{E}_j) \\ &\quad - (\alpha_j - \alpha_{\min}) v(\hat{\mathbf{E}}) + c(\mathbf{E}^*) \\ &\stackrel{v(\mathbf{E}^*) \geq v(\hat{\mathbf{E}})}{\geq} \alpha_j v(\hat{\mathbf{E}}) - c(\hat{E}_j) + c(\hat{E}_j) \\ &\quad - (\alpha_j - \alpha_{\min}) v(\mathbf{E}^*) + c(\mathbf{E}^*) \\ &\stackrel{\text{Ineq. 14}}{\geq} \alpha_j v(\mathbf{E}^*) - c(E_j^*) + c(\hat{E}_j) \\ &\quad - (\alpha_j - \alpha_{\min}) v(\mathbf{E}^*) + c(\mathbf{E}^*) \\ &= \alpha_{\min} v(\mathbf{E}^*) + c(\mathbf{E}_{-j}^*) + c(\hat{E}_j) \\ &\geq \alpha_{\min} v(\mathbf{E}^*) \end{aligned}$$

After that we can prove the bound in the same way as in the proof. Unfortunately, it may be that all Nash equilibria of the game induced by \mathbf{q}^α are strictly Pareto-dominated by \mathbf{E}^* .

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 Section 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 Section 7.

Lemma 1 also gives us the following result, which is a generalization to the multi-expert case of a result shown by Chasang [2013, Theorem 1.i] and Carroll [2015, Section 2.3].

Proposition 3. *For all configurations of available evidence, the Nash equilibria $\hat{\mathbf{E}}$ of Lemma 1 for the linear scoring rule \mathbf{q}^α satisfy*

$$\text{EU}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) \geq \left(1 - \sum_{i=1}^n \alpha_i\right) \max_{\mathbf{E}} \left(v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(E_i)\right). \quad (15)$$

Proof.

$$\begin{aligned} \text{EU}_{\mathbf{q}^\alpha}(\hat{\mathbf{E}}) &= \left(1 - \sum_{i=1}^n \alpha_i\right) v(\hat{\mathbf{E}}) \\ &\geq \left(1 - \sum_{i=1}^n \alpha_i\right) \left(v(\hat{\mathbf{E}}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(\hat{E}_i)\right) \\ &\stackrel{\text{def. } \hat{\mathbf{E}}}{=} \left(1 - \sum_{i=1}^n \alpha_i\right) \max_{\mathbf{E}} \left(v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(E_i)\right) \end{aligned}$$

□

Proposition 3 is essentially a competitive-ratio-type result, except that the benchmark is a little weaker than the first-best solution. The term $\max_{\mathbf{E}} (v(\mathbf{E}) - \sum_{i=1}^n c(E_i)/\alpha_i)$ is the utility obtained in a first-best solution where the cost of i 's information is scaled up by α_i^{-1} . And of that, using \mathbf{q}^α guarantees a fraction of $1 - \sum_{i=1}^n \alpha_i$.

Chassang [2013, Theorem 1.ii] shows how a single-expert version of this result can be used to figure out which value of α to use when the principal knows something about the cost-to-value ratio of information. If information is known to be cheap, then α can be low. Chassang's proof only operates on (the $n = 1$ special case of) Ineq. 15. 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. For instance, the prediction accuracy for some part of the environment could give the principal strong evidence about the costs paid by the experts. 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 4. *Let $0 < H < \max_{\omega \in \Omega} u(\omega)$ and let s be a scoring rule. Then if for all configurations with $v(\mathbf{E}^*) = H$, $\text{REGRET}_s \leq nH/(n+1)$, then it must be that for all $j = 1, \dots, m$, $s_j(\hat{P}, \omega) = u(\omega)/(n+1)$, whenever $\omega \in \text{supp}(\hat{P})$.*

We briefly give a sketch of the proof, which consists of two parts. In the first part, we identify ‘‘critical cases’’ for any 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.

Proof. Define for each $i \in \{1, \dots, n\}$ and each distribution Q_a over Ω , $\phi_i(Q_a) := \mathbb{E}_{O \sim Q_a} [s_i(Q_a, O)]$ to be the expected payoff that expert i obtains after an action is recommended with an honestly reported distribution Q_a . In abuse of notation, define $\phi_i(e) := \phi_i(P(\cdot | a_e, E = e))$ for some value e of some random variable E , where a_e is again the best action for the principal according to e .

We consider the class of cases in which only one of the experts $j \in \{1, \dots, n\}$ has access to some piece of (non-trivial) information E with $v(E) = H$ and where for each action $a \in A$ there is at most one value e of E that identifies a as optimal. This last restriction ensures that honest reporting is the only way for the expert to get the principal to take the optimal action.

Now consider two types of cases for the cost of E to j :

- Imagine that $c_j(E) = 0$. Then to achieve $\text{REGRET} \leq nH/(n+1)$, expert j must weakly prefer acquiring and honestly reporting E and the overall expected payment to the experts for acquiring E must be at most $nH/(n+1)$, i.e.

$$\sum_{i=1}^n \mathbb{E}_E [\phi_i(E)] \leq \frac{n}{n+1} H. \quad (16)$$

- Imagine that $c_j(E) = H/(n+1) - \epsilon$ for any $\epsilon > 0$. Then for $\text{REGRET} \leq nH/(n+1)$, expert j still has to prefer acquiring and honestly reporting E . That is, it has to be the case that

$$\mathbb{E}_E [\phi_j(E)] - \phi_j(P) \geq c_j(E) = H/(n+1) - \epsilon. \quad (17)$$

Since this is true for all $\epsilon > 0$, it must be the case that

$$\mathbb{E}_E [\phi_j(E)] - \phi_j(P) \geq H/(n+1). \quad (18)$$

From Inequalities 16, 18 (for all j) and ex ante individual rationality it follows that for all j

$$\mathbb{E}_E [\phi_j(E)] = H/(n+1) \quad (19)$$

and

$$\phi_j(P) = 0. \quad (20)$$

Next we show that $\phi_j(Q_a) = \mathbb{E}_{O \sim Q_a} [u(O)]/(n+1)$ for all j and Q_a . The main challenge is to show that ϕ_j is affine in the expected utility of Q_a , i.e., that $\phi_j(Q_a) = \lambda \mathbb{E}_{O \sim Q_a} [u(O)] + C$ for some $\lambda, C \in \mathbb{R}$. Having shown that, it will follow immediately (from Equations 19 and 20) that $\lambda = 1/(n+1)$ and $C = 0$.

Remember that [see, e.g. Schneider and Eberly, 2003, Sect. 3.4] $\phi_j(Q_a)$ being affine in the expected value can be characterized by stating that for any two random variables R_a, R_b with the same (expected) expected utility, it is the case that $\mathbb{E}_{R_a} [\phi_j(R_a)] = \mathbb{E}_{R_b} [\phi_j(R_b)]$. Note that because in our case

ϕ_j is a function of probability distributions over Ω , R_a, R_b are random variables whose values are such probability distributions. Now notice that Eq. 19 already implies the desired equation, except it does so only for R_a, R_b that both have an expected utility of H . We now show how we can extend this to random variables with arbitrary expected utilities.

Consider any two random variables R_a, R_b over distributions over $\Delta(\Omega)$ with equal expected utility

$$\mathbb{E}_{R_a} [\mathbb{E}_{O \sim R_a} [u(O)]] = \mathbb{E}_{R_b} [\mathbb{E}_{O \sim R_b} [u(O)]] \quad (21)$$

less than $\max_{\omega \in \Omega} u(\omega)$. Then consider new random variables $\tilde{R}_a = p * R_a + (1-p) * \tilde{R}$ and $\tilde{R}_b = p * R_b + (1-p) * \tilde{R}$ s.t.

$$\mathbb{E}_{\tilde{R}_a} [\mathbb{E}_{O \sim \tilde{R}_a} [u(O)]] = H = \mathbb{E}_{\tilde{R}_b} [\mathbb{E}_{O \sim \tilde{R}_b} [u(O)]] \quad (22)$$

$0 < p \leq 1$ and \tilde{R} is some random variable over distributions over Ω . (Note that if we had allowed $H = \max_{\omega} u(\omega)$, then such $\tilde{R}_a, \tilde{R}_b, \tilde{R}, p$ might not exist.) We have already shown that it must be the case that $\mathbb{E}_{\tilde{R}_a} [\phi_j(\tilde{R}_a)] = \mathbb{E}_{\tilde{R}_b} [\phi_j(\tilde{R}_b)]$

(Eq. 19) since \tilde{R}_a, \tilde{R}_b might be the distributions arising from some evidence E with $v(E) = H$. By the definition of expected value this implies $\mathbb{E}_{R_a} [\phi_j(R_a)] = \mathbb{E}_{R_b} [\phi_j(R_b)]$. We have now shown that for any R_a, R_b , if

$$\mathbb{E}_{R_a} [\mathbb{E}_{O \sim R_a} [u(O)]] = \mathbb{E}_{R_b} [\mathbb{E}_{O \sim R_b} [u(O)]] \quad (23)$$

then $\mathbb{E}_{R_a} [\phi_j(R_a)] = \mathbb{E}_{R_b} [\phi_j(R_b)]$. This is exactly the characterization of $\phi_j(Q)$ being affine in $\mathbb{E}_{O \sim Q} [u(O)]$. As noted earlier, this shows (with eq.s 19 and 20) that $\phi_j(Q_a) = \mathbb{E}_{O \sim Q_a} [u(O)] / (n+1)$. By definition of ϕ_j , this means that for all true distributions Q_a

$$\mathbb{E}_{O \sim Q_a} [s(Q_a, O)] = \mathbb{E}_{O \sim Q_a} [u(O)] / (n+1) \quad (24)$$

All that is left now is to get rid of the expectation $\mathbb{E}_{O \sim Q_a}$. Before we can do that, we show that for all distributions Q_a, Q'_a over Ω

$$\phi_i(Q_a) \geq \mathbb{E}_{O \sim Q_a} [s_i(Q'_a, O)] \quad (25)$$

i.e., that s_i is *proper* and thus experts cannot increase their expected payoff by misreporting the distribution of the recommended action. We prove this by contradiction. Imagine there were Q'_a, Q_a, i violating inequality 25. Now consider for each expert $k \neq i$ the evidence structure where at price $H/(n+1) - \epsilon$ (for some $\epsilon > 0$), expert k can acquire a piece of evidence E_k with $v(E_k) = H$ that reveals the optimal action. Expert i can acquire evidence for free which reveals the distribution that the optimal action gives rise to, with the two possibilities being Q'_a and Q_a . Then it must be an equilibrium for k to acquire E_k . Considering $\epsilon \rightarrow 0$, this implies expert k receives at least $H/(n+1)$. Since this is true for all $k \neq i$, the principal pays (at least) $(n-1)H/(n+1)$ overall in expectation to the experts other than i . But by assumption, if i acquires E_i and reports accurately except for misreporting Q_a (when it is observed) as Q'_a , he gets in expectation strictly more than $\mathbb{E}[\phi_i(E_i, E_k)] = H/(n+1)$. Hence, the principal pays more than $nH/(n+1)$ overall in expectation – which means that the regret does not satisfy the bound.

Finally, we show that as long as $\omega \in \text{supp}(Q'_a)$, $s_i(Q'_a, \omega)$ only depends on (the utility of) the outcome [in a way that resembles the proof of Lemma 2 of Oosterheld and Conitzer, 2019]. From that it will follow that $s(Q'_a, \omega) = u(\omega)/(n+1)$ whenever $\omega \in \text{supp}(Q_a)$, as claimed. So take any distributions Q_a, Q'_a over Ω with $\text{supp}(Q_a) \subseteq \text{supp}(Q'_a)$. Then we have $Q'_a = pQ_a + (1-p)\tilde{Q}_a$ for some $p \in (0, 1]$, $\tilde{Q}_a \in \Delta(\Omega)$. Now consider the following:

$$\begin{aligned} & \mathbb{E}_{O \sim Q'_a} [s_i(Q'_a, O)] \\ &= p \mathbb{E}_{O \sim Q_a} [s_i(Q'_a, O)] + (1-p) \mathbb{E}_{O \sim \tilde{Q}_a} [s_i(Q'_a, O)] \\ &\stackrel{\text{Ineq. 25}}{\leq} p \mathbb{E}_{O \sim Q_a} [s_i(Q_a, O)] + (1-p) \mathbb{E}_{O \sim \tilde{Q}_a} [s_i(Q'_a, O)] \\ &\stackrel{\text{Ineq. 25}}{\leq} p \mathbb{E}_{O \sim Q_a} [s_i(Q_a, O)] + (1-p) \mathbb{E}_{O \sim \tilde{Q}_a} [s_i(\tilde{Q}_a, O)] \\ &\stackrel{\text{Eq. 24}}{=} p \frac{\mathbb{E}_{O \sim Q_a} [u(O)]}{n+1} + (1-p) \frac{\mathbb{E}_{O \sim \tilde{Q}_a} [u(O)]}{n+1} \\ &= \frac{\mathbb{E}_{O \sim Q'_a} [u(O)]}{n+1} \stackrel{\text{Eq. 24}}{=} \mathbb{E}_{O \sim Q'_a} [s_i(Q'_a, O)] \end{aligned}$$

Since the first and the last term are the same, the weak inequalities in the middle must be equalities. Since $p > 0$, we have $\mathbb{E}_{O \sim Q_a} [s_i(Q_a, O)] = \mathbb{E}_{O \sim Q_a} [s_i(Q'_a, O)]$ whenever $\text{supp}(Q_a) \subseteq \text{supp}(Q'_a)$. In particular, for any $\omega \in \text{supp}(Q'_a)$ if we let Q_a be the distribution which puts all probability on ω we obtain

$$\begin{aligned} s_i(Q'_a, \omega) &= \mathbb{E}_{O \sim Q_a} [s_i(Q'_a, O)] = \mathbb{E}_{O \sim Q_a} [s_i(Q_a, O)] \\ &= u(\omega)/(n+1), \end{aligned} \quad (26)$$

as claimed. \square

The different aspects of this result depend on the details of our setup to different extents. In particular, 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. For instance, imagine a case (similar to the one in the above proof) where everyone needs to acquire information for the principal to obtain a utility of H . One expert has to pay $H/(n+1) - \epsilon$ for his piece of information and the others can obtain their information for free. Because the principal does not know which expert pays a cost, she must still pay everyone $H/(n+1)$ in expectation.

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 4 holds true for the competitive ratio-based bound. We omit the proof for brevity, as we can use the same ideas as in the proof of Theorem 4.

Theorem 5. Let $\alpha \in (0, 1)^n$ with $\sum_{j=1}^n \alpha_j < 1$ and s be a scoring rule. Then if for all configurations there is Nash equilibrium $\hat{\mathbf{E}}$

$$\text{EU}_s(\hat{\mathbf{E}}) \geq \left(1 - \sum_{i=1}^n \alpha_i\right) \max_{\mathbf{E}} \left(v(\mathbf{E}) - \sum_{i=1}^n \frac{1}{\alpha_i} c(E_i)\right), \quad (27)$$

then it must be the case that for all $j = 1, \dots, m$, $s_j(\hat{P}, \omega) = \alpha_j u(\omega)$, whenever $\omega \in \text{supp}(\hat{P})$.

7 Restrictions 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 [2013, Theorem 1.ii] result, mentioned at the end of Section 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 4 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 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 4. 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.³

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

³A natural loosening of this assumption is to require that the same piece of information must cost two different experts roughly the same, but two different evidence variables can vary arbitrarily in their costs. The bound under this assumption is only slightly worse, but the analysis and bounds appear to become much more complicated.

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 (Sections 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 7) 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 of Theorem 7 makes clear where its parameters come from. Note that the proof of the positive result is somewhat different from the proofs of Theorem 2 and Proposition 3, because – as noted earlier – it does not use the Lemma 1 equilibria (which in general do not satisfy the regret bound of this theorem).

Theorem 6. Let $n \in \mathbb{N}$ be the number of experts. Given Assumptions 1, 2 and 3, define

$$B_{\Lambda, n} := 1 - \frac{1}{1 + \sum_{i=1}^n \frac{1}{1+(i-1)\Lambda}}. \quad (28)$$

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

$$\alpha_j = \frac{1}{(1 + (j-1)\Lambda) \left(1 + \sum_{i=1}^n \frac{1}{1+(i-1)\Lambda}\right)}. \quad (29)$$

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

Proof. We distinguish two cases. (1) If there is some equilibrium \mathbf{E}' of the game induced by $(\alpha_j)_{j=1, \dots, n}$ with $v(\mathbf{E}') = v(\mathbf{E}^*)$, then the regret in that equilibrium satisfies the bound:

$$\text{REGRET}(\mathbf{E}') \leq \left(\sum_{j=1}^n \alpha_j\right) v(\mathbf{E}^*) = B_{\Lambda, n} v(\mathbf{E}^*)$$

(2) Now consider the case where there is no equilibrium in which the equivalent of the information in \mathbf{E}^* is acquired. Assume that in \mathbf{E}^* , $k \geq 1$ experts acquire non-trivial information (i.e., evidence variables other than E^0). WLOG assume that \mathbf{E}^* is minimal in the sense that $v(\mathbf{E}_{-i}^*) = 0$ for all i who obtain non-trivial information. Now consider the \mathbf{E}' which arises from moving all acquisition to the first k experts (the ones with the highest α_i). If \mathbf{E}' is not an equilibrium, this means that there is some expert $j \in \{1, \dots, k\}$ s.t. $\alpha_j v(\mathbf{E}') < c(E'_j)$. This is because by Assumption 3, none of the currently not acquiring experts $k+1, \dots, n$ can increase v by acquiring. Therefore,

$$\begin{aligned} c(\mathbf{E}^*) &= c(\mathbf{E}') \geq (k-1)\Lambda c(E'_j) + c(E'_j) \\ &> (1 + (k-1)\Lambda)\alpha_j v(\mathbf{E}') \\ &= (1 + (k-1)\Lambda)\alpha_j v(\mathbf{E}^*). \end{aligned} \quad (30)$$

Hence, even if the experts acquire no relevant information at all, the regret is low:

$$\begin{aligned}
\text{REGRET}_{q^\alpha}(\mathbf{E}^0) &= v(\mathbf{E}^*) - c(\mathbf{E}^*) \\
&\stackrel{\text{Ineq. 30}}{<} (1 - (1 + (k-1)\Lambda)\alpha_j)v(\mathbf{E}^*) \\
&\leq (1 - (1 + (k-1)\Lambda)\alpha_k)v(\mathbf{E}^*) \\
&= B_{\Lambda,n}v(\mathbf{E}^*)
\end{aligned}$$

□

We now prove that the scoring rule of Theorem 6 is the only one that achieves its regret bound. Our strategy is the same as the strategy behind the proof of Theorem 4 and the omitted proof of Theorem 5. 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 4 – 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 4 that to obtain the bound, one has to use this linear rule.

Theorem 7. *Let $0 < H < \max_{\omega \in \Omega} u(\omega)$, and s be a scoring rule. Then, if for all configurations with $v(\mathbf{E}^*) = H$ that satisfy Assumptions 1, 2 and 3, we have $\text{REGRET}_s \leq B_{\Lambda,n}H$, then it is the case that – 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. 29.*

Proof. We consider two kinds of cases.

- Imagine that k different pieces of information must be acquired to achieve an expected utility of H . To $n-k+1$ experts all evidence variables cost $\alpha_k H - \epsilon$ for some small $\epsilon > 0$, where the α_k are defined as in the theorem. To the other $k-1$ experts, the evidence variables all cost $\Lambda \alpha_k H$. The regret of not acquiring information, i.e., of acquiring \mathbf{E}^0 , would be too high:

$$\begin{aligned}
\text{REGRET}_s(\mathbf{E}^0) &= H - c(\mathbf{E}^*) \\
&= (1 - (1 + (k-1)\Lambda)\alpha_k)H + \epsilon \\
&= B_{\Lambda,n}H + \epsilon
\end{aligned}$$

Hence, acquiring k evidence variables must be an equilibrium. Note that this is true for all partitions of the experts into $k-1$ experts who pay the lower price of $\Lambda \alpha_k H$ and $n-k+1$ experts who pay the higher price of $\alpha_k H - \epsilon$. To ensure that in all of these cases k experts acquire information, the scoring rule needs to ensure that at least k experts overall are willing to pay the higher price; if there were only $k-1$ who are willing to pay the higher price, then in the partition where those $k-1$ are exactly the ones who can pay the lower price,




Figure 1: Plots of the the regret bound $B_{\Lambda,n}$ for different values of Λ .

no other expert would acquire information. That is, for k experts j it must be the case that

$$\mathbb{E}[\phi_j(\mathbf{E}^*)] - \mathbb{E}[\phi_j(\mathbf{E}^0)] \geq \alpha_k H - \epsilon. \quad (31)$$

Here, ϕ_j is defined as in the proof of Theorem 4. Since this is for all $\epsilon > 0$, it must be the case that

$$\mathbb{E}[\phi_j(\mathbf{E}^*)] - \mathbb{E}[\phi_j(\mathbf{E}^0)] \geq \alpha_k H. \quad (32)$$

- Imagine that the same k different pieces of information must be acquired to achieve an expected utility of H but that all information costs ϵ . Then for regret to be at most $B_{\Lambda,n}H$ as $\epsilon \rightarrow 0$ it must be the case that

$$\sum_{i=1}^n \mathbb{E}[\phi_i(\mathbf{E}^*)] \leq B_{\Lambda,n}H. \quad (33)$$

Since $\sum_{i=1}^n \alpha_i = B_{\Lambda,n}$, inequalities 32 and 33, together with individual rationality, imply that for all j , $\mathbb{E}[\phi_j(\mathbf{E}^0)] = 0$ and whenever $v(\mathbf{E}^*) = H$ there is for each $k = 1, \dots, n$ an expert i s.t. $\mathbb{E}[\phi_i(\mathbf{E}^*)] = \alpha_k H$. By a simple continuity argument, this mapping (i.e., which expert gets which α_k) must be the same for all \mathbf{E}^* . So, up to permutation of the experts, it must be the case that for all $k = 1, \dots, n$ and all \mathbf{E} with $v(\mathbf{E}) = H$, $\mathbb{E}[\phi_k(\mathbf{E})] = \alpha_k H$. The rest of this proof proceeds just like the proof of Theorem 4. □

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 2 and 4. 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. Plots of $B_{\Lambda,n}$ for different values of Λ can be found in Figure 1.

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 7 therefore provides another (and quite different) demonstration of a point made by Winter [2004], 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 [Hurwicz and Shapiro, 1977; Chassang, 2013; Carroll, 2015; Carroll, 2019]; 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 [Lambert, 1986; Stoughton, 1993; Core and Qian, 2002; Barron and Waddell, 2003; Feess and Walzl, 2004; Zermeño, 2011; Häfner and Taylor, 2019]. 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 [Barron and Waddell, 2003; Gromb and Martimort, 2007], 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 3 – was already given in the literature for the single-expert case), Section 7 discusses issues that are very specific to the multi-expert case. To our knowledge, our main optimality arguments (the proofs of Theorems 4, 5 and 7) 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 [Hurwicz and Shapiro, 1977; Stoughton, 1993; Diamond, 1998, Carroll, 2015; Chassang, 2013; Carroll, 2019; Dütting *et al.*, 2019; Oesterheld and Conitzer, 2019, Section 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 [Othman and Sandholm, 2010; Chen *et al.*, 2014; Oesterheld and Conitzer, 2019]. 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 decision [Holmström, 1980; Crawford and Sobel, 1982; Boutilier, 2012; Milgrom and Roberts, 1986]; cf. the literature on Bayesian persuasion [Kamenica and Gentzkow, 2011].

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. 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 [cf. Barron and Waddell, 2003; Core and Qian, 2002; Lambert, 1986; Stoughton, 1993]?

References

- [Anshelevich *et al.*, 2008] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Éva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008. A preliminary version of this paper appeared in “Proc. 45th Annual Symposium on Foundations of Computer Science”, 2004.
- [Barron and Waddell, 2003] John M. Barron and Glen R. Waddell. Executive rank, pay and project selection. *Journal of Financial Economics*, 67:305–349, 2003.
- [Boutilier, 2012] Craig Boutilier. Eliciting forecasts from self-interested experts: Scoring rules for decision makers. In Conitzer, Winikoff, Padgham, and van der Hoek, editors, *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*, pages 737–744, Richland, South Carolina, 2012. International Foundation for Autonomous Agents and Multiagent Systems.
- [Carroll, 2015] Gabriel Carroll. Robustness and linear contracts. *American Economic Review*, 105(2):536–563, 2015.
- [Carroll, 2019] Gabriel Carroll. Robust incentives for information acquisition. *Journal of Economic Theory*, 181:382–420, 5 2019.
- [Chassang, 2013] Sylvain Chassang. Calibrated incentive contracts. *Econometrica*, 81(5):1935–1971, 9 2013.
- [Chen and Waggoner, 2017] Yiling Chen and Bo Waggoner. Informational substitutes, 3 2017.
- [Chen *et al.*, 2014] Yiling Chen, Ian A. Kash, Mike Ruberry, and Victor Shnayder. Eliciting predictions and recommendations for decision making. In *ACM Transactions on Economics and Computation*, volume 2, chapter 6, pages 6:1–6:27. Association for Computing Machinery, New York, 6 2014.
- [Core and Qian, 2002] John E. Core and Jun Qian. Project selection, production, uncertainty, and incentives, 1 2002.
- [Crawford and Sobel, 1982] Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6), 11 1982.
- [Diamond, 1998] Peter Diamond. Managerial incentives: on the near linearity of optimal compensation. *Journal of Political Economy*, 106(5):931–957, 10 1998.
- [Dütting *et al.*, 2019] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. Simple versus optimal contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 369–387. Association for Computing Machinery, New York, 2019.
- [Feess and Walzl, 2004] Eberhard Feess and Markus Walzl. Delegated expertise – when are good projects bad news? *Economic Letters*, 82(1):77–82, 1 2004.
- [Gromb and Martimort, 2007] Denis Gromb and David Martimort. Collusion and the organization of delegated expertise. *Journal of Economic Theory*, 137(1):271–299, 11 2007.
- [Häfner and Taylor, 2019] Samuel Häfner and Curtis R. Taylor. On young turks and yes men: Optimal contracting for advice, 9 2019.
- [Holmström, 1980] Bengt Holmström. On the theory of delegation. Discussion Papers 438, Northwestern University, Center for Mathematical Studies in Economics and Management Science, 6 1980.
- [Hurwicz and Shapiro, 1977] Leonid Hurwicz and Leonard Shapiro. Incentive structures maximizing residual gain under incomplete information. Technical Report 77–83, 4 1977.
- [Kamenica and Gentzkow, 2011] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101:2590–2615, 2011.
- [Koutsoupiasa and Papadimitriou, 2009] Elias Koutsoupiasa and Christos Papadimitriou. Worst-case equilibria. 3(2):65–69, 5 2009.
- [Lambert, 1986] Richard A. Lambert. Executive effort and selection of risky projects. *The RAND Journal of Economics*, 17(1):77–88, 1986.
- [Milgrom and Roberts, 1986] Paul Milgrom and John Roberts. Relying on the information of interested parties. *The RAND Journal of Economics*, 17(1):18–32, 1986.
- [Myerson, 1986] Roger B. Myerson. Multistage games with communication. *Econometrica*, 54(2):323–358, 3 1986.
- [Oosterheld and Conitzer, 2019] Caspar Oosterheld and Vincent Conitzer. Eliciting information for decision making from individual and multiple experts. 2019.
- [Othman and Sandholm, 2010] Abraham Othman and Tuomas Sandholm. Decision rules and decision markets. In *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pages 625–632. 2010.
- [Papadimitriou, 2001] Christos H. Papadimitriou. Algorithms, games, and the internet, 2001. Talk at STOC’01, July 6-8, 2001, Hersonissos, Crete, Greece.
- [Roughgarden and Tardos, 2007] Tim Roughgarden and Éva Tardos. Introduction to the inefficiency of equilibria. In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, editors, *Algorithmic Game Theory*, chapter 17, pages 443–459. Cambridge University Press, 2007.
- [Schneider and Eberly, 2003] Philip J. Schneider and David H. Eberly. *Geometric Tools for Computer Graphics*. Morgan Kaufmann, 2003.
- [Stoughton, 1993] Neal M. Stoughton. Moral hazard and the portfolio management problem. *The Journal of Finance*, 48(5):2009–2028, 12 1993.
- [Winter, 2004] Eyal Winter. Incentives and discrimination. *The American Economic Review*, 94(3):764–773, 6 2004.
- [Zermeño, 2011] Luis Zermeño. A principal-expert model and the value of menus. 11 2011.
- [Zermeño, 2012] Luis Zermeño. The role of authority in a general principal-expert model. 3 2012.