Incentivizing Truth Exploration and Honest Reporting: A Contract Design Approach

Yuming Shao
Tsinghua University
Beijing, China
shaoym21@mails.tsinghua.edu.cn

Zhixuan Fang
Tsinghua University
Beijing, China
Shanghai Qi Zhi Institute
Shanghai, China
zfang@mail.tsinghua.edu.cn

ABSTRACT

In this paper, we study a Principal-Agent problem in which a principal incentivizes an agent by establishing a payment contract that encourages the agent to exert costly effort in exploring the true state of the environment, which is of interest to the principal, and then report the findings. We consider two feedback setups: (1) the true state is ultimately observable by the principal, and (2) only some noisy feedback related to the true state is observable. In the first setup, we demonstrate that the optimal contract is the one that pays the agent only when the report matches the ground truth, and we derive an efficient algorithm to compute this optimal contract. In the second setup, we design a BDD contract and show its approximate optimality with respect to the optimal honest-reporting incentivizing contract, both theoretically and empirically. Furthermore, we introduce a sufficient condition under which the optimal contract encourages honest reporting.

KEYWORDS

Contract Theory; Stackelberg Game; Information Acquisition

ACM Reference Format:
Yuming Shao and Zhixuan Fang. 2025. Incentivizing Truth Exploration and Honest Reporting: A Contract Design Approach. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 9 pages.

1 INTRODUCTION

In real life, delegating the task of truth investigation is common. In business consulting, business owners may hire firms like McKinsey or Boston Consulting Group to conduct market research and develop strategic plans for them. In human resource management, companies lacking expertise in recruitment or performance evaluation may engage HR management firms to evaluate their candidates and employees. These scenarios can be modeled as a **Principal-Agent** interaction problem, where the principal values certain information that she cannot access directly due to a lack of expertise. She seeks assistance from an agent capable of uncovering this information. The difficulty she faces when implementing this delegation is that the information acquisition process typically remains opaque to her, which may result in insufficient incentives for the agent to conduct diligent research (known as the *moral hazard* issue, [16]). Since the truth investigation can be quite costly for the agent, such as requiring detailed industry surveys, he may choose to reduce his level of effort and be strategic about what to report, e.g., merely providing superficial, unsubstantiated guesses while claiming they were obtained from thorough work.

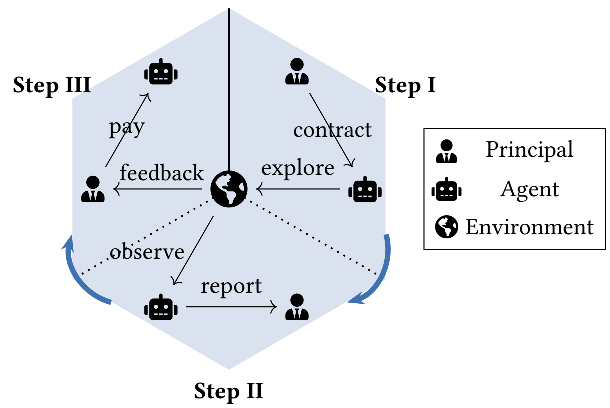

Figure 1: The Principal-Agent interaction process considered in this paper.

In this paper, we investigate the problem of incentivizing the agent to exert effort in the truth exploration process and to make a valid report from the perspective of **contract design**, following a recent trend of using contract theory to address various particular Principal-Agent problems [1, 13, 23]. Specifically, we consider a Stackelberg game [3] where the principal announces a policy of payment to the agent in advance, which specifies the exact payment she delivers to him given his actual report and some feedback that is closely related to the ground truth. The agent then selects his strategy of effort investment and reporting to maximize his net utility. We summarize the complete interaction procedure as follows (See Figure 1 for an illustration):

- **Step I.** The principal designs a contract and presents it to the agent. The agent then chooses the level of effort and begins exploring the truth.
- **Step II.** The agent observes information about the environment and then determines what to report to the principal.
- **Step III.** The principal receives feedback from the environment and then pays the agent according to the contract.

Corresponding authors: Zhixuan Fang (zfang@mail.tsinghua.edu.cn).

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

We present a real-world scenario that our model can capture in the following example.

Example 1 (Delegated Market Research). Financial markets change rapidly, and only a few professionals can accurately access internal market information, which is valuable for making trading decisions. An investor (Principal) may hire a financial expert (Agent) to investigate market dynamics. The expert conducts research and subsequently reports the findings, which the investor can use as a reference for investment decisions. The investment outcome can be characterized by public information, such as stock prices (Feedback). The investor can utilize this public information, or even the market state at the time of decision-making if available, to assess the accuracy of the expert's report and adjust payment accordingly.

In our model, the ground truth is represented by an environment state $\theta \in \Theta$, where Θ is the set of all possible environment states. We assume the agent's exploration process may not always succeed. If successful, the agent observes the correct θ; if not, no information is gained. Different levels of agent effort correlate with varying success rates and costs, with the reasonable assumption that higher success probabilities incur greater costs for the agent. Regardless of whether the exploration is successful, the form of the agent's report is an element of Θ representing his belief about the truth.

In this paper, we analyze the problem from the principal's perspective, assuming both players are utility maximizers. Our focus is on designing an optimal contract to maximize the principal's utility while considering additional factors. For example, a desirable property of a contract is that it incentivizes the agent to report findings honestly. In our model, when exploration is successful, the strategic agent is aware of the payment associated with each reporting choice and may opt for the one that maximizes payment, regardless of its truthfulness. Thus, when evaluating a specific contract, we assess its ability to promote truth-telling, as accurate reporting is typically most beneficial to the principal, while the mutual trust built by honesty increases the likelihood of long-term cooperation between the principal and agent. Consequently, the ideal contract is not only a payment policy that optimizes the principal's utility but also serves as a mechanism that encourages the agent's honesty.

1.1 Our Contributions

In this research, we explore two distinct models that differ primarily in the type of feedback received from the environment. Accordingly, our contributions are twofold:

- In the model where the true state of the environment is ultimately usable by the principal to evaluate the agent's report, we prove that the optimal contract pays only when the report matches the ground truth. Based on this, we present an algorithm that efficiently computes the optimal contract by solving a polynomial number of linear programs.
- In the model where the principal can only observe some related feedback instead of the ground truth, we design a BDD contract and prove its approximate optimality relative to the ideal truth-telling incentivizing contract, both theoretically and empirically. Furthermore, we prove that if a clear distinction exists between the benefits to the principal from truthful versus non-truthful reports, an optimal contract can be constructed to encourage honest reporting.

2 RELATED WORK

2.1 Contract Theory

Contract Theory is a traditional branch of economics, dating back to at least 1979 [14, 19]. In recent years, a series of studies has emerged that studies contracts through the lens of the theory of computation, led by [10]. The authors of [10] raise concerns regarding the complexity and unintuitiveness of optimal contracts, while exploring the approximation ratios of simpler contract forms, such as linear contracts. A line of research [15–17] on the menu of contracts is particularly relevant to our work, as the agent reporting process in our model is similar to the agent's selection of a contract from a proposed menu in their models. The difference, for example, is that in these models, the agent knows his private type in advance and does not need to exert effort to explore the environmental information, and the menu is constructed to either accommodate various types of agents for greater benefits [15, 16] or to assist in learning the private information [17]. There are other papers studying combinatorial contracts from a computational perspective [8, 9, 11, 12], but they are less related to our topic. Recently, an increasing number of works have focused on extending contract theory to other research areas. For instance, [13] studies the delegation problem of sequential probing, while [1, 23] consider contracts for machine learning tasks.

2.2 Proper Scoring Rules

Scoring rules are payment policies designed to incentivize risk-neutral experts to provide their probability assessment for an uncertain event [5]. A scoring rule is proper if the expert's optimal strategy, under any belief he might possess, is to report that belief [21]. Some recent studies [6, 18, 20] examine how to incentivize agents, using proper scoring rules, to access a costly signal that refines their beliefs. The setups of these papers differ from ours, for example, in the following aspects: (1) The form of the reports: in our model, the agent submits only a prediction of the true state, whereas scoring rules typically require a believed probability distribution over all possibilities. (2) The optimization problems: in our model, the objective is principal utility maximization, while these papers have different goals. For instance, some aim to maximize the agent's additional benefit when accessing the signal, treating the principal's preferences as budget constraints [6, 20]. The authors of [22] establish a connection between the menu of contracts and proper scoring rules within a classic hidden-action principal-agent model, featuring the novel aspect that the agent can choose to observe a costly signal correlated with the outcome. They investigate the problem of incentivizing signal acquisition at minimal cost.

3 CONTRACTING ON TRUE ENVIRONMENT STATES

3.1 The Model

We introduce the Principal-Agent model used in this section. In this model, there is a variable $\theta \in \Theta$ representing the true state of the environment. We assume Θ, the set of all possible states, to be discrete and let $m = |\Theta|$, where $|\cdot|$ denotes the size of a finite set. The state θ is randomly sampled by the Nature from a prior distribution $\mathbb{P}_\theta$. We let P_θ be the probability that the true state is realized to be θ and define $\underline{P} = \min_{\theta\in\Theta} P_\theta$. Two players interact in this environment: A **principal** (referred to as she/her) wants to uncover the true state θ and seeks to delegate this exploration task to an **agent** (referred to as he/him) with expertise. They both know the prior distribution $\mathbb{P}_\theta$. The agent's actions have two stages. In the first stage, he chooses a proper level of his exploration effort. Formally, say the agent has n actions. Each action $a \in [n]$ (throughout this paper, we let $[N]$ be the set $\{1, 2, ..., N\}$ for any $N \in \mathbb{N}^+$) is associated with a cost $c_a \geq 0$ and a success rate $q_a \in [0, 1]$ such that, with probability q_a, he successfully identifies the true state θ, while with probability $1 - q_a$, he gains no additional information about it. Without loss of generality, we assume that

(1) The more effort the agent puts in, the more likely the investigation is successful: $c_a < c_{a'} \iff q_a < q_{a'}, \forall a, a'$.
(2) The agent can invest no effort: $c_1 = 0$ and $q_1 = 0$.

In the second stage of the agent's actions, he reports his result $\hat{\theta}$, representing his prediction of the true state, to the principal. Of course, this report $\hat{\theta}$ is not necessarily identical to the truth θ. As in the classical contract theory literature [16], we assume that the principal cannot observe the agent's choice of action a (i.e., the hidden-action model).

To incentivize the agent to put in real effort in his exploration, the principal designs a contract p and offers it to the agent. We assume that the principal is trustworthy and will always fulfill the contract honestly. Both the principal and the agent are risk-neutral utility maximizers. Now, we introduce the definition of their utilities, which serve as the guidelines for contract design in this model. Define the mapping $v : \Theta \times \Theta \to \mathbb{R}^+$ such that $v(\hat{\theta}, \theta)$ is the benefit the principal gains from receiving the report $\hat{\theta}$ when the true state is θ. Note that v can be seen as an $m \times m$-dimensional non-negative matrix, so we also refer to it as a value matrix. In this section, we make the following assumption on the values of v:

Assumption 1. *The value matrix v satisfies*

$$v(\theta, \theta) - v(\theta', \theta) \geq 0, \quad \forall \theta, \theta' \in \Theta.$$

Under Assumption 1, the truthful report is the most beneficial to the principal among all possible reports. The principal attempts to elicit an accurate report via a carefully designed contract p. We consider the contract p to also be an $m \times m$-dimensional matrix, where $p(\hat{\theta}, \theta)$ is the payment delivered to the agent for his report $\hat{\theta}$ when the true state is θ. A payment mechanism contingent on both the report $\hat{\theta}$ and the truth θ is reasonable since in this section, we consider a scenario where both players can observe the true state θ after the agent submits his report $\hat{\theta}$. Following the literature on contract design [10], we focus on non-negative contracts, i.e., $p(\hat{\theta}, \theta) \geq 0, \forall \hat{\theta}, \theta \in \Theta$ (limited liability).

Given the offered contract p, the agent chooses his action and report to maximize his utility. The agent's strategy involves an effort investment strategy $s \in \Delta_{[n]}$ (we define Δ_Ω as the probability simplex on a finite set Ω) and a reporting strategy $r \in \mathcal{R}_0$, where $\mathcal{R}_0$ is the set of all possible reporting strategies and will be defined later. s_a is the probability of choosing action a. Let $\mathsf{suc} \in \{0, 1\}$ be an indicator variable of whether the exploration is successful. $\mathsf{suc} = 1$ if it is successful and $\mathsf{suc} = 0$ if not. The strategy r maps from a, suc, and the true state θ to a distribution over all possible reports. Noting that the report should be independent of θ when the exploration fails, we define

$$\mathcal{R}_0 = \Big\{ r : [n] \times \{0, 1\} \times \Theta \to \Delta_\Theta \,\Big|\, r(a, 0, \theta) = r(a, 0, \theta'), \forall a \in [n], \theta, \theta' \in \Theta \Big\}.$$

Let $\mathrm{Bin}(q)$ denote the Bernoulli distribution with mean q. The agent's expected utility u_{0A} depending on his strategy (s, r) can be written as

$$u_{0A}(s, r) = \mathbb{E}_{\theta\sim\mathbb{P}_\theta, a\sim s, \mathsf{suc}\sim\mathrm{Bin}(q_a), \hat{\theta}\sim r(a,\mathsf{suc},\theta)} \left[p(\hat{\theta}, \theta) - c_a \right].$$

We will show in Proposition 1 that the following pure strategy is the agent's optimal strategy given p, breaking ties towards the benefit of the principal:

$$a = \arg\max_{a'\in[n]} - c_{a'} + q_{a'} \sum_{\theta\in\Theta} P_\theta p(\bar{r}(\theta), \theta) + (1 - q_{a'}) \sum_{\theta\in\Theta} P_\theta p(\underline{r}, \theta), \tag{1}$$

$$\bar{r}(\theta) = \arg\max_{\theta'\in\Theta} p(\theta', \theta), \quad \forall \theta \in \Theta, \tag{2}$$

$$\underline{r} = \arg\max_{\theta'\in\Theta} \sum_{\theta\in\Theta} P_\theta p(\theta', \theta). \tag{3}$$

In this strategy, when the exploration is successful, he reports $\bar{r}(\theta)$ which maximizes the payment given the true state θ. When the exploration fails, he reports $\underline{r}$ which maximizes the expected payment given the prior distribution of θ. The agent's action a is the one that maximizes his expected utility.

Proposition 1. *The pure strategy $(a, \{\bar{r}(\theta)\}_{\theta\in\Theta}, \underline{r})$ defined in (1), (2), and (3) achieves an expected agent utility of $\max_{(s,r)\in\Delta_{[n]}\times\mathcal{R}_0} u_{0A}(s, r)$, given any contract p.*

Now that the agent's behavior is determined by the contract p, we can write the principal's expected utility as

$$u_{0P}(p) = \mathbb{E}_{\theta\sim\mathbb{P}_\theta, \mathsf{suc}\sim\mathrm{Bin}(q_a)} \left[v(\hat{\theta}, \theta) - p(\hat{\theta}, \theta) \right] = \sum_{\theta\in\Theta} P_\theta \Big[q_a \big(v(\bar{r}(\theta), \theta) - p(\bar{r}(\theta), \theta) \big) + (1 - q_a) \big(v(\underline{r}, \theta) - p(\underline{r}, \theta) \big) \Big],$$

noting that the agent's report $\hat{\theta} = \bar{r}(\theta)\mathbb{I}\{\mathsf{suc} = 1\} + \underline{r}\mathbb{I}\{\mathsf{suc} = 0\}$. In summary, the contract design problem faced by the principal can be formalized as a program:

$$\text{(P1)} \quad \max_p \quad u_{0P}(p)$$

$$s.t. \quad \sum_{\theta\in\Theta} P_\theta \big[p(\bar{r}(\theta), \theta) - p(\underline{r}, \theta) \big] (q_a - q_{a'}) \geq c_a - c_{a'}, \forall a' \in [n] \tag{4}$$

$$p(\bar{r}(\theta), \theta) \geq p(\theta', \theta), \quad \forall \theta, \theta' \in \Theta, \tag{5}$$

$$\sum_{\theta\in\Theta} P_\theta p(\underline{r}, \theta) \geq \sum_{\theta\in\Theta} P_\theta p(\theta', \theta), \quad \forall \theta' \in \Theta, \tag{6}$$

$$p(\theta, \theta') \geq 0, \quad \forall \theta, \theta' \in \Theta,$$

$$\underline{r} \in \Theta, a \in [n], \bar{r}(\theta) \in \Theta, \quad \forall \theta \in \Theta.$$

Constraints (4), (5), and (6) correspond to the agent's pure strategy (1), (2), and (3), respectively. This program has $m + 2$ discrete variables, namely a, $\underline{r}$, and $\{\bar{r}(\theta)\}_{\theta\in\Theta}$. Although P1 becomes a linear program when a realization of these variables is fixed, there are

Algorithm 1: Computing an Optimal Solution to P1

Initialize p to be an $m \times m$-dimensional zero matrix

`saved-z` $\leftarrow$ `NULL`, `saved-obj` $\leftarrow$ `-Inf`

for $a \in [n]$ *and* $\underline{r} \in \Theta$ **do**

  Solve the following linear program: $\vec{z} \leftarrow$

$$\begin{aligned} \min_{\vec{h}} \quad & q_a \sum_{\theta \in \Theta} P_\theta h_\theta + (1 - q_a) P_{\underline{r}} h_{\underline{r}}, \\ s.t. \quad & P_{\underline{r}} h_{\underline{r}} \geq P_\theta h_\theta, \forall \theta \neq \underline{r}, \\ & (q_a - q_{a'}) \sum_{\theta : \theta \neq \underline{r}} P_\theta h_\theta \geq c_a - c_{a'}, \forall a' \neq a, \\ & h_\theta \geq 0, \forall \theta \in \Theta. \end{aligned}$$

  `obj` $\leftarrow q_a \sum_{\theta \in \Theta} P_\theta \big(v(\theta, \theta) - z_\theta\big) + (1 - q_a)\big[\sum_{\theta \in \Theta} P_\theta v(\underline{r}, \theta) - P_{\underline{r}} z_{\underline{r}}\big]$

  if `obj` > `saved-obj` **then**

    `saved-obj`, `saved-z` $\leftarrow$ `obj`, $\vec{z}$

  end

end

$\vec{z} \leftarrow$ `saved-z`

for $\theta \in \Theta$ **do**

  $p(\theta, \theta) \leftarrow z_\theta$

end

Output: p the solution to P1

$n \cdot m^{m+1}$ possible realizations. Thus, it is prohibitive to solve this program by enumerating all realizations of the discrete variables.

3.2 Computing the Optimal Contract Efficiently

We introduce an efficient solution for P1 and present it in Algorithm 1. It solves $n \cdot m$ linear programs. Each program has m variables and $2m + n - 2$ inequality constraints. We present the following result justifying the correctness of our method.

THEOREM 2. *Algorithm 1 computes an optimal solution to P1.*

The key ingredient of its proof is demonstrating that there exists an optimal solution to P1 within a small but reasonable family of non-negative $m \times m$-dimensional matrices.

Definition 3 (Diagonal Contract). We say a contract p is diagonal if and only if it satisfies

$$p(\theta', \theta) = 0, \quad \forall \theta, \theta' \in \Theta \quad s.t. \quad \theta' \neq \theta.$$

For any diagonal contract, the entries are zero except for the ones exactly on the diagonal, so the agent only gets paid when his prediction is correct. In the following lemma, we show that to solve the program P1, it suffices to consider only the diagonal contracts.

LEMMA 1. *Let $p_0 : \Theta \times \Theta \to \mathbb{R}^+$ be any non-negative contract. Then there exists a non-negative diagonal contract whose objective value is no smaller than that of p_0.*

PROOF SKETCH OF LEMMA 1. We show a transformation from p_0 into a non-negative diagonal contract p_3 that yields non-decreasing principal utility, involving three steps:

(i) $p_0 \longrightarrow p_1$. For any $\hat{\theta}, \theta \in \Theta$, if $\hat{\theta}$ is neither $\underline{r}$ nor $\bar{r}(\theta)$, we set $p_1(\hat{\theta}, \theta) = 0$. Otherwise, we set $p_1(\hat{\theta}, \theta) = p_0(\hat{\theta}, \theta)$. This modification does not change the agent's choice of effort or his reporting strategy. We can also show that the principal's utility remains unchanged.

(ii) $p_1 \longrightarrow p_2$. We shift the highest payment on each column of p_1 to its diagonal and reduce all of p_1's non-diagonal elements in the row indexed by $\underline{r}$ to 0. This gives a diagonal contract p_2. Let $\theta^* = \arg\max_{\theta' \in \Theta} \sum_{\theta \in \Theta} P_\theta p_2(\theta', \theta) = \arg\max_{\theta \in \Theta} P_\theta p_2\big(\theta, \theta\big)$ denote agent's actual report at exploration failure in terms of p_2. Then, let $\theta_0 \in \arg\max_{\theta' \in \Theta} \sum_{\theta \in \Theta} P_\theta v(\theta', \theta)$ denote the principal's most desired report when agent's exploration fails.

(iii) $p_2 \longrightarrow p_3$. We need to adjust p_2 to match θ^* with θ_0. We 'swap' $p_2\big(\theta^*, \theta^*\big)$ and $p_2\big(\theta_0, \theta_0\big)$ to induce p_3, such that

$$p_3(\theta_0, \theta_0) = \frac{P_{\theta^*}}{P_{\theta_0}} p_2\big(\theta^*, \theta^*\big),$$

$$p_3(\theta^*, \theta^*) = \frac{P_{\theta_0}}{P_{\theta^*}} p_2\big(\theta_0, \theta_0\big),$$

to incentivize the agent to report θ_0 when his exploration fails. It can be shown that our modification will not reduce the principal's utility.

Figure 2 illustrates an example of this transformation process. □

The proof of Theorem 2 is straightforward with Lemma 1, thus we defer it to the appendix [24]. An important implication of Lemma 1 is that, in the model where the true state is finally accessible to the principal, the objectives of optimizing the principal's utility and encouraging the agent to tell the truth are fully compatible. When the principal can incentivize the agent to both work hard and provide honest information, her utility will be maximized.

4 CONTRACTING ON NOISY FEEDBACK

In the previous section, we assumed that the principal can ultimately observe the true state of the environment, θ, and make payments to the agent accordingly. However, accessing this information is often infeasible in practice. For instance, consider a client (the Principal) seeking advice from an expert (the Agent) about the financial market. While she can easily gather public data like stock prices, it is impractical for her to immediately verify the expert's recommendations due to either her limited financial expertise or the extensive time and effort required for validation. In this section, we explore the possibility of incentivizing the agent to investigate the true state and truthfully present his findings, even when the agent's report cannot be directly validated.

4.1 The Model

The Principal-Agent model used here is similar to that in the last section. There are m possible environment states, i.e., $m = |\Theta|$. The Nature samples θ from a prior distribution $\mathbb{P}_\theta$. A principal delegates the truth exploration task to an agent. They both know the prior distribution $\mathbb{P}_\theta$. The agent first chooses his level of effort and then presents his report $\hat{\theta} \in \Theta$ to the principal.

Unlike our previous model, we consider a scenario where the principal faces a noisy feedback $X \in \mathbb{R}$ that is jointly determined by

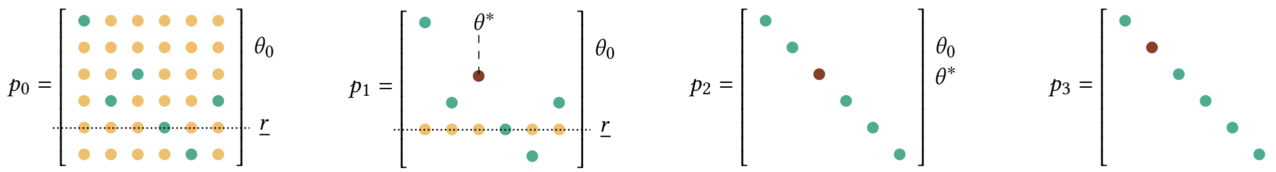

Figure 2: An example that demonstrates how to transform p_0 into a diagonal contract p_3 in the proof of Lemma 1. The elements in p_0 at positions $(\bar{r}(\theta), \theta)$ for all $\theta \in \Theta$ are shown as green nodes, while the others are shown as yellow nodes. The element $p_1(\bar{r}(\theta^*), \theta^*)$ is shown as a brown node. Each empty slot in the matrices indicates that the corresponding payment is 0.

the agent's report $\hat{\theta}$ and the true state θ, but can never directly observe θ. For example, X could be the market value of the principal's portfolio (adjusted after considering the agent's market research report) on a particular trading day, while θ can be the true state of the market. We assume that the random feedback admits an expectation captured by the value matrix v. Specifically,

$$X = v(\hat{\theta}, \theta) + E, \quad E \sim F_\eta,$$

where E is a zero-mean stochastic noise with a cumulative distribution function F_η. Due to the presence of random noise, it is difficult for the principal to infer the true underlying state solely from observing X. Here, we make a slightly stronger assumption (Assumption 2) about the value matrix compared to Section 3. This assumption implies that the principal's profit from receiving truthful reports is at least δ higher than that from receiving non-truthful reports, where δ is a known constant.

Assumption 2. *For a constant $\delta \geq 0$, the value matrix v satisfies*

$$v(\theta, \theta) - v(\theta', \theta) \geq \delta, \quad \forall \theta, \theta' \in \Theta.$$

As in the previous model, the principal is interested in incentivizing the agent to exert effort in his truth investigation via a contract p. Since the only information the principal has is the agent report $\hat{\theta}$ and the random feedback X, it is reasonable to focus on the set of contracts $\mathcal{P} := \{p : \Theta \times \mathbb{R} \to \mathbb{R}^+\}$. In this section, we consider a simpler setting where the agent has only two levels of effort: $c_1 = 0, q_1 = 0$ and $c_2 = c, q_2 = 1$ for a constant $c > 0$. That is, the agent can either uncover the truth θ at a price c or choose to receive no information without any charge.

Given the offered contract p, the agent chooses his action and report to maximize his utility. $s \in \Delta_{[2]}$ and $r \in \mathcal{R}$ are the agent's effort investment strategy for truth exploration and reporting strategy, respectively. In this section, we define the set of all possible agent reporting strategies

$$\mathcal{R} = \Big\{ r : [2] \times \Theta \to \Delta_\Theta \,\Big|\, r(1, \theta) = r(1, \theta'), \forall \theta, \theta' \in \Theta \Big\}.$$

The agent's expected utility u_A can be written as

$$u_A(s, r) = \mathbb{E}_{\theta \sim \mathbb{P}_\theta, a \sim s, \hat{\theta} \sim r(a, \theta), E \sim F_\eta} \big[p(\hat{\theta}, X) - c_a \big].$$

Let $F_{X|\hat{\theta},\theta}$ be the distribution function of X when the report is $\hat{\theta}$ and the truth is θ. It can be shown that the following pure strategy is the agent's optimal strategy given p, breaking ties towards the benefit of the principal:

$$a = \begin{cases} 2, & \text{if } \sum_{\theta \in \Theta} P_\theta \int_{\mathbb{R}} p(\bar{r}(\theta), x) dF_{X|\bar{r}(\theta),\theta} - c \geq \\ & \quad \sum_{\theta \in \Theta} P_\theta \int_{\mathbb{R}} p(\underline{r}, x) dF_{X|\underline{r},\theta}, \\ 1, & \text{otherwise}, \end{cases} \tag{7}$$

$$\bar{r}(\theta) = \arg\max_{\theta' \in \Theta} \int_{\mathbb{R}} p(\theta', x) dF_{X|\theta',\theta}, \quad \forall \theta \in \Theta, \tag{8}$$

$$\underline{r} = \arg\max_{\theta' \in \Theta} \sum_{\theta \in \Theta} P_\theta \int_{\mathbb{R}} p(\theta', x) dF_{X|\theta',\theta}, \tag{9}$$

where the integrals are w.r.t. the variable x. In this strategy, the agent reports $\bar{r}(\theta)$ when he invests effort in exploration and reports $\underline{r}$ when he does not exert any effort. The following result justifies the optimality of this strategy.

Proposition 2. *The pure strategy $\big(a, \{\bar{r}(\theta)\}_{\theta \in \Theta}, \underline{r}\big)$ defined in (7), (8) and (9) achieves an expected agent utility of $\max_{(s,r) \in \Delta_{[2]} \times \mathcal{R}} u_A(s, r)$, given any contract $p \in \mathcal{P}$.*

Now that the agent's behavior is determined by the contract p, we can write the principal's expected utility as

$$\begin{aligned} u_P(p) &= \mathbb{E}_{\theta \sim \mathbb{P}_\theta, E \sim F_\eta} \big[v(\hat{\theta}, \theta) - p(\hat{\theta}, X) \big] \\ &= \sum_{\theta \in \Theta} P_\theta \Big[v(\hat{\theta}, \theta) - \int_{\mathbb{R}} p(\hat{\theta}, x) dF_{X|\hat{\theta},\theta} \Big], \end{aligned}$$

where the agent's report is given by $\hat{\theta} = \bar{r}(\theta)\mathbb{I}\{a = 2\} + \underline{r}\mathbb{I}\{a = 1\}$. In summary, the contract design problem faced by the principal can be formalized as a program:

$$\begin{aligned} \text{(P2)} \quad & \max_{p \in \mathcal{P}} \quad u_P(p) \\ s.t. \; & \sum_{\theta \in \Theta} P_\theta \Big[\int_{\mathbb{R}} p(\bar{r}(\theta), x) dF_{X|\bar{r}(\theta),\theta} - \int_{\mathbb{R}} p(\underline{r}, x) dF_{X|\underline{r},\theta} \Big] \\ & \quad \times (q_a - q_{a'}) \geq c_a - c_{a'}, \quad \forall a' \in [2] \\ & \int_{\mathbb{R}} p(\bar{r}(\theta), x) dF_{X|\bar{r}(\theta),\theta} \geq \int_{\mathbb{R}} p(\theta', x) dF_{X|\theta',\theta}, \quad \forall \theta, \theta' \in \Theta, \\ & \sum_{\theta \in \Theta} P_\theta \Big[\int_{\mathbb{R}} p(\underline{r}, x) dF_{X|\underline{r},\theta} - \int_{\mathbb{R}} p(\theta', x) dF_{X|\theta',\theta} \Big] \geq 0, \forall \theta' \in \Theta, \\ & \underline{r} \in \Theta, a \in [2], \bar{r}(\theta) \in \Theta, \quad \forall \theta \in \Theta. \end{aligned}$$

To simplify this program, we make the following observation. For any contract $p \in \mathcal{P}$, if it incentivizes zero agent effort, i.e., $a = 1$, the principal's utility is at most

$$u_P(p) \leq \sum_{\theta \in \Theta} P_\theta v(\underline{r}, \theta) \leq \max_{\hat{\theta} \in \Theta} \sum_{\theta \in \Theta} P_\theta v(\hat{\theta}, \theta) =: \underline{u}_P.$$

This utility upper bound can be achieved by a zero-payment contract p $s.t.$ $p(\hat{\theta}, x) = 0, \forall \hat{\theta} \in \Theta, x \in \mathbb{R}$. One interpretation of the zero-payment contract is that when the utility derived from delegating the truth exploration task is insufficient, it is more advantageous for the principal to make a prediction herself based on her prior knowledge $\mathbb{P}_\theta$. This observation implies that solving P2 only requires finding an optimal positive-effort-incentivizing contract and comparing its utility against the baseline principal utility $\underline{u}_P$.

4.2 Our Contract Design

In this part, we propose a feasible solution to P2. We aim to design a contract that encourages the agent to truthfully report his findings while maximizing the principal's net utility. Before presenting our design in detail, we first introduce a lower bound on the expected payment, as it provides insight into our contract design. This lower bound characterizes the maximum principal utility that any contract can generate under the truthful reporting constraint.

Proposition 3. *Suppose the noise distribution F_η has a probability density function ϕ_η. A lower bound for the payment that the principal has to make to incentivize positive effort and truth-telling is given by the following linear program*

$$(L1) \quad LB := \min_{\vec{t}} \sum_{\theta \in \Theta} t_\theta$$

$$s.t. \quad \sum_{\theta \in \Theta} t_\theta - c \geq \Big[\inf_{s \in \mathbb{R}} \sum_{\theta \in \Theta} \alpha_{\theta, \theta'}(s) \Big] t_{\theta'}, \quad \forall \theta' \in \Theta, \quad (10)$$

$$t_\theta \geq \Big[\inf_{s \in \mathbb{R}} \alpha_{\theta, \theta'}(s) \Big] t_{\theta'}, \quad \forall \theta, \theta' \in \Theta, \quad (11)$$

$$t_\theta \geq 0, \quad \forall \theta \in \Theta, \quad (12)$$

where $\alpha_{\theta,\theta'}(s) := \big[P_\theta \phi_\eta\big(s - v(\theta', \theta)\big)\big] / \big[P_{\theta'} \phi_\eta\big(s - v(\theta', \theta')\big)\big]$.

We explain this result. For any contract p, define $\vec{y}$ $s.t.$ $y_\theta = P_\theta \int_{\mathbb{R}} p(\theta, x) dF_{X|\theta,\theta}, \forall \theta \in \Theta$. (10) and (11) can be shown to be relaxed versions of the positive effort and truth-telling constraints, respectively. Thus, if p incentivizes positive effort and truth-telling, $\vec{y}$ must satisfy (10) and (11). This leads to $\sum_{\theta \in \Theta} y_\theta \geq$ LB, as $\vec{y}$ is a feasible solution to this program. We complete the proof by observing that $\sum_{\theta \in \Theta} y_\theta$ is precisely the expected payment of p. Intuitively, for p to generate principal utility close to the optimum, the original positive effort and truth-telling constraints it induces should be close to the relaxed versions (10) and (11).

We propose contracts of the form:

$$p(\hat{\theta}, x) = \begin{cases} B_{\hat{\theta}}, & \text{if } |x - v(\hat{\theta}, \hat{\theta})| \leq \rho_{\hat{\theta}}, \\ 0, & \text{otherwise,} \end{cases} \quad \forall x \in \mathbb{R}, \forall \hat{\theta} \in \Theta.$$

B_θ, ρ_θ are non-negative parameters for any $\theta \in \Theta$. Since $p(\hat{\theta}, \cdot)$ approaches a Dirac pulse at $x = v(\hat{\theta}, \hat{\theta})$ when $\rho_{\hat{\theta}} \to 0$, we refer to this type of contract as **Bounded Dirac Delta** (BDD) contracts. The detailed computation process is given in Algorithm 2. This algorithm first computes the weighted expected payment associated with each true state, represented as the vector $\vec{z}$. It then sets the values of B_θ to align the expected payments of p with $\vec{z}$. Our proposed contract has three advantages:

(1) **Ex-Post Boundedness.** There exists an ex-post budget constraint B such that $B_\theta \leq B, \forall \theta \in \Theta$, which is necessary for a contract to be feasible in practice.

Algorithm 2: Bounded Dirac Delta (BDD) Contract

Input: $\{\rho_\theta\}_{\theta \in \Theta}$ payment radius

Initialize $p(\theta, x) = 0, \forall \theta \in \Theta, x \in \mathbb{R}$

Solve the following linear program: $\vec{z} \leftarrow$

$$\min_{\vec{t}} \sum_{\theta \in \Theta} t_\theta$$

$$s.t. \quad \sum_{\theta \in \Theta} t_\theta - c \geq \Bigg[\sum_{\theta \in \Theta} \frac{P_\theta \int_{v(\theta',\theta') - v(\theta',\theta) - \rho_{\theta'}}^{v(\theta',\theta') - v(\theta',\theta) + \rho_{\theta'}} dF_\eta}{P_{\theta'} \int_{-\rho_{\theta'}}^{\rho_{\theta'}} dF_\eta} \Bigg] \cdot t_{\theta'}, \forall \theta' \in \Theta,$$

$$t_\theta \geq \Bigg[\frac{P_\theta \int_{v(\theta',\theta') - v(\theta',\theta) - \rho_{\theta'}}^{v(\theta',\theta') - v(\theta',\theta) + \rho_{\theta'}} dF_\eta}{P_{\theta'} \int_{-\rho_{\theta'}}^{\rho_{\theta'}} dF_\eta} \Bigg] t_{\theta'}, \forall \theta, \theta' \in \Theta,$$

$$t_\theta \geq 0, \forall \theta \in \Theta.$$

if $\sum_{\theta \in \Theta} z_\theta \leq \sum_{\theta \in \Theta} P_\theta v(\theta, \theta) - \underline{u}_P$ **then**
  for $\theta \in \Theta$ **do**
    Compute $B_\theta \leftarrow z_\theta / \big[P_\theta \int_{-\rho_\theta}^{\rho_\theta} dF_\eta\big]$
    Update the contract

$$p(\theta, x) \leftarrow \begin{cases} B_\theta, & \text{if } |x - v(\theta, \theta)| \leq \rho_\theta, \\ 0, & \text{otherwise,} \end{cases} \quad \forall x \in \mathbb{R}$$

  end
end

Output: p the contract

---

(2) **Near Optimality.** It can be shown that, under mild assumptions on F_η and when B is large, the positive effort and truth-telling constraints approach (10) and (11).

(3) **Computational Simplicity.** We note that for BDD contracts, the positive effort and truth-telling constraints are linear constraints without relaxation. Thus, the optimal payment can be efficiently computed by solving a linear program in Algorithm 2, which is similar to L1 in Proposition 3.

Define $\bar{l} = \max_{\theta,\theta' \in \Theta} |v(\theta, \theta) - v(\theta, \theta')|$, $\underline{l} = \min_{\theta \neq \theta'} |v(\theta, \theta) - v(\theta, \theta')|$. To facilitate our theoretical analysis of this method, we make the following assumptions on the noise distribution F_η.

Assumption 3. *The noise distribution F_η has a probability density function ϕ_η satisfying:*

- *ϕ_η is symmetric: $\phi_\eta(x) = \phi_\eta(-x), \forall x \in \mathbb{R}$,*
- *ϕ_η is monotonically non-increasing in $\mathbb{R}^+$: $\phi_\eta(x_1) \geq \phi_\eta(x_2)$, $\forall 0 \leq x_1 \leq x_2$,*
- *ϕ_η is L-Lipschitz on $[-\bar{l}, \bar{l}]$,*
- *$\phi_\eta(x - d)/\phi_\eta(x) \geq \phi_\eta(-d)/\phi_\eta(0), \forall x \leq 0, d \geq 0$.*

It can be validated that the family of Laplace distributions satisfies this assumption.

Definition 4 (Laplace Distribution). For any constant $\lambda > 0$, a random variable has a zero-mean Laplace distribution, denoted Laplace$(0, 1/\lambda)$, if its probability density function is

$$\phi(x) = \frac{\lambda}{2} \exp\big(-\lambda |x|\big), \quad \forall x \in \mathbb{R}.$$

We formally present our theoretical result for the proposed contract in Theorem 5. Say p is a BDD contract computed by Algorithm

2, we show that the gap between the highest possible utility generated by a truth-telling incentivizing contract and the utility yielded by p converges to 0 at a speed of $O(1/B)$ when B increases. That is, our contract is a good approximation of the optimal truth-telling incentivizing contract. In the remainder of this subsection, we outline the proof idea for this theorem.

Theorem 5. *Let p be a contract generated by Algorithm 2 with inputs*

$$\rho_\theta = F_\eta^{-1}\left[\frac{1}{2} + \frac{cP_\theta^{-1}B^{-1}}{1 - \max_{\theta' \in \Theta} \sum_{\tilde{\theta} \in \Theta} P_{\tilde{\theta}} \frac{\phi_\eta(v(\theta', \theta') - v(\theta', \tilde{\theta}))}{\phi_\eta(0)}}\right]$$

for any $\theta \in \Theta$, where $F_\eta^{-1}(y) = \min\{x \in \mathbb{R} \mid F_\eta(x) = y\}, \forall y \in (0, 1)$. Suppose Assumption 2, 3 hold. Then $\exists B_0 \in \mathbb{R}^+, \forall B \geq B_0$, contract p has the following properties:

(1) The budget constraint is never violated: $B_\theta \leq B, \forall \theta \in \Theta$.
(2) The agent is incentivized to report the true state that he observes after the exploration.
(3) The difference between the principal utility generated by any truth-telling incentivizing contract p_0 and the utility induced by contract p is upper bounded as

$$u_P(p_0) - u_P(p) \leq \frac{12\phi_\eta(0)^2 Lc^2}{[\phi_\eta(0) - \phi_\eta(\underline{l})]^4 \underline{P}^2 B}.$$

We introduce the proof idea of Theorem 5. The first two claims in Theorem 5 can be easily validated. For the third one, we notice that since both p_0 and p incentivize the agent to report honestly, they generate the same expected profit $\sum_{\theta \in \Theta} P_\theta v(\theta, \theta)$ for the principal. Thus, it suffices to consider the difference between their expected payments. The key to upper-bounding this difference is as follows:

Since the payment of p is defined by the optimal value of another linear program (in Algorithm 2), which can be shown to be similar to L1, we derive a sensitivity analysis result to demonstrate that the optimal values of both programs are close to each other. This completes the proof of Theorem 5.

Lemma 2. *Consider the following two linear programs,*

$$\begin{aligned} (O1) \quad OPT := \min_{\vec{t}} \quad & \sum_{\theta \in \Theta} t_\theta \\ s.t. \quad & b_{\theta'} t_{\theta'} - \sum_{\theta \in \Theta} t_\theta + c \leq 0, \quad \forall \theta' \in \Theta, \\ & a_{\theta,\theta'} t_{\theta'} - t_\theta \leq 0, \quad \forall \theta, \theta' \in \Theta, \theta \neq \theta', \\ & -t_\theta \leq 0, \quad \forall \theta \in \Theta \end{aligned} \qquad (13)$$

and O2 (OPT' is its optimal value). O2 has the same structure as O1, except that in O2, $a'_{\theta,\theta'}$ replaces $a_{\theta,\theta'}$ and $b'_{\theta'}$ replaces $b_{\theta'}$ for any $\theta, \theta' \in \Theta$. We require $0 \leq a_{\theta,\theta'} \leq a'_{\theta,\theta'} < P_\theta / P_{\theta'}, \forall \theta, \theta' \in \Theta, \theta \neq \theta'$ and $0 \leq b_\theta \leq b'_\theta < 1/P_\theta, \forall \theta \in \Theta$. If there exist series of small positive constants $\{\epsilon_{\theta,\theta'}\}_{\theta,\theta' \in \Theta}, \{\kappa_\theta\}_{\theta \in \Theta}$ such that $0 \leq a'_{\theta,\theta'} - a_{\theta,\theta'} \leq \epsilon_{\theta,\theta'}, \forall \theta, \theta' \in \Theta, \theta \neq \theta'$ and $0 \leq b'_\theta - b_\theta \leq \kappa_\theta, \forall \theta \in \Theta$, then we have that

$$\frac{OPT'}{OPT} \leq \left[1 - \frac{(1 - \lambda) + \max_{\theta \in \Theta} \kappa_\theta}{1 - \max_{\theta \in \Theta} b_\theta P_\theta}\right]^{-1},$$

where $\lambda \in (0, 1)$ is a constant such that

$$1 - \lambda \leq \max_{\theta \neq \theta'} \frac{\epsilon_{\theta,\theta'}}{[P_\theta - a'_{\theta,\theta'} P_{\theta'}] + \epsilon_{\theta,\theta'}}.$$

Proof Sketch of Lemma 2. Assume that $\vec{t}^*$ is an optimal solution to O1. Our goal is to construct a solution to O2 whose sum is as close to $\sum_{\theta \in \Theta} t_\theta^*$ as possible. The first step is to construct an intermediate solution $\vec{y}$ from $\vec{t}^*$ such that $\vec{y}$ satisfies the corresponding constraint of (13) in O2. We achieve this by finding a vector $\vec{t}'$ that satisfies this constraint and defining $\vec{y} = \lambda \vec{t}^* + (1 - \lambda)\vec{t}'$ for a constant $\lambda \in (0, 1)$. We show that a $\lambda \longrightarrow 1$ can guarantee a $\vec{y}$ we want. The second step is to find a constant $\mu > 1$ such that $\mu \vec{y}$ is a feasible solution to O2. We conclude this proof by showing that there exists a $\mu \longrightarrow 1$ which satisfies our requirement. □

Remark. One may notice that the above analysis does not apply to the Gaussian distribution due to its violation of Assumption 3. In the appendix, we present an analysis for another family of noise distributions, including Gaussian distributions. We prove an upper bound for $u_P(p_0) - u_P(p)$ that converges to 0 as the Gaussian variance σ^2 decreases, which is verified by our experiments.

4.3 Dropping the Honest Reporting Constraint

In Section 4.2, our proposed BDD contract was evaluated against the best possible contract that must encourage the agent to report honestly, rather than the optimal contract that purely maximizes the principal's net utility. In this context, one might naturally ask the following question:

Does the optimal contract necessarily incentivize truth-telling, so that the two benchmarks above are simply equivalent?

We demonstrate that the answer is negative by presenting the following counterexample. The principal's value matrix is set to be

$$v = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 100 & 0 \\ 0 & 0 & 0 & 100 \end{bmatrix}$$

with the dimension $m = 4$. Let θ_i denote the true state corresponding to the i-column of the value matrix. We set the prior distribution $P_\theta = \frac{1}{m}, \forall \theta \in \Theta$. The agent's cost of positive effort $c = 1$. The noise distribution F_η is a Laplace distribution with a sufficiently large λ.

The solution to the linear program L1 for this problem instance is $\vec{y} = \left[\frac{1}{6}, \frac{1}{6}, \frac{2}{3}, \frac{2}{3}\right]$, indicating that the expected payment must be at least $\frac{5}{3}$ to incentivize both positive effort and truth-telling. This solution implies that the expected payment when the true state is θ_1 and the agent tells the truth is equal to that when the true state is θ_2 and the agent tells the truth. Such equality is necessary for the adopted contract to incentivize truth-telling, given the structure of the value matrix v. If the payment associated with θ_1 were lower than that with θ_2, the agent would gain more by reporting $\hat{\theta} = \theta_2$ when the true state is θ_1, thereby violating the truth-telling constraint. However, we can construct a contract that breaks this constraint while delivering less payment, thus showing that the optimal contract does not incentivize truth-telling. The constructed contract p implements $\vec{y}' = \left[\frac{1}{6} - \frac{13}{500}, \frac{1}{6} + \frac{3}{50}, \frac{2}{3} - \frac{1}{10}, \frac{2}{3} - \frac{1}{10}\right]$, a slightly perturbed version of $\vec{y}$. This contract does not incentivize truth-telling (the agent reports $\hat{\theta} = \theta_2$ when the true state is θ_1) but reduces the expected payment to $\frac{5}{3} - \frac{2}{25} < \frac{5}{3}$.

Although the answer is negative in general, we do identify a sufficient condition for the two benchmarks to be equivalent: Assumption 2 holds with a sufficiently large δ. Recall that given this assumption, for any realized true state θ, the principal profit decreases by at least δ if the agent is incentivized to report any $\hat{\theta} \neq \theta$. When the reduced cost can never offset the decreased profit, the optimal contract must incentivize both positive effort and truth-telling simultaneously. We formalize this intuition in Theorem 6 under the following assumption. Define $\bar{v} = \max_{\theta \in \Theta} v(\theta, \theta)$ and $\underline{v} = \min_{\theta \in \Theta} v(\theta, \theta)$.

Assumption 4. *The noise distribution F_η has a probability density function ϕ_η satisfying:*

- *ϕ_η is symmetric: $\phi_\eta(x) = \phi_\eta(-x)$, $\forall x \in \mathbb{R}$,*
- *ϕ_η is monotonically non-increasing in $\mathbb{R}^+$: $\phi_\eta(x_1) \geq \phi_\eta(x_2)$, $\forall 0 \leq x_1 \leq x_2$.*

Theorem 6. *For any constants $\lambda, \nu \in (0, 1)$ s.t. $\lambda > \nu$, suppose Assumption 2 with a sufficiently large δ such that the following conditions holds: (a) $\delta \geq \lambda \bar{v} \bigvee \frac{1}{2\underline{P}}\left[c + \sqrt{c^2 + 16c\bar{v}\underline{P}}\right]$, (b) $\phi_\eta(\delta - \nu\bar{v}) \leq \frac{8}{8+\lambda}\frac{\underline{v}}{\bar{v}}\phi_\eta(0)$, (c) $\phi_\eta(\delta/2 - \nu\bar{v}) \leq \frac{\lambda}{8+\lambda}\underline{P}\phi_\eta(0)$ and Assumption 4 hold. Given any contract p_0 such that $\exists \theta \in \Theta, \bar{r}(\theta) \neq \theta$, there exists a contract p satisfying that*

(1) p is ex-post upper bounded in its value,
(2) p incentivizes honest reporting,
(3) $u_P(p) \geq u_P(p_0)$.

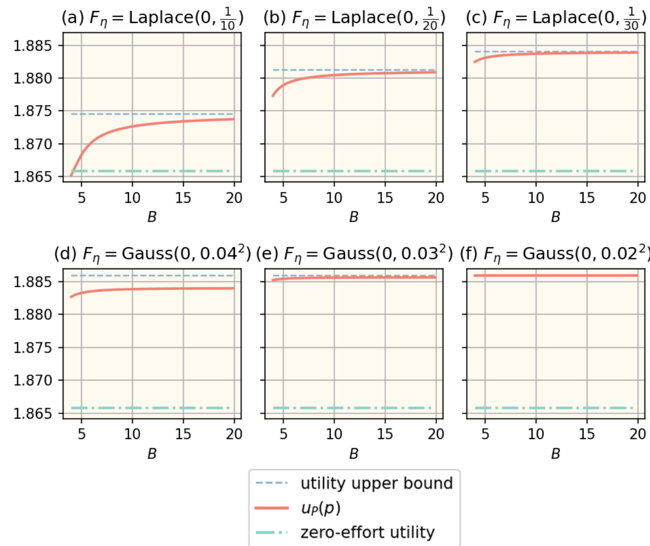

Figure 3: Expected principal utility generated by the proposed BDD contract. We consider two families of noise distributions. In (a)-(c) we set F_η to be the Laplace distribution, while in (d)-(f) we set it to be the Gaussian distribution. As the budget constraint B grows sufficiently large, the generated principal utilities approach the upper bounds. This demonstrates the near-optimality of our proposed contract.

5 EMPIRICAL RESULTS

In this section, we present our experiment results. Our machine is a PC running Windows 11, equipped with an AMD Ryzen 9 5900X 12-Core Processor and an NVIDIA GeForce RTX 3060 GPU.

5.1 Evaluating the BDD Contract

In this experiment, we evaluate the principal utility generated by the proposed BDD contract. We set $m = 5$, $c = 0.2$ and randomly generate a 5×5-dimensional value matrix v. We set a uniform prior distribution over Θ. We consider various settings of the noise distribution F_η:

– Laplace distribution with $\lambda = 10, 20, 30$.
– Gaussian distribution with $\sigma = 0.04, 0.03, 0.02$.

The results are shown in Figure 3. In each subfigure, the horizontal axis is the budget constraint B. In addition to the utility generated by the positive-effort incentivizing BDD contract, we also illustrate the utility upper bound $\sum_{\theta \in \Theta} P_\theta v(\theta, \theta) - \text{LB}$ and the zero-effort utility $\underline{u}_P$ for comparison. Note that although $u_P(p)$ can be lower than $\underline{u}_P$ in subfigure (a), the utility generated by our Algorithm 2 is always at least $\underline{u}_P$. Our experiment demonstrates that the proposed BDD contract is a good approximation of the optimal truth-telling incentivizing contract, especially when the budget constraint B is large.

6 CONCLUSION

In this paper, we propose a principal-agent problem in which a principal incentivizes an agent to undertake a costly exploration of the truth and report the findings through a payment contract. For different setups of feedback information that the principal can use to assess the quality of the agent's report, we demonstrate the importance and efficiency of encouraging the agent to submit honest reports and design our contract solutions accordingly. All omitted proofs can be found in the technical appendix [24].

Future Work. We believe there is still much to explore within the delegated truth exploration framework: (1) In this paper, we assume a principal with extensive knowledge, such as the prior distribution of the environment state, the agent's action set (including success rates and costs), and so on. Inspired by a recent work [6] on proper scoring rules, we find it interesting to extend our discussion to a partial knowledge setting. Moreover, several recent works have focused on the repeated principal-agent interaction setting, e.g., [2, 4, 7, 17]. It would also be valuable to extend our framework to a multi-round variant where the game starts with unknown parameters. (2) In this paper, we assume strict incentive compatibility, meaning the agent is always a utility maximizer. However, the agent may also be willing to follow the principal's suggestion, provided the utility is sufficiently close to optimal. We term this the approximate incentive compatibility assumption. Contract design problems in this context remain unexplored.

ACKNOWLEDGMENTS

This work is supported by Tsinghua University Dushi Program and Shanghai Qi Zhi Institute Innovation Program SQZ202312.

REFERENCES

- [1] Nivasini Ananthakrishnan, Stephen Bates, Michael Jordan, and Nika Haghtalab. 2024. Delegating data collection in decentralized machine learning. In *International Conference on Artificial Intelligence and Statistics*. PMLR, 478–486.
- [2] Francesco Bacchiocchi, Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. 2024. Learning Optimal Contracts: How to Exploit Small Action Spaces. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net. https://openreview.net/forum?id=WKuimaBj4I
- [3] Curtis Bechtel and Shaddin Dughmi. 2021. Delegated Stochastic Probing. *Innovations in Theoretical Computer Science (ITCS)* (2021).
- [4] Yurong Chen, Zhaohua Chen, Xiaotie Deng, and Zhiyi Huang. 2024. Are Bounded Contracts Learnable and Approximately Optimal? *arXiv preprint arXiv:2402.14486* (2024).
- [5] Yiling Chen and Ian A. Kash. 2011. Information elicitation for decision making. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1* (Taipei, Taiwan) *(AAMAS '11)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 175–182.
- [6] Yiling Chen and Fang-Yi Yu. 2023. Optimal Scoring Rule Design under Partial Knowledge. arXiv:2107.07420 [cs.GT] https://arxiv.org/abs/2107.07420
- [7] Natalie Collina, Varun Gupta, and Aaron Roth. 2024. Repeated Contracting with Multiple Non-Myopic Agents: Policy Regret and Limited Liability. *CoRR* abs/2402.17108 (2024). https://doi.org/10.48550/ARXIV.2402.17108 arXiv:2402.17108
- [8] Ramiro Deo-Campo Vuong, Shaddin Dughmi, Neel Patel, and Aditya Prasad. 2024. On supermodular contracts and dense subgraphs. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 109–132.
- [9] Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. 2022. Combinatorial contracts. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 815–826.
- [10] Paul Dütting, Tim Roughgarden, and Inbal Talgam-Cohen. 2019. Simple versus optimal contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation*. 369–387.
- [11] Paul Dutting, Tim Roughgarden, and Inbal Talgam-Cohen. 2021. The complexity of contracts. *SIAM J. Comput.* 50, 1 (2021), 211–254.
- [12] Tomer Ezra, Michal Feldman, and Maya Schlesinger. 2024. On the (in) approximability of combinatorial contracts. In *15th Innovations in Theoretical Computer Science Conference (ITCS 2024)*. Schloss-Dagstuhl-Leibniz Zentrum für Informatik.
- [13] Tomer Ezra, Michal Feldman, and Maya Schlesinger. 2024. Sequential Contracts. *arXiv preprint arXiv:2403.09545* (2024).
- [14] Sanford J. Grossman and Oliver D. Hart. 1983. An Analysis of the Principal-Agent Problem. *Econometrica* 51, 1 (1983), 7–45. http://www.jstor.org/stable/1912246
- [15] Guru Guruganesh, Jon Schneider, Joshua Wang, and Junyao Zhao. 2023. The power of menus in contract design. In *Proceedings of the 24th ACM Conference on Economics and Computation*. 818–848.
- [16] Guru Guruganesh, Jon Schneider, and Joshua R Wang. 2021. Contracts under moral hazard and adverse selection. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 563–582.
- [17] Minbiao Han, Michael Albert, and Haifeng Xu. 2024. Learning in online principal-agent interactions: The power of menus. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 17426–17434.
- [18] Jason D. Hartline, Liren Shan, Yingkai Li, and Yifan Wu. 2023. Optimal Scoring Rules for Multi-dimensional Effort. In *Proceedings of Thirty Sixth Conference on Learning Theory (Proceedings of Machine Learning Research, Vol. 195)*, Gergely Neu and Lorenzo Rosasco (Eds.). PMLR, 2624–2650. https://proceedings.mlr.press/v195/hartline23a.html
- [19] Bengt Holmström. 1979. Moral hazard and observability. *The Bell journal of economics* (1979), 74–91.
- [20] Yingkai Li, Jason D. Hartline, Liren Shan, and Yifan Wu. 2022. Optimization of Scoring Rules. In *Proceedings of the 23rd ACM Conference on Economics and Computation* (Boulder, CO, USA) *(EC '22)*. Association for Computing Machinery, New York, NY, USA, 988–989. https://doi.org/10.1145/3490486.3538338
- [21] Caspar Oesterheld and Vincent Conitzer. 2020. Decision Scoring Rules. In *Web and Internet Economics - 16th International Conference, WINE 2020, Beijing, China, December 7-11, 2020, Proceedings (Lecture Notes in Computer Science, Vol. 12495)*, Xujin Chen, Nikolai Gravin, Martin Hoefer, and Ruta Mehta (Eds.). Springer, 468. https://link.springer.com/content/pdf/10.1007/978-3-030-64946-3.pdf#page=471
- [22] Maneesha Papireddygari and Bo Waggoner. 2022. Contracts with Information Acquisition, via Scoring Rules. In *Proceedings of the 23rd ACM Conference on Economics and Computation* (Boulder, CO, USA) *(EC '22)*. Association for Computing Machinery, New York, NY, USA, 703–704. https://doi.org/10.1145/3490486.3538261
- [23] Eden Saig, Inbal Talgam-Cohen, and Nir Rosenfeld. 2024. Delegated classification. *Advances in Neural Information Processing Systems* 36 (2024).
- [24] Yuming Shao and Zhixuan Fang. 2025. Incentivizing Truth Exploration and Honest Reporting: A Contract Design Approach. https://people.iiis.tsinghua.edu.cn/~fang/paper/contract_2025_full. Accessed: February 21, 2025.