

Optimal Scoring Rules for Multi-dimensional Effort

Jason D. Hartline*

Department of Computer Science, Northwestern University.

HARTLINE@NORTHWESTERN.EDU

Liren Shan*

Department of Computer Science, Northwestern University.

LIRENSHAN2023@U.NORTHWESTERN.EDU

Yingkai Li*

Cowles Foundation for Research in Economics, Yale University.

YINGKAI.LI@YALE.EDU

Yifan Wu*

Department of Computer Science, Northwestern University.

YIFAN.WU@U.NORTHWESTERN.EDU

Editors: Gergely Neu and Lorenzo Rosasco

Abstract

This paper develops a framework for the design of scoring rules to optimally incentivize an agent to exert a multi-dimensional effort. This framework is a generalization to strategic agents of the classical knapsack problem (cf. [Briest, Krysta, and Vöcking, 2005](#); [Singer, 2010](#)) and it is foundational to applying algorithmic mechanism design to the classroom. The paper identifies two simple families of scoring rules that guarantee constant approximations to the optimal scoring rule. The truncated separate scoring rule is the sum of single dimensional scoring rules that is truncated to the bounded range of feasible scores. The threshold scoring rule gives the maximum score if reports exceed a threshold and zero otherwise. Approximate optimality of one or the other of these rules is similar to the bundling or selling separately result of [Babaioff, Immorlica, Lucier, and Weinberg \(2014\)](#). Finally, we show that the approximate optimality of the best of those two simple scoring rules is robust when the agent’s choice of effort is made sequentially.

Keywords: Scoring rules, mechanism design, loss function

1. Introduction

This paper considers mechanism design for the classroom. An instructor aims to design a grading mechanism that incentivizes learning, learning comes from costly effort on the part of a student, and the student aims to optimize their grade less the costs of effort. Two key aspects of this model for mechanism design are that effort is multi-dimensional over a set of assigned tasks and that effort may lead to only partial understanding of each task, i.e., effort does not generally guarantee the student gets an answer correct. The paper formulates this problem as a multi-dimensional strategic version of the knapsack problem and solves it by giving a simple and computationally efficient scoring rule that incentivizes effort on an approximately optimal set of tasks.

Strategic versions of the knapsack problem and multi-dimensional mechanism design are of central interest in algorithmic mechanism design. For example, classic models describe knapsack mechanisms for allocation (e.g., [Briest, Krysta, and Vöcking, 2005](#)) and for procurement (e.g., [Singer, 2010](#)). An important new frontier for algorithmic mechanism design is in incentivizing

* The order of the authors are certified random. The records are available in <https://www.aeaweb.org/journals/policies/random-author-order/search?RandomAuthorsSearch%5Bsearch%5D=ZyH1Ob3fJMgs>. This work is supported by NSF CCF-2229162 and NSF CCF-1733860. Arxiv version: <https://arxiv.org/pdf/2211.03302.pdf>.

private effort, e.g., to impact states as in contract theory (Dütting, Ezra, Feldman, and Kesselheim, 2022), or to collect information as in scoring rules (this paper). Optimization of scoring rules for single-dimensional effort was considered by Li, Hartline, Shan, and Wu (2022). This paper considers multi-dimensional effort where key steps in the analysis resemble those of the well studied bundling-or-selling-separately result of the multi-dimensional mechanism design literature (Babaioff, Immorlica, Lucier, and Weinberg, 2014, 2020).

Mechanism design for the classroom has the potential to address a key challenge for the two decade old field of algorithmic mechanism design. To test the theories of mechanism design in practice, the mechanisms must be run in practice. Unlike in classical algorithm design, where new algorithms can be empirically evaluated on canonical data sets; empirical validation of mechanisms fundamentally requires that their inputs be from agents that are strategically responding to (other agents and) the new mechanism. Researchers of algorithmic mechanism design do not generally have opportunities to test the classical models of allocation or procurement. Due to this challenge most mechanisms of the algorithmic mechanism design literature have never been empirically tested. The classroom applications of mechanism design, as proposed by this paper, provide immediate opportunities for a dialogue between theory and practice; and their advances can lead to better learning outcomes for students. For example, Li, Hartline, Shan, and Wu (2022) motivate their work on optimizing scoring rules for single-dimensional effort by an empirical failure of the classical quadratic scoring rule to provide sufficient incentives of effort for peer grading.

The *knapsack scoring problem* formulated and solved in this paper is as follows. There is a universe of tasks that an instructor could assign to a student. Effort of the student on each task is binary. Each task has a fixed learning value and a fixed cost of effort. The instructor aims to maximize the sum of values of the tasks that the student puts effort on. If effort were directly observable, then this problem would be identical to the knapsack problem: the optimal set of tasks to assign is the solution to the knapsack problem with knapsack capacity equal to the maximum grade and the student receives this maximum grade if effort is exerted on all of the assigned tasks (zero otherwise). Our instructor cannot directly observe effort, but can instead administer a binary test for each task where the student’s belief about the answer to the test improves with effort. The instructor aims to select the set of tasks that the student should perform and design a scoring rule with bounded total score that incentivizes the student to perform these tasks.

How does the instructor select the tasks? And how should the instructor score the student in aggregate? The paper shows that there are two main cases that must be considered. Consider the case that scores from individual scoring rules for the optimal set of tasks concentrate, e.g., because the student is successful at many of them. In this case then a good set of tasks to incentivize can be found by greedily selecting tasks by the ratio of value to cost and a *truncated separate scoring rule* can incentivize effort on these tasks. If the scores do not concentrate then approximately optimal effort can be incentivized by the *threshold scoring rule* and the tasks for this scoring rule can be identified by greedily selecting tasks by the ratio of value to probability that the student’s effort is informative. This observation is robust to whether the agent exerts effort simultaneously or sequentially.

Machine Learning Applications We are motivated by the education example, but the work in this paper also has applications in outsourced computation. In outsourced computation, the principal (she) delegates computation to a strategic agent (he) (Belenkiy et al., 2008; Gennaro et al., 2010; Dong et al., 2017). In our problem, the principal not only incentivizes the agent to train the algo-

rithm, but also to obtain costly samples. The principal has a value over the learning outcome, which is submodular in samples. The principal faces a budget constraint. Our paper designs payment rules for the principal to incentivize the agent to acquire costly samples and to report the learning outcome.

Related Work Prior work has considered mechanism design problems based on strategic versions of the knapsack problem. One framing is that of single-minded multi-unit demand agents as buyers with a seller with a multi-unit supply constraint. In this model, only the values of the agents can be strategically manipulated. [Briest, Krysta, and Vöcking \(2005\)](#) considered welfare maximization with this framing and gave a general method for converting polynomial time approximation schemes (including the one for knapsack) into incentive compatible mechanisms (with the same approximation guarantees). [Aggarwal and Hartline \(2006\)](#) considers the same framing with the goal of revenue maximization and a natural prior-free benchmark.

Another knapsack framing reverses the buyer and seller roles: The agents are sellers with private costs (object sizes in knapsack) and the buyer aims to hire a team (set of sellers) to maximize value but has a budget constraint (capacity of the knapsack). [Singer \(2010\)](#) posed this question and gave prior-free approximation mechanisms when the buyers value function is submodular (generalizing the linear value function of the traditional knapsack problem). [Bei, Chen, Gravin, and Lu \(2012, 2017\)](#) considered the budget-feasibility question in the Bayesian and prior-independent models of mechanism design and give constant approximations. [Balkanski and Hartline \(2016\)](#) consider the Bayesian budget feasibility problem and showed that posted pricing mechanisms give good approximation to the Bayesian optimal mechanism. In comparison to the literature on budget feasibility, this paper’s model of scoring rule optimization has a single agent (resp. multiple agents) with a multi-dimensional strategy space (resp. single-dimensional), the costs are public (resp. private), but effort is private (resp. public). With private effort, the principal optimizing a scoring rule can only validate the agent’s effort in so far as the agent’s posterior information from effort improves over her prior information.

Multi-dimensional mechanism design problems are notoriously difficult. In the classical setting of selling multiple items to a single agent with multi-dimensional preferences, the algorithmic mechanism design literature has identified simple constant-approximation mechanisms in a number of canonical settings. [Babaioff, Immorlica, Lucier, and Weinberg \(2014, 2020\)](#) show that for an agent with independent additive values for multiple items then the better of bundling or selling separately is a constant approximation. [Rubinstein and Weinberg \(2015, 2018\)](#) extend this approximation result to agents with subadditive valuations. See [Babaioff, Immorlica, Lucier, and Weinberg \(2020\)](#) for discussion of the extensive literature generalizing these results. These bundling versus selling separately results are paralleled by this paper’s result showing that the better of truncated separate scoring or threshold scoring is a constant approximation.

[Chen and Waggoner \(2016\)](#) consider a setting where a principal selects signal structures with knapsack constraints on the set of realizable signals. They show that when signals are substitutes, there exists a constant approximation algorithm for signal selection. However, in the general case, no algorithm can achieve a constant approximation with subexponentially many queries to the value of a signal. In our paper, we focus on the setting where the value function is submodular, which can be seen as a special case of substitutional signals if the agent’s incentive is ignored. However, we prove that when the principal faces the task of designing an incentive scheme for the agent to select the set of signals, finding the optimal solution is NP-hard.

This work builds on the general framework for optimizing scoring rules for effort that was initiated by [Li, Hartline, Shan, and Wu \(2022\)](#). Their main result considers binary effort and multi-dimensional state. In contrast, the model of this paper is for multi-dimensional effort and multi-dimensional state, but with a 1-to-1 correspondence between the dimension of effort and state.

[Chen and Yu \(2021\)](#) consider the design of scoring rules for maximizing a binary effort in a max-min design framework. For example, complementing a prior-independent result from [Li, Hartline, Shan, and Wu \(2022\)](#), they show that the quadratic scoring rule is max-min optimal over a large family of distributional settings. [Kong \(2022\)](#) apply the framework of effort-maximization to multi-agent peer prediction where the principal does not have access to the ground truth state and instead must compare reports across several agents.

Several papers look at optimizing for multiple levels of a single-dimensional effort with the objective of accuracy of the forecast (i.e., the posterior from effort which is reported in a proper scoring rule). [Osband \(1989\)](#) considers optimization of quadratic scoring rules with a continuous level of effort. [Zermeno \(2011\)](#) characterizes the optimal single-dimensional scoring rule when the states are partially verifiable. [Neyman, Noarov, and Weinberg \(2021\)](#) consider optimization of scoring rules for integral levels of effort where the effort corresponds to a number of costly samples drawn. [Papireddygari and Waggoner \(2022\)](#) characterizes the optimal scoring rule that maximizes revenue subject to a information cost, with limited liability constraint.

Optimization of effort in scoring rules has similarities to the problem of optimizing effort in contracts, the main difference being that, in the classical model of contract design the distribution over states for each action is common knowledge. In contract for scoring rules, on taking an action the agent receives a signal that gives the agent private information about the distribution of states. For the contract design problems, [Castiglioni, Marchesi, and Gatti \(2022\)](#) show that the optimal contract can be computed in time polynomial in the number of potential actions of the agent even when the costs of actions are private information. For the multi-dimensional effort model, the number of actions is exponential in the size of the dimensions, and [Dütting, Ezra, Feldman, and Kesselheim \(2022\)](#) show that with binary states, the optimal contract can be computed in polynomial time if the function mapping the action choices to the state distributions satisfies the gross substitutes property, but is NP-hard when the function is more generally submodular.

Future Directions The approach of the paper is one of Bayesian mechanism design where the prior distribution is known to both the principal (instructor) and agent (student). Within the Bayesian model there are three main directions for future work. First, the positive results of this paper are restricted to simplistic distributions over posteriors. As discussed in [Appendix D](#), generalizing the results beyond this case necessitates better upper bounds and richer families of approximation mechanisms. Second, our multi-dimensional effort-to-state mapping is one-to-one. It is an open direction to combine results for multi-dimensional effort with the model of [Li, Hartline, Shan, and Wu \(2022\)](#) for single-dimensional effort with multi-dimensional state. Third, for our motivating application in the classroom, the cost of effort varies across students. It is an open direction to combine our model for optimizing scoring rules with the model of budget feasibility where the cost of effort is private.

Bayesian mechanism design is the first model in which to consider novel mechanism design problems. To obtain practical mechanisms, however, it is important to consider robust versions of the problem. The two canonical frameworks are that of prior-independence and sample complexity. Prior-independent framework looks to identify one mechanism that has the best approximation to the

Bayesian optimal mechanism in worst case over distributions. The sample complexity framework looks to bound the number of samples necessary to obtain a $1 + \epsilon$ approximation to the Bayesian optimal mechanism. Li, Hartline, Shan, and Wu (2022), for example, gave such results for the problem of designing scoring rules for a single-dimensional effort. These are open directions for optimizing multi-dimensional effort via scoring rules.

2. Preliminaries

This paper considers the problem of incentivizing effort from an agent to learn about an unknown state. There are n tasks with state space $\Omega = \times_{i=1}^n \Omega_i$ where $\Omega_i = \{0, 1\}$. For each task $i \in [n]$, state $\omega_i \in \Omega_i = \{0, 1\}$ is realized independently according to prior distribution D which is the uniform distribution on Ω_i . Exerting effort on task i induces cost c_i to the agent. The agent can choose to exert effort on a set $\Psi \subseteq [n]$ of tasks at a cost $\sum_{i \in \Psi} c_i$. Let Σ be the signal space where $\perp \in \Sigma$ is an uninformative signal. If the agent does not exert effort on task i , i.e. $i \notin \Psi$, with probability 1, the agent receives an uninformative signal $\sigma_i = \perp$ regardless of the realized state. If the agent exerts effort on task i , i.e. $i \in \Psi$, the agent receives a signal $\sigma_i \in \Sigma$ according to a signal structure, which is a random mapping from the states to the signal space. Note that the signal structure on task i induces a distribution f_i over posterior $\mu_i \in \Delta(\Omega)$.

A special case that is of particular interest for our paper is when $\Sigma = \{0, 1, \perp\}^n$ and the posterior belief is supported on $\{0, 1, 1/2\}^n$. In this case, if the agent exerts effort on task i , i.e. $i \in \Psi$, with probability p_i , the agent receives an informative signal $\sigma_i = \omega_i$, and with probability $1 - p_i$, the agent receives an uninformative signal $\sigma_i = \perp$ regardless of the realized state. We call p_i the state revelation probability of each task i . In the main body of the paper, we will focus on this special model, and discuss the extensions to general information structures in Appendix D.

Given the set of tasks Ψ that the agent exerts effort on, the value of the principal is $v(\Psi)$. We assume that the valuation function v is *submodular*: for every $\Psi' \subseteq \Psi \subseteq [n]$ of assignments, the principal's marginal value decreases, i.e.

$$\forall i \in [n] \setminus \Psi, v(\Psi' \cup \{i\}) - v(\Psi') \geq v(\Psi \cup \{i\}) - v(\Psi).$$

A special case of the submodular valuation is additive valuation, where $v(\Psi) = \sum_{i \in \Psi} v_i$ for given profile of $\{v_i\}_{i \in [n]}$. The goal of the principal is to design a mechanism that maximizes her value subject to the budget constraint, i.e., the payment to the agent is bounded between 0 and 1. Note that if the effort choice of the agent can be observed by the principal, this problem reduces to the classical knapsack problem. The novel feature in our model is that effort is unobservable, and the principal can only score the agent according to the reported signals and realized states.

2.1. Static Effort Model

In the static effort model, we assume that the agent makes the effort choice on all tasks simultaneously, and after the effort choice, the agent receives the signals on all tasks simultaneously. By the revelation principle, it is without loss to restrict attention to mechanisms that recommend a set of tasks Ψ for the agent to exert effort, and after exerting effort, incentivize the agent to truthfully report the received signal to the principal. Let $\Pr_{\sigma \sim \Psi}[\cdot]$ and $\mathbf{E}_{\sigma \sim \Psi}[\cdot]$ be the probability and expectation with respect to the distribution over signals conditional on exerting effort on set Ψ , and let $\Pr_{\omega \sim \sigma}[\cdot]$ and $\mathbf{E}_{\omega \sim \sigma}[\cdot]$ be the probability and expectation with respect to the posterior belief of the agent conditional on receiving signal $\sigma \in \Sigma$.

Definition 1 A scoring rule $S : \Sigma \times \Omega \rightarrow [0, 1]$ is proper if for any $\sigma, \sigma' \in \Sigma$,

$$\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)] \geq \mathbf{E}_{\omega \sim \sigma}[S(\sigma', \omega)].$$

Note that our definition of properness relies on the information structure and the set of signal realizations Σ . In principle, a scoring rule that satisfies our definition of properness may incentive the agent to misreport his belief that cannot be induced by those signal realizations. This may raise a concern for the robustness of the implemented scoring rule. In Appendix B, we show that it is without loss of generality to focus on scoring rules that are only proper for signal realizations in the support, by converting any such scoring rule to one that is proper for all possible beliefs without performance loss.

Definition 2 A mechanism composed by a scoring rule $S : \Sigma \times \Omega \rightarrow [0, 1]$ and a corresponding recommendation set Ψ is incentive compatible if S is proper and for any $\bar{\Psi} \subseteq [n]$,¹

$$\mathbf{E}_{\sigma \sim \bar{\Psi}}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] - \sum_{i \in \bar{\Psi}} c_i \geq \mathbf{E}_{\sigma \sim \bar{\Psi}}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] - \sum_{i \in \bar{\Psi}} c_i.$$

The reward of the agent should be non-negative and the principal has a budget of 1 for rewarding the agent. Thus, the score is ex-post bounded in $[0, 1]$. Given the incentive constraints and reward constraints, the timeline of our model is as follows:

1. The principal commits to an incentive compatible mechanism with scoring rule $S : \Sigma \times \Omega \rightarrow [0, 1]$ and recommendation set Ψ .
2. The agent chooses a set $\bar{\Psi}$ of tasks on which to exert effort and pays cost $\sum_{i \in \bar{\Psi}} c_i$.
3. States $\omega = \{\omega_i\}_{i=1}^n$ are realized, and the agent receives the signals $\sigma \in \Sigma$.
4. The agent reports σ' and receives score $S(\sigma', \omega)$.
5. The principal receives utility $v(\bar{\Psi})$.

Note that the agent is incentivized to choose $\bar{\Psi} = \Psi$ and truthfully reveal the signals in an incentive compatible mechanism. The *knapsack scoring problem* for value function v , costs $\{c_i\}_{i=1}^n$ and state revelation probabilities $\{p_i\}_{i=1}^n$ is formally defined as the following optimization program:

$$\begin{aligned} \text{IC-OPT}(v, \{c_i\}_{i=1}^n, \{p_i\}_{i=1}^n) &= \max_{S, \Psi} v(\Psi) \\ \text{s.t. } & (S, \Psi) \text{ is incentive compatible for } \{c_i\}_{i=1}^n \text{ and } \{p_i\}_{i=1}^n, \\ & S(\sigma, \omega) \in [0, 1], \quad \forall \sigma, \omega. \end{aligned}$$

We use the *knapsack problem* for value function v and costs c_i without incentive constraints as an upper bound on the knapsack scoring problem:

$$\begin{aligned} \text{ALG-OPT}(v, \{c_i\}_{i=1}^n) &= \max_{\Psi \subseteq [n]} v(\Psi) \\ \text{s.t. } & \sum_{i \in \Psi} c_i \leq 1. \end{aligned}$$

1. An alternative formulation of the mechanism is to only specify the scoring rule and delegate the computation of the optimal effort choice to the agent. However, the computation of the optimal effort choice may be NP-hard. The main advantage of our formulation is that we can ensure that the computation of the agent is simple.

It is easy to see that $\text{ALG-OPT}(v, \{c_i\}_{i=1}^n) \geq \text{IC-OPT}(v, \{c_i\}_{i=1}^n, \{p_i\}_{i=1}^n)$ for any $v, \{c_i\}_{i=1}^n$ and $\{p_i\}_{i=1}^n$.

The following characterization shows the budget-minimal scoring rule for incentivizing a single task. To minimize the budget, the agent is indifferent between: 1) reporting truthfully and non-truthfully; 2) exerting effort and not exerting the effort on the task.

Lemma 3 (Li, Hartline, Shan, and Wu, 2022) *With minimal budget $\frac{2c_i}{p_i}$, the agent can be incentivized to exert effort on a single task $\Psi = \{i\}$ with cost c_i and probability p_i of revealing. Moreover, the budget-minimal scoring rule for incentivizing effort is²*

$$S_i(\sigma_i, \omega_i) = \begin{cases} \frac{c_i}{p_i} & \sigma = \perp \\ \frac{2c_i}{p_i} \mathbf{1}[\sigma_i = \omega_i] & \text{otherwise.} \end{cases}$$

By Theorem 3, with budget 1, the agent can be incentivized to exert effort on a single task if and only if $\frac{2c_i}{p_i} \leq 1$. Theorem 4 shows that for multiple tasks there is a monotonicity property for the set of incentivizable tasks.

Lemma 4 (Monotonicity in tasks) *For any set of assignments $\Psi \subseteq [n]$, if there exists a proper scoring rule S such that the agent exerts effort on tasks Ψ , for any subset $\Psi' \subseteq \Psi$, there exists a proper scoring rule S' such that the agent exerts effort on tasks $i \in \Psi'$.*

Proof To incentivize effort on Ψ' , we construct S' by simulating the agent's effort on the set $\Psi \setminus \Psi'$. For any reported signal profile σ' , let $S'(\sigma', \omega) = \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) \mid \sigma_i = \sigma'_i, \forall i \in \Psi']]$ be the score that ignores the report in set $\Psi \setminus \Psi'$, and takes expected score over this set by simulating the signals assuming effort.

The proof follows by showing that exerting effort on set Ψ' is the optimal strategy for the agent with scoring rule S' . Since the new scoring rule S' only depends on reported signals in set Ψ' , the agent has no incentive to exert effort on any task outside Ψ' . For any subset $\hat{\Psi} \subseteq \Psi'$, the expected utility difference between exerting effort in sets $\hat{\Psi}$ and Ψ' given scoring rule S' is identical to the expected utility difference between exerting effort in sets $\hat{\Psi} \cup (\Psi \setminus \Psi')$ and Ψ given scoring rule S . Since exerting effort on all tasks in set Ψ is optimal for scoring rule S , exerting effort on all tasks in set Ψ' is optimal for scoring rule S' . ■

By Theorems 3 and 4, it is without loss to assume that $p_i \geq 2c_i$ for all tasks $i \in [n]$, and we will maintain this assumption throughout the paper.

There are two families of scoring rules that will arise in our analysis, *truncated separate scoring rules* and *threshold scoring rules*. Intuitively, the truncated separate scoring rules specify a scoring rule for each task, and the total score is the sum of scores on each task, truncated between 0 and the budget.

Definition 5 *A scoring rule S is a truncated separate scoring rule with budget $B > 0$ if there exists single-dimensional scoring rules S_1, \dots, S_n and shifting parameter $d \geq 0$ such that $S(\sigma, \omega) = \min \left\{ B, \max \left\{ 0, -d + \sum_{i \in [n]} S_i(\sigma_i, \omega_i) \right\} \right\}$.*

2. By Theorem 17, $\frac{2c_i}{p_i}$ is also the minimum budget required for any scoring rule proper for belief elicitation in order to incentivize the agent to exert effort on single task $\{i\}$.

Note that due to the truncation to $[0, B]$, scoring rule S may not be proper in general even if the individual single-dimensional scoring rules are proper. In later sections, we will properly design the parameter d and single-dimensional scoring rules such that the aggregated scoring rule will remain proper.

Definition 6 A scoring rule S is a threshold scoring rule if there exist a recommendation set $\Psi \subseteq [n]$ and a threshold $\eta \geq 0$ on the number of tasks for the agent to predict correctly, such that:

- the score is 0 if there exists task $i \in \Psi$ such that the reported signal is informative but wrong, i.e. $\sigma_i \neq \perp$ and $\sigma_i \neq \omega_i$;
- let $k \triangleq \#\{i \in \Psi : \sigma_i = \omega_i\}$ be the number of tasks that the agent predicts correctly. The score is 1 if the agent's correct prediction exceeds the threshold, i.e. $k \geq \eta$; and $\frac{1}{2^{\eta-k}}$ otherwise.

The threshold scoring rule in Theorem 6 is proper. In Appendix C.1, to help with the understanding, we provide an equivalent formulation of threshold scoring rules in the special case of threshold 1 such that it is also proper for eliciting the belief. Here we show that it is also equivalent to the following non-proper scoring rule with the same recommendation set Ψ and threshold η :

- the score is 1 if both (1) the number of reported informative signal exceeds the threshold, i.e., $\#\{i \in \Psi : \sigma_i \neq \perp\} \geq \eta$; and (2) any task $i \in \Psi$ such that the reported signal is informative is correct, i.e., $\sigma_i \neq \omega_i$ if $\sigma_i \neq \perp$;
- the score is 0 otherwise.

Conditioning on the agent receiving $k \leq \eta$ informative signals, his best response is to guess the rest $\eta - k$ signals, with a probability $\frac{1}{2^{\eta-k}}$ that he can guess all correctly and receive score 1. His expected utility is thus $\frac{1}{2^{\eta-k}}$, which implies this non-proper scoring rule is equivalent to the proper scoring rule in Theorem 6.

2.2. Sequential Effort Model

In the sequential effort model, we assume that the agent can sequentially exert effort on different tasks before the interaction with the seller, and the agent can make effort decisions based on the signals he has received on previous tasks. Formally, at any moment, let $\hat{\Psi}$ be the set of tasks that the agent has exerted effort, and let $\sigma_{\hat{\Psi}}$ be the set of signals on those tasks. The agent's strategy $\tau(\hat{\Psi}, \sigma_{\hat{\Psi}}) \in [n] \cup \{\emptyset\} \setminus \hat{\Psi}$ specifies a new task to exert effort on or to stop exerting more effort (represented by \emptyset) based on historical observations. The timeline of our model is as follows:

1. The principal commits to a proper scoring rule $S : \Sigma \times \Omega \rightarrow [0, 1]$.
2. The agent adopts a sequential strategy τ for exerting effort on tasks.
3. States $\omega = \{\omega_i\}_{i=1}^n$ are realized. The agent receives signals σ and pays cost $\sum_{i \in \bar{\Psi}} c_i$ where $\bar{\Psi}$ is the set of tasks that the agent has exerted effort on before stopping.
4. The agent reports σ and receives score $S(\sigma, \omega)$.

Let $\mathbf{E}^\tau[\cdot]$ be the expectation when the agent follows strategy τ for exerting effort.

Note that for the sequential effort setting, we do not require the principal to make strategy recommendations to the agent. The main reason is because the agent's optimal search problem might be computationally hard given the designed scoring rules. In this case, it would be unreasonable to prove the performance guarantee of our proposed scoring rules assuming that the agent can best respond to the mechanism. Instead, we make a weak assumption on agent's behavior, and show that for any reasonable response of the agent, the expected value of the set of tasks that the agent has exerted effort on is large enough.

Definition 7 A strategy τ is obviously dominated if there exists $\hat{\Psi}$, signal $\sigma_{\hat{\Psi}}$ and task $i \notin \hat{\Psi}$ such that $\tau(\hat{\Psi}, \sigma_{\hat{\Psi}}) = \emptyset$ and the agent increases his expected utility by exerting effort on task i compared to stopping, i.e.,

$$\mathbf{E}_{\sigma_i \sim \{i\}} \left[\mathbf{E}_{\omega \sim \sigma_i \cup \sigma_{\hat{\Psi}}} [S(\sigma, \omega)] \right] - c_i \geq \mathbf{E}_{\omega \sim \sigma_{\hat{\Psi}}} [S(\sigma, \omega)].$$

Requiring the agent's strategy to be not obviously dominated is in the same spirit of undominated strategies in [Babaioff, Lavi, and Pavlov \(2009\)](#) and advised strategies in [Cai, Thomas, and Weinberg \(2020\)](#), with the adaption to sequential environments (c.f., [Li, 2017](#)). In Section 6, we will show that the principal's payoff is approximately optimal given our proposed scoring rules if the agent's strategy is not obviously dominated.

3. Computational Hardness

In this section, we show that the design of the optimal mechanism for maximizing the principal's value is computationally hard by reduction from the NP-hard integer valued subset sum problem.

Integer valued subset sum. Given n integers z_1, \dots, z_n and a target $Z > z_i$ for all $i \in [n]$, does there exist a set $\Psi \subseteq [n]$ such that $\sum_{i \in \Psi} z_i = Z$?

Our proof idea is similar to the reduction from the subset sum to the knapsack problem. The main challenge for reduction to our problem is that, in order to prevent the agent from randomly guessing the states of the tasks, there is a specific incentive constraint that determines the set of incentivizable tasks. The incentive compatibility constraint potentially generates a much smaller value than the optimal set of tasks with total costs below the budget. To avoid this randomly guessing issue, we add additional tasks to the scoring rule design problem such that the agent's utility from making any random guess is sufficiently low, and that the optimal objective value of the principal exceeds a given value if and only if the objective Z of the subset sum problem can be achieved.

Theorem 8 *Computing the optimal mechanism in the knapsack scoring problem is NP-hard even if the valuation function is additive.*

4. Bicriteria Approximation: Inflating the Budget

In this section, we show that there exists a proper truncated separate scoring rule with a constant budget that achieves higher value for the principal than the optimal mechanism with budget 1. Specifically, we show that by inflating the budget by a constant factor, the principal is able to attain at least the optimal objective value given budget 1 with relaxed incentive constraints.

Truncated Scoring Mechanism for additive values with budget 11

 Post the truncated scoring rule on a recommendation set Ψ

- For each assignment $i \in \Psi$, let the budget-minimal scoring rule be \hat{S}_i .

Posting single dimensional scoring rules:

$$S_i(\sigma_i, \omega_i) = \frac{9}{8} \hat{S}_i(\sigma_i, \omega_i) = \begin{cases} \frac{9c_i}{8p_i} & \sigma = \perp \\ \frac{9c_i}{4p_i} \cdot \mathbf{1}[\sigma_i = \omega_i] & \text{otherwise} \end{cases}$$

- Sum over the single dimensional scores, and truncate back to $[0, 11]$:

$$S = \max \left\{ 0, \min \left\{ 11, \sum_i S_i - d \right\} \right\},$$

 where $d = -\frac{11}{2} + \frac{9}{8} \sum_{i \in \Psi} \frac{c_i}{p_i}$ is the shift on the sum.

Figure 1: Truncated Scoring Mechanism.

Recommendation set Ψ for truncated scoring mechanism

 Input: ground set G

 Output: set Ψ

 Greedily include tasks from G to Ψ , by value-cost ratio with a budget $\frac{3}{2}$ on the total cost.

Figure 2: Procedure for identifying optimal recommendation set for truncated scoring mechanism.

The approximation mechanism we design for the knapsack scoring problem uses the (approximately) optimal solution for the knapsack problem as a blackbox. Note that for general submodular valuations, computing the optimal solution for ALG-OPT is NP-hard. The following lemma shows that there exists a polynomial time algorithm to get an $e/(e-1)$ -approximation.

Lemma 9 (Sviridenko, 2004) *For submodular valuation v , there exists a polynomial time algorithm that computes a feasible solution Ψ such that $v(\Psi) \geq (1 - 1/e)\text{ALG-OPT}$.*

Theorem 10 *The truncated scoring mechanism (Figure 1) with a budget $B = 11$ guarantees value at least the optimal knapsack value (ALG-OPT). Moreover, for submodular values, there is a polynomial time algorithm for computing the recommendation set Ψ (Figure 2), which attains an $e/(e-1)$ -approximation.*

The main idea is that with multiple tasks, the sum of the scores on different tasks concentrates around its expectation. Therefore, we can take the sum of the scores and shift it such that the expected score of not exerting any effort is only one half of the budget 11. Moreover, with an inflated budget, we can ensure that the ex post shifted sum remains in the range of $[0, 11]$ with high probability, and hence the agent's incentive is almost aligned with his incentive in separate scoring rules without the truncation. This allows us to show that the designed truncated separate scoring

rule is proper, and the agent has the incentive to follow the recommendation. The detailed proof of the theorem is provided in Appendix C.3.

5. Value Approximation

In this section, we show that the better of a truncated separate scoring rule and a threshold scoring rule is a constant approximation to the optimal value of the knapsack scoring problem (IC-OPT). The idea is to divide the set of tasks into two subsets based on whether the sum of optimal individual single-dimensional scoring rule concentrates, and then design approximately optimal scoring rule for each subset separately. This approach is analogous to the core-tail decomposition adopted for multi-item auctions (Babaioff, Immorlica, Lucier, and Weinberg, 2020), while the details for proving the results are quite different.

The first case is to consider tasks such that their costs are small compared to their probabilities of revealing the state when the agent exerts effort. In this case, the budget required for incentivizing each single task is small. Thus, analogous to Theorem 10, the variance of the score for incentivizing each task separately is small and the sum of the scores concentrates well given the total budget 1. This implies that the ex post sum is close to its expectation with high probability. By truncating the sum of optimal single-dimensional scoring rules to comply with the ex post budget constraint, the incentives of the agent for exerting effort are barely affected, and we obtain a constant approximation to the knapsack solution in this case.

The second case is to consider tasks such that their costs are large compared to their probabilities of revealing the states when the agent exerts effort. Unlike the traditional knapsack problem where large costs on the tasks indicate the existence of a single task with valuation close to the optimal, in the effort incentivization problem, there still exists the hard case where in the optimal mechanism, the agent need to be incentivized to exert effort on a large number of tasks and each task only contributes to a small fraction of the optimal objective value. Moreover, since the probabilities of revealing the states are small, the expected number of tasks on which the agent receives informative signals is small and hence the sum of scores may not concentrate. Alternatively, we show that in this case, the score of the agent has to be close to the budget if he receives an informative signal on any task. Therefore, to incentivize the agent to exert effort on any task i , the total probability that the agent gets an informative signal on any task $i' \neq i$ cannot be too large because otherwise the principal will not have enough budget to incentivize task i after rewarding the agent for acquiring an informative signal on task i' . Thus, an upper bound is imposed on the sum of probabilities for the set of incentivizable tasks. A greedy algorithm (Figure 4) on the ratio of the value to the probability finds a set of tasks that can be incentivized by the threshold scoring rule (Figure 3). We show that the value of this set is a constant approximation to the value given by the optimal scoring rule.

Theorem 11 *The better of a truncated separate scoring rule and a threshold scoring rule is a 1091-approximation to the optimal value of the knapsack scoring problem (IC-OPT). Moreover, for additive values, the parameters of such mechanism can be computed in polynomial time, and for submodular values, there is a polynomial time algorithm for computing the parameters that loses an additional multiplicative factor of $e/(e-1)$ in approximation ratio.*

The proof of this theorem is deferred to Appendix C.4.

Threshold Scoring Mechanism for additive values with budget 1

 Post the threshold scoring rule on a recommendation set Ψ

- Score 1 if both (a) at least one reported signal in Ψ is informative; and (b) all task reported signals that are informative are correct.
- Score 0, otherwise.

Figure 3: Threshold Scoring Mechanism.

Recommendation set Ψ for threshold scoring mechanism

 Input: ground set G .

 For each task j in the ground set G :

- initialize by adding j into the recommendation $\Psi^j = \{j\}$;
- update the ground set G : $G^j = \{i \in G \mid 1 - \frac{2c_j}{p_j} + p_j \leq 1 - \frac{2c_i}{p_i} + p_i\}$;
- greedily include tasks from G^j by the value-probability ratio $\frac{v_i}{p_i}$ with a budget $\sum_{i \in \Psi^j} p_i \leq 1 - \frac{2c_j}{p_j} + p_j$;
- Consider set $\Psi'^j = \{j, j^*\}$, where $j^* = \operatorname{argmax}_{i \in G^j} v(i)$ is the most valuable task. Take the better of the knapsack solution and the set Ψ'^j

 Output the set with maximum value: $\Psi = \operatorname{argmax}_{\Psi^j} v(\Psi^j)$.

Figure 4: Procedure for identifying approximately optimal recommendation set.

6. Sequential Effort

In this section, we show that the value approximation results for the static effort model can be generalized to the model where the effort choice is made sequentially by applying the same family of scoring rules.

In the sequential search model, the agent makes the effort choice on tasks sequentially with the order of his choice. Our designed scoring rule is robust against the strategy the agent adopts for exerting efforts on the recommendation set as long as the strategy is not obviously dominated.

Theorem 12 *The better of a truncated separate scoring rule and a threshold scoring rule is a 561 -approximation to the optimal value of the knapsack scoring problem (IC-OPT) when the agent does not adopt obviously dominated strategies. Moreover, for additive values, the parameters of such mechanism can be computed in polynomial time, and for submodular values, there is a polynomial time algorithm for computing the parameters that loses an additional multiplicative factor of $e/(e-1)$ in approximation ratio.*

The proof of this theorem is deferred to Appendix C.5.

References

- Gagan Aggarwal and Jason D Hartline. Knapsack auctions. In *Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1083–1092, 2006.
- Moshe Babaioff, Ron Lavi, and Elan Pavlov. Single-value combinatorial auctions and algorithmic implementation in undominated strategies. *Journal of the ACM (JACM)*, 56(1):1–32, 2009.
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. In *55th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2014*, pages 21–30. IEEE Computer Society, 2014.
- Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S Matthew Weinberg. A simple and approximately optimal mechanism for an additive buyer. *Journal of the ACM (JACM)*, 67(4):1–40, 2020.
- Eric Balkanski and Jason D Hartline. Bayesian budget feasibility with posted pricing. In *Proceedings of the 25th International Conference on World Wide Web*, pages 189–203, 2016.
- Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. Budget feasible mechanism design: from prior-free to bayesian. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 449–458, 2012.
- Xiaohui Bei, Ning Chen, Nick Gravin, and Pinyan Lu. Worst-case mechanism design via bayesian analysis. *SIAM Journal on Computing*, 46(4):1428–1448, 2017.
- Mira Belenkiy, Melissa Chase, C Chris Erway, John Jannotti, Alptekin Küpçü, and Anna Lysyanskaya. Incentivizing outsourced computation. In *Proceedings of the 3rd international workshop on Economics of networked systems*, pages 85–90, 2008.
- Sergei Bernstein. Theory of probability. (*Russian*), Moscow, 1927.
- Patrick Briest, Piotr Krysta, and Berthold Vöcking. Approximation techniques for utilitarian mechanism design. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 39–48, 2005.
- Linda Cai, Clay Thomas, and S Matthew Weinberg. Implementation in advised strategies: Welfare guarantees from posted-price mechanisms when demand queries are np-hard. In *11th Innovations in Theoretical Computer Science Conference, ITCS 2020*, page 61, 2020.
- Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. Designing menus of contracts efficiently: The power of randomization. *arXiv preprint arXiv:2202.10966*, 2022.
- Yiling Chen and Bo Waggoner. Informational substitutes. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 2016.
- Yiling Chen and Fang-Yi Yu. Optimal scoring rule design. *arXiv preprint arXiv:2107.07420*, 2021.
- Changyu Dong, Yilei Wang, Amjad Aldweesh, Patrick McCorry, and Aad van Moorsel. Betrayal, distrust, and rationality: Smart counter-collusion contracts for verifiable cloud computing. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 211–227, 2017.

- Paul Dütting, Tomer Ezra, Michal Feldman, and Thomas Kesselheim. Combinatorial contracts. In *62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 815–826. IEEE, 2022.
- Rosario Gennaro, Craig Gentry, and Bryan Parno. Non-interactive verifiable computing: Outsourcing computation to untrusted workers. In *Advances in Cryptology—CRYPTO 2010: 30th Annual Cryptology Conference, Santa Barbara, CA, USA, August 15-19, 2010. Proceedings 30*, pages 465–482. Springer, 2010.
- Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Yuqing Kong. More dominantly truthful multi-task peer prediction with a finite number of tasks. In *13th Innovations in Theoretical Computer Science Conference (ITCS 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.
- Shengwu Li. Obviously strategy-proof mechanisms. *American Economic Review*, 107(11):3257–87, 2017.
- Yingkai Li, Jason D Hartline, Liren Shan, and Yifan Wu. Optimization of scoring rules. *Proceedings of the 23rd ACM Conference on Economics and Computation, to appear*, 2022.
- Eric Neyman, Georgy Noarov, and S Matthew Weinberg. Binary scoring rules that incentivize precision. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 718–733, 2021.
- Kent Osband. Optimal forecasting incentives. *Journal of Political Economy*, 97(5):1091–1112, 1989.
- Maneesha Papireddygari and Bo Waggoner. Contracts with information acquisition, via scoring rules. In *Proceedings of the 23rd ACM Conference on Economics and Computation, EC '22*, page 703–704, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450391504. doi: 10.1145/3490486.3538261. URL <https://doi.org/10.1145/3490486.3538261>.
- Mark S Pinsker. *Information and information stability of random variables and processes*. Holden-Day, 1964.
- Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 377–394, 2015.
- Aviad Rubinstein and S Matthew Weinberg. Simple mechanisms for a subadditive buyer and applications to revenue monotonicity. *ACM Transactions on Economics and Computation (TEAC)*, 6(3-4):1–25, 2018.
- Yaron Singer. Budget feasible mechanisms. In *2010 IEEE 51st Annual Symposium on foundations of computer science*, pages 765–774. IEEE, 2010.

Maxim Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Operations Research Letters*, 32(1):41–43, 2004.

Luis Zerneno. A principal-expert model and the value of menus. *unpublished paper, Massachusetts Institute of Technology*, 4, 2011.

Appendix A. Probability Tools

Lemma 13 (Hoeffding, 1963) Suppose X_1, \dots, X_n are independent random variables such that $X_i \in [a_i, b_i]$. Let $X = \sum_i X_i$. For any $\delta > 0$,

$$\Pr[X - \mathbf{E}[X] \geq \delta] \leq \exp\left(-\frac{2\delta^2}{\sum_i (b_i - a_i)^2}\right).$$

Lemma 14 (Bernstein, 1927) Suppose X_1, \dots, X_n are independent zero-mean random variables such that $|X_i| \leq M$. Let $X = \sum_i X_i$. For any $\delta > 0$,

$$\Pr[|X| \geq \delta] \leq 2 \exp\left(-\frac{\frac{1}{2}\delta^2}{\sum_i \mathbf{E}[X_i^2] + \frac{M}{3}}\right).$$

Lemma 15 (Pinsker, 1964) If P and Q are two probability distributions on a measurable space (X, Σ_X) , then for any measurable event $\mathcal{E} \in \Sigma_X$, it holds that

$$|P(\mathcal{E}) - Q(\mathcal{E})| \leq \sqrt{\frac{1}{2} \text{KL}(P\|Q)},$$

where

$$\text{KL}(P\|Q) = \int_X \left(\ln \frac{dP}{dQ} \right) dP$$

is the Kullback–Leibler divergence.

Appendix B. Properness for Belief Elicitation

The main idea of converting any scoring rule that is potentially not proper for some beliefs to a scoring rule that is proper for all beliefs is to apply the taxation principle and let the agent chooses his best option given the original scoring rule.

Definition 16 A scoring rule $S : \Delta(\Omega) \times \Omega \rightarrow [0, 1]$ is proper for belief elicitation if for any $\mu, \mu' \in \Delta(\Omega)$,

$$\mathbf{E}_{\omega \sim \mu}[S(\mu, \omega)] \geq \mathbf{E}_{\omega \sim \mu}[S(\mu', \omega)].$$

Claim 17 For any proper scoring rule S , there exists another scoring rule \hat{S} that is proper for belief elicitation such that

$$S(\sigma, \omega) = \hat{S}(\mu(\sigma), \omega)$$

for any $\sigma \in \Sigma$ and $\omega \in \Omega$ where $\mu(\sigma)$ is the posterior belief of the agent when receiving signal σ .

Proof Consider the following scoring rule for belief elicitation:

$$\hat{S}(\mu, \omega) = S(\sigma^*, \omega), \text{ where } \sigma^* \in \operatorname{argmax}_{\sigma} \mathbf{E}_{\omega \sim \mu}[S(\sigma, \omega)]. \quad (1)$$

Next, we will show that

1. \hat{S} is proper for belief elicitation. Let $\sigma^*(\mu) = \operatorname{argmax}_{\sigma} \mathbf{E}_{\omega \sim \mu}[S(\sigma, \omega)]$ be the best responding signal when the agent has to choose a signal to report. For any belief μ and μ' , we have

$$\mathbf{E}_{\omega \sim \mu}[S(\mu', \omega)] = \mathbf{E}_{\omega \sim \mu}[S(\sigma^*(\mu'), \omega)] \leq \mathbf{E}_{\omega \sim \mu}[S(\sigma^*(\mu), \omega)] = \mathbf{E}_{\omega \sim \mu}[S(\mu, \omega)]$$

which implies that \hat{S} is proper for belief elicitation.

2. \hat{S} is an extension of S , i.e. $S(\sigma, \omega) = \hat{S}(\mu(\sigma), \omega)$. This follows directly from the properness of the original scoring rule S .

■

Note that given an arbitrary scoring rule S , computing the best response strategy σ^* given his belief μ as in Equation (1) may be NP-hard. Therefore, even though such proper scoring rule exists, it might be challenging to provide its exact form in polynomial time given our designed scoring rules. Fortunately, for our purpose of incentivizing effort, we can adopt a similar solution concept in our sequential effort model (c.f., Theorem 7 and Section 6) by allowing the agent to approximately best response to the scoring rule. More specifically, given any scoring rule S for eliciting the signals, the principal can offer this original scoring rule S to the agent, ask the agent to report his belief, and let the agent choose the best possible signal he can find in polynomial time as input to the scoring rule S for computing his score based on his belief. This protocol disentangles the incentives between reporting beliefs and maximizing the expected score, and hence the agent has no incentive to misreport his true belief. Moreover, since the scoring rule is proper for all signals in the support, for any belief induced by those signals, the agent's best response is to simply report those signals truthfully. For any belief that cannot be induced by those signals, the agent can adopt any polynomial time algorithm for finding an approximately optimal solution. However, as those events happen with probability measure 0, it would not affect the agent's incentives for exerting effort in our model, and all of our results extend naturally.

Appendix C. Missing Proofs and Constructions

C.1. Alternative Formulation of Threshold Scoring Rules

Here we present an alternative formulation of the threshold scoring rule in the special case of threshold 1 given outcome space $\Omega = \{0, 1\}^n$.

Definition 18 (Li, Hartline, Shan, and Wu, 2022) Consider the n -dimensional outcome space $\Omega = \{0, 1\}^n$. Given single-dimensional scoring rules

$$S_i(\mu_i, \omega_i) = \begin{cases} 1 & \mu_i \leq 1/2 \text{ and } \omega_i = 0, \text{ or } \mu_i > 1/2 \text{ and } \omega_i = 1, \\ 0 & \text{otherwise.} \end{cases}$$

The canonical max-over-separate scoring rule S is defined as

$$S(\mu, \omega) = S_i(\mu_i, \omega_i), \text{ where } i = \operatorname{argmax}_{i} \mathbf{E}_{\omega_i \sim \mu_i}[S_i(\mu_i, \omega_i)],$$

By Li, Hartline, Shan, and Wu (2022), the canonical max-over-separate scoring rule is proper for belief elicitation. Moreover, it is easy to verify that it coincides with the threshold scoring rule with threshold 1 given any belief of the agent.

C.2. Proof of Theorem 8

Proof Given an integer valued subset sum instance with integer parameters z_1, \dots, z_n and Z , we construct a knapsack scoring problem. Let $\bar{v} = 1 + \sum_{i \in [n]} z_n$ and $\bar{c} = 1 + \max_{i \in [n]} z_i$. Let k be the minimum integer such that $2^{kn} > Z + 2kn\bar{c} + 1$. It is easy to see that the value of k is polynomial in the number of digits to represent Z and $\max_{i \in [n]} z_i$. Construct a knapsack scoring problem with $(2k + 1)n$ tasks such that if the agent exerts effort on any task i , he observes the state ω_i with probability 1. The values and costs of the tasks are defined in the following way:

- for each task $i \leq n$, let value and cost be $v_i = c_i = z_i$;
- for each task $n + 1 \leq i \leq (2k + 1)n$, let $v_i = \bar{v}$ and $c_i = \bar{c}$.

The budget of the principal is $Z + 2kn\bar{c} + 1$. Note that this instance can be easily converted to our problem with budget 1 by re-scaling the budget and the costs by the same factor. We claim that the subset sum problem is true if and only if the optimal objective value for the knapsack scoring problem is $Z + 2kn\bar{v}$.

If the optimal objective value for the knapsack scoring problem is $Z + 2kn\bar{v}$, this implies that in the optimal solution, the agent is incentivized to exert effort on all tasks $n + 1 \leq i \leq (2k + 1)n$, which has a total contribution of $2kn\bar{v}$. Thus the agent must exert effort on a subset $\Psi \subseteq [n]$ such that $\sum_{i \in \Psi} v_i = Z$. Since $v_i = z_i$ for all $i \in [n]$, Ψ is a solution for the integer valued subset sum problem.

If there exists a set of integers $\mathcal{Z} \subseteq [n]$ such that $\sum_{i \in \mathcal{Z}} z_i = Z$, consider the threshold scoring rule with recommendation set $\Psi = \mathcal{Z} \cup \{n + 1, \dots, (2k + 1)n\}$ and threshold $\eta = |\Psi|$, which scores budget $Z + 2kn\bar{c} + 1$ if the agents predicts all tasks in recommendation set Ψ correctly. It is easy to verify that the utility of the agent for exerting effort on all tasks $i \in \Psi$ is 1. The utility of the principal on recommendation set Ψ is $Z + 2kn\bar{v}$. We are going to show this threshold scoring rule is incentive compatible and optimal.

To prove this threshold scoring rule is incentive compatible, we divide agent's deviation into two cases: 1) the agent exerts effort on a small subset, so that he has to randomly guess on a large number of tasks, which reduces his utility; 2) the agent exerts effort on a large subset, which induces a high total cost.

- If the agent chooses to exert effort on a subset with size at most $|\mathcal{Z}| + kn$, he has to make random guess on at least kn tasks. The utility of the agent is at most $2^{-kn} \cdot (Z + 2kn\bar{c} + 1) < 1$, which is strictly smaller than his utility for exerting effort on all tasks $i \in \Psi$.
- If the agent chooses to exert effort on a subset with size between $|\mathcal{Z}| + kn$ and $|\mathcal{Z}| + 2kn - 1$, the cost of effort for the agent is at least $Z + kn\bar{c} \geq \frac{1}{2}(Z + 2kn\bar{c} + 1)$ since $Z \geq 1$. Moreover, the expected payment to the agent is at most $\frac{1}{2}(Z + 2kn\bar{c} + 1)$ since the agent has to make a random guess on at least one task. This implies that the agent's utility is negative given this deviating strategy.

Thus the agent's optimal choice is to exert effort on all tasks $i \in \Psi$.

Finally, we show that the optimal utility of the principal cannot exceed $Z + 2kn\bar{v}$. Note that for the principal to obtain utility at least $Z + 2kn\bar{v}$, the agent must be incentivized to exert effort on all tasks $i \in \{n + 1, \dots, (2k + 1)n\}$ since the sum of value in $[n]$ is strictly below the value of any task $i \in \{n + 1, \dots, (2k + 1)n\}$. Moreover, the total cost of the agent for exerting effort given the

optimal scoring rule is strictly less than $Z + 2kn\bar{c} + 1$ since the agent can obtain strictly positive utility by exerting no effort and randomly guessing. Since the costs are integer valued, the total cost is at most $Z + 2kn\bar{c}$. As the total cost for exerting effort on tasks $i \in \{n + 1, \dots, (2k + 1)n\}$ is $2kn\bar{c}$, the cost of the agent on tasks within subset $[n]$ is at most Z . Since the value coincides with the cost in this case, the value of the principal from incentivizing the agent to exert effort on tasks within $[n]$ is at most Z . Therefore, the optimal utility of the principal is $Z + 2kn\bar{v}$. \blacksquare

C.3. Proof of Theorem 10

We show that the mechanism in Figure 1 is incentive compatible, by first showing that scoring rule S is proper, and then showing that Ψ is the agent's best effort choice.

Proper. For each task i , conditional on receiving signal $\sigma_i \neq \perp$, the score $S_i(\sigma_i, \omega_i)$ first order stochastically dominates $S_i(\sigma'_i, \omega_i)$ for any σ'_i . Thus, the agent has incentives to truthfully report the signal σ_i if $\sigma_i \neq \perp$.

We then show that the agent has no incentives to misreport on tasks with uninformative signal $\sigma_i = \perp$ by contradiction. Suppose that the agent has incentives to misreport given signal \perp on some tasks. We partition the tasks into three sets. Let Z_0 be the set of tasks i such that $\sigma_i \neq \perp$, Z_1 be the set of tasks i such that $\sigma_i = \perp$ and where the agent truthfully reports the signal, and Z_2 be the set of tasks i such that $\sigma_i = \perp$ and where the agent misreports the signal. First note that if $\sum_{i \in Z_0} 2 \cdot 9c_i/8p_i \geq 11$, then by truthful reporting the signals the agent can secure a deterministic score 11, which is the maximum possible score. Hence the agent has no incentive to misreport in this case.

Next we focus on the case when $\sum_{i \in Z_0} 9c_i/8p_i < 11/2$. Let η_i be a Bernoulli random variable with probability $1/2$ drawn independently for each task $i \in Z_2$. We use η_i to indicate whether the agent guesses correctly on the task $i \in Z_2$. Let

$$s = \sum_{i \in Z_0} \frac{9c_i}{8p_i} + \sum_{i \in Z_2} \frac{9c_i}{4p_i} \left(\eta_i - \frac{1}{2} \right) + \frac{11}{2}.$$

Note that s is the random variable corresponding to the score without truncation by the interval $[0, 11]$. Consider an alternative setting where the score is truncated by the interval $[\sum_{i \in Z_0} 9c_i/4p_i, 11]$. Since the distribution of s is symmetric with respect to the mean $(\sum_{i \in Z_0} 9c_i/4p_i + 11)/2$, the score distribution under the truncation by $[\sum_{i \in Z_0} 9c_i/4p_i, 11]$ is also symmetric with respect to the mean. Thus, the utility of the agent for misreporting in this alternative setting is exactly the same as the utility for truthful reporting, $(\sum_{i \in Z_0} 9c_i/4p_i + 11)/2$. Since $\sum_{i \in Z_0} 9c_i/4p_i > 0$, the utility of the agent for misreporting with truncation by $[0, 11]$ is strictly less than the utility for misreporting with truncation by $[\sum_{i \in Z_0} 9c_i/4p_i, 11]$. Therefore, the agent will not have an incentive to misreport in the original setting when the lower truncation is 0.

Effort Set. We prove that the agent's optimal choice is to exert effort in tasks Ψ . First note that we set the score to be zero for $i \notin \Psi$. This immediately implies that the agent will not exert effort on task $i \notin \Psi$. Fix the agent's effort choice in Ψ . Suppose there exists a task $\hat{i} \in \Psi$ such that the agent's effort on task \hat{i} is 0. Let $\hat{\mathcal{E}}_{\hat{i}}$ be the event that $-d + \sum_{i \in \Psi \setminus \{\hat{i}\}} S_i(\sigma_i, \omega_i) \in [0, 11 - 9c_i/8p_i]$.

Let $\hat{Z} \subseteq \Psi$ be the set on which the agent exerts effort. Therefore,

$$\begin{aligned}
 \Pr[\hat{\mathcal{E}}_i] &= 1 - \Pr\left[-d + \sum_{i \in \Psi \setminus \{\hat{i}\}} S_i(\sigma_i, \omega_i) > 11 - \frac{9c_i}{8p_i}\right] \\
 &= 1 - \Pr\left[\sum_{i \in \hat{Z}} \mathbf{1}[\sigma_i \neq \perp] \cdot \frac{9c_i}{8p_i} > \frac{11}{2} - \frac{9c_i}{4p_i}\right] \\
 &\geq 1 - \exp\left(-\frac{\frac{1}{2}\left(\frac{11}{2} - \frac{9c_i}{4p_i} - \sum_{i \in \hat{Z}} \frac{9c_i}{8}\right)^2}{\frac{1}{4} \sum_{i \in \hat{Z}} \frac{1}{p_i} \cdot \frac{9c_i^2}{4} + \frac{1}{6} \max_{i \in \hat{Z}} \frac{9c_i}{4p_i}}\right) \\
 &\geq 1 - \exp\left(-\frac{\left(11 - \frac{45}{8}\right)^2}{6 \cdot \frac{9}{8} + \frac{3}{2}}\right) \geq \frac{8}{9},
 \end{aligned}$$

where the first inequality holds by applying Bernstein's inequality (Theorem 14). The second inequality holds since (1) $\sum_{i \in \hat{Z}} \frac{9c_i}{4} = \frac{9}{4} \sum_{i \in \hat{Z}} c_i \leq \frac{27}{8}$; (2) $\sum_{i \in \hat{Z}} \sum_{i \in \hat{Z}} \frac{1}{p_i} \cdot \frac{9c_i^2}{4} \leq \sum_{i \in \hat{Z}} 2(9/8)^2 c_i \leq 3(9/8)^2$; and (3) $\max_{i \in \hat{Z}} 9c_i/4p_i \leq 9/8$. Hence, by exerting effort on task \hat{i} , the score of the agent increases by at least $\Pr[\hat{\mathcal{E}}_i] \cdot 9c_i/8 \geq c_i$, which provides a contradiction.

For submodular values, we lose a $e/e - 1$ factor in the value approximation ratio by computing the recommendation set Ψ in polynomial time (Theorem 9). Without computation constraints, we have a scoring rule that achieves the theoretical bound.

C.4. Proof of Theorem 11

We first show an upper bound on the sum of state revelation probabilities for each set of incentivizable tasks when the ratio of the cost to the probability for any task in this set is large.

Lemma 19 *For any set $\Psi \subseteq [n]$ such that $p_i \leq \frac{1}{4}$ and $\frac{2c_i}{p_i} \geq \frac{15}{16}$ for all tasks $i \in \Psi$, if the set Ψ can be incentivized by a proper scoring rule with budget 1, there exists a budget-pivotal task $i^* = \operatorname{argmin}_{i \in \Psi} \frac{16}{3} \left(1 - \frac{2c_i}{p_i}\right) + p_i$, such that the budget over total revealing probabilities is determined by i^* :*

$$\sum_{i \in \Psi} p_i \leq \frac{16}{3} \left(1 - \frac{2c_{i^*}}{p_{i^*}}\right) + p_{i^*}.$$

We first prove the theorem for additive valuations, and then at the end we introduce the details for generalizing our techniques to submodular valuations. Recall that for any task i , we have $p_i \geq 2c_i$ since otherwise that task cannot be incentivized by the principal. Thus, we divide the tasks into two sets X, Y based on the ratio $p/2c_i$ as follows

$$X = \left\{i : \frac{p_i}{2c_i} > 11\right\}; \quad Y = \left\{i : 1 \leq \frac{p_i}{2c_i} \leq 11\right\}.$$

By Theorem 10, there is a truncated separate scoring rule with budget 1 that is an 11-approximation on the set X since this case can be viewed the same as the one in Theorem 10 by scaling the score and the costs by the same constant factor 11.

We divide the set Y into three subsets.

$$Y_1 = \left\{ i : p_i \geq \frac{1}{4}, 1 \leq \frac{p_i}{2c_i} \leq \frac{16}{15} \right\}; Y_2 = \left\{ i : p_i < \frac{1}{4}, 1 \leq \frac{p_i}{2c_i} \leq \frac{16}{15} \right\}; Y_3 = \left\{ i : \frac{16}{15} < \frac{p_i}{2c_i} \leq 11 \right\}.$$

Intuitively, set Y_1 corresponds to the case that the costs of effort are large, and it is sufficient to only incentivize one task with highest value in this set. Both set Y_2 and Y_3 corresponds to the situation where the probabilities of revealing the states are small compared to the costs, and hence the concentration technique cannot be applied. In both cases, we utilize Theorem 19 to bound the sum of probabilities for any set of incentivizable tasks, and hence showing that the set of tasks we identified by our polynomial time algorithm is approximately optimal.

Case 1: $Y_1 = \left\{ i : p_i \geq \frac{1}{4}, 1 \leq \frac{p_i}{2c_i} \leq \frac{16}{15} \right\}$. In this case, $c_i \geq \frac{15p_i}{32} \geq \frac{15}{128}$. Therefore, at most 8 tasks in Y_1 can be incentivized simultaneously in the optimal mechanism. By choosing the task in Y_1 with highest value, the principal attains an 8-approximation by only incentivizing that task.

Case 2: $Y_2 = \left\{ i : p_i < \frac{1}{4}, 1 \leq \frac{p_i}{2c_i} \leq \frac{16}{15} \right\}$. We use the threshold mechanism in Figure 3, with a recommendation set generated by running Figure 4 on set Y_2 .

We prove it is a $\frac{32}{3}$ -approximation by showing: (1) the threshold scoring mechanism is incentive compatible (i.e. the agent's best response is to exert effort on all tasks in the recommendation set); and (2) the total value in the recommendation set Ψ is a 16-approximation of the optimal solution.

- (1) The threshold scoring mechanism is incentive compatible. Specifically, we show that the set Ψ_j can be incentivized for any task $j \in Y_2$. For any task $j \in Y_2$, and any $i' \neq j, i' \in \Psi_j$, according to two constraints used in the construction of Ψ_j , we have

$$\sum_{i \in \Psi_j \setminus \{i'\}} p_i = \sum_{i \in \Psi_j \setminus \{j\}} p_i - p_{i'} + p_j \leq \left(1 - \frac{2c_{i'}}{p_{i'}}\right).$$

Given the threshold scoring rule with threshold $\eta = 1$ on effort set Ψ_j , the expected score increase of exerting effort on task i' is at least the probability of receiving no informative signal on tasks in $\Psi_j \setminus \{i'\}$ times the conditional score increase for exerting effort. By the union bound, we have the probability of receiving no informative signal on tasks in $\Psi_j \setminus \{i'\}$ is at least $\prod_{i \in \Psi_j \setminus \{i'\}} (1 - p_i) \geq 1 - \sum_{i \in \Psi_j \setminus \{i'\}} p_i$. Conditioned on this event, the expected score increase for exerting effort on i' is $p_{i'} + p_{i'}/2 - 1/2 = p_{i'}/2$. Thus, we have the expected score increase of exerting effort on task i' is at least

$$\left(1 - \sum_{i \in \Psi_j \setminus \{i'\}} p_i\right) \cdot \frac{p_{i'}}{2} \geq c_{i'}.$$

Therefore, for all searches $j \in Y_2$, a threshold scoring rule with threshold 1 and recommendation set Ψ_j is incentive compatible.

- (2) The total value in the recommendation set Ψ is a 16-approximation of the optimal solution. By Theorem 19, for any set $\Psi' \subseteq Y_2$ that can be incentivized, and any $i^* \in \Psi'$, we have

$$\sum_{i \in \Psi' \setminus \{i^*\}} p_i \leq \frac{16}{3} \left(1 - \frac{2c_{i^*}}{p_{i^*}}\right).$$

Let Ψ^* be the optimal effort set in the knapsack scoring problem when the set of available tasks is Y_2 . Let $\hat{i} = \operatorname{argmin}_{i \in \Psi^*} \left(1 - \frac{2c_i}{p_i} + p_i\right)$ be the budget-pivotal task. This can be interpreted as a budget over the total probabilities in the optimal set Ψ^* :

$$\sum_{i \in \Psi^*} p_i \leq \frac{16}{3} \left(1 - \frac{2c_{\hat{i}}}{p_{\hat{i}}}\right) + p_{\hat{i}} \leq \frac{16}{3} \left(1 - \frac{2c_{\hat{i}}}{p_{\hat{i}}} + p_{\hat{i}}\right).$$

Suppose we are given an optimal set Ψ^* . Divide it into two sets based on the probability.

$$\Psi_1^* = \left\{i \in \Psi^* \setminus \{\hat{i}\} : p_i > \left(1 - \frac{2c_{\hat{i}}}{p_{\hat{i}}}\right)\right\}; \quad \Psi_2^* = \left\{i \in \Psi^* \setminus \{\hat{i}\} : p_i \leq \left(1 - \frac{2c_{\hat{i}}}{p_{\hat{i}}}\right)\right\}.$$

For the set Ψ_1^* , by Lemma 19, there are at most $16/3$ tasks in Ψ_1^* . By picking the most valuable task among Ψ^* , the set Ψ_1^j achieve a $16/3$ -approximation to the value of Ψ_1^* . For the set Ψ_2^* , we take the knapsack solution with a budget reduced by $\frac{16}{3}$ factor. By enumerating over the budget-pivotal task \hat{i} , the recommendation set in Figure 3 provides a $32/3$ -approximation to the value of Ψ_2^* .

Combining the above two cases, we have

$$\left(\frac{16}{3} + \frac{32}{3}\right) v(\Psi) \geq v(\Psi_1^*) + v(\Psi_2^*) = v(\Psi^*),$$

which implies the recommendation set Ψ is a 16-approximation to the value of Ψ^* .

- Case 3: $Y_3 = \left\{i : \frac{16}{15} < \frac{p_i}{2c_i} \leq 11\right\}$. In this case, for any set $\Psi \subseteq Y_3$ that can be incentivized, and any $i^* \in \Psi$, we have

$$\sum_{i \in \Psi \setminus \{i^*\}} p_i \leq \sum_{i \in \Psi \setminus \{i^*\}} 22c_i \leq 22 \leq 352 \left(1 - \frac{2c_{i^*}}{p_{i^*}}\right)$$

where the last inequality holds since $\frac{2c_{i^*}}{p_{i^*}} \leq \frac{15}{16}$. By the same argument as case 2, the threshold mechanism is a 1056-approximation to the optimal in the knapsack scoring problem when the set of available tasks is Y_3 .

Combining all cases, for additive valuations, the maximum between truncated separate scoring rule and threshold scoring rule is a 1091-approximation to the optimal value IC-OPT, and the parameters can be computed in polynomial time. Finally, for submodular valuation, the only difference

is that the greedy solution we adopted for finding the set of incentivizable tasks loses an additional approximation factor of $e/(e-1)$ in valuations (Sviridenko, 2004). Note that this additional factor can be save if we don't require computational efficiency and brute force search for the optimal set that can be incentivized given our proposed scoring rule.

Proof [Theorem 19] We first define several useful notations. We define \mathcal{E} to be the event that the agent receives no informative signal on all tasks in Ψ . Let $q_0 = \Pr[\mathcal{E}] = \prod_{j \in \Psi} (1 - p_j)$ be the probability that event \mathcal{E} happens. Let $s_0 = \mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \mathcal{E}]$ be the expected score of the agent when he receives no informative signal. We also define \mathcal{E}_i to be the event that the agent receives no informative signal on all tasks in $\Psi \setminus \{i\}$. let $q_i = \Pr[\mathcal{E}_i] = \prod_{j \in \Psi \setminus \{i\}} (1 - p_j)$ be the probability that the event \mathcal{E}_i happens. Let $s_i = \mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \mathcal{E}_i, \sigma_i \neq \perp]$ be the expected score of the agent when he only receives an informative signal on task i .

Next we divide the analysis into two cases: (1) $q_0 \geq 1/2$; and (2) $q_0 < 1/2$.

Case 1: $q_0 \geq 1/2$. In this case, we first show that the expected score for no informative signal s_0 can not be less than $1/4$. Suppose $s_0 < 1/4$, then we show that the incentive constraint for exerting effort on any task i is violated. The utility increase of the agent for exerting effort on task i is

$$\begin{aligned} & \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] - \mathbf{E}_{\sigma \sim \Psi \setminus \{i\}}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] \\ &= p_i (\mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \sigma_i \neq \perp]] - \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \sigma_i = \perp]]) \end{aligned}$$

Then, we bound the expected score increase for receiving an informative signal on task i . Conditioned on event \mathcal{E}_i , the expected score difference is $s_i - s_0$. Since the scoring rule is proper, we have $s_0 \geq s_i/2$, which implies $s_i - s_0 \leq s_0 \leq 1/4$. Conditioned on the complement event $\bar{\mathcal{E}}_i$, by the properness of scoring rule, the expected score difference is at most $1/2$. Thus, the utility increase for exerting effort on task i is at most

$$\mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] - \mathbf{E}_{\sigma \sim \Psi \setminus \{i\}}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] \leq p_i \left(q_i (s_i - s_0) + \frac{1}{2} (1 - q_i) \right) < \frac{3p_i}{8} < c_i,$$

which violates the incentive constraint for exerting effort on task i .

Therefore, we have $s_0 \geq 1/4$. We now lower bound the expected score s_i for receiving only one informative signal on task i . For any task i , the incentive constraint implies that

$$\begin{aligned} c_i &\leq \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] - \mathbf{E}_{\Psi \setminus \{i\}}[\mathbf{E}_{\sigma}[S(\sigma, \omega)]] \\ &= p_i (\mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \sigma_i \neq \perp]] - \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \sigma_i = \perp]]) \\ &\leq p_i \left(q_i (s_i - s_0) + \frac{1}{2} (1 - q_i) \right). \end{aligned}$$

Since $q_i \geq q_0 \geq 1/2$ and $c_i/p_i \geq 15/32$, this further implies that

$$s_i \geq s_0 + \frac{\frac{c_i}{p_i} - \frac{1}{2}(1 - q_i)}{q_i} \geq \frac{11}{16}.$$

Consider any fixed task $i^* \in \Psi$. Let $\hat{s} = \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega) | \sigma_{i^*} = \perp, \bar{\mathcal{E}}_{i^*}]]$ be the expected score of the agent when he has no signal on task i^* , and at least one informative

signal on tasks in $\Psi \setminus \{i^*\}$. Since the scoring rule is proper, $\hat{s} \geq \min_i s_i \geq 11/16$. The incentive constraint on task i^* implies that

$$\begin{aligned} c_{i^*} &\leq p_{i^*} (\mathbf{E}_{\sigma \sim \Psi} [\mathbf{E}_{\omega \sim \sigma} [S(\sigma, \omega) \mid \sigma_{i^*} \neq \perp]] - \mathbf{E}_{\sigma \sim \Psi} [\mathbf{E}_{\omega \sim \sigma} [S(\sigma, \omega) \mid \sigma_{i^*} = \perp]]) \\ &\leq p_{i^*} \left(\frac{q_{i^*}}{2} + (1 - q_{i^*})(1 - \hat{s}) \right), \end{aligned}$$

where the last inequality is due to the expected score difference conditioned on $\bar{\mathcal{E}}_i$ is at most $1 - \hat{s}$. Hence, we have

$$q_{i^*} \geq 1 - \frac{8}{3} \left(1 - \frac{2c_{i^*}}{p_{i^*}} \right).$$

Note that the probability that the agent receives at least one informative signal in $\Psi \setminus \{i^*\}$ is at least the sum of probability that the agent receives an informative signal on task i and zero informative signal on tasks in $\Psi \setminus \{i^*, i\}$. Note that the probability of the latter event is at least $q_0 \geq 1/2$. Thus, it holds that

$$1 - q_{i^*} \geq \frac{1}{2} \sum_{i \in \Psi \setminus \{i^*\}} p_i.$$

By combining the two inequalities above, we have

$$\sum_{i \in \Psi \setminus \{i^*\}} p_i \leq 2(1 - q_{i^*}) \leq \frac{16}{3} \left(1 - \frac{2c_{i^*}}{p_{i^*}} \right).$$

Case 2: Suppose $q_0 < 1/2$. Consider any fixed task $i^* \in \Psi$. In this case, we first show that there exists a subset $\bar{\Psi} \subseteq \Psi$ which satisfies the following three properties: (1) $i^* \in \bar{\Psi}$; (2) $\bar{\Psi}$ can be incentivized by a proper scoring rule; and (3) the probability of no informative signal in $\bar{\Psi} \setminus \{i^*\}$ is between $[1/2, 2/3)$. By case 1, this subset $\bar{\Psi}$ cannot be incentivized, which is a contradiction.

To find such a subset, we remove tasks in $\Psi \setminus \{i^*\}$ from Ψ one by one randomly. Since $p_{i^*} \leq 1/4$ and $q_0 < 1/2$, we have $q_{i^*} = q_0 / (1 - p_{i^*}) < 2/3$. If $q_{i^*} \in [1/2, 2/3)$, then $\bar{\Psi}$ satisfies three properties. We use Ψ' to denote the subset in this deletion process. Let q'_{i^*} be the probability of no informative signal in $\Psi' \setminus \{i^*\}$. If $q_{i^*} < 1/2$, then we have q'_{i^*} increases from q_{i^*} to 1 during this process. If there is no $q'_{i^*} \in [1/2, 2/3)$ in this process, then there exists a task $i \in \Psi$ with $p_i > 1/4$, which contradicts the assumption. Let $\bar{\Psi}$ be the subset with probability $\bar{q}_{i^*} \in [1/2, 2/3)$ during this process. It is easy to see that $\bar{\Psi}$ satisfies other two properties.

However, by union bound,

$$\sum_{i \in \bar{\Psi} \setminus \{i^*\}} p_i \geq 1 - \bar{q}_{i^*} > \frac{1}{3} \geq \frac{16}{3} \left(1 - \frac{2c_{i^*}}{p_{i^*}} \right),$$

which contradicts the assumption that $\bar{\Psi}$ can be incentivized according to the case 1. ■

Recommendation set Ψ for threshold mechanism, with sequentially responding agent
Input: ground set G
Output: set Ψ
Greedy add tasks from G to Ψ , by value-probability ratio $\frac{v_i}{p_i}$ with a budget 0.55 on the total probabilities $\sum_{i \in \Psi} p_i$ of knowing the truth.

Figure 5: Procedure for identifying approximately optimal recommendation set with sequentially responding agent.

C.5. Proof of Theorem 12

Again, we first prove the theorem for additive valuations. Similarly as Theorem 11, we divide the tasks into two sets X, Y based on the ratio $p_i/2c_i$ as follows

$$X = \left\{ i : \frac{p_i}{2c_i} > 11 \right\}; \quad Y = \left\{ i : 1 \leq \frac{p_i}{2c_i} \leq 11 \right\}.$$

On set X , the mechanism in Figure 1 with budget 1 achieves a $\frac{99}{8}$ -approximation. Let the last assignment completed be i . By the same proof of Theorem 10, for any task $i \in \Psi$, the probability that the scoring rule runs out of budget before the agent exerting effort on task i can be bounded by $\frac{8}{9}$. Hence, when adopting strategies that are not obviously dominated, with ex ante probability at least $\frac{8}{9}$, the agent will stop after finishing all the tasks in the recommendation set. The same mechanism loses another $\frac{8}{9}$ factor in the approximation ratio.

On set Y , we divide the tasks into two sets by the probability p_i of knowing the truth.

$$Y_1 = \{i : p_i \geq 0.1\}; \quad Y_2 = \{i : p_i < 0.1\}.$$

On set Y_1 , it is sufficient to pick the highest-value task and post the threshold scoring rule. By the probability-cost ratio $\frac{p_i}{c_i} \leq 22$, each task has $c_i \geq \frac{1}{220}$. At most 440 tasks can be incentivized in Y_1 . Hence a 440-approximation on Y_1 .

On set Y_2 , we use the scoring mechanism in Figure 5. We show this mechanism achieves a 109-approximation, by showing when the adopted strategy is not obviously dominated: (1) with probability at least 0.45, the agent completes all the tasks in the recommendation set; and (2) the total value in the set is a 109-approximation.

- The agent completes tasks in recommendation set with probability at least 0.45. By union bound, the probability that the agent gets any informative signal is $1 - \prod_i (1 - p_i) \leq \sum_i p_i \leq 0.55$. For any order of completing the task, the agent gets no informative signal with probability at least 0.45. The marginal gain of doing one more task is always positive, so the agent will finish the recommendation set with probability at least 0.45.
- The total value in the set is a 49-approximation to the optimal. All tasks in Y_2 has $p_i < 0.1$, so by setting the budget at 0.55, the total probabilities in Ψ is at least the optimal knapsack value with budget 0.45 on total probabilities. Since the probability-cost ratio $\frac{p_i}{c_i} \leq 22$, there is a budget on the total probabilities in any set that can be incentivized: $\sum_i p_i \leq 22$. Hence a 49-approximation.

Combining the claims above, the better of the truncated scoring mechanism and the threshold scoring mechanism achieves a 561-approximation when the agent is responding sequentially.

Appendix D. General Information Structure

In this section, we consider the problem of incentivizing effort with general information structures and illustrate the intrinsic challenges for generalizing our results to general information structures. Here, when the agent exerts effort, instead of assuming that he observes the true state ω_i with probability p_i as in previous sections, the agent receives a signal $\sigma_i \in \Sigma$ given by a signal structure that induces a distribution f_i over posterior $\mu_i \in \Delta(\Omega)$. We show that the optimal value of the knapsack scoring problem can differ a lot under two different information structures even if the optimal scoring rules for the single task problems are the same given those two information structures. Therefore, new ideas for designing approximately optimal scoring rules are required for general information structures.

First, the following lemma characterizes whether a single task can be incentivized by an incentive compatible mechanism under general information structure environments.

Lemma 20 (Li et al., 2022) *For the knapsack scoring problem with general information structures, the agent can be incentivized to exert effort on a single task $\Psi = \{i\}$ with budget 1 if and only if*

$$\mathbf{E}_{\mu_i \sim f_i}[|\mu_i - D|] \geq c_i,$$

where $|\mu_i - D|$ is the difference of the mean between the posterior and the prior.

When there are multiple tasks, a crucial statistic that affects the set of the incentivizable tasks is the expected KL-divergence between the prior and the posterior. Specifically, let

$$\Lambda_i \triangleq \mathbf{E}_{\mu_i \sim f_i}[\text{KL}(D \parallel \mu_i)]$$

where $\text{KL}(D \parallel \mu_i) = \sum_{\omega \in \Omega} D(\omega) \cdot \ln \frac{D(\omega)}{\mu_i(\omega)}$ is the KL-divergence between the prior D and the posterior μ_i . This distance measures how easy for the agent to mimic the signal distributions without exerting effort. The following lemma provides an upper bound on the set of incentivizable tasks given asymmetric and general information structures.

Lemma 21 *For the knapsack scoring problem with general information structures, for any set Ψ^* such that there exists an incentive compatible mechanism where the agent's optimal effort choice is Ψ^* , we have*

$$\sum_{i \in \Psi^*} c_i \leq \sqrt{\frac{1}{2} \sum_{i \in \Psi^*} \Lambda_i}.$$

Proof Note that given any proper scoring rule S , one feasible choice of the agent is to exert no effort, simulate the posterior distribution on set Ψ^* , and report the simulated posterior to the principal. Let P be the distribution over the profile of reports, and states for all tasks in Ψ^* given the simulations on Ψ^* . Let Q be such distribution when the agent exerts effort on all tasks in Ψ^* and get the true informative signals. It is easy to verify that the KL-divergence between P and Q is $\sum_{i \in \Psi^*} \Lambda_i$. Let

\mathcal{E} be the event such that the profile of reports and states does not coincide given the true posterior generating process and the simulated reports. Then we have

$$\begin{aligned} \mathbf{E}_{\sigma \sim \Psi}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] - \mathbf{E}_{\sigma \sim \emptyset}[\mathbf{E}_{\omega \sim \sigma}[S(\sigma, \omega)]] &\leq \mathbf{E}_Q[S(\sigma, \omega)] - \mathbf{E}_P[S(\sigma, \omega)] \\ &\leq |\Pr_P[\mathcal{E}] - \Pr_Q[\mathcal{E}]| \leq \sqrt{\frac{1}{2} \text{KL}(P \| Q)} = \sqrt{\frac{1}{2} \sum_{i \in \Psi^*} \Lambda_i} \end{aligned}$$

where the second inequality holds since the payment of the principal is at most 1, and the third inequality holds by Pinsker's inequality (Theorem 15). \blacksquare

Next we show that given two different information structures such that the design of the optimal scoring rule for both cases are the same in the single task problem, the set of incentivizable tasks may differ a lot when there are multiple tasks.

Specifically, consider the symmetric environment with identical information structures and costs c for all tasks, Theorem 21 implies that $|\Psi^*| \leq \frac{\Lambda}{2c^2}$. Fixing $p > 0$, consider the following two information structures when the agent exerts effort on any single task:

- the agent receives an informative signal $\sigma = \omega$ with probability p , and receives an uninformative signal $\sigma = \perp$ regardless of the realized state with probability $1 - p$;
- the agent receives an informative signal that induces posterior $\mu = \frac{1+p}{2}$ and $\frac{1-p}{2}$ with probability $\frac{1}{2}$ each.

Given both information structures above, in the single task problem, by Theorem 20, we know that the agent can be incentivized to exert effort on the single task if and only if the cost of effort is at most $p/2$.

For the multi-task problem, suppose that the cost of effort on a single task is $c = \frac{p}{4}$. Given the first information structure, it is easy to show that the optimal scoring rule can incentivize the agent to exert effort on $O(\frac{1}{c})$ tasks. By Theorem 11, the agent can be incentivized to exert effort on $O(\frac{1}{c})$ tasks by the threshold scoring rule. In contrast, given the second information structure, we have that $\Lambda = O(p^2)$ and hence by Theorem 21, the size of the incentivizable tasks is at most $\frac{\Lambda}{2c^2} = O(1)$. The gap on the size of the incentivizable tasks between two different information structures are unbounded when p and c are sufficiently small.

The above observation indicates that the design of the (approximately) optimal scoring rules depends on the fine details of different information structures even if they have the same performance on the single task problem. Thus it is unlikely to directly generalize our results for the special case to general information structures, or derive a unified approach for reducing the multi-task knapsack scoring problems to single-task ones. It is an interesting open question to identify tight upper bounds of the optimal solution for the knapsack scoring problem with general information structures, and design approximately optimal scoring rules to approximate the upper bound.