

Optimal Common Contract with Heterogeneous Agents*

Shenke Xiao,¹ Zihe Wang,² Mengjing Chen,¹ Pingzhong Tang,¹ Xiwang Yang³

¹Tsinghua University, ²Shanghai University of Finance and Economics, ³ByteDance
 xsk15@mails.tsinghua.edu.cn, wang.zihe@mail.shufe.edu.cn, ccchmj@qq.com,
 kenshinping@gmail.com, yangxiwang@bytedance.com

Abstract

We consider the principal-agent problem with heterogeneous agents. Previous works assume that the principal signs independent incentive contracts with every agent to make them invest more efforts on the tasks. However, in many circumstances, these contracts need to be identical for the sake of fairness. We investigate the optimal common contract problem. To our knowledge, this is the first attempt to consider this natural and important generalization. We first show this problem is NP-complete. Then we provide a dynamic programming algorithm to compute the optimal contract in $O(n^2m)$ time, where n, m are the number of agents and actions, under the assumption that the agents' cost functions obey increasing difference property. At last, we generalize the setting such that each agent can choose to directly produce a reward in $[0, 1]$. We provide an $O(\log n)$ -approximate algorithm for this generalization.

Introduction

Principal-agent theory is a subfield of mechanism design theory. The principal hires an agent to accomplish a task. The agent is able to take actions on behalf of the principal. Agent's different actions lead to different rewards the principal receives. Moral hazard occurs when the agent acts in his own interest which may be in conflict with the principal's interest. Therefore the principal designs an incentive contract with the agent to maximize the principal's utility subject to the agent's utility being maximized. The contract is a transfer function from the principal to the agent which could depend on the outcome which is affected by the agent's action.

Many economic interactions fit in the principal-agent model. For example, a firm (principal) hires a salesman (agent) to sell products. The salesman invests effort on selling products. More efforts he invests, more products will be sold. To incentivize salesman invest more efforts, the firm can set a bonus depending on the amount of products a salesman has sold. A salesman wants to maximize his utility which is

defined to be his bonus minus his efforts. The firm's utility is the revenue generated from selling products minus the bonus paid to salesmen. The central question in this research field asks: What is the principal's optimal contract?

Due to the wide application, principal-agent model has been extensively studied (Holmstrom and Milgrom 1991; 1987; Chen et al. 2019; Grossman and Hart 1992; Bolton, Dewatripont, and others 2005). The agent takes a hidden action like effort which cannot be observed by the principal directly. The principal can only observe the outcome of this action and the contract is designed to depend on the outcome only. Most works focus on the problem with one principal and one agent (Armstrong and Vickers 2010; Kleinberg and Kleinberg 2018). When the agent takes different actions, there is a different distribution over principal's reward. Given the distribution information, the optimal contract can be computed efficiently through linear programs.

In this paper, we consider the problem when there is one principal and multiple heterogeneous agents. These agents could be good at different tasks and we do not assume any relationship between the cost for different tasks among different agents. For sake of the fairness, we do not allow the principal design personal contracts for different agents. Instead, the principal has to design a common contract that applies to every agents. We assume the mapping from the action played to the outcome is deterministic. So the principal knows every agent's action by observing her outcome. The difficulty in our model stems from the multiple agents. Since the principal can only use a common contract, he needs to balance the incentivization for every agent.

Our Contribution

Our contribution can be summarized as follows.

1. We first show that the optimal contract problem with heterogeneous agents is strongly NP-complete.
2. We then propose an $O(n^2m)$ dynamic programming algorithm, where n is the number of agents and m is the number of actions, to compute an optimal contract under the assumption that the agents' costs obey *increasing differences*.
3. Next, we generalize the discrete-action setting such that each agent can choose to directly produce a reward in $[0, 1]$. We show that this generalization is harder than the

*This work is supported by Science and Technology Innovation 2030 –“New Generation Artificial Intelligence” Major Project No. (2018AAA0100904), Beijing Academy of Artificial Intelligence (BAAI), the Shanghai Sailing Program (Grant No. 18YF1407900), and the National NSFC (Grant No. 61806121).
 Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

original discrete-action version, and provide an $O(\log n)$ -approximate algorithm for this generalization.

Other Related Works

Other works also consider multiple agents but in different angles (Babaioff, Feldman, and Nisan 2006a; 2006b; Babaioff et al. 2012; Bernstein and Winter 2012; Babaioff and Winter 2014; Winter 2004). They assume the union of agents' actions together determines the outcome. The contract is personalized and specifies the payment in every possible outcome. Therefore the payment to an agent depends on both his action and other agents' actions. In contrast, in our paper, the payment to an agent only depends on his own action. In their setting, some works consider that each agent only has a binary action space (Babaioff, Feldman, and Nisan 2006a; Babaioff et al. 2012; Winter 2004), and some works consider the tradeoff between simplicity of the contract and the performance of it (Babaioff and Winter 2014). Bernstein and Winter consider a participation game where each agent has two options representing participating or not (Bernstein and Winter 2012). Each agent's utility from participating is the payment given by the principal plus a value depending on other agents' participation. The principal aims to incentivize all agents to participate in the game with minimal payments.

Alon et al. study how to motivate multiple agents to take desirable actions using a common evaluation mechanism (Alon et al. 2020). They solve multiple problems in different settings. The main difference between their model and ours is the designer's payoff. The mechanism designer in their model cares about the number of agents who have been motivated in admissible ways while the principal in our model is interested in maximizing the reward generated by agents minus the payment paid to agents.

Lavi and Shamash study the model with multiple principals and multiple agents (Lavi and Shamash 2019). Agents do not have cost on actions. This model focuses on the competition between principals. McAfee and McMillan study another totally different problem where multiple agents compete for a principal's contract (McAfee and McMillan 1986). A recent work of Azizan et al. studies a model that is almost the same as ours where each agent can choose to directly produce a real number reward (Azizan et al. 2019). However, they assume the designed payment function can be parameterized by a vector in a given set $\mathcal{A} \subseteq \mathbb{R}^d$, and their algorithm explores the whole set \mathcal{A} , which is not that efficient.

Problem Description

In this paper, we study the Multiple Agents Contract Problem. There is a principal, n agents and m actions. Each agent can take an action $j \in [m]$ and produces a reward $\rho_j \geq 0$ for the principal. The reward only depends on the action, not on the agent. Each agent i also has a cost $c_{i,j} \geq 0$ to take an action j . This cost depends on both the agent and the action. Besides the m actions, there is always a *zero action* with reward 0 such that the cost for each agent to take this action is 0. This action means it is free for each agent to choose to produce nothing. The principal specifies a payment profile (t_1, t_2, \dots, t_m) : each agent taking action j will earn a

payment t_j . The utility for agent i to take action j is $t_j - c_{i,j}$. The agents are self-interested meaning each agent will take an action that maximizes its utility. W.l.o.g., we assume the agents tie-break in favor of the principal. The payoff of the principal is the sum of rewards produced by the agents minus the payments given to the agents, i.e., if agent i takes action i^* , the payoff of the principal is $\sum_{i=1}^n (\rho_{i^*} - t_{i^*})$. Our goal is to design the payment profile (t_1, t_2, \dots, t_m) to maximize the payoff of the principal.

Example 1. *Suppose there are two agents and two actions. The rewards for the two actions are 8 and 10 respectively. For action 1, agent 1 has a cost 5 and agent 2 has a cost 4. For action 2, agent 1 has a cost 9 and agent 2 has a cost 2. Without agent 2, we can set the payments for the two actions to 5 and 0 respectively, which brings a payoff of 3 to the principal. Without agent 1, we can set the payments for the two actions to 0 and 2 respectively, which brings a payoff of 8 to the principal. However, when the two agents both exist, no matter how we set the payments, the payoff of the principal cannot achieve $3 + 8 = 11$. It is optimal to the payments for the two actions to 5 and 3 respectively, which brings a payoff of 10 to the principal.*

Hardness

The problem defined in the previous section is very hard. To see its hardness, let us consider its decision version, i.e., the problem of determining whether there is a payment profile (t_1, t_2, \dots, t_m) such that the payoff of the principal is no less than a given number r . For convenience, we call this decision problem MAC. We will show in the following theorem that MAC is strongly NP-complete.

Theorem 1. *MAC is strongly NP-complete.*

Proof. MAC obviously belongs to NP. In the following proof, we reduce the well-known NP-complete problem Not-All-Equal 3-Satisfiability (NAE3SAT) to MAC to show that MAC is strongly NP-complete.

Given an instance of NAE3SAT with n variables and m clauses (we assume the variables in one clause are different without loss of generality), we build an instance of MAC as follows. For any variable x in an instance of NAE3SAT, we define x^0 as its negation and define $x^1 = x$.

- Agents
 - For each variable x_i , we have an agent A_i .
 - For each literal x_i^b and each clause c_j , we have an agent $T_{i,j}^b$.
 - For each clause c_j , we have 6 agents $V_{j,1}, V_{j,2}, \dots, V_{j,6}$.
- Actions
 - We have a zero action `zero` with reward 0.
 - For each literal x_i^b , we have an action `variable $_i^b$` with reward ρ_1 .
 - For each clause c_j , we have 6 actions `clause $_{j,1}, \dots, \text{clause}_{j,6}$` with reward ρ_2 .
- Costs
 - For the zero action `zero`, each agent has a cost 0.

- For action variable b_i , agent A_i has a cost δ , and $T_{i,j}^b$ has a cost 0 for each j .
- For action clause $_{j,k}$, agent $V_{j,k}$ has a cost 0.
- For action clause $_{j,k}$ where $c_j = x_{i_1}^{b_1} \vee x_{i_2}^{b_2} \vee x_{i_3}^{b_3}$, the costs vary for different k 's and are summarized in Table 1. Note there are exactly 3 agents with cost 1 to take this action. We call the three agents the *associated agents* of this action.
- For each action and agent, if we do not mention the cost above, the cost is greater than the reward of the action.

The parameters ρ_1, ρ_2, δ satisfy the following constraints.¹

$$\rho_1 - \delta > m((2n - 3)(\rho_2 - \rho_1) + n\delta + 4), \quad (1)$$

$$\delta > 3(\rho_2 - \rho_1 - 1), \quad (2)$$

$$\rho_2 - \rho_1 > 2. \quad (3)$$

Figure 1 shows an example instance of MAC corresponding to an instance of NAE3SAT with 4 variables and 2 clauses $x_1 \vee x_2 \vee \neg x_3$ and $x_1 \vee \neg x_2 \vee x_4$.

Then we ask whether we can set the payments to the agents so that the optimal payoff of the principal is no less than

$$n(\rho_1 - \delta) + m(6\rho_2 - 1) + m(n(\rho_1 - \delta) + (n - 3)\rho_1 + 3(\rho_2 - 1)). \quad (4)$$

In an optimal solution, the payment for an action will not exceed the reward, so an agent will never be incentivized to take an action whose cost is greater than the reward. The actions that an agent will be potentially incentivized to take in an optimal solution is summarized as follows (since all agents can take the zero action, we omit it in the following list).

- Agent A_i will potentially take action variable 0_i or variable 1_i .
- Agent $T_{i,j}^b$ will potentially take action variable b_i or, if variable x_i appears in clause c_j , clause $_{j,k}$ for some k .
- Agent $V_{j,k}$ will potentially take action clause $_{j,k}$.

Suppose the instance of NAE3SAT has a valid solution, then we set the payments in the instance of MAC as follows.

- For each action variable b_i , if the value of x_i^b is True, we set the payment to 0; otherwise we set the payment to δ .
- For each action clause $_{j,k}$ with associated agents $T_{i_1,j}^{h_1}, T_{i_2,j}^{h_2}, T_{i_3,j}^{h_3}$, we set the payment to 1 if the values of $x_{i_1}^{h_1}, x_{i_2}^{h_2}, x_{i_3}^{h_3}$ are all True; otherwise we set the payment to 0.

Under these payments, agent $V_{j,k}$ will always take action clause $_{j,k}$. If the value of x_i is True, agent A_i will take action variable 1_i ; otherwise she will take action variable 0_i . For agent $T_{i,j}^b$, if the value of x_i^b is True and variable x_i appears in clause c_j , she will take action clause $_{j,k}$ for some k ; otherwise she will take action

variable b_i . Hence the total payoff of the principal is exactly (4).

Now suppose there exists a payment setting such that the optimal payoff of the principal is no less than (4). We first show that agent A_i will take one of the actions variable 0_i and variable 1_i . Otherwise, the payoff of the principal cannot exceed $(n - 1)(\rho_1 - \delta) + 6m\rho_2 + 2mn\rho_2$, which is less than (4) by (1).

We define b_i such that A_i takes action variable $^{b_i}_i$, then the payment for action variable $^{b_i}_i$ must be no less than δ —the cost for agent A_i to take this action. If the payment is greater than δ , we can adjust it to δ . After this adjustment, some agent $T_{i,j}^{b_i}$ that takes action variable $^{b_i}_i$ before the adjustment may turn out to take action clause $_{j,k}$ for some k . This is the only possible cause of payoff loss of the principal. Suppose the payments for variable $^{b_i}_i$ (before the adjustment) and clause $_{j,k}$ are t_1, t_2 respectively, since agent $T_{i,j}^{b_i}$ chooses to take action variable $^{b_i}_i$ before the adjustment, we have $t_1 \geq t_2 - 1$, so $\rho_2 - t_2 \geq \rho_2 - t_1 - 1 > \rho_1 - t_1$ by (3). This means the adjustment does not reduce the payoff of the principal, hence we can assume the payment for action variable $^{b_i}_i$ is exactly δ . By an analogous argument, we can also assume the payment for action variable $^{1-b_i}_i$ is exactly 0.

If there exist some i, j such that agent $T_{i,j}^{b_i}$ does not take action variable $^{b_i}_i$, it must take action clause $_{j,k}$ for some k , and the payment t for action clause $_{j,k}$ must incentivize agent $T_{i,j}^{b_i}$ to take action clause $_{j,k}$, i.e., it must satisfy

$$t - 1 \geq \delta. \quad (5)$$

Then we adjust the payment for action clause $_{j,k}$ to 1 so that agent $T_{i,j}^{b_i}$ is incentivized to take action variable $^{b_i}_i$. After this adjustment, at most three agents that take action clause $_{j,k}$ before the adjustment deviate to take actions of the form variable $^{k'}_i$. Each of these agent brings a payoff of $\rho_2 - t$ to the principal before the adjustment, and brings a payoff of at least $\rho_1 - \delta$ after the adjustment, so the adjustment reduces the payoff of the principal by at most $3(\rho_2 - t - \rho_1 + \delta)$. On the other hand, the payoff of the principal increases by $t - 1$ due to the contribution of agent $V_{j,k}$. As a result, since $t - 1 \geq 3(\rho_2 - t - \rho_1 + \delta)$ due to (2) and (5), this adjustment does not reduce the payoff of the principal. Hence, we can assume for any i, j , agent $T_{i,j}^{b_i}$ takes action variable $^{b_i}_i$.

Suppose there exists some j such that the payment for action clause $_{j,k}$ is less than 1 for each k . Suppose clause c_j contains three variables $x_{i_1}, x_{i_2}, x_{i_3}$, and action clause $_{j,k_0}$ is an action that the cost for agent $T_{i_1,j}^{1-b_{i_1}}$ to take is 1 (there may exist multiple such k_0 's, and we arbitrarily choose one). We then adjust the payment for clause $_{j,k_0}$ to 1. This adjustment attracts $T_{i_1,j}^{1-b_{i_1}}$ to take action clause $_{j,k_0}$, which increases the payoff of the principal by $\rho_2 - 1 - \rho_1$. On the other hand, the payoff of the principal contributed by V_{j,k_0} is decreased by at most 1, which is the only cause that reduces the payoff of the principal. As a result, since $\rho_2 - 1 - \rho_1 > 1$, the payoff of the principal increases. Hence, we can assume

¹For example, we can set $\delta = 7, \rho_1 = 13mn + 8$ and $\rho_2 = 13mn + 11$.

Table 1: Cost Table

	$T_{i_1}^{b_1}$	$T_{i_2}^{b_2}$	$T_{i_3}^{b_3}$	$T_{i_1}^{1-b_1}$	$T_{i_2}^{1-b_2}$	$T_{i_3}^{1-b_3}$
clause $_{j,1}$	1	1	-	-	-	1
clause $_{j,2}$	1	-	1	-	1	-
clause $_{j,3}$	-	1	1	1	-	-
clause $_{j,4}$	-	-	1	1	1	-
clause $_{j,5}$	-	1	-	1	-	1
clause $_{j,6}$	1	-	-	-	1	1

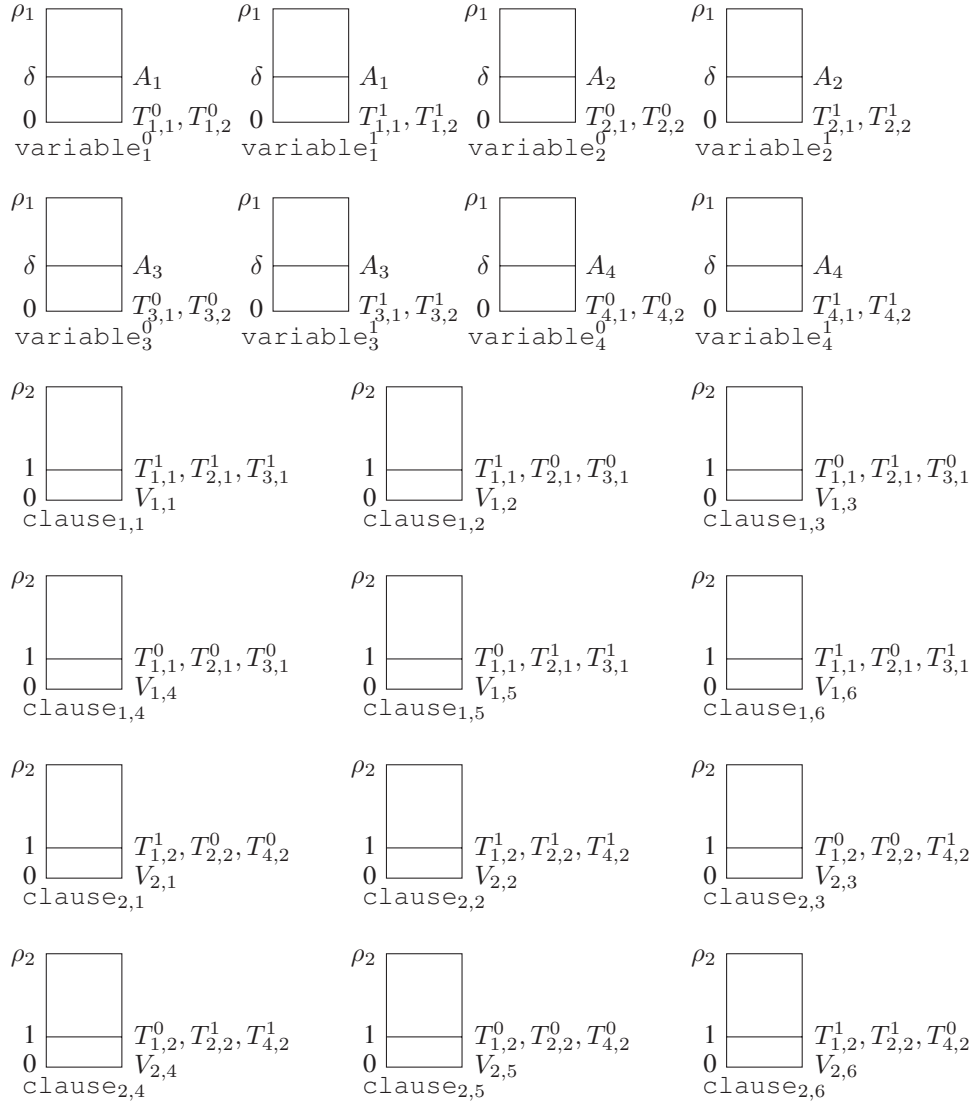


Figure 1: An example instance of MAC corresponding to an instance of NAE3SAT with 4 variables and 2 clauses $x_1 \vee x_2 \vee \neg x_3$ and $x_1 \vee \neg x_2 \vee x_4$: each rectangular represents an action; each line (including the bottom line) in a rectangular represents one or multiple agents, whose names are recorded to the right of the line; the number to the left of a line represents the cost for the agents to take this action; in particular, the number to the left of the top line of a rectangular represents the reward for the action.

for any j , there exists at least one k such that the payment for action $\text{clause}_{j,k}$ is no less than 1.

Under the assumptions above, the maximum payoff of the principal is exactly (4). To achieve this optimal payoff, for any j , say $c_j = x_{i_1}^{h_1} \vee x_{i_2}^{h_2} \vee x_{i_3}^{h_3}$, there exists exactly one k such that the payment for $\text{clause}_{j,k}$ is 1, and its three associated agents $T_{i_1,j}^{1-b_{i_1}}, T_{i_2,j}^{1-b_{i_2}}, T_{i_3,j}^{1-b_{i_3}}$ take this action. According to Table 1, $(1-b_{i_1}) \oplus h_1, (1-b_{i_2}) \oplus h_2, (1-b_{i_3}) \oplus h_3$ do not have the same value. So we can set variable x_i to the value $1-b_i$ (0 represents False and 1 represents True), then all clauses are satisfied. \square

Increasing Differences

In this section, we consider the case where agents have different abilities. Roughly speaking, the agents can be ordered from weak to strong, i_1, i_2, \dots, i_n , in the sense that it takes less cost for a stronger agent to produce a certain amount reward. We have for each $j \in [m]$,

$$c_{i_1,j} > c_{i_2,j} > \dots > c_{i_n,j}, \quad (6)$$

Additionally, we assume the costs obey *increasing differences*.

Definition 1. *Given an instance of the Multiple Agents Contract Problem, we call the costs obey increasing differences if there exists a permutation j_1, j_2, \dots, j_m of $1, 2, \dots, m$ and a permutation i_1, i_2, \dots, i_n of $1, 2, \dots, n$ such that for any pair of (k, k') such that $k < k', 0 < c_{i_k, j_1} - c_{i_{k'}, j_1} < c_{i_k, j_2} - c_{i_{k'}, j_2} < \dots < c_{i_k, j_m} - c_{i_{k'}, j_m}$.*

Though MAC is proved to be hard, we give a dynamic programming algorithm to solve the Multiple Agents Contract Problem under the assumption that the costs obey increasing differences.

The permutation i_1, i_2, \dots, i_n can be found in $O(n \log n)$ time by sorting $c_{1,j}, c_{2,j}, \dots, c_{n,j}$ for an arbitrary j , then the permutation j_1, j_2, \dots, j_m can be found in $O(m \log m)$ time by sorting $c_{i_2,1} - c_{i_1,1}, c_{i_2,2} - c_{i_1,2}, \dots, c_{i_2,m} - c_{i_1,m}$. For convenience, we assume the actions and agents are already ordered without loss of generality, i.e., $i_k = j_k = k$ for each k . The zero action is also considered action 0.

Before describing the algorithm, we first show the following lemmas.

Lemma 1. *If the costs obey increasing differences, then for any payment profile (t_1, t_2, \dots, t_m) , if agent i and i' take actions j and j' respectively, then $i < i' \Rightarrow j \leq j'$.*

Proof. Suppose $i < i'$ but $j > j'$. Since agent i prefers action j to j' , we have

$$t_j - c_{i,j} \geq t_{j'} - c_{i,j'}. \quad (7)$$

Similarly, since agent i' prefers action j' to j , we have

$$t_{j'} - c_{i',j'} \geq t_j - c_{i',j}. \quad (8)$$

By combining (7) and (8), we have

$$c_{i,j} - c_{i',j} \leq c_{i,j'} - c_{i',j'}. \quad (9)$$

However, since $i < i'$ and $j > j'$, by increasing differences we have $c_{i,j} - c_{i',j} > c_{i,j'} - c_{i',j'}$, which contradicts to (9). Therefore, we must have $i < i' \Rightarrow j \leq j'$. \square

Lemma 2. *Given any $0 \leq j_1 \leq \dots \leq j_m \leq m$, we have*

1. *Under the constraint that agent i is incentivized to take action j_i , the optimal payoff of the principal cannot exceed*

$$\sum_{i=1}^{n-1} (\rho_{j_i} - c_{i,j_i} - (n-i)(c_{i,j_i} - c_{i+1,j_i})) + \rho_{j_n} - c_{n,j_n}. \quad (10)$$

2. *If the costs obey increasing differences, and we set the payment profile (t_1, t_2, \dots, t_m) such that*

$$t_j = \begin{cases} \sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + c_{i,j_i}, & \text{if there exists} \\ & i \text{ such that} \\ & j = j_i^2, \\ 0, & \text{otherwise,} \end{cases} \quad (11)$$

then the payoff of the principal is no less than (10).

Proof. Since agent i' prefers action $j_{i'}$ to $j_{i'-1}$, we have

$$t_{j_{i'}} - c_{i',j_{i'}} \geq t_{j_{i'-1}} - c_{i',j_{i'-1}}, \quad (12)$$

and for $i' = 1$ we have $t_{j_1} - c_{1,j_1} \geq 0$ since agent 1 prefers action j_1 to the zero action. By summing up (12) for $i' = 1, 2, \dots, i$, we have

$$t_{j_i} \geq \sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + c_{i,j_i}.$$

Hence,

$$\begin{aligned} \sum_{i=1}^n t_{j_i} &\geq \sum_{i=1}^n \left(\sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + c_{i,j_i} \right) \\ &= \sum_{i=1}^{n-1} (c_{i,j_i} + (n-i)(c_{i,j_i} - c_{i+1,j_i})) + c_{n,j_n}, \end{aligned}$$

so the payoff of the principal cannot exceed (10).

On the other hand, suppose the costs obey increasing differences and we set t_j according to (11). For any agent i and any action j , there are three cases.

1. If there does not exist some k such that $j = j_k$, then

$$\begin{aligned} t_j - c_{i,j} &\leq 0 \\ &\leq \sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) \\ &= t_{j_i} - c_{i,j_i}, \end{aligned} \quad (13)$$

where the inequality (13) holds due to (6).

²If there exist multiple such i 's, we arbitrarily choose one, because if, for example, $j = j_{k'} = j_{k'+1} = \dots = j_k$, then the value of $\sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + c_{i,j_i}$ is the same for $i = k', k' + 1, \dots, k$.

2. If there exists some $k \leq i$ such that $j = j_k$, we have

$$\begin{aligned}
& t_j - c_{i,j} \\
&= \sum_{i'=1}^{k-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + c_{k,j_k} - c_{i,j_k} \\
&= \sum_{i'=1}^{k-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + \sum_{i'=k}^{i-1} (c_{i',j_k} - c_{i'+1,j_k}) \\
&\leq \sum_{i'=1}^{k-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) + \sum_{i'=k}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) \\
&\hspace{15em} (14) \\
&= \sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) \\
&= t_{j_i} - c_{i,j_i},
\end{aligned}$$

where the inequality (14) holds due to increasing differences: for any $i' \geq k$, $c_{i',j_k} - c_{i'+1,j_k} \leq c_{i',j_{i'}} - c_{i'+1,j_{i'}}$.

3. If there exists some $k > i$ such that $j = j_k$, we have

$$\begin{aligned}
& t_j - c_{i,j} \\
&= \sum_{i'=1}^{k-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) - (c_{i,j_k} - c_{k,j_k}) \\
&= \sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) \\
&\quad + \sum_{i'=i}^{k-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) - \sum_{i'=i}^{k-1} (c_{i',j_k} - c_{i'+1,j_k}) \\
&\leq \sum_{i'=1}^{i-1} (c_{i',j_{i'}} - c_{i'+1,j_{i'}}) \\
&= t_{j_i} - c_{i,j_i},
\end{aligned} \hspace{15em} (15)$$

where the inequality (15) holds due to increasing differences: for any $i' < k$, $c_{i',j_{i'}} - c_{i'+1,j_{i'}} \leq c_{i',j_k} - c_{i'+1,j_k}$.

Anyway, we have $t_j - c_{i,j} \leq t_{j_i} - c_{i,j_i}$, which means taking action j_i maximizes agent i 's utility. Note if agent i takes action j_i for each i , the payoff of the principal is exactly (10). Recall that the agents tie-break in favor of the principal, so the payoff of the principal is no less than (10). \square

Lemma 1 and 2 show that we can find $0 \leq j_1 \leq j_2 \leq \dots \leq j_m \leq m$ that maximizes (10), then an optimal payment profile is given by (11). To find the optimal j_1, j_2, \dots, j_m , we use a dynamic programming algorithm. For convenience, we define $\phi(i, j) = \rho_j - c_{i,j} - (n-i)(c_{i,j} - c_{i+1,j})$ for $i = 1, 2, \dots, n-1$, and define $\phi(n, j) = \rho_j - c_{n,j}$. We define the subproblem $\text{OPT}(i, j) = \max_{0 \leq j_1 \leq j_2 \leq \dots \leq j_i \leq j} \sum_{i'=1}^i \phi(i', j_{i'})$. We can see the optimal value of (10) is $\text{OPT}(n, m)$, and we have the recursion

formula

$$\begin{aligned}
& \text{OPT}(i, j+1) = \\
& \max_{0 \leq k \leq i} \left(\text{OPT}(k, j) + \sum_{i'=k+1}^i \phi(i', j+1) \right)
\end{aligned}$$

with $\text{OPT}(i, 0) = 0$ for each i . Hence, the optimal value of (10), as well as the optimal j_1, j_2, \dots, j_m , can be computed in $O(n^2m)$ time. We conclude the result above as the following theorem.

Theorem 2. *If the costs obey increasing differences, there is an $O(n^2m)$ algorithm solving the Multiple Agents Contract Problem.*

Real Number Actions

In previous sections, we considered the Multiple Agents Contract Problem with discrete actions (DA). A natural generalization is to consider the problem where each agent can choose to produce an arbitrary reward in $[0, 1]$. We call this generalization Multiple Agents Contract Problem with Real Number Actions (RNA), and formalize it as follows.

There is a principal and n agents. Each agent chooses to produce a reward $x \in [0, 1]$ for the principal. To take such an action, each agent has a cost which may differ from each other. We define $c_i(x) \geq 0$ as the cost for agent i to produce a reward x . We assume without loss of generality that $c_i(0) = 0$ for all i , which means it is free for each agent to choose to produce nothing. To incentivize these agents to produce rewards, the principal specifies a payment function $t(x)$: each agent taking this action will earn a payment $t(x)$. The utility for agent i to produce x is $t(x) - c_i(x)$. Agents are self-interested, meaning each agent will produce a reward that maximizes her utility. We assume agents tie-break in favor of the principal. The payoff of the principal is the sum of the rewards produced by these agents minus the payments given to the agents, i.e., if agent i produces a reward x_i , the payoff of the principal is $\sum_{i=1}^n (x_i - t(x_i))$. Our goal is to design the payment function to maximize the payoff of the principal.

Note in this paper, the functions t and c_i 's are not necessarily continuous. To guarantee every agent has an optimal action we only concern the payment function t where for all i , $t(x) - c_i(x)$ and $x - t(x)$ (in case of tie-breaking) are able to attain their maximums on $[0, 1]$.

Hardness

We first show that this generalization is harder than our original problem by a reduction from DA to RNA. Given an instance of DA, we can construct an instance of RNA by

letting

$$c_i(x) = \begin{cases} 0, & \text{if } x = 0, \\ \frac{c_{i,1}+M}{\rho_m+mM}, & \text{if } 0 < x \leq \frac{\rho_1+M}{\rho_m+mM}, \\ \vdots \\ \frac{c_{i,j}+jM}{\rho_m+mM}, & \text{if } \frac{\rho_{j-1}+(j-1)M}{\rho_m+mM} < x \leq \frac{\rho_j+M}{\rho_m+mM}, \\ \vdots \\ \frac{c_{i,m}+mM}{\rho_m+mM}, & \text{if } \frac{\rho_{m-1}+(m-1)M}{\rho_m+mM} < x \leq 1, \end{cases} \quad (16)$$

for each i , where M is a large enough number³. We will show how to construct an optimal payment profile of the DA instance from an optimal payment function of the RNA instance. For convenience, we define $z_j = (\rho_j + jM)/(\rho_m + mM)$ and $z_0 = 0$.

Given an optimal payment function $t(x)$ of the RNA instance, suppose agent i chooses to produce x_i and define j_i such that $z_{j_i-1} < x_i \leq z_{j_i}$ (if $x_i = 0$, then $j_i = 0$). Now consider a fixed i . If $x_i < z_{j_i}$, we adjust the value of $t(x)$ at $x = z_{j_i}$ to $t(x_i)$. Before this adjustment, agent i produces x_i , and after this adjustment, agent i has the same utility to produce z_{j_i} as to produce x_i , so agent i will produce z_{j_i} after the adjustment (recall the agent tie-breaks in favor of the principal), which increases the payoff of the principal. On the other hand, for any other agent i' , $t(x_i) - c_{i'}(z_{j_i}) \leq t(x_i) - c_{i'}(x_i)$ (since the cost function is weakly increasing), which means the utility of producing z_{j_i} after the adjustment does not exceed that of producing x_i . Hence, for any agent except i , changing her produced value to z_{j_i} due to the adjustment does not decrease the payoff of the principal (recall again that the agents tie-break in favor of the principal). As a result, the payoff of the principal is increased by this adjustment, which contradicts to the fact that $t(x)$ is optimal. Therefore, we can assume $x_i = z_{j_i}$. The payoff of the principal under the payment function $t(x)$ in the RNA instance is $p_{\text{RNA}} = \sum_{i=1}^n (z_{j_i} - t(z_{j_i}))$.

Now we construct a payment profile (t_1, t_2, \dots, t_m) of the DA instance where

$$t_j = t(z_j)(\rho_m + mM) - jM. \quad (17)$$

Under this payment profile, for each agent i and each j , the utility of agent i to take action j is $t_j - c_{i,j}$, which is exactly $(\rho_m + mM)$ times the utility of agent i to produce z_j under the payment function $t(x)$ in the RNA instance. Also, agent i brings a payoff of $\rho_j - t_j$ to the principal by taking action j , which is exactly $(\rho_m + mM)$ times the payoff of the principal brought by agent i by producing z_j under the payment function $t(x)$ in the RNA instance. Since agent i produces z_{j_i} under payment function $t(x)$ in the RNA instance, she will take action j_i under payment profile (t_1, t_2, \dots, t_m) in the DA instance. The payoff of the principal under the payment profile (t_1, t_2, \dots, t_m) in the DA instance is $p_{\text{DA}} = \sum_{i=1}^n (\rho_{j_i} - t_{j_i}) = (\rho_m + mM)p_{\text{RNA}}$.

To show the payment profile (t_1, t_2, \dots, t_m) is optimal, we compare it to another arbitrary payment profile $(t'_1, t'_2, \dots, t'_m)$. Suppose agent i takes action j'_i under the

payment profile $(t'_1, t'_2, \dots, t'_m)$, then the payoff of the principal under the payment profile $(t'_1, t'_2, \dots, t'_m)$ in the DA instance is $p'_{\text{DA}} = \sum_{i=1}^n (\rho_{j'_i} - t'_{j'_i})$.

Let

$$t'(x) = \begin{cases} 0, & \text{if } x = 0, \\ \frac{t'_1+M}{\rho_m+mM}, & \text{if } 0 < x \leq z_1, \\ \vdots \\ \frac{t'_j+jM}{\rho_m+mM}, & \text{if } z_{j-1} < x \leq z_j, \\ \vdots \\ \frac{t'_m+mM}{\rho_m+mM}, & \text{if } z_{m-1} < x \leq z_m = 1. \end{cases}$$

We can see for all i , $x - t'(x)$ and $t'(x) - c_i(x)$ are able to attain their maximum on $[0, 1]$, so $t'(x)$ is a valid payment function. Under this payment function, agent i has the same utility for producing a reward on $(z_j, z_{j+1}]$, thus she will produce z_j for some j in favor of the principal. Observe, again, that under the payment profile $(t'_1, t'_2, \dots, t'_m)$, for each agent i and each j , the utility of agent i to take action j is $t'_j - c_{i,j}$, which is exactly $(\rho_m + mM)$ times the utility of agent i to produce z_j under the payment function $t'(x)$ in the RNA instance. Also, agent i brings a payoff of $\rho_j - t'_j$ to the principal by taking action j , which is exactly $(\rho_m + mM)$ times the payoff of the principal brought by agent i by producing z_j under the payment function $t'(x)$ in the RNA instance. Hence, agent i will produce $z_{j'_i}$ under the payment function $t'(x)$ in the RNA instance. The payoff of the principal under the payment function $t'(x)$ in the RNA instance is $p'_{\text{RNA}} = \sum_{i=1}^n (z_{j'_i} - t'(z_{j'_i})) = p'_{\text{DA}}/(\rho_m + mM)$.

Hence, $p'_{\text{DA}} = (\rho_m + mM)p'_{\text{RNA}} \leq (\rho_m + mM)p_{\text{RNA}} = p_{\text{DA}}$, which means (t_1, t_2, \dots, t_m) is indeed an optimal payoff profile of the DA instance.

An Approximate Contract

Knowing the RNA problem is hard, we are going to design an approximate contract. We assume for all i , $x - c_i(x)$ is able to attain its maximum on $[0, 1]$. Let $x_i \in \arg \max_{x \in [0,1]} (x - c_i(x))$ (if there are multiple x_i 's achieving the maximum value, we arbitrarily choose one), $y_i = \max_{x \in [0,1]} (x - c_i(x))$, we have immediately

$$y_i = x_i - c_i(x_i) \leq x_i. \quad (18)$$

Let

$$t_i(x) = \begin{cases} 0, & \text{if } 0 \leq x \leq y_i, \\ x - y_i, & \text{if } y_i < x \leq 1. \end{cases}$$

We assume without loss of generality that $y_1 \leq y_2 \leq \dots \leq y_n$. We first show that $t_i(x)$ is a valid payment function, i.e. for all i' , $x - t_i(x)$ and $t_i(x) - c_{i'}(x)$ are able to attain their maximum on $[0, 1]$. The former is trivial. For $t_i(x) - c_{i'}(x)$, if $0 \leq x \leq y_i$, then $t_i(x) - c_{i'}(x) = -c_{i'}(x) \leq 0$; if $y_i < x \leq 1$, then $t_i(x) - c_{i'}(x) = x - y_i - c_{i'}(x) \leq y_{i'} - y_i$, so $t_i(x) - c_{i'}(x) \leq \max\{0, y_{i'} - y_i\}$. In addition, $t_i(0) - c_{i'}(0) = 0$ and $t_i(x_{i'}) - c_{i'}(x_{i'}) \geq x_{i'} - y_i - c_{i'}(x_{i'}) = y_{i'} - y_i$. This means the maximum value of $t_i(x) - c_{i'}(x)$

³It is sufficient to choose $M = \max_{i,j} \{c_{i,j}, \rho_j\} + 1$.

is $\max\{0, y_{i'} - y_i\}$, and is achievable at $x = 0$ or $x = x_{i'}$. Hence, $t_i(x)$ is indeed a valid payment function.

Note the argument above also shows that for any $i' \geq i$, $t_i(x) - c_{i'}(x)$ attains its maximum at $x = x_{i'}$. By (18), we have $x_{i'} \geq y_{i'} \geq y_i$, so if agent i' chooses to produce $x_{i'}$, she brings a payoff of $x_{i'} - t_i(x_{i'}) = x_{i'} - (x_{i'} - y_i) = y_i$ to the principal. Recall that the agents tie-break in favor of the principal, agent i' brings a payoff of at least y_i to the principal. Hence, under the payment function $t_i(x)$, the payoff of the principal is at least $(n - i + 1)y_i$. Let $i^* \in \arg \max_i (n - i + 1)y_i$, then we have for all i ,

$$y_i \leq \frac{(n - i^* + 1)y_{i^*}}{n - i + 1}.$$

On the other hand, let OPT denote the optimal payoff of the principal. Since agent i brings a payoff of at most $\max_{x \in [0,1]}(x - c_i(x)) = y_i$ to the principal, we have $\text{OPT} \leq \sum_{i=1}^n y_i$. Hence,

$$\text{OPT} \leq \sum_{i=1}^n y_i \leq (n - i^* + 1)y_{i^*} \sum_{i=1}^n \frac{1}{n - i + 1}.$$

This means the payment function $t_{i^*}(x)$ is an $\sum_{i=1}^n (1/(n - i + 1))$ -approximate solution, i.e. an $O(\log n)$ -approximate solution.

In conclusion, we have the following algorithm.

1. For any i , find $y_i = \max_{x \in [0,1]}(x - c_i(x))$ and sort them such that $y_1 \geq y_2 \geq \dots \geq y_n$.
2. Let $i^* \in \arg \max_{1 \leq i \leq n} (n - i^* + 1)y_{i^*}$.
3. Output the payment function

$$t(x) = \begin{cases} 0, & \text{if } 0 \leq x \leq y_{i^*}, \\ x - y_{i^*}, & \text{if } y_{i^*} < x \leq 1. \end{cases}$$

Note this algorithm can also be applied to our Multiple Agents Contract Problem with discrete actions. We first construct an instance of the problem with real number actions using (16), obtain a payment function using the algorithm above, then we can get a payment profile of the problem with discrete actions using (17). This payment profile is still an $O(\log n)$ -approximate solution.

References

Alon, T.; Dobson, M. R. C.; Procaccia, A.; Talgam-Cohen, I.; and Tucker-Foltz, J. 2020. Multiagent evaluation mechanisms. In *Proceedings of the AAAI Conference on Artificial Intelligence*, forthcoming.

Armstrong, M., and Vickers, J. 2010. A model of delegated project choice. *Econometrica* 78(1):213–244.

Azizan, N.; Su, Y.; Dvijotham, K.; and Wierman, A. 2019. Optimal pricing in markets with non-convex costs. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 595–595. ACM.

Babaioff, M., and Winter, E. 2014. Contract complexity. In *EC*, 911.

Babaioff, M.; Feldman, M.; Nisan, N.; and Winter, E. 2012. Combinatorial agency. *Journal of Economic Theory* 147(3):999–1034.

Babaioff, M.; Feldman, M.; and Nisan, N. 2006a. Combinatorial agency. In *Proceedings of the 7th ACM conference on Electronic commerce*, 18–28. ACM.

Babaioff, M.; Feldman, M.; and Nisan, N. 2006b. Mixed strategies in combinatorial agency. In *International Workshop on Internet and Network Economics*, 353–364. Springer.

Bernstein, S., and Winter, E. 2012. Contracting with heterogeneous externalities. *American Economic Journal: Microeconomics* 4(2):50–76.

Bolton, P.; Dewatripont, M.; et al. 2005. *Contract theory*. MIT press.

Chen, M.; Tang, P.; Wang, Z.; Xiao, S.; and Yang, X. 2019. Optimal mechanisms with budget for user generated contents. *arXiv preprint arXiv:1907.04740*.

Grossman, S. J., and Hart, O. D. 1992. An analysis of the principal-agent problem. In *Foundations of Insurance Economics*. Springer. 302–340.

Holmstrom, B., and Milgrom, P. 1987. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica: Journal of the Econometric Society* 303–328.

Holmstrom, B., and Milgrom, P. 1991. Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *JL Econ. & Org.* 7:24.

Kleinberg, J., and Kleinberg, R. 2018. Delegated search approximates efficient search. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, 287–302. ACM.

Lavi, R., and Shamash, E. 2019. Principal-agent vcg contracts. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 783–783. ACM.

McAfee, R. P., and McMillan, J. 1986. Bidding for contracts: a principal-agent analysis. *The RAND Journal of Economics* 326–338.

Winter, E. 2004. Incentives and discrimination. *American Economic Review* 94(3):764–773.